

Speaker: Timothy Canning (University of Edinburgh)

Date: 12/10/2021 on Microsoft Teams

Title: Adaptive Transfer Learning

Abstract:

In transfer learning, we wish to make inference about a target population when we have access to data both from the distribution itself, and from a different but related source distribution. We introduce a flexible framework for transfer learning in the context of binary classification, allowing for covariate dependent relationships between the source and target distributions that are not required to preserve the Bayes decision boundary. Our main contributions are to derive the minimax optimal rates of convergence (up to polylogarithmic factors) in this problem, and show that the optimal rate can be achieved by an algorithm that adapts to key aspects of the unknown transfer relationship, as well as the smoothness and tail parameters of our distributional classes. This optimal rate turns out to have several regimes, depending on the interplay between the relative sample sizes and the strength of the transfer relationship, and our algorithm achieves optimality by careful, decision tree-based calibration of local nearest-neighbour procedures.

This is joint work with Henry Reeve (Bristol) and Richard Samworth (Cambridge), and the associated paper is available at <https://arxiv.org/abs/2106.04455>.