



## Faster cyclic loess: normalizing RNA arrays via linear models

Karla V. Ballman\*, Diane E. Grill, Ann L. Oberg and Terry M. Therneau

Division of Biostatistics, Mayo Clinic College of Medicine, Rochester, MN 55905, USA

Received on January 9, 2004; revised and accepted on February 20, 2004

Advance Access publication May 27, 2004

### ABSTRACT

**Motivation:** Our goal was to develop a normalization technique that yields results similar to cyclic loess normalization and with speed comparable to quantile normalization.

**Results:** Fastlo yields normalized values similar to cyclic loess and quantile normalization and is fast; it is at least an order of magnitude faster than cyclic loess and approaches the speed of quantile normalization. Furthermore, fastlo is more versatile than both cyclic loess and quantile normalization because it is model-based.

**Availability:** The Splus/R function for fastlo normalization is available from the authors.

**Contact:** ballman@mayo.edu

### 1 INTRODUCTION

High-density gene expression array technology allows investigators to obtain quantitative measurement of the expression levels for tens of thousands of genes in a biological specimen. There are two major types of microarray technologies: spotted cDNA and oligonucleotide arrays. Expression data obtained from either type of microarray technology has measurement error or variation. One type of variability is due to biological differences between specimen samples. This is what is of interest to the investigator; the investigator would like to know which genes are differentially expressed among different biological samples (e.g. between cancerous and normal kidney tissues, among B-cells challenged with different agents in culture.). Systematic variation also affects the measured gene expression level. Many sources contribute to this type of variation and are found in every microarray experiment. Sources include, but are not limited to, the array manufacturing process, the preparation of the biological sample, the hybridization of the sample to the array, and the quantification of the spot intensities. (Hartemink *et al.*, 2001) provide a more complete discussion of the sources of systematic variation in the microarray experimental

process. The purpose of normalization is to minimize the systematic variations in the measured gene expression levels among different array hybridizations to allow the comparison of expression levels across arrays and so that biological differences can be more easily identified.

There are numerous methods for normalizing gene expression data. Generally, normalization methods use a scaling function to correct for experimental variation. These functions are applied to the raw intensities of the spots (gene sequence on spotted arrays or oligonucleotide probes on the Affymetrix GeneChip array) on the microarray to produce normalized or scaled intensities. Types of normalization techniques include mean correction (Richmond and Somerville, 2000), non-linear models (Yang *et al.*, 2002), linear combination of factors (Alter *et al.*, 2000), and Bayesian methods (Newton *et al.*, 2001). There is compelling evidence that non-linear normalization methods, which are not dependent upon the choice of a baseline array, perform the best (Bolstad *et al.*, 2003). We tend to favor two commonly used non-linear methods: cyclic loess normalization and quantile normalization. Both techniques are non-linear and perform normalization on the set of arrays as a whole without specifying a reference array. Overall, we prefer cyclic loess because it is not as aggressive in its normalization as is quantile normalization; however, cyclic loess is relatively slow for even a moderate sized set of arrays. Quantile normalization, on the other hand, is very fast for even large sets of arrays. The goal of our investigation was to develop a method that produced normalized values similar to that of cyclic loess but would be considerably faster—on the order of the speed of quantile normalization. The result is a new normalization method called fastlo. Since normalization is performed on the raw intensity values of the spots on the arrays, the methods discussed here are applicable to both major types of array technology: spotted cDNA and oligonucleotide. However, our focus is on data arising from GeneChip arrays and it should be noted that there are additional considerations when normalizing two-color spotted cDNA arrays (Yang *et al.*, 2002).

In the next section we discuss cyclic loess normalization and useful insights that can be gained from viewing it as a

\*To whom correspondence should be addressed.




Fig. 1. MA plot for 10 000 randomly selected probes.

parallel algorithm. Section 3 describes a new type of normalization, fast linear loess (fastlo), arising from the observation that cyclic loess is essentially a smoothing function coupled with a very simple linear model. In Section 4, we review quantile normalization. The performances of the three different normalization techniques are compared in Section 5. Comparisons are made using simulated data and real data, one of the benchmark sets for Affymetrix GeneChip expression measures (Cope *et al.*, 2003). The simple linear model underlying the fastlo method is extended in Section 6. We close with a discussion of our findings in Section 7.

## 2 CYCLIC NORMALIZATION

Most methods of normalization, including the methods discussed below, assume that the vast majority of the genes do not change expression levels under the conditions being studied. In other words, they assume that the average (geometric mean) ratio of expression values between two conditions is one or equivalently, the average (arithmetic mean) log ratio of expression is zero for a typical gene. This is biologically plausible for many studies. However, if there is good reason to believe this assumption is not true for a particular study, then the normalization methods described here are not appropriate.

### 2.1 Cyclic loess

A fundamental graphical tool for the analysis of gene expression array data is the M versus A plot (MA plot); here M is the difference in log expression values and A is the average of log expression values (Dudoit *et al.*, 2002). Figure 1 contains an MA plot for a random sample of 10 000 probes from two (unnormalized) GeneChip arrays with a loess smoother superimposed. The MA plot for ideally normalized data would show a point cloud scattered about the  $M = 0$  axis. In other words, the loess smoother would be a horizontal line at 0 for ideally normalized data.

