

# AllelicImbalance Vignette

*Jesper R. Gådin and Lasse Folkersen*

2015-05-17

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>ASEset</b>	<b>2</b>
2.1	Simple example of building an ASEset object . . . . .	2
2.2	Building an ASEset object using Bcf or Vcf files . . . . .	2
2.3	Using strand information . . . . .	3
2.4	Two useful helper functions . . . . .	4
2.5	Adding phenotype data . . . . .	4
2.6	Adding phase information . . . . .	5
2.7	Adding reference and alternative allele information . . . . .	5
<b>3</b>	<b>Tests</b>	<b>6</b>
3.1	Statistical analysis of an ASEset object . . . . .	7
<b>4</b>	<b>Base graphics</b>	<b>7</b>
4.1	Plotting of an ASEset object . . . . .	7
4.2	Plot with annotation . . . . .	10
4.3	locationplot . . . . .	11
4.4	Top allele criteria . . . . .	12
<b>5</b>	<b>Grid graphics</b>	<b>13</b>
5.1	gbarplot . . . . .	14
5.2	glocationplot . . . . .	15
5.3	Custom location plots . . . . .	17
<b>6</b>	<b>Summary functions</b>	<b>18</b>
<b>7</b>	<b>Conclusion</b>	<b>19</b>
<b>8</b>	<b>Links</b>	<b>19</b>
<b>9</b>	<b>Session Info</b>	<b>19</b>

## 1 Introduction

---

This *AllelicImbalance* package contains functions for investigating allelic imbalance effects in RNA-seq data. Maternal and paternal alleles could be expected to show identical transcription rate, resulting in a 50%-50% mix of maternal and paternal mRNA in a sample. However, this turns out to sometimes not be the case. The most extreme example is the X-chromosome inactivation in females, but many autosomal genes also have deviations from identical transcription rate. The causes of this are not always known, but one likely cause is the difference in DNA, namely heterozygous SNPs, affecting enhancers, promoter regions, splicing and stability. Identifying this allelic imbalance is therefore of interest to the characterization of the genome and the aim of the *AllelicImbalance* package is to facilitate this.

Load *AllelicImbalance*

```
library(AllelicImbalance)
```

## 2 ASEset

---

The ASEset object is the central class of objects in the AllelicImbalance package. The ASEset object has the SummarizedExperiment from the [GenomicRanges](#) package as parent class, and all functions you can apply on this class you can also apply on an ASEset.

### 2.1 Simple example of building an ASEset object

In this section we will walk through the various ways an ASEset object can be created. Although the preprocessing of RNA-seq data is not the primary focus of this package, it is a necessary step before analysis. There exists several different methods for obtaining a bam file, and this section should just be considered an example. For further details we refer to the web-pages of tophat, bowtie, bwa and samtools found in the links section at the end of this document.

```
wget ftp://ftp.sra.ebi.ac.uk/vol1/fastq/ERR009/ERR009135/*
bowtie -q --best --threads 5 --sam hg19 +
> -1 ERR009135_1.fastq.gz -2 ERR009135_2.fastq.gz "ERR009135.sam"
samtools view -S -b ERR009135.sam > ERR009135.bam
```

In the above code one paired-end RNA sequencing sample is downloaded and aligned to the human genome, then converted to bam using samtools. The resulting bam files can be the direct input to the AllelicImbalance package. Other aligners can be used as well, as long as bam files are provided as input. The example code following illustrates how to use the import mechanism on a chromosome 17-located subset of 20 RNA-seq experiments of HapMap samples. The output is an ASEset object containing allele counts for all heterozygote coding SNPs in the region.

```
searchArea <- GRanges(seqnames = c("17"), ranges = IRanges(79478301, 79478361))
pathToFiles <- system.file("extdata/ERP000101_subset", package = "AllelicImbalance")
reads <- impBamGAL(pathToFiles, searchArea, verbose = FALSE)
heterozygotePositions <- scanForHeterozygotes(reads, verbose = FALSE)
countList <- getAlleleCounts(reads, heterozygotePositions, verbose = FALSE)
a.simple <- ASEsetFromCountList(heterozygotePositions, countList)
a.simple
```

```
## class: ASEset
## dim: 3 20
## exptData(0):
## assays(3): countsUnknown mapBias phase
## rownames(3): chr17_79478287 chr17_79478331 chr17_79478334
## rowRanges metadata column names(0):
## colnames(20): ERR009097.bam ERR009102.bam ... ERR009160.bam
## ERR009167.bam
## colData names(0):
```

