

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

Schedule Highlights

June 14, 2019

- 101, **ICML 2019 Workshop on Computational Biology** *Pe'er, Prabhakaran, Azizi, Diallo, Kundaje, Engelhardt, Dhifli, MEPHU NGUIFO, Tansey, Vogt, Listgarten, Burdziak*
- 102, **ICML 2019 Time Series Workshop** *Kuznetsov, Yang, Yu, Tang, Wang*
- 103, **Human In the Loop Learning (HILL)** *Wang, Wang, Yu, Zhang, Gonzalez, Jia, Bird, Varshney, Kim, Weller*
- 104 A, **Climate Change: How Can AI Help?** *Rolnick, Lacoste, Maharaj, Chayes, Bengio*
- 104 B, **Workshop on the Security and Privacy of Machine Learning** *Papernot, Tramer, Li, Boneh, Evans, Jha, Liang, McDaniel, Steinhardt, Song*
- 104 C, **Theoretical Physics for Deep Learning** *Lee, Lee, Pennington, Pennington, Bahri, Bahri, Welling, Welling, Ganguli, Ganguli, Bruna, Bruna*
- 201, **AI in Finance: Applications and Infrastructure for Multi-Agent Learning** *Reddy, Balch, Wellman, Kumar, Stoica, Elkind*
- 202, **The Third Workshop On Tractable Probabilistic Modeling (TPM)** *Lowd, Vergari, Molina, Rahman, Domingos*
- 203, **Joint Workshop on On-Device Machine Learning & Compact Deep Neural Network Representations (ODML-CDNNR)** *Ravi, Kozareva, Fan, Welling, Chen, Bailer, Kulis, Hu,*

Dekhtiar, Lin, Marculescu

- 204, **Negative Dependence: Theory and Applications in Machine Learning** *Gartrell, Gillenwater, Kulesza, Mariet*

Grand Ballroom A, **1st Workshop on Understanding and Improving Generalization in Deep Learning** *Krishnan, Mobahi, Neyshabur, Bartlett, Song, Srebro*

Grand Ballroom B, **6th ICML Workshop on Automated Machine Learning (AutoML 2019)** *Hutter, Vanschoren, Eggenberger, Feurer*

Hall A, **Generative Modeling and Model-Based Reasoning for Robotics and AI** *Rajeswaran, Todorov, Mordatch, Agnew, Zhang, Pineau, Chang, Erhan, Levine, Stachenfeld, Zhang*

Hall B, **Robustness and Uncertainty Estimation in Deep Learning** *Li, Lakshminarayanan, Hendrycks, Dietterich, Gilmer*

Seaside Ballroom, **Reinforcement Learning for Real Life** *Li, Geramifard, Li, Szepesvari, Wang*

Seaside Ballroom, **Real-world Sequential Decision Making: Reinforcement Learning and Beyond** *Le, Yue, Swaminathan, Boots, Cheng*

June 15, 2019

101, **Workshop on AI for autonomous driving** *Choromanska, Jackel, Li, Niebles, Gaidon, Chao, Posner*

102, **Workshop on Multi-Task and Lifelong Reinforcement Learning** *Chandar, Sodhani, Khetarpal, Zahavy, Mankowitz, Mannor, Ravindran, Precup, Finn, Gupta, Zhang, Cho, Rusu, Rob Fergus*

103, **Invertible Neural Networks and Normalizing**

Flows *Huang, Krueger, Van den Berg, Papamakarios, Gomez, Cremer, Chen, Courville, J. Rezende*

104 A, **Stein's Method for Machine Learning and Statistics**

Briol, Mackey, Oates, Liu, Goldstein

104 B, **AI For Social Good (AISG)**

Luck, Sankaran, Sylvain, McGregor, Penn, Tadesse, Sylvain, Côté, Mackey, Ghani, Bengio

104 C, **Synthetic Realities: Deep Learning for Detecting AudioVisual Fakes**

Biggio, Korshunov, Mensink, Patrini, Sadhu

201, **ICML Workshop on Imitation, Intent, and Interaction (I3)**

Rhinehart, Levine, Finn, He, Kostrikov, Fu, Reddy

202, **Coding Theory For Large-scale Machine Learning**

Cadambe, Grover, Papailiopoulos, Joshi

203, **The How2 Challenge: New Tasks for Vision & Language**

Metze, Specia, Elliot, Barrault, Sanabria, Palaskar

204, **Machine Learning for Music Discovery**

Schmidt, Nieto, Gouyon, Kinnaird, Lanckriet

Grand Ballroom A, **Workshop on Self-Supervised Learning**

van den Oord, Aytar, Doersch, Vondrick, Radford, Sermanet, Zamir, Abbeel

Grand Ballroom B, **Learning and Reasoning with Graph-Structured Representations**

Fetaya, Hu, Kipf, Li, Liang, Liao, Urtasun, Wang, Welling, Xing, Zemel

Hall A, **Exploration in Reinforcement Learning Workshop**

Bhupatiraju, Eysenbach, Gu, Edwards, White, Oudeyer, Stanley, Brunskill

Hall B, **Identifying and Understanding Deep**

Learning Phenomena *Sedghi, Bengio, Hata, Madry, Morcos, Neyshabur, Raghu, Rahimi, Schmidt, Xiao*

Seaside Ballroom, **Adaptive and Multitask**

Learning: Algorithms & Systems *Al-Shedivat, Platanios, Stretcu, Andreas, Talwalkar, Caruana, Mitchell, Xing*

June 14, 2019

ICML 2019 Workshop on Computational Biology

Donna Pe'er, Sandhya Prabhakaran, Elham Azizi, Abdoulaye Baniré Diallo, Anshul Kundaje, Barbara Engelhardt, Wajdi Dhifli, Engelbert MEPHU NGUIFO, Wesley Tansey, Julia Vogt, Jennifer Listgarten, Cassandra Burdzyak

101, Fri Jun 14, 08:30 AM

The workshop will showcase recent research in the field of Computational Biology. There has been significant development in genomic sequencing techniques and imaging technologies. These approaches not only generate huge amounts of data but provide unprecedented resolution of single cells and even subcellular structures. The availability of high dimensional data, at multiple spatial and temporal resolutions has made machine learning and deep learning methods increasingly critical for computational analysis and interpretation of the data. Conversely, biological data has also exposed unique challenges and problems that call for the development of new machine learning methods. This workshop aims to bring together researchers working at the intersection of Machine Learning and Biology to present recent advances and open questions in Computational Biology to the ML community.

The workshop is a sequel to the WCB workshops we organized in the last three years Joint ICML and IJCAI 2018, Stockholm, ICML 2017, Sydney and ICML 2016, New York as well as Workshop on Bioinformatics and AI at IJCAI 2015 Buenos Aires,

IJCAI 2016 New York, IJCAI 2017 Melbourne which had excellent line-ups of talks and were well-received by the community. Every year, we received 60+ submissions. After multiple rounds of rigorous reviewing around 50 submissions were selected from which the best set of papers were chosen for Contributed talks and Spotlights and the rest were invited as Posters. We have a steadfast and growing base of reviewers making up the Program Committee. For the past edition, a special issue of Journal of Computational Biology will be released in the following weeks with extended versions of 14 accepted papers.

