The University Of Sheffield.

Automatic Control and Systems Engineering

The Department of Automatic Control & Systems Engineering
is pleased to announce the following seminar:

Optimal Analyses of Differential Privacy Using Divergences

Dr Borja Balle
Machine Learning Scientist at Amazon Research Cambridge

Wednesday, 14 November 2018 at 14:00
Sir Henry Stephenson Building, LT02

Abstract

Differential privacy provides a robust mathematical formulation of privacy in the context of data analysis algorithms. To protect the privacy of the individuals whose data is being analyzed, differential privacy injects carefully crafted noise into the computation. Higher levels of noise provide stronger privacy, but also lead to worse accuracies. This tension motivates the study of optimal privacy-accuracy trade-offs.

In this talk I will present a powerful formulation of differential privacy in terms of f-divergences, and discuss two recent applications of this point of view. The first result is an exact method for calibrating the well-known Gaussian output perturbation mechanism [Balle & Wang, ICML'18]. The second result is a tight analysis of the "privacy amplification by subsampling" phenomenon [Balle, Barthe & Gaboardi, NIPS'18], which is a fundamental building block of differentially private SGD algorithms. Our results highlight the practical importance of using optimal constants (as opposed to only optimal rates) when balancing privacy-accuracy trade-offs.

Biography

Borja Balle is currently a Machine Learning Scientist at Amazon Research in Cambridge. Before joining Amazon, Borja was a lecturer at Lancaster University (2015-2017), a postdoctoral fellow at McGill University (2013-2015), and a graduate student at Universitat Politecnica de Catalunya where he obtained his PhD in 2013. His main research interest is in privacy-preserving machine learning, including the use of differential privacy and multi-party computation in distributed learning problems, and the foundations of privacy-aware data science. More info: https://borjaballe.github.io