

BGX: a Bioconductor package for the Bayesian integrated analysis of Affymetrix GeneChips

Ernest Turro^{*1}, Natalia Bochkina², Anne-Mette K. Hein³, Sylvia Richardson¹

¹Centre for Biostatistics, Imperial College London, UK

³Department of Mathematics, University of Edinburgh, UK

²Molecular Diagnostic Laboratory, Skejby Hospital, Aarhus University, Denmark

Email: Ernest Turro* - ernest.turro@ic.ac.uk; Natalia Bochkina - n.bochkina@ed.ac.uk; Anne-Mette K. Hein - anne.mette.hein@ki.au.dk; Sylvia Richardson - sylvia.richardson@imperial.ac.uk;

*Corresponding author

Abstract

Background: Affymetrix 3' GeneChip microarrays are widely used to profile the expression of thousands of genes simultaneously. They differ from many other microarray types in that GeneChips are hybridised using a single labelled extract and because they contain multiple 'match' and 'mismatch' sequences for each transcript. Most algorithms extract the signal from GeneChip experiments in a sequence of separate steps, including background correction and normalisation, which inhibits the simultaneous use of all available information. They provide a point estimate of gene expression without an associated measure of uncertainty, and consequently biologists using these algorithms cannot take into account the reliability of expression estimates when ranking genes. This problem is particularly acute when only a small number of replicate microarrays per experimental condition are available.

Results: BGX is a new Bioconductor R package that implements an integrated Bayesian approach to the analysis of 3' GeneChip data. The software takes into account additive and multiplicative error, non-specific hybridisation and replicate summarisation in the spirit of the model outlined in [1]. It also provides a posterior distribution for the expression of each gene. Moreover, BGX can take into account probe affinity effects from probe sequence information where available. The package employs a novel adaptive Markov chain Monte Carlo (MCMC) algorithm that raises considerably the efficiency with which the posterior distributions are sampled from. Finally, BGX incorporates various ways to analyse the results, such as ranking genes by expression level as well as

statistically based methods for estimating the amount of up and down regulated genes between two conditions.

Conclusions: BGX performs well relative to other widely used methods at estimating expression levels and fold changes. It has the advantage that it provides a measure of uncertainty for its estimates. BGX includes various analysis functions to visualise and exploit the rich output that is produced by the Bayesian model.

Background

Oligonucleotide microarrays allow biomedical researchers to estimate the expression of thousands of genes simultaneously through their mRNA transcripts. A labelled, fragmented version of the RNA may be hybridised onto an array containing hundreds of thousands of complementary oligonucleotides and then scanned. Affymetrix 3' GeneChip arrays represent genes by sets of probe pairs, each of which consists of an oligonucleotide of length 25 which matches a corresponding RNA subsequence perfectly (PM) and an identical probe with an inverted oligonucleotide on position 13 (MM) that is intended to measure non-specific hybridisation.

The BGX model [1] is an integrated approach to the analysis of GeneChip microarrays in which correction for non-specific hybridisation and gene expression level estimation is performed simultaneously. Posterior distributions of parameters in the model may be obtained numerically. Of particular interest is the availability of posterior distributions of differences in expression between two conditions even when only a single array is available in either or both of the groups compared. Based on these distributions, a powerful method for detecting differential expression has been developed [2].

The probes on Affymetrix GeneChips have been found to exhibit varying propensities to “shine” according to the base composition of their sequences [3] and methods for estimating expression levels from GeneChips that incorporate probe affinity effects have shown demonstrable advances over methods in which these effects are ignored (see, e.g. [4]). We present a new Bioconductor [5] package that implements the BGX model, includes an extension to incorporate probe affinity effects, employs novel algorithmic techniques to sample effectively from posterior distributions, and provides various analysis and plotting functions.

