

R-package “simctest” R-class “mmctest” A Short Introduction

Axel Gandy and Georg Hahn

September 17, 2012

This document describes briefly how to use the class “mmctest”, included in the R-package “simctest”. It implements the methods from “A Sequential Sampling Procedure for Multiple Testing Problems with Bounded Risk of Classification Errors”, based on Gandy and Hahn [2012].

The class can be used to evaluate the statistical significance of each hypothesis in a multiple testing problem.

1 Installation

The functions described in this document are included in the R-package “simctest”. Please see the documentation of “simctest” on how to install the package.

2 Usage

The package is loaded by typing

```
> library(simctest)
```

This document can be accessed via

```
> vignette("simctest-mmctest-intro")
```

Documentation of the most useful commands can be obtained as follows:

```
> ? simctest  
> ? mcp  
> ? mmctest
```

2.1 Implementing a Monte Carlo multiple testing problem

The following is an artificial example. Implementing a Monte Carlo multiple testing problem consists of two stages.

Firstly, an interface to draw samples has to be provided. This can be done in two ways, either by implementing the generic class `mmctSamplerGeneric` or by directly providing the number m of hypotheses and a function f which generates samples. Both ways are described in the next section.

Secondly, an object of type `mmctest` has to be created. It provides a “run” method which uses the `mmctest`-object and an `mmctSamplerGeneric`-object to evaluate the multiple testing problem.

The algorithm used in class `mmctest` is the one introduced in Gandy and Hahn [2012]. The multiple testing problem is evaluated until at least one of four stopping criteria is satisfied, see below for a detailed description.

Stopped tests can be resumed with the “cont” function.

Printing an object of type `mmctest` or `mmctestres` will display the number of already rejected and non-rejected hypotheses, the number of undecided hypotheses and the total number of samples drawn up to the current stage.

2.1.1 Implementing the sampling interface

An interface for drawing new samples has to be provided for each multiple testing problem.

If new samples are simply generated by a function f , the derived class `mmctSampler` provided in `simctest` can be used as a shortcut. It works as follows: Any function f used to draw new samples has to be able to accept the arguments “ind”, a vector with indices of hypotheses and a number “n” of trials. The function f has to return a vector containing the number of significant test statistics for each hypothesis specified in “ind” based on “n” additional samples. For instance, passing a vector “ind” of (2,5) and a number “n” of 10 as arguments means that 10 more samples are requested for the hypotheses with index 2 and 5. The function f might need further data to evaluate the tests. Such data can be passed on to f as third argument `data`. For instance,

```
> fun <- function(ind, n, data) sapply(1:length(ind), function(i) sum(runif(n) <=
+   data[ind[i]]))
```

is a function which draws samples from hypotheses having p-values given in vector `data`.

The package `mmctest` provides a shortcut which can be used to easily specify the interface. Given a function `fun` to draw samples and the number `num` of hypotheses, `fun` and `num` (and additional data `data`) can be passed on to the class `mmctSampler` which returns a derived object of the generic class `mmctSamplerGeneric`. For example,

```
> s <- mmctSampler(fun, num = 500, data = c(rep(0, 100), runif(400)))
```

returns an sampler interface `s` for the function `fun` defined above and 500 p-values used to draw new samples.

The class `mmctSamplerGeneric` can also directly be overwritten with an own sampler interface. Any sampler has to implement the two generic functions `getSamples` and `getNumber`:

```
> setClass("mmctSampler1", contains = "mmctSamplerGeneric", representation = representation(data = "num
```

```
[1] "mmctSampler1"
```

```
> setMethod("getSamples", signature(obj = "mmctSampler1"), function(obj,
+   ind, n) {
+   sapply(1:length(ind), function(i) {
+     return(sum(runif(n) <= obj@data[ind[i]]))
+   })
+ })
```

```
[1] "getSamples"
```

```
> setMethod("getNumber", signature(obj = "mmctSampler1"), function(obj) {
+   return(length(obj@data))
+ })
```

```
[1] "getNumber"
```

In this case, the sampler will be

```
> s <- new("mmctSampler1", data = c(rep(0, 100), runif(400)))
```

2.1.2 A simple run of the algorithm

After having specified the sampler, the main algorithm can be executed. This is done by creating an object of type `mmctest` using the pseudo-constructor

```
mmctest(epsilon=0.01, threshold=0.1, h),
```

where `epsilon` is the overall error one is willing to spend (see Gandy and Hahn [2012]), `threshold` is the multiple testing threshold and h is the multiple testing procedure.