We have two confirmed invited speakers and we will invite at least one more leading researcher in the field. Similar to previous years, we plan to request partial funding from Microsoft Research, Google, Python, Swiss Institute of Bioinformatics that we intend to use for student travel awards. In past years, we have also been able to provide awards for best poster/paper and partially contribute to travel expenses for at least 8 students per year.

The Workshop proceedings will be available through CEUR proceedings. We would also have an extended version to be included in a special issue of the Journal of Computational Biology (JCB) for which we already have an agreement with JCB.

ICML 2019 Time Series Workshop

Vitaly Kuznetsov, Scott Yang, Rose Yu, Cheng Tang, Yuyang Wang

102, Fri Jun 14, 08:30 AM

Time series data is both quickly growing and already ubiquitous. In domains spanning as broad a range as climate, robotics, entertainment, finance, healthcare, and transportation, there has been a significant shift away from parsimonious, infrequent

measurements to nearly continuous monitoring and recording. Rapid advances in sensing technologies, ranging from remote sensors to wearables and social sensing, are generating rapid growth in the size and complexity of time series data streams. Thus, the importance and impact of time series analysis and modelling techniques only continues to grow.

At the same time, while time series analysis has been extensively studied by econometricians and statisticians, modern time series data often pose significant challenges for the existing techniques both in terms of their structure (e.g., irregular sampling in hospital records and spatiotemporal structure in climate data) and size. Moreover, the focus on time series in the machine learning community has been comparatively much smaller. In fact, the predominant methods in machine learning often assume i.i.d. data streams, which is generally not appropriate for time series data. Thus, there is both a great need and an exciting opportunity for the machine learning community to develop theory, models and algorithms specifically for the purpose of processing and analyzing time series data.

We see ICML as a great opportunity to bring together theoretical and applied researchers from around the world and with different backgrounds who are interested in the development and usage of time series analysis and algorithms. This includes methods for time series prediction, classification, clustering, anomaly and change point detection, causal discovery, and dimensionality reduction as well as general theory for learning and analyzing stochastic processes. Since time series have been studied in a variety of different fields and have many broad applications, we plan to host leading academic researchers and industry experts with a range of perspectives and interests as invited speakers. Moreover, we also invite researchers from the related areas of batch and online learning,

deep learning, reinforcement learning, data analysis and statistics, and many others to both contribute and participate in this workshop.

Human In the Loop Learning (HILL)

Xin Wang, Xin Wang, Fisher Yu, Shanghang Zhang, Joseph Gonzalez, Yangqing Jia, Sarah Bird, Kush Varshney, Been Kim, Adrian Weller

103, Fri Jun 14, 08:30 AM

<https://sites.google.com/view/hill2019/home>

This workshop is a joint effort between the 4th ICML Workshop on Human Interpretability in Machine Learning (WHI) and the ICML 2019 Workshop on Interactive Data Analysis System (IDAS). We have combined our forces this year to run Human in the Loop Learning (HILL) in conjunction with ICML 2019!

The workshop will bring together researchers and practitioners who study interpretable and interactive learning systems with applications in large scale data processing, data annotations, data visualization, human-assisted data integration, systems and tools to interpret machine learning models as well as algorithm designs for active learning, online learning, and interpretable machine learning algorithms. The target audience for the workshop includes people who are interested in using machines to solve problems by having a human be an integral part of the process. This workshop serves as a platform where researchers can discuss approaches that bridge the gap between humans and machines and get the best of both worlds.

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

Systems for online and interactive learning algorithms,
 Active/Interactive machine learning algorithm design,
 Systems for collecting, preparing, and managing machine learning data,
 Model understanding tools (verifying, diagnosing, debugging, visualization, introspection, etc),
 Design, testing and assessment of interactive systems for data analytics,
 Psychology of human concept learning,
 Generalized additive models, sparsity and rule learning,
 Interpretable unsupervised models (clustering, topic models, etc.),
 Interpretation of black-box models (including deep neural networks),
 Interpretability in reinforcement learning.

Climate Change: How Can AI Help?

David Rolnick, Alexandre Lacoste, Tegan Maharaj, Jennifer Chayes, Yoshua Bengio

104 A, Fri Jun 14, 08:30 AM

Climate Change: How Can AI Help?

We invite submission of extended abstracts applying machine learning to the problems of climate change. There will be three tracks (Deployed, Research, and Ideas).

Workshop website:

Submission website:

*Submission deadline: * **April 30, 11:59 PM**
 Pacific Time

*Notification: * May 15 (early notification possible upon request)

Contact:

Summary

Climate change is widely agreed to be one of the greatest challenges facing humanity. We already observe increased incidence and severity of storms, droughts, fires, and flooding, as well as significant changes to global ecosystems, including the natural resources and agriculture on which humanity depends. The 2018 UN report on climate change estimates that the world has only thirty years to eliminate greenhouse emissions completely if we are to avoid catastrophic consequences.

Many in the machine learning community want to address climate change but feel their skills are inapplicable. This workshop will showcase the many settings in which machine learning can be applied to reducing greenhouse emissions and helping society adapt to the effects of climate change. Climate change is a complex problem requiring simultaneous action from many directions. While machine learning is not a silver bullet, there is significant potential impact for research and implementation.

Call for submissions

We invite submission of extended abstracts on machine learning applied to problems in climate mitigation, adaptation, or modeling, including but not limited to the following topics:

- Power generation and grids
- Transportation
- Smart buildings and cities
- Industrial optimization
- Carbon capture and sequestration
- Agriculture, forestry and other land use
- Climate modeling
- Extreme weather events
- Disaster management and relief
- Societal adaptation
- Ecosystems and natural resources
- Data presentation and management

- Climate finance

Accepted submissions will be invited to give poster presentations at the workshop, of which some will be selected for spotlight talks. Please contact with questions, or if visa considerations make earlier notification important.

Dual-submissions are allowed, and the workshop does not record proceedings. Submissions will be reviewed double-blind; do your best to anonymize your submission, and do not include identifying information for authors in the PDF. We encourage, but do not require, use of the ICML style template (please do not use the “Accepted” format).

Submission tracks

Extended abstracts are limited to 3 pages for the Deployed and Research tracks, and 2 pages for the Ideas track, in PDF format. An additional page may be used for references. All machine learning techniques are welcome, from kernel methods to deep learning. Each submission should make clear why the application has (or could have) positive impacts regarding climate change. There are three tracks for submissions:

Deployed

Work that is already having an impact

Submissions for the Deployed track are intended for machine learning approaches which are impacting climate-relevant problems through consumers or partner institutions. This could include implementations of academic research that have moved beyond the testing phase, as well as results from startups/industry. Details of methodology need not be revealed if they are proprietary, though transparency is encouraged.

Research

Work that will have an impact when deployed

Submissions for the Research track are intended for machine learning research applied to climate-relevant problems. Submissions should provide experimental or theoretical validation of the method proposed, as well as specifying what gap the method fills. Algorithms need not be novel from a machine learning perspective if they are applied in a novel setting.

Datasets may be submitted to this track that are designed to permit machine learning research (e.g. formatted with clear benchmarks for evaluation). In this case, baseline experimental results on the dataset are preferred but not required.

Ideas

Future work that could have an impact

Submissions for the Ideas track are intended for proposed applications of machine learning to solve climate-relevant problems. While the least constrained, this track will be subject to a very high standard of review. No results need be demonstrated, but ideas should be justified as extensively as possible, including motivation for the problem being solved, an explanation of why current tools are inadequate, and details of how tools from machine learning are proposed to fill the gap.

