

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

Schedule Highlights

Dec. 7, 2018

Critiquing and Correcting Trends in Machine Learning *Rainforth, Kusner, Bloem-Reddy, Paige, Caruana, Teh*

NIPS 2018 Competition Track Day 1 *Escalera, Herbrich*

Deep Reinforcement Learning *Abbeel, Silver, Singh, Pineau, Achiam, Houthoof, Srinivas*

All of Bayesian Nonparametrics (Especially the Useful Bits) *Cai, Campbell, Hughes, Broderick, Foti, Williamson*

MLSys: Workshop on Systems for ML and Open Source Software *Lakshmiratan, Bird, Sen, Gonzalez, Crankshaw*

Imitation Learning and its Challenges in Robotics *Mukadam, Choudhury, Srinivasa*

The second Conversational AI workshop – today's practice and tomorrow's potential *Geramifard, Williams, Boureau, Eskenazi, Gasic, Glass, Hakkani-Tur, Heck, Polymenakos, Young*

2nd Workshop on Machine Learning on the Phone and other Consumer Devices (MLPCD 2) *Ravi, Chai, Jia, Aradhye, Jain*

Challenges and Opportunities for AI in Financial Services: the Impact of Fairness, Explainability, Accuracy, and Privacy *Veloso, Kallus, Shah, Kumar, Moulinier, Chen, Paisley*

Causal Learning *Arjovsky, Heinz-Deml, Klimovskaia, Oquab, Bottou, Lopez-Paz*

NIPS 2018 workshop on Compact Deep Neural Networks with industrial applications *Fan, Lin, Welling, Chen, Bailer*

Smooth Games Optimization and Machine Learning *Lacoste-Julien, Mitliagkas, Gidel, Syrgkanis, Tardos, Bottou, Nowozin*

Modeling and decision-making in the spatiotemporal domain *Senanayake, Jean, Ramos, Chowdhary*

Workshop on Security in Machine Learning *Papernot, Tramer, Chaudhuri, Fredrikson, Steinhardt*

Workshop on Ethical, Social and Governance Issues in AI *Bakalar, Bird, Caetano, Felten, Garcia, Kloumann, Lattimore, Mullainathan, Sculley*

Machine Learning for Geophysical & Geochemical Signals *Pyrak-Nolte, Rustad, Baraniuk*

Continual Learning *Pascanu, Teh, Pickett, Ring*

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

Visually grounded interaction and language *Strub, de Vries, Wijmans, Datta, Perez, Malinowski, Lee, Anderson, Courville, MARY, Batra, Parikh, Pietquin, HORI, Marks, Cherian*

Modeling the Physical World: Learning, Perception, and Control *Wu, Allen, Smith, Hamrick, Dupoux, Toussaint, Tenenbaum*

Dec. 8, 2018

NIPS 2018 Competition Track Day 2 *Herbrich, Escalera*

Wordplay: Reinforcement and Language Learning in Text-based Games *Trischler, Lazaridou, Bisk, Tay, Kushman, Côté, Sordani,*

Ricks, Zahavy, Daumé III

Integration of Deep Learning Theories *Baraniuk, Anandkumar, Mallat, Patel, H*

Machine Learning for Systems *Goldie, Mirhoseini, Raiman, Swersky, Hashemi*

Relational Representation Learning *Grover, Varma, Sala, Holtzen, Neville, Ermon, Ré*

NIPS Workshop on Machine Learning for Intelligent Transportation Systems 2018 *Li, Dragan, Niebles, Savarese*

Machine Learning for Health (ML4H): Moving beyond supervised learning in healthcare *Beam, Naumann, Ghassemi, McDermott, Fiterau, Chen, Beaulieu-Jones, Hughes, Shamout, Chivers, Kandola, Yahi, Finlayson, Jedynak, Schulam, Antropova, Fries, Dalca*

Second Workshop on Machine Learning for Creativity and Design *Elliott, Dieleman, Fiebrink, Engel, Roberts, White*

Machine Learning for Molecules and Materials *Hernández-Lobato, Müller, Paige, Kusner, Chmiela, Schütt*

Medical Imaging meets NIPS *Konukoglu, Glocker, Lombaert, de Bruijine*

Machine Learning for the Developing World (ML4D): Achieving sustainable impact *Herlands, De-Arteaga*

Reinforcement Learning under Partial Observability *Pajarinen, Amato, Poupart, Hsu*

Privacy Preserving Machine Learning *Bellet, Gascon, Kilbertus, Ohrimenko, Raykova, Weller*

AI for social good *Luck, Sylvain, Cohen, Fansi Tchango, Goddard, Helouis, Bengio, Greydanus, Wild, Kucherenko, Farahi, Penn, McGregor, Crowley, Gupta, Chen, Côté*

NIPS 2018 Workshop on Meta-Learning *Grant, Hutter, Ravi, Vanschoren, Wang*

CiML 2018 - Machine Learning competitions "in the wild": Playing in the real world or in real time *Guyon, Viegas, Escalera, Abernethy*

Infer to Control: Probabilistic Reinforcement Learning and Structured Control *Kaelbling, Riedmiller, Toussaint, Mordatch, Fox, Haarnoja*

Emergent Communication Workshop *Foerster, Lazaridou, Lowe, Mordatch, Kiela, Cho*

Interpretability and Robustness in Audio, Speech, and Language *Ravanelli, Serdyuk, Vairani, Ramabhadran*

Machine Learning Open Source Software 2018: Sustainable communities *Strathmann, Gal, Curtin, Lisitsyn, Honkela, Ong*

Learning by Instruction *Srivastava, Labutov, Yang, Azaria, Mitchell*

Dec. 7, 2018

Critiquing and Correcting Trends in Machine Learning

Tom Rainforth, Matt Kusner, Ben Bloem-Reddy, Brooks Paige, Rich Caruana, Yee Whye Teh

Fri Dec 07, 08:00 AM

Machine learning has become wildly successful in many application areas. Recently there have been calls for making machine learning more reproducible, less hand-tailored, fair, and generally more thoughtful about how research is conducted and put into practice. These are hallmarks of a mature scientific field, and will be crucial for machine learning to have the wide-ranging, positive impact it is expected to have. Without careful consideration, we as a field risk inflating expectations beyond what is possible. To address this, this workshop aims to better understand and to improve all stages of the research process in machine learning.

There have been a number of recent works that have carefully considered current trends in machine learning as well as the needs of the field when used in real-world scenarios [1-16]. Each of these works introspectively analyzes what we often take for granted as a field. Further, many propose solutions for moving forward. The goal of this workshop is to bring together researchers from all subfields of machine learning to address current shortcomings in the field, and crucially, to propose solutions. We hope to highlight issues and propose solutions in areas such as:

- Common experimental practices [1, 8]
- Implicit technical and empirical assumptions that go unquestioned [2, 3, 5, 7, 11, 12, 13]
- Shortfalls in publication and reviewing setups [15, 16]
- Disconnects between research focus and application requirements [9, 10, 14]
- Surprising observations that make us rethink our research priorities [4, 6]

In a highly interactive format, we will outline the current challenges and ways forward. The program is a collection of invited talks, alongside contributed posters and talks. For these talks, we plan a unique open format of 10 minutes of talk + 10 minutes of follow up discussion. Additionally, a separate panel discussion will collect researchers with a diverse set of viewpoints on the current challenges and potential solutions. During the panel we will also open the conversation to the audience. The discussion will further be open to an online Q&A which will be solicited prior to the workshop. An expected outcome of this workshop is a list of open problems and problematic trends at all levels of machine learning research, and some possible solutions.

Call for Papers:

The NIPS 2018 Workshop: Critiquing and Correcting Trends in Machine Learning calls for papers that critically examine current common

practices and/or trends in methodology, datasets, empirical standards, publication models, or any other aspect of machine learning research. In particular, we seek constructive papers, which propose a solution or indicate a way forward. Examples of such papers include (but are not limited to):

- Papers that argue for broad structural changes, such as: reviewer guidelines and standards; the field's emphasis on short research reports; the conference publication model; standards for code and result reproducibility.
- Proposition papers that outline important and long term challenges in specific subfields.
- Papers that define desirable operating characteristics for machine learning algorithms in real-world application settings.
- Papers that propose to eliminate or replace particular methodological norms. For example, papers which question common implicit assumptions or experimental practices.

Papers should motivate their arguments by describing gaps in the field. Crucially, this is not a venue for settling scores or character attacks, but for moving machine learning forward as a scientific discipline.

The workshop will include a poster session, giving the opportunity to present novel ideas and ongoing projects. Submissions should be 2-4 pages in the NIPS format. Please email all submissions to: nips2018.guidelines.workshop@gmail.com

Deadline: October 18th, 2018, 11:59 UTC

References

- [1] Mania, H., Guy, A., & Recht, B. (2018). Simple random search provides a competitive approach to reinforcement learning. arXiv preprint arXiv:1803.07055.
- [2] Rainforth, T., Kosiorek, A. R., Le, T. A., Maddison, C. J., Igl, M., Wood, F., & Teh, Y. W. (2018). Tighter variational bounds are not necessarily better. ICML.
- [3] Torralba, A., & Efros, A. A. (2011). Unbiased look at dataset bias. In Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on (pp. 1521-1528). IEEE.
- [4] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., & Fergus, R. (2013). Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199.
- [5] Mescheder, L., Geiger, A., Nowozin S. (2018) Which Training Methods for GANs do actually Converge? ICML
- [6] Daumé III, H. (2009). Frustratingly easy domain adaptation. arXiv preprint arXiv:0907.1815
- [7] Urban, G., Geras, K. J., Kahou, S. E., Wang, O. A. S., Caruana, R., Mohamed, A., ... & Richardson, M. (2016). Do deep convolutional nets really need to be deep (or even convolutional)?
- [8] Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., & Meger, D. (2017). Deep reinforcement learning that matters. arXiv preprint arXiv:1709.06560.
- [9] Narayanan, M., Chen, E., He, J., Kim, B., Gershman, S., & Doshi-Velez, F. (2018). How do Humans Understand Explanations from Machine Learning Systems? An Evaluation of the Human-Interpretability of Explanation. arXiv preprint arXiv:1802.00682.
- [10] Schulam, S., Saria S. (2017). Reliable Decision Support using Counterfactual Models. NIPS.
- [11] Rahimi, A. (2017). Let's take machine learning from alchemy to electricity. Test-of-time award presentation, NIPS.
- [12] Lucic, M., Kurach, K., Michalski, M., Gelly, S., Bousquet, O. (2018).

Are GANs Created Equal? A Large-Scale Study. arXiv preprint arXiv:1711.10337.

[13] Le, T.A., Kosiorek, A.R., Siddharth, N., Teh, Y.W. and Wood, F., (2018). Revisiting Reweighted Wake-Sleep. arXiv preprint arXiv:1805.10469.

[14] Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J. and Mané, D., (2016). Concrete problems in AI safety. arXiv preprint arXiv:1606.06565.

[15] Sutton, C. (2018) Making unblinding manageable: Towards reconciling prepublication and double-blind review.

<http://www.theexclusive.org/2017/09/arxiv-double-blind.html>

[16] Langford, J. (2018) ICML Board and Reviewer profiles. <http://hunch.net/?p=8962378>

This workshop aims to highlight recent advances in modeling and computation through the lens of applied, domain-driven problems that require the infinite flexibility and interpretability of BNP. In this workshop, we will explore new BNP methods for diverse applied problems, including cutting-edge models being developed by application domain experts. We will also discuss the limitations of existing methods and discuss key problems that need to be solved. A major focus of the workshop will be to expose participants to practical software tools for performing Bayesian nonparametric analyses. In particular, we plan to host hands-on tutorials to introduce workshop participants to some of the software packages that can be used to easily perform posterior inference for BNP models. On the software panel, we will have researchers who have experience with BNP and development experience with popular software systems, such as TensorFlow, Edward, Stan, and Autograd.

NIPS 2018 Competition Track Day 1

Sergio Escalera, Ralf Herbrich

Fri Dec 07, 08:00 AM

coming soon.

We expect workshop participants to come from a variety of fields, including but not limited to machine learning, statistics, engineering, the social sciences, and biological sciences. The workshop will be relevant both to BNP experts as well as those interested in learning how to apply BNP models. There will be a special emphasis on novel application areas and computational developments that make BNP more accessible to the broader machine learning audience. Participants will leave the workshop with (i) exposure to recent advances in the field, (ii) hands-on experience with software implementing BNP methods, and (iii) an idea of the current major challenges in the field. These goals will be accomplished through a series of invited and contributed talks, a poster session, and at least one hands-on tutorial session where participants can get their hands dirty with BNP methods.

Deep Reinforcement Learning

Pieter Abbeel, David Silver, Satinder Singh, Joelle Pineau, Joshua Achiam, Rein Houthoofd, Aravind Srinivas

Fri Dec 07, 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.

This workshop builds off of:

1. NIPS 2015: "Bayesian Nonparametrics: The Next Generation":

<https://sites.google.com/site/nipsbnp2015/>, and

2. NIPS 2016: "Practical Bayesian Nonparametrics":

<https://sites.google.com/site/nipsbnp2016/>,

which have spanned various areas of BNP, such as theory, applications and computation. This year's workshop will have a fresh take on recent developments in BNP in connection to the broader range of research in statistics, machine learning, and application domains.

The 2018 workshop has received an endorsement from the International Society of Bayesian Analysis (ISBA) and sponsorship from Google.

All of Bayesian Nonparametrics (Especially the Useful Bits)

Diana Cai, Trevor Campbell, Mike Hughes, Tamara Broderick, Nick Foti, Sinead A Williamson

Fri Dec 07, 08:00 AM

Bayesian nonparametric (BNP) methods are well suited to the large data sets that arise in a wide variety of applied fields. By making use of infinite-dimensional mathematical structures, BNP methods allow the complexity of a learned model to grow as the size of a data set grows, exhibiting desirable Bayesian regularization properties for small data sets and allowing the practitioner to learn ever more from larger data sets. These properties have resulted in the adoption of BNP methods across a diverse set of application areas---including, but not limited to, biology, neuroscience, the humanities, social sciences, economics, and finance.

Organizing Committee:

Diana Cai (Princeton)

Trevor Campbell (MIT/UBC)

Mike Hughes (Harvard/Tufts)

Tamara Broderick (MIT)

Nick Foti (U Washington)

Sinead Williamson (UT Austin)

Advisory Committee:

Emily Fox (U Washington)

Antonio Lijoi (Bocconi U)

Sonia Petrone (Bocconi U)

Igor Prünster (Bocconi U)

Erik Sudderth (UC Irvine)

MLSys: Workshop on Systems for ML and Open Source Software

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

Fri Dec 07, 08:00 AM

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.

The ML for Systems workshop focuses on developing ML to optimize systems while we focus on designing systems to enable large scale ML with Systems for ML. Both fields are mature enough to warrant a dedicated workshop. Organizers on both sides are open to merging in the future, but this year we plan to run them separately on two different days.

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, GPU processing, and other topics.

This workshop will follow the successful model we have previously run at ICML, NIPS and SOSP 2017.

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.

Schedule

09:00 AM **Welcome**

Imitation Learning and its Challenges in Robotics

Mustafa Mukadam, Sanjiban Choudhury, Siddhartha Srinivasa

Fri Dec 07, 08:00 AM

Many animals including humans have the ability to acquire skills, knowledge, and social cues from a very young age. This ability to imitate by learning from demonstrations has inspired research across many disciplines like anthropology, neuroscience, psychology, and artificial intelligence. In AI, imitation learning (IL) serves as an essential tool for learning skills that are difficult to program by hand. The applicability of IL to robotics in particular, is useful when learning by trial and error (reinforcement learning) can be hazardous in the real world. Despite the many recent breakthroughs in IL, in the context of robotics there are several challenges to be addressed if robots are to operate freely and interact with humans in the real world.

Some important challenges include: 1) achieving good generalization and sample efficiency when the user can only provide a limited number of demonstrations with little to no feedback; 2) learning safe behaviors in human environments that require the least user intervention in terms of safety overrides without being overly conservative; and 3) leveraging data from multiple sources, including non-human sources, since limitations in hardware interfaces can often lead to poor quality demonstrations.

In this workshop, we aim to bring together researchers and experts in robotics, imitation and reinforcement learning, deep learning, and human robot interaction to

- Formalize the representations and primary challenges in IL as they pertain to robotics
- Delineate the key strengths and limitations of existing approaches with respect to these challenges
- Establish common baselines, metrics, and benchmarks, and identify open questions

Schedule

08:50 AM	Introduction	<i>Mukadam, Choudhury, Srinivasa</i>
09:00 AM	Peter Stone	<i>Stone</i>
09:30 AM	Sonia Chernova	<i>Chernova</i>
10:00 AM	Industry Spotlight	
10:15 AM	Contributed Spotlights	
10:30 AM	Coffee Break and Poster Session I	
11:00 AM	Stefan Schaal	<i>Schaal</i>
11:30 AM	Anca Dragan	<i>Dragan</i>
12:00 PM	Lunch Break	
02:00 PM	Byron Boots	<i>Boots</i>
02:30 PM	Ingmar Posner	<i>Posner</i>
03:00 PM	Coffee Break and Poster Session II	
03:30 PM	Yisong Yue	<i>Yue</i>
04:00 PM	TBD	
04:30 PM	Panel Discussion	

The second Conversational AI workshop – today’s practice and tomorrow’s potential

Alborz Geramifard, Jason D. Williams, Y-Lan Boureau, Maxine Eskenazi, Milica Gasic, Jim Glass, Dilek Hakkani-Tur, Larry Heck, Lazaros C. Polymenakos, Steve Young

Fri Dec 07, 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. 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 year NIPS workshop (9 invited talks, 26 submissions, 3 orals papers, 13 accepted papers, 37 PC members, and couple of hundreds of participants), 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, share findings from large-scale real-world deployments, and generate new ideas for future lines of research.

