

Platinum Sponsors



Gold Sponsors



Silver Sponsors



Supporting Sponsors



Exhibitors

The following exhibitors will have booths in Lille Métropole (where the coffee breaks and poster sessions are) during the conference:

- 1000mercis • Alibaba • Amazon • Crédit Agricole • Criteo • CUP • Disney Research • EF • Euratechnologies • Facebook • Future of life • G-research • Google • MFG Labs • Microsoft/MSR • MIT Press • Nvidia • Panasonic Silicon Valley Laboratory • Springer • Technicolor • Winton • Xilinx

Contents

Sponsors and exhibitors	3
Introduction to ICML 2015	5
ICML 2015 at a glance	7
Local information	9
Organizing committee	13
Keynote speakers	15
Awards	17
Tutorials	18
Main conference	21
Workshops	37
Reviewer awards	57
Authors	59
Sponsor Scholars	65

Front cover: The size of a node represents the number of accepted papers in that field.

Back cover: The map denotes the location of authors. Links denote co-authorship.

Introduction to ICML 2015

Dear ICML attendees,

Welcome to Lille and the 32nd International Conference on Machine Learning (ICML 2015). We are excited to bring the premiere machine learning conference to France for the first time. The ICML committee has planned a wonderful conference.

Technical Program: We have 270 outstanding articles, selected from 1037 submissions. Each author will present their article to the community in a 20-minute talk, and present a poster at one of the poster sessions for discussion in smaller groups. All accepted articles are published in the Journal of Machine Learning Research (JMLR) as Volume 37 of their Workshop and Conference Proceedings series.

Keynote Speakers: We have three invited keynote speeches from some of the world's intellectual leaders: Léon Bottou (Facebook), Jon Kleinberg (Cornell University), and Susan Murphy (University of Michigan).

Tutorials: Six invited tutorials spanning some of the most vital subjects in machine learning: structured prediction, time series modeling, natural language understanding, reinforcement learning, optimization, and computational social sciences.

Workshops: 20 focused workshops for presenting late-breaking research and exploring new areas of machine learning.

Awards: We will present two best paper awards to honor some of the most promising research from the technical program. We will also present the ICML-2015 test of time award. This award is for the paper from the 2005 ICML (Bonn, Germany) that has retrospectively had a significant impact on our field.

We would like to acknowledge all the people who made exceptional efforts and dedicated their time to bring this conference together; we were honored to work with them.

Reviewing and selecting papers for the technical program was a mammoth task. We worked with 105 wonderful area chairs and 786 dedicated reviewers to give each paper three high-quality reviews and make an informed (if sometimes difficult) decision. The entire program committee generously offered their time and expertise to the machine learning community, and we thank them. Some reviewers offered extra dedication; 52 are recognized with an ICML Reviewer Award (page 57). The complete list of the program committee is available on the ICML web site.

In addition to the program committee, we would like to recognize and thank the entire organizing committee (page 13) who put the conference together. Planning for the tutorials, workshops, volunteers, publications, and sponsorship was ably organized and executed by this team. Their efforts over the past year are the backbone of this fantastic event.

We would like to offer special recognition to several people. First, we thank Joelle Pineau, the General Chair, who provided leadership, direction, and advice throughout the planning process. Second, we thank Philippe Preux, the local organizer. Philippe and his team gave their time and energy to see to the many details around the day-to-day of this year's ICML. Last, we thank Alp Kucukelbir, the workflow chair. Alp's help was invaluable in nearly every aspect of our planning process; neither of us can imagine performing this task without him.

Finally, we want to acknowledge our sponsors (page 3) and the IMLS board. ICML 2015 is not possible without their continued support.

On behalf of all of us at ICML 2015, enjoy the conference!

Francis Bach and David Blei
ICML 2015 Program Co-Chairs

Dear ICML attendees,

I am very honored to welcome you to Lille for the 32nd International Conference in Machine Learning. I am very excited that the first ICML in France is held here in Lille.

It has been a great pleasure for me to collaborate with Joelle, David, Francis, and all other chairs to organize ICML. I would like to acknowledge the friendly, efficient, and dedicated atmosphere in which we have worked together.

Organizing such a large scientific conference is a big endeavor. I have relied on many people and I would like to acknowledge their efforts.

First, I would like to thank Lille Grand Palais (LGP) and in particular Soazig Jacquart, who has managed the organization of ICML at LGP. Many others at LGP have also been involved and I want to thank them all.

Second, I would like to thank Inria Lille Nord Europe and, more specifically, the director David Simplot-Ryl. David has made so many things much easier for the local organization of ICML. I would also like to recognize the continuous and enthusiastic support of Nicolas Roussel, the great involvement of Marie-Bénédicte Dernoncourt, and also Amélie Supervielle and Marie-Agnès Énard for their daily help and advice.

All members of the Sequel research group have continuously supported this organization. In particular, I would like to thank Jérémie Mary, officially the ICML 2015 webmaster, but actually much more than that; I discussed many aspects of the organization with him; he is a computer wizard, a talented designer (Jérémie designed all the graphical aspects of ICML 2015, including the cover of this booklet), and so many other things. Romaric Gaudel also deserves special thanks for managing the troop of 120 student volunteers. I am sure they will be efficient and very helpful during the conference. I also wish to thank the local student volunteers, most of them belonging to Sequel.

Finally, I would like to very sincerely thank the Métropole Européenne de Lille, the Région Nord-Pas de Calais, the Université de Lille 1, and the CRISAL for their financial support.

On behalf of the organization team, welcome to ICML, welcome to Lille, and enjoy your stay!

Philippe PREUX
ICML 2015 Local Organization Chair

ICML 2015 at a glance

The registration desk will be open everyday from 8am to 6pm.

Monday, July 6		
08:00	Registration opens	
09:15 – 10:15	Tutorials 1 and 2	(various rooms)
10:15 – 10:45	<i>Coffee break</i>	Lille Métropole
10:45 – 11:45	Tutorials 1 and 2 (continued)	(various rooms)
11:45 – 13:30	<i>Lunch break (on your own)</i>	
13:30 – 15:40	Tutorials 3 and 4	(various rooms)
15:40 – 16:10	<i>Coffee break</i>	Lille Métropole
16:10 – 18:20	Tutorials 5 and 6	(various rooms)
18:30 – 21:00	Welcome cocktail	Jeanne de Flandres

Tuesday, July 7		
08:30 – 08:40	Welcome	Vauban
08:40 – 09:40	Invited talk: Léon Bottou	Vauban
09:40 – 10:00	Best paper award	Vauban
10:00 – 10:30	<i>Coffee break</i>	Lille Métropole
10:30 – 12:10	Session 1	(various rooms)
12:10 – 14:10	<i>Lunch break (on your own)</i>	
14:10 – 16:10	Session 2	(various rooms)
16:10 – 16:40	<i>Coffee break</i>	Lille Métropole
16:40 – 18:00	Session 3	(various rooms)
18:00 – 22:00	Poster session: papers from sessions 1–5	Lille Métropole

Wednesday, July 8		
08:30 – 09:50	Session 4	(various rooms)
09:50 – 10:20	<i>Coffee break</i>	Lille Métropole
10:20 – 12:00	Session 5	(various rooms)
12:00 – 14:00	<i>Lunch break (on your own)</i>	
14:00 – 15:00	Invited talk: Susan Murphy	Vauban
15:00 – 15:20	Best paper award	Vauban
15:30 – 16:30	Session 6	(various rooms)
16:30 – 17:00	<i>Coffee break</i>	Lille Métropole
17:00 – 18:00	Session 7	(various rooms)
18:00 – 22:00	Poster session: papers from sessions 6–11	Lille Métropole

Thursday, July 9		
08:30 – 09:50	Session 8	(various rooms)
09:50 – 10:20	<i>Coffee break</i>	Lille Métropole
10:20 – 12:00	Session 9	(various rooms)
12:00 – 14:00	<i>Lunch break (on your own)</i>	
14:00 – 15:00	Invited talk: Jon Kleinberg	Vauban
15:00 – 15:20	Test of Time award	Vauban
15:30 – 16:30	Session 10	(various rooms)
16:30 – 17:00	<i>Coffee break</i>	Lille Métropole
17:00 – 18:00	Session 11	(various rooms)
18:00 – 19:00	IMLS business meeting	Van Gogh
19:30 – 23:00	ICML banquet	Hall Londres

Friday, July 10		
08:30 – 10:00	Workshop sessions	(various rooms)
10:00 – 10:30	<i>Coffee break</i>	Lille Métropole
10:30 – 12:00	Workshop sessions	(various rooms)
12:00 – 14:00	<i>Lunch break (on your own)</i>	
14:00 – 16:00	Workshop sessions	(various rooms)
16:00 – 16:30	<i>Coffee break</i>	Lille Métropole
16:30 – 18:00	Workshop sessions	(various rooms)

Saturday, July 11		
08:30 – 10:00	Workshop sessions	(various rooms)
10:00 – 10:30	<i>Coffee break</i>	Lille Métropole
10:30 – 12:00	Workshop sessions	(various rooms)
12:00 – 14:00	<i>Lunch break (on your own)</i>	
14:00 – 16:00	Workshop sessions	(various rooms)
16:00 – 16:30	<i>Coffee break</i>	Lille Métropole
16:30 – 18:00	Workshop sessions	(various rooms)

Local information

Wifi. All ICML participants can get Wifi access by connecting to SSID `icml2015` with password `lille2015`. (The dot marks the end of the sentence, it is not part of the password.)

Lunches. All lunches are “on your own” (the cost is not covered by your ICML registration fee).

There is a possibility of having a simple lunch at the bar of Lille Grand Palais. Sandwiches, pasta, hamburgers, salad, and the like will be served. The bar will be open from 11am to 3pm.

There are many possibilities for lunch outside Lille Grand Palais, within a 5–10 minute walk. Head towards Lille-Flandres station and you’ll find many places. Head further 5-10 minutes, towards the “Grand Place” or the old city, and you’ll find much more.

Lille. Beyond lunches themselves, we propose a selection of areas of interest in the center of Lille, as well as a few hints about Lille. See the map below where these areas are indicated.

About safety: the center of Lille is safe all day long (particularly, the part which is shown on the map below). However, as in all large towns in Europe, you are advised to remain cautious about your belongings.

Lille Flandres - Euralille - Lille Europe 

The closest lunch spots to the conference center are the mall Euralille and the surrounding of the train station Lille Flandres. In this area you will mostly find fast food restaurants and supermarkets. Better quality food can be found near Rue Faidherbe and Rue du Molinel, close to the train station. Most of the eating places are located in the smaller neighbouring streets.

Rue de Solférino (M1 - Rep. Beaux Arts) 

Rue de Solférino is known for its nightlife. In addition to bars and pubs, you will find various fast food restaurants (kebab, friterie, etc.). Place des Halles Centrales gathers most of the bars. It is located at the intersection of Rue Solférino and Rue Masséna. The neighboring places like Place Sébastopole, Rue des Stations, or Rue Henri Kolb are also full of bars, clubs and restaurants. Most bars close around 2 a.m.

Vieux Lille (M1 - Rihour) 

Vieux Lille is the must-go historical center. You will find souvenir shops and fine food stores, like bakeries. Restaurants in the historical center propose traditional and gastronomic cooking. Bars have a wide selection of Belgian draft beers and French wine. To get the most of summer evenings you should definitely hang around next to the Cathedrale Notre-Dame-de-la-Treille where there are a few terraces on a public square. For bars, you should go to Rue Doudin. For restaurants, we recommend Rue Royale, Rue de Pétérinck, or Rue des Vieux Murs. For a more affordable dinner, Rue Saint Etienne is a good option.

Rihour-Opera (M1 - Opera) 

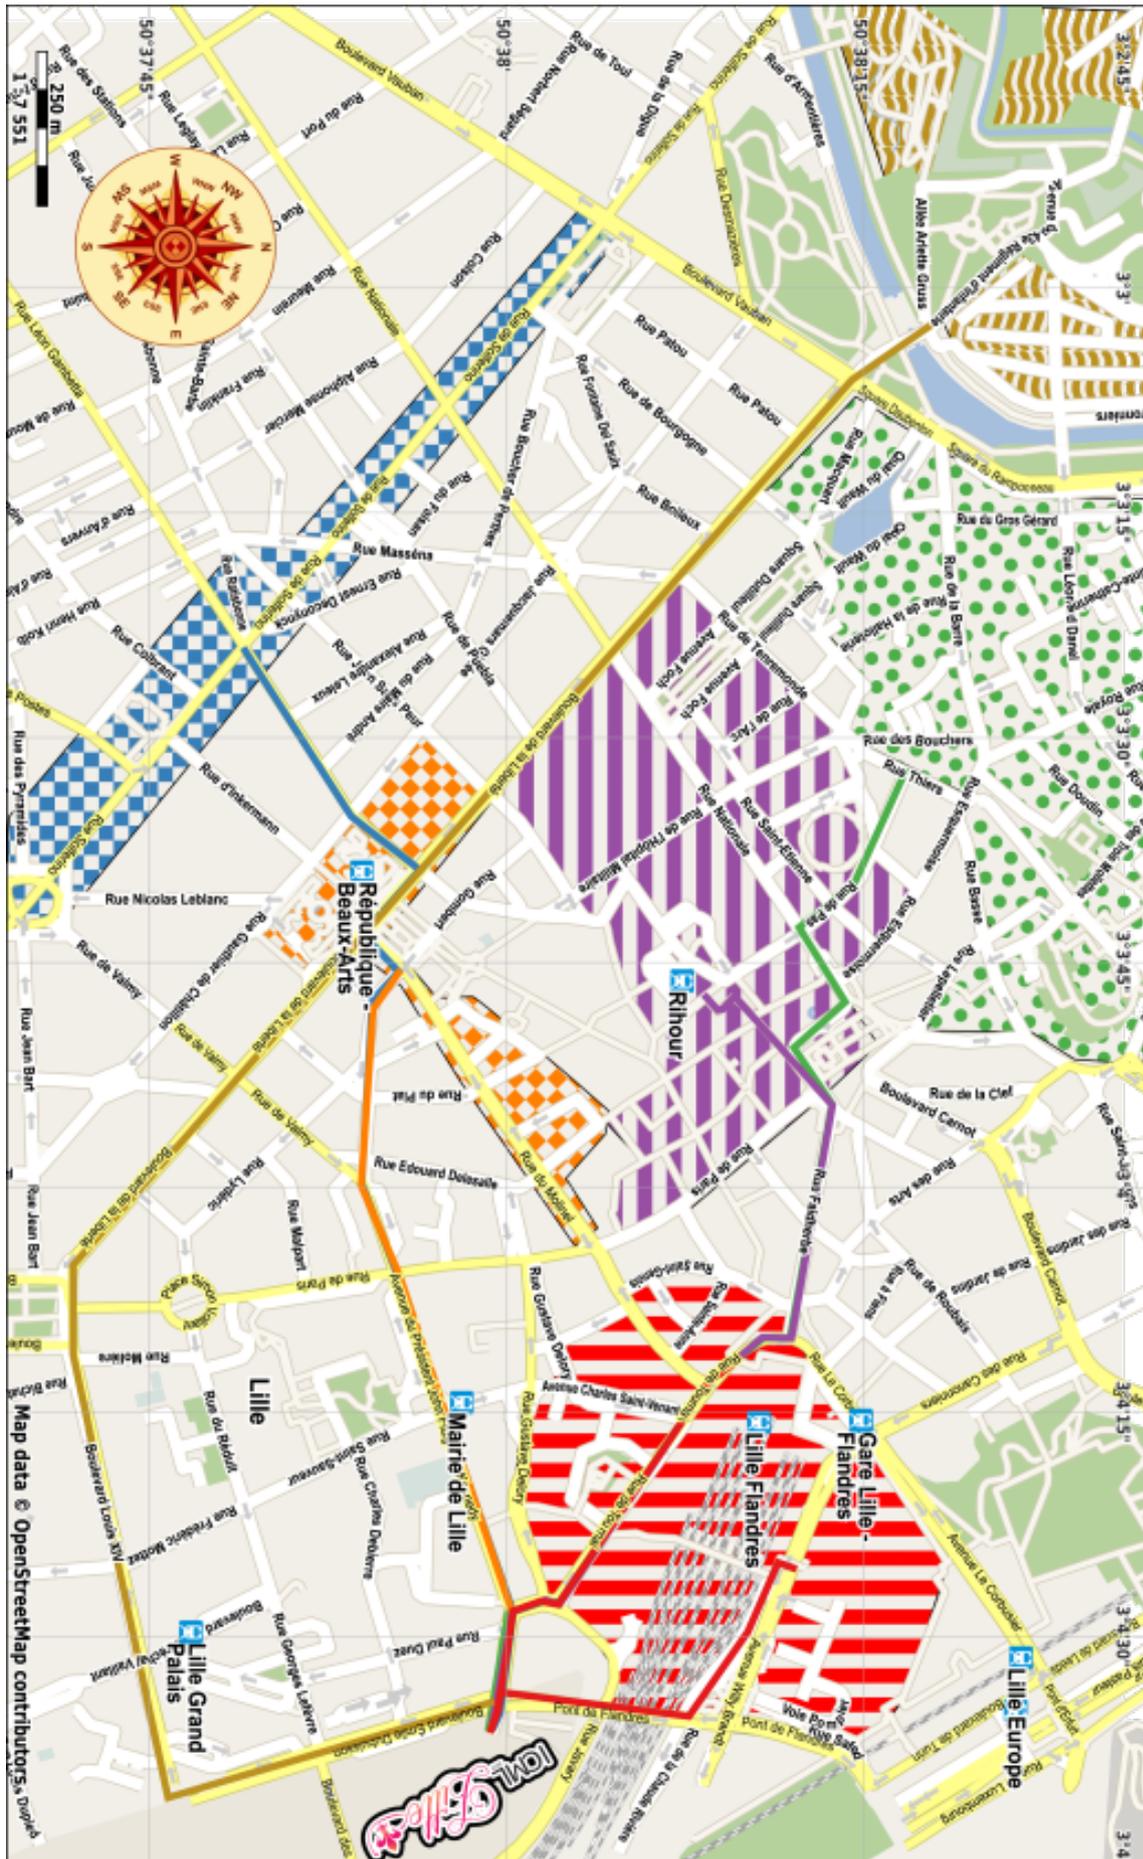
Rihour-Opera is the heart of the city. Around the main public square is Grand Place, where you will find trade-mark buildings of Lille, such as “La Vieille Bourse” or the building of the French newspaper “La Voix du Nord”. You will also find dozens of restaurants and terrace stands on Grand Place or in adjacent streets: Place Rihour, Rue des 7 Agaches, or Rue de Pas.

