10-25-2006

ESTIMATING GENOME-WIDE COPY NUMBER USING ALLELE SPECIFIC MIXTURE MODELS

Wenyi Wang

Johns Hopkins Bloomberg School of Public Health, Department of Biostatistics

Benilton Carvalho

Johns Hopkins Bloomberg School of Public Health, Department of Biostatistics

Nate Miller

Johns Hopkins University, School of Medicine, Department of Neuroscience

Jonathan Pevsner

The Johns Hopkins Medical Institutions, Kennedy Krieger Institute

Aravinda Chakravarti

McKusick-Nathans Institute of Genetic Medicine, Johns Hopkins University School of Medicine

See next page for additional authors

Suggested Citation

Wang, Wenyi; Carvalho, Benilton; Miller, Nate; Pevsner, Jonathan; Chakravarti, Aravinda; and Irizarry, Rafael A., "ESTIMATING GENOME-WIDE COPY NUMBER USING ALLELE SPECIFIC MIXTURE MODELS" (October 2006). *Johns Hopkins University, Dept. of Biostatistics Working Papers*. Working Paper 122.
<http://biostats.bepress.com/jhubiostat/paper122>

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

Copyright © 2011 by the authors

Authors

Wenyi Wang , Benilton Carvalho, Nate Miller, Jonathan Pevsner, Aravinda Chakravarti, and Rafael A. Irizarry

Estimating Genome-wide Copy Number using Allele Specific Mixture Models

Wenyi Wang¹, Benilton Carvalho², Nate Miller³, Jonathan Pevsner⁴, Aravinda Chakravarti⁵, and Rafael A. Irizarry⁶

¹ Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health,
615 North Wolfe Street, Baltimore, MD 21205, USA,
wwang2@jhsph.edu,

² Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health,
615 North Wolfe Street, Baltimore, MD 21205, USA,
bcarvalh@jhsph.edu,

³ Department of Neurology, Kennedy Krieger Institute, 707 North Broadway,
Baltimore, MD 21205, USA,
natemiller@jhmi.edu

⁴ Department of Neurology, Kennedy Krieger Institute, 707 North Broadway,
Baltimore, MD 21205, USA,
pevsner@kennedykrieger.org

⁵ McKusick-Nathans Institute of Genetic Medicine, Johns Hopkins University School
of Medicine Broadway Research Building, Suite 579, 733 N. Broadway, Baltimore,
MD 21205, USA,
aravinda@jhmi.edu

⁶ To whom correspondence should be addressed:
Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, 615
North Wolfe Street, Baltimore, MD 21205, USA,
rafa@jhu.edu

Abstract. Genomic changes such as copy number alterations are thought to be one of the major underlying causes of human phenotypic variation among normal and disease subjects [23, 11, 25, 26, 5, 4, 7, 18]. These include chromosomal regions with so-called copy number alterations: instead of the expected two copies, a section of the chromosome for a particular individual may have zero copies (homozygous deletion), one copy (hemizygous deletions), or more than two copies (amplifications). The canonical example is Down syndrome which is caused by an extra copy of chromosome 21. Identification of such abnormalities in smaller regions has been of great interest, because it is believed to be an underlying cause of cancer.

More than one decade ago comparative genomic hybridization (CGH) technology was developed to detect copy number changes in a high-throughput fashion. However, this technology only provides a 10 MB resolution which limits the ability to detect copy number alterations spanning small regions. It is widely believed that a copy number alteration as small as one base can have significant downstream effects, thus microarray manufacturers have developed technologies that provide much higher resolution. Unfortunately, strong probe effects and variation introduced by sample preparation procedures have made single-point copy number

estimates too imprecise to be useful. CGH arrays use a two-color hybridization, usually comparing a sample of interest to a reference sample, which to some degree removes the probe effect. However, the resolution is not nearly high enough to provide single-point copy number estimates. Various groups have proposed statistical procedures that pool data from neighboring locations to successfully improve precision. However, these procedures need to average across relatively large regions to work effectively thus greatly reducing the resolution. Recently, regression-type models that account for probe-effect have been proposed and appear to improve accuracy as well as precision. In this paper, we propose a mixture model solution specifically designed for single-point estimation, that provides various advantages over the existing methodology. We use a 314 sample database, constructed with public datasets, to motivate and fit models for the conditional distribution of the observed intensities given allele specific copy numbers. With the estimated models in place we can compute posterior probabilities that provide a useful prediction rule as well as a confidence measure for each call. Software to implement this procedure will be available in the Bioconductor `oligo` package (<http://www.bioconductor.org>).

1 Introduction

The demand for technologies that provide high-resolution measurements for copy number estimates has driven microarray manufacturers such as Illumina and Nimblegen to develop CGH-SNP microarrays [20, 6, 24]. Although Affymetrix has not yet developed a product specific for copy number analysis, various groups, including Affymetrix, have described statistical methodology that make use of their SNP chips, originally developed for genotyping [28, 1, 9, 19, 13, 15, 8, 16], to provide successful copy number estimation algorithms. An advantage of the SNP chip technology is that, given the popularity of the genotyping application, the protocols used to prepare the genomic DNA samples is well developed and tested. Another advantage is that we can obtain genotype calls which permits allele specific copy number estimation which in turn can be used to predict parent specific copy number [16].

