Generalization and Learnability under Differential Privacy and its Variants

Yu-Xiang Wang (CMU)

I will be talking about two of our recent work that investigate the connections between notions of privacy and quantities such as stability, generalization and learnability in statistical learning theory.

In particular, I will present our investigation on the following two questions of interest:

1. What problems are still learnable under a privacy constraint?
2. To what extent can we weaken differential privacy, but still preserve its property: "Privacy => Generalization”?

For the first question, we present a characterization of the learnability under differential privacy and (ε, δ)-approximate differential privacy and discuss the separation of private learnability and non-private learnability. For the second question, we propose On-Average KL-Privacy and show that it characterizes the generalization for a wide class of algorithms that sample from a Gibbs distribution. On-Average KL-Privacy also preserves other properties of differential privacy including: small group privacy, adaptive composition and closeness to post-processing. We experimentally illustrate the orders-of-magnitudely more favorable privacy-utility tradeoff under this new notion of privacy than DP.

References:

"Learning with Differential Privacy: Stability, Learnability and the Sufficiency and Necessity of ERM Principle”, http://arxiv.org/abs/1502.06309

"On-Average KL-Privacy and its equivalence to Generalization for Max-Entropy Mechanisms”, http://arxiv.org/abs/1605.02277