République - Rue du Molinel (M1 - rep. Beaux Arts) 

République/Rue du Molinel is between rue de Solférino and Rihour. In this place you can visit the Museum of Fine Art. There are two smaller squares near République where you will find restaurants, namely Place Richebé and Place de Béthune. Pushing North on Rue de Bethune you will find cinemas, brasseries. . .

Citadelle de Lille (M1 - Rihour) 

The Citadel is an ancient military fort built by Vauban, who is a famous French military architect. This place is still used by the military, but it is also one of the largest green areas in Lille. Inside the park, there is a small zoo that you can visit for free. You can also walk along the river La Deule. A lot of people go there to jog. It is also possible to rent a boat or a kayak.



Transportation *Bus and Metro*

The center of Lille is not big and you would not walk for more than half an hour to reach any place within this area. If you only wish to stay in the center of the city, a car is definitely not a good option since it is hard to park in some areas (especially in Vieux Lille). Public transportation works well and Lille Flandres is where buses stop and the two subway lines intersect. To take the public transportation you will need to buy a card on which you can charge trips. The cards are sold in vending machines, inside every subway station, or you can buy them directly in the bus.

Bike Rental

There are 2,000 bicycles in Lille available 24/7, in the 210 stations located in the city. Access to the service ranges from 1,40€ for the day to 3€ per month. The use is free during the first 30 minutes, then each additional half hour costs 1€. It takes only 2 minutes to subscribe from the terminal and take a bicycle. However, you will be requested to enter your credit card details, as a guarantee for the value of the bicycle.

Train

Be careful when leaving the city: there are two train stations. Fortunately they are at a walking distance from each other (taking the metro actually takes longer). Lille Europe serves big European cities like London, Brussels and Paris. From Lille Flandres you can reach closer destinations like Paris, Arras, Dunkerque, or Lens.

TGVs to Charles de Gaulle airport depart from either station. Check your ticket!

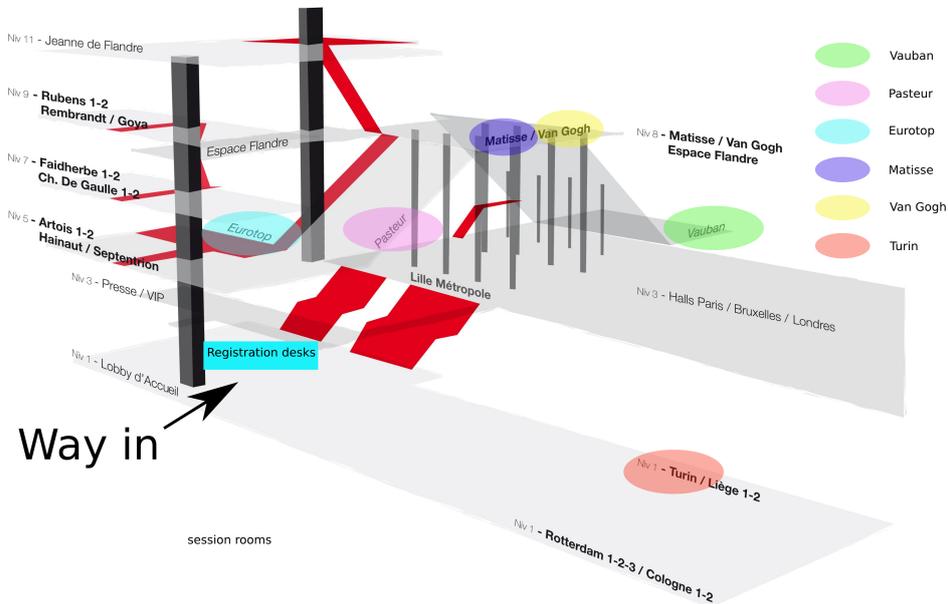
Local dishes Lille lives up to French (and Belgian) food reputation. Northern-French cooking is mainly based on cheese, beer and French fries. For those willing to try the typical local dishes, here is what we recommend:

- La carbonnade flamande: Traditional sweet-sour beef and onion stew made with beer and gingerbread.
- La Flamiche au maroilles: Beware if you are on a diet. This northern pie is made of maroilles, a tasty but (very) smelly cheese.
- Le Welsh: Belgian/British meal that mixes melted cheddar, Belgian beer, eggs, and ham. The best way to eat it is to dip some French fries inside.
- Les Moules-Frites: Dish of mussels and fries. Some are cooked with white wine (Moules marinières), other with Maroilles or Cream. Go to Grand Place to find such restaurants.
- Waffles: Lille's waffle are filled with sugar, butter and vanilla! This was Charles de Gaulle's favorite dessert.
- Le filet américain: Northern version of the steak tartare.

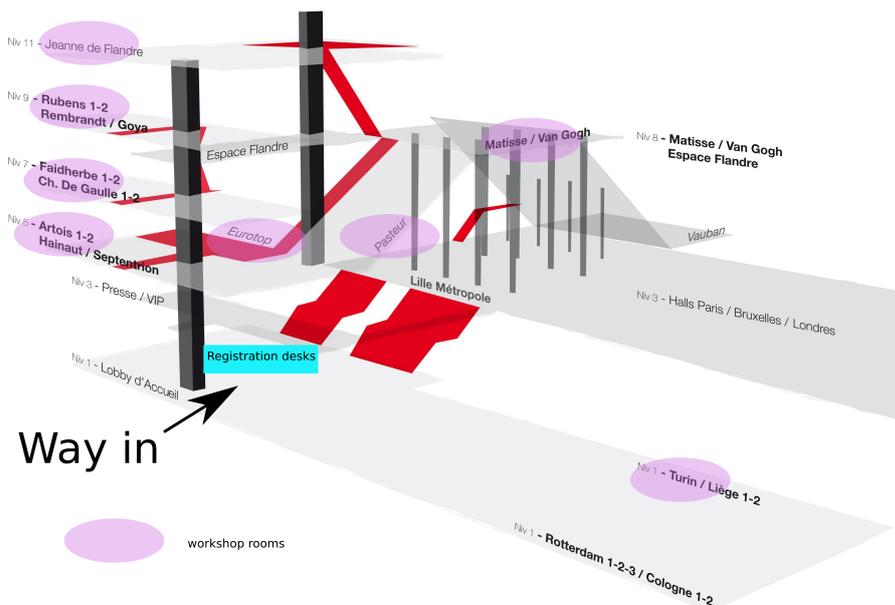
Map of Lille Grand Palais.

Coffee breaks and poster sessions are located in the Lille Métropole platform. Exhibitor booths are also located in this area.

Session rooms (in varied colors).



Workshop rooms (in purple).



Organizing committee

General chair:	Joelle Pineau (McGill University)
Program chairs:	Francis Bach (INRIA – Ecole Normale Supérieure) David Blei (Columbia University)
Local organization chair:	Philippe Preux (University of Lille, France)
Tutorial chair:	Jennifer Neville (Purdue University)
Workshop chair:	Marco Cuturi (Kyoto University)
Volunteer chair:	Ben Marlin (University of Massachusetts Amherst)
Publication chairs:	Percy Liang (Stanford University) David Sontag (New York University)
Publicity chair:	Jingrui He (Arizona State University)
Financial Chairs:	Gert Lanckriet (University of California, San Diego) Kilian Weinberger (Washington University in St. Louis)
Workflow chair:	Alp Kucukelbir (Columbia University)
Webmaster:	Jérémie Mary (University of Lille, France)
Local volunteer chair:	Romaric Gaudel (University of Lille, France)

Senior Program Committee

- Abernethy, Jacob, U. of Michigan
- Adams, Ryan, Harvard
- Agarwal, Alekh, Microsoft
- Airoidi, Edoardo, Harvard
- Anandkumar, Animashree, UCI
- Archambeau, Cedric, Amazon
- Balcan, Maria-Florina, CMU
- Banerjee, Arindam, U. of Minnesota
- Bengio, Samy, Google
- Bengio, Yoshua, U. of Montreal
- Bennett, Paul, Microsoft
- Bhattacharyya, Chiranjib, Indian Institute of Science
- Bilmes, Jeff, U. of Washington
- Bordes, Antoine, Facebook
- Bouchard, Guillaume, Xerox
- Boutillier, Craig, U. of Toronto
- Boyd-Graber, Jordan, U. of Maryland
- Brunskill, Emma, CMU
- Bubeck, Sebastien, Princeton
- Cappe, Olivier, CNRS
- Caramanis, Constantine, U. of Texas
- Cortes, Corinna, Google
- Courville, Aaron, U. Montreal
- Cramer, Koby, Technion
- Cuturi, Marco, Kyoto
- d'Alche-Buc, Florence, IBISC
- De Freitas, Nando, University of Oxford
- Duchi, John, Stanford
- Elkan, Charles, UC San Diego
- Engelhardt, Barbara, Duke
- Fergus, Rob, Facebook
- Fox, Emily, University of Washington
- Fukumizu, Kenji, Institute of Statistical Mathematics
- Globerson, Amir, Hebrew University
- Grauman, Kristen, U. of Texas
- Harchaoui, Zaid, INRIA
- Hein, Matthias, Saarland
- Heller, Katherine, Duke University
- Hoffman, Matthew, Adobe
- Hsu, Daniel, Columbia
- Ihler, Alex, UC Irvine
- Kappen, Bert, Radboud University
- Kolar, Mladen, U of Chicago
- Krause, Andreas, ETH Zurich
- Kulis, Brian, Ohio
- Kwok, James, HKUST
- Lacoste-Julien, Simon, INRIA
- Langford, John, Microsoft
- Larochelle, Hugo, U Sherbrooke
- Laviolette, Francois, U of Laval
- Le Roux, Nicolas, Criteo
- Lee, Honglak, University of Michigan
- Lescovec, Jure, Stanford

- Li, Lihong, Microsoft
- Liang, Percy, Stanford
- Nguyen, Long, U. of Michigan
- Lozano, Aurelie, IBM Research
- Mackey, Lester, Stanford
- Mairal, Julien, INRIA
- Marlin, Ben, U. of Massachussetts
- McAuliffe, Jon, Berkeley
- Meila, Marina, U. of Washington
- Mimno, David, Cornell
- Monteleoni, Claire, George Washington U.
- Munos, Remi, INRIA
- Murray, Iain, University of Edinburgh
- Nowozin, Sebastian, Microsoft
- Obozinski, Guillaume, Ecole des Ponts
- Paisley, John, Columbia
- Pontil, Massimiliano, UCL
- Poupart, Pascal, University of Waterloo
- Precup, Doina, McGill
- Ravikumar, Pradeep, UT Austin
- Rish, Irina, IBM Research
- Rosasco, Lorenzo, MIT
- Roth, Dan, UIUC
- Rudin, Cynthia, MIT
- S. V. N., Vishwanathan, UCSC
- Salakhudinov, Ruslan, U. of Toronto
- Scheffer, Tobias, Potsdam
- Schmidt, Mark, U. of British Columbia
- Schuurmans, Dale, U of Alberta
- Sha, Fei, U of South. Cal.
- Shalev-Shwartz, Shai, Hebrew University of Jerusalem
- Singh, Aarti, CMU
- Singh, Satinder, U. of Michigan
- Smola, Alex, CMU
- Sontag, David, NYU
- Sra, Suvrit, Max Planck Institute
- Srebro, Nathan, TTI-Chicago
- Stone, Peter, U. of Texas
- Sudderth, Erik, Brown
- Sutton, Rich, U. of Alberta
- Szepesvari, Csaba, U of Alberta
- Teh, Yee-Whye, Oxford
- Urtasun, Raquel, U. of Toronto
- Varma, Manik, Microsoft
- Vasconcelos, Nuno, UCSD
- Vayatis, Nicolas, ENS Cachan
- Vert, Jean-Philippe, Paris Tech
- Wallach, Hanna, Microsoft Research
- Weinberger, Kilian, Washington U. in St. Louis
- Wiliamson, Sinead, U. of Texas
- Wingate, David, Lyric Labs
- Zhang, Nevin, Hong Kong University of Science and Technology
- Zhang, Tong, Rutgers
- Zheng, Alice, GraphLab
- Zhu, Jun, Tsinghua University
- Zhu, Xiaojin, University of Wisconsin

Keynote speakers

Léon Bottou (Facebook AI Research): Two high stakes challenges in machine learning

Tuesday July 7, 8:40 – 9:40 (Vauban)

Bio: Léon Bottou received the Diplôme d'Ingénieur de l'École Polytechnique (X84) in 1987, the Magistère de Mathématiques Fondamentales et Appliquées et d'Informatique from École Normale Supérieure in 1988, and a Ph.D. in Computer Science from Université de Paris-Sud in 1991. His research career took him to AT&T Bell Laboratories, AT&T Labs Research, NEC Labs America and Microsoft. He joined Facebook AI Research in 2015.

The long term goal of Léon's research is to understand how to build human-level intelligence. Although reaching this goal requires conceptual advances that cannot be anticipated at this point, it certainly entails clarifying how to learn and how to reason. Leon Bottou best known contributions are his work on "deep" neural networks in the 90s, his work on large scale learning in the 2000's, and possibly his more recent work on causal inference in learning systems. Léon is also known for the DjVu document compression technology.

Abstract: This presentation describes and discusses two serious challenges:

Machine learning technologies are increasingly used in complex software systems such as those underlying internet services today or self-driving vehicles tomorrow. Despite famous successes, there is more and more evidence that machine learning components tend to disrupt established software engineering practices. I will present examples and offer an explanation of this annoying and often very costly effect. Our first high-stake challenge consists therefore in formulating sound and efficient engineering principles for machine learning applications.

Machine learning research can often be viewed as an empirical science. Unlike nearly all other empirical sciences, progress in machine learning has largely been driven by a single experimental paradigm: fitting a training set and reporting performance on a testing set. Three forces may terminate this convenient state of affairs: the first one is the engineering challenge outlined above, the second one arises from the statistics of large-scale datasets, and the third one is our growing ambition to address more serious AI tasks. Our second high-stakes challenge consists therefore in enriching our experimental repertoire, redefining our scientific processes, and still maintain our progress speed.

Susan Murphy (University of Michigan): Learning Treatment Policies in Mobile Health

Wednesday July 8, 14:00 – 15:00 (Vauban)

Bio: Susan Murphy's research focuses on improving sequential, individualized, decision making in health, in particular on clinical trial design and data analysis to inform the development of adaptive interventions (e.g. treatment policies). She is a leading developer of the Sequential Multiple Assignment Randomized Trial (SMART) design which has been and is being used by clinical researchers to develop adaptive interventions in depression, alcoholism, treatment of ADHD, substance abuse, HIV treatment, obesity, diabetes, autism and drug court programs. She collaborates with clinical scientists, computer scientists and engineers and mentors young clinical scientists on developing adaptive interventions.

Susan is currently working as part of several interdisciplinary teams to develop clinical trial designs and learning algorithms to settings in which patient information is collected in real time (e.g. via smart phones or other wearable devices) and thus sequences of interventions can be individualized online. She is a Fellow of the College on Problems in Drug Dependence, a former editor of the *Annals of Statistics* and is a 2013 MacArthur Fellow.

Abstract: We describe a sequence of steps that facilitate effective learning of treatment policies in mobile health. These include a clinical trial with associated sample size calculator and data analytic methods. An off-policy Actor-Critic algorithm is developed for learning a treatment policy from this clinical trial data. Open problems abound in this area, including the development of a variety of online predictors of risk of health problems, missing data and disengagement.

Jon Kleinberg (Cornell University): Social Interaction in Global Networks Cornell University**Thursday July 9, 14:00 – 15:00 (Vauban)**

Bio: Jon Kleinberg is the Tisch University Professor in the Departments of Computer Science and Information Science at Cornell University. His research focuses on issues at the interface of networks and information, with an emphasis on the social and information networks that underpin the Web and other on-line media.

He is a member of the National Academy of Sciences, the National Academy of Engineering, and the American Academy of Arts and Sciences; and he is the recipient of research fellowships from the MacArthur, Packard, Simons, and Sloan Foundations, as well as awards including the Nevanlinna Prize, the Harvey Prize, the ACM SIGKDD Innovation Award, and the ACM-Infosys Foundation Award in the Computing Sciences.

Abstract: With an increasing amount of social interaction taking place in the digital domain, and often in public on-line settings, we are accumulating enormous amounts of data about phenomena that were once essentially invisible to us: the collective behavior and social interactions of hundreds of millions of people, recorded at unprecedented levels of scale and resolution. Analyzing this data computationally offers new insights into the design of on-line applications, as well as a new perspective on fundamental questions in the social sciences. We will review some of the basic issues around these developments; these include the problem of designing information systems in the presence of complex social feedback effects, and the emergence of a growing research interface between computing and the social sciences, facilitated by the availability of large new datasets on human interaction.

Awards

**Best paper award: Chinmay Hegde, Piotr Indyk, Ludwig Schmidt,
A Nearly-Linear Time Framework for Graph-Structured Sparsity**

Tuesday July 7, 9:40 – 10:00 (Vauban)

Abstract: We introduce a framework for sparsity structures defined via graphs. Our approach is flexible and generalizes several previously studied sparsity models. Moreover, we provide efficient projection algorithms for our sparsity model that run in nearly-linear time. In the context of sparse recovery, we show that our framework achieves an information-theoretically optimal sample complexity for a wide range of parameters. We complement our theoretical analysis with experiments demonstrating that our algorithms improve on prior work also in practice.

**Best paper award: Alina Beygelzimer, Satyen Kale, Haipeng Luo,
Optimal and Adaptive Algorithms for Online Boosting**

Wednesday July 8, 15:00 – 15:20 (Vauban)

Abstract: We study online boosting, the task of converting any weak online learner into a strong online learner. Based on a novel and natural definition of weak online learnability, we develop two online boosting algorithms. The first algorithm is an online version of boost-by-majority. By proving a matching lower bound, we show that this algorithm is essentially optimal in terms of the number of weak learners and the sample complexity needed to achieve a specified accuracy. The second algorithm is adaptive and parameter-free, albeit not optimal.

**Test of time award: Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, Greg Hullender,
Learning to Rank Using Gradient Descent**

Thursday July 9, 15:00 – 15:20 (Vauban)

Abstract: We investigate using gradient descent methods for learning ranking functions; we propose a simple probabilistic cost function, and we introduce RankNet, an implementation of these ideas using a neural network to model the underlying ranking function. We present test results on toy data and on data from a commercial internet search engine.

