

Samory Kpotufe

(Columbia University, New York, USA)

Title:

“Some Recent Insights on Transfer-Learning”

Abstract:

A common situation in Machine Learning is one where training data is not fully representative of a target population due to bias in the sampling mechanism or due to prohibitive target sampling costs. In such situations, we aim to 'transfer' relevant information from the training data (a.k.a. source data) to the target application. How much information is in the source data about the target application? Would some amount of target data improve transfer? These are all practical questions that depend crucially on 'how far' the source domain is from the target. However, how to properly measure 'distance' between source and target domains remains largely unclear.

In this talk we will argue that much of the traditional notions of 'distance' (e.g. KL-divergence, extensions of TV such as D_A discrepancy, density-ratios, Wasserstein distance) can yield an over-pessimistic picture of transferability. Instead, we show that some new notions of 'relative dimension' between source and target (which we simply term 'transfer-exponents') capture a continuum from easy to hard transfer. Transfer-exponents uncover a rich set of situations where transfer is possible even at fast rates; they encode relative benefits of source and target samples, and have interesting implications for related problems such as 'multi-task or multi-source learning'.

In particular, in the case of transfer from multiple sources, we will discuss (if time permits) a strong dichotomy between minimax and adaptive rates: no adaptive procedure exists that can achieve the same rates as minimax (oracle) procedures.

The talk is based on earlier work with Guillaume Martinet, and ongoing work with Steve Hanneke.