

Speeding up Greedy Forward Selection for Regularized Least-Squares

Tapio Pahikkala, Antti Airola, and Tapio Salakoski
Department of Information Technology
University of Turku and Turku Centre for Computer Science,
Turku, Finland,
e-mail: firstname.lastname@utu.fi.

Abstract

We propose a novel algorithm for greedy forward feature selection for regularized least-squares (RLS) regression and classification, also known as the least-squares support vector machine or ridge regression. The algorithm, which we call greedy RLS, starts from the empty feature set, and on each iteration adds the feature whose addition provides the best leave-one-out cross-validation performance. Our method is considerably faster than the previously proposed ones, since its time complexity is linear in the number of training examples, the number of features in the original data set, and the desired size of the set of selected features. Therefore, as a side effect we obtain a new training algorithm for learning sparse linear RLS predictors which can be used for large scale learning. This speed is possible due to matrix calculus based shortcuts for leave-one-out and feature addition. We experimentally demonstrate the scalability of our algorithm compared to previously proposed implementations.