Implementation

Basic model

BGX [1] explicitly models probe intensities as arising partly from specific hybridisation, S (the signal), and partly from non-specific hybridisation, H , with only a fraction, ϕ , of the signal occurring at a PM probe also occurring at the corresponding MM probe. The S s and H s are gene (g), probe (j), condition (c) and replicate (r) specific, and the intensities are assumed to be affected by an additive array-specific noise:

$$\text{PM}_{gjcr} \sim N(S_{gjcr} + H_{gjcr}, \tau_{cr}^2), \quad (1)$$

$$\text{MM}_{gjcr} \sim N(\phi S_{gjcr} + H_{gjcr}, \tau_{cr}^2). \quad (2)$$

The log-transformed signal parameter, $\log(S_{gjcr} + 1)$, is assumed to follow a gene and condition specific distribution, while the log-transformed non-specific hybridisation term, $\log(H_{gjcr} + 1)$, is assumed to arise from an array-specific distribution:

$$\log(S_{gjcr} + 1) \sim TN(\mu_{gc}, \sigma_{gc}^2), \quad (3)$$

$$\log(H_{gjcr} + 1) \sim TN(\lambda_{cr}, \eta_{cr}^2), \quad (4)$$

where TN denotes the truncated normal distribution, truncated to the positive axis. The central parameter of equation (3), μ_{gc} , acts as the BGX expression measure, and equations (1) to (4) represent the basic BGX model.

The core of the model is implemented in the C++ programming language for efficiency and uses MCMC to sample from the full posterior distributions of each parameter. Parameters are estimated using Gibbs sampling where possible (ϕ and τ) and a Random Walk Metropolis-Hastings algorithm elsewhere (S , H , μ , σ , λ and η). Three C++ class templates are used to instantiate zero, one and two-dimensional MCMC update objects for each parameter according to the dimensionality of the corresponding suffixes. Each of the instantiated objects is updated in sequence using references to all other necessary parameters during a burn-in period, which is discarded, and a sampling period, which is used for the posterior distributions.

Probe affinity extension

It has been observed that the propensity of probes to hybridise to mRNA is affected by their base composition [3]. In particular, probes with a high number of cytosine bases have a high propensity to hybridise while probes with a high number of adenine bases exhibit the opposite tendency. Moreover, the

nearer the bases are to the centre of the oligonucleotide, the greater the effect. We account for this in an extension to the core model that incorporate affinity effects in the modelling of non-specific hybridisation. We categorise probes in the following way: let α be a function which, for each gene and probe pair, (g, j) , gives the affinity category of a given probe: $\alpha : (g, j) \rightarrow \{1, \dots, K\}$ (defined below). We refine equation (4) by allowing for a category and array specific distribution of the non-specific hybridisation parameter:

$$\log(H_{gjc_r} + 1) \sim TN(\lambda_{c_r}^{\alpha(g,j)}, \eta_{c_r}^2). \quad (5)$$

The extended model, which we denote GCBGX, is based on equations (1), (2), (3) and (5).

The probes on the arrays are, prior to analysis, grouped into a number of probe affinity categories. This is done by: (a) calculating the probe affinities using the *gcrma* Bioconductor package [4], (b) rounding them to the first decimal place, (c) assigning each value to a preliminary probe affinity category and (d) ensuring that the final categories contain a sufficient number of probes by collapsing small preliminary categories together. We enumerate the resulting probe affinity categories $1, \dots, K$ by increasing affinity. Once categorised, the affinity-specific parameters are estimated from the data, simultaneously with all other parameters.

In some cases, Affymetrix do not directly provide the sequences for all probesets due to licensing restrictions and, consequently, there are Bioconductor probe packages that do not contain complete sequence information. For example, *hgu95aprobe* (version $\leq 1.16.2$) lacks sequences for probes belonging to 172 probesets. We tackled this problem by treating $\alpha(g, j)$ as a random variable, taking values from 1 to K , with prior probability equal to the observed frequency of the categories, $p_k = \frac{N_k}{N}$, where N_k is the number of probes in category k and $N = \sum_{k=1}^K N_k$.

Adaptive MCMC

The full conditional distributions of S , H , μ , σ , λ and η are updated by drawing new values from a proposal distribution, typically a Random Walk (RW) Gaussian proposal centred on the current value with a chosen variance. A typical experiment consists of several hundred thousand probes, resulting in potentially millions of S and H components and tens of thousands of μ and σ components. Each component of a given parameter has a different support and consequently a different optimal RW proposal variance. Using a fixed variance for all components results in excessively low or high acceptance ratios for a

large proportion of components, leading to highly autocorrelated chains.

