

ICML 2013 Atlanta Conference Program

The 30th International Conference on
Machine Learning

June 16 – June 21, 2013

Atlanta, Georgia, USA

Contents

Conference at a Glance	5
Sponsors	10
Welcome from the General Chair	13
Organizing Committee	15
Area Chairs	16
Information	17
ICML Tutorials, Sunday June 16	19
NLP-Related Workshops, Sunday June 16	20
ICML Main Conference, Monday June 17	21
ICML Main Conference, Tuesday June 18	29
ICML Main Conference, Wednesday June 19	39
ICML Workshops, Thursday June 20	47
ICML Workshops, Friday June 21	48
Program Committee	49
Maps	55

The 30th International Conference on Machine Learning: ICML 2013 Atlanta

Saturday, June 15, 2013	
Joint NAACL/ICML Symposium on Natural Language Processing and Machine	
The Joint NAACL/ICML Symposium is Located at the Westin Peachtree Plaza Hotel. All other events are at the Marriott Marquis	
1600-2000	ICML Registration Open 1600-2000
Sunday, June 16, 2013	
ICML Tutorials and NLP-Related Workshops	
0730-1800	ICML Registration Open 0730-1800
0830-1000	AM Tutorial: Deep Learning <i>International 7,8,9</i>
	AM Tutorial: Submodularity in Machine Learning: New Directions <i>Marquis 103,104,105</i>
	AM Tutorial: Tensor Decomposition Algorithms for Latent Variable Model Estimation <i>International 1,2,3</i>
	AM Tutorial: Multi-Target Prediction <i>International 10</i>
	All Day Workshop: Deep Learning For Audio Speech and Language Processing <i>International 6</i>
	All Day Workshop: Structured Learning: Inferring Graphs from Structured and Unstructured Inputs <i>International 4,5</i>
1000-1030	Coffee Break (provided)
1030-1200	AM Tutorials Continue
	All Day Workshops Continue
1200-1400	Lunch (on your own)
1400-1530	PM Tutorial: Discovering Multiple Clustering Solutions: Grouping Objects in Different Views <i>International 1,2,3</i>
	PM Tutorial: Copulas in Machine Learning <i>International 7,8</i>
	PM Tutorial: Music Information Research Based on Machine Learning <i>International 10</i>
	PM Tutorial: Topological Data Analysis <i>International 9</i>
	All Day Workshops Continue
1530-1600	Coffee Break (provided)
1600-1730	PM Tutorials Continue
	All Day Workshops Continue
Monday, June 17, 2013	
ICML Main Conference	
0730-1800	Registration Open 0730-1800
0830-1000	Keynote Speaker: Carlos Guestrin <i>International 7,8,9,10</i>
1000-1030	Coffee Break (provided)
1030-1210	Track A: Deep Learning 1 <i>International 7,8,9,10</i>
	Track B: Compressed Sensing 1 <i>International 4,5</i>
	Track C: Reinforcement Learning 1 <i>International 1,2,3</i>
	Track D: Social Networks <i>International 6</i>
1210-1400	Lunch (on your own)
1400-1540	Track A: Deep Learning 2 <i>International 7,8,9,10</i>
	Track B: Compressed Sensing 2 <i>International 4,5</i>
	Track C: Reinforcement Learning 2 <i>International 1,2,3</i>
	Track D: Topic Modeling 1 <i>International 6</i>
1540-1600	Coffee Break (provided)

The 30th International Conference on Machine Learning: ICML 2013 Atlanta

1600-1740	Track A: Deep Learning and Neuroscience	<i>International Rooms 7,8,9,10</i>
	Track B: Compressed Sensing 3	<i>International Rooms 4,5</i>
	Track C: Reinforcement Learning and Time Series	<i>International Rooms 1,2,3</i>
	Track D: Topic Modeling 2	<i>International 6</i>
1930-2000	Poster Setup	<i>Skyline - 10th Floor</i>
2000-2200	Poster Session	<i>Skyline -10th Floor</i>
Tuesday, June 18, 2013		
ICML Main Conference		
0700-0830	Women in ML Breakfast by Google and Georgia <i>Marquis Ballroom Salon A</i>	
0800-1200	Registration Open 0730-1200	
0830-1000	Keynote Speaker: Santosh Vempala	<i>International 7,8,9,10</i>
1000-1030	Coffee Break (provided)	
1000-1210	Track A: Online Learning 1	<i>International Rooms 7,8,9,10</i>
	Track B: Feature Learning	<i>International Rooms 4,5</i>
	Track C: General SVM & Decision Tree Methods	<i>International Rooms 1,2,3</i>
	Track D: Spectral Learning & Tensors	<i>International 6</i>
1210-1400	Lunch (on your own)	
1400-1540	Track A: Online Learning 2	<i>International Rooms 7,8,9,10</i>
	Track B: Structured Labeling	<i>International Rooms 4,5</i>
	Track C: Dimensionality Reduction	<i>International Rooms 1,2,3</i>
	Track D: Statistical Methods	<i>International 6</i>
1540-1600	Coffee Break (provided)	
1600-1740	Track A: Nearest Neighbor & Metric Learning	<i>International Rooms 7,8,9,10</i>
	Track B: General Methods	<i>International Rooms 4,5</i>
	Track C: Transfer Learning	<i>International Rooms 1,2,3</i>
	Track D: Statistical Learning and Inference	<i>International 6</i>
1600-1800	Registration Open 1600-1800	
1800-1900	IMLS Annual Business Meeting	<i>International 7,8,9,10</i>
1930-2000	Poster Setup	<i>Skyline - 10th Floor</i>
2000-2200	Poster Session	<i>Skyline -10th Floor</i>
Wednesday, June 19, 2013		
ICML Main Conference		
0800-1200	Registration Open 0730-1200	
0830-1000	Keynote Speaker: Vincent Vanhoucke	<i>International 7,8,9,10</i>
1000-1030	Coffee Break (provided)	
1030-1210	Track A: Invited Orals	<i>International Rooms 7,8,9,10</i>
	Track B: Optimization	<i>International Rooms 4,5</i>
	Track C: Clustering	<i>International Rooms 1,2,3</i>
	Track D: Learning Theory 1	<i>International 6</i>
1210-1400	Lunch (on your own)	
1400-1540	Track A: Dimensionality Reduction and Semi-Supervised Learning	<i>International Rooms 7,8,9,10</i>
	Track B: Optimization and Integration	<i>International Rooms 4,5</i>
	Track C: Vision	<i>International Rooms 1,2,3</i>
	Track D: Learning Theory 2	<i>International 6</i>
1540-1600	Coffee Break (provided)	

The 30th International Conference on Machine Learning: ICML 2013 Atlanta

1600-1740	Track A: Crowd Sourcing and Large Scale Learning	<i>International Rooms 7,8,9,10</i>
	Track B: Kernel Methods	<i>International Rooms 4,5</i>
	Track C: Matrix Factorization	<i>International Rooms 1,2,3</i>
	Track D: Learning Theory 3	<i>International 6</i>
1600-1800	Registration Open 1600-1800	
1800-2000	ICML Banquet	
1930-2000	Poster Setup	<i>Skyline - 10th Floor</i>
2000-2200	Poster Session	<i>Skyline -10th Floor</i>

Thursday, June 20, 2013

ICML Workshops

0800-1200	Registration Open 0730-1200	
0830-1000	Role of Machine Learning in Transforming Healthcare: Recent progress, Challenges and Opportunities (WHEALTH)	<i>Lobby 508</i>
	Machine Learning for System Identification (WSYSID)	<i>Lobby 405-6</i>
	Machine Learning for Bioacoustics (WBIOAC)	<i>Lobby 503</i>
	Peer Reviewing and Publishing Models (WPEER)	<i>Lobby 504-5</i>
	Robot Learning (WROBL)	<i>Marquis 105</i>
	Divergences and Divergence Learning (WDIV)	<i>Lobby 404</i>
	Numerical Linear Algebra in Machine Learning (WLINALG)	<i>Marquis 103-4</i>
	Infering: Interactions between Inference and Learning (WINFERN)	<i>Lobby 401-3</i>
	Machine Learning with Test-Time Budgets (WTBUDG)	<i>Lobby 506-7</i>
1000-1030	Coffee Break	
1030-1200	Workshops Continue	
1200-1400	Lunch (on your own)	
1400-1530	Workshops Continue	
1530-1600	Coffee Break	
1600-1730	Workshops Continue	

Friday, June 21, 2013

ICML Workshops

0800-1200	Registration Open 0730-1200	
0830-1000	Role of Machine Learning in Transforming Healthcare: Recent progress, Challenges and Opportunities (WHEALTH)	<i>Lobby 508</i>
	Machine Learning for System Identification (WSYSID)	<i>Lobby 405-6</i>
	Machine Learning for Bioacoustics (WBIOAC)	<i>Lobby 503</i>
	Challenges in Representation Learning (WREPL)	<i>Lobby 401-3</i>
	Spectral Learning (WSPECT)	<i>Marquis 103-4</i>
	Machine Learning Meets Crowdsourcing (WCRWD)	<i>Lobby 506-7</i>
	Prediction with Sequential Models (WSEQ)	<i>Lobby 504-5</i>

The 30th International Conference on Machine Learning: ICML 2013 Atlanta

	Reinforcement Learning Competition 2013 (WRLCOMP)	<i>Marquis 105</i>
	Theoretically Grounded Transfer Learning (WTRANS)	<i>Lobby 404</i>
1000-1030	Coffee Break	
1030-1200	Workshops Continue	
1200-1400	Lunch (on your own)	
1400-1530	Workshops Continue	
1530-1600	Coffee Break	
1600-1730	Workshops Continue	

Platinum Sponsor



Gold Sponsors



Silver Sponsors



Sponsors



Supporting Organization



BROWN

Best Paper Awards



Exhibitors and Other Supporters



Welcome from the General Chair

ICML attendees,

Welcome to Atlanta and the International Conference on Machine Learning! This year's conference is the 30th in the series and 2013 marks the 20th anniversary of the "ICML" name.

This year's conference included, for the first time, three reviewing cycles. By spreading the reviewing process over a longer window, we hoped to encourage researchers to bring us their best work as it was ready for dissemination. We believe we were successful, as this year's proceedings includes 282 outstanding articles from all areas of the machine learning field. All articles are published in the Journal of Machine Learning Research (JMLR) as Volume 28 of their Workshop and Conference Proceedings series.

In addition to the top-quality technical program including both oral and poster presentations, the conference includes the full gamut of special events ICML is known for:

- Invited speakers Carlos Guestrin, Santosh Vempala, and Vincent Vanhoucke sharing their insights on exciting machine-learning topics.
- A tutorial program with eight tutorials in cutting-edge areas of machine learning by the experts in these subfields.
- A workshop program with 17 workshops designed to encourage in-depth exploration of emerging applications and techniques.
- A sponsored breakfast and networking opportunity for Women in Machine Learning.
- A shared workshop with the NAACL conference (North American Chapter of the Association for Computational Linguistics).
- An attendee banquet.
- Best paper awards.
- Student scholarships, including sponsored named scholarships for select student participants.

I'd like to acknowledge the efforts of some of the people who dedicated a tremendous amount of their time and vision to bringing this conference together. Sanjoy Dasgupta and David McAllester served as program chairs and created the new 3-cycle system. They brought their world-class research sensibilities to overseeing the construction of a terrific technical program. The 79 area chairs, listed elsewhere, bore the brunt of the adjustment to the new reviewing system. Nonetheless, they handled the task professionally and ensured high quality reviews in conjunction with the 721 members of the program committee. Amir Globerson and Fei Sha, as publications co-chairs, were in charge of producing the proceedings volume in conjunction with JMLR. Peter Stone and Geoff Gordon were responsible

for tutorials and workshops, respectively. In essence, each ran a top-tier mini-machine learning conference, making sure the breadth of the field was well represented in each presentation format.