Cyclic loess normalizes two arrays at a time by applying a correction factor obtained from a loess curve fit through the MA plot of the two arrays, call the curve  $f(x)$ . For example, consider the circled point in Figure 1. This point corresponds

to a spot  $i$  on each array. On one array, the observed expression level at this spot would be reduced by  $1/2$  the distance of  $f(x)$  from the  $y = 0$  line; in other words,  $f(x)/2$  would be subtracted from the expression level of this spot on one of the arrays. The expression value for this spot on the other array would be increased by  $f(x)/2$ . After correction, the MA plot for this particular pair of arrays would be horizontal. One pass of the cyclic loess algorithm consists of performing this pairwise normalization on all distinct pairs of arrays. Passes of the algorithm continue until the computed corrections of a completed pass are essentially zero.

In summary, to perform the cyclic loess algorithm begin with the  $\log_2$  of the spot expression intensities arranged as a matrix with one column per array and one row per array spot and proceed through the steps below.

- (1) Choose two arrays and generate an MA plot of the data. The  $x$ -axis is the mean probe expression value of the two arrays and the  $y$ -axis is the difference (one point for each spot).
- (2) Fit a smooth loess curve  $f(x)$  through the data.
- (3) Subtract  $f(x)/2$  from the first array and add  $f(x)/2$  to the second.
- (4) Repeat until all distinct pairs have been compared.
- (5) Repeat until the algorithm converges.

In practice, the pairs of arrays are chosen by a method that systematically cycles through all pairs. A drawback of cyclic loess is the amount of time required to normalize a set of data; the time grows exponentially as the number of arrays increases. Typically, two or three passes through the complete cycle are required for convergence. Likely, cyclic loess would converge faster if pairings went in a more balanced order. However, the time savings would not be considerable because a loess smooth would still be required for a relatively large number of array pairs.

Further examination of the cyclic loess algorithm reveals some interesting facts. First, the algorithm preserves the row means of the data matrix,  $Y$ . At any given step, one number in a row is increased by  $f(x)/2$  and another is decreased by the same amount. Second, if all the values in one of the columns are increased or decreased by a constant, the final results of the algorithm (the scaled intensities on each array) are changed only by the addition of a constant. This is because any one of the plots on which the smooths are based is identical but for the labeling of its axes, and thus any given smooth is changed only by a constant. One pass through the algorithm requires  $C_2^n$  loess smooths on all the spots on the array.

### 2.2 Parallel loess

Cyclic loess is inherently parallel in nature. Viewing it in this manner may provide insights that allow computational time savings. Imagine that we had a parallel machine, so that all the pairwise normalizations could be done simultaneously. Once

each of the pairwise corrections is obtained, then the correction for spot  $i$  on array  $j$  would be the ‘average’ correction for this spot across all array pairs containing array  $j$ . Essentially, the correction to the spots on array  $j$  would be the average of all computed corrections of the spots for pairs containing array  $j$ . To simplify the logic, we consider the pairing of each array with itself, as well as both orderings of the pairs. This means that for  $n$  arrays, there are  $n^2$  pairings. Here, and throughout the remaining text,  $Y$  is used to denote the matrix of the  $\log_2$  intensity values where the column  $j$  corresponds to an array and row  $i$  corresponds to a spot.


As the simplest example, consider the noiseless case where each column of  $Y$ , corresponding to an array, differs from any other by a constant and array ‘0’ is an imaginary reference representing the true expression level. In other words, the intensity of the spot  $i, j$ ,  $y_{ij}$ , can be expressed as  $y_{ij} = y_{i0} + c_j$ . The ideal correction of all spots for a given chip is  $c_j - \bar{c}$  (a horizontal loess curve) and the average correction for array 1 from the parallel algorithm is

$$(1/n) \sum_{j=1}^n (c_1 - c_j)/2 = (c_1 - \bar{c})/2.$$

That is, the average correction is 1/2 of what it should be. As a result, we define the correction step for the  $i$ -th chip in parallel loess to be  $(2/n) \sum_j f_{ij}$ , where  $f_{ij}$  is the smooth for the plots of chips  $i$  and  $j$ .

Now consider a simple simulation where each column of the data matrix (each array) differs from any other by a constant plus a symmetrically distributed noise term. Array ‘0’ is the set of 5000 true intensities, which were randomly selected from a uniform distribution with range from 0 to 10. The  $\log_2$  intensity levels for the four arrays in the experiment,  $y_{ij}$ , are derived from array ‘0’ as follows:  $y_{ij} = y_{i0} + j + e_{ij}$ , where  $e_{ij} \sim t_8$  ( $i = 1, 2, \dots, 5000$  and  $j = 1, 2, 3, 4$ ). A  $t$ -distribution was selected for the error term (noise distribution) because it is symmetric and produces more outliers than a normal distribution. In this case, the true correction for array  $j$  is  $j - \bar{j} = j - 2.5$ . The four pairwise corrections for a particular spot involving array 1 have expectations of 0,  $-1.0$ ,  $-2.0$  and  $-3.0$ . The average of the four corrections for a particular spot is  $-1.5$ . Figure 2 shows the four computed corrections for each spot from parallel loess involving array 1 as well as the average of the corrections across arrays. This suggests averaging the corrections to get an overall update, rather than applying each of the separate corrections to the data in turn. Figure 2 also includes the corrections for array 1 after applying cyclic loess. For this case performing the smooths cyclically or in parallel produces equivalent results. For this figure, and all remaining figures, a subset of randomly selected points are plotted to thin the plot and better display the relevant structure.