In order to tackle this problem, we implemented the novel Adaptive Metropolis-Within-Gibbs algorithm recently proposed by Roberts and Rosenthal [6, 7]. We used a unique proposal variance for each object, which adapts to its optimal value after successive batches of 50 iterations. The aim is to achieve an acceptance ratio of around 0.44, which has been shown to be optimal for one-dimensional proposals in certain settings [8, 9], and is commonly accepted as being a sensible benchmark. An acceptance rate that is close to zero implies inefficient mixing, while an acceptance rate that is close to one implies the probability space is not efficiently explored. The algorithm proceeds as follows:

- For each component c of parameter p , assign a parameter-specific starting value to the corresponding proposal variance, σ_{cp}^2
- Choose a sequence $\delta(n) \rightarrow 0$. We chose $\delta(n) = \min(0.01, n^{-1/2})$
- Start the MCMC simulation
- After the n^{th} batch of 50 iterations, calculate the acceptance ratio over the last batch
- If the acceptance ratio is less than the optimal value of 0.44, increase $\log(\sigma_{cp}^2)$ by $\delta(n)$, else decrease it by $\delta(n)$

The algorithm preserves ergodicity as long as each kernel has the right stationary distribution; the total variation distance between successive kernels tends to zero in probability; and the convergence time of each kernel is bounded in probability [6].

R package

The C++ component of BGX is compiled as a shared object which is loaded and executed automatically from within the R package [10]. BGX integrates standard Bioconductor classes such as `AffyBatch` to store raw microarray data and `ExpressionSet` to store processed gene expression measures. Users interested in running BGX programmatically from a shell script, for instance, or in a more memory-efficient manner, also have the choice to run a standalone binary version of the program.

Results and Discussion

Usage

The *bgx* package and its dependencies, *affy* and *germa*, may be installed automatically from the Bioconductor repository from an R shell. The package contains documentation and executable examples in a “vignette” file available using `openVignette()`. Users who wish to compile *bgx* from source will require the Boost C++ libraries [11] and the `hgu95av2cdf` Bioconductor package. The core functionality of the package is contained in the `bgx` function, which takes an `AffyBatch` object instantiated from one or more GeneChip CEL files as its first argument and returns an `ExpressionSet` object containing expression values for each gene and condition:

```
aData <- readAffy("chip1.CEL","chip2.CEL")
eset <- bgx(aData)
assayData(eset)$exprs # Returns expression values
assayData(eset)$se.exprs # Returns standard errors for expression values
```

Optional arguments include `samplesets`, which specifies the experimental design; `genes`, which specifies a subset of genes to analyse; `burnin` and `iter`, which specify the number of iterations for the burn-in and post burn-in phases of the algorithm respectively; `probeAff`, which specifies whether or not to use the probe affinity extension to the original BGX model, and `adaptive`, which specifies whether or not to use Metropolis-Within-Gibbs step adaptation. Full documentation for the `bgx` function is available by running `help(bgx)`.

Although the point measures returned in the `ExpressionSet` object are useful, the distinctive power of the BGX method is that it provides samples from the full posterior distributions of the expression parameter, μ_{gc} . These samples are, by default, saved in directories named `run.1`, `run.2`, *etc.* in R’s current working directory, although this may be overridden with the `rundir` argument. They may be read into R in order to analyse the results of a simulation as follows:

```
bgxOutput <- readOutput.bgx("run.1")
```

The `bgxOutput` object is assigned to a list containing values from the full posterior distributions of μ_{gc} (`bgxOutput$mu`), their expected values (`bgxOutput$muave`) and the gene names (`bgxOutput$geneNames`). This object can be passed to a number of functions to analyse the results. The difference in expression between two conditions, $\mu_{g2} - \mu_{g1}$, may be visualised with the `plotExpressionDensity` and

`plotDEDensity` functions (see Figure 1). `plotDEHistogram` plots a histogram of $P(\mu_{g2} - \mu_{g1} < 0)$ (Figure 2), which can be used to estimate the number of differentially expressed genes [2]. `rankByDE` returns a matrix that ranks genes by differential expression between two conditions. It specifies each gene’s name, index and differential expression measure. More information on each function is available via `help(analysis.bgx)`.

