Algorithmic Fairness: A Post-Processing Approach

Christophe Denis

Université Paris 1 Panthéon-Sorbonne

10 septembre 2026

Algorithmic fairness has become a major area of machine learning, aiming to mitigate the impact of biases present in data and algorithmic decision-making. In this talk, we consider post-processing methods for achieving group fairness in the settings of multiclass classification and regression. We derive explicit expressions for the corresponding optimal fair predictors and present a plug-in post-processing algorithm to estimate them. Finally, we discuss both the theoretical guarantees of the proposed approach and its empirical performance through numerical experiments.