Workshop on the Security and Privacy of Machine Learning

Nicolas Papernot, Florian Tramèr, Bo Li, Dan Boneh, David Evans, Somesh Jha, Percy Liang, Patrick McDaniel, Jacob Steinhardt, Dawn Song

104 B, Fri Jun 14, 08:30 AM

As machine learning has increasingly been deployed in critical real-world applications, the dangers of manipulation and misuse of these models has become of paramount importance to

public safety and user privacy. In applications such as online content recognition to financial analytics to autonomous vehicles all have shown the be vulnerable to adversaries wishing to manipulate the models or mislead models to their malicious ends. This workshop will focus on recent research and future directions about the security and privacy problems in real-world machine learning systems. We aim to bring together experts from machine learning, security, and privacy communities in an attempt to highlight recent work in these area as well as to clarify the foundations of secure and private machine learning strategies. We seek to come to a consensus on a rigorous framework to formulate adversarial attacks targeting machine learning models, and to characterize the properties that ensure the security and privacy of machine learning systems. Finally, we hope to chart out important directions for future work and cross-community collaborations.

Theoretical Physics for Deep Learning

Jaehoon Lee, Jaehoon Lee, Jeffrey Pennington, Jeffrey Pennington, Yasaman Bahri, Yasaman Bahri, Max Welling, Max Welling, Surya Ganguli, Surya Ganguli, Joan Bruna, Joan Bruna

104 C, Fri Jun 14, 08:30 AM

Though the purview of physics is broad and includes many loosely connected subdisciplines, a unifying theme is the endeavor to provide concise, quantitative, and predictive descriptions of the often large and complex systems governing phenomena that occur in the natural world. While one could debate how closely deep learning is connected to the natural world, it is undeniably the case that deep learning systems are large and complex; as such, it is reasonable to consider whether the rich body of ideas and powerful tools from theoretical physicists could be harnessed to improve our understanding of

deep learning. The goal of this workshop is to investigate this question by bringing together experts in theoretical physics and deep learning in order to stimulate interaction and to begin exploring how theoretical physics can shed light on the theory of deep learning.

We believe ICML is an appropriate venue for this gathering as members from both communities are frequently in attendance and because deep learning theory has emerged as a focus at the conference, both as an independent track in the main conference and in numerous workshops over the last few years. Moreover, the conference has enjoyed an increasing number of papers using physics tools and ideas to draw insights into deep learning.

AI in Finance: Applications and Infrastructure for Multi-Agent Learning

Prashant Reddy, Tucker Balch, Michael Wellman, Senthil Kumar, Ion Stoica, Edith Elkind

201, Fri Jun 14, 08:30 AM

Finance is a rich domain for AI and ML research. Model-driven strategies for stock trading and risk assessment models for loan approvals are quintessential financial applications that are reasonably well-understood. However, there are a number of other applications that call for attention as well.

In particular, many finance domains involve ecosystems of interacting and competing agents. Consider for instance the detection of financial fraud and money-laundering. This is a challenging multi-agent learning problem, especially because the real world agents involved evolve their strategies constantly. Similarly, in algorithmic trading of stocks, commodities, etc., the actions of

any given trading agent affects, and is affected by, other trading agents -- many of these agents are constantly learning in order to adapt to evolving market scenarios. Further, such trading agents operate at such a speed and scale that they must be fully autonomous. They have grown in sophistication to employ advanced ML strategies including deep learning, reinforcement learning, and transfer learning.

Financial institutions have a long history of investing in technology as a differentiator and have been key drivers in advancing computing infrastructure (e.g., low-latency networking). As more financial applications employ deep learning and reinforcement learning, there is consensus now on the need for more advanced computing architectures--for training large machine learning models and simulating large multi-agent learning systems--that balance scale with the stringent privacy requirements of finance.

Historically, financial firms have been highly secretive about their proprietary technology developments. But now, there is also emerging consensus on the need for (1) deeper engagement with academia to advance a shared knowledge of the unique challenges faced in FinTech, and (2) more open collaboration with academic and technology partners through intellectually sophisticated fora such as this proposed workshop.

The Third Workshop On Tractable Probabilistic Modeling (TPM)

Daniel Lowd, Antonio Vergari, Alejandro Molina, TAHRIMA Rahman, Pedro Domingos

202, Fri Jun 14, 08:30 AM

Probabilistic modeling has become the de facto framework to reason about uncertainty in Machine

Learning and AI. One of the main challenges in probabilistic modeling is the trade-off between the expressivity of the models and the complexity of performing various types of inference, as well as learning them from data.

This inherent trade-off is clearly visible in powerful -- but intractable -- models like Markov random fields, (restricted) Boltzmann machines, (hierarchical) Dirichlet processes and Variational Autoencoders. Despite these models' recent successes, performing inference on them resorts to approximate routines. Moreover, learning such models from data is generally harder as inference is a sub-routine of learning, requiring simplifying assumptions or further approximations. Having guarantees on tractability at inference and learning time is then a highly desired property in many real-world scenarios.

Tractable probabilistic modeling (TPM) concerns methods guaranteeing exactly this: performing exact (or tractably approximate) inference and/or learning. To achieve this, the following approaches have been proposed: i) low or bounded-treewidth probabilistic graphical models and determinantal point processes, that exchange expressiveness for efficiency; ii) graphical models with high girth or weak potentials, that provide bounds on the performance of approximate inference methods; and iii) exchangeable probabilistic models that exploit symmetries to reduce inference complexity. More recently, models compiling inference routines into efficient computational graphs such as arithmetic circuits, sum-product networks, cutset networks and probabilistic sentential decision diagrams have advanced the state-of-the-art inference performance by exploiting context-specific independence, determinism or by exploiting latent variables. TPMs have been successfully used in numerous real-world applications: image classification, completion and generation, scene understanding, activity recognition, language and

speech modeling, bioinformatics, collaborative filtering, verification and diagnosis of physical systems.

The aim of this workshop is to bring together researchers working on the different fronts of tractable probabilistic modeling, highlighting recent trends and open challenges. At the same time, we want to foster the discussion across similar or complementary sub-fields in the broader probabilistic modeling community. In particular, the rising field of neural probabilistic models, such as normalizing flows and autoregressive models that achieve impressive results in generative modeling. It is an interesting open challenge for the TPM community to keep a broad range of inference routines tractable while leveraging these models' expressiveness. Furthermore, the rising field of probabilistic programming promises to be the new lingua franca of model-based learning. This offers the TPM community opportunities to push the expressiveness of the models used for general-purpose universal probabilistic languages, such as Pyro, while maintaining efficiency.

We want to promote discussions and advance the field both by having high quality contributed works, as well as high level invited speakers coming from the aforementioned tangent sub-fields of probabilistic modeling.

Joint Workshop on On-Device Machine Learning & Compact Deep Neural Network Representations (ODML-CDNNR)

Sujith Ravi, Zornitsa Kozareva, Lixin Fan, Max Welling, Yurong Chen, Werner Bailer, Brian Kulis, Haoji Hu, Jonathan Dekhtiar, Yingyan Lin, Diana Marculescu

203, Fri Jun 14, 08:30 AM

This joint workshop aims to bring together researchers, educators, practitioners who are interested in techniques as well as applications of on-device machine learning and compact, efficient neural network representations. One aim of the workshop discussion is to establish close connection between researchers in the machine learning community and engineers in industry, and to benefit both academic researchers as well as industrial practitioners. The other aim is the evaluation and comparability of resource-efficient machine learning methods and compact and efficient network representations, and their relation to particular target platforms (some of which may be highly optimized for neural network inference). The research community has still to develop established evaluation procedures and metrics.