Using the parallel version would probably be faster than cyclic loess. However, we did not intend to obtain computational speed savings by developing a parallel version of cyclic loess.




**Fig. 2.** The computed corrections for array 1 from the 4 pairwise smooths for a parallel loess (1v1, 1v2, 1v3, 1v4), for the average of the pairwise corrections (avg), and the for cyclic loess (cyclic).

Rather, we used this construct to gain insight into how cyclic loess works and how it might be made faster for non-parallel machines.

Parallel loess can be shown to be unbiased for this simple case. Let  $y_{ij} = \alpha_i + \beta_j + \epsilon_{ij}$ ;  $\alpha_i$  are the true probe values ( $y_{i0}$  in the simulation),  $\beta_j$  the simple array effects, and  $\epsilon_{ij}$  the error. Assume that  $a, b$  are two vectors of constants, and a smooth of  $Yb$  on  $Ya$  is to be used to estimate the correction. Since the correction should be constant, we want  $Ya$  and  $Yb$  to be uncorrelated, i.e. that there be no linear bias in the smooth. The covariance matrix of  $Y$  has elements of  $\sigma_\alpha^2 + \sigma_\epsilon^2$  on the diagonal, and  $\sigma_\alpha^2$  off the diagonal, where these are the variances of  $\alpha$  and  $\epsilon$  above. Simple algebra shows the covariance of  $Ya$  and  $Yb$  to be  $\sigma_\epsilon^2 \sum a_j b_j + \sigma_\alpha^2 (\sum a_j) (\sum b_j)$ . Since  $\sigma_\alpha^2$  and  $\sigma_\epsilon^2$  are unknown for a given problem, linear bias in the plot is avoided by choosing  $a$  and  $b$  so that both terms are zero. For parallel loess, the plot for arrays 1 and 2 has  $a = (1/2, 1/2, 0, \dots, 0)$  and  $b = (1, -1, 0, \dots, 0)$ , clearly satisfying the criteria, and likewise for any other pair of arrays  $i, j$ . Other choices of  $a$  and  $b$  will be explored below.

**2.2.1 Parallel loess variant one** An obvious variant of the parallel cyclic loess algorithm is to replace the  $n$  loess smooths, each on  $p$  points, with a single loess smooth on  $np$  points. In other words, the corrections to be applied to spots on array  $j$  would be obtained by placing all the points in the MA plots containing array  $j$  ( $n$  of them) into a single plot and performing a single loess smooth on this plot. Recall, the corrections for array  $j$  are obtained in parallel loess by averaging the  $n$  pairwise corrections, obtained from the  $n$  MA plots; and recall that the corrections are just twice the smooth value. So if the smoother is a linear operator, which loess is other than outlier rejection passes, then the average of the smooths will be equal to a single smooth applied to a plot with all the derived data points, i.e. the points from the  $n$  MA plots involving array  $j$  all placed on a single plot. Because of the very large amount of data on microarray gene expression arrays, outlier rejection



**Fig. 3.** The computed corrections for arrays 2 and 3 from the example of 4 arrays. The corrections are from a single smooth on all the derived 20 000 data points (parallel variant 1), from parallel loess (parallel), and from cyclic loess (cyclic).

is not normally an important issue. Specifically, when there is a large quantity of data over a given range, no point, even if extreme, can exert much influence.

Figure 3 compares the computed corrections for the array spots on two arrays from our four array example introduced above: true expression values for the array shifted by a constant plus symmetrically distributed error. The computed corrections displayed are those obtained from cyclic loess, from the pure parallel version of cyclic loess, and from a single loess smooth on the plot containing all the points from the four MA plots involving the array to be normalized (parallel variant 1). Recall that for the pure parallel version of cyclic loess, the correction for a particular point is found by averaging the corrections obtained from the  $n$  MA plots. As expected, all three methods produced equivalent results.

This parallel variant 1 reduces the number of loess smooths that are computed. One pass through the data with parallel variant 1 requires only  $n$  loess smooths compared to  $C_2^n$  loess smooths required by cyclic loess. However, each smooth for this parallel variant 1 is performed on  $np$  versus  $p$  points. Whether this version of parallel loess is faster than cyclic loess likely depends on the implementation details of each algorithm.