For the purposes of this paper, BGX was run with the “gold standard” of 16k burn-in iterations and 64k sampling iterations. However, the recommended 8k burn-in iterations and 16k sampling iterations are sufficient to provide good estimates of μ_g (Additional File 1). Under these settings, BGX takes approximately one hour per array on a standard 64-bit 3GHz computer. Analyses of up to 100 arrays ought to “fit” in a computer equipped with 4GB of memory. However, BGX may be run separately for each condition, and the output subsequently combined in R by passing multiple output directories to the `readOutput.bgx` function.

Estimation of non-specific hybridisation

The GCBGX model groups MM probes into categories that have similar probe affinities based on their oligonucleotide content. Instead of using a single array-specific parameter, λ_{cr} (Equation (4)), we associate an appropriate $\lambda_{cr}^{\alpha(g,j)}$ component with each probe (Equation (5)), which should correlate positively with probe affinity categories. Figure 3 shows a colour-coded density plot of the $\lambda_{cr}^{\alpha(g,j)}$ parameter obtained in an analysis of the Golden Spike data [12]. As probe affinity categories increase, the distributions shift from left to right. The black density line corresponds to the λ_{cr} distribution in the original BGX model and highlights the discriminatory power of the probe affinity extension.

Probes with unknown sequences have their affinity categories estimated from the data. In order to check the effectiveness of this approach, we performed cross-validation on one out of every 100 genes from the Golden Spike data set, and compared the median estimated category to its true value. Figure 4 shows a positive correlation between estimated and true categories, particularly for high-affinity probes.

Performance of adaptive MCMC

BGX is a computationally intensive program and it is therefore desirable to use an MCMC algorithm that mixes efficiently. Anderson [13] has shown that the empirical variance of an autocorrelated sample X of size n , $E[S_n^2(X)]$ is on average smaller than the true variance, σ^2 . Law [14] has extended Anderson’s result

to show that the variance of the sample mean may be expressed as follows:

$$\begin{aligned} \mathbb{E} \left[\frac{S^2(X)}{n} \right] &= \frac{(n/a(X)) - 1}{n - 1} \cdot \text{var}(\bar{X}) \\ \Rightarrow \text{var}(\bar{X}) &= \frac{n - 1}{n} \cdot \frac{a(X) \cdot \mathbb{E}(S^2(X))}{n - a(X)} \end{aligned} \quad (6)$$

where X is an autocorrelated sample of size n and empirical mean \bar{X} , and $a(n)$ is the integrated autocorrelation time (IACT). The IACT is defined as

$$a(X) = \left[1 + 2 \sum_{j=1}^{n-1} \left(1 - \frac{j}{n} \right) P_j \right]$$

where

$$P_j = \frac{\text{cov}(X_i, X_{i+j})}{\sqrt{S^2(X_i)S^2(X_{i+j})}}$$

Evidently, if the sample is not autocorrelated at all, then $a(X)$ collapses to 1 and $\text{var}(\bar{X})$ becomes equal to the familiar expression $S^2(X)/n$. From (6), the integrated autocorrelation time of a chain relates positively with the variability of its mean, and thus highly autocorrelated chains lead to poor estimates of our gene expression measure.

One solution is to increase the number of iterations while maintaining a fixed subsample size. This translates to subsamples being further apart on the original chain and therefore less correlated. It is faster and more attractive, however, to use an adaptive algorithm that explores the probability space more efficiently. Using the Golden Spike data set [12] for our investigation, we found that the adaptive method led to a range of optimal proposal magnitudes for the Metropolis-Hastings parameters. Figure 5 illustrates this with a histogram of the optimal log variance for S proposals on one array and the original fixed step size overlaid in black. Figure 6 (left & centre) shows a dramatic reduction in the IACT of the S parameters and a milder improvement on the μ_g parameters of expressed genes. A similar improvement was observed for the IACT of the H parameters, this time for all genes (Figure 6 right).

