Future Leaders of AI Retreat (FLAIR)

December 20-21, NYU Shanghai

Conference Handbook

Program

December 20, Wednesday

Time	Event/Talk	Speaker
8:30-9:00am	Check-in	
9:00-9:10am	Opening	
Session I "Theory and Core Algorithms"		
9:10-9:40am	On the Power of Random Initialized Gradient Descent for Learning Convolutional Neural Networks	Shaolei Du
9:40-10:10am	Opening the black of deep neural networks via information	Ravid Schwartz-Ziv
10:10-10:30am	Tea Break (Room 1502)	
10:30-11:00am	Nonconvex Sparse Blind Deconvolution: Global Geometry and Efficient Methods	Yuqian Zhang
11:00-11:30am	Adversarial regularized autoencoders and error encoding network	Junbo Zhao
11:30am-12:00pm	Enhancing Iterative Algorithms Using Neural Networks	David Wipf (Keynote)
12:00-12:15pm	Group Q&A	
12:15-1:15pm	Lunch (B1 canteen)	
Session II "Application: Game and Generalizable Intelligence"		
1:15-1:45pm	Sample-Efficient Deep RL for Robotics: Generalizing On-policy, Off-policy, and Model-based Approaches	Shixiang Gu
1:45-2:15pm	Towards a wholistic approach to dexterous and cluttered manipulation	Jiaji Zhou
2:15-2:45pm	Towards Generalizable Robot Learning with Perception	Yuke Zhu

2:45-3:15pm	Effective Master-Slave Communication for Multi-Agent Deep Reinforcement Learning	Xiangyu Kong
3:15-3:45pm	AI in Games: Achievements and Challenges	Yuandong Tian (Keynote)
3:45-4:00pm	Group Q&A	
4:00-4:15pm	Tea Break (Room 1502)	
Session III "Application: Natural Language Processing and Computer Vision"		
4:15-4:35pm	Present and Future of Natural Language Processing - With Dialogue as Example	Hang Li (Keynote)
4:35-5:05pm	Think out of black-box	Zhengdong Lu
5:05-5:35pm	Learning to Reason with End-to-End Neural Module Networks	Ronghang Hu
5:35-6:05pm	On unifying deep generative models	Zhiting Hu
6:05-6:35pm	Deep Multitask Learning for Semantic Dependency Parsing	Hao Peng
6:35-6:50pm	Group Q&A	

December 21, Thursday

Time	**Event/Talk**	**Speaker**
8:30-9:00am	Check-in	
Session IV "System"		
9:00-9:30am	TVM: An End to End IR Stack for Deep Learning Systems	Tianqi Chen
9:30-10:00am	Tofu: Distributing Tensor Computation Automatically for Deep Learning Systems	Minjie Wang
10:00-10:15am	Group Q&A	

10:15-10:30am	Tea break (Room 1502)	
Session V "Creativity and Music with AI"		
10:30-11:00am	Empowering music intelligent systems with expression and interactivity	Gus Xia (Keynote)
11:00am-11:30pm	Creativity in Musical Intelligence	Zhengshan Shi
11:30am-12:00pm	Crowd Sourcing Clothes Design Directed by Adversarial Neural Networks	Hiroyuki Osone & Daitetsu Sato
12:00-12:15pm	Group Q&A	

FLAIR Speakers
(by track)

Session I "Theory and Core Algorithms"

Shaolei Du (杜少雷), PhD at Carnegie Mellon University

Bio: Simon Shaolei Du is a PhD student in the Machine Learning Department at the School of Computer Science, Carnegie Mellon University, advised by Professor Aarti Singh and Professor Barnabás Póczos. His research interests broadly include topics in theoretical machine learning and statistics, such as deep learning, matrix factorization, convex/non-convex optimization, transfer learning, reinforcement learning, non-parametric statistics and robust statistics. Currently he is also developing methods for precision agriculture. In 2011, he earned his high school degree from The Experimental High School Attached to Beijing Normal University. In 2015, he obtained his B.S. in Engineering Math & Statistics and B.S. in Electrical Engineering & Computer Science from University of California, Berkeley. He has also spent time working at research labs of Microsoft and Facebook.