**2.2.2 Parallel loess variant two** The next idea is to reduce the number of points in the plot containing all the points from the MA plots of parallel loess involving array  $j$  from  $np$  to  $p$ . Since there are  $p$  spots on an array, the intent is to replace each collection of  $n$  points per spot  $i$  (one per each pairwise MA plot—array  $j$  with array 1, array  $j$  with array 2, ..., array  $j$  with array  $n$ ) with a single point, which summarizes the collection of  $n$  points for spot  $i$ . The new plot would contain  $p$  ‘summary’ points and we would perform a smooth on this plot to obtain the spot corrections for array  $j$ .

A way to produce a point for spot  $i$  that summarizes the  $n$  points for spot  $i$  that involve array  $j$  is to set the  $x$ -coordinate equal to the average of the  $n$   $x$ -coordinates:


$$\frac{1}{n} \left[ \frac{(y_{i1} + y_{i1})}{2} + \frac{(y_{i1} + y_{i2})}{2} + \dots + \frac{(y_{i1} + y_{in})}{2} \right] = \frac{y_{i1} + \bar{y}_i}{2}$$

Here  $\bar{y}_i$  is the row mean for the  $i$ th row of the data, which is equal to the mean expression value of spot  $i$  across the arrays. The vertical position for spot  $i$  on array  $j$  is the average of the  $y$ -coordinates:

$$\frac{1}{n} [(y_{i1} - y_{i1}) + (y_{i1} - y_{i2}) + \dots + (y_{i1} - y_{in})] = y_{i1} - \bar{y}_i$$

The next step is to determine corrections for the spots on chip  $j$  by fitting the loess smooth on this set of  $p$  points. This would be repeated for each of the  $n - 1$  other chips; a total of  $n$  smooths need to be computed for one pass through all the data, similar to the parallel loess variant 1 above. However, the number of points for each loess smooth is  $p$  for this variant compared to  $np$  for variant 1.

Repeating the bias computation found in Section 2.2, the plot for array 1 versus array 2 has vertical and horizontal axes  $aY$  and  $bY$ , respectively, with  $a = (1 - 1/n, -1/n, \dots, -1/n)$  and  $b = (1 + 1/n, 1/n, \dots, 1/n)/2$ . This yields  $\sum a_j = 0$ ,  $\sum b_j = 1$  and  $\sum a_j b_j = (n-1)/n$  leading to positive correlation between  $aY$  and  $bY$  and a potentially biased correction. Figure 4 compares the computed corrections produced by this variant of parallel loess (parallel variant two), parallel loess, and cyclic loess for two of the four arrays in our simple data example. Clearly, the bias is severe. In the next section we use a linear models argument to motivate a similar plot, but with  $x$ -coordinate of  $\bar{y}_i$ , leading to  $b = (1/n, \dots, 1/n)$  and  $\sum a_j b_j = 0$  and an unbiased estimate.



**Fig. 4.** The computed corrections for arrays 2 and 4 of the simulated example of 4 arrays. The corrections are from the smooth of the 5000 derived points (parallel var 2), from parallel loess (parallel), and from cyclic loess (cyclic).

### 3 FAST LINEAR LOESS

Cyclic loess can be conceptualized as a smooth (loess), coupled with a very simple linear model. Consider the case of two arrays, 1 and 2, with gene expression levels represented by  $y_{i1}$  and  $y_{i2}$  for  $i = 1, \dots, p$ , and the simplest possible linear model for the data  $y_{ij} = \alpha_i + \epsilon_{ij}$ , an intercept for each spot. The solution is of course  $\hat{\alpha}_i = \bar{y}_{i.}$ . In the MA plot for the two arrays, the  $x$ -coordinate is  $(y_{i1} + y_{i2})/2 = \bar{y}_{i.} = \hat{y}_{i.}$ , the  $y$ -coordinate is  $y_{i1} - y_{i2} = 2(y_{i1} - \hat{y}_{i.})$ , and the adjustment to spot  $i$  on the the first chip is  $1/2$  the height of the loess smooth on this plot at the  $x$ -coordinate corresponding to spot  $i$ .

Extending this idea to the  $n$  array case suggests creating a modified MA plot for each array  $j$ . The modified MA plot consists of array  $j$  and a constructed ‘average’ array; the intensity level of spot  $i$  on the ‘average’ array is equal to the average intensity of spot  $i$  across all arrays for  $i = 1, \dots, p$ . The computed corrections for array  $j$  would be obtained from the loess smooth on the points of this plot. This would be done for each of the  $j = 1, \dots, n$  arrays. This variant of cyclic loess, called fastlo, requires  $n$  loess smooths each on  $p$  points for one pass through the data. Obviously, this is considerably faster than cyclic loess, which requires  $C_2^n$  loess smooths.

The steps of fastlo given below are done on the  $\log_2$  of the spot intensities:

- (1) Create the vector  $\hat{y}_{i.} =$  the row mean of  $Y$ . Note that this is the same as creating an ‘average’ array.
- (2) Plot  $\hat{y}$  versus  $(y_i - \hat{y})$  for each array  $j$ . This plot has one point for each spot (modified MA plot).
- (3) Fit a loess curve  $f(x)$  through the data.
- (4) Subtract  $f(x)$  from array  $j$ .

