Foundations of Computational Science

August 29-30, 2019

Thursday, August 29

Time	Speaker	Title/Abstract
8:30 - 9:30am	**Breakfast**	
9:30 - 9:40am	**Opening**	
9:30 - 10:05am	Tat Seng Chua	Explainable AI and Multimedia Research **Abstract:** AI as a concept has been around since the 1950's. With the recent advancements in machine learning algorithms, and the availability of big data and large computing resources, the scene is set for AI to be used in many more systems and applications which will profoundly impact society. The current deep learning based AI systems are mostly in black box form and are often non-explainable. Though it has high performance, it is also known to make occasional fatal mistakes. This has limited the applications of AI, especially in mission critical applications. In this talk, I will present the current state-of-the arts in explainable AI, which holds promise to helping humans better understand and interpret the decisions made by the black-box AI models. This is followed by our preliminary research on explainable recommendation, relation inference in videos, as well as leveraging prior domain knowledge, information theoretic principles, and adversarial algorithms to achieving explainable framework. I will also discuss future research towards quality, fairness and robustness of explainable AI.
10:05 - 10:35am	**Group Photo and Coffee Break**	
10:35 - 10:50am	Maosong Sun	Deep Learning-based Chinese Language Computation at Tsinghua University: Progress and Challenges
10:50 - 11:05am	Minlie Huang	Controllable Text Generation
11:05 - 11:20am	Peng Cui	Stable Learning: The Convergence of Causal Inference and Machine Learning
11:20 - 11:45am	Yike Guo	Data Efficiency in Machine Learning
11:45 - 12:10pm	Zuowei Shen	Deep Approximation via Deep Learning **Abstract:** The primary task of many applications is approximating/estimating a function through samples drawn from a probability distribution on the input space. The deep approximation is to approximate a function by compositions of many layers of simple functions, that can be viewed as a series of nested feature extractors. The key idea of deep learning network is to convert layers of compositions to layers of tunable parameters that can be adjusted through a learning process, so that it achieves a good approximation with respect to the input data. In this talk, we shall discuss mathematical foundation behind this new approach of approximation; how it differs from the classic approximation theory, and how this new theory can be applied to understand and design deep learning network.

12:10 - 1:45pm	**Lunch**	
1:45 - 2:00pm	Wenwu Zhu	Explainable media and network representation
2:00 - 2:15pm	Wee Sun Lee	Neuralizing Algorithms **Abstract:** Most interesting AI problems are computationally intractable to solve in the worst case. We argue that we should be solving AI problems for the typical or average case instead, and that machine learning provides a good set of tools to do so. We illustrate the approach on approximate inference algorithms on probabilistic graphical models using our recent works on factor graph neural networks and particle filter recurrent neural networks.
2:15 - 2:30pm	Jun Zhu	Particle-based Inference for Bayesian deep learning
2:30 - 3:00pm	**Coffee Break**	
3:00 - 3:15pm	Yuanchun Shi	From Human Action Data To User Input Intention
3:15 - 3:30pm	Ke Deng	Understanding complicated patterns of Chinese texts with very weak training
3:30 - 4:00pm	David Gu	A Geometric View to Optimal Transportation and Generative Adversarial Models
4:00 - 4:30pm	Donald Rubin	Relevant Statistical Evaluations When Comparing Procedures for Analyzing Data

Friday, August 30

Time	Speaker	Title/Abstract
8:30 - 9:00am	**Breakfast**	
9:00-9:25am	Qianxiao Li	A mean-field optimal control formulation of deep learning **Abstract:** In this talk, we discuss formulating, through a continuous-time approximation, deep supervised learning as a mean field optimal control problem. This allows us to derive necessary conditions for optimality in deep learning in the form of a mean-field Pontryagin's maximum principle, as well as global characterizations of optimality using Hamilton-Jacobi Bellman equations. This forms a connection between deep learning on the one hand, and partial differential equations and the calculus of variations on the other. We also discuss interesting numerical algorithms and generalization estimates that can be derived from this viewpoint, as well as some results on function approximation using flows of dynamical systems.
9:25-10:15am	Sarah Adel Bargal	Grounding Deep Models for Improved Decision Making **Abstract**: Deep models are state-of-the-art for many computer vision tasks including object classification, action recognition, and captioning. As Artificial Intelligence systems that utilize deep models are becoming ubiquitous, it is becoming crucial to explain (ground) why they make certain decisions, and utilize such explanations (grounding) to further improve model performance. In this talk, I will present: (1) Frameworks in which grounding guides decision-making on the fly at test time by questioning whether the utilized evidence is 'reasonable', and during learning through the exploitation of new pathways in deep models. (2) A

