

CNV detection in exome sequencing data using **exomeCopy**

Michael Love
love@molgen.mpg.de

March 31, 2012

Abstract

exomeCopy is an R package implementing a hidden Markov model for predicting copy number variants (CNVs) from exome sequencing experiments without paired control experiments as in tumor/normal sequencing.¹ It models read counts in genomic ranges using negative binomial emission distributions depending on a hidden state of the copy number and on positional covariates such as GC-content and background read depth. Normalization and segmentation are performed simultaneously, eliminating the need for preprocessing of the raw read counts.

Contents

1	Introduction	1
2	Importing experiment data	2
2.1	Subdividing targeted regions	2
2.2	Counting reads in genomic ranges	4
2.3	Calculating GC-content	5
2.4	Exome sequencing data from 1000 Genomes Project	5
2.5	Generating background read depth	6
3	Introduction of the model	7
4	Simulating CNVs	8
5	Running exomeCopy	8
5.1	Running exomeCopy on one sample, one chromosome	8
5.2	Looping exomeCopy over multiple samples and chromosomes	9
5.3	Compiling results, filtering and plotting CNVs	10
5.4	Inspecting model parameters	12
6	Session info	13

1 Introduction

The **exomeCopy** package was designed to address the following situation:

- Target enrichment, such as exome enrichment, leads to non-uniform read depth, which is often correlated across samples.
- CNVs overlapping enriched regions can be detected as increases or decreases in read counts relative to “background” read depth, generated by averaging over a control set.

¹The model implemented in this package is described in [Love et al., 2011]. R packages which can be used for CNV detection in whole genome or tumor/normal paired exome sequencing data include **ReadDepth** and **ExomeCNV** respectively, available on CRAN.

- Individual samples can be more or less correlated with background read depth and have different dependencies on GC-content.

While exome sequencing is not designed for CNV genotyping, it can nevertheless be used for finding CNVs which overlap exons and are not common in the control set. It can provide an independent data source to be used in combination with higher resolution array-based methods. In this vignette we show how to import experiment data, generate background read depth, and recover CNVs using `exomeCopy`.

`exomeCopy` is designed to run on read counts from consecutive genomic ranges on a single chromosome, as it tries to identify higher or lower read depth relative to a baseline. We will use a wrapper function to loop the main function over multiple chromosomes and samples. Then these results can be compiled across chromosomes and samples.

2 Importing experiment data

The necessary genomic range information, read counts and positional covariates (background read depth and GC-content) should be stored in a *RangedData* object. The user must provide genomic ranges of targeted enrichment. For exome sequencing, one can use exon annotations. The `exomeCopy` package provides the following convenience functions: `subdivideGRanges` for subdividing the target ranges into nearly-equal sized ranges, `countBamInGRanges` for counting reads from a BAM file in the target ranges, `getGCcontent` for calculating GC content in the target ranges, and `generateBackground` for calculating the median read depth over a set of background samples.

In this vignette we will use a prepackaged *RangedData* object containing real exome sequencing read counts. To minimize memory size, we demonstrate with toy data how to construct a *RangedData* object from a BED file describing the targeted region, BAM files for the read mapping and a FASTA file for the reference sequence.

2.1 Subdividing targeted regions

`subdivideGRanges` divides the targeted genomic ranges into a set of ranges of nearly equal width, which exactly cover the original ranges. While `exomeCopy` can use range width as a covariate for modeling read counts, we find it useful to break apart the largest input ranges into multiple ranges of comparable width to the average input range. `subdivideGRanges` requires an input *GRanges* object and returns a *GRanges* object.

```
> library(exomeCopy)
> gr <- GRanges(seqname = "seq1", IRanges(start = 1, end = 345))
> subdivideGRanges(gr)
```

