

An Introduction to *VariantTools*

Michael Lawrence, Jeremiah Degenhardt

November 13, 2013

Contents

1	Introduction	2
2	Calling single-sample variants	2
2.1	Basic usage	2
2.2	Step by step	4
2.3	Diagnosing the filters	5
2.4	Extending and customizing the workflow	7
3	Comparing variant sets across samples	7
3.1	Calling sample-specific variants	7
4	Exporting the calls as VCF	8
5	Finding Wildtype and No-call Regions	8

1 Introduction

This vignette outlines the basic usages of the *VariantTools* package and the general workflow for loading data, calling single sample variants and tumor-specific somatic mutations or other sample-specific variant types (eg RNA editing). Most of the functions operate on alignments (BAM files) or datasets of called variants. The user is expected to have already aligned the reads with a separate tool, e.g., GSNAP via *gmapR*.

2 Calling single-sample variants

2.1 Basic usage

For our example, we take paired-end RNA-seq alignments from two lung cancer cell lines from the same individual. H1993 is derived from a metastasis and H2073 is derived from the primary tumor.

Below, we call variants from a region around the p53 gene:

```
> library(VariantTools)
> p53 <- gmapR:::exonsOnTP53Genome("TP53") + 1e5
> bams <- LungCancerLines::LungCancerBamFiles()
> bam <- bams$H1993
> tally.param <- TallyVariantsParam(gmapR::TP53Genome(),
+                                 high_base_quality = 23L,
+                                 which = range(p53),
+                                 indels = TRUE)
> called.variants <- callVariants(bam, tally.param)
```

In the above, we load the genome corresponding to the human p53 gene region and the H1993 BAM file (stripped down to the same region). We pass the BAM, genome, read length and quality cutoff to the `callVariants` workhorse. The read length is not strictly required, but it is necessary for one of the QA filters. The value given for the high base quality cutoff is appropriate for Sanger and Illumina 1.8 or above. By default, the high quality counts are used by the likelihood ratio test during calling.

The returned `called_variants` is a variant *GRanges*, in the same form as that returned by `bam_tally` in the *gmapR* package. Unsurprisingly, `callVariants` uses `bam_tally` internally to generate the per-nucleotide counts (pileup) from the BAM file. The result is then filtered to generate the variant calls. The *VCF* class holds similar information; however, we favor the simple tally *GRanges*, because it has a separate record for each ALT, at each position. *VCF*, the class and the file format, has a single record for a position, collapsing over multiple ALT alleles, and this is much less convenient for our purposes.

If we subset the variants by those in an actual p53 exon (not an intron), we find two: one with strong evidence for a homozygous mutation, and another with much weaker evidence (low coverage).

```
> subsetByOverlaps(called.variants, p53, ignore.strand = TRUE)
```

VRanges with 36102 ranges and 15 metadata columns:

	seqnames	ranges	strand	ref
	<Rle>	<IRanges>	<Rle>	<character>
[1]	TP53	[906407, 906407]	+	G
[2]	TP53	[906408, 906408]	+	G
[3]	TP53	[906414, 906414]	+	A
[4]	TP53	[906416, 906416]	+	C
[5]	TP53	[906417, 906417]	+	T
...
[36098]	TP53	[1125618, 1125618]	+	A
[36099]	TP53	[1125619, 1125619]	+	G

