

# CQN (Conditional Quantile Normalization)

Kasper Daniel Hansen  
khansen@jhsph.edu

Zhijin Wu  
zhijin\_wu@brown.edu

Modified: August 8, 2012. Compiled: April 30, 2024

## Introduction

This package contains the CQN (conditional quantile normalization) method for normalizing RNA-seq datasets. This method is described in [1].

```
> library(cqn)
> library(scales)
```

## Data

As an example we use ten samples from Montgomery [2]. The data has been processed as described in [1]. First we have the region by sample count matrix

```
> data(montgomery.subset)
> dim(montgomery.subset)
```

```
[1] 23552    10
```

```
> montgomery.subset[1:4, 1:4]
```

|                  | NA06985 | NA06994 | NA07037 | NA10847 |
|------------------|---------|---------|---------|---------|
| ENSG000000000419 | 69      | 54      | 67      | 70      |
| ENSG000000000457 | 53      | 37      | 27      | 41      |
| ENSG000000000460 | 12      | 25      | 33      | 22      |
| ENSG000000000938 | 168     | 270     | 140     | 103     |

```
> colnames(montgomery.subset)
```

```
[1] "NA06985" "NA06994" "NA07037" "NA10847" "NA11920" "NA11918"
[7] "NA11931" "NA12003" "NA12006" "NA12287"
```

Because of (disc) space issues, We have removed all genes that have zero counts in all 10 samples. Next we have the *sizeFactors* which simply tells us how deep each sample was sequenced:

```
> data(sizeFactors.subset)
> sizeFactors.subset[1:4]

NA06985 NA06994 NA07037 NA10847
3107420 2388948 3087234 2852972
```

Finally, we have a matrix containing length and GC-content for each gene.

```
> data(uCovar)
> head(uCovar)

      length gccontent
ENSG000000000419   1207 0.3976802
ENSG000000000457   2861 0.4606781
ENSG000000000460   4912 0.4338355
ENSG000000000938   3524 0.5749149
ENSG000000000971   8214 0.3613343
ENSG000000001036   2590 0.4312741
```

Note that the row ordering of the count matrix is the same as the row ordering of the matrix containing length and GC-content and that the *sizeFactor* vector has the same column order as the count matrix. We can formally check this

```
> stopifnot(all(rownames(montgomery.subset) == rownames(uCovar)))
> stopifnot(colnames(montgomery.subset) == names(sizeFactors.subset))
```

## Normalization

The methodology is described in [1]. The main workhorse is the function `cqn` which fits the following model

$$\log_2(\text{RPM}) = s(x) + s(\log_2(\text{length}))$$

where  $x$  is some covariate,  $s$  are smooth functions (specifically natural cubic splines with 5 knots), and RPM are “reads per millions”. It is also possible to just fit a model like

$$\log_2(\text{RPKM}) = s(x)$$

In this model gene length is included as a known offset. This is done by using the `cqn` (`lengthMethod = "fixed"`). If this is done, and `lengths` is equal to 1000, it is equivalent to not using gene length at all.

The basic call to `cqn` is relatively easy, we need the count matrix, a vector of lengths, a vector of GC content and a vector of sizeFactors. Make sure that they all have the same ordering.

```
> cqn.subset <- cqn(montgomery.subset, lengths = uCovar$length,
+                   x = uCovar$gccontent, sizeFactors = sizeFactors.subset,
+                   verbose = TRUE)
```

```
RQ fit .....
SQN .
```

```
> cqn.subset
```

Call:


```
cqn(counts = montgomery.subset, x = uCovar$gccontent, lengths = uCovar$lengths,
     sizeFactors = sizeFactors.subset, verbose = TRUE)
```

```
Object of class 'cqn' with
  23552 regions
  10 samples
fitted using smooth length
```

This normalized matrix is similar, but not equivalent, to the data examined in [1]. The main differences are (1) in [1] we normalize 60 samples together, not 10 and (2) we have removed all genes with zero counts in all 10 samples.