### 2.2 Building an ASEset object using Bcf or Vcf files

If more than a few genes and a few samples are analyzed we recommend that a SNP-call is instead made using the samtools mpileup function (see links section). The scanForHeterozygotes function is merely a simple SNP-caller and it is not as computationally optimized as e.g. mpileup. In this bash code we download reference sequence for chromosome 17 and show how to generate mpileup calls on one of the HapMap samples that were described above.

```
#download the reference chromosome in fasta format
wget ftp://hgdownload.cse.ucsc.edu/goldenPath/hg19/chromosomes/chr17.fa.gz
```

```
#Vcf format
samtools mpileup -uf hg19.fa ERR009135.bam > ERR009135.vcf
```

```
#Bcf format
samtools mpileup -uf hg19.fa ERR009135.bam | bcftools view -bvcg - > ERR009135.bcf
```

Samtools mpileup generates by default a Vcf file which contains SNP and short INDEL positions. Piping the output to bcftools we get its binary equivalent (Bcf), which takes less space and can be queried more effectively.

The Vcf file is a text file that stores information about positions in the genome. In addition to the location, stored information could for example be genotype, phase, reference and alternative alleles for a collection of samples. More detailed information can be found following this [link](#).

In the *VariantAnnotation* package there is a lot of useful tools to handle Vcf files. The `readVcf` function reads Vcf data into R, which can be subset to only ranges by `granges` to get a `GRanges` object that is the object type required by the `getAlleleCounts` function.

```
pathToVcf <- paste(pathToFiles, "/ERP000101.vcf", sep="")
VCF <- readVcf(pathToVcf, "hg19")
gr <- granges(VCF)

#only use bi-allelic positions
gr.filt <- gr[width(mcols(gr)[,"REF"]) == 1 |
             unlist(lapply(mcols(gr)[,"ALT"], width)) == 1]

countList <- getAlleleCounts(reads, gr.filt, verbose=FALSE)
a.vcf <- ASEsetFromCountList(rowRanges = gr, countList)
```

With the Bcf files the process of generating an ASEset object starts with a call to the `impBcfGR` function instead. This function will import the Bcf file containing all SNP calls that were generated with the `samtools mpileup` function.

```
BcfGR <- impBcfGR(pathToFiles, searchArea, verbose=FALSE)
countListBcf <- getAlleleCounts(reads, BcfGR, verbose=FALSE)
a.bcf <- ASEsetFromCountList(BcfGR, countListBcf)
```

## 2.3 Using strand information

Many RNA-seq experiments do not yield useful information on the strand from which a given read was made. This is because they involve a step in which a double-stranded cDNA is created without tracking strand-information. Some RNA-seq setups do however give this information and in those cases it is important to keep track of strand in the ASE-experiment. The example data from above is from an experiment which created double-stranded cDNA before labelling and so the '+' and '-' information in it is arbitrary. However, if we assume that the information has strand information, then the correct procedure is as follows:

```
plus <- getAlleleCounts(reads, heterozygotePositions, strand="+", verbose=F)
minus <- getAlleleCounts(reads, heterozygotePositions, strand="-", verbose=F)

a.stranded <-
ASEsetFromCountList(
heterozygotePositions,
countListPlus=plus,
countListMinus=minus)
```

```
)
a.stranded

## class: ASEset
## dim: 3 20
## exptData(0):
## assays(5): countsPlus countsMinus countsUnknown mapBias phase
## rownames(3): chr17_79478287 chr17_79478331 chr17_79478334
## rowRanges metadata column names(0):
## colnames(20): ERR009097.bam ERR009102.bam ... ERR009160.bam
##   ERR009167.bam
## colData names(0):
```

The main effect of doing this, is in the plotting functions which will separate reads from different strands if they are specified as done here. It is important, however, to make sure that the imported RNA-seq experiment does in fact have proper labeling and tracking of strand information before proceeding with this method.

## 2.4 Two useful helper functions

At this stage it is worth highlighting two useful helper functions that both uses existing BioC annotation objects. One is the `getAreaFromGeneNames` which quickly retrieves the above mentioned `searchArea` when given just genesymbols as input, and relies on [org.Hs.eg.db](#). The other other is the `getSnpldFromLocation` function which attempts to rename location-based SNP names to established rs-IDs in case they exist. These functions work as follows:

```
#Getting searchArea from genesymbol
library(org.Hs.eg.db )
searchArea<-getAreaFromGeneNames("ACTG1",org.Hs.eg.db)

#Getting rs-IDs
library(SNPlocs.Hsapiens.dbSNP.20120608)
gr <- rowRanges(a.simple)
updatedGRanges<-getSnpldFromLocation(gr, SNPlocs.Hsapiens.dbSNP.20120608)
rowRanges(a.simple)<-updatedGRanges
```

