

University of Rochester
Department of Electrical and Computer Engineering Colloquium

Learning, Mining, and Interacting with Visual Data on the Web

Jiebo Luo
Professor
University of Rochester, Computer Science

Wednesday, October 19th
11:00 AM – 12:00 PM
Computer Studies Building (CSB) 209

Abstract: Increasingly rich and large-scale image related data are being posted to social network and media sharing websites. Researchers from multidisciplinary areas, including machine learning, computer vision, data mining, and human machine interaction, are developing methods for employing such multi-modality data for various applications. We present several recent advances in this arena of opportunities and challenges. First, we address the multi-modality feature issue by developing new machinery called Heterogeneous Feature Machines (HFM), which builds a kernel logistic regression model based on similarities that combine different features and distance metrics with group LASSO constraints. Its power is demonstrated across a wide variety of visual recognition tasks including scene, event, and action recognition. Second, we examine the recently popular data driven approach that has seemingly diminished the need for machine learning in favor of simply relying on large scale data. We believe it is important to address crucial machine learning issues, in particular cross-domain learning, in order to intelligently leverage large scale web data to solve problems such as searching personal images by keywords and recognizing events in personal videos. Next, we explore the global trends and sentiments that can be drawn by mining the sharing patterns of uploaded and downloaded social multimedia. We consider that each time an image or video is uploaded or shared, it constitutes an implicit yet trustworthy vote for (or against) the subject of the image. By aggregating such votes across millions of Internet users, we reveal the wisdom that is embedded in social multimedia sites for prediction and forecast in politics, economics, and marketing. Finally, we briefly discuss an algorithm for interactive co-segmentation of a foreground object from a group of related images. This algorithm can intelligently recommend, through active learning, where the user should scribble next in what image, and enable users to quickly achieve good quality cutouts.

Bio: Jiebo Luo joined the University of Rochester in Fall 2011 after over fifteen years at Kodak Research Laboratories, where he was a Senior Principal Scientist leading research and advanced development. He has been involved in numerous technical conferences, including serving as the program co-chair of ACM Multimedia 2010 and IEEE CVPR 2012. He is the Editor-in-Chief of the Journal of Multimedia, and has served on the editorial boards of the IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Multimedia, IEEE Transactions on Circuits and Systems for Video Technology, Pattern Recognition Machine Vision and Applications and Journal of Electronic Imaging. He is a Fellow of the SPIE, IEEE and IAPR. His research spans image processing, computer vision, machine learning, data mining, medical imaging, and ubiquitous computing. He has been an advocate for contextual inference in semantic understanding of visual data, and continues to push the frontiers in this area by incorporating geo-location context and social context. A recent research thrust focuses on exploiting social media for machine learning, data mining, and human-computer interaction, for example, mining the wisdom of crowds for social, political, and economic prediction and forecasting. He has published extensively with over 180 papers and 60 US patents. He received his PhD in Electrical Engineering from the University of Rochester in 1995.

Light refreshments will be provided.