Differential expression is estimated by the standardised BGX differences between two conditions:

$$z_g = \frac{\mathbb{E}(d_g)}{\text{sd}(\mathbb{E}(d_g))} = \frac{\mathbb{E}(d_g) \cdot \sqrt{N}}{a(d_g)^2 \cdot \text{sd}(d_g)},$$

where N is the number of values in the posterior distributions and $d_g = \mu_{g2} - \mu_{g1}$, $g = 1 \dots, G$, that is, samples from the posterior distribution of the difference in the BGX expression measure for each gene. This z-statistic differs from the measure used in [2] by a factor of $\sqrt{N}/a(d_g)^2$, which takes into account

autocorrelation. Since the MCSE for expressed genes decreases while the MCSE for non-expressed genes remains approximately constant, our adaptive algorithm increases the ranking of expressed genes and, consequently, our capacity to detect differential expression.

Performance on spike-in datasets

We illustrate the performance of *bgx* by presenting detailed results from analyses of arrays from the Affymetrix Latin Square data [15] and the Golden Spike data set [12].

Latin Square data

Affymetrix published two data sets for assessing the performance of expression algorithms on their microarrays. The HGU95A data set consists of 16 genes spiked in at known concentrations ranging from 0 to 1024pM and arrayed in a Latin Square format. We considered 16 instead of the original 14 genes described by Affymetrix because we included two extra spike-ins, 546_at and 33818_at, as reported in [16]. We used two replicates and 14 unique concentration configurations labelled A to M and Q. 2716 of the probes in this data set had no sequence information and therefore their probe affinity categories were estimated from the data as part of the model. The HGU133A data set consists of 64 genes spiked in at known concentrations ranging from 0 to 512pM. We considered 64 instead of the original 42 genes described by Affymetrix because we included 22 extra spike-ins, as reported in [17]. We used all 3 replicates for each of the 14 concentration groups.

The data from these experiments were analysed using BGX, GCBGX, RMA [18] and GCRMA [4], and the average expression for each concentration level was recorded. Figure 7 (left) shows a steeper gradient at levels lower than 4pM in the HGU95A data set using GCBGX instead of BGX, pointing to an increased ability to detect concentration changes. For both data sets, BGX and GCBGX are more sensitive to changes within the low range than RMA or GCRMA (Figure 7 right).

Golden Spike data set

The Golden Spike data set consists of six DrosGenome1 GeneChips, with three technical replicates from two conditions: C and S. There are 14010 probe sets in each array representing 14010 genes. 2535 of these are expressed equally under both conditions while 1331 genes are up-regulated in S relative to C. The data is highly valuable for comparing chip analysis methods because it is fully controlled and contains very realistic noise. Due to the asymmetry of the spike-ins, a normalisation of the posterior distributions similar

to that advocated in [12] was carried out by fitting a loess curve to the MA plot [19] of the posterior mean values of μ_{gc} for the non-differentially expressed genes, predicting a curve from the fit for all genes, and subtracting the curve from the posterior distributions of the differences in expression. The RMA and GCRMA expression measures were similarly adjusted using loess normalisation instead of the default quantile normalisation.

Receiver operating characteristic (ROC) curves depict the observed false discovery rate vs. the true positive rate as the cut-off of a ranked gene list is varied. The ROC curves in Figure 8 show that GCBGX has a small advantage over BGX and that both models perform well. For false discovery rates below 40%, BGX and GCBGX outperform both RMA and GCRMA. The number of differentially expressed genes (DEGs) were estimated by running `plotDEHistogram` on the output of the nine comparisons involving one array from condition C versus one array from condition S. The number of genes detected as up-regulated with GCBGX ranged from 681 to 883 (mean 783) and for BGX from 560 to 867 (mean 742). Both methods produced an average true positive rate across the nine comparisons of over 97% for up-regulated genes. Finally, an analysis of a three-replicate comparison yielded 1002 and 958 DEGs with 96.6% and 95.7% true positive rates using GCBGX and BGX respectively.

Conclusions

BGX is a new Bioconductor R package for analysing 3' Affymetrix GeneChips. BGX implements a fully integrated Bayesian hierarchical model with the option to take into account sequence-dependent probe affinities. BGX uses a novel adaptive MCMC algorithm that improves the efficiency with which the posterior distributions of parameters are sampled from. BGX compares favourably to RMA and GCRMA at detecting differential expression, particularly at low concentration levels.