GRanges with 3 ranges and 0 elementMetadata cols:

	seqnames	ranges	strand
	<Rle>	<IRanges>	<Rle>
[1]	seq1	[1, 115]	*
[2]	seq1	[116, 230]	*
[3]	seq1	[231, 345]	*

seqlengths:


```
seq1
NA
```

The default setting of `subdivideGRanges` is to divide an input range into ranges around $s = 100\text{bp}$, which is slightly less than the average exon width. Specifically, an input range of width w will be divided evenly into $\max(1, \lfloor w/s \rfloor)$ regions. We can visualize the effect of `subdivideGRanges` on ranges of increasing width.

```

> plot(0, 0, xlim = c(0, 500), ylim = c(0, 25), type = "n",
+      yaxt = "n", ylab = "", xlab = "width of input GRanges object",
+      main = "Effect of subdivideGRanges")
> abline(v = 1:5 * 100, col = "grey")
> for (i in 1:24) {
+   gr <- GRanges(seqname = "chr1", IRanges(start = 1,
+     width = (i * 20)))
+   sbd.gr <- subdivideGRanges(gr)
+   arrows(start(sbd.gr), rep(i, length(sbd.gr)), end(sbd.gr),
+     rep(i, length(sbd.gr)), length = 0.04, angle = 90,
+     code = 3)
+ }

```



Here we demonstrate reading in a targeted region BED file, converting to a *GRanges* object and the result from calling `subdivideGRanges`. Note that if the targeted region is read in from a BED file, one should add 1 to the starting position for representation as a *GRanges* object.

```

> target.file <- system.file("extdata", "targets.bed",
+   package = "exomeCopy")
> target.df <- read.delim(target.file, header = FALSE,
+   col.names = c("seqname", "start", "end"))
> target <- GRanges(seqname = target.df$seqname, IRanges(start = target.df$start +
+   1, end = target.df$end))
> target

```

GRanges with 5 ranges and 0 elementMetadata cols:

	seqnames	ranges	strand
	<Rle>	<IRanges>	<Rle>
[1]	seq1	[101, 250]	*
[2]	seq1	[301, 650]	*
[3]	seq2	[1, 150]	*
[4]	seq2	[401, 550]	*
[5]	seq2	[701, 750]	*

```

seqlengths:
  seq1 seq2
    NA  NA

> target.sub <- subdivideGRanges(target)
> target.sub

GRanges with 7 ranges and 0 elementMetadata cols:
      seqnames      ranges strand
      <Rle>    <IRanges> <Rle>
[1]      seq1 [101, 250]      *
[2]      seq1 [301, 417]      *
[3]      seq1 [418, 533]      *
[4]      seq1 [534, 650]      *
[5]      seq2 [  1, 150]      *
[6]      seq2 [401, 550]      *
[7]      seq2 [701, 750]      *
---
seqlengths:
  seq1 seq2
    NA  NA

```

2.2 Counting reads in genomic ranges

`countBamInGRanges` allows the user to count reads from a BAM read mapping file in genomic ranges covering the targeted region. The function takes as input the BAM filename and a *GRanges* object. It returns a vector of counts, representing the number of sequenced read starts (leftmost position regardless of strand) with mapping quality above a minimum threshold (default of 1) for each genomic range. It is also possible to count stranded read starts, all overlapping reads and only reads mapping to unique positions (see `?countBamInGRanges`).

Users should make sure the sequence names in the *GRanges* object are the same as the sequence names in the BAM file (which can be listed using `scanBamHeader` in the *Rsamtools* package). The BAM file requires a associated index file (see the man page for `indexBam` in the *Rsamtools* package). We will count reads using the subdivided genomic ranges in `target.sub` and store the counts as a new value column, `sample1`.

```

> bam.file <- system.file("extdata", "mapping.bam", package = "exomeCopy")
> scanBamHeader(bam.file)[[1]]$targets

seq1 seq2
  800  800

> seqlevels(target.sub)

[1] "seq1" "seq2"

> rdata <- RangedData(space = seqnames(target.sub), ranges = ranges(target.sub))
> sample.df <- data.frame(samples = "sample1", bam.files = bam.file,
+   stringsAsFactors = FALSE)
> for (i in 1:nrow(sample.df)) {
+   rdata[[sample.df$samples[i]]] <- countBamInGRanges(sample.df$bam.files[i],
+     target.sub)
+ }
> rdata

RangedData with 7 rows and 1 value column across 2 spaces
      space      ranges |   sample1

```

	<factor>	<IRanges>		<integer>
1	seq1	[101, 250]		73
2	seq1	[301, 417]		59
3	seq1	[418, 533]		61
4	seq1	[534, 650]		54
5	seq2	[1, 150]		80
6	seq2	[401, 550]		69
7	seq2	[701, 750]		31

