

Regularization Paths for Generalized Linear Models via Coordinate Descent

Jerome Friedman¹, Trevor Hastie^{1,2,*}, Rob Tibshirani^{2,1}

1. Statistics Department, Stanford University

2. Department of Health, Research and Policy, Stanford University

* Contact author: hastie@stanford.edu

Keywords: regularization, lasso, path, coordinate descent, logistic regression

We develop fast algorithms for estimation of generalized linear models with convex penalties. The models include linear regression, two-class logistic regression, and multinomial regression problems while the penalties include ℓ_1 (the lasso), ℓ_2 (ridge regression) and mixtures of the two (the elastic net). The algorithms use cyclical coordinate descent, computed along a regularization path. The methods can handle large problems and can also deal efficiently with sparse features. In comparative timings we find that the new algorithms are considerably faster than competing methods. The software is on CRAN in the package `glmnet`.

References

- Jerome Friedman, Trevor Hastie, and Robert Tibshirani, *Pathwise coordinate optimization*, Annals of Applied Statistics **1** (2007), no. 2, 302–332.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani, *Regularization paths for generalized linear models via coordinate descent*, Tech. report, Stanford University, 2008.
- Trevor Hastie, Rob Tibshirani, and Jerome Friedman, *The elements of statistical learning: Prediction, inference and data mining (second edition)*, Springer Verlag, New York, 2009.