Distinguished Theme Seminar Series
Deep Learning Methods and Theory, Fall Semester, 2020

August 28, Friday, 3:30pm – 4:30pm, *On Demystifying Adversarial Learning*, Professor Lawrence Carin, Department of Electrical and Computer Engineering & Department of Statistical Science, Duke University

September 4, Friday, 10:30am – 11:30am, *Distributed Machine Learning*, Professor Heng Huang, Department of Electrical and Computer Engineering & Department of Biomedical Informatics, University of Pittsburgh

September 11, Friday, 10:30pm – 11:30pm, *A Representational Model of Grid Cells Based on Matrix Lie Algebras*, Professor Yingnian Wu, Department of Statistics, UCLA

September 18, Friday, 10:30am – 11:30am, *Integrating Domain-Knowledge into Deep Learning*, Professor Ruslan Salakahutdinov, Machine Learning Department, School of Computer Science, Carnegie Mellon University.

Titles, Abstracts and Speakers:

1. *Title*: On Demystifying Adversarial Learning

Abstract: There has been significant recent interest within the machine learning community on adversarial learning. In this setup an "actor," or model, seeks to synthesize samples that are similar to those in the training set, while a "critic" seeks to distinguish the synthesized samples from the real data. The actor and critic play an adversarial "game," and if this game is run effectively, highly realistic samples are manifested. In this talk we will derive such an adversarial learning setup from statistical first principles, and via this foundational perspective, we gain insight into failure mechanisms of such learning, and develop methods to mitigate them. We also show the application of adversarial learning to a diverse set of applications.

Dr. Lawrence Carin earned the BS, MS, and PhD degrees in electrical engineering at the University of Maryland, College Park, in 1985, 1986, and 1989, respectively. In 1989 he joined the Electrical Engineering Department at Polytechnic University (Brooklyn) as an Assistant Professor, and became an Associate Professor there in 1994. In September 1995 he joined the Electrical and Computer Engineering (ECE) Department at Duke University, where he is now a Professor. He was ECE Department Chair from 2011-2014, the Vice Provost for Research from

2014-2019, and since 2019 he has served as Duke's Vice President for Research. From 2003-2014 he held the William H. Younger Distinguished Professorship, and since 2018 he has held the James L. Meriam Distinguished Professorship. Dr. Carin's research focuses on Machine Learning (ML), Artificial Intelligence (AI) and Applied Statistics. He publishes widely in the main ML/AI conferences, and he has also engaged in translation of research to practice. He was co-founder of the small business Signal Innovations Group, which was acquired by BAE Systems in 2014, and in 2017 he co-founded the company Infinia ML. He is an IEEE Fellow.

2. *Title*: Distributed Machine Learning

Abstract: Machine learning is gaining fresh momentum, and has helped us to enhance not only many industrial and professional processes but also our everyday living. The recent success of machine learning relies heavily on the surge of big data, big models, and big computing. However, inefficient algorithms restrict the applications of machine learning to big data mining tasks. In terms of big data, serious concerns, such as communication overhead and data privacy, should be rigorously addressed when we train models using large amounts of data located on multiple devices. In terms of the big model, it is still an underexplored research area if a model is too big to train on a single device. To address these challenging problems, we focused on designing new large-scale machine learning models, efficiently optimizing and training methods for big data mining, and studying new discoveries in both theory and applications.

For the challenges raised by big data, we proposed several new asynchronous distributed stochastic gradient descent or coordinate descent methods for efficiently solving convex and non-convex problems. We also designed new large-batch training methods for deep learning models to reduce the computation time significantly with better generalization performance. For the challenges raised by the big model, we scaled up the deep learning models by parallelizing the layer-wise computations with a theoretical guarantee, which is the first algorithm breaking the lock of backpropagation mechanism such that the large-scale deep learning models can be dramatically accelerated.

