

Workshop organizers make last-minute changes to their schedule. Download this document again to get the latest changes, or use the [NeurIPS mobile application](#).

Schedule Highlights

Dec. 13, 2019

East Ballroom A, **Safety and Robustness in Decision-making** *Ghavamzadeh, Mannor, Yue, Petrik, Chow*

East Ballroom B, **Learning Meaningful Representations of Life** *Wood, Reshef, Bloom, Snoek, Engelhardt, Linderman, Saria, Wiltschko, Greene, Liu, Lindorff-Larsen, Marks*

East Ballroom C, **Optimal Transport for Machine Learning** *Cuturi, Peyré, Flamary, Suvorikova*

East Exhibition Hall A, **Information Theory and Machine Learning** *Zhao, Song, Han, Choi, Kalluri, Poole, Dimakis, Jiao, Weissman, Ermon*

East Meeting Rooms 11 + 12, **MLSys: Workshop on Systems for ML** *Lakshmiratan, Sen, Gonzalez, Crankshaw, Bird*

East Meeting Rooms 1 - 3, **Perception as generative reasoning: structure, causality, probability** *Rosenbaum, Garnelo, Battaglia, Allen, Yildirim*

East Meeting Rooms 8 + 15, **Minding the Gap: Between Fairness and Ethics** *Rubinov, Kondor, Poulson, Warmuth, Moss, Hagerty*

West 109 + 110, **KR2ML - Knowledge Representation and Reasoning Meets Machine Learning** *Thost, Muise, Talamadupula, Singh, Ré*

West 114 + 115, **Retrospectives: A Venue for Self-Reflection in ML Research** *Lowe, Bengio,*

Pineau, Paganini, Forde, Sodhani, Gupta, Lehman, Henderson, Madan

West 116 + 117, **Competition Track Day 1**
Escalante

West 118 - 120, **Workshop on Federated Learning for Data Privacy and Confidentiality**
Fan, Konečný, Liu, McMahan, Smith, Yu

West 121 + 122, **Machine Learning for the Developing World (ML4D): Challenges and Risks** *De-Arteaga, Coston, Afonja*

West 202 - 204, **Visually Grounded Interaction and Language** *Strub, Das, Wijmans, de Vries, Lee, Suhr, Arad Hudson*

West 205 - 207, **Robust AI in Financial Services: Data, Fairness, Explainability, Trustworthiness, and Privacy** *Oprea, Gal, Moulinier, Chen, Veloso, Kumar, Faruque*

West 208 + 209, **Learning with Rich Experience: Integration of Learning Paradigms** *Hu, Wilson, Finn, Lee, Berg-Kirkpatrick, Salakhutdinov, Xing*

West 211 - 214, **Beyond first order methods in machine learning systems** *Kyrillidis, Berahas, Roosta, Mahoney*

West 215 + 216, **CiML 2019: Machine Learning Competitions for All** *Mendrik, Tu, Guyon, Viegas, LI*

West 217 - 219, **AI for Humanitarian Assistance and Disaster Response** *Gupta, Murphy, Darrell, Heim, Wang, Goodman, Bilinski*

West 220 - 222, **Shared Visual Representations in Human and Machine Intelligence** *Deza, Peterson, Murty, Griffiths*

West 223 + 224, **Workshop on Human-Centric Machine Learning** *Angelov, Oliver, Weller, Rodriguez, Valera, Chiappa, Heidari, Kilbertus*

West 301 - 305, **Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications** *Heckel, Hand, Baraniuk, Bruna, Dimakis, Needell*

West 306, **EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition)** *Parihar, Goldfarb, Srivastava, SHENG*

West Ballroom A, **Machine Learning for Health (ML4H): What makes machine learning in medicine different?** *Beam, Naumann, Beaulieu-Jones, Chen, Finlayson, Alsentzer, Dalca, McDermott*

West Ballroom B, **Meta-Learning** *Calandra, Clavera Gilaberte, Hutter, Vanschoren, Wang*

West Ballroom C, **Biological and Artificial Reinforcement Learning** *Chua, Zannone, Behbahani, Ponte Costa, Clopath, Richards, Precup*

West Exhibition Hall A, **Graph Representation Learning** *Hamilton, van den Berg, Bronstein, Jegelka, Kipf, Leskovec, Liao, Sun, Velickovic*

West Exhibition Hall C, **Bayesian Deep Learning** *Gal, Hernández-Lobato, Louizos, Nalisnick, Ghahramani, Murphy, Welling*

Dec. 14, 2019

East Ballroom A, **Real Neurons & Hidden Units: future directions at the intersection of neuroscience and AI** *Lajoie, Shlizerman, Puelma Touzel, Thompson, Kording*

East Ballroom B, **Fair ML in Healthcare** *Joshi, Chen, Obermeyer, Mullainathan*

East Ballroom C, **Tackling Climate Change with ML** *Rolnick, Donti, Kaack, Lacoste, Maharaj, Ng, Platt, Chayes, Bengio*

East Meeting Rooms 11 + 12, **Joint Workshop on AI for Social Good** *Fang, Bullock, Dilhac, Green, saltiel, Adjodah, Clark, McGregor, Luck, Penn, Sylvain, Boucher, Swaine-Simon, Tadesse, Côté, Bethke, Bengio*

East Meeting Rooms 1 - 3, **Machine Learning for Autonomous Driving** *McAllister, Rhinehart, Yu, Li, Dragan*

East Meeting Rooms 8 + 15, **Privacy in Machine Learning (PriML)** *Balle, Chaudhuri, Honkela, Koskela, Meehan, Park, Smart, Weller*

West 109 + 110, **Machine Learning and the Physical Sciences** *Baydin, Carrasquilla, Ho, Kashinath, Paganini, Thais, Anandkumar, Cranmer, Melko, Prabhat, Wood*

West 114 + 115, **Program Transformations for ML** *Lamblin, Baydin, Wiltchko, van Merriënboer, Fertig, Pearlmutter, Duvenaud, Hascoet*

West 116 + 117, **Competition Track Day 2** *Escalante*

West 118 - 120, **Emergent Communication: Towards Natural Language** *Gupta, Noukhovitch, Resnick, Jaques, Filos, Ossenkopf, Lazaridou, Foerster, Lowe, Kiela, Cho*

West 121 + 122, **Science meets Engineering of Deep Learning** *Sagun, Gulcehre, Romero, Rostamzadeh, de Freitas*

West 202 - 204, **ML For Systems** *Hashemi, Mirhoseini, Goldie, Swersky, XU, Raiman*

West 205 - 207, **The third Conversational AI workshop – today's practice and tomorrow's potential** *Geramifard, Williams, Byrne, Celikyilmaz, Gasic, Hakkani-Tur, Henderson, Lastras, Ostendorf*

West 208 + 209, **Document Intelligence** *Duffy, Akkiraju, Bedrax Weiss, Bennett, Motahari-Nezhad*

West 211 - 214, **Learning Transferable Skills** *Mattar, Juliani, Lange, Crosby, Beyret*

West 215 + 216, **Sets and Partitions** *Monath, Zaheer, McCallum, Kobren, Oliva, Poczos, Salakhutdinov*

West 217 - 219, **Context and Compositionality in Biological and Artificial Neural Systems** *Turek, Jain, Huth, Wehbe, Strubell, Yuille, Linzen, Honey, Cho*

West 220 - 222, **Robot Learning: Control and Interaction in the Real World** *Calandra, Rakelly, Kamthe, Kragic, Schaal, Wulfmeier*

West 223 + 224, **NeurIPS Workshop on Machine Learning for Creativity and Design 3.0** *Elliott, Dieleman, Roberts, Engel, White, Fiebrink, Mital, Payne, Tokui*

West 301 - 305, **Medical Imaging meets NeurIPS** *Lombaert, Glocker, Konukoglu, de Bruijne, Feragen, Oguz, Teuwen*

West 306, **Learning with Temporal Point Processes** *Rodriguez, Song, Valera, Liu, De, Zha*

West Ballroom A, **The Optimization Foundations of Reinforcement Learning** *Dai, He, Le Roux, Li, Schuurmans, White*

West Ballroom B, **Machine Learning with Guarantees** *London, Dziugaite, Roy, Joachims,*

Madry, Shawe-Taylor

West Ballroom C, **“Do the right thing”: machine learning and causal inference for improved decision making** *Santacatterina, Joachims, Kallus, Swaminathan, Sontag, Zhou*

West Exhibition Hall A, **Bridging Game Theory and Deep Learning** *Mitliagkas, Gidel, He, Askari Hemmat, Haghtalab, Lacoste-Julien*

West Exhibition Hall C, **Deep Reinforcement Learning** *Abbeel, Finn, Pineau, Silver, Singh, Achiam, Florensa, Grimm, Tang, Veeriah*

Safety and Robustness in Decision-making

Mohammad Ghavamzadeh, Shie Mannor, Yisong Yue, Marek Petrik, Yinlam Chow

East Ballroom A, Fri Dec 13, 08:00 AM

Interacting with increasingly sophisticated decision-making systems is becoming more and more a part of our daily life. This creates an immense responsibility for designers of these systems to build them in a way to guarantee safe interaction with their users and good performance, in the presence of noise and changes in the environment, and/or of model misspecification and uncertainty. Any progress in this area will be a huge step forward in using decision-making algorithms in emerging high stakes applications, such as autonomous driving, robotics, power systems, health care, recommendation systems, and finance.

This workshop aims to bring together researchers from academia and industry in order to discuss main challenges, describe recent advances, and highlight future research directions pertaining to develop safe and robust decision-making systems. We aim to highlight new and emerging theoretical and applied research opportunities for the community that arise from the evolving needs for decision-making systems and algorithms that guarantee safe interaction and good performance under a wide range of uncertainties in the environment.

Schedule

08:35 AM	Opening	<i>Ghosh, Shafiee, Boopathy, Tamkin</i>
09:00 AM	Aviv Tamar	
09:30 AM	Scott Niekum	
10:00 AM	Break and Posters	
11:00 AM	Marco Pavone	
11:30 AM	Dimitar Filev	
12:00 PM	Lunch	
02:00 PM	Nathan Kallus	
02:30 PM	Finale Doshi-Velez	
03:00 PM	Thorsten Joachim	
03:30 PM	Coffee Break	
03:45 PM	Daniel Kuhn	<i>Kuhn</i>
04:15 PM	Andy Sun	
04:45 PM	Panel Discussion	
05:30 PM	Concluding Remarks	

05:40
PM**Poster Session**

*Ghosh, Shafiee,
Boopathy, Tamkin,
Vasiloudis, Nanda,
Baheri, Fieguth,
Bennett, Shi, Liu,
Jain, Tyo, Wang,
Chen, Wainwright,
Shama Sastry, Tang,
Brown, Inouye,
Venuto, Ramani,
Diochnos, Madaan,
Krashenikov, Oren,
Lee, Quint, amirloo,
Pirota, Hartnett,
Dubourg-Felonneau,
Swamy, Chen,
Bogunovic, Carter,
Garcia-Barcos,
Mohapatra, Zhang,
Qian, Martin, Richter,
Zaiter, Weng, ,
Polymenakos,
Hoang, abbasi,
Gallieri, Seurin,
Papini, Turchetta,
Sotoudeh,
Hosseinzadeh,
Fulton, Uehara,
Prasad, Camburu,
Kolaric, Renz,
Jaiswal, Russel,
Islam, Agarwal,
Aldrick, Vernekar,
Lale,
Narayanaswami,
Daulton, Garg, East,
Zhang, Dsidbari,
Goodwin, Krakovna,
Luo, Chung, Shi,
Wang, Jin, Xu*

Learning Meaningful Representations of Life

Liz Wood, Yakir Reshef, Jon Bloom, Jasper Snoek, Barbara Engelhardt, Scott Linderman, Suchi Saria, Alexander Wiltschko, Casey Greene, Chang Liu, Kresten Lindorff-Larsen, Debora Marks

East Ballroom B, Fri Dec 13, 08:00 AM

The last decade has seen both machine learning and biology transformed: the former by the ability to train complex predictors on massive labelled data sets; the latter by the ability to perturb and measure biological systems with staggering throughput, breadth, and resolution. However, fundamentally new ideas in machine learning are needed to translate biomedical data at scale into a mechanistic understanding of biology and disease at a level of abstraction beyond single genes. This challenge has the potential to drive the next decade of creativity in machine learning as the field grapples with how to move beyond prediction to a regime that broadly catalyzes and accelerates scientific discovery.

To seize this opportunity, we will bring together current and future leaders within each field to introduce the next generation of machine learning specialists to the next generation of biological problems. Our full-day workshop will start a deeper dialogue with the goal of Learning Meaningful Representations of Life (LMRL), emphasizing interpretable representation learning of structure and principles. The workshop will address this challenge at five layers of biological abstraction (genome, molecule, cell, system, phenome) through interactive breakout sessions led by a diverse team of experimentalists and computational scientists to facilitate substantive discussion.

We are calling for short abstracts from computer scientists and biological scientists. Submission deadline is Friday, September 20. Significant travel

support is also available. Details here:

<https://lmrl-bio.github.io/call>

<https://lmrl-bio.github.io/travel>

Schedule

08:45 AM	Opening Remarks	<i>Yeshwant</i>
09:00 AM	Keynote - Bio	<i>Regev</i>
09:30 AM	Keynote - ML	<i>Welling</i>
10:00 AM	Keynote - ML/Bio	<i>Koller</i>
10:30 AM	Coffee Break	
10:45 AM	Molecules and Genomes	<i>Morris, Haussler, Noe, Clevert, Keiser, Aspuru-Guzik, Duvenaud, Huang, Jones</i>
12:00 PM	Synthetic Systems	<i>Silver, Marks, Liu</i>
12:30 PM	GWAS Discussion	<i>Wang, D'Amour</i>
01:30 PM	Phenotype	<i>HaCohen, Reshef, Johnson, Morris, Nagy, Eraslan, Singer, Van Allen, Krishnaswamy, Greene, Linderman, Bloemendal, Wiltschko, Kotliar, Zou, Bulik-Sullivan</i>
03:15 PM	Coffee Break	

03:30 PM	Cell	<i>Carpenter, Zhou, Chikina, Tong, Lengerich, Abdelkareem, Eraslan, Blumberg, Ra, Burkhardt, Matsen IV, Moses, Chen, Haghighi, Lu, Schau, Nivala, Shiffman, Harbrecht, Masengo Wa Umba</i>
05:00 PM	Closing Remarks	<i>Sander, Fiete, Peer</i>
06:00 PM	Posters and Social Hour	

Abstracts (10):

Abstract 1: **Opening Remarks in Learning Meaningful Representations of Life**, *Yeshwant*
08:45 AM

Opening remarks by Francis Collins (Director, National Institutes of Health) via video and Krishna Yeshwant, General Partner at Google Ventures.

Abstract 2: **Keynote - Bio in Learning Meaningful Representations of Life**, *Regev* 09:00 AM

Aviv Regev. Professor of Biology; Core Member, Broad Institute; Investigator, Howard Hughes Medical Institute. Aviv Regev pioneers the use of single-cell genomics and other techniques to dissect the molecular networks that regulate genes, define cells and tissues, and influence health and disease.

Abstract 3: **Keynote - ML in Learning Meaningful Representations of Life**, *Welling* 09:30 AM

Max Welling is a research chair in Machine Learning at the University of Amsterdam and a VP Technologies at Qualcomm.

Abstract 4: **Keynote - ML/Bio in Learning Meaningful Representations of Life**, *Koller* 10:00 AM

Daphne Koller is the Rajeev Motwani Professor in the Computer Science Department at Stanford University and founder of insitro.

Abstract 6: **Molecules and Genomes in Learning Meaningful Representations of Life**, *Morris, Haussler, Noe, Clevert, Keiser, Aspuru-Guzik, Duvenaud, Huang, Jones* 10:45 AM

Quaid Morris, Anna Goldenberg, David Haussler, Frank Noe, Djork-Arne Clevert, Michael Keiser, Alan Asparu-Guzik, David Duvenaud, Possu Huang and David Jones present.

Abstract 7: **Synthetic Systems in Learning Meaningful Representations of Life**, *Silver, Marks, Liu* 12:00 PM

Pamela Silver, Debora Marks, and Chang Liu in conversation.

Abstract 8: **GWAS Discussion in Learning Meaningful Representations of Life**, *Wang, D'Amour* 12:30 PM

Yixin Wang and Alex D'Amour in conversation.

Abstract 9: **Phenotype in Learning Meaningful Representations of Life**, *HaCohen, Reshef, Johnson, Morris, Nagy, Eraslan, Singer, Van Allen, Krishnaswamy, Greene, Linderman, Bloemendal, Wiltschko, Kotliar, Zou, Bulik-Sullivan* 01:30 PM

Nir Hacohen, David Reshef, Matt Johnson, Samantha Morris, Aurel Nagy, Gokcen Eraslan, Meromit Singer, Eli van Allen, Smita Krishnaswamy, Casey Greene, Scott Linderman, Alex Bloemendal, Alex Wiltschko, Dylan Kotliar, James Zou, and Brendan Bulik-Sullivan participate.

Abstract 11: **Cell in Learning Meaningful Representations of Life**, *Carpenter, Zhou, Chikina, Tong, Lengerich, Abdelkareem, Eraslan, Blumberg, Ra, Burkhardt, Matsen IV, Moses, Chen, Haghighi, Lu, Schau, Nivala, Shiffman, Harbrecht, Masengo Wa Umba* 03:30 PM

Anne Carpenter, Hui Ting Grace Yeo, Jian Zhou, Maria Chikina, Alexander Tong, Benjamin Lengerich, Aly O. Abdelkareem, Gokcen Eraslan, Andrew Blumberg, Stephen Ra, Daniel Burkhardt, Emanuel Flores Bautista, Frederick Matsen, Alan Moses, Zhenghao Chen, Marzieh Haghighi, Alex Lu, Geoffrey Schau, Jeff Nivala, Luke O'Connor, Miriam Shiffman, Hannes Harbrecht and Shimbi Masengo Wa Umba Papa Levi present in a lightning round.

Abstract 12: **Closing Remarks in Learning Meaningful Representations of Life**, *Sander, Fiete, Peer* 05:00 PM

Chris Sander, Ila Fiete, and Dana Pe'er present.

Optimal Transport for Machine Learning

Marco Cuturi, Gabriel Peyré, Rémi Flamary, Alexandra Suvorikova

East Ballroom C, Fri Dec 13, 08:00 AM

Optimal transport(OT) provides a powerful and flexible way to compare, interpolate and morph probability measures. Originally proposed in the eighteenth century, this theory later led to Nobel Prizes for Koopmans and Kantorovich as well as C. Villani and A. Figalli Fields' Medals in 2010 and 2018. OT is now used in challenging learning problems that involve high-dimensional data such as the inference of individual trajectories by looking at population snapshots in biology, the estimation of generative models for images, or more generally transport maps to transform samples in one space

into another as in domain adaptation. With more than a hundred papers mentioning Wasserstein or transport in their title submitted at NeurIPS this year, and several dozens appearing every month across ML/stats/imaging and data sciences, this workshop's aim will be to federate and advance current knowledge in this rapidly growing field.

Schedule

08:00 AM	Facundo Memoli	<i>Mémoli</i>
08:40 AM	Karren Dai	
09:00 AM	Jon Weed	<i>Niles-Weed</i>
10:30 AM	Stefanie Jegelka	<i>Jegelka</i>
11:10 AM	SPOTLIGHTS 5 x 10	
12:00 PM	Poster Session	
02:00 PM	Geoffrey Schiebinger	<i>Schiebinger</i>
02:40 PM	Charlie Frogner	<i>Frogner</i>
03:00 PM	Aude Genevay	<i>Genevay</i>
04:20 PM	Daniel Kuhn	<i>Kuhn</i>
05:00 PM	Alexei Kroshnin	<i>Kroshnin</i>

05:20 PM

Poster Session

, *Kroshnin, Delalande, Carr, Tompkins, Pooladian, Robert, Makkuva, Genevay, Liu, Zeng, Frogner, Cazelles, Tabak, Ramos, PATY, Balikas, Trigila, Wang, Mahler, Nielsen, Lounici, Swanson, Bhutani, Bréchet, Indyk, cohen, Jegelka, Wu, Sejourne, Manole, zhao, Wang, Wang, Dukler, Wang, Dong*

Information Theory and Machine Learning

Shengjia Zhao, Jiaming Song, Yanjun Han, Kristy Choi, Pratyusha Kalluri, Ben Poole, Alex Dimakis, Jiantao Jiao, Tsachy Weissman, Stefano Ermon

East Exhibition Hall A, Fri Dec 13, 08:00 AM

Information theory is deeply connected to two key tasks in machine learning: prediction and representation learning. Because of these connections, information theory has found wide applications in machine learning tasks, such as proving generalization bounds, certifying fairness and privacy, optimizing information content of unsupervised/supervised representations, and proving limitations to prediction performance. Conversely, progress in machine learning have been successfully applied to classical information theory tasks such as compression and transmission.

These recent progress have lead to new open questions and opportunities: to marry the simplicity and elegance of information theoretic analysis with the complexity of modern high dimensional machine learning setups. However, because of the diversity of information theoretic research, different communities often progress independently despite shared questions and tools. For example, variational bounds to mutual information are concurrently developed in information theory, generative model, and learning theory communities.

This workshop hopes to bring together researchers from different disciplines, identify common grounds, and spur discussion on how information theory can apply to and benefit from modern machine learning setups.

Schedule

09:00 AM	Invited Talk: Aaron van den Oord	<i>van den Oord</i>
09:30 AM	Invited Talk: Po-Ling Loh	
10:00 AM	Invited Talk: Alexander A Alemi	<i>Alemi</i>
11:00 AM	Invited Talk: Stefano Soatto and Alessandro Achille	<i>Soatto, Achille</i>
11:30 AM	Invited Talk: Maxim Raginsky	<i>Raginsky</i>
02:00 PM	Invited Talk: Varun Jog	<i>Jog</i>
02:30 PM	Invited Talk: Jelani Nelson	
03:30 PM	Invited Talk: Irena Fischer-Hwang	

04:10 PM

Poster Session

Flamich, Ubaru, Zheng, Djolonga, Wickstrøm, Granziol, Pitas, Li, Williamson, Yoon, Lee, Zilly, Petrini, Fischer, Dong, Alemi, Nguyen, Brekelmans, Wu, Mahajan, Li, Shiragur, Carmon, Adilova, LIU, An, Dash, Gunluk, Mazumdar, Motani, Rosenzweig, Kamp, Havasi, Barnes, Zhou, Hao, Foster, Benjamini, Srebro, Tschannen, Rubenstein, Gelly, Duchi, Sidford, Ru, Zohren, Dalal, Osborne, Roberts, Charikar, Subramanian, Fan, Schwarzer, Roberts, Lacoste-Julien, Prabhu, Galstyan, Ver Steeg, Sankar, Noh, Dasarathy, Park, Cheung, Tran, Yang, Poole, Censi, Sylvain, Hjelm, Liu, Gallego, Sypherd, Yang, Morshuis

MLSys: Workshop on Systems for ML

Aparna Lakshmiratan, Siddhartha Sen, Joseph Gonzalez, Dan Crankshaw, Sarah Bird

East Meeting Rooms 11 + 12, Fri Dec 13, 08:00 AM

A new area is emerging at the intersection of artificial intelligence, machine learning, and systems design. This has been accelerated by the explosive growth of diverse applications of ML in production, the continued growth in data volume, and the complexity of large-scale learning systems. The goal of this workshop is to bring together experts working at the crossroads of machine learning, system design and software engineering to explore the challenges faced when building large-scale ML systems. In particular, we aim to elicit new connections among these diverse fields, identifying theory, tools and design principles tailored to practical machine learning workflows. We also want to think about best practices for research in this area and how to evaluate it. The workshop will cover state of the art ML and AI platforms and algorithm toolkits (e.g. TensorFlow, PyTorch1.0, MXNet etc.), as well as dive into machine learning-focused developments in distributed learning platforms, programming languages, data structures, hardware accelerators, benchmarking systems and other topics.

This workshop will follow the successful model we have previously run at ICML, NeurIPS and SOSP.

Our plan is to run this workshop annually co-located with one ML venue and one Systems venue, to help build a strong community which we think will complement newer conferences like SysML targeting research at the intersection of systems and machine learning. We believe this dual approach will help to create a low barrier to participation for both communities.

This workshop is part two of a two-part series with one day focusing on ML for Systems and the other on Systems for ML. Although the two workshops are being led by different organizers, we are

coordinating our call for papers to ensure that the workshops complement each other and that submitted papers are routed to the appropriate venue.

Schedule

09:00 AM	Welcome	
09:10 AM	Keynote	
09:30 AM	Contributed Talk	
10:00 AM	Contributed Talk	
10:30 AM	Coffee Break	
		<i>Kumar, Kornuta, Bakhteev, Guan, Dong, Cho, Laue, Vasiloudis, Anghel, Wijmans, Shang, Kuchaiev, Lin, Zhang, Zhu, Chen, Joseph, Ding, Raiman, Shin, Thangarasa, Sankaran, Mathur, Dazzi, Löning, Ho, Zraggen, Nakandala, Kornuta, Kuznetsova</i>
10:30 AM	Posters and Coffee	
	Open source software showcase (TF, PyTorch, TVM...)	
12:30 PM		
02:00 PM	Coffee Break	

02:10 PM	Keynote
02:30 PM	Contributed Talk
02:50 PM	Contributed Talk
03:10 PM	Contributed Talk
03:30 PM	Posters and Coffee
04:30 PM	Panel

Perception as generative reasoning: structure, causality, probability

*Dan Rosenbaum, Marta Garnelo, Peter Battaglia,
Kelsey Allen, Ilker Yildirim*

East Meeting Rooms 1 - 3, Fri Dec 13, 08:00 AM

Many perception tasks can be cast as ‘inverse problems’ where the input signal is the outcome of a causal process and perception is to invert that process. For example in visual object perception, the image is caused by an object and perception is to infer which object gave rise to that image. Following an analysis-by-synthesis approach, modelling the forward and causal direction of the data generation process is a natural way to capture the underlying scene structure, which typically leads to broader generalisation and better sample efficiency. Such a forward model can be applied to solve the inverse problem (inferring the scene structure from an input image) using Bayes rule, for example. This workflow stands in contrast to common approaches in deep learning, where typically one first defines a task, and then optimises a deep model end-to-end to solve it. In this

workshop we propose to revisit ideas from the generative approach and advocate for learning-based analysis-by-synthesis methods for perception and inference. In addition, we pose the question of how ideas from these research areas can be combined with and complement modern deep learning practices.

Schedule

08:50 AM	Opening Remarks	<i>Rosenbaum, Garnelo, Battaglia, Allen, Yildirim</i>
09:00 AM	Sanja Fidler	<i>Fidler</i>
09:35 AM	Spotlights 1	<i>Chorowski, Deng, Chang</i>
10:30 AM	Jiajun Wu	<i>Wu</i>
11:05 AM	Tatiana Lopez-Guevara	<i>López-Guevara</i>
11:40 AM	Spotlights 2	
01:30 PM	Niloy Mitra	<i>Mitra</i>
02:05 PM	Danilo Rezende	<i>Jimenez Rezende</i>
02:40 PM	Posters	<i>Graber, Hu, Fang, Hamrick, Giannone, Co-Reyes, Deng, Crawford, Dittadi, Karkus, Dirks, TRIVEDI, Raj, Felip Leon, Chan, Chorowski, Orchard, Stanislaw, Kortylewski, Zinberg, Zhou, Sun, Mansinghka, Li, Cusumano-Towner</i>

04:15 PM	Invited talk	
04:50 PM	Panel	<i>Fidler, Wu, Tenenbaum, López-Guevara, Jimenez Rezende</i>



Minding the Gap: Between Fairness and Ethics

Igor Rubinov, Risi Kondor, Jack Poulson, Manfred K. Warmuth, Emanuel Moss, Alexa Hagerty

East Meeting Rooms 8 + 15, Fri Dec 13, 08:00 AM

When researchers and practitioners, as well as policy makers and the public, discuss the impacts of deep learning systems, they draw upon multiple conceptual frames that do not sit easily beside each other. Questions of algorithmic fairness arise from a set of concerns that are similar, but not identical, to those that circulate around AI safety, which in turn overlap with, but are distinct from, the questions that motivate work on AI ethics, and so on. Robust bodies of research on privacy, security, transparency, accountability, interpretability, explainability, and opacity are also incorporated into each of these frames and conversations in variable ways. These frames reveal gaps that persist across both highly technical and socially embedded approaches, and yet collaboration across these gaps has proven challenging.

Fairness, Ethics, and Safety in AI each draw upon different disciplinary prerogatives, variously centering applied mathematics, analytic philosophy, behavioral sciences, legal studies, and the social sciences in ways that make conversation between these frames fraught with misunderstandings.

These misunderstandings arise from a high degree of linguistic slippage between different frames, and reveal the epistemic fractures that undermine valuable synergy and productive collaboration. This workshop focuses on ways to translate between these ongoing efforts and bring them into necessary conversation in order to understand the profound impacts of algorithmic systems in society.

Schedule

08:00 AM	Opening Remarks	<i>Warmuth</i>
08:15 AM	Invited Talk	<i>Bengio</i>
08:45 AM	Approaches to Understanding AI	<i>Bengio, Dobbe, Elish, Kroll, Metcalf</i>
09:45 AM	Spectrogram	
10:00 AM	Coffee Break	
10:30 AM	Detecting and Documenting AI Impacts	<i>Christian, Hagerty, Rogers, Schuur, Snow</i>
11:30 AM	Responsibilities	<i>Chowdhury, Kim, Schuur, Siegel, Smart</i>
12:30 PM	Lunch	
02:00 PM	A Conversation with Meredith Whittaker	<i>Sloane, Whittaker</i>
02:45 PM	Global implications	<i>Chowdhury, Malliaraki, Poulson, Sloane</i>
03:45 PM	Coffee Break	
04:30 PM	Solutions	<i>Christian, Hu, Kondor, Marshall, Rogers, Schuur, Warmuth</i>

Abstracts (5):

Abstract 3: Approaches to Understanding AI in Minding the Gap: Between Fairness and Ethics, *Bengio, Dobbe, Elish, Kroll, Metcalf* 08:45 AM

The stakes of AI certainly alter how we relate to each other as humans - how we know what we know about reality, how we communicate, how we work and earn money, and about how we think of ourselves as human. But in grappling with these changing relations, three fairly concrete approaches have dominated the conversation: ethics, fairness, and safety. These approaches come from very different academic backgrounds, draw attention to very different aspects of AI, and imagine very different problems and solutions as relevant, leading us to ask:

- What are the commonalities and differences between ethics, fairness, and safety as approaches to addressing the challenges of AI?
- How do these approaches imagine different problems and solutions for the challenges posed by AI?
- How can these approaches work together, or are there some areas where they are mutually incompatible?

Abstract 6: Detecting and Documenting AI Impacts in Minding the Gap: Between Fairness and Ethics, *Christian, Hagerty, Rogers, Schuur, Snow* 10:30 AM

Algorithmic systems are being widely used in key social institutions and while they promise radical improvements in fields from public health to energy allocation, they also raises troubling issues of bias, discrimination, and “automated inequality.” They also present irresolvable challenges related to the dual-use nature of these technologies, secondary effects that are difficult to anticipate, and alter power relations between individuals, companies, and

governments.

- How should we delimit the scope of AI impacts? What can properly be considered an AI impact, as opposed to an impact arising from some other cause?
- How do we detect and document the social impacts of AI?
- What tools, processes, and institutions ought to be involved in addressing these questions?

Abstract 7: Responsibilities in Minding the Gap: Between Fairness and Ethics, *Chowdhury, Kim, Schuur, Siegel, Smart* 11:30 AM

While there is a great deal of AI research happening in academic settings, much of that work is operationalized within corporate contexts. Some companies serve as vendors, selling AI systems to government entities, some sell to other companies, some sell directly to end-users, and yet others sell to any combination of the above.

- What set of responsibilities does the AI industry have w.r.t. AI impacts?
- How do those responsibilities shift depending on a B2B, B2G, B2C business model?
- What responsibilities does government have to society, with respect to AI impacts arising from industry?
- What role does civil society organizations have to play in this conversation?

Abstract 10: Global implications in Minding the Gap: Between Fairness and Ethics, *Chowdhury, Malliaraki, Poulson, Sloane* 02:45 PM

The risks and benefits of AI are unevenly distributed within societies and across the globe. Governance regimes are drastically different in various regions of the world, as are the political and ethical implications of AI technologies.

- How do we better understand how AI technologies operate around the world and the range of risks they carry for different societies?
- Are there global claims about the implications of AI

that can apply everywhere around the globe? If so, what are they?

- What can we learn from AI's impacts on labor, environment, public health and agriculture in diverse settings?

Abstract 12: **Solutions in Minding the Gap: Between Fairness and Ethics**, *Christian, Hu, Kondor, Marshall, Rogers, Schuur, Warmuth* 04:30 PM

With the recognition that there are no fully sufficient steps that can be taken to addressing all AI impacts, there are concrete things that ought to be done, ranging across technical, socio-technical, and legal or regulatory possibilities.

- What are the technical, social, and/or regulatory solutions that are necessary to address the riskiest aspects of AI?
- What are key approaches to minimize the risks of AI technologies?

KR2ML - Knowledge Representation and Reasoning Meets Machine Learning

Veronika Thost, Christian Muise, Kartik Talamadupula, Sameer Singh, Chris Ré

West 109 + 110, Fri Dec 13, 08:00 AM

Machine learning (ML) has seen a tremendous amount of recent success and has been applied in a variety of applications. However, it comes with several drawbacks, such as the need for large amounts of training data and the lack of explainability and verifiability of the results. In many domains, there is structured knowledge (e.g., from electronic health records, laws, clinical guidelines, or common sense knowledge) which can be leveraged for reasoning in an informed way (i.e., including the information encoded in the knowledge representation itself) in order to obtain high quality answers. Symbolic approaches for knowledge

representation and reasoning (KRR) are less prominent today - mainly due to their lack of scalability - but their strength lies in the verifiable and interpretable reasoning that can be accomplished. The KR2ML workshop aims at the intersection of these two subfields of AI. It will shine a light on the synergies that (could/should) exist between KRR and ML, and will initiate a discussion about the key challenges in the field.

