

The `quantro` user's guide

Stephanie C. Hicks^{*} and Rafael A. Irizarry[†]

Dana-Farber Cancer Institute / Harvard T.H. Chan School of Public Health

^{*}shicks@jimmy.harvard.edu [†]rafa@jimmy.harvard.edu

Modified: April 7, 2017. Compiled: October 30, 2018

Contents

1	Introduction	1
2	Getting Started	2
3	Data	3
3.1	flowSorted Data Example	3
3.2	Plot distributions	4
4	Using the <code>quantro()</code> function	5
4.1	Input for <code>quantro()</code>	5
4.2	Running <code>quantro()</code>	5
4.3	eSets	7
4.4	Output from <code>quantro()</code>	7
5	Assessing the statistical significance.	8
6	Visualizing the statistical significance from permutation tests.	9
7	SessionInfo	10

1 Introduction

Multi-sample normalization techniques such as quantile normalization [1, 2] have become a standard and essential part of analysis pipelines for high-throughput data. Although it was originally developed for gene expression microarrays, it is now used across many different high-throughput applications including genotyping arrays, DNA Methylation, RNA Sequencing (RNA-Seq) and Chromatin Immunoprecipitation Sequencing (ChIP-Seq). These techniques

transform the original raw data to remove unwanted technical variation. However, quantile normalization and other global normalization methods rely on assumptions about the data generation process that are not appropriate in some context. Until now, it has been left to the researcher to check for the appropriateness of these assumptions.

Quantile normalization assumes that the statistical distribution of each sample is the same. Normalization is achieved by forcing the observed distributions to be the same and the average distribution, obtained by taking the average of each quantile across samples, is used as the reference. This method has worked very well in practice but note that when the assumptions are not met, global changes in distribution, that may be of biological interest, will be wiped out and features that are not different across samples can be artificially induced. These types of assumptions are justified in many biomedical applications, for example in gene expression studies in which only a minority of genes are expected to be differentially expressed. However, if, for example, a substantially higher percentage of genes are expected to be expressed in only one group of samples, it may not be appropriate to use global adjustment methods.

The `quantro` R-package can be used to test for global differences between groups of distributions which assesses whether global normalization methods such as quantile normalization should be applied. Our method uses the raw unprocessed high-throughput data to test for global differences in the distributions across a set of groups. The main function `quantro()` will perform two tests:

1. An ANOVA to test if the medians of the distributions are different across groups. Differences across groups could be attributed to unwanted technical variation (such as batch effects) or real global biological variation. This is a helpful step for the user to verify if there is any technical variation unaccounted for.
2. A test for global differences between the distributions across groups which returns a test statistic called `quantroStat`. This test statistic is a ratio of two variances and is similar to the idea of ANOVA. The main idea is to compare the variability of distributions within groups relative to between groups. If the variability between groups is sufficiently larger than the variability within groups, then this suggests global adjustment methods may not be appropriate. As a default, we perform this test on the median normalized data, but the user may change this option.

2 Getting Started

Load the package in R

```
library(quantro)
```

3 Data

3.1 flowSorted Data Example

To explore how to use `quantro()`, we use the `FlowSorted.DLPFC.450k` data package in Bioconductor [3]. This data set contains raw data objects of 58 Illumina 450K DNA methylation microarrays, formatted as `RGset` objects. The samples represent two different cellular populations of brain tissues on the same 29 individuals extracted using flow sorting. For more information on this data set, please see the `FlowSorted.DLPFC.450k` User's Guide. For the purposes of this vignette, a `MethylSet` object from the `minfi` Bioconductor package [4] was created which is a subset of the rows from the original `FlowSorted.DLPFC.450k` data set. This `MethylSet` object is found in the `/data` folder and the script to create the object is found in `/inst`.

