

The bumphunter user's guide

Kasper Daniel Hansen `khansen@jhsph.edu`
Martin Aryee `aryee.martin@mgh.harvard.edu`
Rafael A. Irizarry `rafa@jhu.edu`

Modified: November 23, 2012. Compiled: May 3, 2016

Introduction

This package implements the statistical procedure described in [2] (with some small modifications). Notably, batch effect removal and the application of the bootstrap to linear models of Efron and Tibshirani [1] need additional code.

For any given type of data, it is usually necessary to make a number of choices and/or transformations before the bump hunting methodology is ready to be applied. Typically, these modifications resides in other packages. Examples are `charm` (for CHARM-like methylation microarrays), `bsseq` (for whole-genome bisulfite sequencing data) and `minfi` (for Illumina 450k methylation arrays). In some cases (specifically `bsseq`) only parts of the methodology as implemented in the `bumphunter` package is applied, although the conceptual approach is still build on bump hunting.

In other words, this package is mostly intended for developers wishing to adapt the general methodology to their specific applications.

The core of the package is encapsulated in the `bumphunter` method which uses the underlying `bumphunterEngine` to do the heavy lifting. However, `bumphunterEngine` consists of a number of useful functions that does much of the specific tasks involved in bump hunting. This document attempts to describe the overall workflow as well as the specific functions. The relevant functions are `clusterMaker`, `getSegments`, `findRegions`.

```
> library(bumphunter)
```

Note that this package is written with genomic data as an illustrative example but most of it is easily generalizable to other data types.

Other functions

Most of the `bumphunter` package is code for bump hunting. But we also include a number of convenience functions we have found useful, which are not necessarily part of the bump hunting exercise. At the time of writing, this include `annotateNearest`.

The Statistical Model

The bump hunter methodology is meant to work on data with several biological replicates, similar to the `lmFit` function in `limma`. While the package is written using genomic data as an illustrative example, most of it is generalizable to other data types (with some one-dimensional location information).

We assume we have data Y_{ij} where i represents (biological) replicate and l_j represents genomic location. The use of j and l_j is a convenience notation, allowing us to discuss the “next” observation $j + 1$ which may be some distance $l_{j+1} - l_j$ away. Note that we assume in this notation that all replicates have been observed at the same set of genomic locations.

The basic statistical model is the following:

$$Y_{ij} = \beta_0(l_j) + \beta_1(l_j)X_j + \varepsilon_{ij}$$

with i representing subject, l_j representing the j th location, X_j is the covariate of interest (for example $X_j = 1$ for cases and $X_j = 0$ otherwise), ε_{ij} is measurement error, $\beta_0(l)$ is a baseline function, and $\beta_1(l)$ is the parameter of interest, which is a function of location. We assume that $\beta_1(l)$ will be equal to zero over most of the genome, and we want to identify stretched where $\beta_1(l) \neq 0$, which we call *bumps*.

We want to share information between nearby locations, typically through some form of smoothing.

Creating clusters

For many genomic applications the locations are clustered. Each cluster is a distinct unit where the model fitting will be done separately, and each cluster does not depend on the data, only on the locations l_j . Typically there is some maximal distance, and we do not want to smooth between observations separated by more than this distance. The choice of distance is very application dependent.

“Clusters” are simply groups of locations such that two consecutive locations in the cluster are separated by less than some distance `maxGap`. For genomic applications, the biggest possible clusters are chromosomes.

The function `clusterMaker` defines such grouping locations.

