

FunChIP: A Functional Data Analysis approach to cluster ChIP-Seq peaks according to their shapes

Alice Parodi alicecarla.parodi@polimi.it

Laura M. Sangalli
Piercesare Secchi
Simone Vantini
Marco J. Morelli

May 4, 2023

Contents

```
> library(FunChIP)
```

1 Introduction

The *FunChIP* package provides a set of methods for the *GRanges* class of the package *GenomicRanges* to cluster ChIP-Seq peaks according to their shapes, starting from a BAM file containing the aligned reads and a *GRanges* object with the corresponding enriched regions.

2 Input and Preprocessing

ChIP-Seq enriched regions are provided by the user in a *GRanges* object `GR`. The user must provide the BAM file containing the reads aligned on the positive and negative strands of the DNA. From the BAM file we can compute, for each region of the *GRanges* (let N be the total number of regions), the base-level coverage separately for positive and negative reads. These two count vectors are used to estimate the distance d_{pn} between positive and negative reads and then the total length of the fragments of the ChIP-Seq experiment d . In particular, we assume that the positive and negative counts measure the same signal, shifted by d_{pn} , as they are computed from the two ends of the sequencing fragments. The global length of the fragment is the sum between the length of the reads of the BAM file, r^1 , and the distance between the positive and negative coverage d_{pn}

$$d = d_{pn} + r.$$

The function `compute_fragments_length` computes, from the *GRanges* object and the BAM file, the estimated length of the fragments. Given a range for d_{pn} : $[d_{\min}; d_{\max}]$, the optimum distance d_{pn} is

$$d_{pn} = \operatorname{argmin}_{\delta \in [d_{\min}; d_{\max}]} \sum_{n=1}^N D(f_{n+}, f_{n-}^{\delta}),$$

¹If in the BAM file multiple length are present, r is estimated as the average length.

where f_{n+} is the positive coverage function of the n -th region, and f_{n-}^{δ} is the negative coverage of the n th region, shifted by δ . The distance D is the square of the L^2 distance between the coverages, normalized by the width of the region. The definition of the L^2 distance is detailed in Section ??.

```
> # load the GRanges object associated
> # to the ChIP-Seq experiment on the
> # transcription factor c-Myc in murine cells
>
> data(GR100)
> # name of the .bam file (the
> # .bam.bai index file must also be present)
>
> bamf <- system.file("extdata", "test.bam",
+                      package="FunChIP", mustWork=TRUE)
> # compute d
>
> d <- compute_fragments_length(GR, bamf, min.d = 0, max.d = 300)

[1] "estimated distance positive - negative read 148"
[1] "estimated read length 51"

> d

[1] 199

>
```

In Figure ?? the distance function is shown varying the parameter δ , and the minimum value d_{pn} is computed.




Figure 1: Identification of d . optimal value of d_{pn} is presented. It is the minimum of the global distance function.

Once we have correctly identified the fragment length we can compute the final coverage function to obtain the shape of the peaks. The `pileup_peak` method for the `GRanges` class uses the BAM file to compute the base-level coverage on these regions, once the reads are extended up to their final length d . `pileup_peak` adds to the `GRanges` a `counts` metadata column, containing for each region a vector with length equal to the width of the region storing the coverage function.

```
> # associate to each peak
> # of the GRanges object the correspondent
> # coverage function
>
> peaks <- pileup_peak(GR, bamf, d = d)
> peaks

GRanges object with 100 ranges and 1 metadata column:
      seqnames      ranges strand |      counts
      <Rle>        <IRanges> <Rle> |      <list>
[1]   chr18  3337524-3338025     * |    7,8,8,...
[2]   chr18  4369126-4369352     * |    7,9,9,...
[3]   chr18  4375448-4375883     * |    8,8,8,...
[4]   chr18  4715744-4716162     * |    5,5,5,...
[5]   chr18  4716374-4716597     * |   15,15,15,...
...     ...                ...   ...   ...
[96]  chr18 35112325-35112593     * |   12,12,12,...
[97]  chr18 35113538-35114826     * |    9,8,8,...
[98]  chr18 35118063-35118304     * |   10,10,10,...
[99]  chr18 35182164-35182425     * |    8,8,8,...
[100] chr18 35205390-35205649     * |   14,14,14,...
-----
seqinfo: 20 sequences from an unspecified genome; no seqlengths
>
```

Additional information can be found in the help page of the `pileup_peak` method.

3 Smoothing

The `counts` metadata is approximated by a combination of splines to guarantee the smoothness and regularity needed for further analysis, as described in the following Sections.

The preprocessing steps carried out in the `smooth_peak` method are the following:

- *Removal of the background and extension.* In ChIP-Seq experiments, peaks may have an additive noisy background, and the removal of this background is mandatory to compare different peaks. The background is estimated as a constant value "raising" the peak and equal to the minimum value the coverage assumes. Consequently, once the background has been removed, each peak has zero as minimum value, thus allowing the peak to be indefinitely extended with zeros, if necessary. In Section ??, how this choice affects the algorithm will be discussed.
- *Smoothing.* In order to be regular enough to computed derivatives, a peak has to be transformed in a suitable functional object, as described in Section ??. The smoothing of the count vector c is performed through the projection of c on a cubic B-spline basis $\Phi = \{\phi_1, \dots, \phi_K\}$ with a penalization on the second derivative [?]. The result is a spline approximation of the data, which is continuous on the whole domain, together with its first order derivatives. Moreover, the penalization on the second derivative allows to control the global regularity of the function avoiding over-fitting and a consequent noisy spline definition. The spline approximation $s = \sum_{k=1}^K \theta_k \phi_k$ of the count vector

$c = \{c_j\}$ is defined minimizing

$$S(\lambda) = \sum_{j=1}^n [c_j - s(x_j)]^2 + \lambda \int [s''(x)]^2 dx,$$

with x_j being the relative genomic coordinate the counts. The multiplying coefficient λ quantifies the penalization on the second derivative and is chosen through the Generalized Cross Validation criteria. For each peak i the GCV_i index is computed with a leave-one-out cross validation

$$GCV_i = \left(\frac{n}{n - df(\lambda)} \right) \left(\frac{SSE_i}{n - df(\lambda)} \right)$$

and then it is summed on the whole data set to obtain the global GCV . The number of degrees of freedom $df(\lambda)$ is automatically computed from the definition of the basis Φ .

The error SSE_i can be computed either on the data (SSE_i^0) or on the derivatives (SSE_i^1), to control the regularity of the function or the regularity of the derivatives, respectively:

$$SSE_i^0 = \sqrt{\sum_{j=1}^n (c_j - s(x_j))^2} \text{ or } SSE_i^1 = \sqrt{\sum_{j=1}^{n-1} (\nabla c_j - s'(x_j))^2},$$

with ∇c_j being the finite-difference approximation of the derivative of the `counts` vector c for the data i : $c = c(i)$, while $s'(x_i)$ is the evaluation of the first derivative $s' = s'(i)$ on the genomic coordinates. For further details on the spline definition see the `spline` function of the `fda` package.

