

Detection and accurate false discovery rate control of differentially methylated regions from whole genome bisulfite sequencing

KEEGAN KORTHAUER

Department of Biostatistics & Computational Biology, Dana-Farber Cancer Institute, 450 Brookline Ave, Boston, MA 02215, USA and Department of Biostatistics, Harvard T.H. Chan School of Public Health, 677 Huntington Ave, Boston, MA 02115, USA

SUTIRTHA CHAKRABORTY

Novartis, Inorbit Mall Rd, Silpa Gram Craft Village, HITEC City, Hyderabad, Telangana 500081, India

YUVAL BENJAMINI

The Statistics Department, Hebrew University, Mount Scopus, Jerusalem 91905, Israel

RAFAEL A. IRIZARRY*

Department of Biostatistics & Computational Biology, Dana-Farber Cancer Institute, 450 Brookline Ave, Boston, MA 02215, USA and Department of Biostatistics, Harvard T.H. Chan School of Public Health, 677 Huntington Ave, Boston, MA 02115, USA

rafa@jimmy.harvard.edu

SUMMARY

With recent advances in sequencing technology, it is now feasible to measure DNA methylation at tens of millions of sites across the entire genome. In most applications, biologists are interested in detecting differentially methylated regions, composed of multiple sites with differing methylation levels among populations. However, current computational approaches for detecting such regions do not provide accurate statistical inference. A major challenge in reporting uncertainty is that a genome-wide scan is involved in detecting these regions, which needs to be accounted for. A further challenge is that sample sizes are limited due to the costs associated with the technology. We have developed a new approach that overcomes these challenges and assesses uncertainty for differentially methylated regions in a rigorous manner. Region-level statistics are obtained by fitting a generalized least squares regression model with a nested autoregressive correlated error structure for the effect of interest on transformed methylation proportions. We develop an inferential approach, based on a pooled null distribution, that can be implemented even when as few as two samples per population are available. Here, we demonstrate the advantages of our method using both experimental data and Monte Carlo simulation. We find that the new method improves the specificity and sensitivity of lists of regions and accurately controls the false discovery rate.

Keywords: Bisulfite sequencing; Differential methylation; False discovery rate; Generalized least squares.

*To whom correspondence should be addressed.

1. INTRODUCTION

DNA methylation is an important epigenetic modification that plays a role in a wide variety of biological processes. Numerous studies have been carried out to locate CpG loci where DNA methylation may be involved in gene regulation, differentiation, and cancer. With recent advances in sequencing technology such as whole genome bisulfite sequencing (WGBS), it is now possible to measure DNA methylation at single base resolution across all CpGs in the genome. Even though the most common application of the technology is to detect differentially methylated regions (DMRs) between populations, most methods for analysis of WGBS experiments focus on statistical differences for CpG loci one at a time (*Akalin and others, 2012; Dolzhenko and Smith, 2014; Lee and Morris, 2016; Park and others, 2014; Park and Wu, 2016*). While useful, approaches for identification of differentially methylated loci (DML) have many practical limitations in both implementation and interpretation. Here, we discuss these limitations as well as outline the challenges of performing inference at the region level. Finally, we introduce a rigorous statistical approach that overcomes these challenges to construct de novo DMRs with accurate false discovery rate (FDR) control.

Methods to identify DMLs in WGBS experiments are greatly hindered by the high-dimensionality and low sample size setting that is common in high-throughput genomics studies. The number of tests performed is equal to the number of loci analyzed, which is very large in typical WGBS studies. In the human genome, for example, there are close to 30 million CpG loci (*Smith and Meissner, 2013*). Further, DML methods generally do not account for the well-known fact that measurements are spatially correlated across the genome (*Leek and others, 2010*) and instead treat measurements from all loci as independent. Correcting for multiple comparisons without taking into account these correlations can result in a loss of power.