- (5) Repeat for all remaining arrays.
- (6) Repeat until the algorithm converges.

A further interesting aspect of fastlo is that it requires only 1 or at most 2 iterations to converge. If there is no outlier down weighting, then the loess smoother will be a linear operator and the average of the smooths will be the smooth of all the points:

$$\frac{1}{n} \sum Sm(y_n - \bar{y}) \approx Sm \left( \sum_n y - \bar{y} \right) = Sm(0) = 0$$

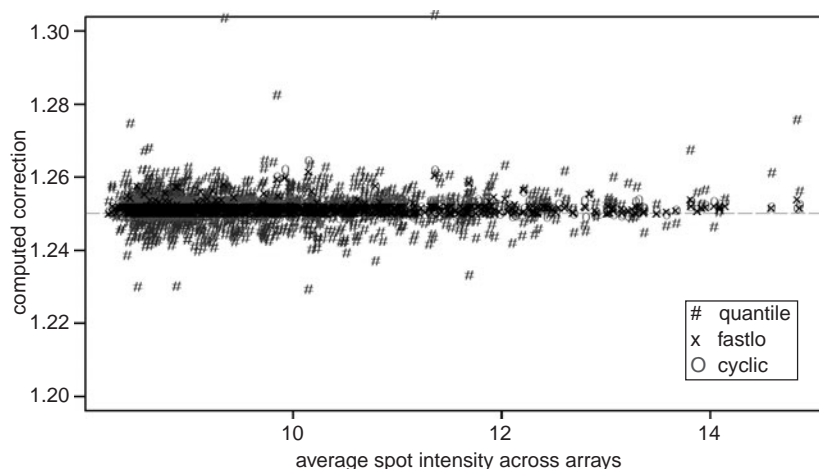
where  $Sm$  is the smoother. If this holds, then for any given row of  $Y$ , some elements increase and some decrease, but the mean stays the same. If the row means do not change, the algorithm has converged.

### 4 QUANTILE NORMALIZATION

Quantile normalization makes the overall distribution of values for each array identical, while preserving the overall distribution of the values. It consists of two steps.

- (1) Create a mapping between ranks and values. For rank 1 find the  $n$  values, one per array, that are the smallest value on the array, and save their average. Similarly for rank 2 and the second smallest values, and on up to the  $n$  largest values, one per array.
- (2) For each array, replace the actual values with these averages.

As mentioned, this produces identical distributions of values on each array; quite an aggressive normalization process. On the other hand, quantile normalization is extremely fast—it only requires a sort of the arrays and a computation of



**Fig. 5.** The computed corrections for array 8 of simulated dataset 2. The corrections are from cyclic loess (cyclic), quantile (quantile) and fastlo (fastlo). The dashed line represents the true correction value for each spot.

means across sorted rows. It completes in a single pass through the data.

## 5 RESULTS

As a final step, we compare the performance of the three normalization algorithms: cyclic, quantile and fastlo. Results of applying these algorithms to a simulated and real dataset are presented. Descriptions of the datasets and a comparison of the results follow below.

### 5.1 Simulated dataset

For the simulated data introduced in Section 2.2, we assumed that the spot intensities were uniformly distributed and the non-systematic errors were symmetrically distributed. Both assumptions are arguably unrealistic and were relaxed. The simulated data in this section were generated from a set of five human liver tissue (total cRNA concentration of  $5 \mu\text{g}$ ) arrays from the GeneLogic dilution experiment ([www.genelogic.com/media/studies/dilution.cfm](http://www.genelogic.com/media/studies/dilution.cfm)); this set will be denoted as liver5. The simulated dataset has 10 arrays and each array has 20 000 spots. The 20 000 spots were randomly selected from among the PM spots on the Affymetrix U95Av2 GeneChip. The spot  $i$  intensity on array  $j$  was generated as  $y_{ij} = y_{i,\text{rand}(\text{liver5})_j} + j/2$ . The value  $y_{i,\text{rand}(\text{liver5})_j}$  is randomly selected from among the group of five spot  $i$  values from the arrays in the liver5 set. A constant,  $j/2$ , is added to spot  $i$  on array  $j$  for all  $i = 1, 2, \dots, 20\,000$  spots and all  $j = 1, 2, \dots, 10$  arrays.

