Estimation of S-shaped functions and beyond

Yining Chen

LSE

Abstract:

We consider the nonparametric estimation of an S-shaped regression function. The least squares estimator provides a very natural, tuning-free approach, but results in a non-convex optimisation problem, since the inflection point is unknown. We show that the estimator may nevertheless be regarded as a projection onto a finite union of convex cones, which allows us to propose a mixed primal-dual bases algorithm for its efficient, sequential computation. Our main theoretical results provide sharp oracle inequalities that yield worst-case and adaptive risk bounds for the estimation of the regression function, as well as a rate of convergence for the estimation of the inflection point. These results reveal that the estimator achieves the minimax or optimal rate of convergence for both the estimation of the regression function and its inflection point (up to a logarithmic factor in the latter case). We then discuss how our approach can be extended to handle additive models, as well as how it can be used for multiple feature detection in a change-point analysis framework.

This is joint work with Raymond Carroll, Oliver Feng, Piotr Fryzlewicz, Qiyang Han and Richard Samworth.