Tutorials

Tutorial 1: Natural Language Understanding: Foundations and State-of-the-Art

Percy Liang (Stanford University)

Monday July 7, 09:15–11:45 (Pasteur)

Building systems that can understand human language—being able to answer questions, follow instructions, carry on dialogues—has been a long-standing challenge since the early days of AI. Due to recent advances in machine learning, there is again renewed interest in taking on this formidable task. A major question is how one represents and learns the semantics (meaning) of natural language, to which there are only partial answers. The goal of this tutorial is (i) to describe the linguistic and statistical challenges that any system must address; and (ii) to describe the types of cutting edge approaches and the remaining open problems. Topics include distributional semantics (e.g., word vectors), frame semantics (e.g., semantic role labeling), model-theoretic semantics (e.g., semantic parsing), the role of context, grounding, neural networks, latent variables, and inference. The hope is that this unified presentation will clarify the landscape, and show that this is an exciting time for the machine learning community to engage in the problems in natural language understanding.

Tutorial 2: Policy Search: Methods and Applications

Gerhard Neumann (Technische Universität Darmstadt)

Jan Peters (Technische Universität Darmstadt & Max Planck Institute for Intelligent Systems, Tübingen)

Monday July 7, 09:15–11:45 (Eurotop)

Policy search is a subfield in reinforcement learning which focuses on finding good parameters for a given policy parametrization. It is well suited for robotics as it can cope with high-dimensional state and action spaces, one of the main challenges in robot learning. We review recent successes of both model-free and model-based policy search in robot learning.

Model-free policy search is a general approach to learn policies based on sampled trajectories. We classify model-free methods based on their policy evaluation strategy, policy update strategy, and exploration strategy and present a unified view on existing algorithms. Learning a policy is often easier than learning an accurate forward model, and, hence, model-free methods are more frequently used in practice. However, for each sampled trajectory, it is necessary to interact with the robot, which can be time consuming and challenging in practice. Model-based policy search addresses this problem by first learning a simulator of the robot's dynamics from data. Subsequently, the simulator generates trajectories that are used for policy learning. For both model-free and model-based policy search methods, we review their respective properties and their applicability to robotic systems.

Tutorial 3: Advances in Structured Prediction

Hal Daumé III (University of Maryland)

John Langford (Microsoft Research)

Monday July 7, 13:30–15:40 (Pasteur)

Structured prediction is the problem of making a joint set of decisions to optimize a joint loss. There are two families of algorithms for such problems: Graphical model approaches and learning to search approaches. Graphical models include Conditional Random Fields and Structured SVMs and are effective when writing down a graphical model and solving it is easy. Learning to search approaches, explicitly predict the joint set of decisions incrementally, conditioning on past and future decisions. Such models may be particularly useful when the dependencies between the predictions are complex, the loss is complex, or the construction of an explicit graphical model is impossible.

We will describe both approaches, with a deeper focus on the latter learning-to-search paradigm, which has less tutorial support. This paradigm has been gaining increasing traction over the past five years, making advances in natural language processing (dependency parsing, semantic parsing), robotics (grasping and path planning), social network analysis and computer vision (object segmentation).

Tutorial 4: Bayesian Time Series Modeling: Structured Representations for Scalability

Emily Fox (University of Washington)

Monday July 7, 13:30–15:40 (Eurotop)

Time series of increasing complexity are being collected in a variety of fields ranging from neuroscience, genomics, and environmental monitoring to e-commerce based on technologies and infrastructures previously unavailable. These datasets can be viewed either as providing a single, high-dimensional time series or as a massive collection of time series with intricate and possibly evolving relationships between them. For scalability, it is crucial to discover and exploit sparse dependencies between the data streams or dimensions. Such representational structures for independent data sources have been extensively explored in the machine learning community. However, in the conversation on big data, despite the importance and prevalence of time series, the question of how to analyze such data at scale has received limited attention and represents an area of research opportunities.

For these time series of interest, there are two key modeling components: the dynamic and relational models, and their interplay. In this tutorial, we will review some foundational time series models, including the hidden Markov model (HMM) and vector autoregressive (VAR) process. Such dynamical models and their extensions have proven useful in capturing complex dynamics of individual data streams such as human motion, speech, EEG recordings, and genome sequences. However, a focus of this tutorial will be on how to deploy scalable representational structures for capturing sparse dependencies between data streams. In particular, we consider clustering, directed and undirected graphical models, and low-dimensional embeddings in the context of time series. An emphasis is on learning such structure from the data. We will also provide some insights into new computational methods for performing efficient inference in large-scale time series.

Throughout the tutorial we will highlight Bayesian and Bayesian nonparametric approaches for learning and inference. Bayesian methods provide an attractive framework for examining complex data streams by naturally incorporating and propagating notions of uncertainty and enabling integration of heterogeneous data sources; the Bayesian nonparametric aspect allows the complexity of the dynamics and relational structure to adapt to the observed data.

Tutorial 5: Modern Convex Optimization Methods for Large-scale Empirical Risk Minimization

Peter Richtárik (University of Edinburgh)

Mark Schmidt (University of British Columbia)

Monday July 7, 16:10–18:20 (Pasteur)

This tutorial reviews recent advances in convex optimization for training (linear) predictors via (regularized) empirical risk minimization. We exclusively focus on practically efficient methods which are also equipped with complexity bounds confirming the suitability of the algorithms for solving huge-dimensional problems (a very large number of examples or a very large number of features).

The first part of the tutorial is dedicated to modern primal methods (belonging to the stochastic gradient descent variety), while the second part focuses on modern dual methods (belonging to the randomized coordinate ascent variety). While we make this distinction, there are very close links between the primal and dual methods, some of which will be highlighted. We shall also comment on mini-batch, parallel and distributed variants of the methods as this is an important consideration for applications involving big data.

Tutorial 6: Computational Social Science

Hanna Wallach (Microsoft Research & University of Massachusetts Amherst)

Monday July 7, 16:10–18:20 (Eurotop)

From interactions between friends, colleagues, or political leaders to the activities of corporate or government organizations, complex social processes underlie almost all human endeavor. The emerging field of computational social science is concerned with the development of new mathematical models and computational tools for understanding and reasoning about such processes from noisy, missing, or uncertain information. Computational social science is an inherently interdisciplinary area, situated at the intersection of computer science, statistics, and the social sciences, with researchers from traditionally disparate backgrounds working together to answer questions arising in sociology, political science, economics, public policy, journalism, and beyond.

In the first half of this tutorial, I will provide an overview of computational social science, emphasizing recent research that moves beyond the study of small-scale, static snapshots of networks, and onto nuanced, data-driven analyses of the structure, content, and dynamics of large-scale social processes. I will focus on commonalities of these social processes, as well as differences between the types of modeling tasks typically prioritized by computer scientists and social scientists. I will then discuss Bayesian latent variable modeling as a methodological framework for understanding and reasoning about complex social processes, and provide a brief overview of Bayesian inference.

In the second half of this tutorial, I will concentrate specifically on political science. I will discuss data sources, acquisition methods, and research questions, as well as the mathematical details of several models recently developed by the political methodology community. These models, which draw upon research in machine learning and natural language processing, not only serve as examples of the outstanding methodological work being done in the social sciences, but also demonstrate how ideas originally developed by computer scientists can be adapted and used to answer substantive questions that further our understanding of society.

Main conference

Each paper is presented orally during a session and in the poster session. Papers presented orally before Wednesday noon (sessions 1–5) are presented during the **Tuesday poster session**. Papers presented orally after Wednesday noon (sessions 6–11) are presented during the **Wednesday poster session**. Poster sessions are located in Lille Métropole.

Tuesday	Vauban	Pasteur	Eurotop	Matisse	Van Gogh	Turin
08:30-08:40	Welcome – Vauban					
08:40-09:40	Invited talk: Léon Bottou – Vauban					
09:40-10:00	Best paper award – Vauban					
10:00-10:30	Coffee break – Lille Métropole					
10:30-12:10	1A: Deep Learning I	1B: Distributed Optimization	1C: Bayesian Nonparametrics I	1D: Reinforcement Learning I	1E: Learning Theory I	1F: Supervised Learning I
12:10-14:10	Lunch – on your own					
14:10-16:10	2A: Bandit Learning	2B: Deep Learning Computations	2C: Causality	2D: Matrix Factorization	2E: Sparsity I	2F: Ranking Learning
16:10-16:40	Coffee break – Lille Métropole					
16:40-18:00	3A: Optimization I	3B: Structured Prediction I	3C: Reinforcement Learning II	3D: Gaussian Processes I	3E: Transfer Learning I	3F: Networks and Graphs I
18:00-22:00	Poster session (papers from sessions 1–5) – Lille Métropole					
Wednesday	Vauban	Pasteur	Eurotop	Matisse	Van Gogh	Turin
08:30-09:50	4A: Deep Learning II	4B: Networks and Graphs II	4C: Online Learning I	4D: Probabilistic Models I	4E: Learning Theory II	4F: Transfer Learning II
09:50-10:20	Coffee break – Lille Métropole					
10:20-12:00	5A: Deep Learning and Vision	5B: Bayesian Optimization	5C: Variational Inference	5D: Large Scale Learning	5E: Structured Prediction II	5F: Manifold Learning
12:00-14:00	Lunch – on your own					
14:00-15:00	Invited talk: Susan Murphy – Vauban					
15:00-15:20	Best paper award – Vauban					
15:30-16:30	6A: Gaussian Processes II	6B: Monte Carlo Methods	6C: Hashing	6D: Feature Selection I	6E: Kernel Methods	6F: Computational Advertising and Social Science
16:30-17:00	Coffee break – Lille Métropole					
17:00-18:00	7A: Optimization II	7B: Approximate Inference I	7C: Bayesian Nonparametrics II	7D: Time Series Analysis I	7E: Feature Selection II	7F: Unsupervised Learning I
18:00-22:00	Poster session (papers from sessions 6–11) – Lille Métropole					
Thursday	Vauban	Pasteur	Eurotop	Matisse	Van Gogh	Turin
08:30-09:50	8A: Optimization III	8B: Probabilistic Models II	8C: Natural Language Processing I	8D: Submodularity	8E: Online Learning II	8F: Unsupervised Learning II
09:50-10:20	Coffee break – Lille Métropole					
10:20-12:00	9A: Deep Learning III	9B: Privacy	9C: Topic Models	9D: Reinforcement Learning III	9E: Gaussian Processes III	9F: Supervised Learning II
12:00-14:00	Lunch – on your own					
14:00-15:00	Invited talk: Jon Kleinberg – Vauban					
15:00-15:20	Test of Time award – Vauban					
15:30-16:30	10A: Reinforcement Learning IV	10B: Optimization IV	10C: Bayesian Nonparametrics III	10D: Sparsity II	10E: Clustering I	10F: Vision
16:30-17:00	Coffee break – Lille Métropole					
17:00-18:00	11A: Sparse optimization	11B: Approximate Inference II	11C: Natural Language Processing II	11D: Time Series Analysis II	11E: Clustering II	11F: Learning Theory III
18:00-19:00					IMLS business meeting	
19:30-23:00	ICML banquet – Grand Palais					

Session 1: Tuesday 10:30-12:10Poster session: **Tuesday July 7, 18:00 – 22:00** (Lille Métropole)**1A: Deep Learning I (session chair: Aaron Courville)****Tuesday July 7, 10:30-12:10 (Vauban)**

Learning Program Embeddings to Propagate Feedback on Student Code.

Chris Piech, Jonathan Huang, Andy Nguyen, Mike Phulsuksombati, Mehran Sahami, Leonidas Guibas

BilBOWA: Fast Bilingual Distributed Representations without Word Alignments.

Stephan Gouws, Yoshua Bengio, Greg Corrado

Modeling Order in Neural Word Embeddings at Scale.

Andrew Trask, David Gilmore, Matthew Russell

Gated Feedback Recurrent Neural Networks.

Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, Yoshua Bengio

An Empirical Exploration of Recurrent Network Architectures.

*Rafal Jozefowicz, Wojciech Zaremba, Ilya Sutskever***1B: Distributed Optimization (session chair: Shai Shalev-Shwartz)****Tuesday July 7, 10:30-12:10 (Pasteur)**

PASSCoDe: Parallel ASynchronous Stochastic dual Co-ordinate Descent.

Cho-Jui Hsieh, Hsiang-Fu Yu, Inderjit Dhillon

An Asynchronous Distributed Proximal Gradient Method for Composite Convex Optimization.

Necdet Aybat, Zi Wang, Garud Iyengar

DiSCO: Distributed Optimization for Self-Concordant Empirical Loss.

Yuchen Zhang, Xiao Lin

Distributed Box-Constrained Quadratic Optimization for Dual Linear SVM.

Ching-Pei Lee, Dan Roth

Adding vs. Averaging in Distributed Primal-Dual Optimization.

*Chenxin Ma, Virginia Smith, Martin Jaggi, Michael Jordan, Peter Richtarik, Martin Takac***1C: Bayesian Nonparametrics I (session chair: Cedric Archambeau)****Tuesday July 7, 10:30-12:10 (Eurotop)**

A Bayesian nonparametric procedure for comparing algorithms.

Alessio Benavoli, Giorgio Corani, Francesca Mangili, Marco Zaffalon

Manifold-valued Dirichlet Processes.

Hyunwoo Kim, Jia Xu, Baba Vemuri, Vikas Singh

JUMP-Means: Small-Variance Asymptotics for Markov Jump Processes.

Jonathan Huggins, Karthik Narasimhan, Ardavan Saeedi, Vikash Mansinghka

Atomic Spatial Processes.

Sean Jewell, Neil Spencer, Alexandre Bouchard-Côté

DP-space: Bayesian Nonparametric Subspace Clustering with Small-variance Asymptotics.

*Yining Wang, Jun Zhu***1D: Reinforcement Learning I (session chair: Pascal Poupart)****Tuesday July 7, 10:30-12:10 (Matisse)**

Large-Scale Markov Decision Problems with KL Control Cost and its Application to Crowdsourcing.

Yasin Abbasi-Yadkori, Peter Bartlett, Xi Chen, Alan Malek

Off-policy Model-based Learning under Unknown Factored Dynamics.

*Assaf Hallak, Francois Schnitzler, Timothy Mann, Shie Mannor*On the Rate of Convergence and Error Bounds for LSTD(λ).*Manel Tagorti, Bruno Scherrer*

A Relative Exponential Weighing Algorithm for Adversarial Utility-based Dueling Bandits.

Pratik Gajane, Tanguy Urvoy, Fabrice Clérot

On TD(0) with function approximation: Concentration bounds and a centered variant with exponential convergence.

Nathaniel Korda, Prashanth La

1E: Learning Theory I (session chair: Nati Srebro)

Tuesday July 7, 10:30-12:10 (Van Gogh)

Entropy-Based Concentration Inequalities for Dependent Variables.

Liva Ralaivola, Massih-Reza Amini

The Ladder: A Reliable Leaderboard for Machine Learning Competitions.

Avrim Blum, Moritz Hardt

Theory of Dual-sparse Regularized Randomized Reduction.

Tianbao Yang, Lijun Zhang, Rong Jin, Shenghuo Zhu

A Theoretical Analysis of Metric Hypothesis Transfer Learning.

Michaël Perrot, Amaury Habrard

Optimizing Non-decomposable Performance Measures: A Tale of Two Classes.

Harikrishna Narasimhan, Purushottam Kar, Prateek Jain

1F: Supervised Learning I (session chair: Guillaume Bouchard)

Tuesday July 7, 10:30-12:10 (Turin)

Learning from Corrupted Binary Labels via Class-Probability Estimation.

Aditya Menon, Brendan Van Rooyen, Cheng Soon Ong, Bob Williamson

Multi-instance multi-label learning in the presence of novel class instances.

Anh Pham, Raviv Raich, Xiaoli Fern, Jesús Pérez Arriaga

Entropy evaluation based on confidence intervals of frequency estimates : Application to the learning of decision trees.

Mathieu Serrurier, Henri Prade

Support Matrix Machines.

Luo Luo, Yubo Xie, Zihua Zhang, Wu-Jun Li

Attribute Efficient Linear Regression with Distribution-Dependent Sampling.

Doron Kukliansky, Ohad Shamir

Session 2: Tuesday 14:10-16:10

Poster session: **Tuesday July 7, 18:00 – 22:00** (Lille Métropole)

2A: Bandit Learning (session chair: Olivier Cappé)

Tuesday July 7, 14:10-16:10 (Vauban)

Optimal Regret Analysis of Thompson Sampling in Stochastic Multi-armed Bandit Problem with Multiple Plays.

Junpei Komiyama, Junya Honda, Hiroshi Nakagawa

Simple regret for infinitely many armed bandits.

Alexandra Carpentier, Michal Valko

Efficient Learning in Large-Scale Combinatorial Semi-Bandits.

Zheng Wen, Branislav Kveton, Azin Ashkan

Cascading Bandits: Learning to Rank in the Cascade Model.

Branislav Kveton, Csaba Szepesvari, Zheng Wen, Azin Ashkan

Cheap Bandits.

Manjesh Hanawal, Venkatesh Saligrama, Michal Valko, Remi Munos

Qualitative Multi-Armed Bandits: A Quantile-Based Approach.

Balazs Szorenyi, Robert Busa-Fekete, Paul Weng, Eyke Hüllermeier

2B: Deep Learning Computations (session chair: Yoshua Bengio)

Tuesday July 7, 14:10-16:10 (Pasteur)

On Deep Multi-View Representation Learning.

Weiran Wang, Raman Arora, Karen Livescu, Jeff Bilmes

Scaling up Natural Gradient by Sparsely Factorizing the Inverse Fisher Matrix.

Roger Grosse, Ruslan Salakhudinov

Compressing Neural Networks with the Hashing Trick.

Wenlin Chen, James Wilson, Stephen Tyree, Kilian Weinberger, Yixin Chen

Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift.

Sergey Ioffe, Christian Szegedy

Optimizing Neural Networks with Kronecker-factored Approximate Curvature.

James Martens, Roger Grosse

Deep Learning with Limited Numerical Precision.

Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, Pritish Narayanan

2C: Causality (session chair: Emily Fox)

Tuesday July 7, 14:10-16:10 (Eurotop)

Removing systematic errors for exoplanet search via latent causes.

Bernhard Schölkopf, David Hogg, Dun Wang, Dan Foreman-Mackey, Dominik Janzing, Carl-Johann Simon-Gabriel, Jonas Peters

Towards a Learning Theory of Cause-Effect Inference.

David Lopez-Paz, Krikamol Muandet, Bernhard Schölkopf, Iliya Tolstikhin

Counterfactual Risk Minimization: Learning from Logged Bandit Feedback.

Adith Swaminathan, Thorsten Joachims

Telling cause from effect in deterministic linear dynamical systems.

Naji Shajarisales, Dominik Janzing, Bernhard Schoelkopf, Michel Besserve

Causal Inference by Identification of Vector Autoregressive Processes with Hidden Components.

Philipp Geiger, Kun Zhang, Bernhard Schoelkopf, Mingming Gong, Dominik Janzing

Discovering Temporal Causal Relations from Subsampled Data.