2.3 Calculating GC-content

exomeCopy can model read counts from samples which are not perfectly correlated with background read depth using GC-content (ratio of G and C bases to total number of bases). The GC-content of DNA fragments is known to be a factor in the efficiency of high-throughput sequencing. To obtain the GC-content information, we only need the targeted ranges and a reference FASTA file.

```
> reference.file <- system.file("extdata", "reference.fa",
+   package = "exomeCopy")
> rdata[["GC"]] <- getGCcontent(target.sub, reference.file)
> rdata
```

RangedData with 7 rows and 2 value columns across 2 spaces

	space	ranges		sample1	GC
	<factor>	<IRanges>		<integer>	<numeric>
1	seq1	[101, 250]		73	0.4533333
2	seq1	[301, 417]		59	0.4273504
3	seq1	[418, 533]		61	0.5689655
4	seq1	[534, 650]		54	0.5213675
5	seq2	[1, 150]		80	0.5733333
6	seq2	[401, 550]		69	0.4800000
7	seq2	[701, 750]		31	0.3600000

2.4 Exome sequencing data from 1000 Genomes Project

For demonstrating the use of exomeCopy, we have provided a sample dataset, `exomecounts`, of exome sequencing read counts. The genomic ranges used in this dataset are generated from a small subset of the CCDS regions on chromosome 1 [Pruitt et al., 2009]. The regions are downloaded from the hg19 tables of the UCSC Genome Browser (<http://genome.ucsc.edu/cgi-bin/hgGateway>). Alternatively, one could use the `GenomicFeatures` package to download the CCDS regions. The original CCDS regions are subdivided using `subdivideGRanges` with default settings as in the example code above. The CCDS regions are convenient to use as genomic ranges for CNV detection in exome sequencing data, as they are often in the center of the targeted region in exome enrichment protocols and tend to have less variable coverage than the flanking regions.

The read counts are taken from exome enriched, paired-end sequencing data of the 1000 Genomes Project for 16 samples of the PUR population [1000 Genomes Project Consortium, 2010]. The BAM read mapping files and descriptions of the experiments are available at the 1000 Genomes Project website (<http://www.1000genomes.org/data>). The ftp addresses used are listed in the file `1000Genomes_files.txt` in the `extdata` directory. The sample names are included as column names in the provided dataset.

```
> data(exomecounts)
> dim(exomecounts)
```


```
[1] 1000 17
```

```
> exomecounts[1:5, 1:3]
```

```
RangedData with 5 rows and 3 value columns across 1 space
  space      ranges |      GC  HG00551  HG00641
  <factor>  <IRanges> | <numeric> <numeric> <numeric>
1   chr1 [861322, 861393] |  0.6389    139    223
2   chr1 [865535, 865716] |  0.6484     45     90
3   chr1 [866419, 866469] |  0.5882     77    123
4   chr1 [871152, 871276] |  0.6480    254    285
5   chr1 [874420, 874509] |  0.6111     24     40
```

The genomic ranges in `exomecounts` have been filtered such that only ranges with nonzero read count over the 16 samples are retained. The range of the 1000 genomic ranges is from 0.8 to 7.8 Mb on chromosome 1. Plotting the counts for one sample in a region of 1 Mb, one can observe both the irregular spacing of the ranges as well as the non-uniformities in read counts per range.

```
> plot(start(exomecounts), exomecounts$HG00551, xlim = c(8e+05,
+ 1800000), xlab = "genomic position", ylab = "counts",
+ main = "HG00551 read counts in exonic ranges")
```



2.5 Generating background read depth

In order to run `exomeCopy`, we first generate background read depth. For demonstration of working with multiple chromosomes, we will copy the chr1 data as a new chromosome.

```
> chr1a <- exomecounts
> chr1a.ranges <- ranges(chr1a)
> names(chr1a.ranges) <- "chr1a"
> chr1a <- RangedData(chr1a.ranges, as.data.frame(exomecounts)[,
+ -c(1:4)])
> exomecounts.plus <- c(exomecounts, chr1a)
```

The background is generated by a simple function which performs 3 simple steps: 1) given a vector of names of samples to be used as background, extract the read counts data frame

from the *RangedData* object, 2) divide each sample by its mean read count (column means), 3) calculate the median of these normalized read counts (row medians). The function can also be used to calculate any function across the rows of normalized read counts. For example, we also show how to calculate the standard deviation of the normalized read counts for a measure of read depth deviation at each range, but we will not use covariate in this vignette.

