

Introduction: microarray quality assessment with arrayQualityMetrics

Audrey Kauffmann, Wolfgang Huber

April 4, 2013

Contents

1	Basic use	3
1.1	Affymetrix data - before preprocessing	3
1.2	Affymetrix data - after preprocessing	3
1.3	ExpressionSet and ExpressionSetIllumina	4
1.4	Two colour arrays, NChannelSet, RGList, MAList	4
1.5	Loading data from ArrayExpress	4
2	Making the report more informative by adding a factor of interest	4
3	Extended use	5
3.1	Spatial layout of the array	5
3.2	Mapping of the reporters	5
3.3	RNA quality	6

Introduction

The *arrayQualityMetrics* package produces, through a single function call, a comprehensive HTML report of *quality metrics* about a microarray dataset [1, 2, 3]. The quality metrics are mainly on the *per array* level, i.e. they can be used to assess the relative quality of different arrays within a dataset. Some of the metrics can also be used to diagnose batch effects, and thus the quality of the overall dataset.

The report can be extended to contain further diagnostics through additional arguments, and we will see examples for this in Section 3.

The aim of the *arrayQualityMetrics* package is to produce information that is relevant for your decision making - not, to make the decision. It will often be applied to two, somewhat distinct, use cases: (i) assessing quality of a “raw” dataset, in order to get feedback on the experimental procedures that produced the data; (ii) assessing quality of a normalised dataset, in order to decide whether and how to use the dataset (or subsets of arrays in it) for subsequent data analysis.

Different types of microarray data (one colour, two colour, Affymetrix, Illumina) are represented by different object classes in Bioconductor. The function `arrayQualityMetrics` will work in the same way for all of them. Further information about its arguments can be found in its manual page.




Figure 1: *Left*: An example plot from the report. The plot shows the arrays (points) in a two-dimensional plot area spanned by the first two axes of a principal component analysis (PCA). By moving the mouse over the points, the corresponding array’s metadata is displayed in the table to the right of the plot. By clicking on a point, it can be selected or deselected. Selected arrays are indicated by larger points or wider lines in the plots and by ticked checkboxes in the array table shown in the *right* panel. Arrays can also be (de)selected by clicking the checkboxes. Initially, when the report is loaded (or reloaded) by the browser, all arrays are selected that were called outliers by at least one criterion.

When the function `arrayQualityMetrics` is finished, a report is produced in the directory specified by the function’s `outdir` argument. By default, a directory with a suitable name is created in the current working directory. This directory contains an HTML page `index.html` that can be opened by a browser. The report contains a series of plots explained by text. Some of the plots are interactive (see Figure 1). Technically, this is achieved by the use of SVG (scalable vector graphics) and JavaScript, and it requires that you use a recent (HTML5 capable) web browser¹. Other plots, where interactivity is less relevant, are provided as bitmaps (PNG format) and are also linked to PDF files that provide high resolution versions e. g. for publication.

Plus (+) or minus (-) symbols at the begin of different section headings of the report (as in the left panel of Figure 1) indicate that you can show or hide these sections by clicking on the heading. After (re)loading, all sections are shown except for the *Outlier detection* barplots, which are hidden and can be expanded by clicking on them.

Metadata about the arrays is shown at the top of the report as a table (see Figure 1). It is extracted from the `phenoData` slot of the data object supplied to `arrayQualityMetrics`. It can be useful to adjust the contents this slot before producing the report, and to make sure it contains the right quantity of information to make an informative report - not too much, not too little.

In the case of *AffyBatch* input, some Affymetrix specific sections are added to the standard report. Also for other types of arrays, sections can be added to the standard report if certain metadata are present in the input object (see Section 3).

The function `arrayQualityMetrics` also produces an R object (essentially, a big list) with all the information contained in the report, and this object can be used by downstream tools

¹If in doubt, please see the notes about browser compatibility at the top of the report; or contact the Bioconductor mailing list.

for programmatic analysis of the report. This is discussed in the vignette *Advanced topics: Customizing arrayQualityMetrics reports and programmatic processing of the output*

1 Basic use

1.1 Affymetrix data - before preprocessing

If you are working with Affymetrix GeneChips, an *AffyBatch* object is the most appropriate way to import your raw data into Bioconductor. Starting from CEL files, this is typically done using the function `ReadAffy` from the *affy* package². Here, we use the dataset *MLL.A*, an object of class *AffyBatch* provided in the data package *ALLMLL*.

```
> library("ALLMLL")
> data("MLL.A")
```

Now that the data are loaded, we can call `arrayQualityMetrics`³.

```
> library("arrayQualityMetrics")
> arrayQualityMetrics(expressionset = MLL.A[, 1:5],
+                      outdir = "Report_for_MLL_A",
+                      force = TRUE,
+                      do.logtransform = TRUE)
```

This is the simplest way of calling the function. We give a name to the directory (`outdir`) and we overwrite the possibly existing files of this directory (`force`). Finally, we set `do.logtransform` to logarithm transform the intensities. You can then view the report by directing your browser to the file `index.html` in the directory whose name is indicated by `outdir`.