Tucker Balch took on the task of local arrangements chair; every element of the organization of the conference passed through him at some point in the process. He did a fantastic job of guiding his team---including Sharon Crouch (event manager), April Foster (web), and Maria Hybinette (registration)---as they handled a constant stream of administrative challenges. Charles Isbell worked with Tucker as local arrangements co-chair, taking on additional duties as sponsorship chair. In this role, he worked closely with governmental and industrial partners to guarantee a successful event. Jacob Eisenstein (Student Volunteer Chair) made sure we could fund as many students as possible so they could participate in the conference.

Jingrui He (Publicity) was responsible for spreading the word and keeping the community informed about the conference. I'd also like to acknowledge William Cohen, IMLS President, for serving as the connection between ICML-2013 and past and future ICML conferences. Among his many talents is the ability to recruit former and current colleagues to help run the conference.

All of these people worked hard because they believe that a healthy machine learning conference is important for a healthy machine learning research community. On behalf of all of us, we hope you enjoy the conference!

Michael L. Littman
General Chair, ICML-2013

Organizing Committee

Michael Littman, General Chair
Brown University

Sanjoy Dasgupta, Program Co-Chair
University of California, San Diego

David McAllester, Program Co-Chair
Toyota Technological Institute at Chicago

Tucker Balch, Local Co-Chair
Georgia Institute of Technology

Charles Isbell, Local Co-Chair
Georgia Institute of Technology

Geoff Gordon, Workshop Chair
Carnegie Mellon University

Peter Stone, Tutorial Chair
The University of Texas at Austin

Maria Hybinette, Registration Chair
University of Georgia

Jingrui He, Publicity Chair
Stevens Institute of Technology

Jacob Eisenstein, Student Volunteer Chair
Georgia Institute of Technology

Amir Globerson, Publications Co-Chair
The Hebrew University of Jerusalem

Fei Sha, Publications Co-Chair
University of Southern California

Sharon Crouch, Event Manager

Alicia Richart, Administrative Support

April Foster, Webmaster

Area Chairs

Aarti Singh	Gert Lanckriet	Peter Auer
Alan Fern	Gregory Shakhnarovich	Peter Grunwald
Alekh Agarwal	Honglak Lee	Phil Long
Alex Ihler	Jeff Bilmes	Ran El-Yaniv
Amir Globerson	Jerry Zhu	Ran Gilad-Bachrach
Andreas Krause	Joelle Pineau	Raquel Urtasun
Animashree Anandkumar	John Platt	Rich Caruana
Arindam Banerjee	Jure Leskovec	Rich Sutton
Ashutosh Saxena	Kamalika Chaudhuri	Rich Zemel
Ben Taskar	Karsten Borgwardt	Risi Kondor
Brian Kulis	Kilian Weinberger	Ron Parr
Charles Elkan Koby Crammer		Ruslan Salakhutdinov
Constantine Caramanis	Lihong Li	SVN Vishwanathan
Corinna Cortes	Lise Getoor	Satyen Kale
Csaba Szepesvari	Luke Zettlemoyer	Shai Ben-David
Cynthia Rudin	Marc'Aurelio Ranzato	Shie Mannor
Dale Schuurmans	Marina Meila	Thorsten Joachims
Dan Roth	Maya Gupta	Tobias Scheffer
Daniel Hsu	Mehryar Mohri	Tong Zhang
David Sontag	Miroslav Dudik	Ulrike von Luxburg
Drew Bagnell	Nati Srebro	Yann Le Cun
Edoardo Airoldi	Nicolo Cesa-Bianchi	Yisong Yue
Elad Hazan	Nina Balcan	Yoav Freund
Emily Fox	Ofer Dekel	Yoshua Bengio
Eyal Amir	Ohad Shamir	Zico Kolter
Fei Sha	Pascal Poupart	
Francis Bach	Percy Liang	

Information

Spotlight Talks

To fit all of our presentations into the time allotted we continue the tradition of spotlight talks. You will notice that some papers have full length talks (20 minutes) while other papers have spotlight talks (5 minutes). Please be aware of this if you wish to switch sessions midway through.

Poster Sessions

Poster sessions will be held on the evenings of Monday, Tuesday, and Wednesday, starting at 8PM. All papers will have an associated poster, on the same day as their talk. The posters will be presented on the 10th floor, Skyline Level. A cash bar will be available. Please note that you must be 21 years old to purchase alcoholic beverages in Georgia, and identification will be checked.

Internet Access

Free wireless Internet access is provided to ICML attendees. Please look in your goodie bag for a small card labeled "Internet Access" for details.

Workshops

This year there will be three days of workshops: On Sunday, June 16 we are offering two "NLP-related" workshops for the convenience of folks attending NAACL just before ICML. All registrants may attend these workshops.

The remaining workshops will be held on Thursday June 20 and Friday June 21. In order to attend these workshops you must have paid for workshop attendance. You were asked to select workshops when you registered online, but this choice is not binding. If you have paid the fee for the workshops, you may attend any workshops that you like.

Banquet

The conference banquet will be held on Wednesday June 19. If you registered for the main conference or the workshops you may attend the banquet. Otherwise you must purchase a separate ticket to the banquet at the registration desk.

ICML Tutorials, Sunday June 16

Deep Learning

Yann Lecun, Marc'Aurelio Ranzato
0830-1200: International 7, 8, 9

Multi-Target Prediction

Willem Waegeman, Krzysztof Dembczynski and Eyke Hullermeier
0830-1200: International 10

Submodularity in Machine Learning: New Directions

Andreas Krause, Stefanie Jegelka
0830-1200: Marqius 103, 104, 105
Tutorial website: submodularity.org

Tensor Decomposition Algorithms for Latent Variable Model Estimation

Anima Anandkumar, Daniel Hsu, Sham M. Kakade
0830-1200: International 1, 2, 3
Tutorial website: <http://cseweb.ucsd.edu/~djhsu/tensor-tutorial/>

Copulas in Machine Learning

Gal Elidan
1400-1730: International 7, 8

Music Information Research Based on Machine Learning

Masataka Goto and Kazuyoshi Yoshii
1400-1730: International 10
Tutorial website: <http://staff.aist.go.jp/m.goto/ICML2013/tutorial.htm>

Topological Data Analysis

Primož Skraba and Sayan Mukherjee
1400-1730: International 9

Discovering Multiple Clustering Solutions: Grouping Objects in Different Views of Data

Emmanuel Muller, Stephan Gunnemann, Ines Farber, Thomas Seidl
1400-1730: international 1, 2, 3
Tutorial website: <http://dme.rwth-aachen.de/en/DMCS>

NLP-Related Workshops, Sunday June 16

Note: See also the workshop program on Thursday and Friday

Deep Learning for Audio, Speech, and Language Processing (WDLASL)

Brian Kingsbury, Tara N. Sainath, Li Deng, Andrew Senior

0830-1730: International 6

Workshop website: <https://sites.google.com/site/deeplearningicml2013/>

Structured Learning: Inferring Graphs from Structured and Unstructured Inputs (WSTRUC)

Hal Daumé III, Evgeniy Gabrilovich, Lise Getoor, Kevin Murphy

0830-1730: International 4, 5

Workshop website: <https://sites.google.com/site/slgworkshop2013/>

Keynote: Carlos Guestrin

Monday, June 17, 2013, 8:30 to 10:00

Machine Learning at Scale with GraphLab

International 7, 8, 9, 10

Abstract: Today, machine learning (ML) methods play a central role in industry and science. The growth of the Web and improvements in sensor data collection technology have been rapidly increasing the magnitude and complexity of the ML tasks we must solve. This growth is driving the need for scalable, parallel ML algorithms that can handle "Big Data." In this talk, we will focus on:

1. Examining common algorithmic patterns in distributed ML methods.
2. Qualifying the challenges of implementing these algorithms in real distributed systems.
3. Describing computational frameworks for implementing these algorithms at scale.

In the latter part, we will focus mainly on the GraphLab framework, which naturally expresses asynchronous, dynamic graph computations that are key for state-of-the-art ML algorithms. When these algorithms are expressed in our higher-level abstraction, GraphLab will effectively address many of the underlying parallelism challenges, including data distribution, optimized communication, and guaranteeing sequential consistency, a property that is surprisingly important for many ML algorithms. On a variety of large-scale tasks, GraphLab provides 20-100x performance improvements over Hadoop. In recent months, GraphLab has received many tens of thousands of downloads, and is being actively used by a number of startups, companies, research labs and universities.

Bio: Carlos Guestrin is the Amazon Professor of Machine Learning in Computer Science & Engineering at the University of Washington. He is also the co-founder of GGideaLab, a start up focused on monetizing social networks. Previously, he was a senior researcher at the Intel Research Lab in Berkeley. Carlos received his MSc and PhD in Computer Science from Stanford University in 2000 and 2003, respectively, and a Mechatronics Engineer degree from the Polytechnic School of the University of Sao Paulo, Brazil, in 1998. Carlos' work received awards at a number of conferences and a journal: KDD 2007 and 2010, IPSN 2005 and 2006, VLDB 2004, NIPS 2003 and 2007, UAI 2005, ICML 2005, AISTATS 2010, JAIR in 2007, and JWRPM in 2009. He is also a recipient of the ONR Young Investigator Award, NSF Career Award, Alfred P. Sloan Fellowship, IBM Faculty Fellowship, the Siebel Scholarship and the Stanford Centennial Teaching Assistant Award. Carlos was named one of the 2008 'Brilliant 10' by Popular Science Magazine, received the IJCAI Computers and Thought Award and the Presidential Early Career Award for Scientists and Engineers (PECASE). He is a former member of the Information Sciences and Technology (ISAT) advisory group for DARPA.

Technical Sessions

Monday, June 17, 10:30 to 12:10

International Level

Track A: Deep Learning 1

Session Chair: Marc'Aurelio Ranzato

Monday 1030-1210

- 853, On autoencoder scoring, Hanna Kamyshanska; Roland Memisevic
- 1124, On the difficulty of training Recurrent Neural Networks, Razvan Pascanu; Tomas Mikolov; Yoshua Bengio
- 1125, Maxout Networks, Ian Goodfellow; David Warde-Farley; Mehdi Mirza; Aaron Courville; Yoshua Bengio
- 576, Collaborative hyperparameter tuning, Rémi Bardenet; Mátyás Brendel; Balazs Kegl; Michele Sebag

Spotlight Presentations:

- 136, Learning mid-level representations of objects by harnessing the aperture problem, Roland Memisevic; Georgios Exarchakis
- 274, Approximation properties of DBNs with binary hidden units and real-valued visible units, Oswin Krause; Asja Fischer; Tobias Glasmachers; Christian Igel
- 375, Better Mixing via Deep Representations, Yoshua Bengio; Gregoire Mesnil; Yann Dauphin; Salah Rifai
- 532, Fast dropout training, Sida Wang; Christopher Manning

Track B: Compressed Sensing 1

Session Chair: Zico Kolter

1030-1210

- 210, Feature Selection in High-Dimensional Classification, Mladen Kolar; Han Liu
- 105, Markov Network Estimation From Multi-attribute Data, Mladen Kolar; Han Liu; Eric Xing
- 875, Exact Rule Learning via Boolean Compressed Sensing, Dmitry Malioutov; Kush Varshney
- 118, Sparse Recovery under Linear Transformation, Ji Liu; Lei Yuan; Jieping Ye
- 246, Noisy and Missing Data Regression: Distribution-Oblivious Support Recovery, Yudong Chen; Constantine Caramanis

Track C: Reinforcement Learning 1

Session Chair: Csaba Szepesvari

Monday 1030-1210

- 1149, Learning Policies for Contextual Submodular Prediction, Stephane Ross; Jiaji Zhou; Yisong Yue; Debadepta Dey; Drew Bagnell

- 412, Learning an Internal Dynamics Model from Control Demonstration, Matthew Golub; Steven Chase; Byron Yu
- 423, Safe Policy Iteration, Matteo Pirodda; Marcello Restelli; Alessio Pecorino; Daniele Calandriello
- 755, Temporal Difference Methods for the Variance of the Reward To Go, Aviv Tamar; Dotan Di Castro; Shie Mannor

Spotlight Presentations:

