

Knowledge-Augmented Deep Neural Networks for Facial Action Detection

Qiang Ji, jiq@rpi.edu

Rensselaer Polytechnic Institute

Existing facial action unit (AU) detection methods cannot generalize well to novel datasets and subjects, partially because of lack of sufficient quality AU annotations. To address this problem, we propose to systematically exploit generic prior knowledge that governs the facial expression generation and behaviors, and integrate the identified generic knowledge with data to yield data-efficient and generalizable AU detection. Specifically, we first identify AU prior knowledge from two sources: the facial muscle activation principles that govern AU dependencies and the existing psychological studies that reveal the dependencies between AU and facial expressions. We then propose to use probabilistic constraints to represent the identified AU dependencies and their uncertainties. A constraint optimization method is then introduced to learn a Bayesian Network (BN) to compactly and rigorously encode the structured and uncertain dependencies among AU and expression. The BN is then integrated into a deep learning framework as a weak supervision to train an AU detection model. Evaluations on benchmark datasets demonstrate the effectiveness of our AU detection model in both data efficiency and its generalization performance. In particular, for within-dataset evaluation, the proposed AU detection model achieves competitive performance without any AU annotations, and for cross-dataset evaluation, it significantly generalizes better than state of the art fully supervised, semi and weakly supervised methods.

Bio for Qiang Ji

Qiang Ji is a Professor with the Department of Electrical, Computer, and Systems Engineering at Rensselaer Polytechnic Institute (RPI). He received his Ph.D degree in Electrical Engineering from the University of Washington. He was a program director at the National Science Foundation, where he managed NSF's computer vision and machine learning programs. He also held teaching and research positions at University of Illinois at Urbana-Champaign, Carnegie Mellon University, and University of Nevada at Reno.

His research interests are in human-centered computer vision, probabilistic graphical models, probabilistic deep learning, and their applications. He has published over 300 papers, received multiple awards for his work, serve as an editor for multiple international journals, and organize numerous international conferences/workshops. He is a fellow of the IEEE and the IAPR.