UIC ELECTRICAL AND COMPUTER ENGINEERING

Friday, January 27, 2017
Dr. Matt Nokleby, Ph.D.
Assistant Professor, Wayne State University
11AM, LC D5

Title: Information-theoretic Limits on Machine Learning Performance

The previous decade has witnessed an explosion of sophisticated techniques and an availability of massive datasets, leading to tremendous advances in the empirical performance of machine learning systems. However, theoretical understanding of these systems and their performance has not kept pace with these advances. Simple and critical questions, such as the number of signal types that a classifier can discern or the number of training samples needed to learn a data model, remain unsolved in general. This talk lays out an analytical framework that harnesses insights from information theory to provide preliminary answers to these questions. The framework has two components: the classification capacity and the information-generalization trade-off. The classification capacity characterizes the fundamental scaling law on the classification of signals lying near low-dimensional subspaces from a small number of linear measurements. By analogy to the Shannon capacity, which limits the number of codewords a receiver can distinguish in noise, the classification capacity specifies how many unique low-dimensional structures a classifier can distinguish. The information-generalization trade-off characterizes the fundamental relationship between the number of available training samples and how well one can learn a classifier in a Bayesian parametric environment. By analogy to the rate-distortion trade-off, which dictates the number of bits needed to describe a random source, the information-generalization trade-off bounds the number of training samples needed for learning in terms of information-theoretic quantities such as entropy and Fisher information. Particularly when there are few training samples, the bounds are tighter than those resulting from the well-known probably approximately correct (PAC) framework. This is joint work with Miguel Rodrigues at University College London, Ahmad Beirami at Harvard, and Robert Calderbank at Duke.

Bio:
Dr. Matthew Nokleby received the B.S. and M.S. degrees from Brigham Young University, Provo, UT, in 2006 and 2008, respectively, and the Ph.D. degree from Rice University, Houston, TX, in 2012, all in electrical engineering. From 2013-2015 he was a post-doctoral research associate at the Information Initiative @ Duke, a multidisciplinary data science initiative spanning engineering, mathematics, statistics, and computer science. In 2015 he joined Wayne State University, where he is now an assistant professor of Electrical and Computer Engineering. He has authored or coauthored numerous of journal and conference papers on the topics of distributed signal processing, wireless communications, network information theory, and machine learning. Dr. Nokleby received the Texas Instruments Distinguished Fellowship (2008-2012) and the Best Dissertation Award (2012) from the Department of Electrical and Computer Engineering at Rice University. He is a member of the IEEE Signal Processing, Information Theory, and Communications Societies.

Host: Dr. Natasha Devroye, Associate Professor
devroye@uic.edu

Electrical & Computer Engineering | University of Illinois at Chicago
851 S. Morgan St., Chicago, IL 60607
Phone: 312 996 3423 | Fax: 312 996 6465