- 465, Value Iteration with incremental representation learning for continuous POMDPs, Sebastian Brechtel; Tobias Gindele; Rüdiger Dillmann
- 39, The Sample-Complexity of General Reinforcement Learning, Tor Lattimore; Marcus Hutter; Peter Sunehag
- 338, Online Feature Selection for Model-based Reinforcement Learning, Trung Nguyen; Zhuoru Li; Tomi Silander; Tze Yun Leong
- 1073, Bayesian Learning of Recursively Factored Environments, Marc Bellemare; Joel Veness; Michael Bowling

Track D: Social Networks

Session Chair: Andreas Krause

Monday 1030-1210

- 1062, Copy or Coincidence? A Model for Detecting Social Influence and Duplication Events, Lisa Friedland; David Jensen; Michael Lavine
- 475, Mixture of Mutually Exciting Processes for Viral Diffusion, Shuang-Hong Yang; Hongyuan Zha
- 172, Dynamic Probabilistic Models for Latent Feature Propagation in Social Networks, Creighton Heaukulani; Ghahramani Zoubin
- 828, Modeling Information Propagation with Survival Theory, Manuel Gomez-Rodriguez; Jure Leskovec; Bernhard Schölkopf

Spotlight Presentations:

- 1123, Learning Triggering Kernels for Multi-dimensional Hawkes Processes, Ke Zhou; Le Song; Hongyuan Zha
- 369, Causal Estimation of Peer Influence Effects, Edward Kao; Panos Toulis; Edoardo Airoldi; Donald Rubin
- 974, Modeling Temporal Evolution and Multiscale Structure in Networks, Tue Herlau; Morten Mørup; Mikkel Schmidt
- 544, Scalable Optimization of Neighbor Embedding for Visualization, Zhirong Yang; Jaakko Peltonen; Samuel Kaski

Technical Sessions

Monday, June 17, 14:00 to 15:40

International Level

Track A: Deep Learning 2

Session Chair: Ruslan Salakhutdinov

Monday 1400-1540

- 925, Learning the Structure of Sum-Product Networks, Robert Gens; Domingos Pedro
- 1129, Deep learning with COTS HPC systems, Adam Coates; Brody Huval; Tao Wang; David Wu; Bryan Catanzaro; Ng Andrew
- 93, Learning and Selecting Features Jointly with Point-wise Gated Boltzmann Machines, Kihyuk Sohn; Guanyu Zhou; Chansoo Lee; Honglak Lee
- 1026, Regularization of Neural Networks using DropConnect, Li Wan; Matthew Zeiler; Sixin Zhang; Yann Le Cun; Rob Fergus

Spotlight Presentations:

- 502, Thurstonian Boltzmann Machines: Learning from Multiple Inequalities, Truyen Tran; Dinh Phung; Svetha Venkatesh
- 279, Iterative Learning and Denoising in Convolutional Neural Associative Memories, Amin Karbasi; Amir Hesam Salavati; Amin Shokrollahi,
- 457, No more pesky learning rates, Tom Schaul; Sixin Zhang; Yann LeCun
- 73, Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures, James Bergstra; Daniel Yamins; David Cox

Track B: Compressed Sensing 2

Session Chair: Tong Zhang

Monday 1400-1540

- 680, Learning Heteroscedastic Models by Convex Programming under Group Sparsity, Arnak Dalalyan; Mohamed Hebiri; Katia Meziani; Joseph Salmon
- 58, Noisy Sparse Subspace Clustering, Yu-Xiang Wang; Huan Xu
- 263, One-Bit Compressed Sensing: Provable Support and Vector Recovery, Sivakant Gopi; Praneeth Netrapalli; Prateek Jain; Aditya Nori
- 403, Smooth Sparse Coding via Marginal Regression for Learning Sparse Representations, Krishnakumar Balasubramanian; Kai Yu; Guy Lebanon

Spotlight Presentations:

- 599, Sparse projections onto the simplex, Anastasios Kyrillidis; Stephen Becker; Volkan Cevher; Christoph Koch
- 1056, Intersecting singularities for multi-structured estimation, Emile Richard; Francis BACH; Jean-Philippe Vert
- 29, Sparse Uncorrelated Linear Discriminant Analysis, Xiaowei Zhang; Delin Chu
- 350, Estimating Unknown Sparsity in Compressed Sensing, Miles Lopes

Track C: Reinforcement learning 2

Session Chair: Lihong Li

Monday 1400-1540

- 955, Concurrent Reinforcement Learning from Customer Interaction Sequences, David Silver
- 100, Modelling Sparse Dynamical Systems with Compressed Predictive State Representations, William Hamilton; Mahdi Milani Fard; Joelle Pineau,
- 1199, Coco-Q: Learning in Stochastic Games with Side Payments, Elizabeth Hilliard; Eric Sodomka; Michael Littman; Amy Greenwald
- 26, Guided Policy Search, Sergey Levine; Vladlen Koltun
- 1069, The Cross-Entropy Method Optimizes for Quantiles, Sergiu Goschin; Ari Weinstein; Michael Littman

Track D: Topic Modeling 1

Session Chair: Emily Fox

Monday 1400-1540

- 617, A Practical Algorithm for Topic Modeling with Provable Guarantees, Sanjeev Arora; Rong Ge; Yonatan Halpern; David Mimno; Ankur Moitra; David Sontag; Yichen Wu; Michael Zhu
- 376, Online Latent Dirichlet Allocation with Infinite Vocabulary, KE ZHAI; Jordan Boyd-Graber
- 76, Gibbs Max-Margin Topic Models with Fast Sampling Algorithms,
- Jun Zhu; Ning Chen; Hugh Perkins; Bo Zhang

Spotlight Presentations:

- 606, Modeling Musical Influence with Topic Models, Uri Shalit; Daphna Weinshall; Gal Chechik
- 1184, Nested Chinese Restaurant Franchise Process: Applications to User Tracking and Document Modeling, Amr Ahmed; Liangjie Hong; Alexander Smola
- 61, Parallel Markov Chain Monte Carlo for Nonparametric Mixture Models, Sinead Williamson; Avinava Dubey; Eric Xing
- 354, MAD-Bayes: MAP-based Asymptotic Derivations from Bayes, Tamara Broderick; Brian Kulis; Michael Jordan
- 801, Topic Model Diagnostics: Assessing Domain Relevance via Topical Alignment, Jason Chuang; Sonal Gupta; Christopher Manning; Jeffrey Heer

Technical Sessions

Monday, June 17, 16:00 to 17:40

International Level

Track A: Deep Learning and Neuroscience

Session Chair: Yoshua Bengio

Monday 1600-1740

- 1051, On the importance of initialization and momentum in deep learning, Ilya Sutskever; James Martens; George Dahl; Geoffrey Hinton
- 1055, A non-IID Framework for Collaborative Filtering with Restricted Boltzmann Machines, Kostadin Georgiev; Preslav Nakov
- 219, Parsing epileptic events using a Markov switching process model for correlated time series, Drausin Wulsin; Emily Fox; Brian Litt
- 611, Exploring the Mind: Integrating Questionnaires and fMRI, Esther Salazar; Ryan Bogdan; Adam Gorka; Ahmad Hariri; Lawrence Carin

Spotlight Presentations:

- 552, Gated Autoencoders with Tied Input Weights, Alain Droniou; Olivier Sigaud
- 696, Simple Sparsification Improves Sparse Denoising Autoencoders in Denoising Highly Corrupted Images, Kyunghyun Cho
- 983, Natural Image Bases to Represent Neuroimaging Data, Ashish Gupta; Murat Ayhan; Anthony Maida
- 658, Direct Modeling of Complex Invariances for Visual Object Features, Ka Yu Hui

Track B: Compressed Sensing 3

Session Chair: Alekh Agarwal

Monday 1400-1540

- 693, Spectral Compressed Sensing via Structured Matrix Completion, Yuxin Chen; Yuejie Chi
- 870, Sparse PCA through Low-rank Approximations, Dimitris Papailiopoulos; Alexandros Dimakis; Stavros Korokythakis
- 179, Efficient Sparse Group Feature Selection via Nonconvex Optimization, Shuo Xiang; Xiaotong Shen; Jieping Ye
- 500, A General Iterative Shrinkage and Thresholding Algorithm for Non-convex Regularized Optimization Problems, Pinghua Gong; Changshui Zhang; Zhaosong Lu; Jianhua Huang; Jieping Ye
- 876, Robust Sparse Regression under Adversarial Corruption, Yudong Chen; Constantine Caramanis; Shie Mannor

Track C: Reinforcement Learning and Time Series

Session Chair: Joelle Pineau

Monday 1600-1740

- 840, ABC Reinforcement Learning, Christos Dimitrakakis; Nikolaos Tziortziotis
- 393, Mean Reversion with a Variance Threshold, Marco Cuturi; Alexandre d'Aspremont
- 1029, Gaussian Process Kernels for Pattern Discovery and Extrapolation, Andrew Wilson; Ryan Adams

Spotlight Presentations:

- 207, Average Reward Optimization Objective In Partially Observable Domains, Yuri Grinberg; Doina Precup
- 463, Planning by Prioritized Sweeping with Small Backups, Harm van Seijen; Rich Sutton
- 780, Dynamic Covariance Models for Multivariate Financial Time Series, Yue Wu; Jose Miguel Hernandez-Lobato; Ghahramani Zoubin
- 300, Learning Sparse Penalties for Change-point Detection using Max Margin Interval Regression, Toby Hocking; Guillem Rigall; Jean-Philippe VERT; Francis BACH
- 670, Hierarchically-coupled hidden Markov models for learning kinetic rates from single-molecule data, Jan-Willem Van de Meent; Jonathan Bronson; Frank Wood; Ruben Gonzalez, Jr.; Chris Wiggins
- 529, Learning Connections in Financial Time Series, Gartheeban Ganeshapillai; John Guttag; Andrew Lo
- 1042, The Extended Parameter Filter, Yusuf Bugra Erol; Lei Li; Bharath Ramsundar; Russell Stuart
- 563, Transition Matrix Estimation in High Dimensional Time Series, Fang Han; Han Liu

Track D: Topic Modeling 2

Session Chair: Elad Hazan

Monday 1600-1740

- 977, Dependent Normalized Random Measures, Changyou Chen; Vinayak Rao; Yee Whye Teh; Wray Buntine
- 1070, Topic Discovery through Data Dependent and Random Projections, Weicong Ding; Mohammad Hossein Rohban; Prakash Ishwar; Venkatesh Saligrama
- 821, Factorial Multi-Task Learning : A Bayesian Nonparametric Approach, Sunil Gupta; Dinh Phung; Svetha Venkatesh

Spotlight Presentations:

- 1003, Scaling the Indian Buffet Process via Submodular Maximization, Colorado Reed; Ghahramani Zoubin
- 506, A Variational Approximation for Topic Modeling of Hierarchical Corpora, Do-kyum Kim; Geoffrey Voelker; Lawrence Saul
- 1156, Manifold Preserving Hierarchical Topic Models for Quantization and Approximation, Minje Kim; Paris Smaragdis
- 607, Subtle Topic Models and Discovering Subtly Manifested Software Concerns Automatically, Mrinal Das; Suparna Bhattacharya; Chiranjib Bhattacharyya; Gopinath Kanchi
- 852, Latent Dirichlet Allocation Topic Model with Soft Assignment of Descriptors to Words, Daphna Weinshall; Gal Levi; Dmitri Hanukaev
- 692, Efficient Multi-label Classification with Many Labels, Wei Bi; James Kwok
- 242, A Randomized Mirror Descent Algorithm for Large Scale Multiple Kernel Learning, Arash Afkanpour; Andras Gyorgy; Csaba Szepesvari; Michael Bowling
- 112, MILEAGE: Multiple Instance LEarning with Global Embedding, Dan Zhang; Jingrui He; Luo Si; Richard Lawrence

Poster Session in Skyline Room, 10th Floor at 20:00

Keynote: Santosh Vempala
Tuesday, June 18, 2013, 8:30 to 10:00

High-dimensional Sampling Algorithms and their Applications

International 7, 8, 9, 10

Abstract: How efficiently can we solve fundamental problems such as Optimization, Integration, Rounding and Sampling in high dimension? Under appropriate convexity assumptions, these general problems can be solved in time polynomial in the dimension, with sampling playing a central role. In this talk, we survey the state-of-the-art and the main ideas that led to it, including geometric random walks, simulated annealing, isoperimetric inequalities and concentration of measure.

