

Analysis of bead-level data using `beadarray`

Mark Dunning and Matt Ritchie

April 21, 2008

Introduction

Illumina have created an alternative microarray technology (BeadArray) based on randomly arranged beads. A specific oligonucleotide sequence is assigned to each *bead type*, which is replicated about 30 times on an array. A series of decoding hybridisations is used to identify each bead on the array. The high degree of replication makes robust measurements for each bead type possible.

BeadArrays are used in many applications, including gene expression studies, SNP genotyping and methylation profiling and are processed in parallel as a Sentrix Array Matrix (SAM) or BeadChip. A SAM is a plate of 96 uniquely prepared hexagonal BeadArrays, each of which contains around 1,500 bead types. The BeadChip technology comprises a series of rectangular strips on a slide with each strip containing about 24,000 bead types. For example, there are six pairs of strips on each Human-6 BeadChip. Depending on the particular assay used, the data from a BeadArray may be single channel or two-colour.

The data from Illumina BeadArrays is available in different formats. In this guide we describe how to read raw *bead-level data* from an expression array. Although this example is from a single-channel technology, the same procedure can be applied to data from two-colour platforms.

The second format is produced by Illumina's BeadStudio software. We refer to this output as *bead-summary data* as these files contain summary intensities for each bead type on each array. For details on how to use the `beadarray` package to read in and process this kind of data, refer to the bead-summary user's guide which can be launched with the following command

```
> library(beadarray)
> beadarrayUsersGuide(topic = "beadsummary")
```

1 Obtaining bead-level data

The raw images and text files required to perform a bead-level analysis are produced by Illumina's BeadScan software. To modify BeadScan's default settings to obtain bead-level data, edit the `settings.xml` file in the main BeadScan directory to include the lines

```
<SaveTextFiles>true</SaveTextFiles>
<ExcludeOutliers>>false</ExcludeOutliers>
and
<IncludeXY>true</IncludeXY>.
```

For further details see the instructions at <http://www.compbio.group.cam.ac.uk/Resources/illumina/>

2 Importing bead-level data

The example data used in this guide (`BeadLevelExample.zip`, 800MB) can be downloaded from:

<http://www.compbio.group.cam.ac.uk/Resources/illumina/BeadLevelExample.zip>

These arrays were part of a pilot project using Human-6 version 1 BeadChips. In this example data set we have 4 different samples, two of which were supplied by Illumina (IC and IH), and two from cell lines (P and Norm). This BeadChip is a subset of the experiment analysed in the bead-summary user's guide.

2.1 Description of files

To read in bead-level data using `beadarray` you will need several files produced by Illumina's BeadScan software. We briefly describe these files below.

- text files (*required*) - a `.txt` or `.csv` file for each strip or hexagon which stores the position, identity and intensity (after background correction) of each bead. These files are usually named `ChipNumber_Array_Strip.csv` (e.g. `1475542113_A_1.csv`) and are required because of the random arrangement of probes on the array surface, which is unique for each BeadArray.
- TIFFs (*optional*) - 1 (single channel) or 2 (two-colour) for each strip on a BeadChip, These are usually named using the convention `ChipNumber_Array_Strip_Channel.tif`, e.g. `1475542113_A_1_Grn.tif`, is the Cy3 (green) image for strip 1 from array A on BeadChip 1475542113. Cy5 (red) images end with the extension `_Red.tif`.
- bead manifest file (*optional*) - contains annotation information about each bead type on the array. See the file `Human_WG6.csv` for an example. This file can be used to obtain further information about the transcript each bead type targets. Different arrays have different manifest files. Version 2 and 3 BeadChips store this information in a slightly different `.bgx` format. Other platforms have analogous files, such as the `.opa` for SNP BeadArrays.
- targets file (*recommended*) - contains sample information for each strip/array. See the file `targets.txt` for an example.
- metrics file (*optional*) - one for each each BeadChip or SAM, usually named `Metrics.txt` which contains summary information about intensity, the amount of saturation, focus and registration on the image(s) from each strip or hexagon. This data from this file can be useful for quality control purposes. In this example, the metrics file is not available.

