"Comprestimation": Microarray Images in Abundance.

Rebecka Jörnsten and Bin Yu[1]
Department of Statistics
University of California at Berkeley
Berkeley, CA, 94720
rebecka/binyu@stat.berkeley.edu

Abstract — **Microarray analysis is a powerful technique for the classification and identification of gene expressions. Recently significant progress has been made in the development of microarray instruments, and as the popularity of the technique is soaring, an impressive amount of gene expression data is collected at numerous labs across the world. One of the most powerful microarray instrument on the market today can, in the space of twelve hours, extract information from 10000 genes, corresponding to two 30 Mb files. Often considerable data reduction is used, extracting information from the microarray images, and forming a single table of summary statistics for each gene on the array. However, there is still debate as to what the *informative* summary statistics are. The recent explosion of interest in the technique opens questions to determination of standards, data sharing and optimal exchange of information between researchers in the field. A compression routine that either fully preserves all data features, or nearly so, can enable researchers to modify their method of analysis of gene expression data at a later stage, as well as simplify distribution of data. The components that constitute our efficient compression routine also provide information on data quality, and can be thought of as a multiscale description of microarray images. We have constructed a lossless compression scheme with compression ratios of 1:25. Preliminary results of lossy schemes, that provide near faithful representation of data features, give compression ratios of 1:35.**

I. Introduction

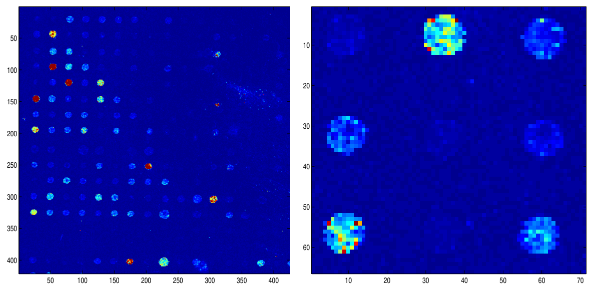

Figure 1: Subset of a microarray, and close-up.

As the Human Genome Project is drawing to a conclusion, the quest for identification of genes move on to identification of their functions. Microarray analysis is a powerful tool for this purpose. An array is prepared by automatic printing, or "spotting", of thousands of cDNA (complimentary DNA) fragments on a glass slide. These spots provide gene specific targets. Samples of mRNA (messenger) from a cell of interest (e.g. cancerous) , and a baseline cell, can be labeled using fluor tags that emit light of different wavelengths during excitation. The mRNA fragments can then be allowed to hybridize onto the array. Excitation of the fluor-labeled targets are done with a laser scan of set wavelengths (dual wavelength scan, or two separate scans), and two intensity images (referred to as Red and Green (R, G)) are recorded. The intensity at each spot location is proportional to the amount of hybridization that occurred, and therefore related to the gene expression level or activity. Intensity ratios, R/G, are the quantity of interest, taken as a measurement of the differential gene expression between the baseline cell and the cell under investigation. Ratios instead of differences are preferred since the overall response level may vary between experimental settings. Researchers are especially interested in detecting small differential expression levels (intensity ratios near 1).

The data consists of two 16 bpp intensity images. Within each image are located higher intensity regions corresponding to the array genes (Figure 1). Due to experimental variation in the spotting procedure, the mRNA extraction, the hybridization and post-processing clean-up, the regions of interest (ROI's) are often irregular, noisy and corrupted by smoothly varying background. Some standard analysis tools use a fixed segmentation of the data into regions of interest and surrounding background. However, for small differential gene expressions improper segmentation can seriously limit the chances of detection. Given a segmentation of the intensity images, differential expression levels are calculated as follows. Pixels in each image are summed up within an ROI. Since each scan results in a specific background noise level, that may also vary locally, a background correction is applied. The differential expression level, R/G (referring to Red and Green fluor tags), is then calculated as the ratio of the summed ROI intensities;

$$R/G = \frac{\sum_{y_i \in ROI} Y_i - BgX}{\sum_{x_i \in ROI} X_i - BgY},$$

where Bg refers to the estimate of the local background and X and Y refers to the individual intensity images. Simulation studies show that even slight contamination of the ROI with pure noise pixels can greatly reduce the power of detection of small differential expressions levels.