Bio: Vempala attended Carnegie Mellon University, where he received his Ph.D. in 1997 under professor Avrim Blum. In 1997, he was awarded a Miller Fellowship at Berkeley. Subsequently, he was a Professor at MIT in the Mathematics Department, until he moved to Georgia Tech in 2006. His main work has been in the area of theoretical computer science, with particular activity in the fields of algorithms, randomized algorithms, computational geometry, and computational learning theory, including the authorship of books on random projection and spectral methods. Vempala has received numerous awards, including a Guggenheim Fellowship, Sloan Fellowship, and being listed in Georgia Trend's 40 under 40. In 2008, he co-founded the Computing for Good (C4G) program at Georgia Tech.

Technical Sessions

Tuesday, June 18, 10:30 to 12:10

International Level

Track A: Online Learning 1

Session Chair: Ofer Dekel

Tuesday 1030-1210

- 805, Online Kernel Learning with a Near Optimal Sparsity Bound, Lijun Zhang; Rong Jin; Xiaofei He
- 710, On the Generalization Ability of Online Learning Algorithms for Pairwise Loss Functions, Prateek Jain; Bharath Sriperumbudur; Purushottam Kar; Harish Karnick
- 178, Thompson Sampling for Contextual Bandits with Linear Payoffs, Shipra Agrawal; Navin Goyal
- 1189, Online Learning under Delayed Feedback, Pooria Joulani; Andras Gyorgy; Csaba Szepesvari
- 1102, Almost Optimal Exploration in Multi-Armed Bandits, Zohar Karnin; Tomer Koren; Oren Somekh

Track B: Feature Learning

Session Chair: Yoshua Bengio

Tuesday 1030-1210

- 508, Forecastable Component Analysis, Georg Goerg
- 107, Discriminatively Activated Sparselets, Ross Girshick; Hyun Oh Song; Trevor Darrell
- 458, Multi-View Clustering and Feature Learning via Structured Sparsity, Hua Wang; Feiping Nie; Heng Huang
- 126, Connecting the Dots with Landmarks: Discriminatively Learning Domain-Invariant Features for Unsupervised Domain Adaptation, Boqing Gong; Kristen Grauman; Fei Sha

Spotlight Presentations:

- 1202, On Nonlinear Generalization of Sparse Coding and Dictionary Learning, Jeffrey Ho; Yuchen Xie; Baba Vemuri
- 891, Feature Multi-Selection among Subjective Features, Sivan Sabato; Adam Kalai
- 21, Sparsity-Based Generalization Bounds for Predictive Sparse Coding, Nishant Mehta; Alexander Gray
- 788, A Unified Robust Regression Model for Lasso-like Algorithms, Wenzhuo Yang; Huan Xu

Track C: General SVM and Decision Tree Methods

Session Chair: Rich Caruana

Tuesday 1030-1210

- 74, Multi-Class Classification with Maximum Margin Multiple Kernel, Corinna Cortes; Mehryar Mohri; Afshin Rostamizadeh
- 397, Top-down particle filtering for Bayesian decision trees, Balaji Lakshminarayanan; Daniel Roy; Yee Whye Teh
- 77, Cost-Sensitive Tree of Classifiers, Zhixiang Xu; Matt Kusner; Kilian Weinberger; Minmin Chen

Spotlight Presentations:

- 794, On the Statistical Consistency of Algorithms for Binary Classification under Class Imbalance, Aditya Menon; Harikrishna Narasimhan; Shivani Agarwal; Sanjay Chawla
- 790, Quickly Boosting Decision Trees — Pruning Underachieving Features Early, Ron Appel; Thomas Fuchs; Piotr Dollar; Pietro Perona
- 1185, Tree-Independent Dual-Tree Algorithms, Ryan Curtin; William March; Parikshit Ram; David Anderson; Alexander Gray; Charles Isbell
- 767, Loss-Proportional Subsampling for Subsequent ERM, Paul Mineiro; Nikos Karampatziakis
- 961, Saving Evaluation Time for the Decision Function in Boosting: Representation and Reordering Base Learner, Peng Sun; Jie Zhou
- 1157, Safe Screening of Non-Support Vectors in Pathwise SVM Computation, Kohei Ogawa; Yoshiki Suzuki; Ichiro Takeuchi
- 95, Convex formulations of radius-margin based Support Vector Machines, Huyen Do; Alexandros Kalousis
- 111, The Pairwise Piecewise-Linear Embedding for Efficient Non-Linear Classification, Ofir Pele; Ben Taskar; Amir Globerson; Michael Werman

Track D: Spectral Learning and Tensors

Session Chair: Daniel Hsu

Tuesday 1030-1210

- 806, Spectral Learning of Hidden Markov Models from Dynamic and Static Data, Tzu-Kuo Huang; Jeff Schneider
- 146, Learning Linear Bayesian Networks with Latent Variables, Animashree Anandkumar; Daniel Hsu; Adel Javanmard; Sham Kakade
- 1018, Spectral Experts for Estimating Mixtures of Linear Regressions, Arun Tejasvi Chaganty; Percy Liang
- 850, On learning parametric-output HMMs, Aryeh Kontorovich; Boaz Nadler; Roi Weiss

Spotlight Presentations:

- 283, Tensor Analyzers, Yichuan Tang; Ruslan Salakhutdinov ; Geoffrey Hinton
- 439, Unfolding Latent Tree Structures using 4th Order Tensors, Mariya Ishteva; Haesun Park; Le Song

- 448, Hierarchical Tensor Decomposition of Latent Tree Graphical Models, Le Song; Mariya Ishteva; Ankur Parikh; Eric Xing; Haesun Park
- 786, Infinite Positive Semidefinite Tensor Factorization with Application to Music Signal Analysis, Kazuyoshi Yoshii; Ryota Tomioka; Daichi Mochihashi; Masataka Goto

Technical Sessions

Tuesday, June 18, 14:00 to 15:40

International Level

Track A: Online Learning 2

Session Chair: Satyen Kale

Tuesday 1400-1540

- 367, Optimal Regret Bounds for Selecting the State Representation in Reinforcement Learning, Odalric-Ambrym Maillard; Phuong Nguyen; Ronald Ortner; Daniil Ryabko
- 89, Combinatorial Multi-Armed Bandit: General Framework, Results and Applications, Wei Chen; Yajun Wang; Yang Yuan
- 387, Dynamical Models and tracking regret in online convex programming, Eric Hall; Rebecca Willett
- 833, Better Rates for Any Adversarial Deterministic MDPs, Ofer Dekel; Elad Hazan

Spotlight Presentations:

- 169, Multiple Identifications in Multi-Armed Bandits, Sebastian Bubeck; Tengyao Wang; Nitin Viswanathan
- 37, Gossip-based distributed stochastic bandit algorithms, Balazs Szorenyi; Robert Busa-Fekete; Istvan Hegedus; Robert Ormandi; Mark Jelasi; Balazs Kegl
- 247, Dual Averaging and Proximal Gradient Descent for Online Alternating Direction Multiplier Method, Taiji Suzuki

Track B: Structured Labeling

Session Chair: Simon Lacoste-Julien

Tuesday 1400-1540

- 306, Learning from Human List Production, Kwang-Sung Jun; Jerry Zhu; Burr Settles; Timothy Rogers
- 346, A Structural SVM Based Approach for Optimizing Partial AUC, Harikrishna Narasimhan; Shivani Agarwal
- 106, A Machine Learning Framework for Programming by Example, Aditya Menon; Omer Tamuz; Sumit Gulwani; Butler Lampson; Adam Kalai
- 454, Convex Adversarial Collective Classification, Mohamad Ali Torkamani; Daniel Lowd

Spotlight Presentations:

- 954, Learning Convex QP Relaxations for Structured Prediction, Jeremy Jancsary; Sebastian Nowozin; Carsten Rother
- 117, Fixed-Point Model For Structured Labeling, Quannan Li; Jingdong Wang; David Wipf; Zhuowen Tu
- 310, A Generalized Kernel Approach to Structured Output Learning, Hachem Kadri; Mohammad Ghavamzadeh; Philippe Preux

- 1049, Optimizing the F-measure in Multi-label Classification: Plug-in Rule Approach versus Structured Loss Minimization, Krzysztof Dembczynski; Wojciech Kotlowski; Arkadiusz Jachnik; Willem Waegeman; Eyke Huellermeier

Track C: Dimensionality Reduction

Session Chair: Nina Balcan

Tuesday 1400-1540

- 141, Principal Component Analysis on non-Gaussian Dependent Data, Fang Han; Han Liu
- 1103, Deep Canonical Correlation Analysis, Galen Andrew; Jeff Bilmes; Raman Arora; Karen Livescu
- 654, Canonical Correlation Analysis based on Hilbert-Schmidt Independence Criterion and Centered Kernel Target Alignment, Billy Chang; Uwe Kruger; Rafal Kustra; Junping Zhang
- 408, Vanishing Component Analysis, Roi Livni; David Lehavi; Sagi Schein; Hila Nachliely; Shai Shalev-Shwartz; Amir Globerson

Spotlight Presentations:

- 1183, Fast algorithms for sparse principal component analysis based on Rayleigh quotient iteration, Volodymyr Kuleshov
- 215, Efficient Dimensionality Reduction for Canonical Correlation Analysis, Haim Avron; Christos Boutsidis ; Sivan Toledo ; Anastasios Zouzias
- 205, Adaptive Sparsity in Gaussian Graphical Models , Eleanor Wong; Suyash Awate; P. Thomas Fletcher
- 357, The Most Generative Maximum Margin Bayesian Networks, Robert Peharz; Sebastian Tschatschek; Franz Pernkopf

Track D: Statistical Methods

Session Chair: Anima Anandkumar

Tuesday 1400-1540

- 872, Computation-Risk Tradeoffs for Covariance-Thresholded Regression, Dinah Shender; John Lafferty
- 769, Scalable Simple Random Sampling and Stratified Sampling, Xiangrui Meng
- 1015, The lasso, persistence, and cross-validation Darren Homrighausen; Daniel McDonald
- 889, Consistency versus Realizable H-Consistency for Multiclass Classification, Phil Long; Rocco Servedio

Spotlight Presentations:

- 738, Two-Sided Exponential Concentration Bounds for Bayes Error Rate and Shannon Entropy, Jean Honorio; Jaakkola Tommi
- 1137, Scale Invariant Conditional Dependence Measures, Sashank J Reddi; Barnabas Poczos
- 865, Infinite Markov-Switching Maximum Entropy Discrimination Machines, Sotirios Chatzis

- 1020, Distribution to Distribution Regression,
- Junier Oliva; Barnabas Poczos; Jeff Schneider

Technical Sessions

Tuesday, June 18, 16:00 to 17:40

International Level

Track A: Nearest Neighbor and Metric Learning

Session Chair: Killian Weinberger

Tuesday 1600-1740

- 746, Entropic Affinities: Properties and Efficient Numerical Computation, Max Vladymyrov; Miguel Carreira-Perpinan
- 86, Learning Hash Functions Using Column Generation, Xi Li; Guosheng Lin; Chunhua Shen; Anton van den Hengel; Anthony Dick
- 415, Robust Structural Metric Learning, Daryl Lim; Gert Lanckriet; Brian McFee
- 785, Revisiting the Nystrom method for improved large-scale machine learning, Alex Gittens; Michael Mahoney

Spotlight Presentations:

- 743, That was fast! Speeding up NN search of high dimensional distributions., Emanuele Coviello; Adeel Mumtaz; Antoni Chan; Gert Lanckriet
- 339, Stochastic k-Neighborhood Selection for Supervised and Unsupervised Learning, Daniel Tarlow; Kevin Swersky; Ilya Sutskever; Laurent Charlin; Rich Zemel
- 1126, Predictable Dual-View Hashing, Mohammad Rastegari; Jonghyun Choi; Shobeir Fakhraei; Daume Hal; Larry Davis
- 526, A unifying framework for vector-valued manifold regularization and multi-view learning, Minh Ha Quang; Loris Bazzani; Vittorio Murino