Schedule

08:00 AM	Opening Remarks	
08:05 AM	Invited Talk (William W. Cohen)	<i>Cohen</i>
	Contributed Talk: Neural-Guided Symbolic Regression with Asymptotic Constraints	
08:35 AM		<i>Singh</i>
08:50 AM	Contributed Talk: Towards Finding Longer Proofs	<i>Zombori</i>
09:05 AM	Contributed Talk: Neural Markov Logic Networks	<i>Kuzelka</i>
		<i>Bahn, Xu, Su, Cunnington, Hwang, Dash, Camacho, Khurana, Li, Zhang,</i>
09:20 AM	Poster Spotlights A (23 posters)	<i>Naderi, Zeng, Khosravi, Colon-Hernandez, Diochnos, Windridge, Percy, Manhaeve, Belle, Juba</i>
09:45 AM	Coffee Break + Poster Session	

10:30 AM	Invited Talk (Xin Luna Dong)	<i>Dong</i>
	Contributed Talk: Layerwise Knowledge Extraction from Deep Convolutional Networks	
11:00 AM		<i>Odense</i>
	Contributed Talk: Ontology-based Interpretable Machine Learning with Learnable Anchors	
11:15 AM		<i>Lai</i>
	Contributed Talk: Learning multi-step spatio-temporal reasoning with Selective Attention Memory Network	
11:30 AM		<i>Jayram</i>
	Contributed Talk: MARLeME: A Multi-Agent Reinforcement Learning Model Extraction Library	
11:45 AM		<i>Kazhdan</i>
12:00 PM	Invited Talk (Vivek Srikumar)	<i>Srikumar</i>
12:30 PM	Lunch Break	
02:00 PM	Invited Talk (Francesca Rossi)	<i>Rossi</i>
	Contributed Talk: TP-N2F: Tensor Product Representation for Natural To Formal Language Generation	
02:30 PM		<i>Chen</i>

	Contributed Talk: TabFact: A Large-scale Dataset for Table-based Fact Verification	
02:45 PM		<i>Chen</i>
	Contributed Talk: LeDeepChef: Deep Reinforcement Learning Agent for Families of Text-Based Games	
03:00 PM		<i>Adolphs</i>
	Poster Spotlights B (14 posters)	<i>Belle, Govindarajulu, Bockhorst, Gunel, UCEDA-SOSA, Klassen, Weyde, Ghalwash, Arora, Illanes, Raiman, Wang, Lew, Min</i>
03:15 PM		
03:30 PM	Coffee Break + Poster Session	
04:15 PM	Invited Talk (Guy Van den Broeck)	<i>Van den Broeck</i>
04:45 PM	Invited Talk (Yejin Choi)	<i>Choi</i>
05:15 PM	Discussion Panel	
05:55 PM	Closing Remarks	

Retrospectives: A Venue for Self-Reflection in ML Research

Ryan Lowe, Yoshua Bengio, Joelle Pineau, Michela Paganini, Jessica Forde, Shagun Sodhani, Abhishek Gupta, Joel Lehman, Peter Henderson, Kanika Madan

West 114 + 115, Fri Dec 13, 08:00 AM

The NeurIPS Workshop on Retrospectives in Machine Learning will kick-start the exploration of a new kind of scientific publication, called retrospectives. The purpose of a retrospective is to answer the question:

“What should readers of this paper know now, that is not in the original publication?”

Retrospectives provide a venue for authors to reflect on their previous publications, to talk about how their intuitions have changed, to identify shortcomings in their analysis or results, and to discuss resulting extensions that may not be sufficient for a full follow-up paper. A retrospective is written about a single paper, by that paper's author, and takes the form of an informal paper. The overarching goal of retrospectives is to improve the science, openness, and accessibility of the machine learning field, by widening what is publishable and helping to identifying opportunities for improvement. Retrospectives will also give researchers and practitioners who are unable to attend top conferences access to the author’s updated understanding of their work, which would otherwise only be accessible to their immediate circle.

Schedule

09:00 AM	Opening Remarks
09:15 AM	Invited talk #1
09:35 AM	Invited talk #2
09:55 AM	Invited talk #3
10:15 AM	Coffee break + poster set-up

10:30 AM	Invited talk #4
10:50 AM	Panel discussing how to increase transparency and dissemination of ‘soft knowledge’ in ML
12:00 PM	Lunch break
01:50 PM	Retrospective: An Intriguing Failing of Convolutional Neural Networks and the CoordConv Solution <i>Liu</i>
01:55 PM	Retrospective: Learning the structure of deep sparse graphical models <i>Ghahramani</i>
02:00 PM	Retrospective: Lessons Learned from The Lottery Ticket Hypothesis <i>Frankle</i>
02:05 PM	Retrospective: FiLM: Visual Reasoning with a General Conditioning Layer
02:10 PM	Retrospective: Deep Ptych: Subsampled Fourier Ptychography via Generative Priors
02:15 PM	Retrospective: Markov games that people play <i>Littman</i>

02:20 PM	Retrospective: DLPaper2Code: Auto-Generation of Code from Deep Learning Research Papers	<i>Sankaran</i>
02:25 PM	Retrospective; Deep Reinforcement Learning That Matters	<i>Islam</i>
02:30 PM	Smarter prototyping for neural learning	<i>Pradhan</i>
02:35 PM	Advances in deep learning for skin cancer detection	
02:40 PM	Unsupervised Minimax: Adversarial Curiosity, Generative Adversarial Networks, and Predictability Minimization	<i>Schmidhuber</i>
02:45 PM	Posters + Coffee Break	
04:10 PM	Invited talk #5	
04:30 PM	Retrospectives brainstorming session: how do we produce impact?	



Competition Track Day 1

Hugo Jair Escalante

West 116 + 117, Fri Dec 13, 08:00 AM

<https://nips.cc/Conferences/2019/CallForCompetitions>

Schedule

08:10 AM	Disentanglement Challenge - Disentanglement and Results of the Challenge Stages 1 & 2	<i>Miladinovic, Bauer, Keysers</i>
08:40 AM	Robot open-Ended Autonomous Learning: a challenge	<i>Cartoni, Baldassarre</i>
09:00 AM	The Pommerman competition	<i>Resnick, Bouhenguel</i>
09:30 AM	The CellSignal challenge	<i>Mabey, Sypetkowski, Haque</i>
10:00 AM	Learn to Move: Walk Around	<i>Zhou, Zeng, Wang</i>
10:30 AM	Coffee break	
02:15 PM	MicroNet Challenge	<i>Gale, Wang, Leng, Cheng, Wang, Elsen, Yan</i>
03:15 PM	Reconnaissance Blind Chess competition	<i>Llorens, Gardner, Perrotta, Highley, Shen, Goldbloom, Earnshaw, Clark, Perrotta, Bernardoni, Jordan, Wang</i>
04:15 PM	The 3D Object Detection over HD Maps for Autonomous Cars Challenge	<i>Vincent, Jain</i>

Abstracts (4):

Abstract 1: Disentanglement Challenge - Disentanglement and Results of the Challenge

Stages 1 & 2 in Competition Track Day 1,*Miladinovic, Bauer, Keysers* 08:10 AM

Stefan Bauer: Learning Disentangled Representations

Djordje Miladinovic: Disentanglement in the Real-World

Daniel Keysers: Disentanglement_lib

Bernhard Schölkopf: Hand-out-of-certificates

Abstract 4: **The CellSignal challenge in Competition Track Day 1**, *Mabey, Sypetkowski, Haque* 09:30 AM

- * Opening remarks, description of competition, summary of results.
- * Description of first prize solution.
- * Description of second prize solution.
- * Mention of third prize solution
- * congratulations to winners and description of AutoML solution.
- * Prize ceremony.

Abstract 7: **MicroNet Challenge in Competition Track Day 1**, *Gale, Wang, Leng, Cheng, Wang, Elsen, Yan* 02:15 PM

Trevor Gale and Erich Elsen. Introduction to the competition and overview of results.

Peisong Wang, Cong Leng, and Jian Cheng. An Empirical Study of Network Compression for Image Classification.

Trevor Gale and Erich Elsen. Highlights of other notable entries.

Zhongxia Yan and Hanrui Wang. Efficient Memory-Augmented Language Models with Network Compression

Trevor Gale and Erich Elsen. Updates and improvements for the 2020 MicroNet Challenge.

Abstract 8: **Reconnaissance Blind Chess competition in Competition Track Day 1**, *Llorens, Gardner, Perrotta, Highley, Shen, Goldbloom, Earnshaw, Clark, Perrotta, Bernardoni, Jordan, Wang* 03:15 PM

- * Chair: I-Jeng Wang
- * Competition and Game Overview (Ashley Llorens)
- * Challenges of the Game (Ryan Gardner)
- * Competition Results (Casey Richardson)
- * Overview of the StrangeFish Bot (Gino Perrotta and Robert Perrotta)
- * Overview of the LaSalle Bot (T.J. Highley)
- * Overview of the penumbra Bot (Gregory Clark)
- * Overview of the wbernar5 Bot (William Bernardoni)
- * Overview of the MBot Bot (Mark Jordan)

Workshop on Federated Learning for Data Privacy and Confidentiality

Lixin Fan, Jakub Konečný, Yang Liu, Brendan McMahan, Virginia Smith, Han Yu

West 118 - 120, Fri Dec 13, 08:00 AM

Overview

Privacy and security have become critical concerns in recent years, particularly as companies and organizations increasingly collect detailed information about their products and users. This information can enable machine learning methods that produce better products. However, it also has the potential to allow for misuse, especially when private data about individuals is involved. Recent research shows that privacy and utility do not necessarily need to be at odds, but can be addressed by careful design and analysis. The need

for such research is reinforced by the recent introduction of new legal constraints, led by the European Union's General Data Protection Regulation (GDPR), which is already inspiring novel legislative approaches around the world such as Cyber-security Law of the People's Republic of China and The California Consumer Privacy Act of 2018.

An approach that has the potential to address a number of problems in this space is federated learning (FL). FL is an ML setting where many clients (e.g., mobile devices or whole organizations) collaboratively train a model under the orchestration of a central server (e.g., service provider), while keeping the training data decentralized.

Organizations and mobile devices have access to increasing amounts of sensitive data, with scrutiny of ML privacy and data handling practices increasing correspondingly. These trends have produced significant interest in FL, since it provides a viable path to state-of-the-art ML without the need for the centralized collection of training data – and the risks and responsibilities that come with such centralization. Nevertheless, significant challenges remain open in the FL setting, the solution of which will require novel techniques from multiple fields, as well as improved open-source tooling for both FL research and real-world deployment

This workshop aims to bring together academic researchers and industry practitioners with common interests in this domain. For industry participants, we intend to create a forum to communicate what kind of problems are practically relevant. For academic participants, we hope to make it easier to become productive in this area. Overall, the workshop will provide an opportunity to share the most recent and innovative work in FL, and discuss open problems and relevant approaches. The technical issues encouraged to be submitted include general computation based on decentralized data (i.e., not only machine learning), and how such

computations can be combined with other research areas, such as differential privacy, secure multi-party computation, computational efficiency, coding theory, etc. Contributions in theory as well as applications are welcome, including proposals for novel system design. Work on fully-decentralized (peer-to-peer) learning will also be considered, as there is significant overlap in both interest and techniques with federated learning.

Call for Contributions

We welcome high quality submissions in the broad area of federated learning (FL). A few (non-exhaustive) topics of interest include:

- . Optimization algorithms for FL, particularly communication-efficient algorithms tolerant of non-IID data
- . Approaches that scale FL to larger models, including model and gradient compression techniques
- . Novel applications of FL
- . Theory for FL
- . Approaches to enhancing the security and privacy of FL, including cryptographic techniques and differential privacy
- . Bias and fairness in the FL setting
- . Attacks on FL including model poisoning, and corresponding defenses
- . Incentive mechanisms for FL
- . Software and systems for FL
- . Novel applications of techniques from other fields to the FL setting: information theory, multi-task learning, model-agnostic meta-learning, and etc.
- . Work on fully-decentralized (peer-to-peer) learning will also be considered, as there is significant overlap in both interest and techniques with FL.

Submissions in the form of extended abstracts must be at most 4 pages long (not including references), be anonymized, and adhere to the NeurIPS 2019 format. Submissions will be accepted as contributed talks or poster presentations. The workshop will not have formal proceedings, but accepted papers will

be posted on the workshop website.

We support reproducible research and will sponsor a prize to be given to the best contribution that provides code to reproduce their results.

Submission link:
<https://easychair.org/conferences/?conf=fneurips2019>

Important Dates (2019)
Submission deadline: Sep 9
Author notification: Sep 30
Camera-Ready Papers Due: TBD
Workshop: Dec 13

Organizers:
Lixin Fan, WeBank
Jakub Konečný, Google
Yang Liu, WeBank
Brendan McMahan, Google
Virginia Smith, CMU
Han Yu, NTU

Invited Speakers:
Francoise Beaufays, Principal Researcher, Google
Shahrokh Daijavad, Distinguished Research, IBM
Dawn Song, Professor, University of California, Berkeley
Ameet Talwalkar, Assistant Professor, CMU; Chief Scientist, Determined AI
Max Welling, Professor, University of Amsterdam; VP Technologies, Qualcomm
Qiang Yang, Hong Kong University of Science and Technology, Hong Kong; Chief AI Officer, WeBank

FAQ
Can supplementary material be added beyond the 4-page limit and are there any restrictions on it?
Yes, you may include additional supplementary material, but you should ensure that the main paper is self-contained, since looking at supplementary material is at the discretion of the reviewers. The supplementary material should also follow the same

NeurIPS format as the paper and be limited to a reasonable amount (max 10 pages in addition to the main submission).

Can a submission to this workshop be submitted to another NeurIPS workshop in parallel?

We discourage this, as it leads to more work for reviewers across multiple workshops. Our suggestion is to pick one workshop to submit to.

Can a paper be submitted to the workshop that has already appeared at a previous conference with published proceedings?

We won't be accepting such submissions unless they have been adapted to contain significantly new results (where novelty is one of the qualities reviewers will be asked to evaluate).

Can a paper be submitted to the workshop that is currently under review or will be under review at a conference during the review phase?

It is fine to submit a condensed version (i.e., 4 pages) of a parallel conference submission, if it also fine for the conference in question. Our workshop does not have archival proceedings, and therefore parallel submissions of extended versions to other conferences are acceptable.

=====

- Accepted papers:
1. Paul Pu Liang, Terrance Liu, Liu Ziyin, Russ Salakhutdinov and Louis-Philippe Morency. Think Locally, Act Globally: Federated Learning with Local and Global Representations
 2. Xin Yao, Tianchi Huang, Rui-Xiao Zhang, Ruiyu Li and Lifeng Sun. Federated Learning with Unbiased Gradient Aggregation and Controllable Meta Updating
 3. Daniel Peterson, Pallika Kanani and Virendra Marathe. Private Federated Learning with Domain Adaptation
 4. Daliang Li and Junpu Wang. FedMD: Heterogenous Federated Learning via Model

Distillation

5. Sebastian Caldas, Jakub Konečný, H. Brendan McMahan and Ameet Talwalkar. Mitigating the Impact of Federated Learning on Client Resources
6. Jianyu Wang, Anit Sahu, Zhouyi Yang, Gauri Joshi and Soumya Kar. MATCHA: Speeding Up Decentralized SGD via Matching Decomposition Sampling
7. Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H. Brendan McMahan, Virginia Smith and Ameet Talwalkar. Leaf: A Benchmark for Federated Settings
8. Yihan Jiang, Jakub Konečný, Keith Rush and Sreeram Kannan. Improving Federated Learning Personalization via Model Agnostic Meta Learning
9. Zhicong Liang, Bao Wang, Stanley Osher and Yuan Yao. Exploring Private Federated Learning with Laplacian Smoothing
10. Tribhuvanesh Orekondy, Seong Joon Oh, Yang Zhang, Bernt Schiele and Mario Fritz. Gradient-Leaks: Understanding Deanonimization in Federated Learning
11. Yang Liu, Yan Kang, Xinwei Zhang, Liping Li and Mingyi Hong. A Communication Efficient Vertical Federated Learning Framework
12. Ahmed Khaled, Konstantin Mishchenko and Peter Richtárik. Better Communication Complexity for Local SGD
13. Yang Liu, Xiong Zhang, Shuqi Qin and Xiaoping Lei. Differentially Private Linear Regression over Fully Decentralized Datasets
14. Florian Hartmann, Sunah Suh, Arkadiusz Komarzewski, Tim D. Smith and Ilana Segall. Federated Learning for Ranking Browser History Suggestions
15. Aleksei Triastcyn and Boi Faltings. Federated Learning with Bayesian Differential Privacy
16. Jack Goetz, Kshitiz Malik, Duc Bui, Seungwhan Moon, Honglei Liu and Anuj Kumar. Active Federated Learning
17. Kartikeya Bhardwaj, Wei Chen and Radu Marculescu. FedMAX: Activation Entropy Maximization Targeting Effective Non-IID Federated Learning
18. Mingshu Cong, Zhongming Ou, Yanxin Zhang, Han Yu, Xi Weng, Jiabao Qu, Siu Ming Yiu, Yang Liu and Qiang Yang. Neural Network Optimization for a VCG-based Federated Learning Incentive Mechanism
19. Kai Yang, Tao Fan, Tianjian Chen, Yuanming Shi and Qiang Yang. A Quasi-Newton Method Based Vertical Federated Learning Framework for Logistic Regression
20. Suyi Li, Yong Cheng, Yang Liu and Wei Wang. Abnormal Client Behavior Detection in Federated Learning
21. Songtao Lu, Yawen Zhang, Yunlong Wang and Christina Mack. Learn Electronic Health Records by Fully Decentralized Federated Learning
22. Shicong Cen, Huishuai Zhang, Yuejie Chi, Wei Chen and Tie-Yan Liu. Convergence and Regularization of Distributed Stochastic Variance Reduced Methods
23. Zhaorui Li, Zhicong Huang, Chaochao Chen and Cheng Hong. Quantification of the Leakage in Federated Learning

24. Tzu-Ming Harry Hsu, Hang Qi and Matthew Brown. Measuring the Effects of Non-Identical Data Distribution for Federated Visual Classification

25. Boyue Li, Shicong Cen, Yuxin Chen and Yuejie Chi. Communication-Efficient Distributed Optimization in Networks with Gradient Tracking

26. Khaoula El Mekkaoui, Paul Blomstedt, Diego Mesquita and Samuel Kaski. Towards federated stochastic gradient Langevin dynamics

27. Felix Sattler, Klaus-Robert Müller and Wojciech Samek. Clustered Federated Learning

28. Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh and Brendan McMahan. Backdoor Attacks on Federated Learning and Corresponding Defenses

29. Neta Shoham, Tomer Avidor, Aviv Keren, Nadav Israel, Daniel Benditkis, Liron Mor-Yosef and Itai Zeitak. Overcoming Forgetting in Federated Learning on Non-IID Data

30. Ahmed Khaled and Peter Richtárik. Gradient Descent with Compressed Iterates

31. Jiahuan Luo, Xueyang Wu, Yun Luo, Anbu Huang, Yunfeng Huang, Yang Liu and Qiang Yang. Real-World Image Datasets for Federated Learning

32. Ahmed Khaled, Konstantin Mishchenko and Peter Richtárik. First Analysis of Local GD on Heterogeneous Data

33. Dashan Gao, Ce Ju, Xiguang Wei, Yang Liu, Tianjian Chen and Qiang Yang. HHHFL: Hierarchical Heterogeneous Horizontal Federated Learning for Electroencephalography

The workshop schedule (tentative):

Schedule

08:45 AM	Opening remarks	<i>Fan</i>
08:50 AM	Contributed talk #0	
09:00 AM	Qiang Yang Talk	<i>Yang</i>
09:30 AM	Ameet Talwalkar Talk	<i>Talwalkar</i>
10:00 AM	Coffee break and poster	
10:30 AM	Contributed talk #1	
10:40 AM	Contributed talk #2	
10:50 AM	Max Welling Talk	<i>Welling</i>
11:20 AM	Contributed talk #3	
11:30 AM	Contributed talk #4	
11:40 AM	Dawn Song Talk	<i>Song</i>

		<i>Sattler, El Mekkaoui, Shoham, Hong, Hartmann, Li, Li, Caldas Rivera, Wang, Bhardwaj, Orekondy, KANG, Gao, Cong, Yao, Lu, LUO, Cen, Kairouz, Jiang, Hsu, Triastcyn, Liu, Khaled Ragab Bayoumi, Liang, Faltings, Moon, Li, Fan, Huang, Miao, Qi, Brown, Glass, Wang, Chen, Marculescu, avidor, Wu, Hong, Ju, Rush, Zhang, ZHOU, Beaufays, Zhu, Xia</i>
12:10 PM	Lunch break and poster	
01:30 PM	Dan Ramage Talk	<i>Ramage</i>
02:00 PM	Contributed talk #5	
02:10 PM	Contributed talk #6	
02:20 PM	Francoise Beaufays Talk	
02:50 PM	Contributed talk #7	
03:00 PM	Contributed talk #8	
03:10 PM	Raluca Popa Talk	<i>Popa</i>
03:40 PM	Coffee break and poster	
04:15 PM	Contributed talk #9	
04:25 PM	Contributed talk #10	

	FOCUS: Federate	
04:35 PM	Opportunity Computing for Ubiquitous System	<i>Chen</i>
05:00 PM	Panel disucssion	
06:00 PM	Closing Remark	



Machine Learning for the Developing World (ML4D): Challenges and Risks

Maria De-Arteaga, Amanda Coston, Tejumade Afonja

West 121 + 122, Fri Dec 13, 08:00 AM

As the use of machine learning becomes ubiquitous, there is growing interest in understanding how machine learning can be used to tackle global development challenges. The possibilities are vast, and it is important that we explore the potential benefits of such technologies, which has driven the agenda of the ML4D workshop in the past. However, there is a risk that technology optimism and a categorization of ML4D research as inherently “social good” may result in initiatives failing to account for unintended harms or deviating scarce funds towards initiatives that appear exciting but have no demonstrated effect. Machine learning technologies deployed in developing regions have often been created for different contexts and are trained with data that is not representative of the new deployment setting. Most concerning of all, companies sometimes make the deliberate choice to deploy new technologies in countries with little regulation in order to experiment.

This year’s program will focus on the challenges and risks that arise when deploying machine

learning in developing regions. This one-day workshop will bring together a diverse set of participants from across the globe to discuss essential elements for ensuring ML4D research moves forward in a responsible and ethical manner. Attendees will learn about potential unintended harms that may result from ML4D solutions, technical challenges that currently prevent the effective use of machine learning in vast regions of the world, and lessons that may be learned from other fields.

The workshop will include invited talks, a poster session of accepted papers and panel discussions. We welcome paper submissions featuring novel machine learning research that characterizes or tackle challenges of ML4D, empirical papers that reveal unintended harms of machine learning technology in developing regions, and discussion papers that examine the current state of the art of ML4D and propose paths forward.

Schedule

08:45 AM	Opening Remarks	
<hr/>		
09:00 AM	Deborah Raji	<i>Raji</i>
<hr/>		
09:30 AM	Anubha Sinha	<i>Sinha</i>
<hr/>		
10:00 AM	Kentaro Toyama	<i>Toyama</i>
<hr/>		
10:30 AM	Coffee Break	
<hr/>		
11:00 AM	Breakout sessions	

11:30 AM	Poster session	<p><i>Melese Woldeyohannis, Duvenhage, Waigama, Senay, Babirye, Ayalew, Ogueji, Prabhu, Ravindran, Wahab, Nwokoye, Duckworth, Abera, Mideksa, Benabbou, Sinha, Kiskin, Soden, Isagah, Mwawado, Hussien, Wilder, Omeiza, Rane, Mgaya, Knight, Gonzalez Villarreal, Beyene, Obrocka Tulinska, Cantu Diaz de Leon, Aro, Smith, Famoroti, Vepakomma, Raskar, Bhowmick, Nwokoye, Noriega Campero, Mbelwa</i></p>
<hr/>		
12:30 PM	Lunch	
<hr/>		
02:00 PM	Data sharing in and for Africa	
<hr/>		
02:15 PM	Unsupervised Neural Machine Translation from West African Pidgin to English	<i>Ogueji</i>
<hr/>		
02:30 PM	A Noxious Market for Personal Data	<i>Abdulrahim</i>

02:45 PM	Image Based Identification of Ghanaian Timbers Using the XyloTron: Ravindran Opportunities, Risks and Challenges	
03:00 PM	Coffee and Posters	
03:30 PM	Grace Mutung'u	<i>Mutung'u</i>
04:00 PM	Elisa Celis	<i>Celis</i>
04:30 PM	Rockefeller Foundation and ML4D	<i>Gjekmarkaj</i>
04:35 PM	Partnership on AI and ML4D	<i>Xiang</i>
04:40 PM	Wadhvani AI and ML4D	<i>Mahale</i>
04:45 PM	Panel Discussion: Risks and Challenges in ML4D	
05:35 PM	Closing Remarks and Town Hall	

Abstracts (3):

Abstract 10: Unsupervised Neural Machine Translation from West African Pidgin to English in Machine Learning for the Developing World (ML4D): Challenges and Risks, Ogueji 02:15 PM

Over 1000 languages are spoken across West and Central Africa. Despite the obvious diversity amongst these languages, one language significantly unifies them all - Pidgin English. There are over 75 million speakers of Pidgin English in Nigeria alone, however, there is no known Natural

Language Processing work on this language. This work has five major contributions. First, the provision of a Pidgin English corpus of over 56000 sentences, which is the largest there is so far. Secondly, the training of the first ever Pidgin word vectors. Thirdly, the provision of a Pidgin English to English dictionary of over 1000 words. Fourthly, the alignment of Pidgin English word vectors with English word vectors which achieves a Nearest Neighbor accuracy of 0.1282. This aligned vector will be helpful in the performance of various downstream tasks and transfer of models from English to Pidgin. Finally, the creation of an Unsupervised Neural Machine Translation model between Pidgin English and English which achieves a BLEU score of 20.82 from English to Pidgin and 21.59 from Pidgin to English. In all, this work greatly reduces the barrier of entry for future works on Pidgin English.

Abstract 11: A Noxious Market for Personal Data in Machine Learning for the Developing World (ML4D): Challenges and Risks, Abdulrahim 02:30 PM

Many policymakers, academics and governments have advocated for exchangeable property rights over information as it presents a market solution to what could be considered a market failure. Particularly in jurisdictions such as Africa, Asia or South America, where weaker legal protections and fleeting regulatory enforcement leaves data subjects vulnerable or exploited regardless of the outcome. We argue that whether we could achieve this personal data economy in which individuals have ownership rights akin to property rights over their data should be approached with caution as a solution to ensuring individuals have agency over their data across different legal landscapes.

We present an objection to the use of property rights, a market solution, due to the noxious nature of personal data - which is founded on Satz and Sandell's objection to markets. Ultimately, our rights

over personal data and privacy are borne out of our basic human rights and are a precondition for the self-development, personal fulfilment and the free enjoyment of other fundamental human rights - and putting it up for sale risks corrupting its essence and value.

Abstract 12: Image Based Identification of Ghanaian Timbers Using the XyloTron: Opportunities, Risks and Challenges in Machine Learning for the Developing World (ML4D): Challenges and Risks, Ravindran 02:45 PM

Computer vision systems for wood identification have the potential to empower both producer and consumer countries to combat illegal logging if they can be deployed effectively in the field. In this work, carried out as part of an active international partnership with the support of UNIDO, we constructed and curated a field-relevant image data set to train a classifier for wood identification of \$15\$ commercial Ghanaian woods using the XyloTron system. We tested model performance in the laboratory, and then collected real-world field performance data across multiple sites using multiple XyloTron devices. We present efficacies of the trained model in the laboratory and in the field, discuss practical implications and challenges of deploying machine learning wood identification models, and conclude that field testing is a necessary step - and should be considered the gold-standard - for validating computer vision wood identification systems.

Visually Grounded Interaction and Language

Florian Strub, Abhishek Das, Erik Wijmans, Harm de Vries, Stefan Lee, Alane Suhr, Drew Arad Hudson

West 202 - 204, Fri Dec 13, 08:00 AM

The dominant paradigm in modern natural language understanding is learning statistical language models from text-only corpora. This approach is founded on a distributional notion of semantics, i.e. that the "meaning" of a word is based only on its relationship to other words. While effective for many applications, this approach suffers from limited semantic understanding -- symbols learned this way lack any concrete groundings into the multimodal, interactive environment in which communication takes place. The symbol grounding problem first highlighted this limitation, that "meaningless symbols (i.e. words) cannot be grounded in anything but other meaningless symbols".

On the other hand, humans acquire language by communicating about and interacting within a rich, perceptual environment -- providing concrete groundings, e.g. to objects or concepts either physical or psychological. Thus, recent works have aimed to bridge computer vision, interactive learning, and natural language understanding through language learning tasks based on natural images or through embodied agents performing interactive tasks in physically simulated environments, often drawing on the recent successes of deep learning and reinforcement learning. We believe these lines of research pose a promising approach for building models that do grasp the world's underlying complexity.

The goal of this third ViGIL workshop is to bring together scientists from various backgrounds - machine learning, computer vision, natural language processing, neuroscience, cognitive science, psychology, and philosophy - to share their perspectives on grounding, embodiment, and interaction. By providing this opportunity for cross-discipline discussion, we hope to foster new ideas about how to learn and leverage grounding in machines as well as build new bridges between the science of human cognition and machine learning.

Schedule

08:20 AM	Opening Remarks	<i>Strub, de Vries, Das, Lee, Wijmans, Arad Hudson, Suhr</i>
08:30 AM	Grasping Language	<i>Baldrige</i>
09:10 AM	From Human Language to Agent Action	<i>Thomason</i>
09:50 AM	Coffee Break	
10:30 AM	Spotlight	
10:50 AM	Why language understanding is not a solved problem	<i>McClelland</i>
11:30 AM	Louis-Philippe Morency	<i>Morency</i>
12:10 PM	Poster session	<i>Ross, Mrabet, Subramanian, Cideron, Mu, Bhooshan, Okur Kavil, Delbrouck, Kuo, Lair, Ilharco, Jayram, Herrera Palacio, Fujiyama, Tieleman, Potapenko, Chao, Sutter, Kovaleva, Lai, Wang, Sharma, Cangea, Krishnaswamy, Tsuboi, Kuhnle, Nguyen, Yu, Saha, Xiang, Venkataraman, Kalra, Xie, Doran, Goodwin, Kadav, Daghaghi, Baldrige, Wu</i>

01:50 PM	Lisa Anne Hendricks	<i>Hendricks</i>
02:30 PM	Linda Smith	<i>Smith</i>
03:10 PM	Poster Session	
04:00 PM	Timothy Lillicrap	<i>Lillicrap</i>
04:40 PM	Josh Tenenbaum	<i>Tenenbaum</i>
05:20 PM	Panel Discussion	<i>Smith, Tenenbaum, Hendricks, McClelland, Lillicrap, Thomason, Baldrige, Morency</i>
06:00 PM	Closing Remarks	

Abstracts (7):

Abstract 2: **Grasping Language in Visually Grounded Interaction and Language**, *Baldrige*
08:30 AM

There is a usability gap between manipulation-capable robots and helpful in-home digital agents. Dialog-enabled smart assistants have recently seen widespread adoption, but these cannot move or manipulate objects. By contrast, manipulation-capable and mobile robots are still largely deployed in industrial settings and do not interact with human users. Language-enabled robots can bridge this gap---natural language interfaces help robots and non-experts collaborate to achieve their goals. Navigation in unexplored environments to high-level targets like "Go to the room with a plant" can be facilitated by enabling agents to ask questions and react to human clarifications on-the-fly. Further, high-level instructions like "Put a plate of toast on the table"

require inferring many steps, from finding a knife to operating a toaster. Low-level instructions can serve to clarify these individual steps. Through two new datasets and accompanying models, we study human-human dialog for cooperative navigation, and high- and low-level language instructions for cooking, cleaning, and tidying in interactive home environments. These datasets are a first step towards collaborative, dialog-enabled robots helpful in human spaces.

Abstract 3: From Human Language to Agent Action in Visually Grounded Interaction and Language, *Thomason* 09:10 AM

There is a usability gap between manipulation-capable robots and helpful in-home digital agents. Dialog-enabled smart assistants have recently seen widespread adoption, but these cannot move or manipulate objects. By contrast, manipulation-capable and mobile robots are still largely deployed in industrial settings and do not interact with human users. Language-enabled robots can bridge this gap---natural language interfaces help robots and non-experts collaborate to achieve their goals. Navigation in unexplored environments to high-level targets like "Go to the room with a plant" can be facilitated by enabling agents to ask questions and react to human clarifications on-the-fly. Further, high-level instructions like "Put a plate of toast on the table" require inferring many steps, from finding a knife to operating a toaster. Low-level instructions can serve to clarify these individual steps. Through two new datasets and accompanying models, we study human-human dialog for cooperative navigation, and high- and low-level language instructions for cooking, cleaning, and tidying in interactive home environments. These datasets are a first step towards collaborative, dialog-enabled robots helpful in human spaces.

Abstract 6: Why language understanding is not a solved problem in Visually Grounded Interaction

and Language, *McClelland* 10:50 AM

Over the years, periods of intense excitement about the prospects of machine intelligence and language understanding have alternated with periods of skepticism, to say the least. It is possible to look back over the ~70 year history of this effort and see great progress, and I for one am pleased to see how far we have come. Yet from where I sit we still have a long way to go, and language understanding may be one of those parts of intelligence that will be the hardest to solve. In spite of recent breakthroughs, humans create and comprehend more structured discourse than our current machines. At the same time, psycholinguistic research suggests that humans suffer from some of the same limitations as these machines. How can humans create and comprehend structured arguments given these limitations? Will it be possible for machines to emulate these aspects of human achievement as well?

