

VanillaICE: Hidden Markov Models for the Assessment of Chromosomal Alterations using High-throughput SNP Arrays

Robert Scharpf

May 21, 2014

Abstract

This package provides an implementation of a hidden Markov Model for high throughput SNP arrays. Users of this package should already have available locus-level estimates of copy number and B allele frequencies or genotype calls. Copy number estimates can be relative or absolute.

1 Overview

This vignette requires that you have

- an absolute estimate of the *total* copy number organized such that rows correspond to loci and columns correspond to samples
and / or
- a matrix of B allele frequencies or genotype calls (1=AA, 2 = AB, 3= BB): rows correspond to loci and columns correspond to samples

If Affymetrix CEL files or Illumina Idat files were processed with the R package *crlmm*, see the vignette `crlmmDownstream` included with this package.

Other HMM implementations for the joint analysis of copy number and genotype include QuantiSNP [1] and PennCNV [3].

Data considerations. The HMM implemented in this package is most relevant for heritable diseases for which integer copy numbers are expected. For somatic cell diseases such as cancer, we suggest circular binary segmentation, as implemented in the R package DNACopy [2].

Citing this software. Robert B Scharpf, Giovanni Parmigiani, Jonathan Pevsner, and Ingo Ruczinski. Hidden Markov models for the assessment of chromosomal alterations using high-throughput SNP arrays. *Annals of Applied Statistics*, 2(2):687–713, 2008.

2 Organizing the locus-level data

This package includes simulated genotype and copy number data for 9165 SNPs.

```
> library(oligoClasses)
> library(VanillaICE)
> library2(SNPchip)
> library2(Biobase)
> library2(IRanges)
> data(locusLevelData)
```

The copy number estimates in the `locusLevelData` object were multiplied by 100 and saved as an integer. An object of class `oligoSnpSet` is instantiated from the simulated data in the next code chunk. Note that the assay data elements for the `oligoSnpSet` object should be integers. The `integerMatrix` scales the copy number matrix by a factor of 100 (by default) and returns a matrix of integers. When available, B allele frequencies should be scaled by a factor of 1000. For example,

```
integerMatrix(bAlleleFreq, scale=1000) \verb
> cn <- log2(locusLevelData[["copynumber"]]/100)
> oligoSet <- new("oligoSnpSet",
+               copyNumber=integerMatrix(cn),
+               call=locusLevelData[["genotypes"]],
+               callProbability=locusLevelData[["crlmmConfidence"]],
+               annotation=locusLevelData[["platform"]],
+               genome="hg19")
> oligoSet <- oligoSet[!is.na(chromosome(oligoSet)), ]
> oligoSet <- oligoSet[chromosome(oligoSet) < 3, ]
> oligoSet <- chromosomePositionOrder(oligoSet)
```

3 Fitting the HMM

When jointly modeling the copy number and genotypes/allele frequencies, we assume that the genotype and copy number estimates are independent conditional on the underlying hidden state. The method `hmm` is defined for several classes, including objects of class `oligoSnpSet` and `BeadStudioSet`.

The viterbi algorithm is used to obtain the most likely sequence of hidden states given the observed data. The value returned is an object of class `GRangesList` with genomic coordinates of the normal and altered regions. We also return the log-likelihood ratio (LLR) of the predicted sequence in an interval versus the likelihood of diploid copy number. For intervals inferred to have diploid copy number, the LLR is zero.

```
> fit.van <- hmm(oligoSet)
> fit.van <- fit.van[[1]]
```

Figure 1 plots the data for chromosome 1 and the predictions from the hidden markov model:

To find which markers are included in each genomic interval returned by the `hmm` method, one can use the `findOverlaps` method in the `oligoClasses`. For example, the second interval in the `RangedDataHMM` object `fit.van` contains 102 markers.

```
> fit.van[2, ]
```