II. Lossless and Lossy Compression schemes.

Image compression tools for natural images provide virtually no visual distortion. However, these tools cannot be expected to be efficient for compression of microarray images.

[1]Bell Laboratories, Lucent Technologies, Murray Hill, NJ

Microarray images consist of a fairly regular grid of many regions of high, and strictly local, variability. Within are features we want to preserve with high accuracy, without we are happy to capture overall trends and smooth variation. In medical imaging preservation of such *important* features is often not only desirable, but a condition set by researchers. The routines used here have the flavor of medical image compression, though novel tools have to be developed to take advantage of particularly unique features of microarray images. In general, wavelet image compression tools are not appropriate for microarray data due to the highly irregular shapes, and limited size of the ROI's.

The schematic description of our compression routine is as follows.

- **Registration**: Due to inaccuracy of the automatic printing procedure, the ROI's are not placed on a regular grid. A search for local maxima is performed on the image to estimate the experimental grid.
- **Segmentation**: The irregular shapes of the ROI's necessitates segmentation of the data near a located grid point. The shape of the determined ROI is encoded using chain coding.
- **Background Estimation**: The distribution of the local background needs to be estimated for proper background correction. The local background, with ROI's removed, is encoded with a lossy or lossless procedure.
- **Denoising**: The local background distribution provides information about the additive noise (scanning). We apply soft thresholding within the ROI in order to reduce noise while also saving on encoding cost.
- **Predictive Coding**: There is strong spatial correlation within the ROI's, such that predictive coding provides substantial savings for encoding purposes. For ROI's with sharp transition from the background a *boundary predictive coder* is used.
- **Quantization**: For ROI's with substantial signal-to-noise ratio (SNR), and with large estimated differential expression levels, a lossy scheme is employed to further reduce the coding cost.

Methods used for registration and segmentation will not be discussed in detail here. The registration is done automatically on the whole microarray using input parameters from the array design, such as spot distance and spot width ([11]).

For segmentation, two methods are used in conjunction. A seeded region growing algorithm provide a coarse initial segmentation. Seeds for background and ROI are available from the estimation of the experimental grid. For a fully automatic procedure to work unsupervised, the segmentation needs be overly generous initially. This leads to some inclusion of pure noise, i.e. background intensity pixels, among the signal, high intensity pixels, within the ROI. Improvement on the segmentation performance is done locally by fitting a mixture of Gaussians to the coarsely defined ROI. If the separation between the mixtures is significant, we classify the dominant mixture as *signal*, and the lower mixture as noise. We remove pixel points on the boundary of the ROI if they are classified as noise, and update our definition of the ROI accordingly. Noise pixels may also be located within the ROI due to improper mixing of the mRNA fragments, and other experimental factors. The occurrence of noise pixels within the ROI varies from data set to data set. The Gaussian mixture model seemed very accurate at identifying these pixels, while also providing a natural signal-noise threshold. When intensity ratios are calculated these internal noise pixels are not included. For encoding purposes, the ROI is encoded in full, including internal noise pixels.

The local background distributions are estimated from the regions the furthest removed from any ROI ("valleys between peaks") ([11]). The median of the identified local background is used as an estimate of the local noise level. This correction factor is removed from the corresponding ROI in order to put signal measurements of both images on the same scale.

The **background** with ROI's removed is encoded with a lossy or lossless scheme. A 3-level S+P transform ([8]) is used, keeping only the two lowest levels of the transform for reconstruction in the lossy case. This approach could obviously be modified using quantization of higher subbands, or zerotree encoding. The simplistic approach preserves the large scale features of the background that experimentalists use for quality assessment, while removing speckle dust and other noisy features. We describe the transform in the one dimensional case, extensions to two dimensions is equivalent to performing the transform sequentially on rows and columns. The S transform of data string $c[n]$, $n = 0, \cdots, N-1$ is:

$$l[n] = \lfloor (c[2n] - c[2n+1])/2 \rfloor$$

$$h[n] = c[2n] - c[2n+1].$$

To further reduce first order entropy, a prediction of the high-pass h is obtained (P transform). Referring to the predictor of h as $\hat{h}$ we can replace h in the transform by the difference

$$h_d[n] = h[n] - \lfloor \hat{h}[n] + 1/2 \rfloor.$$

The estimate $\hat{h}[n]$ is given by

$$\hat{h}[n] = .25\Delta l[n] + .25\Delta l[n+1],$$

where $\Delta l[n] = l[n-1] - l[n]$. Other prediction filters may also be applied ([8]). The inverse transform is designed for perfect reconstruction.