The workshop also aims at reproducibility and comparability of methods for compact and efficient neural network representations, and on-device machine learning. Contributors are thus encouraged to make their code available. The workshop organizers plan to make some example tasks and datasets available, and invite contributors to use them for testing their work. In order to provide comparable performance evaluation conditions, the use of a common platform (such as Google Colab) is intended.

Negative Dependence: Theory and Applications in Machine Learning

Mike Gartrell, Jennifer Gillenwater, Alex Kulesza, Zelda Mariet

204, Fri Jun 14, 08:30 AM

Models of negative dependence are increasingly important in machine learning. Whether selecting training data, finding an optimal experimental design, exploring in reinforcement learning, or

making suggestions with recommender systems, selecting high-quality but diverse items has become a core challenge. This workshop aims to bring together researchers who, using theoretical or applied techniques, leverage negative dependence in their work. We will delve into the rich underlying mathematical theory, understand key applications, and discuss the most promising directions for future research.

1st Workshop on Understanding and Improving Generalization in Deep Learning

Dilip Krishnan, Hossein Mobahi, Behnam Neyshabur, Peter Bartlett, Dawn Song, Nati Srebro

Grand Ballroom A, Fri Jun 14, 08:30 AM

[Workshop's Website](#)

The 1st workshop on Generalization in Deep Networks: Theory and Practice will be held as part of ICML 2019. Generalization is one of the fundamental problems of machine learning, and increasingly important as deep networks make their presence felt in domains with big, small, noisy or skewed data. This workshop will consider generalization from both theoretical and practical perspectives. We welcome contributions from paradigms such as representation learning, transfer learning and reinforcement learning. The workshop invites researchers to submit working papers in the following research areas:

- Implicit regularization: the role of optimization algorithms in generalization
- Explicit regularization methods
- Network architecture choices that improve generalization
- Empirical approaches to understanding generalization

Generalization bounds; empirical evaluation criteria to evaluate bounds

Robustness: generalizing to distributional shift a.k.a dataset shift

Generalization in the context of representation learning, transfer learning and deep reinforcement learning: definitions and empirical approaches

6th ICML Workshop on Automated Machine Learning (AutoML 2019)

Frank Hutter, Joaquin Vanschoren, Katharina Eggenberger, Matthias Feurer

Grand Ballroom B, Fri Jun 14, 08:30 AM

Workshop website:

Submission website:

Submission deadline: May 2, anywhere on earth

Notification: May 17

Machine learning has achieved considerable successes in recent years, but this success often relies on human experts, who construct appropriate features, design learning architectures, set their hyperparameters, and develop new learning algorithms. Driven by the demand for off-the-shelf machine learning methods from an ever-growing community, the research area of AutoML targets the progressive automation of machine learning aiming to make effective methods available to everyone. The workshop targets a broad audience ranging from core machine learning researchers in different fields of ML connected to AutoML, such as neural architecture search, hyperparameter optimization, meta-learning, and learning to learn, to domain experts aiming to apply machine learning to new types of problems.

We invite submissions on the topics of:

- * Model selection, hyper-parameter optimization, and model search

- * Neural architecture search
- * Meta-learning and transfer learning
- * Learning to learn new algorithms and strategies
- * Automation of any element of the ML pipeline, including:
 - * feature extraction / construction
 - * data cleaning
 - * generation of workflows / workflow reuse
 - * problem "ingestion" (from raw data and miscellaneous formats)
 - * acquisition of new data (active learning, experimental design)
 - * report generation (providing insight on automated data analysis)
 - * selection of evaluation metrics / validation procedures
 - * selection of algorithms under time/space/power constraints
 - * construction of fair and unbiased machine learning models
 - * semi-supervised and unsupervised machine learning
 - * Extending the scope of AutoML towards automated data science
 - * Human-in-the-loop approaches for AutoML
 - * Demos of existing AutoML systems
 - * Robustness of AutoML systems (w.r.t. randomized algorithms, data, hardware etc.)
 - * Hyperparameter agnostic algorithms

We welcome submissions up to 6 pages in JMLR format (+ references). We strongly encourage attachments of code to foster reproducibility; reproducibility of results and easy availability of code will be taken into account in the decision making process. All accepted papers will be presented as posters. We may invite the best 2-3 papers for an oral plenary presentation. Unless indicated by the authors, we will provide PDFs of all accepted papers on . There will be no archival proceedings. For submission details please see the submission page.

Generative Modeling and Model-Based Reasoning for Robotics and AI

Aravind Rajeswaran, Emanuel Todorov, Igor Mordatch, William Agnew, Amy Zhang, Joelle Pineau, Michael Chang, Dumitru Erhan, Sergey Levine, Kimberly Stachenfeld, Marvin Zhang

Hall A, Fri Jun 14, 08:30 AM

Workshop website:

<https://sites.google.com/view/mbri-icml2019>

In the recent explosion of interest in deep RL, "model-free" approaches based on Q-learning and actor-critic architectures have received the most attention due to their flexibility and ease of use. However, this generality often comes at the expense of efficiency (statistical as well as computational) and robustness. The large number of required samples and safety concerns often limit direct use of model-free RL for real-world settings.

Model-based methods are expected to be more efficient. Given accurate models, trajectory optimization and Monte-Carlo planning methods can efficiently compute near-optimal actions in varied contexts. Advances in generative modeling, unsupervised, and self-supervised learning provide methods for learning models and representations that support subsequent planning and reasoning. Against this backdrop, our workshop aims to bring together researchers in generative modeling and model-based control to discuss research questions at their intersection, and to advance the state of the art in model-based RL for robotics and AI. In particular, this workshops aims to make progress on questions related to:

1. How can we learn generative models efficiently? Role of data, structures, priors, and uncertainty.
2. How to use generative models efficiently for

planning and reasoning? Role of derivatives, sampling, hierarchies, uncertainty, counterfactual reasoning etc.

3. How to harmoniously integrate model-learning and model-based decision making?
4. How can we learn compositional structure and environmental constraints? Can this be leveraged for better generalization and reasoning?

Schedule

08:30 AM	Introduction	<i>Rajeswaran</i>
----------	---------------------	-------------------

Abstracts (1):

Abstract 1: **Introduction in Generative Modeling and Model-Based Reasoning for Robotics and AI**, *Rajeswaran* 08:30 AM

Opening remarks by workshop organizers.

Robustness and Uncertainty Estimation in Deep Learning

Yixuan (Sharon) Li, Balaji Lakshminarayanan, Dan Hendrycks, Tom Dietterich, Justin Gilmer

Hall B, Fri Jun 14, 08:30 AM

There has been growing interest in rectifying deep neural network vulnerabilities. Challenges arise when models receive samples drawn from outside the training distribution. For example, a neural network tasked with classifying handwritten digits may assign high confidence predictions to cat images. Anomalies are frequently encountered when deploying ML models in the real world. Well-calibrated predictive uncertainty estimates are indispensable for many machine learning applications, such as self-driving cars and medical diagnosis systems. Generalization to unseen and

worst-case inputs is also essential for robustness to distributional shift. In order to have ML models reliably predict in open environment, we must deepen technical understanding in the following areas: (1) learning algorithms that are robust to changes in input data distribution (e.g., detect out-of-distribution examples); (2) mechanisms to estimate and calibrate confidence produced by neural networks and (3) methods to improve robustness to adversarial and common corruptions, and (4) key applications for uncertainty such as in artificial intelligence (e.g., computer vision, robotics, self-driving cars, medical imaging) as well as broader machine learning tasks.