Mingming Gong, Kun Zhang, Bernhard Schoelkopf, Dacheng Tao, Philipp Geiger

2D: Matrix Factorization (session chair: Anima Anandkumar)

Tuesday July 7, 14:10-16:10 (Matisse)

Guaranteed Tensor Decomposition: A Moment Approach.

Gongguo Tang, Parikshit Shah

PU Learning for Matrix Completion.

Cho-Jui Hsieh, Nagarajan Natarajan, Inderjit Dhillon

Low-Rank Matrix Recovery from Row-and-Column Affine Measurements.

Or Zuk, Avishai Wagner

Intersecting Faces: Non-negative Matrix Factorization With New Guarantees.

Rong Ge, James Zou

CUR Algorithm for Partially Observed Matrices.

Miao Xu, Rong Jin, Zhi-Hua Zhou

An Explicit Sampling Dependent Spectral Error Bound for Column Subset Selection.

Tianbao Yang, Lijun Zhang, Rong Jin, Shenghuo Zhu

2E: Sparsity I (session chair: Guillaume Obozinski)

Tuesday July 7, 14:10-16:10 (Van Gogh)

Swept Approximate Message Passing for Sparse Estimation.

Andre Manoel, Florent Krzakala, Eric Tramel, Lenka Zdeborová

Towards a Lower Sample Complexity for Robust One-bit Compressed Sensing.

Rongda Zhu, Quanquan Gu

Sparse Subspace Clustering with Missing Entries.

Congyuan Yang, Daniel Robinson, Rene Vidal

Safe Subspace Screening for Nuclear Norm Regularized Least Squares Problems.

Qiang Zhou, Qi Zhao

Complete Dictionary Recovery Using Nonconvex Optimization.

Ju Sun, Qing Qu, John Wright

2F: Ranking Learning (session chair: Alekh Agarwal)

Tuesday July 7, 14:10-16:10 (Turin)

Ranking from Stochastic Pairwise Preferences: Recovering Condorcet Winners and Tournament Solution Sets at the Top.

Arun Rajkumar, Suprovat Ghoshal, Lek-Heng Lim, Shivani Agarwal

MRA-based Statistical Learning from Incomplete Rankings.

Eric Sibony, Stéphane Clemençon, Jérémie Jakubowicz

Preference Completion: Large-scale Collaborative Ranking from Pairwise Comparisons.

Dohyung Park, Joe Neeman, Jin Zhang, Sujay Sanghavi, Inderjit Dhillon

Distributional Rank Aggregation, and an Axiomatic Analysis.

Adarsh Prasad, Harsh Pareek, Pradeep Ravikumar

Spectral MLE: Top-K Rank Aggregation from Pairwise Comparisons.

Yuxin Chen, Changho Suh

Generalization error bounds for learning to rank: Does the length of document lists matter?.

Ambuj Tewari, Sougata Chaudhuri

Session 3: Tuesday 16:40-18:00

Poster session: **Tuesday July 7, 18:00 – 22:00** (Lille Métropole)

3A: Optimization I (session chair: Mark Schmidt)

Tuesday July 7, 16:40-18:00 (Vauban)

A Lower Bound for the Optimization of Finite Sums.

Alekh Agarwal, Leon Bottou

Stochastic Optimization with Importance Sampling for Regularized Loss Minimization.

Peilin Zhao, Tong Zhang

Stochastic Dual Coordinate Ascent with Adaptive Probabilities.

Dominik Csiba, Zheng Qu, Peter Richtarik

Global Convergence of Stochastic Gradient Descent for Some Non-convex Matrix Problems.

Christopher De Sa, Christopher Re, Kunle Olukotun

3B: Structured Prediction I (session chair: Simon Lacoste-Julien)

Tuesday July 7, 16:40-18:00 (Pasteur)

How Hard is Inference for Structured Prediction?.

Amir Globerson, Tim Roughgarden, David Sontag, Cafer Yildirim

Paired-Dual Learning for Fast Training of Latent Variable Hinge-Loss MRFs.

Stephen Bach, Bert Huang, Jordan Boyd-Graber, Lise Getoor

Learning Submodular Losses with the Lovasz Hinge.

Jiaqian Yu, Matthew Blaschko

PeakSeg: constrained optimal segmentation and supervised penalty learning for peak detection in count data.

Toby Hocking, Guillem Rigai, Guillaume Bourque

3C: Reinforcement Learning II (session chair: Pascal Poupart)

Tuesday July 7, 16:40-18:00 (Eurotop)

Non-Stationary Approximate Modified Policy Iteration.

Boris Lesner, Bruno Scherrer

Trust Region Policy Optimization.

John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, Philipp Moritz

Improved Regret Bounds for Undiscounted Continuous Reinforcement Learning.

K. Lakshmanan, Ronald Ortner, Daniil Ryabko

Approximate Dynamic Programming for Two-Player Zero-Sum Markov Games.

Julien Perolat, Bruno Scherrer, Bilal Piot, Olivier Pietquin

3D: Gaussian Processes I (session chair: Cedric Archambeau)

Tuesday July 7, 16:40-18:00 (Matisse)

Fast Kronecker Inference in Gaussian Processes with non-Gaussian Likelihoods.

Seth Flaxman, Andrew Wilson, Daniel Neill, Hannes Nickisch, Alex Smola

A Unifying Framework of Anytime Sparse Gaussian Process Regression Models with Stochastic Variational Inference for Big Data.

Trong Nghia Hoang, Quang Minh Hoang, Bryan Kian Hsiang Low

Controversy in mechanistic modelling with Gaussian processes.

Benn Macdonald, Catherine Higham, Dirk Husmeier

Kernel Interpolation for Scalable Structured Gaussian Processes (KISS-GP).

Andrew Wilson, Hannes Nickisch

3E: Transfer Learning I (session chair: Nati Srebro)

Tuesday July 7, 16:40-18:00 (Van Gogh)

Active Nearest Neighbors in Changing Environments.

Christopher Berlind, Ruth Urner

Asymmetric Transfer Learning with Deep Gaussian Processes.

Melih Kandemir

An embarrassingly simple approach to zero-shot learning.

Bernardino Romera-Paredes, Philip Torr

Safe Screening for Multi-Task Feature Learning with Multiple Data Matrices.

Jie Wang, Jieping Ye

3F: Networks and Graphs I (session chair: Florence d'Alché)

Tuesday July 7, 16:40-18:00 (Turin)

A Divide and Conquer Framework for Distributed Graph Clustering.

Wenzhuo Yang, Huan Xu

Scalable Model Selection for Large-Scale Factorial Relational Models.

Chunchen Liu, Lu Feng, Ryohei Fujimaki, Yusuke Muraoka

Consistent estimation of dynamic and multi-layer block models.

Qiuyi Han, Kevin Xu, Edoardo Airoldi

Community Detection Using Time-Dependent Personalized PageRank.

Haim Avron, Lior Horesh

Session 4: Wednesday 08:30-09:50

Poster session: **Tuesday July 7, 18:00 – 22:00** (Lille Métropole)

4A: Deep Learning II (session chair: Yoshua Bengio)

Wednesday July 8, 08:30-09:50 (Vauban)

Variational Generative Stochastic Networks with Collaborative Shaping.

Philip Bachman, Doina Precup

How Can Deep Rectifier Networks Achieve Linear Separability and Preserve Distances?.

Senjian An, Farid Boussaid, Mohammed Bennamoun

Unsupervised Domain Adaptation by Backpropagation.

Yaroslav Ganin, Victor Lempitsky

Learning Transferable Features with Deep Adaptation Networks.

Mingsheng Long, Yue Cao, Jianmin Wang, Michael Jordan

4B: Networks and Graphs II (session chair: Florence d'Alché)

Wednesday July 8, 08:30-09:50 (Pasteur)

Correlation Clustering in Data Streams.

KookJin Ahn, Graham Cormode, Sudipto Guha, Andrew McGregor, Anthony Wirth

An Aligned Subtree Kernel for Weighted Graphs.

Lu Bai, Luca Rossi, Zhihong Zhang, Edwin Hancock

HawkesTopic: A Joint Model for Network Inference and Topic Modeling from Text-Based Cascades.

Xinran He, Theodoros Rekatsinas, James Foulds, Lise Getoor, Yan Liu

Inferring Graphs from Cascades: A Sparse Recovery Framework.

Jean Pouget-Abadie, Thibaut Horel

4C: Online Learning I (session chair: Csaba Szepesvari)

Wednesday July 8, 08:30-09:50 (Eurotop)

Strongly Adaptive Online Learning.

Amit Daniely, Alon Gonen, Shai Shalev-Shwartz

Adaptive Belief Propagation.

Georgios Papachristoudis, John Fisher

The Hedge Algorithm on a Continuum.

Walid Krichene, Maximilian Balandat, Claire Tomlin, Alexandre Bayen

4D: Probabilistic Models I (session chair: David Sontag)

Wednesday July 8, 08:30-09:50 (Matisse)

Inference in a Partially Observed Queuing Model with Applications in Ecology.

Kevin Winner, Garrett Bernstein, Dan Sheldon

Non-Gaussian Discriminative Factor Models via the Max-Margin Rank-Likelihood.

Xin Yuan, Ricardo Henao, Ephraim Tsaliq, Raymond Langley, Lawrence Carin

On the Relationship between Sum-Product Networks and Bayesian Networks.

Han Zhao, Mazen Melibari, Pascal Poupart

Fixed-point algorithms for learning determinantal point processes.

Zelda Mariet, Suvrit Sra

4E: Learning Theory II (session chair: Nina Balcan)

Wednesday July 8, 08:30-09:50 (Van Gogh)

Convex Calibrated Surrogates for Hierarchical Classification.

Harish Ramaswamy, Ambuj Tewari, Shivani Agarwal

Risk and Regret of Hierarchical Bayesian Learners.

Jonathan Huggins, Josh Tenenbaum

Classification with Low Rank and Missing Data.

Elad Hazan, Roi Livni, Yishay Mansour

Surrogate Functions for Maximizing Precision at the Top.

Purushottam Kar, Harikrishna Narasimhan, Prateek Jain

4F: Transfer Learning II (session chair: Honglak Lee)

Wednesday July 8, 08:30-09:50 (Turin)

Non-Linear Cross-Domain Collaborative Filtering via Hyper-Structure Transfer.

Yan-Fu Liu, Cheng-Yu Hsu, Shan-Hung Wu

A Probabilistic Model for Dirty Multi-task Feature Selection.

Daniel Hernandez-Lobato, Jose Miguel Hernandez-Lobato, Zoubin Ghahramani

Convex Learning of Multiple Tasks and their Structure.

Carlo Ciliberto, Youssef Mroueh, Tomaso Poggio, Lorenzo Rosasco

Convex Formulation for Learning from Positive and Unlabeled Data.

Marthinus Du Plessis, Gang Niu, Masashi Sugiyama

Session 5: Wednesday 10:20-12:00

Poster session: **Tuesday July 7, 18:00 – 22:00** (Lille Métropole)

5A: Deep Learning and Vision (session chair: Honglak Lee)**Wednesday July 8, 10:20-12:00 (Vauban)**

Unsupervised Learning of Video Representations using LSTMs.

Nitish Srivastava, Elman Mansimov, Ruslan Salakhudinov

Deep Edge-Aware Filters.

Li Xu, Jimmy Ren, Qiong Yan, Renjie Liao, Jiaya Jia

DRAW: A Recurrent Neural Network For Image Generation.

Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Rezende, Daan Wierstra

Show, Attend and Tell: Neural Image Caption Generation with Visual Attention.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, Yoshua Bengio

Online Tracking by Learning Discriminative Saliency Map with Convolutional Neural Network.

Seunghoon Hong, Tackgeun You, Suha Kwak, Bohyung Han

5B: Bayesian Optimization (session chair: Emily Fox)**Wednesday July 8, 10:20-12:00 (Pasteur)**

Gradient-based Hyperparameter Optimization through Reversible Learning.

Dougal Maclaurin, David Duvenaud, Ryan Adams

Scalable Bayesian Optimization Using Deep Neural Networks.

Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, Mr Prabhat, Ryan Adams

High Dimensional Bayesian Optimisation and Bandits via Additive Models.

Kirthevasan Kandasamy, Jeff Schneider, Barnabas Poczos

Safe Exploration for Optimization with Gaussian Processes.

Yanan Sui, Alkis Gotovos, Joel Burdick, Andreas Krause

Predictive Entropy Search for Bayesian Optimization with Unknown Constraints.

Jose Miguel Hernandez-Lobato, Michael Gelbart, Matthew Hoffman, Ryan Adams, Zoubin Ghahramani

5C: Variational Inference (session chair: David Blei)**Wednesday July 8, 10:20-12:00 (Eurotop)**

A trust-region method for stochastic variational inference with applications to streaming data.

Lucas Theis, Matt Hoffman

Variational Inference with Normalizing Flows.

Danilo Rezende, Shakir Mohamed

Markov Chain Monte Carlo and Variational Inference: Bridging the Gap.

Tim Salimans, Diederik Kingma, Max Welling

An Empirical Study of Stochastic Variational Inference Algorithms for the Beta Bernoulli Process.

Amar Shah, David Knowles, Zoubin Ghahramani

Scalable Variational Inference in Log-supermodular Models.

Josip Djolonga, Andreas Krause

5D: Large Scale Learning (session chair: Zaid Harchaoui)**Wednesday July 8, 10:20-12:00 (Matisse)**

Coresets for Nonparametric Estimation - the Case of DP-Means.

Olivier Bachem, Mario Lucic, Andreas Krause

K-hyperplane Hinge-Minimax Classifier.

Margarita Osadchy, Tamir Hazan, Daniel Keren

Faster cover trees.

Mike Izbicki, Christian Shelton

Finding Linear Structure in Large Datasets with Scalable Canonical Correlation Analysis.

Zhuang Ma, Yichao Lu, Dean Foster

Large-scale log-determinant computation through stochastic Chebyshev expansions.

Insu Han, Dmitry Malioutov, Jinwoo Shin

5E: Structured Prediction II (session chair: Simon Lacoste-Julien) Wednesday July 8, 10:20-12:00 (Van Gogh)

Algorithms for the Hard Pre-Image Problem of String Kernels and the General Problem of String Prediction.

Sébastien Giguère, Amélie Rolland, Francois Laviolette, Mario Marchand

Learning Scale-Free Networks by Dynamic Node Specific Degree Prior.

Qingming Tang, Siqi Sun, Jinbo Xu

Learning Deep Structured Models.

Liang-Chieh Chen, Alexander Schwing, Alan Yuille, Raquel Urtasun

Learning to Search Better than Your Teacher.

Kai-Wei Chang, Akshay Krishnamurthy, Alekh Agarwal, Hal Daume, John Langford

Learning Fast-Mixing Models for Structured Prediction.

Jacob Steinhardt, Percy Liang

5F: Manifold Learning (session chair: Kilian Weinberger) Wednesday July 8, 10:20-12:00 (Turin)

Unsupervised Riemannian Metric Learning for Histograms Using Aitchison Transformations.

Tam Le, Marco Cuturi

Learning Local Invariant Mahalanobis Distances.

Ethan Fetaya, Shimon Ullman

Landmarking Manifolds with Gaussian Processes.

Dawen Liang, John Paisley

Subsampling Methods for Persistent Homology.

Frederic Chazal, Brittany Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, Larry Wasserman

Bipartite Edge Prediction via Transductive Learning over Product Graphs.

Hanxiao Liu, Yiming Yang

Session 6: Wednesday 15:30-16:30

Poster session: **Wednesday July 8, 18:00 – 22:00** (Lille Métropole)

6A: Gaussian Processes II (session chair: Raquel Urtasun) Wednesday July 8, 15:30-16:30 (Vauban)

Latent Gaussian Processes for Distribution Estimation of Multivariate Categorical Data.

Yarin Gal, Yutian Chen, Zoubin Ghahramani

Scalable Nonparametric Bayesian Inference on Point Processes with Gaussian Processes.

Yves-Laurent Kom Samo, Stephen Roberts

Finding Galaxies in the Shadows of Quasars with Gaussian Processes.

Roman Garnett, Shirley Ho, Jeff Schneider

6B: Monte Carlo Methods (session chair: Olivier Cappé) Wednesday July 8, 15:30-16:30 (Pasteur)

Exponential Integration for Hamiltonian Monte Carlo.

Wei-Lun Chao, Justin Solomon, Dominik Michels, Fei Sha

Nested Sequential Monte Carlo Methods.

Christian Naesseth, Fredrik Lindsten, Thomas Schon

The Fundamental Incompatibility of Scalable Hamiltonian Monte Carlo and Naive Data Subsampling.

Michael Betancourt

6C: Hashing (session chair: Manik Varma)**Wednesday July 8, 15:30-16:30 (Eurotop)**

Binary Embedding: Fundamental Limits and Fast Algorithm.

Xinyang Yi, Constantine Caramanis, Eric Price

Hashing for Distributed Data.

Cong Leng, Jiaxiang Wu, Jian Cheng, Xi Zhang, Hanqing Lu

On Symmetric and Asymmetric LSHs for Inner Product Search.

*Behnam Neyshabur, Nathan Srebro***6D: Feature Selection I (session chair: Kilian Weinberger)****Wednesday July 8, 15:30-16:30 (Matisse)**

Multiview Triplet Embedding: Learning Attributes in Multiple Maps.

Ehsan Amid, Antti Ukkonen

Low Rank Approximation using Error Correcting Coding Matrices.

Shashanka Ubaru, Arya Mazumdar, Yousef Saad

A Unified Framework for Outlier-Robust PCA-like Algorithms.

*Wenzhuo Yang, Huan Xu***6E: Kernel Methods (session chair: Julien Mairal)****Wednesday July 8, 15:30-16:30 (Van Gogh)**

A low variance consistent test of relative dependency.

Wacha Bounliphone, Arthur Gretton, Arthur Tenenhaus, Matthew Blaschko

Double Nyström Method: An Efficient and Accurate Nyström Scheme for Large-Scale Data Sets.

Woosang Lim, Minhwan Kim, Haesun Park, Kyomin Jung

The Kendall and Mallows Kernels for Permutations.

*Yunlong Jiao, Jean-Philippe Vert***6F: Computational Advertising and Social Science (session chair: Lihong Li) (Turin)****Wednesday July 8, 15:30-16:30**

Threshold Influence Model for Allocating Advertising Budgets.

Atsushi Miyachi, Yuni Iwamasa, Takuro Fukunaga, Naonori Kakimura

Approval Voting and Incentives in Crowdsourcing.

Nihar Shah, Dengyong Zhou, Yuval Peres

Budget Allocation Problem with Multiple Advertisers: A Game Theoretic View.