## 2.5 Adding phenotype data

Typically an RNA-seq experiment will include additional information about each sample. It is an advantage to include this information when creating an ASEset because it can be used for subsequent highlights or subsetting in plotting and analysis functions.

```
#simulate phenotype data
pdata <- DataFrame(
  Treatment=sample(c("ChIP", "Input"),length(reads),replace=TRUE),
  Gender=sample(c("male", "female"),length(reads),replace=TRUE),
  row.names=paste("individual",1:length(reads),sep=""))

#make new ASEset with pdata
a.new <- ASEsetFromCountList(
  heterozygotePositions,
  countList,
  colData=pdata)

#add to existing object
colData(a.simple) <- pdata
```

## 2.6 Adding phase information

For some functionality phase information is necessary. Phasing can be obtained from many external sources e.g. samtools. The phase information is often present in VCF-files. The lines below show how to access that information and transfer it to an ASEset. The ASEset follows the VCF-conventions on how to describe the phase, i.e. each patients phase will be described by the established notation of the form “1|0”, “1|1” or “1/0”. There the left number is the description of the maternal allele and the right number is the description of the paternal allele. If it is 0 the allele is the same as the reference allele, and if it is 1 it is the alternative allele. For “|” the phase is known and for “/” the phase is not known. Note, that in the AllelicImbalance package only bi-allelic expression is allowed.

The phase can be manually added by constructing a user-generated matrix, or transforming the data into a matrix from an external source. The most convenient way of importing phase information is probably by reading it from a Vcf file, which is commonly used to store phase information. The readGT function from the [VariantAnnotation](#) package will return a matrix ready to just attach to an ASEset object.

```
#construct an example phase matrix
set.seed(1)
rows <- nrow(a.simple)
cols <- ncol(a.simple)
p1 <- matrix(sample(c(1,0),replace=TRUE, size=rows*cols), nrow=rows, ncol=cols)
p2 <- matrix(sample(c(1,0),replace=TRUE, size=rows*cols), nrow=rows, ncol=cols)
phase.mat <- matrix(paste(p1,sample(c("|","|","/"), size=rows*cols, replace=TRUE), p2, sep=""),
  nrow=rows, ncol=cols)
phase(a.simple) <- phase.mat

#use the phase information from a Vcf file
pathToVcf <- system.file("extdata/ERP000101_subset/ERP000101.vcf", package = "AllelicImbalance")
p <- readGT(pathToVcf)
#The example Vcf file contains only 19 out of our 20 samples
#So we have to subset and order
a.subset <- a.simple[,colnames(a.simple) %in% colnames(p)]
p.subset <- p[, colnames(p) %in% colnames(a.subset)]
p.ordered <- p.subset[, match(colnames(a.subset), colnames(p.subset))]
```

## 2.7 Adding reference and alternative allele information

Having the information of reference and alternative allele is important to investigate any presence of mapping bias. It is also important to be able to use phasing information. The reference and alternative alleles are stored in the meta-columns and can be accessed and set through the mcols() function. All functions that require reference or alternative allele will visit the meta-columns “ref” or “alt” to extract that information.

```
#from simulated data
ref(a.simple) <- c("G", "T", "C")

#infer and set alternative allele
alt <- inferAltAllele(a.simple)
alt(a.simple) <- alt

#from reference genome
data(ASEset.sim)
fasta <- system.file('extdata/hg19.chr17.subset.fa', package='AllelicImbalance')
```

```
ref <- refAllele(ASEset.sim,fasta=fasta)
ref(ASEset.sim) <- ref
```

Using reference allele information when measuring the impact of mapping bias globally, can be done by measuring the average reference allele fraction from a representative set of SNPs. Below the simulated example uses 1000 SNPs to assess any presence of mapping bias. Any deviations from 0.5 suggests a bias favouring one or the other allele. The most likely outcome is a value higher than 0.5 and is probably due to mapping bias, i.e., the reference allele actually has mapped more often than the alternative allele.

```
#make an example countList including a global sample of SNPs
set.seed(1)
countListUnknown <- list()
samples <- paste(rep("sample",10),1:10,sep="")
snps <- 1000
for(snp in 1:snps){
  count<-matrix(0, nrow=length(samples), ncol=4, dimnames=list(samples, c('A','T','G','C')))
  alleles <- sample(1:4, 2)
  for(sample in 1:length(samples)){
    count[sample, alleles] <- as.integer(rnorm(2,mean=50,sd=10))
  }
  countListUnknown[[snp]] <- count
}

#make example rowRanges for the corresponding information
gr <- GRanges(
  seqnames = Rle(sample(paste("chr",1:22,sep=""),snps, replace=TRUE)),
  ranges = IRanges(1:snps, width = 1, names= paste("snp",1:snps,sep="")),
  strand="*"
)

#make ASEset
a <- ASEsetFromCountList(gr, countListUnknown=countListUnknown)

#set a random allele as reference
genotype(a) <- inferGenotypes(a)
ref(a) <- randomRef(a)

#get the fraction of the reference allele
refFrac <- fraction(a, top.fraction.criteria="ref")

#check mean
mean(refFrac)
```

```
## [1] 0.5013641
```