Here we will explore the distributions of these two cellular populations of brain tissue (`NeuN_pos` and `NeuN_neg`) and then test if there are global differences in the distributions across groups. First, load the `MethylSet` object (`flowSorted`) and compute the Beta values using the function `getBeta()` in the `minfi` Bioconductor package. We use an offset of 100 as this is the default used by Illumina.

```
library(quantro)
library(minfi)
data(flowSorted)
p <- getBeta(flowSorted, offset = 100)
pd <- pData(flowSorted)
dim(p)

## [1] 10000    58

head(pd)

## DataFrame with 6 rows and 11 columns
##      Sample_Name SampleID CellType Sentrix.Barcode Sample.Section
##      <character> <character> <character>      <numeric>      <character>
## 813_N      813_N      813    NeuN_pos      7766130090      R02C01
## 1740_N     1740_N     1740    NeuN_pos      7766130090      R02C02
## 1740_G     1740_G     1740    NeuN_neg      7766130090      R04C01
## 1228_G     1228_G     1228    NeuN_neg      7766130090      R04C02
## 813_G      813_G      813    NeuN_neg      7766130090      R06C01
## 1228_N     1228_N     1228    NeuN_pos      7766130090      R06C02
##      diag      sex ethnicity      age      PMI
##      <character> <character> <character> <integer> <integer>
## 813_N      Control      Female  Caucasian      30      14
## 1740_N      Control      Female  African        13      17
## 1740_G      Control      Female  African        13      17
## 1228_G      Control      Male    Caucasian      47      13
## 813_G      Control      Female  Caucasian      30      14
## 1228_N      Control      Male    Caucasian      47      13
##
## BasePath
```


```
##                                                                 <character>
## 813_N /dcs01/lieber/ajaffe/450k/Zack_450k/IDAT/7766130090/7766130090_R02C01
## 1740_N /dcs01/lieber/ajaffe/450k/Zack_450k/IDAT/7766130090/7766130090_R02C02
## 1740_G /dcs01/lieber/ajaffe/450k/Zack_450k/IDAT/7766130090/7766130090_R04C01
## 1228_G /dcs01/lieber/ajaffe/450k/Zack_450k/IDAT/7766130090/7766130090_R04C02
## 813_G /dcs01/lieber/ajaffe/450k/Zack_450k/IDAT/7766130090/7766130090_R06C01
## 1228_N /dcs01/lieber/ajaffe/450k/Zack_450k/IDAT/7766130090/7766130090_R06C02
```

3.2 Plot distributions


quantro contains two functions to view the distributions of the samples of interest: `matdensity()` and `matboxplot()`. `matdensity()` computes the density for each sample (columns) and uses the `matplot()` function to plot all the densities. `matboxplot()` orders and colors the samples by a group level variable. These two functions use the `RColorBrewer` package and the brewer palettes can be changed using the arguments `brewer.n` and `brewer.name`.

The distributions of the two groups of cellular populations are shown here. The NeuN_neg samples are colored in green and the NeuN_pos are colored in red.

```
matdensity(p, groupFactor = pd$CellType, xlab = " ", ylab = "density",
           main = "Beta Values", brewer.n = 8, brewer.name = "Dark2")
legend('top', c("NeuN_neg", "NeuN_pos"), col = c(1, 2), lty = 1, lwd = 3)
```



```
matboxplot(p, groupFactor = pd$CellType, xaxt = "n", main = "Beta Values")
```



4 Using the `quantro()` function

4.1 Input for `quantro()`

The `quantro()` function must have two objects as input:

- an `object` which is a data frame or matrix with observations (e.g. probes or genes) on the rows and samples as the columns.
- a `groupFactor` which represents the group level information about each sample. For example if the samples represent tumor and normal samples, provide `quantro()` with a factor representing which columns in the `object` are normal and tumor samples.

4.2 Running `quantro()`

In this example, the groups we are interested in comparing are contained in the `CellType` column in the `pd` dataset. To run the `quantro()` function, input the data object and the object containing the phenotypic data. Here we use the `flowSorted` data set as an example.

```
qtest <- quantro(object = p, groupFactor = pd$CellType)
```

The `quantro` user's guide

```
## [quantro] Average medians of the distributions are
##                               not equal across groups.

## [quantro] Calculating the quantro test statistic.

## [quantro] No permutation testing performed.
##                               Use B > 0 for permutation testing.

qtest

## quantro: Test for global differences in distributions
##   nGroups: 2
##   nTotSamples: 58
##   nSamplesinGroups: 29 29
##   anovaPval: 0.01206
##   quantroStat: 8.80735
##   quantroPvalPerm: Use B > 0 for permutation testing.
```

The details related to the experiment can be extracted using the `summary` accessor function:

```
summary(qtest)

## $nGroups
## [1] 2
##
## $nTotSamples
## [1] 58
##
## $nSamplesinGroups
## NeuN_neg NeuN_pos
##      29      29
```

To assess if the medians of the distributions different across groups, we perform an ANOVA on the medians from the samples. Those results can be found using `anova`:

```
anova(qtest)

## Analysis of Variance Table
##
## Response: objectMedians
##           Df   Sum Sq   Mean Sq F value   Pr(>F)
## groupFactor 1 0.006919 0.0069194   6.7327 0.01206 *
## Residuals  56 0.057553 0.0010277
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The full output can be seen The test statistic produced from `quantro()` testing for global differences between distributions is given by `quantroStat`:

```
quantroStat(qtest)

## [1] 8.807348
```

4.3 `eSets`

`quantro()` also can accept objects that inherit `eSets` such as an `ExpressionSet` or `MethylSet`. The `groupFactor` must still be provided.