Note that if the control set is a mix of males and females, the non-pseudoautosomal portion of chrX will be a mixture of copy counts 1 and 2. In this case, one should create separate *RangedData* objects for chrX and chrY and run males and females separately with the appropriate expected copy count.

```
> exome.samples <- grep("HG.", colnames(exomecounts.plus),
+   value = TRUE)
> exomecounts.plus[["bg"]] <- generateBackground(exome.samples,
+   exomecounts.plus, median)
> exomecounts.plus[["bg.sd"]] <- generateBackground(exome.samples,
+   exomecounts.plus, sd)
```

The relationship between read counts and GC-content over the ranges varies across protocols and samples. It can be roughly approximated per sample using second-order polynomial terms of GC-content. We store the square of GC-content as a new value column. Other functions of GC-content could be used as well. We also store the width of the ranges as a value column.

```
> exomecounts.plus[["GC.sq"]] <- exomecounts.plus$GC^2
> exomecounts.plus[["width"]] <- width(exomecounts.plus)
```

3 Introduction of the model

exomeCopy models the sample read counts on one chromosome, O , as emitted observations of a hidden Markov model (HMM), where the hidden state is the copy number of the sample. The emission distributions f are modeled with negative binomial distributions, as the read counts from high-throughput sequencing are often overdispersed for the Poisson distribution.

$$f \sim \text{NB}(O_t, \mu_{ti}, \phi)$$

$$\mu_{ti} = \frac{S_i}{d}(x_{t*}\beta)$$

The mean parameter, μ_{ti} , for genomic range t and hidden state i is a product of the possible copy number state S_i over the expected copy number d and an estimate of the positional effects on read depth ($x_{t*}\beta$), where x_{t*} is the t -th row of X and β is a column vector of coefficients. The estimated positional effect is a linear combination of background read depth, GC-content, range width, and any other useful covariates which are stored in the matrix X , with a row for each range and a column for each covariate. The coefficients β are fit by the model, using the forward equations to assess the likelihood of the HMM over all hidden state paths. In this way, the normalization and segmentation steps are combined into one step of maximizing the likelihood of the parameters given the data. The Viterbi algorithm is then applied to provide the most likely path.

The base model uses a scalar estimate for the dispersion parameter ϕ (equivalent to `1/size` in the function `dnbinom`). An extension of this model also tries to fit the variance using positional information such as the standard deviation of the background read depth stored in a matrix Y .

$$f \sim \text{NB}(O_t, \mu_{ti}, \phi_t)$$

$$\phi_t = y_{t*}\gamma$$

Both μ_{ti} and ϕ_t must be positive, so negative estimates are replaced with a small positive number (`1e-8`).

Table 1: Summary of notation

O_t	observed count of reads in the t -th genomic range
f	the emission distribution for read counts
μ_{ti}	the mean parameter for f at range t in copy state i
ϕ	the dispersion parameter for f
S_i	the copy number for state i ($S_i \in \{0, 1, 2, \dots\}$)
d	the expected background copy number (2 for diploid, 1 for haploid)
X	the matrix of covariates for estimating μ
Y	the matrix of covariates for estimating ϕ
β	the fitted coefficients for estimating μ
γ	the fitted coefficients for estimating ϕ

4 Simulating CNVs

Next we simulate CNVs in 8 samples on both chromosomes representing regions with a copy number of 0,1,3 and 4, relative to a background copy number of 2. This is accomplished by selecting a fraction of the reads contained in the CNV bounds and removing or doubling them.

```
> simulateCNV <- function(x, indices, multiply, prob) {
+   x[indices] <- x[indices] + multiply * rbinom(length(indices),
+     prob = prob, size = x[indices])
+   return(x)
+ }
> rdata <- exomecounts.plus
> set.seed(2)
> cnv.probs <- rep(c(0.99, 0.5, 0.5, 0.95), each = 2)
> cnv.mult <- rep(c(-1, 1), each = 4)
> for (i in 1:8) {
+   rdata[[exome.samples[i]]] <- simulateCNV(rdata[[exome.samples[i]]],
+     (i * 100 + 1):((i + 1) * 100), multiply = cnv.mult[i],
+     prob = cnv.probs[i])
+   rdata[[exome.samples[i]]] <- simulateCNV(rdata[[exome.samples[i]]],
+     1000 + (i * 100 + 1):((i + 1) * 100), multiply = cnv.mult[i],
+     prob = cnv.probs[i])
+ }
```