- *Scaling of the peaks.* This optional preprocessing step makes all the curves having the same width and area. In particular all the abscissa grid are scaled to become equal to the smallest grid throughout the data, while y -values are scaled to make areas of all the curves equal to 1.

The `smooth_peak` method approximates the `counts` metadata by removing the background, computing the spline and potentially defining the scaled approximation. Focusing on the spline approximation, `smooth_peak` automatically chooses the optimal λ parameter according to the GCV criteria; the user can decide whether to consider the data or the derivatives to compute the SSE .

```
> # the method smooth_peak
> # removes the background and defines the spline
> # approximation from the previously computed peaks
> # with lambda estimated from the
> # GCV on derivatives. The method spans a non-uniform
> # grid for lambda from 10^-4 to 10^12.
> # ( the grid is uniform for log10(lambda) )
>
> peaks.smooth <- smooth_peak(peaks, lambda = 10^(-4:12),
+                             subsample.data = 50, GCV.derivatives = TRUE,
+                             plot.GCV = TRUE, rescale = FALSE )

[1] "iteration on lambda: 1"
[1] "iteration on lambda: 2"
```

```
[1] "iteration on lambda: 3"
[1] "iteration on lambda: 4"
[1] "iteration on lambda: 5"
[1] "iteration on lambda: 6"
[1] "iteration on lambda: 7"
[1] "iteration on lambda: 8"
[1] "iteration on lambda: 9"
[1] "iteration on lambda: 10"
[1] "iteration on lambda: 11"
[1] "iteration on lambda: 12"
[1] "iteration on lambda: 13"
[1] "iteration on lambda: 14"
[1] "iteration on lambda: 15"
[1] "iteration on lambda: 16"
[1] "iteration on lambda: 17"

>
```




Figure 2: Generalized Cross Validation index. GCV computed on data (left), and on the derivatives (right), as a function of λ .

In Figure ??, the plot of the GCV for both data and derivatives is shown. From this Figure we see that the optimum value of λ , which minimizes the GCV for the derivatives, is also associated to a small value of the GCV for the data thus supporting the automatic choice.

```
> # the automatic choice is lambda = 10^6
>
> peaks.smooth <- smooth_peak(peaks, lambda = 10^6,
+                             plot.GCV = FALSE)
> head(peaks.smooth)
```

GRanges object with 6 ranges and 6 metadata columns:

	seqnames	ranges	strand	counts	spline
	<Rle>	<IRanges>	<Rle>	<list>	<list>
[1]	chr18	3337524-3338025	*	7,8,8,... 0.107577,0.115944,0.124571,...	
[2]	chr18	4369126-4369352	*	7,9,9,... 0.106531,0.125616,0.145441,...	
[3]	chr18	4375448-4375883	*	8,8,8,... 0.103757,0.127649,0.152502,...	
[4]	chr18	4715744-4716162	*	5,5,5,... 0.108025,0.121182,0.134814,...	
[5]	chr18	4716374-4716597	*	15,15,15,... 0.109476,0.126596,0.144360,...	
[6]	chr18	4921270-4921506	*	8,8,8,... 0.104672,0.120304,0.136527,...	

```

      spline_der width_spline start_spline end_spline
      <list>      <integer>      <numeric> <numeric>
[1] 0.00823719,0.00849664,0.00875919,...      593      3337483      3338075
[2] 0.0187191,0.0194523,0.0201997,...      335      4369075      4369409
[3] 0.0234173,0.0243690,0.0253395,...      508      4375417      4375924
[4] 0.0129218,0.0133933,0.0138718,...      523      4715698      4716220
[5] 0.0168009,0.0174399,0.0180908,...      315      4716303      4716617
[6] 0.0153411,0.0159259,0.0165217,...      332      4921234      4921565
-----
seqinfo: 20 sequences from an unspecified genome; no seqlengths

> # maintaining this choice of lambda smooth_peak
> # can also define the scaled approximation
> # of the spline
>
> peaks.smooth.scaled <- smooth_peak(peaks, lambda = 10^6,
+                                   plot.GCV = FALSE, rescale = TRUE)
> head(peaks.smooth.scaled)

GRanges object with 6 ranges and 8 metadata columns:
      seqnames      ranges strand |      counts      spline
      <Rle>      <IRanges> <Rle> |      <list>      <list>
[1] chr18 3337524-3338025      * | 7,8,8,... 0.107577,0.115944,0.124571,...
[2] chr18 4369126-4369352      * | 7,9,9,... 0.106531,0.125616,0.145441,...
[3] chr18 4375448-4375883      * | 8,8,8,... 0.103757,0.127649,0.152502,...
[4] chr18 4715744-4716162      * | 5,5,5,... 0.108025,0.121182,0.134814,...
[5] chr18 4716374-4716597      * | 15,15,15,... 0.109476,0.126596,0.144360,...
[6] chr18 4921270-4921506      * | 8,8,8,... 0.104672,0.120304,0.136527,...
      spline_der width_spline start_spline end_spline
      <list>      <integer>      <numeric> <numeric>
[1] 0.00823719,0.00849664,0.00875919,...      593      3337483      3338075
[2] 0.0187191,0.0194523,0.0201997,...      335      4369075      4369409
[3] 0.0234173,0.0243690,0.0253395,...      508      4375417      4375924
[4] 0.0129218,0.0133933,0.0138718,...      523      4715698      4716220
[5] 0.0168009,0.0174399,0.0180908,...      315      4716303      4716617
[6] 0.0153411,0.0159259,0.0165217,...      332      4921234      4921565
      spline_rescaled      spline_der_rescaled
      <list>      <list>
[1] 0.0026361:0.00304401:0.00347677:... 0.000395668:0.000420244:0.000445387:...
[2] 0.00122834:0.00147221:0.00172654:... 0.000238709:0.000249059:0.000259633:...
[3] 0.000864564:0.00120333:0.00156496:... 0.000327572:0.000350081:0.000373281:...
[4] 0.000788311:0.000956446:0.00113505:... 0.000162986:0.000173327:0.000183923:...
[5] 0.00235269:0.0027355:0.0031333:... 0.000375411:0.000390254:0.000405393:...
[6] 0.00139356:0.00162209:0.00186007:... 0.000223859:0.00023322:0.000242774:...
-----
seqinfo: 20 sequences from an unspecified genome; no seqlengths
>

```

Now the *GRanges* object contains, besides *counts*, 5 new metadata columns with the spline approximation evaluated on the base-level grid, its derivatives, the width of the spline and the new starting and ending points (see Figure ??). For a more detailed description of the metadata columns, see the help page of the *smooth_peak* method.