	alt	totalDepth	refDepth	altDepth
	<characterOrRle>	<integerOrRle>	<integerOrRle>	<integerOrRle>
[36100]	TP53 [1125630, 1125630]	+	A	
[36101]	TP53 [1125631, 1125631]	+	A	
[36102]	TP53 [1125632, 1125632]	+	A	
[1]	A	2	0	2
[2]	T	2	0	2
[3]	G	2	0	2
[4]	A	2	0	2
[5]	A	2	0	2
...
[36098]	T	2	0	2
[36099]	A	2	0	2
[36100]	C	2	0	2
[36101]	C	2	0	2
[36102]	C	2	0	2
sampleNames	softFilterMatrix	n.read.pos	n.read.pos.ref	
<factorOrRle>	<matrix>	<integer>	<integer>	
[1]	<NA>	2	0	
[2]	<NA>	2	0	
[3]	<NA>	2	0	
[4]	<NA>	2	0	
[5]	<NA>	2	0	
...
[36098]	<NA>	2	0	
[36099]	<NA>	2	0	
[36100]	<NA>	2	0	
[36101]	<NA>	2	0	
[36102]	<NA>	2	0	
raw.count	raw.count.ref	raw.count.total	mean.quality	
<integer>	<integer>	<integer>	<numeric>	
[1]	2	0	2	34.5
[2]	2	0	2	33.5
[3]	2	0	2	34.5
[4]	2	0	2	34
[5]	2	0	2	34.5
...
[36098]	2	0	2	33
[36099]	2	0	2	26.5
[36100]	2	0	2	33
[36101]	2	0	2	32.5
[36102]	2	0	2	31
mean.quality.ref	count.pos	count.pos.ref	count.neg	
<numeric>	<integer>	<integer>	<integer>	
[1]	NaN	1	0	1
[2]	NaN	1	0	1
[3]	NaN	1	0	1
[4]	NaN	1	0	1
[5]	NaN	1	0	1
...
[36098]	NaN	1	0	1

```

[36099]      NaN      1      0      1
[36100]      NaN      1      0      1
[36101]      NaN      1      0      1
[36102]      NaN      1      0      1
      count.neg.ref read.pos.mean read.pos.mean.ref read.pos.var
      <integer>      <numeric>      <numeric>      <numeric>
[1]      0      91      NaN      162
[2]      0      91      NaN      128
[3]      0      91      NaN      8
[4]      0      91      NaN      0
[5]      0      91      NaN      2
...      ...      ...      ...      ...
[36098]      0      75.5      NaN      24.5
[36099]      0      75.5      NaN      40.5
[36100]      0      75.5      NaN      480.5
[36101]      0      75.5      NaN      544.5
[36102]      0      75.5      NaN      612.5
      read.pos.var.ref
      <numeric>
[1]      <NA>
[2]      <NA>
[3]      <NA>
[4]      <NA>
[5]      <NA>
...      ...
[36098]      <NA>
[36099]      <NA>
[36100]      <NA>
[36101]      <NA>
[36102]      <NA>

```

```

---
seqlengths:
  TP53
  2025772
hardFilters(4): nonRef nonNRef readCount likelihoodRatio

```

The next section goes into further detail on the process, including the specific filtering rules applied, and how one might, for example, tweak the parameters to avoid calling low-coverage variants, like the one above.

2.2 Step by step

The `callVariants` method for BAM files, introduced above, is a convenience wrapper that delegates to several low-level functions to perform each step of the variant calling process: generating the tallies, basic QA filtering and the actual variant calling. Calling these functions directly affords the user more control over the process and provides access to intermediate results, which is useful e.g. for diagnostics and for caching results. The workflow consists of three function calls that rely on argument defaults to achieve the same result as our call to `callVariants` above. Please see their man pages for the arguments available for customization.

The first step is to tally the variants from the BAM file. By default, this will return observed differences from the reference, excluding N calls and only counting reads above 13 in mapping quality (MAPQ) score. There are three read position bins: the first 10 bases, the final 10 bases, and the stretch between them (these will be used in the QA step).

```
> raw.variants <- tallyVariants(bam, tally.param, mc.cores=1)
```

Unless one is running a variant caller in a routine fashion over familiar types of data, we highly recommend performing detailed QC of the tally results. *VariantTools* provides several QA filters that aim to expose artifacts, especially those generated during alignment. These filters are *not* designed for filtering during actual calling; rather, they are meant for annotating the variants during exploratory analysis. The filters include a minimum read count (2) check, minimum unique read position count (2) check, and Fisher Exact Test on the per-strand counts vs. reference for strand bias (p-value cutoff: 0.001). If there are at least three read position bins in the tallies, at least one read must present the variant in an internal read position bin. The intent is to ensure that we have sufficient data and that the data are not due to strand-specific nor read position-specific artifacts.