Track B: General Methods

Session Chair: Shai Ben-David

Tuesday 1600-1740

- 183, Domain Adaptation for Sequence Labeling Tasks with a Probabilistic Language Adaptation Model, Min Xiao; Yuhong Guo
- 186, Maximum Variance Correction with Application to A* Search, Wenlin Chen; Kilian Weinberger; Yixin Chen
- 270, Learning with Marginalized Corrupted Features, Laurens van der Maaten; Minmin Chen; Stephen Tyree; Kilian Weinberger

Spotlight Presentations:

- 285, Scaling Multidimensional Gaussian Processes using Projected Additive Approximations, Elad Gilboa; Yunus Saatci; John Cunningham
- 1131, Nonparametric Mixture of Gaussian Processes with Constraints, James Ross; Jennifer Dy
- 967, Fast Dual Variational Inference for Non-Conjugate Latent Gaussian Models, Mohammad Emteyaz Khan; Aleksandr Aravkin; Michael Friedlander; Matthias Seeger

- 484, Gaussian Process Vine Copulas for Multivariate Dependence, David Lopez-Paz; Jose Miguel Hernandez-Lobato; Ghahramani Zoubin
- 1060, Structure Discovery in Nonparametric Regression through Compositional Kernel Search, David Duvenaud; James Lloyd; Roger Grosse; Joshua Tenenbaum; Ghahramani Zoubin
- 594, Sequential Bayesian Search, Zheng Wen; Branislav Kveton; Brian Eriksson; Sandilya Bhamidipati
- 923, Kernelized Bayesian Matrix Factorization, Mehmet Gönen; Suleiman Khan; Samuel Kaski
- 583, SADA: A General Framework to Support Robust Causation Discovery, Ruichu Cai; Zhenjie Zhang; Zhifeng Hao

Track C: Transfer Learning

Session Chair: Tobias Scheffer

Tuesday 1600-1740

- 13, Domain Generalization via Invariant Feature Representation, Krikamol Muandet; David Balduzzi; Bernhard Schölkopf
- 868, A PAC-Bayesian Approach for Domain Adaptation with Specialization to Linear Classifiers, Pascal Germain; Amaury Habrard; François Laviolette; Emilie Morvant
- 657, Sparse coding for multitask and transfer learning, Andreas Maurer; Massi Pontil; Bernardino Romera-Paredes
- 91, Bayesian Games for Adversarial Regression Problems,
- Michael Groóhans; Christoph Sawade; Michael Brückner; Tobias Scheffer

Spotlight Presentations:

- 372, Joint Transfer and Batch-mode Active Learning, Rita Chattopadhyay; Wei Fan; Ian Davidson; Sethuraman Panchanathan; Jieping Ye
- 1187, Multilinear Multitask Learning, Bernardino Romera-Paredes; Hane Aung; Nadia Bianchi-Berthouze; Massimiliano Pontil
- 965, Stability and Hypothesis Transfer Learning, Ilja Kuzborskij; Francesco Orabona
- 695, Multi-Task Learning with Gaussian Matrix Generalized Inverse Gaussian Model, Ming Yang; Li Yingming; Zhang Zhongfei (Mark)

Track D: Statistical Learning and Inference

Session Chair: Amir Globerson

Tuesday 1600-1740

- 360, Convex Relaxations for Learning Bounded-Treewidth Decomposable Graphs, Sesh Kumar K. S.; Francis Bach
- 760, ∞ SVM for Learning with Label Proportions, Felix Yu; Dong Liu; Sanjiv Kumar; Jebara Tony; Shih-Fu Chang
- 1105, Consistency of Online Random Forests, Misha Denil; David Matheson; De Freitas Nando

Spotlight Presentations:

- 259, Inference algorithms for pattern-based CRFs on sequence data, Rustem Takhanov; Vladimir Kolmogorov
- 389, Relaxed expectation propagation based on l1-penalized KL minimization, Yuan Qi; Yandong Guo
- 324, A Fast and Exact Energy Minimization Algorithm for Cycle MRFs, Huayan Wang; Koller Daphne
- 571, Subproblem-Tree Calibration: A Unified Approach to Max-Product Message Passing, Huayan Wang; Koller Daphne
- 993, Approximate Inference in Collective Graphical Models, Daniel Sheldon; Tao Sun; Akshat Kumar; Tom Dietterich
- 628, An Adaptive Learning Rate for Stochastic Variational Inference, Rajesh Ranganath; Chong Wang; Blei David; Eric Xing
- 1101, The Bigraphical Lasso, Alfredo Kalaitzis; John Lafferty; Neil Lawrence
- 1030, Anytime Representation Learning, Zhixiang Xu; Matt Kusner; Gao Huang; Kilian Weinberger

IMLS Annual Business Meeting in Track A Room at 18:00

Poster Session in Skyline Room, 10th Floor at 20:00

Keynote: Vincent Vanhoucke
Wednesday, June 18, 2013, 8:30 to 10:00

Acoustic Modeling and Deep Learning for Speech Recognition

International 7,8,9,10

Abstract: Over the past few years, advances in deep learning have triggered a mini-revolution in the field of acoustic modeling for automatic speech recognition. Acoustic modeling has evolved largely independently from machine learning for many years, developing its own set of unique techniques in an increasingly complex and specialized ecosystem. The success of deep learning has forced the community to rethink many long held assumptions about what matters to speech recognition accuracy: what are the roles of discriminative learning, speaker adaptation, noise robustness and feature engineering? Can we perform unsupervised, semi-supervised, and transfer learning effectively? How much and what type of data can we really use? More importantly, this development is providing the machine learning and speech recognition communities with an opportunity to reconnect around a familiar set of basic tools and methods. In this talk, I will provide an overview of these recent developments and attempt to paint a picture of what new opportunities lie ahead.

Bio: Vincent Vanhoucke is a Research Scientist at Google. He leads the speech recognition quality effort for Google Search by Voice. He holds a Ph.D. in Electrical Engineering from Stanford University and a Diplôme d'Ingénieur from the Ecole Centrale Paris.

Technical Sessions

Wednesday, June 19, 10:30 to 12:10

International Level

Track A: Invited Orals

Session Chair: Michael Littman

Wednesday 1030-1210

- Classic Paper Prize Talk: Semi-supervised learning using gaussian fields and harmonic functions, ICML 2003, by Xiaojin (Jerry) Zhu, Zoubin Ghahramani, and John Lafferty
- Classic Paper Prize Talk: Online convex programming and generalized infinitesimal gradient ascent, ICML 2003, by Martin Zinkevich
- What do we learn from Kaggle competitions? Ben Hamner
- Machine Learning and Natural Language Processing, Percy Liang
- 394, Large-Scale Bandit Problems and KWIK Learning, Jacob Abernethy; Kareem Amin; Michael Kearns; Moez Draief

Track B: Optimization

Session Chair: Fei Sha

Wednesday 1030-1210

- 48, Stochastic Gradient Descent for Non-smooth Optimization: Convergence Results and Optimal Averaging Schemes, Ohad Shamir; Tong Zhang
- 232, Optimal rates for stochastic convex optimization under Tsybakov noise condition, Aaditya Ramdas; Aarti Singh
- 918, Fast Semidifferential-based Submodular Function Optimization, Rishabh Iyer; Stefanie Jegelka; Jeff Bilmes
- 612, A proximal Newton framework for composite minimization: Graph learning without Cholesky decompositions and matrix inversions, Quoc Tran Dinh; Anastasios Kyrillidis; Volkan Cevher

Spotlight Presentations:

- 1008, Mini-Batch Primal and Dual Methods for SVMs, Martin Takac; Avleen Bijral; Peter Richtarik; Nati Srebro
- 53, Stochastic Alternating Direction Method of Multipliers, Hua Ouyang; Niao He; Long Tran; Alexander Gray
- 884, Optimization with First-Order Surrogate Functions, Julien Mairal
- 41, Fast Probabilistic Optimization from Noisy Gradients, Philipp Hennig

Track C: Clustering

Session Chair: Maya Gupta

Wednesday 1030-1210

- 691, A Local Algorithm for Finding Well-Connected Clusters, Silvio Lattanzi; Vahab Mirrokni; Zeyuan Allen Zhu
- 550, Monochromatic Bi-Clustering, Sharon Wulff; Ruth Uerner; Shai Ben-David

- 424, Constrained fractional set programs and their application in local clustering and community detection, Thomas Bleher; Shyam Sundar Rangapuram; Simon Setzer; Matthias Hein
- 987, Breaking the Small Cluster Barrier of Graph Clustering, Nir Ailon; Yudong Chen; Huan Xu

Spotlight Presentations:

- 555, Strict Monotonicity of Sum of Squares Error and Normalized Cut in the Lattice of Clusterings, Nicola Rebagliati
- 190, Clustering and Learning Behaviors using a Sparse Latent Space, Lui Montesano; Manuel Lopes; Javier Almingol
- 545, Precision-recall space to correct external indices for biclustering, Blaise Hanczar; Mohamed Nadif
- 1171, Semi-supervised Clustering by Input Pattern Assisted Pairwise Similarity Matrix Completion, Jinfeng Yi; Rong Jin; Qi Qian; Anil Jain

Track D: Learning Theory 1

Session Chair: Phil Long

Wednesday 1030-1210

- 644, Margins, Shrinkage and Boosting, Matus Telgarsky
- 849, Sharp Generalization Error Bounds for Randomly-projected Classifiers, Robert Durrant; Ata Kaban
- 63, Risk Bounds and Learning Algorithms for the Regression Approach to Structured Output Prediction, Sébastien Giguère; Francois Laviolette; Mario Marchand; Khadidja Sylla
- 898, Collective Stability and Structured Prediction: Generalization from One Example, Ben London; Bert Huang; Ben Taskar; Lise Getoor

Spotlight Presentations:

- 54, Hierarchical Regularization Cascade for Joint Learning, Alon Zweig; Daphna Weinshall
- 443, Learning Fair Representations, Rich Zemel; Yu Wu; Kevin Swersky; Toniann Pitassi; Cynthia Dwork
- 173, Differentially Private Learning with Kernels, Prateek Jain; Abhradeep Thakurta
- 461, Rounding Methods for Discrete Linear Classification, Yann Chevaleyre; Frederick Koriche; Jean-Daniel Zucker

Technical Sessions

Wednesday, June 19, 14:00 to 15:40

International Level

Track A: Dimensionality Reduction and Semi-Supervised Learning

Session Chair: Jerry Zhu

Wednesday 1400-1540

- 27, Squared-loss Mutual Information Regularization, Gang Niu; Wittawat Jitkrittum; Bo Dai, ; Hirotaka Hachiya; Masashi Sugiyama
- 509, Ellipsoidal Multiple Instance Learning, Gabriel Krumpal; Cheng Soon Ong; Joachim Buhmann
- 933, Infinitesimal Annealing for Training Semi-Supervised Support Vector Machine, Kohei Ogawa; Motoki Imamura; Ichiro Takeuchi; Masashi Sugiyama
- 1108, Sparse Gaussian Conditional Random Fields: Algorithms, and Application to Energy Forecasting, Matt Wytoczek; Zico Kolter
- 1198, Adaptive Hamiltonian and Riemann Manifold Monte Carlo, Ziyu Wang; Shakir Mohamed; De Freitas Nando

Track B: Optimization and Integration

Session Chair: David Sontag

Wednesday 1400-1540

- 487, Stochastic Simultaneous Optimistic Optimization, Michal Valko; Alexandra Carpentier; Remi Munos
- 36, Block-Coordinate Frank-Wolfe Optimization for Structural SVMs, Simon Lacoste-Julien; Martin Jaggi; Mark Schmidt; Patrick Pletscher
- 656, Taming the Curse of Dimensionality: Discrete Integration by Hashing and Optimization, Stefano Ermon; Carla Gomes; Ashish Sabharwal; Bart Selman