This workshop will include invited talks from academia and industry, contributed work, and open discussion. In these talks, senior technical leaders from many of the most popular conversational services will give insights into real usage and challenges at scale. An open call for papers will be issued, and we will prioritize forward-looking papers that propose interesting and impactful contributions. We will end the day with an open discussion, including a panel consisting of academic and industrial researchers.

Schedule

01:25 PM	Alex’s talk	<i>Rudnicky</i>
03:30 PM	Removing Natural Interaction Frictions from Conversational AI Systems	<i>Sarikaya</i>

2nd Workshop on Machine Learning on the Phone and other Consumer Devices (MLPCD 2)

Sujith Ravi, Wei Chai, Yangqing Jia, Hrishikesh Aradhye, Prateek Jain

Fri Dec 07, 08:00 AM

The 2nd Workshop on Machine Learning on the Phone and other Consumer Devices (MLPCD 2) aims to continue the success of the 1st

MLPCD workshop held at NIPS 2017 in Long Beach, CA.

Previously, the first MLPCD workshop edition, held at NIPS 2017 was successful, attracted over 200+ attendees and led to active research & panel discussions as well as follow-up contributions to the open-source community (e.g., release of new inference libraries, tools, models and standardized representations of deep learning models). We believe that interest in this space is only going to increase, and we hope that the workshop plays the role of an influential catalyst to foster research and collaboration in this nascent community.

After the first workshop where we investigated initial directions and trends, the NIPS 2018 MLPCD workshop focuses on theory and practical applications of on-device machine learning, an area that is highly relevant and specializes in the intersection of multiple topics of interest to NIPS and broader machine learning community -- efficient training & inference for deep learning and other machine learning models; interdisciplinary mobile applications involving vision, language & speech understanding; and emerging topics like Internet of Things.

We plan to incorporate several new additions this year -- inspirational opening Keynote talk on "future of intelligent assistive & wearable experiences"; two panels including a lively closing panel debate discussing pros/cons of two key ML computing paradigms (Cloud vs. On-device); solicited research papers on new & recent hot topics (e.g., theoretical & algorithmic work on low-precision models, compression, sparsity, etc. for training and inference), related challenges, applications and recent trends; demo session showcasing ML in action for real-world apps.

Description & Topics:

Deep learning and machine learning, in general, has changed the computing paradigm. Products of today are built with machine intelligence as a central attribute, and consumers are beginning to expect near-human interaction with the appliances they use. However, much of the Deep Learning revolution has been limited to the cloud, enabled by popular toolkits such as Caffe, TensorFlow, and MxNet, and by specialized hardware such as TPUs. In comparison, mobile devices until recently were just not fast enough, there were limited developer tools, and there were limited use cases that required on-device machine learning. That has recently started to change, with the advances in real-time computer vision and spoken language understanding driving real innovation in intelligent mobile applications. Several mobile-optimized neural network libraries were recently announced (CoreML, Caffe2 for mobile, TensorFlow Lite), which aim to dramatically reduce the barrier to entry for mobile machine learning. Innovation and competition at the silicon layer has enabled new possibilities for hardware acceleration. To make things even better, mobile-optimized versions of several state-of-the-art benchmark models were recently open sourced. Widespread increase in availability of connected “smart” appliances for consumers and IoT platforms for industrial use cases means that there is an ever-expanding surface area for mobile intelligence and ambient devices in homes. All of these advances in combination imply that we are likely at the cusp of a rapid increase in research interest in on-device machine learning, and in particular, on-device neural computing.

Significant research challenges remain, however. Mobile devices are even more personal than “personal computers” were. Enabling machine

learning while simultaneously preserving user trust requires ongoing advances in the research of differential privacy and federated learning techniques. On-device ML has to keep model size and power usage low while simultaneously optimizing for accuracy. There are a few exciting novel approaches recently developed in mobile optimization of neural networks. Lastly, the newly prevalent use of camera and voice as interaction models has fueled exciting research towards neural techniques for image and speech/language understanding. This is an area that is highly relevant to multiple topics of interest to NIPS -- e.g., core topics like machine learning & efficient inference and interdisciplinary applications involving vision, language & speech understanding as well as emerging area (namely, Internet of Things).

With this emerging interest as well as the wealth of challenging research problems in mind, we are proposing the second NIPS 2018 workshop dedicated to on-device machine learning for mobile and ambient home consumer devices.

Areas/topics of interest include, but not limited to:

- * Model compression for efficient inference with deep networks and other ML models
- * Privacy preserving machine learning
- * Low-precision training/inference & Hardware acceleration of neural computing on mobile devices
- * Real-time mobile computer vision
- * Language understanding and conversational assistants on mobile devices
- * Speech recognition on mobile and smart home devices
- * Machine intelligence for mobile gaming
- * ML for mobile health other real-time prediction scenarios
- * ML for on-device applications in the automotive industry (e.g., computer vision for self-driving cars)
- * Software libraries (including open-source) optimized for on-device ML

Target Audience:

The next wave of ML applications will have significant processing on mobile and ambient devices. Some immediate examples of these are single-image classification, depth estimation, object recognition and segmentation running on-device for creative effects, or on-device recommender and ranking systems for privacy-preserving, low-latency experiences. This workshop will bring ML practitioners up to speed on the latest trends for on-device applications of ML, offer an overview of the latest HW and SW framework developments, and champion active research towards hard technical challenges emerging in this nascent area. The target audience for the workshop is both industrial and academic researchers and practitioners of on-device, native machine learning. The workshop will cover both "informational" and "aspirational" aspects of this emerging research area for delivering ground-breaking experiences on real-world products.

Given the relevance of the topic, target audience (mix of industry + academia & related parties) as well as the timing (confluence of research ideas + practical implementations both in industry as well as through publicly available toolkits), we feel that NIPS 2018 would continue to be a great venue for this workshop.

Challenges and Opportunities for AI in Financial Services: the Impact of Fairness, Explainability, Accuracy, and Privacy

Manuela Veloso, Nathan Kallus, Sameena Shah, Senthil Kumar, Isabelle Moulinier, Jiahao Chen, John Paisley

Fri Dec 07, 08:00 AM

The adoption of artificial intelligence in the financial service industry, particularly the adoption of machine learning, presents challenges and opportunities. Challenges include algorithmic fairness, explainability, privacy, and requirements of a very high degree of accuracy. 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 doesn't cause humans to miss critical pieces of data. For some use cases, the operating standards require nothing short of perfect accuracy.

Privacy issues around collection and use of consumer and proprietary data require high levels of scrutiny. Many machine learning models are deemed unusable if they are not supported by appropriate levels of explainability. 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 and other types of risks.

The aim of this workshop is to bring together researchers and practitioners to discuss challenges for AI in financial services, and the opportunities such challenges represent to the community. The workshop will consist of a series of sessions, including invited talks, panel discussions and short paper presentations, which will showcase ongoing research and novel algorithms.

Causal Learning

Martin Arjovsky, Christina Heinze-Deml, Anna Klimovskaia, Maxime Oquab, Leon Bottou, David Lopez-Paz

Fri Dec 07, 08:00 AM

Site for the workshop:

<https://sites.google.com/view/nips2018causallearning/home>

The route from machine learning to artificial intelligence remains uncharted. Recent efforts describe some of the conceptual problems that lie along this route [4, 9, 12]. The goal of this workshop is to investigate how much progress is possible by framing these problems beyond learning correlations, that is, by uncovering and leveraging causal relations:

1. Machine learning algorithms solve statistical problems (e.g. maximum likelihood) as a proxy to solve tasks of interest (e.g. recognizing objects). Unfortunately, spurious correlations and biases are often easier to learn than the task itself [14], leading to unreliable or unfair predictions. This phenomenon can be framed as causal confounding.

2. Machines trained on large pools of i.i.d. data often crash confidently when deployed in different circumstances (e.g., adversarial examples, dataset biases [18]). In contrast, humans seek prediction rules robust across multiple conditions. Allowing machines to learn robust rules from

multiple environments can be framed as searching for causal invariances [2, 11, 16, 17].

3. Humans benefit from discrete structures to reason. Such structures seem less useful to learning machines. For instance, neural machine translation systems outperform those that model language structure. However, the purpose of this structure might not be modeling common sentences, but to help us formulate new ones. Modeling new potential sentences rather than observed ones is a form of counterfactual reasoning [8, 9].

4. Intelligent agents do not only observe, but also shape the world with actions. Maintaining plausible causal models of the world allows to build intuitions, as well as to design intelligent experiments and interventions to test them [16, 17]. Is causal understanding necessary for efficient reinforcement learning?

5. Humans learn compositionally; after learning simple skills, we are able to recombine them quickly to solve new tasks. Such abilities have so far eluded our machine learning systems. Causal models are compositional, so they might offer a solution to this puzzle [4].

6. Finally, humans are able to digest large amounts of unsupervised signals into a causal model of the world. Humans can learn causal affordances, that is, imagining how to manipulate new objects to achieve goals, and the outcome of doing so. Humans rely on a simple blueprint for a complex world: models that contain the correct causal structures, but ignore irrelevant details [16, 17].

We cannot address these problems by simply performing inference on known causal graphs. We need to learn from data to discover plausible causal models, and to construct predictors that are robust to distributional shifts. Furthermore, much prior work has focused on estimating explicit causal structures from data, but these methods are often unscalable, rely on untestable assumptions like faithfulness or acyclicity, and are difficult to incorporate into high-dimensional, complex and nonlinear machine learning pipelines. Instead of considering the task of estimating causal graphs as their final goal, learning machines may use notions from causation indirectly to ignore biases, generalize across distributions, leverage structure to reason, design efficient interventions, benefit from compositionality, and build causal models of the world in an unsupervised way.

Call for papers

Submit your anonymous, NIPS-formatted manuscript here [https://easychair.org/cfp/NIPSCL2018]. All accepted submissions will require a poster presentation. A selection of submissions will be awarded a 5-minute spotlight presentation. We welcome conceptual, thought-provoking material, as well as research agendas, open problems, new tasks, and datasets.

Submission deadline: 28 October 2018

Acceptance notifications: 9 November 2018

Schedule:

See <https://sites.google.com/view/nips2018causallearning/home> for the up-to-date schedule.

Speakers:

Judea Pearl
David Blei
Nicolai Meinshausen
Bernhard Schölkopf
Isabelle Guyon
Csaba Szepesvari
Pietro Perona

References

1. Krzysztof Chalupka, Pietro Perona, Frederick Eberhardt (2015): Visual Causal Feature Learning [https://arxiv.org/abs/1412.2309]
2. Christina Heinze-Deml, Nicolai Meinshausen (2018): Conditional Variance Penalties and Domain Shift Robustness [https://arxiv.org/abs/1710.11469]
3. Fredrik D. Johansson, Uri Shalit, David Sontag (2016): Learning Representations for Counterfactual Inference [https://arxiv.org/abs/1605.03661]
4. Brenden Lake (2014): Towards more human-like concept learning in machines: compositionality, causality, and learning-to-learn [https://dspace.mit.edu/handle/1721.1/95856]
5. Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, Samuel J. Gershman (2016): Building Machines That Learn and Think Like People [https://arxiv.org/abs/1604.00289]
6. David Lopez-Paz, Krikamol Muandet, Bernhard Schölkopf, Ilya Tolstikhin (2015): Towards a Learning Theory of Cause-Effect Inference [https://arxiv.org/abs/1309.6779]
7. David Lopez-Paz, Robert Nishihara, Soumith Chintala, Bernhard Schölkopf, Léon Bottou (2017): Discovering Causal Signals in Images [https://arxiv.org/abs/1605.08179]
8. Judea Pearl (2009): Causality: Models, Reasoning, and Inference [http://bayes.cs.ucla.edu/BOOK-2K/]
9. Judea Pearl (2018): The Seven Pillars of Causal Reasoning with Reflections on Machine Learning [http://ftp.cs.ucla.edu/pub/stat_ser/r481.pdf]
10. Jonas Peters, Joris Mooij, Dominik Janzing, Bernhard Schölkopf (2014): Causal Discovery with Continuous Additive Noise Models [https://arxiv.org/abs/1309.6779]
11. Jonas Peters, Peter Bühlmann, Nicolai Meinshausen (2016): Causal inference using invariant prediction: identification and confidence intervals [https://arxiv.org/abs/1501.01332]
12. Jonas Peters, Dominik Janzing, Bernhard Schölkopf (2017): Elements of Causal Inference: Foundations and Learning Algorithms [https://mitpress.mit.edu/books/elements-causal-inference]
13. Peter Spirtes, Clark Glymour, Richard Scheines (2001): Causation, Prediction, and Search [http://cognet.mit.edu/book/causation-prediction-and-search]
14. Bob L. Sturm (2016): The HORSE conferences [http://c4dm.eecs.qmul.ac.uk/horse2016/, http://c4dm.eecs.qmul.ac.uk/horse2017/]
15. Dustin Tran, David M. Blei (2017): Implicit Causal Models for Genome-wide Association Studies [https://arxiv.org/abs/1710.10742]
16. Michael Waldmann (2017): The Oxford Handbook of Causal Reasoning [https://global.oup.com/academic/product/the-oxford-handbook-of-causal-reasoning-9780198719871?cc=us]
17. James Woodward (2005): Making Things Happen: A Theory of Causal Explanation [https://global.oup.com/academic/product/making-things-happen-9780195189537?cc=us]
18. Antonio Torralba, Alyosha Efros (2011): Unbiased look at dataset

bias.
[\[http://people.csail.mit.edu/torralba/publications/datasets_cvpr11.pdf\]](http://people.csail.mit.edu/torralba/publications/datasets_cvpr11.pdf)

NIPS 2018 workshop on Compact Deep Neural Networks with industrial applications

Lixin Fan, Zhouchen Lin, Max Welling, Yurong Chen, Werner Bailer

Fri Dec 07, 08:00 AM

This workshop aims to bring together researchers, educators, practitioners who are interested in techniques as well as applications of making compact and efficient neural network representations. One main theme of the workshop discussion is to build up consensus in this rapidly developed field, and in particular, to establish close connection between researchers in Machine Learning community and engineers in industry. We believe the workshop is beneficial to both academic researchers as well as industrial practitioners.

News:

- . the workshop NIPS webpage:
<https://nips.cc/Conferences/2018/Schedule?showEvent=10941>
- . the workshop webpage is online at
<https://lixinfan01.wixsite.com/cdnria>
- . the workshop OpenReview submission site is online at
<https://openreview.net/group?id=NIPS.cc/2018/Workshop/CDNRIA>
- . Notice: the workshop submission deadline is changed to 20 Oct 2018 (see below) !
- . Please submit your abstract as soon as you can, so that reviewers will comment and discuss with you on OpenReview early.
- . A best paper award will be presented to the contribution selected by reviewers, who will also take into account active discussions on OpenReview. One FREE NIPS ticket will be awarded to the best paper.

Call for submissions:

We invite you to submit original work in, but not limited to, following areas:

Neural network compression techniques:

- . Binarization, quantization, pruning, thresholding and coding of neural networks
- . Efficient computation and acceleration of deep convolutional neural networks
- . Deep neural network computation in low power consumption applications (e.g., mobile or IoT devices)
- . Differentiable sparsification and quantization of deep neural networks
- . Benchmarking of deep neural network compression techniques

Neural network representation and exchange:

- . Exchange formats for (trained) neural networks
- . Efficient deployment strategies for neural networks
- . Industrial standardization of deep neural network representations
- . Performance evaluation methods of compressed networks in application context (e.g., multimedia encoding and processing)

Video & media compression methods using DNNs such as those developed in MPEG group:

- . To improve video coding standard development by using deep neural

networks
 . To increase practical applicability of network compression methods

An extended abstract (3 pages long using NIPS style, see <https://nips.cc/Conferences/2018/PaperInformation/StyleFiles>) in PDF format should be submitted for evaluation of the originality and quality of the work. The evaluation is double-blind and the abstract must be anonymous. References may extend beyond the 3 page limit, and parallel submissions to a journal or conferences (e.g. AAAI or ICLR) are permitted.

Submissions will be accepted as contributed talks (oral) or poster presentations. Extended abstract should be submitted through OpenReview (<https://openreview.net/group?id=NIPS.cc/2018/Workshop/CDNRIA>) by 20 Oct 2018. All accepted abstracts will be posted on the workshop website and archived.

Selection policy: all submitted abstracts will be evaluated based on their novelty, soundness and impacts. At the workshop we encourage DISCUSSION about NEW IDEAS, each submitter is thus expected to actively respond on OpenReview webpage and answer any questions about his/her ideas. The willingness to respond in OpenReview Q/A discussions will be an important factor for the selection of accepted oral or poster presentations.