The reference fraction mean `mean(refFrac)` is in this case very close to 0.5, and suggests that the mapbias globally is low if present. In this example we randomly assigned the reference genome to one of the two most expressed allele by the `randomRef` function, so it should not be any reference allele bias (mapbias).

### 3 Tests

---

### 3.1 Statistical analysis of an ASEset object

One of the simplest statistical test for use in allelic imbalance analysis is the chi-square test. This test assumes that the uncertainty of ASE is represented by a normal distribution around an expected mean (i.e 0.5 for equal expression). A significant result suggests an ASE event. Every strand is tested independently.

```
#use a subset for tests
```

```
a2 <- a.stranded[,5:8]
```

```
#binomial test
```

```
binom.test(a2, "+")
```

```
##      chr17_79478287 chr17_79478331 chr17_79478334
## [1,]              NA              NA              NA
## [2,]      0.7265625      0.006610751      0.2668457
## [3,]              NA      0.107752144              NA
## [4,]              NA              NA              NA
```

```
#chi-square test
```

```
chisq.test(a2, "-")
```

```
##      chr17_79478287 chr17_79478331 chr17_79478334
## [1,]              NA              NA              NA
## [2,]      0.7962534              NA              NA
## [3,]              NA              NA              NA
## [4,]              NA              NA              NA
```


## 4 Base graphics

---

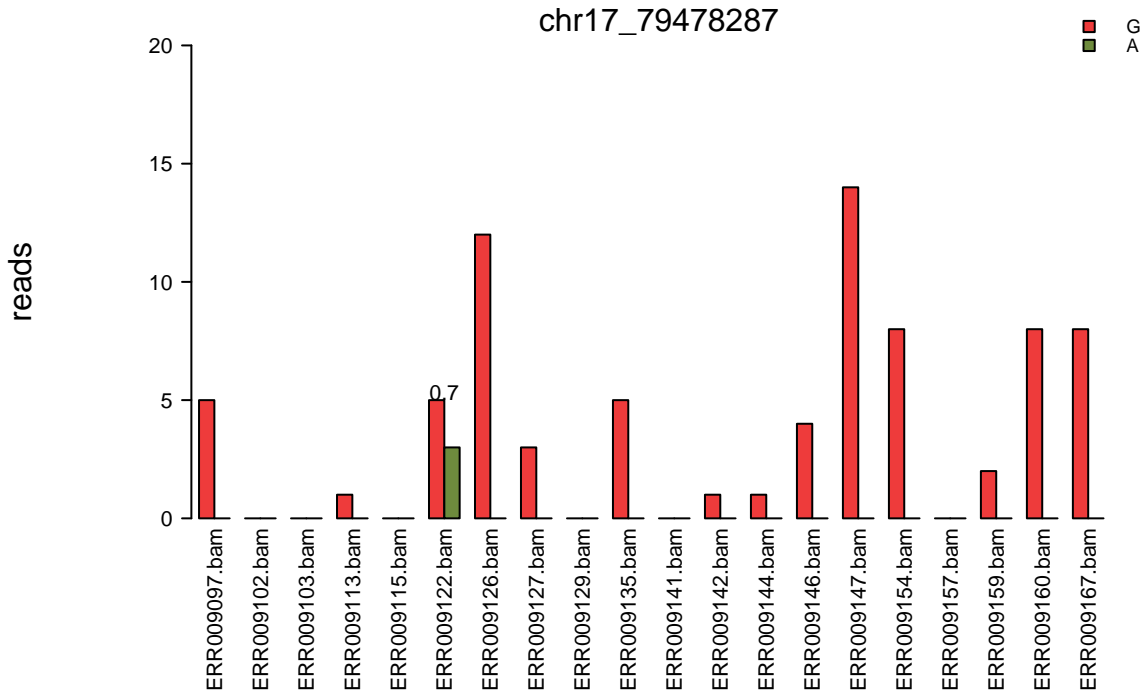
### 4.1 Plotting of an ASEset object

The barplot function for ASEset objects plots the read count of each allele in each sample. This is useful for getting a very detailed view of individual SNPs in few samples. As can be seen below, four samples from the HapMap data contains a strong imbalance at the chr17:79478331 position on the plus strand. By default the p-value is calculated by a chi-square test. To use other test results the arguments testValue and testValue2 can be used. When the counts for one allele are below 5 for one allele the chi-square test returns NA. This is why there is no P-value above the first bar in the example below.

```
barplot(a.stranded[1], strand="+", xlab.text="", legend.interspace=2)
```




```
#use another source of p-values  
btp <- binom.test(a.stranded[1], "+")  
barplot(a.stranded[1], strand="+", testValue=btp, xlab.text="", legend.interspace=2)
```