With the introduction of the smoothing, counts at the edges of the peak are connected with regularity to 0, and therefore new values different from zeros may be introduced. In order to maintain regularity, the grid is extended up to the new boundaries.

Adding to `smooth_peak` the option `rescale = TRUE` the method, beside the 5 metadata columns previously introduced, returns 2 more metadata columns with the scaled approximation of the spline and its derivatives.

Once the spline approximation is defined, the summit of the smoothed peak (or even of the scaled peak), i.e. of its spline approximation, can be detected. The summit will be used to initialize the peak alignment procedure, described in Section ??, and it can either be a user-defined parameter, stored in a vector of the same length of the `GR`, or automatically computed as the maximum height of the spline. The summit is stored in the new metadata column `summit_spline`. If the `rescale` option is set to `TRUE` the summit of the scaled approximation is also returned in the metadata column `summit_spline_rescaled`.

```
> # peaks.summit identifies the maximum point
> # of the smoothed peaks
>
> peaks.summit <- summit_peak(peaks.smooth)
> head(peaks.summit)
```

GRanges object with 6 ranges and 7 metadata columns:

	seqnames	ranges	strand	counts	spline
	<Rle>	<IRanges>	<Rle>	<list>	<list>
[1]	chr18	3337524-3338025	*	7,8,8,... 0.107577,0.115944,0.124571,...	
[2]	chr18	4369126-4369352	*	7,9,9,... 0.106531,0.125616,0.145441,...	
[3]	chr18	4375448-4375883	*	8,8,8,... 0.103757,0.127649,0.152502,...	
[4]	chr18	4715744-4716162	*	5,5,5,... 0.108025,0.121182,0.134814,...	
[5]	chr18	4716374-4716597	*	15,15,15,... 0.109476,0.126596,0.144360,...	
[6]	chr18	4921270-4921506	*	8,8,8,... 0.104672,0.120304,0.136527,...	

	spline_der	width_spline	start_spline	end_spline
	<list>	<integer>	<numeric>	<numeric>
[1]	0.00823719,0.00849664,0.00875919,...	593	3337483	3338075
[2]	0.0187191,0.0194523,0.0201997,...	335	4369075	4369409
[3]	0.0234173,0.0243690,0.0253395,...	508	4375417	4375924
[4]	0.0129218,0.0133933,0.0138718,...	523	4715698	4716220
[5]	0.0168009,0.0174399,0.0180908,...	315	4716303	4716617
[6]	0.0153411,0.0159259,0.0165217,...	332	4921234	4921565

	summit_spline
	<integer>
[1]	444
[2]	186
[3]	174
[4]	310
[5]	121
[6]	169

seqinfo: 20 sequences from an unspecified genome; no seqlengths

```
> # peaks.summit can identify also the maximum
> # point of the scaled approximation
>
> peaks.summit.scaled <- summit_peak(peaks.smooth.scaled,
```

```

+                               rescale = TRUE)
> head(peaks.summit.scaled)

GRanges object with 6 ranges and 10 metadata columns:
      seqnames      ranges strand |      counts      spline
      <Rle>      <IRanges> <Rle> |      <list>      <list>
[1] chr18 3337524-3338025      * | 7,8,8,... 0.107577,0.115944,0.124571,...
[2] chr18 4369126-4369352      * | 7,9,9,... 0.106531,0.125616,0.145441,...
[3] chr18 4375448-4375883      * | 8,8,8,... 0.103757,0.127649,0.152502,...
[4] chr18 4715744-4716162      * | 5,5,5,... 0.108025,0.121182,0.134814,...
[5] chr18 4716374-4716597      * | 15,15,15,... 0.109476,0.126596,0.144360,...
[6] chr18 4921270-4921506      * | 8,8,8,... 0.104672,0.120304,0.136527,...
      spline_der width_spline start_spline end_spline
      <list>      <integer>      <numeric> <numeric>
[1] 0.00823719,0.00849664,0.00875919,...      593      3337483      3338075
[2] 0.0187191,0.0194523,0.0201997,...      335      4369075      4369409
[3] 0.0234173,0.0243690,0.0253395,...      508      4375417      4375924
[4] 0.0129218,0.0133933,0.0138718,...      523      4715698      4716220
[5] 0.0168009,0.0174399,0.0180908,...      315      4716303      4716617
[6] 0.0153411,0.0159259,0.0165217,...      332      4921234      4921565
      spline_rescaled spline_der_rescaled
      <list>      <list>
[1] 0.0026361:0.00304401:0.00347677:... 0.000395668:0.000420244:0.000445387:...
[2] 0.00122834:0.00147221:0.00172654:... 0.000238709:0.000249059:0.000259633:...
[3] 0.000864564:0.00120333:0.00156496:... 0.000327572:0.000350081:0.000373281:...
[4] 0.000788311:0.000956446:0.00113505:... 0.000162986:0.000173327:0.000183923:...
[5] 0.00235269:0.0027355:0.0031333:... 0.000375411:0.000390254:0.000405393:...
[6] 0.00139356:0.00162209:0.00186007:... 0.000223859:0.00023322:0.000242774:...
      summit_spline_rescaled summit_spline
      <integer>      <integer>
[1]      227      444
[2]      168      186
[3]      104      174
[4]      180      310
[5]      117      121
[6]      155      169
-----
seqinfo: 20 sequences from an unspecified genome; no seqlengths

```

4 The k-mean alignment algorithm and the `cluster_peak` method

The k-mean alignment algorithm is an efficient method to classify functional data allowing for general transformation of abscissae [?]; this general method is implemented in the package `fdakma` and various applications to real dataset are introduced in [?], [?], [?].

In particular, given

- a set of curves s_1, \dots, s_n ,
- the number of clusters K ,

- a distance function $d(s_i, s_j)$ between two curves s_i and s_j , as for example the integral of the difference $s_i - s_j$,
- a family of warping functions \mathcal{W} to transform the abscissae of the curves and therefore align the peaks. Generally, \mathcal{W} is the set of shifts or dilations or affine transformations (shift + dilation),

the algorithm, presented in Algorithm ??, is an iterative procedure to split the curves into K clusters. The introduction of the warping function $h \in \mathcal{W}$ allows each curve to be shifted, dilated, or both, to define the minimum distance between curves. The new curve $s \circ h$ has the same values of s , but its abscissa grid is modified. For example, in Figure ?? two peaks




Figure 3: Alignment procedure. Representation of two smoothed peaks. In the left panel they are not aligned, while in the right panel they are aligned with an integer shift.