The `qaVariants` function will *soft* filter the variants via `softFilter`. No variants are removed; the filter results are added to the `softFilterMatrix` component of the object.

```
> qa.variants <- qaVariants(raw.variants)
> head(softFilterMatrix(qa.variants))
```

```
FilterMatrix (6 x 4)
  readPosCount fisherStrand readPosBin readPosTTest
[1]      FALSE          TRUE         TRUE         TRUE
[2]      FALSE          TRUE         TRUE         TRUE
[3]      FALSE          TRUE         TRUE         TRUE
[4]      FALSE          TRUE         TRUE         TRUE
[5]      FALSE          TRUE         TRUE         TRUE
[6]      FALSE          TRUE         TRUE         TRUE
```

The final step is to actually call the variants. The `callVariants` function uses a binomial likelihood ratio test for this purpose. The ratio is $P(D|p = p_{lower})/P(D|p = p_{error})$, where $p_{lower} = 0.2$ is the assumed lowest variant frequency and $p_{error} = 0.001$ is the assumed error rate in the sequencing (default: 0.001).

```
> called.variants <- callVariants(qa.variants)
```

The `callVariants` function applies an additional set of filters after the actual variant calling. These are known as “post” filters and consider the putative variant calls as a set, independent of the calling algorithm. Currently, there is only one post filter by default, and it discards variants that are clumped together along the chromosome, as these often result from mapping difficulties.

2.3 Diagnosing the filters

The calls to `qaVariants` and `callVariants` are essentially filtering the tallies, so it is important to know, especially when faced with a new dataset, the effect of each filter and the effect of the individual parameters on each filter.

The filters are implemented as modules and are stored in a *FilterRules* object from the *IRanges* package. We can create those filters directly and rely on some *FilterRules* utilities to diagnose the filtering process.

Here we construct the *FilterRules* that implements the `qaVariants` function. Again, we rely on the argument defaults to generate the same answer.

```
> qa.filters <- VariantQAFilters()
```

We can now ask for a summary of the filtering process, which gives the number of variants that pass each filter, separately and then combined:

```
> summary(qa.filters, raw.variants)
```

```

<initial> readPosCount fisherStrand  readPosBin readPosTTest
65978      43520      65671      65978      63174
<final>
40809

```

Now we retrieve only the variants that pass the filters:

```
> qa.variants <- subsetByFilter(raw.variants, qa.filters)
```

We could do the same, except modify a filter parameter, such as the p-value cutoff for the Fisher Exact Test for strand bias:

```
> qa.filters.custom <- VariantQAFilters(fisher.strand.p.value = 1e-4)
> summary(qa.filters.custom, raw.variants)
```

```

<initial> readPosCount fisherStrand  readPosBin readPosTTest
65978      43520      65671      65978      63174
<final>
40809

```

To get a glance at the additional variants we are discarding compared to the previous cutoff, we can subset the filter sets down to the Fisher strand filter, evaluate the old and new filter, and compare the results:

```
> fs.original <- eval(qa.filters["fisherStrand"], raw.variants)
> fs.custom <- eval(qa.filters.custom["fisherStrand"], raw.variants)
> raw.variants[fs.original != fs.custom]
```

VRanges with 0 ranges and 15 metadata columns:

```

seqnames  ranges strand      ref      alt
  <Rle> <IRanges> <Rle> <character> <characterOrRle>
  totalDepth  refDepth  altDepth  sampleNames
<integerOrRle> <integerOrRle> <integerOrRle> <factorOrRle>
softFilterMatrix | n.read.pos n.read.pos.ref raw.count raw.count.ref
  <matrix> | <integer> <integer> <integer> <integer>
raw.count.total mean.quality mean.quality.ref count.pos
  <integer> <numeric> <numeric> <integer>
count.pos.ref count.neg count.neg.ref read.pos.mean
  <integer> <integer> <integer> <numeric>
read.pos.mean.ref read.pos.var read.pos.var.ref
  <numeric> <numeric> <numeric>
---

```

```
seqlengths:
```

```
TP53
```

```
2025772
```

```
hardFilters(0):
```

Below, we demonstrate how one might add a mask to e.g. filter out variants in low complexity regions, where mapping errors tend to dominate:

```
> tally.param@mask <- GRanges("TP53", IRanges(1010000, width=10000))
> raw.variants.masked <- tallyVariants(bam, tally.param)
```

We can also diagnose the filters for calling variants after basic QA checks.

```
> calling.filters <- VariantCallingFilters()
> summary(calling.filters, qa.variants)
```

<initial>	nonRef	nonNRef	readCount
40809	40809	40809	33447
likelihoodRatio	<final>		
33082	32422		

Check how the post filter would perform prior to variant calling:

```
> post.filters <- VariantPostFilters()
> summary(post.filters, qa.variants)
```

<initial>	avgNborCount	<final>
40809	4	4

What about if we preserved the ones we have already called?

```
> post.filters <- VariantPostFilters(whitelist = called.variants)
> summary(post.filters, qa.variants)
```

<initial>	avgNborCount	<final>
40809	37961	37961

2.4 Extending and customizing the workflow

Since the built-in filters are implemented using *FilterRules*, it is easy to mix and match different filters, including those implemented externally to the *VariantTools* package. This is the primary means of extending and customizing the variant calling workflow.