*Takanori Maehara, Akihiro Yabe, Ken-ichi Kawarabayashi***Session 7: Wednesday 17:00-18:00**Poster session: **Wednesday July 8, 18:00 – 22:00** (Lille Métropole)**7A: Optimization II (session chair: Shai Shalev-Shwartz)****Wednesday July 8, 17:00-18:00 (Vauban)**

Coordinate Descent Converges Faster with the Gauss-Southwell Rule Than Random Selection.

Julie Nutini, Mark Schmidt, Issam Laradji, Michael Friedlander, Hoyt Koepke

A Stochastic PCA and SVD Algorithm with an Exponential Convergence Rate.

Ohad Shamir

Online Learning of Eigenvectors.

*Dan Garber, Elad Hazan, Tengyu Ma***7B: Approximate Inference I (session chair: John Paisley)****Wednesday July 8, 17:00-18:00 (Pasteur)**

A Hybrid Approach for Probabilistic Inference using Random Projections.

Michael Zhu, Stefano Ermon

Reified Context Models.

Jacob Steinhardt, Percy Liang

Rebuilding Factorized Information Criterion: Asymptotically Accurate Marginal Likelihood.

Kohei Hayashi, Shin-ichi Maeda, Ryohei Fujimaki

7C: Bayesian Nonparametrics II (session chair: Guillaume Bouchard) Wednesday July 8, 17:00-18:00 (Eurotop)

Ordered Stick-Breaking Prior for Sequential MCMC Inference of Bayesian Nonparametric Models.

Mrinal Das, Trapit Bansal, Chiranjib Bhattacharyya

Bayesian and Empirical Bayesian Forests.

Taddy Matthew, Chun-Sheng Chen, Jun Yu, Mitch Wyle

Metadata Dependent Mondrian Processes.

Yi Wang, Bin Li, Yang Wang, Fang Chen

7D: Time Series Analysis I (session chair: Nicolas Le Roux) Wednesday July 8, 17:00-18:00 (Matisse)

Accelerated Online Low Rank Tensor Learning for Multivariate Spatiotemporal Streams.

Rose Yu, Dehua Cheng, Yan Liu

A Multitask Point Process Predictive Model.

Wenzhao Lian, Ricardo Henao, Vinayak Rao, Joseph Lucas, Lawrence Carin

Moderated and Drifting Linear Dynamical Systems.

Jinyan Guan, Kyle Simek, Ernesto Brau, Clayton Morrison, Emily Butler, Kobus Barnard

7E: Feature Selection II (session chair: Marco Cuturi) Wednesday July 8, 17:00-18:00 (Van Gogh)

Streaming Sparse Principal Component Analysis.

Wenzhuo Yang, Huan Xu

Pushing the Limits of Affine Rank Minimization by Adapting Probabilistic PCA.

Bo Xin, David Wipf

Statistical and Algorithmic Perspectives on Randomized Sketching for Ordinary Least-Squares.

Garvesh Raskutti, Michael Mahoney

7F: Unsupervised Learning I (session chair: Aaron Courville) Wednesday July 8, 17:00-18:00 (Turin)

Deep Unsupervised Learning using Nonequilibrium Thermodynamics.

Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, Surya Ganguli

Structural Maxent Models.

Corinna Cortes, Vitaly Kuznetsov, Mehryar Mohri, Umar Syed

Hidden Markov Anomaly Detection.

Nico Goernitz, Mikio Braun, Marius Kloft

Session 8: Thursday 08:30-09:50

Poster session: **Wednesday July 8, 18:00 – 22:00** (Lille Métropole)

8A: Optimization III (session chair: Julien Mairal) Thursday July 8, 08:30-09:50 (Vauban)

Faster Rates for the Frank-Wolfe Method over Strongly-Convex Sets.

Dan Garber, Elad Hazan

$\ell_{1,p}$ -Norm Regularization: Error Bounds and Convergence Rate Analysis of First-Order Methods.

Zirui Zhou, Qi Zhang, Anthony Man-Cho So

Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization.

Roy Frostig, Rong Ge, Sham Kakade, Aaron Sidford

Blitz: A Principled Meta-Algorithm for Scaling Sparse Optimization.

Tyler Johnson, Carlos Guestrin

8B: Probabilistic Models II (session chair: David Blei)

Thursday July 8, 08:30-09:50 (Pasteur)

Harmonic Exponential Families on Manifolds.

Taco Cohen, Max Welling

Celeste: Variational inference for a generative model of astronomical images.

Jeffrey Regier, Andrew Miller, Jon McAuliffe, Ryan Adams, Matt Hoffman, Dustin Lang, David Schlegel, Prabhat

Vector-Space Markov Random Fields via Exponential Families.

Wesley Tansey, Oscar Hernan Madrid Padilla, Arun Sai Suggala, Pradeep Ravikumar

Dealing with small data: On the generalization of context trees.

Ralf Eggeling, Mikko Koivisto, Ivo Grosse

8C: Natural Language Processing I (session chair: Percy Liang)

Thursday July 8, 08:30-09:50 (Eurotop)

Learning Word Representations with Hierarchical Sparse Coding.

Dani Yogatama, Manaal Faruqui, Chris Dyer, Noah Smith

A Linear Dynamical System Model for Text.

David Belanger, Sham Kakade

From Word Embeddings To Document Distances.

Matt Kusner, Yu Sun, Nicholas Kolkin, Kilian Weinberger

A Fast Variational Approach for Learning Markov Random Field Language Models.

Yacine Jernite, Alexander Rush, David Sontag

8D: Submodularity (session chair: Andreas Krause)

Thursday July 8, 08:30-09:50 (Matisse)

On Greedy Maximization of Entropy.

Dravyansh Sharma, Ashish Kapoor, Amit Deshpande

Submodularity in Data Subset Selection and Active Learning.

Kai Wei, Rishabh Iyer, Jeff Bilmes

Random Coordinate Descent Methods for Minimizing Decomposable Submodular Functions.

Alina Ene, Huy Nguyen

The Power of Randomization: Distributed Submodular Maximization on Massive Datasets.

Rafael Barbosa, Alina Ene, Huy Nguyen, Justin Ward

8E: Online Learning II (session chair: Zaid Harchaoui)

Thursday July 8, 08:30-09:50 (Van Gogh)

An Online Learning Algorithm for Bilinear Models.

Yuanbin Wu, Shiliang Sun

Online Time Series Prediction with Missing Data.

Oren Anava, Elad Hazan, Assaf Zeevi

On Identifying Good Options under Combinatorially Structured Feedback in Finite Noisy Environments.

Yifan Wu, Andras Gyorgy, Csaba Szepesvari

Following the Perturbed Leader for Online Structured Learning.

Alon Cohen, Tamir Hazan

8F: Unsupervised Learning II (session chair: Marco Cuturi)

Thursday July 8, 08:30-09:50 (Turin)

Entropic Graph-based Posterior Regularization.

Maxwell Libbrecht, Michael Hoffman, Jeff Bilmes, William Noble

Information Geometry and Minimum Description Length Networks.

Ke Sun, Jun Wang, Alexandros Kalousis, Stephan Marchand-Maillet

Context-based Unsupervised Data Fusion for Decision Making.

Erfan Soltanmohammadi, Mort Naraghi-Pour, Mihaela van der Schaar

Alpha-Beta Divergences Discover Micro and Macro Structures in Data.

Karthik Narayan, Ali Punjani, Pieter Abbeel

Session 9: Thursday 10:20-12:00

Poster session: **Wednesday July 8, 18:00 – 22:00** (Lille Métropole)

9A: Deep Learning III (session chair: **Samy Bengio**)

Thursday July 8, 10:20-12:00 (Vauban)

Weight Uncertainty in Neural Network.

Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, Daan Wierstra

MADE: Masked Autoencoder for Distribution Estimation.

Mathieu Germain, Karol Gregor, Iain Murray, Hugo Larochelle

Generative Moment Matching Networks.

Yujia Li, Kevin Swersky, Rich Zemel

Probabilistic Backpropagation for Scalable Learning of Bayesian Neural Networks.

Jose Miguel Hernandez-Lobato, Ryan Adams

Boosted Categorical Restricted Boltzmann Machine for Computational Prediction of Splice Junctions.

Taehoon Lee, Sungroh Yoon

9B: Privacy (session chair: **Alekh Agarwal**)

Thursday July 8, 10:20-12:00 (Pasteur)

Is Feature Selection Secure against Training Data Poisoning?.

Huang Xiao, Battista Biggio, Gavin Brown, Giorgio Fumera, Claudia Eckert, Fabio Roli

Privacy for Free: Posterior Sampling and Stochastic Gradient Monte Carlo.

Yu-Xiang Wang, Stephen Fienberg, Alex Smola

Rademacher Observations, Private Data, and Boosting.

Richard Nock, Giorgio Patrini, Arik Friedman

The Composition Theorem for Differential Privacy.

Peter Kairouz, Sewoong Oh, Pramod Viswanath

Differentially Private Bayesian Optimization.

Matt Kusner, Jacob Gardner, Roman Garnett, Kilian Weinberger

9C: Topic Models (session chair: **Anima Anandkumar**)

Thursday July 8, 10:20-12:00 (Eurotop)

Markov Mixed Membership Models.

Aonan Zhang, John Paisley

Latent Topic Networks: A Versatile Probabilistic Programming Framework for Topic Models.

James Foulds, Shachi Kumar, Lise Getoor

Scalable Deep Poisson Factor Analysis for Topic Modeling.

Zhe Gan, Changyou Chen, Ricardo Henao, David Carlson, Lawrence Carin

Ordinal Mixed Membership Models.

Seppo Virtanen, Mark Girolami

Efficient Training of LDA on a GPU by Mean-for-Mode Estimation.

Jean-Baptiste Tristan, Joseph Tassarotti, Guy Steele

9D: Reinforcement Learning III (session chair: **Lihong Li**)

Thursday July 8, 10:20-12:00 (Matisse)

Training Deep Convolutional Neural Networks to Play Go.

Christopher Clark, Amos Storkey

Safe Policy Search for Lifelong Reinforcement Learning with Sublinear Regret.

Haitham Bou Ammar, Rasul Tutunov, Eric Eaton

Robust partially observable Markov decision process.

Takayuki Osogami

Fictitious Self-Play in Extensive-Form Games.

Johannes Heinrich, Marc Lanctot, David Silver

A Deeper Look at Planning as Learning from Replay.

Harm Vanseljen, Rich Sutton

9E: Gaussian Processes III (session chair: Andreas Krause)

Thursday July 8, 10:20-12:00 (Van Gogh)

Variational Inference for Gaussian Process Modulated Poisson Processes.

Chris Lloyd, Tom Gunter, Michael Osborne, Stephen Roberts

Distributed Gaussian Processes.

Marc Deisenroth, Jun Wei Ng

Enabling scalable stochastic gradient-based inference for Gaussian processes by employing the Unbiased Linear System Solver (ULISSE).

Maurizio Filippone, Raphael Engler

Improving the Gaussian Process Sparse Spectrum Approximation by Representing Uncertainty in Frequency Inputs.

Yarin Gal, Richard Turner

Sparse Variational Inference for Generalized GP Models.

Rishit Sheth, Yuyang Wang, Roni Khardon

9F: Supervised Learning II (session chair: Nicolas Le Roux)

Thursday July 8, 10:20-12:00 (Turin)

On the Optimality of Multi-Label Classification under Subset Zero-One Loss for Distributions Satisfying the Composition Property.

Maxime Gasse, Alexandre Aussem, Haytham Elghazel

Consistent Multiclass Algorithms for Complex Performance Measures.

Harikrishna Narasimhan, Harish Ramaswamy, Aadirupa Saha, Shivani Agarwal

Feature-Budgeted Random Forest.

Feng Nan, Joseph Wang, Venkatesh Saligrama

A New Generalized Error Path Algorithm for Model Selection.

Bin Gu, Charles Ling

Dynamic Sensing: Better Classification under Acquisition Constraints.

Oran Richman, Shie Mannor

Session 10: Thursday 15:30-16:30

Poster session: **Wednesday July 8, 18:00 – 22:00** (Lille Métropole)

10A: Reinforcement Learning IV (session chair: Csaba Szepesvari)

Thursday July 8, 15:30-16:30 (Vauban)

Universal Value Function Approximators.

Tom Schaul, Daniel Horgan, Karol Gregor, David Silver

High Confidence Policy Improvement.

Philip Thomas, Georgios Theodorou, Mohammad Ghavamzadeh

Abstraction Selection in Model-based Reinforcement Learning.

Nan Jiang, Alex Kulesza, Satinder Singh

10B: Optimization IV (session chair: Mark Schmidt)

Thursday July 8, 15:30-16:30 (Pasteur)

Stochastic Primal-Dual Coordinate Method for Regularized Empirical Risk Minimization.

Yuchen Zhang, Xiao Lin

Adaptive Stochastic Alternating Direction Method of Multipliers.

Peilin Zhao, Jinwei Yang, Tong Zhang, Ping Li

A General Analysis of the Convergence of ADMM.

Robert Nishihara, Laurent Lessard, Ben Recht, Andrew Packard, Michael Jordan

10C: Bayesian Nonparametrics III (session chair: Jon McAuliffe)

Thursday July 8, 15:30-16:30 (Eurotop)

A Convex Exemplar-based Approach to MAD-Bayes Dirichlet Process Mixture Models.

En-Hsu Yen, Xin Lin, Kai Zhong, Pradeep Ravikumar, Inderjit Dhillon

Distributed Inference for Dirichlet Process Mixture Models.

Hong Ge, Yutian Chen, Moquan Wan, Zoubin Ghahramani

Large-scale Distributed Dependent Nonparametric Trees.

Zhiting Hu, Ho Qirong, Avinava Dubey, Eric Xing

10D: Sparsity II (session chair: Francis Bach)

Thursday July 8, 15:30-16:30 (Matisse)

A Deterministic Analysis of Noisy Sparse Subspace Clustering for Dimensionality-reduced Data.

Yining Wang, Yu-Xiang Wang, Aarti Singh

Stay on path: PCA along graph paths.

Megasthenis Asteris, Anastasios Kyrillidis, Alex Dimakis, Han-Gyol Yi, Bharath Chandrasekaran

Geometric Conditions for Subspace-Sparse Recovery.

Chong You, Rene Vidal

10E: Clustering I (session chair: Raquel Urtasun)

Thursday July 8, 15:30-16:30 (Van Gogh)

A Convex Optimization Framework for Bi-Clustering.

Shiau Hong Lim, Yudong Chen, Huan Xu

Multi-Task Learning for Subspace Segmentation.

Yu Wang, David Wipf, Qing Ling, Wei Chen, Ian Wassell

Multi-view Sparse Co-clustering via Proximal Alternating Linearized Minimization.

Jiangwen Sun, Jin Lu, Tingyang Xu, Jinbo Bi

10F: Vision (session chair: Manik Varma)

Thursday July 8, 15:30-16:30 (Turin)

Complex Event Detection using Semantic Saliency and Nearly-Isotonic SVM.

Xiaojun Chang, Yi Yang, Eric Xing, Yaoliang Yu

Log-Euclidean Metric Learning on Symmetric Positive Definite Manifold with Application to Image Set Classification.

Zhiwu Huang, Ruiping Wang, Shiguang Shan, Xianqiu Li, Xilin Chen

Bayesian Multiple Target Localization.

Purnima Rajan, Weidong Han, Raphael Sznitman, Peter Frazier, Bruno Jedynak

Session 11: Thursday 17:00-18:00

Poster session: **Wednesday July 8, 18:00 – 22:00** (Lille Métropole)

11A: Sparse optimization (session chair: Guillaume Obozinski)

Thursday July 8, 17:00-18:00 (Vauban)

Mind the duality gap: safer rules for the Lasso.

Olivier Fercoq, Alexandre Gramfort, Joseph Salmon

Tracking Approximate Solutions of Parameterized Optimization Problems over Multi-Dimensional (Hyper-)Parameter Domains.

Katharina Blechschmidt, Joachim Giesen, Soeren Laue

A Modified Orthant-Wise Limited Memory Quasi-Newton Method with Convergence Analysis.

Pinghua Gong, Jieping Ye

11B: Approximate Inference II (session chair: David Sontag)

Thursday July 8, 17:00-18:00 (Pasteur)

The Benefits of Learning with Strongly Convex Approximate Inference.

Ben London, Bert Huang, Lise Getoor

Message Passing for Collective Graphical Models.

Tao Sun, Dan Sheldon, Akshat Kumar

Proteins, Particles, and Pseudo-Max-Marginals: A Submodular Approach.

Jason Pacheco, Erik Sudderth

11C: Natural Language Processing II (session chair: Samy Bengio)

Thursday July 8, 17:00-18:00 (Eurotop)

Phrase-based Image Captioning.

Remi Lebre, Pedro Pinheiro, Ronan Collobert

Long Short-Term Memory Over Recursive Structures.

Xiaodan Zhu, Parinaz Sobihani, Hongyu Guo

Bimodal Modelling of Source Code and Natural Language.

Miltos Allamanis, Daniel Tarlow, Andrew Gordon, Yi Wei

11D: Time Series Analysis II (session chair: Percy Liang)

Thursday July 8, 17:00-18:00 (Matisse)

Learning Parametric-Output HMMs with Two Aliased States.

Roi Weiss, Boaz Nadler

Robust Estimation of Transition Matrices in High Dimensional Heavy-tailed Vector Autoregressive Processes.

Huitong Qiu, Sheng Xu, Fang Han, Han Liu, Brian Caffo

Functional Subspace Clustering with Application to Time Series.

Mohammad Taha Bahadori, David Kale, Yingying Fan, Yan Liu

11E: Clustering II (session chair: Francis Bach)

Thursday July 8, 17:00-18:00 (Van Gogh)

Yinyang K-Means: A Drop-In Replacement of the Classic K-Means with Consistent Speedup.

Yufei Ding, Yue Zhao, Xipeng Shen, Madanlal Musuvathi, Todd Mytkowicz

A Provable Generalized Tensor Spectral Method for Uniform Hypergraph Partitioning.

Debarghya Ghoshdastidar, Ambedkar Dukkipati

Spectral Clustering via the Power Method - Provably.

Christos Boutsidis, Prabhajan Kambadur, Alex Gittens

11F: Learning Theory III (session chair: Nina Balcan)

Thursday July 8, 17:00-18:00 (Turin)

Distributed Estimation of Generalized Matrix Rank: Efficient Algorithms and Lower Bounds.

Yuchen Zhang, Martin Wainwright, Michael Jordan

Deterministic Independent Component Analysis.

Ruitong Huang, Andras Gyorgy, Csaba Szepesvári

Convergence rate of Bayesian tensor estimator and its minimax optimality.