We can examine plots of systematic effects by using `cqnplot`. The `n` argument refers to the systematic effect, `n=1` is always the covariate specified by the `x` argument above, while `n=2` is lengths.

```
> par(mfrow=c(1,2))
> cqnplot(cqn.subset, n = 1, xlab = "GC content", lty = 1, ylim = c(1,7))
> cqnplot(cqn.subset, n = 2, xlab = "length", lty = 1, ylim = c(1,7))
```



The normalized expression values are

```
> RPKM.cqn <- cqn.subset$y + cqn.subset$offset
> RPKM.cqn[1:4, 1:4]
```

|                  | NA06985  | NA06994  | NA07037  | NA10847  |
|------------------|----------|----------|----------|----------|
| ENSG000000000419 | 5.762584 | 5.569507 | 5.548028 | 5.976474 |
| ENSG000000000457 | 4.436529 | 4.109895 | 3.393383 | 4.139069 |
| ENSG000000000460 | 2.603214 | 3.443656 | 3.776678 | 3.067667 |
| ENSG000000000938 | 5.152593 | 6.084818 | 4.698305 | 4.281571 |

These values are on the  $\log_2$ -scale.

We can do a MA plot of these fold changes, and compare it to fold changes based on standard RPKM. First we compute the standard RPKM (on a  $\log_2$  scale):

```
> RPM <- sweep(log2(montgomery.subset + 1), 2, log2(sizeFactors.subset/10^6))
> RPKM.std <- sweep(RPM, 1, log2(uCovar$length / 10^3))
```

We now look at differential expression between two groups of samples. We use the same grouping as in [1], namely

```
> grp1 <- c("NA06985", "NA06994", "NA07037", "NA10847", "NA11920")
> grp2 <- c("NA11918", "NA11931", "NA12003", "NA12006", "NA12287")
```

We now do an MA-plot, but we only choose to plot genes with average standard  $\log_2$ -RPKM of  $\log_2(5)$  or greater, and we also form the M and A values:

```

> whGenes <- which(rowMeans(RPKM.std) >= 2 & uCovar$length >= 100)
> M.std <- rowMeans(RPKM.std[whGenes, grp1]) - rowMeans(RPKM.std[whGenes, grp2])
> A.std <- rowMeans(RPKM.std[whGenes,])
> M.cqn <- rowMeans(RPKM.cqn[whGenes, grp1]) - rowMeans(RPKM.cqn[whGenes, grp2])
> A.cqn <- rowMeans(RPKM.cqn[whGenes,])


```

Now we do the MA plots, with alpha-blending

```

> par(mfrow = c(1,2))
> plot(A.std, M.std, cex = 0.5, pch = 16, xlab = "A", ylab = "M",
+      main = "Standard RPKM", ylim = c(-4,4), xlim = c(0,12),
+      col = alpha("black", 0.25))
> plot(A.cqn, M.cqn, cex = 0.5, pch = 16, xlab = "A", ylab = "M",
+      main = "CQN normalized RPKM", ylim = c(-4,4), xlim = c(0,12),
+      col = alpha("black", 0.25))

```




We can also color the genes according to whether they have high/low GC-content. Here one needs to be careful, because of overplotting. One solution is to leave out all genes with intermediate GC content. We define high/low GC content as the 10% most extreme genes:

```

> par(mfrow = c(1,2))
> gccontent <- uCovar$gccontent[whGenes]
> whHigh <- which(gccontent > quantile(gccontent, 0.9))
> whLow <- which(gccontent < quantile(gccontent, 0.1))
> plot(A.std[whHigh], M.std[whHigh], cex = 0.2, pch = 16, xlab = "A",
+      ylab = "M", main = "Standard RPKM",

```

```
+      ylim = c(-4,4), xlim = c(0,12), col = "red")
> points(A.std[whLow], M.std[whLow], cex = 0.2, pch = 16, col = "blue")
> plot(A.cqn[whHigh], M.cqn[whHigh], cex = 0.2, pch = 16, xlab = "A",
+      ylab = "M", main = "CQN normalized RPKM",
+      ylim = c(-4,4), xlim = c(0,12), col = "red")
> points(A.cqn[whLow], M.cqn[whLow], cex = 0.2, pch = 16, col = "blue")
```



Note that genes/regions with very small counts should not be relied upon, even if the CQN normalized fold change are big. They should be filtered out using some kind of statistical test, good packages for this are *DESeq*[\[3\]](#) and *edgeR*[\[4, 5\]](#).