are presented: in the left panel, they are not aligned, while the right panel shows the effects of alignment; the transformation of the abscissae (shift transformation) makes the two peaks more similar, and the distance d is not anymore affected by artificial phase distance. The code generating Figure ?? calls `cluster_peak` and `plot_peak`, which are described in Section ?? and Section ??.

```
> # representation of two peaks
>
> par (mfrow = c(1,2))
> plot_peak(peaks.summit, index = c(6,7), col=c('red',2), cex.main = 2, cex.lab = 2, cex.axis = 1.5, lwd = 2)
> aligned.peaks <- cluster_peak(peaks.summit[c(6,7)], parallel = FALSE,
+                               n.clust = 1, seeds = 1, shift.peak = TRUE,
+                               weight = 1, alpha = 1, p = 2, t.max = 2,
+                               plot.graph.k = FALSE, verbose = FALSE)
> aligned.peaks
```

GRanges object with 2 ranges and 10 metadata columns:

	seqnames	ranges	strand	counts	spline
	<Rle>	<IRanges>	<Rle>	<list>	<list>
[1]	chr18	4921270-4921506	*	8,8,8,...	0.104672,0.120304,0.136527,...
[2]	chr18	5078473-5078803	*	11,11,11,...	0.100641,0.110866,0.121442,...

	spline_der	width_spline	start_spline	end_spline
	<list>	<integer>	<numeric>	<numeric>
[1]	0.0153411,0.0159259,0.0165217,...	332	4921234	4921565
[2]	0.0100516,0.0103994,0.0107528,...	420	5078418	5078837

	summit_spline	cluster_shift	coef_shift	dist_shift
	<integer>	<list>	<list>	<list>
[1]				
[2]				

```
[1]      169      1     -15  0.447149
[2]      283      1      16      0
-----
seqinfo: 20 sequences from an unspecified genome; no seqlengths

> # shift coefficients
> aligned.peaks$coef_shift

[[1]]
[1] -15

[[2]]
[1] 16

> plot_peak(aligned.peaks, col = 'forestgreen',
+           shift = TRUE, k = 1, cluster.peak = TRUE,
+           line.plot = 'spline',
+           cex.main = 2, cex.lab = 2, cex.axis = 1.5, lwd = 2)
```

For the specific case of ChIP-Seq data, the admitted warping functions for the k-mean alignment algorithm (in the `cluster_peak` method), are integer shifts:

$$\mathcal{W} = \{h : h(t) = t + q \text{ with } q \in \mathbb{Z}\}. \quad \mathbf{1}$$

In other words, with this choice, peaks can be shifted by integer values in the *alignment* procedure of the algorithm.

Algorithm 1: k-mean alignment algorithm

Given a set of functions s_1, \dots, s_n and a number K of clusters

Template: random choice (if not provided) of the initial centers of the clusters $c_1 \dots, c_k$

while decrease of the distance higher than a fixed threshold **do**

foreach $i \in 1 : n$ **do**

Alignment: s_i is aligned to each template c_k : the optimal warping function $h_{i,k}^*$ in \mathcal{W} is detected

$$h_{i,k}^* = \operatorname{argmin}_{h \in \mathcal{W}} d(c_k, x_i \circ h)$$

 with the corresponding distance $d_{i,k}^* = \min_{h \in \mathcal{W}} d(c_k, x_i \circ h)$

Assignment: s_i is assigned to the best cluster

$$k_i^* = \operatorname{argmin}_{k \in 1:K} d_{i,k}^*$$

end

foreach $k \in 1 : K$ **do**

Template: identification of the new template of the cluster c_k

Normalization: the average warping function of the curves belonging to k is set to be the identity transformation

$$h(s) = s$$

end

end

In the `cluster_peak` method the distance between two curves s_1 and s_2 is defined as

$$\begin{aligned} d(s_1, s_2) &= (1 - \alpha) d_0(s_1, s_2) + \alpha w d_1(s_1, s_2) = \\ &= (1 - \alpha) \|s_1^e - s_2^e\|_p + \alpha w \|(s_1^e)' - (s_2^e)'\|_p, \end{aligned} \quad 2$$

where

- $\|f\|_p$ is the p norm of f . In particular, for $p = 0$, $\|\cdot\|_p$ is the L^∞ norm

$$\|f\|_0 = \|f\|_{L^\infty} = \max_{x \in U} |f(x)|,$$

with U being the domain of f .

For $p = 1$, $\|\cdot\|_p$ is the L^1 norm

$$\|f\|_1 = \|f\|_{L^1} = \int_U |f(x)| dx.$$

And for $p = 2$, $\|\cdot\|_p$ is the L^2 norm

$$\|f\|_2 = \|f\|_{L^2} = \int_U (f(x))^2 dx.$$

- s_1^e and s_2^e are the functions s_1 and s_2 extended with zeros where not defined, after their backgrounds have been removed (see Section ??). The distance function is computed on the union of the domains of s_1 and s_2 (U); s_1 and s_2 need to be extended to cover the whole U .
- $\alpha \in [0, 1]$ is a coefficient tuning the contributions of the norm of the data and the norm of the derivatives. If $\alpha = 0$, the distance is computed on the data, while if $\alpha = 1$ it is based on the derivatives. Intermediate values balance these two contributions: increasing the relevance given to the derivatives emphasizes the shapes of the peaks, while data are more related to the height.
- w is a weight coefficient, essential to make the norm of the data and of the derivatives comparable. It can be user defined or computed inside the `cluster_peak` method. A suggestion for computing the weight w is given in Section ??.

4.1 Definition of weight in the distance function

If not provided, the method `cluster_peak` defines w as

$$w = \text{median} \left(\frac{d_0(s_i, s_j)}{d_1(s_i, s_j)} \right)$$

where $d_0(i, j) = \|s_i^e - s_j^e\|_p$ and $d_1(i, j) = \|(s_i^e)' - (s_j^e)'\|_p$. These matrices can be automatically computed with the `distance_peak` function.

```
> # compute the weight from the first 10 peaks
>
> dist_matrix <- distance_peak(peaks.summit)
> # dist matrix contains the two matrices d_0(i,j)
> # and d_1(i,j), used to compute w
> names(dist_matrix)

[1] "dist_matrix_d0" "dist_matrix_d1"
```

```
> ratio_norm <- dist_matrix$dist_matrix_d0 / dist_matrix$dist_matrix_d1
> ratio_norm_upper_tri <- ratio_norm[upper.tri(ratio_norm)]
> summary(ratio_norm_upper_tri)

   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
 28.21  91.36 117.79 124.07 148.17 330.95

> # suggestion: use the median as weight
> w <- median(ratio_norm_upper_tri)
> w

[1] 117.7907
>
```

4.2 The `cluster_peak` method

The two main characteristics of the k-mean alignment algorithm used in *FunChIP* are the distance function d (defined in Equation (??)), used to compute the distance between curves, and the set of warping functions \mathcal{W} (defined in Equation (??)) considered for the alignment. The `cluster_peak` method applies the k-mean alignment algorithm with these specifications to the set of peaks stored in the *GRanges* object. In particular, the parameters `weight`², `alpha` and `p` define the distance used in the algorithm, while `t.max` sets the maximum shift of each peak in each iteration (in this particular case, q of Equation (??) does not vary in the whole \mathbb{Z} but $q \in \{-t.max \cdot |U|, \dots, +t.max \cdot |U|\}$, with $|U|$ being the maximum width of the spline approximation of the peaks).