The following code can be used to read the example data into R (provided that the contents of `BeadLevelExample.zip` have been extracted to the current working directory).

```
> library(beadarray)
> targets = read.table("targets.txt", sep = "\t",
+   header = TRUE, as.is = TRUE)
> targets
> BLData = readIllumina(arrayNames = targets$ArrayName,
+   textType = ".csv", targets = targets,
+   backgroundMethod = "none")
```

The function `readIllumina` implements the image processing steps used by Illumina when `useImages=TRUE`, however, both the sharpening and background correction steps are optional.

We estimate a background for each bead by taking the average of the 5 dimmest pixels in a local area around each bead centre. The background value is stored separately and is not subtracted from the foreground value automatically (as occurs in BeadScan output) when `backgroundMethod="none"`. When `useImages=FALSE`, the image analysis is not reproduced in R, with the intensities read-in directly from the text files. This option has the benefit of saving on memory, although the background intensities are no longer available, as the intensities stored in these text files have already been background corrected by BeadScan.

By default, `readIllumina` will read all arrays in the current working directory for which both text files and TIFFs can be found. For two-colour experiments you will need to set `singleChannel=FALSE` and have both red and green images available.

On version 1 chips, the 2 strips for each array have a different set of bead types attached. Version 2 and 3 BeadChips have each bead type represented on each strip. Access to the raw data allows the two strips to be analysed separately.

3 The BLData object

Once imported, the bead-level data is stored in an object of class *BeadLevelList*. This class can handle raw data from both single channel and two-colour BeadArrays. Due to the random nature of the technology, each array generally has a variable number of rows of intensity data, and we use an R environment variable to store this information in a memory efficient way.

The *BeadLevelList* class contains a number of slots useful for describing Illumina data. The data from each strip/array can be accessed by subsetting the *beadData* slot by the name of the array. The `$` operator can then be used to extract the appropriate column from the *data.frame*. Alternatively, the `getArrayData` function can be used.

```
> class(BLData)

[1] "BeadLevelList"
attr(,"package")
[1] "beadarray"

> slotNames(BLData)

[1] "beadData"      "phenoData"      "arrayInfo"
[4] "annotation"    "beadAnno"       "scanMetrics"

> an = arrayNames(BLData)
> an

[1] "1475542113_A_1" "1475542113_A_2"
[3] "1475542113_B_1" "1475542113_B_2"
[5] "1475542113_C_1" "1475542113_C_2"
[7] "1475542113_D_1" "1475542113_D_2"
[9] "1475542113_E_1" "1475542113_E_2"
[11] "1475542113_F_1" "1475542113_F_2"

> names(BLData@beadData[[an[1]]])

[1] "ProbeID" "G"      "Gb"      "GrnX"
[5] "GrnY"
```

```

> BLData[[an[1]]]$G[1:5]

[1] 647.1598 1291.8708 4646.7958 994.2587
[5] 716.0407

> BLData[[an[2]]]$Gb[1:5]

[1] 636 634 635 637 639

> pData(BLData)

      ArrayName SampleID  Origin
1 1475542113_A_1      IC Illumina
2 1475542113_A_2      IC Illumina
3 1475542113_B_1      IH Illumina
4 1475542113_B_2      IH Illumina
5 1475542113_C_1      IC Illumina
6 1475542113_C_2      IC Illumina
7 1475542113_D_1      P   Breast
8 1475542113_D_2      P   Breast
9 1475542113_E_1      P   Breast
10 1475542113_E_2     P   Breast
11 1475542113_F_1     Norm Normal
12 1475542113_F_2     Norm Normal

> getArrayData(BLData, array = 1, what = "G",
+   log = FALSE)[1:5]

[1] 647.1598 1291.8708 4646.7958 994.2587
[5] 716.0407

> getArrayData(BLData, array = 2, what = "Gb",
+   log = FALSE)[1:5]

[1] 636 634 635 637 639