Important dates:

- . Extended abstract submission deadline: 20 Oct 2018, (3:00 PM UTC)
- . Acceptance notification: 29 Oct. 2018, (1:00 PM UTC)
- . Camera ready submission: 12 November 2018, (3:00 PM UTC)
- . Workshop: 7 December 2018

Submission:

Please submit your extended abstract through OpenReview system (<https://openreview.net/group?id=NIPS.cc/2018/Workshop/CDNRIA>).

NIPS Complimentary workshop registration

We will help authors of accepted submissions to get access to a reserve pool of NIPS tickets. So please register to the workshop early.

Schedule

09:00 AM	Opening and Introduction	
09:05 AM	TBD 1	
09:30 AM	Bandwidth efficient deep learning by model compression	<i>Han</i>
09:55 AM	Neural network compression in the wild: why aiming for high compression factors is not enough	<i>Genewein</i>
10:20 AM	TBD 2	
10:45 AM	Coffee break (morning)	

11:00 AM	Network compression via differentiable pruning and quantization	<i>Louizos</i>
11:25 AM	Deep neural networks for multimedia processing, coding and standardization	
11:50 AM	TBD 3	
12:15 PM	Lunch break (on your own)	
02:00 PM	Efficient Computation of Deep Convolutional Neural Networks: A Quantization Perspective	<i>Cheng</i>
02:25 PM	Deep neural network compression and acceleration	<i>Yao</i>
02:50 PM	TBD 4	
03:15 PM	Coffee break (afternoon)	
03:30 PM	Poster presentations TBD	
04:30 PM	Panel discussion	
05:30 PM	Invited talk (TBD)	
05:55 PM	Closing	

Abstracts (6):

Abstract 3: **Bandwidth efficient deep learning by model compression in NIPS 2018 workshop on Compact Deep Neural Networks with industrial applications**, *Han* 09:30 AM

In the post-ImageNet era, computer vision and machine learning researchers are solving more complicated AI problems using larger datasets driving the demand for more computation. However, we are in the post-Moore's Law world where the amount of computation per unit cost and power is no longer increasing at its historic rate. This mismatch between supply and demand for computation highlights the need for co-designing efficient algorithms and hardware. In this talk, I will talk about bandwidth efficient deep learning by model compression, together with efficient hardware architecture support, saving memory bandwidth, networking bandwidth, and engineer bandwidth.

Abstract 4: **Neural network compression in the wild: why aiming for high compression factors is not enough in NIPS 2018 workshop on Compact Deep Neural Networks with industrial applications**, *Genewein* 09:55 AM

Abstract: the widespread use of state-of-the-art deep neural network models in the mobile, automotive and embedded domains is often hindered by the steep computational resources that are required for running such models. However, the recent scientific literature proposes a plethora of ways to alleviate the problem, either on the level of efficient network architectures, efficiency-optimized hardware or via network compression methods. Unfortunately, the usefulness of a network compression method strongly depends on the other aspects (network architecture and target hardware) as well as the task itself (classification, regression, detection, etc.), but very few publications consider this

interplay. This talk highlights some of the issues that arise from the strong interplay between network architecture, target hardware, compression algorithm and target task. Additionally some shortcomings in the current literature on network compression methods are pointed-out, such as incomparability of results (different base-line networks, different training-/data-augmentation schemes, etc.), lack of results on tasks other than classification, or use of very different (and perhaps not very informative) quantitative performance indicators such as naive compression rate, operations-per-second, size of stored weight matrices, etc. The talk concludes by proposing some guidelines and best-practices for increasing practical applicability of network compression methods and a call for standardizing network compression benchmarks.

Abstract 7: **Network compression via differentiable pruning and quantization in NIPS 2018 workshop on Compact Deep Neural Networks with industrial applications**, *Louizos* 11:00 AM

Abstract: neural network compression has become an important research area due to its great impact on deployment of large models on resource constrained devices. In this talk, we will introduce two novel techniques that allow for differentiable sparsification and quantization of deep neural networks; both of these are achieved via appropriate smoothing of the overall objective. As a result, we can directly train architectures to be highly compressed and hardware-friendly via off-the-self stochastic gradient descent optimizers.

Abstract 8: **Deep neural networks for multimedia processing, coding and standardization in NIPS 2018 workshop on Compact Deep Neural Networks with industrial applications**, 11:25 AM

Deep neural networks have proven to be an effective approach for many applications such as in computer vision, data mining and so on. Recent advances show that they may also be helpful for video, image and rich media processing and compression. This talk will introduce some work being carried in Tencent Media Lab which utilize deep neural networks to help improving audio/visual experiences, and some recent advances and on-going activities of using neural networks to help video compression in ITU-T and ISO/IEC video coding standard development and in general. Opportunities and challenges will be discussed from the perspective of industry applications.

Abstract 11: **Efficient Computation of Deep Convolutional Neural Networks: A Quantization Perspective in NIPS 2018 workshop on Compact Deep Neural Networks with industrial applications**, *Cheng* 02:00 PM

Deep neural networks have evolved remarkably over the past few years and they are currently the fundamental tools of many intelligent systems. At the same time, the computational complexity and resource consumption of these networks also continue to increase. This will pose a significant challenge to the deployment of such networks, especially in real-time applications or resource-limited devices. It is becoming a critical issue how to efficiently compute those networks, such as acceleration, compression. In this talk, we will first provide a brief introduction to network acceleration and compression, and then emphasize the efficient computation by quantization approach. Finally, we will introduce and discuss a few possible future directions

Abstract 12: **Deep neural network compression and acceleration in NIPS 2018 workshop on Compact Deep Neural Networks with industrial applications**, *Yao* 02:25 PM

In the past several years, Deep Neural Networks (DNNs) have demonstrated record-breaking accuracy on a variety of artificial intelligence tasks. However, the intensive storage and computational costs of DNN models make it difficult to deploy them on the mobile and embedded systems for real-time applications. In this technical talk, Dr. Yao will introduce their recent works on deep neural network compression and acceleration, showing how they achieve impressive compression performance without noticeable loss of model prediction accuracy, from the perspective of pruning and quantization.

Smooth Games Optimization and Machine Learning

Simon Lacoste-Julien, Ioannis Mitliagkas, Gauthier Gidel, Vasilis Syrgkanis, Eva Tardos, Leon Bottou, Sebastian Nowozin

Fri Dec 07, 08:00 AM

Overview

Advances in generative modeling and adversarial learning gave rise to a recent surge of interest in smooth two-players games, specifically in the context of learning generative adversarial networks (GANs). Solving these games raise intrinsically different challenges than the minimization tasks the machine learning community is used to. The goal of this workshop is to bring together the several communities interested in such smooth games, in order to present what is known on the topic and identify current open questions, such as how to handle the non-convexity appearing in GANs.

Background and objectives

A number of problems and applications in machine learning are formulated as games. A special class of games, smooth games, have come into the spotlight recently with the advent of GANs. In a two-players smooth game, each player attempts to minimize their differentiable cost function which depends also on the action of the other player. The dynamics of such games are distinct from the better understood dynamics of optimization problems. For example, the Jacobian of gradient descent on a smooth two-player game, can be non-symmetric and have complex eigenvalues. Recent work by ML researchers has identified these dynamics as a key challenge for efficiently solving similar problems.

A major hurdle for relevant research in the ML community is the lack of interaction with the mathematical programming and game theory communities where similar problems have been tackled in the past, yielding useful tools. While ML researchers are quite familiar with the convex optimization toolbox from mathematical programming, they are less familiar with the tools for solving games. For example, the extragradient algorithm to solve variational inequalities has been known in the mathematical programming literature for decades, however the ML community has until recently mainly appealed to gradient descent to optimize adversarial objectives.

The aim of this workshop is to provide a platform for both theoretical and applied researchers from the ML, mathematical programming and game theory community to discuss the status of our understanding on the interplay between smooth games, their applications in ML, as well existing tools and methods for dealing with them.

We also encourage, and will devote time during the workshop, on work

that identifies and discusses open, forward-looking problems of interest to the NIPS community.

Examples of topics of interest to the workshop are as follow:

- * Other examples of smooth games in machine learning (e.g. actor-critic models in RL).
- * Standard or novel algorithms to solve smooth games.
- * Empirical test of algorithms on GAN applications.
- * Existence and unicity results of equilibria in smooth games.
- * Can approximate equilibria have better properties than the exact ones ? [Arora 2017, Lipton and Young 1994].
- * Variational inequality algorithms [Harker and Pang 1990, Gidel et al. 2018].
- * Handling stochasticity [Hazan et al. 2017] or non-convexity [Gmarova et al. 2018] in smooth games.
- * Related topics from mathematical programming (e.g. bilevel optimization) [Pfau and Vinyals 2016].

Modeling and decision-making in the spatiotemporal domain

Ransalu Senanayake, Neal Jean, Fabio Ramos, Girish Chowdhary

Fri Dec 07, 08:00 AM

Abstract Understanding the evolution of a process over space and time is fundamental to a variety of disciplines. To name a few, such phenomena that exhibit dynamics in both space and time include propagation of diseases, variations in air pollution, dynamics in fluid flows, and patterns in neural activity. In addition to these fields in which modeling the nonlinear evolution of a process is the focus, there is also an emerging interest in decision-making and controlling of autonomous agents in the spatiotemporal domain. That is, in addition to learning what actions to take, when and where to take actions is crucial for an agent to efficiently and safely operate in dynamic environments. Although various modeling techniques and conventions are used in different application domains, the fundamental principles remain unchanged. Automatically capturing the dependencies between spatial and temporal components, making accurate predictions into the future, quantifying the uncertainty associated with predictions, real-time performance, and working in both big data and data scarce regimes are some of the key aspects that deserve our attention. Establishing connections between Machine Learning and Statistics, this workshop aims at;

- (1) raising open questions on challenges of spatiotemporal modeling and decision-making,
- (2) establishing connections among diverse application domains of spatiotemporal modeling, and
- (3) encouraging conversation between theoreticians and practitioners to develop robust predictive models.

Keywords

Theory: deep learning/convolutional LSTM, kernel methods, chaos theory, reinforcement learning for dynamic environments, dynamic policy learning, biostatistics, epidemiology, geostatistics, climatology, neuroscience, etc.

Applications:

Natural phenomena: disease propagation and outbreaks, environmental monitoring, climate modeling, etc.

Social and economics: predictive policing, population mapping, poverty mapping, food resources, agriculture, etc.

Engineering/robotics: active data collection, traffic modeling, motion prediction, fluid dynamics, spatiotemporal prediction for safe autonomous driving, etc.

Web: <https://sites.google.com/site/nips18spatiotemporal/>

Schedule

08:30 AM	Introduction	<i>Wikle</i>
03:00 PM	Decision-making	<i>Hsieh</i>

Workshop on Security in Machine Learning

Nicolas Papernot, Florian Tramer, Kamalika Chaudhuri, Matt Fredrikson, Jacob Steinhardt

Fri Dec 07, 08:00 AM

There is growing recognition that ML exposes new vulnerabilities in software systems. Some of the threat vectors explored so far include training data poisoning, adversarial examples or model extraction. Yet, the technical community’s understanding of the nature and extent of the resulting vulnerabilities remains limited. This is due in part to (1) the large attack surface exposed by ML algorithms because they were designed for deployment in benign environments---as exemplified by the IID assumption for training and test data, (2) the limited availability of theoretical tools to analyze generalization, (3) the lack of reliable confidence estimates. In addition, the majority of work so far has focused on a small set of application domains and threat models.

This workshop will bring together experts from the computer security and machine learning communities in an attempt to highlight recent work that contribute to address these challenges. Our agenda will complement contributed papers with invited speakers. The latter will emphasize connections between ML security and other research areas such as accountability or formal verification, as well as stress social aspects of ML misuses. We hope this will help identify fundamental directions for future cross-community collaborations, thus charting a path towards secure and trustworthy ML.

Schedule

09:15 AM	Invited Talk by Aditi Raghunathan	
11:00 AM	Keynote by danah boyd	<i>boyd</i>
01:30 PM	Invited Talk by Been Kim	<i>Kim</i>
02:15 PM	Invited Talk by Moustapha Cisse	
04:15 PM	Invited Talk by Marta Kwiatkowska	<i>Kwiatkowska</i>
04:45 PM	Invited Talk by Somesh Jha	<i>Jha</i>

Workshop on Ethical, Social and Governance Issues in AI

Chloe Bakalar, Sarah Bird, Tiberio Caetano, Edward W Felten, Dario Garcia, Isabel Kloumann, Finnian Lattimore, Sendhil Mullainathan, D. Sculley

Fri Dec 07, 08:00 AM

Abstract

Ethics is the philosophy of human conduct: It addresses the question “how should we act?” Throughout most of history the repertoire of actions available to us was limited and their consequences constrained in scope and impact through dispersed power structures and slow trade. Today, in our globalised and networked world, a decision can affect billions of people instantaneously and have tremendously complex repercussions. Machine learning algorithms are replacing humans in making many of the decisions that affect our everyday lives. How can we decide how machine learning algorithms and their designers should act? What is the ethics of today and what will it be in the future?

In this one day workshop we will explore the interaction of AI, society, and ethics through three general themes.

Advancing and Connecting Theory: How do different fairness metrics relate to one another? What are the trade-offs between them? How do fairness, accountability, transparency, interpretability and causality relate to ethical decision making? What principles can we use to guide us in selecting fairness metrics within a given context? Can we connect these principles back to ethics in philosophy? Are these principles still relevant today?

Tools and Applications: Real-world examples of how ethical considerations are affecting the design of ML systems and pipelines. Applications of algorithmic fairness, transparency or interpretability to produce better outcomes. Tools that aid identifying and or alleviating issues such as bias, discrimination, filter bubbles, feedback loops etc. and enable actionable exploration of the resulting trade-offs.

Regulation: With the GDPR coming into force in May 2018 it is the perfect time to examine how regulation can help (or hinder) our efforts to deploy AI for the benefit of society. How are companies and organisations responding to the GDPR? What aspects are working and what are the challenges? How can regulatory or legal frameworks be designed to continue to encourage innovation, so society as a whole can benefit from AI, whilst still providing protection against its harms.

This workshop is designed to be focused on some of the larger ethical issues related to AI and can be seen as a complement to the FATML proposal, which is focused more on fairness, transparency and accountability. We would be happy to link or cluster the workshops together, but we (us and the FATML organizers) think that there is more than 2 day worth of material that the community needs to discuss in the area of AI and ethics, so it would be great to have both workshops if possible.

Machine Learning for Geophysical & Geochemical Signals

Laura Pyrak-Nolte, Jim R Rustad, Richard Baraniuk

Fri Dec 07, 08:00 AM

Motivation

The interpretation of Earth's subsurface evolution from full waveform analysis requires a method to identify the key signal components related to the evolution in physical properties from changes in stress, fluids, geochemical interactions and other natural and anthropogenic processes. The analysis of seismic waves and other geophysical/geochemical signals remains for the most part a tedious task that geoscientists may perform by visual inspection of the available seismograms. The complexity and noisy nature of a broad array of geoscience signals combined with sparse and irregular sampling make this analysis difficult and imprecise. In addition, many signal components are ignored in tomographic imaging and continuous signal analysis that may prevent discovery of previously unrevealed signals that may point to new physics.

Ideally a detailed interpretation of the geometric contents of these data sets would provide valuable prior information for the solution of corresponding inverse problems. This unsatisfactory state of affairs is indicative of a lack of effective and robust algorithms for the computational parsing and interpretation of seismograms (and other geoscience data sets). Indeed, the limited frequency content, strong nonlinearity, temporally scattered nature of these signals make their analysis with standard signal processing techniques difficult and insufficient.

Once important seismic phases are identified, the next challenge is determining the link between a remotely-measured geophysical response and a characteristic property (or properties) of the fractures and fracture system. While a strong laboratory-based foundation has established a link between the mechanical properties of simple fracture systems (i.e. single fractures, parallel sets of fractures) and elastic wave scattering, bridging to the field scale faces additional complexity and a range of length scales that cannot be achieved from laboratory insight alone. This fundamental knowledge gap at the critical scale for long-term monitoring and risk assessment can only be narrowed or closed with the development of appropriate mathematical and numerical representations at each scale and across scales using multiphysics models that traverse spatial and temporal scales.

Topic

Major breakthroughs in bridging the knowledge gaps in geophysical sensing are anticipated as more researchers turn to machine learning (ML) techniques; however, owing to the inherent complexity of machine learning methods, they are prone to misapplication, may produce uninterpretable models, and are often insufficiently documented. This combination of attributes hinders both reliable assessment of model validity and consistent interpretation of model outputs. By providing documented datasets and challenging teams to apply fully documented workflows for ML approaches, we expect to accelerate progress in the application of data science to longstanding research issues in geophysics.

The goals of this workshop are to:

- (1) bring together experts from different fields of ML and geophysics to explore the use of ML techniques related to the identification of the physics contained in geophysical and chemical signals, as well as from images of geologic materials (minerals, fracture patterns, etc.); and
- (2) announce a set of geophysics machine learning challenges to the community that address earthquake detection and the physics of rupture

and the timing of earthquakes.

Target Audience

We aim to elicit new connections among these diverse fields, identify novel tools and models that can be transferred from one to the other, and explore novel ML applications that will benefit from ML algorithms paradigm. We believe that a successful workshop will lead to new research directions in a variety of areas and will also inspire the development of novel theories and tools.

