segmentSeq: methods for identifying small RNA loci from high-throughput sequencing data

Thomas J. Hardcastle

August 4, 2015

1 Introduction

High-throughput sequencing technologies allow the production of large volumes of short sequences, which can be aligned to the genome to create a set of *matches* to the genome. By looking for regions of the genome which to which there are high densities of matches, we can infer a segmentation of the genome into regions of biological significance. The methods we propose allows the simultaneous segmentation of data from multiple samples, taking into account replicate data, in order to create a consensus segmentation. This has obvious applications in a number of classes of sequencing experiments, particularly in the discovery of small RNA loci and novel mRNA transcriptome discovery.

We approach the problem by considering a large set of potential *segments* upon the genome and counting the number of tags that match to that segment in multiple sequencing experiments (that may or may not contain replication). We then adapt the empirical Bayesian methods implemented in the `baySeq` package [1] to establish, for a given segment, the likelihood that the count data in that segment is similar to background levels, or that it is similar to the regions to the left or right of that segment. We then rank all the potential segments in order of increasing likelihood of similarity and reject those segments for which there is a high likelihood of similarity with the background or the regions to the left or right of the segment. This gives us a large list of overlapping segments. We reduce this list to identify non-overlapping loci by choosing, for a set of overlapping segments, the segment which has the lowest likelihood of similarity with either background or the regions to the left or right of that segment and rejecting all other segments that overlap with this segment. For fuller details of the method, see Hardcastle *et al.* [2].

2 Preparation

We begin by loading the `segmentSeq` package.

```
>    library(segmentSeq)
```

Note that because the experiments that `segmentSeq` is designed to analyse are usually massive, we should use (if possible) parallel processing as implemented by the `parallel` package. If using this approach, we need to begin by define a *cluster*. The following command will use eight processors on a single machine; see the help page for 'makeCluster' for more information. If we don't want to parallelise, we can proceed anyway with a `NULL` cluster.

```
> if(require("parallel"))
+ {
+     numCores <- min(8, detectCores())
+     cl <- makeCluster(numCores)
+ } else {
+     cl <- NULL
+ }
```

The `readGeneric` function is able to read in tab-delimited files which have appropriate column names, and create an `alignmentData` object. Alternatively, if the appropriate column names are not present, we can specify which columns to use for the data. In either case, to use this function we pass a character vector of files, together with information on which data are to be treated as replicates to the function. We also need to define the lengths of

the chromosome and specifiy the chromosome names as a character. The data here, drawn from text files in the 'data' directory of the segmentSeq package are taken from the first million bases of an alignment to chromosome 1 and the first five hundred thousand bases of an alignment to chromosome 2 of *Arabidopsis thaliana* in a sequencing experiment where libraries 'SL9' and 'SL10' are replicates, as are 'SL26' and 'SL32'. Libraries 'SL9' and 'SL10' are sequenced from an Argonaute 6 IP, while 'SL26' and 'SL32' are an Argonaute 4 IP.

A similar function, readBAM performs the same operation on files in the BAM format. Please consult the help page for further details.

```
> chrlens <- c(1e6, 2e5)
> datadir <- system.file("extdata", package = "segmentSeq")
> libfiles <- c("SL9.txt", "SL10.txt", "SL26.txt", "SL32.txt")
> libnames <- c("SL9", "SL10", "SL26", "SL32")
> replicates <- c("AGO6", "AGO6", "AGO4", "AGO4")
> aD <- readGeneric(files = libfiles, dir = datadir,
+                   replicates = replicates, libnames = libnames,
+                   chrs = c(">Chr1", ">Chr2"), chrlens = chrlens,
+                   polyLength = 10, header = TRUE, gap = 200)
> aD

An object of class "alignmentData"
13765 rows and 4 columns

Slot "libnames":
[1] "SL9"  "SL10" "SL26" "SL32"

Slot "replicates":
[1] AGO6 AGO6 AGO4 AGO4
Levels: AGO4 AGO6

Slot "alignments":
GRanges object with 13765 ranges and 2 metadata columns:
          seqnames             ranges strand  |                  tag multireads
             <Rle>          <IRanges>  <Rle>  |          <character>  <numeric>
     [1]     >Chr1         [265, 284]      -  |   AAATGAAGATAAACCATCCA          1
     [2]     >Chr1         [405, 427]      -  | AAGGAGTAAGAATGACAATAAAT          1
     [3]     >Chr1         [406, 420]      -  |      AAGAATGACAATAAA          1
     [4]     >Chr1         [600, 623]      +  | AAGGATTGGTGGTTTGAAGACACA          1
     [5]     >Chr1         [665, 688]      +  | ATCCTTGTAGCACACATTTTGGCA          1
     ...       ...                ...    ... ...                  ...        ...
 [13761]     >Chr2 [179972, 179993]      +  |   ATGAATGGCTCTCTCTAGCGGA          1
 [13762]     >Chr2 [179978, 180000]      -  |   GAGATTCTCCGCTAGAGAGAGCC          1
 [13763]     >Chr2 [179999, 180022]      -  | ATTAATATTAATTCATCGGGAAGA          1
 [13764]     >Chr2 [180002, 180022]      -  |    ATTAATATTAATTCATCGGGA          1
 [13765]     >Chr2 [180014, 180037]      +  | AATATTAATGGTATTTGTGGAAAA          1
 -------
 seqinfo: 2 sequences from an unspecified genome

Slot "data":
Matrix with  13765  rows.
      SL9 SL10 SL26 SL32
1       1    0    0    0
2       0    0    0    2
3       0    1    0    0
4       0    1    0    0
5       7    1    0    0
...   ...  ...  ...  ...
13761   2    7    0    0
13762   0    1    0    0
13763   0    1    0    0
```

```
13764    0    1    0    0
13765    1    0    0    0
```

```
Slot "libsizes":
[1] 4447 6531 9666 6675
```