The noise and signal data are well approximated by discretized and truncated Gaussians. We use this approximation of the data distribution to design a **denoising** procedure. We assume that the observed signal pixels Y_i are a noisy version of the true signal T_i, $i = 1, \cdot, N^2$, with mean μ and variance σ_T^2. The additive noise ϵ_i, $i = 1, \cdot, N^2$ has mean 0 and variance σ^2. We want to compute the optimal soft threshold δ such that the thresholded estimate $\hat{T}_i$ has minimum Bayes risk, i.e. we want to find δ such that

$$E_T E_{Y|T}(\hat{T} - T)^2$$

is minimized, where

$$\hat{T}_i = sign(Y_i - \bar{Y}) max(0, |Y_i - \bar{Y}| - \delta).$$

An approximation of the optimal threshold is given by ([3])

$$\delta^* = \sigma^2 / \sigma_T.$$

The noise variance, σ^2, is estimated from the local background distribution. The discretization of the data is taken into account by parametrizing bin probabilities to avoid serious bias of the estimator. The signal variance, σ_T^2, is estimated by the difference between the variance of the signal component of the fitted Gaussian mixture component in the ROI, minus the estimated noise variance. Denoising compresses the scale of the data within the ROI, as well as creating values in a "zero-zone". This reduces the entropy of the ROI and enables more efficient coding, while also having an intuitive appeal for removing excess variation in the ROI.

To encode the ROI losslessly we employ a **predictive coding** scheme in the signal domain. The causal neighborhood of each ROI pixel provides us with numerous possible predictors. A median of predictor values based on context pixels seems to provide sufficient predictive power for the lossless encoder, i.e.

$$\hat{y}_i = median\ \lfloor p_{i,j} \rfloor\ j = 1, \cdots, J,$$

where $p_{i,j}$ refers to the prediction of pixel i using predictor model j. The models we use are of the type

$$p_{i,j} = \sum_{k \in N_i} \alpha_{k,j} y_k,$$

where N_i refers to a context neighborhood of pixel i, and $\alpha_{.,j}$ refers to weights allotted to context pixels in predictor model j. E.g. the mean predictor model would have equal weights such that $\alpha_{k,j} = \alpha_{k',j}, \forall k, k' \in N_i, \sum \alpha_{k,j} = 1$.
We also try more refined schemes with learning, where the accuracy of a predictor is traced throughout the ROI, followed by re-weighting of the predictors accordingly. However, though offering slight improvements in bit-rate, no one prediction scheme was found to be overall best for compression performance. Success of a model varies between ROI's and we are so far unable to define specific markers, or selectors, other than testing each ROI individually. For pixels on or near the boundary a weight of the overall ROI mean and background median is used, greatly reducing encoding cost.

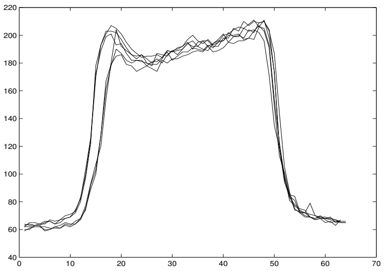

Figure 2: Profile of ROI

For ROIs with sharp transitions from the background level standard context based prediction coding schemes fail, resulting in the overall coding cost of the ROI being completely dominated by its boundary. In order to remedy this behaviour we design a **boundary predictive coder**. The predictive coder traces out the boundary, predicting along the trace, working in a spiral towards the center of the ROI. When the inner smooth region of the ROI is reached, as determined by estimates of the local gradient, the boundary predictive coder is replaced by a context based scheme (above) until the entire ROI is encoded. Considerable saving is achieved with this method. By training on the data we are able to determine a simple threshold value of the signal-to-noise ratio of the ROI's, above which gain is made using a boundary prediction scheme.