Additionally, methods for assessing the significance of DMLs typically require large sample sizes due to reliance on large sample approximations (*Dolzhenko and Smith, 2014; Hansen and others, 2012; Hebestreit and others, 2013; Lee and Morris, 2016*). Although WGBS is the current gold standard for estimating whole genome methylation profiles (*Marx, 2016*), cost limitations are still a barrier to acquiring more than a few individuals per biological condition (*Ziller and others, 2015*). This is reflected in the study design of major consortia that aim to characterize the epigenome. For example, WGBS experiments in murine embryos carried out as part of the ENCODE project are limited to two biological replicates per tissue type and developmental time point combination (*He and others, 2017*). In addition, the number of biological replicates measured with WGBS in the UCSD Human Reference Epigenome Mapping Project (*Schultz and others, 2015*) is also limited to 2–3 per tissue type. As such, we aim to maximize power while controlling the FDR even with sample sizes as small as two samples per condition.

Methods for identifying DMLs also need to properly account for the statistical properties of count data that do not conform to standard Gaussian models. This is in contrast to methylation array analysis, where Gaussian models performed well (*Jaffe and others, 2012*). One option is to assume that methylation proportions, defined as the number of methylated reads divided by the number of total reads covering a given CpG locus, follow a normal distribution (*Hansen and others, 2012*). However this assumption clearly does not hold when the total reads covering the CpG, referred to as the coverage, is small, a common occurrence in these data sets. The approach also ignores that variance of this proportion depends on the coverage. To overcome these limitations, DML approaches have also modeled WGBS count data using Binomial models (*Saito and others, 2014*). However, Binomial models on their own cannot account for biological variability within sample groups. In order to account for biological variability in count data, Beta-Binomial models (*Park and others, 2014; Sun and others, 2014*) are a natural extension. However, they come at the cost of increased computational burden when testing millions of loci.

Beyond implementation challenges, DML approaches also suffer from limited interpretability. In general, identifying DMRs is more biologically relevant than reporting DMLs. Apart from the so-called CpG

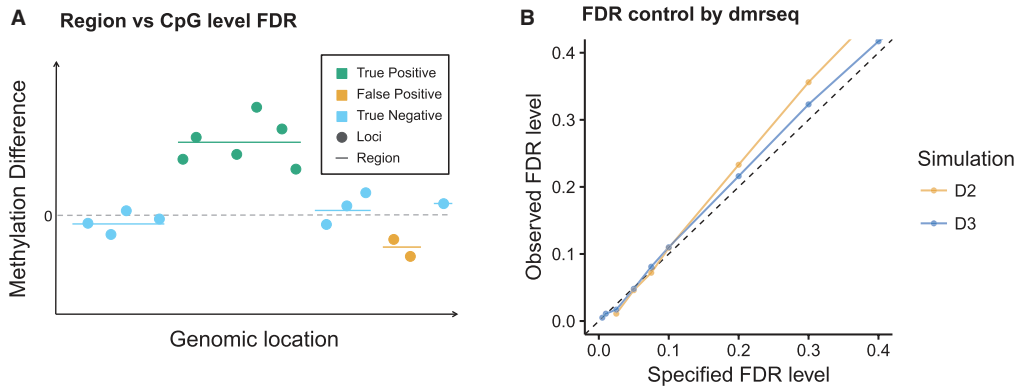


Fig. 1. False discovery rate (FDR) control at the region level. (A) Illustration of why FDR at the loci level is not the same as FDR at the region level. This schematic shows a plot of genomic location versus methylation difference estimates at several neighboring loci. The individual CpGs (points) are shaded by whether they are a true or false positive. Regions are denoted by lines. The loci FDR is $FDR_{\text{loci}} = (\# \text{False Positive Loci}) / (\text{Total } \# \text{ of Significant Loci})$, which is equal to 0.25 in this example. The region FDR is $FDR_{\text{region}} = (\# \text{False Positive Regions}) / (\text{Total } \# \text{ of Significant Regions})$, which is equal to 0.50 in this example. (B) dmrseq provides accurate FDR control of regions. Specified versus observed region-level FDR level is plotted for two different sample size settings from simulated data for dmrseq. Note that region-level FDR cannot be specified for BSmooth or DSS, and results for metilene are shown in Figure S3 of supplementary material available at *Biostatistics* online.

traffic lights (Khamis *and others*, 2017), most individual CpG loci likely do not have a large impact on epigenetic function on their own, but rather through a biochemical modification that involves several loci. Most notably, regional DNA methylation levels are correlated with the expression levels of nearby genes. Specifically, methylation gain is associated with stable transcriptional silencing of nearby genes (Bird, 2002). In the context of differential methylation analysis, Aryee *and others* (2014) found that differentially expressed genes were consistently more likely to be located near DMRs than DMLs.