GRanges with 1 range and 3 metadata columns:

seqnames	ranges	strand	numberProbes
<Rle>	<IRanges>	<Rle>	<integer>
NA06993 chr1	[49825223, 54637167]	*	102
	sample state		
	<character>	<integer>	
NA06993	NA06993	4	

seqlengths:			
	chr1	chr2	
	249250621	243199373	

To find the names of the 102 markers that are included in this genomic interval, one could do the following




Figure 1: Plot of artificial data for one chromosome.

```
> frange <- makeFeatureGRanges(oligoSet)
> markersInterval2 <- subsetByOverlaps(frange, fit.van[2, ])
> markersInterval2
```

GRanges with 102 ranges and 0 metadata columns:

	seqnames	ranges	strand
	<Rle>	<IRanges>	<Rle>
SNP_A-1753140	chr1	[49825223, 49825223]	*
SNP_A-1668144	chr1	[49829596, 49829596]	*
SNP_A-1754711	chr1	[49832529, 49832529]	*
SNP_A-1684573	chr1	[49928063, 49928063]	*
SNP_A-1686257	chr1	[49937160, 49937160]	*
...
SNP_A-1694924	chr1	[54517324, 54517324]	*
SNP_A-1662412	chr1	[54529526, 54529526]	*
SNP_A-1696791	chr1	[54542362, 54542362]	*
SNP_A-1645299	chr1	[54636977, 54636977]	*
SNP_A-1645411	chr1	[54637167, 54637167]	*

seqlengths:

chr1	chr2	chr3	...	chr22	chrX	chrY
249250621	243199373	198022430	...	51304566	155270560	59373566

Multipanel displays can be useful for visualizing the low-level data for copy number alterations. We extend the `xyplot` method in the R package `lattice` for two common use-cases: by-locus and by-sample.

By locus. To plot the genomic data for a set of ranges at a given locus, we provide an unevaluated code chunk below (our example contains only a single sample):

```
> xyplot2(cn ~ x | id, oligoSet, range=RangedDataObject, ylim=c(-0.5,4), panel=xypanel)
```

Note that the default in the above command is to use a common x- and y-scale for each panel. To allow the x-axes to change for each panel, one could set the x-scales to 'free':

```
> xyplot2(cn ~ x | id, oligoSet, range=RangedDataObject, ylim=c(-0.5,4), scales=list(x="free"),
+         panel=xypanel)
```

The function `xypanel` provides default colors for annotating the plotting symbols by genotype and by whether the markers are polymorphic. The `GenomicRanges` must be passed by the name `range` to the `xyplot2` method.

By sample. To plot the low-level data for multiple alterations occurring in a single sample, one can again pass a `GenomicRanges` object with name `range` to the `xyplot2`. For example, see Figure 2. The code for producing Figure 2 is in the following code chunk. Note that the formula in the example below conditions on `range` instead of `id`. The conditioning variable for displaying multiple panels of a single sample must be called 'range'. We plot a 2 Mb window surrounding each of the alterations in the simulated data by specifying `frame=2e6`.

```
> ranges.altered <- fit.van[!state(fit.van) %in% c(3,4) & numberProbes(fit.van) > 5, ]
> ##elementMetadata(ranges.altered)$sampleId <- sampleNames(oligoSet)
> xy.example <- xyplot2(cn ~ x | range, oligoSet, range=ranges.altered, frame=2e6,
+                       scales=list(x="free"), ylim=c(-1,3),
+                       panel=xypanel, cex.pch=0.4, pch=21, border="grey",
+                       ylab=expression(log[2]("copy number")))
```

See `?xyplot` for additional details.

Note also that `scales` must be set to `free` in the above call to `xyplot`.

3.1 ICE HMM

This will likely be phased out in favor of using B allele frequencies instead of genotype calls / call probabilities.

