For favour of Posting

THE HONG KONG UNIVERSITY OF SCIENCE & TECHNOLOGY

Department of Mathematics

SEMINAR ON STATISTICS AND DATA SCIENCE

Gaussian Differential Privacy, with Applications to Deep Learning

By

Prof. Weijie SU & Dr. Jinshuo DONG

University of Pennsylvania

Abstract

Privacy-preserving data analysis has been put on a firm mathematical foundation since the introduction of differential privacy (DP) in 2006. This privacy definition, however, has some well-known weaknesses: notably, it does not tightly handle composition. This weakness has inspired several recent relaxations of differential privacy based on the Renyi divergences. We propose an alternative relaxation we term "f-DP", which has a number of nice properties and avoids some of the difficulties associated with divergence based relaxations. First, f-DP preserves the hypothesis testing interpretation of differential privacy, which makes its guarantees easily interpretable. It allows for lossless reasoning about composition and post-processing, and notably, a direct way to analyze privacy amplification by subsampling. We define a canonical single-parameter family of definitions within our class that is termed "Gaussian Differential Privacy", based on hypothesis testing of two shifted normal distributions. We prove that this family is focal to f-DP by introducing a central limit theorem, which shows that the privacy guarantees of any hypothesis-testing based definition of privacy (including differential privacy) converge to Gaussian differential privacy in the limit under composition. This central limit theorem also gives a tractable analysis tool. We demonstrate the use of the tools we develop by giving an improved analysis of the privacy guarantees of noisy stochastic gradient descent. This is joint work with Jinshuo Dong and Aaron Roth.

Biography: *Weijie Su is an Assistant Professor in the Wharton Statistics Department at the University of Pennsylvania, where he co-directs the Penn Research in Machine Learning. Prior to joining Penn, he received his Ph.D. in Statistics from Stanford University in 2016 and his B.S. in Mathematics from Peking University in 2011. His research interests span mathematical optimization, high-dimensional statistics, privacy-preserving data analysis, multiple hypothesis testing, and deep learning theory. He is a recipient of the Stanford Theodore W. Anderson Dissertation Award in 2016, an NSF CAREER Award in 2019, and an Alfred P. Sloan Research Fellowship in 2020.*

Date : 10 April, 2020 (Friday)
Time : 9:30am – 11:00am
Zoom Meeting : https://hkust.zoom.com.cn/j/5616960008

All are Welcome!