Spotlight Presentations:

- 1116, Expensive Function Optimization with Stochastic Binary Outcomes, Matthew Tesch; Jeff Schneider; Howie Choset
- 1047, $O(\log T)$ Projections for Stochastic Optimization of Smooth and Strongly Convex Functions, Lijun Zhang; Tianbao Yang; Rong Jin; Xiaofei He
- 275, Revisiting Frank-Wolfe: Projection-Free Sparse Convex Optimization, Martin Jaggi
- 1036, Algorithms for Direct 0-1 Loss Optimization in Binary Classification, Tan Nguyen; Scott Sanner
- 488, Toward Optimal Stratification for Stratified Monte-Carlo Integration, Alexandra Carpentier; Remi Munos

Track C: Vision

Session Chair: Brian Kulis

Wednesday 1400-1540

- 10, An Optimal Policy for Target Localization with Application to Electron Microscopy, Raphael Sznitman; Aurelien Lucchi; Peter Frazier; Bruno Jedynak; Pascal Fua
- 1115, Fast Image Tagging, Minmin Chen; Alice Zheng; Kilian Weinberger
- 348, An Efficient Posterior Regularized Latent Variable Model for Interactive Sound Source Separation, Nicholas Bryan; Gautham Mysore

Spotlight Presentations:

- 904, Max-Margin Multiple-Instance Dictionary Learning, Xinggang Wang; Zhuowen Tu
- 761, Parameter Learning and Convergent Inference for Dense Random Fields, Philipp Kraehenbuehl; Vladlen Koltun
- 418, Can We Recognize Tiger by Bus Images? _ Robust and Discriminative Self-Taught Image Categorization, Hua Wang; Feiping Nie; Heng Huang
- 886, Learning Spatio-Temporal Structure from RGB-D Videos for Human Activity Detection and Forecasting, Hema Koppula; Ashutosh Saxena
- 15, A Spectral Learning Approach to Range-Only SLAM, Byron Boots; Geoff Gordon
- 726, Non-Linear Stationary Subspace Analysis with Application to Video Classification, Mahsa Baktashmotlagh; Mehrtash Harandi; Abbas Bigdeli; Brian Lovell; Mathieu Salzmann
- 776, On Compact Codes for Spatially Pooled Features, Yangqing Jia; Oriol Vinyals; Trevor Darrell
- 812, Analogy-preserving Semantic Embedding for Visual Object Categorization, Sung Ju Hwang; Kristen Grauman; Fei Sha

Track D: Learning Theory 2

Session Chair: Mehryar Mohri

Wednesday 1400-1540

- 1043, Exploiting Ontology Structures and Unlabeled Data for Learning, Nina Balcan; Avrim Blum; Yishay Mansour
- 952, One-Pass AUC Optimization, Wei Gao; Rong Jin; Shenghuo Zhu; Zhi-Hua Zhou
- 16, Near-Optimal Bounds for Cross-Validation via Loss Stability, Ravi Kumar; Daniel Lokshtanov; Sergei Vassilvitskii; Andrea Vattani,
- 820, Algebraic classifiers: a generic approach to fast cross-validation, parallel training, Michael Izbicki

Spotlight Presentations:

- 1041, Top-k Selection based on Adaptive Sampling of Noisy Preferences, Robert Busa-Fekete; Weiwei Cheng; Paul Weng; Eyke Huellermeier
- 319, Enhanced statistical rankings via targeted data collection, Braxton Osting; Christoph Brune; Stanley Osher

The 30th International Conference on Machine Learning: ICML 2013 Atlanta

- 138, Efficient Ranking from Pairwise Comparisons, Fabian Wauthier; Michael Jordan; Nebojsa Jojic
- 903, Stable Coactive Learning via Perturbation, Karthik Raman; Thorsten Joachims; Pannaga Shivaswamy; Tobias Schnabel

Technical Sessions

Wednesday, June 19, 16:00 to 17:40

International Level

Track A: Crowd Sourcing and Large Scale Learning

Session Chair: Gert Lanckriet

Wednesday 1600-1740

- 96, Optimistic Knowledge Gradient Policy for Optimal Budget Allocation in Crowdsourcing, Xi Chen; Qihang Lin; Dengyong Zhou
- 926, Quantile Regression for Large-scale Applications, Michael Mahoney; Jiyan Yang; Xiangrui Meng
- 621, Distributed training of Large-scale Logistic models, Siddharth Gopal; Yiming Yang
- 570, Label Partitioning For Sublinear Ranking, Jason Weston; Ameesh Makadia; Hector Yee

Spotlight Presentations:

- 213, Human Boosting, Harsh Pareek; Pradeep Ravikumar
- 655, Large-Scale Learning with Less RAM via Randomization, Daniel Golovin; D. Sculley; Brendan McMahan; Michael Young
- 930, Robust Regression on MapReduce, Michael Mahoney; Xiangrui Meng
- 366, Adaptive Task Assignment for Crowdsourced Classification, Chien-Ju Ho; Shahin Jabbari; Jennifer Wortman Vaughan

Track B: Kernel Methods

Session Chair: Percy Liang

Wednesday, 1600-1740

- 751, Local Deep Kernel Learning for Efficient Non-linear SVM Prediction, Cijo Jose; Praseon Goyal; Parv Aggrwal; Manik Varma
- 361, Fastfood – Computing Hilbert Space Expansions in loglinear time, Quoc Le; Tamas Sarlos; Alexander Smola
- 1063, Smooth Operators and an RKHS Integration Approach, Steffen Grunewalder; Gretton Arthur; John Shawe-Taylor
- 895, Domain Adaptation under Target and Conditional Shift, Kun Zhang; Bernhard Schoelkopf; Krikamol Muandet; Zhikun Wang

Spotlight Presentations:

- 171, Learning Optimally Sparse Support Vector Machines, Andrew Cotter; Shai Shalev-Shwartz; Nati Srebro
- 262, A New Frontier of Kernel Design for Structured Data, Kilho Shin
- 384, Characterizing the Representer Theorem, Yaoliang Yu; Hao Cheng; Dale Schuurmans; Csaba Szepesvari
- 683, Covariate Shift in Hilber Space: A Solution Via Sorrogate Kernels, Kai Zhang; Vincent Zheng; Qiaojun Wang; James Kwok; Qiang Yang

Track C: Matrix Factorization

Session Chair: Rich Zemel

Wednesday 1600-1740

- 135, Fast Conical Hull Algorithms for Near-separable Non-negative Matrix Factorization, Abhishek Kumar; Vikas Sindhwani; Prabhanjan Kambadur
- 278, General Functional Matrix Factorization Using Gradient Boosting, Tianqi Chen; Hang Li; Qiang Yang; Yong Yu
- 981, Fast Max-Margin Matrix Factorization with Data Augmentation, Minjie Xu; Jun Zhu; Bo Zhang
- 513, Local Low-Rank Matrix Approximation, Joonseok Lee; Seungyeon Kim; Guy Lebanon; Yoram Singer

Spotlight Presentations:

- 1174, Learning the beta-Divergence in Tweedie Compound Poisson Matrix Factorization Models, Umut Simsekli; Yusuf Kenan Yilmaz; Ali Taylan Cemgil
- 343, ELLA: An Efficient Lifelong Learning Algorithm, Paul Ruvolo; Eric Eaton
- 772, Riemannian Similarity Learning, Li Cheng
- 1120, Multiple-Source Cross Validation, Krzysztof Geras; Charles Sutton
-

Track D: Learning Theory 3

Session Chair: Jeff Bilmes

Wednesday 1600-1740

- 675, Activized Learning with Uniform Classification Noise, Liu Yang; Steve Hanneke
- 316, Efficient Active Learning of Halfspaces: an Aggressive Approach, Alon Gonen; Sivan Sabato; Shai Shalev-Shwartz
- 1091, Selective sampling algorithms for cost-sensitive multiclass prediction, Alekh Agarwa
- 521, Generic Exploration and K-armed Voting Bandits, Tanguy Urvoy; Fabrice Clerot; Raphael Feraud; Sami Naamane

Spotlight Presentations:

- 433, Efficient Semi-supervised and Active Learning of Disjunctions, Nina Balcan; Christopher Berlind; Steven Ehrlich; Yingyu Liang
- 1158, Cost-sensitive Multiclass Classification Risk Bounds, Bernardo Pires; Csaba Szepesvari; Mohammad Ghavamzadeh
- 308, Active Learning for Multi-Objective Optimization, Marcela Zuluaga; Guillaume Sergent; Andreas Krause; Markus Poeschel
- 90, Near-optimal Batch Mode Active Learning and Adaptive Submodular Optimization, Yuxin Chen; Andreas Krause

ICML Banquet at 18:00

Poster Session in Skyline Room, 10th Floor at 20:00

ICML Workshops, Thursday June 20

Role of Machine Learning in Transforming Healthcare: Recent progress, Challenges and Opportunities (WHEALTH)

Noémie Elhadad, Faisal Farooq, Misha Pavel, Suchi Saria, Jimeng Sun, Shipeng Yu
Two days – Continues to Friday: Lobby 508

Machine Learning for System Identification (WSYSID)

Francesco Dinuzzo, Abdeslam Boularias, Lennart Ljung
Two days – Continues to Friday: Lobby 405-6

Machine Learning for Bioacoustics (WBIOAC)

Pr. H. Glotin, Pr. Y. LeCun, Dr. C. Clark, Dr. X. Halkias, Dr. Peter Dugan, Associate Pr. Jérôme Sueur
Two Days – Continues to Friday: Lobby 503

Peer Reviewing and Publishing Models (WPEER)

Andrew McCallum, Aaron Courville
Lobby 504-5

Robot Learning (WROBL)

Anca Dragan, Baris Akgun, Brian Ziebart
Marquis 105

Divergences and Divergence Learning (WDIV)

Rong Jin, Meizhu Liu, Chunhua Shen, Jieping Ye, and Zhi-Hua Zhou
Lobby 404

Numerical Linear Algebra in Machine Learning (WLINALG)

Haim Avron, Christos Boutsidis, Vikas Sindhwani
Marquis 103-4

Inferning: Interactions between Inference and Learning (WINFERN)

Janardhan Rao (Jana) Doppa, Pawan Kumar, Michael Wick, Sameer Singh, Ruslan Salakhutdinov
Lobby 401-3

Machine Learning with Test-Time Budgets (WTBUDG)

Minmin Chen, Matthew Kusner, Venkatesh Saligrama, Kirill Trapeznikov, Kilian Weinberger, Zhixiang (Eddie) Xu
Lobby 506-7

ICML Workshops, Friday June 21

Role of Machine Learning in Transforming Healthcare: Recent progress, Challenges and Opportunities (WHEALTH)

Noémie Elhadad, Faisal Farooq, Misha Pavel, Suchi Saria, Jimeng Sun, Shipeng Yu
Continued from Thursday: Lobby 508

Machine Learning for System Identification (WSYSID)

Francesco Dinuzzo, Abdeslam Boularias, Lennart Ljung
Continued from Thursday: Lobby 405-6

Machine Learning for Bioacoustics (WBIOAC)

Pr. H. Glotin, Pr. Y. LeCun, Dr. C. Clark, Dr. X. Halkias, Dr. Peter Dugan, Associate Pr. Jérôme Sueur
Continued from Thursday: Lobby 503

Challenges in Representation Learning (WREPL)

Ian Goodfellow, Dumitru Erhan, Yoshua Bengio
Lobby 401-3

Spectral Learning (WSPECT)

Byron Boots, Daniel Hsu, Borja Balle, Ankur Parikh
Marquis 103-4

Machine Learning Meets Crowdsourcing (WCRWD)

Paul Bennett, Xi Chen, Qihang Lin, Qiang Liu, John Platt, Dengyong Zhou
Lobby 506-7

Prediction with Sequential Models (WSEQ)

Djalel Benbouzid, Patrick Gallinari, Ludovic Denoyer, Balazs Kegl, Gabriel Dulac-Arnold, Michele Sebag
Lobby 504-5

Reinforcement Learning Competition 2013 (WRLCOMP)

