




VANCOUVER 2019

Dec 8th - 14th



CONFERENCE AT A GLANCE

MONDAY DECEMBER 9th

8:30 - 10:30 AM	Tutorials Session 1
10:30 - 11:15 AM	Coffee break
11:15 - 1:15 PM	Tutorials Session 2
1:15 - 2:45 PM	Lunch on your own
2:45 - 4:45 PM	Tutorials Session 3
4:45 - 5:00 PM	Break
5:00 - 5:45 PM	Opening Remarks
5:45 - 6:35 AM	Invited talk, Celeste Kidd: <i>How To Know</i>
6:35 - 8:30 PM	Opening Reception

TUESDAY DECEMBER 10th

8:30 - 9:20 AM	Invited talk, Bin Yu: <i>Three principles of data science and interpretable machine learning with case studies</i>
9:20 - 10:05 AM	Coffee break
10:05 - 10:45 PM	Parallel Tracks
10:45 - 12:45 PM	Poster Sessions A and Demonstrations
12:45 - 2:15 PM	Lunch on your own
2:15 - 3:05 PM	Invited talk, Dana Pe'er: <i>Machine learning meets single-cell biology: insights and challenges</i>
3:05 - 3:25 AM	Test Of Time Award: <i>Lin Xiao</i> <i>Dual Averaging Method for Regularized Stochastic Learning and Online Optimization</i>
3:25 - 4:10 PM	Coffee break
4:10 - 5:30 PM	Parallel Tracks
5:30 - 7:30 PM	Poster Sessions B and Demonstrations
7:00 - 10:00 PM	NeurIPS Socials

WEDNESDAY DECEMBER 11th

8:30 - 9:20 AM	Invited talk, Blaise Aguerre y Arcas: <i>Social Intelligence</i>
9:20 - 10:05 AM	Coffee break
10:05 - 10:45 PM	Parallel Tracks
10:45 - 12:45 PM	Poster Sessions A and Demonstrations
12:45 - 2:15 PM	Lunch on your own
2:15 - 3:05 PM	Invited talk, Yoshua Bengio: <i>From System 1 Deep Learning to System 2 Deep Learning</i>
3:05 - 3:50 PM	Coffee break
3:50 - 5:00 PM	Parallel Tracks
5:00 - 7:00 PM	Poster Sessions B and Demonstrations
7:00 - 10:00 PM	NeurIPS Socials

THURSDAY DECEMBER 12th

8:30 - 9:20 AM	Invited talk, Kafui Dzirasa: <i>Mapping emotions: Discovering structure in mesoscale electrical brain recordings</i>
9:20 - 10:05 AM	Coffee break
10:05 - 10:45 AM	Parallel Tracks
10:45 - 12:45 PM	Poster Sessions A
12:45 - 2:15 PM	Lunch on your own
2:15 - 3:05 PM	Invited talk, Jeff Heer: <i>Agency + Automation: Designing Artificial Intelligence into Interactive Systems</i>
3:05 - 3:50 PM	Coffee break
3:50 - 5:00 PM	Parallel Tracks
5:00 - 7:00 PM	Poster Sessions B
7:00 - 10:00 PM	NeurIPS Socials

FRIDAY & SATURDAY DECEMBER 13th & 14th

Each workshop has its own schedule, check the website

8:00 - 6:00 PM	Workshops
9:45 - 10:30 AM	Coffee break
12:00 - 2:00 PM	Lunch on your own
3:30 - 4:15 PM	Coffee Break
6:00 - 10:00 PM	Saturday night reception

Table Of Contents

Conference At A Glance	2
Committees & NIPS Officers	3
General Information	4
Conference Maps	5 - 6
NeurIPS socials	7
Monday Schedule	8
Monday Tutorial Sessions	8
Tuesday Full Schedule	12
Tuesday Invited Speakers	13
Tuesday Poster Sessions	14
Tuesday Demos	15
Wednesday Full Schedule	16
Wednesday Invited Speakers	17
Wednesday Poster Sessions	18
Wednesday Demos	19
Thursday Full Schedule	20
Thursday Invited Speakers	21
Thursday Poster Sessions	22
Workshops (Fri & Sat)	23
Sponsors	24

ORGANIZING COMMITTEE

General Chairs

Hanna Wallach, Microsoft Research

Program Chair

Hugo Larochelle, Google Brain

Program Co-Chairs

Alina Beygelzimer, Yahoo! Research

Florence d'Alché-Buc, Télécom ParisTech

Emily Fox, University of Washington & Apple

Diversity and Inclusion Chairs

Katherine Heller, Duke University & Google

Charles Isbell, Georgia Tech

Tutorials Chairs

Alice Oh, KAIST

Danielle Belgrave, Microsoft Research

Workshop Chair

Jenn Wortman Vaughan, Microsoft Research

Workshop Co-Chairs

Marzyeh Ghassemi, University of Toronto

Shakir Mohammed, DeepMind

Bob Williamson, ANU & DATA61

Demonstration and Competition Chairs

Hugo Jair Escalante, INAOE

Raia Hadsell, DeepMind

Communication Chairs

Katherine Gorman, Collective Next

Michael Littman, Brown University

Reproducibility Chairs

Joelle Pineau, McGill & FAIR

Koustuv Sinha, McGill

Social Chairs

Maria Skoularidou, University of Cambridge

Gergely Neu, Universitat Pompeu Fabra

Meetup Chairs

Nicolas Le Roux, Google Brain

Anoush Najarian, MathWorks

Publication Chair

Roman Garnett, Washington U. in St. Louis

Party Chairs

Pablo Samuel Castro, Google Brain

Hassan Kane, Sela Labs

Expo Chair

Neil Lawrence, Amazon

FOUNDATION OFFICERS & BOARD MEMBERS

PRESIDENT

Terrence Sejnowski, The Salk Institute

TREASURER

Marian Bartlett, Apple

SECRETARY

Michael Mozer, UC Boulder

EXECUTIVE DIRECTOR

Mary Ellen Perry, The Salk Institute

LEGAL ADVISOR

David Kirkpatrick

IT DIRECTOR

Lee Campbell, The Salk Institute

SUPPORT TEAM

The organization and management of NeurIPS would not be possible without the help of many volunteers, students, researchers and administrators who donate their valuable time and energy to assist the conference in various ways. The support staffs' tireless efforts make the conference run smoothly and efficiently every year. NeurIPS would particularly like to acknowledge the exceptional work of:

Brad Brockmeyer

Sherry Xue

Susan Westphal

Darren Nelson

Terrance Gaines

Jen Perry

Kevin Waldrop

EXECUTIVE BOARD

Samy Bengio, Google Research

Corinna Cortes, Google Research

Isabelle Guyon, U. Paris-Saclay & ChaLearn

Neil D. Lawrence, Cambridge University

Daniel D. Lee, University of Pennsylvania

Ulrike von Luxburg, University of Tübingen

Masashi Sugiyama, RIKEN & The University of Tokyo

Hanna Wallach, Microsoft Research

Max Welling, University of Amsterdam

ADVISORY BOARD

Peter Bartlett, Queensland U.,
UC Berkeley

Sue Becker, McMaster U.
Ontario, Canada

Yoshua Bengio, U. of Montreal,
Canada

Léon Bottou, Microsoft
Research

Chris J.C. Burges, Microsoft
Research

Jack Cowan, U. of Chicago

Thomas G. Dietterich, Oregon
State U.

Zoubin Ghahramani, U. of
Cambridge

Stephen Hanson, Rutgers U.

Michael I. Jordan, UC Berkeley

Michael Kearns, U. of
Pennsylvania

Scott Kirkpatrick, Hebrew U.
Jerusalem

Daphne Koller, Stanford U.

John Lafferty, Yale U.

Todd K. Leen

Richard Lippmann, MIT

Bartlett Mel, U. of Southern
California

John Moody, JEM

John C. Platt, Google

Fernando Pereira, Google
Research

Gerald Tesauro, IBM Watson
Labs

Sebastian Thrun, Stanford U.

Dave Touretzky, Carnegie
Mellon U.

Lawrence Saul, UC San Diego

Bernhard Schölkopf, MPI,
Tübingen/Stuttgart

Dale Schuurmans, U. of
Alberta, Canada

John Shawe-Taylor, U. College
London

Sara A. Solla, Northwestern U.

Yair Weiss, Hebrew U. of
Jerusalem

Chris Williams, U. of Edinburgh

Rich Zemel, U. of Toronto

GENERAL INFORMATION

REGISTRATION DESK

Sunday, Dec. 8th: 8 am – 8 pm
Monday: 6:30 am – 7 pm
Tuesday – Friday: 7 am – 6 pm
Saturday: 7 am – 12 pm

RECEPTIONS

To accommodate all attendees, F&B stations are located in the East AND West buildings. If one location is too crowded, there are other locations available. Coffee stations will be in both East and West. See the Catering Map on the next page.

Opening Reception - Monday, Dec 9th, 6:35- 8:30pm

East Exhibition Hall A & C, Ballrooms B & C

West Exhibition Hall B1+B2 & Ballroom D

****The Affinity Groups Poster Sessions...**

East Exhibition Hall B

Closing Reception - Saturday, Dec 14th, 6 - 10pm

East Exhibition Hall A-C, Ballrooms A-C (music provided).

East Meeting Level 1-3 - Designated quiet area

POSTER SESSIONS

Tuesday - Thursday

Poster A Session: 10:45 pm – 12:45 pm

Poster B Session: 5:30 pm – 7:30 pm (Tues)

Poster B Session: 5 pm – 7 pm (Wed + Thurs)

Posters still up after 8:30 pm will be discarded.

COAT AND LUGGAGE CHECK, NO CHARGE

East, Meeting Level 18,19, 20.

Luggage check on Monday, Friday and Saturday

WIFI

SSID: neurips

Password: conference

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. You're all set!

Now you will be able to:

- View the event agenda and plan your schedule.
- If you set up your own profile, you can send in-app messages and exchange contact information
- Receive update notifications from organizers.
- Access agenda, maps, and 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: **neurips**

CHARGING TABLES

Located throughout the venue

SPONSOR BOOTHS & BOOK PUBLISHERS

Sponsor booths located in East Exhib. Hall A, Ballrooms B + C

Coffee served all day Monday - Wednesday, 9am - 5pm

Book publishers in the East Foyer

NeurIPS would like to especially thank Slides Live for streaming services.



NeurIPS would like to especially thank Microsoft Research for their donation of Conference Management Toolkit (CMT) software and server space.



FUTURE CONFERENCES

2020 - Vancouver, Canada




2021 - Sydney Australia



CONFERENCE MAP - WEST LEVELS




WEST LEVEL 1




- Workshops
- Food & Beverage
- Competitions
- Socials
- Interview Rooms

WEST LEVEL 3




WEST LEVEL 2




CONFERENCE MAP - EAST LEVELS


EAST EXHIBITION LEVEL



EAST MEETING LEVEL



SPONSOR MAPS



SOCIALS AT NEURIPS

We are very happy to inaugurate Social Events at NeurIPS 2019. Given the steadily increasing amount of attendees of NeurIPS, it's becoming more and more difficult to meet colleagues with similar interests each year. Social events address this issue by providing a meeting place for attendees with similar interests. The 15 socials we selected to host this year capture the diverse interests of the NeurIPS community, addressing topics ranging from core technical content to societal matters within the community and beyond. We are proud to announce the program below.

TUESDAY 7-10 PM

Well-Being In AI

Room 202 - 204

Contact: Alfredo Kalaitzis
freddie@element.ai

The Well-Being in ML (WBiML) event is an opportunity to make the ML community mindful of well-being at the very event that epitomises its science.

AI For Good

Room 205 - 207

Contact: Devin Krotman
Devin.Krotman@xprize.org

A networking happy hour, followed by a series of lightning talks from practitioners in the AI for Good applications space.

Learning Theory

Room 217 - 219

Contact: Vianney Perchet
vianney.perchet@gmail.com

A non-random friendly meeting point for people working on or being interested in learning theory.

Reinforcement Learning Social

Room 220 - 222

Contact: Yuxi Li
yuxili@attain.ai

An event for people working on or being interested in RL in order to 1) socialise with experts and 2) PDP their meetups.

Inverse Problems Social

Room 223 - 224

Contact: Ajil Jalal
ajiljalal@utexas.edu

An informal event for networking and socializing amongst (but not restricted to) the attendees of the "Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications" workshop.

WEDNESDAY 7-10 PM

Find your allies: how to be an individual champion of ethical AI practices at your company & meet collaborators you need

Room 202 - 204

Contact: Jingying Yang
jingying@partnershiponai.org

An inviting event to catalyze collaborations between individuals in the NeurIPS community who care about AI ethics and opening that conversation up to the many more individuals who find themselves with a similar desire to make a positive contribution to people and society through their work in the ML/AI field.

Women in AI Ignite

Room 205 - 207

Contact: Anoush Najarian
anoushn@mathworks.com

This event aims at building a platform for Women in AI Ignite speakers, giving participants the opportunity to see the power of tools like Ignite to bring about change and building a plan, rooted in social science and intersectionality.

ML 4 Space Social

Room 217 - 219

Contact: Jodie Hughes
jodie@frontierdevelopmentlab.org

An event to invite people to explore the opportunities for applied ML in space exploration and stewardship of our home planet.

Probabilistic Programming Social

Room 220 - 222

Contact: Vikash K. Mansinghka
vkm@mit.edu

An informal, inclusive event for networking and socializing among probabilistic programming students and researchers.

T-PRIME

Room 223 - 224

Contact: Martin Trapp
trapp.martin@gmail.com

T-PRIME (Tractable PProbabilistic Inference MEeting) is the first social event of a series for researchers and practitioners working on (or interested in) tractable inference to join forces.

THURSDAY 7-10 PM

Social Event for Budding Researchers

Room 202 - 204

Contact: Prabhu Pradhan
prabhuspradhan@gmail.com

Buds@NeurIPS is an event that aims at gathering young researchers (especially Pre-PhD participants) and mixing them with domain-experts to revitalize grey cells after a busy conference day.

