UB University at Buffalo *The State University of New York*
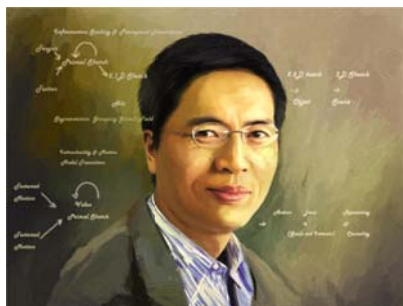
COMPUTER SCIENCE AND ENGINEERING

Distinguished Speakers Series Presents

Song-Chun, UCLA

Understanding Video and Text by Joint Spatial, Temporal and Causal Inference

Images/videos contain a vast amount of spatio-temporal patterns for objects, scenes, and events, as well as abstract knowledge for cognition and reasoning, such as causality, intents etc. In this talk, I will present a series of work done by myself and collaborators in pursuing a unified representation and model, and joint inference with video and text. It will include three parts: i) A spatial, temporal, and causal And-or Graph (STC-AoG) for representing objects, scenes, events, and causal-effects; ii) Joint spatial, temporal and causal inference which outputs a STC-parse graph as plausible interpretation of the semantic contents; and iii) Coreference resolution between the parse graph from video and the parse graph from narrative text description surrounding the video via NLP parsing. Then this joint parse graph is used to generate text description and answer human queries of who, what, where, when, and why.

Bio-sketch: Song-Chun Zhu received a Ph.D. degree from Harvard University in 1996 (advised by Dr. David Mumford). He worked at Brown, Stanford, Ohio State, before he joined in UCLA in 2002, where he is currently a professor of Statistics and Computer Science, and director of the Center for Vision, Learning, Cognition and Arts. He has published over 180 papers and received a number of honors, including the David Marr Prize in 2003 for image parsing with Z. Tu et al., the Marr Prize honorary nominations in 1999 for texture modeling and 2007 for object modeling with Y. Wu et al. As a junior faculty he received in 2001 the Sloan Fellow in Computer Science, NSF Career Award, and ONR Young Investigator Award. In 2008 he received the Aggarwal prize from the Intl Association of Pattern Recognition for "contributions to a unified foundation for visual pattern conceptualization, modeling, learning, and inference". In 2013 he received the Helmholtz Test-of-time prize at ICCV. He is a fellow of IEEE Computer Society.

February 6, 2014

3:30–4:30 PM

University at Buffalo - North Campus – Davis 101
This talk is free and open to the public