

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

MLSys: Workshop on Systems for ML *Lakshmiratan, Sen, Gonzalez, Crankshaw, Bird*

Minding the Gap: Between Fairness and Ethics *Rubinov, Kondor, Poulson, Warmuth, Moss, Hagerty*

CiML 2019: Machine Learning Competitions for All *Mendrik, Tu, Guyon, Viegas*

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

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

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

EMC2: Energy Efficient Machine Learning and Cognitive Computing for Embedded Applications (5th edition) *Parihar, Goldfarb, Srivastava*

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

Shared Visual Representations in Human and Machine Intelligence *Deza, Peterson, Murty, Griffiths*

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

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

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

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

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

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

Retrospectives: A Venue for Self-Reflection in ML Research *Lowe, Bengio, Pineau, Paganini, Forde, Sodhani, Gupta, Lehman, Henderson, Madan*

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

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

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

Perception as generative reasoning: structure, causality, probability *Rosenbaum, Garnelo, Battaglia, Allen, Yildirim*

Optimal Transport for Machine Learning *Cuturi, Peyré, Flamary, Suvorikova*

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

Beyond first order methods in machine learning systems *Kyriillidis, Berahas, Roosta, Mahoney*

Bayesian Deep Learning *Gal, Hernández-Lobato, Louizos, Nalisnick, Ghahramani, Murphy, Welling*

AI for Humanitarian Assistance and Disaster Response *Gupta, Murphy, Darrell, Heim, Wang, Goodman*

Meta-Learning *Calandra, Clavera Gilaberte, Hutter, Vanschoren, Wang*

Dec. 14, 2019

Machine Learning with Guarantees *London, Dziugaite, Roy, Joachims, Madry, Shawe-Taylor*

Science meets Engineering of Deep Learning *Sagun, Gulcehre, Romero, Rostamzadeh, de Freitas*

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*

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

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

Program Transformations for ML *Lamblin, Baydin, Wiltschko, van Merriënboer, Fertig, Pearlmutter, Duvenaud, Hascoet*

Machine Learning for Autonomous Driving *McAllister, Rhinehart, Yu, Li, Dragan*

Tackling Climate Change with ML *Rolnick, Lacoste, Maharaj, Donti, Kaack, Platt, Chayes, Bengio*

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

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

Sets and Partitions *Monath, Zaheer, McCallum, Kobren, Oliva, Poczos, Salakhutdinov*

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

Privacy in Machine Learning (PriML) *Balle, Chaudhuri, Honkela, Koskela, Meehan, Park, Smart, Weller*

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

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

Document Intelligence *Duffy, Akkiraju, Bedrax Weiss, Bennett, Motahari-Nezhad*

Deep Reinforcement Learning *Abbeel, Finn, Pineau, Silver, Singh, Achiam, Florensa, Grimm, Tang, Veeriah*

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

Fair ML in Healthcare *Joshi, Chen, Obermeyer, Mullainathan*

Learning Transferable Skills *Mattar, Juliani, Lange, Crosby, Beyret*

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

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

ML For Systems *Hashemi, Mirhoseini, Goldie, Swersky, XU, Raiman*

Medical Imaging meets NeurIPS *Lombaert, Glocker, Konukoglu, de Bruijne, Feragen, Oguz, Teuwen*

Bridging Game Theory and Deep Learning *Mitliagkas, Gidel, He, Askari Hemmat, H, Lacoste-Julien*

MLSys: Workshop on Systems for ML

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

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.

Minding the Gap: Between Fairness and Ethics

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

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.

CiML 2019: Machine Learning Competitions for All

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

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 the 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 propose to 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 will provide room for discussions 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 having

experience in this domain.

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

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

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.

Visually Grounded Interaction and Language

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

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.

Workshop on Federated Learning for Data Privacy and Confidentiality

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

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=flneurips2019>

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.

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

Raj Parihar, Michael Goldfarb, Satyam Srivastava

Fri Dec 13, 08:00 AM

Sample Call for Papers:

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).

More details in the attached pdf

Workshop on Human-Centric Machine Learning

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

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.

Shared Visual Representations in Human and Machine Intelligence

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

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.