Abstract 7: Louis-Philippe Morency in Visually Grounded Interaction and Language, *Morency* 11:30 AM

Note that the schedule is not final, and may change.

Abstract 9: Lisa Anne Hendricks in Visually Grounded Interaction and Language, *Hendricks* 01:50 PM

Note that the schedule is not final, and may change.

Abstract 10: Linda Smith in Visually Grounded Interaction and Language, *Smith* 02:30 PM

Note that the schedule is not final, and may change.

Abstract 13: Josh Tenenbaum in Visually Grounded Interaction and Language, *Tenenbaum* 04:40 PM

Note that the schedule is not final, and may change.

Robust AI in Financial Services: Data, Fairness, Explainability, Trustworthiness, and Privacy

Alina Oprea, Avigdor Gal, Isabelle Moulinier, Jiahao Chen, Manuela Veloso, Senthil Kumar, Tanveer Faruque

West 205 - 207, Fri Dec 13, 08:00 AM

The financial services industry has unique needs for robustness when adopting artificial intelligence and machine learning (AI/ML). Many challenges can be described as intricate relationships between algorithmic fairness, explainability, privacy, data management, and trustworthiness. For example, there are ethical and regulatory needs to prove that models used for activities such as credit decisioning and lending are fair and unbiased, or that machine reliance does not cause humans to miss critical pieces of data. The use and protection of customer data necessitates secure and privacy-aware computation, as well as explainability around the use of sensitive data. Some challenges like entity resolution are exacerbated because of scale, highly nuanced data points and missing information.

On top of these fundamental requirements, the financial industry is ripe with adversaries who purport fraud, resulting in large-scale data breaches and loss of confidential information in the financial industry. The need to counteract malicious actors therefore calls for robust methods that can tolerate noise and adversarial corruption of data. However, recent advances in adversarial attacks of AI/ML systems demonstrate how often generic solutions for robustness and security fail, thus highlighting the need for further advances. The challenge of robust AI/ML is further complicated by constraints on data privacy and fairness, as imposed by ethical and regulatory concerns like GDPR.

This workshop aims to bring together researchers and practitioners to discuss challenges for AI/ML in

financial services, and the opportunities such challenges represent to research communities. The workshop will consist of invited talks, panel discussions and short paper presentations, which will showcase ongoing research and novel algorithms resulting from collaboration of AI/ML and cybersecurity communities, as well as the challenges that arise from applying these ideas in domain-specific contexts.

Schedule

08:00 AM	Opening Remarks	<i>Chen, Veloso, , Moulinier, Gal, Oprea, Faruque, Kurshan</i>
08:15 AM	In search of predictability	<i>Perlich</i>
08:45 AM	Oral highlight presentations for selected contributed papers (10 min x 3)	
10:30 AM	Invited Talk by Louiqa Raschid (University of Maryland)	<i>Raschid</i>
11:00 AM	Oral highlight presentations for selected contributed papers (10 min x 6)	
01:30 PM	Understanding equilibrium properties of multi-agent systems	<i>Wooldridge</i>
02:00 PM	Oral highlight presentations for selected contributed papers (10 min x 6)	

02:30 PM	Discussion Panel	<i>Veloso</i>
03:00 PM	Putting Ethical AI to the Vote	<i>Procaccia</i>
		<i>Baracaldo Angel, Neel, Le, Philips, Tao, Chatzis, Suzumura, Wang, BAO, Barocas, Raghavan, Maina, Bryant, Varshney, Speakman, Gill, Schmidt, Compher, Govindarajulu, Sharma, Vepakomma, Swedish, Kalpathy-Cramer, Raskar, Zheng, Pechenizkiy, Schreyer, Ling, Nagpal, Tillman, Veloso, Chen, Wang, Wellman, van Adelsberg, Wood, Buehler, Mahfouz, Alexos, Shearer, Polychroniadou, Stavarache, Efimov, Hall, Zhang, Diana, Ganesh, Ravi, , panda, Renard, Jagielski, Shavit, Williams, Wei, Zhai, Li, Shen, Matsunaga, Choi, Laignelet, Guler, Roa Vicens, Desai, Aigrain, Samoilescu</i>
03:30 PM	Poster Session	
05:15 PM	Invited Talk by Yuan (Alan) Qi (Ant Financial)	<i>Qi</i>

05:45 PM	Closing remarks	
Learning with Rich Experience: Integration of Learning Paradigms		
<i>Zhiting Hu, Andrew Wilson, Chelsea Finn, Lisa Lee, Taylor Berg-Kirkpatrick, Ruslan Salakhutdinov, Eric Xing</i>		
West 208 + 209, Fri Dec 13, 08:00 AM		
<p>Machine learning is about computational methods that enable machines to learn concepts and improve performance from experience. Here, experience can take diverse forms, including data examples, abstract knowledge, interactions and feedback from the environment, other models, and so forth. Depending on different assumptions on the types and amount of experience available there are different learning paradigms, such as supervised learning, active learning, reinforcement learning, knowledge distillation, adversarial learning, and combinations thereof. On the other hand, a hallmark of human intelligence is the ability to learn from all sources of information. In this workshop, we aim to explore various aspects of learning paradigms, particularly theoretical properties and formal connections between them, and new algorithms combining multiple modes of supervisions, etc.</p>		
Schedule		
08:50 AM	Opening Remarks	
09:00 AM	Contributed Oral	
09:10 AM	Invited Talk	<i>Hadsell</i>

09:45 AM	Coffee Break	
10:30 AM	Invited Talk	<i>Mitchell</i>
11:05 AM	Invited Talk	<i>Bilmes</i>
11:40 AM	1min Lightning Talks - I	
11:55 AM	Poster Session	<i>Chourasia, Xu, Cortes, Chang, Nagano, Min, Boecking, Tran, Seyed Ghasemipour, Ding, Mani, Voleti, Fakoor, Xu, Marino, Lee, Tresp, Kagy, Zhang, Poczos, Khandelwal, Bardes, Shelhamer, Zhu, Li, Li, Krasheninnikov, Wang, Jaiswal, Barsoum, Sanjeev, Wattanavekin, Xie, Wu, Yoshida, Kanaa, Khoshfetrat Pakazad, Maasoumy</i>
12:30 PM	Lunch	
02:00 PM	1min Lightning Talks - II	
02:20 PM	Invited Talk	<i>Abbeel</i>
02:55 PM	Invited Talk	<i>Choi</i>
03:30 PM	Coffee Break	
04:20 PM	Invited Talk	<i>Griffiths</i>

04:55 PM	Contributed Oral
05:10 PM	Panel Discussion
05:50 PM	Closing Remarks

Beyond first order methods in machine learning systems

Anastasios Kyrillidis, Albert Berahas, Fred Roosta, Michael W Mahoney

West 211 - 214, Fri Dec 13, 08:00 AM

Optimization lies at the heart of many exciting developments in machine learning, statistics and signal processing. As models become more complex and datasets get larger, finding efficient, reliable and provable methods is one of the primary goals in these fields.

In the last few decades, much effort has been devoted to the development of first-order methods. These methods enjoy a low per-iteration cost and have optimal complexity, are easy to implement, and have proven to be effective for most machine learning applications. First-order methods, however, have significant limitations: (1) they require fine hyper-parameter tuning, (2) they do not incorporate curvature information, and thus are sensitive to ill-conditioning, and (3) they are often unable to fully exploit the power of distributed computing architectures.

Higher-order methods, such as Newton, quasi-Newton and adaptive gradient descent methods, are extensively used in many scientific and engineering domains. At least in theory, these methods possess several nice features: they exploit

local curvature information to mitigate the effects of ill-conditioning, they avoid or diminish the need for hyper-parameter tuning, and they have enough concurrency to take advantage of distributed computing environments. Researchers have even developed stochastic versions of higher-order methods, that feature speed and scalability by incorporating curvature information in an economical and judicious manner. However, often higher-order methods are “undervalued.”

This workshop will attempt to shed light on this statement. Topics of interest include --but are not limited to-- second-order methods, adaptive gradient descent methods, regularization techniques, as well as techniques based on higher-order derivatives.

Schedule

08:00 AM	Opening Remarks	<i>Kyrillidis, Berahas, Roosta, Mahoney</i>
08:15 AM	Talk 1	
09:00 AM	Spotlight talks	<i>Granziol, Pedregosa, Asi, Adolphs, Meng</i>

09:45 AM	Poster Session	<i>Gorbunov, d'Aspremont, Wang, Wang, Ginsburg, Quaglino, Castera, Adya, Granziol, Das, Bollapragada, Pedregosa, Takac, Jahani, Karimireddy, Asi, Daroczy, Adolphs, Rawal, Brandt, Li, Ughi, Romero, Skorokhodov, Scieur, Bae, Mishchenko, Anil, Sharan, Balu, Chen, Yao, Ergen, Grigas, Li, Ba, Roberts, Vaswani, Eftekhari, Sharma</i>
----------	-----------------------	--

10:30 AM	Economical use of second-order information in training machine learning models	<i>Goldfarb</i>
11:15 AM	Adaptive gradient methods: efficient implementation and generalization	
12:00 PM	Lunch break	
02:00 PM	K-FAC: Extensions, improvements, and applications	<i>Martens</i>
02:45 PM	Spotlight talks	<i>Grigas, Yao, Anil, Mishchenko, Scieur</i>
03:30 PM	Poster Session (same as above)	

04:15 PM	Analysis of linear search methods for various gradient approximation schemes for noisy derivative free optimization.	<i>Scheinberg</i>
05:00 PM	Second-order methods for nonconvex optimization with complexity guarantees	<i>Wright</i>
05:45 PM	Final remarks	<i>Kyrillidis, Berahas, Roosta, Mahoney</i>

Abstracts (12):

Abstract 1: **Opening Remarks in Beyond first order methods in machine learning systems**, *Kyrillidis, Berahas, Roosta, Mahoney* 08:00 AM

Opening remarks for the workshop

Abstract 2: **Talk 1 in Beyond first order methods in machine learning systems**, 08:15 AM

Plenary talk #1

Abstract 3: **Spotlight talks in Beyond first order methods in machine learning systems**, *Granzio, Pedregosa, Asi, Adolphs, Meng* 09:00 AM

How does mini-batching affect Curvature information for second order deep learning optimization? Diego Granzio (Oxford); Stephen Roberts (Oxford); Xingchen Wan (Oxford University); Stefan Zohren (University of Oxford); Binxin Ru (University of Oxford); Michael A. Osborne (University of Oxford); Andrew Wilson (NYU); sebastien ehrhardt (Oxford); Dmitry P Vetrov (Higher School of Economics); Timur Garipov

(Samsung AI Center in Moscow)

Acceleration through Spectral Modeling. Fabian Pedregosa (Google); Damien Scieur (Princeton University)

Using better models in stochastic optimization. Hilal Asi (Stanford University); John Duchi (Stanford University)

Ellipsoidal Trust Region Methods for Neural Nets. Leonard Adolphs (ETHZ); Jonas Kohler (ETHZ)

Sub-sampled Newton Methods Under Interpolation. Si Yi Meng (University of British Columbia); Sharan Vaswani (Mila, Université de Montréal); Issam Laradji (University of British Columbia); Mark Schmidt (University of British Columbia); Simon Lacoste-Julien (Mila, Université de Montréal)

Abstract 4: **Poster Session in Beyond first order methods in machine learning systems**, *Gorbunov, d'Aspremont, Wang, Wang, Ginsburg, Quaglino, Castera, Adya, Granzio, Das, Bollapragada, Pedregosa, Takac, Jahani, Karimireddy, Asi, Daroczy, Adolphs, Rawal, Brandt, Li, Ughi, Romero, Skorokhodov, Scieur, Bae, Mishchenko, Anil, Sharan, Balu, Chen, Yao, Ergen, Grigas, Li, Ba, Roberts, Vaswani, Eftekhari, Sharma* 09:45 AM

An Accelerated Method for Derivative-Free Smooth Stochastic Convex Optimization. Eduard Gorbunov (Moscow Institute of Physics and Technology); Pavel Dvurechenskii (WIAS Germany); Alexander Gasnikov (Moscow Institute of Physics and Technology)

Fast Bregman Gradient Methods for Low-Rank Minimization Problems. Radu-Alexandru Dragomir (Université Toulouse 1); Jérôme Bolte (Université Toulouse 1); Alexandre d'Aspremont (Ecole Normale Supérieure)

Gluster: Variance Reduced Mini-Batch SGD with Gradient Clustering. Fartash Faghri (University of Toronto); David Duvenaud (University of Toronto); David Fleet (University of Toronto); Jimmy Ba (University of Toronto)

Neural Policy Gradient Methods: Global Optimality and Rates of Convergence. Lingxiao Wang (Northwestern University); Qi Cai (Northwestern University); Zhuoran Yang (Princeton University); Zhaoran Wang (Northwestern University)

A Gram-Gauss-Newton Method Learning Overparameterized Deep Neural Networks for Regression Problems. Tianle Cai (Peking University); Ruiqi Gao (Peking University); Jikai Hou (Peking University); Siyu Chen (Peking University); Dong Wang (Peking University); Di He (Peking University); Zhihua Zhang (Peking University); Liwei Wang (Peking University)

Stochastic Gradient Methods with Layerwise Adaptive Moments for Training of Deep Networks. Boris Ginsburg (NVIDIA); Oleksii Hrinchuk (NVIDIA); Jason Li (NVIDIA); Vitaly Lavrukhin (NVIDIA); Ryan Leary (NVIDIA); Oleksii Kuchaiev (NVIDIA); Jonathan Cohen (NVIDIA); Huyen Nguyen (NVIDIA); Yang Zhang (NVIDIA)

Accelerating Neural ODEs with Spectral Elements. Alessio Quaglino (NNAISENSE SA); Marco Gallieri (NNAISENSE); Jonathan Masci (NNAISENSE); Jan Koutnik (NNAISENSE)

An Inertial Newton Algorithm for Deep Learning. Camille Castera (CNRS, IRIT); Jérôme Bolte (Université Toulouse 1); Cédric Févotte (CNRS, IRIT); Edouard Pauwels (Toulouse 3 University)

Nonlinear Conjugate Gradients for Scaling Synchronous Distributed DNN Training. Saurabh Adya (Apple); Vinay Palakkode (Apple Inc.); Oncel Tuzel (Apple Inc.)

* How does mini-batching affect Curvature information for second order deep learning optimization? Diego Granzio (Oxford); Stephen Roberts (Oxford); Xingchen Wan (Oxford University); Stefan Zohren (University of Oxford); Binxin Ru (University of Oxford); Michael A. Osborne (University of Oxford); Andrew Wilson (NYU); sebastien ehrhardt (Oxford); Dmitry P Vetrov (Higher School of Economics); Timur Garipov (Samsung AI Center in Moscow)

On the Convergence of a Biased Version of Stochastic Gradient Descent. Rudrajit Das (University of Texas at Austin); Jiong Zhang (UT-Austin); Inderjit S. Dhillon (UT Austin & Amazon)

Adaptive Sampling Quasi-Newton Methods for Derivative-Free Stochastic Optimization. Raghu Bollapragada (Argonne National Laboratory); Stefan Wild (Argonne National Laboratory)

* Acceleration through Spectral Modeling. Fabian Pedregosa (Google); Damien Scieur (Princeton University)

Accelerating Distributed Stochastic L-BFGS by sampled 2nd-Order Information. Jie Liu (Lehigh University); Yu Rong (Tencent AI Lab); Martin Takac (Lehigh University); Junzhou Huang (Tencent AI Lab)

Grow Your Samples and Optimize Better via Distributed Newton CG and Accumulating Strategy. Majid Jahani (Lehigh University); Xi He (Lehigh University); Chenxin Ma (Lehigh University); Aryan Mokhtari (UT Austin); Dheevatsa Mudigere (Intel Labs); Alejandro Ribeiro (University of Pennsylvania); Martin Takac (Lehigh University)

Global linear convergence of trust-region Newton's method without strong-convexity or smoothness.

Sai Praneeth Karimireddy (EPFL); Sebastian Stich (EPFL); Martin Jaggi (EPFL)

FD-Net with Auxiliary Time Steps: Fast Prediction of PDEs using Hessian-Free Trust-Region Methods. Nur Sila Gulgec (Lehigh University); Zheng Shi (Lehigh University); Neil Deshmukh (MIT BeaverWorks - Medlytics); Shamim Pakzad (Lehigh University); Martin Takac (Lehigh University)

* Using better models in stochastic optimization. Hilal Asi (Stanford University); John Duchi (Stanford University)

Tangent space separability in feedforward neural networks. Bálint Daróczy (Institute for Computer Science and Control, Hungarian Academy of Sciences); Rita Aleksziev (Institute for Computer Science and Control, Hungarian Academy of Sciences); Andras Benczur (Hungarian Academy of Sciences)

* Ellipsoidal Trust Region Methods for Neural Nets. Leonard Adolphs (ETHZ); Jonas Kohler (ETHZ)

Closing the K-FAC Generalisation Gap Using Stochastic Weight Averaging. Xingchen Wan (University of Oxford); Diego Granzio (Oxford); Stefan Zohren (University of Oxford); Stephen Roberts (Oxford)

* Sub-sampled Newton Methods Under Interpolation. Si Yi Meng (University of British Columbia); Sharan Vaswani (Mila, Université de Montréal); Issam Laradji (University of British Columbia); Mark Schmidt (University of British Columbia); Simon Lacoste-Julien (Mila, Université de Montréal)

Learned First-Order Preconditioning. Aditya Rawal (Uber AI Labs); Rui Wang (Uber AI); Theodore Moskovitz (Gatsby Computational Neuroscience Unit); Sanyam Kapoor (Uber); Janice Lan (Uber AI);

Jason Yosinski (Uber AI Labs); Thomas Miconi (Uber AI Labs)

Iterative Hessian Sketch in Input Sparsity Time. Charlie Dickens (University of Warwick); Graham Cormode (University of Warwick)

Nonlinear matrix recovery. Florentin Goyens (University of Oxford); Coralia Cartis (Oxford University); Armin Eftekhari (EPFL)

Making Variance Reduction more Effective for Deep Networks. Nicolas Brandt (EPFL); Farnood Salehi (EPFL); Patrick Thiran (EPFL)

Novel and Efficient Approximations for Zero-One Loss of Linear Classifiers. Hiva Ghanbari (Lehigh University); Minhan Li (Lehigh University); Katya Scheinberg (Lehigh)

A Model-Based Derivative-Free Approach to Black-Box Adversarial Examples: BOBYQA. Giuseppe Ughi (University of Oxford)

Distributed Accelerated Inexact Proximal Gradient Method via System of Coupled Ordinary Differential Equations. Chhavi Sharma (IIT Bombay); Vishnu Narayanan (IIT Bombay); Balamurugan Palaniappan (IIT Bombay)

Finite-Time Convergence of Continuous-Time Optimization Algorithms via Differential Inclusions. Orlando Romero (Rensselaer Polytechnic Institute); Mouhacine Benosman (MERL)

Loss Landscape Sightseeing by Multi-Point Optimization. Ivan Skorokhodov (MIPT); Mikhail Burtsev (NI)

* Symmetric Multisecant quasi-Newton methods. Damien Scieur (Samsung AI Research Montreal); Thomas Pumar (Princeton University); Nicolas Boumal (Princeton University)

Does Adam optimizer keep close to the optimal point? Kiwook Bae (KAIST)*; Heechang Ryu (KAIST); Hayong Shin (KAIST)

* Stochastic Newton Method and its Cubic Regularization via Majorization-Minimization. Konstantin Mishchenko (King Abdullah University of Science & Technology (KAUST)); Peter Richtarik (KAUST); Dmitry Koralev (KAUST)

* Full Matrix Preconditioning Made Practical. Rohan Anil (Google); Vineet Gupta (Google); Tomer Koren (Google); Kevin Regan (Google); Yoram Singer (Princeton)

Memory-Sample Tradeoffs for Linear Regression with Small Error. Vatsal Sharan (Stanford University); Aaron Sidford (Stanford); Gregory Valiant (Stanford University)

On the Higher-order Moments in Adam. Zhanhong Jiang (Johnson Controls International); Aditya Balu (Iowa State University); Sin Yong Tan (Iowa State University); Young M Lee (Johnson Controls International); Chinmay Hegde (Iowa State University); Soumik Sarkar (Iowa State University)

h-matrix approximation for Gauss-Newton Hessian. Chao Chen (UT Austin)

* Hessian-Aware trace-Weighted Quantization. Zhen Dong (UC Berkeley); Zhewei Yao (University of California, Berkeley); Amir Gholami (UC Berkeley); Yaohui Cai (Peking University); Daiyaan Arfeen (UC Berkeley); Michael Mahoney ("University of California, Berkeley"); Kurt Keutzer (UC Berkeley)

Random Projections for Learning Non-convex Models. Tolga Ergen (Stanford University); Emmanuel Candes (Stanford University); Mert Pilanci (Stanford)

* New Methods for Regularization Path Optimization via Differential Equations. Paul Grigas (UC Berkeley); Heyuan Liu (University of California, Berkeley)

Hessian-Aware Zeroth-Order Optimization. Haishan Ye (HKUST); Zhichao Huang (HKUST); Cong Fang (Peking University); Chris Junchi Li (Tencent); Tong Zhang (HKUST)

Higher-Order Accelerated Methods for Faster Non-Smooth Optimization. Brian Bullins (TTIC)

Abstract 5: **Economical use of second-order information in training machine learning models in Beyond first order methods in machine learning systems**, *Goldfarb* 10:30 AM

Stochastic gradient descent (SGD) and variants such as Adagrad and Adam, are extensively used today to train modern machine learning models. In this talk we will discuss ways to economically use second-order information to modify both the step size (learning rate) used in SGD and the direction taken by SGD. Our methods adaptively control the batch sizes used to compute gradient and Hessian approximations and ensure that the steps that are taken decrease the loss function with high probability assuming that the latter is self-concordant, as is true for many problems in empirical risk minimization. For such cases we prove that our basic algorithm is globally linearly convergent. A slightly modified version of our method is presented for training deep learning models. Numerical results will be presented that show that it exhibits excellent performance without the need for learning rate tuning. If there is time, additional ways to efficiently make use of second-order information will be presented.

Abstract 6: **Adaptive gradient methods: efficient implementation and generalization in Beyond first order methods in machine learning**

systems, 11:15 AM

Adaptive gradient methods have had a transformative impact in deep learning. We will describe recent theoretical and experimental advances in their understanding, including low-memory adaptive preconditioning, and insights into their generalization ability.

Abstract 8: **K-FAC: Extensions, improvements, and applications in Beyond first order methods in machine learning systems**, *Martens* 02:00 PM

Second order optimization methods have the potential to be much faster than first order methods in the deterministic case, or pre-asymptotically in the stochastic case. However, traditional second order methods have proven ineffective or impractical for neural network training, due in part to the extremely high dimension of the parameter space. Kronecker-factored Approximate Curvature (K-FAC) is second-order optimization method based on a tractable approximation to the Gauss-Newton/Fisher matrix that exploits the special structure present in neural network training objectives. This approximation is neither low-rank nor diagonal, but instead involves Kronecker-products, which allows for efficient estimation, storage and inversion of the curvature matrix. In this talk I will introduce the basic K-FAC method for standard MLPs and then present some more recent work in this direction, including extensions to CNNs and RNNs, both of which requires new approximations to the Fisher. For these I will provide mathematical intuitions and empirical results which speak to their efficacy in neural network optimization. Time permitting, I will also discuss some recent results on large-batch optimization with K-FAC, and the use of adaptive adjustment methods that can eliminate the need for costly hyperparameter tuning.

Abstract 9: **Spotlight talks in Beyond first order methods in machine learning systems**, *Grigas*,

Yao, Anil, Mishchenko, Scieur 02:45 PM

Symmetric Multisecant quasi-Newton methods. Damien Scieur (Samsung AI Research Montreal); Thomas Pumir (Princeton University); Nicolas Boumal (Princeton University)

Stochastic Newton Method and its Cubic Regularization via Majorization-Minimization. Konstantin Mishchenko (King Abdullah University of Science & Technology (KAUST)); Peter Richtarik (KAUST); Dmitry Koralev (KAUST)

Full Matrix Preconditioning Made Practical. Rohan Anil (Google); Vineet Gupta (Google); Tomer Koren (Google); Kevin Regan (Google); Yoram Singer (Princeton)

Hessian-Aware trace-Weighted Quantization. Zhen Dong (UC Berkeley); Zhewei Yao (University of California, Berkeley); Amir Gholami (UC Berkeley); Yaohui Cai (Peking University); Daiyaan Arfeen (UC Berkeley); Michael Mahoney ("University of California, Berkeley"); Kurt Keutzer (UC Berkeley)

New Methods for Regularization Path Optimization via Differential Equations. Paul Grigas (UC Berkeley); Heyuan Liu (University of California, Berkeley)

Abstract 10: **Poster Session (same as above) in Beyond first order methods in machine learning systems**, 03:30 PM

An Accelerated Method for Derivative-Free Smooth Stochastic Convex Optimization. Eduard Gorbunov (Moscow Institute of Physics and Technology); Pavel Dvurechenskii (WIAS Germany); Alexander Gasnikov (Moscow Institute of Physics and Technology)

Fast Bregman Gradient Methods for Low-Rank Minimization Problems. Radu-Alexandru Dragomir (Université Toulouse 1); Jérôme Bolte (Université

Toulouse 1); Alexandre d'Aspremont (Ecole Normale Supérieure)

Gluster: Variance Reduced Mini-Batch SGD with Gradient Clustering. Fartash Faghri (University of Toronto); David Duvenaud (University of Toronto); David Fleet (University of Toronto); Jimmy Ba (University of Toronto)

Neural Policy Gradient Methods: Global Optimality and Rates of Convergence. Lingxiao Wang (Northwestern University); Qi Cai (Northwestern University); Zhuoran Yang (Princeton University); Zhaoran Wang (Northwestern University)

A Gram-Gauss-Newton Method Learning Overparameterized Deep Neural Networks for Regression Problems. Tianle Cai (Peking University); Ruiqi Gao (Peking University); Jikai Hou (Peking University); Siyu Chen (Peking University); Dong Wang (Peking University); Di He (Peking University); Zhihua Zhang (Peking University); Liwei Wang (Peking University)

Stochastic Gradient Methods with Layerwise Adaptive Moments for Training of Deep Networks. Boris Ginsburg (NVIDIA); Oleksii Hrinchuk (NVIDIA); Jason Li (NVIDIA); Vitaly Lavrukhin (NVIDIA); Ryan Leary (NVIDIA); Oleksii Kuchaiev (NVIDIA); Jonathan Cohen (NVIDIA); Huyen Nguyen (NVIDIA); Yang Zhang (NVIDIA)

Accelerating Neural ODEs with Spectral Elements. Alessio Quaglino (NNAISENSE SA); Marco Gallieri (NNAISENSE); Jonathan Masci (NNAISENSE); Jan Koutnik (NNAISENSE)

An Inertial Newton Algorithm for Deep Learning. Camille Castera (CNRS, IRIT); Jérôme Bolte (Université Toulouse 1); Cédric Févotte (CNRS, IRIT); Edouard Pauwels (Toulouse 3 University)

Nonlinear Conjugate Gradients for Scaling

Synchronous Distributed DNN Training. Saurabh Adya (Apple); Vinay Palakkode (Apple Inc.); Oncel Tuzel (Apple Inc.)

* How does mini-batching affect Curvature information for second order deep learning optimization? Diego Granzio (Oxford); Stephen Roberts (Oxford); Xingchen Wan (Oxford University); Stefan Zohren (University of Oxford); Binxin Ru (University of Oxford); Michael A. Osborne (University of Oxford); Andrew Wilson (NYU); sebastien ehrhardt (Oxford); Dmitry P Vetrov (Higher School of Economics); Timur Garipov (Samsung AI Center in Moscow)

On the Convergence of a Biased Version of Stochastic Gradient Descent. Rudrajit Das (University of Texas at Austin); Jiong Zhang (UT-Austin); Inderjit S. Dhillon (UT Austin & Amazon)

Adaptive Sampling Quasi-Newton Methods for Derivative-Free Stochastic Optimization. Raghu Bollapragada (Argonne National Laboratory); Stefan Wild (Argonne National Laboratory)

* Acceleration through Spectral Modeling. Fabian Pedregosa (Google); Damien Scieur (Princeton University)

Accelerating Distributed Stochastic L-BFGS by sampled 2nd-Order Information. Jie Liu (Lehigh University); Yu Rong (Tencent AI Lab); Martin Takac (Lehigh University); Junzhou Huang (Tencent AI Lab)

Grow Your Samples and Optimize Better via Distributed Newton CG and Accumulating Strategy. Majid Jahani (Lehigh University); Xi He (Lehigh University); Chenxin Ma (Lehigh University); Aryan Mokhtari (UT Austin); Dheevatsa Mudigere (Intel Labs); Alejandro Ribeiro (University of Pennsylvania); Martin Takac (Lehigh University)

Global linear convergence of trust-region Newton's method without strong-convexity or smoothness. Sai Praneeth Karimireddy (EPFL); Sebastian Stich (EPFL); Martin Jaggi (EPFL)

FD-Net with Auxiliary Time Steps: Fast Prediction of PDEs using Hessian-Free Trust-Region Methods. Nur Sila Gulgec (Lehigh University); Zheng Shi (Lehigh University); Neil Deshmukh (MIT BeaverWorks - Medlytics); Shamim Pakzad (Lehigh University); Martin Takac (Lehigh University)

* Using better models in stochastic optimization. Hilal Asi (Stanford University); John Duchi (Stanford University)

Tangent space separability in feedforward neural networks. Bálint Daróczy (Institute for Computer Science and Control, Hungarian Academy of Sciences); Rita Aleksziev (Institute for Computer Science and Control, Hungarian Academy of Sciences); Andras Benczur (Hungarian Academy of Sciences)

* Ellipsoidal Trust Region Methods for Neural Nets. Leonard Adolphs (ETHZ); Jonas Kohler (ETHZ)

Closing the K-FAC Generalisation Gap Using Stochastic Weight Averaging. Xingchen Wan (University of Oxford); Diego Granzio (Oxford); Stefan Zohren (University of Oxford); Stephen Roberts (Oxford)

* Sub-sampled Newton Methods Under Interpolation. Si Yi Meng (University of British Columbia); Sharan Vaswani (Mila, Université de Montréal); Issam Laradji (University of British Columbia); Mark Schmidt (University of British Columbia); Simon Lacoste-Julien (Mila, Université de Montréal)

Learned First-Order Preconditioning. Aditya Rawal

(Uber AI Labs); Rui Wang (Uber AI); Theodore Moskovitz (Gatsby Computational Neuroscience Unit); Sanyam Kapoor (Uber); Janice Lan (Uber AI); Jason Yosinski (Uber AI Labs); Thomas Miconi (Uber AI Labs)

Iterative Hessian Sketch in Input Sparsity Time. Charlie Dickens (University of Warwick); Graham Cormode (University of Warwick)

Nonlinear matrix recovery. Florentin Goyens (University of Oxford); Coralía Cartis (Oxford University); Armin Eftekhari (EPFL)

Making Variance Reduction more Effective for Deep Networks. Nicolas Brandt (EPFL); Farnood Salehi (EPFL); Patrick Thiran (EPFL)

Novel and Efficient Approximations for Zero-One Loss of Linear Classifiers. Hiva Ghanbari (Lehigh University); Minhan Li (Lehigh University); Katya Scheinberg (Lehigh)

A Model-Based Derivative-Free Approach to Black-Box Adversarial Examples: BOBYQA. Giuseppe Ughi (University of Oxford)

Distributed Accelerated Inexact Proximal Gradient Method via System of Coupled Ordinary Differential Equations. Chhavi Sharma (IIT Bombay); Vishnu Narayanan (IIT Bombay); Balamurugan Palaniappan (IIT Bombay)

Finite-Time Convergence of Continuous-Time Optimization Algorithms via Differential Inclusions. Orlando Romero (Rensselaer Polytechnic Institute); Mouhacine Benosman (MERL)

Loss Landscape Sightseeing by Multi-Point Optimization. Ivan Skorokhodov (MIPT); Mikhail Burtsev (NI)

* Symmetric Multisecant quasi-Newton methods.