The genotyping platform provided by Affymetrix interrogates hundreds of thousands of human single nucleotide polymorphisms (SNPs) on a microarray. DNA is obtained and fragmented at known locations so that the SNPs are far from the ends of these fragments, the fragmented DNA is amplified with a polymerase chain reaction (PCR), and the sample is labeled and hybridized to an array containing probes designed to interrogate the resulting fragments. We refer to the measurements obtained from these probes as the *feature intensities*. There are currently three products available from Affymetrix: an array covering approximately 10,000 SNPs (GeneChip Human Mapping 10K), a pair of arrays covering approximately 100,000 SNPs (GeneChip Human Mapping 50K Xba and Hind Array), and a pair of arrays covering approximately 500,000 SNPs (GeneChip Human Mapping 250K Nsp Array and Sty Array). These are referred to as the 10K, 100K, and 500K chips respectively.

Fig. 1. Log (base 2) intensity, $S_{i,j}$ plotted against chromosomal position of SNP i . Chromosomes 20, 21, 22, and X are shown for male with Down syndrome. Notice that we expect values from chromosome 21 and X to be higher (3 copies) and lower (1 copy) respectively.

To motivate the model and estimation procedures described here we need to understand the basics of the feature-level data. We provide the essential details here and refer readers to Kennedy et al. [14] for a complete description. Each SNP on the array is represented by a collection of probe quartets. As with Affymetrix expression arrays, the probes are defined by 25-mer oligonucleotide molecules referred to as perfect match (*PM*) probes⁷. A difference with expression arrays is that *PM* probes differ in three important ways. First, two alleles are interrogated (for most SNPs only two alleles are observed in nature). These are denoted by *A* and *B* and divide the probes into two groups of equal size. For each *PM* probe representing the *A* allele there is an allele *B* that differs by just one base pair (the SNP). Second, features are included to represent the sense and antisense strands. This difference divides the probes into two groups that are not necessarily of the same size. Finally, for each allele/strand combination, various features are added by shifting the position of the SNP within the probe. The position shift ranges only from -4 to 4 bases, therefore within each strands the probes are relatively similar.

Most copy number algorithms can be divided into three main steps which we refer to as 1) the *preprocessing* step, 2) the copy number *estimation* step, and 3) the *smoothing* across the chromosome step. In the preprocessing step we summarize feature intensities into two quantities, representative of allele *A* and *B*. We refer to this step as *preprocessing*. In this paper, we use the following notation to denote the preprocessed data: $\theta_{A,i,j}$ and $\theta_{B,i,j}$ are the logarithms (base 2) of quantities proportional to the amount of DNA in target sample j associated with alleles *A* and *B* for SNP i . In the estimation step, we use these θ s

⁷ There are also mismatch probes *MM* which we completely ignore because the manufacture has plans of no longer using them.

Fig. 2. CARAT regression prediction on a SNP from Chromosome X. The dashed grey lines indicate allele specific copy number = 0,1,2 (from low to high). Figure A) and B) shows the allele specific prediction of copy number (log base 2) as compared to the real copy number (log base 2). We can see that even though the middle of each boxplot lies closely to the intersection of the dashed lines and the diagonal line, the ranges of boxplots overlap. Figure C) shows a scatterplot of preprocessed log intensities for allele A and B.

to estimate the true copy number, which we denote with $C_{i,j}$. The allele specific copy number are denoted with $C_{A,i,j}$ and $C_{B,j,i}$. Notice that the total copy number is the sum of the allele specific copy numbers, i.e. $C_{i,j} = C_{A,i,j} + C_{B,i,j}$. As we demonstrate here (see Figure 1), estimates of θ_A and θ_B are, in general, not precise enough to provide useful copy number calls. Therefore, most copy number estimation algorithms include the smoothing step in which copy number estimates from neighboring regions are averaged to improve the signal to noise ratio. These techniques range from the simple methods such as running median to more complicated ones such as hidden Markov models (HMM). In Section 3 we review some of the existing methods and motivate our mixture model approach. In Section 4 we describe the mixture model approach. In Sections 5 we present results and discussion respectively. Throughout this extended abstract we use data obtained from collaborators and public repositories which we briefly describe in Section 2.

2 Control Data

In Section 4 we describe a model that is trained using a reference set of 314 normal samples hybridized to Affymetrix's 100K array. We screened out samples not achieving the quality standard described by Carvalho et al [2]. Our reference set consists of 86 Hapmap samples, 124 samples from the Coriell Repositories (42 African American, 20 Asians, 40 Caucasians and 22 samples from the polymorphisms discovery panel) [3], and 104 from Chakravarti's lab. The test data was sampled from 20 Trisomy samples from Pevsner's lab.