This workshop will bring together researchers and practitioners from the machine learning communities, and highlight recent work that contribute to address these challenges. Our agenda will feature contributed papers with invited speakers. Through the workshop we hope to help identify fundamentally important directions on robust and reliable deep learning, and foster future collaborations.

Reinforcement Learning for Real Life

Yuxi Li, Alborz Geramifard, Lihong Li, Csaba Szepesvari, Tao Wang

Seaside Ballroom, Fri Jun 14, 08:30 AM

Reinforcement learning (RL) is a general learning, predicting, and decision making paradigm. RL provides solution methods for sequential decision making problems as well as those can be transformed into sequential ones. RL connects deeply with optimization, statistics, game theory, causal inference, sequential experimentation, etc., overlaps largely with approximate dynamic programming and optimal control, and applies broadly in science, engineering and arts.

RL has been making steady progress in academia recently, e.g., Atari games, AlphaGo, visuomotor policies for robots. RL has also been applied to real world scenarios like recommender systems and neural architecture search. See a recent collection about RL applications at <https://medium.com/@yuxili/rl-applications-73ef685c07eb>.

It is desirable to have RL systems that work in the real world with real benefits. However, there are many issues for RL though, e.g. generalization, sample efficiency, and exploration vs. exploitation dilemma. Consequently, RL is far from being widely deployed. Common, critical and pressing questions for the RL community are then: Will RL have wide deployments? What are the issues? How to solve them?

The goal of this workshop is to bring together researchers and practitioners from industry and academia interested in addressing practical and/or theoretical issues in applying RL to real life scenarios, review state of the arts, clarify impactful research problems, brainstorm open challenges, share first-hand lessons and experiences from real life deployments, summarize what has worked and what has not, collect tips for people from industry looking to apply RL and RL experts interested in applying their methods to real domains, identify potential opportunities, generate new ideas for future lines of research and development, and promote awareness and collaboration. This is not "yet another RL workshop": it is about how to successfully apply RL to real life applications. This is a less addressed issue in the RL/ML/AI community, and calls for immediate attention for sustainable prosperity of RL research and development.

Hoang Le, Yisong Yue, Adith Swaminathan, Byron Boots, Ching-An Cheng

Seaside Ballroom, Fri Jun 14, 14:00 PM

Workshop website:

This workshop aims to bring together researchers from industry and academia in order to describe recent advances and discuss future research directions pertaining to real-world sequential decision making, broadly construed. We aim to highlight new and emerging research opportunities for the machine learning community that arise from the evolving needs for making decision making theoretically and practically relevant for realistic applications.

Research interest in reinforcement and imitation learning has surged significantly over the past several years, with the empirical successes of self-playing in games and availability of increasingly realistic simulation environments. We believe the time is ripe for the research community to push beyond simulated domains and start exploring research directions that directly address the real-world need for optimal decision making. We are particularly interested in understanding the current theoretical and practical challenges that prevent broader adoption of current policy learning and evaluation algorithms in high-impact applications, across a broad range of domains.

This workshop welcomes both theory and application contributions.

**Real-world Sequential Decision Making:
Reinforcement Learning and Beyond**

June 15, 2019

Workshop on AI for autonomous driving

Anna Choromanska, Larry Jackel, Li Erran Li, Juan Carlos Niebles, Adrien Gaidon, Wei-Lun Chao, Ingmar Posner

101, Sat Jun 15, 08:30 AM

A diverse set of methods have been devised to develop autonomous driving platforms. They range from modular systems, systems that perform manual decomposition of the problem, systems where the components are optimized independently, and a large number of rules are programmed manually, to end-to-end deep-learning frameworks. Today's systems rely on a subset of the following: camera images, HD maps, inertial measurement units, wheel encoders, and active 3D sensors (LIDAR, radar). There is a general agreement that much of the self-driving software stack will continue to incorporate some form of machine learning in any of the above mentioned systems in the future.

Self-driving cars present one of today's greatest challenges and opportunities for Artificial Intelligence (AI). Despite substantial investments, existing methods for building autonomous vehicles have not yet succeeded, i.e., there are no driverless cars on public roads today without human safety drivers. Nevertheless, a few groups have started working on extending the idea of learned tasks to larger functions of autonomous driving. Initial results on learned road following are very promising.

The goal of this workshop is to explore ways to create a framework that is capable of learning

autonomous driving capabilities beyond road following, towards fully driverless cars. The workshop will consider the current state of learning applied to autonomous vehicles and will explore how learning may be used in future systems. The workshop will span both theoretical frameworks and practical issues especially in the area of deep learning.

Workshop on Multi-Task and Lifelong Reinforcement Learning

Sarath Chandar, Shagun Sodhani, Khimya Khetarpal, Tom Zahavy, Daniel J. Mankowitz, Shie Mannor, Balaraman Ravindran, Doina Precup, Chelsea Finn, Abhishek Gupta, Amy Zhang, Kyunghyun Cho, Andrei Rusu, Facebook Rob Fergus

102, Sat Jun 15, 08:30 AM

****Website****

<https://sites.google.com/view/mtlr>

****Abstract****

Significant progress has been made in reinforcement learning, enabling agents to accomplish complex tasks such as Atari games, robotic manipulation, simulated locomotion, and Go. These successes have stemmed from the core reinforcement learning formulation of learning a single policy or value function from scratch. However, reinforcement learning has proven challenging to scale to many practical real world problems due to problems in learning efficiency and objective specification, among many others. Recently, there has been emerging interest and research in leveraging structure and information across multiple reinforcement learning tasks to more efficiently and effectively learn complex behaviors. This includes:

- * curriculum and lifelong learning, where the problem requires learning a sequence of tasks, leveraging their shared structure to enable knowledge transfer
- * goal-conditioned reinforcement learning techniques that leverage the structure of the provided goal space to learn many tasks significantly faster
- * meta-learning methods that aim to learn efficient learning algorithms that can learn new tasks quickly
- * hierarchical reinforcement learning, where the reinforcement learning problem might entail a compositions of subgoals or subtasks with shared structure

Multi-task and lifelong reinforcement learning has the potential to alter the paradigm of traditional reinforcement learning, to provide more practical and diverse sources of supervision, while helping overcome many challenges associated with reinforcement learning, such as exploration, sample efficiency and credit assignment. However, the field of multi-task and lifelong reinforcement learning is still young, with many more developments needed in terms of problem formulation, algorithmic and theoretical advances as well as better benchmarking and evaluation.

The focus of this workshop will be on both the algorithmic and theoretical foundations of multi-task and lifelong reinforcement learning as well as the practical challenges associated with building multi-tasking agents and lifelong learning benchmarks. Our goal is to bring together researchers that study different problem domains (such as games, robotics, language, and so forth), different optimization approaches (deep learning, evolutionary algorithms, model-based control, etc.), and different formalisms (as mentioned above) to discuss the frontiers, open problems and meaningful next steps in multi-task and lifelong reinforcement learning.