5 Running exomeCopy

5.1 Running exomeCopy on one sample, one chromosome

We can now run `exomeCopy` using the read counts for one of the simulated CNV samples. Later we will show how to write a simple wrapper function to loop the `exomeCopy` function over multiple chromosomes and samples. We specify the possible copy number values with `S` from 0 to 6 and the expected copy number state with `d` of 2.

```
> fit <- exomeCopy(rdata["chr1"], sample.name = exome.samples[3],
+   X.names = c("bg", "GC", "GC.sq", "width"), S = 0:6,
+   d = 2)
> show(fit)
```

```
ExomeCopy object
sample name: HG00731
percent normal state: 90%
```



After fitting, we call the function `copyCountSegments` on the *ExomeCopy* object, which provides the segmentation with the predicted copy number, the log odds of read counts being emitted from predicted copy count over normal copy count, the number of input genomic ranges contained within each segment, the number of targeted basepairs contained in the segments, and the name of the sample to help compile the segments across samples. The columns `log.odds` and `nranges` can be useful for filtering, see `?copyCountSegments` for details.

```
> copyCountSegments(fit)
```

```
RangedData with 3 rows and 5 value columns across 1 space
  space      ranges | copy.count log.odds  nranges
  <factor>   <IRanges> | <integer> <numeric> <numeric>
1   chr1 [ 861322, 1421615] |      2     0.00     300
2   chr1 [1423243, 2125213] |      1    768.87     100
3   chr1 [2125214, 7838229] |      2     0.00     600
  targeted.bp sample.name
  <integer> <character>
1      34275    HG00731
2      11524    HG00731
3      69553    HG00731
```

Calling `plot` on the *ExomeCopy* object draws segments of constant predicted copy number as colored horizontal lines and normalized read counts as points. If no colors are provided, the function will use red for copy counts less than `d` and blue for copy counts higher than `d`.

```
> cnv.cols <- c("red", "orange", "black", "deepskyblue",
+              "blue", "blue2", "blue4")
> plot(fit, col = cnv.cols)
```



5.2 Looping exomeCopy over multiple samples and chromosomes

In order to apply *exomeCopy* to a full dataset of multiple chromosomes and samples, we write a wrapper function `runExomeCopy`. We will only model duplications up to 4 copies. The wrapper construction allows for distribution of fitting each samples across workstations. Keep in mind that the non-PAR region of chrX, in male samples should be fit separately, with `d = 1`.

```
> runExomeCopy <- function(sample.name, seqs) {
+   lapply(seqs, function(seq.name) exomeCopy(rdata[seq.name],
+       sample.name, X.names = c("bg", "GC", "GC.sq",
+       "width"), S = 0:4, d = 2))
+ }
```

We can now run `exomeCopy` using either `lapply` or using a function like `clusterApplyLB` from the `snow` package. We apply `runExomeCopy` for 8 samples over 2 chromosomes and the result is stored as a named list of named lists.

```
> seqs <- c("chr1", "chr1a")
> names(seqs) <- seqs
> samples <- exome.samples[1:8]
> names(samples) <- samples
> t0 <- as.numeric(proc.time()[3])
> fit.list <- lapply(samples, runExomeCopy, seqs)
> t1 <- as.numeric(proc.time()[3])
> time.elapsed <- as.numeric(t1 - t0)
> paste(round(time.elapsed), "seconds for", length(samples),
+   "samples,", round(sum(width(rdata))/1000), "kb of target")

[1] "25 seconds for 8 samples, 231 kb of target"

> paste("~", round(time.elapsed/60/(8 * sum(width(rdata)))) *
+   3.2e+07, 1), " minutes for 1 sample, 32 Mb of target",
+   sep = "")

[1] "~7.3 minutes for 1 sample, 32 Mb of target"
```