Schedule

08:30 AM	Introduction	<i>Pyrak-Nolte, Rustad, Baraniuk</i>
08:40 AM	Paul Johnson	<i>Johnson</i>
09:05 AM	Greg Beroza, Mostafa Mousavi, and Weiqiang Zhu.	<i>Beroza</i>
09:30 AM	Maarten de Hoop	
09:55 AM	Karianne Jodine Bergen	<i>Bergen</i>
10:20 AM	Spotlight Posters	
10:40 AM	Coffee Break	
11:00 AM	Poster Session	
12:00 PM	Lunch	
02:00 PM	Bertrand Rouet-Leduc	<i>Rouet-Leduc</i>
02:20 PM	Joan Bruna	<i>Bruna</i>
02:40 PM	Claudia Hulbert	<i>Hulbert</i>
03:00 PM	Coffee Break	
03:30 PM	Ivan Dokmanic	<i>Dokmanic</i>
03:50 PM	Joe Morris	<i>Morris</i>
04:10 PM	Youzou Lin	<i>Lin</i>
04:30 PM	Panel Discussion	<i>Baraniuk</i>

Abstracts (9):

Abstract 1: **Introduction in Machine Learning for Geophysical & Geochemical Signals**, *Pyrak-Nolte, Rustad, Baraniuk* 08:30 AM

Introductory comments by organizers

Abstract 2: **Paul Johnson in Machine Learning for Geophysical & Geochemical Signals**, *Johnson* 08:40 AM

Probing Earthquake Fault Slip using Machine Learning

Earthquakes take place when two juxtaposed fault blocks are stressed sufficiently to overcome the frictional force holding them in place and they abruptly slip relative to each other. Earthquake faults exhibit a continuum of behaviors ranging from stick slip associated with strong shaking, to slow slip which is primarily aseismic, to very slow slip that is both aseismic and can take place over hours to months. We are characterizing faulting physics by analyzing with machine learning continuous acoustic data streams in the laboratory and continuous

seismic data streams in Earth. We use as labels characteristics of the measured fault slip behavior in the laboratory such as the fault friction, shear displacement and fault thickness. In Earth, we use surface displacement as determined by Global Positioning Systems (GPS). Other data data such as INSAR can be used as well. We find that the laboratory acoustic data and the Earth seismic data are a type of Rosetta Stone revealing fault characteristics at all times and fault displacements. This is a surprising observation because previously we believed most or much of the signal was noise. Here we describe an overview of recent work in this area and also describe recent efforts on parallel problems such as volcanoes and geysers.

Abstract 5: Karianne Jodine Bergen in Machine Learning for Geophysical & Geochemical Signals, *Bergen* 09:55 AM

Towards data-driven earthquake detection: Extracting weak seismic signals with locality-sensitive hashing

Extracting weak earthquake signals from continuous waveform data recorded by sensors in a seismic network is a fundamental and challenging task in seismology. In this talk, I will present Fingerprint and Similarity Thresholding (FAST; Yoon et al, 2015), a computationally efficient method for large-scale earthquake detection. FAST adapts technology used for rapid audio identification to the problem of extracting weak earthquake signals in continuous seismic data. FAST uses locality-sensitive hashing, a data mining technique for efficiently identifying similar items in large data sets, to detect similar waveforms (candidate earthquakes) in continuous seismic data. A distinguishing feature of our approach is that FAST is an unsupervised detector; FAST can discover new sources without any template waveforms or waveform characteristics available as training data – a common situation for seismic data sets. In our recent work, we have extended FAST to enable earthquake detection using data from multiple sensors spaced tens or hundreds of kilometers apart (Bergen and Beroza, 2018), and optimized the FAST software for detection at scale (Rong et al., 2018). FAST can now detect earthquakes with previously unknown sources in 10-year, multi-sensor seismic data sets without training data – a capability that was not previously available for seismic data analysis.

Abstract 10: Bertrand Rouet-Leduc in Machine Learning for Geophysical & Geochemical Signals, *Rouet-Leduc* 02:00 PM

Estimating the State of Faults from the Full Continuous Seismic Data Using Machine Learning

Nearly all aspects of earthquake rupture are controlled by the friction along the fault that progressively increases with tectonic forcing, but in general cannot be directly measured. Using machine learning, we show that instantaneous statistical characteristics of the seismic data are a fingerprint of the fault zone frictional state in laboratory experiments. Using a similar methodology in Earth, where we rely on other geophysical datasets as labels in order to extract informative signals from raw seismic waves, we show that subduction zones are continuously broadcasting a tremor-like signal that precisely informs of fault displacement rate throughout their slow earthquake slip cycle. We posit that this signal provides indirect, real-time access to frictional properties of megathrusts and may ultimately reveal a connection between slow slip and megaquakes

Abstract 11: Joan Bruna in Machine Learning for Geophysical & Geochemical Signals, *Bruna* 02:20 PM

Geometric Deep Learning for Many-Particle and non Euclidean Systems

Across many areas of science, one is required to process data defined on irregular and non-Euclidean domains. For example, in particle physics, measurements in the LHC are highly variable particle collisions with cylindrical calorimeters, whereas the IceCube detector looks for neutrinos using an irregular 3d array of sensors. Despite such non-Euclidean structure, many of these tasks satisfy essential geometric priors, such as stability to deformations. In this talk, I will describe a broad family of neural architectures that leverage such geometric priors to learn efficient models with provable stability. I will also describe recent and current progress on several applications including particle physics and inverse problems.

Abstract 12: Claudia Hulbert in Machine Learning for Geophysical & Geochemical Signals, *Hulbert* 02:40 PM

Machine Learning Reveals the Coupling Between Slow Slips and Major Earthquakes

The potential connection between slow slips and earthquakes of large magnitude in subduction zones remains an open question in seismology. Slow slips (earthquakes releasing energy over long periods of times, up to several months) have been observed preceding major earthquake ruptures, suggesting that they may couple to or evolve into a megaquake.

We rely on supervised machine learning algorithms to analyze vast amounts of continuous seismic data, with the goal of identifying hidden signals preceding earthquakes. We find that continuous seismic signals identified in our previous studies of slow slip events carry information about the timing of impending earthquakes of large magnitude. Our results suggest that large earthquakes occur almost systematically in the same phase of the slow slip cycle, and point to a systematic, large-scale coupling between slow slip events and major earthquakes.

Abstract 14: Ivan Dokmanic in Machine Learning for Geophysical & Geochemical Signals, *Dokmanic* 03:30 PM

I will present a new learning-based approach to ill-posed inverse problems. Instead of directly learning the ill-posed inverse mapping, we learn an ensemble of simpler mappings from the data to the projections of the unknown model into random low-dimensional subspaces. We choose structured subspaces of piecewise-constant images on random Delaunay triangulations. With this choice, the projected inverse maps are simpler to learn in terms of robustness and generalization error. We form the reconstruction by combining the estimated subspace projections. This allow us to address inverse problems with extremely sparse data and still get good reconstructions of the unknown geometry; it also makes our method robust against arbitrary data corruptions not seen during training. Further, it marginalizes the role of the training dataset which is essential for applications in geophysics where ground-truth datasets are exceptionally scarce.

Abstract 15: Joe Morris in Machine Learning for Geophysical & Geochemical Signals, *Morris* 03:50 PM

Towards Realtime Hydraulic Fracture Monitoring using Machine Learning and Distributed Fiber Sensing

Joseph Morris, Christopher Sherman, Robert Mellors, Frederick Ryerson, Charles Yu, Michael Messerly

Abstract: Hydraulic fracturing operations (“pumping jobs”) are typically planned well in advance and do not allow for on-the-fly modification of control parameters, such as pumping rate and viscosity enhancement, that can be used to optimize the efficacy of the operation. Monitoring technologies, such as microseismic, have enabled an iterative cycle where observations of one pumping job may influence the selection of parameters of subsequent jobs. However, the significant time lag introduced by data processing and interpretation means that the iterative cycle may take weeks. We seek to enable a future where data collected during a job enables actionable, realtime decision making. Recent advances in distributed acoustic sensor (DAS) technology have produced a source of abundant new data for monitoring processes in the subsurface. Because of the massive dataset size (TB per day), developing a machine learning approach for interpreting DAS data is essential for effective use, such as in operational situations, which require near-realtime results. In our work, we use the massively parallel multi-physics code GEOS to generate a catalog of synthetic DAS measurements that are typical of those recorded during the stimulation of a hydraulic fracture. We then relate physical observables in the model such as the extents of the generated fractures, fluid flow, and interactions with pre-existing rock fractures to the DAS. These data quantify the potential of DAS measurements for revealing subsurface processes in realtime. Determining how best to construct and train a neural network is challenging. We will present our specific approach to building a deep neural network, including the nature of the training data and subsequent success of the network in identifying features. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344.

Abstract 16: **Youzou Lin in Machine Learning for Geophysical & Geochemical Signals**, *Lin* 04:10 PM

Accurate and Efficient Seismic Waveform-Inversion with Convolutional Neural Networks

Seismic full-waveform inversion has become a promising tool for velocity estimation in complex geological structures. The traditional seismic full-waveform inversion problems are usually posed as nonlinear optimization problems. Solving full-waveform inversion can be computationally challenging for two major reasons. One is the expensive computational cost and the other is the issue of local minima. In this work, we develop an end-to-end data-driven inversion technique, called “InversionNet”, to learn a regression relationship from seismic waveform datasets to subsurface models. Specifically, we build a novel deep convolutional neural network with an encoder-decoder structure, where the encoder learns an abstract representation of the seismic data, which is then used by the decoder to produce a subsurface model. We further incorporate atrous convolutions in our network structure to account for contextual information from the subsurface model. We evaluate the performance of our InversionNet with synthetic seismic waveform data. The experiment results demonstrate that our InversionNet not only yields accurate inversion results but also produces almost real-time inversion.

Continual Learning

Razvan Pascanu, Yee Teh, Marc Pickett, Mark Ring

Fri Dec 07, 08:00 AM

Continual learning (CL) is the ability of a model to learn continually from a stream of data, building on what was learnt previously, hence exhibiting positive transfer, as well as being able to remember previously seen tasks. CL is a fundamental step towards artificial intelligence, as it allows the agent to adapt to a continuously changing environment, a hallmark of natural intelligence. It also has implications for supervised or unsupervised learning. For example, when the dataset is not properly shuffled or there exists a drift in the input distribution, the model overfits the recently seen data, forgetting the rest -- phenomena referred to as catastrophic forgetting, which is part of CL and is something CL systems aim to address.

Continual learning is defined in practice through a series of desiderata. A non-complete lists includes:

- * Online learning -- learning occurs at every moment, with no fixed tasks or data sets and no clear boundaries between tasks;
- * Presence of transfer (forward/backward) -- the model should be able to transfer from previously seen data or tasks to new ones, as well as possibly new task should help improve performance on older ones;
- * Resistance to catastrophic forgetting -- new learning does not destroy performance on previously seen data;
- * Bounded system size -- the model capacity should be fixed, forcing the system use its capacity intelligently as well as gracefully forgetting information such to ensure maximising future reward;
- * No direct access to previous experience -- while the model can remember a limited amount of experience, a continual learning algorithm should not have direct access to past tasks or be able to rewind the environment;

In the previous edition of the workshop the focus has been on defining a complete list of desiderata, of what a continual learning (CL) enabled system should be able to do. We believe that in this edition we should further constrain the discussion with a focus on how to evaluate CL and how it relates to other existing topics (e.g. life-long learning, transfer learning, meta-learning) and how ideas from these topics could be useful for continual learning.

Different aspects of continual learning are in opposition of each other (e.g. fixed model capacity and not-forgetting), which also raises the question of how to evaluate continual learning systems. One one hand, what are the right trade-offs between these different opposing forces? How do we compare existing algorithms given these different dimensions along which we should evaluate them (e.g. forgetting, positive transfer)? What are the right metrics we should report? On the other hand, optimal or meaningful trade-offs will be tightly defined by the data or at least type of tasks we use to test the algorithms. One prevalent task used by many recent papers is PermutedMNIST. But as MNIST is not a reliable dataset for classification, so PermutedMNIST might be extremely misleading for continual learning. What would be the right benchmarks, datasets or tasks for fruitfully exploiting this topic?

Finally, we will also encourage presentation of both novel approaches to CL and implemented systems, which will help concretize the discussion of what CL is and how to evaluate CL systems.

Bayesian Deep Learning

Yarin Gal, Jose Miguel Hernández-Lobato, Christos Louizos, Andrew Wilson, Zoubin Ghahramani, Kevin P Murphy, Max Welling

Fri Dec 07, 08:00 AM

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 over the past few years, with the introduction of new deep learning models that take advantage of Bayesian techniques, as well as Bayesian models that incorporate deep learning elements [1-11]. In fact, the use of Bayesian techniques in deep learning can be traced back to the 1990s', in seminal works by Radford Neal [12], David MacKay [13], and Dayan et al. [14]. These gave us tools to reason about deep models' confidence, and achieved state-of-the-art performance on many tasks. However earlier tools did not adapt when new needs arose (such as scalability to big data), and were consequently forgotten. Such ideas are now being revisited in light of new advances in the field, yielding many exciting new results.

Extending on the workshop's success from the past couple of years, this workshop will again study the advantages and disadvantages of the ideas above, and 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. The main theme this year will be applications of Bayesian deep learning in the real world, highlighting the requirements of practitioners from the research community. 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 an almost 100% increase in the number of submissions (75 submissions in total), attracting sponsorship from Google, Microsoft Ventures, Uber, and Qualcomm in the form of student travel awards.

Topics:

Probabilistic deep models for classification and regression (such as extensions and application of Bayesian neural networks),
 Generative deep models (such as variational autoencoders),
 Incorporating explicit prior knowledge in deep learning (such as posterior regularization with logic rules),
 Approximate inference for Bayesian deep learning (such as variational Bayes / expectation propagation / etc. in Bayesian neural networks),
 Scalable MCMC inference in Bayesian deep models,
 Deep recognition models for variational inference (amortized inference),
 Model uncertainty in deep learning,
 Bayesian deep reinforcement learning,
 Deep learning with small data,
 Deep learning in Bayesian modelling,
 Probabilistic semi-supervised learning techniques,
 Active learning and Bayesian optimization for experimental design,
 Applying non-parametric methods, one-shot learning, and Bayesian deep learning in general,
 Implicit inference,
 Kernel methods in Bayesian deep learning.

Call for papers:

A submission should take the form of an extended abstract (3 pages long) in PDF format using the NIPS style. Author names do not need to be anonymized and references (as well as appendices) may extend as far as needed beyond the 3 page upper limit. If research has previously appeared in a journal, workshop, or conference (including NIPS 2017 conference), the workshop submission should extend that previous work. Submissions will be accepted as contributed talks or poster presentations.

Related previous workshops:

Bayesian Deep Learning (NIPS 2017)
 Principled Approaches to Deep Learning (ICML 2017)
 Bayesian Deep Learning (NIPS 2016)
 Data-Efficient Machine Learning (ICML 2016)
 Deep Learning Workshop (ICML 2015, 2016)
 Deep Learning Symposium (NIPS 2015 symposium)
 Advances in Approximate Bayesian Inference (NIPS 2015)
 Black box learning and inference (NIPS 2015)
 Deep Reinforcement Learning (NIPS 2015)
 Deep Learning and Representation Learning (NIPS 2014)
 Advances in Variational Inference (NIPS 2014)

Visually grounded interaction and language

Florian Strub, Harm de Vries, Erik T Wijmans, Samyak Datta, Ethan Perez, Mateusz Malinowski, Stefan Lee, Peter Anderson, Aaron Courville, Jeremie MARY, Dhruv Batra, Devi Parikh, Olivier Pietquin, Chiori HORI, Tim Marks, Anoop Cheria

Fri Dec 07, 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, methods in this family suffer from limited semantic understanding, as they miss learning from the multimodal and interactive environment in which communication often takes place - the symbols of language thus are not grounded in anything concrete. The symbol grounding problem first highlighted this limitation, that "meaningless symbols (i.e.) words cannot be grounded in anything but other meaningless symbols" [18].

On the other hand, humans acquire language by communicating about and interacting within a rich, perceptual environment. This behavior provides the necessary grounding for symbols, i.e. to concrete objects or concepts (i.e. physical or psychological). Thus, recent work has aimed to bridge vision, interactive learning, and natural language understanding through language learning tasks based on natural images (ReferIt [1], GuessWhat?! [2], Visual Question Answering [3,4,5,6], Visual Dialog [7], Captioning [8]) or through embodied agents performing interactive tasks [13,14,17,22,23,24,26] in physically simulated environments (DeepMind Lab [9], Baidu XWorld [10], OpenAI Universe [11], House3D [20], Matterport3D [21], GIBSON [24], MINOS [25], AI2-THOR [19], StreetLearn [17]), often drawing on the recent successes of deep learning and reinforcement learning. We believe this line of research poses a promising, long-term solution to the grounding problem faced by

current, popular language understanding models.

