User's Guide to the Robust FDR Routine
Stan Pounds
June 13, 2006

This document describes how to use the R/S-plus routines to implement the robust method of FDR estimation proposed by Pounds and Cheng (2006).

Getting Started

Download the library code file (*robust-fdr.ssc* or *robust-fdr.R*) from www.stjuderesearch.org/depts/biostats. Save it to a desired location. Then use the source() command to make the routines available for use.

Using the Library

The end-user will be primarily interested in the routine `robust.fdr`. The `robust.fdr` routine accepts up to 5 arguments. The first argument, `p`, is the vector of p-values from the analysis. The second argument, `sides`, indicates whether those p-values are one-sided (set `sides=1`) or two-sided (set `sides=2`). By default, `sides=1`. The third argument, `p2`, is optional and may be provided if the p-values in `p` are one-sided (default=1-`p`). If provided, `p2` should be a vector of p-values for testing the "other alternative." The default for `p2` is 1-`p`, following arguments in the paper. Providing `p2` may slightly improve power; however, this assertion has not been thoroughly explored at this time. The fourth argument, `discrete`, is a logical (T/F) that indicates whether the p-values are discrete (default=F). The fifth argument, `use8`, indicates whether the constant 8 should be used in computing the estimate of the null proportion (default=T; see paper for more details). The `use8` argument is relevant only if discrete=T.

The `robust.fdr` routine returns a list with the following components:

`p`: the vector of p-values provided by the user
`fdr`: the smoothed local FDR estimates
`q`: the q-value computed from the smoothed local FDR estimates
`cdf`: the p-value empirical distribution function
`loc.fdr`: the unsmoothed local FDR estimates
`fp`: the estimated number of false positives at the p-value threshold in p
`fn`: the estimated number of false negatives at the p-value threshold in p
`te`: the estimated total number of errors in the inference at the p-value threshold in p
`pi`: the estimated proportion of tests with a true null hypothesis
`ord`: a vector to order results by ascending p-value

Example

Below is an example of how to use the library.

```
# Use source to make routines available
source("robust-fdr.R")

# Generate some p-values for purposes of illustration
p<-rbeta(1000,0.6,1)

# Call robust.fdr and store results in an object called res
res<-robust.fdr(p,sides=2,discrete=F)

# Plot the q-value curve
plot(res$p,res$q)

# Create a result table
restab<-cbind(p=res$p,q=res$q)

# Order the result table
restab<-restab[res$ord,]

# View first 10 entries of the ordered result table
restab[1:10,]
```