Another example of plotting that is useful is the one invoked with the plotting type argument "fraction". This plotting mechanism is useful to illustrate more SNPs and more samples in less space than the standard plot. As can be seen here several other samples are not heterozygote at the chr17:79478331 location.




```
barplot(a.simple[1], type="fraction", cex.main = 0.7)
```



samples

A typical question would be to ask why certain heterozygote samples have allele specific expression. The arguments `sampleColour.top` and `sampleColour.bot` allows for different highlights such as illustrated here below for gender. This could also be used to highlight based on genotype of proximal non-coding SNPs if available.


```
#top colour
sampleColour.top <- rep("palevioletred", ncol(a.simple))
sampleColour.top[colData(a.simple)[,"Gender"]%in%"male"] <- "darkgreen"
#bottom colour
sampleColour.bot <- rep("blue", ncol(a.simple))
sampleColour.bot[colData(a.simple)[,"Gender"]%in%"male"] <- "seashell2"
barplot(a.simple[1], type="fraction", sampleColour.top=sampleColour.top,
        sampleColour.bot=sampleColour.bot, cex.main = 0.7)
```



## 4.2 Plot with annotation

It is often of interest to combine the RNA sequencing data with genomic annotation information from online databases. For this purpose there is a function to extract variant specific annotation such as gene, exon, transcript and CDS.


```
library(org.Hs.eg.db)
barplot(a.simple[1],OrgDb=org.Hs.eg.db,
        cex.main = 0.7,
        xlab.text="",
        ypos.annotation = 0.08,
        annotation.interspace=1.3,
        legend.interspace=1.5
        )
```




### 4.3 locationplot

Finally a given gene or set of proximal genes will often have several SNPs close to each other. It is of interest to investigate all of them together, in connection with annotation information. This can be done using the `locationplot` function. This function in its simplest form just plot all the SNPs in an `ASEset` distributed by genomic location. Additionally it contains methods for including gene-map information through the arguments `OrgDb` and `TxDb`.

```
library(TxDb.Hsapiens.UCSC.hg19.knownGene)
#using count type
locationplot(a.stranded, type="count", cex.main=0.7, cex.legend=0.4)
```



```
#use annotation
a.stranded.sorted <- sort(a.stranded, decreasing=FALSE)
locationplot(a.stranded.sorted, OrgDb=org.Hs.eg.db, TxDb=TxDb.Hsapiens.UCSC.hg19.knownGene)
```




#### 4.4 Top allele criteria


The barplot with type='function' can be visualized according to three different allele sorting criteria. The default behaviour is just to show the allele with highest overall abundance in the upper half of the plot. This works well for most single SNP investigations. For more complex situations, however, it can be essential to keep track of phase information. This is done either through the reference allele sorting function, or even better, through consistently showing the maternal allele on top. When phase is known, this is essential to compare effect-directions of different coding SNPs

```
#load data
data(ASEset)

#using reference and alternative allele information
alt(ASEset) <- inferAltAllele(ASEset)
barplot(ASEset[1], type="fraction", strand="+", xlab.text="", top.fraction.criteria="ref")
```



```
#using phase information
phase(ASEset) <- phase.mat
barplot(ASEset[1], type="fraction", strand="+", xlab.text="", top.fraction.criteria="phase")
```



## 5 Grid graphics


To make use of the extra graphical flexibility that comes with grid graphics, the AllelicImbalance package has now functions that can be integrated in the grid environment. The low level grid functions are located in the [grid package](#), but higher level functions based on grid can be found in popular packages such as e.g. [lattice](#), [ggplot2](#). The benefits using grid graphics is a better control over different graphical regions and coordinate systems, how to construct and redirect graphical output, such as lines, points, etc and their appearance.

One particular reason for the AllelicImbalance to use grid graphics is the [Gviz](#) package, which uses grids functionality to construct tracks to visualize genomic regions in a smart way. The AllelicImbalance package has grid based barplots and functionality to create Gviz tracks. Observe the extra “g” prefix for the bar- and locationplot versions using the grid, see examples below.