Availability and requirements

Project name: BGX

Project homepage: <http://bgx.org.uk>

Operating systems: Platform independent

Programming language: C++, R

Other requirements: R, Bioconductor

License: GNU GPL

Any restrictions to use by non-academics: No

Authors contributions

ET implemented and tested the code. AKH and SR extended the BGX model to take into account probe affinity effects. NB and ET modelled and implemented probe affinity estimation of probes with missing sequences. ET, AKH, NB and SR provided comments, discussion and wrote the paper.

Acknowledgements

We thank Jeffrey S. Rosenthal, Gareth O. Roberts for advice on adaptive Metropolis-Hastings algorithms, and Alex Lewin and Marta Blangiardo for valuable discussion. This work was supported by BBSRC ‘Exploiting Genomics’ grant 28EGM16093 and the John and Birthe Meyer Foundation.

References

1. Hein AMK, Richardson S, Causton HC, Ambler GK, Green PJ: **BGX: a fully Bayesian integrated approach to the analysis of Affymetrix GeneChip data.** *Biostatistics* 2005, **6**(3):349–373.
2. Hein AMK, Richardson S: **A powerful method for detecting differentially expressed genes from GeneChip arrays that does not require replicates.** *BMC Bioinformatics* 2006, **7**:353.
3. Naef F, Magnasco MO: **Solving the riddle of the bright mismatches: labeling and effective binding in oligonucleotide arrays.** *Phys Rev E Stat Nonlin Soft Matter Phys* 2003, **68**(1 Pt 1):011906.
4. Wu Z, Irizarry RA, Gentleman R, Martinez-Murillo F, Spencer F: **A Model-Based Background Adjustment for Oligonucleotide Expression Arrays.** *Journal of the American Statistical Association* 2004, **99**:909–917(9).
5. **Bioconductor** [<http://bioconductor.org>].
6. Roberts GO, Rosenthal JS: **Coupling and Ergodicity of Adaptive MCMC.** *Journal of Applied Probability* 2005, [<http://www.probability.ca/jeff/ftplib/adapt.pdf>]. [To appear].
7. Roberts GO, Rosenthal JS: **Examples of Adaptive MCMC** 2006, [<http://probability.ca/jeff/ftplib/adaptex.pdf>].
8. Roberts GO, Gelman A, Gilks WR: **Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithms.** *The Annals of Applied Probability* 1997, **7**:110–120.
9. Roberts GO, Rosenthal JS: **Optimal Scaling for Various Metropolis-Hastings Algorithms.** *Statistical Science* 2001, **16**(4):351–367.
10. **The R Project for Statistical Computing** [<http://r-project.org>].
11. **Boost C++ Libraries** [<http://boost.org>].
12. Choe SE, Boutros M, Michelson AM, Church GM, Halfon MS: **Preferred analysis methods for Affymetrix GeneChips revealed by a wholly defined control dataset.** *Genome Biol* 2005, **6**(2):R16.
13. Anderson T: *The Statistical Analysis of Time Series.* New York: John Wiley and Sons 1971.
14. Law AM: **Confidence intervals in discrete event simulation: a comparison of replication and batch means.** *Naval Res. Logistics Quart.* 1977, **23**:667–678.

15. **Affymetrix - Latin Square Data**
[http://www.affymetrix.com/support/technical/sample_data/datasets.affx].
16. Cope LM, Irizarry RA, Jaffee HA, Wu Z, Speed TP: **A benchmark for Affymetrix GeneChip expression measures**. *Bioinformatics* 2004, **20**(3):323–331.
17. McGee M, Chen Z: **New Spiked-In Probe Sets for the Affymetrix HGU-133A Latin Square Experiment**. *COBRA Preprint Series* 2006.
18. Irizarry RA, Hobbs B, Collin F, Beazer-Barclay YD, Antonellis KJ, Scherf U, Speed TP: **Exploration, normalization, and summaries of high density oligonucleotide array probe level data**. *Biostatistics* 2003, **4**(2):249–264.
19. Yang YH, Dudoit S, Luu P, Lin DM, Peng V, Ngai J, Speed TP: **Normalization for cDNA microarray data: a robust composite method addressing single and multiple slide systematic variation**. *Nucl. Acids Res.* 2002, **30**(4):e15–.