Any function

```
h <- function(p, threshold) ...
```

can be used as a multiple testing procedure as long as it takes a vector p of p-values and a threshold `threshold` as arguments and returns the indices of all rejected hypotheses as vector.

The Benjamini-Hochberg procedure `hBH`, its modification by Pounds and Cheng [2006] `hPC` and the Bonferroni correction `hBonferroni` are available by default:

```
> s <- mmctSampler(fun, num = 500, data = c(rep(0, 100), runif(400)))
> m <- mmctest(h = hBH)
```

The algorithm can now be started by calling

```
run(alg, gensample, maxsteps=list(maxit=0, maxnum=0, undecided=0, elapsedsec=0))
```

which takes an object `alg` of type `mmctest`, a sampler object `gensample` to generate samples and a list `maxsteps` as stopping criterion. The list `maxsteps` can include a maximal number of iterations `maxit` after which the algorithm stops, a maximal total number of samples `maxnum` drawn until stopping, a number `undecided` of undecided hypotheses one is willing to tolerate or a time constraint `elapsedsec` in seconds.

Specifying other items in list `maxsteps` will lead to an error message and an empty list will be reset to the default list

```
list(maxit=0, maxnum=0, undecided=0, elapsedsec=0).
```

As an example, the following lines evaluate the previously created multiple testing problem `m` using the Benjamini-Hochberg procedure `hBH` and the previous sampler `s`. The algorithm stops before reaching more than a total of 1000000 samples or after all but 20 hypotheses are classified:

```
> m <- run(m, s, maxsteps = list(maxnum = 1e+06, undecided = 20))
> m
```

```
Number of rejected hypotheses: 100
Number of non-rejected hypotheses: 380
Number of unclassified hypotheses: 20
Total number of samples: 109562
```

Printing the object displays the number of already rejected and non-rejected hypotheses, the number of undecided hypotheses and the total number of samples drawn up to the current stage.

A formatted summary of the indices belonging to rejected and undecided hypotheses can be printed via `summary.mmctestres`. All indices not printed belong to non-rejected hypotheses.

```
> summary.mmctestres(m)

Number of hypotheses: 500
Indices of rejected hypotheses: 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39
40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62
63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85
86 87 88 89 90 91 92 93 94 95 96 97 98 99 100
Indices of unclassified hypotheses: 105 113 157 184 214 243 255 267
276 279 291 298 309 316 331 382 385 419 434 482
All hypotheses not listed are classified as not rejected.
```

2.1.3 Continuing a run of the algorithm

Each run can be continued with the `cont` function using a new stopping criterion:

```
> m <- cont(m, steps = list(undecided = 10))
> m
```

```
Number of rejected hypotheses: 104
Number of non-rejected hypotheses: 386
Number of unclassified hypotheses: 10
Total number of samples: 161971
```

Here, the algorithm has been applied again to the previously stopped multiple testing problem `m`. It has been resumed until all but 10 hypotheses were classified.

As before, `maxit`, `maxnum`, `undecided` and `elapsedsec` are valid stopping criteria for parameter `steps` of function `cont`.

2.1.4 Requesting the test result

The current test result can be requested from any `mmctestres` object. Calling `testResult` of class `mmctestres` will return a list containing indices of rejected hypotheses (vector ‘`$rejected`’), nonrejected hypotheses (vector ‘`$nonrejected`’) and undecided hypotheses (vector ‘`$undecided`’). For the previously continued run of object `m`, the result of the computation can be requested as follows:

```
> res <- testResult(m)
> res$undecided

[1] 105 113 184 255 276 279 291 331 419 434

> length(res$rejected)

[1] 104

> length(res$nonrejected)

[1] 386
```

In the example above, the current computation result of object `m` is stored in variable `res`. For object `m`, the algorithm has been run until all but (at least) 10 hypotheses have been classified. The indices of the undecided hypotheses as well as the number of rejected and nonrejected hypotheses (i.e. the length of the vectors containing their indices) are displayed.

2.1.5 Confidence intervals and estimates of p-values

At any stage, p-values can be estimated based on the total number of samples drawn for each hypothesis during initial or continued runs:

```
> estimate <- pEstimate(m)
> lastindex <- length(estimate)
> estimate[lastindex]

[1] 0.5185185
```

The function `pEstimate` takes an object of type `mmctest` as argument and returns a vector containing estimates of all p-values.