While machine learning research exploring visually-grounded language learning may be in its earlier stages, it may be possible to draw insights from the rich research literature on human language acquisition. In neuroscience, recent progress in fMRI technology has enabled to better understand the interleave between language, vision and other modalities [15,16] suggesting that the brains shares neural representation of concepts across vision and language. Differently, developmental cognitive scientists have also argued that children acquiring various words is closely linked to them learning the underlying concept in the real world [12].

This workshop thus aims to gather people from various backgrounds - machine learning, computer vision, natural language processing, neuroscience, cognitive science, psychology, and philosophy - to share and debate their perspectives on why grounding may (or may not) be important in building machines that truly understand natural language.

We invite you to submit papers related to the following topics:

- language acquisition or learning through interactions
- visual captioning, dialog, and question-answering
- reasoning in language and vision
- visual synthesis from language
- transfer learning in language and vision tasks
- navigation in virtual worlds via natural-language instructions or multi-agent communication
- machine translation with visual cues
- novel tasks that combine language, vision and actions
- modeling of natural language and visual stimuli representations in the human brain
- position papers on grounded language learning
- audio visual scene-aware dialog

Submissions should be up to 4 pages excluding references, acknowledgements, and supplementary material, and should be in the NIPS format. We also welcome published papers that are within the scope of the workshop (without re-formatting). Please email your submission to vigilworkshop2018@gmail.com before or on November 1st 2018.

- [1] Sahar Kazemzadeh et al. "ReferItGame: Referring to Objects in Photographs of Natural Scenes." EMNLP, 2014.
- [2] Harm de Vries et al. "GuessWhat?! Visual object discovery through multi-modal dialogue." CVPR, 2017.
- [3] Stanislaw Antol et al. "Vqa: Visual question answering." ICCV, 2015.
- [4] Mateusz Malinowski et al. "Ask Your Neurons: A Neural-based Approach to Answering Questions about Images." ICCV, 2015.
- [5] Mateusz Malinowski et al. "A Multi-World Approach to Question Answering about Real-World Scenes based on Uncertain Input." NIPS, 2014.
- [6] Geman Donald, et al. "Visual Turing test for computer vision systems." PNAS, 2015.
- [7] Abhishek Das et al. "Visual dialog." CVPR, 2017.
- [8] Anna Rohrbach et al. "Generating Descriptions with Grounded and Co-Referenced People." CVPR, 2017.
- [9] Charles Beattie et al. Deepmind lab. arXiv, 2016.
- [10] Haonan Yu et al. "Guided Feature Transformation (GFT): A Neural Language Grounding Module for Embodied Agents." arXiv, 2018.
- [11] Openai universe. <https://universe.openai.com>, 2016.
- [12] Alison Gopnik et al. "Semantic and cognitive development in 15- to

21-month-old children." Journal of Child Language, 1984.

- [13] Abhishek Das et al. "Learning Cooperative Visual Dialog Agents with Deep Reinforcement Learning." ICCV, 2017.
- [14] Karl Moritz Hermann et al. "Grounded Language Learning in a Simulated 3D World." arXiv, 2017.
- [15] Alexander G. Huth et al. "Natural speech reveals the semantic maps that tile human cerebral cortex." Nature, 2016.
- [16] Alexander G. Huth, et al. "Decoding the semantic content of natural movies from human brain activity." Frontiers in systems neuroscience, 2016.
- [17] Piotr Mirowski et al. "Learning to Navigate in Cities Without a Map." arXiv, 2018.
- [18] Stevan Harnad. "The symbol grounding problem." CNLS, 1989.
- [19] E Kolve, R Mottaghi, D Gordon, Y Zhu, A Gupta, A Farhadi. "AI2-THOR: An Interactive 3D Environment for Visual AI." arXiv, 2017.
- [20] Yi Wu et al. "House3D: A Rich and Realistic 3D Environment." arXiv, 2017.
- [21] Angel Chang et al. "Matterport3D: Learning from RGB-D Data in Indoor Environments." arXiv, 2017.
- [22] Abhishek Das et al. "Embodied Question Answering." CVPR, 2018.
- [23] Peter Anderson et al. "Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments." CVPR, 2018.
- [24] Fei Xia et al. "Gibson Env: Real-World Perception for Embodied Agents." CVPR, 2018.
- [25] Manolis Savva et al. "MINOS: Multimodal indoor simulator for navigation in complex environments." arXiv, 2017.
- [26] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, Ali Farhadi. "IQA: Visual Question Answering in Interactive Environments." CVPR, 2018.

Modeling the Physical World: Learning, Perception, and Control

Jiajun Wu, Kelsey Allen, Kevin Smith, Jessica Hamrick, Emmanuel Dupoux, Marc Toussaint, Josh Tenenbaum

Fri Dec 07, 08:00 AM

Despite recent progress, AI is still far from achieving common-sense scene understanding and reasoning. A core component of this common sense is a useful representation of the physical world and its dynamics that can be used to predict and plan based on how objects interact. This capability is universal in adults, and is found to a certain extent even in infants. Yet despite increasing interest in the phenomenon in recent years, there are currently no models that exhibit the robustness and flexibility of human physical reasoning.

There have been many ways of conceptualizing models of physics, each with their complementary strengths and weaknesses. For instance, traditional physical simulation engines have typically used symbolic or analytic systems with "built-in" knowledge of physics, while recent connectionist methods have demonstrated the capability to learn approximate, differentiable system dynamics. While more precise, symbolic models of physics might be useful for long-term prediction and physical inference; approximate, differentiable models might be more practical for inverse dynamics and system identification. The design of a physical dynamics model fundamentally affects the ways in which that model can, and should, be used.

This workshop will bring together researchers in machine learning, computer vision, robotics, computational neuroscience, and cognitive psychology to discuss artificial systems that capture or model the physical world. It will also explore the cognitive foundations of physical representations, their interaction with perception, and their applications in planning and control. There will be invited talks from world leaders in the fields, presentations and poster sessions based on contributed papers, and a panel discussion.

Topics of discussion will include

- Building and learning physical models (deep networks, structured probabilistic generative models, physics engines)
- How to combine model-based and model-free approaches to physical prediction
- How to use physics models in higher-level tasks such as navigation, video prediction, robotics, etc.
- How perception and action interact with physical representations
- How cognitive science and computational neuroscience may inform the design of artificial systems for physical prediction
- Methodology for comparing models of infant learning with artificial systems
- Development of new datasets or platforms for physics and visual common sense

Dec. 8, 2018

NIPS 2018 Competition Track Day 2*Ralf Herbrich, Sergio Escalera***Sat Dec 08, 08:00 AM**

coming soon

Wordplay: Reinforcement and Language Learning in Text-based Games*Adam Trischler, Angeliki Lazaridou, Yonatan Bisk, Wendy Tay, Nate Kushman, Marc-Alexandre Côté, Alessandro Sordani, Daniel Ricks, Tom Zahavy, Hal Daumé III***Sat Dec 08, 08:00 AM**

Video games, via interactive learning environments like ALE [Bellemare et al., 2013], have been fundamental to the development of reinforcement learning algorithms that work on raw video inputs rather than featurized representations. Recent work has shown that text-based games may present a similar opportunity to develop RL algorithms for natural language inputs [Narasimhan et al., 2015, Haroush et al., 2018]. Drawing on insights from both the RL and NLP communities, this workshop will explore this opportunity, considering synergies between text-based and video games as learning environments as well as important differences and pitfalls.

Video games provide infinite worlds of interaction and grounding defined by simple, physics-like dynamics. While it is difficult, if not impossible, to simulate the full and social dynamics of linguistic interaction (see, e.g., work on user simulation and dialogue [Georgila et al., 2006, El Asri et al., 2016]), text-based games nevertheless present complex, interactive simulations that ground language in world and action semantics. Games like Zork [Infocom, 1980] rose to prominence in the age before advanced computer graphics. They use simple language to describe the state of the environment and to report the effects of player actions. Players interact with the environment through text commands that respect a predefined grammar, which, though simplistic, must be discovered in each game. Through sequential decision making, language understanding, and language generation, players work toward goals that may or may not be specified explicitly, and earn rewards (points) at completion or along the way.

Text-based games present a broad spectrum of challenges for learning algorithms. In addition to language understanding, successful play generally requires long-term memory and planning, exploration/experimentation, affordance extraction [Fulda et al., 2017], and common sense. Text games also highlight major open challenges for RL: the action space (text) is combinatorial and compositional, while game states are partially observable, since text is often ambiguous or underspecific. Furthermore, in text games the set of actions that affect

the state is not known in advance but must be learned through experimentation, typically informed by prior world/linguistic knowledge.

There has been a host of recent work towards solving text games [Narasimhan et al., 2015, Fulda et al., 2017, Kostka et al., 2017, Zhilin, et al., 2017, Haroush et al., 2018]. Nevertheless, commercial games like Zork remain beyond the capabilities of existing approaches. We argue that addressing even a subset of the aforementioned challenges would represent important progress in machine learning. Agents that solve text-based games may further learn functional properties of language; however, it is unclear what limitations the constraints and simplifications of text games (e.g., on linguistic diversity) impose on agents trained to solve them.

This workshop will highlight research that investigates existing or novel RL techniques for text-based settings, what agents that solve text-based games (might) learn about language, and more generally whether text-based games provide a good testbed for research at the intersection of RL and NLP. The program will feature a collection of invited talks alongside contributed posters and spotlight talks, curated by a committee with broad coverage of the RL and NLP communities. Panel discussions will highlight perspectives of influential researchers from both fields and encourage open dialogue. We will also pose a text-based game challenge several months in advance of the workshop (a similar competition is held annually at the IEEE Conference on Computational Intelligence and Games). This optional component will enable participants to design, train, and test agents in a carefully constructed, interactive text environment. The best-performing agent(s) will be recognized and discussed at the workshop. In addition to the exchange of ideas and the initiation of collaboration, an expected outcome is that text-based games emerge more prominently as a benchmark task to bridge RL and NLP research.

Relevant topics to be addressed at the workshop include (but are not limited to):

- RL in compositional, combinatorial action spaces
- Open RL problems that are especially pernicious in text-based games, like (sub)goal identification and efficient experimentation
- Grounded language understanding
- Online language acquisition
- Affordance extraction (on the fly)
- Language generation and evaluation in goal-oriented settings
- Automatic or crowdsourcing methods for linguistic diversity in simulations
- Use of language to constrain or index RL policies [Andreas et al., 2017]

Integration of Deep Learning Theories*Richard Baraniuk, Anima Anandkumar, Stephane Mallat, Ankit B Patel, nhmt H***Sat Dec 08, 08:00 AM**

Deep learning has driven dramatic performance advances on numerous difficult machine learning tasks in a wide range of applications. Yet, its theoretical foundations remain poorly understood, with many more questions than answers. For example: What are the modeling assumptions underlying deep networks? How well can we expect deep networks to perform? When a certain network succeeds or fails, can we

determine why and how? How can we adapt deep learning to new domains in a principled way?

While some progress has been made recently towards a foundational understanding of deep learning, most theory work has been disjointed, and a coherent picture has yet to emerge. Indeed, the current state of deep learning theory is like the fable “The Blind Men and the Elephant”.

The goal of this workshop is to provide a forum where theoretical researchers of all stripes can come together not only to share reports on their individual progress but also to find new ways to join forces towards the goal of a coherent theory of deep learning. Topics to be discussed include:

- Statistical guarantees for deep learning models
- Expressive power and capacity of neural networks
- New probabilistic models from which various deep architectures can be derived
- Optimization landscapes of deep networks
- Deep representations and invariance to latent factors
- Tensor analysis of deep learning
- Deep learning from an approximation theory perspective
- Sparse coding and deep learning
- Mixture models, the EM algorithm, and deep learning

In addition to invited and contributed talks by leading researchers from diverse backgrounds, the workshop will feature an extended poster/discussion session and panel discussion on which combinations of ideas are most likely to move theory of deep learning forward and which might lead to blind alleys.

Machine Learning for Systems

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

Sat Dec 08, 08:00 AM

This workshop is part two of a two-part series with one day focusing on Machine Learning for Systems and the other on Systems for Machine Learning. 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.

The Systems for Machine Learning workshop focuses on designing systems to enable ML, whereas we focus on developing ML to optimize systems. Both fields are mature enough to warrant a dedicated workshop. Organizers on both sides are open to merging in the future, but this year we plan to run them separately on two different days.

Designing specialized hardware and systems for deep learning is a topic that has received significant research attention, both in industrial and academic settings, leading to exponential increases in compute capability in GPUs and accelerators. However, using machine learning to optimize and accelerate software and hardware systems is a lightly explored but promising field, with broad implications for computing as a whole. Very recent work has outlined a broad scope where deep learning vastly outperforms traditional heuristics, including topics such as:

scheduling [1], data structure design [2], microarchitecture [3], compilers [4], and control of warehouse scale computing systems [5].

The focus of this workshop is to expand upon this recent work and build a community focused on using machine learning in computer systems problems. We seek to improve the state of the art in the areas where learning has already proven to perform better than traditional heuristics, as well as expand to new areas throughout the system stack such as hardware/circuit design and operating/runtime systems.

By forming a community of academic and industrial researchers who are excited about this area, we seek to build towards intelligent, self optimizing systems and answer questions such as: How do we generate and share high quality datasets that span the layers of the system stack? Which learned representations best represent code performance and runtime? Which simulators and simulation methodologies provide a tractable proving ground for techniques like reinforcement learning?

To this end, the target audience for this workshop includes a wide variety of attendees from state-of-the-art researchers in machine learning to domain experts in computer systems design. We have invited a broad set of expert speakers to present the potential for impact of combining machine learning research with computer systems. We hope that providing a formal venue for researchers from both fields to meet and interact will push forward both fundamental research in ML as well as real-world impact to computer systems design and implementation.

The workshop will host 6 speakers/panelists (all confirmed) and we will put out a call for researchers to submit relevant papers, up to 4 pages in the default NIPS style, that will undergo a peer review process. Selected works will be presented as spotlights, contributed talks and/or posters. Speakers will be invited to participate in an interactive panel discussion to conclude the workshop.

The organizers of this workshop span core research in machine learning, computer systems and architecture, as well as their intersection. Jointly, they have published in top-tier systems and machine learning conferences including: NIPS, ICML, ICLR, ISCA, MICRO, DAC, and SIGMETRICS.

References:

- [1] Device Placement Optimization with Reinforcement Learning, <https://arxiv.org/pdf/1706.04972.pdf>
- [2] The Case for Learned Index Structures, <https://arxiv.org/abs/1712.01208>
- [3] Learning Memory Access Patterns, <https://arxiv.org/pdf/1803.02329.pdf>
- [4] End to End Deep Learning of Optimization Heuristics: <https://ieeexplore.ieee.org/document/8091247/?reload=true>
- [5] <https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/>
- [6] Bayesian optimization for tuning the JVM, <https://www.youtube.com/watch?v=YhNI468S8CI>
- [7] Safe Exploration for Identifying Linear Systems via Robust Optimization: <https://arxiv.org/abs/1711.11165>

Relational Representation Learning

Aditya Grover, Paroma Varma, Fred Sala, Steven Holtzen, Jennifer Neville, Stefano Ermon, Chris Ré

Sat Dec 08, 08:00 AM

Relational reasoning, *i.e.*, learning and inference with relational data, is key to understanding how objects interact with each other and give rise to complex phenomena in the everyday world. Well-known applications include knowledge base completion and social network analysis. Although many relational datasets are available, integrating them directly into modern machine learning algorithms and systems that rely on continuous, gradient-based optimization and make strong i.i.d. assumptions is challenging. Relational representation learning has the potential to overcome these obstacles: it enables the fusion of recent advancements like deep learning and relational reasoning to learn from high-dimensional data. Success of such methods can facilitate novel applications of relational reasoning in areas like scene understanding, visual question-answering, reasoning over chemical and biological domains, program synthesis and analysis, and decision-making in multi-agent systems.

How should we rethink classical representation learning theory for relational representations? Classical approaches based on dimensionality reduction techniques such as isoMap and spectral decompositions still serve as strong baselines and are slowly paving the way for modern methods in relational representation learning based on random walks over graphs, message-passing in neural networks, group-invariant deep architectures etc. amongst many others. How can systems be designed and potentially deployed for large scale representation learning? What are promising avenues, beyond traditional applications like knowledge base and social network analysis, that can benefit from relational representation learning?

This workshop aims to bring together researchers from both academia and industry interested in addressing various aspects of representation learning for relational reasoning. Topics include, but are not limited to:

- * Algorithmic approaches. E.g., probabilistic generative models, message-passing neural networks, embedding methods, dimensionality reduction techniques, group-invariant architectures etc. for relational data
- * Theoretical aspects. E.g., when and why do learned representations aid relational reasoning? How does the non-i.i.d. nature of relational data conflict with our current understanding of representation learning?
- * Optimization and scalability challenges due to the inherent discreteness and curse of dimensionality of relational datasets
- * Evaluation of learned relational representations
- * Security and privacy challenges
- * Domain-specific applications
- * Any other topic of interest

NIPS Workshop on Machine Learning for Intelligent Transportation Systems 2018

Li Erran Li, Anca Dragan, Juan Carlos Niebles, Silvio Savarese

Sat Dec 08, 08:00 AM

Our transportation systems are poised for a transformation as we make progress on autonomous vehicles, vehicle-to-vehicle (V2V) and

vehicle-to-everything (V2X) communication infrastructures, and smart road infrastructures (like smart traffic lights). But many challenges stand in the way of this transformation. For example, how do we make perception accurate and robust enough to accomplish safe autonomous driving? How do we generate policies that equip autonomous cars with adaptive human negotiation skills when merging, overtaking, or yielding? How do we decide when a system is safe enough to deploy? And how do we optimize efficiency through intelligent traffic management and control of fleets?