Christos Dimitrakakis, Nikolaos Tziortziotis
Marquit 105

Theoretically Grounded Transfer Learning (WTRANS)

Haitham Bou Ammar, Matthew E. Taylor, Karl Tuyls
Lobby 404

Program Committee

Animashree Anandkumar	Guillaume Desjardins	Oren Anava
Aaron Courville	Guillaum Lecue	Oriol Vinyals
Aarti Singh	Guillaume Obozinski	Simon Osindero
Akshay Balsubramani	Maya Gupta	Takayuki Osogami
Antoni Chan	Graham Taylor	Ozgur Simsek
Abdeslam Boularias	Andras Gyorgy	Pierre Geurts
Abhishek Kumar	Patrick Haffner	Pannaga Shivaswamy
Adam Coates	Haimonti Dutta	Ron Parr
Aditya Gopalan	Hamed Masnadi-Shirazi	Nathan Parrish
Adam Kalai	Jihun Hamm	Partha P. Talukdar
Alan Fern	Hyrum Anderson	Andrea Passerini
Andre Martins	Han Liu	Paul Bennett
Alexander Franks	Hannaneh Hajishirzi	Pavel Laskov
Andrew Gelfand	Hannes Wettig	Dmitry Pechyony
Alborz Geramifard	Harm van Seijen	Percy Liang
Alex Grubb	Matthias Hein	Franz Pernkopf
Arthur Guez	Hema S. Koppula	Peter Grunwald
Andrew Wilson	Heng Luo	Jan Peters
Edoardo Airoldi	Al Hero	Marek Petrik
Ali Jalali	Xuming He	Peter Frazier
Ashish Kapoor	Hui Lin	Peter Gehler
Kareem Amin	Honglak Lee	Phil Long
Aditya K. Menon	Shen-Shyang Ho	Philipp Hennig
Alex Krizhevsky	Hossam Sharara	Philippe Preux
Alain Rakotomamonjy	Huan Wang	Pierre Dupont
Yuan Qi	Hugo Larochelle	Patrick M. Pilarski
Alek Kolcz	Ian Porteous	Ping Li
Alekh Agarwal	Ian Porteus	Piotr Mirowski
Alessandro Lazaric	Ian Lenz	Piotr Mirowski
Alex Graves	Christian Igel	Piyush Rai
Alex Grubb	Alex Ihler	Robert Pless
Alexandre Lacoste	Ilya Sutskever	Po-Ling Loh
Alexandre Bouchard-Côté	Tomas Mikolov	David Poole
Alex Niculescu-Mizil	Ingo Steinwart	Pascal Poupart
Alice Zheng	Givoni Inmar	Prateek Jain
Alon Zweig	Irene Rodriguez	Prakash Balachandran
Alon Gonen	Joris Mooij	Parikshit Shah
Andrew Maas	Jinbo Xu	Paul Vernaza
Amr Ahmed	Jacob D. Abernethy	Qirong Ho
Ameet Talwalkar	Jacob Eisenstein	Qiang Liu

The 30th International Conference on Machine Learning: ICML 2013 Atlanta

Alona Fyshe	Jaesik Choi	Quoc Le
Amir-Massoud Farahmand	Jagadeesh Jagarlamudi	Emile Richard
Amit Daniely	James Bergstra	Razvan Pascanu
Andriy Mnih	James Cussens	Rebecca Hutchinson
Andres Munoz Medina	James Scott	Barbara Rakitsch
Amir Navot	James Kwok	Ran Gilad-Bachrach
Andre Barreto	James F. Saunderson	Ran El-Yaniv
Anirban Dasgupta	Jan Ramon	Marc'Aurelio Ranzato
Ankan Saha	Jascha Sohl-Dickstein	Nathan Ratliff
Ann Nowe	Jason Weston	Richard Wang
Anne-Claire Haury	Jason Ernst	Remi Munos
Antonio Criminisi	Jason K. Johnson	Gholamreza Haffari
Antoine Bordes	Jacob Crandall	Ramon Huerta
Antonio Bahamonde	Jarvis Haupt	Ricardo Silva
Apoorv Agarwal	Jason D. Lee	Richard Socher
Ankur Parikh	John C.	Rina Foygel
Andreas Argyriou	Jean-Philippe Vert	Risi Kondor
Arik Friedman	Enrique Munoz de Cote	Margareta Ackerman
Armand Joulin	Jennifer Listgarten	Ragesh Jaiswal
Art Munson	David Jensen	Rishabh
Arvind Agarwal	Jerry Zhu	Rich Maclin
Anand Sarwate	Jesse Davis	Raman Arora
Andrew Saxe	Jesse Hoey	Richard Nock
Ashutosh Saxena	Jake Hofman	Robby Goetschalckx
Asela Gunawardana	Jun Huan	Roi Livni
Ashique Rupam Mahmood	Jonathan C. Huang	Roland Memisevic
Ashley Llorens	Jiang Su	Romaric Gaudel
Adam D	Jieping Ye	Ron Bekkerman
Amar Subramanya	Jyrki Kivinen	Ronald Ortner
Arthur Asuncion	Jens Kober	Ronan Collobert
Arthur Szlam	James Martens	Rong Ge
Atsuyoshi Nakamura	Jose Miguel Hernandez-Lobato	Rong Jin
Aude Billard	Joseph Modayil	Afshin Rostamizadeh
Peter Auer	Joao Graca	Ryan Adams
Augustin Lefevre	John M. Agosta	Ruslan Salakhutdinov
Austin Abrams	Jon McAuliffe	Rajhans Samdani
Adam White	Jonas Peters	Rich Sutton
Anthony Wirth	Jonathan Chang	Cynthia Rudin
Arthur Choi	Jordan Frank	Raquel Urtasun
Drew Bagnell	Jordan Boyd-Graber	Ruth Urner
Balaji Krishnapuram	John Paisley	Ragupathyraj

Balazs Kegl	Jian Peng	Valluvan
Arindam Banerjee	Joelle Pineau	Rene Vidal
Barnabas Poczos	John Platt	Ryan Tibshirani
Remi Bardenet	James Sharpnack	Steffen Grunewalder
Gabor Bartok	Jo-Anne Ting	Sahand
Borja Balle	Johannes Fuernkranz	Saketha Nath
Brian King	Julia Hockenmaier	Salah Rifai
Byron Boots	Julian McAuley	Samuel Kaski
Benjamin Schrauwen	Julien Mairal	Sujay Sanghavi
Bert Huang	Junming Yin	Sanjiv Kumar
Alina Beygelzimer	Jure Leskovec	Satyen Kale
Bharath	Nikos Karampatziakis	Christoph Sawade
Bianca Zadrozny	Junzhou Huang	Sebastian Bubeck
Jeff A. Bilmes	Kai Yu	Tom Schaul
Jeff A. Bilmes (reviewer)	Mrinal Kalakrishnan	Tobias Scheffer
Blaine Nelson	Karthik Mohan	Bruno Scherrer
John Blitzer	Karsten Borgwardt	Sumit Chopra
Ben London	Karthik Sridharan	Scott Sanner
Brian McFee	Karthik Raman	Scott Yih
Bob Williamson	Aryeh Kontorovich	Michele Sebag
Michael Bowling	Hisashi Kashima	Yevgeny Seldin
Bernardo Pires	Katya Scheinberg	Bart Selman
W. Bradley Knox	Yoshinobu Kawahara	Sergey Feldman
Francesco Orabona	Khalid El-Arini	Sean Gerrish
Brendan O'Connor	Keith Hall	Shai Ben-David
Wolfram Burgard	Kai-Wei Chang	Shakir Mohamed
Robert Busa-Fekete	Kamalika Chaudhuri	Sharon Wulff
Brian Ziebart	Krzysztof Dembczynski	Shaul Markovitch
Chris Amato	Kedar Bellare	Daniel Sheldon
Olivier Cappé	Kee-Eung Kim	Shiau Hong Lim
Constantine Caramanis	Kevin Duh	Shie Mannor
Rich Caruana	Kevin Lai	Shimon Whiteson
Cassio P. de Campos	Kilian Weinberger	Shipeng Yu
Chris Burges	Karen Livescu	Shivani Agarwal
Carlos Diuk	Koby Crammer	Simon Lacoste-Julien
Chris Dyer	Nikos Komodakis	Sinead Williamson
Cedric Archambeau	Koray Kavukcuoglu	Yoram Singer
Brahim Chaib-draa	Wojciech Kotowski	Sivan Sabato
Charles Bergeron	Konstantina Palla	Srivatsan Laxman
Beechung Chen	Kriti Puniyani	Set Sullivant
Cheng Soon Ong	Andreas Krause	Sriraam Natarajan
Alexey Chernov	Kristina Toutanova	Soeren Laue
		Sofus Macskassy

The 30th International Conference on Machine Learning: ICML 2013 Atlanta

David Chiang	Kush Varshney	Sylvie Ong
Chris Hillar	Kaushik Sinha	Sujith Ravi
Christoph H. Lampert	Kevin Swersky	Soumya Ray
Chloé-Agathe Azencott	Brian Kulis	Srujana Merugu
Chong Wang	Kurt Driessens	Alexander Rush
Choon Hui Teo	Kuzman V. Ganchev	Shohei Shimizu
Christian Konig	Lauren Hannah	Taiji Suzuki
Chu Wei	Niels Landwehr	Stan Matwin
Chun-Nam Yu	Laurent El Ghaoui	Stanley Kok
Chun-Nan Hsu	Laurent Charlin	Bin Shen
Chih-Jen Lin	Laurent Jacob	Stefan Wrobel
Claudio Gentile	Lavi Shpigelman	Stefanie Jegelka
Clay Scott	Laurent Charlin	Stephane Canu
Claire Monteleoni	Wee Sun Lee	Stephane Ross
Corinna Cortes	Lek-Heng Lim	Stephen Becker
Christian R. Shelton	Bottou Leon	Stephen Gould
Charles	Hanzhong Liu	Steve Hanneke
John Cunningham	Congcong Li	Masashi Sugiyama
David Hardoon	Lihong Li	Sumit Basu
Daniel Roy	Lily Mihalkova	Sunita Sarawagi
Daniel McDonald	Limin Li	Min Sun
Dale Schuurmans	Lin Xiao	Staal Vinterbo
Alexander D'Amour	Linli Xu	Steve Wright
Daniel Hsu	Christoph Lippert	Csaba Szepesvari
Daniel Zoran	Liu Yang	Istvan Szita
Daniil Ryabko	Liva Ralaivola	Yasuo Tabei
Dan Roth	Lester Mackey	Prasad Tadepalli
David Andrzejewski	Santosh Srivastava	Talya Meltzer
David Balduzzi	Daniel Lowd	Tamir Hazan
David Mimno	Lev Reyzin	Yichuan Tang
David P. Reichert	Le Song	Cheng Tang
David Silver	Luke Zettlemoyer	Vincent Tan
Dayne Freitag	Long Tran-Thanh	Tapas Kanungo
Dhruv Batra	Ulrike von Luxburg	Ben Taskar
David Bradley	Laurens van der Maaten	Gavin Taylor
Jun Zhu	Marco Wiering	Matthew E. Taylor
Daniel Lee	Mark Girolami	Ted Sandler
Doug Downey	Martin Jaggi	Terran Lane
Dean Foster	Sridhar Mahadevan	Ambuj Tewari
		Sivaraman
Deepayan Chakrabarti	Mahdi Milani Fard	Balakrishnan
		Deepak
Delbert Dueck	Mahito Sugiyama	Ramachandran