Talk Title: On the Power of Randomly Initialized Gradient Descent for Learning Convolutional Neural Networks

Talk Abstract:
Convolutional neural networks trained by randomly initialized (stochastic) gradient descent have achieved the state-of-art performances in many applications. However, its theoretical properties remain elusive from an optimization point of view. In this talk, I will present two results on explaining the success of gradient descent.

In the first part, I will show under certain structural conditions of the input distribution, random initialized gradient descent provably learns a convolutional filter with ReLU activation and average pooling. This is the first recovery guarantee of gradient-based algorithms for learning a convolutional filter on general input distributions.

In the second part of the talk, I will show if the input distribution is Gaussian, then random initialized gradient descent with weight-normalization learns a ReLU activated one-hidden-layer convolutional neural network where both the convolutional weights and the output weights are to be optimized. To the best our knowledge, this is the first recovery guarantee of random initialized gradient-based algorithms for neural networks that contain more than one layers to be learned.

*This talk is based on works with Jason D. Lee, Barnabás Póczos, Aarti Singh and Yuandong Tian.

Ravid Shwartz-Ziv, PhD at the Hebrew University of Jerusalem

Bio: Ravid's research interests lie in the intersection of learning, information and optimization, especially in deep neural networks. Ravid's main focus is exploring learning and dynamics via information for both artificial and biological neural network.
Ravid is currently a computational neuroscience PhD candidate at the Hebrew University of Jerusalem at the Machine Learning Lab under the supervision of Prof. Tali Tishby. Ravid received both his B.A and his M.Sc degrees from the Hebrew University of Jerusalem

Talk Title: Opening the black of deep neural networks via information

Talk abstract:
I will present a analytical and numerical study of Deep Neural Networks in the "Information Plane" - the Mutual Information the network layers preserve on the input and the output variables. By exploration of this plane we obtain the following new insights:
1. The training epochs, for each layer, are divided into two phases: (1) fitting the training data - increasing the mutual information on the labels; (2) compressing the representation - reducing the mutual information on the inputs. The layers are learnt hierarchically, from the bottom to the top layer, with some overlaps.
2. Most of the training time is spent on compressing the representation (the second phase) - not on fitting the training data labels, even when the training has no regularization or terms that directly aim at such compression.
3. The convergence point, for every hidden layer, lies on or very close to the Information Bottleneck (IB) theoretical bound. Thus, the mappings from the input to the hidden layer and from the hidden layer to the output obey the IB self-consistent equations for some value of the compression-prediction tradeoff.
4. Stochastic gradient descent, as used in deep learning, achieves this optimal bound - as the compression phase for each layer amounts to relaxation to a maximum conditional entropy state, subject to the proper constraints on the error/information on the labels. In that sense, deep learning is a method for solving the Information Bottleneck problem for large scale learning problems.

I will explain these new results and the benefits of exploring deep learning in the "Information Plane", and discuss some of the theoretical and practical consequences of our analysis.

Yuqian Zhang (张雨倩), PhD at Columbia University

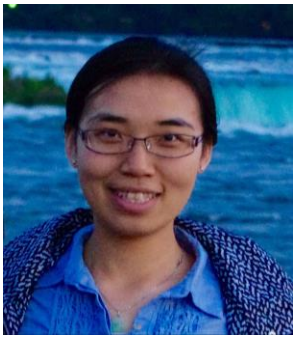

Bio: Yuqian Zhang is a Ph.D. candidate in the Electrical Engineering Department at Columbia University, advised by Professor John Wright. She received her B.S. in Electrical Engineering from Xi'an Jiaotong University. Her research spans across optimization, computer vision, signal processing, and machine learning. Specifically, her primary research interest is to develop efficient, reliable and robust algorithms for applications in computer vision, scientific data analysis, etc.