Damien Scieur (Samsung AI Research Montreal); Thomas Pumis (Princeton University); Nicolas Boumal (Princeton University)

Does Adam optimizer keep close to the optimal point? Kiwook Bae (KAIST)*; Heechang Ryu (KAIST); Hayong Shin (KAIST)

* Stochastic Newton Method and its Cubic Regularization via Majorization-Minimization. Konstantin Mishchenko (King Abdullah University of Science & Technology (KAUST)); Peter Richtarik (KAUST); Dmitry Koralev (KAUST)

* Full Matrix Preconditioning Made Practical. Rohan Anil (Google); Vineet Gupta (Google); Tomer Koren (Google); Kevin Regan (Google); Yoram Singer (Princeton)

Memory-Sample Tradeoffs for Linear Regression with Small Error. Vatsal Sharan (Stanford University); Aaron Sidford (Stanford); Gregory Valiant (Stanford University)

On the Higher-order Moments in Adam. Zhanhong Jiang (Johnson Controls International); Aditya Balu (Iowa State University); Sin Yong Tan (Iowa State University); Young M Lee (Johnson Controls International); Chinmay Hegde (Iowa State University); Soumik Sarkar (Iowa State University)

h-matrix approximation for Gauss-Newton Hessian. Chao Chen (UT Austin)

* Hessian-Aware trace-Weighted Quantization. Zhen Dong (UC Berkeley); Zhewei Yao (University of California, Berkeley); Amir Gholami (UC Berkeley); Yaohui Cai (Peking University); Daiyaan Arfeen (UC Berkeley); Michael Mahoney ("University of California, Berkeley"); Kurt Keutzer (UC Berkeley)

Random Projections for Learning Non-convex

Models. Tolga Ergen (Stanford University); Emmanuel Candes (Stanford University); Mert Pilanci (Stanford)

* New Methods for Regularization Path Optimization via Differential Equations. Paul Grigas (UC Berkeley); Heyuan Liu (University of California, Berkeley)

Hessian-Aware Zeroth-Order Optimization. Haishan Ye (HKUST); Zhichao Huang (HKUST); Cong Fang (Peking University); Chris Junchi Li (Tencent); Tong Zhang (HKUST)

Higher-Order Accelerated Methods for Faster Non-Smooth Optimization. Brian Bullins (TTIC)

Abstract 11: Analysis of linear search methods for various gradient approximation schemes for noisy derivative free optimization. in Beyond first order methods in machine learning systems, Scheinberg 04:15 PM

We develop convergence analysis of a modified line search method for objective functions whose value is computed with noise and whose gradient estimates are not directly available. The noise is assumed to be bounded in absolute value without any additional assumptions. In this case, gradient approximation can be constructed via interpolation or sample average approximation of smoothing gradients and thus they are always inexact and possibly random. We extend the framework based on stochastic methods which was developed to provide analysis of a standard line-search method with exact function values and random gradients to the case of noisy function. We introduce a condition on the gradient which when satisfied with some sufficiently large probability at each iteration, guarantees convergence properties of the line search method. We derive expected complexity bounds for convex, strongly convex and nonconvex functions. We motivate these results with several recent papers related to policy optimization.

Abstract 12: **Second-order methods for nonconvex optimization with complexity guarantees in Beyond first order methods in machine learning systems**, *Wright* 05:00 PM

We consider problems of smooth nonconvex optimization: unconstrained, bound-constrained, and with general equality constraints. We show that algorithms for these problems that are widely used in practice can be modified slightly in ways that guarantees convergence to approximate first- and second-order optimal points with complexity guarantees that depend on the desired accuracy. The methods we discuss are constructed from Newton's method, the conjugate gradient method, log-barrier method, and augmented Lagrangians. (In some cases, special structure of the objective function makes for only a weak dependence on the accuracy parameter.) Our methods require Hessian information only in the form of Hessian-vector products, so do not require the Hessian to be evaluated and stored explicitly. This talk describes joint work with Clement Royer, Yue Xie, and Michael O'Neill.

Abstract 13: **Final remarks in Beyond first order methods in machine learning systems**, *Kyrillidis, Berahas, Roosta, Mahoney* 05:45 PM

Final remarks for the workshop

CiML 2019: Machine Learning Competitions for All

Adrienne Mendrik, Wei-Wei Tu, Isabelle Guyon, Evelyne Viegas, Ming LI

West 215 + 216, Fri Dec 13, 08:00 AM

Challenges in machine learning and data science are open online competitions that address problems by providing datasets or simulated environments. They measure the performance of machine learning

algorithms with respect to a given problem. The playful nature of challenges naturally attracts students, making challenges a great teaching resource. However, in addition to the use of challenges as educational tools, challenges have a role to play towards a better democratization of AI and machine learning. They function as cost effective problem-solving tools and a means of encouraging the development of re-usable problem templates and open-sourced solutions. However, at present, the geographic, sociological repartition of challenge participants and organizers is very biased. While recent successes in machine learning have raised much hopes, there is a growing concern that the societal and economical benefits might increasingly be in the power and under control of a few.

CiML (Challenges in Machine Learning) is a forum that brings together workshop organizers, platform providers, and participants to discuss best practices in challenge organization and new methods and application opportunities to design high impact challenges. Following the success of previous years' workshops, we will reconvene and discuss new opportunities for broadening our community.

For this sixth edition of the CiML workshop at NeurIPS our objective is twofold: (1) We aim to enlarge the community, fostering diversity in the community of participants and organizers; (2) We aim to promote the organization of challenges for the benefit of more diverse communities.

The workshop provides room for discussion on these topics and aims to bring together potential partners to organize such challenges and stimulate "machine learning for good", i.e. the organization of challenges for the benefit of society. We have invited prominent speakers that have experience in this domain.

Schedule

08:00 AM	Welcome and Opening Remarks	<i>Mendrik, Tu, Guyon, Viegas, LI</i>		
08:15 AM	Amir Banifatemi (XPrize) "AI for Good via Machine Learning Challenges"	<i>Banifatemi</i>		
09:00 AM	Emily Bender (University of Washington) "Making Stakeholder Impacts Visible in the Evaluation Cycle: Towards Fairness-Integrated Shared Tasks and Evaluation Metrics"	<i>Bender</i>		
09:45 AM	Coffee Break			
10:30 AM	Dina Machuve (Nelson Mandela African Institution of Science and Technology) "Machine Learning Competitions: The Outlook from Africa"	<i>Machuve</i>		
11:15 AM	Dog Image Generation Competition on Kaggle	<i>Kan</i>		
11:30 AM	Learning To Run a Power Network Competition	<i>Donnot</i>		
11:45 AM	The AI Driving Olympics: An Accessible Robot Learning Benchmark	<i>Walter</i>		
12:00 PM	Conclusion on TrackML, a Particle Physics Tracking Machine Learning Challenge Combining Accuracy and Inference Speed	<i>Rousseau, vlimant</i>		
12:15 PM	Catered Lunch and Poster Viewing (in Workshop Room)	<i>Stolovitzky, Pradhan, Duboue, Tang, Natekin, Bondi, Bouthillier, Milani, Müller, Holzinger, Harrer, Day, Ustyuzhanin, Guss, Mirmomeni</i>		
02:00 PM	Yang Yu (Nanjing University) on Machine Learning Challenges to Advance AI in China	<i>Yu</i>		
02:45 PM	Design and Analysis of Experiments: A Challenge Approach in Teaching	<i>Pavao</i>		
03:00 PM	The model-to-data paradigm: overcoming data access barriers in biomedical competitions	<i>Guinney</i>		

03:15 PM	The Deep Learning Epilepsy Detection Challenge: Design, Implementation, and Test of a New Crowd-Sourced AI Challenge Ecosystem	<i>Kiral</i>
03:30 PM	Coffee Break	
04:15 PM	Frank Hutter (University of Freiburg) "A Proposal for a New Competition Design Emphasizing Scientific Insights"	<i>Hutter</i>
05:00 PM	Open Space Topic "The Organization of Challenges for the Benefit of More Diverse Communities"	<i>Mendrik, Guyon, Tu, Viegas, LI</i>

Abstracts (11):

Abstract 2: Amir Banifatemi (XPrize) "AI for Good via Machine Learning Challenges" in CiML 2019: Machine Learning Competitions for All, Banifatemi 08:15 AM

"AI for Good" efforts (e.g., applications work in sustainability, education, health, financial inclusion, etc.) have demonstrated the capacity to simultaneously advance intelligent system research and the greater good. Unfortunately, the majority of research that could find motivation in real-world "good" problems still center on problems with industrial or toy problem performance baselines.

Competitions can serve as an important shaping

reward for steering academia towards research that is simultaneously impactful on our state of knowledge and the state of the world. This talk covers three aspects of AI for Good competitions. First, we survey current efforts within the AI for Good application space as a means of identifying current and future opportunities. Next we discuss how more qualitative notions of "Good" can be used as benchmarks in addition to more quantitative competition objective functions. Finally, we will provide notes on building coalitions of domain experts to develop and guide socially-impactful competitions in machine learning.

Abstract 3: Emily Bender (University of Washington) "Making Stakeholder Impacts Visible in the Evaluation Cycle: Towards Fairness-Integrated Shared Tasks and Evaluation Metrics" in CiML 2019: Machine Learning Competitions for All, Bender 09:00 AM

In a typical machine learning competition or shared task, success is measured in terms of systems' ability to reproduce gold-standard labels. The potential impact of the systems being developed on stakeholder populations, if considered at all, is studied separately from system `performance'. Given the tight train-eval cycle of both shared tasks and system development in general, we argue that making disparate impact on vulnerable populations visible in dataset and metric design will be key to making the potential for such impact present and salient to developers. We see this as an effective way to promote the development of machine learning technology that is helpful for people, especially those who have been subject to marginalization. This talk will explore how to develop such shared tasks, considering task choice, stakeholder community input, and annotation and metric design desiderata.

Joint work with Hal Daumé III, University of Maryland, Bernease Herman, University of Washington, and Brandeis Marshall, Spelman

College.

Abstract 5: Dina Machuve (Nelson Mandela African Institution of Science and Technology) “Machine Learning Competitions: The Outlook from Africa” in CiML 2019: Machine Learning Competitions for All, Machuve 10:30 AM

The current AI landscape in Africa mainly focuses on capacity building. The ongoing efforts to strengthen the AI capacity in Africa are organized in summer schools, workshops, meetups, competitions and one long-term program at the Masters level. The main AI initiatives driving the AI capacity building agenda in Africa include a) Deep Learning Indaba, b) Data Science Africa, c) Data Science Nigeria, d) Nairobi Women in Machine Learning and Data Science, e) Zindi and f) The African Master's in Machine Intelligence (AMMI) at AIMS. The talk will summarize our experience on low participation of African AI developers at machine learning competitions and our recommendations to address the current challenges.

Abstract 6: Dog Image Generation Competition on Kaggle in CiML 2019: Machine Learning Competitions for All, Kan 11:15 AM

We present a novel format of machine learning competitions where a user submits code that generates images trained on training samples, the code then runs on Kaggle, produces dog images, and user receives scores for the performance of their generative content based on 1. quality of images, 2. diversity of images, and 3. memorization penalty. This style of competition targets the usage of Generative Adversarial Networks (GAN)[4], but is open for all generative models. Our implementation addresses overfitting by incorporating two different pre-trained neural networks, as well as two separate "ground truth" image datasets, for the public and private leaderboards. We also have an enclosed compute environment to prevent submissions of

non-generated images. In this paper, we describe both the algorithmic and system design of our competition, as well as sharing our lessons learned from running this competition [6] in July 2019 with 900+ teams participating and over 37,000 submissions and their code received.

Abstract 7: Learning To Run a Power Network Competition in CiML 2019: Machine Learning Competitions for All, Donnot 11:30 AM

We present the results of the first edition as well as some perspective for a next potential edition of the "Learning To Run a Power Network" (L2RPN) competition to test the potential of Reinforcement Learning to solve a real-world problem of great practical importance: controlling power transportation in power grids while keeping people and equipment safe.

Abstract 8: The AI Driving Olympics: An Accessible Robot Learning Benchmark in CiML 2019: Machine Learning Competitions for All, Walter 11:45 AM

Despite recent breakthroughs, the ability of deep learning and reinforcement learning to outperform traditional approaches to control physically embodied robotic agents remains largely unproven. To help bridge this gap, we have developed the "AI Driving Olympics" (AI-DO), a competition with the objective of evaluating the state-of-the-art in machine learning and artificial intelligence for mobile robotics. Based on the simple and well specified autonomous driving and navigation environment called "Duckietown," AI-DO includes a series of tasks of increasing complexity—from simple lane-following to fleet management. For each task, we provide tools for competitors to use in the form of simulators, data logs, code templates, baseline implementations, and low-cost access to robotic hardware. We evaluate submissions in simulation online, on standardized hardware environments, and finally at the competition events.

We have held successful AI-DO competitions at NeurIPS 2018 and ICRA 2019, and will be holding AI-DO 3 at NeurIPS 2020. Together, these competitions highlight the need for better benchmarks, which are lacking in robotics, as well as improved mechanisms to bridge the gap between simulation and reality.

Abstract 10: Catered Lunch and Poster Viewing (in Workshop Room) in CiML 2019: Machine Learning Competitions for All, *Stolovitzky, Pradhan, Duboue, Tang, Natekin, Bondi, Bouthillier, Milani, Müller, Holzinger, Harrer, Day, Ustyuzhanin, Guss, Mirmomeni* 12:15 PM

Accepted Posters

Kandinsky Patterns: An open toolbox for creating explainable machine learning challenges
Heimo Muller · Andreas Holzinger

MOCA: An Unsupervised Algorithm for Optimal Aggregation of Challenge Submissions
Robert Vogel · Mehmet Eren Ahsen · Gustavo A. Stolovitzky

FDL: Mission Support Challenge
Luís F. Simões · Ben Day · Vinutha M. Shreenath · Callum Wilson

From data challenges to collaborative gig science. Cooperative research process and platform
Andrey Ustyuzhanin · Mikhail Belous · Leyla Khatbullina · Giles Strong

Smart(er) Machine Learning for Practitioners
Prabhu Pradhan

Improving Reproducibility of Benchmarks
Xavier Bouthillier

Guaranteeing Reproducibility in Deep Learning Competitions
Brandon Houghton

Organizing crowd-sourced AI challenges in enterprise environments: opportunities and challenges
Mahtab Mirmomeni · Isabell Kiral · Subhrajit Roy · Todd Mummert · Alan Braz · Jason Tsay · Jianbin Tang · Umar Asif · Thomas Schaffter · Eren Mehmet · Bruno De Assis Marques · Stefan Maetschke · Rania Khalaf · Michal Rosen-Zvi · John Cohn · Gustavo Stolovitzky · Stefan Harrer

WikiCities: a Feature Engineering Educational Resource
Pablo Duboue

Reinforcement Learning Meets Information Seeking: Dynamic Search Challenge
Zhiwen Tang · Grace Hui Yang

AI Journey 2019: School Tests Solving Competition
Alexey Natekin · Peter Romov · Valentin Malykh

A BIRDSAI View for Conservation
Elizabeth Bondi · Milind Tambe · Raghav Jain · Palash Aggrawal · Saket Anand · Robert Hannaford · Ashish Kapoor · Jim Piavis · Shital Shah · Lucas Joppa · Bistra Dilkina

Abstract 12: Design and Analysis of Experiments: A Challenge Approach in Teaching in CiML 2019: Machine Learning Competitions for All, *Pavao* 02:45 PM

Over the past few years, we have explored the benefits of involving students both in organizing and in participating in challenges as a pedagogical tool, as part of an international collaboration. Engaging in the design and resolution of a competition can be seen as a hands-on means of learning proper design and analysis of experiments and gaining a deeper understanding other aspects of Machine Learning. Graduate students of University Paris-Sud (Paris, France) are involved in class projects in creating a challenge end-to-end, from defining the

research problem, collecting or formatting data, creating a starting kit, to implementing and testing the website. The application domains and types of data are extremely diverse: medicine, ecology, marketing, computer vision, recommendation, text processing, etc. The challenges thus created are then used as class projects of undergraduate students who have to solve them, both at University Paris-Sud, and at Rensselaer Polytechnic Institute (RPI, New York, USA), to provide rich learning experiences at scale. New this year, students are involved in creating challenges motivated by “AI for good” and will create re-usable templates to inspire others to create challenges for the benefit of humanity.

Abstract 13: The model-to-data paradigm: overcoming data access barriers in biomedical competitions in CiML 2019: Machine Learning Competitions for All, *Guinney* 03:00 PM

Data competitions often rely on the physical distribution of data to challenge participants, a significant limitation given that much data is proprietary, sensitive, and often non-shareable. To address this, the DREAM Challenges has advanced a challenge framework called model-to-data (MTD), requiring participants to submit re-runnable algorithms instead of model predictions. The DREAM organization has successfully completed multiple MTD-based challenges, and is expanding this approach to unlock highly sensitive and non-distributable human data for use in biomedical data challenges.

Abstract 16: Frank Hutter (University of Freiburg) "A Proposal for a New Competition Design Emphasizing Scientific Insights" in CiML 2019: Machine Learning Competitions for All, *Hutter* 04:15 PM

The typical setup in machine learning competitions is to provide one or more datasets and a performance metric, leaving it entirely up to

participants which approach to use, how to engineer better features, whether and how to pretrain models on related data, how to tune hyperparameters, how to combine multiple models in an ensemble, etc. The fact that work on each of these components often leads to substantial improvements has several consequences: (1) amongst several skilled teams, the one with the most manpower and engineering drive often wins; (2) it is often unclear *why* one entry performs better than another one; and (3) scientific insights remain limited.

Based on my experience in both participating in several challenges and also organizing some, I will propose a new competition design that instead emphasizes scientific insight by dividing the various ways in which teams could improve performance into (largely orthogonal) modular components, each of which defines its own competition. E.g., one could run a competition focussing only on effective hyperparameter tuning of a given pipeline (across private datasets). With the same code base and datasets, one could likewise run a competition focussing only on finding better neural architectures, or only better preprocessing methods, or only a better training pipeline, or only better pre-training methods, etc. One could also run multiple of these competitions in parallel, hot-swapping better components found in one competition into the other competitions. I will argue that the result would likely be substantially more valuable in terms of scientific insights than traditional competitions and may even lead to better final performance.

Abstract 17: Open Space Topic “The Organization of Challenges for the Benefit of More Diverse Communities” in CiML 2019: Machine Learning Competitions for All, *Mendrik, Guyon, Tu, Viegas, LI* 05:00 PM

“Open Space” is a technique for running meetings where the participants create and manage the agenda themselves. Participants can propose ideas that address the open space topic, these will be

divided into various sessions that all other participants can join and brainstorm about. After the open space we will collect all the ideas and see whether we could write a whitepaper on this topic with all participants.

AI for Humanitarian Assistance and Disaster Response

Ritwik Gupta, Robin Murphy, Trevor Darrell, Eric Heim, Zhangyang Wang, Bryce Goodman, Piotr Bilinski

West 217 - 219, Fri Dec 13, 08:00 AM

Natural disasters are one of the oldest threats to not just individuals but to the societies they co-exist in. As a result, humanity has ceaselessly sought way to provide assistance to people in need after disasters have struck. Further, natural disasters are but a single, extreme example of the many possible humanitarian crises. Disease outbreak, famine, and oppression against disadvantaged groups can pose even greater dangers to people that have less obvious solutions.

In this proposed workshop, we seek to bring together the Artificial Intelligence (AI) and Humanitarian Assistance and Disaster Response (HADR) communities in order to bring AI to bear on real-world humanitarian crises.

Through this workshop, we intend to establish meaningful dialogue between the communities.

By the end of the workshop, the NeurIPS research community can come to understand the practical challenges of in aiding those in crisis, while the HADR can understand the landscape that is the state of art and practice in AI.

Through this, we seek to begin establishing a pipeline of transitioning the research created by the NeurIPS community to real-world humanitarian issues.

Schedule

08:00 AM	Introduction and Welcome	<i>Gupta, Sajeev</i>
08:15 AM	Invited Talks (x4)	<i>Matias, Adole, Brown</i>
10:15 AM	Spotlight Talks (x6)	<i>Kruspe, Dalmasso, Schrempf, Oh, Doshi, Lu</i>
11:30 AM	Lunch	
01:00 PM	Invited Talks (x4)	<i>Rasmussen, Stromberg, Darrell</i>
03:00 PM	Spotlight Talks (x6)	<i>Wang, Seo, Veitch-Michaelis, Sidrane, Kapadia, Nevo, Dubey</i>
04:30 PM	Convergence: Two-Way Limitations in Taking Theory to Applications	<i>Dzombak, Yang</i>
05:15 PM	Poster Session	

Abstracts (4):

Abstract 2: Invited Talks (x4) in AI for Humanitarian Assistance and Disaster Response, *Matias, Adole, Brown* 08:15 AM

- * Yossi Matias
- * Tracy Adole
- * Col Jason Brown
- * Yang Cai

Abstract 3: Spotlight Talks (x6) in AI for Humanitarian Assistance and Disaster Response, *Kruspe, Dalmasso, Schrempf, Oh, Doshi, Lu* 10:15 AM

TBD based on accepted papers

Abstract 5: **Invited Talks (x4) in AI for Humanitarian Assistance and Disaster Response**, *Rasmussen, Stromberg, Darrell* 01:00 PM

- * Eric Rasmussen
- * Maj Megan Stromberg
- * TBD
- * TBD

Abstract 6: **Spotlight Talks (x6) in AI for Humanitarian Assistance and Disaster Response**, *Wang, Seo, Veitch-Michaelis, Sidrane, Kapadia, Nevo, Dubey* 03:00 PM

TBD based on accepted papers

Shared Visual Representations in Human and Machine Intelligence

Arturo Deza, Joshua Peterson, Apurva Ratan Murty, Tom Griffiths

West 220 - 222, Fri Dec 13, 08:00 AM

The goal of the Shared Visual Representations in Human and Machine Intelligence workshop is to disseminate relevant, parallel findings in the fields of computational neuroscience, psychology, and cognitive science that may inform modern machine learning methods. In the past few years, machine learning methods---especially deep neural networks---have widely permeated the vision science, cognitive science, and neuroscience communities.

As a result, scientific modeling in these fields has greatly benefited, producing a swath of potentially critical new insights into human learning and intelligence, which remains the gold standard for many tasks. However, the machine learning community has been largely

unaware of these cross-disciplinary insights and analytical tools, which may help to solve many of the current problems that ML theorists and engineers face today (e.g., adversarial attacks, compression, continual learning, and unsupervised learning). Thus we propose to invite leading cognitive scientists with strong computational backgrounds to disseminate their findings to the machine learning community with the hope of closing the loop by nourishing new ideas and creating cross-disciplinary collaborations.

Schedule

08:50 AM	Opening Remarks	<i>Deza, Peterson, Murty, Griffiths</i>
09:00 AM	Olivier Henaff (DeepMind)	<i>Henaff</i>
09:25 AM	Irina Higgins (DeepMind)	<i>Higgins</i>
09:50 AM	Bill Freeman (MIT)	<i>Freeman</i>
10:15 AM	Q&A from the Audience: Ask a Neuro / Cognitive Scientist	<i>Griffiths, DiCarlo, Konkle</i>
10:45 AM	Coffee Break	
11:00 AM	Ruairidh Battleday (Princeton)	<i>Battleday</i>
11:15 AM	Will Xiao (Harvard)	<i>Xiao</i>
11:30 AM	Erin Grant (UC Berkeley)	<i>Grant</i>
11:45 AM	Andrei Barbu (MIT)	<i>Barbu</i>
12:10 PM	Mike Tarr (CMU)	<i>Tarr</i>

12:35 PM	James DiCarlo (MIT) <i>DiCarlo</i>	
01:00 PM	Lunch on your own	
		<i>Harris, White, Choung, Shinozaki, Pal, Hermann, Borowski, Fosco, Firestone, Veerabadran, Lahner, Ryali, Doshi, Singh, Zhou, Besserve, Chang, Newman, Niranjan, Hare, Mihai, Savvides, Kornblith, Funke, Oliva, de Sa, Krotov, Conwell, Alvarez, Kolchinski, Zhao, Gordon, Bernstein, Ermon, Mehrjou, Schölkopf, Co-Reyes, Janner, Wu, Tenenbaum, Levine, Mohsenzadeh, Zhou</i>
02:00 PM	Poster Session	
03:00 PM	Q&A from the Audience. Ask the Grad Students	<i>Grant, Battleday, Sanborn, Chang</i>
03:30 PM	Talia Konkle (Harvard)	<i>Konkle</i>
03:55 PM	Nikolaus Kriegeskorte (Columbia)	<i>Kriegeskorte</i>
04:20 PM	Matthias Bethge (Universität Tübingen)	<i>Bethge</i>
04:45 PM	Eero Simoncelli (NYU)	<i>Simoncelli</i>

	Panel Discussion: What sorts of cognitive or biological (architectural) inductive biases will be crucial for developing effective artificial intelligence?	<i>Higgins, Konkle, Bethge</i>
05:10 PM		
06:00 PM	Concluding Remarks & Prizes Ceremony	<i>Deza, Peterson, Murty, Griffiths</i>
06:10 PM	Evening Reception	

Abstracts (5):

Abstract 5: **Q&A from the Audience: Ask a Neuro / Cognitive Scientist in Shared Visual Representations in Human and Machine Intelligence**, *Griffiths, DiCarlo, Konkle* 10:15 AM

Panelists: Talia Konkle, Thomas Griffiths, James DiCarlo.

Abstract 15: **Q&A from the Audience. Ask the Grad Students in Shared Visual Representations in Human and Machine Intelligence**, *Grant, Battleday, Sanborn, Chang* 03:00 PM

"Cross-disciplinary research experiences and tips for Graduate School Admissions Panelists"

Panelists:
Erin Grant (UC Berkeley)
Nadine Chang (CMU)
Ruairidh Battleday (Princeton)
Sophia Sanborn (UC Berkeley)

Abstract 20: **Panel Discussion: What sorts of cognitive or biological (architectural) inductive**

biases will be crucial for developing effective artificial intelligence? in Shared Visual Representations in Human and Machine Intelligence, *Higgins, Konkle, Bethge* 05:10 PM

Panelists: Irina Higgins (DeepMind), Talia Konkle (Harvard), Nikolaus Kriegeskorte (Columbia), Matthias Bethge (Universität Tübingen)

Abstract 21: Concluding Remarks & Prizes Ceremony in Shared Visual Representations in Human and Machine Intelligence, *Deza, Peterson, Murty, Griffiths* 06:00 PM

Best Paper Award Prize (NVIDIA Titan RTX) and Best Poster Award Prize (Oculus Quest)

Abstract 22: Evening Reception in Shared Visual Representations in Human and Machine Intelligence, 06:10 PM

Sponsored by MIT Quest for Intelligence

Workshop on Human-Centric Machine Learning

Plamen P Angelov, Nuria Oliver, Adrian Weller, Manuel Rodriguez, Isabel Valera, Silvia Chiappa, Hoda Heidari, Niki Kilbertus

West 223 + 224, Fri Dec 13, 08:00 AM

The growing field of Human-centric ML seeks to minimize the potential harms, risks, and burdens of big data technologies on the public, and at the same time, maximize their societal benefits. In this workshop, we address a wide range of challenges from diverse, multi-disciplinary viewpoints. We bring together experts from a diverse set of backgrounds. Our speakers are leading experts in ML, human-computer interaction, ethics, and law. Each of our speakers will focus on one core human-centred challenge (namely, fairness, accountability, interpretability, transparency,

security, and privacy) in specific application domains (such as medicine, welfare programs, governance, and regulation). One of the main goals of this workshop is to help the community understand where it stands after a few years of rapid technical development and identify promising research directions to pursue in the years to come. Our speakers identify in their presentations 3-5 research directions that they consider to be of crucial importance. These directions are further debated in one of our panel discussions.

Schedule

08:20 AM	Welcome and introduction	
08:30 AM	Invited talks #1,2	<i>Gummadi</i>
09:20 AM	Contributed talks	
10:05 AM	Panel #1: On the role of industry, academia, and government in developing HCML	
10:30 AM	Coffe break	
11:00 AM	Invited talk #3	<i>Mulligan</i>
11:25 AM	Spotlight talks	
12:00 PM	Lunch and poster session	
01:30 PM	Invited talks #4,5	<i>Roth, Chouldechova</i>
02:20 PM	Contributed talks	
03:05 PM	Coffee break	

03:30 PM	Invited talks #6,7	<i>Doshi-Velez, Kim</i>
04:20 PM	Contributed talks	
04:50 PM	Panel #2: Future research directions and interdisciplinary collaborations in HCML	
05:15 PM	Poster session	<i>Gu, Xiang, Kasirzadeh, Han, Florez, Harder, Nguyen, Akhavan Rahnama, Donini, Slack, Ali, Koley, Bakker, Hilgard, James-Sorenson, Ramos, Lu, Yang, Boyarskaya, Pawelczyk, Sokol, Jaiswal, Bhatt, Alvarez Melis, Grover, Marx, Yang, Liang, Wang, Çapan, Wang, Grünewälder, Khajehnejad, Patro, Kunes, Deng, Liu, Oneto, Li, Weber, Matthes, Tu</i>
06:15 PM	Closing remarks	

Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications

Reinhard Heckel, Paul Hand, Richard Baraniuk, Joan Bruna, Alex Dimakis, Deanna Needell

West 301 - 305, Fri Dec 13, 08:00 AM

There is a long history of algorithmic development for solving inverse problems arising in sensing and imaging systems and beyond. Examples include medical and computational imaging, compressive sensing, as well as community detection in networks. Until recently, most algorithms for solving inverse problems in the imaging and network sciences were based on static signal models derived from physics or intuition, such as wavelets or sparse representations.

Today, the best performing approaches for the aforementioned image reconstruction and sensing problems are based on deep learning, which learn various elements of the method including i) signal representations, ii) stepsizes and parameters of iterative algorithms, iii) regularizers, and iv) entire inverse functions. For example, it has recently been shown that solving a variety of inverse problems by transforming an iterative, physics-based algorithm into a deep network whose parameters can be learned from training data, offers faster convergence and/or a better quality solution. Moreover, even with very little or no learning, deep neural networks enable superior performance for classical linear inverse problems such as denoising and compressive sensing. Motivated by those success stories, researchers are redesigning traditional imaging and sensing systems.

However, the field is mostly wide open with a range of theoretical and practical questions unanswered. In particular, deep-neural network based approaches often lack the guarantees of the traditional physics based methods, and while typically superior can make drastic reconstruction errors, such as fantasizing a tumor in an MRI reconstruction.

This workshop aims at bringing together theoreticians and practitioners in order to chart out

recent advances and discuss new directions in deep neural network based approaches for solving inverse problems in the imaging and network sciences.

Schedule

08:30 AM	Opening Remarks	<i>Heckel, Hand, Dimakis, Bruna, Needell, Baraniuk</i>
----------	------------------------	--

08:40 AM	The spiked matrix model with generative priors	<i>Zdeborová</i>
----------	---	------------------

09:10 AM	Robust One-Bit Recovery via ReLU Generative Networks: Improved Statistical Rate and Global Landscape Analysis	<i>Qiu, Wei, Yang</i>
----------	--	-----------------------

09:40 AM	Coffee Break	
----------	---------------------	--

10:30 AM	Computational microscopy in scattering media	<i>Waller</i>
----------	---	---------------

11:00 AM	Basis Decomposition of Deep Learning	<i>Sapiro</i>
----------	---	---------------

11:30 AM	Neural Reparameterization Improves Structural Optimization	<i>Hoyer, Sohl-Dickstein, Greydanus</i>
----------	---	---

12:00 PM	Lunch Break	
----------	--------------------	--

02:00 PM	Learning-Based Low-Rank Approximations	<i>Indyk</i>
----------	---	--------------

02:30 PM	Blind Denoising, Self-Supervision, and Implicit Inverse Problems	<i>Batson</i>
----------	---	---------------

03:00 PM	Learning Regularizers from Data	<i>Chandrasekaran</i>
----------	--	-----------------------

04:15
PM**Poster Session**

Scarlett, Indyk,
 Vakilian, Weller,
 Mitra, Aubin,
 Loureiro, Krzakala,
 Zdeborová,
 Monakhova,
 Yurtsever, Waller,
 Sommerhoff, Moeller,
 Anirudh, Qiu, Wei,
 Yang, J. Thiagarajan,
 Asif, Gillhofer,
 Brandstetter,
 Hochreiter, Petersen,
 Patel, Oberai,
 Kamath, Karmalkar,
 Price, Ahmed,
 Kadkhodaie, Mohan,
 Simoncelli,
 Fernandez-Granda,
 Leong, Sakla, Willett,
 Hoyer,
 Sohl-Dickstein,
 Greydanus, Jagatap,
 Hegde, Kellman,
 Tamir, Laanait, Dia,
 Ravanelli, Binas,
 Rostamzadeh, Jalali,
 Fang, Schwing,
 Lachapelle,
 Brouillard, Deleu,
 Lacoste-Julien, Yu,
 Mazumdar, Rawat,
 Zhao, Chen, Li,
 Ramsauer, Rizzuti,
 Mitsakos, Cao,
 Strohmer, Li, Peng,
 Ongie

with deep networks: New architectures, theoretical foundations, and applications,
 Zdeborová 08:40 AM

Using a low-dimensional parametrization of signals is a generic and powerful way to enhance performance in signal processing and statistical inference. A very popular and widely explored type of dimensionality reduction is sparsity; another type is generative modelling of signal distributions. Generative models based on neural networks, such as GANs or variational auto-encoders, are particularly performant and are gaining on applicability. In this paper we study spiked matrix models, where a low-rank matrix is observed through a noisy channel. This problem with sparse structure of the spikes has attracted broad attention in the past literature. Here, we replace the sparsity assumption by generative modelling, and investigate the consequences on statistical and algorithmic properties. We analyze the Bayes-optimal performance under specific generative models for the spike. In contrast with the sparsity assumption, we do not observe regions of parameters where statistical performance is superior to the best known algorithmic performance. We show that in the analyzed cases the approximate message passing algorithm is able to reach optimal performance. We also design enhanced spectral algorithms and analyze their performance and thresholds using random matrix theory, showing their superiority to the classical principal component analysis. We complement our theoretical results by illustrating the performance of the spectral algorithms when the spikes come from real datasets.

Abstract 3: Robust One-Bit Recovery via ReLU Generative Networks: Improved Statistical Rate and Global Landscape Analysis in Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications, Qiu, Wei, Yang 09:10 AM

Abstracts (7):

Abstract 2: The spiked matrix model with generative priors in Solving inverse problems

We study the robust one-bit compressed sensing problem whose goal is to design an algorithm that faithfully recovers any sparse target vector $\theta \in \mathbb{R}^d$ uniformly from m quantized noisy measurements. Under the assumption that the measurements are sub-Gaussian, to recover any k -sparse θ ($\|\theta\|_0 \leq k$) uniformly up to an error ϵ with high probability, the best known computationally tractable algorithm requires $m \geq \tilde{O}(k \log d / \epsilon^4)$. In this paper, we consider a new framework for the one-bit sensing problem where the sparsity is implicitly enforced via mapping a low dimensional representation x_0 through a known n -layer ReLU generative network $G: \mathbb{R}^k \rightarrow \mathbb{R}^d$. Such a framework poses low-dimensional priors on θ_0 without a known basis. We propose to recover the target $G(x_0)$ via an unconstrained empirical risk minimization (ERM) problem under a much weaker ϵ -sub-exponential measurement assumption. For such a problem, we establish a joint statistical and computational analysis. In particular, we prove that the ERM estimator in this new framework achieves an improved statistical rate of $m = \tilde{O}(k n \log d / \epsilon^2)$ recovering any $G(x_0)$ uniformly up to an error ϵ . Moreover, from the lens of computation, despite non-convexity, we prove that the objective of our ERM problem has no spurious stationary point, that is, any stationary point is equally good for recovering the true target up to scaling with a certain accuracy. Our analysis sheds some light on the possibility of inverting a deep generative model under partial and quantized measurements, complementing the recent success of using deep generative models for inverse

problems.