To meet these requirements in safety, efficiency, control, and capacity, the systems must be automated with intelligent decision making. Machine learning will be an essential component of that. Machine learning has made rapid progress in the self-driving domain (e.g., in real-time perception and prediction of traffic scenes); has started to be applied to ride-sharing platforms such as Uber (e.g., demand forecasting); and by crowd-sourced video scene analysis companies such as Nexar (e.g., understanding and avoiding accidents). But to address the challenges arising in our future transportation system, we need to consider the transportation systems as a whole rather than solving problems in isolation, from prediction, to behavior, to infrastructure.

The goal of this workshop is to bring together researchers and practitioners from all areas of intelligent transportations systems to address core challenges with machine learning. These challenges include, but are not limited to pedestrian detection, intent recognition, and negotiation, coordination with human-driven vehicles, machine learning for object tracking, unsupervised representation learning for autonomous driving, deep reinforcement learning for learning driving policies, cross-modal and simulator to real-world transfer learning, scene classification, real-time perception and prediction of traffic scenes, uncertainty propagation in deep neural networks, efficient inference with deep neural networks predictive modeling of risk and accidents through telematics, modeling, simulation and forecast of demand and mobility patterns in large scale urban transportation systems, machine learning approaches for control and coordination of traffic leveraging V2V and V2X infrastructures,

The workshop will include invited speakers, panels, presentations of accepted papers, and posters. We invite papers in the form of short, long, and position papers to address the core challenges mentioned above. We encourage researchers and practitioners on self-driving cars, transportation systems and ride-sharing platforms to participate. Since this is a topic of broad and current interest, we expect at least 150 participants from leading university researchers, auto-companies and ride-sharing companies.

This will be the 3rd NIPS workshop in this series. Previous workshops have been very successful and have attracted large numbers of participants from both academia and industry.

Schedule

09:00 AM	Invited talk: Alfredo Canziani, NYU	<i>Canziani</i>
----------	--	-----------------

09:30 AM	Invited Talk: Drew Bagnell, CMU and Aurora
10:00 AM	Invited Talk: Alex Bayen, UC Berkeley
10:30 AM	Coffee break: morning
11:00 AM	Invited Talk: Nathaniel Fairfield, Waymo
02:00 PM	Invited Talk: Marco Pavone, Stanford <i>Pavone</i>
03:00 PM	Coffee break: afternoon
03:30 PM	Invited Talk: Ingmar Posner, Oxford <i>Posner</i>
04:00 PM	Invited Talk: Dorsa Sadigh, Stanford <i>Sadigh</i>
04:30 PM	Invited Talk: Ekaterina Taralova and Sarah Tariq, Zoox <i>Taralova, Tariq</i>

Machine Learning for Health (ML4H): Moving beyond supervised learning in healthcare

Andrew Beam, Tristan Naumann, Marzyeh Ghassemi, Matthew McDermott, Madalina Fiterau, Irene Y Chen, Brett Beaulieu-Jones, Mike Hughes, Farah Shamout, Corey Chivers, Jaz Kandola, Alexandre Yahi, Sam G Finlayson, Bruno Jedynak, Peter Schulam, Natalia Antropova, Jason Fries, Adrian Dalca

Sat Dec 08, 08:00 AM

Machine learning has had many notable successes within healthcare and medicine. However, nearly all such successes to date have been driven by supervised learning techniques. As a result, many other important areas of machine learning have been neglected and under appreciated in healthcare applications. In this workshop, we will convene a diverse set of leading researchers who are pushing beyond the boundaries of traditional supervised approaches. Attendees at the workshop will gain an appreciation for problems that are unique to healthcare and a better understanding of how machine learning techniques, including clustering, active learning, dimensionality reduction, reinforcement learning, causal inference, and others, may be leveraged to solve important clinical problems.

This year's program will also include spotlight presentations and two poster sessions highlighting novel research contributions at the intersection of machine learning and healthcare. We will invite submission of two page abstracts (not including references) for poster contributions. Topics of interest include but are not limited to models for diseases and clinical data, temporal models, Markov decision processes for clinical decision support, multiscale data-integration, modeling with missing or biased data, learning with non-stationary data, uncertainty and uncertainty propagation, non i.i.d. structure in the data, critique of models, interpretable models, causality, model biases, transfer learning, and incorporation of non-clinical (e.g., socioeconomic) factors.

The broader goal of the NIPS 2018 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 import problems in medicine and healthcare.

Second Workshop on Machine Learning for Creativity and Design

Luba Elliott, Sander Dieleman, Rebecca Fiebrink, Jesse Engel, Adam Roberts, Tom White

Sat Dec 08, 08:00 AM

Over the past few years, generative machine learning and machine creativity have continued 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 sketch-rnn and the Universal Music Translation Network. 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 and new designs, 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 building tools for better "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.

Background

In 2016, DeepMind's AlphaGo made two moves against Lee Sedol that were described by the Go community as "brilliant," "surprising," "beautiful," and so forth. Moreover, there was little discussion surrounding the fact that these very creative moves were actually made by a machine; it was enough that they were great examples of go playing. At the same time, the general public showed more concern for other applications of generative models. Algorithms that allow for convincing voice style transfer (Lyrebird) or puppet-like video face control (Face2Face) have raised ethical concerns that generative ML will be used to make convincing forms of fake news

Balancing this, the arts and music worlds have positively embraced generative models. Starting with DeepDream and expanding with image and video generation advances (e.g. GANs) we've seen lots of new and interesting art and music technologies provided by the machine learning community. We've seen research projects like Google Brain's Magenta, Sony CSL's FlowMachines and IBM's Watson undertake collaborations and attempt to build tools and ML models for use by these communities.

Research

Recent advances in generative models enable new possibilities in art

and music production. Language models can be used to write science fiction film scripts (Sunspring), theatre plays (Beyond the Fence) and even replicate the style of individual authors (Deep Tingle). Generative models for image and video allow us to create visions of people, places and things that resemble the distribution of actual images (GANs etc). Sequence modelling techniques have opened up the possibility of generating realistic musical scores (MIDI generation etc) and even raw audio that resembles human speech and physical instruments (DeepMind's WaveNet, MILA's Char2Wav and Google's NSynth). In addition, sequence modelling allows us to model vector images to construct stroke-based drawings of common objects according to human doodles (sketch-rnn). Lately, domain transfer techniques (FAIR's Universal Music Translation Network) have enabled the translation of music across musical instruments, genres, and styles.

In addition to field-specific research, a number of papers have come out that are directly applicable to the challenges of generation and evaluation such as learning from human preferences (Christiano et al., 2017) and CycleGAN. The application of Novelty Search (Stanley), evolutionary complexification (Stanley - CPPN, NEAT, Nguyen et al - Plug&Play GANs, Innovation Engine) and intrinsic motivation (Oudeyer et al 2007, Schmidhuber on Fun and Creativity) techniques, where objective functions are constantly evolving, is still not common practice in art and music generation using machine learning.

Another focus of the workshop is how to better enable human influence over generative models. This could include learning from human preferences, exposing model parameters in ways that are understandable and relevant to users in a given application domain (e.g., similar to Morris et al. 2008), enabling users to manipulate models through changes to training data (Fiebrink et al. 2011), allowing users to dynamically mix between multiple generative models (Akten & Grierson 2016), or other techniques. Although questions of how to make learning algorithms controllable and understandable to users are relatively nascent in the modern context of deep learning and reinforcement learning, such questions have been a growing focus of work within the human-computer interaction community (e.g., examined in a CHI 2016 workshop on Human-Centred Machine Learning), and the AI Safety community (e.g. Christiano et al. 2017, using human preferences to train deep reinforcement learning systems). Such considerations also underpin the new Google "People + AI Research" (PAIR) initiative.

Artists and Musicians

All the above techniques improve our capabilities of producing text, sound and images and have helped popularise the themes of machine learning and artificial intelligence in the art world with a number of art exhibitions (ZKM's Open Codes, Frankfurter Kunstverein's I am here to learn, NRW Forum's Pendoran Vinci) and media art festivals (Impakt Festival 2018 Algorithmic Superstructures, Retune 2016) dedicated to the topic.

Art and music that stands the test of time however requires more than generative capabilities. Recent research includes a focus on novelty in creative adversarial networks (Elgammal et al., 2017) and considers how generative algorithms can integrate into human creative processes, supporting exploration of new ideas as well as human influence over generated content (Akten & Grierson 2016a, 2016b). Artists including Mario Klingemann, Roman Lipski, Mike Tyka, and Memo Akten have further contributed to this space of work by creating artwork that compellingly demonstrates capabilities of generative algorithms, and by publicly reflecting on the artistic affordances of these new tools. Other

artists such as Mimi Onuoha, Caroline Sindors, and Adam Harvey have explored the ethical dimensions of machine learning technologies, reflecting on the issues of biased datasets and facial recognition.

The goal of this workshop is to bring together researchers interested in advancing art and music generation to present new work, foster collaborations and build networks.

In this workshop, we are particularly interested in how the following can be used in art and music generation: reinforcement learning, generative adversarial networks, novelty search and evaluation as well as learning from user preferences. We welcome submissions of short papers, demos and extended abstracts related to the above.

Like last year, there will be an open call for a display of artworks incorporating machine learning techniques. The exhibited works serve as a separate and more personal forum for collecting and sharing some of the latest creative works incorporating machine learning techniques with the NIPS community.

Schedule

08:45 AM	Kenneth Stanley	<i>Stanley</i>
09:15 AM	Yaroslav Ganin	<i>Ganin</i>
11:00 AM	Allison Parrish	<i>Parrish</i>
02:00 PM	Yaniv Taigman	<i>Taigman</i>

Machine Learning for Molecules and Materials

Jose Miguel Hernández-Lobato, Klaus-Robert Müller, Brooks Paige, Matt Kusner, Stefan Chmiela, Kristof Schütt

Sat Dec 08, 08:00 AM

The success of machine learning has been demonstrated time and time again in classification, generative modelling, and reinforcement learning. This revolution in machine learning has largely been in domains with at least one of two key properties: (1) the input space is continuous, and thus classifiers and generative models are able to smoothly model unseen data that is 'similar' to the training distribution, or (2) it is trivial to generate data, such as in controlled reinforcement learning settings such as Atari or Go games, where agents can re-play the game millions of times.

Unfortunately there are many important learning problems in chemistry, physics, materials science, and biology that do not share these attractive properties, problems where the input is molecular or material data.

Accurate prediction of atomistic properties is a crucial ingredient toward rational compound design in chemical and pharmaceutical industries. Many discoveries in chemistry can be guided by screening large databases of computational molecular structures and properties, but high level quantum-chemical calculations can take up to several days per molecule or material at the required accuracy, placing the ultimate achievement of in silico design out of reach for the foreseeable future. In large part the current state of the art for such problems is the expertise of individual researchers or at best highly-specific rule-based heuristic systems. Efficient methods in machine learning, applied to the prediction

of atomistic properties as well as compound design and crystal structure prediction, can therefore have pivotal impact in enabling chemical discovery and foster fundamental insights.

Because of this, in the past few years there has been a flurry of recent work towards designing machine learning techniques for molecule and material data [1-38]. These works have drawn inspiration from and made significant contributions to areas of machine learning as diverse as learning on graphs to models in natural language processing. Recent advances enabled the acceleration of molecular dynamics simulations, contributed to a better understanding of interactions within quantum many-body system and increased the efficiency of density based quantum mechanical modeling methods. This young field offers unique opportunities for machine learning researchers and practitioners, as it presents a wide spectrum of challenges and open questions, including but not limited to representations of physical systems, physically constrained models, manifold learning, interpretability, model bias, and causality.

The goal of this workshop is to bring together researchers and industrial practitioners in the fields of computer science, chemistry, physics, materials science, and biology all working to innovate and apply machine learning to tackle the challenges involving molecules and materials. In a highly interactive format, we will outline the current frontiers and present emerging research directions. We aim to use this workshop as an opportunity to establish a common language between all communities, to actively discuss new research problems, and also to collect datasets by which novel machine learning models can be benchmarked. The program is a collection of invited talks, alongside contributed posters. A panel discussion will provide different perspectives and experiences of influential researchers from both fields and also engage open participant conversation. An expected outcome of this workshop is the interdisciplinary exchange of ideas and initiation of collaboration.

Call for papers:

The 1 day NIPS 2018 Workshop on Machine Learning for Molecules and Materials is calling for contributions on theoretical models, empirical studies, and applications of machine learning for molecules and materials. We also welcome challenge papers on possible applications or datasets. Topics of interest (though not exhaustive) include: chemoinformatics, applications of deep learning to predict molecular properties, drug-discovery and material design, retrosynthesis and synthetic route prediction, modeling and prediction of chemical reaction data, and the analysis of molecular dynamics simulations. We invite submissions that either address new problems and insights for chemistry and quantum physics or present progress on established problems. The workshop includes a poster session, giving the opportunity to present novel ideas and ongoing projects. Submissions should be no longer than 10 pages in any format. Please email all submissions to: nips2018moleculesworkshop@gmail.com

References

- [1] Behler, J., Lorenz, S., Reuter, K. (2007). Representing molecule-surface interactions with symmetry-adapted neural networks. *J. Chem. Phys.*, 127(1), 07B603.
- [2] Behler, J., Parrinello, M. (2007). Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.*, 98(14), 146401.
- [3] Kang, B., Ceder, G. (2009). Battery materials for ultrafast charging and discharging. *Nature*, 458(7235), 190.
- [4] Bartók, A. P., Payne, M. C., Kondor, R., Csányi, G. (2010). Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Phys. Rev. Lett.*, 104(13), 136403.
- [5] Behler, J. (2011). Atom-centered symmetry functions for constructing high-dimensional neural network potentials. *J. Chem. Phys.*, 134(7), 074106.
- [6] Behler, J. (2011). Neural network potential-energy surfaces in chemistry: a tool for large-scale simulations. *Phys. Chem. Chem. Phys.*, 13(40), 17930-17955.
- [7] Rupp, M., Tkatchenko, A., Müller, K.-R., von Lilienfeld, O. A. (2012). Fast and accurate modeling of molecular atomization energies with machine learning. *Phys. Rev. Lett.*, 108(5), 058301.
- [8] Snyder, J. C., Rupp, M., Hansen, K., Müller, K.-R., Burke, K. (2012). Finding density functionals with machine learning. *Phys. Rev. Lett.*, 108(25), 253002.
- [9] Montavon, G., Rupp, M., Gobre, V., Vazquez-Mayagoitia, A., Hansen, K., Tkatchenko, A., Müller, K.-R., von Lilienfeld, O. A. (2013). Machine learning of molecular electronic properties in chemical compound space. *New J. Phys.*, 15(9), 095003.
- [10] Hansen, K., Montavon, G., Biegler, F., Fazli, S., Rupp, M., Scheffler, M., Tkatchenko, A., Müller, K.-R. (2013). Assessment and validation of machine learning methods for predicting molecular atomization energies. *J. Chem. Theory Comput.*, 9(8), 3404-3419.
- [11] Bartók, A. P., Kondor, R., Csányi, G. (2013). On representing chemical environments. *Phys. Rev. B*, 87(18), 184115.
- [12] Schütt K. T., Glawe, H., Brockherde F., Sanna A., Müller K.-R., Gross E. K. U. (2014). How to represent crystal structures for machine learning: towards fast prediction of electronic properties. *Phys. Rev. B.*, 89(20), 205118.
- [13] Ramsundar, B., Kearnes, S., Riley, P., Webster, D., Konerding, D., Pande, V. (2015). Massively multitask networks for drug discovery. *arXiv preprint arXiv:1502.02072*.
- [14] Rupp, M., Ramakrishnan, R., & von Lilienfeld, O. A. (2015). Machine learning for quantum mechanical properties of atoms in molecules. *J. Phys. Chem. Lett.*, 6(16), 3309-3313.
- [15] V. Botu, R. Ramprasad (2015). Learning scheme to predict atomic forces and accelerate materials simulations., *Phys. Rev. B*, 92(9), 094306.
- [16] Hansen, K., Biegler, F., Ramakrishnan, R., Pronobis, W., von Lilienfeld, O. A., Müller, K.-R., Tkatchenko, A. (2015). Machine learning predictions of molecular properties: Accurate many-body potentials and nonlocality in chemical space. *J. Phys. Chem. Lett.*, 6(12), 2326-2331.
- [17] Alipanahi, B., Delong, A., Weirauch, M. T., Frey, B. J. (2015). Predicting the sequence specificities of DNA-and RNA-binding proteins by deep learning. *Nat. Biotechnol.*, 33(8), 831-838.
- [18] Duvenaud, D. K., Maclaurin, D., Aguilera-Iparraguirre, J., Gomez-Bombarelli, R., Hirzel, T., Aspuru-Guzik, A., Adams, R. P. (2015). Convolutional networks on graphs for learning molecular fingerprints. *NIPS*, 2224-2232.
- [19] Faber F. A., Lindmaa A., von Lilienfeld, O. A., Armiento, R. (2016). Machine learning energies of 2 million elpasolite (A B C 2 D 6) crystals. *Phys. Rev. Lett.*, 117(13), 135502.
- [20] Gomez-Bombarelli, R., Duvenaud, D., Hernandez-Lobato, J. M., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., Aspuru-Guzik, A. (2016). Automatic chemical design using a data-driven continuous representation of molecules. *arXiv preprint arXiv:1610.02415*.
- [21] Wei, J. N., Duvenaud, D., Aspuru-Guzik, A. (2016). Neural networks for the prediction of organic chemistry reactions. *ACS Cent. Sci.*, 2(10), 725-732.
- [22] Sadowski, P., Fooshee, D., Subrahmanya, N., Baldi, P. (2016). Synergies between quantum mechanics and machine learning in

reaction prediction. *J. Chem. Inf. Model.*, 56(11), 2125-2128.