Talk Title: Nonconvex Sparse Blind Deconvolution: Global Geometry and Efficient Methods

Talk Abstract:

Blind deconvolution is a ubiquitous problem aiming to recover a convolution kernel **a** (length k) and an activation signal **x** (length m) from their convolution **y**. This is an ill-posed problem in general. This talk focuses on the short and sparse blind deconvolution problem, where the convolution kernel is short (k<<m) and the activation signal is sparsely and randomly supported. This variant models convolutional signals in several important practical application scenarios. The observation is invariant to some mutual scaling and shift of the convolutional pairs. Such scaled-shift symmetry is intrinsic to the convolution operator and imposes challenges for reliable algorithm design. We normalize the convolution kernel to have unit Frobenius norm and then cast the blind deconvolution problem as a nonconvex optimization problem over the kernel sphere. We demonstrate that (i) under conditions, every local optimum is close to some shift truncation of the ground truth, and (ii) for a generic filter **a** on the sphere, when the sparsity of activation signal satisfies theta < O(k^{-2/3}) and number of measurements m > poly(k), provable recovery of some shift truncation of the ground truth kernel can be obtained.

Junbo Zhao (赵俊博), PhD at New York University

Bio: Junbo (Jake) Zhao is currently a 2nd year PhD student at CILVR lab at NYU, under the supervision of Professor Yann LeCun. His recent main research interests include deep learning and unsupervised learning, on both domains of vision and language. In recent years, Jake has interned at Facebook AI research team, Clarifai engineering team, NVIDIA autonomous driving team. He graduates from Wuhan University majoring in electrical engineering in 2014 and holds a master degree in data science from NYU.

Talk Title: Adversarial regularized autoencoders and error encoding network

Talk Abstract:
Auto-encoding is a key technique towards representation learning. One often needs to prevent the degenerate solution in which the auto-encoder falls into an meaningless copy function. Several techniques have been developed within the recent years, such as imposing sparsity prior, variational auto-encoders or injecting noise and etc. Here, we term these “regularization”. In this talk, I would present some recent development on regularizing autoencoders -- adversarially regularized autoencoder (ARAE) and error-encoding network (EEN). With some newly developed regularizers, both model have empirically been showed that they learn meaningful representations. Thanks to a good representation, we show that ARAE can be adopted to achieve state-of-the-art results on unaligned text style transfer and EEN is able to predict into future with multimodality.

David Wipf, Lead Researcher at Microsoft Research
(Keynote)

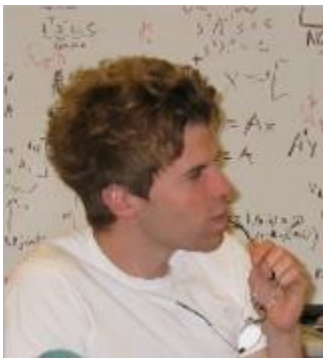

Bio: David Wipf is a researcher with the Visual Computing Group at Microsoft Research in Beijing, where he has been employed full-time since 2011. Prior to this position, he received the B.S. degree with highest honors in electrical engineering from the University of Virginia, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of California, San Diego. He was later an NIH Postdoctoral Fellow in the Biomagnetic Imaging Lab at the University of California, San Francisco. His research interests include developing and analyzing deep network models, Bayesian learning algorithms, and non-convex optimization techniques across problems in signal/image processing. He is the recipient of numerous fellowships and awards including the 2012 Signal Processing Society Best Paper Award, the Biomag 2008 Young Investigator Award, and the 2006 NIPS Outstanding Student Paper Award. He is currently an Action Editor for the Journal of Machine Learning Research and a member of the IEEE Machine Learning for Signal Processing Technical Committee. He also served/will serve as an Area Chair for NIPS 2014, NIPS 2017, ICCV 2017, and ICML 2018, as well as Program Chair for SPCOM 2018.

Talk Title: Enhancing Iterative Algorithms Using Neural Networks

Talk Abstract:
The iterations of many first-order algorithms, when applied to minimizing common regularized regression functions, often resemble neural network layers with pre-specified weights. This observation has prompted the development of learning-based approaches that purport to replace these iterations with enhanced surrogates forged as DNN models from available training data. For example, important NP-hard sparse estimation problems have recently benefited from this genre of upgrade, with simple feedforward or recurrent networks ousting proximal gradient-based iterations. Analogously, this paper demonstrates that more powerful Bayesian algorithms for promoting sparsity, which rely on complex multi-loop majorization-minimization techniques, mirror the structure of more sophisticated long short-term memory (LSTM) networks, or alternative gated feedback networks previously designed for sequence prediction. As part of this development, we examine the parallels between latent variable trajectories operating across multiple time-scales during optimization, and the activations within deep network structures designed to adaptively model such characteristic sequences. The resulting insights lead to a novel sparse estimation system that, when granted training data, can estimate optimal solutions efficiently in regimes where other algorithms fail, including practical direction-of-arrival (DOA) and 3D geometry recovery problems. The underlying principles we expose are also suggestive of a learning process for a richer class of multi-loop algorithms in other domains.