Graph Representation Learning

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

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.

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

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.

KR2ML - Knowledge Representation and Reasoning Meets Machine Learning

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

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.

Safety and Robustness in Decision-making

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

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.

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

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.

Learning with Rich Experience: Integration of Learning Paradigms

Zhiting Hu, Andrew Wilson, Chelsea Finn, Lisa Lee, Taylor Berg-Kirkpatrick, Ruslan Salakhutdinov, Eric Xing

Fri Dec 13, 08:00 AM

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.

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

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.

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

Andrew Beam, Tristan Naumann, Brett Beaulieu-Jones, Madalina Fiterau, Irene Y Chen, Sam Finlayson, Emily Alsentzer

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 vitro)
 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.

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

María De-Arteaga, Amanda Coston, Tejumade Afonja

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.

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 S. Marks

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>

Perception as generative reasoning: structure, causality, probability

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

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.

Optimal Transport for Machine Learning

Marco Cuturi, Gabriel Peyré, Rémi Flamary, Alexandra Suvorikova

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.

Biological and Artificial Reinforcement Learning

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

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.

Beyond first order methods in machine learning systems

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

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.

Bayesian Deep Learning

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

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.

The BDL workshop was the second largest workshop at NIPS over the past couple of years, with last year's workshop seeing again a 2x increase in the number of submissions (136 accepted abstracts in total), attracting sponsorship from Google, Microsoft, Uber, and Qualcomm in the form of student travel awards.

Speakers (all confirmed):

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

We would like to highlight the extraordinary gender balance this year with 3 of the 7 invited speakers women in ML.

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.

AI for Humanitarian Assistance and Disaster Response

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

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.



Meta-Learning

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

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.

Machine Learning with Guarantees

Ben London, Gintare Karolina Dziugaite, Dan Roy, Thorsten Joachims, Aleksander Madry, John Shawe-Taylor

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.

Science meets Engineering of Deep Learning

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

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 -- \textbf{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 \textbf{insightfully harness the tremendous collective experimental capability of the field}. As a community, some primary aims would be to (i) \emph{identify} obstacles to better models and algorithms, (ii) \emph{identify} the general trends that are potentially important which we wish to understand scientifically and potentially theoretically and; (iii) \emph{careful design} of scientific experiments whose purpose is to clearly \emph{resolve and pinpoint} the origin of mysteries (so-called "smoking-gun" experiments).

Joint Workshop on AI for Social Good

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

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.

"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

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.

Robot Learning: Control and Interaction in the Real World

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

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.

Program Transformations for ML

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

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

Machine Learning for Autonomous Driving

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

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.

Tackling Climate Change with ML

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

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.

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

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

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 AAAI 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.

The Optimization Foundations of Reinforcement Learning

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

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.

Sets and Partitions

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

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

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

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

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.

Privacy in Machine Learning (PriML)

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

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 researchers, which should result in mutually beneficial new long-term collaborations.

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

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.

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

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.

Document Intelligence

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

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.

Deep Reinforcement Learning

Pieter Abbeel, Chelsea Finn, Joelle Pineau, David Silver, Satinder Singh, Joshua Achiam, Carlos Florensa, Christopher Grimm, Haoran Tang, Vivek Veeriah

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.

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

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

09:00 AM	Invited Talk - 1	<i>Gibson</i>
----------	-------------------------	---------------

10:30 AM	Invited Talk - 2	<i>Zaslavsky</i>
02:00 PM	Invited Talk - 3	<i>Eisner</i>
03:00 PM	Invited Talk - 4	<i>Andreas</i>
04:15 PM	Invited Talk - 5	<i>Lee</i>

Fair ML in Healthcare

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

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.

Learning Transferable Skills

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

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.

Learning with Temporal Point Processes

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

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.

Machine Learning and the Physical Sciences

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

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.

ML For Systems

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

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.

Medical Imaging meets NeurIPS

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

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.

Bridging Game Theory and Deep Learning

Ioannis Mitliagkas, Gauthier Gidel, Niao He, Reyhane Askari Hemmat, N H, Simon Lacoste-Julien

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.