Example: We first generate an example of typical genomic locations

```
> pos <- list(pos1=seq(1,1000,35),
+            pos2=seq(2001,3000,35),
+            pos3=seq(1,1000,50))
> chr <- rep(paste0("chr",c(1,1,2)), times = sapply(pos,length))
> pos <- unlist(pos, use.names=FALSE)
```

Now we run the function to obtain the three clusters from the positions. We use the default gap of 300 base pairs (bps), i.e. any two points more than 300 bps away are put in a new cluster. Also, locations from different chromosomes are separated.

```
> cl <- clusterMaker(chr, pos, maxGap = 300)
> table(cl)
```

```
cl
 1  2  3
29 29 20
```

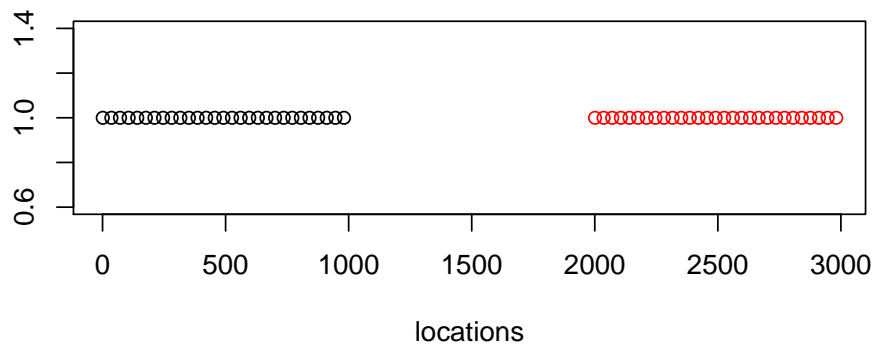
The output is an indexing variable telling us which cluster each location belongs to. Locations on different chromosomes are always on different clusters.

Note that data from the first chromosome has been split into two clusters:

```

> ind <- which(chr=="chr1")
> plot(pos[ind], rep(1,length(ind)), col=c1[ind],
+       xlab="locations", ylab="")

```



Breaking into segments

The function `getSegments` is used to find segments that are positive, near zero, and negative. Specifically we have a vector of numbers θ_j with each number associated with a genomic location l_j (think either test statistics or estimates of $\beta_i(l)$). A segment is a list of consecutive locations such that all θ_i in the segment are either “positive”, “near zero” or “negative”. In order to define “positive” etc we need a `cutoff` which is one number L (in which case “near zero” is $[-L, L]$) or two numbers L, U (in which case “near zero” is $[L; U]$).

Example: we are going to create a simulated $\beta_1(l)$ with a couple of real bumps.

```

> Indexes <- split(seq_along(c1), c1)
> beta1 <- rep(0, length(pos))
> for(i in seq(along=Indexes)){
+   ind <- Indexes[[i]]
+   x <- pos[ind]
+   z <- scale(x, median(x), max(x)/12)
+   beta1[ind] <- i*(-1)^(i+1)*pmax(1-abs(z)^3,0)^3 ##multiply by i to vary size
+ }


```

We now find bumps of this functions by

```
> segs <- getSegments(beta1, cl, cutoff=0.05)
```

Now we can make, for example, a plot of all the positive bumps

```
> par(mfrow=c(1,2))
> for(ind in segs$upIndex){
+   index <- which(cl==cl[ind[1]])
+   plot(pos[index], beta1[index],
+        xlab=paste("position on", chr[ind[1]]),
+        ylab="beta1")
+   points(pos[ind], beta1[ind], pch=16, col=2)
+   abline(h = 0.05, col = "blue")
+ }
```



This function is used by regionFinder which is described next.

regionFinder

This function packages up the results of `getSegments` into a table of regions with the location and characteristics of bumps.

```
> tab <- regionFinder(beta1, chr, pos, cl, cutoff=0.05)
> tab
```

	chr	start	end	value	area	cluster	indexStart	indexEnd	L
3	chr1	2281	2701	-1.2636037	16.426848	2	38	50	13
2	chr2	451	501	2.7262463	5.452493	3	68	69	2
1	chr1	421	561	0.5336474	2.668237	1	13	17	5

	clusterL
3	29
2	20
1	29

In the plot in the preceding section we show two of these regions in red.

Note that `regionFinder` and `getSegments` do not really contain any statistical model. All it does is finding regions based on segmenting a vector θ_j associated with genomic locations l_j .