Given a *GRanges* `GR` containing the metadata columns computed from the `smooth_peak` method, `cluster_peak` applies the k-mean alignment algorithm for all the values of k between 1 and `n.clust` (parameter of the function).

The algorithm can be run in parallel, setting to `TRUE` the `parallel` argument of the method and providing the number of cores `num.cores`. With these settings, the different applications of the algorithm, corresponding to different numbers of clusters, are executed in parallel.

As detailed in the help, the `cluster_peak` method has 2 outputs:

- The *GRanges* object, updated with new metadata columns associated to the classification. In particular, in the general case of classification with and without alignment, columns with information on the clustering of the peaks (`cluster_shift` and `cluster_NOshift`), the corresponding shifts (`coef_shift`) and the distances from the template of the clusters (`dist_shift` and `dist_NOshift`) are added.
- The graph of the global distance within clusters³ as a function of the number of clusters (if `plot.graph.k = TRUE`). This plot can be used to identify the optimal number of clusters of the partition of the data set and the effect of the alignment procedure. In particular, if `shift = NULL`, the algorithm is run both with and without alignment and two trend lines are plotted: the black line corresponds to the global distance without the shift, and the red line corresponds to the distance obtained with alignment. If `shift` is set to `TRUE` or `FALSE`, just one type of algorithm is run and the correspondent curve is plotted. For each trend line, this graph allows the identification of the optimal

²`weight` can be also set to `NULL` and it will be automatically computed as specified in Section ?? . To save computational time, it is generally computed on a random sub-sample of data, whose size is set by the `sub.sample.weight` parameter.

³sum over all the peaks of the distance of each peak from the corresponding template.

value of the number of clusters: for this value, the distance significantly decreases with respect to the lower values of k , and negligibly increases with respect to higher values of k (elbow in the line). The gap between the red and the black line, instead, shows the decrease of the distance when the shift is introduced.

It is relevant to point out that the algorithm can be run both on the original data and on the scaled peaks, depending on the focus of the analysis. The logic parameter `rescale` allows the user to choose.

```
> # classification of the smooth peaks in different
> # numbers of clusters, from 1 ( no distinction, only shift )
> # to 4.
>
> # here the analysis is run on the spline approximation
> # without scaling
> peaks.cluster <- cluster_peak(peaks.summit, parallel = FALSE, seeds=1:4,
+                               n.clust = 1:4, shift = NULL,
+                               weight = 1, alpha = 1, p = 2, t.max = 2,
+                               plot.graph.k = TRUE, verbose = FALSE)
> head(peaks.cluster)
```

GRanges object with 6 ranges and 12 metadata columns:

	seqnames <Rle>	ranges <IRanges>	strand <Rle>	counts <list>	spline <list>
[1]	chr18	3337524-3338025	*	7,8,8,... 0.107577,0.115944,0.124571,...	
[2]	chr18	4369126-4369352	*	7,9,9,... 0.106531,0.125616,0.145441,...	
[3]	chr18	4375448-4375883	*	8,8,8,... 0.103757,0.127649,0.152502,...	
[4]	chr18	4715744-4716162	*	5,5,5,... 0.108025,0.121182,0.134814,...	
[5]	chr18	4716374-4716597	*	15,15,15,... 0.109476,0.126596,0.144360,...	
[6]	chr18	4921270-4921506	*	8,8,8,... 0.104672,0.120304,0.136527,...	

	spline_der <list>	width_spline <integer>	start_spline <numeric>	end_spline <numeric>
[1]	0.00823719,0.00849664,0.00875919,...	593	3337483	3338075
[2]	0.0187191,0.0194523,0.0201997,...	335	4369075	4369409
[3]	0.0234173,0.0243690,0.0253395,...	508	4375417	4375924
[4]	0.0129218,0.0133933,0.0138718,...	523	4715698	4716220
[5]	0.0168009,0.0174399,0.0180908,...	315	4716303	4716617
[6]	0.0153411,0.0159259,0.0165217,...	332	4921234	4921565

	summit_spline <integer>	cluster_NOshift <list>	dist_NOshift <list>	cluster_shift <list>
[1]	444	1,1,1,... 0.624908,0.225284,0.225284,...		1,1,1,...
[2]	186	1,2,2,... 0.436710,0.517292,0.432370,...		1,2,2,...
[3]	174	1,2,2,... 0.378181,0.485136,0.382850,...		1,2,2,...
[4]	310	1,2,2,... 0.539505,0.554752,0.578859,...		1,2,2,...
[5]	121	1,1,1,... 0.534987,0.217738,0.217738,...		1,1,1,...
[6]	169	1,1,2,... 0.293730,0.414829,0.292960,...		1,1,2,...

	coef_shift <list>	dist_shift <list>
[1]	6,18,23,... 0.611799,0.217936,0.217936,...	
[2]	16,21,22,... 0.296362,0.492958,0.252045,...	
[3]	-11,-3,-5,... 0.371425,0.382139,0.372516,...	
[4]	-13,-6,-10,... 0.531585,0.556766,0.556058,...	
[5]	-14,9,14,... 0.528634,0.217563,0.217563,...	

```
[6]      1,14,6,... 0.273368,0.413050,0.273738,...
-----
seqinfo: 20 sequences from an unspecified genome; no seqlengths
>
```

```
> # here the analysis is run on the spline approximation
> # with scaling
> peaks.cluster.scaled <- cluster_peak(peaks.summit.scaled, parallel = FALSE, seeds=1:4,
+                                     n.clust = 1:4, shift = NULL,
+                                     weight = 1, alpha = 1, p = 2, t.max = 2,
+                                     plot.graph.k = TRUE, verbose = FALSE,
+                                     rescale = TRUE)
> head(peaks.cluster.scaled)
```

GRanges object with 6 ranges and 15 metadata columns:

	seqnames	ranges	strand	counts	spline
	<Rle>	<IRanges>	<Rle>	<list>	<list>
[1]	chr18	3337524-3338025	*	7,8,8,... 0.107577,0.115944,0.124571,...	
[2]	chr18	4369126-4369352	*	7,9,9,... 0.106531,0.125616,0.145441,...	
[3]	chr18	4375448-4375883	*	8,8,8,... 0.103757,0.127649,0.152502,...	
[4]	chr18	4715744-4716162	*	5,5,5,... 0.108025,0.121182,0.134814,...	
[5]	chr18	4716374-4716597	*	15,15,15,... 0.109476,0.126596,0.144360,...	
[6]	chr18	4921270-4921506	*	8,8,8,... 0.104672,0.120304,0.136527,...	