Effective Altruism social

Room 205 - 207

Contact: Claudia Shi
(claudia.j.shi@gmail.com)

An event with informal presentations on how to address pressing issues in the world as ML researchers, open debate about the impact of working on long-term AI safety versus ML for social impact and academic speed friending on high social impact projects.

British Parliamentary style debate

Room 217 - 219

Contact: Jonathan Hunt
j@me.net.nz

An informal event with British parliament style debates amongst several specialists in AI.

{Dis}Ability in AI

Room 220 - 222

Contact: Maria Skoularidou
ms2407@cam.ac.uk

Panel discussion aiming at raising awareness, supporting and advocating for disabled people.

Deep Learning Researchers and Start-Ups: Pros and Cons of Working at Start-Ups vs Large Companies, and How to Fix Diversity Issues

Room 223 - 224

Contact: Colin White
colin@realityengines.ai

The goal of this social is to get researchers to consider start-ups as a viable career option in the industry and to find out more about the pros and cons of working at smaller companies.

MONDAY SCHEDULE

8:30 am - 10:30 am - Tutorials Session 1

Imitation Learning and its Application to Natural Language Generation

Kyunghyun Cho · Hal Daume III

West Exhibition Hall C + B3

Human Behavior Modeling with Machine Learning: Opportunities & Challenges

Nuria M Oliver · Albert Ali Salah

West Ballroom A + B

Deep Learning with Bayesian Principles

Mohammad Emtiyaz Khan

West Exhibition Hall A

10:30 - 11:15 am - Coffee Break

11:15 am -- 1:15 pm - Tutorials Session 2

Efficient Processing of Deep Neural Network: from Algorithms to Hardware Architectures

Vivienne Sze

West Exhibition Hall C + B3

Interpretable Comparison of Distributions and Models

Wittawat Jitkrittum · Dougal J Sutherland · Arthur Gretton

West Ballroom A + B

Machine Learning for Computational Biology and Health

Anna Goldenberg · Barbara Engelhardt

West Exhibition Hall A

1:15 - 2:45 pm - Lunch On Your Own

2:45 - 4:45 pm - Tutorials Session 3

Reinforcement Learning: Past, Present and Future Perspectives

Katja Hofmann

West Exhibition Hall C + B3

Synthetic Control

Alberto Abadie · Vishal Misra · Devavrat Shah

West Ballroom A + B

Representation Learning and Fairness

Moustapha Cisse · Sanmi Koyejo

West Exhibition Hall A

5:00 - 5:45 pm - Opening Remarks

West Exhibition Hall C + B3

5:45 - 6:35 pm - Invited Talk: Celeste Kidd

How To Know

West Exhibition Hall C + B3


6:35 pm - 8:30 pm: Opening Reception

East Exhibition A, Ballrooms B + C



Tutorials Session 1 - 8:30 - 10:30 am

Imitation Learning and its Application to Natural Language Generation



Kyunghyun Cho (New York U.)
Hal Daume III (Microsoft, U. of Maryland)

Location: West Exhibition Hall C + B3

Imitation learning is a learning paradigm that interpolates reinforcement learning on one extreme and supervised learning on the other extreme. In the specific case of generating structured outputs—as in natural language generation—imitation learning allows us to train generation policies with neither strong supervision on the detailed generation procedure (as would be required in supervised learning) nor with only a sparse reward signal (as in reinforcement learning). Imitation learning accomplishes this by exploiting the availability of potentially suboptimal “experts” that provide supervision along an execution trajectory of the policy. In the first part of this tutorial, we overview the paradigm of imitation learning and a suite of practical imitation learning algorithms. We then consider the specific application of natural language generation, framing this problem as a sequential decision making process. Under this view, we demonstrate how imitation learning could be successfully applied to natural language generation and open the door to a range of possible ways to learn policies that generate natural language sentences beyond naive left-to-right autoregressive generation.

Human Behavior Modeling with Machine Learning: Opportunities and Challenges



Nuria M Oliver (Microsoft Research)
Albert Ali Salah (Bogazici U.)

Location: West Ballroom A + B

Human behavior is complex, multi-level, multimodal, culturally and contextually shaped. Computer analysis of human behavior in its multiple scales and settings leads to a steady influx of new applications in diverse domains including human-computer interaction, affective computing, social

signal processing and computational social sciences, autonomous systems, smart healthcare, customer behavior analysis, urban computing and AI for social good. In this tutorial, we will share a proposed taxonomy to understand, model and predict both individual, dyadic and aggregate human behavior from a variety of data sources and using machine learning techniques. We will illustrate this taxonomy through relevant examples from the literature and will highlight existing open challenges and research directions that might inspire attendees to embark in the fascinating and promising area of computational human behavior modeling.

The goal of this tutorial is to provide an introduction to this burgeoning area, describing tools for automatically interpreting complex behavioral patterns generated when humans interact with machines or with others. A second goal is to inspire a new generation of researchers to join forces into realizing the immense potential of machine learning to help build intelligent systems that understand and interact with humans, and contribute to our understanding of human individual and aggregate behavior while always having human interests and wellbeing at their core.

Deep Learning with Bayesian Principles



Emtiyaz Khan (RIKEN)

Location: West Exhibition Hall A

Deep learning and Bayesian learning are considered two entirely different fields often used in complementary settings. It is clear that combining ideas from the two fields would be beneficial, but how can we achieve this given their fundamental differences?

This tutorial will introduce modern Bayesian principles to bridge this gap. Using these principles, we can derive a range of learning-algorithms as special cases, e.g., from classical algorithms, such as linear regression and forward-backward algorithms, to modern deep-learning algorithms, such as SGD, RMSprop and Adam. This view then enables new ways to improve aspects of deep learning, e.g., with uncertainty, robustness, and interpretation. It also enables the design of new methods to tackle challenging problems, such as those arising in active learning, continual learning, reinforcement learning, etc.

Overall, our goal is to bring Bayesians and deep-learners closer than ever before, and motivate them to work together to solve challenging real-world problems by combining their strengths.

Tutorials Session 2 - 11:15 am - 1:15 pm

Efficient Processing of Deep Neural Network: from Algorithms to Hardware Architectures

Vivienne Sze (MIT)



Location: West Exhibition Hall C + B3

This tutorial describes methods to enable efficient processing for deep neural networks (DNNs), which are used in many AI applications including computer vision, speech recognition, robotics, etc. While DNNs deliver best-in-class accuracy and quality of results, it comes at the cost of high computational complexity. Accordingly, designing efficient algorithms and hardware architectures for deep neural networks is an important step towards enabling the wide deployment of DNNs in AI systems (e.g., autonomous vehicles, drones, robots, smartphones, wearables, Internet of Things, etc.), which often have tight constraints in terms of speed, latency, power/energy consumption, and cost.

In this tutorial, we will provide a brief overview of DNNs, discuss the tradeoffs of the various hardware platforms that support DNNs including CPU, GPU, FPGA and ASICs, and highlight important benchmarking/comparison metrics and design considerations for evaluating the efficiency of DNNs. We will then describe recent techniques that reduce the computation cost of DNNs from both the hardware architecture and network algorithm perspective. Finally, we will also discuss how these techniques can be applied to a wide range of image processing and computer vision tasks.

Interpretable Comparison of Distributions and Models

Wittawat Jitkrittum (MPI),
Dougal J Sutherland (TTIC),
Arthur Gretton (UCL)



Location: West Exhibition Hall A

Modern machine learning has seen the development of models of increasing complexity for high-dimensional real-world data, such as documents and images. Some of these models are implicit, meaning they generate samples without specifying a probability distribution function (e.g. GANs), and some are explicit, specifying a distribution function – with a potentially quite complex structure which may not admit efficient sampling or normalization. This tutorial will provide modern nonparametric tools for evaluating and benchmarking both implicit and explicit models. For implicit models, samples from the model are compared with real-world samples; for explicit models, a Stein operator is defined to compare the model to data samples without requiring a normalized probability distribution. In both cases, we also consider relative tests to choose the best of several incorrect models. We will emphasize interpretable tests throughout, where the way in which the model differs from the data is conveyed to the user.

Machine Learning for Computational Biology and Health

Anna Goldenberg (SickKids,
U. of Toronto)
Barbara Engelhardt
(Princeton)



Location: West Ballroom A+B

Questions in biology and medicine pose big challenges to existing ML methods. The impact of creating ML methods to address these questions may positively impact all of us as patients, as scientists, and as human beings. In this tutorial, we will cover some of the major areas of current biomedical research, including genetics, the microbiome, clinical data, imaging, and drug design. We will focus on progress-to-date at the intersection of biology, health, and ML. We will also discuss challenges and open questions. We aim to leave you with thoughts on how to perform meaningful work in this area. It is assumed that participants have a good grasp of ML. Understanding of biology beyond high school level is not required.

Tutorials Session 3 - 2:45 - 4:45 pm



Reinforcement Learning: Past, Present, & Future Perspectives

Katja Hofmann (U. of Amsterdam) Location:
West Exhibition Hall C + B3

Reinforcement learning (RL) is a systematic approach to learning and decision making. Developed and studied for decades, recent combinations of RL with modern deep learning have led to impressive demonstrations of the capabilities of today's RL systems, and have fuelled an explosion of interest and research activity. Join this tutorial to learn about the foundations of RL - elegant ideas that give rise to agents that can learn extremely complex behaviors in a wide range of settings. Broadening out, I give a (subjective) overview of where we currently are in terms of what's possible. I conclude with an outlook on key opportunities - both for future research and for real-world applications of RL.



Representation Learning & Fairness

Moustapha Cisse (Google Brain)
Sanmi Koyejo (UIUC)

Location: West Exhibition Hall A

It is increasingly evident that widely-deployed machine learning models can lead to discriminatory outcomes and can exacerbate disparities in the training data. With the accelerating adoption of machine learning for real-world decision-making tasks, issues of bias and fairness in machine learning must be addressed. Our motivating thesis is that among a variety of emerging approaches, representation learning provides a unique toolset for evaluating and potentially mitigating unfairness. This tutorial presents existing research and proposes open problems at the intersection of representation learning and fairness. We will look at the (im)possibility of learning fair task-agnostic representations, connections between fairness and generalization performance, and the opportunity for leveraging tools from representation learning to implement algorithmic individual and group fairness, among others. The tutorial is designed to be accessible to a broad audience of machine learning practitioners, and the necessary background is a working knowledge of predictive machine learning.

Synthetic Control

Alberto Abadie (MIT)
Vishal Misra (Columbia U.)
Devavrat Shah (MIT)

Location: West Ballroom A+B



The synthetic control method, introduced in Abadie and Gardeazabal (2003), has emerged as a popular empirical methodology for estimating a causal effects with observational data, when the "gold standard" of a randomized control trial is not feasible. Starting from policy evaluation literature,

synthetic controls have found their way more broadly to social sciences, biological sciences, engineering and even sports. However, only recently, synthetic controls have been introduced to the machine learning community through its natural connection to matrix and tensor estimation in Amjad, Shah and Shen (2017) as well as Amjad, Misra, Shah and Shen (2019). In this tutorial, we will survey the rich body of literature on methodical aspects, mathematical foundations and empirical case studies of synthetic controls. We will also discuss how synthetic controls are likely to be instrumental in the next wave of development in reinforcement learning using observational data.

Invited Speaker - 5:45 - 6:30 pm



How to Know

Celeste Kidd (UC Berkeley)

West Exhibition Hall C + B3

This talk will discuss Kidd's research about how people come to know what they know. The world is a sea of information too vast for any one person to acquire entirely. How then do people navigate the information overload, and how do their decisions shape their knowledge and beliefs? In this talk, Kidd will discuss research from her lab about the core cognitive systems people use to guide their learning about the world—including attention, curiosity, and metacognition (thinking about thinking). The talk will discuss the evidence that people play an active role in their own learning, starting in infancy and continuing through adulthood. Kidd will explain why we are curious about some things but not others, and how our past experiences and existing knowledge shape our future interests. She will also discuss why people sometimes hold beliefs that are inconsistent with evidence available in the world, and how we might leverage our knowledge of human curiosity and learning to design systems that better support access to truth and reality.

Celeste Kidd is an Assistant Professor of Psychology at the University of California, Berkeley, where her lab investigates learning and belief formation. The Kidd Lab is one of few in the world that combine technologically sophisticated behavioral experiments with computational models in order to broadly understand knowledge acquisition. Her lab employs a range of methods, including eye-tracking and touchscreen testing with human infants, in order to show how learners sample information from their environment and build knowledge gradually over time. Her work has been published in PNAS, Neuron, Psychological Science, Developmental Science, and elsewhere. Her lab has received funding from NSF, DARPA, Google, the Jacobs Foundation, the Human Frontiers Science Program, and the Templeton Foundation. She is a recipient of the American Psychological Science Rising Star designation, the Glushko Dissertation Prize in Cognitive Science, and the Cognitive Science Society Computational Modeling Prize in Perception/Action. Kidd was also named as one of TIME Magazine's 2017 Persons of the Year as one of the "Silence Breakers" for her advocacy for better protections for students against sexual misconduct.