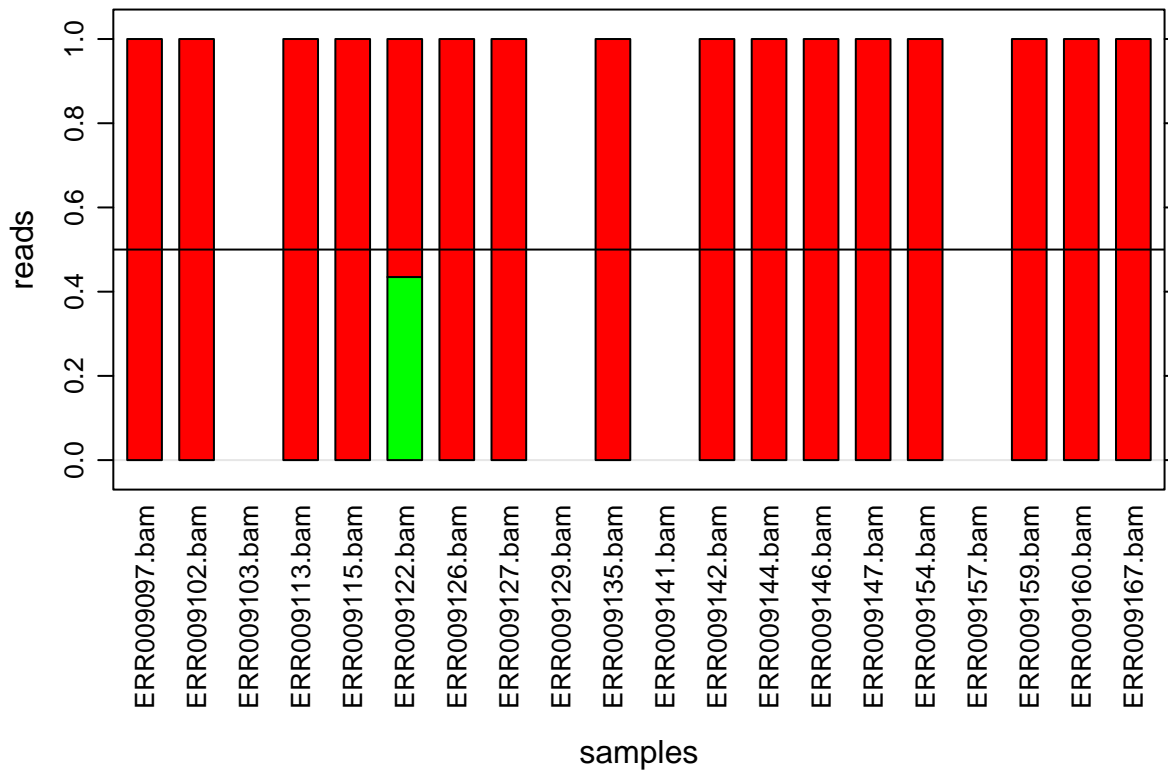
## 5.1 gbarplot

The standard barplot functions for ASEset objects have at the moment much more parameters that can be set by the user than the gbarplot. Base graphics also produce graphs faster than the grid environment. The advantage of the gbarplot is instead the possibility to integrate into the Gviz package and other grid based environments.

```
#gbarplots with type "count"
gbarplot(ASEset, type="count")
```



```
# gbarplots with type "fraction"
gbarplot(ASEset, type="fraction")
```




## 5.2 glocationplot

The glocationplot is a wrapper to quickly get an overview of ASE in a region. The overview offers the possibility to view consistency of SNP fractions over consecutive exons which is important to assess the reliability of the measured SNPs. In the best case all fractions deviates from e.g. 0.3 for every SNP in the same gene for an individual. To be able to inspect that the direction is right, it is possible to use the top.fraction.criteria and use phase information as explained earlier.


```
#remember to set the genome
genome(ASEset) <- "hg19"

glocationplot(ASEset, strand='+')
```



*#for ASEsets with fewer SNPs the 'count' type plot is useful*  
`glocationplot(ASEset[,1:5], type="count")`





