

In modern applications, we are interested in fitting regression models that are fast (i.e. can be fit to high-dimensional data), flexible (e.g. do not assume a linear conditional mean relationship between the response and the predictors), and interpretable (i.e. a person can make sense of the fitted model). I'll present two recent developments towards this end – the fused lasso additive model (FLAM), and convex regression via interpretable sharp partitions (CRISP).