3 Previous Work and Motivation

The first algorithms we describe do not provide allele specific results. We therefore define the total copy number quantity $S_{i,j} = \log_2(2^{\theta_{A,i,j}} + 2^{\theta_{B,i,j}})$. Ideally the $S_{i,j}$ are proportional to the true log-scale copy number. Figure 1A shows data for a male with Down syndrome. The $S_{i,j}$ are highly noisy and differences between chromosome 21 and X are hard to detect unless we smooth along the chromosome (in Figure 1 we show the results of running median). The lessons learned from expression arrays help understand this problem. Various authors have proposed an additive background/multiplicative model for gene expression microarray [22, 10, 27]. Furthermore, for Affymetrix arrays, various authors [12, 17] have clearly shown the existence of a strong multiplicative probe-specific effect. Probe-specific background noise, attributed to non-specific binding, have also been described [27]. Others [21, 16, 2] demonstrate that similar sized effects are seen with SNP chips. These effects are strong enough to be clearly seen even after averaging the various feature intensities associated with each SNP. Extending these findings to the copy number case results in the following model:

$$\theta_{a,i,j} = \log(\beta_{a,i} + \phi_{a,i} c_a \varepsilon_{a,i,j}) \quad (1)$$

with $a = A, B$ denoting allele, i identifies the SNP, j identifies the sample, β represents a positive valued SNP-specific background level, ϕ represents a SNP-specific probe-effect, and ε is multiplicative measurement error. Assuming this model holds, relatively simple calculations demonstrate that large values of $\beta_{a,i}$ result in attenuation of real differences and that large variability of $\phi_{a,i}$ across SNPs explains the large variance seen in Figure 1A.

Affymetrix's Copy Number Analysis Tool (CNAT) [1, 9] deals with the probe-effect using a simple yet effective technique. CNAT does not provide allele specific results and concentrates on estimating the overall copy number. For the preprocessing step, all feature intensities related to the SNP are therefore averaged to form $S_{i,j}$. Using dozens of control subjects CNAT defines a SNP specific average $\bar{S}_{ig}$, standard deviation $\hat{\sigma}_{ig}$, for each genotype $g=AA,AB,BB$. Values $S_{i,j'}$, from any new sample j' that are called genotype g are standardized in the usual manner: $(S_{i,j'} - \bar{S}_{ig})/\hat{\sigma}_{ig}$. A predefined regression equation is then used to transform these standardized values to the copy number scale. The standardized S values are used to obtain p-values from the null hypothesis the $S = 0$ ($C = 2$). Figure 1B shows de-meaned values for the same data shown in Figure 1A. The improvement is clear and it is due to the fact that ϕ is partially removed from the de-mean-ed values. However, notice that the signal to noise ratio still appears to be large: the separation between chromosomes with known differences is far from perfect. To avoid false positives, the third step in CNAT involves looking for strings of consecutive p-values that are smaller than some predefined cut-off. Other smoothing approaches have been used. For example, Zhao et al. [28] proposes the use Hidden Markov models (HMM) to define the procedure implemented by dChip.

Other authors have noted that further improvements can be obtained by reducing the variance at the preprocessing step. For example, several groups [19,

Fig. 3. Conditional joint distributions of $[\theta_{i,j} | \mathbf{C}_{i,j} = \mathbf{c}]$ for a SNP on chromosome X and a SNP on chromosome 21. The X-axis is the log base 2 intensity for allele B. The Y-axis is the log base 2 intensity for allele A. The red dots are from 45 new samples of male normal and 20 Down syndrome patients, respectively. The ellipses show the 95% critical region around the centers. The brown ellipses represent $C = 3$, $\mathbf{C}_{i,j} = (0, 3), (1, 2), (2, 1), (3, 0)$. The orange ellipses represent $C = 2$, $\mathbf{C}_{i,j} = (0, 2), (1, 1), (2, 0)$. The tan ellipses represent $C = 1$, $\mathbf{C}_{i,j} = (0, 1), (1, 0)$. The yellow ellipses represent $C = 0$, $\mathbf{C}_{i,j} = (0, 0)$.

13, 15] have used probe-sequence information, mainly GC content, and fragment length to predict and remove some of the probe-effect related variability. However, even after accounting for such factors the signal to noise ratio remains too low to make single-point copy number calls, thus these authors propose their own versions of the smoothing step.

Huang et al. [8] noted that CNAT's mean removal approach does not fully remove the probe effect because it does not properly deal with the additive background effect β . They propose the Copy Number Analysis with Regression And Tree (CARAT) algorithm which uses a non-linear regression model, based on model (1), to account for the probe-specific effects. To estimate model parameters they use a control dataset composed of dozens of arrays. First, genotype calls are obtained and treated as known. This permits estimation of allele specific parameter estimates. For example, for allele A we have known values of $c_A = 0, c_B = 2$ (BB genotype), $c_A = 1, c_B = 1$ (AB genotype), and $(c_A = 2, c_B = 0)$ (AA genotype) and thus we can estimate $\beta_{a,i}$ and $\phi_{a,i}$ with, for example, least squares, for each SNP i and each allele $a = A, B$. For a new sample, we can predict c_A, c_B using the fitted parameters and equation (1). Calls can then be based on cut-offs for the prediction of $\hat{c}_A + \hat{c}_B$. Huang et al. [8] suggest using $[0, 1.5]$, $[1.5, 2.5]$, $(2.5, \infty)$ for total copy number $< 2, = 2, > 2$ respectively. Figure 2A shows the data used in the regression for allele A from a randomly chosen SNP. The figure demonstrates that the model works reasonably well but that the signal to noise ratio is not large enough to provide perfect accuracy (the boxplots overlap). CARAT utilized a regression tree approach in the smoothing step.