While DML approaches may construct DMRs by chaining together neighboring significant loci, this type of approach will not yield a proper assessment of the statistical significance of the constructed regions, nor will the FDR be properly controlled (Robinson *and others*, 2014). This is because controlling the FDR at the level of individual loci is not the same as controlling FDR of regions, as has been noted in the context of peak calling in ChIP-seq experiments (Lun and Smyth, 2014; Siegmund *and others*, 2011). FDR correction at the level of individual loci means that the proportion of expected false positive (FP) loci is controlled, not the proportion of FP regions. Statistically, this is a critical point since FDR control of DMR detection is not guaranteed under the DML setting. In fact, many discoveries at the loci level may constitute only a single discovery. This means that a large number of correct rejections at the loci level can inflate the denominator in the FDR calculation, which will artificially lower the FDR of loci as compared to regions (Figure 1A). We were motivated to develop a procedure to control FDR at the region level and provide an accurate measure of statistical significance for each region.

Many recent computational approaches have been developed with the goal of identifying DMRs, but most do not provide formal inference for regions (Hansen *and others*, 2012; Saito *and others*, 2014; Wu *and others*, 2015; Yu and Sun, 2016) and instead join together significant DMLs. This type of procedure will suffer from the problems outlined above. Other approaches can perform inference at the region level, but only for predefined regions of interest or fixed sliding windows (Hebestreit *and others*, 2013; Sun *and others*, 2014; Mayo *and others*, 2015). Though useful in targeted settings such as reduced representation bisulfite sequencing (RRBS), or when we have prior knowledge of the DMR size, they are not applicable

to identifying DMRs of arbitrary size from WGBS. Those methods that scan the genome for DMRs and provide inference at the region level do not properly control FDR (Juhling *and others*, 2016; Wen *and others*, 2016). This is evidenced, for example, by the FDRs reported in the simulation studies of Wen *and others* (2016), which were as high as 0.85 and widely varied across scenarios. Juhling *and others* (2016) also do not achieve accurate FDR control in simulation studies (see Section 4.1).

The challenge of performing inference at the region level is complicated by several factors in addition to the challenges already discussed in the context of DML analysis. The first challenge is in defining the region boundaries themselves. Without prior knowledge or predefined regions, we need to construct data-driven regions. Calculating a test statistic for these data-driven regions of varying sizes with a known null distribution is not straightforward. In addition, challenges are presented by the complex statistical dependencies observed in measurements from nearby loci (Benjamini *and others*, 2016), as well as different within group variability across loci (Hansen *and others*, 2012). Some methods ignore correlation across loci (Wen *and others*, 2016) or biological variability from sample to sample (Saito *and others*, 2014; Wu *and others*, 2015). Not properly accounting for both of these sources of variability in DNA methylation data; however, results in misleading conclusions or loss of power. For a full review of DML and DMR methods, see Shafi *and others* (2018).

To overcome the limitations and challenges detailed above, we propose a two-stage approach that first detects candidate regions and then explicitly evaluates statistical significance at the region level while accounting for known sources of variability. Candidate DMRs are defined by segmenting the genome into groups of CpGs that show consistent evidence of differential methylation. Because the methylation levels of neighboring CpGs are highly correlated, we first smooth the signal to combat loss of power due to low coverage as done by Hansen *and others* (2012). In the second stage, we compute a statistic for each candidate DMR that takes into account variability between biological replicates and spatial correlation among neighboring loci. Significance of each region is assessed via a permutation procedure which uses a pooled null distribution that can be generated from as few as two biological replicates, and FDR is controlled using the procedure of Benjamini and Hochberg (1995). Code to reproduce the analyses presented in this article is provided in [supplementary material](#) available at *Biostatistics* online and the open-source R package `dmrseq` that implements the approach is available on GitHub.