Taiji Suzuki

Workshops

Workshop on Deep Learning

Friday and Saturday (Pasteur)

Organizers:

- Geoff Hinton (Google)
- Yann LeCun (Facebook)
- Yoshua Bengio (Université de Montréal)
- Max Welling (University of Amsterdam)
- Kyunghyun Cho (Université de Montréal)
- Durk Kingma (University of Amsterdam)

Abstract:

Deep learning is a fast-growing field of Machine Learning concerned with the study and design of computer algorithms for learning good representations of data, at multiple levels of abstraction. There has been rapid progress in this area in recent years, both in terms of methods and in terms of applications, which are attracting the major IT companies. Many challenges remain, however, in aspects like large-scale (hyper-) parameter optimization, modeling of temporal data with long-term dependencies, generative modeling, efficient Bayesian inference for deep learning, multi-modal data and models, and learning representations for reinforcement learning. The workshop aims at bringing together researchers in the field of deep learning to discuss recent advances, ongoing developments and the road that lies ahead.

Invited speakers:

- Tara Sainath, Google
- Yann Ollivier, Paris-Sud University
- Oriol Vinyals, Google
- Jason Weston, Facebook
- Jorge Nocedal, Northwestern University
- Neil Lawrence, Sheffield University
- Roland Memisevic, University of Montreal
- Rajesh Ranganath, Princeton University
- Ian Goodfellow, Google
- Karol Gregor, Google DeepMind

Website:

<https://sites.google.com/site/deeplearning2015/>

European Workshop on Reinforcement Learning (EWRL)

Friday and Saturday (Matisse)

Organizers:

- Alessandro Lazaric (INRIA)
- Mohammad Ghavamzadeh (Adobe Research / INRIA)
- Rémi Munos (Google DeepMind / INRIA)

Abstract:

Reinforcement learning's (RL) objective is to develop agents able to learn optimal policies in unknown environments by trial-and-error and with limited supervision. Recent developments in exploration-exploitation, online learning, and representation learning are making RL more and more appealing to real-world applications, with promising results in challenging domains such as recommendation systems, computer games, and robotics. The 12th edition of EWRL will serve as a forum to discuss the state-of-the-art and future research directions and opportunities for the growing field of RL. Beyond traditional topics, we will encourage discussions on representation learning, risk-averse learning, apprenticeship and transfer learning, and practical applications.

Invited speakers:

- Marcus Hutter, Australian National University, Canberra, Australia
- Thomas G. Diettrich, Oregon State University, Corvallis, Oregon, USA
- David Silver, Google Deep Mind, London, UK
- Lihong Li, Microsoft Research, USA
- Shie Mannor, Technion, Israel
- Csaba Szepesvari, University of Alberta, Canada

Website:

<https://ewrl.wordpress.com/ewrl12-2015/>

Advances in Active Learning: Bridging Theory and Practice

Friday (Van Gogh)

Organizers:

- Akshay Krishnamurthy (Carnegie Mellon University)
- Aaditya Ramdas (Carnegie Mellon University)
- Nina Balcan (Carnegie Mellon University)
- Aarti Singh (Carnegie Mellon University)

Abstract:

Active learning has been a topic of significant research over the past several decades with much attention devoted to both theoretical and practical considerations. A variety of algorithms and sampling paradigms have been proposed and studied, but roughly speaking, this line of research focuses on how to make feedback driven decisions about data collection, and how to leverage this power for efficient learning. In this workshop, we hope to find future research directions that address the disconnect between active learning theory and practice.

Invited speakers:

- Steve Hanneke
- Adam Kalai (Microsoft Research)
- Andreas Krause (ETH Zurich)
- John Langford (Microsoft Research)
- Maja Temerinac-Ott (Univ. of Freiburg)
- Jeff Schneider (Carnegie Mellon University)
- Burr Settles (Duolingo, FAWM)

Website:

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

Constructive Machine Learning

Friday (Rubens)

Organizers:

- Thomas Gärtner (University of Bonn and Fraunhofer IAIS)
- Andrea Passerini (University of Trento)
- Roman Garnett (Washington University, St. Louis)
- Fabrizio Costa (University of Freiburg)

Abstract:

Constructive machine learning describes a class of machine learning problems where the ultimate goal is not finding a good model of the data but rather one or more particular instances of the domain which are likely to exhibit desired properties. While traditional approaches choose these instances from a given set of unlabeled instances, constructive machine learning is typically iterative and searches an infinite or exponentially large instance space. With this workshop we want to bring together domain experts employing machine learning tools in constructive processes and machine learners investigating novel approaches or theories concerning constructive processes as a whole.

Invited speakers:

- Ryan Adams (Harvard University)
- François Pachet (Sony)
- Javier González (University of Sheffield)
- Michele Sebag (Université Paris-Sud),

Website:

<http://www.kdml-bonn.de/cml2015>

CrowdML – ICML ’15 Workshop on Crowdsourcing and Machine Learning

Friday (Faidherbe)

Organizers:

- Adish Singla (ETH Zurich)
- Matteo Venanzi (University of Southampton)
- Rafael M. Frongillo (Harvard University)

Abstract:

Crowdsourcing and human computation are emerging paradigms in computing impacting the ability of academic researchers to build new systems and run new experiments involving people, and is also gaining a lot of use within industry for collecting training data for the purpose of machine learning. The fundamental question that we plan to explore in this workshop is: How can we build systems that combine the intelligence of humans and the computing power of machines for solving challenging scientific and engineering problems? The goal is to improve the performance of complex human-powered systems by making them more efficient, robust, and scalable.

Invited speakers:

- Alya Abbott and Ioannis Antonellis, Upwork (Elance-oDesk).
- Daoud Clarke, Lumi.do.
- Jeffrey P. Bigham, Carnegie Mellon University.
- Julian Eisenschlos, Facebook’s Crowdsourcing Team.
- Long Tran-Thanh, University of Southampton.
- Matthew Lease, University of Texas, Austin.
- Mausam, Indian Institute of Technology Delhi.
- Victor Naroditskiy, OneMarketData.

Website:

<http://crowdwisdom.cc/icml2015/>

Extreme Classification: Learning with a Very Large Number of Labels

Friday (Eurotop)

Organizers:

- Moustapha Cissé (KAUST)
- Samy Bengio (Google)
- Patrick Gallinari (LIP6/UPMC)
- Paul Mineiro (Microsoft)
- Nicolas Usunier (Facebook)
- Jia Yuan Yu (IBM)
- Xiangliang Zhang (KAUST)

Abstract:

There is an increasing number of real world classification tasks (e.g. web, biology) where the number of labels is very large (thousands or even millions), the labels are noisy, many of them are rare, sub-linear time prediction (in the number of labels) is mandatory, among several other challenges. This problem, called Extreme Classification is the central topic of the workshop: How to solve it efficiently?

Invited speakers:

- John Langford, Microsoft Research
- Jason Weston, Facebook AI Research
- Jia Deng, University of Michigan
- Manik Varma, Microsoft Research

Website:

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

Features and Structures (FEAST 2015)**Friday (Turin)****Organizers:**

- Chloé-Agathe Azencott (Mines ParisTech)
- Veronika Cheplygina (Erasmus Medical Center)
- Aasa Feragen (University of Copenhagen)

Abstract:

Data consisting of an underlying discrete structure associated with continuous attributes is becoming increasingly important. Representing objects by sets, graphs or sequences is relevant for many fields such as natural language processing, medical imaging, computer vision, bioinformatics, or network analysis. In spite of recent efforts on classification and mining of structural data, many open problems remain. These include converting raw data to a structural representation, combining structure with continuous-valued attributes for classification and regression tasks, and the low sample size / high dimensionality situations associated with biomedical applications, among others. FEAST aims to spark discussion and interaction around these problems.

Invited speakers:

- Robert P.W. Duin, Delft University of Technology
- Pierre Vanderghenst, EPFL
- Florence d'Alché-Buc, Institut Mines-Télécom / Télécom ParisTech

Website:

<https://sites.google.com/site/feast2015/>

Greed is Great

Friday (Artois)

Organizers:

- Liva Ralaivola (Aix-Marseille Université)
- Sandrine Anthoine (Aix-Marseille Université)
- Alain Rakotomamonjy (INSA Rouen)

Abstract:

Many problems from machine learning and signal processing have aim at automatically learning sparse representations from data. Aiming at sparsity actually involves an l_0 “norm” regularization/constraint, and the l_1 convex relaxation way is essentially a proxy to induce sparsity. Greedy methods constitute a strategy to tackle the combinatorial optimization problems posed by the issue of learning sparse representations/predictors. This family of methods has been much less investigated than the convex relaxation approach by the ICML community. This is precisely the purpose of this workshop to give a central place to greedy methods for machine learning and discuss the blessings of such methods.

Invited speakers:

- Aurélien Bellet, Post-doctoral Researcher, Statistics and Applications Group, LTCI, CNRS, Télécom ParisTech
- Kveton Branislav, Researcher, Adobe, San Jose, USA
- Cédric Herzet, Researcher, INRIA Rennes, France
- John Shawe-Taylor, Prof., Centre for Computational Statistics and Machine Learning, UCL, London, UK
- Ke Wei, post-doctoral Researcher, Dept. of Mathematics, Hong-Kong University of Science and Technology

Website:

<https://sites.google.com/site/gretaproject/greed-is-great-icml>

ICML Workshop on Statistics, Machine Learning and Neuroscience (Stamlins 2015)**Friday (Jeanne de Flandre 1)****Organizers:**

- Bertrand Thirion (Parietal team, Inria)
- Lars Kai Hansen (Department of Applied Mathematics and Computer Science, DTU)
- Sanmi Koyejo (Poldrack Lab, Stanford University)

Abstract:

In the last decade, machine learning has had a growing influence on neuroimaging data handling and analysis, making it an ubiquitous component of all kinds of data analysis procedures and software. While it is clear that machine learning has the potential to revolutionize both scientific discovery and clinical diagnosis applications, continued progress requires close collaboration between statisticians, machine learning practitioners and neuroscientists. We propose a one day workshop to provide a forum for interaction between these groups. Our workshop goals are to highlight best practices, disseminate the state of the art in high dimensional methods and related tools with a focus on application to neuroimaging data analysis, and to facilitate discussions to identify the key open problems and opportunities for machine learning in neuroscience.

Invited speakers:

- Pradeep Ravikumar (University of Texas at Austin)
- Morten Mørup (Technical University of Denmark)
- Alexandre Gramfort (Telecom ParisTech, CNRS LTCI)

Website:

<https://sites.google.com/site/stamlins2015/home>

Machine Learning for Education

Friday (Charles de Gaulle)

Organizers:

- Richard G. Baraniuk (Rice University)
- Emma Brunskill (Carnegie Mellon University)
- Jonathan Huang (Google)
- Mihaela van der Schaar (University of California Los Angeles)
- Michael C. Mozer (University of Colorado Boulder)
- Christoph Studer (Cornell University)
- Andrew S. Lan (Rice University)

Abstract:

The goal of this workshop is to bring together experts from different fields of machine learning, cognitive science, and education to explore the interdisciplinary nature of research on the topic of machine learning for education. In particular, we aim to elicit new connections among these diverse fields, identify novel tools and models that can be transferred from one to the others, and explore novel machine learning applications that will benefit the education community. Topics of interest of this workshop include learning and content analytics, scheduling, automatic grading systems, cognitive psychology, and experimental design.

Invited speakers:

- Mihaela van der Schaar, UCLA
- Andrew Lan, Rice University
- Mehdi Sajjadi, University of Hamburg
- Thorsten Joachims, Cornell University
- Jerry Zhu, University of Wisconsin Madison
- Lawrence Carin, Duke University
- Zoran Popovic, University of Washington
- Mehran Sahami, Stanford University
- Igor Labutov, Cornell University
- Burr Settles, Duolingo
- Jacob Whitehill, HarvardX
- Varun Aggarwal,
- Alina von Davier, ETS
- Joseph Jay Williams, HarvardX
- Chris Piech, Stanford University

Website:

http://dsp.rice.edu/ML4Ed_ICML2015

Workshop on Machine Learning Open Source Software 2015: Open Ecosystems**Friday (Jeanne de Flandre 2)****Organizers:**

- Gaël Varoquaux (INRIA)
- Antti Honkela (University of Helsinki)
- Cheng Soon Ong (NICTA)

Abstract:

The workshop is about open source software (OSS) in machine learning. Our aim is to bring together developers to share experiences in publishing ML methods as OSS and to foster interoperability between different packages, as well as allowing developers to demonstrate their software to potential users in the machine learning community. Continuing the tradition of previous MLOSS workshops, we will have a mix of invited speakers, contributed talks and discussion/activity sessions. For 2015, we focus on building open ecosystems.

Invited speakers:

- Matthew Rocklin
- John Myles White

Website:

<https://mloss.org/workshop/icml15/>

Automatic Machine Learning (AutoML)

Saturday (Van Gogh)

Organizers:

- Frank Hutter (University of Freiburg)
- Balazs Kégl (Université Paris-Saclay / CNRS)
- Rich Caruana (Microsoft Research)
- Isabelle Guyon (ChaLearn)
- Hugo Larochelle (Université de Sherbrooke)
- Evelyne Viegas (Microsoft Research)

Abstract:

The success of machine learning in many domains crucially relies on human machine learning experts, who select appropriate features, workflows, machine learning paradigms, algorithms, and their hyperparameters. The rapid growth of machine learning applications has created a demand for off-the-shelf machine learning methods that can be used easily and without expert knowledge. We call the resulting research area that targets progressive automation of machine learning AutoML. For example, a recent instantiation of AutoML we'll discuss is the ongoing ChaLearn AutoML challenge (<http://codalab.org/AutoML>).

Invited speakers:

- Rich Caruana: Robust and Efficient Methods for Practical AutoML
- David Duvenaud: Automatic Model Construction with Gaussian Processes
- Matt Hoffman: Bandits and Bayesian optimization for AutoML
- Jürgen Schmidhuber (tentative)
- Michele Sebag: Algorithm Recommendation as Collaborative Filtering
- Joaquin Vanschoren: OpenML: A Foundation for Networked & Automatic Machine Learning

Website:

<http://icml2015.automl.org>

Demand Forecasting and Machine Learning

Saturday (Artois)

Organizers:

- Francesco Dinuzzo (IBM Research)
- Mathieu Sinn (IBM Research)
- Yannig Goude (EDF R&D)
- Matthias Seeger (Amazon Research)

Abstract:

Demand forecasting is the problem of predicting the amount of goods or services demanded by customers during some future time range: a critical application for many businesses. Retailers base in-stock management decisions like ordering and storage, as well as supply chain management, on demand forecasts. Energy utility companies use forecasting for scheduling operations, investment planning and price bidding. The data revolution creates new opportunities to improve forecast accuracy and granularity, given that heterogenous data sources can be integrated. The focus of the workshop is on demand forecasting by means of data-driven techniques, with a specific emphasis on retail, energy, and transportation industries. We hope to identify the most important challenges from a business point of view, and to start a focussed discussion on how to formalize and solve them by means of machine learning techniques, or on which tools are missing and require additional research efforts.

Invited speakers:

- Nicolas Chapados, ApSTAT Technologies
- Gregory Duncan, Amazon and University of Washington
- Shie Mannor, Technion
- Jean Michel Poggi, Univ. Paris Descartes and Univ. Paris-Sud Orsay, LMO
- Brian Seaman, Walmart Labs
- Gilles Stoltz, CNRS
- Rafal Weron, Wroclaw University of Technology
- Felix Wick, Blue Yonder

Website:

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

Fairness, Accountability, and Transparency in Machine Learning

Saturday (Charles de Gaulle)

Organizers:

- Solon Barocas (Princeton)
- Sorelle Friedler (Haverford)
- Joshua Kroll (Princeton)
- Moritz Hardt (IBM)
- Carlos Scheidegger (U. Arizona)
- Suresh Venkatasubramanian (U. Utah)
- Hanna Wallach (UMass-Amherst)

Abstract:

This interdisciplinary workshop will consider the issues of fairness, accountability, and transparency in machine learning. It will address growing anxieties about the role that automated machine learning systems play in consequential decision-making in such areas as commerce, employment, healthcare, education, and policing.

Invited speakers:

- Nick Diakopoulos, University of Maryland, College Park
- Kazuto Fukuchi, University of Tsukuba
- Manuel Gomez Rodriguez, Max Planck Institute for Software Systems
- Krishna Gummadi, Max Planck Institute for Software Systems
- Sara Hajian, Eurecat Technology Center, Barcelona
- Zubin Jelveh, New York University
- Toshihiro Kamishima, National Institute of Advanced Industrial Science and Technology
- Jeremy Kun, University of Illinois, Chicago
- Isabel Valera Martinez, Universidad Carlos III de Madrid
- Salvatore Ruggieri, Università di Pisa
- Muhammad Bilal Zafar, Max Planck Institute for Software Systems
- Indrė Žliobaitė, Aalto University, University of Helsinki

Website:

<http://fatml.org>

Large-Scale Kernel Learning: Challenges and New Opportunities

Saturday (Eurotop)

Organizers:

- Dino Sejdinovic (Oxford)
- Fei Sha (USC)
- Le Song (Georgia Tech)
- Andrew Gordon Wilson (CMU)
- Zhiyun Lu (USC)

Abstract:

Kernel methods, such as SVMs and Gaussian processes, provide a flexible and expressive learning framework – they are widely applied to small to medium-size datasets, but have been perceived as lacking in scalability. Recently, there has been a resurgence of interest in various fast approximation techniques for scaling up those methods. Indeed, characterizing tradeoffs between their statistical and computational efficiency is growing into an active research topic with encouraging results: it has been reported that appropriately scaled up kernel methods are competitive with deep neural networks. The workshop will overview recent advances and discuss opportunities and challenges within the field.

Invited speakers:

- Francis Bach (ENS)
- Zaid Harchaoui (INRIA & NYU)
- Marius Kloft (Humbold U. of Berlin)
- Neil Lawrence (Sheffield)
- Ruslan Salakhutdinov (Toronto)

Website:

<https://sites.google.com/site/largescalekernelwsicml15/>

Machine Learning for Music Recommendation

Saturday (Jeanne de Flandre 1)

Organizers:

- Erik Schmidt (Pandora)
- Fabien Gouyon (Pandora)
- Gert Lanckriet (University of California San Diego)

Abstract:

The ever-increasing size and accessibility of vast music libraries has created a demand more than ever for machine learning systems that are capable of understanding and organizing this complex data. Collaborative filtering provides excellent music recommendations when the necessary user data is available, but these approaches also suffer heavily from the cold-start problem. Furthermore, defining musical similarity directly is extremely challenging as myriad features play some role (e.g., cultural, emotional, timbral, rhythmic). The topics discussed will span a variety of music recommender systems challenges including cross-cultural recommendation, content-based audio processing and representation learning, automatic music tagging, and evaluation.