Similarly, the current confidence limits for the exact (Clopper-Pearson) confidence intervals can be requested:

```
> l <- confidenceLimits(m)
> l$lowerLimits[lastindex]

[1] 0.1077857

> l$upperLimits[lastindex]

[1] 0.9103009
```

The function `confidenceLimits` takes an object of type `mmctest` as argument and returns a list containing lower confidence limits (vector ‘`lowerLimits`’) and upper confidence limits (vector ‘`upperLimits`’) for each p-value.

2.1.6 An extended example

In this extended example a permutation test is used to determine if two groups A and B have equal means. This is done in $n_{groups} = 20$ cases. Each group has size 4 and both groups A and B are stored together in one row of length $n = 8$ in a matrix G .

```
> n <- 8
> ngroups <- 20
> G <- matrix(rep(0, n * ngroups), nrow = ngroups)
> for (j in 1:(ngroups/2)) G[j, ] <- c(rnorm(n/2, mean = 0, sd = 0.55),
+   rnorm(n/2, mean = 1, sd = 0.55))
> for (j in (ngroups/2 + 1):ngroups) G[j, ] <- rnorm(n, mean = 0,
+   sd = 3)
```

To implement this test as a Monte-Carlo test, we start by overwriting the generic class `mmctSamplerGeneric` to specify the sampler. The data stored in an `ExSampler` object is the matrix G .

```
> setClass("ExSampler", contains = "mmctSamplerGeneric", representation = representation(data = "matrix"))
[1] "ExSampler"

> setMethod("getSamples", signature(obj = "ExSampler"), function(obj,
+   ind, n) {
+   sapply(1:length(ind), function(i) {
+     v <- obj@data[ind[i], ]
+     s <- matrix(rep(v, n + 1), byrow = T, ncol = length(v))
+     for (j in 1:n) s[j + 1, ] <- sample(v)
+     means <- abs(rowMeans(s[, 1:(length(v)/2)]) - rowMeans(s[,
+       (length(v)/2 + 1):length(v)]))
+     return(sum(means > means[1]))
+   })
+ })
[1] "getSamples"

> setMethod("getNumber", signature(obj = "ExSampler"), function(obj) {
+   return(length(obj@data[, 1]))
+ })
[1] "getNumber"
```

The `getSamples` method generates n permutations of each row i in the indices vector `ind` and counts how many times the generated means exceeded the data mean (stored in row 1).

The sampler is then

```
> exsampler <- new("ExSampler", data = G)
```

As before, the multiple testing problem is set up by creating an object of type `mmctest` using `hBH` as multiple testing procedure and the `exsampler` object as sampler interface. The constructor `mmctest` uses a default threshold of 0.1.

```
> m <- mmctest(h = hBH)
> m <- run(m, exsampler, maxsteps = list(undecided = 0))
```

Algorithm `mmctest` has been run until all hypotheses were classified. Based on this run, the following hypotheses are rejected:

```
> testResult(m)$rejected
[1] 1 3 6
```

Estimates for p-values are

```
> pEstimate(m)

[1] 0.00000000 0.18085106 0.00000000 0.03339090 0.19148936 0.00000000
[7] 0.13690476 0.06159170 0.09722222 0.02623047 1.00000000 0.83333333
[13] 0.23880597 0.21276596 0.40000000 0.75000000 0.04969630 0.44444444
[19] 0.48148148 1.00000000
```

To verify this result, exact p-values are computed by enumerating all permutations of each row. This is done using algorithm `QuickPerm`.

Based on the exact p-values, given by

```
> pexact

[1] 0.00000000 0.20000000 0.00000000 0.02857143 0.08571429 0.00000000
[7] 0.11428571 0.05714286 0.08571429 0.02857143 0.91428571 0.74285714
[13] 0.40000000 0.20000000 0.34285714 0.37142857 0.05714286 0.54285714
[19] 0.60000000 0.94285714
```

the Benjamini-Hochberg procedure at threshold 0.1 will give the following set of rejections:

```
> which(hBH(pexact, threshold = 0.1))

[1] 1 3 6
```

References

- A. Gandy and G. Hahn. MMCTest - A Safe Algorithm for Implementing Multiple Monte Carlo Tests. arXiv, 2012.
- S. Pounds and C. Cheng. Robust estimation of the false discovery rate. *Bioinformatics*, 22(16):1979–1987, 2006. doi: 10.1093/bioinformatics/btl328.