5.3 Compiling results, filtering and plotting CNVs

Using `compileCopyCountSegments` on the list of lists, we can compile the segments across samples. We remove the segments with copy count 2, leaving the predicted CNVs for autosomal ranges.

```
> compiled.segments <- compileCopyCountSegments(fit.list)
> CNV.segments <- compiled.segments[compiled.segments$copy.count !=
+   2, ]
```

Some non-CNV ranges with too high or low read counts might be predicted to be CNV, despite the use of positional covariates and overdispersion in the emission distributions of the HMM. Often the majority of these cases can be removed by requiring a minimum number of ranges (for example requiring the `nranges` column returned by `copyCountSegments` to be greater than 5). The filtered CNV segments can be exported to tab-delimited files using `as.data.frame` and then `write.table`.

`copyCountSegments` also provides a column of log odds for each segment. This is the log ratio of the probability density function for the fitted predicted distribution over the fitted normal state distribution, then summed over all ranges in a segment (details of these are given in the help for `copyCountSegments`.) By definition these will be zero for the segments predicted normal and positive for the CNV segments.

```
> CNV.segments[1:6, ]
```

```
RangedData with 6 rows and 5 value columns across 2 spaces
  space      ranges | copy.count log.odds  nranges
<factor> <IRanges> | <integer> <numeric> <numeric>
1   chr1 [ 986106, 1249745] |      0  1200.92     100
2   chr1 [4715486, 4772259] |      3   12.22       4
3   chr1 [1250204, 1421615] |      0  1121.99     100
```

```

4   chr1 [1635263, 1650894] |      3    21.08      5
5   chr1 [1423243, 2125213] |      1    766.84    100
6   chr1 [2125214, 2939075] |      1    864.79    100

```

```

targeted.bp sample.name
  <integer> <character>
1      11484    HG00551
2       329    HG00551
3      10784    HG00641
4       608    HG00641
5      11524    HG00731
6      10717    HG00732

```

```

> CNV.segments <- CNV.segments[CNV.segments$nranges > 5,
+ ]


```

plotCompiledCNV will plot the CNVs for a given sequence across samples if given an object returned by compileCopyCountSegments and the sequence name.

```

> cnv.cols <- c("red", "orange", "black", "deepskyblue",
+ "blue")
> plotCompiledCNV(CNV.segments = CNV.segments, seq.name = "chr1",
+ col = cnv.cols)


```



```

> plotCompiledCNV(CNV.segments = CNV.segments, seq.name = "chr1a",
+ col = cnv.cols)

```



5.4 Inspecting model parameters

Finally, we can inspect a number of fitted parameters in the *ExomeCopy* object, such as the β vector which is adjusted in optimizing the likelihood of the HMM. The β vector is initialized using a linear regression of counts on the covariates, and this and other initial parameter values can be accessed in the `init.par` slot of the *ExomeCopy* object. These coefficients are adjusted while optimizing the likelihood of the HMM. The final β vector and other final parameter values are accessible in the `final.par` slot.

```
> fit.list[[1]][[1]]@init.par$beta.hat
```

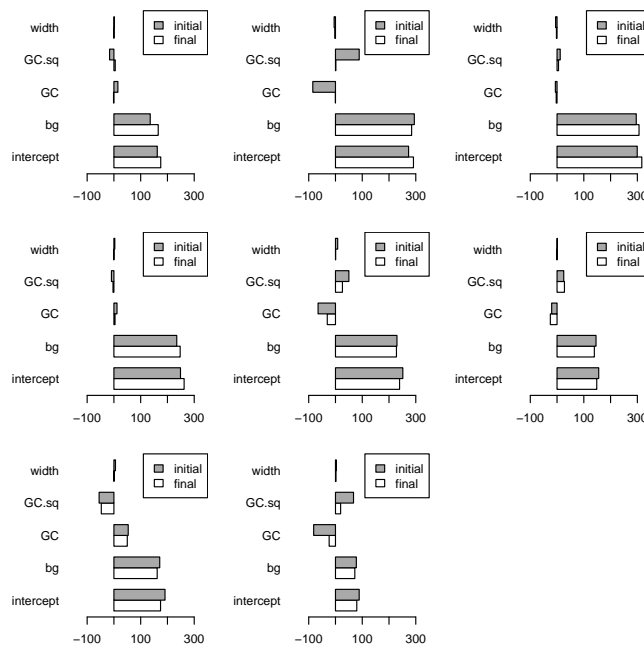
intercept	bg	GC	GC.sq	width
161.980000	135.739695	14.853306	-16.206500	1.766409

```
> fit.list[[1]][[1]]@final.par$beta
```

intercept	bg	GC	GC.sq	width
174.9950817	165.5299325	-1.4539701	4.9982643	-0.2111414

