Using Perception to Supervise Language Learning and Language to Supervise Perception

Raymond Mooney
Department of Computer Sciences
University of Texas at Austin
mooney@cs.utexas.edu

Abstract

Using Perception to Supervise Language Learning and Language to Supervise Perception less... Learning to understand and generate language and learning to interpret perceptual input are both difficult problems that usually require human-annotated training data. Our current research focuses on developing methods that can use perceptual data to supervise language learning and linguistic data to supervise the interpretion of images and video. This talk will survey three of our recent projects: 1) Learning by example to linguistically describe simulated robot soccer games, 2) Using co-training with separate linguistic and visual views to perform semi-supervised classification of captioned images and videos. and 3) Using closed-captions to train activity recognizers that improve video retrieval.

Biography

Raymond J. Mooney is a Professor in the Department of Computer Sciences at the University of Texas at Austin. He received his Ph.D. in 1988 from the University of Illinois at Urbana/Champaign. He is an author of over 150 published research papers, primarily in the areas of machine learning and natural language processing. He is the current President of the International Machine Learning Society, was program co-chair for the 2006 AAAI Conference on Artificial Intelligence, general chair of the 2005 Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, and co-chair of the 1990 International Conference on Machine Learning. He is a Fellow of the American Association for Artificial Intelligence and recipient of best paper awards from the National Conference on Artificial Intelligence, the SIGKDD International Conference on Knowledge Discovery and Data Mining, the International Conference on Machine Learning, and the Annual Meeting of the Association for Computational Linguistics. His recent research has focused on learning for natural-language processing, text mining for bioinformatics, statistical relational learning, and transfer learning.