Session II "Application: Game and Generalizable Intelligence"

Shixiang Gu (顾世翔), PhD at University of Cambridge

Bio: Shixiang (Shane) Gu is a PhD candidate at University of Cambridge and Max Planck Institute for Intelligent Systems, where he is jointly co-supervised by Richard E. Turner, Zoubin Ghahramani, and Bernhard Schoelkopf. He holds BASc. in Engineering Science from University of Toronto, where he completed this thesis with Professor Geoffrey Hinton. His research interests span deep reinforcement learning, deep learning, robotics, approximate inference and causality, and his research has been featured by MIT Technology Review and Google Research Blog. He also collaborates closely with Sergey Levine from UC Berkeley/Google and Tim Lillicrap from DeepMind.

Talk Title: Sample-Efficient Deep RL for Robotics: Generalizing On-policy, Off-policy, and Model-based Approaches

Talk Abstract:
Deep reinforcement learning (RL) has shown promising results for learning complex sequential decision-making behaviors in various environments. However, most successes have been exclusively in simulation, and results in real-world applications such as robotics are limited, largely due to poor sample efficiency of typical deep RL algorithms. In this talk, I will present methods to improve sample efficiency of these algorithms, blurring the boundaries among classic model-based RL, off-policy and on-policy model-free RL. The first part of the talk will discuss Q-Prop, a control variate technique for policy gradient that combines on-policy and off-policy learning and discusses empirical results and theoretical variance reduction. The second part of the talk focuses on temporal difference models (TDMs), an extension of goal-conditioned value functions that enables multi time resolution model-base planning. TDMs generalize traditional predictive models, bridge the gap between model-based and off-policy model-free RL, and empirically lead to substantial improvements in sample efficiency with vectorized implementation.

Jiaji Zhou (周佳骥), PhD at Carnegie Mellon University

Bio: Jiaji Zhou is a PhD student in the Manipulation Lab of the Robotics Institute at Carnegie Mellon University, co-advised by Matt Mason and Drew Bagnell. His work won the ICRA 2016 Best Conference Paper Award. He has interned at GoogleX self-driving car team, Dato and Toyota Research Institute Manipulation Group.

Talk Title: Towards a wholistic approach to dexterous and cluttered manipulation

Talk Abstract:
In this talk, I will present our work on tackling different aspects of robotic manipulation: mechanics modeling, non-prehensile manipulation planning and continuously adjusted motion and perception for cluttered grasping. However, many pieces are missing including some system level issues often ignored by academia. Boston dynamics is undoubtedly the leader in the field of robotic locomotion. I argue that there are rich opportunities for manipulation researchers to learn from Boston Dynamics and adopt a wholistic approach to build systems from high level algorithms to low level camera and actuator design.

Yuke Zhu (朱玉可), PhD at Stanford University

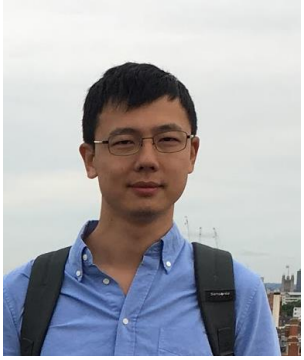

Bio: Yuke Zhu is a fifth-year Ph.D. student in Computer Science at Stanford University, advised by Professor Fei-Fei Li and Professor Silvio Savarese. His research focuses on the principles and applications of computer vision, machine learning, and robotics, in particular, visual knowledge and deep reinforcement learning. Prior to coming to Stanford, he received a BEng. degree from Zhejiang University and a BSc. degree from Simon Fraser University, working with Professor Greg Mori. He also collaborates with research labs including Snap Research, Allen Institute for Artificial Intelligence, and Google DeepMind.