	spline_der	width_spline	start_spline	end_spline
	<list>	<integer>	<numeric>	<numeric>
[1]	0.00823719,0.00849664,0.00875919,...	593	3337483	3338075
[2]	0.0187191,0.0194523,0.0201997,...	335	4369075	4369409
[3]	0.0234173,0.0243690,0.0253395,...	508	4375417	4375924
[4]	0.0129218,0.0133933,0.0138718,...	523	4715698	4716220
[5]	0.0168009,0.0174399,0.0180908,...	315	4716303	4716617
[6]	0.0153411,0.0159259,0.0165217,...	332	4921234	4921565

	spline_rescaled	spline_der_rescaled
	<list>	<list>
[1]	0.0026361:0.00304401:0.00347677:... 0.000395668:0.000420244:0.000445387:...	
[2]	0.00122834:0.00147221:0.00172654:... 0.000238709:0.000249059:0.000259633:...	
[3]	0.000864564:0.00120333:0.00156496:... 0.000327572:0.000350081:0.000373281:...	
[4]	0.000788311:0.000956446:0.00113505:... 0.000162986:0.000173327:0.000183923:...	
[5]	0.00235269:0.0027355:0.0031333:... 0.000375411:0.000390254:0.000405393:...	
[6]	0.00139356:0.00162209:0.00186007:... 0.000223859:0.00023322:0.000242774:...	

	summit_spline_rescaled	summit_spline	cluster_N0shift
	<integer>	<integer>	<list>
[1]	227	444	1,1,1,...
[2]	168	186	1,2,2,...
[3]	104	174	1,2,3,...
[4]	180	310	1,2,2,...
[5]	117	121	1,2,3,...
[6]	155	169	1,2,2,...

	dist_N0shift	cluster_shift	coef_shift
	<list>	<list>	<list>
[1]	0.01354655,0.00634872,0.00634872,...	1,1,1,...	1,36,41,...
[2]	0.00314305,0.00377609,0.00260030,...	1,2,2,...	7,8,10,...

```
[3] 0.00668285,0.00627910,0.00460132,...    1,2,3,... -17,-13,-4,...
[4] 0.00461631,0.00470065,0.00449050,...    1,2,2,...  -6,-2,-1,...
[5] 0.00446076,0.00382455,0.00269385,...    1,2,3,... -13,-11,2,...
[6] 0.00130264,0.00127949,0.00132011,...    1,2,2,...  -3,-3,-2,...
      dist_shift
      <list>
[1] 0.01360197,0.00529869,0.00529869,...
[2] 0.00288555,0.00288555,0.00196126,...
[3] 0.00556534,0.00561814,0.00411762,...
[4] 0.00469498,0.00474231,0.00488327,...
[5] 0.00330416,0.00331770,0.00579343,...
[6] 0.00121593,0.00123025,0.00000000,...
-----
seqinfo: 20 sequences from an unspecified genome; no seqlengths
>
```




Figure 4: Global distance within clusters. Global distance of the peaks from the corresponding template, as a function of the number of clusters k . In the left panel the graph for the original spline approximation, while in the right panel results are relative to the scaled approximation.

The particular case of k -mean alignment with $k = 1$ clusters can be used to highlight the effects of the alignment of the peaks: no grouping is performed, just the shifts are computed. Therefore, the decrease of the global distance is solely due to a change of the abscissae of the functions, as Figure ?? shows. Moreover, focusing for example on the first panel of Figure ??, we can deduce that, for this case

- the alignment can effectively decrease the distance, for example for $k = 6$, the gap between red and black line is significant;
- the alignment may change the optimal k : looking at the black line, one would have chosen $k = 4$, while the red line suggests $k = 3$ is the best choice. With the introduction of the shifts, data which are originally different becomes more similar and therefore one less cluster is needed; it has to be noted that the distance obtained with $k = 3$ and alignment is very similar to the one obtained with $k = 4$ and no alignment.

Therefore, for this case, one possible classification is the one associated to $k = 3$ with shift. On the contrary for the scaled peaks the value of k we can identify as crucial is $k = 2$ and shift is relevant since it reduces a lot the global distance. In Section ?? we present two quantitative methods to isolate the most suitable value of k : the first quantifies the elbow

rule we have presented here and the second that computes the Silhouette index to consider the global structure of the dataset. The results for this specific number of clusters can then be selected with the `choose_k` method:

```
> # select the results for k = 3 with alignment
> peaks.classified.short <- choose_k(peaks.cluster, k = 3,
+                                   shift = TRUE, cleaning = TRUE)
> head(peaks.classified.short)
```

GRanges object with 6 ranges and 1 metadata column:

	seqnames	ranges	strand	cluster
	<Rle>	<IRanges>	<Rle>	<numeric>
[1]	chr18	3337524-3338025	*	1
[2]	chr18	4369126-4369352	*	2
[3]	chr18	4375448-4375883	*	2
[4]	chr18	4715744-4716162	*	2
[5]	chr18	4716374-4716597	*	1
[6]	chr18	4921270-4921506	*	2

seqinfo: 20 sequences from an unspecified genome; no seqlengths

```
> peaks.classified.extended <- choose_k(peaks.cluster, k = 3,
+                                       shift = TRUE, cleaning = FALSE)
> # and for the scaled version for k =2 and alignment
>
> peaks.classified.scaled.short <- choose_k(peaks.cluster.scaled, k = 2,
+                                           shift = TRUE, cleaning = TRUE)
> head(peaks.classified.scaled.short)
```

GRanges object with 6 ranges and 4 metadata columns:

	seqnames	ranges	strand	spline_rescaled
	<Rle>	<IRanges>	<Rle>	<list>
[1]	chr18	3337524-3338025	*	0.0026361:0.00304401:0.00347677:...
[2]	chr18	4369126-4369352	*	0.00122834:0.00147221:0.00172654:...
[3]	chr18	4375448-4375883	*	0.000864564:0.00120333:0.00156496:...
[4]	chr18	4715744-4716162	*	0.000788311:0.000956446:0.00113505:...
[5]	chr18	4716374-4716597	*	0.00235269:0.0027355:0.0031333:...
[6]	chr18	4921270-4921506	*	0.00139356:0.00162209:0.00186007:...

	spline_der_rescaled	summit_spline_rescaled	cluster
	<list>	<integer>	<numeric>
[1]	0.000395668:0.000420244:0.000445387:...	227	1
[2]	0.000238709:0.000249059:0.000259633:...	168	2
[3]	0.000327572:0.000350081:0.000373281:...	104	2
[4]	0.000162986:0.000173327:0.000183923:...	180	2
[5]	0.000375411:0.000390254:0.000405393:...	117	2
[6]	0.000223859:0.00023322:0.000242774:...	155	2

seqinfo: 20 sequences from an unspecified genome; no seqlengths

```
> peaks.classified.scaled.extended <- choose_k(peaks.cluster.scaled, k = 2,
+                                              shift = TRUE, cleaning = FALSE)
```


The `choose_k` method allows, respectively, to remove all the metadata columns computed by `FunChIP` and obtain a `GRanges` equivalent to the initial one, with an extra the metadata column `cluster` containing the classification labels (`cleaning = TRUE`), or a `GRanges` retaining all the details of the preprocessing and clustering (all the previously described metadata columns), with the extra column `cluster` (`cleaning = FALSE`).