Confirmed Speakers

- * Jacob Andreas
- * Jeff Clune
- * Karol Hausman
- * Nicolas Heess
- * Sergey Levine
- * Natalia Rodriguez
- * Benjamin Rosman
- * Peter Stone
- * Martha White

Invertible Neural Networks and Normalizing Flows

Chin-Wei Huang, David Krueger, Rianne Van den Berg, George Papamakarios, Aidan Gomez, Chris Cremer, Ricky T. Q. Chen, Aaron Courville, Danilo J. Rezende

103, Sat Jun 15, 08:30 AM

Invertible neural networks have been a significant thread of research in the ICML community for several years. Such transformations can offer a range of unique benefits:

- (1) They preserve information, allowing perfect reconstruction (up to numerical limits) and obviating the need to store hidden activations in memory for backpropagation.
- (2) They are often designed to track the changes in probability density that applying the transformation induces (as in normalizing flows).
- (3) Like autoregressive models, normalizing flows can be powerful generative models which allow exact likelihood computations; with the right architecture, they can also allow for much cheaper sampling than autoregressive models.

While many researchers are aware of these topics and intrigued by several high-profile papers, few are

familiar enough with the technical details to easily follow new developments and contribute. Many may also be unaware of the wide range of applications of invertible neural networks, beyond generative modelling and variational inference.

Stein's Method for Machine Learning and Statistics

Francois-Xavier Briol, Lester Mackey, Chris Oates, Qiang Liu, Larry Goldstein

104 A, Sat Jun 15, 08:30 AM

Stein's method is a technique from probability theory for bounding the distance between probability measures using differential and difference operators. Although the method was initially designed as a technique for proving central limit theorems, it has recently caught the attention of the machine learning (ML) community and has been used for a variety of practical tasks. Recent applications include generative modeling, global non-convex optimisation, variational inference, de novo sampling, constructing powerful control variates for Monte Carlo variance reduction, and measuring the quality of (approximate) Markov chain Monte Carlo algorithms. Stein's method has also been used to develop goodness-of-fit tests and was the foundational tool in one of the NeurIPS 2017 Best Paper awards.

Although Stein's method has already had significant impact in ML, most of the applications only scratch the surface of this rich area of research in probability theory. There would be significant gains to be made by encouraging both communities to interact directly, and this inaugural workshop would be an important step in this direction. More precisely, the aims are: (i) to introduce this emerging topic to the wider ML community, (ii) to highlight the wide range of existing applications in

ML, and (iii) to bring together experts in Stein's method and ML researchers to discuss and explore potential further uses of Stein's method.

Further details can be found on the [workshop's website](https://steinworkshop.github.io/).

AI For Social Good (AISG)

Margaux Luck, Kris Sankaran, Tristan Sylvain, Sean McGregor, Jonnie Penn, Girmaw Abebe Tadesse, Virgile Sylvain, Myriam Côté, Lester Mackey, Rayid Ghani, Yoshua Bengio

104 B, Sat Jun 15, 08:30 AM

#AI for Social Good

##Important information

Contact information:

aisg2019.icml.contact@gmail.com

Submission deadline: **EXTENDED to April 26th 2019 11:59PM ET**

**[Workshop

website](https://aiforsocialgood.github.io/icml2019/index.htm)**

**[Submission

website](https://cmt3.research.microsoft.com/AISGW2019/)**

Poster Information:

* Poster Size - **36W x 48H inches or 90 x 122 cm**

* Poster Paper - **lightweight paper - not laminated**

##Abstract

This workshop builds on our AI for Social Good workshop at [NeurIPS 2018](https://aiforsocialgood.github.io/2018/) and

[ICLR 2019](https://aiforsocialgood.github.io/iclr2019/).

****Introduction:**** The rapid expansion of AI research presents two clear conundrums:

- the comparative lack of incentives for researchers to address social impact issues and
- the dearth of conferences and journals centered around the topic. Researchers motivated to help often find themselves without a clear idea of which fields to delve into.

****Goals:**** Our workshop address both these issues by bringing together machine learning researchers, social impact leaders, stakeholders, policy leaders, and philanthropists to discuss their ideas and applications for social good. To broaden the impact beyond the convening of our workshop, we are partnering with [AI Commons](http://www.aicommons.com) to expose accepted projects and papers to the broader community of machine learning researchers and engineers. The projects/research may be at varying degrees of development, from formulation as a data problem to detailed requirements for effective deployment. We hope that this gathering of talent and information will inspire the creation of new approaches and tools by the community, help scientists access the data they need, involve social and policy stakeholders in the framing of machine learning applications, and attract interest from philanthropists invited to the event to make a dent in our shared goals.

****Topics:**** The UN Sustainable Development Goals (SDGs), a set of seventeen objectives whose completion is set to lead to a more equitable, prosperous, and sustainable world. In this light, our main areas of focus are the following: health, education, the protection of democracy, urban planning, assistive technology, agriculture, environmental protection and sustainability, social

welfare and justice, developing world. Each of these themes presents unique opportunities for AI to reduce human suffering and allow citizens and democratic institutions to thrive.

Across these topics, we have dual goals: recognizing high-quality work in machine learning motivated by or applied to social applications, and creating meaningful connections between communities dedicated to solving technical and social problems. To this extent, we propose two research tracks:

- ****Short Papers Track** (Up to four page papers + unlimited pages for citations)** for oral and/or poster presentation. The short papers should focus on past and current research work, showcasing actual results and demonstrating beneficial effects on society. We also accept short papers of recently published or submitted journal contributions to give authors the opportunity to present their work and obtain feedback from conference attendees.
- ****Problem Introduction Track** (Application form, up to five page responses + unlimited pages for citations)** which will present a specific solution that will be shared with stakeholders, scientists, and funders. The workshop will provide a suite of questions designed to: (1) estimate the feasibility and impact of the proposed solutions, and (2) estimate the importance of data in their implementation. The application responses should highlight ideas that have not yet been implemented in practice but can lead to real impact. The projects may be at varying degrees of development, from formulation as a data problem to structure for effective deployment. The workshop provides a supportive platform for developing these early-stage or hobby proposals into real projects. This process is designed to foster sharing different points of view ranging from the scientific assessment of feasibility, discussion of practical constraints that may be encountered, and attracting interest from philanthropists invited to the event. Accepted

submissions may be promoted to the wider AI solutions community following the workshop via the [AI Commons](http://www.aicommons.com), with whom we are partnering to promote the longer-term development of projects.

Synthetic Realities: Deep Learning for Detecting AudioVisual Fakes

Battista Biggio, Pavel Korshunov, Thomas Mensink, Giorgio Patrini, Arka Sadhu

104 C, Sat Jun 15, 08:30 AM

With the latest advances of deep generative models, synthesis of images and videos as well as of human voices have achieved impressive realism. In many domains, synthetic media are already difficult to distinguish from real by the human eye and ear. The potential of misuses of these technologies is seldom discussed in academic papers; instead, vocal concerns are rising from media and security organizations as well as from governments. Researchers are starting to experiment on new ways to integrate deep learning with traditional media forensics and security techniques as part of a technological solution. This workshop will bring together experts from the communities of machine learning, computer security and forensics in an attempt to highlight recent work and discuss future effort to address these challenges. Our agenda will alternate contributed papers with invited speakers. The latter will emphasize connections among the interested scientific communities and the standpoint of institutions and media organizations.