As for the simulated data set in Section 2.2, the true correction of the systematic bias is a constant for array  $j$ ,  $j = 1, 2, \dots, 10$ . In this case, the constant for array  $j$  is equal to  $j/2 - 2.75$ . Again, all three algorithms yielded similar results in terms of correcting for the known constant offset (Fig. 5, arrays not shown yielded similar results). The distributions of normalized values of the 10 arrays yielded by each of the three

**Table 1.** Means ( $\bar{x}$ ) and standard deviations ( $s$ ) of the 20 000 spot intensity corrections produced by each of the three normalization methods when applied to the 10 arrays of simulated dataset 2

	$\bar{x}$			$s$		
	Quantile	Cyclic	Fastlo	Quantile	Cyclic	Fastlo
Array 1	-2.25	-2.25	-2.25	0.007	0.002	0.002
Array 2	-1.75	-1.75	-1.75	0.008	0.003	0.003
Array 3	-1.25	-1.25	-1.25	0.008	0.003	0.004
Array 4	-0.75	-0.75	-0.75	0.006	0.002	0.002
Array 5	-0.25	-0.25	-0.25	0.009	0.005	0.004
Array 6	0.25	0.25	0.25	0.008	0.003	0.002
Array 7	0.75	0.75	0.75	0.009	0.005	0.004
Array 8	1.25	1.25	1.25	0.007	0.002	0.001
Array 9	1.75	1.75	1.75	0.007	0.004	0.003
Array 10	2.25	2.25	2.25	0.008	0.002	0.003

methods were also essentially equivalent (results not shown). Quantile normalization was the fastest followed by fastlo and the slowest was cyclic normalization; the ratios of CPU time required were 1.0:1.1:16.1 for quantile:fastlo:cyclic loess.

For this simulated setting, we know the true amount of systematic error needed to be removed from each array via normalization. A normalization technique that performed well would yield unbiased corrections and corrections with the smallest standard deviation. The results in Table 1 reveal all three methods to be unbiased. Fastlo appears to yield the smallest standard deviations, in general, followed closely by cyclic loess and quantile normalization having the largest. For practical purposes, the differences are so small as not to matter.

### 5.2 GeneLogic dilution data

A final comparison of the three methods was done on real data from the GeneLogic dilution experiment. The data

**Table 2.** Five number summary plus the mean for the distribution of standard deviations of 201 800 normalized PM spot intensity values across the 5 arrays within each concentration level

		Min	Q1	Median	Q3	Max	Mean
1.25	quantile	0.0035	0.0797	0.1066	0.1377	4.0684	0.1135
	cyclic	0.0052	0.0788	0.1055	0.1361	4.1000	0.1123
	fastlo	0.0050	0.0794	0.1062	0.1370	4.2428	0.1129
	fastlo2	0.0042	0.0797	0.1067	0.1377	4.3301	0.1135
2.5	quantile	0.0000	0.0802	0.1080	0.1406	2.6932	0.1172
	cyclic	0.0030	0.0795	0.1071	0.1394	2.7350	0.1163
	fastlo	0.0034	0.0797	0.1072	0.1394	2.6653	0.1162
	fastlo2	0.0062	0.0798	0.1075	0.1397	2.7690	0.1166
5.0	quantile	0.0058	0.0774	0.1039	0.1352	3.6135	0.1121
	cyclic	0.0046	0.0773	0.1037	0.1346	3.7939	0.1118
	fastlo	0.0022	0.0774	0.1038	0.1348	3.6643	0.1119
	fastlo2	0.0046	0.0773	0.1038	0.1349	3.6019	0.1120
7.5	quantile	0.0060	0.0845	0.1134	0.1479	3.5267	0.1234
	cyclic	0.0009	0.0827	0.1115	0.1462	3.5527	0.1223
	fastlo	0.0029	0.0835	0.1122	0.1466	3.6194	0.1229
	fastlo2	0.0055	0.0829	0.1118	0.1470	3.7071	0.1233
10.0	quantile	0.0055	0.0744	0.1002	0.1302	2.8603	0.1062
	cyclic	0.0035	0.0743	0.0998	0.1294	3.3200	0.1058
	fastlo	0.0043	0.0743	0.0997	0.1292	3.3048	0.1058
	fastlo2	0.0035	0.0743	0.0998	0.1296	3.3075	0.1061
20.0	quantile	0.0027	0.0822	0.1110	0.1457	3.0393	0.1199
	cyclic	0.0038	0.0819	0.1105	0.1451	3.0374	0.1195
	fastlo	0.0034	0.0820	0.1105	0.1447	3.0400	0.1191
	fastlo2	0.0033	0.0820	0.1106	0.1447	2.9952	0.1191

The four methods compared are quantile, cyclic loess, fastlo, and a second version of fastlo, which normalizes all 30 arrays simultaneously using a group variable (fastlo2).

consists of six groups of human liver cRNA at concentrations of 1.25, 2.5, 5.0, 7.5, 10.0 and 20.0  $\mu\text{g}$  total cRNA. Five replicate arrays were made for each concentration level. See Irizarry *et al.* (2003) for more details; the data can be obtained through the GeneLogic website (<http://www.genelogic.com/media/studies/dilution.cfm>). These data were normalized by each of the three normalization methods under investigation: cyclic loess, quantile, and fastlo. Each group of five arrays were normalized together. A comparison of the boxplots of the normalized values (data not shown) revealed essentially equivalent distributions of the normalized spot intensities generated by the three methods.


