

Large-scale multi-task learning with gauge regularization penalty

Zaid Archaoui

We study multi-task learning problems with general sparsity-inducing matrix regularization penalties. We formulate the matrix regularizers as gauge functions, and, using their structure, we lift the optimization problem in a higher space where we propose to apply a coordinate descent algorithm. Our framework allows to efficiently tackle difficult matrix-regularized objectives, e.g. with a trace-norm regularization penalty. We present experimental results on synthetic datasets and on real-world large-scale computer vision datasets. Our algorithm is competitive and often outperforms existing approaches on those problems.