Abstract 5: Computational microscopy in scattering media in Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications, Waller
10:30 AM

Computational imaging involves the joint design of imaging system hardware and software, optimizing across the entire pipeline from acquisition to reconstruction. Computers can replace bulky and expensive optics by solving computational inverse problems. This talk will describe new microscopes that use computational imaging to enable 3D fluorescence and phase measurement using image reconstruction algorithms that are based on large-scale nonlinear non-convex optimization combined with unrolled neural networks. We further discuss engineering of data capture for computational microscopes by end-to-end learned design.

Abstract 6: Basis Decomposition of Deep Learning in Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications, Sapiro 11:00 AM

Ordinary convolutional neural networks (CNNs) learn non-parametric filters, applied in multiple layers, leading to need to learn tens of millions of variables with large training data. In this talk we show how such filters can be replaced by basis, not only reducing the number of parameters and needed training samples by orders of magnitudes but also intrinsically and naturally achieving invariance, domain adaptation, and stochasticity. We present the basic plug-and-play framework; its natural incorporation into virtually any existing CNN; theoretical results; and applications in numerous areas, including invariant classification, domain shift, domain-invariant learning, diverse generative networks, and stochastic networks. This is joint work with Ze Wang, Qiang Qiu, and

Xiuyuan Cheng.

Abstract 7: Neural Reparameterization Improves Structural Optimization in Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications, *Hoyer, Sohl-Dickstein, Greydanus*
11:30 AM

Structural optimization is a popular method for designing objects such as bridge trusses, airplane wings, and optical devices. Unfortunately, the quality of solutions depends heavily on how the problem is parameterized. In this paper, we propose using the implicit bias over functions induced by neural networks to improve the parameterization of structural optimization. Rather than directly optimizing densities on a grid, we instead optimize the parameters of a neural network which outputs those densities. This reparameterization leads to different and often better solutions. On a selection of 116 structural optimization tasks, our approach produces an optimal design 50% more often than the best baseline method.

Abstract 10: Blind Denoising, Self-Supervision, and Implicit Inverse Problems in Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications, *Batson* 02:30 PM

We will discuss a self-supervised approach to the foundational inverse problem of denoising (Noise2Self). By taking advantage of statistical independence in the noise, we can estimate the mean-square error for a large class of deep architectures without access to ground truth. This allows us to train a neural network to denoise from noisy data alone, and also to compare between architectures, selecting one which will produce images with the lowest MSE. However, architectures with the same MSE performance can produce qualitatively different results, i.e., the hypersurface of images with fixed MSE is very

heterogeneous. We will discuss ongoing work in understanding the types of artifacts which different denoising architectures give rise to.

Abstract 11: Learning Regularizers from Data in Solving inverse problems with deep networks: New architectures, theoretical foundations, and applications, *Chandrasekaran* 03:00 PM

Regularization techniques are widely employed in the solution of inverse problems in data analysis and scientific computing due to their effectiveness in addressing difficulties due to ill-posedness. In their most common manifestation, these methods take the form of penalty functions added to the objective in variational approaches for solving inverse problems. The purpose of the penalty function is to induce a desired structure in the solution, and these functions are specified based on prior domain-specific expertise. We consider the problem of learning suitable regularization functions from data in settings in which precise domain knowledge is not directly available; the objective is to identify a regularizer to promote the type of structure contained in the data. The regularizers obtained using our framework are specified as convex functions that can be computed efficiently via semidefinite programming. Our approach for learning such semidefinite regularizers combines recent techniques for rank minimization problems along with the Operator Sinkhorn procedure.
(Joint work with Yong Sheng Soh)

EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition)

Raj Purihar, Michael Goldfarb, Satyam Srivastava, TAO SHENG

West 306, Fri Dec 13, 08:00 AM

A new wave of intelligent computing, driven by recent advances in machine learning and cognitive algorithms coupled with process technology and new design methodologies, has the potential to usher unprecedented disruption in the way modern computing systems are designed and deployed. These new and innovative approaches often provide an attractive and efficient alternative not only in terms of performance but also power, energy, and area. This disruption is easily visible across the whole spectrum of computing systems -- ranging from low end mobile devices to large scale data centers and servers including intelligent infrastructures.

A key class of these intelligent solutions is providing real-time, on-device cognition at the edge to enable many novel applications including computer vision and image processing, language understanding, speech and gesture recognition, malware detection and autonomous driving. Naturally, these applications have diverse requirements for performance, energy, reliability, accuracy, and security that demand a holistic approach to designing the hardware, software, and intelligence algorithms to achieve the best power, performance, and area (PPA).

Topics:

- Architectures for the edge: IoT, automotive, and mobile
- Approximation, quantization reduced precision computing
- Hardware/software techniques for sparsity
- Neural network architectures for resource constrained devices

- Neural network pruning, tuning and automatic architecture search
- Novel memory architectures for machine learning
- Communication/computation scheduling for better performance and energy
- Load balancing and efficient task distribution techniques
- Exploring the interplay between precision, performance, power and energy
- Exploration of new and efficient applications for machine learning
- Characterization of machine learning benchmarks and workloads
- Performance profiling and synthesis of workloads
- Simulation and emulation techniques, frameworks and platforms for machine learning
- Power, performance and area (PPA) based comparison of neural networks
- Verification, validation and determinism in neural networks
- Efficient on-device learning techniques
- Security, safety and privacy challenges and building secure AI systems

Schedule

08:00 AM	TBD	<i>LeCun</i>
08:45 AM	Efficient Computing for AI and Robotics	<i>Sze</i>
09:30 AM	Abandoning the Dark Arts: New Directions in Efficient DNN Design	<i>Keutzer</i>
10:00 AM	Poster Session 1	<i>Spasov, Nayak, Diego Andilla, Zhang, Trivedi</i>

10:30 AM	Putting the “Machine” Back in Machine Learning: The Case for Hardware-ML Model Co-design	<i>Marculescu</i>
11:00 AM	Adaptive Multi-Task Neural Networks for Efficient Inference	<i>Feris</i>
11:30 AM	Oral Session 1	<i>Yu, Hartmann, Li, Shafiee, Yang, Zafrir</i>
12:00 PM	Qualcomm Industry Talk	<i>Lee</i>
12:30 PM	Lunch	
02:00 PM	Cheap, Fast, and Low Power Deep Learning: I need it now!	<i>Delp</i>
02:45 PM	Advances and Prospects for In-memory Computing	<i>Verma</i>
03:15 PM	Algorithm-Accelerator Co-Design for Neural Network Specialization	<i>Zhang</i>
03:45 PM	Poster Session 2	<i>Prato, Thakker, Galindez Olascoaga, Zhang, Partovi Nia, Adamczewski</i>
04:15 PM	Efficient Algorithms to Accelerate Deep Learning on Edge Devices	<i>Han</i>
04:45 PM	Oral Session 2	<i>Liao, McKinstry, Izsak, Li, Huang, Mordido</i>

05:30 PM **Microsoft Industry Talk** *Darvish Rouhani*

06:00 PM **EMCC Results**

Abstracts (9):

Abstract 1: **TBD in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition)**, *LeCun* 08:00 AM

TBD

Abstract 2: **Efficient Computing for AI and Robotics in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition)**, *Sze* 08:45 AM

Computing near the sensor is preferred over the cloud due to privacy and/or latency concerns for a wide range of applications including robotics/drones, self-driving cars, smart Internet of Things, and portable/wearable electronics. However, at the sensor there are often stringent constraints on energy consumption and cost in addition to the throughput and accuracy requirements of the application. In this talk, we will describe how joint algorithm and hardware design can be used to reduce energy consumption while delivering real-time and robust performance for applications including deep learning, computer vision, autonomous navigation/exploration and video/image processing. We will show how energy-efficient techniques that exploit correlation and sparsity to reduce compute, data movement and storage costs can be applied to various tasks including image classification, depth estimation, super-resolution, localization and mapping.

Abstract 3: **Abandoning the Dark Arts: New Directions in Efficient DNN Design in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition)**, *Keutzer* 09:30

AM

Deep Neural Net models have provided the most accurate solutions to a very wide variety of problems in vision, language, and speech; however, the design, training, and optimization of efficient DNNs typically requires resorting to the “dark arts” of ad hoc methods and extensive hyperparameter tuning. In this talk we present our progress on abandoning these dark arts by using Differential Neural Architecture Search to guide the design of efficient DNNs and by using Hessian-based methods to guide the processes of training and quantizing those DNNs.

Abstract 5: Putting the “Machine” Back in Machine Learning: The Case for Hardware-ML Model Co-design in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition), *Marculescu* 10:30 AM

Machine learning (ML) applications have entered and impacted our lives unlike any other technology advance from the recent past. Indeed, almost every aspect of how we live or interact with others relies on or uses ML for applications ranging from image classification and object detection, to processing multi-modal and heterogeneous datasets. While the holy grail for judging the quality of a ML model has largely been serving accuracy, and only recently its resource usage, neither of these metrics translate directly to energy efficiency, runtime, or mobile device battery lifetime. This talk will uncover the need for building accurate, platform-specific power and latency models for convolutional neural networks (CNNs) and efficient hardware-aware CNN design methodologies, thus allowing machine learners and hardware designers to identify not just the best accuracy NN configuration, but also those that satisfy given hardware constraints. Our proposed modeling framework is applicable to both high-end and mobile platforms and achieves 88.24% accuracy for latency, 88.34% for power, and 97.21% for energy prediction. Using similar

predictive models, we demonstrate a novel differentiable neural architecture search (NAS) framework, dubbed Single-Path NAS, that uses one single-path over-parameterized CNN to encode all architectural decisions based on shared convolutional kernel parameters. Single-Path NAS achieves state-of-the-art top-1 ImageNet accuracy (75.62%), outperforming existing mobile NAS methods for similar latency constraints (~80ms) and finds the final configuration up to 5,000x faster compared to prior work. Combined with our quantized CNNs (Flexible Lightweight CNNs or FLightNNs) that customize precision level in a layer-wise fashion and achieve almost iso-accuracy at 5-10x energy reduction, such a modeling, analysis, and optimization framework is poised to lead to true co-design of hardware and ML model, orders of magnitude faster than state of the art, while satisfying both accuracy and latency or energy constraints.

Abstract 6: Adaptive Multi-Task Neural Networks for Efficient Inference in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition), *Feris* 11:00 AM

Very deep convolutional neural networks have shown remarkable success in many computer vision tasks, yet their computational expense limits their impact in domains where fast inference is essential. While there has been significant progress on model compression and acceleration, most methods rely on a one-size-fits-all network, where the same set of features is extracted for all images or tasks, no matter their complexity. In this talk, I will first describe an approach called BlockDrop, which learns to dynamically choose which layers of a deep network to execute during inference, depending on the image complexity, so as to best reduce total computation without degrading prediction accuracy. Then, I will show how this approach can be extended to design compact multi-task networks, where a different set of layers is executed depending on the task complexity, and the level of

feature sharing across tasks is automatically determined to maximize both the accuracy and efficiency of the model. Finally, I will conclude the talk presenting an efficient multi-scale neural network model, which achieves state-of-the-art results in terms of accuracy and FLOPS reduction on standard benchmarks such as the ImageNet dataset.

Abstract 10: Cheap, Fast, and Low Power Deep Learning: I need it now! in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition), *Delp* 02:00 PM

In this talk I will describe the need for low power machine learning systems. I will motivate this by describing several current projects at Purdue University that have a need for energy efficient deep learning and in some cases the real deployment of these methods will not be possible without lower power solutions. The applications include precision farming, health care monitoring, and edge-based surveillance.

Abstract 11: Advances and Prospects for In-memory Computing in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition), *Verma* 02:45 PM

Edge AI applications retain the need for high-performing inference models, while driving platforms beyond their limits of energy efficiency and throughput. Digital hardware acceleration, enabling 10-100x gains over general-purpose architectures, is already widely deployed, but is ultimately restricted by data-movement and memory accessing that dominates deep-learning computations. In-memory computing, based on both SRAM and emerging memory, offers fundamentally new tradeoffs for overcoming these barriers, with the potential for 10x higher energy efficiency and area-normalized throughput demonstrated in recent designs. But, those tradeoffs instate new challenges, especially affecting scaling to the level

of computations required, integration in practical heterogeneous architectures, and mapping of diverse software. This talk examines those tradeoffs to characterize the challenges. It then explores recent research that provides promising paths forward, making in-memory computing more of a practical reality than ever before.

Abstract 12: Algorithm-Accelerator Co-Design for Neural Network Specialization in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition), *Zhang* 03:15 PM

In recent years, machine learning (ML) with deep neural networks (DNNs) has been widely deployed in diverse application domains. However, the growing complexity of DNN models, the slowdown of technology scaling, and the proliferation of edge devices are driving a demand for higher DNN performance and energy efficiency. ML applications have shifted from general-purpose processors to dedicated hardware accelerators in both academic and commercial settings. In line with this trend, there has been an active body of research on both algorithms and hardware architectures for neural network specialization.

This talk presents our recent investigation into DNN optimization and low-precision quantization, using a co-design approach featuring contributions to both algorithms and hardware accelerators. First, we review static network pruning techniques and show a fundamental link between group convolutions and circulant matrices – two previously disparate lines of research in DNN compression. Then we discuss channel gating, a dynamic, fine-grained, and trainable technique for DNN acceleration. Unlike static approaches, channel gating exploits input-dependent dynamic sparsity at run time. This results in a significant reduction in compute cost with a minimal impact on accuracy. Finally, we present outlier channel splitting, a technique to improve DNN weight quantization by removing outliers from the weight distribution without

retraining.

Abstract 14: Efficient Algorithms to Accelerate Deep Learning on Edge Devices in EMC2: Energy Efficient Machine Learning and Cognitive Computing (5th edition), *Han* 04:15 PM

Efficient deep learning computing requires algorithm and hardware co-design to enable specialization. However, the extra degree of freedom creates a much larger design space. We propose AutoML techniques to architect efficient neural networks. We investigate automatically designing small and fast models (ProxylessNAS), auto channel pruning (AMC), and auto mixed-precision quantization (HAQ). We demonstrate such learning-based, automated design achieves superior performance and efficiency than rule-based human design. Moreover, we shorten the design cycle by 200x than previous work to efficiently search efficient models, so that we can afford to design specialized neural network models for different hardware platforms. We accelerate computation-intensive AI applications including (TSM) for efficient video recognition and PVCNN for efficient 3D recognition on point clouds. Finally, we'll describe scalable distributed training and the potential security issues of efficient deep learning.

Machine Learning for Health (ML4H): What makes machine learning in medicine different?

Andrew Beam, Tristan Naumann, Brett Beaulieu-Jones, Irene Y Chen, Sam Finlayson, Emily Alsentzer, Adrian Dalca, Matthew McDermott

West Ballroom A, Fri Dec 13, 08:00 AM

The goal of the NeurIPS 2019 Machine Learning for Health Workshop (ML4H) is to foster collaborations that meaningfully impact medicine by bringing

together clinicians, health data experts, and machine learning researchers. Attendees at this workshop can also expect to broaden their network of collaborators to include clinicians and machine learning researchers who are focused on solving some of the most important problems in medicine and healthcare. The organizers of this proposal have successfully run NeurIPS workshops in the past and are well-equipped to run this year's workshop should this proposal be accepted.

This year's theme of "What makes machine learning in medicine different?" aims to elucidate the obstacles that make the development of machine learning models for healthcare uniquely challenging. To speak to this theme, we have received commitments to speak from some of the leading researchers and physicians in this area. Below is a list of confirmed speakers who have agreed to participate.

Luke Oakden-Raynor, MBBS (Adelaide)
 Russ Altman, MD/PhD (Stanford)
 Lilly Peng, MD/PhD (Google)
 Daphne Koller, PhD (in situ)
 Jeff Dean, PhD (Google)

Attendees at the workshop will gain an appreciation for problems that are unique to the application of machine learning for healthcare and a better understanding of how machine learning techniques may be leveraged to solve important clinical problems. This year's workshop builds on the last two NeurIPS ML4H workshops, which were both attended by more than 500 people each year, and helped form the foundations of an emerging research community.

Please see the attached document for the full program.

Schedule

08:45 AM	Daphne Koller Talk	
09:15 AM	Emily Fox Talk	<i>Fox</i>
10:15 AM	Luke Oakden-Rayner Talk	<i>Oakden-Rayner</i>
10:45 AM	Paper spotlight talks	

11:15
AM

Poster Session I

Zheng, Kapur, Asif, Rozenberg, Gilet, Sidorov, Kumar, Van Steenkiste, Boag, Ouyang, Jaeger, Liu, Balagopalan, Rajan, Skreta, Pattisapu, Goschenhofer, Prabhu, Jin, Gardiner, Li, kumar, Hu, Motani, Lovelace, Roshan, Wang, Valmianski, Lee, Mallya, Chaibub Neto, Kemp, Charpignon, Nigam, Weng, Boughorbel, Bellot, Gondara, Zhang, Bahadori, Zech, Shao, Choi, Seyyed-Kalantari, Aiken, Bica, Shen, Chin-Cheong, Roy, Baldini, Min, Deschrijver, Marttinen, Pascual Ortiz, Nagesh, Rindtorff, Mulyar, Hoebel, Shaka, Machart, Gatys, Ng, Hüser, Taylor, Barbour, Martinez, McCreery, Eyre, Natarajan, Yi, Ma, Nagpal, Du, Gao, Tuladhar, Shleifer, Ren, Mashouri, Lu, Bagherzadeh-Khiabani, Choudhury, Raghu, Fleming, Jain, YANG, Harley, Pfohl, Rumetshofer, Fedorov, Dash, Pfau, Tomkins, Targonski, Brudno, Li, Yu

02:45 PM	Lily Peng talk	<i>Peng</i>
03:15 PM	Anna Goldenberg Talk	<i>Goldenberg</i>
03:45 PM	Poster Session II	
04:45 PM	Deepmind Talk	<i>Tomasev</i>
05:15 PM	Panel Discussion	
06:15 PM	Message from sponsor	<i>Baranov, Ji</i>

Meta-Learning

**Roberto Calandra, Ignasi Clavera Gilaberte,
Frank Hutter, Joaquin Vanschoren, Jane Wang**

West Ballroom B, Fri Dec 13, 08:00 AM

Recent years have seen rapid progress in metalearning methods, which learn (and optimize) the performance of learning methods based on data, generate new learning methods from scratch, and learn to transfer knowledge across tasks and domains. Metalearning can be seen as the logical conclusion of the arc that machine learning has undergone in the last decade, from learning classifiers, to learning representations, and finally to learning algorithms that themselves acquire representations and classifiers. The ability to improve one's own learning capabilities through experience can also be viewed as a hallmark of intelligent beings, and there are strong connections with work on human learning in neuroscience. The goal of this workshop is to bring together researchers from all the different communities and topics that fall under the umbrella of metalearning. We expect that the presence of these different

communities will result in a fruitful exchange of ideas and stimulate an open discussion about the current challenges in metalearning, as well as possible solutions.

Schedule

09:10 AM	Invited Talk 1	<i>Abbeel</i>
09:40 AM	Invited Talk 2	<i>Clune</i>
10:10 AM	Poster Spotlights 1	<i>Takagi, Javed, Sommer, Sharaf, D'Oro, Wei, Doveh, White, Gonzalez, Nguyen, li, Yu, Ramalho, Nomura, Alvi, Ton, Huang, Lee, Flennerhag, Zhang, Friesen, Blomstedt, Dubatovka, Bartunov, Yi, Shcherbatyi, Simon, Shang, MacLeod, Liu, Fowl, Parente Paiva Mesquita, Quillen</i>
10:30 AM	Coffee/Poster session 1	
11:30 AM	Invited Talk 3	<i>Grant</i>
12:00 PM	Discussion 1	
02:00 PM	Invited Talk 4	<i>Abel</i>
02:30 PM	Invited Talk 5	<i>Hadsell</i>
03:00 PM	Poster Spotlights 2	

03:20 PM	Coffee/Poster session 2	<i>Song, Mangla, Salinas, Zhuang, Feng, Hu, Puri, Maddox, Raghu, Tossou, Yin, Dasgupta, Lee, Alet, Xu, Franke, Harrison, Warrell, Dhillon, Zela, Qiu, Siems, Mendonca, Schlessinger, Li, Manolache, Dutta, Glass, Singh, Koehler</i>
<hr/>		
04:30 PM	Contributed Talk 1	
<hr/>		
04:45 PM	Contributed Talk 2	
<hr/>		
05:00 PM	Invited Talk 6	<i>Lake</i>
<hr/>		
05:30 PM	Discussion 2	

Biological and Artificial Reinforcement Learning

Raymond Chua, Sara Zannone, Feryal Behbahani, Rui Ponte Costa, Claudia Clopath, Blake Richards, Doina Precup

West Ballroom C, Fri Dec 13, 08:00 AM

Reinforcement learning (RL) algorithms learn through rewards and a process of trial-and-error. This approach was strongly inspired by the study of animal behaviour and has led to outstanding achievements in machine learning (e.g. in games, robotics, science). However, artificial agents still struggle with a number of difficulties, such as sample efficiency, learning in dynamic environments and over multiple timescales, generalizing and

transferring knowledge. On the other end, biological agents excel at these tasks. The brain has evolved to adapt and learn in dynamic environments, while integrating information and learning on different timescales and for different duration. Animals and humans are able to extract information from the environment in efficient ways by directing their attention and actively choosing what to focus on. They can achieve complicated tasks by solving sub-problems and combining knowledge as well as representing the environment in efficient ways and plan their decisions off-line. Neuroscience and cognitive science research has largely focused on elucidating the workings of these mechanisms. Learning more about the neural and cognitive underpinnings of these functions could be key to developing more intelligent and autonomous agents. Similarly, having a computational and theoretical framework, together with a normative perspective to refer to, could and does contribute to elucidate the mechanisms used by animals and humans to perform these tasks. Building on the connection between biological and artificial reinforcement learning, our workshop will bring together leading and emergent researchers from Neuroscience, Psychology and Machine Learning to share: (i) how neural and cognitive mechanisms can provide insights to tackle challenges in RL research and (ii) how machine learning advances can help further our understanding of the brain and behaviour.

Schedule

09:00 AM	Opening Remarks	<i>Chua, Behbahani, Zannone, Ponte Costa, Clopath, Precup, Richards</i>
<hr/>		
09:15 AM	Invited Talk #1: Jane Wang	<i>Wang</i>

		<i>Mohinta, Agostinelli, Moringen, Lee, Lo, Maass, Sheffer, Bredenberg, Eysenbach, Xia, Markou, Lichtenberg, Richemond, Zhang, Lanier, Lin, Fedus, Berseth, Sarrico, Crosby, McAleer, Ghiassian, Scherr, Bellec, Salaj, Kolbeinsson, Rosenberg, Shin, Lee, Cecchi, Rish, Hajek</i>
09:45 AM	Coffee Break & Poster Session	
10:30 AM	Contributed Talk #1: Humans flexibly transfer options at multiple levels of abstractions	<i>Xia</i>
10:45 AM	Contributed Talk #2: Slow processes of neurons enable a biologically plausible approximation to policy gradient	<i>Maass</i>
11:00 AM	Invited Talk #2: Jacqueline Gottlieb	<i>Gottlieb</i>
11:30 AM	Invited Talk #3 Emma Brunskill	<i>Brunskill</i>
12:00 PM	Lunch Break & Poster Session	
01:30 PM	Invited Talk #5: Ida Momennejad	
02:00 PM	Invited Talk #4: Igor Mordatch	<i>Mordatch</i>
02:30 PM	Invited Talk: #6 Jeff Clune	<i>Clune</i>

03:00 PM	Invited Talk #7: Angela Yu	<i>Yu</i>
03:30 PM	Coffee Break & Poster Session	
04:15 PM	Contributed Talk #3 MEMENTO: Further Progress Through Forgetting	<i>Fedus</i>
04:30 PM	Invited Talk #8: Richard Sutton	<i>Sutton</i>
05:00 PM	Panel Discussion led by Grace Lindsay (part 1)	<i>Lindsay, Richards, Precup</i>
05:30 PM	Panel Discussion led by Grace Lindsay (part 2)	<i>Lindsay, Richards, Precup</i>



Graph Representation Learning

Will Hamilton, Rianne van den Berg, Michael Bronstein, Stefanie Jegelka, Thomas Kipf, Jure Leskovec, Renjie Liao, Yizhou Sun, Petar Veličković

West Exhibition Hall A, Fri Dec 13, 08:00 AM

Graph-structured data is ubiquitous throughout the natural and social sciences, from telecommunication networks to quantum chemistry. Building relational inductive biases into deep learning architectures is crucial if we want systems that can learn, reason, and generalize from this kind of data. Furthermore, graphs can be seen as a natural generalization of simpler kinds of structured data (such as images), and therefore, they represent a natural avenue for the next breakthroughs in machine learning.

Recent years have seen a surge in research on

graph representation learning, including techniques for deep graph embeddings, generalizations of convolutional neural networks to graph-structured data, and neural message-passing approaches inspired by belief propagation. These advances in graph neural networks and related techniques have led to new state-of-the-art results in numerous domains, including chemical synthesis, 3D-vision, recommender systems, question answering, and social network analysis.

The workshop will consist of contributed talks, contributed posters, and invited talks on a wide variety of methods and problems related to graph representation learning. We will welcome 4-page original research papers on work that has not previously been published in a machine learning conference or workshop. In addition to traditional research paper submissions, we will also welcome 1-page submissions describing open problems and challenges in the domain of graph representation learning. These open problems will be presented as short talks (5-10 minutes) immediately preceding a coffee break to facilitate and spark discussions.

The primary goal for this workshop is to facilitate community building; with hundreds of new researchers beginning projects in this area, we hope to bring them together to consolidate this fast-growing area of graph representation learning into a healthy and vibrant subfield.

Schedule

08:45 AM	Opening remarks	<i>Hamilton</i>
09:00 AM	Invited talk: Marco Gori	<i>Gori</i>
09:30 AM	Invited talk: Marinka Zitnik	<i>Zitnik</i>

10:00 AM	Open Challenges - Spotlight Presentations	<i>Sumba Toral, Maron, Kolbeinsson</i>
10:30 AM	Coffee Break	
11:00 AM	Invited talk: Andrew McCallum	<i>McCallum</i>
11:30 AM	Poster Session #1	<i>Sanborn, Yao, Cai, Duan, Shao, Belli, Boyarski, Ye, Sarkar, KHADEMI, Wang, Wang, Wu, Joshi, Bose, Ma, Brockschmidt, Zambon, Graber, Glorot, Cameron, Huang, Li, Sun, SALHA, Vlastelica Poganj, Huang</i>
12:30 PM	Lunch	
01:30 PM	Outstanding Contribution Talk #1	
01:45 PM	Outstanding Contribution Talk #2	
02:00 PM	Outstanding Contribution Talk #3	
02:15 PM	Invited talk: Tommi Jaakkola	<i>Jaakkola</i>
02:45 PM	Discussion Panel: Graph Neural Networks and Combinatorial Optimization	

03:15 PM	Poster Session #2	<i>Frasca, Velivkovi, Xu, Chen, Cheng, Chami, Kim, Gomes, Maziarka, Hoffmann, Levie, Gogoglou, Gong, Wang, Leng, Hu, TIAO, Flam-Shepherd, Zhang, KHADEMI, Su, vargas vieyra, Malik, Verma, Bose, Stanic, Liu, Ghalebi, D'Oro</i>
04:15 PM	Invited talk: Bistra Dilkina	<i>Dilkina</i>
04:45 PM	Invited talk: Peter Battaglia	<i>Battaglia</i>
05:15 PM	Closing Remarks	<i>Hamilton, Liao, Sun, Velivkovi, Leskovec, Jegelka, Bronstein, Kipf, van den Berg</i>

Bayesian Deep Learning

Yarin Gal, Jose Miguel Hernández-Lobato, Christos Louizos, Eric Nalisnick, Zoubin Ghahramani, Kevin Murphy, Max Welling

West Exhibition Hall C, Fri Dec 13, 08:00 AM

Extending on the workshop's success from the past 3 years, this workshop will study the developments in the field of Bayesian deep learning (BDL) over the past year. The workshop will be a platform to host the recent flourish of ideas using Bayesian approaches in deep learning, and using deep learning tools in Bayesian modelling. The program includes a mix of invited talks, contributed talks, and contributed posters. Future directions for the field

will be debated in a panel discussion.

Speakers:

- * Andrew Wilson
- * Deborah Marks
- * Jasper Snoek
- * Roger Grosse
- * Chelsea Finn
- * Yingzhen Li
- * Alexander Matthews

Workshop summary:

While deep learning has been revolutionary for machine learning, most modern deep learning models cannot represent their uncertainty nor take advantage of the well studied tools of probability theory. This has started to change following recent developments of tools and techniques combining Bayesian approaches with deep learning. The intersection of the two fields has received great interest from the community, with the introduction of new deep learning models that take advantage of Bayesian techniques, and Bayesian models that incorporate deep learning elements. Many ideas from the 1990s are now being revisited in light of recent advances in the fields of approximate inference and deep learning, yielding many exciting new results.

Schedule

08:00 AM	Opening remarks	
08:05 AM	Invited talk	
08:25 AM	Contributed talk	
08:40 AM	Invited talk 2	<i>Matthews</i>
09:00 AM	Contributed talk 2	

09:20

AM

Poster spotlights

09:35

AM

Poster session

*Farquhar, Daxberger,
 Look, Benatan, Zhang,
 Havasi, Gustafsson,
 Brofos, Seedat, Livne,
 Ustyuzhaninov, Cobb,
 McGregor, McClure,
 Davidson, Hiranandani,
 Arora, Itkina, Nielsen,
 Harvey,
 Valdenegro-Toro,
 Peluchetti, Moriconi, Cui,
 Smidl, Cemgil,
 Fitzsimons, Zhao, ,
 vargas vieyra,
 Bhattacharyya, Sharma,
 Dubourg-Felonneau,
 Warrell, Voloshynovskiy,
 Rosca, Song, Ross,
 Fashandi, Gao, Shokri
 Razaghi, Chang, Xiao,
 Boehm, Giannone,
 Krishnan, Davison,
 Ashukha, Liu, Huang,
 Nikishin, Park, Ahuja,
 Subedar, , Gadetsky,
 Arias Figueroa, Rudner,
 Aslam, Csiszárík,
 Moberg, Hebbal, Grosse,
 Marttinen, An, Jónsson,
 Kessler, Kumar,
 Figurnov, ,
 Saemundsson, Heljakka,
 Varga, Heim, Rossi,
 Laves, Gharbieh,
 Roberts, Pérez Rey,
 Willetts, Chakrabarty,
 Ghaisas, Shneider,
 Buntine, Adamczewski,
 Gitiaux, Lin, Fu, Rättsch,
 Gomez, Bodin, Phung,
 Svensson, Tusi Amaral
 Laganá Pinto, Alizadeh,
 Du, Murphy, Benk■,
 Vattikutu, Gordon, Kanan,
 Ihler, Graham, Teng,
 Kirsch, Pevny, Holotyak*

10:35
AM **Invited talk 3**

10:55
AM **Contributed talk 3**

11:10
AM **Invited talk 4**

11:30
AM **Contributed talk 4**

01:20
PM **Invited talk 5**

01:40
PM **Contributed talk 5**

01:55
PM **Invited talk 6**

02:10
PM **Contributed talk 6**

02:30
PM **Poster session 2**

03:30
PM **Contributed talk 7**

03:50
PM **Invited talk 7**

04:05
PM **Contributed talk 8**

04:30
PM **Panel session**

05:30
PM **Poster session 3**

Real Neurons & Hidden Units: future directions at the intersection of neuroscience and AI

Guillaume Lajoie, Eli Shlizerman, Maximilian Puelma Touzel, Jessica Thompson, Konrad Kording

East Ballroom A, Sat Dec 14, 08:00 AM

Recent years have witnessed an explosion of progress in AI. With it, a proliferation of experts and practitioners are pushing the boundaries of the field without regard to the brain. This is in stark contrast with the field's transdisciplinary origins, when interest in designing intelligent algorithms was shared by neuroscientists, psychologists and computer scientists alike. Similar progress has been made in neuroscience where novel experimental techniques now afford unprecedented access to brain activity and function. However, it is unclear how to maximize them to truly advance an end-to-end understanding of biological intelligence. The traditional neuroscience research program, however, lacks frameworks to truly advance an end-to-end understanding of biological intelligence. For the first time, mechanistic discoveries emerging from deep learning, reinforcement learning and other AI fields may be able to steer fundamental neuroscience research in ways beyond standard uses of machine learning for modelling and data analysis. For example, successful training algorithms in artificial networks, developed without biological constraints, can motivate research questions and hypotheses about the brain. Conversely, a deeper understanding of brain

computations at the level of large neural populations may help shape future directions in AI. This workshop aims to address this novel situation by building on existing AI-Neuro relationships but, crucially, outline new directions for artificial systems and next-generation neuroscience experiments. We invite contributions concerned with the modern intersection between neuroscience and AI and in particular, addressing questions that can only now be tackled due to recent progress in AI on the role of recurrent dynamics, inductive biases to guide learning, global versus local learning rules, and interpretability of network activity. This workshop will promote discussion and showcase diverse perspectives on these open questions.