TUESDAY SCHEDULE

TIME	DESCRIPTION	LOCATION
8:30 - 9:20 AM	Invited Talk: Bin Yu <i>Three principles of data science and interpretable machine learning with case studies</i>	West Exhibition Hall C + B3
9:20 - 10:05 AM	Coffee break	
10:05 - 10:45 AM	Parallel Tracks:	
	Track 1	West Exhibition Hall C + B3
	Track 2	West Exhibition Hall A
	Track 3	West Ballrooms A + B
	Track 4	West Ballroom C
10:45 - 12:45 PM	Poster A Sessions	East Exhibition Hall B + C
12:45 - 2:15 PM	Lunch on your own	
2:15 - 3:05 PM	Invited Talk: Dana Pe'er <i>Machine learning meets single-cell biology: insights and challenges</i>	West Exhibition Hall C + B3
3:05 - 3:25 PM	Test Of Time Award: Dual Averaging Method for Regularized Stochastic Learning and Online Optimization <i>Lin Xiao</i>	West Exhibition Hall C + B3
3:25 - 4:10 PM	Coffee break	
4:10 - 5:30 PM	Parallel Tracks:	
	Track 1	West Exhibition Hall C + B3
	Track 2	West Exhibition Hall A
	Track 3	West Ballrooms A + B
	Track 4	West Ballroom C
5:30 - 7:30 PM	Poster B Sessions Demonstrations	East Exhibition Hall B + C East Exhibition Hall B + C
7:00 - 10:00 pm	NeurIPS Socials	West Level 2



Invited Speaker

8:30 - 9:20 am

Veridical Data Science

Bin Yu
UC Berkeley

Location: West Exhibition C

Data science is a field of evidence-seeking that combines data with domain information to generate new knowledge. It addresses key considerations in AI regarding when and where data-driven solutions are reliable and appropriate. Such considerations require involvement from humans who collectively understand the domain and tools used to collect, process, and model data. Throughout the data science life cycle, these humans make judgment calls to extract information from data. Veridical data science seeks to ensure that this information is reliable, reproducible, and clearly communicated so that empirical evidence may be evaluated in the context of human decisions. Three core principles: predictability, computability, and stability (PCS) provide the foundation for veridical data science. In this talk we will present a unified PCS framework for data analysis, consisting of both a workflow and documentation, illustrated through iterative random forests and case studies from genomics and precision medicine.

Bin Yu is Chancellor's Professor in the Departments of Statistics and of Electrical Engineering & Computer Sciences at the University of California at Berkeley and a former chair of Statistics at UC Berkeley. Her research focuses on practice, algorithm, and theory of statistical machine learning and causal inference. Her group is engaged in interdisciplinary research with scientists from genomics, neuroscience, and precision medicine.

She is a member of the U.S. National Academy of Sciences and Fellow of the American Academy of Arts and Sciences. She was a Guggenheim Fellow in 2006, and the Tukey Memorial Lecturer of the Bernoulli Society in 2012. She was President of IMS (Institute of Mathematical Statistics) in 2013-2014 and the Rietz Lecturer of IMS in 2016. She received the E. L. Scott Award from COPSS (Committee of Presidents of Statistical Societies) in 2018. Moreover, Yu was a founding co-director of the Microsoft Research Asia (MSR) Lab at Peking University and is a member of the scientific advisory board at the UK Alan Turing Institute for data science and AI.



Invited Speaker

2:15 - 3:05 pm

Machine learning meets single-cell biology: insights and challenges

Dana Pe'er
Sloan Kettering Institute

Location: West Exhibition C

Biology is becoming a data science. Recent single-cell profiling technologies are creating a data deluge, wherein thousands of variables are measured for each of hundreds of thousands to millions of cells in a single dataset. The proliferation of single-cell genomic and imaging data is creating opportunities to apply machine learning approaches in order to construct a human cell atlas with enormous potential to uncover new biology—by describing the incredible diversity of our constituent cell populations, how they function, how this diversity emerges from a single cell and how processes go awry in disease. We will present success stories and computational challenges raised by these new data modalities, in both health and disease settings. Examples will include methods from manifold learning, probabilistic graphical models and deep learning.

Dana Pe'er is Chair of Computational and Systems Biology program, Sloan Kettering Institute and Director of Alan and Sandra Gerry Center for Metastasis and Tumor Ecosystems. The Pe'er lab develops machine learning approaches for the analysis and interpretation of single cell data and uses these to study Cancer, Development and Immunology. Dana is member of Human Cell Atlas Organizing Committee and co-chair of its Analysis Working Group, recipient of the Burroughs Wellcome Fund Career Award, NIH Director's New Innovator Award, NSF CAREER award, Stand Up To Cancer Innovative Research Grant, Packard Fellow in Science and Engineering, Overton award, NIH Director's Pioneer award, Lenfest Distinguished Faculty Award and Ernst W. Bertner Memorial Award



Poster Sessions A

10:45 - 12:45 pm - East Exhibition Hall B + C

--- Algorithms ---

- | | |
|----------------------------------|-----------------|
| • Adversarial Learning | Posters 1-16 |
| • Bandit Algorithms | Posters 17 - 27 |
| • Clustering | Posters 28 - 42 |
| • Components Analysis | Posters 43 - 47 |
| • Density Estimation | Posters 48 - 54 |
| • Dynamical Systems | Posters 55 - 56 |
| • Kernel Methods | Posters 57 - 62 |
| • Missing Data | Poster 63 |
| • Representation Learning | Posters 64 - 73 |
| • Similarity & Distance Learning | Posters 74 - 80 |

--- Applications ---

- | | |
|---|-----------------|
| • Communication or
Memory-Bounded Learning | Posters 81 - 82 |
| • Dialog or Communication-
Based Learning | Poster 83 |
| • Game Playing | Poster 84 |
| • Privacy, Anonymity & Security | Posters 85 - 96 |
| • Recommender Systems | Poster 97 - 99 |
| • Web Applications and Internet Data | Poster 100 |

--- Deep Learning ---

- | | |
|---|-------------------|
| • Biologically Plausible Deep Networks | Posters 101 - 106 |
| • Deep Autoencoders | Posters 107 - 109 |
| • Efficient Inference Methods | Posters 110 - 117 |
| • Generative Models | Posters 118 - 133 |
| • Interaction-Based Deep Networks | Poster 134 |
| • Optimization for Deep Networks | Posters 135 - 145 |
| • Predictive Models | Posters 146 - 149 |
| • Recurrent Networks | Posters 150 - 158 |
| • Visualization or Exposition
Techniques for Deep Networks | Posters 159 - 173 |

--- Optimization ---

- | | |
|------------------------------|-------------------|
| • Combinatorial Optimization | Posters 174 - 178 |
|------------------------------|-------------------|

--- Probabilistic Methods ---

- | | |
|--------------------|-------------------|
| • Causal Inference | Posters 179 - 187 |
|--------------------|-------------------|

--- Reinforcement Learning & Planning ---

- | | |
|-----------------------------|-------------------|
| • Decision and Control | Posters 188 - 191 |
| • Exploration | Posters 192 - 198 |
| • Markov Decision Processes | Posters 199 - 207 |
| • Navigation | Posters 208 - 209 |

--- Theory ---

- | | |
|---|-------------------|
| • Computational Complexity | Posters 210 - 211 |
| • Frequentist Statistics | Posters 212 - 213 |
| • Hardness of Learning & Approximations | Posters 214 - 218 |
| • Learning Theory | Posters 219 - 229 |

Poster Sessions B

5:30 - 7:30 pm - East Exhibition Hall B + C

--- Algorithms ---

- | | |
|---|-----------------|
| • AutoML | Posters 1 - 10 |
| • Bandit Algorithms | Posters 11 - 23 |
| • Large Scale Learning | Posters 24 - 40 |
| • Meta-Learning | Posters 41 - 49 |
| • Regression | Posters 50 - 53 |
| • Structured Prediction | Posters 54 - 58 |
| • Unsupervised Learning | Posters 59 - 64 |
| • Body Pose, Face, and Gesture Analysis | Posters 65 - 70 |

--- Deep Learning ---

- | | |
|----------------------------------|-----------------|
| • Generative Models | Posters 71 - 87 |
| • Optimization for Deep Networks | Posters 88 - 99 |

--- Optimization ---

- | | |
|---------------------------|-------------------|
| • Convex Optimization | Posters 100 - 115 |
| • Non-Convex Optimization | Posters 116 - 127 |
| • Stochastic Optimization | Posters 128 - 133 |

--- Probabilistic Methods ---

- | | |
|-------------------------|-------------------|
| • Causal Inference | Posters 134 - 143 |
| • Distributed Inference | Posters 144 - 145 |
| • Gaussian Processes | Posters 146 - 152 |
| • Hierarchical Models | Posters 153 - 154 |
| • MCMC | Posters 155 - 163 |
| • Variational Inference | Posters 164 - 173 |

--- Reinforcement Learning & Planning ---

- | | |
|--------------------------|-------------------|
| • Decision and Control | Posters 174 - 177 |
| • Exploration | Posters 178 - 184 |
| • Model-Based RL | Posters 185 - 193 |
| • Multi-Agent RL | Posters 194 - 200 |
| • Reinforcement Learning | Posters 201 - 215 |

--- Theory ---

- | | |
|-------------------|-------------------|
| • Control Theory | Poster 216 |
| • Learning Theory | Posters 217 - 228 |

This year the poster sessions have been ordered by classification and topic. We hope this will make it easier to view the posters that are relevant to you.

Demonstrations - 5:30 - 7:30 pm - East Exhibition Hall B + C

801 **exBERT: A Visual Analysis Tool to Explain BERT's Learned Representations**

Ben Hoover · Hendrik Strobelt · Sebastian Gehrmann

802 **Streamlit, a new app framework for machine learning tools**

Adrien Treuille · Amanda Kelly

803 **Discovering Neural Wirings Neural Network Visualizer**

Alvaro Herrasti · Mitchell Wortsman

804 **"How Can This Paper Get In?" - A Game To Advise researchers when writing for a top AI conference**

Aabhas Sharma · Narendra Nath Joshi · Michael Muller · Casey Dugan

805 **Robot-Assisted Hair-Brushing**

Eura Shin · Hejia Zhang · Rey J Pocius · Nathaniel Dennler · Heather Culbertson · Naghmeh Zamani · Stefanos Nikolaidis

806 **Learning Machines can Curl - Adaptive Deep Reinforcement Learning enables the robot Curly to win against human players in an icy world**

Dong-Ok Won · Sang-Hoon Lee · Klaus-Robert Müller · Seong-Whan Lee

807 **Human Gesture Recognition using Spiking Input on Akida Neuromorphic Platform**

Sounak Dey · Arijit Mukherjee · Gilles BEZARD · Douglas McLelland

808 **GENO -- Optimization for Classical Machine Learning Made Fast and Easy**

Soeren Laue · Matthias Mitterreiter · Joachim Giesen

805 **Toronto Annotation Suite**

Amlan Kar · Sanja Fidler · Jun Gao · Seung Wook Kim · huan ling

810 **SCC: Deep Reinforcement Learning Agent plays StarCraft II at Human Master Level**

XJ Wang · Peng Peng

811 **AI in Two-sided Ride-sharing Marketplace**

Zhiwei Qin · Shikai Luo · lingyu zhang · yan jiao · Xiaocheng Tang · Lulu Zhang · hongtu zhu · Jieping Ye

812 **NNgen: A Model-Specific Hardware Synthesis Compiler for Deep Neural Network**


Shinya Takamaeda-Yamazaki · Shinya Fujisawa · Shuichi Fujisaki

813 **Realtime Modeling and Anomaly Detection in Multivariate Data Streams**

Christopher Hannon · Andrey Lokhov · Deep Deka

814 **Empathy based Affective Portrait Painter**

Steve DiPaola · Ozge Nilay YALCIN · Nouf Abukhodair



WEDNESDAY SCHEDULE

TIME	DESCRIPTION	LOCATION
8:30 - 9:20 AM	Invited Talk: Blaise Aguera y Arcas <i>Social Intelligence</i>	West Exhibition C + B3
9:20 - 10:05 AM	Coffee break	
10:05 - 10:45 AM	Parallel Tracks:	
	Track 1	West Exhibition Hall C + B3
	Track 2	West Exhibition Hall A
	Track 3	West Ballrooms A + B
	Track 4	West Ballroom C
10:45 - 12:45 PM	Poster A Sessions	East Exhibition Hall B + C
12:45 - 2:15 PM	Lunch on your own	
2:15 - 3:05 PM	Invited Talk: Yoshua Bengio <i>From System 1 Deep Learning to System 2 Deep Learning</i>	West Exhibition C + B3
3:05 - 3:50 PM	Coffee break	
3:50 - 5:00 PM	Parallel Tracks:	
	Track 1	West Exhibition Hall C + B3
	Track 2	West Exhibition Hall A
	Track 3	West Ballrooms A + B
	Track 4	West Ballroom C
5:00 - 7:00 PM	Poster B Sessions Demonstrations	East Exhibition Hall B + C East Exhibition Hall B + C
7:00 - 10:00 pm	NeurIPS Socials	West Level 2



Invited Speaker

8:30 - 9:20 am

Social Intelligence

Blaise Aguera y Arcas
Google

Location: West Exhibition C



In the past decade, we've figured out how to build artificial neural nets that can achieve superhuman performance at almost any task for which we can define a loss function and gather or create a sufficiently large dataset. While this is unlocking a wealth of valuable applications, it also raises questions: how can we make fair and ethical models? How can we have privacy in a world where our data are the fuel for training all of these models? Does AI at scale increase or curtail human agency? And are intelligences really just function approximators?

This talk will be technically grounded, but will also address these big questions and some non-obvious interconnections between them. We will begin with privacy and federated computation, then delve deeper into the limitations of the optimization framework for ML, exploring alternative approaches involving meta-learning, evolution strategies, populations, sociality, and cultural accumulation. We hypothesize that this relatively underexplored approach to general intelligence may be both fruitful in the near term and more optimistic in its long-term outlook.

Blaise leads an organization at Google AI working on both basic research and new products. Among the team's public contributions are MobileNets, Federated Learning, Coral, and many Android and Pixel AI features. They also founded the Artists and Machine Intelligence program, and collaborate extensively with academic researchers in a variety of fields. Until 2014 Blaise was a Distinguished Engineer at Microsoft, where he worked in a variety of roles, from inventor to strategist, and led teams with strengths in machine learning, interaction design, prototyping, augmented reality, wearable computing, and graphics. Blaise has given TED talks on Seadragon and Photosynth (2007, 2012), Bing Maps (2010), and machine creativity (2016). In 2008, he was awarded MIT's TR35 prize.