```
is(flowSorted, "MethylSet")

## [1] TRUE

qtest <- quantro(flowSorted, groupFactor = pData(flowSorted)$CellType)

## [quantro] Average medians of the distributions are
##                               not equal across groups.

## [quantro] Calculating the quantro test statistic.

## [quantro] No permutation testing performed.
##                               Use B > 0 for permutation testing.

qtest

## quantro: Test for global differences in distributions
##   nGroups: 2
##   nTotSamples: 58
##   nSamplesinGroups: 29 29
##   anovaPval: 0.01206
##   quantroStat: 8.80735
##   quantroPvalPerm: Use B > 0 for permutation testing.
```

4.4 Output from `quantro()`

Elements in the S4 object from `quantro()` include:

Element	Description
<code>summary</code>	Returns a list of three elements related to a summary of the experiment: <code>nGroups</code> : number of groups <code>nTotSamples</code> : total number of samples <code>nSamplesinGroups</code> : number of samples in each group
<code>anova</code>	Results from an ANOVA to test if the average medians of the distributions are different across groups
<code>MSbetween</code>	Mean squared error between groups
<code>MSwithin</code>	Mean squared error within groups
<code>quantroStat</code>	A test statistic which is a ratio of the mean squared error between groups of distributions (<code>MSbetween</code>) to the mean squared error within groups of distributions (<code>MSwithin</code>)
<code>quantroStatPerm</code>	If B is not equal to 0, then a permutation test was performed to assess the statistical significance of <code>quantroStat</code> . These are the test statistics resulting from the permuted samples
<code>quantroPvalPerm</code>	If B is not equal to 0, then this is the p -value associated with the proportion of times the test statistics resulting from the permuted samples were larger than <code>quantroStat</code>

5 Assessing the statistical significance

To assess statistical significance of the test statistic, we use permutation testing. We use the `foreach` package which distribute the computations across multiple cross in a single machine or across multiple machines in a cluster. The user must pick how many permutations to perform where B is the number of permutations. At each permutation of the samples, a test statistic is calculated. The proportion of test statistics (`quantroStatPerm`) that are larger than the `quantroStat` is reported as the `quantroPvalPerm`. To use the `foreach` package, we first register a backend, in this case a machine with 1 cores.

```
library(doParallel)
registerDoParallel(cores=1)
qtestPerm <- quantro(p, groupFactor = pd$CellType, B = 1000)

## [quantro] Average medians of the distributions are
##                               not equal across groups.

## [quantro] Calculating the quantro test statistic.

## [quantro] Starting permutation testing.

## [quantro] Using a single core (backend: doParallelMC,
##                               version: 1.0.14).

qtestPerm

## quantro: Test for global differences in distributions
##   nGroups: 2
##   nTotSamples: 58
##   nSamplesinGroups: 29 29
```




```
## anovaPval: 0.01206
## quantroStat: 8.80735
## quantroPvalPerm: 0.005
```

6 Visualizing the statistical significance from permutation tests

If permutation testing was used (i.e. specifying $B > 0$), then there is a second function in the package called `quantroPlot()` which will plot the test statistics of the permuted samples. The plot is a histogram of the null test statistics `quantroStatPerm` from `quantro()` and the red line is the observed test statistic `quantroStat` from `quantro()`.

```
quantroPlot(qtestPerm)
```



Additional options in the `quantroPlot()` function include:

Element	Description
<code>xLab</code>	the x-axis label
<code>yLab</code>	the y-axis label
<code>mainLab</code>	title of the histogram
<code>binWidth</code>	change the binwidth

7 SessionInfo

```

sessionInfo()