Invited talks

- Brian McFee, New York University
- Philippe Hamel, Google
- Sander Dieleman, Aäron van den Oord, Joni Dambre, Ghent University
- Arthur Flexer, Austrian Research Institute for Artificial Intelligence
- Geoffroy Peeters, Frederic Cornu, David Doukhan, Enrico Marchetto, Remi Mignot, Kevin Perros, Lise Regnier, UMR STMS IRCAM, CNRS UMPC
- Bob L. Sturm, Hugo Maruri-Aguilar, Ben Parker, Heiko Grossmann, Queen Mary University of London, Institute for Mathematical Stochastics

Website:

<http://sites.google.com/site/ml4md2015/>

Machine Learning meets Medical Imaging

Saturday (Jeanne de Flandre 2)

Organizers:

- Kanwal Bhatia (Imperial College London)
- Hervé Lombaert (Microsoft Research – INRIA Joint Center)

Abstract:

We aim to present original methods and applications on the interface between Machine Learning and Medical Imaging. Developments in machine learning have created novel opportunities in knowledge discovery, analysis, visualisation and reconstruction of medical image datasets. However, medical images also pose several particular challenges for standard approaches, for instance, lack of data availability, poor image quality or dedicated training requirements, giving rise to questions such as how to better exploit smaller datasets, or understand fundamentals on image spaces or generative models. The workshop will cover both theoretical aspects as well as effective applications of machine learning and medical imaging.

Invited speakers:

- Bertrand Thirion, INRIA France
- John Ashburner, Functional Imaging Laboratory, University College London, UK
- Marleen de Bruijne, Erasmus Medical Center, Rotterdam, NL and DIKU, CS Dept., Copenhagen, DK
- Ben Glocker, Imperial College London, UK

Website:

<https://sites.google.com/site/icml2015mi/>

Mining Urban Data (MUD2)

Saturday (Faidherbe)

Organizers:

- Ioannis Katakis (National & Kapodistrian University of Athens)
- François Schnitzler (The Technion)
- Thomas Liebig (TU Dortmund University)
- Gennady Andrienko (Fraunhofer IAIS and City University London)
- Dimitrios Gunopulos (National & Kapodistrian University of Athens)
- Shie Mannor (The Technion)
- Katharina Morik (TU Dortmund University)

Abstract:

We are gradually moving towards a smart city era. Technologies that apply machine learning algorithms to urban data will have significant impact in a lot of aspects of the citizens' everyday life. Unfortunately, urban data have some characteristics that hinder the state of the art in machine learning algorithms, such as diversity, privacy, lack of labels, noise, complementarity of multiple sources and requirement for online learning. Many smart city applications require to tackle all these problems at once. This workshop aims at discussing a set of new Machine Learning applications and paradigms emerging from the smart city environment.

Invited speakers:

- Eleni Pratsini, Smarter Cities Technology Center, IBM Research, Ireland
- Kristian Kersting, Fraunhofer IAIS and Technical University of Dortmund, Germany
- Sharad Mehrotra, University of California, Irvine

Website:

<http://insight-ict.eu/mud2/>

Resource-Efficient Machine Learning

Saturday (Turin)

Organizers:

- Ralf Herbrich (Amazon)
- Venkatesh Saligrama (Boston University)
- Kilian Q. Weinberger (Washington University in St. Louis)
- Joe Wang (Boston University)
- Tolga Bolukbasi (Boston University)
- Matt Kusner (Washington University in St. Louis)

Abstract:

The aim of this workshop is to bring together researchers in the emerging topics related to learning and decision-making under budget constraints in uncertain and dynamic environments. These problems introduce a new trade off between prediction accuracy and prediction cost. Studying this tradeoff is an inherent challenge that needs to be investigated in a principled fashion in order to invent practically relevant machine learning algorithms.

Invited speakers:

- Jacob Abernethy (University of Michigan)
- Rich Caruana (Microsoft Research) [tentative]
- Csaba Szepesvari (University of Alberta)
- Yoshua Bengio (Université de Montréal)
- Balazs Kegl (CNRS-University of Paris)
- Manik Varma (Microsoft Research)
- Ofer Dekel (Microsoft Research)

Website:

<https://sites.google.com/site/icml2015budgetedml/>

4th Workshop on Machine Learning for Interactive Systems

Saturday (Rubens)

Organizers:

- Heriberto Cuayáhuatl (Heriot-Watt University)
- Nina Dethlefs (University of Hull)
- Lutz Frommberger (University of Bremen)
- Martijn van Otterlo (Radboud University Nijmegen)
- Manuel Lopes (INRIA)
- Olivier Pietquin (University Lille 1)

Abstract:

Learning systems or robots that interact with their environment by perceiving, acting and communicating often face a challenge in how to bring these different concepts together. The challenge arises because core concepts are still predominantly studied in their core communities, such as the computer vision, robotics or natural language processing communities, without much interdisciplinary exchange. Machine learning lies at the core of these communities, and can therefore act as a unifying factor in bringing them closer together. This will be highly important for understanding how state-of-the-art approaches from different disciplines can be combined, refined, and applied to form generally intelligent interactive systems. It will also open a channel for communication and collaboration across research communities.

Invited speakers:

- Jürgen Schmidhuber (tbc), IDSIA, Switzerland
- Ruslan Salakhutdinov, University of Toronto, Canada
- Björn Schuller, Imperial College, United Kingdom
- Ashutosh Saxena, Cornell University, United States

Website:

<http://mlis-workshop.org/2015>

Reviewer awards

The program and area chairs would like to acknowledge the following individuals for their excellent work reviewing for the conference:

Akshay Balsubramani	Philipp Hennig	Rajhans Samdani
Stephen Becker	Daniel Hernandez-Lobato	Suchi Saria
Michael Betancourt	Bert Huang	Lawrence Saul
Y-Lan Boureau	Jonathan Huang	Alexander Schwing
Kyunghyun Cho	Jonathan Huggins	Dafna Shahaf
Marc Deisenroth	Adam Johansen	Martin Slawski
Amir-massoud Farahmand	Hachem Kadri	Bharath Sriperumbdur
Rina Foygel	Purushottam Kar	Jacob Steinhardt
Peter Frazier	Mark Kozdoba	Daniel Tarlow
Thomas Gartner	Tobias Lang	Michal Valko
Rong Ge	Tor Lattimor	Jan-Willem Vandemeent
Michael Gelbart	Alessandro Lazaric	Nakul Verma
Ran Gilad-Bachrach	Sahand Negahban	Huahua Wang
Steffen Grunewalder	Willie Neiswanger	Adam White
Marek Grzes	Brendan O'Connor	Martha White
Andras Gyorgy	Pedro Ortega	Lin Xiao
Amaury Habrard	Razvan Pascanu	
Yoni Halpern	Patrick Pletscher	

Authors

- Abbasi-Yadkori, Yasin, 22
Abbeel, Pieter, 25, 33
Adams, Ryan, 28, 32, 33
Agarwal, Alekh, 25, 29
Agarwal, Shivani, 25, 27, 34
Agrawal, Ankur, 24
Ahn, KookJin, 27
Airoldi, Edoardo, 26
Allamanis, Miltos, 36
Amid, Ehsan, 30
Amini, Massih-Reza, 23
Ammar, Haitham Bou, 33
An, Senjian, 26
Anava, Oren, 32
Arora, Raman, 23
Arriaga, Jesús Pérez, 23
Ashkan, Azin, 23
Asteris, Megasthenis, 35
Aussem, Alexandre, 34
Avron, Haim, 26
Aybat, Necdet, 22
- Ba, Jimmy, 28
Bach, Stephen, 25
Bachem, Olivier, 28
Bachman, Philip, 26
Bahadori, Mohammad Taha, 36
Bai, Lu, 27
Balandat, Maximilian, 27
Bansal, Trapit, 31
Barbosa, Rafael, 32
Barnard, Kobus, 31
Bartlett, Peter, 22
Bayen, Alexandre, 27
Belanger, David, 32
Benavoli, Alessio, 22
Bengio, Yoshua, 22, 28
Bennamoun, Mohammed, 26
Berlind, Christopher, 26
Bernstein, Garrett, 27
Besserve, Michel, 24
Betancourt, Michael, 29
Bhattacharyya, Chiranjib, 31
Bi, Jinbo, 35
Biggio, Battista, 33
Bilmes, Jeff, 23, 32
Blaschko, Matthew, 25, 30
- Blechschmidt, Katharina, 35
Blum, Avrim, 23
Blundell, Charles, 33
Bottou, Leon, 25
Bouchard-Côté, Alexandre, 22
Bounliphone, Wacha, 30
Bourque, Guillaume, 25
Boussaid, Farid, 26
Boutsidis, Christos, 36
Boyd-Graber, Jordan, 25
Brau, Ernesto, 31
Braun, Mikio, 31
Brown, Gavin, 33
Burdick, Joel, 28
Busa-Fekete, Robert, 23
Butler, Emily, 31
- Caffo, Brian, 36
Cao, Yue, 26
Caramanis, Constantine, 30
Carin, Lawrence, 27, 31, 33
Carlson, David, 33
Carpentier, Alexandra, 23
Chandrasekaran, Bharath, 35
Chang, Kai-Wei, 29
Chang, Xiaojun, 35
Chao, Wei-Lun, 29
Chaudhuri, Sougata, 25
Chazal, Frederic, 29
Chen, Changyou, 33
Chen, Chun-Sheng, 31
Chen, Fang, 31
Chen, Liang-Chieh, 29
Chen, Wei, 35
Chen, Wenlin, 24
Chen, Xi, 22
Chen, Xilin, 35
Chen, Yixin, 24
Chen, Yudong, 35
Chen, Yutian, 29, 35
Chen, Yuxin, 25
Cheng, Dehua, 31
Cheng, Jian, 30
Cho, Kyunghyun, 22, 28
Chung, Junyoung, 22
Ciliberto, Carlo, 27
Clérot, Fabrice, 22
- Clark, Christopher, 33
Clemençon, Stéphan, 25
Cohen, Alon, 32
Cohen, Taco, 32
Collobert, Ronan, 36
Corani, Giorgio, 22
Cormode, Graham, 27
Cornebise, Julien, 33
Corrado, Greg, 22
Cortes, Corinna, 31
Courville, Aaron, 28
Csiba, Dominik, 25
Cuturi, Marco, 29
- Daniely, Amit, 27
Danihelka, Ivo, 28
Das, Mrinal, 31
Daume, Hal, 29
Deisenroth, Marc, 34
Deshpande, Amit, 32
Dhillon, Inderjit, 22, 24, 25, 35
Dimakis, Alex, 35
Ding, Yufei, 36
Djlonga, Josip, 28
Dubey, Avinava, 35
Dukkipati, Ambedkar, 36
Duvenaud, David, 28
Dyer, Chris, 32
- Eaton, Eric, 33
Eckert, Claudia, 33
Eggeling, Ralf, 32
Elghazel, Haytham, 34
Ene, Alina, 32
Engler, Raphael, 34
Ermon, Stefano, 30
- Fan, Yingying, 36
Faruqui, Manaal, 32
Fasy, Brittany, 29
Feng, Lu, 26
Fercoq, Olivier, 35
Fern, Xiaoli, 23
Fetaya, Ethan, 29
Fienberg, Stephen, 33
Filippone, Maurizio, 34
Fisher, John, 27
Flaxman, Seth, 26

Foreman-Mackey, Dan, 24
 Foster, Dean, 28
 Foulds, James, 27, 33
 Frazier, Peter, 35
 Friedlander, Michael, 30
 Friedman, Arik, 33
 Frostig, Roy, 31
 Fujimaki, Ryohei, 26, 31
 Fukunaga, Takuro, 30
 Fumera, Giorgio, 33

 Gajane, Pratik, 22
 Gal, Yarin, 29, 34
 Gan, Zhe, 33
 Ganguli, Surya, 31
 Ganin, Yaroslav, 26
 Garber, Dan, 30, 31
 Gardner, Jacob, 33
 Garnett, Roman, 29, 33
 Gasse, Maxime, 34
 Ge, Hong, 35
 Ge, Rong, 24, 31
 Geiger, Philipp, 24
 Gelbart, Michael, 28
 Germain, Mathieu, 33
 Getoor, Lise, 25, 27, 33, 36
 Ghahramani, Zoubin, 27–29, 35
 Ghavamzadeh, Mohammad, 34
 Ghoshal, Suprovat, 25
 Ghoshdastidar, Debarghya, 36
 Giesen, Joachim, 35
 Giguère, Sébastien, 29
 Gilmore, David, 22
 Girolami, Mark, 33
 Gittens, Alex, 36
 Globerson, Amir, 25
 Goernitz, Nico, 31
 Gonen, Alon, 27
 Gong, Mingming, 24
 Gong, Pinghua, 36
 Gopalakrishnan, Kailash, 24
 Gordon, Andrew, 36
 Gotovos, Alkis, 28
 Gouws, Stephan, 22
 Gramfort, Alexandre, 35
 Graves, Alex, 28
 Gregor, Karol, 28, 33, 34
 Gretton, Arthur, 30
 Grosse, Ivo, 32

 Grosse, Roger, 23, 24
 Gu, Bin, 34
 Gu, Quanquan, 24
 Guan, Jinyan, 31
 Guestrin, Carlos, 31
 Guha, Sudipto, 27
 Guibas, Leonidas, 22
 Gulcehre, Caglar, 22
 Gunter, Tom, 34
 Guo, Hongyu, 36
 Gupta, Suyog, 24
 Gyorgy, Andras, 32, 36

 Hüllermeier, Eyke, 23
 Habrard, Amaury, 23
 Hallak, Assaf, 22
 Han, Bohyung, 28
 Han, Fang, 36
 Han, Insu, 29
 Han, Qiuyi, 26
 Han, Weidong, 35
 Hanawal, Manjesh, 23
 Hancock, Edwin, 27
 Hardt, Moritz, 23
 Hayashi, Kohei, 31
 Hazan, Elad, 27, 30–32
 Hazan, Tamir, 28, 32
 He, Xinran, 27
 Heinrich, Johannes, 34
 Henao, Ricardo, 27, 31, 33
 Hernandez-Lobato, Daniel, 27
 Hernandez-Lobato, Jose Miguel,
 27, 28, 33
 Higham, Catherine, 26
 Ho, Shirley, 29
 Hoang, Quang Minh, 26
 Hoang, Trong Nghia, 26
 Hocking, Toby, 25
 Hoffman, Matt, 28, 32
 Hoffman, Matthew, 28
 Hoffman, Michael, 32
 Hogg, David, 24
 Honda, Junya, 23
 Hong, Seunghoon, 28
 Horel, Thibaut, 27
 Horesh, Lior, 26
 Horgan, Daniel, 34
 Hsieh, Cho-Jui, 22, 24
 Hsu, Cheng-Yu, 27

 Hu, Zhiting, 35
 Huang, Bert, 25, 36
 Huang, Jonathan, 22
 Huang, Ruitong, 36
 Huang, Zhiwu, 35
 Huggins, Jonathan, 22, 27
 Husmeier, Dirk, 26

 Ioffe, Sergey, 24
 Iwamasa, Yuni, 30
 Iyengar, Garud, 22
 Iyer, Rishabh, 32
 Izbicki, Mike, 28

 Jaggi, Martin, 22
 Jain, Prateek, 23, 27
 Jakubowicz, Jérémie, 25
 Janzing, Dominik, 24
 Jedynak, Bruno, 35
 Jernite, Yacine, 32
 Jewell, Sean, 22
 Jia, Jiaya, 28
 Jiang, Nan, 34
 Jiao, Yunlong, 30
 Jin, Rong, 23, 24
 Joachims, Thorsten, 24
 Johnson, Tyler, 31
 Jordan, Michael, 22, 25, 26, 35,
 36
 Jozefowicz, Rafal, 22
 Jung, Kyomin, 30

 Kairouz, Peter, 33
 Kakade, Sham, 31, 32
 Kakimura, Naonori, 30
 Kale, David, 36
 Kalousis, Alexandros, 32
 Kambadur, Prabhanjan, 36
 Kandasamy, Kirthivasan, 28
 Kandemir, Melih, 26
 Kapoor, Ashish, 32
 Kar, Purushottam, 23, 27
 Kavukcuoglu, Koray, 33
 Kawarabayashi, Ken-ichi, 30
 Keren, Daniel, 28
 Khardon, Roni, 34
 Kim, Hyunwoo, 22
 Kim, Minhwan, 30
 Kingma, Diederik, 28
 Kiros, Ryan, 28

Kloft, Marius, 31
 Knowles, David, 28
 Koepke, Hoyt, 30
 Koivisto, Mikko, 32
 Kolkin, Nicholas, 32
 Komiyama, Junpei, 23
 Korda, Nathaniel, 22
 Krause, Andreas, 28
 Krichene, Walid, 27
 Krishnamurthy, Akshay, 29
 Krzakala, Florent, 24
 Kukliansky, Doron, 23
 Kulesza, Alex, 34
 Kumar, Akshat, 36
 Kumar, Shachi, 33
 Kusner, Matt, 32, 33
 Kuznetsov, Vitaly, 31
 Kveton, Branislav, 23
 Kwak, Suha, 28
 Kyrillidis, Anastasios, 35

 La, Prashanth, 22
 Lakshmanan, K., 25
 Lanctot, Marc, 34
 Lang, Dustin, 32
 Langford, John, 29
 Langley, Raymond, 27
 Laradji, Issam, 30
 Larochele, Hugo, 33
 Laue, Soeren, 35
 Laviolette, Francois, 29
 Le, Tam, 29
 Lebret, Remi, 36
 Lecci, Fabrizio, 29
 Lee, Ching-Pei, 22
 Lee, Taehoon, 33
 Lempitsky, Victor, 26
 Leng, Cong, 30
 Lesner, Boris, 25
 Lessard, Laurent, 35
 Levine, Sergey, 25
 Li, Bin, 31
 Li, Ping, 35
 Li, Wu-Jun, 23
 Li, Xianqiu, 35
 Li, Yujia, 33
 Lian, Wenzhao, 31
 Liang, Dawen, 29
 Liang, Percy, 29, 30

 Liao, Renjie, 28
 Libbrecht, Maxwell, 32
 Lim, Lek-Heng, 25
 Lim, Shiau Hong, 35
 Lim, Woosang, 30
 Lin, Xiao, 22, 35
 Lin, Xin, 35
 Lindsten, Fredrik, 29
 Ling, Charles, 34
 Ling, Qing, 35
 Liu, Chunchen, 26
 Liu, Han, 36
 Liu, Hanxiao, 29
 Liu, Yan, 27, 31, 36
 Liu, Yan-Fu, 27
 Livescu, Karen, 23
 Livni, Roi, 27
 Lloyd, Chris, 34
 London, Ben, 36
 Long, Mingsheng, 26
 Lopez-Paz, David, 24
 Low, Bryan Kian Hsiang, 26
 Lu, Hanqing, 30
 Lu, Jin, 35
 Lu, Yichao, 28
 Lucas, Joseph, 31
 Lucic, Mario, 28
 Luo, Luo, 23