Next, we process this `alignmentData` object to produce a `segData` object. This `segData` object contains a set of potential segments on the genome defined by the start and end points of regions of overlapping alignments in the `alignmentData` object. It then evaluates the number of tags that hit in each of these segments.

```
> sD <- processAD(aD, gap = 100, cl = cl)
> sD
```

```
An object of class "segData"
14444 rows and 4 columns
```

```
Slot "replicates":
[1] AGO6 AGO6 AGO4 AGO4
Levels: AGO4 AGO6
```

```
Slot "coordinates":
GRanges object with 14444 ranges and 0 metadata columns:
           seqnames             ranges strand
              <Rle>          <IRanges>  <Rle>
      [1]      >Chr1         [265, 284]      *
      [2]      >Chr1         [405, 427]      *
      [3]      >Chr1         [600, 623]      *
      [4]      >Chr1         [600, 688]      *
      [5]      >Chr1         [600, 830]      *
      ...        ...                ...    ...
  [14440]      >Chr2 [179708, 179872]      *
  [14441]      >Chr2 [179708, 180037]      *
  [14442]      >Chr2 [179738, 179872]      *
  [14443]      >Chr2 [179738, 180037]      *
  [14444]      >Chr2 [179923, 180037]      *
  -------
  seqinfo: 2 sequences from an unspecified genome
```

```
Slot "locLikelihoods" (stored on log scale):
Matrix with  0  rows.
<0 x 0 matrix>
```

```
Slot "data":
Matrix with  0  rows.Matrix with  0  rows.
     SL9 SL10 SL26 SL32
```

```
Slot "libsizes":
[1] 4447 6531 9666 6675
```

We can now construct a segment map from these potential segments.

Segmentation by heuristic methods

A fast method of segmentation can be achieved by exploiting the bimodality of the densities of small RNAs in the potential segments. In this approach, we assign each potential segment to one of two clusters for each replicate group, either as a segment or a null based on the density of sequence tags within that segment. We then combine these clusterings for each replicate group to gain a consensus segmentation map.

```
> hS <- heuristicSeg(sD = sD, aD = aD, RKPM = 1000, largeness = 1e8, getLikes = TRUE, cl = cl)
```

......

Segmentation by empirical Bayesian methods

A more refined approach to the problem uses an existing segment map (or, if not provided, a segment map defined by the hS function) to acquire empirical distributions on the density of sequence tags within a segment. We can then estimate posterior likelihoods for each potential segment as being either a true segment or a null. We then identify all potential segments in the with a posterior likelihood of being a segment greater than some value 'lociCutoff' and containing no subregion with a posterior likelihood of being a null greater than 'nullCutoff'. We then greedily select the longest segments satisfying these criteria that do not overlap with any other such segments in defining our segmentation map.

```
> classSegs <- classifySeg(sD = sD, aD = aD, cD = hS, cl = cl)
```