Invited Speaker

2:15 - 3:05 pm

From System 1 Deep Learning to System 2 Deep Learning

Yoshua Bengio
U. of Montreal

Location: West Exhibition C



Past progress in deep learning has concentrated mostly on learning from a static dataset, mostly for perception tasks and other System 1 tasks which are done intuitively and unconsciously by humans. However, in recent years, new tools such as meta-learning, soft-attention and progress in deep reinforcement learning are opening the door to novel deep architectures and training frameworks for addressing System 2 tasks (which are done consciously), such as reasoning, planning, capturing causality and obtaining systematic generalization in natural language processing. This expansion from System 1 to System 2 tasks is important to achieve the old deep learning goal of disentangling high-level abstract representations: we argue that meta-learning, the modularization aspect of the consciousness prior and an agent perspective on representation learning should put pressure on top-level representation learning to capture semantic concepts and facilitate re-use of learned components in novel ways, enabling more powerful forms of compositional and out-of-distribution generalization.

Yoshua Bengio is Full Professor in the computer science and operations research department at U. Montreal, scientific director and founder of Mila and of IVADO, Turing Award 2018 recipient, Canada Research Chair in Statistical Learning Algorithms, as well as a Canada AI CIFAR Chair. He pioneered deep learning and has been getting the most citations per day in 2018 among all computer scientists, worldwide. He is an officer of the Order of Canada, member of the Royal Society of Canada, was awarded the Killam Prize, the Marie-Victorin Prize and the Radio-Canada Scientist of the year in 2017, and he is a member of the NeurIPS advisory board and co-founder of the ICLR conference, as well as program director of the CIFAR program on Learning in Machines and Brains. His goal is to contribute to uncover the principles giving rise to intelligence through learning, as well as favour the development of AI for the benefit of all.

Poster Sessions A

10:45 - 12:45 pm - East Exhibition Hall B + C

--- Algorithms ---

- | | |
|--|-----------------|
| • Adaptive Data Analysis | Posters 1 - 4 |
| • Boosting and Ensemble Methods | Posters 5 - 12 |
| • Model Selection & Structure Learning | Posters 13 - 19 |
| • Regression | Posters 20 - 23 |
| • Semi-Supervised Learning | Posters 24 - 34 |
| • Uncertainty Estimation | Posters 34 - 54 |
| • Unsupervised Learning | Posters 55 - 61 |

--- Applications ---

- | | |
|--|-----------------|
| • Computer Vision | Posters 62 - 75 |
| • Image Segmentation | Posters 76 - 84 |
| • Object Detection | Posters 85 - 89 |
| • Robotics | Posters 90 - 94 |
| • Visual Scene Analysis & Interpretation | Posters 95 - 98 |

--- Deep Learning ---

- | | |
|-------------------------------|-------------------|
| • Efficient Inference Methods | Posters 99 - 107 |
| • Generative Models | Posters 108 - 124 |
| • Supervised Deep Networks | Posters 125 - 131 |

--- Neuroscience and Cognitive Science ---

- | | |
|---|-------------------|
| • Brain Imaging | Posters 132 - 133 |
| • Brain Mapping | Poster 134 |
| • Brain-Computer Interfaces and Neural Prostheses | Posters 135 - 138 |
| • Connectomics | Poster 139 |
| • Human or Animal Learning | Posters 140 - 142 |
| • Language for Cognitive Science | Poster 143 |
| • Memory | Poster 144 |
| • Neuroscience | Posters 145 - 149 |
| • Perception | Poster 150 |
| • Problem Solving | Poster 151 |
| • Visual Perception | Posters 152 - 156 |

--- Optimization ---

- | | |
|---------------------------|------------------|
| • Stochastic Optimization | Poster 157 - 163 |
| • Submodular Optimization | Poster 164 - 170 |

--- Probabilistic Methods ---

- | | |
|---------------------------|------------------|
| • Bayesian Nonparametrics | Poster 171 - 173 |
| • Belief Propagation | Poster 174 - 175 |
| • Graphical Models | Poster 176 - 184 |
| • Latent Variable Models | Poster 185 - 187 |
| • Topic Models | Poster 188 - 190 |

--- Reinforcement Learning & Planning ---

- | | |
|--------------------------|-------------------|
| • Decision and Control | Posters 191 - 194 |
| • Hierarchical RL | Posters 195 - 199 |
| • Reinforcement Learning | Posters 200 - 214 |

--- Theory ---

- | | |
|---|-------------------|
| • Game Theory & Computational Economics | Posters 215 - 220 |
| • Learning Theory | Posters 221 - 231 |
| • Regularization | Posters 232 - 236 |

Poster Sessions B

5:00 - 7:00 pm - East Exhibition Hall B + C

--- Algorithms ---

- | | |
|-----------------------------------|-----------------|
| • Bandit Algorithms | Posters 1 - 12 |
| • Classification | Posters 13 - 15 |
| • Collaborative Filtering | Posters 16 - 17 |
| • Few-Shot Learning | Posters 18 - 25 |
| • Meta-Learning | Posters 26 - 35 |
| • Metric Learning | Posters 36 - 38 |
| • Multitask and Transfer Learning | Posters 39 - 47 |
| • Online Learning | Posters 48 - 56 |
| • Ranking and Preference Learning | Posters 57 - 59 |
| • Relational Learning | Posters 60 - 70 |
| • Spectral Methods | Posters 71 - 73 |

--- Applications ---

- | | |
|-----------------------------------|-------------------|
| • Audio and Speech Processing | Posters 74 - 77 |
| • Computer Vision | Posters 78 - 91 |
| • Information Retrieval | Posters 92 - 94 |
| • Matrix and Tensor Factorization | Posters 95 - 99 |
| • Natural Language Processing | Posters 100 - 109 |
| • Signal Processing | Posters 110 - 112 |
| • Visual Question Answering | Posters 113 - 120 |

Data, Challenges, Implementations & Software

- | | |
|----------------------------------|-------------------|
| • Benchmarks | Posters 121 - 122 |
| • Data Sets or Data Repositories | Posters 123 - 124 |
| • Virtual Environments | Poster 125 |

--- Deep Learning ---

- | | |
|---------------------|-------------------|
| • Attention Models | Posters 126 - 134 |
| • Generative Models | Posters 135 - 151 |

--- Optimization ---

- | | |
|-----------------------|-------------------|
| • Convex Optimization | Posters 152 - 167 |
|-----------------------|-------------------|

--- Probabilistic Methods ---

- | | |
|-------------------------|-------------------|
| • Gaussian Processes | Posters 168 - 174 |
| • Hierarchical Models | Poster 175 |
| • MCMC | Posters 176 - 183 |
| • Variational Inference | Posters 184 - 194 |

--- Reinforcement Learning and Planning ---

- | | |
|--------------------------|-------------------|
| • Multi-Agent RL | Posters 195 - 201 |
| • Reinforcement Learning | Posters 202 - 216 |

--- Theory ---

- | | |
|--|-------------------|
| • Game Theory & Computational Economics | Posters 217 - 222 |
| • Large Deviations & Asymptotic Analysis | Posters 223 - 224 |
| • Learning Theory | Posters 225 - 235 |
| • Statistical Physics of Learning | Posters 236 - 242 |

Demonstrations - 5:00 - 7:00 pm - East Exhibition Hall B + C

- 801 **Melody Slot Machine**
Masatoshi Hamanaka
- 802 **Smart Home Appliances: Chat with your Fridge**
Denis Gudovskiy · Alec Hodgkinson · Stefano Alletto · Luca Rigazio
- 803 **Shared Mobile-Cloud Inference for Collaborative Intelligence**
Mateen Ulhaq · Ivan Bajić
- 804 **Project BB: Bringing AI to the Command Line**
Tathagata Chakraborti · Mayank Agarwal
- 805 **Passcode: A cooperative word guessing game between a human and AI agent**
Katy Gero · Maria Ruis · Zahra Ashktorab · J Johnson · Sadhana Kumaravel · Wei Zhang · Qian Pan · Murray Campbell · Casey Dugan · David Millen · Sarah Miller · Werner Geyer
- 806 **Immersion - How Does Music Sound to Artificial Ears?**
Vincent Herrmann
- 807 **The Option Keyboard: Combining Skills in Reinforcement Learning**
Daniel Toyama · Shaobo Hou · Gheorghe Comanici · Andre Barreto · Doina Precup · Shihab Mourad · Eser Aygün · Philippe Hamel
- 808 **AllenNLP Interpret: Explaining Predictions of NLP Models**
Jens Tuyls · Eric Wallace · Matt Gardner · Junlin Wang · Sameer Singh · Sanjay Subramanian
- 809 **Deep Space-Time Prior for Realtime Mobile Novel View Synthesis**
Zainul Shah
- 810 **AIDeme: An active learning based system for interactive exploration of large datasets**
Enhui Huang · Luciano Di Palma · Laurent Cetinsoy · Yanlei Diao · Anna Liu
- 811 **BIM-GAN: a sketch to layout, 3D, and VR tool for architectural floor plan design**
Chin-Yi Cheng
- 812 **One-on-one fitness training with an AI avatar**
Roland Memisevic · Guillaume Berger · Tippi Puar · David Greenberg
- 813 **Real Time CFD simulations with 3D Mesh Convolutional Networks**
Pierre Baque · Pascal Fua · François Fleuret
- 814 **F1/10: An open-source 1/10th scale platform for autonomous racing and reinforcement learning**
Matthew O'Kelly · Dhruv Karthik · Hongrui Zheng · Joseph Auckley · Siddharth Singh · Shashank D Prasad · Kim Luong · Matthew R Lebermann · Rahul Mangharam

See page 15 for Demo Location Maps

Outstanding Paper Award

- **Distribution-Independent PAC Learning of Halfspaces with Massart Noise**
Ilias Diakonikolas · Themis Gouleakis · Christos Tzamos

Honorable Mentions:

- **Non-parametric Density Estimation Convergence Rates for GANs under Besov IPM Losses**
Ananya Uppal · Shashank Singh · Barnabas Poczos
- **Fast and Accurate Least-Mean-Squares Solvers**
Ibrahim Jubran · Alaa Maalouf · Dan Feldman

Outstanding New Directions Paper Award

- **Uniform Convergence May Be Unable To Explain Generalization in Deep Learning**
Vaishnavh Nagarajan · J. Zico Kolter

Honorable Mentions:

- **Putting An End to End-to-End: Gradient-Isolated Learning of Representations**
Sindy Löwe · Peter O'Connor · Bastiaan Veeling
- **Scene Representation Networks: Continuous 3D-Structure-Aware Neural Scene Representations**
Vincent Sitzmann · Michael Zollhoefer · Gordon Wetzstein

THURSDAY SCHEDULE

TIME	DESCRIPTION	LOCATION
8:30 - 9:20 AM	Invited Talk: Kafui Dzirasa <i>Mapping emotions: Discovering structure in mesoscale electrical brain recordings</i>	West Exhibition C + B3
9:20 - 10:05 AM	Coffee break	
10:05 - 10:45 AM	Parallel Tracks:	
	Track 1	West Exhibition Hall C + B3
	Track 2	West Exhibition Hall A
	Track 3	West Ballrooms A + B
	Track 4	West Ballroom C
10:45 - 12:45 PM	Poster A Sessions	East Exhibition Hall B + C
12:45 - 2:15 PM	Lunch on your own	
2:15 - 3:05 PM	Invited Talk: Jeff Heer <i>Agency + Automation: Designing Artificial Intelligence into Interactive Systems</i>	West Exhibition C + B3
3:05 - 3:50 PM	Coffee break	
3:50 - 5:00 PM	Parallel Tracks:	
	Track 1	West Exhibition Hall C + B3
	Track 2	West Exhibition Hall A
	Track 3	West Ballrooms A + B
	Track 4	West Ballroom C
5:00 - 7:00 PM	Poster B Sessions	East Exhibition Hall B + C
7:00 - 10:00 pm	NeurIPS Socials	West Level 2



Invited Speaker

8:30 - 9:20 am

Mapping emotions: Discovering structure in mesoscale electrical brain recordings

Kafui Dzirasa
Duke University

Location: West Exhibition C



Brain-wide fluctuations in local field potential oscillations reflect emergent network-level signals that mediate behavior. Cracking the code whereby these oscillations coordinate in time and space (spatiotemporal dynamics) to represent complex behaviors would provide fundamental insights into how the brain signals emotional pathology. Using machine learning, we discover a spatiotemporal dynamic network that predicts the emergence of major depressive disorder (MDD)-related behavioral dysfunction in mice subjected to chronic social defeat stress. Activity patterns in this network originate in prefrontal cortex and ventral striatum, relay through amygdala and ventral tegmental area, and converge in ventral hippocampus. This network is increased by acute threat, and it is also enhanced in three independent models of MDD vulnerability. Finally, we demonstrate that this vulnerability network is biologically distinct from the networks that encode dysfunction after stress. Thus, these findings reveal a convergent mechanism through which MDD vulnerability is mediated in the brain.

Kafui Dzirasa completed a PhD in Neurobiology at Duke University. His research interests focus on understanding how changes in the brain produce neurological and mental illness, and his graduate work has led to several distinctions including: the Somjen Award for Most Outstanding Dissertation Thesis, the Ruth K. Broad Biomedical Research Fellowship, the UNCF-Merck Graduate Science Research Fellowship, and the Wakeman Fellowship. Kafui obtained an MD from the Duke University School of Medicine in 2009, and he completed residency training in General Psychiatry in 2016.

Kafui received the Charles Johnson Leadership Award in 2007, and he was recognized as one of Ebony magazine's 30 Young Leaders of the Future in February 2008. He has also been awarded the International Mental Health Research Organization Rising Star Award, the Sydney Baer Prize for Schizophrenia Research, and his laboratory was featured on CBS 60 Minutes in 2011. In 2016, he was awarded the inaugural Duke Medical Alumni Emerging Leader Award and the Presidential Early Career Award for Scientists and Engineers: The Nation's highest award for scientists and engineers in the early stages of their independent research careers. In 2017, he was recognized as 40 under 40 in Health by the National Minority Quality Forum, and the Engineering Alumni of the Year from UMBC. He was inducted into the American Society for Clinical Investigation in 2019.