1.2 Affymetrix data - after preprocessing

We can call the RMA algorithm on *MLL.A* to obtain a preprocessed dataset. The preprocessing includes background correction, between array intensity adjustment (normalisation) and probeset summarisation. The resulting object `nMLL` is of class *ExpressionSet* and contains one value (expression estimate) for each gene for each array.

```
> nMLL = rma(MLL.A)
```

We can then call again the function `arrayQualityMetrics`.

```
> arrayQualityMetrics(expressionset = nMLL,
+                      outdir = "Report_for_nMLL",
+                      force = TRUE)
```

We do not need to set `do.logtransform` as after `rma` the data are already logarithm transformed.

²For more information on how to produce an *AffyBatch* from your data, please see the documentation of the *affy* package.

³For this vignette, in order to save computation time, we only call the function on the first 5 arrays; in your own application, you can call it on the complete data object.

1.3 ExpressionSet and ExpressionSetIllumina

If you are working on one colour arrays other than Affymetrix genechips, you can load your data into Bioconductor as an *ExpressionSet* object ⁴, or if you work with Illumina data and the *beadarray* package, as an *ExpressionSetIllumina* object. You can then proceed exactly as above.

1.4 Two colour arrays, NChannelSet, RGList, MAList

The package *limma* imports a wide range of data formats used for two colour arrays and produces objects of class *RGList* or *MAList*. When presented with an object of these classes, *arrayQualityMetrics* tries to convert them into an *NChannelSet* and then proceeds with calling its *NChannelSet* method.

Alternatively, you can create an *NChannelSet* to contain your data “from scratch”. The documentation of the *Biobase* package gives instructions on how to do so.

The *arrayQualityMetrics* function expects the *assayData* slot of the *NChannelSet* object to contain the elements **R** and **G**, for the “red” and the “green” intensities. Optionally, it can contain elements **Rb** and **Gb** for associated “background” intensities. As an alternative to all that, the *arrayQualityMetrics* function also accepts *NChannelSet* objects with a single slot *exprs*, and will then simply behave like it does for (single-colour) *ExpressionSet* objects.

As an example, we consider the dataset **CC14** from the data package *CC14* and normalize it using the variance stabilization method available in the package *vsn*.

```
> library("vsnp")
> library("CC14")
> data("CC14")
> nCC14 = justvsnp(CC14, subsample = 15000)
> arrayQualityMetrics(expressionset = nCC14,
+                       outdir = "Report_for_nCC14",
+                       force = TRUE)
```

1.5 Loading data from ArrayExpress

You can use the *ArrayExpress* package [4] to download datasets from the EBI’s ArrayExpress database. The resulting *ExpressionSet*, *AffyBatch* or *NChannelSet* objects can be directly fed into *arrayQualityMetrics*.

2 Making the report more informative by adding a factor of interest

A useful feature of *arrayQualityMetrics* is the possibility to show the results in the context of an experimental factor of interest, i.e. a categorical variable associated with the arrays such as *hybridisation date*, *treatment level* or *replicate number*. Specifying a factor does *not* change how the quality metrics are computed. By setting the argument *intgroup* to contain the names of one or multiple columns of the data object’s *phenoData* slot⁵, a bar on the

⁴See the documentation of the *Biobase* package.

⁵This is where Bioconductor objects store array annotation

side of the heatmap with colours representing the respective factors is added. Similarly, the colours of the boxplots and density plots reflect the levels of the first of the factors named by `intgroup`.

We use the *nMLL* example again, and create artificial array metadata factors `condition` and `batch` (see Section 3.3 for a more realistic example).

```
> pData(nMLL)$condition = rep(letters[1:4], times = 5)
> pData(nMLL)$batch = rep(paste(1:4), each = 5)

> arrayQualityMetrics(expressionset = nMLL,
+                       outdir = "Report_for_nMLL_with_factors",
+                       force = TRUE,
+                       intgroup = c("condition", "batch"))
```

3 Extended use

Some of the quality metrics that the package can compute require specific information about the features on the arrays. To use these, you need to make sure that this information is provided in your input object. We use the *nCC14* example again.

3.1 Spatial layout of the array

To plot the spatial distributions of the intensities of the arrays, `arrayQualityMetrics` needs the spatial coordinates of the features on the chip. For *AffyBatch* or *BeadLevelList*, this information is automatically available without further user input. For the other types of objects, two columns corresponding to *X* and *Y* coordinates of the features are required in the `featureData` slot of the object. These columns should be named "X" and "Y". If the arrays are split into blocks, rows and columns, then the function `addXYfromGAL` (please check its manual page for details) can be used to convert the row, column and blocks indices into absolute "X" and "Y" coordinates on the array. In the example of the dataset *CC14*, the coordinates of the spots are in the columns named "Row" and "Column" of the `featureData` (the slot of the object containing the annotation of the probes). We copy this information into columns named "X" and "Y" respectively

```
> featureData(nCC14)$X = featureData(nCC14)$Row
> featureData(nCC14)$Y = featureData(nCC14)$Column
```

The next call to `arrayQualityMetrics` with this refined version of *nCC14* (see Section 3.3) will now include this information in the report, and the spatial distribution of the intensities will be shown.

