

Oracle inequality of the Exponentially Weighted Aggregate with Laplace prior

Edwin Grappin

We study the statistical behaviour of the Exponentially Weighted Aggregate (EWA) in the problem of high-dimensional regression with fixed design. Under the assumption that the underlying regression vector is sparse, it is reasonable to use the Laplace distribution as prior. The resulting estimator and, specifically, a particular instance of it referred to as the Bayesian lasso, was already used in the statistical literature because of its computational convenience. This work establishes sharp oracle inequalities for the EWA with the Laplace prior. These inequalities show that if the temperature parameter is small, the EWA with the Laplace prior satisfies the same type of oracle inequality as the lasso estimator does, as long as the quality of estimation is measured by the prediction loss. Joint work with Arnak Dalalyan and Quentin Paris.