CS Seminar

Deep Learning for Visual Perception: Algorithms, Applications and Theory

Dr Ping LUO
Department of Electronic Engineering
The Chinese University of Hong Kong

Date:
January 30, 2019
Wednesday
10:00 am

Venue:
Room 308
Chow Yei Ching Building
The University of Hong Kong

Abstract:

This talk covers both background and frontier of deep learning in three aspects, including an introduction of deep neural networks, image and video analyses via deep learning, and theoretical understanding of deep learning. In the first part, I will briefly introduce convolutional neural network, which is the core algorithm to solve many problems in computer vision. In the second part, I will go through multiple vision-based systems, including holistic face analyses, object detection and segmentation, AI in Fashion, video understanding, as well as autonomous driving and robot arms. In the last part, I will cover advanced topics by introducing a new family of normalization techniques for deep learning and their importance to theoretically understand deep learning. The following questions will be answered. (1) Why do deep neural networks work so well? (2) What do normalization methods help in deep learning? (3) Is reinforcement learning necessary to solve Neural Architecture Search (NAS)?

About the Speaker:

Ping Luo received his PhD degree in 2014 from Information Engineering, the Chinese University of Hong Kong (CUHK), supervised by Xiaoou Tang and Xiaogang Wang. He was a Postdoctoral Fellow in CUHK from 2014 to 2016. He joined SenseTime Research as a Principal Research Scientist from 2017 to 2018. Now he joins the department of Electronic Engineering of CUHK as a Research Assistant Professor. His research interests are machine learning and computer vision. He has published 60+ peer-reviewed articles in top-tier conferences and journals such as TPAMI, IJCV, ICML, ICLR, CVPR, and NIPS. His work has high impact with 5500 citations according to Google Scholar. For instance, he has presented the first ConvNet for face image synthesis by using deep learning (CVPR12, ICCV13, NIPS14), the first ConvNet to solve Markov Random Field (MRF) for image segmentation (ICCV15), and the first multitask deep learning for facial landmark detection (ECCV14). He also built popular benchmarks for the research community such as DeepFashion, CelebA, and WIDERFace. His recent researches also focus on theoretical understanding normalization techniques in deep learning such as GWNN (ICML17), EigenNet (IJCAI17), Kalman Normalization (NIPS18), Switchable Normalization (ICLR19), and Regularization of BN (ICLR19). He won a number of competitions and awards such as the first runner up in 2014 ImageNet ILSVRC Challenge, the first place in 2017 DAVIS Challenge on Video Object Segmentation, Gold medal in 2017 Youtube-8M Video Classification Challenge, the first place in 2018 Drivable Area Segmentation Challenge for Autonomous Driving, 2011 HK PhD Fellow Award, and 2013 Microsoft Research Fellow Award (ten PhDs in Asia).