Invited Speaker

2:15 - 3:05 pm

Agency + Automation: Designing Artificial Intelligence into Interactive Systems

Jeff Heer
University of Washington

Location: West Exhibition C



Much contemporary rhetoric regards the prospects and pitfalls of using artificial intelligence techniques to automate an increasing range of tasks, especially those once considered the purview of people alone. These accounts are often wildly optimistic, understating outstanding challenges while turning a blind eye to the human labor that undergirds and sustains ostensibly “automated” services. This long-standing focus on purely automated methods unnecessarily cedes a promising design space: one in which computational assistance augments and enriches, rather than replaces, people’s intellectual work. This tension between agency and automation poses vital challenges for design, engineering, and society at large. In this talk we will consider the design of interactive systems that enable adaptive collaboration among people and computational agents. We seek to balance the often complementary strengths and weaknesses of each, while promoting human control and skillful action. We will review case studies in three arenas—data wrangling, exploratory visualization, and natural language translation—that integrate proactive computational support into interactive systems. To improve outcomes and support learning by both people and machines, I will describe the use of shared representations of tasks augmented with predictive models of human capabilities and actions.

Jeffrey Heer is the Jerre D. Noe Endowed Professor of Computer Science & Engineering at the University of Washington, where he directs the Interactive Data Lab and conducts research on data visualization, human-computer interaction, and social computing. The visualization tools developed by Jeff and his collaborators (Vega, D3.js, Protovis, Prefuse) are used by researchers, companies, and thousands of data enthusiasts around the world. Jeff's research papers have received awards at the premier venues in Human-Computer Interaction and Visualization (ACM CHI, ACM UIST, IEEE InfoVis, IEEE VAST, EuroVis). Other honors include MIT Technology Review's TR35 (2009), a Sloan Fellowship (2012), the ACM Grace Murray Hopper Award (2016), and the IEEE Visualization Technical Achievement Award (2017). Jeff holds B.S., M.S., and Ph.D. degrees in Computer Science from UC Berkeley, whom he then “betrayed” to join the Stanford faculty (2009–2013). He is also a co-founder of Trifacta, a provider of interactive tools for scalable data transformation.

Poster Sessions A

10:45 - 12:45 pm - East Exhibition Hall B + C

--- Algorithms ---

- | | |
|-----------------------------------|-----------------|
| • Active Learning | Posters 1 - 9 |
| • Adversarial Learning | Posters 10 - 26 |
| • AutoML | Posters 27 - 36 |
| • Bandit Algorithms | Posters 37 - 47 |
| • Classification | Posters 48 - 51 |
| • Multitask and Transfer Learning | Posters 52 - 61 |
| • Representation Learning | Posters 62 - 71 |

--- Applications ---

- | | |
|---|-------------------|
| • Fairness, Accountability & Transparency | Posters 72 - 86 |
| • Privacy, Anonymity & Security | Posters 87 - 100 |
| • Quantitative Finance & Econometrics | Poster 101 |
| • Time Series Analysis | Posters 102 - 114 |

--- Deep Learning ---

- | | |
|----------------------------------|-------------------|
| • Adversarial Networks | Posters 115 - 128 |
| • CNN Architectures | Posters 129 - 151 |
| • Efficient Training Methods | Posters 152 - 164 |
| • Optimization for Deep Networks | Posters 165 - 175 |

--- Neuroscience and Cognitive Science ---

- | | |
|---------------------|-------------------|
| • Cognitive Science | Posters 176 - 179 |
| • Neural Coding | Posters 180 - 187 |
| • Neuroscience | Posters 188 - 191 |
| • Reasoning | Posters 192 - 193 |

--- Optimization ---

- | | |
|---------------------------|-------------------|
| • Non-Convex Optimization | Posters 194 - 205 |
| • Stochastic Optimization | Posters 206 - 212 |

--- Reinforcement Learning and Planning ---

- | | |
|--------------------------|-------------------|
| • Planning | Posters 213 - 220 |
| • Reinforcement Learning | Posters 221 - 235 |

--- Theory ---

- | | |
|----------------------|-------------------|
| • Information Theory | Posters 236 - 243 |
| • Regularization | Posters 244 - 248 |

Poster Sessions B

5:00 - 7:00 pm - East Exhibition Hall B + C

--- Algorithms ---

- | | |
|---|-----------------|
| • Components Analysis
(e.g., CCA, ICA, LDA, PCA) | Posters 1 - 5 |
| • Kernel Methods | Posters 6 - 12 |
| • Nonlinear Dimensionality Reduction
and Manifold Learning | Posters 13 - 18 |
| • Online Learning | Posters 19 - 27 |
| • Representation Learning | Posters 28 - 38 |
| • Sparse Coding & Dimensionality
Expansion | Poster 39 |
| • Sparsity & Compressed Sensing | Posters 40 - 56 |
| • Stochastic Methods | Posters 57 - 59 |
| • Structured Prediction | Posters 60 - 64 |
| • Unsupervised Learning | Posters 65 - 71 |

--- Applications ---

- | | |
|---|-------------------|
| • Activity and Event Recognition | Posters 72 - 75 |
| • Computational Biology & Bioinformatics | Posters 76 - 81 |
| • Computational Photography | Posters 82 - 84 |
| • Computational Social Science | Posters 86 - 86 |
| • Computer Vision | Posters 87 - 101 |
| • Denoising | Posters 102 - 103 |
| • Fairness, Accountability & Transparency | Posters 104 - 119 |
| • Hardware and Systems | Posters 120 - 126 |
| • Health | Posters 127 - 130 |
| • Natural Language Processing | Posters 131 - 141 |
| • Network Analysis | Posters 142 - 148 |
| • Object Recognition | Posters 149 - 151 |
| • Privacy, Anonymity & Security | Posters 152 - 163 |
| • Program Understanding & Generation | Posters 164 - 170 |
| • Sustainability | Poster 171 |
| • Tracking and Motion in Video | Posters 172 - 174 |
| • Video Analysis | Posters 175 - 176 |

--- Data, Challenges, Implementations & Software ---

- | | |
|----------------------------------|-------------------|
| • Data Sets or Data Repositories | Posters 177 - 179 |
| • Software Toolkits | Posters 180 - 182 |

--- Deep Learning ---

- | | |
|------------------------------------|-------------------|
| • Embedding Approaches | Posters 183 - 190 |
| • Memory-Augmented Neural Networks | Posters 191 - 194 |
| • Optimization for Deep Networks | Posters 195 - 206 |
| • Predictive Models | Poster 246 |

--- Optimization ---

- | | |
|---------------------------|-------------------|
| • Non-Convex Optimization | Posters 207 - 217 |
|---------------------------|-------------------|

--- Theory ---

- | | |
|--|-------------------|
| • Frequentist Statistics | Posters 218 - 220 |
| • Game Theory & Computational Economics | Posters 221 - 225 |
| • Large Deviations & Asymptotic Analysis | Posters 226 - 227 |
| • Learning Theory | Posters 228 - 239 |
| • Spaces of Functions and Kernels | Posters 240 - 245 |

Friday Workshops - 8 am - 6:40 pm

- Information Theory and Machine Learning** **E Hall A**
 Shengjia Zhao · Jiaming Song · Yanjun Han · Kristy Choi · Pratyusha Kalluri · Ben Poole · Alexandros Dimakis · Jiantao Jiao · Tsachy Weissman · Stefano Ermon
- Beyond First Order Methods in Machine Learning Systems** **W 211-214**
 Anastasios Kyrillidis · Albert Berahas · Fred Roosta · Michael W Mahoney
- AI for Humanitarian Assistance and Disaster Response** **W 217-219**
 Ritwik Gupta · Robin Murphy · Trevor Darrell · Eric Heim · Zhangyang Wang · Bryce Goodman · Piotr Biliński
- KR2ML - Knowledge Representation and Reasoning Meets Machine Learning** **W 109-110**
 Veronika Thost · Christian Muise · Kartik Talamadupula · Sameer Singh · Christopher Ré
- EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition)** **W 306**
 Raj Parihar · Raj Parihar · Michael Goldfarb · Michael Goldfarb · Satyam Srivastava · Tao Sheng
- MLSys: Workshop on Systems for ML** **E MR 11-12**
 Aparna Lakshmiratan · Siddhartha Sen · Joseph Gonzalez · Dan Crankshaw · Sarah Bird
- Optimal Transport for Machine Learning** **E Brm C**
 Marco Cuturi · Gabriel Peyré · Rémi Flamary · Alexandra Suvorikova
- Workshop on Federated Learning for Data Privacy and Confidentiality** **W 118-120**
 Lixin Fan · Jakub Konečný · Yang Liu · Brendan McMahan · Virginia Smith · Han Yu
- Bayesian Deep Learning** **W Ex. Hall C**
 Yarin Gal · José Miguel Hernández-Lobato · Christos Louizos · Eric Nalisnick · Zoubin Ghahramani · Kevin Murphy · Max Welling
- Learning Meaningful Representations of Life** **E Brm B**
 Elizabeth Wood · Yakir Reshef · Jonathan Bloom · Jasper Snoek · Barbara Engelhardt · Scott Linderman · Suchi Saria · Alexander Wiltchko · Casey Greene · Chang Liu · Kresten Lindorff-Larsen · Debora Marks
- Retrospectives: A Venue for Self-Reflection in ML Research** **W 114-115**
 Ryan Lowe · Yoshua Bengio · Joelle Pineau · Michela Paganini · Jessica Forde · Shagun Sodhani · Abhishek Gupta · Joel Lehman · Peter Henderson · Kanika Madan
- Visually Grounded Interaction and Language** **W 202-204**
 Florian Strub · Abhishek Das · Erik Wijmans · Harm de Vries · Stefan Lee · Alane Suhr · Dor Arad Hudson
- Machine Learning for the Developing World (ML4D): Challenges and Risks** **W 121-122**
 Maria De-Arteaga · Amanda Coston · Tejumade Afonja
- Minding the Gap: Between Fairness & Ethics** **E MR 8+15**
 Igor Rubinov · Risi Kondor · Jack Poulson · Manfred K. Warmuth · Emanuel Moss · Alexa Hagerty
- Graph Representation Learning** **E Ex Hall A**
 Will Hamilton · Rianne van den Berg · Michael Bronstein · Stefanie Jegelka · Thomas Kipf · Jure Leskovec · Renjie Liao · Yizhou Sun · Petar Veličković
- Solving Inverse Problems with Deep Networks: New Architectures, Theoretical Foundations, and Applications** **W 301-305**
 Reinhard Heckel · Paul Hand · Richard Baraniuk · Joan Bruna · Alexandros Dimakis · Deanna Needell
- Biological and Artificial Reinforcement Learning** **W Brm C**
 Raymond Chua · Sara Zannone · Feryal Behbahani · Rui Ponte Costa · Claudia Clopath · Blake Richards · Doina Precup
- Robust AI in Financial Services: Data, Fairness, Explainability, Trustworthiness, and Privacy** **W 205-207**
 Alina Oprea · Avigdor Gal · Eren Kurshan · Isabelle Moulinier · Jiahao Chen · Manuela Veloso · Senthil Kumar · Tanveer Faruque
- CiML 2019: Machine Learning Competitions for All** **W 215-216**
 Adrienne Mendrik · Wei-Wei Tu · Wei-Wei Tu · Isabelle Guyon · Evelyne Viegas · Ming Li
- Meta-Learning** **W Brm B**
 Roberto Calandra · Ignasi Clavera Gilaberte · Frank Hutter · Joaquin Vanschoren · Jane Wang
- Learning with Rich Experience: Integration of Learning Paradigms** **W 208-209**
 Zhiting Hu · Andrew Wilson · Chelsea Finn · Lisa Lee · Taylor Berg-Kirkpatrick · Ruslan Salakhutdinov · Eric Xing
- Shared Visual Representations in Human and Machine Intelligence** **W 220-222**
 Arturo Deza · Joshua Peterson · Apurva Ratan Murty · Tom Griffiths
- Safety & Robustness in Decision-making** **E Brm A**
 Mohammad Ghavamzadeh · Shie Mannor · Yisong Yue · Marek Petrik · Yinlam Chow
- Machine Learning for Health (ML4H): What makes machine learning in medicine different?** **W Brm A**
 Andrew Beam · Tristan Naumann · Brett Beaulieu-Jones · Madalina Fiterau · Irene Y Chen · Samuel Finlayson · Emily Alsentzer · Adrian Dalca · Matthew McDermott
- Competition Track Day 1** **W 116-117**
 Hugo Jair Escalante
- Workshop on Human-Centric Machine Learning** **W 223-224**
 Plamen P Angelov · Nuria Oliver · Adrian Weller · Manuel Rodriguez · Isabel Valera · Silvia Chiappa · Hoda Heidari · Niki Kilbertus
- Perception as generative reasoning: structure, causality, probability** **E MR 1-3**
 Dan Rosenbaum · Marta Garnelo · Peter Battaglia · Kelsey Allen · Ilker Yildirim