For ROI's with high signal-to-noise ratio and estimates of the differential expression level far from 1, further savings can be achieved if we allow a **lossy compression** of the ROI's. The high signal to noise ratio ROI's are inherently grainy and noisy, which makes these costly to encode with a lossless scheme. However, a lossy scheme can be designed such that the overall features of the ROI's are preserved.

Two approaches are investigated here. The saving in bit-rate is achieved through quantization. We quantize each image separately. The level of quantization is selected based on methods as presented in [7]. Firstly, we design a quantizer based on the assumed approximate Gaussian distribution of the data. The optimal quantizer (Lloyd-Max) for Gaussian data with variance σ^2 is known to be a scaled version of the quantizer for variance equal to 1. We compute a rate-distortion curve for quantization of Gaussian data with variance 1 and pick a chosen acceptable distortion level on the curve, at which the slope is denoted by λ and the quantizer by $Q(1) = [B(1), C(1)]$. B denotes the quantizer partition and C the corresponding centroids, or reconstruction levels. For each ROI we quantize, we match a point with slope λ/σ^2 on the curve to a scaled quantizer $Q(\sigma) = [\sigma B, \sigma C]$, thus providing overall the same distortion level for each quantized ROI ([7] and references therein). We refer to this lossy scheme as quantization in the signal domain, and set the level of distortion such that the data within the ROI is near faithfully represented.

A second approach consists of using an S+P transform, as described above, within the ROI. The transform was actually performed on the minimum square containing the ROI fully. Since the shape of the ROI is encoded separately we only need to encode the transform coefficients that have an influence on the reconstruction of the pixels within the ROI. The non-LL subbands were quantized using a similar rate-distortion optimal scheme as described above, but now using a Laplacian distribution for the quantizer design. The Laplacian distribution more accurately describes the distribution of the non-LL subbands of the S+P transformed images. The lowest level was not quantized. We refer to this approach as quantization in the transform domain. The two methods are compared for the same bit-rate, which also gives near equal distortion, as calculated for each image.

III. Results

We here give some compression results for the microarray image shown in Figure 1. The test data consisted of 16×16 spots. A complete microarray can be viewed on `http://www.stat.berkeley.edu/~rebecka`. The full image file usually contains a larger margin of background such that a smaller proportion of a full array is taken up by ROI's than in our test image. One can therefore assume that the results provided here are fairly conservative estimates of the final compression ratio. The lossless compression ratio we attain is 1:25 such that a 30 Mb file can be reduced to 1.2Mb.

Table 1 shows the compression ratios using predictive coding, boundary prediction, denoising, and the two methods of quantization (SD for quantization in the signal domain, and TD for quantization in the transform domain). We provide a total bit-rate for both images (32 bpp in original), and the two images separately (16 bpp original). The total bit-rate includes the cost of encoding the ROI's of each image, the backgrounds and the shapes of the ROI's. For the individual images the cost of the ROI's and the background is included. The two lossy schemes described in the previous section were run in order to produce the same distortion as computed by mean squared error, such that their output is directly comparable. In the remainder of this section we will show some features of each method in figures.

Method	Total	R Scan	G scan	Compr. Ratio
Lossless	1.31	0.61	0.67	1:24
Boundary	1.29	0.59	0.65	1:25
Denoising	1.21	0.57	0.62	1:27
Lossy SD	0.83	0.38	0.43	1:35
Lossy TD	0.84	0.39	0.43	1:35

Table 1: Bit-rate for different compression schemes.