ICML Workshop on Imitation, Intent, and Interaction (I3)

Nicholas Rhinehart, Sergey Levine, Chelsea Finn, He He, Ilya Kostrikov, Justin Fu, Siddharth Reddy

201, Sat Jun 15, 08:30 AM

****Website**:**

[https://sites.google.com/view/icml-i3](https://sites.google.com/

****Abstract**:** A key challenge for deploying interactive machine learning systems in the real world is the ability for machines to understand human intent. Techniques such as imitation learning and inverse reinforcement learning are popular data-driven paradigms for modeling agent intentions and controlling agent behaviors, and have been applied to domains ranging from robotics and autonomous driving to dialogue systems. Such techniques provide a practical solution to specifying objectives to machine learning systems when they are difficult to program by hand.

While significant progress has been made in these areas, most research effort has concentrated on modeling and controlling single agents from dense demonstrations or feedback. However, the real world has multiple agents, and dense expert data collection can be prohibitively expensive. Surmounting these obstacles requires progress in frontiers such as:

- 1) the ability to infer intent from multiple modes of data, such as language or observation, in addition to traditional demonstrations.
- 2) the ability to model multiple agents and their intentions, both in cooperative and adversarial settings.
- 3) handling partial or incomplete information from the expert, such as demonstrations that lack dense action annotations, raw videos, etc..

The workshop on Imitation, Intention, and Interaction (I3) seeks contributions at the interface of these frontiers, and will bring together researchers from multiple disciplines such as

robotics, imitation and reinforcement learning, cognitive science, AI safety, and natural language understanding. Our aim will be to reexamine the assumptions in standard imitation learning problem statements (e.g., inverse reinforcement learning) and connect distinct application disciplines, such as robotics and NLP, with researchers developing core imitation learning algorithms. In this way, we hope to arrive at new problem formulations, new research directions, and the development of new connections across distinct disciplines that interact with imitation learning methods.

Coding Theory For Large-scale Machine Learning

Viveck Cadambe, Pulkit Grover, Dimitris Papailiopoulos, Gauri Joshi

202, Sat Jun 15, 08:30 AM

Coding Theory For Large-scale Machine Learning
Coding theory involves the art and science of how to add redundancy to data to ensure that a desirable output is obtained at despite deviations from ideal behavior from the system components that interact with the data. Through a rich, mathematically elegant set of techniques, coding theory has come to significantly influence the design of modern data communications, compression and storage systems. The last few years have seen a rapidly growing interest in coding theory based approaches for the development of efficient machine learning algorithms towards robust, large-scale, distributed computational pipelines.

The CodML workshop brings together researchers developing coding techniques for machine learning, as well as researchers working on systems implementations for computing, with cutting-edge presentations from both sides. The goal is to learn about non-idealities in system components as well

as approaches to obtain reliable and robust learning despite these non-idealities, and identify problems of future interest.

The workshop is co-located with ICML 2019, and will be held in Long Beach, California, USA on June 14th or 15th, 2019.

Please see the [website](<https://sites.google.com/view/codml2019>) for more details:

Call for Posters

Scope of the Workshop

In this workshop we solicit research papers focused on the application of coding and information-theoretic techniques for distributed machine learning. More broadly, we seek papers that address the problem of making machine learning more scalable, efficient, and robust. Both theoretical as well as experimental contributions are welcome. We invite authors to submit papers on topics including but not limited to:

- * Asynchronous Distributed Training Methods
- * Communication-Efficient Training
- * Model Compression and Quantization
- * Gradient Coding, Compression and Quantization
- * Erasure Coding Techniques for Straggler Mitigation
- * Data Compression in Large-scale Machine Learning
- * Erasure Coding Techniques for ML Hardware Acceleration
- * Fast, Efficient and Scalable Inference
- * Secure and Private Machine Learning
- * Data Storage/Access for Machine Learning Jobs
- * Performance evaluation of coding techniques

Submission Format and Instructions

The authors should prepare extended abstracts in the ICML paper format and submit via [CMT](https://cmt3.research.microsoft.com/CODMLW2019/). Topics include the study of multi-modal representations, translation between modalities, bootstrapping of labels from one modality into another, visually-grounded question answering, segmentation and storytelling, and grounding the meaning of language in visual data. An ever-increasing number of tasks and datasets are appearing around this recently-established field.

Submitted papers may not exceed three (3) single-spaced double-column pages excluding references. All results, proofs, figures, tables must be included in the 3 pages. The submitted manuscripts should include author names and affiliations, and an abstract that does not exceed 250 words. The authors may include a link to an extended version of the paper that includes supplementary material (proofs, experimental details, etc.) but the reviewers are not required to read the extended version.

Dual Submission Policy

Accepted submissions will be considered non-archival and can be submitted elsewhere without modification, as long as the other conference allows it. Moreover, submissions to CodML based on work recently accepted to other venues are also acceptable (though authors should explicitly make note of this in their submissions).

Key Dates

****Paper Submission:**** May 3rd, 2019, 11:59 PM anywhere on earth

****Decision Notification:**** May 12th, 2019.

****Workshop date:**** June 14 or 15, 2019

The How2 Challenge: New Tasks for Vision & Language

Florian Metze, Lucia Specia, Desmond Elliot, Loic Barrault, Ramon Sanabria, Shruti Palaskar

203, Sat Jun 15, 08:30 AM

Research at the intersection of vision and language has been attracting a lot of attention in recent years. Topics include the study of multi-modal representations, translation between modalities, bootstrapping of labels from one modality into another, visually-grounded question answering, segmentation and storytelling, and grounding the meaning of language in visual data. An ever-increasing number of tasks and datasets are appearing around this recently-established field.

At NeurIPS 2018, we released the How2 data-set, containing more than 85,000 (2000h) videos, with audio, transcriptions, translations, and textual summaries. We believe it presents an ideal resource to bring together researchers working on the previously mentioned separate tasks around a single, large dataset. This rich dataset will facilitate the comparison of tools and algorithms, and hopefully foster the creation of additional annotations and tasks. We want to foster discussion about useful tasks, metrics, and labeling techniques, in order to develop a better understanding of the role and value of multi-modality in vision and language. We seek to create a venue to encourage collaboration between different sub-fields, and help establish new research directions and collaborations that we believe will sustain machine learning research for years to come.

[Workshop Homepage](https://srvk.github.io/how2-challenge/)

Schedule

08:30 AM **Welcome**

Machine Learning for Music Discovery

Erik Schmidt, Oriol Nieto, Fabien Gouyon, Katherine Kinnaid, Gert Lanckriet

204, Sat Jun 15, 08:30 AM

The ever-increasing size and accessibility of vast music libraries has created a demand more than ever for artificial systems that are capable of understanding, organizing, or even generating such complex data. While this topic has received relatively marginal attention within the machine learning community, it has been an area of intense focus within the community of Music Information Retrieval (MIR). While significant progress has been made, these problems remain far from solved.

Furthermore, the recommender systems community has made great advances in terms of collaborative feedback recommenders, but these approaches suffer strongly from the cold-start problem. As such, recommendation techniques often fall back on content-based machine learning systems, but defining musical similarity is extremely challenging as myriad features all play some role (e.g., cultural, emotional, timbral, rhythmic). Thus, for machines must actually understand music to achieve an expert level of music recommendation.