[23] Lee, A. A., Brenner, M. P., Colwell L. J. (2016). Predicting protein-ligand affinity with a random matrix framework. *Proc. Natl. Acad. Sci.*, 113(48), 13564-13569.

[24] Behler, J. (2016). Perspective: Machine learning potentials for atomistic simulations. *J. Chem. Phys.*, 145(17), 170901.

[25] De, S., Bartók, A. P., Csányi, G., Ceriotti, M. (2016). Comparing molecules and solids across structural and alchemical space. *Phys. Chem. Chem. Phys.*, 18(20), 13754-13769.

[26] Schütt, K. T., Arbabzadah, F., Chmiela, S., Müller, K.-R., Tkatchenko, A. (2017). Quantum-chemical insights from deep tensor neural networks. *Nat. Commun.*, 8, 13890.

[27] Segler, M. H., Waller, M. P. (2017). Neural-symbolic machine learning for retrosynthesis and reaction prediction. *Chem. Eur. J.*, 23(25), 5966-5971.

[28] Kusner, M. J., Paige, B., Hernández-Lobato, J. M. (2017). Grammar variational autoencoder. *arXiv preprint arXiv:1703.01925*.

[29] Coley, C. W., Barzilay, R., Jaakkola, T. S., Green, W. H., Jensen K. F. (2017). Prediction of organic reaction outcomes using machine learning. *ACS Cent. Sci.*, 3(5), 434-443.

[30] Altae-Tran, H., Ramsundar, B., Pappu, A. S., Pande, V. (2017). Low data drug discovery with one-shot learning. *ACS Cent. Sci.*, 3(4), 283-293.

[31] Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., Dahl, G. E. (2017). Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*.

[32] Chmiela, S., Tkatchenko, A., Sauceda, H. E., Poltavsky, Igor, Schütt, K. T., Müller, K.-R. (2017). Machine learning of accurate energy-conserving molecular force fields. *Sci. Adv.*, 3(5), e1603015.

[33] Ju, S., Shiga T., Feng L., Hou Z., Tsuda, K., Shiomi J. (2017). Designing nanostructures for phonon transport via bayesian optimization. *Phys. Rev. X*, 7(2), 021024.

[34] Ramakrishnan, R, von Lilienfeld, A. (2017). Machine learning, quantum chemistry, and chemical space. *Reviews in Computational Chemistry*, 225-256.

[35] Hernandez-Lobato, J. M., Requeima, J., Pyzer-Knapp, E. O., Aspuru-Guzik, A. (2017). Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. *arXiv preprint arXiv:1706.01825*.

[36] Smith, J., Isayev, O., Roitberg, A. E. (2017). ANI-1: an extensible neural network potential with DFT accuracy at force field computational cost. *Chem. Sci.*, 8(4), 3192-3203.

[37] Brockherde, F., Li, L., Burke, K., Müller, K.-R. By-passing the Kohn-Sham equations with machine learning. *Nat. Commun.*, 8, 872.

[38] Schütt, K. T., Kindermans, P. J., Sauceda, H. E., Chmiela, S., Tkatchenko, A., Müller, K. R. (2017). SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. *NIPS* 30.

Schedule

08:40 AM	Invited Talk Session 1	<i>Noe</i>
11:00 AM	Invited Talk Session 2	<i>Marks, Isayev, Smidt, Thomas</i>
02:00 PM	Invited Talk Session 3	<i>Tkatchenko, Jaakkola, Wei</i>
03:30 PM	Invited Talk Session 4	<i>Clementi</i>

Medical Imaging meets NIPS

Ender Konukoglu, Ben Glocker, Hervé Lombaert, Marleen de Bruijine

Sat Dec 08, 08:00 AM

Medical imaging and radiology are facing a major crisis with an ever-increasing complexity and volume of data and immense economic pressure. With the current advances in imaging technologies and their widespread use, interpretation of medical images pushes human abilities to the limit with the risk of missing critical patterns of disease. 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 NIPS" 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 NIPS 2017

(<https://sites.google.com/view/med-nips-2017>). 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.

Different from last year and based on feedback from participants, we propose to implement two novelties.

1. The workshop will accept paper submissions and have oral presentations with a format that aims to foster in depth discussions of a few selected articles. We plan to implement a Program Committee who will be responsible for reviewing articles and initiating discussions. The abstract track organized last year has brought a significant number of submission and has clearly demonstrated an appetite for more.
2. Along the workshop, we will host a challenge on outlier detection in brain Magnetic Resonance Imaging (MRI), which is one of the main applications of advanced unsupervised learning algorithms and generative models in medical imaging. The challenge will highlight a problem where the machine learning community can have a huge impact. To facilitate the challenge and potential further research, we provide necessary pre-processed datasets to simplify the use of medical imaging data and lower data-related entry barrier. Data collection for this challenge is finalized and ethical approval for data sharing is in place. We plan to open the challenge as soon as acceptance of the workshop is confirmed.

Schedule

09:00 AM	Making the Case for using more Inductive Bias in Deep Learning	<i>Welling</i>
09:45 AM	The U-net does its job – so what next?	<i>Ronneberger</i>
02:00 PM	To be determined	<i>Roth</i>
02:45 PM	TBD	<i>Arbel</i>

Abstracts (2):

Abstract 2: The U-net does its job – so what next? in Medical Imaging meets NIPS, *Ronneberger* 09:45 AM

U-net based architectures have demonstrated very high performance in a wide range of medical image segmentation tasks, but a powerful segmentation architecture alone is only one part of building clinically applicable tools. In my talk I'll present three projects from the DeepMind Health Research team that address these challenges.

The first project, a collaboration with University College London Hospital, deals with the challenging task of the precise segmentation of radiosensitive head and neck anatomy in CT scans, an essential input for radiotherapy planning [1]. With a 3D U-net we reach a performance similar to human experts on the majority of anatomical classes. Beside some minor architectural adaptations, e.g. to tackle the large imbalance of foreground to background voxels, a substantial focus of the project was in generating a high-quality test set [2] where each scan was manually segmented by two independent experts. Furthermore we introduced a new surface based performance metric, the surface DSC [3], designed to be a better proxy for the expected performance in a real-world radiotherapy setting than existing metrics.

The second project, together with Moorfields Eye Hospital, developed a system that analyses 3D OCT (optical coherence tomography) eye scans to provide referral decisions for patients [4]. The performance was on par with world experts with over 20 years experience. We use two network ensembles to decouple the variations induced by the imaging system from the patient-to-patient variations. The first ensemble of 3D U-nets creates clinically interpretable device-independent tissue map hypotheses; the second (3D dense-net based) ensemble maps the tissue map hypotheses to the diagnoses and referral recommendation. Adaptation to a new scanning device type only needed sparse manual segmentations on 152 scans, while the diagnosis model (trained with 14,884 OCT scans) could be reused without changes.

The third project deals with the segmentation of ambiguous images [5]. This is of particular relevance in medical imaging where ambiguities can often not be resolved from the image context alone. We propose a combination of a U-net with a conditional variational autoencoder that is capable of efficiently producing an unlimited number of plausible segmentation map hypotheses for a given ambiguous image. We show that each hypothesis provides an overall consistent segmentation, and that the probabilities of these hypotheses are well calibrated.

[1] Nikolov et al. (2018) "Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy" (soon available on ArXiv)

[2] Dataset will be soon available at <https://github.com/deepmind/tcia-ct-scan-dataset>

[3] Implementation available at <https://github.com/deepmind/surface-distance>

[4] De Fauw, et al. (2018) "Clinically applicable deep learning for diagnosis and referral in retinal disease" Nature Medicine (in press). <https://doi.org/10.1038/s41591-018-0107-6> (fulltext available from <https://deepmind.com/blog/moorfields-major-milestone/>)

[5] Kohl, et al. (2018) "A Probabilistic U-Net for Segmentation of Ambiguous Images". NIPS 2018 (accepted). Preprint available at <https://arxiv.org/abs/1806.05034>

Abstract 3: **To be determined in Medical Imaging meets NIPS**, *Roth* 02:00 PM

TBD

Machine Learning for the Developing World (ML4D): Achieving sustainable impact

William Herlands, Maria De-Arteaga

Sat Dec 08, 08:00 AM

Global development experts are beginning to employ ML for diverse problems such as aiding rescue workers allocate resources during natural disasters, providing intelligent educational and healthcare services in regions with few human experts, and detecting corruption in government contracts. While ML represents a tremendous hope for accelerated development and societal change, it is often difficult to ensure that machine learning projects provide their promised benefit. The challenging reality in developing regions is that pilot projects disappear after a few years or do not have the same effect when expanded beyond the initial test site, and prototypes of novel methodologies are often never deployed.

At the center of this year's program is how to achieve sustainable impact of Machine Learning for the Developing World (ML4D). This one-day workshop will bring together a diverse set of participants from across the globe to discuss major roadblocks and paths to action. Practitioners and development experts will discuss essential elements for ensuring successful deployment and maintenance of technology in developing regions. Additionally, the workshop will feature cutting edge research in areas such as transfer learning, unsupervised learning, and active learning that can help ensure long-term ML system viability. Attendees will learn about contextual components to ensure effective projects, development challenges that can benefit from machine learning solutions, and how these problems can inspire novel machine learning research.

The workshop will include invited and contributed talks, a poster session of accepted papers, panel discussions, and breakout sessions tailored to the workshop theme. We welcome paper submissions focussing on core ML methodology addressing ML4D roadblocks, application papers that showcase successful examples of ML4D, and research that evaluates the societal impact of ML.

Reinforcement Learning under Partial Observability

Joni Pajarinen, Chris Amato, Pascal Poupart, David Hsu

Sat Dec 08, 08:00 AM

Reinforcement learning (RL) has succeeded in many challenging tasks such as Atari, Go, and Chess and even in high dimensional continuous domains such as robotics. Most impressive successes are in tasks where the agent observes the task features fully. However, in real world problems, the agent usually can only rely on partial observations. In real time games the agent makes only local observations; in robotics the agent has to cope with noisy sensors, occlusions, and unknown dynamics. Even more fundamentally, any agent without a full a priori

world model or without full access to the system state, has to make decisions based on partial knowledge about the environment and its dynamics.

Reinforcement learning under partial observability has been tackled in the operations research, control, planning, and machine learning communities. One of the goals of the workshop is to bring researchers from different backgrounds together. Moreover, the workshop aims to highlight future applications. In addition to robotics where partial observability is a well known challenge, many diverse applications such as wireless networking, human-robot interaction and autonomous driving require taking partial observability into account.

Partial observability introduces unique challenges: the agent has to remember the past but also connect the present with potential futures requiring memory, exploration, and value propagation techniques that can handle partial observability. Current model-based methods can handle discrete values and take long term information gathering into account while model-free methods can handle high-dimensional continuous problems but often assume that the state space has been created for the problem at hand such that there is sufficient information for optimal decision making or just add memory to the policy without taking partial observability explicitly into account.

In this workshop, we want to go further and ask among others the following questions.

- * How can we extend deep RL methods to robustly solve partially observable problems?
- * Can we learn concise abstractions of history that are sufficient for high-quality decision-making?
- * There have been several successes in decision making under partial observability despite the inherent challenges. Can we characterize problems where computing good policies is feasible?
- * Since decision making is hard under partial observability do we want to use more complex models and solve them approximately or use (inaccurate) simple models and solve them exactly? Or not use models at all?
- * How can we use control theory together with reinforcement learning to advance decision making under partial observability?
- * Can we combine the strengths of model-based and model-free methods under partial observability?
- * Can recent method improvements in general RL already tackle some partially observable applications which were not previously possible?
- * How do we scale up reinforcement learning in multi-agent systems with partial observability?
- * Do hierarchical models / temporal abstraction improve RL efficiency under partial observability?

Schedule

08:30 AM	Opening Remarks
08:40 AM	Invited Talk 1
09:05 AM	Invited Talk 2
09:30 AM	Contributed Talk 1
09:45 AM	Invited Talk 3
11:00 AM	Contributed Talk 2
11:15 AM	Invited Talk 4

11:40 AM	Spotlights & Poster Session
02:00 PM	Invited Talk 5
02:25 PM	Contributed Talk 3
02:40 PM	Invited Talk 6
03:35 PM	Invited Talk 7
04:00 PM	Panel Discussion
05:30 PM	Poster Session

Privacy Preserving Machine Learning

Aurélien Bellet, Adria Gascon, Niki Kilbertus, Olga Ohrimenko, Mariana Raykova, Adrian Weller

Sat Dec 08, 08:00 AM

Our workshop will focus on privacy preserving techniques for training, inference, and disclosure in large scale data analysis, both in the distributed and centralized settings.

We have observed increasing interest of the ML community in leveraging cryptographic techniques such as Multi-Party Computation (MPC) and Homomorphic Encryption (HE) for privacy preserving training and inference, as well as Differential Privacy (DP) for disclosure.

Simultaneously, the systems security and cryptography community has proposed various secure frameworks for ML. We will encourage both theory and application-oriented submissions exploring a range of approaches, including cryptographic, hardware-based, and DP-like techniques.

Topics of interest include

- secure multi-party computation techniques for ML
- homomorphic encryption techniques for ML
- hardware-based approaches to privacy preserving ML
- centralized and decentralized protocols for learning on encrypted data
- differential privacy: theory, applications, and implementations
- statistical notions of privacy including relaxations of differential privacy
- empirical and theoretical comparisons between different notions of privacy
- trade-offs between privacy and utility

We think it will be very valuable to have a forum to unify different perspectives and start a discussion about the relative merits of each approach. The workshop will also serve as a venue for networking people from different communities interested in this problem, and hopefully foster fruitful long-term collaboration.

The one day workshop will include talks by world-renowned experts both from the machine learning and the cryptography communities, who have made remarkable contributions to problems at the intersection of both fields. We have secured a £3000 sponsorship from The Alan Turing Institute to provide travel and registration stipends for some participants as well as to cover a dinner for invited speakers.

We started forming the program committee right before this submission with an emphasis on building an inclusive and diverse PC, and got a very positive reception. The following are our confirmed PC members so far:

- Pauline Anthony (Google)
- Borja Balle (Amazon)

- Keith Bonawitz (Google)
- Emiliano de Cristofaro (UCL)
- David Evans (University of Virginia)
- Irene Giacomelli (University of Wisconsin)
- Kim Laine (MSR)
- Catuscia Palamidessi (Ecole Polytechnique)
- Mijung Park (MPI for Intelligent Systems)
- Benjamin Rubinstein (University of Melbourne)
- Anand Sawarte (Rutgers University)
- Nigel Smart (KU Leuven)

AI for social good

Margaux Luck, Tristan Sylvain, Joseph Paul Cohen, Arsene Fansi Tchango, Valentine Goddard, Aurelie Helouis, Yoshua Bengio, Sam Greydanus, Cody Wild, Taras Kucherenko, Arya Farahi, Jonnie Penn, Sean McGregor, Mark Crowley, Abhishek Gupta, Kenny Chen, Myriam Côté

Sat Dec 08, 08:00 AM

AI for Social Good

Important information

[Workshop website](<https://aiforsocialgood.github.io/2018/>)

[Submission

website](<https://cmt3.research.microsoft.com/User/Login?ReturnUrl=%2FAISG2018>)

Abstract

The "AI for Social Good" will focus on social problems for which artificial intelligence has the potential to offer meaningful solutions. The problems we chose to focus on are inspired by the United Nations Sustainable Development Goals (SDGs), a set of seventeen objectives that must be addressed in order to bring the world to a more equitable, prosperous, and sustainable path. In particular, we will focus on the following areas: health, education, protecting democracy, urban planning, assistive technology for people with disabilities, agriculture, environmental sustainability, economic inequality, social welfare and justice. Each of these themes present opportunities for AI to meaningfully impact society by reducing human suffering and improving our democracies.

The AI for Social Good workshop divides the in-focus problem areas into thematic blocks of talks, panels, breakout planning sessions, and posters. Particular emphasis is given to celebrating recent achievements in AI solutions, and fostering collaborations for the next generation of solutions for social good.