Method	R Scan	G scan
Lossless	2.15 (1:7)	2.02 (1:8)
Denoising	1.91 (1:8)	1.71 (1:9)
Lossy	1.47 (1:11)	1.38 (1:12)

Table 2: Bit-rate (compr. ratio) within the ROI's

Within the ROI, the compression ratios are of course much lower than for the whole image. Considerable part of the saving achieved results from proper segmentation of background and ROI's. Table 2 gives the compression ratios within ROI's for the two images. For a lossless scheme we attain average compression ratios of 1:8, while a lossy scheme can further increase this to 1:12. Lossy schemes thus offer little additional bit-rate reduction, though improvement of the lossy design is expected as our investigation continues. The ROI's constitute 25 % of the microarray image, but roughly 50 % of the total encoding cost. The background bit-rate for the lossy scheme is about .35 bpp, and preserves all large scale features.

a) b)

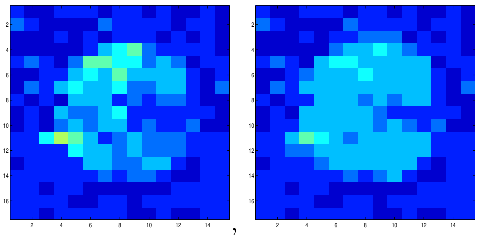

Figure 3: a) Raw. b) Denoised.

Denoising has a marked effect on lower signal-to-noise ratio ROI's. In Figure 3 we see an example of this. Denoising lifts out the ROI, and the grainy structure is reduced. As can be seen from table 1, denoising also increases the compression ratio from 1:25 to 1:27. This seemingly marginal reduction corresponds to a saving of .1 Mb for these very large images.

a) b)

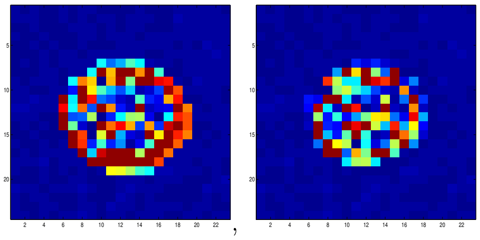

Figure 4: a) Errors with Predictive coding. b) Errors with Boundary prediction.

The boundary predictor also provides a reduction of encoding cost. The sharper the boundary, the more gain a boundary predictor achieves over a context based predictive coder, as figure 4 clearly indicates. For other microarray images (see URL), with higher overall signal intensities, the boundary predictor becomes an even more important tool, providing substantial saving in encoding cost.

For lossy compression it was found that the first quantization method (SD) works better for lower signal-to-noise ratio and complex shaped, or small ROI's. For high signal-to-noise ratio ROI's working in the transform domain (TD) was preferable, but only when the ROI was of considerable size, and reasonably smooth. For lower bit-rates, downsampling and quantizing created a blocky effect in the ROI's quantized in the transform domain, whereas working in the signal domain emphasized the spread of the signal measurements in the ROI on a pixel level but created large patches of over smoothed regions (see Figure 5).

a) b) c)

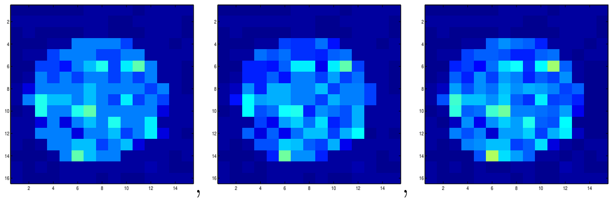

Figure 5: a) Quantizing, SD. b) Quantizing, TD. c) No quantization.

However, although both methods provide near identical results in terms of MSE (mean squared error) and bit-rate, we are forgetting that the goal of compression here is to preserve all important features of the data set. When the two quantized images are compared we find that separate quantization leads to a distortion *between* the images, in the sense that pixel cross-correlation estimates based on quantized data are not similar to those based on the full data. It was found that quantizing in the signal domain preserves the cross-correlation structure very poorly. Quantizing in the transform domain performs better, but still tends to seriously over- or underestimate this parameter. How we should design an optimal lossy compression scheme that preserves these cross-image features is a question for future work. The goal is to compress the data such that we are enabling not only data quality assessment on compressed data, but also ensuring that inferences, statistical modeling and estimation based on compressed data will not differ much, or at all, from inferences based on the full data. This should also uphold for a wide range of possible analyses.