On the other side of this problem sits the recent explosion of work in the area of machine creativity. Relevant examples are both Google Magenta and the startup Jukedeck, who seek to develop algorithms capable of composing and performing completely original (and compelling) works of music. These algorithms require a similar deep understanding of music and present challenging new problems for the machine learning and AI community at large.

This workshop proposal is timely in that it will bridge these separate pockets of otherwise very related research. And in addition to making progress on the challenges above, we hope to engage the wide AI and machine learning community with our nebulous

problem space, and connect them with the many available datasets the MIR community has to offer (e.g., Audio Set, AcousticBrainz, Million Song Dataset), which offer near commercial scale to the academic research community.

Workshop on Self-Supervised Learning

Aaron van den Oord, Yusuf Aytar, Carl Doersch, Carl Vondrick, Alec Radford, Pierre Sermanet, Amir Zamir, Pieter Abbeel

Grand Ballroom A, Sat Jun 15, 08:30 AM

Self-supervised learning is a promising alternative where proxy tasks are developed that allow models and agents to learn without explicit supervision in a way that helps with downstream performance on tasks of interest. One of the major benefits of self-supervised learning is increasing data efficiency: achieving comparable or better performance with less labeled data or fewer environment steps (in Reinforcement learning / Robotics).

The field of self-supervised learning (SSL) is rapidly evolving, and the performance of these methods is creeping closer to the fully supervised approaches. However, many of these methods are still developed in domain-specific sub-communities, such as Vision, RL and NLP, even though many similarities exist between them. While SSL is an emerging topic and there is great interest in these techniques, there are currently few workshops, tutorials or other scientific events dedicated to this topic.

This workshop aims to bring together experts with different backgrounds and applications areas to share inter-domain ideas and increase cross-pollination, tackle current shortcomings and explore new directions. The focus will be on the machine learning point of view rather than the domain side.

<https://sites.google.com/corp/view/self-supervised-icml2019>

Learning and Reasoning with Graph-Structured Representations

Ethan Fetaya, Zhiting Hu, Thomas Kipf, Yujia Li, Xiaodan Liang, Renjie Liao, Raquel Urtasun, Hao Wang, Max Welling, Eric Xing, Richard Zemel

Grand Ballroom B, Sat Jun 15, 08:30 AM

Graph-structured representations are widely used as a natural and powerful way to encode information such as relations between objects or entities, interactions between online users (e.g., in social networks), 3D meshes in computer graphics, multi-agent environments, as well as molecular structures, to name a few. Learning and reasoning with graph-structured representations is gaining increasing interest in both academia and industry, due to its fundamental advantages over more traditional unstructured methods in supporting interpretability, causality, transferability, etc. Recently, there is a surge of new techniques in the context of deep learning, such as graph neural networks, for learning graph representations and performing reasoning and prediction, which have achieved impressive progress. However, it can still be a long way to go to obtain satisfactory results in long-range multi-step reasoning, scalable learning with very large graphs, flexible modeling of graphs in combination with other dimensions such as temporal variation and other modalities such as language and vision. New advances in theoretical foundations, models and algorithms, as well as empirical discoveries and applications are therefore all highly desirable.

The aims of this workshop are to bring together researchers to dive deeply into some of the most

promising methods which are under active exploration today, discuss how we can design new and better benchmarks, identify impactful application domains, encourage discussion and foster collaboration. The workshop will feature speakers, panelists, and poster presenters from machine perception, natural language processing, multi-agent behavior and communication, meta-learning, planning, and reinforcement learning, covering approaches which include (but are not limited to):

- Deep learning methods on graphs/manifolds/relational data (e.g., graph neural networks)
- Deep generative models of graphs (e.g., for drug design)
- Unsupervised graph/manifold/relational embedding methods (e.g., hyperbolic embeddings)
- Optimization methods for graphs/manifolds/relational data
- Relational or object-level reasoning in machine perception
- Relational/structured inductive biases for reinforcement learning, modeling multi-agent behavior and communication
- Neural-symbolic integration
- Theoretical analysis of capacity/generalization of deep learning models for graphs/manifolds/relational data
- Benchmark datasets and evaluation metrics

Exploration in Reinforcement Learning Workshop

Surya Bhupatiraju, Benjamin Eysenbach, Shixiang Gu, Harrison Edwards, Martha White, Pierre-Yves Oudeyer, Kenneth Stanley, Emma Brunskill

Hall A, Sat Jun 15, 08:30 AM

Exploration is a key component of reinforcement learning (RL). While RL has begun to solve relatively simple tasks, current algorithms cannot complete complex tasks. Our existing algorithms often endlessly dither, failing to meaningfully explore their environments in search of high-reward states. If we hope to have agents autonomously learn increasingly complex tasks, these machines must be equipped with machinery for efficient exploration.

The goal of this workshop is to present and discuss exploration in RL, including deep RL, evolutionary algorithms, real-world applications, and developmental robotics. Invited speakers will share their perspectives on efficient exploration, and researchers will share recent work in spotlight presentations and poster sessions.

Identifying and Understanding Deep Learning Phenomena

Hanie Sedghi, Samy Bengio, Kenji Hata, Aleksander Madry, Ari Morcos, Behnam Neyshabur, Maithra Raghu, Ali Rahimi, Ludwig Schmidt, Ying Xiao

Hall B, Sat Jun 15, 08:30 AM

Our understanding of modern neural networks lags behind their practical successes. As this understanding gap grows, it poses a serious challenge to the future pace of progress because fewer pillars of knowledge will be available to designers of models and algorithms. This workshop aims to close this understanding gap in deep learning. It solicits contributions that view the behavior of deep nets as a natural phenomenon to investigate with methods inspired from the natural sciences, like physics, astronomy, and biology. We solicit empirical work that isolates phenomena in deep nets, describes them quantitatively, and then

replicates or falsifies them.

As a starting point for this effort, we focus on the interplay between data, network architecture, and training algorithms. We are looking for contributions that identify precise, reproducible phenomena, as well as systematic studies and evaluations of current beliefs such as “sharp local minima do not generalize well” or “SGD navigates out of local minima”. Through the workshop, we hope to catalogue quantifiable versions of such statements, as well as demonstrate whether or not they occur reproducibly.

Adaptive and Multitask Learning: Algorithms & Systems

Maruan Al-Shedivat, Anthony Platanios, Otilia Stretcu, Jacob Andreas, Ameet Talwalkar, Rich Caruana, Tom Mitchell, Eric Xing

Seaside Ballroom, Sat Jun 15, 08:30 AM

Driven by progress in deep learning, the machine learning community is now able to tackle increasingly more complex problems—ranging from multi-modal reasoning to dexterous robotic manipulation—all of which typically involve solving nontrivial combinations of tasks. Thus, designing adaptive models and algorithms that can efficiently learn, master, and combine multiple tasks is the next frontier. AMTL workshop aims to bring together machine learning researchers from areas ranging from theory to applications and systems, to explore and discuss:

* advantages, disadvantages, and applicability of different approaches to learning in multitask settings,

* formal or intuitive connections between methods developed for different problems that help better understand the landscape of multitask learning

techniques and inspire technique transfer between research lines,

* fundamental challenges and open questions that the community needs to tackle for the field to move forward.

Webpage:

[www.amtl-workshop.org](<https://www.amtl-workshop.org>)