### 5.3 Custom location plots

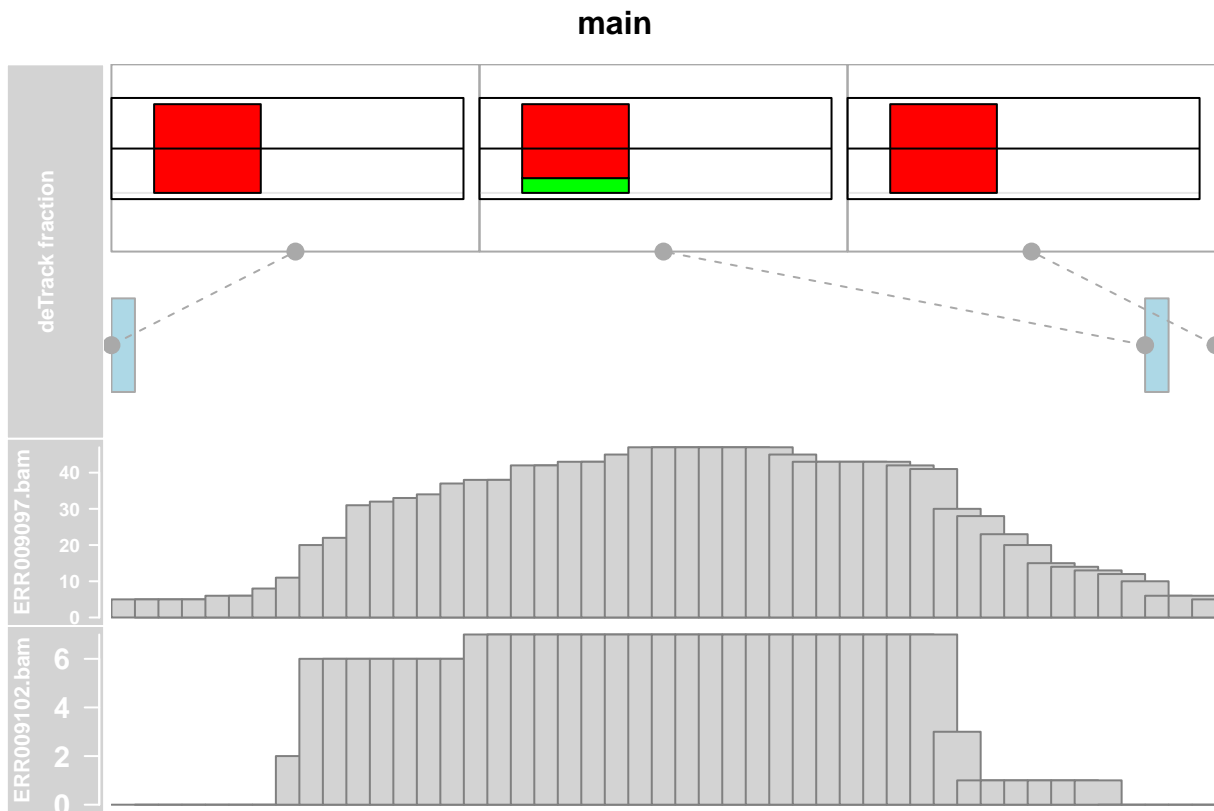
More flexibility and functionality from the Gviz package is accessed if the tracks are constructed separately. This is useful to evaluate AI in respect to for example read coverage or transcript annotation. An inconsistency of AI could potentially be explained by an uneven coverage over the exons, or lack of coverage. Below is an example using two samples from the ASEset and the corresponding reads. Ensuring and setting the seqlevels equal is a security measure, for the internal overlap calculations that otherwise might throw a warning. The GR object is defining the region that we want to plot. The ASEAnnotationTrack will create a track based on the gbarplots and the CoverageDataTrack a track based on read coverage. The different tracks are then collected in a list with the first element appearing at the top in the final plotting by the plotTracks function. The sizes vector defines the vertical space assigned to each track. To use transcript annotation more detailed information can be found in the Gviz [vignette](#)

```
#subset ASEset and reads
x <- ASEset[,1:2]
r <- reads[1:2]
seqlevels(r, force=TRUE) <- seqlevels(x)

GR <- GRanges(seqnames=seqlevels(x),
              ranges=IRanges(start=min(start(x)),end=max(end(x))),
              strand='+', genome=genome(x))

deTrack <- ASEAnnotationTrack(x, GR=GR, type='fraction',strand='+')
covTracks <- CoverageDataTrack(x, BamList=r, strand='+')
lst <- c(deTrack,covTracks)
sizes <- c(0.5,rep(0.5/length(covTracks),length(covTracks)))
```

```
plotTracks(lst, from=min(start(x)), to=max(end(x)), sizes=sizes, col.line = NULL, showId = FALSE, main = '')
```