```

4 Quality assessment

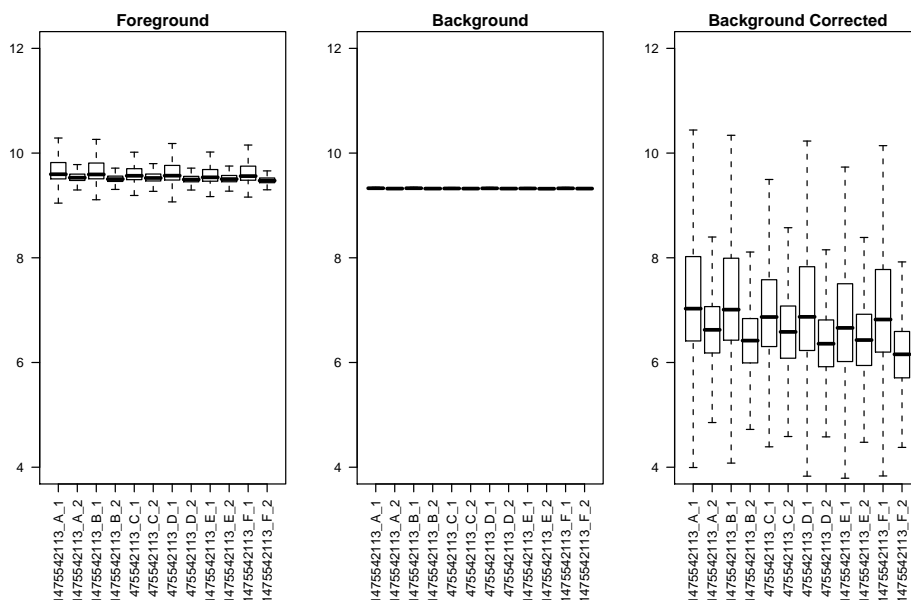
Boxplots can be used to compare foreground and background intensities between strips. In this example we can see that the first strip from each array has a systematically higher intensity than the second strip. This is related to the design of version 1 chips, which contain probes for RefSeq transcripts on the first strip, and probes for less well annotated transcripts on the second strip. Besides this obvious effect, there appears to be a subtle decrease in intensity from top to bottom on this BeadChip (array A is at the top of the chip, through to array F at the bottom of the chip). Notice that the background level appears to be virtually constant both for beads on the same array and between arrays.

Background correction can be performed by the `backgroundCorrect` function if it has not already been carried out when the data was read into R by `readIllumina`. The default setting in both functions is to subtract the local background estimate from the foreground of each bead, as per BeadScan. Other options are available by changing the `method` argument in `backgroundCorrect` or the `background-Method` argument in `readIllumina`.

```

> par(mfrow = c(1, 3), mai = c(2, 0.5, 0.2,
+   0.1))
> boxplotBeads(BLData, las = 2, outline = FALSE,
+   ylim = c(4, 12), main = "Foreground")
> boxplotBeads(BLData, las = 2, whatToPlot = "Gb",
+   outline = FALSE, ylim = c(4, 12),
+   main = "Background")
> BLData.bc = backgroundCorrect(BLData,
+   method = "subtract")
> boxplotBeads(BLData.bc, las = 2, outline = FALSE,
+   ylim = c(4, 12), main = "Background Corrected")

```



The `whatToPlot` argument of `boxplotBeads` controls which intensities are plotted for each bead. Options are `G`, `Gb` and `residG` (Cy3 residuals) for single channel data and `R`, `Rb`, `residR`, `M` (log-ratios) `residM` or `A` (average log-intensities) for two-colour data.


We can also produce boxplots of bead intensities using `plotBeadIntensities`. This function takes a list of `ProbeIDs` and `strips/arrays` as arguments and produces a boxplot for each bead type on each strip/array. `ProbeIDs` on the same strip/array are grouped together.