The lossless and lossy compression routines for microarray data offer a form of *multiscale* description of the data. On the coarsest scale, we reduce the microarrays to equi-sized ROI's with a constant intensity level. This corresponds to the table output of some software used in the field, here displayed as an image. For data quality assessment the *shapes* of the ROI's can be added. Complex shapes indicate poor quality of data. Information such as signal-to-noise ratios, local background levels and variances are also diagnostics that can be provided. The within ROI variability can be illustrated by adding the lossy, or losslessly encoded information to the display.

IV. Future Work

This work is only in its preliminary stages and a lot of highly interesting questions have revealed themselves during the investigation of microarray images. Lossy quantization of the data should be developed further, considering both progressive schemes and fixed bit-rate schemes. The question of major interest that is facing us is whether compressed data can still provide us with sufficient information to perform statistical modeling and testing on the data. Quantization is done separately on the two images in this initial work, yet the final analysis is concerned with combining information from both images. It is not necessary that a separate scheme, optimal for compression of each image, is optimal for the purpose of inference based on the output of the two quantizers together, as we found concerning estimates of the cross correlation in the previous section. The cross-correlation between the pixels in the individual images is of interest to some researchers for data quality assessment. *Alternative estimators* of the differential expression levels, such as regression models, mean of ratio intensities (as opposed to ratio of the mean intensities), involve evaluation of this parameter. A first order approximation of $Var(R/G)$ is given by

$$Var(R/G) = (\ Var(R) + (\mu_R/\mu_G)^2 Var(G)$$
$$-2(\mu_R/\mu_G)\rho\sqrt{Var(R)Var(G)}\ )/(n\mu_G^2),$$

where ρ denotes the cross-correlation between the two images, and n the number of pixels within the ROI. Simulations with Gaussian distributions show that this estimate of the variance of the differential expression level holds for a wide range of signal-to-noise ratios and values of expression levels. Construction of confidence levels for differential expression levels, and testing, should ideally provide the same results for compressed data and the full data. Quantizing the two images separately, however, may affect the estimate of the cross-correlation and thereby construction of confidence levels and testing, thus *not* preserving one very important data feature. We need to investigate a range of testing situations, and quantizer designs in order to make an informed judgement on how inference is affected by compression.

We can also perceive wanting to preserve the features of an image of ratio intensities. The distortion of such an image under quantization can be approximated as follows.

$$E(\frac{\hat{R}}{\hat{G}} - \frac{R}{G})^2 =$$

$$(E(\hat{R}-R)^2 + (\mu_R/\mu_G)^2 E(\hat{G}-G)^2)/(Var(G)+\mu_G^2).$$

This expression has an intuitive appeal. When the mean ratio $\frac{\mu_R}{\mu_G}$ is near 1 the distortion of each image has near equal influence on the distortion of the ratio image. When the mean ratio is large, the distortion of the denominator dominates the distortion of the ratio image. When the mean ratio is small, the numerator distortion dominates. In order to minimize overall distortion of the ratio image we should therefore allow different distortion of the individual images.

Since separate quantization of the images may lead to distortion of important data features such as cross-correlation, joint quantization schemes should be defined. However, should separate quantization be performed on the images we may still investigate how this affects the estimates of parameters of interest, or inference such as testing of significance of differential expression levels. The *Multiterminal data compression* framework applies in this setting. Multiterminal data compression is a new and exciting field, in which very few papers have been published up to date ([2] and references therein). Some results exist for testing and estimation in the case of discrete data distributions, and separate quantization, although many important questions are still unanswered. We aim to develop results that reveal loss of estimation efficiency of a parameter due to separate quantization, and investigate optimal design of quantizers in this setting. For the application to microarray data, assuming the data is well approximated by discretized Gaussians, we may use parametrized probabilities to form a discrete distribution. We hope to be able to extend previous results for discrete distributions to our current application. The question of continuous data is still an open problem that we hope to address at a later date. With results such as these, a simple and separate quantization scheme might be adopted, while we could also provide bounds on loss of estimation efficiency, and decrease of power of testing. These bounds can be considered as giving an assessment of how faithfully the compressed data preserves important data features. Future work will go into investigation of such efficiency bounds.