Schedule

08:15 AM	Opening Remarks	<i>Lajoie, Thompson, Puelma Touzel, Shlizerman, Kording</i>
08:30 AM	Blake's Talk	<i>Richards</i>
09:00 AM	Tim's Talk	<i>Lillicrap</i>
09:30 AM	Contributed Talk #1	
09:45 AM	Coffee Break + Posters	
10:30 AM	Cristina's Talk	<i>Savin</i>
11:00 AM	David's Talk	<i>Sussillo</i>
11:30 AM	Contributed Talk #2	
11:45 AM	Contributed Talk #3	
12:00 PM	Lunch Break	

02:00 PM	Ila's Talk	<i>Fiete</i>
02:30 PM	Surya's Talk	<i>Ganguli</i>
03:00 PM	Contributed Talk #4	
03:15 PM	Contributed Talk #5	
03:30 PM	Coffee Break + Posters	

04:15
PM **Poster Session**

Sainath, Akrouf, Delahunt, Kutz, Yang, Marino, Abbott, Vecoven, Ernst, warrington, , Kagan, Cho, Harris, Grinberg, Hopfield, Krotov, Muhammad, Cobos, Walker, Reimer, Tolia, Ecker, Sheth, Zhang, Woźnyk, Tabor, Maszke, Pogodin, Corneil, Gerstner, Lin, Cecchi, Reinen, Rish, Bellec, Salaj, Subramoney, Maass, Wang, Pakman, Lee, Paninski, Tripp, Graber, Schwing, Prince, Ocker, Buice, Lansdell, Kording, Lindsey, Sejnowski, Farrell, Shea-Brown, Farrugia, Nepveu, Im, Branson, Hu, Iyer, Mihalas, Aenugu, Hazan, Dai, Nguyen, Tsao, Baraniuk, Anandkumar, Tanaka, Nayebi, Baccus, Ganguli, Pospisil, Muller, Cheng, Varoquaux, Dadi, Gklezakos, Rao, Louis, Papadimitriou, Vempala, Yadati, Zdeblick, Witten, Roberts, Prabhu, Bellec, Ramesh, Macke, Cadena, Bellec, Scherr, Marschall, Kim, Rapp, Fonseca, Armitage, Im, Hardcastle, Sharma, Bair, Valente, Shang, Stern, Patil, Wang, Gorantla, Stratton, Edwards, Lu, Ester, Vlasov

05:00 PM	Doina's Talk	<i>Precup</i>
05:30 PM	Panel Session: A new hope for neuroscience	<i>Bengio, Richards, Lillicrap, Fiete, Sussillo, Precup, Kording, Ganguli</i>



Fair ML in Healthcare

Shalmali Joshi, Irene Y Chen, Ziad Obermeyer, Sendhil Mullainathan

East Ballroom B, Sat Dec 14, 08:00 AM

Clinical healthcare has been a natural application domain for ML with a few modest success stories of practical deployment. Inequity and healthcare disparity has long been a concern in clinical and public health for decades. However, the challenges of fair and equitable care using ML in health has largely remained unexplored. While a few works have attempted to highlight potential concerns and pitfalls in recent years, there are massive gaps in academic ML literature in this context. The goal of this workshop is to investigate issues around fairness that are specific to ML based healthcare. We hope to investigate a myriad of questions via the workshop.

Schedule

09:00 AM	Check in	<i>Wang, Kinyanjui, Zhang, d'Almeida, Tulabandhula, Bayeleygne</i>
09:15 AM	Opening Remarks	
09:30 AM	Keynote - Milind Tambe	

10:00 AM	Invited Talk - Ziad Obermeyer	<i>Obermeyer</i>
		<i>Panda, Sattigeri, Varshney, Natesan Ramamurthy, Singh, Mhasawade, Joshi, Seyyed-Kalantari, McDermott, Yona, Atwood, Srinivasan, Halpern, Sculley, Babaki, Carvalho, Williams, Razavian, Zhang, Lu, Chen, Mao, Zhou, Kallus</i>
10:30 AM	Coffee Break and Poster Session	

11:00 AM	Breakout Sessions	
12:45 PM	Lunch Break	
02:00 PM	Invited Talk - Sharad Goel	
02:30 PM	Invited Talk - Noam Barda/Noam Barda	<i>Barda, Dagan</i>
03:00 PM	Invited Talk - Chelsea Barabas	
03:30 PM	Coffee Break and Poster Session	
04:00 PM	Discussion Panel - All invited speakers will be panelists	
05:00 PM	Spotlight Talks and Poster Session	



Tackling Climate Change with ML

David Rolnick, Priya Donti, Lynn Kaack, Alexandre Lacoste, Tegan Maharaj, Andrew Ng, John Platt, Jennifer Chayes, Yoshua Bengio

East Ballroom C, Sat Dec 14, 08:00 AM

Climate change is one of the greatest problems society has ever faced, with increasingly severe consequences for humanity as natural disasters multiply, sea levels rise, and ecosystems falter. Since climate change is a complex issue, action takes many forms, from designing smart electric grids to tracking greenhouse gas emissions through satellite imagery. While no silver bullet, machine learning can be an invaluable tool in fighting climate change via a wide array of applications and techniques. These applications require algorithmic innovations in machine learning and close collaboration with diverse fields and practitioners. This workshop is intended as a forum for those in the machine learning community who wish to help tackle climate change.

Schedule

08:15 AM	Welcome and Opening Remarks	
08:30 AM	Jeff Dean (Google AI)	<i>Dean</i>
09:05 AM	Spotlight talks	
09:45 AM	Coffee Break + Poster Session	
10:30 AM	Felix Creutzig (TU Berlin, MCC)	<i>Creutzig</i>
11:05 AM	Spotlight talks	
11:15 AM	Panel Discussion	<i>Bengio, Gomes, Ng, Dean, Mackey</i>

12:00 PM **Lunch + Poster Session**

Diehl, Cai, Hoedt, Kochanski, Kim, Lee, Park, Zhou, Gauch, Wilson, Chatterjee, Shrotriya, Papadimitriou, Schön, Zantedeschi, Baasch, Waegeman, Cosne, Farrell, Lucier, Mones, Robinson, Chitsiga, Kristof, Das, Min, Puchko, Luccioni, Story, Hickey, Hu, Lütjens, Wang, Jing, Flaspohler, Wang, Sinha, Tang, Tiihonen, Glatt, Komurcu, Drgona, Gomez-Romero, Kapoor, Fitzpatrick, Rezvanifar, Albert, Irzak, Lamb, Mahesh, Maeng, Kratzert, Friedler, Dalmaso, Robson, Malobola, Maystre, Lin, Mukkavili, Hutchinson, Lacoste, Wang, Wang, Zhang, Preston, Pettit, Vrabie, Molina-Solana, Buonassisi, Annex, Marques, Voss, Rausch, Evans

02:00 PM **Carla Gomes (Cornell)** *Gomes*

02:40 PM **Spotlight talks**

03:30 PM **Coffee Break + Poster Session**

04:15 PM	Lester Mackey (Microsoft Research <i>Mackey</i> and Stanford)
04:40 PM	Spotlight talks
05:00 PM	Panel Discussion <i>Chayes, Platt, Creutzig, Gonzalez, Miller</i>

Joint Workshop on AI for Social Good

Fei Fang, Joseph Bullock, Marc-Antoine Dilhac, Brian Green, natalie saltiel, Dhaval Adjodah, Jack Clark, Sean McGregor, Margaux Luck, Jonnie Penn, Tristan Sylvain, Geneviève Boucher, Sydney Swaine-Simon, Girmaw Abebe Tadesse, Myriam Côté, Anna Bethke, Yoshua Bengio

East Meeting Rooms 11 + 12, Sat Dec 14, 08:00 AM

The accelerating pace of intelligent systems research and real world deployment presents three clear challenges for producing "good" intelligent systems: (1) the research community lacks incentives and venues for results centered on social impact, (2) deployed systems often produce unintended negative consequences, and (3) there is little consensus for public policy that maximizes "good" social impacts, while minimizing the likelihood of harm. As a result, researchers often find themselves without a clear path to positive real world impact.

The Workshop on AI for Social Good addresses these challenges by bringing together machine learning researchers, social impact leaders, ethicists, and public policy leaders to present their ideas and applications for maximizing the social good. This workshop is a collaboration of three

formerly separate lines of research (i.e., this is a "joint" workshop), including researchers in applications-driven AI research, applied ethics, and AI policy. Each of these research areas are unified into a 3-track framework promoting the exchange of ideas between the practitioners of each track.

We hope that this gathering of research talent will inspire the creation of new approaches and tools, provide for the development of intelligent systems benefiting all stakeholders, and converge on public policy mechanisms for encouraging these goals.

Schedule

08:00 AM	Opening remarks	<i>Bengio</i>
08:05 AM	Track 1: Producing Good Outcomes	<i>Dietterich, Gomes, Luengo-Oroz, Dilkina, Cornebise</i>
10:30 AM	Break	

		<i>Abebe, Kleinberg, Lucier, Cuthbertson, Mathewson, Schumann, Chi, Babirye, Lim, Rane, Owoeye, Da San Martino, Kimura, Rutkowski, Fruehwirt, Rho, Charpignon, Konya, Ben Daya, Thomas, Abdulrahim, Ssendiwala, Namanya, Akera, Manandhar, Greeff, Verma, Nyman, Kermode, , Narain, Johnson, Yanagihara, Sugiyama, Sharma, Dey, Sarbajna, Govindaraj, Cornebise, Dulhanty, Deglint, Bilich, Masood, Varga, Gomes, Dietterich, Luengo-Oroz, Dilkina, Mironova, Yu, Srikanth, Clifton, Larson, Levin, Adams-Cohen, Dean</i>
11:00 AM	Track 1: Producing Good Outcomes	
12:00 PM	Lunch - on your own	
02:00 PM	Track 2: From Malicious Use to Responsible AI	<i>Dobbe, Shamout, Clifton, Whittlestone, Kinsey, Elhalal, Bajaj, Wall, Tomasev, Green</i>
03:00 PM	Break	
03:30 PM	Track 2: From Malicious Use to Responsible AI	<i>Yang, Lin, Tomasev, Raicu, Vincent</i>

04:00 PM	Track 3: Public Policy	<i>Sun, Veeramachaneni, Ramirez Diaz, Cuesta-Infante, Elzayn, Gamper, Schim van der Loeff, Green</i>
----------	-------------------------------	--

Machine Learning for Autonomous Driving

Rowan McAllister, Nick Rhinehart, Fisher Yu, Li Erran Li, Anca Dragan

East Meeting Rooms 1 - 3, Sat Dec 14, 08:00 AM

Autonomous vehicles (AVs) provide a rich source of high-impact research problems for the machine learning (ML) community at NeurIPS in diverse fields including computer vision, probabilistic modeling, gesture recognition, pedestrian and vehicle forecasting, human-machine interaction, and multi-agent planning. The common goal of autonomous driving can catalyze discussion between these subfields, generating a cross-pollination of research ideas. Beyond the benefits to the research community, AV research can improve society by reducing road accidents; giving independence to those unable to drive; and inspiring younger generations towards ML with tangible examples of ML-based technology clearly visible on local streets.

As many NeurIPS attendees are key drivers behind AV-applied ML, the proposed NeurIPS 2019 Workshop on Autonomous Driving intends to bring researchers together from both academia and industries to discuss machine learning applications in autonomous driving. Our proposal includes regular paper presentations, invited speakers, and technical benchmark challenges to present the current state of the art, as well as the limitations and

future directions for autonomous driving.

Schedule

08:45 AM	Welcome	<i>McAllister, Rhinehart, Dragan</i>
09:00 AM	Invited Talk	<i>Urtasun</i>
09:30 AM	Contributed Talks	<i>Chen, Gählert, Leurent, Lehner, Bhattacharyya, Behl, Lim, Kim, Novosel, Osinski, Das, Shen, Hawke, Sicking, Shahian Jahromi, Tulabandhula, Michaelis, Rusak, BAO, Rashed, Chen, Ansari, Cha, Zahran, Reda, Kim, Dohyun, Suk, Jhung, Kister, Fahrland, Jakubowski, Mio, Mercat, Arsenal, Homoceanu, Liu, Torr, El Sallab, Sobh, Arnab, Galias</i>
09:45 AM	Coffee + Posters	
10:30 AM	Invited Talk	<i>Rus</i>
11:00 AM	Invited Talk	<i>Karpathy</i>
11:30 AM	Invited Talk	<i>Koltun</i>
12:00 PM	Lunch + Posters	
01:30 PM	Invited Talk	<i>Frazzoli</i>

02:00 PM	Invited Talk	<i>Wu</i>
02:30 PM	Invited Talk	<i>Fernández Fisac</i>
03:00 PM	Contributed Talks	<i>Caine, Wang, Sakib, Otawara, Kaushik, amirloo, Djuric, Rock, Agarwal, Filos, Tigkas, Lee, Jeon, Jaipuria, Wang, Zhao, Zhang, Singh, Banijamali, Rohani, Sinha, Joshi, Chan, Abdou, Chen, Kim, mohamed, OKelly, Singhanian, Tsukahara, Keyaki, Palanisamy, Norden, Marchetti-Bowick, Gu, Arora, Deshpande, Schneider, Jui, Aggarwal, Gangopadhyay, Yan</i>
03:30 PM	Coffee + Posters	
04:30 PM	Invited Talk	<i>Gilitschenski</i>
05:00 PM	Invited Talk	<i>Baker</i>
05:30 PM	Competition	<i>Chang, Singh, Hartnett, Cebren</i>



Privacy in Machine Learning (PriML)

Borja Balle, Kamalika Chaudhuri, Antti Honkela, Antti Koskela, Casey Meehan, Mi Jung Park, Mary Anne Smart, Adrian Weller

East Meeting Rooms 8 + 15, Sat Dec 14, 08:00 AM

The goal of our workshop is to bring together privacy experts working in academia and industry to discuss the present and the future of privacy-aware technologies powered by machine learning. The workshop will focus on the technical aspects of privacy research and deployment with invited and contributed talks by distinguished researchers in the area. The programme of the workshop will emphasize the diversity of points of view on the problem of privacy. We will also ensure there is ample time for discussions that encourage networking between researches, which should result in mutually beneficial new long-term collaborations.

Schedule

08:15 AM	TBA: Brendan McMahan	<i>McMahan</i>
10:30 AM	TBA: Ashwin Machanavajjhala	<i>Machanavajjhala</i>

11:30 AM

Poster Session

Canonne, Jun, Neel, Wang, vietri, Song, Lebensold, Zhang, Gondara, Li, Miresghallah, Dong, Sarwate, Koskela, Jälkö, Kusner, Chen, Park, Machanavajjhala, Kalpathy-Cramer, , Feldman, Tomkins, Phan, Esfandiari, Jaiswal, Sharma, Druce, Meehan, Zhao, Hsu, Railsback, Flaxman, , Adebayo, Korolova, Xu, Holohan, Basu, Joseph, Thai, Yang, Vitercik, Hutchinson, Wang, Yauney, Tao, Jin, Lee, McMillan, Izmailov, Guo, Swaroop, Orekondy, Esmailzadeh, Procopio, Polyzotis, Mohammadi

02:00 PM

TBA: Lalitha Sankar Sankar

04:15 PM

TBA: Philip Leclerc Leclerc

Abstracts (5):

Abstract 1: **TBA: Brendan McMahan in Privacy in Machine Learning (PriML)**, *McMahan* 08:15 AM

Tentative schedule, details TBA.

Abstract 2: **TBA: Ashwin Machanavajjhala in Privacy in Machine Learning (PriML)**, *Machanavajjhala* 10:30 AM

Tentative schedule, details TBA.

Abstract 3: Poster Session in Privacy in Machine Learning (PriML), *Canonne, Jun, Neel, Wang,*

vietri, Song, Lebensold, Zhang, Gondara, Li, Mireshghallah, Dong, Sarwate, Koskela, Jälkö, Kusner, Chen, Park, Machanavajjhala, Kalpathy-Cramer, , Feldman, Tomkins, Phan, Esfandiari, Jaiswal, Sharma, Druce, Meehan, Zhao, Hsu, Railsback, Flaxman, , Adebayo, Korolova, Xu, Holohan, Basu, Joseph, Thai, Yang, Vitercik, Hutchinson, Wang, Yauney, Tao, Jin, Lee, McMillan, Izmailov, Guo, Swaroop, Orekondy, Esmailzadeh, Procopio, Polyzotis, Mohammadi
11:30 AM

Schedule is not final. Details TBA.

Abstract 4: TBA: Lalitha Sankar in Privacy in Machine Learning (PriML), *Sankar* 02:00 PM

Tentative schedule, details TBA.

Abstract 5: TBA: Philip Leclerc in Privacy in Machine Learning (PriML), *Leclerc* 04:15 PM

Tentative schedule, details TBA.

Machine Learning and the Physical Sciences

Atilim Gunes Baydin, Juan Carrasquilla, Shirley Ho, Karthik Kashinath, Michela Paganini, Savannah Thais, Anima Anandkumar, Kyle Cranmer, Roger Melko, Mr. Prabhat, Frank Wood

West 109 + 110, Sat Dec 14, 08:00 AM

Machine learning methods have had great success in learning complex representations that enable them to make predictions about unobserved data. Physical sciences span problems and challenges at all scales in the universe: from finding exoplanets in trillions of sky pixels, to finding machine learning inspired solutions to the quantum many-body

problem, to detecting anomalies in event streams from the Large Hadron Collider. Tackling a number of associated data-intensive tasks including, but not limited to, segmentation, 3D computer vision, sequence modeling, causal reasoning, and efficient probabilistic inference are critical for furthering scientific discovery. In addition to using machine learning models for scientific discovery, the ability to interpret what a model has learned is receiving an increasing amount of attention.

In this targeted workshop, we would like to bring together computer scientists, mathematicians and physical scientists who are interested in applying machine learning to various outstanding physical problems, in particular in inverse problems and approximating physical processes; understanding what the learned model really represents; and connecting tools and insights from physical sciences to the study of machine learning models. In particular, the workshop invites researchers to contribute papers that demonstrate cutting-edge progress in the application of machine learning techniques to real-world problems in physical sciences, and using physical insights to understand what the learned model means.

By bringing together machine learning researchers and physical scientists who apply machine learning, we expect to strengthen the interdisciplinary dialogue, introduce exciting new open problems to the broader community, and stimulate production of new approaches to solving open problems in sciences. Invited talks from leading individuals in both communities will cover the state-of-the-art techniques and set the stage for this workshop.

Schedule

08:30 AM	Opening Remarks	<i>Baydin, Carrasquilla, Ho, Kashinath, Paganini, Thais, Anandkumar, Cranmer, Melko, Prabhat, Wood</i>	
08:40 AM	Bernhard Schölkopf	<i>Schölkopf</i>	
09:20 AM	JAX, M.D.: End-to-End Differentiable, Hardware Accelerated, Molecular Dynamics in Pure Python	<i>Schoenholz</i>	
09:40 AM	Morning Coffee Break & Poster Session		<i>Metodiev, Zhang, Stoye, Churchill, Sarkar, Cranmer, Brehmer, Jimenez Rezende, Harrington, Nigam, Thuerey, Maziarka, Sanchez Gonzalez, Okan, Ritchie, Erichson, Cheng, Jiang, Pahng, Koelle, Khairy, Pol, Anirudh, Born, Sanchez-Lengeling, Timar, Goodall, Kriváchy, Lu, Adler, Trask, Cherrier, Konno, Kasim, Golling, Alperstein, Ustyuzhanin, Stokes, Golubeva, Char, Korovina, Cho, Chatterjee, Westerhout, Muñoz-Gil, Zamudio-Fernandez, Wei, Lee, Kofler, Power, Kazeev, Ustyuzhanin, Maevskiy, Friederich, Tavakoli, Neiswanger, Kulchytskyy, hari, Leu, Atzberger</i>
10:40 AM	Maria Schuld	<i>Schuld</i>	
11:20 AM	Metric Methods with Open Collider Data	<i>Metodiev</i>	
11:40 AM	Equivariant Hamiltonian Flows	<i>Jimenez Rezende</i>	

12:00 PM	Towards physics-informed deep learning for turbulent flow prediction	<i>Yu</i>
12:20 PM	Lunch Break	
02:00 PM	Katie Bouman	<i>Bouman</i>
02:40 PM	Lenka Zdeborova	<i>Zdeborová</i>
03:20 PM	Afternoon Coffee Break & Poster Session	<i>Komkov, Fort, Wang, Yu, Park, Schoenholz, Cheng, Griffiths, Shimmin, Mukkavili, Schwaller, Knoll, Sun, Kisamori, Graham, Portwood, Huang, Novello, Munchmeyer, Jungbluth, Levine, Ayed, Atkinson, Hermann, Grönquist, Saha, Glaser, Li, Iiyama, Anirudh, Koch-Janusz, Sundar, Lanusse, Köhler, Yip, guo, Ju, Hanuka, Albert, Salvatelli, Verzetti, Duarte, Moreno, de Bézenac, Vlontzos, Singh, Klijnsma, Neuberg, Wright, Mustafa, Schmidt</i>
04:20 PM	Alán Aspuru-Guzik	<i>Aspuru-Guzik</i>
05:00 PM	Yasaman Bahri	<i>Bahri</i>

05:40 PM	Hamiltonian Graph Networks with ODE Integrators	<i>Sanchez Gonzalez</i>
05:55 PM	Learning Symbolic Physics with Graph Networks	<i>Cranmer</i>



Program Transformations for ML

Pascal Lamblin, Atilim Gunes Baydin, Alexander Wiltschko, Bart van Merriënboer, Emily Fertig, Barak Pearlmutter, David Duvenaud, Laurent Hascoet

West 114 + 115, Sat Dec 14, 08:00 AM

Machine learning researchers often express complex models as a program, relying on program transformations to add functionality. New languages and transformations (e.g., TorchScript and TensorFlow AutoGraph) are becoming core capabilities of ML libraries. However, existing transformations, such as automatic differentiation (AD), inference in probabilistic programming languages (PPL), and optimizing compilers are often built in isolation, and limited in scope. This workshop aims at viewing program transformations in ML in a unified light, making these capabilities more accessible, and building entirely new ones. Program transformations are an area of active study. AD transforms a program performing numerical computation into one computing the gradient of those computations. In PPL, a program describing a sampling procedure can be modified to perform inference on model parameters given observations. Other examples are vectorizing a program expressed on one data point, and learned transformations where ML models use programs as inputs or outputs.

This workshop will bring together researchers in the fields of AD, programming languages, compilers,

and ML, with the goal of understanding the commonalities between disparate approaches and views, and sharing ways to make these techniques broadly available. It would enable ML practitioners to iterate faster on novel models and architectures (e.g., those naturally expressed through high-level constructs like recursion).

Topics:

- Abstractions and syntax (beyond meta-programming and operator overloading) to naturally express a program (expression, or procedure) as an object to be manipulated.
- Techniques from AD and PPL the ML community could adopt to enable research on new models
- How to overcome challenges due to the ML’s specific hardware (GPUs, specialized chips) and software (Python) stacks, and the particular demands of practitioners for their tools
- Greater collaboration between ML and programming languages communities

Schedule

08:30 AM	Opening statements	
08:40 AM	Jan-Willem van de Meent - TBA	<i>van de Meent</i>
09:30 AM	Applications of a disintegration transformation	<i>Narayanan</i>
09:50 AM	Coffee break	
10:30 AM	Christine Tasson - TBA	<i>Tasson</i>
11:20 AM	The Differentiable Curry	<i>Vytiniotis</i>
11:40 AM	Functional Tensors for Probabilistic Programming	<i>Obermeyer</i>

		<i>Considine, Innes, Phan, Maclaurin, Manhaeve, Radul, Gowda, Sharma, Sennesh, Kochurov, Plotkin, Wiecki, Kukreja, Shan, Johnson, Belov, Pradhan, Meert, Kimmig, De Raedt, Patton, Hoffman, A. Saurous, Roy, Bingham, Jankowiak, Carroll, Lao, Paull, Abadi, Rojas Jimenez, Chen</i>
12:00 PM	Lunch break & Poster session	
02:00 PM	Zachary DeVito - TBA	<i>DeVito</i>
02:50 PM	Skye Wanderman-Milne - TBA	<i>Wanderman-Milne</i>
03:40 PM	Coffee break	
04:20 PM	Generalized Abs-Linear Learning	<i>Griewank</i>
04:40 PM	Towards Polyhedral Automatic Differentiation	<i>Hueckelheim</i>
05:00 PM	Taylor-Mode Automatic Differentiation for Higher-Order Derivatives in JAX	<i>Bettencourt</i>
05:20 PM	Panel and general discussion	



Competition Track Day 2

Hugo Jair Escalante**West 116 + 117, Sat Dec 14, 08:00 AM**<https://nips.cc/Conferences/2019/CallForCompetitions>**Schedule**

08:00 AM	Causality for Climate (C4C)	
09:00 AM	The MineRL competition	<i>Ogura, Booth, Sun, Topin, Houghton, Guss, Milani, Vinyals, Hofmann</i>
11:15 AM	The Game of Drones Competition	<i>Toumieh, Vemprala</i>
12:15 PM	The AutoDL Challenge	<i>Treguer, Kim</i>
02:15 PM	Overview of the Live Reinforcement Learning Malaria Challenge	<i>Remy</i>

Abstracts (2):

Abstract 2: **The MineRL competition in Competition Track Day 2**, *Ogura, Booth, Sun, Topin, Houghton, Guss, Milani, Vinyals, Hofmann*
09:00 AM

MineRL Competition on Sample Efficient Reinforcement Learning.

Competition chairs: Brandon Houghton, William Guss, Stephanie Milani, Nicholay Topin

* Overview and highlights of the competition.
Brandon Houghton, William Guss

* Competition Awards. Stephanie Milani

* Special Awards. Oriol Vinyals & advisory board.

* Discussion of future competitions. Katja Hofmann

* Competitors Presentations

Abstract 3: **The Game of Drones Competition in Competition Track Day 2**, *Toumieh, Vemprala*
11:15 AM

Opening/Introduction

Tier 1 Presentations

Tier 2 Presentations

Tier 3 Presentations

Prize Distribution

Emergent Communication: Towards Natural Language

Abhinav Gupta, Michael Noukhovitch, Cinjon Resnick, Natasha Jaques, Angelos Filos, Marie Ossenkopf, Angeliki Lazaridou, Jakob Foerster, Ryan Lowe, Douwe Kiela, Kyunghyun Cho

West 118 - 120, Sat Dec 14, 08:00 AM

Communication is one of the most impressive human abilities but historically it has been studied in machine learning on confined datasets of natural language, and by various other fields in simple low-dimensional spaces. Recently, with the rise of deep RL methods, the questions around the emergence of communication can now be studied in new, complex multi-agent scenarios. Two previous successful workshops (2017, 2018) have gathered the community to discuss how, when, and to what end communication emerges, producing research that was later published at top ML venues such as ICLR, ICML, AAAI. Now, we wish to extend these ideas and explore a new direction: how emergent

communication can become more like natural language, and what natural language understanding can learn from emergent communication.

The push towards emergent natural language is a necessary and important step in all facets of the field. For studying the evolution of human language, emerging a natural language can uncover the requirements that spurred crucial aspects of language (e.g. compositionality). When emerging communication for multi-agent scenarios, protocols may be sufficient for machine-machine interactions, but emerging a natural language is necessary for human-machine interactions. Finally, it may be possible to have truly general natural language understanding if agents learn the language through interaction as humans do. To make this progress, it is necessary to close the gap between artificial and natural language learning.

To tackle this problem, we want to take an interdisciplinary approach by inviting researchers from various fields (machine learning, game theory, evolutionary biology, linguistics, cognitive science, and programming languages) to participate and engaging them to unify the differing perspectives. We believe that the third iteration of this workshop with a novel, unexplored goal and strong commitment to diversity will allow this burgeoning field to flourish.

Schedule

08:00 AM	Posters	<i>LaCroix, Ossenkopf, Lee, Fitzgerald, Mihai, Hare, Zaidi, Cowen-Rivers, Brown, Marzoev, Kharitonov, Yuan, Korbak, Liang, Ren, Dessi, Potash, Guo, Hashimoto, Liang, Zubek, Fu, Zhu</i>
----------	----------------	--

08:55 AM	Intro Remarks	
09:00 AM	Invited Talk - 1	<i>Gibson</i>
09:45 AM	Contributed Talk - 1 <i>Lee</i>	
10:00 AM	Coffee Break / Poster Session	
10:30 AM	Invited Talk - 2	<i>Zaslavsky</i>
11:15 AM	Contributed Talk - 2 <i>Cowen-Rivers</i>	
11:30 AM	Spotlight presentations x5	
02:00 PM	Invited Talk - 3	<i>Eisner</i>
02:45 PM	Contributed Talk - 3 <i>Brown</i>	
03:00 PM	Invited Talk - 4	<i>Andreas</i>
03:45 PM	Coffee Break / Poster Session	
04:15 PM	Invited Talk - 5	<i>Lee</i>
05:00 PM	Panel Discussion	
05:55 PM	Closing Remarks	



Science meets Engineering of Deep Learning

Levent Sagun, CAGLAR Gulcehre, Adriana Romero, Negar Rostamzadeh, Nando de Freitas

West 121 + 122, Sat Dec 14, 08:00 AM

Deep learning can still be a complex mix of art and engineering despite its tremendous success in recent years, and there is still progress to be made before it has fully evolved into a mature scientific discipline. The interdependence of architecture, data, and optimization gives rise to an enormous landscape of design and performance intricacies that are not well-understood. The evolution from engineering towards science in deep learning can be achieved by pushing the disciplinary boundaries. Unlike in the natural and physical sciences -- where experimental capabilities can hamper progress, i.e. limitations in what quantities can be probed and measured in physical systems, how much and how often -- *in deep learning the vast majority of relevant quantities that we wish to measure can be tracked in some way*. As such, a greater limiting factor towards scientific understanding and principled design in deep learning is how to *insightfully harness the tremendous collective experimental capability of the field*. As a community, some primary aims would be to (i) identify obstacles to better models and algorithms, (ii) identify the general trends that are potentially important which we wish to understand scientifically and potentially theoretically and; (iii) careful design of scientific experiments whose purpose is to clearly resolve and pinpoint the origin of mysteries (so-called 'smoking-gun' experiments).

Schedule

		<i>Sagun, Gulcehre,</i>
08:00 AM	Welcoming remarks and introduction	<i>Romero, Rostamzadeh, de Freitas</i>
08:15 AM	Session 1 - Theory	<i>Krzakala, Bahri, Ganguli, Zdeborová, Dieng, Bruna</i>
09:45 AM	Coffee and posters	

10:30 AM	Session 2 - Vision	<i>Schmid, Urtasun, Fidler, Neverova, Radosavovic</i>
		<i>Song, Hoffer, Chang, Cohen, Islam, Blumenfeld, Madsen, Frankle, Goldt, Chatterjee, Panigrahi, Renda, Bartoldson, Birhane, Baratin, Chatterji, Novak, Forde, Jiang, Du, Adilova, Kamp, Weinstein, Hubara, Ben-Nun, Hoefler, Soudry, Yu, Zhong, Yang, Dhillon, Carbonell, Zhang, Gilboa, Brandstetter, Johansen, Dziugaite, Somani, Morcos, Kalaitzis, Sedghi, Xiao, Zech, Yang, Kaur, Ma, Tsai, Salakhutdinov, Yaida, Lipton, Roy, Carbin, Krzakala, Zdeborová, Gur-Ari, Dyer, Krishnan, Mobahi, Bengio, Neyshabur, Netrapalli, Sankaran, Cornebise, Bengio, Michalski, Ebrahimi Kahou, Arefin, Hron, Lee, Sohl-Dickstein, Schoenholz, Schwab, Li, Choe, Petzka, Verma, Lin</i>
12:00 PM	Lunch Break and posters	
02:00 PM	Session 3 - Further Applications	<i>Durand, Cho, Chaudhuri, Dauphin, Firat, Gorur</i>

03:30 PM	Coffee and posters	
04:15 PM	Panel - The Role of Communication at Large	<i>Lakshmiratan, Yakubova, Doshi-Velez, Ganguli, Lipton, Paganini, Anandkumar</i>
05:10 PM	Contributed Session - Spotlight Talks	<i>Frankle, Schwab, Morcos, Ma, Tsai, Salakhutdinov, Jiang, Krishnan, Mobahi, Bengio, Yaida, Yang</i>

Abstracts (1):

Abstract 5: **Lunch Break and posters in Science meets Engineering of Deep Learning**, *Song, Hoffer, Chang, Cohen, Islam, Blumenfeld, Madsen, Frankle, Goldt, Chatterjee, Panigrahi, Renda, Bartoldson, Birhane, Baratin, Chatterji, Novak, Forde, Jiang, Du, Adilova, Kamp, Weinstein, Hubara, Ben-Nun, Hoefler, Soudry, Yu, Zhong, Yang, Dhillon, Carbonell, Zhang, Gilboa, Brandstetter, Johansen, Dziugaite, Somani, Morcos, Kalaitzis, Sedghi, Xiao, Zech, Yang, Kaur, Ma, Tsai, Salakhutdinov, Yaida, Lipton, Roy, Carbin, Krzakala, Zdeborová, Gur-Ari, Dyer, Krishnan, Mobahi, Bengio, Neyshabur, Netrapalli, Sankaran, Cornebise, Bengio, Michalski, Ebrahimi Kahou, Arefin, Hron, Lee, Sohl-Dickstein, Schoenholz, Schwab, Li, Choe, Petzka, Verma, Lin* 12:00 PM

Since we are a small workshop, we will hold the poster sessions during the day, including all the breaks as the authors wish.