## Import into edgeR

First we construct a *DGEList*. In the `groups` argument we use that the first 5 samples (columns) in `montgomery.subset` is what we earlier called `grp1` and the last 5 samples (columns) are `grp2`.

```
> library(edgeR)
> d.mont <- DGEList(counts = montgomery.subset, lib.size = sizeFactors.subset,
+                  group = rep(c("grp1", "grp2"), each = 5), genes = uCova)
```

In this object we cannot (unfortunately, yet) also store the computed offsets. Since we will use the offsets computed by *cqn*, there is no need to normalize using the normalization tools from *edgeR*, such as `calcNormFactors`. Also, as is clearly described in the *edgeR* user's guide, the `lib.size` is unnecessary, since we plan to use the offsets computed from *cqn*.

However, we need to use the component `glm.offset` which is on the natural logarithmic scale and also includes correcting for `sizeFactors`. It is possible to include the offset directly into the `DGEList`, by post-processing the output like

```
> ## Not run
> d.mont$offset <- cqn.subset$glm.offset
```

Using *edgeR* is well described in the user's guide, and we refer to that document for further information. The analysis presented below should be thought of as an example, and not necessarily the best analysis of this data.

The first step is estimating the dispersion parameter(s). Several methods exist, such as `estimateGLMCommonDisp` or `estimateTagwiseDisp`. We also need to setup a design matrix, which is particularly simple for this two group comparison. Further information about constructing design matrices may be found in both the *edgeR* user's guide and the *limma* user's guide.

```
> design <- model.matrix(~ d.mont$sample$group)
> d.mont$offset <- cqn.subset$glm.offset
> d.mont.cqn <- estimateGLMCommonDisp(d.mont, design = design)
```

After fitting the dispersion parameter(s), we need to fit the model, and do a test for significance of the parameter of interest. With this design matrix, there are two coefficients. The first coefficient is just an intercept (overall level of expression for the gene) and it is (usually) not meaningful to test for this effect. Instead, the interesting coefficient is the second one that encodes a group difference.

```
> efit.cqn <- glmFit(d.mont.cqn, design = design)
> elrt.cqn <- glmLRT(efit.cqn, coef = 2)
> topTags(elrt.cqn, n = 2)
```

```
Coefficient: d.mont$sample$groupgrp2
              length gccontent      logFC    logCPM      LR
ENSG00000211642    365 0.5835616 -10.29393  6.362880 126.0812
ENSG00000211660    411 0.5888078 -10.11432  5.999853 120.4073
              PValue      FDR
ENSG00000211642 2.951582e-29 6.951566e-25
ENSG00000211660 5.151951e-28 6.066938e-24
```

`topTags` shows (per default) the "top 10" genes. In this case, since we have biological replicates and just a random group structure, we would expect no differentially expression genes. Instead we get

```
> summary(decideTestsDGE(elrt.cqn))
```

|        | d.mont\$sample\$groupgrp2 |
|--------|---------------------------|
| Down   | 148                       |
| NotSig | 22965                     |
| Up     | 439                       |

significantly differentially expressed at an FDR (false discovery rate) of 5%. We may contrast this with the result of not using *cqn*:

```
> d.mont.std <- estimateGLMCommonDisp(d.mont, design = design)
> efit.std <- glmFit(d.mont.std, design = design)
> elrt.std <- glmLRT(efit.std, coef = 2)
> summary(decideTestsDGE(elrt.std))
```

|        | d.mont\$sample\$groupgrp2 |
|--------|---------------------------|
| Down   | 148                       |
| NotSig | 22965                     |
| Up     | 439                       |

In this evaluation, it is not clear that using CQN is better.

What is arguably as important is that we achieve a much better estimation of the fold change using *cqn*.

## Question and Answers

### Can I run *cqn()* on only 1 sample?

CQN is meant to normalize several samples together. It is not clear that it makes sense at all to use this normalization technique on a single sample. But it is possible.

### Can I use this for small RNA-seq (microRNAs)?

We do not have personal experience with using CQN to normalize small RNA sequencing data. However, we believe it might be beneficial. As always, it is *highly* recommended to evaluate whether it is necessary and beneficial.