4.3 The choice of the final classification with the bending index and the Silhouette plot

In this section we present two analyses which could be useful to identify the most suitable value of the number of clusters.

The first analysis quantifies the elbow rule we proposed in the previous section. The function `bending_index` computes for each value of k in the range $[2 : K - 1]$, with K the maximum number of clusters allowed in `cluster_peak`, an index that quantifies the bending of the curve. The higher this index, the more appropriate is the correspondent value of k . The index is computed as the distance of the point P in k of the global distance function, normalized to its maximum value, ($P = (k, \tilde{D}(k))$) from the line r passing by the point in $k - 1$ and in $k + 1$ ($r = r((k - 1, \tilde{D}(k - 1)), (k + 1, \tilde{D}(k + 1)))$). The normalized distance function is

$$\tilde{D}(k) = \frac{D(k)}{\max_{k=1:K} D(k)}$$

with $D(k)$ being the global distance of the set of classified peaks, when divided into k clusters:

$$D(k) = \sum_{i=1}^n d_{i,k}^*.$$

In particular, $d_{i,k}^*$ is the distance of each peak i from the center of the correspondent cluster $c_{k_i}^*$. Then, the index can be written as:

$$r(k) = \text{dist}((k, \tilde{D}(k)), r((k - 1, \tilde{D}(k - 1)), (k + 1, \tilde{D}(k + 1))))$$

```
> # here we compute the bending index for the classification
> # provided for the non scaled dataset
>
> index <- bending_index(peaks.cluster, plot.graph.k = FALSE)
> index

$index_shift
      2      3
0.009578264 0.044113662

$index_N0shift
      2      3
0.01664078 0.03811411

> # and then for the scaled dataset
> index_scaled <- bending_index(peaks.cluster.scaled, plot.graph.k = FALSE)
> index_scaled

$index_shift
      2      3
```

```
0.04202466 0.02561751

$index_N0shift
      2      3
0.065931403 0.002706642

>
```

The function returns the value of the bending index varying k for both the classification with `index_shift` and without `index_N0shift` alignment, if both classifications are stored in the object. It can also show the plot of the global distance curves of Figure ??, if `plot.graph.k = TRUE`.

Moreover, we present another method to evaluate the choice the proper number of clusters, which takes into account the global set of data. This method is based on the definition of the Silhouette index of [?] and, for the peak i , the index is computed as

$$s(i) = \frac{a(i) - b(i)}{\max(a(i), b(i))},$$

with $a(i)$ being the average dissimilarity of peak i with all other data within the same cluster and $b(i)$ the lowest average dissimilarity of i to any other cluster, of which i is not a member. To compute the $b(i)$ value it is necessary to compare peaks belonging to different clusters. In case of classification with alignment, to define the distance of peaks of different clusters it is necessary first to align the two clusters. To perform this global alignment in an efficient way, we align the centers of the clusters and then use the estimated shift coefficient to align all the peaks of the two clusters.

```
> sil <- silhouette_plot(peaks.cluster, p=2, weight = 1, alpha = 1,
+                        rescale = FALSE, t.max = 2)

[1] "Silhouette for the non aligned peaks"
[1] "Silhouette for the aligned peaks"
```

```
> sil <- silhouette_plot(peaks.cluster.scaled, p=2, weight = 1, alpha = 1,
+                        rescale = FALSE, t.max = 2)

[1] "Silhouette for the non aligned peaks"
[1] "Silhouette for the aligned peaks"
```

Th plots shown in Figure ?? represent the Silhouette index for all the peaks divided in the clusters, together with the global average Silhouette index. The average Silhouette index is maximum when the clustering is separating the data in the most suited manner and then this index has to be maximized to choose the best classification. In the case here presented here, the best classification according to this criteria is the classification with alignment and $k = 2$ or $k = 3$. Considering both Silhouette and the bending criteria, in the next section we present the results of the classification with alignment and $k = 3$ clusters.

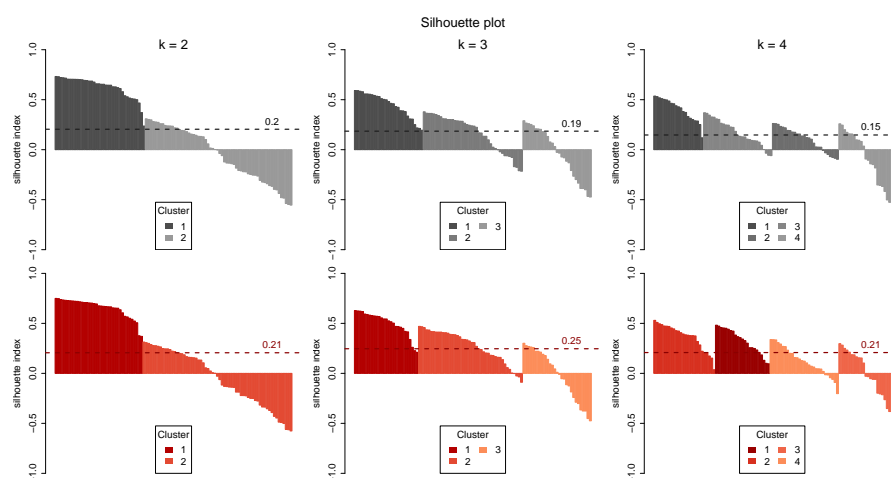


Figure 5: Silhouette graph for the classification of unscaled peaks. In the top panel, Silhouette plots for the classification without alignment, as function of the number of clusters; in the bottom panel, Silhouette plots for the classification with alignment.