First, the workshop will feature a series of invited talks and panels on agriculture and environmental protection, education, health and assistive technologies, urban planning and social services. Secondly, it will bring together ML researchers, leaders of social impact, people who see the needs in the field as well as philanthropists in a forum to present and discuss interesting research ideas and applications with the potential to address social issues. Indeed, the rapidly expanding field of AI has the potential to transform many aspects of our lives. However, two main problems arise when attempting to tackle social issues. There are few venues in which to share successes and failures in research at the intersection of AI and social problems, an absence this workshop is

designed to address by showcasing these marginalized but impactful works of research. Also, it is difficult to find and evaluate problems to address for researchers with an interest on having a social impact. We hope this will inspire the creation of new tools by the community to tackle these important problems. Also, this workshop promotes the sharing of information about datasets and potential projects which could interest machine learning researchers who want to apply their skills for social good.

The workshop also explores how artificial intelligence can be used to enrich democracy, social welfare, and justice. A focus on these topics will connect researchers to civil society organizations, NGOs, local governments, and other organizations to enable applied AI research for beneficial outcomes. Various case-studies and discussions are introduced around these themes: summary of existing AI for good projects and key issues for the future, AI's impact on economic inequality, AI approaches to social sciences, and civil society organizations.

The definition of what constitutes social good being essential to this workshop, we will have panel discussions with leading social scholars to frame how contemporary AI/ML applications relate to public and philosophical notions of social good. We also aim to define new, quantifiable, and impactful research questions for the AI/ML community. Also, we would like as an outcome of this event the creation of a platform to share data, a pact with leading tech companies to support research staff sabbaticals with social progress organizations, and the connection of researchers to on-the-ground problem owners and funders for social impact.

We invite contributions relating to any of the workshop themes or more broadly any of the UN SDGs. The models or approaches presented do not necessarily need to be of outstanding theoretical novelty, but should demonstrate potential for a strong social impact. We invite two types of submissions. First, we invite research work as short papers (4 page limit) for oral and/or poster presentation. Second, we invite two page abstracts presenting a specific solution that would, if accepted, be discussed during round-table events. The short papers should focus on past and current work, showcasing actual results and ideally demonstrated beneficial effect on society, whereas the two page abstracts could highlight ideas that have not yet been applied in practice. These are designed to foster sharing different points of view ranging from the scientific assessment of feasibility, to discussion of practical constraints that may be encountered when they are deployed, also attracting interest from philanthropists invited to the event. The workshop provides a platform for developing these two page abstracts into real projects with a platform to connect with stakeholders, scientists, and funders.

NIPS 2018 Workshop on Meta-Learning

Erin Grant, Frank Hutter, Sachin Ravi, Joaquin Vanschoren, Jane Wang

Sat Dec 08, 08:00 AM

Recent years have seen rapid progress in meta-learning 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. Meta-learning 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.

Meta-learning methods are also of substantial practical interest, since they have, e.g., been shown to yield new state-of-the-art automated machine learning methods, novel deep learning architectures, and substantially improved one-shot learning systems.

Some of the fundamental questions that this workshop aims to address are:

- What are the fundamental differences in the learning "task" compared to traditional "non-meta" learners?
- Is there a practical limit to the number of meta-learning layers (e.g., would a meta-meta-meta-learning algorithm be of practical use)?
- How can we design more sample-efficient meta-learning methods?
- How can we exploit our domain knowledge to effectively guide the meta-learning process?
- What are the meta-learning processes in nature (e.g. in humans), and how can we take inspiration from them?
- Which ML approaches are best suited for meta-learning, in which circumstances, and why?
- What principles can we learn from meta-learning to help us design the next generation of learning systems?

The goal of this workshop is to bring together researchers from all the different communities and topics that fall under the umbrella of meta-learning. 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 meta-learning, as well as possible solutions.

In terms of prospective participants, our main targets are machine learning researchers interested in the processes related to understanding and improving current meta-learning algorithms. Specific target communities within machine learning include, but are not limited to: meta-learning, AutoML, reinforcement learning, deep learning, optimization, evolutionary computation, and Bayesian optimization. Our invited speakers also include researchers who study human learning, to provide a broad perspective to the attendees.

CiML 2018 - Machine Learning competitions "in the wild": Playing in the real world or in real time

Isabelle Guyon, Evelyne Viegas, Sergio Escalera, Jacob D Abernethy

Sat Dec 08, 08:00 AM

Challenges in machine learning and data science are competitions running over several weeks or months to resolve problems using provided datasets or simulated environments. The playful nature of challenges naturally attracts students, making challenge a great teaching resource. For this fifth edition of the CiML workshop at NIPS we want to go beyond simple data science challenges using canned data. We will explore the possibilities offered by challenges in which code submitted

by participants are evaluated "in the wild", directly interacting in real time with users or with real or simulated systems. Organizing challenges "in the wild" is not new. One of the most impactful such challenge organized relatively recently is the DARPA grant challenge 2005 on autonomous navigation, which accelerated research on autonomous vehicles, leading to self-driving cars. Other high profile challenge series with live competitions include RoboCup, which has been running from the past 22 years. Recently, the machine learning community has started being interested in such interactive challenges, with last year at NIPS the learning to run challenge, an reinforcement learning challenge in which a human avatar had to be controlled with simulated muscular contractions, and the ChatBot challenge in which humans and robots had to engage into an intelligent conversation. Applications are countless for machine learning and artificial intelligence programs to solve problems in real time in the real world, by interacting with the environment. But organizing such challenges is far from trivial

The workshop will give a large part to discussions around two principal axes: (1) Design principles and implementation issues; (2) Opportunities to organize new impactful challenges.

Our objectives include bringing together potential partner to organize new such challenges and stimulating "machine learning for good", i.e. the organization of challenges for the benefit of society.

CiML 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 challenges "in the wild", one of the hottest topics in challenge organization. We have invited prominent speakers having experience in this domain.

The audience of this workshop is targeted to workshop organizers, participants, and anyone with scientific problem involving machine learning, which may be formulated as a challenge. The emphasis of the workshop is on challenge design. Hence it complements nicely the workshop on the NIPS 2018 competition track and will help paving the way toward next year's competition program.

Submit abstract (up to 2 pages) before October 10 by sending email to nips2018@chalearn.org. See <http://ciml.chalearn.org/ciml2018#CALL>.

Infer to Control: Probabilistic Reinforcement Learning and Structured Control

Leslie Kaelbling, Martin Riedmiller, Marc Toussaint, Igor Mordatch, Roy Fox, Tuomas Haarnoja

Sat Dec 08, 08:00 AM

Reinforcement learning and imitation learning are effective paradigms for learning controllers of dynamical systems from experience. These fields have been empowered by recent success in deep learning of differentiable parametric models, allowing end-to-end training of highly nonlinear controllers that encompass perception, memory, prediction, and decision making. The aptitude of these models to represent latent dynamics, high-level goals, and long-term outcomes is unfortunately curbed by the poor sample complexity of many current algorithms for learning these models from experience.

Probabilistic reinforcement learning and inference of control structure are emerging as promising approaches for avoiding prohibitive amounts of

controller–system interactions. These methods leverage informative priors on useful behavior, as well as controller structure such as hierarchy and modularity, as useful inductive biases that reduce the effective size of policy search space and shape the optimization landscape. Intrinsic and self-supervised signals can further guide the training process of distinct internal components — such as perceptual embeddings, predictive models, exploration policies, and inter-agent communication — to break down the hard holistic problem of control into more efficiently learnable parts.

Effective inference methods are crucial for probabilistic approaches to reinforcement learning and structured control. Approximate control and model-free reinforcement learning exploit latent system structure and priors on policy structure, that are not directly evident in the controller–system interactions, and must be inferred by the learning algorithm. The growing interest of the reinforcement learning and optimal control community in the application of inference methods is synchronized well with the development by the probabilistic learning community of powerful inference techniques, such as probabilistic programming, variational inference, Gaussian processes, and nonparametric regression.

This workshop is a venue for the inference and reinforcement learning communities to come together in discussing recent advances, developing insights, and future potential in inference methods and their application to probabilistic reinforcement learning and structured control. The goal of this workshop is to catalyze tighter collaboration within and between the communities, that will be leveraged in upcoming years to rise to the challenges of real-world control problems.

Schedule

08:30 AM **Introduction**

Emergent Communication Workshop

Jakob Foerster, Angeliki Lazaridou, Ryan Lowe, Igor Mordatch, Douwe Kiela, Kyunghyun Cho

Sat Dec 08, 08:00 AM

Abstract

Communication is one of the most impressive human abilities. The question of how communication arises has been studied for many decades, if not centuries. However, due to computational and representational limitations, past work was restricted to low dimensional, simple observation spaces. With the rise of deep reinforcement learning methods, this question can now be studied in complex multi-agent settings, which has led to flourishing activity in the area over the last two years. In these settings agents can learn to communicate in grounded multi-modal environments and rich communication protocols emerge.

Last year at NIPS 2017 we successfully organized the inaugural workshop on emergent communication (<https://sites.google.com/site/emecom2017/>). We had a number of interesting submissions looking into the question of how language can emerge using evolution (see this Nature paper that was also presented at the workshop last year, <https://www.nature.com/articles/srep34615>)

and under what conditions emerged language exhibits compositional properties, while others explored specific applications of agents that can communicate (e.g., answering questions about textual inputs, a paper presented by Google that was subsequently accepted as an oral presentation at ICLR this year, etc.).

While last year’s workshop was a great success, there are a lot of open questions. In particular, the more challenging and realistic use cases come from situations where agents do not have fully aligned interests and goals, i.e., how can we have credible communication amongst self-interested agents where each agent maximizes its own individual rewards rather than a joint team reward? This is a new computational modeling challenge for the community and recent preliminary results (e.g. “Emergent Communication through Negotiation”, Cao et al., ICLR 2018.) reinforce the fact that it is no easy feat.

Since machine learning has exploded in popularity recently, there is a tendency for researchers to only engage with recent machine learning literature, therefore at best reinventing the wheel and at worst recycling the same ideas over and over, increasing the probability of being stuck in local optima. For these reasons, just like last year, we want to take an interdisciplinary approach on the topic of emergent communication, inviting researchers from different fields (machine learning, game theory, evolutionary biology, linguistics, cognitive science, and programming languages) interested in the question of communication and emergent language to exchange ideas.

This is particularly important for this year’s focus, since the question of communication in general-sum settings has been an active topic of research in game theory and evolutionary biology for a number of years, while it’s a nascent topic in the area of machine learning.

Interpretability and Robustness in Audio, Speech, and Language

Mirco Ravanelli, Dmitriy Serdyuk, Ehsan Variani, Bhuvana Ramabhadran

Sat Dec 08, 08:00 AM

Domains of natural and spoken language processing have a rich history deeply rooted in information theory, statistics, digital signal processing and machine learning. With the rapid rise of deep learning (“deep learning revolution”), many of these systematic approaches have been replaced by variants of deep neural methods, that often achieve unprecedented performance levels in many fields. With more and more of the spoken language processing pipeline being replaced by sophisticated neural layers, feature extraction, adaptation, noise robustness are learnt inherently within the network. More recently, end-to-end frameworks that learn a mapping from speech (audio) to target labels (words, phones, graphemes, sub-word units, etc.) are becoming increasingly popular across the board in speech processing in tasks ranging from speech recognition, speaker identification, language/dialect identification, multilingual speech processing, code switching, natural language processing, speech synthesis and much much more.

A key aspect behind the success of deep learning lies in the discovered low and high-level representations, that can potentially capture relevant

underlying structure in the training data. In the NLP domain, for instance, researchers have mapped word and sentence embeddings to semantic and syntactic similarity and argued that the models capture latent representations of meaning. Nevertheless, some recent works on adversarial examples have shown that it is possible to easily fool a neural network (such as a speech recognizer or a speaker verification system) by just adding a small amount of specially constructed noise. Such a remarkable sensibility towards adversarial attacks highlights how superficial the discovered representations could be, rising crucial concerns on the actual robustness, security, and interpretability of modern deep neural networks. This weakness naturally leads researchers to ask very crucial questions on what these models are really learning, how we can interpret what they have learned, and how the representations provided by current neural networks can be revealed or explained in a fashion that modeling power can be enhanced further. These open questions have recently raised the interest towards interpretability of deep models, as witness by the numerous works recently published on this topic in all the major machine learning conferences. Moreover, some workshops at NIPS 2016, NIPS 2017 and Interspeech 2017 have promoted research and discussion around this important issue.

With our initiative, we wish to further foster some progresses on interpretability and robustness of modern deep learning techniques, with a particular focus on audio, speech and NLP technologies. The workshop will also analyze the connection between deep learning and models developed earlier for machine learning, linguistic analysis, signal processing, and speech recognition. This way we hope to encourage a discussion amongst experts and practitioners in these areas with the expectation of understanding these models better and allowing to build upon the existing collective expertise.

The workshop will feature invited talks, panel discussions, as well as oral and poster contributed presentations. We welcome papers that specifically address one or more of the leading questions listed below:

1. Is there a theoretical/linguistic motivation/analysis that can explain how nets encapsulate the structure of the training data it learns from?
2. Does the visualization of this information (MDS, t-SNE) offer any insights to creating a better model?
3. How can we design more powerful networks with simpler architectures?
4. How can we can exploit adversarial examples to improve the system robustness?
5. Do alternative methods offer any complimentary modeling power to what the networks can memorize?
6. Can we explain the path of inference?
7. How do we analyze data requirements for a given model? How does multilingual data improves learning power?

Schedule

08:45 AM	Workshop Opening	
09:00 AM	Invited Speaker 1	<i>Caruana</i>
09:30 AM	Invited Speaker 2	<i>Yosinski</i>
10:00 AM	Contributed Talks 1	
10:30 AM	Coffee break + posters 1	
11:00 AM	Invited Speaker 3	<i>Hermansky</i>
11:30 AM	Invited Speaker 4	<i>Bacchiani</i>

12:00 PM	Contributed Talks 2	
12:30 PM	Lunch Break	
01:30 PM	Invited Speaker 5	<i>Schlüter</i>
02:00 PM	Invited Speaker 6	<i>Rush</i>
02:30 PM	Contributed Talks 3	
03:00 PM	Coffee break + posters 2	
03:30 PM	Invited Speaker 7	<i>Schuster</i>
04:00 PM	Invited Speaker 8	
04:30 PM	Contributed Talks 4	
05:00 PM	Panel Discussion	
06:00 PM	Workshop Closing	

Machine Learning Open Source Software 2018: Sustainable communities

Heiko Strathmann, Viktor Gal, Ryan Curtin, Sergey Lisitsyn, Antti Honkela, Cheng Soon Ong

Sat Dec 08, 08:00 AM

Machine learning open source software (MLOSS) is one of the cornerstones of open science and reproducible research. Once a niche area for ML research, MLOSS today has gathered significant momentum, fostered both by scientific community, and more recently by corporate organizations. Along with open access and open data, it enables free reuse and extension of current developments in ML. The past mloss.org workshops at NIPS06, NIPS08, ICML10, NIPS13, and ICML15 successfully brought together researchers and developers from both fields, to exchange experiences and lessons learnt, to encourage interoperability between people and projects, and to demonstrate software to users in the ML community.

Continuing the tradition in 2018, we plan to have a workshop that is a mix of invited speakers, contributed talks and discussion/activity sessions. This year's headline aims to give an insight of the challenges faced by projects as they seek long-term sustainability, with a particular focus on community building and preservation, and diverse teams. In the talks, we will cover some of the latest technical innovations as done by established and new projects. The main focus, however, will be on insights on project sustainability, diversity, funding and attracting new developers, both from academia and industry. We will discuss various strategies that helps promoting gender diversity in projects (e.g. implementing quotas etc.) and how to promote developer growth within a project.

We aim to make this workshop as diverse as possible within the field. This includes a gender balanced speakers, focussing on programming languages from different scientific communities, and in particular most of our invited speakers represent umbrella projects with a hugely diverse set of applications and users (NumFOCUS, openML, tidyverse).

With a call for participation for software project demos, we aim to provide improved outreach and visibility, especially for smaller OSS projects as

typically present in academia. In addition, our workshop will serve as a gathering of OSS developers in academia, for peer to peer exchange of learnt lessons, experiences, and sustainability and diversity tactics.

The workshop will include an interactive session to produce general techniques for driving community engagement and sustainability, such as application templates (Google Summer of Code, etc), "getting started" guides for new developers, and a collection of potential funding sources. We plan to conclude the workshop with a discussion on the headline topic.

Learning by Instruction

Shashank Srivastava, Igor Labutov, Bishan Yang, Amos Azaria, Tom Mitchell

Sat Dec 08, 08:00 AM

Today machine learning is largely about pattern discovery and function approximation. But as computing devices that interact with us in natural language become ubiquitous (e.g., Siri, Alexa, Google Now), and as computer perceptual abilities become more accurate, they open an exciting possibility of enabling end-users to teach machines similar to the way in which humans teach one another. Natural language conversation, gesturing, demonstrating, teleoperating and other modes of communication offer a new paradigm for machine learning through instruction from humans. This builds on several existing machine learning paradigms (e.g., active learning, supervised learning, reinforcement learning), but also brings a new set of advantages and research challenges that lie at the intersection of several fields including machine learning, natural language understanding, computer perception, and HCI.

The aim of this workshop is to engage researchers from these diverse fields to explore fundamental research questions in this new area, such as:

How do people interact with machines when teaching them new learning tasks and knowledge?

What novel machine learning models and algorithms are needed to learn from human instruction?

What are the practical considerations towards building practical systems that can learn from instruction?