As in the simulated data, we expect that the true spot intensity levels should be equal across the replicate arrays for a given dilution concentration. For this situation, a desirable quality of a normalization technique would be one that yields the smallest standard deviation at a spot  $i$  across the five arrays, for all  $i$ . Table 2 summarizes the distribution of the standard deviations of the PM spots. These summaries are remarkably similar for all three methods. Quantile normalization tended to produce the largest standard deviation on average (either measured by the median value or the mean value) and cyclic loess and fastlo having the smallest in many of the sets. Overall, the differences in these distributions are probably not of practical

significance. Interestingly, the distributions of the standard deviations did not appear to differ in identifiable ways among the different concentration levels; specifically, the standard deviations in spot intensities did not appear to increase as the concentration levels of the total cRNA increased.

## 6 SIMPLE EXTENSION OF FASTLO

When normalizing arrays that are to be compared, a decision is made as to whether to normalize the entire set or to normalize separately within biologically identifiable subsets. For example, consider a situation where expression levels are to be compared between two distinct groups such as cancerous tissues and normal tissues. In this case, the entire set of arrays could be normalized together or separate normalizations could be done for the set of cancerous tissue arrays and for the set of normal tissue arrays. Determination of the most appropriate normalization approach depends upon which process is responsible for most of the systematic variation in expression levels among the arrays: the biological process of interest or the microarray experimental process. For instance, assume that the global spot average was 12.0 in the arrays from cancer tissues and 11.4 from the normal tissues. Do we ascribe this overall difference to the treatment (e.g. more genes are expressed in the cancer tissues) or to the arrays (variation in experimental conditions that produced the arrays)? If most of the systematic variation is believed to arise from biological mechanisms of interest, then it is reasonable to perform a separate normalization for each group being compared. On the other hand, if it is believed that most of the systematic variability is due to the experimental processes (and so not of interest), then it is reasonable to perform a single normalization on the entire group. Unfortunately, it is often not clear which normalization approach is the most appropriate.

Both cyclic loess and quantile normalization require the analyst to decide whether the biologic mechanism or microarray experimental process dominates the systematic variation in spot intensity levels among arrays. On the other hand, fastlo can be extended so that the analyst does not have to commit to one or the other of these assumptions. This is a consequence of the model-based nature of fastlo. Recall that the ideal plot for performing normalization for an array would be one with true spot intensity on one axis, say the  $x$ -axis, and the observed—true spot intensity on the  $y$ -axis. Since we do not know the true spot intensity, we could use a linear model to estimate it. Fastlo uses the simplest linear model, the mean spot intensity value across the arrays. Other linear models could be used instead such as one that also incorporates a grouping variable, which indicates the group to which an array belongs. This allows the data to determine the relative contributions of the biologic mechanisms of interest (represented by a grouping variable) and of the microarray experimental process to the systematic variation of spot intensity levels among the arrays.



**Fig. 6.** Boxplots of the normalized spot intensity values for the five arrays of the GeneLogic dilution data at total cRNA concentration of  $10 \mu\text{g}$ . The normalizations are from cyclic loess, quantile, fastlo, and fastlo2, which uses a linear model with spot and grouping variables.

Consider the set of GeneLogic liver dilution data (Section 5.2). The grouping variable for this data represents the concentration level of the liver tissue cRNA. It seems reasonable to assume the systematic variability in spot intensity levels among the 30 arrays is dominated by the grouping variable or the concentration level. This suggests the best normalization approach would be to separately normalize the six sets of five arrays separately. This can be done with each of the three normalization techniques: quantile, cyclic loess and fastlo. Extending fastlo allows us to perform a single normalization for the entire set of 30 arrays. There are three basic models that could be used to accomplish this.

Let  $k$  index the treatment group and  $g$  the genes; the array index  $j$  is nested within  $k$  and the spot index  $i$  is nested within  $g$  (11–20 PM spots per gene for the Affymetrix GeneChip platform). Simple fastlo on all 30 arrays together corresponds to the linear model  $y_{ij} = \alpha_i + \epsilon_{ij}$ . Conversely, the model  $y_{ij} = \alpha_i + \beta_k + (\alpha\beta)_{ik(j)} + \epsilon_{ij}$ , containing all treatment and spot interactions, is equivalent to separate fastlo fits to each treatment group. Intermediate models include  $y_{ij} = \alpha_i + \beta_{k(j)} + \epsilon_{ij}$ , a simple main effect for treatment or  $y_{ij} = \alpha_i + \beta_{k(j)} + (\alpha\beta)_{g(i)k(j)}$ , which has a separate treatment effect per gene. [In the model,  $k(j)$  is treatment  $k$  for the sample hybridized to array  $j$  and  $g(i)$  is defined similarly.] The latter model may be more plausible, in that treatment would affect a gene, not an individual probe.

As proof of principle, we modified the fastlo algorithm to fit this last model and applied it to all 30 arrays in the GeneLogic liver dilution dataset. The modified fastlo algorithm is called fastlo2. Figure 6 compares the normalized  $\log_2$  spot intensity values for the set of five arrays with a total cRNA of  $10 \mu\text{g}$  produced by four different normalization approaches: cyclic loess, quantile, fastlo, and fastlo2. Cyclic loess, quantile, and fastlo were done separately on each set of five arrays

corresponding to the six different concentration levels, and fastlo2 was performed on the entire set of 30 arrays simultaneously. The distribution of the normalized spot intensities produced by the four methods are essentially equivalent. Similar results (data not shown) were observed for the five other concentration levels. Table 2 summarizes the distribution of the standard deviations of all the spots for the four methods. The summary values for fastlo2 were remarkably similar to the other three methods. In particular, the results produced by fastlo2 were generally smaller than those of quantile but larger than cyclic. As observed earlier, the differences in these distributions are so small that they are not of practical significance.