The Probe-level allele-specific quantitation (PLASQ) [16] procedure is similar to CARAT. Two major difference is that PLASQ fits model (1) to the feature-level data and that PLASQ does not rely on external genotype calls. Although, in our opinion, PLASQ provides a superior model-based framework than any other approach, it is computationally challenging to implement. This is because a non-linear estimation procedure is performed at the feature-level for every SNP. Furthermore, it is difficult to adapt it to be robust to outliers and to take probe-sequence and fragment size into account.

Model based approaches such as CARAT and PLASQ provide a great advantage over previous ones: reliable confidence intervals can be computed for single-point copy number estimates. Huang et al. [8] point out that their uncertainty assessment permits one to call a relatively large group of SNPs and keep the false positive rate relatively low. We now briefly describe a simple adaptation of these methods that provides further improvements.

Notice that all of the above described algorithms use regression-type approaches to give a continuous prediction of copy number. The current approaches rely on three assumptions that we believe are not exactly true. The first is that the linear relationship predicted by model 1. Figure 2A and 2B show that there are small but significant deviations from these models. Other SNPs (not shown) show slightly larger deviations. The second is that θ_A and θ_B are independent. This assumption is clearly not true as demonstrated by Figures 2C and 3. The third assumption is that the variance of the measurement error term does not

depend on allele-specific copy number values. Figure 3 also shows this is not the case. In general, we are making convenience assumptions regarding the conditional probabilities of $(\theta_A, \theta_B)'$ given allele-specific copy number that might be hurting our bottom-line results.

Theoretically, one can show that the best predictor of discrete classes given continuous covariates is Bayes classifier. Bayes classifier is a function of the conditional distribution of the predictors given the classes which we can not always estimate effectively. In the next section we describe how we can use the large amount of public data and equation (1) to obtain useful estimates of these conditional distributions and therefore improved copy number calls.

Another problem with the above mentioned approaches is that they do not take into account the strong lab/study effects that are seen in SNP chip data [2]. Huang et al. [8] notice this problem and propose an ad-hoc solution that relies on having genotype calls for many samples from each study. This approach is clearly not usable in cases where we wish to obtain copy number calls for a small batch of new samples. Recently, Carvalho et al. [2] proposed a preprocessing algorithm, SNP-RMA, that remove much of the lab-effect. We use SNP-RMA preprocessed values throughout this extended abstract. These authors also describe the CRLMM procedure which we use to obtain genotype calls for the control dataset.

4 Allele-specific Mixture Model

Figure 2C shows a scatterplot of θ_A vs. θ_B values from hundreds of control arrays for a particular SNP. The three genotypes are clearly seen. Furthermore each of the three “clouds” looks bivariate normal with the AB genotype showing signs of correlation. Most SNPs show very similar plots and motivate the following model:

$$[\theta_{i,j} | \mathbf{C}_{i,j} = \mathbf{c}] = \begin{pmatrix} \gamma_{A,c_A,i} \\ \gamma_{B,c_B,i} \end{pmatrix} + \begin{pmatrix} \epsilon_{A,c_A,i,j} \\ \epsilon_{A,c_B,i,j} \end{pmatrix} \quad (2)$$

with $\theta_{i,j} = (\theta_A, \theta_B)'$, $\mathbf{C}_{i,j} = (C_{A,i,j}, C_{B,i,j})'$ representing the un-observed true copy number of alleles A and B for SNP i on sample j , $\mathbf{c} = (c_A, c_B)'$ are the possible values $\mathbf{C}_{i,j}$ can take, $(\gamma_{A,c_A,i}, \gamma_{B,c_B,i})'$ accounts for the shifts in location caused by the probe-effect, and $(\epsilon_{A,c_A,i,j}, \epsilon_{B,c_B,i,j})'$ is a bivariate normal error with mean 0 and copy-number-specific covariance matrix $\Sigma_{c,i}$ which is defined by the allele-copy-number-specific standard deviations $\sigma_{c_A,i}$, $\sigma_{c_B,i}$ and copy-number-pair-specific correlation $\rho_{c_A,c_B,i}$.

Remember that the sequence composition of the sense and antisense probes are quite different. Carvalho et al. [2] point out that, for a few hundred SNPs, one of the two strands does not appear to be sensitive to genotype changes. This difference is also observed with probes within the same strand. Huang et al. [8] used arbitrary cutoffs to select probes with strong allelic dosage response for further analysis. For this reason, we propose fitting the above model for probe-level data. Currently, the SNP-RMA preprocessing provides, by default, separate

values of θ_A and θ_B for sense and antisense. However, without loss of generality we describe the procedure as if there was only one probe.