In the following code we show how to plot the intensities of ten different bead types on two separate arrays in the experiment. For this particular example we have to remember that all odd-numbered strips in the experiment contain different probes to the even-numbered strips, therefore we plot the intensities of the beads on the first and third strips (corresponding to the first and second arrays).

```

> ids = unique(BLData.bc[[an[1]]]$ProbeID)[2:11]
> ids
> ProbeCols = rainbow(10)
> plotBeadIntensities(BLData.bc, arrays = c(1,
+   3), ProbeIDs = ids, ProbeCols = ProbeCols,
+   log = TRUE, ylim = c(5, 16))


```



Spatial artefacts on the array surface can occur from mis-handling or scanning problems. Image plots can be used to identify these artefacts. This kind of visualisation is not possible when using the summarised BeadStudio output, as the summary values are averaged over spatial positions. Image plots in R are also more convenient than scrutinising the original TIFFs, as multiple arrays can be visualised on the one page.

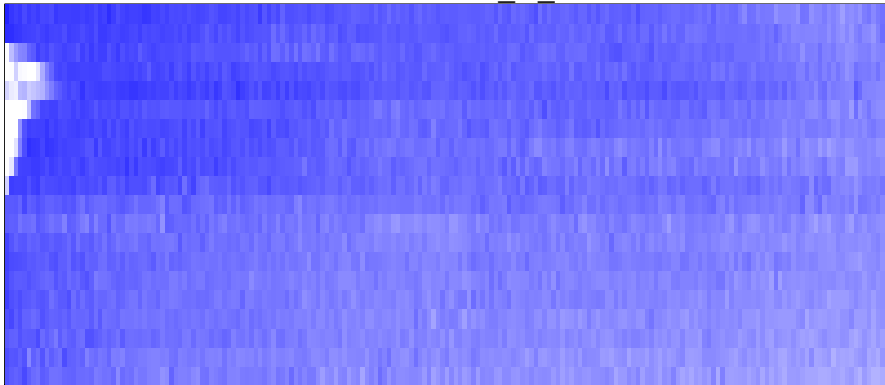
```
> par(mfrow = c(2, 1))
> for (i in 1:2) {
+   imageplot(BLData.bc, array = i, nrow = 20,
+             ncol = 200, zlim = range(5, 8),
+             what = "G", main = an[i])
+ }
```

1475542113_A_1



z-range 6.4 to 8.4 (saturation 5, 8)

1475542113_A_2



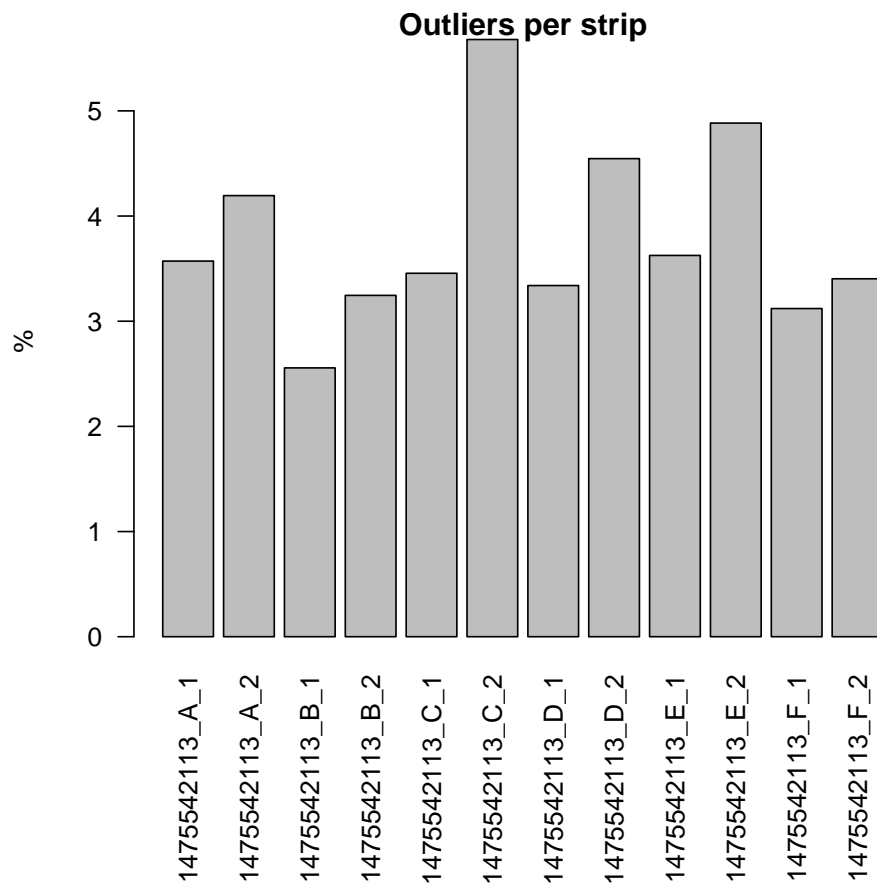
z-range 4.2 to 7.6 (saturation 5, 8)