ML For Systems

Milad Hashemi, Azalia Mirhoseini, Anna Goldie, Kevin Swersky, Xinlei XU, Jonathan Raiman

West 202 - 204, Sat Dec 14, 08:00 AM

Compute requirements are growing at an exponential rate, and optimizing these computer systems often involves complex high-dimensional combinatorial problems. Yet, current methods rely heavily on heuristics. Very recent work has outlined a broad scope where machine learning vastly outperforms these traditional heuristics: including scheduling, data structure design, microarchitecture, compilers, circuit design, and the control of warehouse scale computing systems. In order to continue to scale these computer systems, new learning approaches are needed. The goal of this workshop is to develop novel machine learning methods to optimize and accelerate software and hardware systems.

Machine Learning for Systems is an interdisciplinary workshop that brings together researchers in computer architecture and systems and machine learning. This workshop is meant to serve as a platform to promote discussions between researchers in the workshops target areas.

This workshop is part two of a two-part series with one day focusing on ML for Systems and the other on Systems for ML. Although the two workshops are being led by different organizers, we are coordinating our call for papers to ensure that the workshops complement each other and that submitted papers are routed to the appropriate venue.

Schedule

09:00 AM	Opening	
09:10 AM	Invited Speaker 1	<i>Bakshy</i>
09:45 AM	Break	

10:30 AM	Poster Session 1	<i>Mao, Nathan, Baldini, Sivakumar, Wang, Magalle Hewa, Shi, Kaufman, Fang, Zhou, Ding, He, Lubin</i>
11:00 AM	Contributed Talk 1	
11:15 AM	Contributed Talk 2	
11:30 AM	Contributed Talk 3	
11:45 AM	Contributed Talk 4	
12:00 PM	Lunch	
01:30 PM	Invited Speaker 2	<i>Stoica</i>
02:00 PM	Invited Speaker 3	<i>Jain</i>
02:30 PM	Contributed Talk 5	
02:45 PM	Contributed Talk 6	
03:00 PM	Poster Session 2	<i>Wang, Lin, Duan, Paliwal, Haj-Ali, Marcus, Hope, Xu, Le, Sun, Cutler, Nathan, Sun</i>
03:30 PM	Break	
04:15 PM	Invited Speaker 4	
04:45 PM	Keynote Speaker	<i>Dean</i>
05:15 PM	Panel	

The third Conversational AI workshop – today's practice and tomorrow's potential

Alborz Geramifard, Jason Williams, Bill Byrne, Asli Celikyilmaz, Milica Gasic, Dilek Hakkani-Tur, Matt Henderson, Luis Lastras, Mari Ostendorf

West 205 - 207, Sat Dec 14, 08:00 AM

In the span of only a few years, conversational systems have become commonplace. Every day, millions of people use natural-language interfaces such as Siri, Google Now, Cortana, Alexa and others via in-home devices, phones, or messaging channels such as Messenger, Slack, Skype, among others. At the same time, interest among the research community in conversational systems has blossomed: for supervised and reinforcement learning, conversational systems often serve as both a benchmark task and an inspiration for new ML methods at conferences which don't focus on speech and language per se, such as NIPS, ICML, IJCAI, and others. Such movement has not been unnoticed by major publications. This year in collaboration with AAIL community, the AI magazine will have a special issue on conversational AI (<https://tinyurl.com/y6shq2ld>). Moreover, research community challenge tasks are proliferating, including the seventh Dialog Systems Technology Challenge (DSTC7), the Amazon Alexa prize, and the Conversational Intelligence Challenge live competitions at NIPS (2017, 2018).

Following the overwhelming participation in our last two NeurIPS workshops:
 2017: 9 invited talks, 26 submissions, 3 oral papers, 13 accepted papers, 37 reviewers
 2018: 4 invited talks, 42 submission, 6 oral papers, 23 accepted papers, 58 reviewers, we are excited to continue promoting cross-pollination of ideas between academic research centers and industry.

The goal of this workshop is to bring together researchers and practitioners in this area, to clarify impactful research problems, understand well-founded methods, share findings from large-scale real-world deployments, and generate new ideas for future lines of research.

This one day workshop will include invited talks and a panel from academia and industry, contributed work, and open discussion.

Schedule

08:30 AM	Opening	<i>Geramifard, Williams</i>
08:40 AM	Invited talk - Gabriel Skantze	<i>Skantze</i>
09:10 AM	Invited talk - Zhou Yu	<i>Yu</i>
09:40 AM	Poster lighting round	<i>Zheng, Søgaard, Saleh, Jang, Gong, Florez, Li, Madotto, Nguyen, Kulikov, einolghozati, Wang, Eric, Hansen, Lubis, Wu</i>
09:55 AM	Posters + coffee break	
10:40 AM	Contributed talk 1	
10:55 AM	Contributed talk 2	
11:10 AM	Contributed talk 3	
11:25 AM	Contributed talk 4	
11:40 AM	Invited talk - Alan Ritter	

12:10 PM	Lunch	
01:45 PM	Invited talk - David Traum	<i>Traum</i>
02:15 PM	Invited talk - Y-Lan Boureau	<i>Boureau</i>
02:45 PM	Contributed talk 5	
03:00 PM	Contributed talk 6	
03:15 PM	Contributed talk 7	
03:30 PM	Posters + coffee break	
04:15 PM	Invited Talk - Ryuichiro Higashinaka	<i>Higashinaka</i>
04:45 PM	Contributed talk 8	
05:00 PM	Panel discussion	
05:50 PM	Closing	<i>Geramifard, Williams</i>



Document Intelligence

Nigel Duffy, Rama Akkiraju, Tania Bedrax Weiss, Paul Bennett, Hamid Reza Motahari-Nezhad

West 208 + 209, Sat Dec 14, 08:00 AM

Business documents are central to the operation of business. Such documents include sales agreements, vendor contracts, mortgage terms, loan applications, purchase orders, invoices, financial statements, employment agreements and a wide many more. The information in such business documents is presented in natural

language, and can be organized in a variety of ways from straight text, multi-column formats, and a wide variety of tables. Understanding these documents is made challenging due to inconsistent formats, poor quality scans and OCR, internal cross references, and complex document structure. Furthermore, these documents often reflect complex legal agreements and reference, explicitly or implicitly, regulations, legislation, case law and standard business practices.

The ability to read, understand and interpret business documents, collectively referred to here as “Document Intelligence”, is a critical and challenging application of artificial intelligence (AI) in business. While a variety of research has advanced the fundamentals of document understanding, the majority have focused on documents found on the web which fail to capture the complexity of analysis and types of understanding needed across business documents. Realizing the vision of document intelligence remains a research challenge that requires a multi-disciplinary perspective spanning not only natural language processing and understanding, but also computer vision, knowledge representation and reasoning, information retrieval, and more -- all of which have been profoundly impacted and advanced by neural network-based approaches and deep learning in the last few years. We propose to organize a workshop for AI researchers, academics and industry practitioners to discuss the opportunities and challenges for document intelligence.

Schedule

08:30 AM	Opening Remarks	
08:30 AM	David Lewis	<i>Lewis</i>
09:30 AM	Ndapa Nakashole	<i>Nakashole</i>

10:30 AM	Coffee Break	
11:00 AM	Discussion Session / Posters	
02:00 PM	Rajasekar Krishnamurthy	<i>Krishnamurthy</i>
03:00 PM	Coffee Break	
03:30 PM	Asli Celikyilmaz	<i>Celikyilmaz</i>
		<i>Denk, Androustopoulos, Bakhteev, Kane, Stojanov, Park, Mamidibathula, Liepieshov, Höhne, Feng, Bayraktar, Aruleba, OGALTSOV, Kuznetsova, Bennett, Fadnis, Lastras, Jabbarzadeh Gangeh, Reisswig, Elwany, Chalkidis, DeGange, Zhang, de Oliveira, Koçyigit, Dong, Liao</i>
04:30 PM	Discussion / Posters	
05:00 PM	Closing Remarks	

Learning Transferable Skills

Marwan Mattar, Arthur Juliani, Danny Lange, Matthew Crosby, Benjamin Beyret

West 211 - 214, Sat Dec 14, 08:00 AM

After spending several decades on the margin of AI, reinforcement learning has recently emerged as a

powerful framework for developing intelligent systems that can solve complex tasks in real-world environments. This has had a tremendous impact on a wide range of tasks ranging from playing games such as Go and StarCraft to learning dexterity. However, one attribute of intelligence that still eludes modern learning systems is generalizability. Until very recently, the majority of reinforcement learning research involved training and testing algorithms on the same, sometimes deterministic, environment. This has resulted in algorithms that learn policies that typically perform poorly when deployed in environments that differ, even slightly, from those they were trained on. Even more importantly, the paradigm of task-specific training results in learning systems that scale poorly to a large number of (even interrelated) tasks.

Recently there has been an enduring interest in developing learning systems that can learn transferable skills. This could mean robustness to changing environment dynamics, the ability to quickly adapt to environment and task variations or the ability to learn to perform multiple tasks at once (or any combination thereof). This interest has also resulted in a number of new data sets and challenges (e.g. Obstacle Tower Environment, Animal-AI, CoinRun) and an urgency to standardize the metrics and evaluation protocols to better assess the generalization abilities of novel algorithms. We expect this area to continue to increase in popularity and importance, but this can only happen if we manage to build consensus on which approaches are promising, and, equally important, how to test them.

The workshop will include a mix of invited speakers, accepted papers (oral and poster sessions) and a panel discussion. The workshop welcomes both theoretical and applied research, in addition to novel data sets and evaluation protocols.

Schedule

09:00 AM	Opening Remarks	<i>Mattar, Juliani, Crosby, Beyret, Lange</i>
09:15 AM	Raia Hadsell (DeepMind)	<i>Hadsell</i>
10:00 AM	Environments and Data Sets	<i>Cobbe, De Fabritiis</i>
11:00 AM	Coffee Break	
11:15 AM	Vladlen Koltun (Intel)	<i>Koltun</i>
12:00 PM	Lunch	
01:30 PM	David Ha (Google Brain)	<i>Ha</i>
02:15 PM	Oral Presentations	<i>Grau-Moya, Pascual-Diaz, Marinier, Pietquin, Efros, Isola, Darrell, Lu, Pathak, Ferret</i>

03:15 PM	Poster Presentations	<p><i>Mehta, Lampinen, Chen, Pascual-Diaz, Grau-Moya, Faisal, Tompson, Lu, Khetarpal, Klissarov, Bacon, Precup, Kurutach, Tamar, Abbeel, He, Igl, Whiteson, Boehmer, Marinier, Pietquin, Hausman, Levine, Finn, Yu, Lee, Eysenbach, Parisotto, Xing, Salakhutdinov, Ren, Anandkumar, Pathak, Lu, Darrell, Efros, Isola, Liu, Han, Niu, Sugiyama, Kumar, Petangoda, Ferret, McClelland, Liu, Garg, Lange</i></p>
04:15 PM	Katja Hofmann (Microsoft Research)	<p><i>Hofmann</i></p>
05:00 PM	Woj Zaremba (OpenAI)	<p><i>Zaremba</i></p>
05:45 PM	Closing Remarks	<p><i>Mattar, Juliani, Crosby, Beyret, Lange</i></p>

Sets and Partitions

Nicholas Monath, Manzil Zaheer, Andrew McCallum, Ari Kobren, Junier Oliva, Barnabas Póczos, Ruslan Salakhutdinov

West 215 + 216, Sat Dec 14, 08:00 AM

Classic problems for which the input and/or output is set-valued are ubiquitous in machine learning.

For example, multi-instance learning, estimating population statistics, and point cloud classification are all problem domains in which the input is set-valued. In multi-label classification the output is a set of labels, and in clustering, the output is a partition. New tasks that take sets as input are also rapidly emerging in a variety of application areas including: high energy physics, cosmology, crystallography, and art. As a natural means of succinctly capturing large collections of items, techniques for learning representations of sets and partitions have significant potential to enhance scalability, capture complex dependencies, and improve interpretability. The importance and potential of improved set processing has led to recent work on permutation invariant and equivariant representations (Ravanbakhsh et al, 2016; Zaheer et al, 2017; Ilse et al, 2018; Hartford et al, 2018; Lee et al, 2019, Cotter et al, 2019, Bloom-Reddy & Teh, 2019, and more) and continuous representations of set-based outputs and partitions (Tai and Lin, 2012; Belanger & McCallum, 2015; Wiseman et al, 2016; Caron et al, 2018; Zhang et al, 2019; Vikram et al 2019).

The goal of this workshop is to explore:

- Permutation invariant and equivariant representations; empirical performance, limitations, implications, inductive biases of proposed representations of sets and partitions, as well as rich models of interaction among set elements;
- Inference methods for predicting sets or clusterings; approaches based on gradient-descent, continuous representations, amenable to end-to-end optimization with other models;
- New applications of set and partition-based models.

The First Workshop on Sets and Partitions, to be held as a part of the NeurIPS 2019 conference, focuses on models for tasks with set-based inputs/outputs as well as models of partitions and novel clustering methodology. The workshop

welcomes both methodological and theoretical contributions, and also new applications. Connections to related problems in optimization, algorithms, theory as well as investigations of learning approaches to set/partition problems are also highly relevant to the workshop. We invite both paper submissions and submissions of open problems. We hope that the workshops will inspire further progress in this important field.

Organizing Committee:

Andrew McCallum, UMass Amherst
 Ruslan Salakhutdinov, CMU
 Barnabas Poczos, CMU
 Junier Oliva, UNC Chapel Hill
 Manzil Zaheer, Google Research
 Ari Kobren, UMass Amherst
 Nicholas Monath, UMass Amherst
 with senior advisory support from Alex Smola.

Invited Speakers:

Siamak Ravanbakhsh
 Abhishek Khetan
 Eunsu Kang
 Amr Ahmed
 Stefanie Jegelka

Schedule

08:45 AM	Opening Remarks	<i>Zaheer, Monath, Kobren, Oliva, Poczos, Salakhutdinov, McCallum</i>
09:00 AM	Invited Talk - Siamak Ravanbakhsh	<i>Ravanbakhsh</i>

09:45 AM	Coffee Break & Poster Session 1	<i>Zhang, Hare, Prugel-Bennett, Leung, Flaherty, Wiratchotisatian, Epasto, Lattanzi, Vassilvitskii, Zadimoghaddam, Tulabandhula, Fuchs, Kosiorek, Posner, Hang, Goldie, Ravi, Mirhoseini, Xiong, Ren, Liao, Urtasun, Zhang, Borassi, Luo, Trapp, Dubourg-Felonneau, Kussad, Bender, Zaheer, Oliva, Stypu█kowski, Zieba, Dill, Li, Ge, Kang, Parker Jones, Wong, Payne, Li, Nazi, Erdem, Erdem, O'Connor, Garcia, Zamorski, Chorowski, Sinha, Clifford, Cassidy</i>
10:30 AM	Invited Talk - Stefanie Jegelka	<i>Jegelka</i>
11:15 AM	Contributed Talk - Towards deep amortized clustering	<i>Lee, Lee, Teh</i>
11:30 AM	Contributed Talk - Fair Hierarchical Clustering	<i>Ahmadian, Epasto, Knittel, Kumar, Mahdian, Pham</i>
11:45 AM	Invited Talk - Abhishek Khetan	<i>Khetan</i>
02:00 PM	Contributed Talk - Limitations of Deep Learning on Point Clouds	<i>Bueno</i>

02:15 PM	Contributed Talk - Chirality Nets: Exploiting Structure in Human Pose Regression	<i>Yeh, Hu, Schwing</i>
02:30 PM	Invited Talk - Eunsu Kang	<i>Kang</i>
03:15 PM	Coffee Break & Poster Session 2	<i>Lee, Lee, Teh, Yeh, Hu, Schwing, Ahmadian, Epasto, Knittel, Kumar, Mahdian, Bueno, Sanghi, Jayaraman, Arroyo-Fernández, Hryniowski, Mathur, Singh, Haddadan, Portilheiro, Zhang, Yuksekgonul, Arias Figueroa, Maurya, Ravindran, NIELSEN, Pham, Payan, McCallum, Mehta, Sun</i>
04:15 PM	Invited Talk - Amr Ahmed	<i>Ahmed</i>
05:00 PM	Panel Discussion	
05:40 PM	Closing Remarks	

Abstracts (3):

Abstract 3: Coffee Break & Poster Session 1 in Sets and Partitions, *Zhang, Hare, Prugel-Bennett, Leung, Flaherty, Wiratchotisatian, Epasto, Lattanzi, Vassilvitskii, Zadimoghaddam, Tulabandhula, Fuchs, Kosiorek, Posner, Hang, Goldie, Ravi, Mirhoseini, Xiong, Ren, Liao, Urtasun, Zhang, Borassi, Luo, Trapp, Dubourg-Felonneau, Kussad, Bender, Zaheer, Oliva, Stypankowski, Zieba, Dill,*

Li, Ge, Kang, Parker Jones, Wong, Payne, Li, Nazi, Erdem, Erdem, O'Connor, Garcia, Zamorski, Chorowski, Sinha, Clifford, Cassidy 09:45 AM

Poster Session 1 Paper Titles & Authors:

Deep Set Prediction Networks. Yan Zhang, Jonathon Hare, Adam Prügel-Bennett

Deep Hyperedges: a Framework for Transductive and Inductive Learning on Hypergraphs. Joshua Payne

FSPool: Learning Set Representations with Featurewise Sort Pooling. Yan Zhang, Jonathon Hare, Adam Prügel-Bennett

Deep Learning Features Through Dictionary Learning with Improved Clustering for Image Classification. Shengda Luo, Alex Po Leung, Haici Zhang

Globally Optimal Model-based Clustering via Mixed Integer Nonlinear Programming. Patrick Flaherty, Pitchaya Wiratchotisatian, Andrew C. Trapp

Sliding Window Algorithms for k-Clustering Problems. Michele Borassi, Alessandro Epasto, Silvio Lattanzi, Sergei Vassilvitski, Morteza Zadimoghaddam

Optimized Recommendations When Customers Select Multiple Products. Praseem Patidar, Deeksha Sinha, Theja Tulabandhula

Manipulating Person Videos with Natural Language. Levent Karacan, Mehmet Gunel, Aykut Erdem, Erkut Erdem

Permutation Invariance and Relational Reasoning in Multi-Object Tracking. Fabian B. Fuchs, Adam R. Kosiorek, Li Sun, Oiwi Parker Jones, Ingmar Posner.

Clustering by Learning to Optimize Normalized Cuts. Azade Nazi, Will Hang, Anna Goldie, Sujith Ravi, Azalia Mirhoseini

Deformable Filter Convolution for Point Cloud Reasoning. Yuwen Xiong, Mengye Ren, Renjie Liao, Kelvin Wong, Raquel Urtasun

Learning Embeddings from Cancer Mutation Sets for Classification Tasks. Geoffroy Dubourg-Felonneau, Yasmeeen Kussad, Dominic Kirkham, John Cassidy, Harry W Clifford

Exchangeable Generative Models with Flow Scans. Christopher M. Bender, Kevin O'Connor, Yang Li, Juan Jose Garcia, Manzil Zaheer, Junier Oliva

Conditional Invertible Flow for Point Cloud Generation. Stypulkowski Michal, Zamorski Maciej, Zieba Maciej, Chorowski Jan

Getting Topology and Point Cloud Generation to Mesh. Austin Dill, Chun-Liang Li, Songwei Ge, Eunsu Kang

Distributed Balanced Partitioning and Applications in Large-scale Load Balancing. Aaron Archer, Kevin Aydin, MohammadHossein Bateni, Vahab Mirrokni, Aaron Schild, Ray Yang, Richard Zhuang

Abstract 8: **Contributed Talk - Limitations of Deep Learning on Point Clouds in Sets and Partitions**, *Bueno* 02:00 PM

Limitations of Deep Learning on Point Clouds
Christian Bueno, Alan G. Hylton

Abstract 11: **Coffee Break & Poster Session 2 in Sets and Partitions**, *Lee, Lee, Teh, Yeh, Hu, Schwing, Ahmadian, Epasto, Knittel, Kumar, Mahdian, Bueno, Sanghi, Jayaraman, Arroyo-Fernández, Hryniowski, Mathur, Singh, Haddadan, Portilheiro, Zhang, Yuksekgonul, Arias*

Figueroa, Maurya, Ravindran, NIELSEN, Pham, Payan, McCallum, Mehta, Sun 03:15 PM

Poster Session 2 Paper Titles & Authors:

Towards deep amortized clustering. Juho Lee, Yoonho Lee, Yee Whye Teh

Chirality Nets: Exploiting Structure in Human Pose Regression. Raymond Yeh, Yuan-Ting Hu, Alexander Schwing

Fair Hierarchical Clustering. Sara Ahmadian, Alessandro Epasto, Marina Knittel, Ravi Kumar, Mohammad Mahdian, Philip Pham

Limitations of Deep Learning on Point Clouds.
Christian Bueno, Alan G. Hylton

How Powerful Are Randomly Initialized Pointcloud Set Functions? Aditya Sanghi, Pradeep Kumar Jayaraman

On the Possibility of Rewarding Structure Learning Agents: Mutual Information on Linguistic Random Sets. Ignacio Arroyo-Fernández, Mauricio Carrasco-Ruiz, José Anibal Arias-Aguilar

Modelling Convolution as a Finite Set of Operations Through Transformation Semigroup Theory.
Andrew Hryniowski, Alexander Wong

HCA-DBSCAN: HyperCube Accelerated Density Based Spatial Clustering for Applications with Noise. Vinayak Mathur, Jinesh Mehta, Sanjay Singh

Finding densest subgraph in probabilistically evolving graphs. Sara Ahmadian, Shahrzad Haddadan

Representation Learning with Multisets. Vasco Portilheiro

PairNets: Novel Fast Shallow Artificial Neural Networks on Partitioned Subspaces. Luna Zhang

Fair Correlation Clustering. Sara Ahmadian, Alessandro Epasto, Ravi Kumar, Mohammad Mahdian

Learning Maximally Predictive Prototypes in Multiple Instance Learning. Mert Yuksekogul, Ozgur Emre Sivrikaya, Mustafa Gokce Baydogan

Deep Clustering using MMD Variational Autoencoder and Traditional Clustering Algorithms. Jhosimar Arias

Hypergraph Partitioning using Tensor Eigenvalue Decomposition. Deepak Maurya, Balaraman Ravindran, Shankar Narasimhan

Information Geometric Set Embeddings: From Sets to Distributions. Ke Sun, Frank Nielsen

Document Representations using Fine-Grained Topics. Justin Payan, Andrew McCallum

Context and Compositionality in Biological and Artificial Neural Systems

Javier Turek, Shailee Jain, Alexander Huth, Leila Wehbe, Emma Strubell, Alan Yuille, Tal Linzen, Christopher Honey, Kyunghyun Cho

West 217 - 219, Sat Dec 14, 08:00 AM

The ability to integrate semantic information across narratives is fundamental to language understanding in both biological and artificial cognitive systems. In recent years, enormous strides have been made in NLP and Machine Learning to develop architectures and techniques that effectively capture these effects. The field has moved away from traditional bag-of-words approaches that ignore temporal ordering, and

instead embraced RNNs, Temporal CNNs and Transformers, which incorporate contextual information at varying timescales. While these architectures have led to state-of-the-art performance on many difficult language understanding tasks, it is unclear what representations these networks learn and how exactly they incorporate context. Interpreting these networks, systematically analyzing the advantages and disadvantages of different elements, such as gating or attention, and reflecting on the capacity of the networks across various timescales are open and important questions.

On the biological side, recent work in neuroscience suggests that areas in the brain are organized into a temporal hierarchy in which different areas are not only sensitive to specific semantic information but also to the composition of information at different timescales. Computational neuroscience has moved in the direction of leveraging deep learning to gain insights about the brain. By answering questions on the underlying mechanisms and representational interpretability of these artificial networks, we can also expand our understanding of temporal hierarchies, memory, and capacity effects in the brain.

In this workshop we aim to bring together researchers from machine learning, NLP, and neuroscience to explore and discuss how computational models should effectively capture the multi-timescale, context-dependent effects that seem essential for processes such as language understanding.

We invite you to submit papers related to the following (non-exhaustive) topics:

- * Contextual sequence processing in the human brain
- * Compositional representations in the human brain
- * Systematic generalization in deep learning
- * Compositionality in human intelligence

- * Compositionality in natural language
- * Understanding composition and temporal processing in neural network models
- * New approaches to compositionality and temporal processing in language
- * Hierarchical representations of temporal information
- * Datasets for contextual sequence processing
- * Applications of compositional neural networks to real-world problems

Submissions should be up to 4 pages excluding references, and should be NIPS format and anonymous. The review process is double-blind.

We also welcome published papers that are within the scope of the workshop (without re-formatting). This specific papers do not have to be anonymous. They will only have a very light review process.

Schedule

08:00 AM	Opening Remarks	<i>Huth</i>
08:15 AM	Patricia Churchland	<i>Churchland</i>
09:00 AM	Gina Kuperberg	<i>Kuperberg</i>
09:45 AM	Poster Session + Break	
10:30 AM	Spotlights - TBA	
11:00 AM	Tom Mitchell	<i>Mitchell</i>

		<i>Nye, Kim, St Clere Smithe, Itoh, Florez, G. Djokic, Aenugu, Toneva, Schlag, Schwartz, Sobroza Marques, Sainath, Li, Bommasani, Kim, Soulos, Frankland, Chirkova, Han, Kortylewski, Pang, Rabovsky, Mamou, Kumar, Marra</i>
12:00 PM	Poster Session + Lunch	
01:45 PM	Liina Pylkkanen	<i>Pylkkanen</i>
02:30 PM	Yoshua Bengio - Towards compositional understanding of the world by agent-based deep learning	<i>Bengio</i>
03:30 PM	Poster Session + Break	
04:15 PM	Panel	<i>Willke, Fedorenko, Lee, Smolensky, Marcus</i>
05:55 PM	Closing remarks	<i>Wehbe</i>

Abstracts (7):

Abstract 1: **Opening Remarks in Context and Compositionality in Biological and Artificial Neural Systems**, *Huth* 08:00 AM

Note: schedule not final and may change

Abstract 2: **Patricia Churchland in Context and Compositionality in Biological and Artificial Neural Systems**, *Churchland* 08:15 AM

Note: schedule not final and may change

Abstract 3: **Gina Kuperberg in Context and Compositionality in Biological and Artificial Neural Systems**, *Kuperberg* 09:00 AM

Note: schedule not final and may change

Abstract 6: **Tom Mitchell in Context and Compositionality in Biological and Artificial Neural Systems**, *Mitchell* 11:00 AM

Note: schedule not final and may change

Abstract 8: **Liina Pylikkanen in Context and Compositionality in Biological and Artificial Neural Systems**, *Pylikkanen* 01:45 PM

Note: schedule not final and may change

Abstract 9: **Yoshua Bengio - Towards compositional understanding of the world by agent-based deep learning in Context and Compositionality in Biological and Artificial Neural Systems**, *Bengio* 02:30 PM

Note: schedule not final and may change

Abstract 11: **Panel in Context and Compositionality in Biological and Artificial Neural Systems**, *Willke, Fedorenko, Lee, Smolensky, Marcus* 04:15 PM

Note: schedule not final and may change

Robot Learning: Control and Interaction in the Real World

Roberto Calandra, Kate Rakelly, Sanket Sayaji Kamthe, Danica Kragic, Stefan Schaal, Markus Wulfmeier

West 220 - 222, Sat Dec 14, 08:00 AM

The growing capabilities of learning-based methods in control and robotics has precipitated a shift in the design of software for autonomous systems. Recent successes fuel the hope that robots will increasingly perform varying tasks working alongside humans in complex, dynamic environments. However, the application of learning approaches to real-world robotic systems has been limited because real-world scenarios introduce challenges that do not arise in simulation.

In this workshop, we aim to identify and tackle the main challenges to learning on real robotic systems. First, most machine learning methods rely on large quantities of labeled data. While raw sensor data is available at high rates, the required variety is hard to obtain and the human effort to annotate or design reward functions is an even larger burden. Second, algorithms must guarantee some measure of safety and robustness to be deployed in real systems that interact with property and people. Instantaneous reset mechanisms, as common in simulation to recover from even critical failures, present a great challenge to real robots. Third, the real world is significantly more complex and varied than curated datasets and simulations. Successful approaches must scale to this complexity and be able to adapt to novel situations.

Schedule

09:10 AM	Invited Talk 1	<i>Deisenroth</i>
09:45 AM	Coffee Break	
11:30 AM	Poster 1	<i>Genc, Clavera, Gilaberte, Zimmer, Smith, Xiao, Fu, Ding, Stepputtis, Mallya, Bodapati, Lin</i>
12:00 PM	Lunch Break	

01:30 PM	Invited Talk 2	<i>Veloso</i>
01:30 PM	Invited Talk 3	<i>Osa</i>
02:00 PM	Invited Talk 4	<i>Fazeli</i>
02:30 PM	Invited Talk 5	<i>Schoellig</i>
03:00 PM	Poster Session 2	<i>Tangkaratt, Nair, Di Palo, Yang, Yang, Florensa, Lee, Church, Han, Qi, Zhang, Pan</i>
03:30 PM	Coffee Break	
04:30 PM	Invited Talk 6	<i>Johns</i>
05:00 PM	Panel	

NeurIPS Workshop on Machine Learning for Creativity and Design 3.0

Luba Elliott, Sander Dieleman, Adam Roberts, Jesse Engel, Tom White, Rebecca Fiebrink, Parag Mital, Christine Payne, Nao Tokui

West 223 + 224, Sat Dec 14, 08:00 AM

Generative machine learning and machine creativity have continued to grow and attract a wider audience to machine learning. Generative models enable new types of media creation across images, music, and text - including recent advances such as StyleGAN, MuseNet and GPT-2. This one-day workshop broadly explores issues in the applications of machine learning to creativity and design. We will look at algorithms for generation and creation of new media, engaging researchers

building the next generation of generative models (GANs, RL, etc). We investigate the social and cultural impact of these new models, engaging researchers from HCI/UX communities and those using machine learning to develop new creative tools. In addition to covering the technical advances, we also address the ethical concerns ranging from the use of biased datasets to the use of synthetic media such as “DeepFakes”. Finally, we’ll hear from some of the artists and musicians who are adopting machine learning including deep learning and reinforcement learning as part of their own artistic process. We aim to balance the technical issues and challenges of applying the latest generative models to creativity and design with philosophical and cultural issues that surround this area of research.

Schedule

08:15 AM	Welcome and Introduction	
08:30 AM	Alec Radford	<i>Radford</i>
09:00 AM	Giorgio Patrini	<i>Patrini</i>
09:30 AM	AI Art Gallery Overview	<i>Elliott</i>
10:30 AM	Yann LeCun	<i>LeCun</i>
11:00 AM	Neural Painters: A learned differentiable constraint for generating brushstroke paintings	<i>Nakano</i>

11:10 AM	Transform the Set: Memory Attentive Generation of Guided and Unguided Image Collages	<i>Jetchev, Vollgraf</i>
11:20 AM	Paper Dreams: An Interactive Interface for Generative Visual Expression	<i>Bernal, Zhou</i>
11:30 AM	Deep reinforcement learning for 2D soft body locomotion	<i>Rojas</i>
11:40 AM	Towards Sustainable Architecture: 3D Convolutional Neural Networks for Computational Fluid Dynamics Simulation and Reverse Design Workflow	<i>Musil</i>
11:50 AM	Human and GAN collaboration to create haute couture dress	<i>Seita, Koga</i>
01:30 PM	Poster Session 1	<i>Lee, Saeed, Broad, Gillick, Hertzmann, Aggarwal, Sung, Champandard, Park, Mellor, Herrmann, Wu, Lee, Jieun, Han, jung, Kim</i>
02:30 PM	Sougwen Chung	<i>Chung</i>
03:45 PM	Claire Evans	<i>Evans, Bechtolt, Kieswetter</i>

04:15 PM	MidiMe: Personalizing a MusicVAE model with user data	<i>Dinculescu</i>
04:25 PM	First Steps Towards Collaborative Poetry Generation	<i>Uthus, Voitovich</i>
04:35 PM	Panel Discussion	<i>Saxena, Frosst, Cabannes, Kogan, Dill, Sarkar, Moniz, Thio, Sievert, Coleman, De Bleser, Quanz, Kereliuk, Achlioptas, Elhoseiny, Ge, Gomez, Brew</i>
05:00 PM	Poster Session 2	<i>Sarin, Bourached, Carr, Zukowski, Zhou, Malakhova, Petric, Lorenzo, O'Brien, Wegner, Kishi, Burnam</i>
05:05 PM	Artwork	



Medical Imaging meets NeurIPS

Hervé Lombaert, Ben Glocker, Ender Konukoglu, Marleen de Bruijne, Aasa Feragen, Ipek Oguz, Jonas Teuwen

West 301 - 305, Sat Dec 14, 08:00 AM

Medical imaging and radiology are facing a major crisis with an ever-increasing complexity and volume of data along an immense economic pressure. The current advances and widespread use of imaging technologies now overload the human capacity of interpreting medical images, dangerously posing a risk of missing critical patterns

of diseases. Machine learning has emerged as a key technology for developing novel tools in computer aided diagnosis, therapy and intervention. Still, progress is slow compared to other fields of visual recognition, which is mainly due to the domain complexity and constraints in clinical applications, i.e., robustness, high accuracy and reliability.

“Medical Imaging meets NeurIPS” aims to bring researchers together from the medical imaging and machine learning communities to discuss the major challenges in the field and opportunities for research and novel applications. The proposed event will be the continuation of a successful workshop organized in NeurIPS 2017 and 2018 (<https://sites.google.com/view/med-nips-2018>). It will feature a series of invited speakers from academia, medical sciences and industry to give an overview of recent technological advances and remaining major challenges.