....

```
> classSegs

GRanges object with 254 ranges and 0 metadata columns:
        seqnames            ranges strand
           <Rle>         <IRanges>  <Rle>
    [1]     >Chr1 [    1,    599]      *
    [2]     >Chr1 [  600,    938]      *
    [3]     >Chr1 [  939,    967]      *
    [4]     >Chr1 [  968,  17054]      *
    [5]     >Chr1 [17055,  18728]      *
    ...       ...               ...    ...
  [250]     >Chr2 [169231, 178343]      *
  [251]     >Chr2 [178344, 178636]      *
  [252]     >Chr2 [178637, 179707]      *
  [253]     >Chr2 [179708, 180037]      *
  [254]     >Chr2 [180038, 200000]      *
  -------
  seqinfo: 2 sequences from an unspecified genome
An object of class "lociData"
254 rows and 4 columns

Slot "replicates"
AGO6 AGO6 AGO4 AGO4
Slot "groups":
[[1]]
[1] AGO6 AGO6 AGO4 AGO4
Levels: AGO4 AGO6


Slot "data":
     AGO6.1 AGO6.2 AGO4.1 AGO4.2
[1,]      1      1      0      2
[2,]     41     39     65     83
[3,]     13      7      0      0
[4,]      2      3      0      0
[5,]    682    621   1405   1103
249 more rows...

Slot "annotation":
data frame with 0 columns and 254 rows

Slot "locLikelihoods" (stored on log scale):
Matrix with  254  rows.
```

```
        AGO4      AGO6
1     0.26082 0.098357
2     0.93552  0.96812
3     0.13704  0.99008
4    0.036916 0.032084
5     0.98633  0.99537
...       ...      ...
250   0.32593 0.061916
251   0.92245  0.98141
252  0.056506 0.060477
253   0.94817  0.97325
254  0.036254 0.026154


Expected number of loci in each replicate group
     AGO4      AGO6
113.6243 139.0449
```

By one of these methods, we finally acquire an annotated `lociData` object, with the annotations describing the co-ordinates of each segment.

We can use this `lociData` object, in combination with the `alignmentData` object, to plot the segmented genome.

```
> par(mfrow = c(2,1), mar = c(2,6,2,2))
> plotGenome(aD, hS, chr = ">Chr1", limits = c(1, 1e5),
+          showNumber = FALSE, cap = 50)
> plotGenome(aD, classSegs, chr = ">Chr1", limits = c(1, 1e5),
+          showNumber = FALSE, cap = 50)
```

Given the calculated likelihoods, we can filter the segmented genome by controlling on likelihood, false discovery rate, or familywise error rate

```
> loci <- selectLoci(classSegs, FDR = 0.05)
> loci
GRanges object with 116 ranges and 0 metadata columns:
        seqnames            ranges strand
           <Rle>         <IRanges>  <Rle>
   [1]     >Chr1 [  600,    938]      *
   [2]     >Chr1 [  939,    967]      *
   [3]     >Chr1 [17055,  18728]      *
   [4]     >Chr1 [42217,  42435]      *
   [5]     >Chr1 [44710,  44811]      *
   ...       ...               ...    ...
 [112]     >Chr2 [144202, 144327]      *
 [113]     >Chr2 [152150, 152173]      *
 [114]     >Chr2 [169196, 169230]      *
 [115]     >Chr2 [178344, 178636]      *
 [116]     >Chr2 [179708, 180037]      *
   -------
   seqinfo: 2 sequences from an unspecified genome
An object of class "lociData"
116 rows and 4 columns

Slot "replicates"
AGO6 AGO6 AGO4 AGO4
Slot "groups":
[[1]]
[1] AGO6 AGO6 AGO4 AGO4
Levels: AGO4 AGO6
```

Figure 1: The segmented genome (first 10^5 bases of chromosome 1.

```
Slot "data":
      AGO6.1 AGO6.2 AGO4.1 AGO4.2
[1,]      41     39     65     83
[2,]      13      7      0      0
[3,]     682    621   1405   1103
[4,]      31     11     48     56
[5,]      73     57     47     21
111 more rows...

Slot "annotation":
data frame with 0 columns and 116 rows

Slot "locLikelihoods" (stored on log scale):
Matrix with  116  rows.
        AGO4    AGO6
1    0.93552 0.96812
2    0.13704 0.99008
3    0.98633 0.99537
4    0.94375 0.94224
5    0.95582 0.99689
...      ...     ...
112 0.79255 0.95779
```

```
113  0.1432 0.96747
114 0.13126  0.9877
115 0.92245 0.98141
116 0.94817 0.97325


Expected number of loci in each replicate group
     AGO4       AGO6
 79.48241 112.78859
```

The `lociData` objects can now be examined for differential expression with the `baySeq` package.