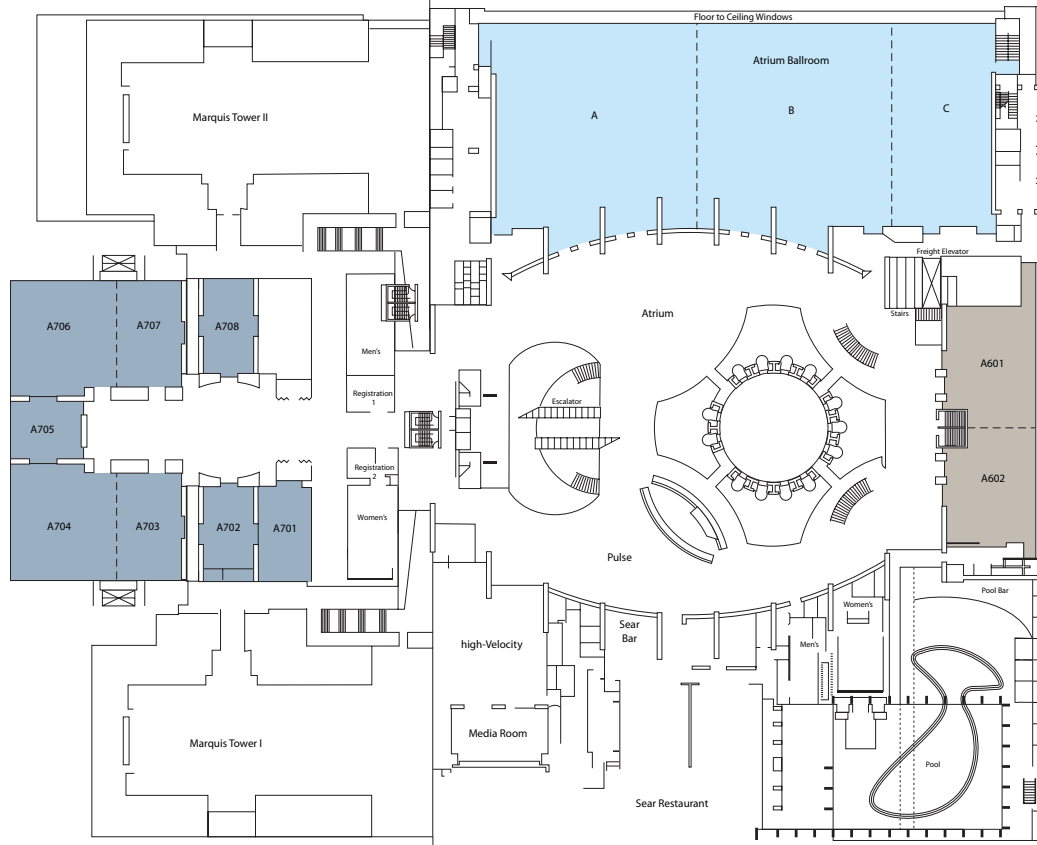
Dengyong Zhou	Maia Fraser	Theodoros Rekatsinas
Paramveer Dhillon	Manuel Gomez-Rodriguez	Thomas Walsh
Dilan Gorur	Marc Toussaint	Tiberio Caetano
David K. Duvenaud	Mark Reid	Tie-Yan Liu
Daniel Lizotte	Mark Schmidt	Tim Van Erven
Daniel Munoz	Ben Marlin	Ivan Titov
Janardhan Rao (Jana)		
Doppa	Martin Zinkevich	Thorsten Joachims
Diane Oyen	Massih-Reza Amini	Tobias Jung
David Helmbold	Mathieu Salzmann	Todd Hester
Doina Precup	Mathieu Sinn	Tom Croonenborghs
Dragos Margineantu	Yiming Ying	Ryota Tomioka
Patrick Nguyen	Matthias Seeger	Tong Zhang
D. Sculley	Matthew Johnson	Lyle Ungar
		Venkat
David Sontag	Mohammad Gheshlaghi Azar	Chandrasekaran
Daniel Tarlow	Misha Bilenko	Veselin Stoyanov
David Yanay	Mustafa Bilgic	Vibhav Gogate
Dumitru Erhan	Marcus	Vikas Sindhwani
Dummy Dummy	Minmin Chen	Vincent Guigue
David B. Dunson	Miguel Carreira-Perpinan	S V N Vishwanathan
Ery Arias-Castro	Andrew McCallum	Vladimir Jojic
Elias Bareinboim	Scott McQuade	Varun Kanade
Emily Fox	Marco Cuturi	Vikash Mansinghka
Elad Hazan	Mark Davenport	Vladimir Kolmogorov
Elad Yom-Tov	Matt Hoffman	Volkan Cevher
Liza Levina	Miroslav Dudik	Volker Roth
Tina Eliassi Rad	Mehmet Gönen	Vladimier Vovk
Charles Elkan	Marina Meila	Vinayak Rao
Elliot Ludvig	Ofer Meshi	Vivek Srikumar
Emine Yilmaz	Marek Grzes	Slobodan Vucetic
Emma Brunskill	Michael Brückner	Vincent Tan
Barbara Engelhardt	Michael Gutmann	Hongning Wang
Eric P. Xing	Michael James	Tomas Werner
Ermin Wei	Michael Rabbat	Rebecca Willett
Tom Erez	Mihajlo Grbovic	Aaron Wilson
Eyal Amir	Michael	David Wingate
Eyal Gofer	Ming-Wei Chang	Kiri Wagstaff
Eyke Huellermeier	Mingyuan Zhou	Wojciech Kotlowski
Fabien Moutarde	Min Xu	Lawrence Cayton
Fabio Vitale	Majid Janzamin	Xiaoli Z. Fern
Fabrice Rossi	Mikko Koivisto	Vojtech Franc
Fei Sha	Marius Kloft	Xinhua Zhang

The 30th International Conference on Machine Learning: ICML 2013 Atlanta

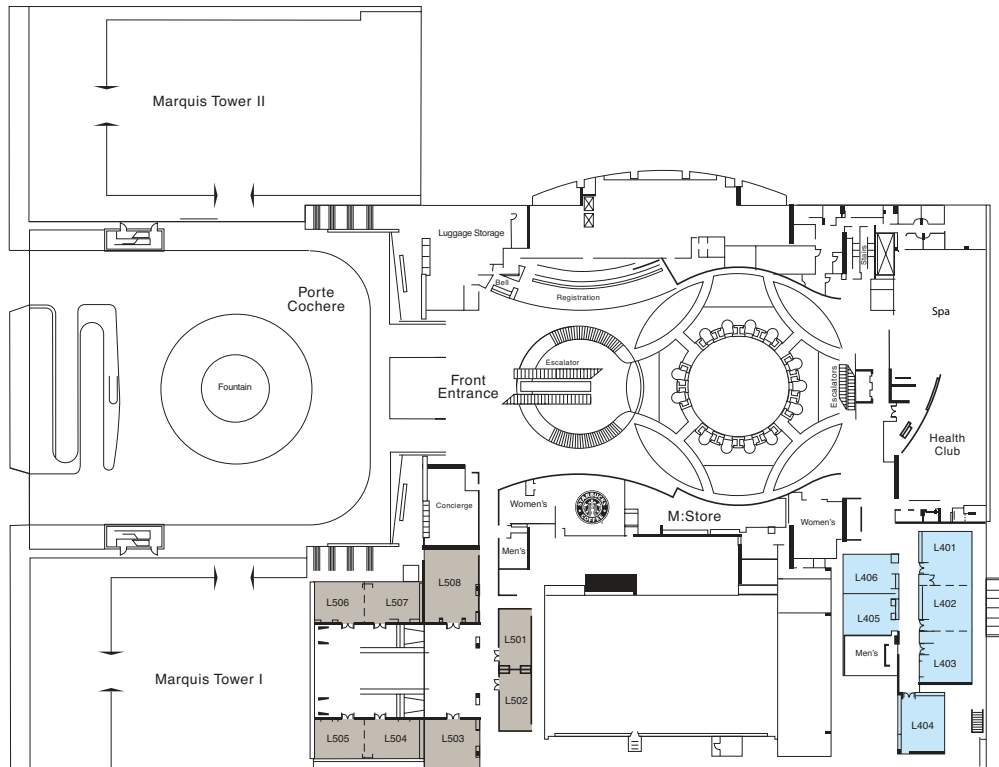
Feng Yan	Michael Mahoney	Xuejun Liao
Fernando Perez-Cruz	Marina Meila	Xiaotong Yuan
Finale Doshi-Velez	Mohammad Ghavamzadeh	Ya Zhang
Daria Sorokina	Mehryar Mohri	Yair Wiener
Florence D'Alche-Buc	Ankur Moitra	Yann Le Cun
Fabian Wauthier	Morteza Alamgir	Yann Dauphin
Francis Bach	Michael Osborne	Yaoliang Yu
Francois Caron	Alessandro Moschitti	Yasin Abbasi-Yadkori
François Laviolette	Huan Xu	Yejin Choi
Bela Frigyük	Matus Telgarsky	Yudong Chen
Frank Wood	Maksims Volkovs	Yoav Freund
Gregoire Montavon	Myunghwan Kim	Yihua Chen
Gal Chechik	Matt Zucker	Yijun Sun
Genevera Allen	Naoki Abe	Yin Lou
Garvesh Raskutti	Nir Ailon	Yisong Yue
Gabor Csardi	Shinichi Nakajima	Yixin Chen
George Konidakis	Noah Smith	Y-Lan Boureau
George P. Burdell		
Greg Druck	Nati Srebro	Yuanqing Lin
Sathya Keerthi	Nakul Verma	Yucheng Low
Gert Lanckriet	Gergely Neu	Yoshua Bengio
Lise Getoor	Nicolas Chapados	Yuchen Zhang
Guido F. Montufar	Nicolas Usunier	Yuhong Guo
Cuartas	Nicolo Cesa-Bianchi	Hyokun Yun
Lee Giles	Shin Matsushima	Yun Jiang
Gilles Blanchard	Nikhil Rasiwasia	Yury Makarychev
Gideon Mann	Nikos Vlassis	Yutian Chen
Jacob Goldberger	Jordi Nin	Matthew D. Zeiler
Daniel Golovin	Nina Balcan	Rich Zemel
Ian Goodfellow	Nan Lin	Teng Zhang
Tsuyoshi Ide	William	Zhao Xu
Amit Goyal	Novi Quadrianto	Xueyuen Zhou
George Papandreou	Sebastian Nowozin	Zhi-Hua Zhou
David Grangier	Odalric-Ambrym Maillard	Zico Kolter
Grigorios Tsoumakas	Ofer Dekel	Shenghuo Zhu
Gregory Shakhnarovich	Ohad Shamir	Elenor Rigby
Guibo Ye	Omid Madani	
Guillaume Bouchard		

The 30th International Conference on Machine Learning: ICML 2013 Atlanta


Atrium Level



Lobby Level



International Level





SEEKING SCIENTIFIC PROGRAMMERS / SOFTWARE ENGINEERS

Brain Corporation (San Diego, CA) seeks talented and highly motivated scientific programmers/software engineers to implement novel algorithms with applications in perception, motor control, and autonomous navigation. The position will require solving interesting novel technical challenges in object recognition, reinforcement learning, complex system design and integration, computer-simulated environments and robotic interfaces.

TEN REASONS TO JOIN BRAIN CORPORATION

Work on the most exciting scientific challenge of the century.

Outstanding team of computational scientists, programmers, and engineers.

Industry-level salary, stock option grant, matching 401K, end of year bonus.

Health, dental, and vision insurance.

Free breakfast and lunch at work; lots of healthy choices.

Free Smart phone for new employees; free gadgets twice per year.

Live 5 minutes from the Pacific Ocean. Surf or dive any day of the year. Bike to work.

Bring your family to sunny San Diego: beaches, zoo, Sea World, Legoland, Wild Animal Park.

Family-friendly events, all-expense-paid holiday retreats, deep-sea fishing.

Attend conferences and we stay an extra day just to ski.

REQUIREMENTS:

Exceptional coding skills.

Experience in software development teams.

Ability to work in collaborative environments.

Clear communication and creativity skills.

Experience implementing algorithms.

Practical experience in machine learning advantageous.

Neural network modeling experience desirable.

Robotics experience helpful.

Embedded systems or GPU experience a plus.

To meet with us at ICML
or to apply, contact us at
jobs+icml@braincorporation.com

We are designing intelligent systems inspired by the secrets of the brain. We are passionate about changing the world with smarter robots that have artificial nervous systems.

PLEASE VISIT WWW.BRAINCORPORATION.COM
FOR MORE INFORMATION