Bumphunting

`Bumphunter` is a more complicated function. In addition to `regionFinder` and `clusterMaker` it also implements a statistical model as well as permutation testing to assess uncertainty.

We start by creating a simulated data set of 10 cases and 10 controls (recall that `beta1` was defined above).

```
> beta0 <- 3*sin(2*pi*pos/720)
> X <- cbind(rep(1,20), rep(c(0,1), each=10))
> error <- matrix(rnorm(20*length(beta1), 0, 1), ncol=20)
> y <- t(X[,1])%x%beta0 + t(X[,2])%x%beta1 + error
```

Now we can run `bumphunter`

```
> tab <- bumphunter(y, X, chr, pos, cl, cutoff=.5)
> tab
```

```
$table
  chr start end value area cluster indexStart indexEnd
10 chr1 2351 2631 -1.9115998 17.2043986 2 40 48
```

2	chr1	456	526	1.4412809	4.3238428	1	14	16
5	chr2	451	501	2.0750719	4.1501439	3	68	69
4	chr1	2946	2981	0.7911790	1.5823579	2	57	58
7	chr1	351	351	-0.9231620	0.9231620	1	11	11
3	chr1	2841	2841	0.9091460	0.9091460	2	54	54
1	chr1	211	211	0.6952925	0.6952925	1	7	7
9	chr1	2281	2281	-0.6697780	0.6697780	2	38	38
8	chr1	2176	2176	-0.6586276	0.6586276	2	35	35
6	chr1	281	281	-0.5429832	0.5429832	1	9	9
11	chr2	101	101	-0.5094585	0.5094585	3	61	61

L clusterL

10	9	29
2	3	29
5	2	20
4	2	29
7	1	29
3	1	29
1	1	29
9	1	29
8	1	29
6	1	29
11	1	20

\$coef

[,1]

[1,]	-0.104780237
[2,]	-0.421982838
[3,]	-0.301904454
[4,]	0.257161047
[5,]	-0.445740544
[6,]	-0.138566272
[7,]	0.695292534
[8,]	0.422951314
[9,]	-0.542983160
[10,]	-0.361848957
[11,]	-0.923161979
[12,]	0.230199581
[13,]	0.048739101
[14,]	1.652696228
[15,]	1.483793624
[16,]	1.187352907
[17,]	-0.270008848

[18,] -0.325464769
[19,] 0.246049397
[20,] 0.366852223
[21,] 0.369002177
[22,] -0.377282325
[23,] -0.198911405
[24,] 0.118387079
[25,] 0.484813231
[26,] -0.194718809
[27,] 0.348990235
[28,] 0.201490951
[29,] -0.336790164
[30,] 0.163567052
[31,] 0.266516635
[32,] 0.178178045
[33,] -0.085101351
[34,] -0.015223280
[35,] -0.658627558
[36,] -0.061251171
[37,] -0.086485897
[38,] -0.669777995
[39,] -0.022824023
[40,] -1.691002575
[41,] -1.731240485
[42,] -3.053800456
[43,] -2.580320720
[44,] -2.468277011
[45,] -1.267287270
[46,] -2.069714039
[47,] -1.090653844
[48,] -1.252102237
[49,] -0.088814501
[50,] -0.031396976
[51,] -0.272288094
[52,] 0.055377959
[53,] 0.050836594
[54,] 0.909145974
[55,] -0.192853857
[56,] 0.049853857
[57,] 0.677834425
[58,] 0.904523521
[59,] 0.273961773

[60,] -0.155861926
[61,] -0.509458524
[62,] -0.027513322
[63,] -0.130221814
[64,] -0.042977789
[65,] -0.091295307
[66,] -0.191852474
[67,] -0.108487076
[68,] 2.300456951
[69,] 1.849686927
[70,] 0.021340434
[71,] 0.227408454
[72,] 0.279847347
[73,] 0.215577020
[74,] 0.381003926
[75,] 0.179245110
[76,] -0.009056165
[77,] 0.171634261
[78,] -0.055192557