Saturday Workshops - 8 am - 6:40 pm

- **Competition Track Day 2** **West 116-117**
Hugo Jair Escalante
- **Machine Learning with Guarantees** **West Brm B**
Ben London · Gintare Karolina Dziugaite · Daniel Roy · Thorsten Joachims · Aleksander Madry · John Shawe-Taylor
- **Machine Learning & the Physical Sciences** **W 109-110**
Atilim Gunes Baydin · Juan Carrasquilla · Shirley Ho · Karthik Kashinath · Michela Paganini · Savannah Thais · Anima Anandkumar · Kyle Cranmer · Roger Melko · Mr. Prabhat · Frank Wood
- **Learning Transferable Skills** **West 211-214**
Marwan Mattar · Arthur Juliani · Danny Lange · Matthew Crosby · Benjamin Beyret
- **Emergent Communication: Towards Natural Language** **West 118-120**
Abhinav Gupta · Michael Noukhovitch · Cinjon Resnick · Natasha Jaques · Angelos Filos · Marie Ossenkopf · Angeliki Lazaridou · Jakob Foerster · Ryan Lowe · Douwe Kiela · Kyunghyun Cho
- **Context and Compositionality in Biological and Artificial Neural Systems** **West 217-219**
Javier Turek · Shailee Jain · Alexander Huth · Leila Wehbe · Emma Strubell · Alan Yuille · Tal Linzen · Christopher Honey · Kyunghyun Cho
- **Privacy in Machine Learning (PriML)** **East MR 8+15**
Borja Balle · Kamalika Chaudhuri · Antti Honkela · Antti Koskela · Casey Meehan · Mi Jung Park · Mary Anne Smart · Mary Anne Smart · Adrian Weller
- **Sets and Partitions** **West 215-216**
Nicholas Monath · Manzil Zaheer · Andrew McCallum · Ari Kobren · Junier Oliva · Barnabas Poczos · Ruslan Salakhutdinov
- **The Third Conversational AI Workshop: Today's Practice & Tomorrow's Potential** **W 205-207**
Alborz Geramifard · Jason Williams · Bill Byrne · Asli Celikyilmaz · Milica Gasic · Dilek Hakkani-Tur · Matt Henderson · Luis Lastras · Mari Ostendorf
- **Deep Reinforcement Learning** **West Ex Hall C**
Pieter Abbeel · Chelsea Finn · Joelle Pineau · David Silver · Satinder Singh · Joshua Achiam · Carlos Florensa · Christopher Grimm · Haoran Tang · Vivek Veeriah
- **Real Neurons & Hidden Units: Future Directions at the intersection of Neuroscience and AI** **East Brm A**
Guillaume Lajoie · Eli Shlizerman · Maximilian Puelma Touzel · Jessica Thompson · Konrad Kording
- **Science meets Engineering of Deep Learning** **West 121-122**
Levent Sagun · Caglar Gulcehre · Adriana Romero · Negar Rostamzadeh · Nando de Freitas
- **Document Intelligence** **West 208-209**
Nigel Duffy · Rama Akkiraju · Tania Bedrax Weiss · Paul Bennett · Hamid Reza Motahari-Nezhad
- **Medical Imaging meets NeurIPS** **West 301-305**
Hervé Lombaert · Ben Glocker · Ender Konukoglu · Marleen de Bruijne · Aasa Feragen · Ipek Oguz · Jonas Teuwen
- **Bridging Game Theory & Deep Learning** **West Ex Hall A**
Ioannis Mitliagkas · Gauthier Gidel · Niao He · Reyhane Askari Hemmat · Nika Haghtalab · N H · Simon Lacoste-Julien
- **Program Transformations for ML** **West 114-115**
Pascal Lamblin · Atilim Gunes Baydin · Alexander Wiltchko · Bart van Merriënboer · Emily Fertig · Barak Pearlmutter · David Duvenaud · Laurent Hascoet
- **ML For Systems** **West 202-204**
Milad Hashemi · Azalia Mirhoseini · Anna Goldie · Kevin Swersky · Jonathan Raiman · Xinlei XU · Jonathan Raiman
- **NeurIPS Workshop on Machine Learning for Creativity and Design 3.0** **West 223-224**
Luba Elliott · Sander Dieleman · Adam Roberts · Jesse Engel · Tom White · Rebecca Fiebrink · Parag Mital · Christine Payne · Nao Tokui
- **Learning with Temporal Point Processes** **West 306**
Manuel Rodriguez · Le Song · Isabel Valera · Yan Liu · Abir De · Hongyuan Zha
- **Machine Learning for Autonomous Driving** **East MR 1-3**
Rowan McAllister · Nicholas Rhinehart · Fisher Yu · Li Erran Li · Anca Dragan
- **Tackling Climate Change with ML** **East Brm C**
David Rolnick · Priya Donti · Lynn Kaack · Alexandre Lacoste · Tegan Maharaj · Andrew Ng · John Platt · Jennifer Chayes · Yoshua Bengio
- **Fair ML in Healthcare** **East Brm B**
Shalmali Joshi · Irene Y Chen · Ziad Obermeyer · Sendhil Mullainathan
- **Robot Learning: Control and Interaction in the Real World** **West 220-222**
Markus Wulfmeier · Roberto Calandra · Kate Rakelly · Sanket Sayaji Kamthe · Danica Kragic · Stefan Schaal · Markus Wulfmeier
- **Joint Workshop on AI for Social Good** **East MR 11-12**
Fei Fang · Joseph Bullock · Marc-Antoine Dilhac · Brian Green · natalie saltiel · Dhaval Adjodah · Jack Clark · Sean McGregor · Margaux Luck · Jonathan Penn · Tristan Sylvain · Geneviève Boucher · Sydney Swaine-Simon · Girmaw Abebe Tadesse · Myriam Côté · Anna Bethke · Yoshua Bengio
- **"Do the right thing": machine learning and causal inference for improved decision making** **West Brm C**
Michele Santacatterina · Thorsten Joachims · Nathan Kallus · Adith Swaminathan · David Sontag · Angela Zhou
- **The Optimization Foundations of Reinforcement Learning** **West Brm A**
Bo Dai · Niao He · Nicolas Le Roux · Lihong Li · Dale Schuurmans · Martha White

2019 SPONSORS

Neural Information Processing Systems gratefully acknowledges the generosity of those individuals and organizations who have provided financial support for the NeurIPS 2019 conference. Their financial support enables us to sponsor student travel and participation, general support to host the conference, and the volunteers who assist us.

DIAMOND SPONSORS



LAMBDA - Lambda is an AI infrastructure company, providing computation to accelerate human progress. Our workstations, servers, and cloud services power engineers and researchers at the forefront of human knowledge. Our customers include Apple, Intel, Microsoft, Amazon Research, Tencent, Kaiser Permanente, MIT, Stanford, Harvard, Caltech, and the Department of Defense. With Lambda, you simply plug the system into the wall and get started making business and scientific breakthroughs.

LYFT - Level 5 is Lyft's self-driving division. We're working toward building the leading self-driving system for ridesharing, because we believe in a future where self-driving cars make transportation safer and more accessible for everyone. With fewer cars on the road and less pollution in the air, we can reshape cities around people, not cars.

MICROSOFT - Microsoft Research is where leading scientists and engineers have the freedom and support to propel discovery and innovation. Here, they pursue and publish curiosity-driven research in a range of scientific and technical disciplines that can be translated into products. With access to vast computing power, global multi-disciplinary teams tackle complex problems that drive breakthrough technologies and improve lives.

DOC.AI - Our vision is to build the most impactful AI platform in the world to deliver products and services that close the healthcare data loop to expand access, create new and better 'medicine', decrease costs and improve the health and wellbeing of billions of people around the world.

CITADEL & CITADEL SECURITIES - Citadel is a leading investor in the world's financial markets. For over a quarter of a century, we have sought to deliver industry-leading investment returns to clients including corporate pensions, endowments, foundations, public institutions, and sovereign wealth funds. For more information, visit Citadel.com. Citadel Securities is a leading global market maker, delivering a broad array of fixed income and equity products to banks, broker-dealers, government agencies, corporations, insurers, and sovereign wealth funds. Through innovation and efficiency, the firm provides liquidity with the goal of driving price discovery in more than 35 countries and making markets more competitive, open, and transparent. For more information, visit CitadelSecurities.com.

GRAPHCORE - Graphcore has created a completely new processor, the Intelligence Processing Unit (IPU), specifically designed for machine intelligence. The IPU's unique architecture means developers can run current machine learning models orders of magnitude faster. More importantly, it lets AI researchers undertake entirely new types of work, not possible using current technologies, to drive the next great breakthroughs in general machine intelligence. We believe our IPU technology will become the worldwide standard for machine intelligence compute. The performance of Graphcore's IPU is going to be transformative across all industries and sectors, whether you are a medical researcher, roboticist or autonomous car manufacturer.

GOOGLE - Google's mission is to organize the world's information and make it universally accessible and useful, and we advance that mission every day in incredible new ways. Research across Google provides new ways of looking at old problems and helps transform how we all work and live, and we think the biggest impact comes when everyone in the world can access it. To that end, we use state-of-the-art computer science techniques to solve problems for our users, our customers and the world, making it easier for you to do things every day, whether it's searching for photos of people you love, breaking down language barriers, or helping you get things done with your own personal digital assistant.

SALESFORCE - Salesforce Research advances state-of-the-art AI techniques, developing models and prototypes that pave the path for innovative products at Salesforce. We tackle real-world problems in areas such as image recognition and natural language processing from Salesforce's enterprise customers by harnessing the latest deep learning techniques.

DIAMOND SPONSORS



APPLE - We're a diverse collection of people, reimagining what's possible to help us do what we love in new ways. The people who work here have reinvented entire industries with the Mac, iPhone, iPad, and Apple Watch, and with services, including iTunes, the App Store, Apple Music, and Apple Pay.

AMAZON - We fundamentally believe that scientific innovation is essential to being the most customer-centric company in the world. It's this ability to have an impact at scale that allows us to attract some of the brightest minds in artificial intelligence and related fields. Our scientists continue to publish, teach, and engage with the academic community, but one of the unique aspects of research at Amazon is our working backwards method applied to science. Each research project starts with a customer pain point and a description of how a scientific solution will improve the customer experience.

ANTHEM.AI - Anthem.ai is building a world class AI team including engineers, clinical professionals, data scientists, and AI mathematicians from Silicon Valley. The group's goal is to improve patient outcomes with advanced technology such as Genomics, Machine Learning, User Experience, and especially Precision Medicine. We are setting strategic direction and working with teams across Anthem and all of Medicine to set a new standard for healthcare and continuously deliver AI capabilities.

OPENAI - OpenAI is an AI research "capped-profit" company based in San Francisco, California. Its goal is to ensure that artificial general intelligence benefits all of humanity, and seeks to achieve this by building safe AGI and sharing the benefits with the world. The company seeks to achieve its mission in line with the principles outlined in the OpenAI Charter.

IBM - IBM Research has been exploring artificial intelligence and machine learning technologies and techniques for decades. We believe AI will transform the world in dramatic ways in the coming years – and we're advancing the field through our

portfolio of research focused on three areas: Advancing AI, Scaling AI, and Trusting AI.

SONY - Sony is a major electronics, entertainment and finance conglomerate. Sony continues to innovate in all industries: Mobile Communications, Games & Network Services, Imaging Products & Solutions, Home Entertainment & Sound, Semiconductor, Components, Movies, Music, Finance and Other Businesses. Artificial Intelligence and Machine Learning are one of the key drivers in Sony's business success and the company has been at the forefront of developing Machine Learning technologies for finance, manufacturing, music and entertainment among many others.

QUALCOMM - Qualcomm invents breakthrough technologies that transform how the world connects, computes and communicates. When we connected the phone to the Internet, the mobile revolution was born. Today, our inventions are the foundation for life-changing products, experiences, and industries.

DEEPMIND - Artificial intelligence could be one of humanity's most useful inventions. DeepMind aims to build advanced AI to expand our knowledge and find new answers. By solving this one thing, we believe we could help people solve thousands of problems. We're a team of scientists, engineers, machine learning experts and more, working together to advance the state of the art in artificial intelligence. We use our technologies for widespread public benefit and scientific discovery, and collaborate with others on critical challenges, ensuring safety and ethics are the highest priority.

WAYMO - Waymo is building the world's most experienced driver. We are a self-driving technology company with a mission to make it safe and easy for people and things to move around. From our beginnings as the Google self-driving car project, we've been working to make our roads safer and increase mobility for the millions of people who cannot drive.

INTEL AI - Intel delivers a versatile mix of HW, enabling SW and optimizations for popular deep learning frameworks, and other tools to bring AI out of theory and into practice. Our wide-ranging HW for compute, acceleration, memory, storage, and interconnect delivers the right AI performance where it's needed, from workstations to clouds to network edge and devices. This is AI on Intel.

FACEBOOK - 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.

PLATINUM SPONSORS



INVENIA LABS - Invenia Labs uses machine learning to optimise the electricity grid, ensure demands are met at least-possible prices and minimum pollution. Our work helps to reduce emissions and pollution, improve the reliability of the grid and increase economic efficiency.

HABANA LABS - Habana Labs is an AI Processor company founded in 2016 to develop from the ground-up processor platforms that are optimized for training deep neural networks and for inference deployment in production environments. We are unlocking the true potential of AI with platforms offering orders of magnitude improvements in processing performance, scalability, cost, and power consumption. Habana is located in Tel-Aviv, San Jose, Beijing and Gdansk, employing over 150 people worldwide.

VOLEON GROUP - Voleon is a technology company that applies state-of-the-art machine learning techniques to real-world problems in finance. For more than a decade, we have led our industry and worked at the frontier of applying machine learning to investment management. We have become a multibillion dollar asset manager, and we have ambitious goals for the future.

QUANTUMBLACK - 'QuantumBlack, a McKinsey Company, helps companies use data to drive decisions. We combine business experience, expertise in large-scale data analysis and visualisation, and advanced software engineering know-how to deliver results. From aerospace to finance to Formula One, we help companies prototype, develop, and deploy bespoke data science and data visualisation solutions to make better decisions.

NAVER LINE - As Korea's No.1 internet company, NAVER Corporation accounts for over 76 percent of the country's search market, and operates a diverse range of services and products related to news, blogging, music, translations, webtoons, video and more. Based in Japan, LINE Corporation launched the LINE messaging app in June 2011 and since then has grown into a diversified platform, offering a variety of services and content for more than 200 million users around the globe. NAVER LABS Europe is the European research division of NAVER and is the largest industrial research organization in AI in France.

UBER - Good things happen when people can move, whether across town or towards their dreams. Opportunities appear, open up, become reality. What started as a way to tap a button to get a ride has led to billions of moments of human connection as people go all kinds of places in all kinds of ways with the help of our technology.

BOREALIS AI - Borealis AI, a RBC Institute for Research, is a curiosity-driven research centre dedicated to achieving state-of-the-art in machine learning. Established in 2016, and with labs in Toronto, Montreal, Edmonton, Waterloo, and Vancouver, we support open academic collaborations and partner with world-class research centres in artificial intelligence. With a focus on ethical AI that will help communities thrive, our machine learning scientists perform integrated research in areas such as reinforcement learning, natural language processing, deep learning, and unsupervised learning to solve ground-breaking problems in diverse fields.