## 6 Summary functions

The `regionSummary` function can be used to investigate if there is a consistent imbalance in the same direction over a region (e.g. a transcript).

```
# in this example every snp is on separate exons
```

```
region <- granges(ASEset)
rs <- regionSummary(ASEset, region)
head(rs)
```

```
##          het hom  mean.fr      sd.fr mean.delta  sd.delta ai.up
## ERR009097.bam  2  1 0.1961538 0.005439283 0.3038462 0.005439283  0
## ERR009102.bam  0  3      NaN           NA      NaN      NA      0
## ERR009103.bam NA NA      NaN           NA      NaN      NA      0
## ERR009113.bam  0  3      NaN           NA      NaN      NA      0
## ERR009115.bam  0  3      NaN           NA      NaN      NA      0
## ERR009122.bam  3  0 0.5567819 0.309060711 0.2277220 0.150229174  2
##
##          ai.down
## ERR009097.bam  2
## ERR009102.bam  0
## ERR009103.bam  0
## ERR009113.bam  0
## ERR009115.bam  0
## ERR009122.bam  1
```

## 7 Conclusion

---

In conclusion we hope that you will find this package useful in the investigation of the genetics of RNA-seq experiments. The various import functions should assist in the task of actually retrieving allele counts for specific nucleotide positions from all RNA-seq reads, including the non-trivial cases of intron-spanning reads. Likewise, the statistical analysis and plotting functions should be helpful in discovering any allele specific expression patterns that might be found in your data.

## 8 Links

---

Bowtie [link](#)

BWA [link](#)

Samtools [link](#)

Samtools pileup [link](#)

Grid graphics [link](#)

## 9 Session Info

---

```
sessionInfo()
```

```
## R version 3.2.0 (2015-04-16)
## 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
##  [3] LC_TIME=en_US.UTF-8      LC_COLLATE=C
##  [5] LC_MONETARY=en_US.UTF-8  LC_MESSAGES=en_US.UTF-8
##  [7] LC_PAPER=en_US.UTF-8    LC_NAME=C
##  [9] LC_ADDRESS=C             LC_TELEPHONE=C
## [11] LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C
##
## attached base packages:
##  [1] stats4    parallel  grid      stats     graphics  grDevices  utils
##  [8] datasets  methods   base
##
## other attached packages:
##  [1] TxDb.Hsapiens.UCSC.hg19.knownGene_3.1.2
##  [2] GenomicFeatures_1.20.1
##  [3] SNPlocs.Hsapiens.dbSNP.20120608_0.99.9
##  [4] BSgenome_1.36.0
##  [5] rtracklayer_1.28.2
##  [6] org.Hs.eg.db_3.1.2
##  [7] RSQLite_1.0.0
##  [8] DBI_0.3.1
##  [9] AnnotationDbi_1.30.1
## [10] Biobase_2.28.0
## [11] AllelicImbalance_1.6.2
## [12] Gviz_1.12.0
```

```

## [13] VariantAnnotation_1.14.1
## [14] GenomicAlignments_1.4.1
## [15] Rsamtools_1.20.2
## [16] Biostrings_2.36.1
## [17] XVector_0.8.0
## [18] GenomicRanges_1.20.3
## [19] GenomeInfoDb_1.4.0
## [20] IRanges_2.2.1
## [21] S4Vectors_0.6.0
## [22] BiocGenerics_0.14.0
## [23] BiocStyle_1.6.0
##
## loaded via a namespace (and not attached):
## [1] reshape2_1.4.1      splines_3.2.0      lattice_0.20-31
## [4] colorspace_1.2-6   htmltools_0.2.6    yaml_2.1.13
## [7] XML_3.98-1.1       survival_2.38-1    foreign_0.8-63
## [10] BiocParallel_1.2.1 RColorBrewer_1.1-2 lambda.r_1.1.7
## [13] matrixStats_0.14.0 plyr_1.8.2         stringr_1.0.0
## [16] zlibbioc_1.14.0    munsell_0.4.2      gtable_0.1.2
## [19] futile.logger_1.4.1 evaluate_0.7        latticeExtra_0.6-26
## [22] knitr_1.10.5       biomaRt_2.24.0     proto_0.3-10
## [25] Rcpp_0.11.6        acepack_1.3-3.3    scales_0.2.4
## [28] formatR_1.2        Hmisc_3.16-0       gridExtra_0.9.1
## [31] ggplot2_1.0.1      digest_0.6.8       stringi_0.4-1
## [34] biovizBase_1.16.0 ade4_1.7-2          tools_3.2.0
## [37] bitops_1.0-6       magrittr_1.5        RCurl_1.95-4.6
## [40] dichromat_2.0-0    Formula_1.2-1      cluster_2.0.1
## [43] futile.options_1.0.0 seqinr_3.1-3       MASS_7.3-40
## [46] rmarkdown_0.6.1    rpart_4.1-9        nnet_7.3-9

```

“Who is is the master if nature provides the techniques to study itself”