Dr. Heng Huang is John A. Jurenko Endowed Professor in Electrical and Computer Engineering at University of Pittsburgh, and also Professor in Biomedical Informatics at University of Pittsburgh Medical Center. Dr. Huang received the PhD degree in Computer Science at Dartmouth College. His research areas include machine learning, big data mining, and biomedical data science. Dr. Huang has published more than 220 papers in top-tier conferences and many papers in premium journals, such as ICML, NeurIPS, KDD, RECOMB, ISMB, ICCV, CVPR, IJCAI, AAAI, Nature Machine Intelligence, Nucleic Acids Research, Bioinformatics, Medical Image Analysis, Neurobiology of Aging, IEEE TMI, TKDE, etc. Based on csrankings.org, for the last ten years, Dr. Huang is ranked 3rd among researchers who published

most top computer science conference papers. As PI, Dr. Huang currently is leading NIH R01, U01, and multiple NSF funded projects on machine learning, neuroimaging, precision medicine, electronic medical record data analysis and privacy-preserving, smart healthcare, and cyber physical system. Over the past 13 years, Dr. Huang received more than $34,000,000 research funding. He is a Fellow of AIBME and serves as the Program Chair of ACM SIGKDD Conference 2020.

3. *Title*: A Representational Model of Grid Cells Based on Matrix Lie Algebras

Abstract: A key perspective of deep learning is representation learning, where concepts are embedded in latent space and are represented by latent vectors whose units can be interpreted as neurons. In this talk, I will explain a representational model in the mammalian brain for navigation that involves grid cells and place cells. The grid cells in the mammalian medial entorhinal cortex exhibit striking hexagon firing patterns when the agent (e.g., a rat or a human) navigates in the open field. It is hypothesized that the grid cells are involved in path integral so that the agent is aware of its self-position by accumulating its self-motion. Assuming the grid cells form a vector representation of self-position, we elucidate a minimally simple recurrent model for path integral, which models the change of the vector representation given the self-motion, and we uncover two matrix Lie algebras and their Lie groups that are naturally coupled together. This enables us to connect the path integral model to the dimension reduction model for place cells via group representation theory of harmonic analysis. By reconstructing the kernel functions for place cells, our model learns hexagon grid patterns that characterize the grid cells. The learned model is capable of near perfect path integral, and it is also capable of error correction. Joint work with Ruiqi Gao, Jianwen Xie, and Song-Chun Zhu.


Dr. Ying Nian Wu is currently a professor in Department of Statistics, UCLA. He received his A.M. degree and Ph.D. degree in statistics from Harvard University in 1994 and 1996 respectively. He was an assistant professor in Department of Statistics, University of Michigan from 1997 to 1999. He joined UCLA in 1999. He has been a full professor since 2006. Wu's research areas include Generative Modeling, Representation Learning, Unsupervised Learning, Computer Vision, Computational Neuroscience, and Bioinformatics.

4. *Title*: Integrating Domain-Knowledge into Deep Learning

Abstract: I will first discuss deep learning models that can find semantically meaningful representations of words, learn to read documents and answer questions about their content. I will introduce methods that can augment neural representation of text with structured data from Knowledge Bases (KBs) for question answering, and show how we can answer complex multi-

hop questions using a text corpus as a virtual KB. In the second part of the talk, I will show how we can design modular hierarchical reinforcement learning agents for visual navigation that can perform tasks specified by natural language instructions, perform efficient exploration and long-term planning, learn to build the map of the environment, while generalizing across domains and tasks.

Dr. Ruslan Salakhutdinov is a UPMC professor of Computer Science in the Machine Learning Department, School of Computer Science at Carnegie Mellon University. He received his M.S. degree and Ph.D. degree both in Computer Science from the University of Toronto in 2003 and 2009 respectively. He works in the field of statistical machine learning. His research interests include Deep Learning, Probabilistic Graphical Models, and Large-scale Optimization. He has received numerous awards including Google Focused Award, Microsoft Research Faculty Fellowship, and Sloan Research Fellowship.