First we define the possible models of differential expression on the data. In this case, the models are of non-differential expression and pairwise differential expression.

```
> groups(classSegs) <- list(NDE = c(1,1,1,1), DE = c("AGO6", "AGO6", "AGO4", "AGO4"))
```

Then we get empirical distributions on the parameter space of the data.

```
> classSegs <- getPriors(classSegs, cl = cl)
```

Then we get the posterior likelihoods of the data conforming to each model. Since the 'classSegs' object contains null regions as well as true loci, we will use the 'nullData = TRUE' option to distinguish between non-differentially expressed loci and non-expressed regions. By default, the loci likelihoods calculated earlier will be used to weight the initial parameter fit in order to detect null data.

```
> classSegs <- getLikelihoods(classSegs, nullData = TRUE, cl = cl)

.
```

We can examine the highest likelihood non-expressed ('null') regions

```
> topCounts(classSegs, NULL, number = 3)

  seqnames  start    end width strand AGO6.1 AGO6.2 AGO4.1 AGO4.2 Likelihood       FDR.
1    >Chr1 754198 758593  4396      *      1      0      1      0  0.9340274 0.06597257
2    >Chr1 950796 958752  7957      *      2      4      3      2  0.9321692 0.06690166
3    >Chr1 148877 222440 73564      *      0      0      0      0  0.9271423 0.06888701
       FWER.
1 0.06597257
2 0.12932836
3 0.19276348
```

The highest likelihood expressed but non-differentially expressed regions

```
> topCounts(classSegs, "NDE", number = 3)

  seqnames start   end width strand AGO6.1 AGO6.2 AGO4.1 AGO4.2 Likelihood    FDR.NDE
1    >Chr2  1554  8147  6594      *   5326  12490  13684   9679  0.9813884 0.01861158
2    >Chr2  8148  8209    62      *     76   2999    643    217  0.9682372 0.02518721
3    >Chr1 446325 447437 1113      *    789    536   1291   1184  0.9609515 0.02980763
    FWER.NDE
1 0.01861158
2 0.04978326
3 0.08688778
```

And the highest likelihood differentially expressed regions

```
> topCounts(classSegs, "DE", number = 3)

  seqnames start   end width strand AGO6.1 AGO6.2 AGO4.1 AGO4.2 Likelihood  ordering
1    >Chr2 52652 53314   663      *     84     86    982    696  0.9988873 AGO4>AGO6
2    >Chr2 58131 59084   954      *     81     83    965    677  0.9986058 AGO4>AGO6
3    >Chr2 49137 50333  1197      *    101    116    992    719  0.9984271 AGO4>AGO6
       FDR.DE     FWER.DE
1 0.001112749 0.001112749
2 0.001253482 0.002505413
3 0.001359954 0.004074370
```

Finally, to be a good citizen, we stop the cluster we started earlier:

```
> if(!is.null(cl))
+     stopCluster(cl)
```

Session Info

```
> sessionInfo()

R version 3.2.1 (2015-06-18)
Platform: x86_64-unknown-linux-gnu (64-bit)
Running under: Ubuntu 14.04.2 LTS

locale:
 [1] LC_CTYPE=en_US.UTF-8       LC_NUMERIC=C               LC_TIME=en_US.UTF-8
 [4] LC_COLLATE=C               LC_MONETARY=en_US.UTF-8    LC_MESSAGES=en_US.UTF-8
 [7] LC_PAPER=en_US.UTF-8       LC_NAME=C                  LC_ADDRESS=C
[10] LC_TELEPHONE=C             LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C

attached base packages:
[1] stats4    parallel  stats     graphics  grDevices utils     datasets  methods
[9] base

other attached packages:
 [1] segmentSeq_2.2.2     ShortRead_1.26.0     GenomicAlignments_1.4.1
 [4] Rsamtools_1.20.4     Biostrings_2.36.2    XVector_0.8.0
 [7] BiocParallel_1.2.19  baySeq_2.2.0         perm_1.0-0.0
[10] abind_1.4-3          GenomicRanges_1.20.5 GenomeInfoDb_1.4.1
[13] IRanges_2.2.5        S4Vectors_0.6.3      BiocGenerics_0.14.0

loaded via a namespace (and not attached):
 [1] zlibbioc_1.14.0     lattice_0.20-33    hwriter_1.3.2       tools_3.2.1
 [5] grid_3.2.1          Biobase_2.28.0     latticeExtra_0.6-26 lambda.r_1.1.7
 [9] futile.logger_1.4.1 RColorBrewer_1.1-2 futile.options_1.0.0 bitops_1.0-6
[13] BiocStyle_1.6.0
```

References

[1] Thomas J. Hardcastle and Krystyna A. Kelly. *baySeq: Empirical Bayesian Methods For Identifying Differential Expression In Sequence Count Data.* BMC Bioinformatics (2010).

[2] Thomas J. Hardcastle and Krystyna A. Kelly and David C. Baulcombe. *Identifying small RNA loci from high-throughput sequencing data.* Bioinformatics (2012).