One special aspect of small RNAs is that they all have very similar length. Fitting a model with a smooth effect of gene length might very well lead to mathematical instability (you get an error). This can be avoided by using the argument `lengthMethod = "fixed"` which just divides the gene counts by the gene length instead of using a smooth function. Additionally, it may be coupled with setting `lengths = 1` which completely removes gene length from the model.



### Could it be true that genes with higher GC content are higher expressed?

It has been suggested that genes that are either extremely high or extremely low expressed are under some form of selection leading to “extreme” GC content. What CQN does, is making the effect of GC content comparable across samples, and we show in [1] that this leads to improved inference. It also flattens the effect of GC content on gene expression, but we believe this is better than having the effect of GC content depend on the sample.

### Does cqn remove batch effects?

No, unless a batch effect only (or mainly) affects your measurements through GC content. We believe that the sample-specific effect of GC content on gene expression is a kind of batch effect, but is unlikely to be the only one. CQN does normalize your RNA-seq data in the same way that say quantile normalization normalizes microarray data, but such normalization does not remove batch effects.

### I don’t understand the difference between offset and glm.offset?

This comes from a historical error. In our paper, we use the quantity

```
> cqn$y + cqn$offset
```

as the CQN-corrected estimated expression measures. However, the offset quantity is on the wrong scale for inclusion into a GLM-type model (like edgeR or DEseq2). For this purpose, use glm.offset. We have kept the original naming in order to achieve backwards compatibility.

## SessionInfo

- R version 4.4.0 Patched (2024-04-24 r86482), aarch64-apple-darwin20
- Locale: C/en\_US.UTF-8/en\_US.UTF-8/C/en\_US.UTF-8/en\_US.UTF-8
- Time zone: America/New\_York
- TZcode source: internal
- Running under: macOS Ventura 13.6.6
- Matrix products: default
- BLAS:  
/Library/Frameworks/R.framework/Versions/4.4-arm64/Resources/lib/libRblas.dylib

- LAPACK:  
/Library/Frameworks/R.framework/Versions/4.4-arm64/Resources/lib/libRlapack.dylib;  
; LAPACK version 3.12.0
- Base packages: base, datasets, grDevices, graphics, methods, splines, stats, utils
- Other packages: SparseM 1.81, cqn 1.49.0, edgeR 4.1.33, limma 3.59.10, mclust 6.1, nor1mix 1.3-3, preprocessCore 1.65.0, quantreg 5.97, scales 1.3.0
- Loaded via a namespace (and not attached): MASS 7.3-60.2, Matrix 1.7-0, MatrixModels 0.5-3, R6 2.5.1, Rcpp 1.0.12, cli 3.6.2, colorspace 2.1-0, compiler 4.4.0, farver 2.1.1, glue 1.7.0, grid 4.4.0, lattice 0.22-6, lifecycle 1.0.4, locfit 1.5-9.9, munsell 0.5.1, rlang 1.1.3, statmod 1.5.0, survival 3.6-4, tools 4.4.0

## References

- [1] KD Hansen, RA Irizarry, and Z Wu. Removing technical variability in RNA-seq data using conditional quantile normalization. *Biostatistics* 2012, **13**(2), 204–216. DOI: [10.1093/biostatistics/kxr054](https://doi.org/10.1093/biostatistics/kxr054).
- [2] SB Montgomery *et al.* Transcriptome genetics using second generation sequencing in a Caucasian population. *Nature* 2010, **464**, 773–777. DOI: [10.1038/nature08903](https://doi.org/10.1038/nature08903)
- [3] S Anders and W Huber. Differential expression analysis for sequence count data. *Genome Biology* 2010, **11**(10), R106. DOI: [10.1186/gb-2010-11-10-r106](https://doi.org/10.1186/gb-2010-11-10-r106).
- [4] MD Robinson, DJ McCarthy, GK Smyth. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* 2010, **26**(1), 139–140. DOI: [10.1093/bioinformatics/btp616](https://doi.org/10.1093/bioinformatics/btp616).
- [5] DJ McCarthy, Y Chen, GK Smyth. Differential expression analysis of multifactor RNA-Seq experiments with respect to biological variation. *Nucleic Acids Research* 2012, **40**, 4288–4297. DOI: [10.1093/nar/gks042](https://doi.org/10.1093/nar/gks042).