\$fitted

[,1]
[1,] -0.104780237
[2,] -0.421982838
[3,] -0.301904454
[4,] 0.257161047
[5,] -0.445740544
[6,] -0.138566272
[7,] 0.695292534
[8,] 0.422951314
[9,] -0.542983160
[10,] -0.361848957
[11,] -0.923161979
[12,] 0.230199581
[13,] 0.048739101
[14,] 1.652696228
[15,] 1.483793624
[16,] 1.187352907
[17,] -0.270008848
[18,] -0.325464769
[19,] 0.246049397
[20,] 0.366852223

[21,] 0.369002177
[22,] -0.377282325
[23,] -0.198911405
[24,] 0.118387079
[25,] 0.484813231
[26,] -0.194718809
[27,] 0.348990235
[28,] 0.201490951
[29,] -0.336790164
[30,] 0.163567052
[31,] 0.266516635
[32,] 0.178178045
[33,] -0.085101351
[34,] -0.015223280
[35,] -0.658627558
[36,] -0.061251171
[37,] -0.086485897
[38,] -0.669777995
[39,] -0.022824023
[40,] -1.691002575
[41,] -1.731240485
[42,] -3.053800456
[43,] -2.580320720
[44,] -2.468277011
[45,] -1.267287270
[46,] -2.069714039
[47,] -1.090653844
[48,] -1.252102237
[49,] -0.088814501
[50,] -0.031396976
[51,] -0.272288094
[52,] 0.055377959
[53,] 0.050836594
[54,] 0.909145974
[55,] -0.192853857
[56,] 0.049853857
[57,] 0.677834425
[58,] 0.904523521
[59,] 0.273961773
[60,] -0.155861926
[61,] -0.509458524
[62,] -0.027513322

```

[63,] -0.130221814
[64,] -0.042977789
[65,] -0.091295307
[66,] -0.191852474
[67,] -0.108487076
[68,]  2.300456951
[69,]  1.849686927
[70,]  0.021340434
[71,]  0.227408454
[72,]  0.279847347
[73,]  0.215577020
[74,]  0.381003926
[75,]  0.179245110
[76,] -0.009056165
[77,]  0.171634261
[78,] -0.055192557

```

```

$pvaluesMarginal
[1] NA

```

```

> names(tab)

```

```

[1] "table"          "coef"           "fitted"
[4] "pvaluesMarginal"

```

```

> tab$table

```

	chr	start	end	value	area	cluster	indexStart	indexEnd
10	chr1	2351	2631	-1.9115998	17.2043986	2	40	48
2	chr1	456	526	1.4412809	4.3238428	1	14	16
5	chr2	451	501	2.0750719	4.1501439	3	68	69
4	chr1	2946	2981	0.7911790	1.5823579	2	57	58
7	chr1	351	351	-0.9231620	0.9231620	1	11	11
3	chr1	2841	2841	0.9091460	0.9091460	2	54	54
1	chr1	211	211	0.6952925	0.6952925	1	7	7
9	chr1	2281	2281	-0.6697780	0.6697780	2	38	38
8	chr1	2176	2176	-0.6586276	0.6586276	2	35	35
6	chr1	281	281	-0.5429832	0.5429832	1	9	9
11	chr2	101	101	-0.5094585	0.5094585	3	61	61

	L	cluster	L
10	9		29
2	3		29
5	2		20
4	2		29
7	1		29
3	1		29
1	1		29
9	1		29
8	1		29
6	1		29
11	1		20

Briefly, the `bumphunter` function fits a linear model for each location (like `lmFit` from the `limma` package), focusing on one specific column (coefficient) of the design matrix. This coefficient of interest is optionally smoothed. Subsequently, a permutation can be used to test is formed for this specific coefficient.