On the practical coding side of the matter, we have seen that considerable compression is possible while preserving most, or all data features. If the compressed data had to be fully reconstructed before researchers could study the nature of a data subset of interest, such as all ROIs with barely significant differential expression levels, this would be unfortunate. We therefore propose to construct a tool that enables direct access to subsets of data, with reconstruction *only* on that subset. This seems a most appealing approach to handling particularly large data sets, such as arrays with thousands of genes. The compression scheme we have designed separates features of the data such as shape, signal-to-noise ratio, mean intensity level, and background. A header of the compressed file can thus allow us to store features for each ROI and indicate locations in the bit-stream where they are located. Compared to the whole compressed file size such a header is negligible. Storing the data in this format enables us to access particular ROI's on a request basis. We can fetch ROI's with complex shapes, with low SNRs, at a specific grid point, with sharp transitions, just to mention a few features of interest.

To fully take advantage of a significant data reduction, such as the lossless scheme we have described here, one should avoid reconstruction whenever possible. We therefore propose to investigate to what extent features of the data can be extracted from the compressed data format itself. As working with orthogonal transforms offers alternate descriptions of data where computations may be performed more efficiently, there may be estimation questions that can be answered while working on compressed data. When we can design compression routines that are efficient, preserve features within *and* between data sets, and allows computations and estimation to be performed on compressed data, then we may truly coin the concept *"Comprestimation"*.

V. Conclusion

We have constructed a highly successful lossless and lossy compression scheme for microarray images. Lossless compression ratios of 1:25 were observed. The compression routine also provides a convenient multiscale description of the data. We think researchers in the field will find this very useful and informative.

We hope that further investigation into compression and estimation of microarray images will prove fruitful on many levels. Obviously there is a great need for efficient compression when data of this quantity is collected daily, and in separate locations. However, an interesting question on a more theoretical level has also been revealed. In many compression situations, compression is done separately for components of data. But how can we assure ourselves that separate treatment of components does not seriously affect the inferences we make? It is of great interest to be able to quantify the loss of estimation efficiency due to compression, and ultimately with confidence be able to state that we have an understanding of the nature of distortion that compression induces.

Acknowledgments

Microarray data was provided by the Human Genome Project group at LLNL, Lawrence Livermore National Labs, Livermore, California, and Elva Siaz at Tito Serafini's lab, MCB department, UC Berkeley, California.

Bell labs, Lucent Technologies were kindly hosting Ms. Jörnsten while the work presented here was conducted.

References

[1] Adams and Bischof "Seeded region growing", IEEE trans on Pattern Analysis and Machine Intelligence 16, 641-647.

[2] S. Amari and T. Han "Statistical Inference Under Multiterminal Data compression", IEEE IT Vol 44(6), 1998

[3] G. Chang, B. Yu and M. Vetterli "Adaptive Wavelet Thresholding for Image Denoising and Compression.", IEEE trans of Image Processing, to appear.

[4] C. Lu and J. Dunham "Highly Efficient Coding Schemes for Contour Lines Based on Chain Code Representation", IEEE trans Comm, Vol 39(10) 1511-1514, 1991

[5] T. Nakachi et al. "Pel Adaptive Predictive Coding Based on Image Segmentation for Lossless Compression", IEICE trans fundamentals, Vol E82(6) 1037-1045, 1999

[6] D. Nister and C. Christopoulos "Lossless Region of Interest Coding", Signal Processing 78 1-17, 1999

[7] A. Ortega and K. Ramchandran, "Image and Video Compression", IEEE Signal Processing Magazine, 1998

[8] A. Said and W. Perlman "Reversible image compression via multiresolution representation and predictive coding", Proc SPIE 2094, 664-674, 1993

[9] L. Shen and R. Rangayyan "A Segmentation-Based Lossless Image Coding Method for High-Resolution Medical Image Compression", Vol 16(3) 301-307, 1997

[10] J. Ström and P. Cosman "Medical Image Compression with Lossless Regions of Interest", Signal Processing 59(2) 155-171, 1997

[11] Y. Yang and T. Speed " 'Zarray', Image processing on microarray data", preprint