For a large set of control data, described in Section 2, we obtain genotype calls, act as if these are known. This implies we know $\mathbf{C}$ for all these samples and we can estimate the γ s by simply using:

$$\hat{\gamma}_{A,c_A,i} = N_{c_A,i}^{-1} \sum_{\{j: C_{A,i,j}=c_A\}} \theta_{A,i,j}, \quad c_A = 0, 1, 2 \quad (3)$$

with $N_{c_A,i}$ the number of samples with genotypes implying $C_{A,i,j} = c_A$. The covariance matrix $\Sigma_{\mathbf{c},i,s}$ is computed in a similar way, namely using the sample covariance matrix of $\theta_{i,j}$ for samples j implying $\mathbf{C}_{i,j} = \mathbf{c}$.⁸ Because we assume the ϵ s are normal, these sets of parameter estimates define the conditional distributions for $\mathbf{C} = (2, 0), (1, 1)$ and $(0, 2)$. Notice that model (1) is not used to form any of these conditional distributions. Next we assume that the behavior of the θ s for $C_A = c_A$ is similar for all values of C_B and vice-versa. We then infer the conditional means for $\mathbf{C} = (0, 0), (0, 1), (1, 0)$. For example the conditional mean for SNP i when $\mathbf{C} = (0, 1)$ will be $(\gamma_{A,0,i}, \gamma_{B,1,i})$. The covariance matrix is inferred in a similar way (described below). To predict $(\gamma_{A,c_A,i}, \gamma_{B,c_B,i})'$ for $\mathbf{C} = (3, 0), (2, 1), (1, 2), (0, 3), (4, 0), (3, 1), \dots$, we use model 1. For example, we first use the estimates of $\gamma_{A,0,i}, \gamma_{A,1,i}, \gamma_{A,2,i}$ as outcomes in model (1) for values of $C_A = 0, 1, 2$ respectively, fit the model, and obtain estimates of $\beta_{A,i}$ and $\phi_{A,i}$, which permit us to predict $\gamma_{A,3,i}$, for $C_A = 3$.

We now describe how we infer $\Sigma_{\mathbf{c},i}$ for cases other than $\mathbf{C} = (2, 0), (1, 1), (0, 2)$ using the estimates we already have. For the A and B variance components (the diagonal entry) of the covariance matrix, we simply assume they depend only on c_A and c_B respectively. For $c_A > 2$ and $c_B > 2$ we assume the same variance as $c_A = 2$ and $c_B = 2$ respectively. We therefore use the estimates of the six parameters: $\sigma_{A,c_A,i}, c_A = 0, 1, 2, \sigma_{B,c_B,i}, c_B = 0, 1, 2$ and do not need to predict any new values. The correlation component is a bit more difficult. We assume that the correlation coefficient when $C_A > 0$ and $C_B > 0$ is the same as $\mathbf{C} = (1, 1)$. The rationale for this is that correlations are due to PCR effects being different from sample to sample. Thus if both allele fragments are present, the resulting quantities will be similar regardless of the starting quantities. When one of the two alleles is not present (PCR no longer makes it grow) we assume that the correlation for case where $C_A > 0$ but $C_B = 0$ is the same as $\mathbf{C} = (2, 0)$ and $C_A = 0$ but $C_B > 0$ the same as $\mathbf{C} = (0, 2)$. For $\mathbf{C} = (0, 0)$ we simply assume independence. With this assumption in place we can produce conditional expectations for any value of $\mathbf{C}$ given the observed θ s, described as follows.

With the model parameter estimates in place we are able to provide posterior probabilities for allele specific copy number. Furthermore, we can compute these posterior probabilities for total copy number:

$$[C_{A,i,j} + C_{B,i,j} = c | \theta_{i,j}]$$

⁸ In this extended abstract we actually use robust (to outliers) versions of these sample means and covariances

$$\propto \sum_{\{\mathbf{c}: c_A + c_B = c\}} [\boldsymbol{\theta}_{i,j} | \mathbf{C}_{i,j} = \mathbf{c}] \times [\mathbf{C}_{i,j} = \mathbf{c}]$$

where $[\boldsymbol{\theta}_{i,j} | \mathbf{C}_{i,j} = \mathbf{c}]$ is the bivariate normal distribution defined by model (3). The marginal probability of the $\mathbf{C}$ pair can be pre-specified and used to control specificity and sensitivity for any copy number value. We can obtain meaningful values by decomposing the probability into: $\Pr(C_{A,i,j} = c_A, C_{B,i,j} = c_B) = \Pr(C_{A,i,j} = c_A, C_{B,i,j} = c_B | C_{A,i,j} + C_{B,i,j} = c) \Pr(C_{A,i,j} + C_{B,i,j} = c)$. The first component relates to the proportion of each genotype in the population and can be computed using the Hardy-Weinberg Equilibrium for diploids ($C_{A,i,j} + C_{B,i,j} = 2$). The second component relates to the probability of each alteration $c = 0, 1, 3, 4, \dots$ which is unknown. We recommend the user define these probabilities to control specificity and sensitivity. For the examples shown in this extended abstract we assigned equal probabilities to $C_{A,i,j} + C_{B,i,j} = 0, 1, \dots, 6$.⁹ Once we have calculated the probabilities above we can provide estimates of copy number by, for example, computing the expected value of $C = C_A + C_B$.