Schedule

08:15 AM	Opening Remarks	<i>Lombaert, Glocker, Konukoglu, de Bruijne, Feragen, Oguz, Teuwen</i>
<hr/>		
08:30 AM	Session 1 (Invited Talk + presentations)	<i>Schnabel, Vidal, Sodickson, Grady, Vidal</i>

10:30 AM	Coffee Break + Poster Session	<i>Glocker, Manakov, Taleb, Abdi, Zhang, Khalvati, Sajda, Sinha, Preuhs, Stern, Mohan, Olatunji, Romero, Razdaibiedina, Weng, Kohl, Zimmerer, Perkonigg, Hofmanninger, Vazquez Romaguera, Chakravarty, Weiss, Gossmann, CHEN, Yao, Albarqouni, Sahoo, Martinez Manzanera, Pinckaers, Dalmis, Trivedi, Gamper, Gong, Haq, Hallgrimsson, Yang, Jeon, De Goyeneche Macaya, Wang, Wang, Garcia, Sidorov, Kames, Soni, Patra, Dutt, Roth, Tamir, Imran, OOTA, Schobs, Eitel, Jin, Fadnavis, Shinde, Ruhe, Paetzold, Kaku, Douglas, Heckel</i>
<hr/>		
11:00 AM	Session 2 (Invited Talk + presentations)	
<hr/>		
12:30 PM	Lunch	
<hr/>		
01:30 PM	Session 3 (Invited Talk + presentations)	

		<i>Liu, Sharan, Abolmaesumi, Parmar, Lei, Gavves, Nabi, Namdar, Chen, Modi, Fels, Rauscher, Li, Chung, Oktay, Gopinath, Selvan, Adiga Vasudeva, Poblenz, Baltatzis, Wei, Velayutham, Garyfallidis, Ellis, Bhatia, Galitz, Muckley, Cai, Prasanna</i>
03:00 PM	Coffee Break + Poster Session	
03:30 PM	Session 4 (Invited Talk + presentations)	
05:00 PM	fastMRI Challenge Talks	<i>Yakubova, Pezzotti, Wang, Zitnick, Karkalousos, Sun, Caan, Murrell</i>
06:00 PM	Closing Remarks	

Abstracts (1):

Abstract 9: fastMRI Challenge Talks in Medical Imaging meets NeurIPS, *Yakubova, Pezzotti, Wang, Zitnick, Karkalousos, Sun, Caan, Murrell*

Yakubova, Pezzotti, Wang, Zitnick, Karkalousos, Sun, Caan, Murrell

05:00 PM

tentative

Learning with Temporal Point Processes

Manuel Rodriguez, Le Song, Isabel Valera, Yan Liu, Abir De, Hongyuan Zha

West 306, Sat Dec 14, 08:00 AM

In recent years, there has been an increasing number of machine learning models and algorithms based on the theory of temporal point processes, which is a mathematical framework to model asynchronous event data. These models and algorithm have found a wide range of human-centered applications, from social and information networks and recommender systems to crime prediction and health. Moreover, this emerging line of research has already established connections to deep learning, deep generative models, Bayesian nonparametrics, causal inference, stochastic optimal control and reinforcement learning. However, despite these recent advances, learning with temporal point processes is still a relatively niche topic within the machine learning community---there are only a few research groups across the world with the necessary expertise to make progress. In this workshop, we aim to popularize temporal point processes within the machine learning community at large. In our view, this is the right time to organize such a workshop because, as algorithmic decisions becomes more consequential to individuals and society, temporal point processes will play a major role on the development of human-centered machine learning models and algorithms accounting for the feedback loop between algorithmic and human decisions, which are inherently asynchronous events. Moreover, it will be a natural follow up of a very successful and well-attended ICML 2018 tutorial on learning with temporal point processes, which two of us recently taught.

Schedule

08:30 AM	Welcome Address and Introduction	
08:35 AM	Invited Talk by Negar Kiyavash	<i>Kiyavash</i>

09:15 AM	Fused Gromov-Wasserstein Alignment for Hawkes Processes	
09:30 AM	Insider Threat Detection via Hierarchical Neural Temporal Point Processes	<i>Wu</i>
09:45 AM	Coffee Break	
10:30 AM	Invited Talk By Niloy Ganguly	<i>ganguly</i>
11:10 AM	Intermittent Demand Forecasting with Deep Renewal Processes	<i>Turkmen</i>
11:25 AM	Temporal Logic Point Processes	
11:40 AM	The Graph Hawkes Network for Reasoning on Temporal Knowledge Graphs	
11:55 AM	Multivariate coupling estimation between continuous signals and point processes	<i>Besserve</i>
12:10 PM	Lunch Break	
01:50 PM	Invited Talk by Walter Dempsey	<i>Dempsey</i>
02:30 PM	Better Approximate Inference for Partial Likelihood Models with a Latent Structure	<i>Setlur</i>

02:45 PM	Deep Point Process Destructors	
03:00 PM	A sleep-wake detection algorithm for memory-constrained wearable devices: Change Point Decoder	<i>Cakmak</i>
03:15 PM	Topics are not Marks: Modeling Text-based Cascades using Multi-network Hawkes Process	<i>Choudhari</i>
03:30 PM	Poster Setup + Coffee Break	
04:15 PM	Temporal point process models vs. discrete time models	
05:00 PM	Poster Session	<i>Cakmak, Zhang, Prabhakarannair, Kusumam, Ahmed, Wu, Choudhari, Inouye, Taylor, Besserve, Turkmen, Islam, Artés, Setlur, Fu, Han, De, Du, Sanchez Martin</i>



The Optimization Foundations of Reinforcement Learning

Bo Dai, Niao He, Nicolas Le Roux, Lihong Li, Dale Schuurmans, Martha White

West Ballroom A, Sat Dec 14, 08:00 AM

Interest in reinforcement learning (RL) has boomed with recent improvements in benchmark tasks that suggest the potential for a revolutionary advance in practical applications. Unfortunately, research in RL remains hampered by limited theoretical understanding, making the field overly reliant on empirical exploration with insufficient principles to guide future development. It is imperative to develop a stronger fundamental understanding of the success of recent RL methods, both to expand the useability of the methods and accelerate future deployment. Recently, fundamental concepts from optimization and control theory have provided a fresh perspective that has led to the development of sound RL algorithms with provable efficiency. The goal of this workshop is to catalyze the growing synergy between RL and optimization research, promoting a rational reconsideration of the foundational principles for reinforcement learning, and bridging the gap between theory and practice.

Schedule

08:00 AM	Opening Remarks	<i>Dai, He, Le Roux, Li, Schuurmans, White</i>
08:10 AM	Plenary Talk	<i>Agrawal</i>
08:50 AM	Contributed Talk	
09:10 AM	Poster Spotlight	

09:30 AM **Poster and Coffee Break**

Sidford, Mahajan, Ribeiro, Lewandowski, Sayed, Tewari, Steger, Anandkumar, Mujika, Kappen, Zhou, Boots, Finn, Wei, Jin, Cheng, Yu, Gehring, Boutilier, Lin, McNamee, Russo, Brandfonbrener, Zhou, Jha, Romeres, Precup, Thalmeier, Gorbunov, Hazan, Smirnova, Dohmatob, Brunskill, Munoz de Cote, Waldie, Meier, Schaefer, Liu, Neu, Kaplan, Sun, Yao, Bhandari, Preiss, Subramanian, LI, Ye, Smith, Bas Serrano, Bruna, Langford, Lee, Arjona-Medina, Zhang, Singh, Luo, Ahmed, Chen, Wang, Li, Yang, Xu, Tang, Mao, Brandfonbrener, Di-Castro Shashua, Islam, Fu, Naik, Kumar, Petit, Kamoutsi, Totaro, Raghunathan, Wu, Lee, Ding, Koppel, Sun, Tjandraatmadja, Karami, Mei, Xiao, Wen, Zhang, Goroshin, Pezeshki, Zhai, Amortila, Huang, Vasileva, Bergou, Ahmadyan, Sun, Zhang, Gruber

10:30 AM	Plenary Talk	<i>Kakade</i>
11:10 AM	Contributed Talk	
11:30 AM	Poster Spotlight	<i>Li</i>
02:00 PM	Plenary Talk	<i>Van Roy</i>
02:40 PM	Plenary Talk	<i>Wang</i>

		<i>Hausman, Dong, Goldberg, Li, Yang, Wang, Shani, Wang, Amdahl-Culleton, Cassano, Dymetman, Bellemare, Tomczak, Castro, Kloft, Dinu, Holzleitner, White, Wang, Jordan, Jovanovic, Yu, Chen, Ryu, Zaheer, Agarwal, Jiang, He, Yasui, Karampatziakis, Vieillard, Nachum, Pietquin, Sener, Xu, Kamalaruban, Mineiro, Rolland, Amortila, Bacon, Panangaden, Cai, Liu, Gu, Seraj, Sutton, Valenzano, Dadashi, Toro Icarte, Shariff, Fox, Wang, Ghadimi, Sokota, Sinclair, Hochreiter, Levine, Valcarcel Macua, Kakade, Zhang, McIlraith, Mannor, Whiteson, Li, Qiu, Li, Banerjee, Luan, Basar, Doan, Yu, Liu, Zahavy, Klassen, Zhao, Gómez, Liu, Cevher, Suttle, Chang, Wei, Liu, Li, Chen, Song, Liu, Jiang, Feng, Du, Chow, Ye, Mansour, , Efroni, Chen, Wang, Dai, Wei, Shrivastava, Zhang, Zheng, SATPATHI, Liu, Vall</i>
03:20 PM	Poster and Coffee Break	

04:20 PM	Plenary Talk	<i>Yu</i>
05:00 PM	Contributed Talk	
05:20 PM	Panel Discussion	<i>Van Roy, Yu, Wang, Kakade, Agrawal</i>
05:45 PM	Closing Remarks	<i>Dai, He, Le Roux, Li, Schuurmans, White</i>

Abstracts (7):

Abstract 2: **Plenary Talk in The Optimization Foundations of Reinforcement Learning**, *Agrawal* 08:10 AM

TBA

Abstract 6: **Plenary Talk in The Optimization Foundations of Reinforcement Learning**, *Kakade* 10:30 AM

TBA

Abstract 9: **Plenary Talk in The Optimization Foundations of Reinforcement Learning**, *Van Roy* 02:00 PM

TBA

Abstract 10: **Plenary Talk in The Optimization Foundations of Reinforcement Learning**, *Wang* 02:40 PM

TBA

Abstract 12: **Plenary Talk in The Optimization Foundations of Reinforcement Learning**, *Yu* 04:20 PM

TBA

Abstract 14: **Panel Discussion in The Optimization Foundations of Reinforcement**

Learning, *Van Roy, Yu, Wang, Kakade, Agrawal* 05:20 PM

TBA

Abstract 15: **Closing Remarks in The Optimization Foundations of Reinforcement Learning**, *Dai, He, Le Roux, Li, Schuurmans, White* 05:45 PM

Awards Announcement

Machine Learning with Guarantees

Ben London, Gintare Karolina Dziugaite, Dan Roy, Thorsten Joachims, Aleksander Madry, John Shawe-Taylor

West Ballroom B, Sat Dec 14, 08:00 AM

As adoption of machine learning grows in high-stakes application areas (e.g., industry, government and health care), so does the need for guarantees: how accurate a learned model will be; whether its predictions will be fair; whether it will divulge information about individuals; or whether it is vulnerable to adversarial attacks. Many of these questions involve unknown or intractable quantities (e.g., risk, regret or posterior likelihood) and complex constraints (e.g., differential privacy, fairness, and adversarial robustness). Thus, learning algorithms are often designed to yield (and optimize) bounds on the quantities of interest. Beyond providing guarantees, these bounds also shed light on black-box machine learning systems.

Classical examples include structural risk minimization (Vapnik, 1991) and support vector machines (Cristianini & Shawe-Taylor, 2000), while more recent examples include non-vacuous risk bounds for neural networks (Dziugaite & Roy, 2017, 2018), algorithms that optimize both the weights and structure of a neural network (Cortes, 2017),

counterfactual risk minimization for learning from logged bandit feedback (Swaminathan & Joachims, 2015; London & Sandler, 2019), robustness to adversarial attacks (Schmidt et al., 2018; Wong & Kolter, 2018), differentially private learning (Dwork et al., 2006, Chaudhuri et al., 2011), and algorithms that ensure fairness (Dwork et al., 2012).

This one-day workshop will bring together researchers in both theoretical and applied machine learning, across areas such as statistical learning theory, adversarial learning, fairness and privacy, to discuss the problem of obtaining performance guarantees and algorithms to optimize them. The program will include invited and contributed talks, poster sessions and a panel discussion. We particularly welcome contributions describing fundamentally new problems, novel learning principles, creative bound optimization techniques, and empirical studies of theoretical findings.

Schedule

08:45 AM	Welcome Address	<i>London</i>
09:00 AM	TBD	<i>Roth</i>
09:45 AM	TBD	

		<i>Marcu, Yang, Gourdeau, Zhu, Lykouris, Chi, Kozdoba, Bhagoji, Wu, Nandy, Smith, Wen, Xie, Pitas, Shit, Andriushchenko, Yu, Letarte, Khodak, Mozannar, Podimata, Foulds, Wang, Zhang, Kuzelka, Levine, Lu, Mhammedi, Viillard, Cai, Gondara, Lucas, Mahdaviyeh, Baratin, Bommasani, Barp, Ilyas, Wu, Behrmann, Rivasplata, Nazemi, Raghunathan, Stephenson, Singla, Gupta, Choi, Kilcher, Lyle, Manino, Bennett, Xu, Chatterji, Barut, Prost, Toro Icarte, Blaas, Yun, Lale, Jiang, Medini, Rezaei, Meinke, Mell, Kazantsev, Garg, Sinha, Lokhande, Rizk, Zhao, Akash, Hou, Ghodsi, Hein, Sypherd, Yang, Pentina, Gillot, Ledent, Gur-Ari, MacAulay, Zhang</i>
10:15 AM	Break / Poster Session 1	
11:00 AM	Mehryar Mohri	<i>Mohri</i>
11:45 AM	TBD	
12:15 PM	Lunch Break	

02:00 PM	Po-Ling Loh	<i>Loh</i>
02:45 PM	TBD	
03:15 PM	Coffee Break / Poster Session 2	
04:00 PM	Emma Brünskill	<i>Brunskill</i>
04:45 PM	TBD	
05:15 PM	Discussion Panel	

Abstracts (2):

Abstract 4: **Break / Poster Session 1 in Machine Learning with Guarantees**, *Marcu, Yang, Gourdeau, Zhu, Lykouris, Chi, Kozdoba, Bhagoji, Wu, Nandy, Smith, Wen, Xie, Pitas, Shit, Andriushchenko, Yu, Letarte, Khodak, Mozannar, Podimata, Foulds, Wang, Zhang, Kuzelka, Levine, Lu, Mhammedi, Viillard, Cai, Gondara, Lucas, Mahdaviyeh, Baratin, Bommasani, Barp, Ilyas, Wu, Behrmann, Rivasplata, Nazemi, Raghunathan, Stephenson, Singla, Gupta, Choi, Kilcher, Lyle, Manino, Bennett, Xu, Chatterji, Barut, Prost, Toro Icarte, Blaas, Yun, Lale, Jiang, Medini, Rezaei, Meinke, Mell, Kazantsev, Garg, Sinha, Lokhande, Rizk, Zhao, Akash, Hou, Ghodsi, Hein, Sypherd, Yang, Pentina, Gillot, Ledent, Gur-Ari, MacAulay, Zhang* 10:15 AM

Presenters without NIPS accounts:
jjindong.gu@siemens.com

Abstract 10: **Coffee Break / Poster Session 2 in Machine Learning with Guarantees**, 03:15 PM

Same presenters as Poster Session 1

“Do the right thing”: machine learning and causal inference for improved decision making

Michele Santacatterina, Thorsten Joachims, Nathan Kallus, Adith Swaminathan, David Sontag, Angela Zhou

West Ballroom C, Sat Dec 14, 08:00 AM

In recent years, machine learning has seen important advances in its theoretical and practical domains, with some of the most significant applications in online marketing and commerce, personalized medicine, and data-driven policy-making. This dramatic success has led to increased expectations for autonomous systems to make the right decision at the right target at the right time. This gives rise to one of the major challenges of machine learning today that is the understanding of the cause-effect connection. Indeed, actions, intervention, and decisions have important consequences, and so, in seeking to make the best decision, one must understand the process of identifying causality. By embracing causal reasoning autonomous systems will be able to answer counterfactual questions, such as “What if I had treated a patient differently?”, and “What if had ranked a list differently?” thus helping to establish the evidence base for important decision-making processes.

The purpose of this workshop is to bring together experts from different fields to discuss the relationships between machine learning and causal inference and to discuss and highlight the formalization and algorithmization of causality toward achieving human-level machine intelligence.

This purpose will guide the makeup of the invited talks and the topics for the panel discussions. The panel discussions will tackle controversial topics, with the intent of drawing out an engaging intellectual debate and conversation across fields.

This workshop will lead to advance and extend knowledge on how machine learning could be used to conduct causal inference, and how causal inference could support the development of machine learning methods for improved decision-making.

Schedule

08:45 AM	Opening Remarks	<i>Joachims, Kallus, Santacatterina, Swaminathan, Sontag, Zhou</i>
09:00 AM	Susan Athey	<i>Athey</i>
09:30 AM	Andrea Rotnitzky	<i>Rotnitzky</i>
10:00 AM	Poster Spotlights	<i>Namkoong, Charpignon, Rudolph, Coston, Saito, Dhillon, Markham</i>
10:15 AM	Coffee break, posters, and 1-on-1 discussions	<i>Lu, Chen, Namkoong, Charpignon, Rudolph, Coston, von Kügelgen, Prasad, Dhillon, Xu, Wang, Markham, Rohde, Singh, Zhang, Hassanpour, Sharma, Lee, Pouget-Abadie, Krijthe, Mahajan, Ke, Wirnsberger, Semenova, Mykhaylov, Shen, Takatsu, Sun, Yang, Franks, Wong, Zaman, Mitchell, kang</i>

11:00 AM	Susan Murphy	<i>Murphy</i>
11:30 AM	Ying-Qi Zhao	<i>Zhao</i>
12:00 PM	Tentative topic: Reasoning about untestable assumptions in the face of unknowable counterfactuals	
12:45 PM	Lunch	
02:30 PM	Susan Shortreed	<i>Shortreed</i>
03:00 PM	Contributed talk 1	<i>Chen, Boehnke, Wang</i>
03:15 PM	Contributed talk 2	<i>Mahajan, Khosravi, D'Amour</i>
03:30 PM	Poster Spotlights	<i>Griveau-Billion, Singh, Zhang, Lee, Krijthe, Charles, Semenova, Ladhania, Opreescu</i>

03:45 PM	Coffee break, posters, and 1-on-1 discussions	<p><i>von Kügelgen, Rohde, Schumann, Charles, Veitch, Semenova, Demirer, Syrgkanis, Nair, Puli, Uehara, Gopalan, Ding, Ng, Khosravi, Sherman, Zeng, Wieczorek, Liu, Gan, Hartford, Oprescu, D'Amour, Boehnke, Saito, Griveau-Billion, Modi, Karimov, Berrevoets, Graham, Josse, Sridhar, Dahabreh, Mishler, Wadsworth, Qureshi, Ladhania, Morishita</i></p>
-------------	--	---

05:00 PM	Closing Remarks
-------------	------------------------

Abstracts (5):

Abstract 4: **Poster Spotlights in “Do the right thing”: machine learning and causal inference for improved decision making**, *Namkoong, Charpignon, Rudolph, Coston, Saito, Dhillon, Markham* 10:00 AM

Poster spotlights ID: 10, 11, 16, 17, 20, 24, 31

Abstract 8: **Tentative topic: Reasoning about untestable assumptions in the face of unknowable counterfactuals in “Do the right thing”: machine learning and causal inference for improved decision making**, 12:00 PM

Tentative topic: How machine learning, and causal inference work together: cross-pollination and new challenges.

Abstract 11: **Contributed talk 1 in “Do the right thing”: machine learning and causal inference**

for improved decision making, *Chen, Boehnke, Wang* 03:00 PM

Oral Spotlights ID: 8,9, 27

Abstract 12: **Contributed talk 2 in “Do the right thing”: machine learning and causal inference for improved decision making**, *Mahajan, Khosravi, D'Amour* 03:15 PM

Oral Spotlights ID: 57, 93, 113

Abstract 13: **Poster Spotlights in “Do the right thing”: machine learning and causal inference for improved decision making**, *Griveau-Billion, Singh, Zhang, Lee, Krijthe, Charles, Semenova, Ladhania, Oprescu* 03:30 PM

Poster Spotlights ID: 34, 35, 39, 50, 56, 68, 75, 111, 112

Bridging Game Theory and Deep Learning

Ioannis Mitliagkas, Gauthier Gidel, Niao He, Reyhane Askari Hemmat, Nika Haghtalab, Simon Lacoste-Julien

West Exhibition Hall A, Sat Dec 14, 08:00 AM

Advances in generative modeling and adversarial learning gave rise to a recent surge of interest in differentiable two-players games, with much of the attention falling on generative adversarial networks (GANs). Solving these games introduces distinct challenges compared to the standard minimization tasks that the machine learning (ML) community is used to. A symptom of this issue is ML and deep learning (DL) practitioners using optimization tools on game-theoretic problems. Our NeurIPS 2018 workshop, "Smooth games optimization in ML", aimed to rectify this situation, addressing theoretical aspects of games in machine learning, their special dynamics, and typical challenges. For this year, we significantly expand our scope to tackle questions

like the design of game formulations for other classes of ML problems, the integration of learning with game theory as well as their important applications. To that end, we have confirmed talks from Éva Tardos, David Balduzzi and Fei Fang. We will also solicit contributed posters and talks in the area.

Schedule

08:15 AM	Opening remarks	
08:30 AM	Invited talk: Eva Tardos (Cornell)	<i>Tardos</i>
09:00 AM	Morning poster Spotlight	
09:30 AM	Morning poster session -- coffee break	<i>Marchesi, Celli, Ohrimenko, Berard, Jain, Lin, Yan, McWilliams, Mishchenko, Çelikok, Abernethy, Liu, Fang, Li, Lee, Fridovich-Keil, Wang, Tsirigotis, Zhang, Lerer, Bondi, Jin, Fiez, Chasnov, Bennett, D'Orazio, Farina, Carmon, Mazumdar, Ibrahim, Zheng</i>
11:00 AM	Invited talk: David Balduzzi (DeepMind)	<i>Balduzzi</i>
11:30 AM	Contributed talk: What is Local Optimality in Nonconvex-Nonconcave Minimax Optimization?	<i>Jin</i>

12:00 PM	Contributed talk: Characterizing Equilibria in Stackelberg Games	<i>Fiez</i>
12:30 PM	Lunch break	
02:00 PM	Invited talk: Fei Fang (CMU)	<i>Fang</i>
02:30 PM	Contributed talk: On Solving Local Minimax Optimization: A Follow-the-Ridge Approach	<i>Wang</i>
03:00 PM	Contributed talk: Exploiting Uncertain Real-Time Information from Deep Learning in Signaling Games for Security and Sustainability	<i>Bondi</i>
03:30 PM	Coffee break	
04:00 PM	Invited talk: Asu Ozdaglar (MIT)	<i>Ozdaglar</i>

		<i>Netrapalli, Jordan, Gatti, Marchesi, Bianchi, Tschitschek, Almahairi, Vincent, Lacoste-Julien, Jain, Lin, Yang, Gemp, Koralev, Richtarik, Peltola, Kaski, Lai, Wibisono, Mroueh, Zhang, Cui, Das, Jelassi, Scieur, Mensch, Bruna, Li, Tang, Wakin, Eysenbach, Parisotto, Xing, Levine, Salakhutdinov, Tomlin, Zhang, Ba, Grosse, Hjelm, Courville, Hu, Foerster, Brown, Xu, Dilkina, Ratliff, Sastry, Chasnov, Fiez, Kallus, Schnabel, Kroer, Sandholm, Jin, Sidford, Tian, Zheng, Anandkumar</i>
04:30 PM	Afternoon poster spotlight	
05:00 PM	Discussion panel	
05:30 PM	Concluding remarks -- afternoon poster session	

Abstracts (4):

Abstract 6: Contributed talk: What is Local Optimality in Nonconvex-Nonconcave Minimax Optimization? in Bridging Game Theory and Deep Learning, Jin 11:30 AM

Minimax optimization has found extensive applications in modern machine learning, in settings such as generative adversarial networks (GANs), adversarial training and multi-agent reinforcement learning. As most of these applications involve continuous nonconvex-nonconcave formulations, a very basic question arises---`what is a proper definition of local optima?"

Most previous work answers this question using classical notions of equilibria from simultaneous games, where the min-player and the max-player act simultaneously. In contrast, most applications in machine learning, including GANs and adversarial training, correspond to sequential games, where the order of which player acts first is crucial (since minimax is in general not equal to maximin due to the nonconvex-nonconcave nature of the problems). The main contribution of this paper is to propose a proper mathematical definition of local optimality for this sequential setting---local minimax, as well as to present its properties and existence results. Finally, we establish a strong connection to a basic local search algorithm---gradient descent ascent (GDA): under mild conditions, all stable limit points of GDA are exactly local minimax points up to some degenerate points.

Abstract 7: Contributed talk: Characterizing Equilibria in Stackelberg Games in Bridging Game Theory and Deep Learning, Fiez 12:00 PM

This paper investigates the convergence of learning dynamics in Stackelberg games on continuous action spaces, a class of games distinguished by the hierarchical order of play between agents. We establish connections between the Nash and Stackelberg equilibrium concepts and characterize conditions under which attractors of simultaneous gradient descent are Stackelberg equilibria in zero-sum games. Moreover, we show that the only stable attractors of the Stackelberg gradient dynamics are Stackelberg equilibria in zero-sum games. Using this insight, we develop two-timescale learning dynamics that converge to Stackelberg

equilibria in zero-sum games and the set of stable attractors in general-sum games.

Abstract 10: Contributed talk: On Solving Local Minimax Optimization: A Follow-the-Ridge Approach in Bridging Game Theory and Deep Learning, *Wang* 02:30 PM

Many tasks in modern machine learning can be formulated as finding equilibria in \emph{sequential} games. In particular, two-player zero-sum sequential games, also known as minimax optimization, have received growing interest. It is tempting to apply gradient descent to solve minimax optimization given its popularity in supervised learning. However, we note that naive application of gradient descent fails to find local minimax -- the analogy of local minima in minimax optimization, since the fixed points of gradient dynamics might not be local minimax. In this paper, we propose \emph{Follow-the-Ridge} (FR), an algorithm that locally converges to and only converges to local minimax. We show theoretically that the algorithm addresses the limit cycling problem around fixed points, and is compatible with preconditioning and \emph{positive} momentum. Empirically, FR solves quadratic minimax problems and improves GAN training on simple tasks.

Abstract 11: Contributed talk: Exploiting Uncertain Real-Time Information from Deep Learning in Signaling Games for Security and Sustainability in Bridging Game Theory and Deep Learning, *Bondi* 03:00 PM

Motivated by real-world deployment of drones for conservation, this paper advances the state-of-the-art in security games with signaling. The well-known defender-attacker security games framework can help in planning for such strategic deployments of sensors and human patrollers, and warning signals to ward off adversaries. However, we show that defenders can suffer significant losses when ignoring real-world uncertainties, such as

detection uncertainty resulting from imperfect deep learning models, despite carefully planned security game strategies with signaling. In fact, defenders may perform worse than forgoing drones completely in this case. We address this shortcoming by proposing a novel game model that integrates signaling and sensor uncertainty; perhaps surprisingly, we show that defenders can still perform well via a signaling strategy that exploits the uncertain real-time information primarily from deep learning models. For example, even in the presence of uncertainty, the defender still has an informational advantage in knowing that she has or has not actually detected the attacker; and she can design a signaling scheme to ``mislead" the attacker who is uncertain as to whether he has been detected. We provide a novel algorithm, scale-up techniques, and experimental results from simulation based on our ongoing deployment of a conservation drone system in South Africa.

Deep Reinforcement Learning

Pieter Abbeel, Chelsea Finn, Joelle Pineau, David Silver, Satinder Singh, Joshua Achiam, Carlos Florensa, Christopher Grimm, Haoran Tang, Vivek Veeriah

West Exhibition Hall C, Sat Dec 14, 08:00 AM

In recent years, the use of deep neural networks as function approximators has enabled researchers to extend reinforcement learning techniques to solve increasingly complex control tasks. The emerging field of deep reinforcement learning has led to remarkable empirical results in rich and varied domains like robotics, strategy games, and multiagent interaction. This workshop will bring together researchers working at the intersection of deep learning and reinforcement learning, and it will help interested researchers outside of the field gain a high-level view about the current state of the art and potential directions for future contributions.

Schedule

08:45 AM	Welcome Comments	
09:00 AM	Invited Talk	<i>Vinyals</i>
09:30 AM	Contributed Talks	<i>Tang, Guo, Hafner</i>
10:00 AM	Invited Talk	<i>Whiteson</i>
10:30 AM	Coffee Break	
11:00 AM	Invited Talk	<i>Brunskill</i>
11:30 AM	Contributed Talks	<i>Lu, Hausknecht, Nachum</i>
12:00 PM	Invited Talk	<i>Fei-Fei</i>
01:30 PM	Invited Talk	<i>Todorov</i>
02:00 PM	Contributed Talks	<i>Agarwal, Gleave, Lee</i>

03:00 PM

Poster Session 1

Sabatelli, Stooke, Abdi, Rauber, Adolphs, Osband, Meisheri, Kurach, Ackermann, Benatan, ZHANG, Tessler, Shen, Samvelyan, Islam, Dalal, Harries, Kurenkov, ■o■na, Dasari, Hartikainen, Nachum, Lee, Holzleitner, Nguyen, Song, Grimm, Silva, Luo, Wu, Lee, Paine, Qu, Graves, Flet-Berliac, Tang, Nair, Hausknecht, Bagaria, Schmitt, Baker, Parmas, Eysenbach, Lee, Lin, Seita, Gupta, Simmons-Edler, Guo, Corder, Kumar, Fujimoto, Lerer, Clavera Gilaberte, Rhinehart, Nair, Yang, Wang, Sohn, Hernandez-Garcia, Lee, Srivastava, Khetarpal, Xiao, Carvalho Melo, Agarwal, Yu, Berseth, Chaplot, Tang, Srinivasan, Medini, Havens, Laskin, Mujika, Saphal, Marino, Ray, Achiam, Mandlekar, Liu, Hafner, Tang, Xiao, Walton, Druce, Alet, Hong, Chan, Nagabandi, Liu, Sun, Liu, Jayaraman, Co-Reyes, Sanborn

	NeurIPS RL	
04:00	Competitions	
PM	Results	
	Presentations	

05:00	Invited Talk	<i>Littman</i>
PM		

05:30	Panel Discussion	
PM		

06:00	Poster Session 2	
PM		

Abstracts (11):

Abstract 2: **Invited Talk in Deep Reinforcement Learning**, *Vinyals* 09:00 AM

(Talk title and abstract TBD.)

Abstract 3: **Contributed Talks in Deep Reinforcement Learning**, *Tang, Guo, Hafner* 09:30 AM

* "Playing Dota 2 with Large Scale Deep Reinforcement Learning" - OpenAI, Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dēbiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique Pondé de Oliveira Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, Susan Zhang

* "Efficient Exploration with Self-Imitation Learning via Trajectory-Conditioned Policy" - Yijie Guo, Jongwook Choi, Marcin Moczulski, Samy Bengio, Mohammad Norouzi, Honglak Lee

* "Efficient Visual Control by Latent Imagination" - Danijar Hafner, Timothy Lillicrap, Jimmy Ba, Mohammad Norouzi

Abstract 4: **Invited Talk in Deep Reinforcement Learning**, *Whiteson* 10:00 AM

(Speaker and details forthcoming.)

Abstract 6: **Invited Talk in Deep Reinforcement Learning**, *Brunskill* 11:00 AM

(Speaker and details forthcoming.)

Abstract 7: **Contributed Talks in Deep Reinforcement Learning**, *Lu, Hausknecht, Nachum* 11:30 AM

* "Adaptive Online Planning for Lifelong Reinforcement Learning" - Kevin Lu, Igor Mordatch, Pieter Abbeel

* "Interactive Fiction Games: A Colossal Adventure" - Matthew Hausknecht, Prithviraj V Ammanabrolu, Marc-Alexandre Côté, Xingdi Yuan

* "Hierarchy is Exploration: An Empirical Analysis of the Benefits of Hierarchy" - Ofir Nachum, Haoran Tang, Xingyu Lu, Shixiang Gu, Honglak Lee, Sergey Levine

Abstract 8: **Invited Talk in Deep Reinforcement Learning**, *Fei-Fei* 12:00 PM

(Speaker and details forthcoming.)

Abstract 9: **Invited Talk in Deep Reinforcement Learning**, *Todorov* 01:30 PM

(Talk title and abstract TBD.)

Abstract 10: **Contributed Talks in Deep Reinforcement Learning**, *Agarwal, Gleave, Lee* 02:00 PM

* "Striving for Simplicity in Off-Policy Deep Reinforcement Learning" - Rishabh Agarwal, Dale Schuurmans, Mohammad Norouzi

* "Adversarial Policies: Attacking Deep Reinforcement Learning" - Adam R Gleave, Michael Dennis, Neel Kant, Cody Wild, Sergey Levine,

Stuart Russell

* "A Simple Randomization Technique for Generalization in Deep Reinforcement Learning" -
Kimin Lee, Kibok Lee, Jinwoo Shin, Honglak Lee

Abstract 12: **NeurIPS RL Competitions Results Presentations in Deep Reinforcement Learning**,
04:00 PM

16:00 - 16:15 Learn to Move: Walk Around
16:15 - 16:30 Animal Olympics
16:30 - 16:45 Robot open-Ended Autonomous Learning (REAL)
16:45 - 17:00 MineRL

Abstract 13: **Invited Talk in Deep Reinforcement Learning**, *Littman* 05:00 PM

(Talk title and abstract TBD.)

Abstract 14: **Panel Discussion in Deep Reinforcement Learning**, 05:30 PM

(Topic and panelists TBA.)