Smooth re-parametrizations for sparse regularization

Clarice Poon
University of Bath
cmhsp20@bath.ac.uk

Many low-complexity regularizers such as ℓ_1 or nuclear norm regularisation admit a quadratic variational form and this is the basis of the celebrated iteratively re-weighted least squares (IRLS) algorithms. In this work, we show how a simple re-parametrization of IRLS coupled with a bilevel formulation can lead to a smooth non-convex optimisation problem that can be handled with robust numerical procedures such as BFGS. Our formulation allows us to efficiently handle many popular optimisation settings, including the Lasso, total variation regularisation problems, sparsity problems with linear constraints, robust losses (such as the ℓ_1 loss), as well as non-convex ℓ_q minimisation. Our numerical benchmarking experiments demonstrate that this approach is highly versatile, leading to substantial performance advantages in a wide range of settings.

In additional to enabling the use of robust smooth optimisation tools, some of the reasons for the favourable behaviour of our proposed method are as follows. Although the resultant problem is non-convex, we show that there are no spurious minima since all saddle points are strict. Moreover, in the "fine-grids" setting where the columns of the matrix can be highly coherent with each other (this often occurs due to discretization in continuous problems), this re-parametrization leads to dimension free convergence rates – this is in stark contrast typical optimisation schemes based on the Euclidean distance such as Forward-Backward where convergence rates depend on the problem dimension.

Joint work with: Gabriel Peyré