		formulation that simultaneously grounds evidence in space and time, in a single pass, using top-down saliency. This visualizes the spatiotemporal cues that contribute to a deep recurrent neural network's classification/captioning output. Based on these spatiotemporal cues, segments within a video that correspond with a specific action, or phrase from a caption, could be localized without explicitly optimizing/training for these tasks.
10:15-10:40am	Xiaoqin Wang	Encoding and decoding auditory information by the brain
10:40-11:00am	**Coffee Break**	
11:00-11:15am	Hang Su	Adversarial attacks in deep learning
11:15-11:30am	Bin Xu	AI Practice for Gaokao: Knowledge Graph Construction for Chinese K12 Education
11:30 - 11:45am	Liu Hanzhong	Penalized regression-adjusted average treatment effect estimates in randomized experiments
11:45 - 1:30pm	**Lunch**	
1:30 - 2:15pm	Cengiz Pehlevan	Building sensory representations through learning
2:15 - 3:00pm	Sergiy Verstyuk	Modeling economic time series using deep learning
3:00 - 3:30pm	**Coffee**	
3:30 - 4:15pm	Xiao-Li Meng	Artificial Bayesian Monte Carlo Integration: A Practical Resolution to the Bayesian (Normalizing Constant) Paradox **Abstract:** Advances in Markov chain Monte Carlo in the past 30 years have made Bayesian analysis a routine practice. However, there is virtually no practice of performing Monte Carlo integration from the Bayesian perspective; indeed, this problem has earned the "paradox" label in the context of computing normalizing constants (Wasserman, 2013). We first use the modeling-what-we-ignore idea of Kong et al. (2003) to explain that the crux of the paradox is not with the likelihood theory, which is essentially the same as for a standard non-parametric probability/density estimation (Vardi, 1985); though via using group theory, it provides a richer framework for modeling the trade-off between statistical efficiency and computational efficiency. But there is a real Bayesian paradox: Bayesian analysis cannot be applied exactly for solving Bayesian computation, because to perform the exact Bayesian Monte Carlo integration would require more computation than needed to solve the original Monte Carlo problem. We then show that there is a practical resolution to this paradox using the profile likelihood obtained in Kong et al. (2006) and that this approximation is second-order valid asymptotically. We also investigate a more computationally efficient approximation via an artificial likelihood of Geyer (1994). This artificial likelihood approach is only first-order valid, but there is a computationally trivial adjustment to render its second-order validity. We demonstrate empirically the efficiency of these approximated Bayesian estimators, compared to the usual frequentist-based Monte Carlo estimators, such as bridge sampling estimators (Meng and Wong, 1996). [This is a joint work with Masatoshi Uehara.]

Saturday, August 31

Time	Speaker	Title/Abstract
8:30 - 9:00am	**Breakfast**	
9:00 - 9:45am	Brian Kulis	New Directions in Metric Learning **Abstract:** Metric learning is a supervised machine learning problem concerned with learning a task-specific distance function from supervised data. It has found numerous applications in problems such as similarity search, clustering, and ranking. Much of the foundational work in this area focused on the class of so-called Mahalanobis metrics, which may be viewed as Euclidean distances after linear transformations of the data. This talk will describe two recent directions in metric learning: deep metric learning and divergence learning. The first replaces the linear transformations with the output of a neural network, while the second considers a broader class than Mahalanobis metrics. I will discuss some of my recent work along both of these fronts, as well as ongoing attempts to combine these approaches together using a novel framework called deep divergences.
9:45 - 10:30am	Justin Solomon	Linking the Theory and Practice of Optimal Transport **Abstract:** Optimal transport is a theory linking probability to geometry, with applications across computer graphics, machine learning, and scientific computing. While transport has long been recognized as a valuable theoretical tool, only recently have we developed the computational machinery needed to apply it to practical computational problems. In this talk, I will discuss efforts with my students to scale up transport and related computations, showing that the best algorithm and model for this task depends on details of the application scenario. In particular, we will consider settings in representation learning using entropically-regularized transport, Bayesian inference using semi-discrete transport, and graphics/PDE using dynamical Eulerian models.
10:30 - 11:00am	**Coffee**	
11:00 - 11:45am	Mirac Suzgun	Towards Understanding the Limitations of Deep Learning Models for Language
11:45 - 1:15pm	**Lunch**	
1:15pm - 2:00pm	Jiafeng Chen & Suproteem Sarkar	Robust and Extensible Deep Learning for Economic and Financial Applications
2:00 - 2:45pm	Scott Kominers	