## 7 DISCUSSION

We have accomplished our goal of developing a method, fastlo, that yields normalized values similar to those produced by cyclic loess normalization but is considerably faster. Results from both the simulated data and the real data indicate that cyclic loess and fastlo produced equivalent normalizations. Fastlo only requires  $n$  loess smooths for one pass through a set of  $n$  arrays compared to  $C_2^n$  loess smooths required by cyclic loess. This results in at least an order of magnitude in computational (and time) savings. In addition, fastlo appears to have the same order of magnitude in computation (and run-time) as quantile normalization; the relative speed depends on the implementation details of fastlo and quantile normalization.

All three methods appear to yield similar normalized values. It is worth noticing that the datasets used above were ideal for quantile normalization; the true spot intensities (and so their distributions) were identical across the arrays making this perfect for a method that forces the spot intensity distributions of the arrays to be identical. In situations where spot intensities



do differ among tissue initially thought to be similar, e.g. a molecular subtype within brain cancer tissue, quantile normalization may be too aggressive and attenuate differences of interest. These issues are yet to be resolved.

An advantage of fastlo not shared by the other methods is that it is model-based. As demonstrated above, this allows us to normalize an entire set of arrays to be compared as a group—i.e. no need to perform separate normalizations for subgroups. This is useful in situations where it is not clear to what degree the biological variation between the groups and the microarray experimental process contribute to the systematic variation among the arrays. This model can be easily extended to incorporate other variables that are part of the underlying experimental design such as subject demographic variables, time, dose level, etc. Other estimation options are also open, e.g. shrinkage of the treatment coefficients toward zero in an experiment where it was felt that only a small percentage of genes are differentially expressed.

Finally, a more robust estimator could be used to generate  $\hat{y}$ . If the number of arrays to be normalized is relatively small, the mean estimate can be significantly influenced by a single outlier. We explored this using the median of spot  $i$  across all arrays to generate the reference array. The results were not promising; they resulted in a considerably biased estimate. Recall the data to be normalized is a matrix of arrays as columns and spot intensities as rows. Fastlo iterates between fits to the rows ( $\hat{y}$ ) and loess smooths on the columns of the array. It appears that the same type of estimator must be used for row and column operations in the algorithm. We have not pursued this further.

Overall, fastlo (1) produces normalized results similar to cyclic loess (as well as quantile) normalization, (2) is considerably faster than cyclic loess, and (3) has added versatility

above quantile and cyclic loess normalization through its connection with linear models.

## REFERENCES

- Alter, O., Brown, P.O. and Botstein, D. (2000) Singular value decomposition for genome-wide expression processing and modeling. *Proc. Natl Acad. Sci., USA*, **97**, 10101–10106.
- Bolstad, B., Irizarry, R., Astrand, M. and Speed, T. (2003) A comparison of normalization methods for high density oligonucleotide array data based on bias and variance. *Bioinformatics*, **19**, 185–193.
- Cope, L.M., Irizarry, R.A., Jaffee, H., Wu, Z. and Speed, T.P. (2003) A benchmark for Affymetrix genechip expression measures. *Bioinformatics*, **1**, 1–13.
- Dudoit, S., Yang, Y.H., Callow, M.J. and Speed, T.P. (2002) Statistical methods for identifying differentially expressed genes in replicated cDNA microarray experiments. *Stat. Sinica*, **12**, 111–139.
- Hartemink, A., Gifford, D., Jaakkola, T. and Young, R. (2001) Maximum likelihood estimation of optimal scaling factors for expression array normalization. In *SPIE BIOS 2001*, San Jose, CA, USA, **4266**, 132–140.
- Irizarry, R., Hobbs, B., Collins, F., Beazer-Barclay, Y., Anntonellis, K., Scherf, U. and Speed, T. (2003) Exploration, normalization, and summaries of high density oligonucleotide array probe level data. *Biostatistics*, **4**, 249–264.
- Newton, M.A., Kendziorski, C.M., Richmond, C.S., Blattner, F.R. and Tsui, K.W. (2001) On differential variability of expression ratios: improving statistical inference about gene expression changes from microarray data. *J. Comput. Biol.*, **8**, 37–52.
- Richmond, T. and Somerville, S. (2000) Chasing the dream: plant EST microarrays. *Curr. Opin. Plant Biol.*, **3**, 108–116.
- Yang, Y.H., Dudoit, S., Luu, P., Lin, D.M., Peng, V., Ngai, J. and Speed, T. (2002) Normalization for cDNA microarray data: a robust composite method addressing single and multiple slide systematic variation. *Nucleic Acids Res.*, **30**, e15.