A summary of the algorithm:

1. For each array, we obtain the pre-processed probe-level log intensities from snpRMA, the pre-processing algorithm used by CRLMM. These resulting measurements are $\theta_{A,+}, \theta_{A,-}, \theta_{B,+}, \theta_{B,-}$ for each SNP.
2. We estimate the conditional probability of these measurements, given allele specific copy number. We assume a bivariate normal for the A and B alleles at each copy number pair. This reduces the number of parameters greatly and we can estimate them precisely using a large training set. We use genotype calls to treat the allele specific copy number as known. We do this independently for sense (+) and antisense (-). More specifically:
3. We assume the prior probability for the joint distribution of C_A and C_B is a uniform distribution.
4. For a new dataset, we use the above estimates to calculate the posterior probability for C_A and C_B being $0, 1, 2, \dots, K$ (K is the maximum copy number permitted). We average the sense and antisense results.
5. Finally, we compute the posterior probability of $C_A + C_B$ being $0, 1, 2, \dots, K$.

5 Results

We now describe some of the applications of the hierarchical model described above. In general we refer to our procedure as the Copy Number-Robust Linear Model and Mixture Model (CN-RLMM) procedure.

Figure 3 gives the SNP-specific bivariate normal distribution of $\boldsymbol{\theta}$ for $\mathbf{C} = (0, 0), (0, 1), (1, 0), (0, 2), (1, 1), (2, 0), (0, 3), (1, 2), (2, 1), (3, 0)$, depicted in ellipses (95% confidence regions). These are estimated from the control data as described

⁹ Remember that we perform the above calculation separately for the sense and antisense values. A final estimate of the posterior probability simply average these two values.

in Section 2. Figure 3A and 3B give the sense and antisense-specific distributions for a SNP on chromosome X. We observe the extrapolated distribution of θ s given $\mathbf{C} = (0, 1), (1, 0)$ coincide with the observed θ s from 45 male samples that were not used in training. Similarly, figure 3C and 3D give the sense and antisense-specific distributions for a SNP on chromosome 21 and we observe that our extrapolated distributions coincide with the observed θ s from 20 Trisomy samples. This demonstrates that our assumptions seem to provide reasonable estimates of the conditional distribution of copy number the cases predicted with mode (1) ($\mathbf{C} = (0, 0), (0, 1), (1, 0), (0, 3), (1, 2), (2, 1), (3, 0)$).

Fig. 4. CN-RLMM results. CN is abbreviation for CN-RLMM, c.CNAT is CRLMM preprocessed probes + CNAT. Figure A) shows the expected copy number given preprocessed log intensities for 817 SNPs on Chromosome 21 of 20 Down syndrome patients (with identified Trisomy 21). Figure B) shows the expected copy number given preprocessed log intensities for 807 SNPs on Chromosome X of 45 male trio samples. Figure C) shows the average true positive rate versus 1-average percentage of call rate for 2 Down syndrome patients. The points demonstrate some of the corresponding posterior probability values used as cut-offs.

In Figure 4A and 4B we demonstrate how our results have much better precision than CNAT and values with and without probe- sequence and fragment length corrections. We achieve this precision without any loss of accuracy. Note that we could not get PLASQ to work with our data and no software is available to implement CARAT. The preprocessing used by dChip is very similar to CNAT and thus we expect results to be the same. Keep in mind the smoothing step is not being assessed. We observe that the degree of improvement is not equivalent for copy numbers 3 and copy number 1. This is expected because it is easier to detect a 2 times difference (copy number 2 versus 1) than to detect a 1.5 times difference (copy number 3 versus 2).

The most useful application of our results is that we provide improved single-point copy number estimates with reliable uncertainty assessment without the need to re-calibrate for new samples. Note that we can easily control our false positive rate by simply restricting calls to SNPs with posterior probabilities close

to 1. Figure 4C demonstrates that we can get usable single-point copy number estimates for a large amount of SNPs. Notice that the worst performance is observed for $CN=3$. This is likely due to the fact that we used model (1) to extrapolate (as done by CARAT).

6 Discussion

We have presented a mixture model approach that permits us to obtain improved copy number estimate as well as reliable single-point copy number calls. A major advantage of our methodology over the best existing one, e.g. CARAT and PLASQ, is that we explicitly model the conditional joint distribution of the intensities given the copy number values. This permits us to model the strong correlation that sometimes exists between A and B and exploit this information to improve bottom-line results. This advantage is best exemplified by Figure 3C where the $C = (2, 1), (1, 1)$, and $(1, 2)$ are usefully separable only if we take this correlation into account. Furthermore, avoiding the linearity assumption made by these other procedures seems to help as well. This is best demonstrated by the fact that we perform worst in cases where we rely on this assumption, i.e. making calls for $CN = 3$. Finally, because we use training data to fit the mixture models, the procedure is entirely linear. Other procedures, such as CARAT and PLASQ rely on non-linear algorithms that present many practical problems.

