

Multiview representation learning

Raman Arora, Ph.D.

Department of Computer Science, Johns Hopkins University

Sept. 26, 2016 @ 3:30 – 4:45 PM, CSB - Room 210

Dr. Arora is an assistant professor in the Department of Computer Science at Johns Hopkins University (JHU), where he has been since 2014. He is also affiliated with the Center of Language and Speech Processing (CLSP) and the Institute for Data Intensive Engineering and Science (IDIES). Prior to joining JHU, he was a Research Assistant Professor at Toyota Technological Institute at Chicago (TTIC), a visiting researcher at Microsoft Research Redmond and a postdoctoral research associate at the University of Washington in Seattle. He received his M.S. and Ph.D. degrees in Electrical and Computer Engineering from the University of Wisconsin-Madison. His research interests include machine learning, stochastic optimization, and statistical signal processing, with emphasis on representation learning techniques including subspace learning, multiview learning, deep learning and spectral learning. Central to his research is the theory and application of stochastic approximation algorithms that can scale to big data. (JHU)

Abstract: Unsupervised learning of useful features, or representations, is one of the most basic challenges of machine learning. Unsupervised representation learning techniques capitalize on unlabeled data which is often cheap and abundant and sometimes virtually unlimited. The goal of these ubiquitous techniques is to learn a representation that reveals intrinsic low-dimensional structure in data, disentangles underlying factors of variation, and is useful across multiple tasks and domains. This talk will focus on multiview representation learning that uses multiple "views" of data to learn improved representations for each of the views. The views can be multiple measurement modalities (audio + video, text + images, etc.) but also different information extracted from the same source (words + context, document text + links). The different views often contain complementary information, and multiview representation learning methods can take advantage of this information to learn features that are useful for understanding the structure of the data and that are beneficial for downstream tasks. Multiple views can help by reducing noise (what is noise in one view is not in the other) or improving confidence (when one view is more confident than the other). In this talk, we will focus on multiview learning of representations (features) for speech and language processing and social media analytics.

Contact Dr. Soo-Yeon Ji (sji@bowiestate.edu) if you have any question.