 Ma, Chenxin, 22
 Ma, Tengyu, 30
 Ma, Zhuang, 28
 Macdonald, Benn, 26
 Maclaurin, Dougal, 28
 Maeda, Shin-ichi, 31
 Maehara, Takanori, 30
 Maheswaranathan, Niru, 31
 Mahoney, Michael, 31
 Malek, Alan, 22
 Malioutov, Dmitry, 29
 Mangili, Francesca, 22
 Mann, Timothy, 22
 Mannor, Shie, 22, 34
 Manoel, Andre, 24
 Mansimov, Elman, 28
 Mansinghka, Vikash, 22
 Mansour, Yishay, 27
 Marchand, Mario, 29
 Marchand-Maillet, Stephan, 32

 Mariet, Zelda, 27
 Martens, James, 24
 Matthew, Taddy, 31
 Mazumdar, Arya, 30
 McAuliffe, Jon, 32
 McGregor, Andrew, 27
 Melibari, Mazen, 27
 Menon, Aditya, 23
 Michel, Bertrand, 29
 Michels, Dominik, 29
 Miller, Andrew, 32
 Miyauchi, Atsushi, 30
 Mohamed, Shakir, 28
 Mohri, Mehryar, 31
 Moritz, Philipp, 25
 Morrison, Clayton, 31
 Mroueh, Youssef, 27
 Muandet, Krikamol, 24
 Munos, Remi, 23
 Muraoka, Yusuke, 26
 Murray, Iain, 33
 Musuvathi, Madanlal, 36
 Mytkowicz, Todd, 36

 Nadler, Boaz, 36
 Naesseth, Christian, 29
 Nakagawa, Hiroshi, 23
 Nan, Feng, 34
 Naraghi-Pour, Mort, 32
 Narasimhan, Harikrishna, 23, 27, 34
 Narasimhan, Karthik, 22
 Narayan, Karthik, 33
 Narayanan, Pritish, 24
 Natarajan, Nagarajan, 24
 Neeman, Joe, 25
 Neill, Daniel, 26
 Neyshabur, Behnam, 30
 Ng, Jun Wei, 34
 Nguyen, Andy, 22
 Nguyen, Huy, 32
 Nickisch, Hannes, 26
 Nishihara, Robert, 35
 Niu, Gang, 27
 Noble, William, 32
 Nock, Richard, 33
 Nutini, Julie, 30

 Oh, Sewoong, 33

Olukotun, Kunle, 25
 Ong, Cheng Soon, 23
 Ortner, Ronald, 25
 Osadchy, Margarita, 28
 Osborne, Michael, 34
 Osogami, Takayuki, 34

 Pacheco, Jason, 36
 Packard, Andrew, 35
 Padilla, Oscar Hernan Madrid, 32
 Paisley, John, 29, 33
 Papachristoudis, Georgios, 27
 Pareek, Harsh, 25
 Park, Dohyung, 25
 Park, Haesun, 30
 Patrini, Giorgio, 33
 Patwary, Mostofa, 28
 Peres, Yuval, 30
 Perolat, Julien, 25
 Perrot, Michaël, 23
 Peters, Jonas, 24
 Pham, Anh, 23
 Phulsuksombati, Mike, 22
 Piech, Chris, 22
 Pietquin, Olivier, 25
 Pinheiro, Pedro, 36
 Piot, Bilal, 25
 Plessis, Marthinus Du, 27
 Poczoz, Barnabas, 28
 Poggio, Tomaso, 27
 Pouget-Abadie, Jean, 27
 Poupert, Pascal, 27
 Prabhat, , 32
 Prabhat, Mr, 28
 Prade, Henri, 23
 Prasad, Adarsh, 25
 Precup, Doina, 26
 Price, Eric, 30
 Punjani, Ali, 33

 Qirong, Ho, 35
 Qiu, Huitong, 36
 Qu, Qing, 24
 Qu, Zheng, 25

 Raich, Raviv, 23
 Rajan, Purnima, 35
 Rajkumar, Arun, 25
 Ralaivola, Liva, 23
 Ramaswamy, Harish, 27, 34

 Rao, Vinayak, 31
 Raskutti, Garvesh, 31
 Ravikumar, Pradeep, 25, 32, 35
 Re, Christopher, 25
 Recht, Ben, 35
 Regier, Jeffrey, 32
 Rekatsinas, Theodoros, 27
 Ren, Jimmy, 28
 Rezende, Danilo, 28
 Richman, Oran, 34
 Richtarik, Peter, 22, 25
 Rigai, Guillem, 25
 Rinaldo, Alessandro, 29
 Rippel, Oren, 28
 Roberts, Stephen, 29, 34
 Robinson, Daniel, 24
 Roli, Fabio, 33
 Rolland, Amélie, 29
 Romera-Paredes, Bernardino, 26
 Rooyen, Brendan Van, 23
 Rosasco, Lorenzo, 27
 Rossi, Luca, 27
 Roth, Dan, 22
 Roughgarden, Tim, 25
 Rush, Alexander, 32
 Russell, Matthew, 22
 Ryabko, Daniil, 25

 Sa, Christopher De, 25
 Saad, Yousef, 30
 Saeedi, Ardavan, 22
 Saha, Aadirupa, 34
 Sahami, Mehran, 22
 Salakhudinov, Ruslan, 23, 28
 Saligrama, Venkatesh, 23, 34
 Salimans, Tim, 28
 Salmon, Joseph, 35
 Samo, Yves-Laurent Kom, 29
 Sanghavi, Sujay, 25
 Satish, Nadathur, 28
 Schölkopf, Bernhard, 24
 Schaar, Mihaela van der, 32
 Schaul, Tom, 34
 Scherrer, Bruno, 22, 25
 Schlegel, David, 32
 Schmidt, Mark, 30
 Schneider, Jeff, 28, 29
 Schnitzler, Francois, 22
 Schoelkopf, Bernhard, 24

 Schon, Thomas, 29
 Schulman, John, 25
 Schwing, Alexander, 29
 Serrurier, Mathieu, 23
 Sha, Fei, 29
 Shah, Amar, 28
 Shah, Nihar, 30
 Shah, Parikshit, 24
 Shajarisales, Naji, 24
 Shalev-Shwartz, Shai, 27
 Shamir, Ohad, 23, 30
 Shan, Shiguang, 35
 Sharma, Dravyansh, 32
 Sheldon, Dan, 27, 36
 Shelton, Christian, 28
 Shen, Xipeng, 36
 Sheth, Rishit, 34
 Shin, Jinwoo, 29
 Sibony, Eric, 25
 Sidford, Aaron, 23
 Silver, David, 34
 Simek, Kyle, 31
 Simon-Gabriel, Carl-Johann, 24
 Singh, Aarti, 35
 Singh, Satinder, 34
 Singh, Vikas, 22
 Smith, Noah, 32
 Smith, Virginia, 22
 Smola, Alex, 26, 33
 Snoek, Jasper, 28
 So, Anthony Man-Cho, 31
 Sobihani, Parinaz, 36
 Sohl-Dickstein, Jascha, 31
 Solomon, Justin, 29
 Soltanmohammadi, Erfan, 32
 Sontag, David, 25, 32
 Spencer, Neil, 22
 Sra, Suvrit, 27
 Srebro, Nathan, 30
 Srivastava, Nitish, 28
 Steele, Guy, 33
 Steinhardt, Jacob, 29, 30
 Storkey, Amos, 33
 Sudderth, Erik, 36
 Suggala, Arun Sai, 32
 Sugiyama, Masashi, 27
 Suh, Changho, 25
 Sui, Yanan, 28
 Sun, Jiangwen, 35

Sun, Ju, 24
 Sun, Ke, 32
 Sun, Shiliang, 32
 Sun, Siqi, 29
 Sun, Tao, 36
 Sun, Yu, 32
 Sundaram, Narayanan, 28
 Sutskever, Ilya, 22
 Sutton, Rich, 34
 Suzuki, Taiji, 36
 Swaminathan, Adith, 24
 Swersky, Kevin, 28, 33
 Syed, Umar, 31
 Szegedy, Christian, 24
 Szepesvári, Csaba, 36
 Szepesvari, Csaba, 23, 32
 Sznitman, Raphael, 35
 Szorenyi, Balazs, 23

 Tagorti, Manel, 22
 Takac, Martin, 22
 Tang, Gongguo, 24
 Tang, Qingming, 29
 Tansey, Wesley, 32
 Tao, Dacheng, 24
 Tarlow, Daniel, 36
 Tassarotti, Joseph, 33
 Tenenbaum, Josh, 27
 Tenenhaus, Arthur, 30
 Tewari, Ambuj, 25, 27
 Theis, Lucas, 28
 Theocharous, Georgios, 34
 Thomas, Philip, 34
 Tolstikhin, Iliya, 24
 Tomlin, Claire, 27
 Torr, Philip, 26
 Tramel, Eric, 24
 Trask, Andrew, 22
 Tristan, Jean-Baptiste, 33
 Tsalik, Ephraim, 27
 Turner, Richard, 34
 Tutunov, Rasul, 33
 Tyree, Stephen, 24

 Ubaru, Shashanka, 30
 Ukkonen, Antti, 30
 Ullman, Shimon, 29
 Urner, Ruth, 26
 Urtasun, Raquel, 29

 Urvoy, Tanguy, 22

 Valko, Michal, 23
 Vanseijen, Harm, 34
 Vemuri, Baba, 22
 Vert, Jean-Philippe, 30
 Vidal, Rene, 24, 35
 Virtanen, Seppo, 33
 Viswanath, Pramod, 33

 Wagner, Avishai, 24
 Wainwright, Martin, 36
 Wan, Moquan, 35
 Wang, Dun, 24
 Wang, Jianmin, 26
 Wang, Jie, 26
 Wang, Joseph, 34
 Wang, Jun, 32
 Wang, Ruiping, 35
 Wang, Weiran, 23
 Wang, Yang, 31
 Wang, Yi, 31
 Wang, Yining, 22, 35
 Wang, Yu, 35
 Wang, Yu-Xiang, 33, 35
 Wang, Yuyang, 34
 Wang, Zi, 22
 Ward, Justin, 32
 Wassell, Ian, 35
 Wasserman, Larry, 29
 Wei, Kai, 32
 Wei, Yi, 36
 Weinberger, Kilian, 24, 32, 33
 Weiss, Eric, 31
 Weiss, Roi, 36
 Welling, Max, 28, 32
 Wen, Zheng, 23
 Weng, Paul, 23
 Wierstra, Daan, 28, 33
 Williamson, Bob, 23
 Wilson, Andrew, 26
 Wilson, James, 24
 Winner, Kevin, 27
 Wipf, David, 31, 35
 Wirth, Anthony, 27
 Wright, John, 24
 Wu, Jiaxiang, 30
 Wu, Shan-Hung, 27
 Wu, Yifan, 32

 Wu, Yuanbin, 32
 Wyle, Mitch, 31

 Xiao, Huang, 33
 Xie, Yubo, 23
 Xin, Bo, 31
 Xing, Eric, 35
 Xu, Huan, 26, 30, 31, 35
 Xu, Jia, 22
 Xu, Jinbo, 29
 Xu, Kelvin, 28
 Xu, Kevin, 26
 Xu, Li, 28
 Xu, Miao, 24
 Xu, Sheng, 36
 Xu, Tingyang, 35

 Yabe, Akihiro, 30
 Yan, Qiong, 28
 Yang, Congyuan, 24
 Yang, Jinwei, 35
 Yang, Tianbao, 23, 24
 Yang, Wenzhuo, 26, 30, 31
 Yang, Yi, 35
 Yang, Yiming, 29
 Ye, Jieping, 26, 36
 Yen, En-Hsu, 35
 Yi, Han-Gyol, 35
 Yi, Xinyang, 30
 Yildirim, Cafer, 25
 Yogatama, Dani, 32
 Yoon, Sungroh, 33
 You, Chong, 35
 You, Tackgeun, 28
 Yu, Hsiang-Fu, 22
 Yu, Jiaqian, 25
 Yu, Jun, 31
 Yu, Rose, 31
 Yu, Yaoliang, 35
 Yuan, Xin, 27
 Yuille, Alan, 29

 Zaffalon, Marco, 22
 Zaremba, Wojciech, 22
 Zdeborová, Lenka, 24
 Zeevi, Assaf, 32
 Zemel, Rich, 28, 33
 Zhang, Aonan, 33
 Zhang, Jin, 25
 Zhang, Kun, 24

Zhang, Lijun, 23, 24
Zhang, Qi, 31
Zhang, Tong, 25, 35
Zhang, Xi, 30
Zhang, Yuchen, 22, 35, 36
Zhang, Zhihong, 27
Zhang, Zhihua, 23
Zhao, Han, 27

Zhao, Peilin, 25, 35
Zhao, Qi, 24
Zhao, Yue, 36
Zhong, Kai, 35
Zhou, Dengyong, 30
Zhou, Qiang, 24
Zhou, Zhi-Hua, 24
Zhou, Zirui, 31

Zhu, Jun, 22
Zhu, Michael, 30
Zhu, Rongda, 24
Zhu, Shenghuo, 23, 24
Zhu, Xiaodan, 36
Zou, James, 24
Zuk, Or, 24

Sponsor Scholars

ICML platinum sponsors (Alibaba, Baidu Research, Euratechnologies, NSF) and gold sponsors (Amazon, Crédit Agricole S.A., Facebook, Google, Intel, MFG labs, Microsoft research, Panasonic, Winton) have participated to our “sponsor scholar” program. A part of their sponsorship has been allocated to scholarships to help 102 young researchers be able to travel to ICML and present their work to their community.

We thank the sponsors very much, and we congratulate the sponsor scholars.

- Miltiadis Allamanis, University of Edinburgh
- Ehsan Amid, Aalto University
- Stephen Bach, University of Maryland, College Park
- Olivier Bachem, ETH Zurich
- Mohammad Taha Bahadori, University of Southern California
- Sadaf Behlim, National University of Computer and Emerging Sciences
- David Belanger, University of Massachusetts Amherst
- Christopher Berling, Georgia Tech
- Xin Bo, Peking University
- Liang-Chieh Chen, University of California- Los Angeles
- Junyoung Chung, University of Montreal
- Lidia Contreras Ochando, Polytechnic University of Valencia
- Mrinal Das, Indian Institute of Science
- Josip Djolonga, ETH Zurich
- Csiba Dominik, University of Edinburgh
- Alexandre Drouin, Laval University
- Mehrdad Farajtabar, Georgia Institute of Technology
- Cristina I Font Julian, Polytechnic University of Valencia
- Dylan Foster, Cornell University
- Yaroslav Ganin, Skoltech
- Jacob Gardner, Washington University in St. Louis
- Bernstein Garrett, University of Massachusetts Amherst
- Maxime Gasse, Université Lyon 1
- Mathieu Germain, Université de Sherbrooke
- Alon Gonen, Hebrew University of Jerusalem
- Alkis Gotovos, ETH Zurich
- Jinyan Guan, University of Arizona
- Assaf Hallak, Technion
- Xinran He, University of Southern California
- Quang Minh Hoang, National University of Singapore
- Seunghoon Hong, POSTECH
- Thibaut Horel, Harvard University
- Cho-Jui Hsieh, UT Austin
- Xiao Huang, Technische Universität München
- Han Insu, KAIST
- Mike Izbicki, UC Riverside
- Zubin Jelveh, New York University
- Sean Jewell, University of British Columbia
- Nan Jiang, University of Michigan
- Pooria Joulani, University of Alberta
- Kirthevasan Kandasamy, Carnegie Mellon University
- Gil Keren, Ben Gurion University
- Hyunwoo Kim, University of Wisconsin-Madison
- Ramya Korlakai VInayak, California Institute of Technology
- Matt Kusner, Washington University in St. Louis
- Ilja Kuzborskij, École Polytechnique Fédérale de Lausanne
- Vitaly Kuznetsov, New York University
- Taehoon Lee, Seoul National University

- Ching-Pei Lee, University of Illinois at Urbana Champaign
- Cong Leng, Institute of Automation, Chinese Academy of Sciences
- Yujia Li, University of Toronto
- Max Libbrecht, University of Washington
- Xin Lin, The University of Texas at Austin
- Yan-Fu (Keith) Liu, National Tsing Hua University
- Hanxiao Liu, Carnegie Mellon University
- Ben London, University of Maryland
- Mario Lucic, ETH Zurich
- Luo Luo, Shanghai Jiao Tong University
- Haipeng Luo, Princeton University
- Zhuang Ma, University of Pennsylvania
- Benn Macdonald, University of Glasgow
- Zelda Mariet, Massachusetts Institute of Technology
- Julie Nutini, University of British Columbia
- Cem Orhan, Bilkent University
- Giorgio Patrini, Australian National University / NICTA
- Matteo Pirota, Politecnico di Milano
- Adarsh Prasad, UT Austin
- Qing Qu, Columbia University
- Siddharth Reddy, Cornell University
- Amélie Rolland, Université Laval
- Flaxman Seth, Carnegie Mellon University
- Nihar Shah, University of California, Berkeley
- Naji Shajarisales, Max Planck Institute of Intelligent Systems
- Dravyansh Sharma, Indian Institute of Technology Delhi
- Rishit Sheth, Tufts University
- Neil Spencer, University of British Columbia
- Nitish Srivastava, University of Toronto
- Arun Sai Suggala, University of Texas at Austin
- Tao Sun, University of Massachusetts Amherst
- Adith Swaminathan, Cornell University
- Wesley Tansey, University of Texas at Austin
- Shashank, Ubaru, University of Minnesota-Twin Cities
- Avishai Wagner, Hebrew University of Jerusalem
- Yu Wang, University of Cambridge
- Yu-Xiang Wang, Carnegie Mellon University
- Kai Wei, University of Washington
- Kevin Winner, University of Massachusetts, Amherst
- Lim Woosang, KAIST
- Yifan Wu, University of Alberta
- Kelvin Xu, Université de Montréal
- Jia Xu, University of Wisconsin-Madison
- Sheng Xu, Johns Hopkins University
- Congyuan Yang, Johns Hopkins University
- Chong You, Johns Hopkins University
- Qi (Rose) Yu, University of Southern California
- Qi Zhang, Chinese University of Hong Kong
- Yuchen Zhang, University of California Berkeley
- Han Zhao, University of Waterloo
- Hu Zhiting, Carnegie Mellon University
- Qiang Zhou, National University of Singapore
- Rongda Zhu, University of Illinois at Urbana Champaign
- Michael Zhu, Stanford University