Talk Title: Towards Generalizable Robot Learning with Perception

Talk Abstract:
Robots have demonstrated remarkable dexterity of backflipping in Boston Dynamics' latest videos. Nevertheless, the DARPA Robotics Challenges show robots failing abysmally at seemingly simple tasks such as turning a door handle. While it is possible to engineer useful robots for narrow tasks, they often fail to generalize as soon as we modify the task. This talk will present methods that improve the generalizability of robot controllers through a tight integration between perception and embodiment. I will discuss generalization across task goals with two case studies of visual navigation and visual semantic planning. I will also talk about neural task programming for generalizing across task structures, and present a reinforcement and imitation model that generalizes across large visual and physical variations in learning visuomotor skills.

Xiangyu Kong (孔祥宇), PhD at Peking University

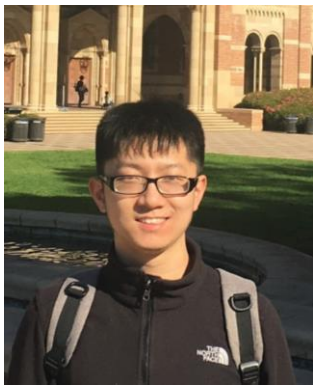

Bio: Xiangyu Kong is a fifth-year Ph.D candidate in Computer Science at Peking University, under the supervision of Prof. Yizhou Wang. He also works very closely with Dr. Bo Xin of Microsoft Research Asia. Prior to that, He graduated from Harbin Institute of Technology with a Bachelor of Computer Science. His current research interest includes computer vision, machine learning (in particular, multi-agent deep reinforcement learning) and their applications in video game playing.

Talk Title: Effective Master-Slave Communication for Multi-Agent Deep Reinforcement Learning

Talk Abstract:
Many challenging practical problems require multiple agents to solve collaboratively. However, communication becomes a bottleneck when a multi-agent system (MAS) scales. This is particularly true when a MAS is deployed to autonomous learning (e.g. reinforcement learning), where massive interactive communication is required. We argue that the effectiveness of communication is a key factor to determine the intelligence level of a multi-agent learning system. In this regard, we propose to adapt the classical hierarchical master-slave architecture to facilitate efficient multi-agent communication during the interactive reinforcement learning (RL) process implemented on a deep neural network. The master agent aggregates messages uploaded from the slaves and generates unique message to each slave according to the aggregated information and the states of the slave. Each slave incorporates both the instructive message from the master and its own to take actions to fulfill the goal. In this way, the joint action-state space of the agents grows only linearly instead of geometrically with the number of agents compared to the peer-to-peer architecture. In experiments, we show that with effective communication, the proposed multi-agent learning system consistently outperforms latest competing methods both in synthetics experiments and the challenging StarCraft micromanagement tasks.

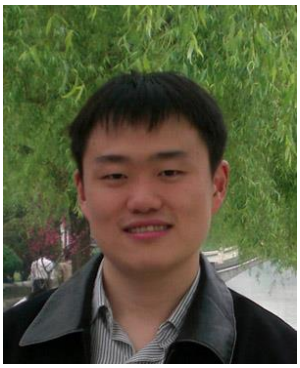

Yuandong Tian (田渊栋), Research Scientist at Facebook AI Research
(Keynote)

Bio: Yuandong Tian is a Research Scientist and Manager in Facebook AI Research, working on deep reinforcement learning, its applications in games and theoretical analysis of deep models. Prior to that, he was a Software Engineer/Researcher in Google Self-driving Car team during 2013-2014. He received Ph.D in Robotics Institute, Carnegie Mellon University on 2013, Bachelor and Master degree of Computer Science in Shanghai Jiao Tong University. He is the recipient of 2013 ICCV Marr Prize Honorable Mentions.

Talk Title: AI in Games: Achievements and Challenges

Talk Abstract:
Recently, substantial progress of AI has been made in applications that require advanced pattern reading, including computer vision, speech recognition and natural language processing. However, it remains an open problem whether AI will make the same level of progress in tasks that require sophisticated reasoning, planning and decision making in complicated game environments similar to the real-world. In this talk, I present the state-of-the-art approaches to build such an AI, our recent contributions in terms of designing more effective algorithms and building extensive and fast general environments, as well as issues and challenges.