3 Comparing variant sets across samples

So far, we have called variants for the metastatic H1993 sample. We leave the processing of the primary tumor H2073 sample as an exercise to the reader and instead turn our attention to detecting the variants that are specific to the metastatic sample, as compared to the primary tumor.

3.1 Calling sample-specific variants

The function `callSampleSpecificVariants` takes the case (e.g., tumor) sample and control (e.g., matched normal) sample as input. In our case, we are comparing the metastatic line (H1993) to the primary tumor line (H2073) from the same patient, a smoker. To avoid inconsistencies, it is recommended to pass BAM files as input, for which tallies are automatically generated, subjected to QA, and called as variants vs. reference, prior to determining the sample-specific variants.

Here, we find the somatic mutations from a matched tumor/normal pair. Since we are starting from BAM files, we have to provide `tally.param` for the tally step.

```
> tally.param@bamTallyParam@indels <- FALSE
> somatic <- callSampleSpecificVariants(bams$H1993, bams$H2073, tally.param)
```

This can be time-consuming for the entire genome, since the tallies need to be computed. To avoid repeated computation of the tallies, the user can pass the raw tally *GRanges* objects instead of the BAM files. This is less safe, because anything could have happened to those *GRanges* objects.

The QA and initial calling are optionally controlled by passing *FilterRules* objects, typically those returned by `VariantQAFilters` and `VariantCallingFilters`, respectively. For controlling the final step, determining the sample-specific variants, one may pass filter parameters directly to `callSampleSpecificVariants`. Here is an example of customizing some parameters.

```

> calling.filters <- VariantCallingFilters(read.count = 3L)
> somatic <- callSampleSpecificVariants(bams$H1993, bams$H2073, tally.param,
+                                     calling.filters = calling.filters,
+                                     power = 0.9, p.value = 0.001)

```

4 Exporting the calls as VCF

VCF is a common file format for communicating variants. To export our variants to a VCF file, we first need to coerce the *GRanges* to a *VCF* object. Then, we use `writeVcf` from the *VariantAnnotation* package to write the file (indexing is highly recommended for large files). Note that the sample names need to be non-missing to generate the VCF. Also, for simplicity and scalability, we typically do not want to output all of our metadata columns, so we remove all of them here.

```

> sampleNames(called.variants) <- "H1993"
> mcols(called.variants) <- NULL
> vcf <- asVCF(called.variants)

> writeVcf(vcf, "H1993.vcf", index = TRUE)

```

5 Finding Wildtype and No-call Regions

So far, our analysis has yielded a set of positions that are likely to be variants. We have not made any claims about the status of the positions outside of that set. For this, we need to decide, for each position, whether there was sufficient coverage to detect a variant, if one existed. The following call carries out a power test to decide whether a region is variant, wildtype or is unable to be called due to lack of coverage. The variants must have been called using the filters returned by `VariantCallingFilters`. The algorithm depends on the filter parameter settings, so it is possible and indeed required for the user to pass filter object used for calling the variants. This requirement is an attempt to ensure consistency and will be made more convenient in the future. To request the calls for a particular set of positions, one can pass a *GenomicRanges* (where all the ranges are of width 1) as the `pos` argument. When `pos` is specified, each element of the result corresponds to an element in `pos`.

```

> called.variants <- called.variants[!isIndel(called.variants)]
> pos <- c(called.variants, shift(called.variants, 3))
> wildtype <- callWildtype(bam, called.variants, VariantCallingFilters(),
+                          pos = pos, power = 0.85)

```

The returned object is a logical *RleList* object, with `TRUE` for wildtype, `FALSE` for variant and `NA` for no-call. Thus, we could calculate the fraction called as follows:

```

> mean(!is.na(wildtype))

[1] 0.9131044

```

Sometimes it is desirable for the wildtype calls to be returned as simple vector, with a logical value for each position (range of width one) in `bamWhich`. Such a vector is returned when `global = FALSE` is passed to `callWildtype`. This is the same as extracting the positions from the ordinary *Rle* return value, but it is implemented more efficiently, at least for a relatively small number of positions.