Mingming Gong: From Statistical to Causal Learning

Erfani and Ma: Deep learning security: adversarial attack and defense

Ehsan Abbasnejad: Deep learning methods, practices and applications

Details:

From Statistical to Causal Learning

Speaker: Mingming Gong

Abstract:
Statistical dependence is the main driving force for current machine learning systems. With the learned dependence, one can estimate beliefs or probabilities when the data generating process is fixed. However, we humans are able to adapt ourselves to new situations and make accurate predictions and decisions. This is because we have a good causal understanding of the world and therefore, we could manipulate our experience at will to imagine hypothetical situations for learning. It is crucial to equip machines with causal learning capabilities, so that they can predict the effects of interventions and generalize to new environments. In this tutorial, we will survey the rich body of literature on mathematical foundations and methodology development of causal learning. Specifically, we will introduce how we can infer causal relations from various kinds of observational data, such as time series, noisy data, and nonstationary/heterogeneous data. In addition, we will discuss how causal information facilitates understanding and tackling complex transfer learning problems in which the data distributions change across domains.

Biography:
https://mingming-gong.github.io/

Deep learning security: adversarial attack and defense

Speakers: Dr Sarah Erfani and Dr Xingjun Ma

Abstract:
Deep learning has become increasingly popular in the past few years. This is largely attributed to a family of powerful models called deep neural networks (DNNs). With many stacked layers, and millions of neurons, DNNs are capable of learning complex non-linear mappings, and have demonstrated near or even surpassing human-level performance in a wide range of applications such as image classification, object detection, natural language processing, speech recognition self-driving cars, playing games or medical diagnosis. Despite their great success, DNNs have recently been found vulnerable to adversarial examples (or attacks), which are input instances slightly modified in a way that is intended to fool the model. Such a surprising weakness of DNNs has raised security and reliability concerns on the development of deep learning systems in safety-critical scenarios such as face recognition, autonomous driving and medical diagnosis. Since the first discovery, this has attracted a huge volume of work on either attacking or defending DNNs against these attacks. In this

tutorial, we will introduce this adversarial phenomenon, explanations to this phenomenon, and techniques that have been developed for both attack and defense.

Biography:
Dr Sarah Erfani is a Senior Lecturer in the School of Computing and Information Systems at The University of Melbourne. Her research interests include machine learning, large-scale data mining, cyber security, and data privacy. Find out more about her at https://people.eng.unimelb.edu.au/smonazam/

Dr Xingjun Ma currently works as a research fellow at the School of Computing and Information Systems, The University of Melbourne, where he obtained his PhD. Prior to his PhD, Xingjun received his M.Eng. and B.Eng. degrees from Tsinghua University and Jilin University successively. Xingjun is a passionate researcher in the field of AI, machine learning and deep learning. His works have been published at prestigious international conferences such as ICML, ICLR, CVPR, ICCV, IJCAI, and AAAI. Find more about him at http://xingjunma.com/

Deep learning methods, practices and applications

Speaker: Ehsan Abbasnejad

Abstract:
In this talk we will discuss the principles of machine learning--in particular, deep learning--and their applications. We will introduce the convolutional and recurrent networks and briefly discuss how they are used in computer vision and natural language processing tasks.

Biography:
Dr. Ehsan Abbasnejad is a Lecturer in the University of Adelaide and is a member of Australian Institute for Machine Learning. He was awarded his PhD degree in 2015 in Computer Science from the Australian National University. He has extensive experience in deep learning for vision and language applications. His past experience in industry includes Microsoft, Xerox and NEC. He has been a research scientist and founding member of DeepSightX who won the second prize in the global mineral discovery challenge.