We have plans to extend and improve our approach in various ways. First, we plan to implement it for the 500K chips. Second, we believe this approach can be used with Illumina's SNP array and thus plan to try it with data from this platform. Third, we plan to add another level to the model that will permit us to borrow strength across the thousands of SNPs to better estimate the parameters of the conditional probabilities. We plan to use an approach similar to that of CRLMM. Fourth, we plan to look for ways to avoid using the linearity assumption to infer the parameter of conditional distributions when $C > 2$. We plan to use general regression approaches that predict these parameters from the known parameters $C \leq 2$. We can train this regression model with Trisomy data ($C = 3$) and design experiments to be able to train for $C > 3$. Fifth, we plan to look for better ways of combining the results from sense and antisense probes. It is desirable to detect and ignore misbehaving strands. Finally, we have observed correlation between parameter estimates coming from proximal locations on the chromosome. This could be due to the fact that various SNPs are on each of the fragments that are amplified. We will explore ways to exploit this finding.

It is possible that the reference set we use has an influence on our results. We plan to study this problem in more detail in the near future. We also plan to substantially increase the size of the reference set to reduce the effect of outlier samples. By combining various publicly available assessment experiments, we plan to develop a comparison protocol for analysis methods. This will help us determine not only which methods work better, but to explore if subsets of the reference set provide better results.

Notice that we did not offer any solutions for the smoothing step as we are more interested in developing techniques for single-point estimates. We expect some of the existing techniques to work well when applied to our estimates of copy number. However, because we explicitly model the conditional probabilities it is possible to develop new methods that impose the across-chromosome correlation through those probabilities instead of the actual copy number estimates.

References

1. G. R. Bignell, J. Huang, J. Greshock, S. Watt, A. Butler, S. West, M. Grigорова, K. W. Jones, W. Wei, M. R. Stratton, P. A. Futreal, B. Weber, M. H. Shapero, and R. Wooster. High-resolution analysis of DNA copy number using oligonucleotide microarrays. *Genome Res*, 14(2):287–95, Feb 2004.
2. B. Carvalho, T. P. Speed, and R. A. Irizarry. Exploration, normalization, and genotype calls of high density oligonucleotide snp array data. *Johns Hopkins University, Dept. of Biostatistics Working Papers*, (111), 2006.
3. F. S. Collins, L. D. Brooks, and A. Chakravarti. A DNA polymorphism discovery resource for research on human genetic variation. *Genome Res*, 8(12):1229–31, Dec 1998.
4. D. F. Conrad, T. D. Andrews, N. P. Carter, M. E. Hurles, and J. K. Pritchard. A high-resolution survey of deletion polymorphism in the human genome. *Nat Genet*, 38(1):75–81, Jan 2006.
5. L. Feuk, A. R. Carson, and S. W. Scherer. Structural variation in the human genome. *Nat Rev Genet*, 7(2):85–97, Feb 2006.
6. S. M. Gribble, D. Kalaitzopoulos, D. C. Burford, E. Prigmore, R. R. Selzer, B. L. Ng, N. S. W. Matthews, K. M. Porter, R. Curley, S. J. Lindasy, J. Baptista, T. A. Richmond, and N. P. Carter. Ultra-high resolution array painting facilitates breakpoint sequencing. *J Med Genet*, Sep 2006.
7. D. A. Hinds, A. P. Klok, M. Jen, X. Chen, and K. A. Frazer. Common deletions and SNPs are in linkage disequilibrium in the human genome. *Nat Genet*, 38(1):82–5, Jan 2006.
8. J. Huang, W. Wei, J. Chen, J. Zhang, G. Liu, X. Di, R. Mei, S. Ishikawa, H. Aburatani, K. W. Jones, and M. H. Shapero. CARAT: a novel method for allelic detection of DNA copy number changes using high density oligonucleotide arrays. *BMC Bioinformatics*, 7:83, 2006.
9. J. Huang, W. Wei, J. Zhang, G. Liu, G. R. Bignell, M. R. Stratton, P. A. Futreal, R. Wooster, K. W. Jones, and M. H. Shapero. Whole genome DNA copy number changes identified by high density oligonucleotide arrays. *Hum Genomics*, 1(4):287–99, May 2004.
10. W. Huber, A. von Heydebreck, H. Suetmann, A. Poutska, and M. Vingron. Variance stabilization applied to microarray data calibration and to the quantification of differential expression. *Bioinformatics*, 1, 2002.
11. A. Iafrate, L. Feuk, M. Rivera, M. Listewnik, P. Donahoe, Y. Qi, S. Scherer, and C. Lee. Detection of large-scale variation in the human genome. *Nature Genetics*, 36(9):949–951, 2004.
12. R. Irizarry, F. C. B. Hobbs, Y. Beazer-Barclay, K. Antonellis, U. Scherf, and T. Speed. Exploration, normalization, and summaries of high density oligonucleotide array probe level data. *Biostatistics*, 4:249–264, 2003.