3.2 Mapping of the reporters

The report can also include an assessment of the effect of the target mapping of the reporters. You can define a `featureData` column named `hasTarget` that indicates, by logical `TRUE`, if the reporter matches a known transcript, and by `FALSE`, if not. In the *CC14* example, many of the reporter names are RefSeq identifiers, while others are not. Thus, we let `hasTarget` indicate whether the name begins with "NM".

```
> featureData(nCC14)$hasTarget = (regexpr("^NM", featureData(nCC14)$Name) > 0)
> table(featureData(nCC14)$hasTarget)
```

```
FALSE TRUE
33296 10332
```

The next call to `arrayQualityMetrics` with this refined version of `nCC14` (see Section 3.3) will now include this information in the report, and the spatial distribution of the intensities will be shown.

3.3 RNA quality

The RNA hybridized to the arrays in the *CC14* dataset was intentionally made to good, medium or poor quality, and this is recorded by a so-called RIN value (see *CC14* vignette).

```
> pd = pData(CC14)
> rownames(pd) = NULL
> pd
```

	Cy3	Cy5	RIN.Cy3	RIN.Cy5
1	DMSO	CC14	9.0	9.7
2	DMSO	CC14	9.0	5.0
3	DMSO	CC14	9.0	2.5
4	DMSO	CC14	9.0	2.5
5	CC14	DMSO	9.7	9.0
6	CC14	DMSO	5.0	9.0
7	CC14	DMSO	2.5	9.0
8	DMSO	CC14	9.0	9.7
9	CC14	DMSO	9.7	9.0
10	CC14	DMSO	5.0	9.0
11	CC14	DMSO	2.5	9.0
12	DMSO	CC14	9.0	9.7
13	DMSO	CC14	9.0	5.0
14	DMSO	CC14	9.0	5.0
15	DMSO	CC14	9.0	2.5
16	CC14	DMSO	9.7	9.0
17	CC14	DMSO	5.0	9.0
18	CC14	DMSO	2.5	9.0

The RIN is always 9 for the reference (DMSO), the relevant value is that for the test sample (CC14).

```
> RIN = with(pd, ifelse( Cy3=="CC14", RIN.Cy3, RIN.Cy5))
> fRIN = factor(RIN)
> levels(fRIN) = c("poor", "medium", "good")
> pData(nCC14)$"RNA-integrity" = fRIN
```

Now we can use this to set the argument `intgroup` when calling the function `arrayQualityMetrics`.

```
> arrayQualityMetrics(expressionset = nCC14,
+                       outdir = "Report_for_nCC14_with_RIN",
+                       force = TRUE,
+                       intgroup = "RNA-integrity")
```

Boxplots, PCA plot and heatmap in the report will now indicate the values of the factor RNA-integrity for each array.

Session Info

- R version 3.0.0 (2013-04-03), 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=C, 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: ALLMLL 1.2.13, AnnotationDbi 1.22.0, Biobase 2.20.0, BiocGenerics 0.6.0, CC14 1.0.13, affy 1.38.0, arrayQualityMetrics 3.16.0, hgu133acdf 2.12.0, limma 3.16.0, vsn 3.28.0
- Loaded via a namespace (and not attached): BeadDataPackR 1.12.0, BiocInstaller 1.10.0, Biostrings 2.28.0, Cairo 1.5-2, DBI 0.2-5, Hmisc 3.10-1, IRanges 1.18.0, KernSmooth 2.23-10, RColorBrewer 1.0-5, RSQLite 0.11.2, SVGAnnotation 0.93-1, XML 3.96-1.1, affyPLM 1.36.0, affyio 1.28.0, annotate 1.38.0, beadarray 2.10.0, cluster 1.14.4, colorspace 1.2-1, gcrma 2.32.0, genefilter 1.42.0, grid 3.0.0, hwriter 1.3, lattice 0.20-15, latticeExtra 0.6-24, plyr 1.8, preprocessCore 1.22.0, reshape2 1.2.2, setRNG 2011.11-2, splines 3.0.0, stats4 3.0.0, stringr 0.6.2, survival 2.37-4, tools 3.0.0, xtable 1.7-1, zlibbioc 1.6.0

References

- [1] Audrey Kauffmann, Robert Gentleman, and Wolfgang Huber. arrayQualityMetrics - a Bioconductor package for quality assessment of microarray data. *Bioinformatics*, 25:415–416, 2009.
- [2] Audrey Kauffmann and Wolfgang Huber. Microarray data quality control improves the detection of differentially expressed genes. *Genomics*, 95:138–142, 2010.
- [3] Matthew N. McCall, Peter N. Murakami, and Rafael A. Irizarry. Assessing microarray quality. Technical report, Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, Baltimore, MD, 2010.
- [4] Audrey Kauffmann, Tim F. Rayner, Helen Parkinson, Misha Kapushesky, Margus Lukk, Alivs Brazma, and Wolfgang Huber. Importing ArrayExpress datasets into R/Bioconductor. *Bioinformatics*, 25:2092–2094, 2009.