Universität Basel

Publication

A complete analysis of the l_1,p Group-Lasso

ConferencePaper (Artikel, die in Tagungsbänden erschienen sind)

The Group-Lasso is a well-known tool for joint regularization in machine learning methods. While the $l_{1,2}$ and the $l_{1,\infty}$ version have been studied in detail and efficient algorithms exist, there are still open questions regarding other $l_{1,p}$ variants. We characterize conditions for solutions of the $l_{1,p}$ Group-Lasso for all p-norms with $1 >= p >= \infty$, and we present a unified active set algorithm. For all p-norms, a highly efficient projected gradient algorithm is presented. This new algorithm enables us to compare the prediction performance of many variants of the Group-Lasso in a multi-task learning setting, where the aim is to solve many learning problems in parallel which are coupled via the Group-Lasso constraint. We conduct large-scale experiments on synthetic data and on two real-world data sets. In accordance with theoretical characterizations of the different norms we observe that the weak-coupling norms with p between 1.5 and 2 consistently outperform the strong-coupling norms with p « 2.