13. S. Ishikawa, D. Komura, S. Tsuji, K. Nishimura, S. Yamamoto, B. Panda, J. Huang, M. Fukayama, K. W. Jones, and H. Aburatani. Allelic dosage analysis with genotyping microarrays. *Biochem Biophys Res Commun*, 333(4):1309–14, Aug 2005.
14. G. C. Kennedy, H. Matsuzaki, S. Dong, W. min Liu, J. Huang, G. Liu, X. Su, M. Cao, W. Chen, J. Zhang, W. Liu, G. Yang, X. Di, T. Ryder, Z. He, U. Surti, M. S. Phillips, M. T. Boyce-Jacino, S. P. Fodor, and K. W. Jones. Large-scale genotyping of complex DNA. *Nature Biotechnology*, 21:1233–1237, 2003.
15. D. Komura, K. Nishimura, S. Ishikawa, B. Panda, J. Huang, H. Nakamura, S. Ihara, M. Hirose, K. W. Jones, and H. Aburatani. Noise Reduction from genotyping microarrays using probe level information. *In Silico Biol*, 6(1-2):0009, Feb 2006.
16. T. Laframboise, D. Harrington, and B. A. Weir. PLASQ: A Generalized Linear Model-Based Procedure to Determine Allelic Dosage in Cancer Cells from SNP Array Data. *Biostatistics*, Jun 2006.
17. C. Li and W. Wong. Model-based analysis of oligonucleotide arrays: Expression index computation and outlier detection. *Proceedings of the National Academy of Science U S A*, 98:31–36, 2001.
18. S. A. McCarroll, T. N. Hadnott, G. H. Perry, P. C. Sabeti, M. C. Zody, J. C. Barrett, S. Dallaire, S. B. Gabriel, C. Lee, M. J. Daly, D. M. Altshuler, and I. H. Consortium. Common deletion polymorphisms in the human genome. *Nat Genet*, 38(1):86–92, Jan 2006.
19. Y. Nannya, M. Sanada, K. Nakazaki, N. Hosoya, L. Wang, A. Hangaishi, M. Kurokawa, S. Chiba, D. K. Bailey, G. C. Kennedy, and S. Ogawa. A robust algorithm for copy number detection using high-density oligonucleotide single nucleotide polymorphism genotyping arrays. *Cancer Res*, 65(14):6071–9, Jul 2005.
20. D. A. Peiffer, J. M. Le, F. J. Steemers, W. Chang, T. Jenniges, F. Garcia, K. Haden, J. Li, C. A. Shaw, J. Belmont, S. W. Cheung, R. M. Shen, D. L. Barker, and K. L. Gunderson. High-resolution genomic profiling of chromosomal aberrations using Infinium whole-genome genotyping. *Genome Res*, 16(9):1136–48, Sep 2006.
21. N. Rabbee and T. P. Speed. A genotype calling algorithm for affymetrix SNP arrays. *Bioinformatics*, 22(1):7–12, Jan 2006.
22. D. M. Rocke and B. Durbin. A model for measurement error for gene expression arrays. *J Comput Biol*, 8(6):557–569, 2001.
23. J. Sebat, B. Lakshmi, J. Troge, J. Alexander, J. Young, P. Lundin, S. Maner, H. Massa, M. Walker, M. Chi, N. Navin, R. Lucito, J. Healy, J. Hicks, K. Ye, A. Reiner, T. Gilliam, B. Trask, N. Patterson, A. Zetterberg, and M. Wigler. Large-scale copy number polymorphism in the human genome. *Science*, 305:525–528, 2004.
24. A. J. Sharp, S. Hansen, R. R. Selzer, Z. Cheng, R. Regan, J. A. Hurst, H. Stewart, S. M. Price, E. Blair, R. C. Hennekam, C. A. Fitzpatrick, R. Segraves, T. A. Richmond, C. Guiver, D. G. Albertson, D. Pinkel, P. S. Eis, S. Schwartz, S. J. L. Knight, and E. E. Eichler. Discovery of previously unidentified genomic disorders from the duplication architecture of the human genome. *Nat Genet*, 38(9):1038–42, Sep 2006.
25. A. J. Sharp, D. P. Locke, S. D. McGrath, Z. Cheng, J. A. Bailey, R. U. Vallente, L. M. Pertz, R. A. Clark, S. Schwartz, R. Segraves, V. V. Oseroff, D. G. Albertson, D. Pinkel, and E. E. Eichler. Segmental duplications and copy-number variation in the human genome. *Am J Hum Genet*, 77(1):78–88, Jul 2005.
26. E. Tuzun, A. J. Sharp, J. A. Bailey, R. Kaul, V. A. Morrison, L. M. Pertz, E. Haugen, H. Hayden, D. Albertson, D. Pinkel, M. V. Olson, and E. E. Eichler. Fine-scale structural variation of the human genome. *Nat Genet*, 37(7):727–32, Jul 2005.

27. Z. Wu, R. Irizarry, R. Gentleman, F. Martinez-Murillo, and F. Spencer. A model based background adjustment for oligonucleotide expression arrays. *Journal of the American Statistical Association*, 2004.
28. X. Zhao, C. Li, J. G. Paez, K. Chin, P. A. Jeanne, T.-H. Chen, L. Girard, J. Minna, D. Christiani, C. Leo, J. W. Gray, W. R. Sellers, and M. Meyerson. An integrated view of copy number and allelic alterations in the cancer genome using single nucleotide polymorphism arrays. *Cancer Res*, 64(9):3060–71, May 2004.