Figures

Figure 1 - Expression and differential expression densities

`plotExpressionDensity` plots the density of the posterior distribution of a given gene under a given condition (left). `plotDEDensity` plots the density of the difference in the posterior distributions of a given gene under two conditions (right).

Figure 2 - Estimating the number of differentially expressed genes

`plotDEHistogram` plots a histogram of $P(\mu_{g2} - \mu_{g1} < 0)$, which is used to estimate the number of up and down-regulated genes between two conditions using a routine incorporated in the package.

Figure 3 - Distributions of λ for each probe affinity category

Colour-coded density plot of the $\lambda_{cr}^{\alpha(g,j)}$ parameter of an analysis of the Golden Spike data set. As probe affinity categories increase, the distributions shift from left to right. The black density line is the λ_{cr} distribution from the original BGX model and illustrates the discriminatory power of the probe affinity extension.

Figure 4 - Estimating the affinity of probes with unknown sequence

Sequence information was artificially removed from one out of every 100 genes from the Golden Spike data set and the corresponding probes' affinity categories were estimated from the data. There is a positive correlation between estimated and true categories, particularly for high-affinity probes.

Figure 5 - Adapting the size of MCMC proposal steps

When the adaptive MCMC algorithm is used, the variance of the Metropolis-Hastings proposal step is adapted independently for each S component. The plot shows a histogram of the optimal log variance for S proposals and the fixed step size used in the non-adaptive version overlaid in black, highlighting that a wide range of proposal variances are needed.

Figure 6 - Decreasing the IACT of expressed genes

The plots show a dramatic reduction in the IACT of the S parameter and a milder improvement on the μ parameter of expressed genes (left & centre). A similar improvement was observed for the IACT of the H parameter, for all genes (right).

Figure 7 - Performance on the Latin Square data

There is an increased ability to detect concentration changes at levels lower than 4pM in the HGU95A data set using GCBGX instead of BGX. Both BGX and GCBGX are more sensitive to changes within the low range than RMA or GCRMA.

Figure 8 - Performance on the Golden Spike data

Nine single-array comparisons between condition C and S were performed using BGX, GCBGX, RMA and GCRMA. The ROC curves show that GCBGX has a small advantage over BGX and that both models perform well. For false discovery rates below 50%, BGX and GCBGX outperform both RMA and GCRMA.

Additional Files


Additional file 1 — Varying the number of sampling iterations

Plot of mean posterior values of μ_g obtained using a sampling length of 64k iterations versus those obtained using sampling lengths of 16k, 32k and 64k using an alternative seed for the pseudo random number generator. The values for the 32k and 64k comparisons are shifted upwards for clarity. Runs of more than 16k confer a small increase in stability of estimation for non-expressed genes. At 32k, estimates are as stable as between two 64k runs using different seeds.

Additional file 2 — R script to reproduce plots

R script to reproduce the plots in this paper.

Densities of mu for gene 141200_at



Density of mu for gene 141200_at
condition 2 - condition 1





Figure 1

Histogram of $P(\mu_{g2} - \mu_{g1} < 0)$; $df=7$



Densities of λ^k (GCBGX) and λ (BGX)




Figure 3




Figure 4

Histogram of proposal variances for the S parameters in one array






Figure 5

IACT of $\log(S+1)$ for expressed genes



IACT of μ for expressed genes



IACT of $\log(H+1)$ for subsample of all genes





Figure 6

Concentration plots – HGU95a



Concentration plots – HGU133atag




Figure 7

Golden Spike average single array comparisons




Figure 8

Additional files provided with this submission:

Additional file 1: additional1.pdf, 1192K

<http://www.biomedcentral.com/imedia/3711754931584915/supp1.pdf>

Additional file 2: additional2.r, 16K

<http://www.biomedcentral.com/imedia/7303371841584915/supp2.r>