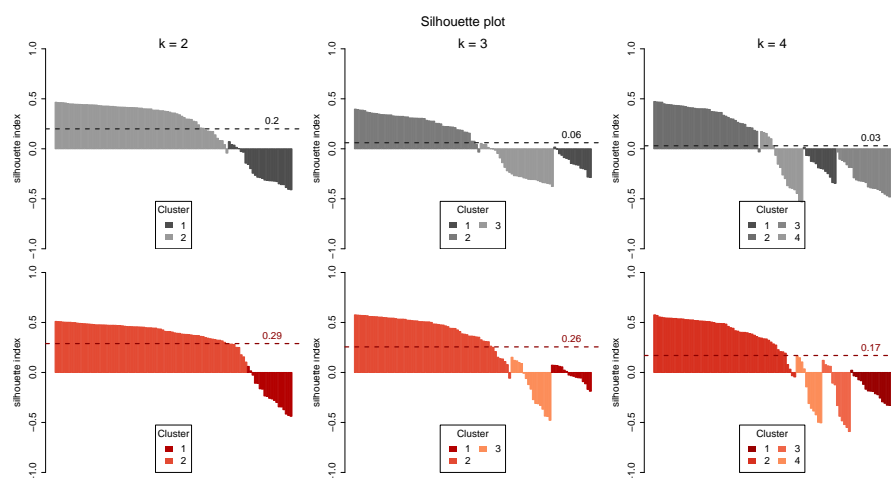


Figure 6: Silhouette graph for the classification of scaled peaks. In the top panel, Silhouette plots for the classification without alignment, as a function of the number of clusters; in the bottom panel, Silhouette plots for the classification with alignment.

5 Visualization of the peaks

The `plot_peak` method is a very flexible function for displaying ChIP-Seq peaks. In particular, it allows to plot the raw counts obtained by the `pileup_peak` method, as in Figure ?? . It can also plot smoothed peaks, possibly centered around the summit, as in Figure ?? , or scaled as in Figure ?? and centered.

```
> # plot of the first 10 peaks (raw data)
> par(mar=c(4.5,5,4,4))
> plot_peak(peaks, index = 1:10, line.plot = 'count',
+          cex.main = 2, cex.lab = 2, cex.axis = 2)
```




Figure 7: 10 peaks: counts. Representation of the original peaks as raw counts (no smoothing).

```
> # plot of the smoothed approximation of the first 10 peaks
> par(mar=c(4.5,5,4,4))
> plot_peak(peaks.smooth, index = 1:10, line.plot = 'spline', cex.main = 2, cex.lab = 2, cex.axis = 2)
>
```

```
> # plot of the smoothed approximation of the first 10 peaks,
> # centering peaks around their summits
> par(mar=c(4.5,5,4,4))
> plot_peak(peaks.summit, index = 1:10, line.plot = 'spline', cex.main = 2, cex.lab = 2, cex.axis = 2)
```

```
> # plot of the smoothed approximation of the first 10 peaks;
> # the scaled functions are plotted.
> #
> par(mar=c(4.5,5,4,4))
> plot_peak(peaks.smooth.scaled, index = 1:10,
+           line.plot = 'spline', rescale = TRUE,
+           cex.main = 2, cex.lab = 2, cex.axis = 2)
>
```

```
> # plot of the scaled approximation of the first 10 peaks,
> # centering peaks around their summits
> par(mar=c(4.5,5,4,4))
> plot_peak(peaks.summit.scaled, index = 1:10,
+           line.plot = 'spline', rescale = TRUE,
+           cex.main = 2, cex.lab = 2, lwd = 2, cex.axis = 2)
```

From the comparison of Figure ?? and Figure ?? it is clear how the scaling affects the shape of splines. Now peaks are no more related to the magnitude, but just to their shapes.

Moreover, plotting both raw counts and spline is also possible: Figure ?? shows a single peak in its raw and smoothed version. This representation is useful to check the accuracy of the smoothing and, if needed, manually set the λ parameter of the spline approximation.




Figure 8: 10 spline-smoothed peaks. In the left panel, smoothed peaks are shown, while in the right panel the same peaks are centered around their summits.




Figure 9: 10 spline-smoothed and scaled peaks. In the left panel, scaled peaks are shown, while in the right panel the same peaks are centered around their summits.

```
> # plot of a peak comparing its raw structure and
> # its spline-smoothed version.
> par(mar=c(4.5,5,4,4))
> plot_peak(peaks.summit, index = 3, lwd =2, line.plot = 'both', col = 'darkblue', cex.main = 2,cex.lab = 2,
```




Figure 10: Read coverage and spline approximation. Plot of the original read coverage of a peak and its smoothing (spline approximation), centered around the summit.

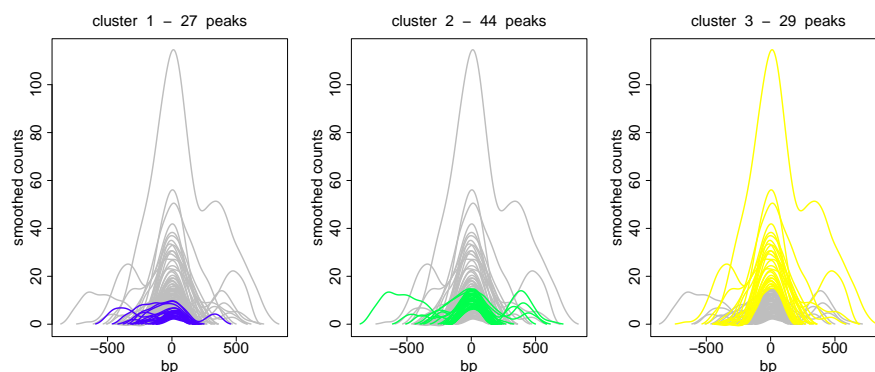


Figure 11: Peaks divided in the three clusters The same spline-smoothed peaks are plotted in grey, and for each panel the peaks in the corresponding cluster are colored to show their different shapes. Peaks are aligned with the shift coefficients obtained by the k-mean alignment algorithm.




Figure 12: Scaled peaks divided in the two clusters The same spline-smoothed scaled peaks are plotted in grey, and for each panel the peaks in the corresponding cluster are colored to show their different shapes. Peaks are aligned with the shift coefficients obtained by the k-mean alignment algorithm.

```
> # plot of the results of the kmean alignment.
> # Peaks are plotted in three different panels
> # according to the clustering results.
>
> plot_peak(peaks.cluster, index = 1:100, line.plot = 'spline',
+          shift = TRUE, k = 3, cluster.peak = TRUE,
+          col = topo.colors(3), cex.main = 2, cex.lab = 2, cex.axis = 2)
```

```
> # plot of the results of the kmean alignment.
> # Scaled peaks are plotted in three different panels
> # according to the clustering results.
>
> plot_peak(peaks.cluster.scaled, index = 1:100, line.plot = 'spline',
+          shift = TRUE, k = 2, cluster.peak = TRUE, rescale = TRUE,
+          col = heat.colors(2), cex.main = 2, cex.lab = 2, cex.axis = 2)
>
```

***FunChIP*: A Functional Data Analysis approach to cluster ChIP-Seq peaks according to their shapes**

Finally, the `plot_peak` method allows to plot the results of the clustering via the k-mean alignment. In Figure ?? and Figure ??, smoothed and scaled peaks are divided into the three clusters and plotted with the optimal shift obtained with the alignment.