8 False discovery rate smoothing

James Scott (University of Texas at Austin) with Wesley Tansey, Sanmi Koyejo and Russ Poldrack

Many approaches for multiple testing begin with the assumption that all tests in a given study should be combined into a global false-discovery-rate analysis. But this may be inappropriate for many of today's large-scale screening problems, where test statistics have a natural spatial lattice structure (voxels in the brain, distance along the chromosome), and where a combined analysis can lead to poorly calibrated error rates. To address this problem, we introduce an empirical-Bayes approach called false-discovery-rate smoothing. FDR smoothing automatically finds spatially localized regions of significant test statistics. It then relaxes the threshold of statistical significance within these regions, and tightens it elsewhere, in a manner that controls the overall false-discovery rate at a given level. This results in increased power and cleaner spatial separation of signals from noise. The approach requires solving a non-standard high-dimensional optimization problem, for which an efficient augmented-Lagrangian algorithm is presented. We demonstrate that FDR smoothing exhibits state-of-the-art performance on simulated examples. We also apply the method to a data set from an fMRI experiment on spatial working memory, where it detects patterns that are much more biologically plausible than those detected by existing FDR-controlling methods. All code for FDR smoothing is publicly available in Python and R.