Session III "Application: Natural Language Processing and Computer Vision"

Hang Li (李航), Director of Toutiao AI Lab

(Keynote)

Bio: Hang Li is director of Toutiao AI Lab, adjunct professors of Peking University and Nanjing University. He is an IEEE Fellow and an ACM Distinguished Scientist. His research areas include information retrieval, natural language processing, machine learning, and data mining. Hang graduated from Kyoto University in 1988 and earned his PhD from the University of Tokyo in 1998. He worked at NEC Research as researcher from 1990 to 2001, Microsoft Research Asia as senior researcher and research manager from 2001 to 2012, and chief scientist and director of Huawei Noah's Ark from 2012 to 2017. He joined Toutiao in 2017.

Talk Title: Present and Future of Natural Language Processing - With Dialogue as Example

Talk Abstract:
In this talk, I will discuss the current trends and future directions of natural language processing (NLP). I will take natural language dialogue as example and introduce recent work on the problem. I will argue that neural symbolic processing, semantic grounding, modular and hierarchical processing, and meta learning will be the key technologies for advancement of NLP in the coming years.

Zhengdong Lu (吕正东), Founder of DeeplyCurious.ai

Bio: Zhengdong Lu obtained his Ph.D degree in Computer Science in Oregon Health & Science University, USA. Before founding DeeplyCurious.ai, he was a senior researcher in Noah's Ark Lab, Huawei, and an associated researcher in MSRA. Dr. Lu is interested in natural language processing and deep learning, and is noted for his work on neural reasoning, neural machine translation, semantic parsing and neural-symbolism.

Talk Title: Think out of black-box

Talk Abstract:
I will talk about our thoughts and recent practice on understanding documents in closed domains. More specifically, I will discuss 1) modeling document reading/understanding as a complex decision process, and 2) combining neural network and symbolic AI (neural-symbolism) in document understanding.

Ronghang Hu (胡戎航), PhD at UC Berkeley

Bio: Ronghang Hu is a 3rd-year Ph.D. student in computer science at UC Berkeley, working with Prof. Trevor Darrell. He has been working on a variety of topics in computer vision, and most notably joint vision and language tasks such as visual question answering. In 2017 summer, he was a research intern in Facebook AI Research (FAIR) working with Dr. Ross Girshick. He obtained his B.E. degree from Tsinghua University in 2015. Previously in 2013 and 2014, He was a research intern at Institute of Computing Technology, Chinese Academy of Science (ICTCAS) and was advised by Prof. Shiguang Shan and Prof. Ruiping Wang.

Talk Title: Learning to Reason with End-to-End Neural Module Networks

Talk Abstract:
Natural language questions are inherently compositional, and many are most easily answered by reasoning about their decomposition into modular sub-problems. For example, to answer "is there an equal number of balls and boxes?" we can look for balls, look for boxes, count them, and compare the results. The recently proposed Neural Module Network (NMN) architecture implements this approach to question answering by parsing questions into linguistic substructures and assembling question-specific deep networks from smaller modules that each solve one subtask. However, existing NMN implementations rely on brittle off-the-shelf parsers, and are restricted to the module configurations proposed by these parsers rather than learning them from data. In this paper, we propose End-to-End Module Networks (N2NMNs), which learn to reason by directly predicting instance-specific network layouts without the aid of a parser. Our model learns to generate network structures (by imitating expert demonstrations) while simultaneously learning network parameters (using the downstream task loss).

Zhiting Hu (胡志挺), PhD at Carnage Mellon University

Bio: Zhiting Hu is a PhD student at Machine Learning Department, Carnage Mellon University. His advisor is Prof. Eric Xing. His research is focusing on knowledge-enriched deep learning, Bayesian modeling and inference, large-scale machine learning, and their applications in natural language processing, esp., text generation. His work on harnessing deep neural networks with logic rules was selected as one of the outstanding papers in ACL2016. He is the recipient of 2017 IBM Fellowship.