PDT PARTNERS - PDT Partners has had significant success over the past 26 years as a quantitative investment manager by using a scientifically rigorous approach to develop and deploy trading strategies. We're looking for people who love digging into challenging, data-driven problems. We have a supportive, collaborative culture, and are committed to being a fantastic place to work!

XPRIZE - XPRIZE is a nonprofit organization that designs and manages public competitions intended to encourage technological development that could benefit humanity. XPRIZE mission is to bring about "radical breakthroughs for the benefit of humanity" through incentivized competition. It fosters high-profile competitions to motivate individuals, companies and organizations across all disciplines to develop innovative ideas and technologies that help solve the grand challenges that restrict humanity's progress

D.E. SHAW & CO - The D. E. Shaw group is trusted by investors across the world to manage their assets by seeking the optimal balance of risk and reward. We use a combination of quantitative and qualitative tools to uncover independent, hard-to-find sources of return across global public and private markets.

TWITTER - Twitter, Inc. (NYSE: TWTR) is what's happening and what people are talking about, all around the world. From breaking news and entertainment to sports and politics, from big events to everyday interests. If it's happening anywhere, it's happening first on Twitter. Twitter is where the full story unfolds with all the live commentary and where live events come to life unlike anywhere else.

BLOOMBERG - Bloomberg is building the world's most trusted information network for financial professionals. Our 5,500+ software engineers, data scientists and researchers are dedicated to advancing and building new systems for the Bloomberg Terminal to solve complex real-world problems.

NETFLIX - Netflix is the world's leading internet entertainment service with over 151 million paid memberships in over 190 countries enjoying TV series, documentaries and feature films across a wide variety of genres and languages. Members can watch as much as they want, anytime, anywhere, on any internet-connected screen. Members can play, pause and resume watching, all without commercials or commitments.

PLATINUM SPONSORS



SIGOPT - SigOpt empowers experts to build models capable of transforming their companies. By designing enterprise-grade software solutions that automate experiment management and hyperparameter optimization, SigOpt accelerates the impact of experimentation on machine learning, deep learning, simulation and other AI models. Enterprise modeling leaders who represent \$500B in market capitalization across six distinct verticals and algorithmic trading leaders with over \$300B assets under management rely on SigOpt to realize 8x faster model tuning, 30% productivity gains and reliable improvement in model performance.

CEREBRAS SYSTEMS - Cerebras Systems is a team of pioneering computer architects, computer scientists, deep learning researchers, and engineers of all types. We have come together to build a new class of computer to accelerate artificial intelligence work. In August 2019, we announced the largest chip ever built – the Cerebras Wafer Scale Engine (WSE). The WSE is 56X larger than the largest GPU and has 78X more compute cores and 3,000 times more on chip memory. In artificial intelligence work, large chips process information more quickly producing answers in less time. As a result, neural networks that in the past took weeks to train, can train in minutes on the Cerebras WSE. Visit us in booth 422.

YOKOGAWA ELECTRONICS COMPANY - Since its founding in Tokyo in 1915, Yokogawa has contributed to the development of a wide range of industries including oil, gas, chemicals, power, and iron & steel, and has established a global reputation as a leader in the control field. Yokogawa delivers comprehensive solutions that make use of machine learning and AI, and creates sustainable value through co-innovation with its customers.

BOSCH - Established in 1886 as a workshop for Precision Mechanics and Electrical Engineering, Bosch has developed into a multinational company with roughly 410,000 associates worldwide and revenues of 78.5 billion Euros. Bosch supplies technologies and services throughout the world for Mobility Solutions, Industrial Technology, Consumer Goods, and Energy Building & Technology. As a forward thinking company, Bosch's aim for the future is to build reliable and robust machines that will be able to learn constantly and act intelligently. Therefore, the Bosch Center for Artificial Intelligence (BCAI) was founded in early 2017 to deploy cutting-edge AI technologies across Bosch products and services creating solutions that are "invented for life".

JP MORGAN - At J.P. Morgan, we are using data science and artificial intelligence to transform financial services. Our data analytics capabilities enable us to serve our clients with greater depth and sophistication, building our reputation as leading technology bank. We look for exceptional talent: those who have the skills and imagination to transform the way we live and work, and change our industry.

INSPUR - As the world's leading AI computing provider, Inspur is fully engaged in the development of AI infrastructures on four layers, including computing platform, management and performance suite, optimized deep-learning frameworks, and application acceleration to deliver end-to-end, agile, cost-efficient and optimized AI solutions for its industry customers. According to IDC's First Half 2018 China AI Infrastructure Market Survey Report, with 51% market share, Inspur ranks first in the AI server market. IDC and Gartner's worldwide server market tracker for 2018 ranked Inspur No. 3 in the worldwide x86 server market. Committed to offering state-of-the-art computing for global customers through innovative design, Inspur has become a business partner of many leading companies in the world.

ALIBABA GROUP - Alibaba's mission is to make it easy to do business anywhere. We enable businesses to transform the way they market, sell and operate and improve their efficiencies. We provide the technology infrastructure and marketing reach to help merchants, brands and other businesses to leverage the power of new technology to engage with their users and customers and operate in a more efficient way.

BAIDU - Baidu is the leading Chinese language Internet search provider and the largest Chinese website globally. Baidu aims to make the complicated world simpler through technology.

DIDI - Didi Chuxing ("DiDi") is the world's leading mobile transportation platform. The company offers a full range of app-based transportation options for 550 million users, including Taxi, Express, Premier, Luxe, Bus, Designated Driving, Enterprise Solutions, Bike Sharing, E-bike Sharing, Car Rental and Sharing and food delivery. Tens of millions of drivers who find flexible work opportunities on the DiDi platform provide 10 billion passenger trips a year. DiDi is committed to collaborating with policymakers, the taxi industry and communities to solve the world's transportation, environmental and employment challenges with smart transportation innovations. The company is working with a growing alliance of car industry players to build out a next-generation auto-solutions and operations platform.

EY - EY is a global leader in assurance, tax, transaction and advisory services. The insights and quality services we deliver help build trust and confidence in the capital markets and in economies the world over. We develop outstanding leaders who team to deliver on our promises to all of our stakeholders. In so doing, we play a critical role in building a better working world for our people, for our clients and for our communities.

TWO SIGMA - At Two Sigma, our mission is to find value in the world's data. We hire scientists, engineers, and academics who tackle the world's most challenging financial problems. Our clients in investment management, insurance, securities, private equity and venture capital benefit from our research-led data science approach and advanced technology. Two Sigma has over 1,600 employees with offices in New York, Hong Kong, Houston, London, Shanghai, and Tokyo. Come seek with us.

PLATINUM SPONSORS



JUMP TRADING - Jump Trading is a leader in automated trading powered by applied machine learning, high performance computing, and custom hardware. Founded in 1999, Jump Trading remains a privately owned and funded company with an entrepreneurial culture and 700 employees across offices in Chicago, New York, London, Amsterdam, Singapore, and Shanghai.

JANE STREET - Jane Street is a proprietary trading firm that operates around the clock and around the globe. We bring a deep understanding of markets, a scientific approach, and innovative technology to bear on the problem of trading profitably in the world's highly competitive financial markets.

HUDSON RIVER TRADING - Hudson River Trading brings a scientific approach to trading financial products. We have built one of the world's most sophisticated computing environments for research and development. Our researchers are at the forefront of innovation in the world of algorithmic trading.

DISNEY RESEARCH - As part of The Walt Disney Company, Disney Research builds upon a rich legacy of innovation and technology leadership in the entertainment industry that continues to this day. Our research covers a broad range of exciting and challenging applications that are experienced daily by millions of people around the world.

UNITY TECHNOLOGIES - Unity democratizes development by enabling success for any creator with a vision. In order to pave the path to success, we solve the hard technical and business-related problems developers face. More game developers use Unity than any other engine. We are also used in a variety of other fields such as Film and Animation, Automotive and Manufacturing, Architecture and Construction, and Augmented/Virtual Reality!

BIOMIND - Hanalytics BioMind Pte Ltd (BioMind) is a deeptech company specialising in healthcare. The company builds advanced AI technology and creates predictive applications to help hospitals diagnose medical conditions and manage healthcare-related problems and is supported by a strong team of in-house deep learning scientists, medical experts, and research advisors from prestigious hospitals and universities. In December 2018, BioMind partnered Beijing Tiantan Hospital, a global leader in neurology and neurosurgery, to establish the world's first and largest AI research centre for neurological diseases.

SPORTLOGIQ - We are the world's sports analytics leader, using cutting-edge AI technology to produce unique sports data and performance insights. These insights provide an edge to teams, leagues, media organizations, betting operators and rights holders. We use a variety of techniques to collect and analyze sports data, allowing our partners to derive value from them in new ways. Sportlogiq partners with professional sports teams, including NHL, NFL and MLS teams, and broadcast networks. We are the Official Data Partner of the Swedish Hockey League, Official Insights Partner of the Canadian Premier League and Official Statistics Partner of the National Lacrosse League.

BYTEDANCE - ByteDance's platforms enable people to discover and create a world of content powered by technology. We inform, educate, entertain and inspire people across languages, cultures, and geographies.

NATIONAL SECURITY AGENCY - The National Security Agency/Central Security Service (NSA/CSS) leads the U.S. Government in cryptology that encompasses both Signals Intelligence (SIGINT) and Information Assurance (IA) products and services, and enables Computer Network Operations (CNO) in order to gain a decision advantage for the Nation and our allies under all circumstances.

YANDEX - Yandex is one of the largest internet companies in Europe, operating Russia's most popular search engine. We provide user-centric products and services based on the latest innovations in information retrieval, machine learning and machine intelligence to a worldwide customer audience on all digital platforms and devices.

CRUISE - Cruise is a San Francisco-based company building the world's most advanced self-driving vehicles to safely connect people with the places, things, and experiences they care about. We're at the beginning of the self-driving car industry, and Cruise is leading in the space. We've secured capital commitments totaling \$7.25B, including funds and accounts advised by T. Rowe Price Associates, Inc., General Motors, SoftBank Vision Fund, and Honda. Today, we lead the industry with fully integrated manufacturing at scale. Our innovative AI research team, deep resources, and progressive technology approach will help us launch all-electric, self-driving vehicles at scale and improve life in our cities.

GOOGLE X THE MOONSHOT FACTORY - X is a moonshot factory. Our builders, innovators and researchers leverage AI that, when partnered human creativity, can unlock solutions to some really hard problems. X's goal is to develop and de-risk early-stage technologies and turn them into products that can be the foundation for large, sustainable businesses, coupled with the riskiness of research and speed of a startup.

GOLD SPONSORS



WEIGHTS & BIASES - Weights & Biases is an experiment tracking platform for deep learning. With a few lines of code, explore how hyperparameters affect model performance in real time, record and visualize every detail of your research, compare across 1000s of runs, reproduce and share results and collaborate more easily.

MATHWORKS - The MATLAB and Simulink product families are fundamental applied math and computational tools at the world's educational institutions. Adopted by more than 5000 universities and colleges, MathWorks products accelerate the pace of learning, teaching, and research in engineering and science. MathWorks products also help prepare students for careers in industry worldwide, where the tools are widely used for data analysis, mathematical modeling, and algorithm development in collaborative research and new product development. Application areas include data analytics, mechatronics, communication systems, image processing, computational finance, and computational biology. mathworks.com

XTX MARKETS - XTX Markets is a leading quantitative-driven electronic market-maker partnering with counterparties, exchanges and e-trading venues globally to provide liquidity in the Equity, FX, Fixed Income and Commodity markets. Machine learning is at the heart of what we do and we are leading the world of financial technology for fair and efficient markets.

SQUAREPOINT - Squarepoint is a systematic investment manager with deep functional expertise in Quantitative Research, Trading & Technology. We design and build our own cutting-edge systems, from high performance trading platforms to large scale data analysis and compute farms. With main offices in New York, London, Singapore & Montreal we emphasize true global collaboration by aligning our teams functionally.

VINAI RESEARCH - VinAI Research is the first AI research lab located in Hanoi, Vietnam, focusing on top-tier research and publications. Funded by VinGroup, the largest enterprise in Vietnam by capitalization, our mission is to conduct high-impact research that pushes the knowledge frontier in AI and to accelerate applications of AI in Vietnam, the Asia Pacific region, and beyond.

G-RESEARCH - G-Research is a leading quantitative research and technology company. By using the latest scientific techniques, we produce world-beating predictive research and build advanced technology to analyse the world's data. Our mission is to develop models to forecast financial time series. This is a challenging and highly competitive space so rather than deploy standard methods off the shelf you will likely need to extend classical methods or develop entirely new techniques. Our problems are well-defined and success is highly measurable and has direct impact on the business. We employ cutting edge machine learning methods drawn from diverse areas such as neural networks and deep learning; non-convex optimisation; Bayesian non-parametrics and approximate inference. We have the freedom to...

HORIZON ROBOTICS - Horizon Robotics is a Chinese company, leading technology powerhouse of embedded Artificial Intelligence. The company is dedicated to providing integrated and open embedded Artificial Intelligence solutions of high performance, low power and low cost. We equip smart cameras and cars with "brains", turning them into intelligent entities that have the ability from perception, understanding to decision-making for convenience, safety and fun. After two years'R&D, Horizon Robotics unveiled Chinese first world-leading, Brain Processing Unit (BPU) based, proprietary Gauss-architecture embedded AI computer vision processors - Journey and Sunrise, powering smart cars and smart cameras, to provide industrial customers with a complete solution including algorithm, chip and cloud.


NEURAL MAGIC - Neural Magic is no hardware AI. With Neural Magic's software, data science teams can use ubiquitous and unconstrained CPU resources to achieve performance breakthroughs without specialized hardware.

PREFERRED NETWORKS - Preferred Networks (PFN) is a Tokyo-based startup that applies deep learning to industry. PFN develops Chainer, a deep learning framework. PFN works with Toyota Motor for autonomous driving, FANUC for manufacturing robots, and National Cancer Center Japan for healthcare. Recently we unveiled a personal robot system by exhibiting autonomous tidying-up robots. A subsidiary is located in California.