To compute emission probabilities that incorporate the `crImm` genotype confidence scores, we set ICE to TRUE. This option is limited too a few platforms and the Affy 100k platforms is not one of them.

```
> VanillaICE:::icePlatforms()
```

```
[1] "pd.genomewidesnp.6"
[2] "genomewidesnp6"
[3] "genomewidesnp6Crlmm"
[4] "pd.mapping250k.nsp"
[5] "pd.mapping250k.sty"
[6] "pd.mapping250k.nsp, pd.mapping250k.sty"
```

For illustration, we assign 'genomewidesnp6' to the annotation slot since this platform is supported.

```

> ann <- annotation(oligoSet)
> annotation(oligoSet) <- "genomewidesnp6"
> fit.ice <- hmm(oligoSet, ICE=TRUE)
> fit.ice <- fit.ice[[1]]
> annotation(oligoSet) <- ann

```

3.2 Other options


Regions of homozygosity. A HMM for genotype-only data can be used to find long stretches of homozygosity. Note that hemizygous deletions are also identified as 'ROH' when copy number is ignored (as the biallelic genotype call in a hemizygous deletions tends to be all homozygous calls).

```

> snpSet <- as(oligoSet, "SnpSet2")
> fit.gt <- hmm(snpSet, prGtHom=c(0.7, 0.99), normalIndex=1L, S=2L)
> fit.gt <- fit.gt[[1]]

```

A suggested visualization:



4 Trouble shooting

Read in previously saved data for chromosome 1:

```

> bset <- readRDS(file.path(system.file("extdata", package="VanillaICE"), "bset.rds"))

```

The hmm is fit independently to each arm. We can determine the arm for each marker in the `bset` object using the function `getArm`.

```

> library(oligoClasses)
> arm <- getArm(bset)
> table(arm)

```

```

arm
chr1p chr1q
12277 10756

```

We will fit the hmm to arm 'chr1p' and monitor the updates to the mean parameters for the copy number states by setting `verbose=TRUE`:

```

> b1p <- bset[arm=="chr1p", ]
> fit1p <- hmm(b1p, verbose=TRUE)

```

```

      A      AAAB      AAB      AB      ABB      ABBB      B
0.000000 0.100000 0.333333 0.500000 0.666667 0.900000 1.000000
$mus
      A      AAAB      AAB      AB      ABB      ABBB      B
0.000000 0.100000 0.333333 0.5414002 0.6554347 0.900000 1.000000

```

```

$sigmas
      A      AAAB      AAB      AB      ABB
0.018255734 0.046266239 0.046266239 0.046266239 0.046266239
      ABBB      B
0.046266239 0.009315174

```

```

$prOutlier
$prOutlier$prOutlier
[1] 0.001

```

```

$prOutlier$max
[1] 0.001

```

```

$prOutlier$maxROH
[1] 1e-05

```

```

[1] -2.0 -0.4 0.0 0.0 0.4 1.0
$mu
[1] -2.00000000 -0.40000000 -0.02014275 -0.02014275 0.19303931
[6] 1.00000000

```

```

$sigmas
[1] 0.2131821 0.1689636 0.1689636 0.1689636 0.1689636 0.1689636

```

```

$p
[1] 0.01523173

```

```

      A      AAAB      AAB      AB      ABB      ABBB      B
0.000000 0.100000 0.333333 0.5414002 0.6554347 0.900000 1.000000
$mus
      A      AAAB      AAB      AB      ABB      ABBB      B
0.000000 0.100000 0.333333 0.5420193 0.6666667 0.900000 1.000000

```

```

$sigmas
      A      AAAB      AAB      AB      ABB

```

0.017274996 0.047609088 0.047609088 0.047609088 0.047609088
 ABBB B
0.047609088 0.008797982

\$prOutlier
\$prOutlier\$prOutlier
[1] 0.0001629195

\$prOutlier\$max
[1] 0.001

\$prOutlier\$maxROH
[1] 1e-05

[1] -2.00000000 -0.40000000 -0.02014275 -0.02014275 0.19303931
[6] 1.00000000
\$mu
[1] -2.00000000 -0.40000000 -0.01452717 -0.01452717 0.25854534
[6] 1.00000000