In the `imageplot` function, the argument `whatToPlot` is used to choose the quantity to display. For single channel data, `whatToPlot` can be set to `G`, `Gb` or `residG` to plot the Cy3 foreground, background or foreground residuals respectively. For two-colour data, the Cy5 foreground (`R`), background (`Rb`), log-ratios (`M`), average log-intensities (`A`) and residuals (`residR`, `residM`) can be plotted by changing `whatToPlot`. Because of the high number of beads on each array, the `imageplot` function maps a grid of size specified by the `nrow` and `ncol` arguments onto the array surface and averages the intensities of the beads within each section of the grid.

Recall that the BeadArray technology includes around 30 replicates of each bead type on every array. By default, BeadStudio removes outliers greater than 3 median absolute deviations (MADs) from the median prior to calculating the bead summary values. The `findAllOutliers` function identifies outlier beads on a given strip/array using this criterion, by returning a vector of row indices of the outliers. Typically we find that the number of outliers on a strip/array is less than 10% and both the number and location of outliers (which can be plotted using the `plotBeadLocations` function) can be used as a diagnostic tool. An excessively high number of outliers can be indicative of a poor quality array.

```
> outliers = NULL
> for (i in 1:12) {
+   outliers[i] = length(findAllOutliers(BLData.bc,
+   array = i))
+ }
> par(mai = c(2, 1, 0.2, 0.1))
```

```
> barplot(outliers/numBeads(BLData.bc) *
+ 100, main = "Outliers per strip",
+ ylab = "%", las = 2, names = an)
```