BENEVOLENT AI - BenevolentAI, founded in 2013, creates and applies AI technologies to transform the way medicines are discovered and developed. The company has developed the Benevolent Platform™ - a discovery platform used by BenevolentAI scientists to find new ways to treat disease and personalise drugs to patients.

PROWLER.IO - Our mission is to help leaders and organisations make better business decisions using trusted, explainable AI. Not in theory, not in the future – but right now - and in the real world. Our dynamic, decision-making AI has become a powerful tool for business, combining branches of mathematics and engineering in ways that have never previously been envisaged. This integrated approach - matched with our industry-leading research credentials - gives us a unique competitive advantage, helping us solve problems across industry sectors.

GOLD SPONSORS



PETUUM - Petuum provides innovative industry solutions with the most advanced artificial intelligence methodologies that have been out of reach for most businesses. Our products solve complex real-world challenges that traditional techniques have failed to solve in a simple, cost-effective manner. These are operationalized by world-class AI experts and deep subject matter specialists, on the Petuum AI Symphony platform, which delivers distributed computing and parallelized data processing as well as machine learning and deep learning workflows, at scale in every environment.

SCALE AI - Scale AI's mission is to accelerate the development of AI by democratizing access to intelligent data. Our suite of managed labeling services such as Sensor Fusion (For LiDAR and RADAR Annotation), 2D Box Annotation, 3D Cuboid Annotation, Semantic Segmentation, and Categorization combine manual labeling with best in class tools and machine driven checks to yield highly accurate training data.

PRYON - Pryon is an AI company that delivers augmented intelligence for the enterprise. Driven by the inventors of core natural language technologies, the company is developing a platform that connects employees to digital transformation, extending their ability to find and use knowledge, drive workflows, and make better decisions from wherever they are.

SPLUNK - Splunk Inc. (NASDAQ: SPLK) turns data into doing with the Data-to-Everything Platform. Splunk technology is designed to investigate, monitor, analyze and act on data at any scale, from any source over any time period. The Data-to-Everything platform removes the barriers between data and action, so our customers -- regardless of size or business -- have the freedom to deliver meaningful outcomes across their entire organization. Our unique approach to data has empowered companies to improve service levels, reduce operations costs, mitigate risk, enhance DevOps collaboration and create new product and service offerings.

SBERBANK - Sberbank is a powerful innovative bank which is rapidly becoming one of the major digital financial institutions. Sberbank is an international bank in the top 20 in terms of capitalization with offices in Switzerland, Austria, England, Turkey and a number of European countries. We are actively using artificial intelligence and machine learning technologies to empower our products and services.

MIPSOLOGY - Mipsology develops state-of-the-art FPGA-based accelerators targeted for deep learning applications in neural networks. It was founded in 2015 by a team of engineers and scientists who created a family of world-class FPGA-based super-computers over the past 20 years.

POINT 72/CUBIST SYSTEM STRATEGIES - Cubist Systematic Strategies, the quantitative investing business of Point72, deploys systematic, computer-driven trading strategies across multiple liquid asset classes, including equities, futures, and foreign exchange. 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.


ZILLOW GROUP - Zillow Group is committed to empowering consumers with unparalleled data, inspiration and knowledge around homes, and connecting them with the right local professionals to help. The company's brands focus on all stages of the home lifecycle: renting, buying, selling, financing and home improvement.

NOVARTIS - Our purpose is to reimagine medicine to improve and extend people's lives. Our strategy is to build a leading, focused medicines company powered by advanced therapy platforms and data science. We address some of society's most challenging healthcare issues. We discover and develop breakthrough treatments and find new ways to deliver them to as many people as possible.

SNAP INC. - Snap Inc. is a camera company. We believe that reinventing the camera represents our greatest opportunity to improve the way people live and communicate. We contribute to human progress by empowering people to express themselves, live in the moment, learn about the world, and have fun together.

CISCO - Cisco is a technology company that solves, creates, inspires, heals, and secures to help build your bridge to possible. <https://www.cisco.com/c/en/us/about/bridge-to-possible.html>

BOOK PUBLISHERS



SILVER SPONSORS



khosla ventures



ELEMENT^{AI}



EDGESTREAM - At Edgestream, we are devoted to fully-systematic, widely-diversified, absolute-return institutional investment management. Operating since 2003, our investment process has attracted some of the world's most discriminating institutional investors, establishing us as a highly regarded quantitative firm within the alternative investment industry.

CENTURION CAPITAL - Centurion Capital is a quantitative investment management company researching in global financial markets, dedicated to producing exceptional returns by strictly adhering to mathematical and statistical methods.

KHOSLA VENTURES - Khosla Ventures provides venture assistance and strategic advice to entrepreneurs working on breakthrough technologies. With over five billion dollars under management, the firm focuses on a broad range of areas including consumer, enterprise, education, advertising, financial services, semiconductors, health, big data, agriculture/food, sustainable energy and robotics.

AVIRA - Avira protects people in the connected world across all devices, both directly and via our OEM partnerships. Machine Learning is core to our products, especially to our threat detection and prevention engines. What makes our approach so successful is being able to combine our expertise in Machine Learning and AI with 30 years of experience in the cybersecurity industry.

GHIELIA CORP - GHIELIA Inc. was established in June 2017 as a joint venture between Sony CSL, UEI Corporation, and WiL, LLC.. GHIELIA aims to utilize AI technologies not only for business enterprises but also for human enhancement. In order to achieve this, we are working on a new AI platform suitable for lay people to easily develop their own AI systems and freely distribute them across the globe.

TENCENT AI LAB - Established in April 2016, Tencent AI Lab is a corporate-level research and application lab of artificial intelligence. AI Lab's vision is to "Make AI Everywhere" for the benefit of people in

all aspects of life. Its research focuses on four key areas: machine learning, computer vision, speech recognition, and natural language processing.

ARM - Arm defines the pervasive computing that's shaping today's connected world. Realized in 125+ billion silicon chips, our device architectures orchestrate the performance of the technology that's transforming our lives — from smartphones to supercomputers, from medical instruments to agricultural sensors, and from base stations to servers.

ELEMENT AI - Element AI is a global AI company that develops AI software products at scale to help people work smarter. Founded in 2016 by seasoned entrepreneur JF Gagné and pioneering AI researcher and A.M. Turing Award recipient, Yoshua Bengio, the company turns cutting-edge research and industry expertise into software solutions that continuously learn and improve.

ACCENTURE - Accenture is a leading global professional services company, providing a broad range of services and solutions in strategy, consulting, digital, technology and operations. With more than 450,000 people serving clients in over 120 countries, Accenture drives innovation to improve the way the world works and lives.

EXPEDIA - We are the world's travel platform. Our purpose is to bring the world within reach. We are among the largest technology companies in the world, and our work is solely dedicated to one of the most socially and economically important activities on the planet – Travel.

JOHNSON & JOHNSON - We're looking for people ready use their invaluable technology skills, expertise, and perspectives to change the trajectory of health for humanity. Data Science careers thrive at Johnson & Johnson because of the options and resources available through the largest and most broadly-based healthcare organization in world. See Data Science and Digital Health jobs at careers.jnj.com. #JNJDataScience.

CYLANCE - Cylance develops AI/ML to deliver prevention-first, predictive security products. Cylance provides full spectrum predictive threat prevention and visibility across the enterprise to combat advanced cybersecurity attacks. Cylance ML models power prevention, threat hunting, and automated detection without increasing staff workload or costs.

MOQI TECHNOLOGIES - Moqi (<https://FingerID.ai>) is an AI technology company dedicated to explore cutting-edge tech for massive unstructured data. Moqi developed a high-speed and high-precision fingerprint image search engine, FingerID. By building a national fingerprint center inside of a national police department, FingerID gained large-scale application to match one fingerprint from 2 billion ones within one second.

CAPITAL GROUP - Founded in 1931 and headquartered in the beautiful southern California area, Capital Group is one of the world's largest and most trusted investment management companies and home to the American Funds. With US\$1.8 trillion in assets and over 7,500 associates around the globe, we make our clients and investors our first priority each and every day.

SILVER SPONSORS



Booz | Allen | Hamilton*



TERRAQUANTA - TerraQuanta is a tech-driven company based in Beijing, China. We teamed up at the end of 2016 and operation began in January 2017. We are dedicated to the development of AI algorithms and data processing systems for satellite remote sensing data. As of now, we are working on products that make impacts on businesses such as environmental conservation, agriculture, finance, energy, etc.

ALEGION - Alegion is an Austin-based technology company that provides the most powerful and flexible annotation platform for training data in market. It accelerates model development for the most sophisticated and subjective use cases. It uses integrated ML and has unique capabilities like conditional logic, iterative tasks, multi-stage and workflows, that are essential for high quality at scale.

VECTRA - Vectra is the world leader in applying artificial intelligence to detect and respond to cyberattacks in cloud, data center and enterprise infrastructures in real time, while empowering security analysts to perform conclusive incident investigations and AI-assisted threat hunting.


BOOZ | ALLEN | HAMILTON - For more than 100 years, business, government, and military leaders have turned to Booz Allen Hamilton to solve their most complex problems. As a consulting firm with experts in analytics, digital, engineering, and cyber, we help organizations transform. To learn more, visit BoozAllen.com.

HAPPY ELEMENTS - Happy Elements is a leading interactive entertainment company with products and services in games, comic & animation, IP affiliated products, virtual idols. Founded in 2009, we have over 900 employees, with offices in Beijing, Shanghai, Tokyo, Kyoto and San Francisco. We apply AI and Data Science in games to optimize the game productivity and generate the best experience for all users.

WALMART LABS - Imagine working in an environment where one experiment can catapult an entire industry toward a smarter future. That's what we do at Walmart Labs. We're a team of 5,000+ software engineers, data scientists, designers and product managers within Walmart, the world's largest retailer, delivering innovations to improve how our customers shop and our enterprise operates. Join us today!

SIEMENS HEALTHINEERS - At Siemens Healthineers, our purpose is to enable healthcare providers to increase value by empowering them on their journey toward expanding precision medicine, transforming care delivery, and improving patient experience, all made possible by digitalizing healthcare.

BRONZE SPONSORS



Wadhvani AI

NEXT AI - NextAI is Canada's premiere startup accelerator in Toronto and Montreal. NextAI is for entrepreneurs, researchers and scientists launching AI-enabled ventures. We provide up to \$150K in seed funding and founder development in the form of in-depth business and technical education taught by award-winning faculty. Plus workspace, mentorship and visa support for international applicants.

SIMON FRASER UNIVERSITY - As Canada's engaged university, Simon Fraser University works with communities, organizations and partners to create, share and embrace knowledge that improves life and generates real change. We connect research and innovation to entrepreneurship and industry to deliver sustainable, relevant solutions to today's problems. Engage with us at sfu.ca.

WADHWANI INSTITUTE - Wadhvani Institute for Artificial Intelligence Foundation is an independent non-profit research institute and global hub developing AI solutions for social good.

GRAMEEN RESEARCH - Grameen Research, Inc. www.GrameenResearch.org is a not-for-profit organization with a mission to provide research, training & other support in the field of microfinance & other services for low income populations. Our purpose is to engage in the business of supporting microlending & related services in the world as a means of assisting the world's low income populations to support themselves.

LAB 41 - Lab41 is a Silicon Valley research lab where experts from the U.S. Intelligence Community, academia, industry, and IQT come together to gain a better understanding of how to work with – and ultimately use – data analytics. Lab41 allows participants from diverse backgrounds to gain access to ideas, talent, and technology to explore what works and what doesn't in data analytics. An open,...

DIAMOND SPONSORS

Lambda
Amazon
Apple
Intel AI
Lyft

Microsoft
Anthem.ai
Facebook
Google
OpenAI

Graphcore
IBM Research
Qualcomm
Citadel & Citadel Securities
Doc.ai

Sony
Deep Mind
Waymo
Salesforce

PLATINUM SPONSORS

Invenia Labs
Naver Line
XPRIZE
SigOpt
J.P. Morgan
Baidu
Jump Trading
BioMind
Cruise

Habana Labs
Uber
D. E. Shaw Group
Cerebras Systems
Netflix
EY
Jane Street
Sportlogiq
National Security Agency

Voleon Group
Borealis AI
Twitter
Yokogawa Electric Corp.
Inspur
DiDi
Unity Technologies
Hudson River Trading
Yandex

QuantumBlack
PDT Partners
Bloomberg
Bosch
Alibaba Group
Two Sigma
Disney Research
ByteDance
Google X
(The moonshot Factory)

GOLD SPONSORS

MathWorks
G-Research
BenevolentAI
Scale AI
Mipsology

Weights & Biases
VinAI Research
Preferred Networks
Pryon
Snap Inc
Point 72/Cubist Systematic Strategies

XTX Markets
Horizon Robotics
Prowler.io
Splunk Inc.
Zillow
Cisco Systems

Squarepoint
Neural Magic
Petuum Inc.
Sberbank
Novartis

SILVER SPONSORS

Avira
Tencent AI Lab
Expedia Groups
Vectra AI
Walmart Labs

Khosla Ventures
Arm Inc
Johnson & Johnson
TerraQuanta
Booz Allen Hamilton
Siemens Medical Solutions

Edgestream Partners
Element AI
Blackberry Cylance
Moqi Technology
Happy Elements

GHELIA Inc.
Accenture
Capital Group
Alegion
Centurion Capital

BRONZE SPONSORS

NextAI (NEXT CANADA) Simon Fraser U. Grameen Research Lab41 Wadhwani AI

30% of your funds support Next Generation Outreach: For 2019, every sponsor will have 30% of their funds directed to a new need-based program to support travel, hotel, food, and registration costs for those in financial need, particularly graduate students and diversity and inclusion efforts. We are adding a new administrative position within the Neural Information Processing Systems Foundation to provide support to our affinity groups (WiML, Black in AI, LatinX in AI, {Dis}Ability in AI, and Queer in AI). We have made conference talks and presentations available via livestream to allow those not traveling to the conference to access remotely.