\$sigmas
[1] 0.2131821 0.1591816 0.1591816 0.1591816 0.1591816 0.1591816

\$p
[1] 0.02940458

 A AAAB AAB AB ABB ABBB B
0.0000000 0.1000000 0.3333333 0.5420193 0.6666667 0.9000000 1.0000000
\$mus

 A AAAB AAB AB ABB ABBB B
0.0000000 0.1000000 0.3333333 0.5420725 0.6666667 0.9000000 1.0000000

\$sigmas
 A AAAB AAB AB ABB
0.017274769 0.047815352 0.047815352 0.047815352 0.047815352
 ABBB B
0.047815352 0.008660548

\$prOutlier
\$prOutlier\$prOutlier
[1] 0.0001629195

\$prOutlier\$max
[1] 0.001

\$prOutlier\$maxROH
[1] 1e-05

[1] -2.00000000 -0.40000000 -0.01452717 -0.01452717 0.25854534

```
[6] 1.00000000
$mu
[1] -2.00000000 -0.40000000 -0.01149589 -0.01149589 0.27206259
[6] 1.00000000
```

```
$sigmas
[1] 0.2131821 0.1559988 0.1559988 0.1559988 0.1559988 0.1559988
```

```
$p
[1] 0.0338845
```

```
      A      AAAB      AAB      AB      ABB      ABBB      B
0.0000000 0.1000000 0.3333333 0.5420725 0.6666667 0.9000000 1.0000000
$mus
```

```
      A      AAAB      AAB      AB      ABB      ABBB      B
0.0000000 0.1000000 0.3333333 0.5420728 0.6666667 0.9000000 1.0000000
```

```
$sigmas
      A      AAAB      AAB      AB      ABB
0.017274652 0.047816280 0.047816280 0.047816280 0.047816280
      ABBB      B
0.047816280 0.008659948
```

```
$prOutlier
$prOutlier$prOutlier
[1] 0.0001629195
```

```
$prOutlier$max
[1] 0.001
```

```
$prOutlier$maxROH
[1] 1e-05
```

```
[1] -2.00000000 -0.40000000 -0.01149589 -0.01149589 0.27206259
[6] 1.00000000
$mu
[1] -2.00000000 -0.40000000 -0.01020573 -0.01020573 0.27415391
[6] 1.00000000
```

```
$sigmas
[1] 0.2131821 0.1547808 0.1547808 0.1547808 0.1547808 0.1547808
```

```
$p
[1] 0.03575792
```

```
      A      AAAB      AAB      AB      ABB      ABBB      B
0.0000000 0.1000000 0.3333333 0.5420728 0.6666667 0.9000000 1.0000000
$mus
```

```
      A      AAAB      AAB      AB      ABB      ABBB      B
0.0000000 0.1000000 0.3333333 0.5420728 0.6666667 0.9000000 1.0000000
```



```

$sigmas
      A      AAAB      AAB      AB      ABB
0.017274601 0.047816277 0.047816277 0.047816277 0.047816277
      ABBB      B
0.047816277 0.008659946

$prOutlier
$prOutlier$prOutlier
[1] 0.0001629195

$prOutlier$max
[1] 0.001

$prOutlier$maxROH
[1] 1e-05

[1] -2.00000000 -0.40000000 -0.01020573 -0.01020573 0.27415391
[6] 1.00000000
$mu
[1] -2.000000000 -0.400000000 -0.009680834 -0.009680834 0.274530298
[6] 1.000000000

$sigmas
[1] 0.2131821 0.1543055 0.1543055 0.1543055 0.1543055 0.1543055

$p
[1] 0.03575792