Talk Title: On unifying deep generative models

Talk Abstract:
Deep generative models have achieved impressive success in recent years. Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs), as powerful frameworks for deep generative model learning, have largely been considered as two distinct paradigms and received extensive independent studies respectively. This paper aims to establish formal connections between GANs and VAEs through a new formulation of them. We interpret sample generation in GANs as performing posterior inference, and show that GANs and VAEs involve minimizing KL divergences of respective posterior and inference distributions with opposite directions, extending the two learning phases of classic wake-sleep algorithm, respectively. The unified view provides a powerful tool to analyze a diverse set of existing model variants, and enables to transfer techniques across research lines in a principled way. For example, we apply the importance weighting method in VAE literatures for improved GAN learning, and enhance VAEs with an adversarial mechanism that leverages generated samples. Experiments show generality and effectiveness of the transferred techniques.

Hao Peng (彭昊), PhD at University of Washington

Bio: Hao Peng is a second year Ph.D. student in Computer Science and Engineering at the University of Washington, advised by Prof. Noah Smith. He works on natural language processing and machine learning, and is particularly interested in broad-coverage semantics. Previously, Hao received B.S. from Peking University in 2016 (with hornor), and also visited University of Edinburgh and Microsoft Research Asia.

Talk Title: Deep Multitask Learning for Semantic Dependency Parsing

Talk Abstract:
In this talk I will present a deep neural architecture that parses sentences into three semantic dependency graph formalisms. By using efficient, nearly arc-factored inference and a bidirectional-LSTM composed with a multi-layer perceptron, our base system is able to significantly improve the state of the art for semantic dependency parsing, without using hand-engineered features or syntax. We then explore two multitask learning approaches---one that shares parameters across formalisms, and one that uses higher-order structures to predict the graphs jointly. We find that both approaches improve performance across formalisms on average, achieving a new state of the art.

Session IV "System"

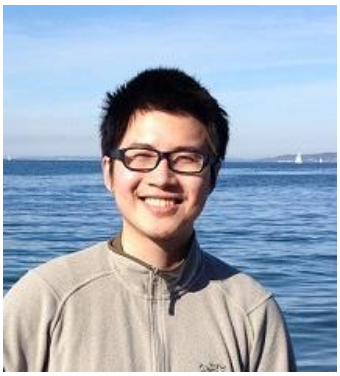

Tianqi Chen (陈天奇), University of Washington

Bio: Tianqi is a PhD student in University of Washington, working on machine learning and systems. He received his bachelor and master degrees from Shanghai Jiao Tong University. He is recipient of a Google PhD Fellowship in Machine Learning

Talk Title: TVM: An End to End IR Stack for Deep Learning Systems

Talk Abstract:
Deep learning has become ubiquitous and indispensable. Part of this revolution has been fueled by scalable deep learning systems. In this talk, I am going to talk about TVM: a unified compilation optimization stack that will close the gap between the productivity-focused deep learning frameworks, and the performance- or efficiency-oriented hardware backends. TVM is a novel framework that can: Represent and optimize the common deep learning computation workloads for CPUs, GPUs, and other specialized hardware; Automatically transform the computation graph to minimize memory utilization, optimize data layout and fuse computation patterns; Provide an end-to-end compilation from existing front-end frameworks down to bare-metal hardware and specialized accelerators. I will talk about the problems and chance of learning system research around TVM.

Minjie Wang (王敏捷), PhD at New York University

Bio: Minjie Wang is a forth year Ph.D. student at New York University and a member of the NYU systems group. Before joining NYU, Minjie got his master's and bachelor's at Shanghai Jiao Tong University. He also spent two years as a research intern in Microsoft Research Asia, where he found his research interests in machine learning systems and built his first deep learning system: Minerva. Minjie was also one of the founding members of the Deep Machine Learning Community. He is one of the main developers of the MXNet, NNVM, and MinPy projects. He is the recipient of 2016 NVIDIA Graduate Fellowship.

Talk Title: Tofu: Distributing Tensor Computation Automatically for Deep Learning Systems

Talk Abstract:
We present Tofu, which improves the scaling performance and programmability of a tensor dataflow-based DNN system by performing automatic distribution. Tofu can explore a spectrum of distribution strategies, including data parallelism, model parallelism and others in between. Such exploration is enabled through the development of tensor description language (TDL), which allows Tofu to discover all feasible ways of distributing an operator by partitioning its tensor along different dimensions. To find the best strategy with minimal communication cost for the overall dataflow graph, Tofu uses a novel search algorithm that exploits the layer-by-layer characteristics of neural network computation. We implement Tofu in MXNet and show its performance benefits for several DNN applications.