If we plot the initial and final β vectors from 8 samples on chr1, we can observe that the β vector varies across patients and that the coefficients which maximize the full HMM are not the same as the initial estimates from regression.

```
> par(mfrow = c(3, 3), mar = c(3, 5, 1, 1))
> for (i in 1:8) {
+   barplot(rbind(fit.list[[i]][[1]]@final.par$beta,
+                 fit.list[[i]][[1]]@init.par$beta.hat), horiz = TRUE,
+           las = 1, beside = TRUE, col = c("white", "darkgrey"),
+           xlim = c(-150, 350))
+   legend("topright", legend = c("initial", "final"),
+         fill = c("darkgrey", "white"))
+ }
```



6 Session info

```
> sessionInfo()
```

```
R version 2.15.0 RC (2012-03-22 r58802)
```

```
Platform: i386-apple-darwin9.8.0/i386 (32-bit)
```

```
locale:
```

```
[1] C
```

```
attached base packages:
```

```
[1] stats      graphics  grDevices  utils      datasets  methods
```

```
[7] base
```

```
other attached packages:
```

```
[1] exomeCopy_1.2.0    Rsamtools_1.8.0    Biostrings_2.24.0
```

```
[4] GenomicRanges_1.8.0 IRanges_1.14.0     BiocGenerics_0.2.0
```

```
loaded via a namespace (and not attached):
```

```
[1] bitops_1.0-4.1 stats4_2.15.0 tools_2.15.0    zlibbioc_1.2.0
```

References

1000 Genomes Project Consortium. A map of human genome variation from population-scale sequencing. *Nature*, 467(7319):1061–1073, October 2010. ISSN 1476-4687. doi: 10.1038/nature09534. URL <http://dx.doi.org/10.1038/nature09534>.

Michael I. Love, Alena Mysířková, Ruping Sun, Vera Kalscheuer, Martin Vingron, and Stefan A. Haas. Modeling Read Counts for CNV Detection in Exome Sequencing Data. *Statistical Applications in Genetics and Molecular Biology*, 10(1), November 2011. ISSN 1544-6115. doi: 10.2202/1544-6115.1732. URL http://cmb.molgen.mpg.de/publications/Love_2011_exomeCopy.pdf.

Kim D. Pruitt, Jennifer Harrow, Rachel A. Harte, Craig Wallin, Mark Diekhans, Donna R. Maglott, Steve Searle, Catherine M. Farrell, Jane E. Loveland, Barbara J. Ruef, Elizabeth Hart, Marie-Marthe M. Suner, Melissa J. Landrum, Bronwen Aken, Sarah Ayling, Robert Baertsch, Julio Fernandez-Banet, Joshua L. Cherry, Val Curwen, Michael Dicuccio, Manolis Kellis, Jennifer Lee, Michael F. Lin, Michael Schuster, Andrew Shkeda, Clara Amid, Garth Brown, Oksana Dukhanina, Adam Frankish, Jennifer Hart, Bonnie L. Maidak, Jonathan Mudge, Michael R. Murphy, Terence Murphy, Jeena Rajan, Bhanu Rajput, Lillian D. Riddick, Catherine Snow, Charles Steward, David Webb, Janet A. Weber, Laurens Wilming, Wenyu Wu, Ewan Birney, David Haussler, Tim Hubbard, James Ostell, Richard Durbin, and David Lipman. The consensus coding sequence (CCDS) project: Identifying a common protein-coding gene set for the human and mouse genomes. *Genome research*, 19(7):1316–1323, July 2009. ISSN 1088-9051. doi: 10.1101/gr.080531.108. URL <http://dx.doi.org/10.1101/gr.080531.108>.