```

Each line beginning with the header "A AAAB " is one iteration of the Viterbi algorithm. The numeric vector below the header is the mean for the different genotype states. Note that we use the same mean for homozygous genotypes (mean of A = mean of AA = mean of AAA). The next line is the mean of the log R ratios for copy number states homozygous deletion, hemizygous deletion, diploid, copy-neutral homozygosity, single copy gain, and two-copy gain, respectively. We can plot the data and see whether the means in the last iteration are consistent with what we observe.

```

> par(mfrow=c(2,1), las=1)
> x <- position(b1p)/1e6
> plot(x, lrr(b1p)/100, pch=".", col="grey")
> abline(h=-0.07547)
> plot(x, baf(b1p)/1000, pch=".", col="grey")
> abline(h=0.542)

```



```
> b1q <- bset[arm=="chr1q", ]
> fit1q <- hmm(b1q, verbose=TRUE)
```

```

      A      AAAB      AAB      AB      ABB      ABBB      B
0.000000 0.100000 0.333333 0.500000 0.666667 0.900000 1.000000
$mus
```

```

      A      AAAB      AAB      AB      ABB      ABBB      B
0.000000 0.100000 0.333333 0.485602 0.666667 0.900000 1.000000
```

```
$sigmas
```

```

      A      AAAB      AAB      AB      ABB      ABBB      B
0.1540716 0.2507250 0.2507250 0.2507250 0.2507250 0.2507250 0.1151923
```

```
$prOutlier
```

```
$prOutlier$prOutlier
[1] 0
```

```
$prOutlier$max
```

```
[1] 0.001
```

```
$prOutlier$maxROH
```

[1] 1e-05

[1] -2.0 -0.4 0.0 0.0 0.4 1.0

\$mu

[1] -2.000000000 -0.400000000 0.008540304 0.008540304 0.400000000

[6] 1.000000000

\$sigmas

[1] 0.1783766 0.1502543 0.1502543 0.1502543 0.1502543 0.1502543

\$p

[1] 0.006600967

	A	AAAB	AAB	AB	ABB	ABBB	B
0.0000000	0.1000000	0.3333333	0.4856020	0.6666667	0.9000000	1.0000000	

\$mus

	A	AAAB	AAB	AB	ABB	ABBB	B
0.0000000	0.1000000	0.3333333	0.5214707	0.6666667	0.9000000	1.0000000	

\$sigmas

	A	AAAB	AAB	AB	ABB	ABBB	B
0.1784226	0.2752851	0.2752851	0.2752851	0.2752851	0.2752851	0.1136589	

\$prOutlier

\$prOutlier\$prOutlier

[1] 0

\$prOutlier\$max

[1] 0.001

\$prOutlier\$maxROH

[1] 1e-05

[1] -2.000000000 -0.400000000 0.008540304 0.008540304 0.400000000

[6] 1.000000000

\$mu

[1] -2.000000000 -0.400000000 0.009844782 0.009844782 0.400000000

[6] 1.000000000

\$sigmas

[1] 0.1783766 0.1427471 0.1427471 0.1427471 0.1427471 0.1427471

\$p

[1] 0.01217925

	A	AAAB	AAB	AB	ABB	ABBB	B
0.0000000	0.1000000	0.3333333	0.5214707	0.6666667	0.9000000	1.0000000	

\$mus

	A	AAAB	AAB	AB	ABB	ABBB	B

```
0.0000000 0.1000000 0.3333333 0.5214707 0.6666667 0.9000000 1.0000000
```

```
$sigmas
```

```
      A      AAAB      AAB      AB      ABB      ABBB      B  
0.1885729 0.2818259 0.2818259 0.2818259 0.2818259 0.2818259 0.1098109
```

```
$prOutlier
```

```
$prOutlier$prOutlier
```

```
[1] 0
```

```
$prOutlier$max
```

```
[1] 0.001
```

```
$prOutlier$maxROH
```

```
[1] 1e-05
```

```
[1] -2.000000000 -0.400000000 0.009844782 0.009844782 0.400000000
```

```
[6] 1.000000000
```

```
$mu
```

```
[1] -2.000000000 -0.400000000 0.009999987 0.009999987 0.400000000
```

```
[6] 1.000000000
```

```
$sigmas
```

```
[1] 0.1783766 0.1404938 0.1404938 0.1404938 0.1404938 0.1404938
```

```
$p
```

```
[1] 0.0158981
```

```
> par(mfrow=c(2,1), las=1)
```



```
> x <- position(b1q)/1e6
```

```
> plot(x, lrr(b1q)/100, pch=".", col="grey")
```

```
> abline(h=c(0, .359778))
```

```
> plot(x, baf(b1q)/1000, pch=".", col="grey")
```

```
> abline(h=c(0.235, 0.790))
```