The simplest way to use permutations to create a null distribution is to set `B`. If the number of samples is large this can be set to a large number, such as 1000. Note that this will be slow and we have therefore provided parallelization capabilities. In cases were the user wants to define the permutations, for example cases in which all possible permutations can be enumerated, these can be supplied via the `permutation` argument.

Note that the function permits the matrix `X` to have more than two columns. This can be useful for those wanting to fit models that try to adjust for confounders, such as age and sex. However, when `X` has columns other than those representing an intercept term and the covariate of interest, the permutation test approach is not recommended. The function will run but give a warning. A method based on the bootstrap for linear models of Efron and Tibshirani [1] may be more appropriate but this is not currently implemented.

Faster bumphunting with multiple cores

`bumphunter` can be speeded up by using multiple cores. We use the `foreach` package which allows different parallel "back-ends" that will distribute the computation across multiple cores in a single machine, or across multiple machines in a cluster. The most straightforward usage, illustrated below, involves multiple cores on a single machine. See the `foreach` documentation for more complex use cases, as well as the packages `doParallel` and `doSNOW` (among others). Finally, we use `doRNG` to ensure reproducibility of setting the seed within the parallel computations.

In order to use the `foreach` package we need to register a backend, in this case a multicore machine with 2 cores.

```
> library(doParallel)
> registerDoParallel(cores = 2)
```

`bumphunter` will now automatically use this backend

```
> tab <- bumphunter(y, X, chr, pos, cl, cutoff=.5, B=250, verbose = TRUE)
> tab
```

a 'bumps' object with 11 bumps

References

- [1] B. Efron and R.J. Tibshirani.
- [2] Andrew E Jaffe, Peter Murakami, Hwajin Lee, Jeffrey T Leek, M Daniele Fallin, Andrew P Feinberg, and Rafael A Irizarry. Bump hunting to identify differentially methylated regions in epigenetic epidemiology studies. *International Journal of Epidemiology*, 41(1):200–209, 2012.

Cleanup

This is a cleanup step for the vignette on Windows; typically not needed for users.

```
> bumphunter:::foreachCleanup()
```

SessionInfo

- R version 3.3.0 (2016-05-03), x86_64-pc-linux-gnu

- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_US.UTF-8, LC_COLLATE=C, LC_MONETARY=en_US.UTF-8, LC_MESSAGES=en_US.UTF-8, LC_PAPER=en_US.UTF-8, LC_NAME=C, LC_ADDRESS=C, LC_TELEPHONE=C, LC_MEASUREMENT=en_US.UTF-8, LC_IDENTIFICATION=C
- Base packages: base, datasets, grDevices, graphics, methods, parallel, stats, stats4, utils
- Other packages: BiocGenerics 0.18.0, GenomeInfoDb 1.8.0, GenomicRanges 1.24.0, IRanges 2.6.0, S4Vectors 0.10.0, bumhunter 1.12.0, doParallel 1.0.10, doRNG 1.6, foreach 1.4.3, iterators 1.0.8, locfit 1.5-9.1, pkgmaker 0.22, registry 0.3, rngtools 1.2.4
- Loaded via a namespace (and not attached): AnnotationDbi 1.34.0, Biobase 2.32.0, BiocParallel 1.6.0, Biostrings 2.40.0, DBI 0.4, GenomicAlignments 1.8.0, GenomicFeatures 1.24.0, RCurl 1.95-4.8, RSQLite 1.0.0, Rsamtools 1.24.0, SummarizedExperiment 1.2.0, XML 3.98-1.4, XVector 0.12.0, biomaRt 2.28.0, bitops 1.0-6, codetools 0.2-14, compiler 3.3.0, digest 0.6.9, grid 3.3.0, lattice 0.20-33, magrittr 1.5, matrixStats 0.50.2, rtracklayer 1.32.0, stringi 1.0-1, stringr 1.0.0, tools 3.3.0, xtable 1.8-2, zlibbioc 1.18.0