## R version 3.5.1 Patched (2018-07-12 r74967)
## Platform: x86_64-apple-darwin15.6.0 (64-bit)
## Running under: OS X El Capitan 10.11.6
##
## Matrix products: default
## BLAS: /Library/Frameworks/R.framework/Versions/3.5/Resources/lib/libRblas.0.dylib
## LAPACK: /Library/Frameworks/R.framework/Versions/3.5/Resources/lib/libRlapack.dylib
##
## locale:
## [1] C/en_US.UTF-8/en_US.UTF-8/C/en_US.UTF-8/en_US.UTF-8
##
## attached base packages:
## [1] stats4      parallel    stats       graphics    grDevices   utils       datasets
## [8] methods     base
##
## other attached packages:
##  [1] doParallel_1.0.14      minfi_1.28.0
##  [3] bumpHunter_1.24.0      locfit_1.5-9.1
##  [5] iterators_1.0.10       foreach_1.4.4
##  [7] Biostrings_2.50.0      XVector_0.22.0
##  [9] SummarizedExperiment_1.12.0 DelayedArray_0.8.0
## [11] BiocParallel_1.16.0    matrixStats_0.54.0
## [13] Biobase_2.42.0         GenomicRanges_1.34.0
## [15] GenomeInfoDb_1.18.0    IRanges_2.16.0
## [17] S4Vectors_0.20.0      BiocGenerics_0.28.0
## [19] quantro_1.16.0
##
## loaded via a namespace (and not attached):
##  [1] colorspace_1.3-2      siggenes_1.56.0
##  [3] mclust_5.4.1          rprojroot_1.3-2
##  [5] base64_2.0            bit64_0.9-7
##  [7] AnnotationDbi_1.44.0  xml2_1.2.0
##  [9] codetools_0.2-15     splines_3.5.1
## [11] knitr_1.20            Rsamtools_1.34.0
## [13] annotate_1.60.0       HDF5Array_1.10.0
## [15] BiocManager_1.30.3    readr_1.1.1
## [17] compiler_3.5.1        httr_1.3.1
## [19] backports_1.1.2       assertthat_0.2.0
## [21] Matrix_1.2-14         lazyeval_0.2.1
## [23] limma_3.38.0          htmltools_0.3.6
## [25] prettyunits_1.0.2     tools_3.5.1
## [27] bindrcpp_0.2.2        gtable_0.2.0
## [29] glue_1.3.0            GenomeInfoDbData_1.2.0
## [31] dplyr_0.7.7           doRNG_1.7.1
## [33] Rcpp_0.12.19          multtest_2.38.0
## [35] preprocessCore_1.44.0 nlme_3.1-137

```

```
## [37] rtracklayer_1.42.0      DelayedMatrixStats_1.4.0
## [39] stringr_1.3.1           rngtools_1.3.1
## [41] XML_3.98-1.16           beanplot_1.2
## [43] zlibbioc_1.28.0         MASS_7.3-51
## [45] scales_1.0.0            BiocStyle_2.10.0
## [47] hms_0.4.2              rhdf5_2.26.0
## [49] GEOquery_2.50.0         RColorBrewer_1.1-2
## [51] yaml_2.2.0              memoise_1.1.0
## [53] ggplot2_3.1.0           pkgmaker_0.27
## [55] biomaRt_2.38.0          reshape_0.8.8
## [57] stringi_1.2.4           RSQLite_2.1.1
## [59] highr_0.7               genefilter_1.64.0
## [61] GenomicFeatures_1.34.0  bibtex_0.4.2
## [63] rlang_0.3.0.1           pkgconfig_2.0.2
## [65] bitops_1.0-6            nor1mix_1.2-3
## [67] evaluate_0.12           lattice_0.20-35
## [69] purrr_0.2.5             Rhdf5lib_1.4.0
## [71] bindr_0.1.1            GenomicAlignments_1.18.0
## [73] labeling_0.3            bit_1.1-14
## [75] tidyselect_0.2.5        plyr_1.8.4
## [77] magrittr_1.5            R6_2.3.0
## [79] DBI_1.0.0               pillar_1.3.0
## [81] withr_2.1.2             survival_2.43-1
## [83] RCurl_1.95-4.11         tibble_1.4.2
## [85] crayon_1.3.4            rmarkdown_1.10
## [87] progress_1.2.0          grid_3.5.1
## [89] data.table_1.11.8       blob_1.1.1
## [91] digest_0.6.18           xtable_1.8-3
## [93] tidyr_0.8.2             illuminaio_0.24.0
## [95] openssl_1.0.2           munsell_0.5.0
## [97] registry_0.5            quadprog_1.5-5
```

References

- [1] B M Bolstad, R A Irizarry, M Astrand, and T P Speed. A comparison of normalization methods for high density oligonucleotide array data based on variance and bias. *Bioinformatics*, 19(2):185–93, Jan 2003.
- [2] Rafael A Irizarry, Bridget Hobbs, Francois Collin, Yasmin D Beazer-Barclay, Kristen J Antonellis, Uwe Scherf, and Terence P Speed. Exploration, normalization, and summaries of high density oligonucleotide array probe level data. *Biostatistics*, 4(2):249–64, Apr 2003. doi:[10.1093/biostatistics/4.2.249](https://doi.org/10.1093/biostatistics/4.2.249).
- [3] A J Jaffe and Z A Kaminsky. *FlowSorted.DLPFC.450k: Illumina HumanMethylation data on sorted frontal cortex cell populations*, r package version 1.0.0 edition.

- [4] Martin J Aryee, Andrew E Jaffe, Hector Corrada-Bravo, Christine Ladd-Acosta, Andrew P Feinberg, Kasper D Hansen, and Rafael A Irizarry. Minfi: a flexible and comprehensive bioconductor package for the analysis of infinium dna methylation microarrays. *Bioinformatics*, 30(10):1363–9, May 2014.
[doi:10.1093/bioinformatics/btu049](https://doi.org/10.1093/bioinformatics/btu049).