Session V "Creativity and Music with AI"

Gus Xia (夏光宇), Assistant Professor of CS at NYU Shanghai **(Keynote)**

Bio: I am an Assistant Professor in Computer Science at NYU Shanghai. I received my Ph.D. in the Machine Learning Departmentat Carnegie Mellon University where I study Machine Learning and Computer Music under the advice of Prof. Roger Dannenberg. I was a Neukom Fellow at Dartmouth from 2016 to 2017. In 2010, I received my undergraduate degree in Information Science with a minor in Psychology at Peking University.

I am also a professional DI and XIAO (Chinese flute and vertical flute) player. I am currently a soloist of the Pitt Carpathian Ensemble. Prior to that, I was the prime soloist of the Chinese Music Institute (CMI) in Peking University, where I also served as the president and assistant conductor. I held my solo concert in 2010.

Talk Title: Empowering music intelligent systems with expression and interactivity

Talk Abstract:
Why expression and interactivity are important for future intelligent systems and what are the possible solutions? Gus will lead the discussion from an artistic perspective while presenting the state-of-the-art interactive performance system. He will also cover some current research projects which aims to combine expressive AI and HCI.

Zhengshan Shi (施正珊), PhD at Stanford University

Bio: Kitty Zhengshan Shi is a current 4th year PhD student at Stanford University in Center for Computer Research in Music and Acoustics (CCRMA). She obtained her bachelor's degree at Shanghai Conservatory of Music, and a master at New York University. She is interested in intelligent music software design as well as creative music information retrieval. She also enjoys playing piano, accordion, violin, and bagpipes. She is a native Shanghainese.

Talk Title: Creativity in Musical Intelligence

Talk Abstract:
This talk highlights computer-mediated Music, Computing and Design research. It leads the discussion of how can computer be creative from the aspect of music, as well as exploring how can we bring human into the loop in the era of artificial intelligence. Talk covers creative computer music applications, music information retrieval, as well as recent works at Stanford Laptop Orchestra.

Hiroyuki Osone, Undergradudate at University of Tsukuba
Daitetsu Sato, Undergradudate at University of Tsukuba

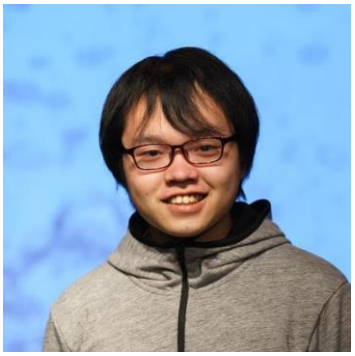

Bio:
Hiroyuki Osone is a sophomore in Digital Nature Group at University of Tsukuba, advised by Associate Prof. Yoichi Ochai. His work accepted the NIPS 2017 workshop Machine Learning for Creativity Design. He is interested in image generation by GAN and application of that image.

Daitetsu Sato is a collaborate research student in Digital Nature Group at University of Tsukuba and a manager of Academic Collaboration Promotional Laboratory at University of Tsukuba. His work accepted the NIPS 2017 workshop Machine Learning for Creativity Design. R & D of web system is specialized, but recently He is interested in making dataset using data on the internet.

Talk Title: Crowd Sourcing Clothes Design Directed by Adversarial Neural Networks

Talk Abstract:
Deep neural networks (DNNs) applications are now increasingly pervasive and powerful. However, fashion designers are lagging behind in leveraging this increasingly common technology. DNNs are not yet a standard part of fashion design practice, in either clothes patterns or prototyping tools. In this paper, we present DeepWear, a method using deep convolutional generative adversarial networks (DCGANs) for clothes design. The DNNs learn the feature of specific brand clothes and generate images, then patterns instructed from the images are made, and an author creates clothes based on that. We evaluated this system by evaluating the credibility of the actual sold clothes on market with our clothes. As the result, we found it is possible to make clothes look like actual products from the generated images. Our findings have implications for collaborative design between machine and human intelligence.