5 Session Information

The version number of R and packages loaded for generating the vignette were:

- R version 3.1.0 (2014-04-10), x86_64-unknown-linux-gnu
- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_US.UTF-8, LC_COLLATE=C, LC_MONETARY=en_US.UTF-8, LC_MESSAGES=en_US.UTF-8, LC_PAPER=en_US.UTF-8, LC_NAME=C, LC_ADDRESS=C, LC_TELEPHONE=C, LC_MEASUREMENT=en_US.UTF-8, LC_IDENTIFICATION=C
- Base packages: base, datasets, grDevices, graphics, methods, parallel, stats, utils
- Other packages: Biobase 2.24.0, BiocGenerics 0.10.0, Biostrings 2.32.0, DBI 0.2-7, GenomeInfoDb 1.0.2, GenomicRanges 1.16.3, IRanges 1.22.7, RColorBrewer 1.0-5, RSQLite 0.11.4, SNPchip 2.10.0, VanillaICE 1.26.1, XVector 0.4.0, matrixStats 0.8.14, oligo 1.28.2, oligoClasses 1.26.0, pd.mapping50k.hind240 1.10.0, pd.mapping50k.xba240 1.10.0
- Loaded via a namespace (and not attached): BiocInstaller 1.14.2, Matrix 1.1-3, R.methodsS3 1.6.1, affxparser 1.36.0, affyio 1.32.0, bit 1.1-12, codetools 0.2-8, compiler 3.1.0, expm 0.99-1.1, ff 2.2-13, foreach 1.4.2, grid 3.1.0, iterators 1.0.7, lattice 0.20-29, msm 1.3, mvtnorm 0.9-99992, preprocessCore 1.26.1, splines 3.1.0, stats4 3.1.0, survival 2.37-7, tools 3.1.0, zlibbioc 1.10.0

References

- [1] Stefano Colella, Christopher Yau, Jennifer M Taylor, Ghazala Mirza, Helen Butler, Penny Clouston, Anne S Bassett, Anneke Seller, Christopher C Holmes, and Jiannis Ragoussis. QuantiSNP: an Objective Bayes Hidden-Markov Model to detect and accurately map copy number variation using SNP genotyping data. *Nucleic Acids Res*, 35(6):2013–2025, 2007.
- [2] Adam B Olshen, E. S. Venkatraman, Robert Lucito, and Michael Wigler. Circular binary segmentation for the analysis of array-based DNA copy number data. *Biostatistics*, 5(4):557–72, Oct 2004.
- [3] Kai Wang, Mingyao Li, Dexter Hadley, Rui Liu, Joseph Glessner, Struan F A Grant, Hakon Hakonarson, and Maja Bucan. Penncnv: an integrated hidden markov model designed for high-resolution copy number variation detection in whole-genome snp genotyping data. *Genome Res*, 17(11):1665–1674, Nov 2007.




Figure 2: The method `xyp1ot` is used to create a multi-panel display of alterations in a single sample. Each panel displays a single copy number alteration detected by the HMM that is boxed by a rectangle. The alteration is framed by specifying the number of basepairs to plot upstream and downstream of the alteration. Here, we used a frame of 2 Mb. Homozygous SNPs with diallelic genotypes ‘AA’ and ‘BB’ are shaded blue; SNPs with diallelic genotype call ‘AB’ are shaded in red.