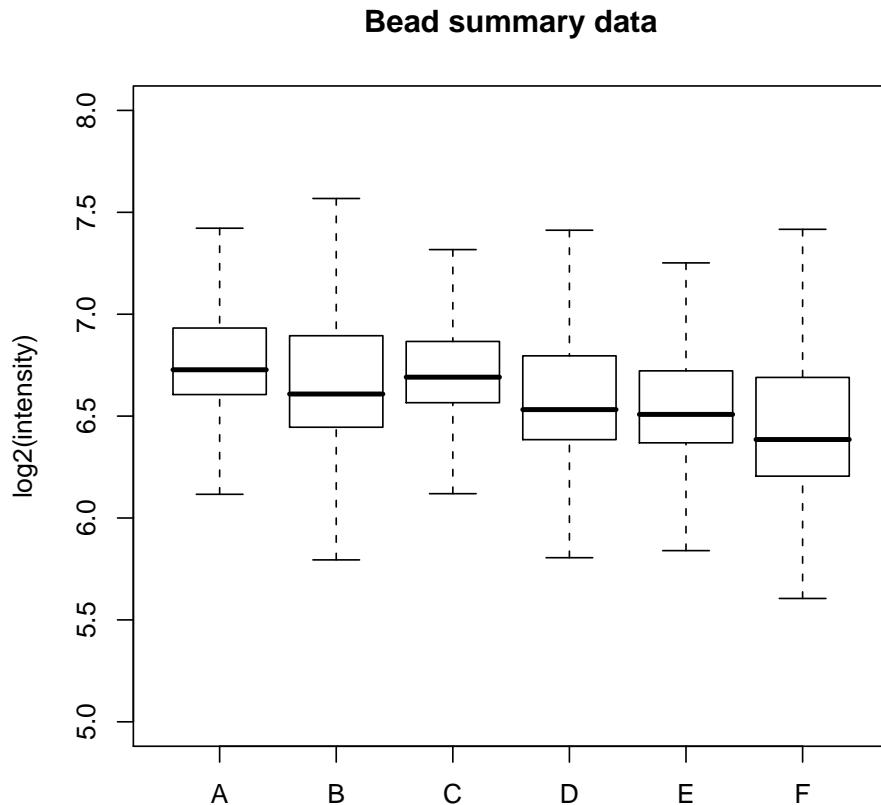
5 Summarisation

The `createBeadSummaryData` function can be used to summarise the replicate intensities from each bead type. By default, outliers are removed using the 3 MADs rule, and the mean of the remaining beads is taken as the summary value. The number of beads used in the calculation and standard error are also stored. Alternative summarisation methods, such as calculating a trimmed mean, median or mean can be applied by changing the `method` argument of `createBeadSummaryData`. If the outliers have already been removed (which will occur if BeadScan's `settings.xml` file includes the line `<ExcludeOutliers>true</ExcludeOutliers>`), then the mean should be calculated with `method="mean"`. At this point we can combine the two strips for each array by using the `imagesPerArray` argument, leaving us with 6 columns now instead of 12. By default, we summarise the values for the green channel. In the case of two-colour data, one may wish to create summary values for the red and green channels separately or summarise the log-ratios for each bead. This can be achieved by setting the `what` argument to `RG` or `M` respectively.


```

> BSData = createBeadSummaryData(BLData.bc,
+   imagesPerArray = 2)
> boxplot(as.data.frame(log2(exprs(BSData))),
+   outline = FALSE, ylab = "log2(intensity)",
+   main = "Bead summary data", ylim = c(5,
+   8), names = c("A", "B", "C", "D",
+   "E", "F"))

```



The default settings for `createBeadSummaryData` assumes that the same probes are to be found on each strip/array in the experiment, which will be true in general. The `BSData` object can be analysed further using the functions described in the bead-summary user's guide.

This user guide was built using the following packages:

```
> sessionInfo()
```

```
R version 2.7.0 Under development (unstable) (2008-02-28 r44635)
x86_64-unknown-linux-gnu
```

```
locale:
```

```
LC_CTYPE=en_GB.UTF-8;LC_NUMERIC=C;LC_TIME=C;LC_COLLATE=C;LC_MONETARY=C;LC_MESSAGES=C;LC_PAPER=C;LC_NAM
```

attached base packages:

```
[1] tools      stats      graphics  grDevices
[5] utils      datasets  methods   base
```

other attached packages:

```
[1] weaver_1.5.0          codetools_0.1-3
[3] digest_0.3.1         illuminaHumanv1_1.4.0
[5] beadarray_1.7.11     affy_1.17.8
[7] preprocessCore_1.1.5 affyio_1.7.13
[9] geneplotter_1.17.7   lattice_0.17-6
[11] annotate_1.17.11     xtable_1.5-2
[13] AnnotationDbi_1.1.24 RSQLite_0.6-8
[15] DBI_0.2-4            Biobase_1.99.1
[17] limma_2.13.5
```

loaded via a namespace (and not attached):

```
[1] KernSmooth_2.22-22 RColorBrewer_1.0-2
[3] grid_2.7.0         tcltk_2.7.0
```

Acknowledgements

We are grateful to Inma Spiteri for providing the data set used in this guide.