In Section 2, we provide a detailed description of the data sets used. We describe the methodological details of the approach and detail the data processing and analysis procedure in Section 3. In Section 4, we present our findings using both experimental data and simulations. We demonstrate that the proposed approach assigns greater statistical significance to regions that have greater biological significance in terms of potential functional roles in the regulation of gene expression. We also evaluate sensitivity and specificity of the approach by analyzing null comparisons of samples from the same biological condition, with and without adding simulated DMRs. We demonstrate that `dmrseq` has higher sensitivity than existing approaches and accurately assesses statistical significance of regions through FDR estimation. A discussion of the advantages and limitations of the method are given in Section 5.

2. DATA DESCRIPTION

`dmrseq` is generally applicable to WGBS data which contains the counts for both methylated and unmethylated reads mapping to each CpG loci. This information can be obtained from raw sequencing reads using mapping software such as Bismark (Krueger and Andrews, 2011), as described in the [supplementary material](#) available at *Biostatistics* online. In this study, we include all CpG loci that are covered by at least one read in every sample. Other methods for analysis of WGBS data recommend only including CpG sites that have several reads in every sample, and while processed data of this form may be analyzed by our approach, it is important to note that this may result in a loss of power to detect regions in low-coverage

areas of the genome (see Section 4.3 of [supplementary material](#) available at *Biostatistics* online). In general, all CpG loci that are covered by at least one read in at least one sample per biological group can be analyzed by dmrseq.

In this study, we use our approach to identify DMRs using publicly available WGBS data from two different case studies, as described below. We also evaluate sensitivity and specificity of DMR methods by applying them to simulated data. Summary of coverage and methylation values for all data sets used can be found in Table S1 and Figure S2 of [supplementary material](#) available at *Biostatistics* online. For more details on data processing, see Section 1 of the [supplementary material](#) available at *Biostatistics* online.

2.1. Simulated data

Two sets of simulated data were constructed: one representing a null comparison (with no DMRs) and another containing simulated DMRs. To ensure that the simulated data sets closely match the characteristics of the observed experimental data, they were generated based on WGBS data from a study of human dendritic cells ([Pacis and others, 2015](#)). This study estimated methylation profiles of human dendritic cells from six donors before and after infection with a pathogen. The null comparison was constructed by randomly partitioning the six control samples (before infection) into two groups of three samples each, denoted Simulation N3. The same is done for a subset of four of the samples to evaluate performance when there are only two samples in each population, denoted Simulation N2.

Starting with the null comparisons, 3000 simulated DMRs were added to each data set in order to evaluate specificity and sensitivity. These are denoted Simulations D2 and D3 for two and three samples per population, respectively. Briefly, a DMR is constructed by sampling a cluster of neighboring CpGs and simulating the number of methylated reads, conditional on observed coverage, for the samples from one population from a binomial distribution. The binomial probabilities are equal to the observed methylation proportions plus or minus an effect size that is randomly sampled from a distribution that represents small to moderate effect sizes (ranging from approximately 0.1 to 0.5). In order to mimic the shape of DMRs detected in the case studies (Figure 2A), this difference is allowed to vary smoothly over the region according to a function similar to the tricube kernel ([Cleveland, 1979](#)) (see Section 2.4 of the [supplementary material](#) available at *Biostatistics* online).

2.2. UCSD Human Reference Epigenome Mapping Project

Data from several human tissue samples from the UCSD Human Reference Epigenome Mapping Project ([Schultz and others, 2015](#)) was used to identify DMRs related to tissue type. Specifically, four tissues were selected for performing pairwise comparisons: (1) heart, left ventricle, (2) heart, right ventricle, (3) sigmoid colon, and (4) small intestine.

2.3. Murine models of leukemia

In this study, marrow or thymus cells from two biological replicates of each of three different murine lines were extracted and genome-wide methylation levels measured with WGBS. One condition consisted of a wild-type control mouse. The other two had alterations in one or both of the DNMT3a or FLT3 loci, both of which have previously demonstrated implications in the development of leukemia ([Yang and others, 2016](#)). The mouse model with a wild-type DNMT3a locus and a duplication of the FLT3 locus has been shown to induce ALL. The mouse model with the same duplication of the FLT3 locus as well as a knock out of DNMT3a has been shown to induce the more lethal and aggressive AML. The DNMT3a also plays a role in promoting DNA methylation, so it is of interest to characterize the resulting differences in methylated regions among the control and two different leukemia models.

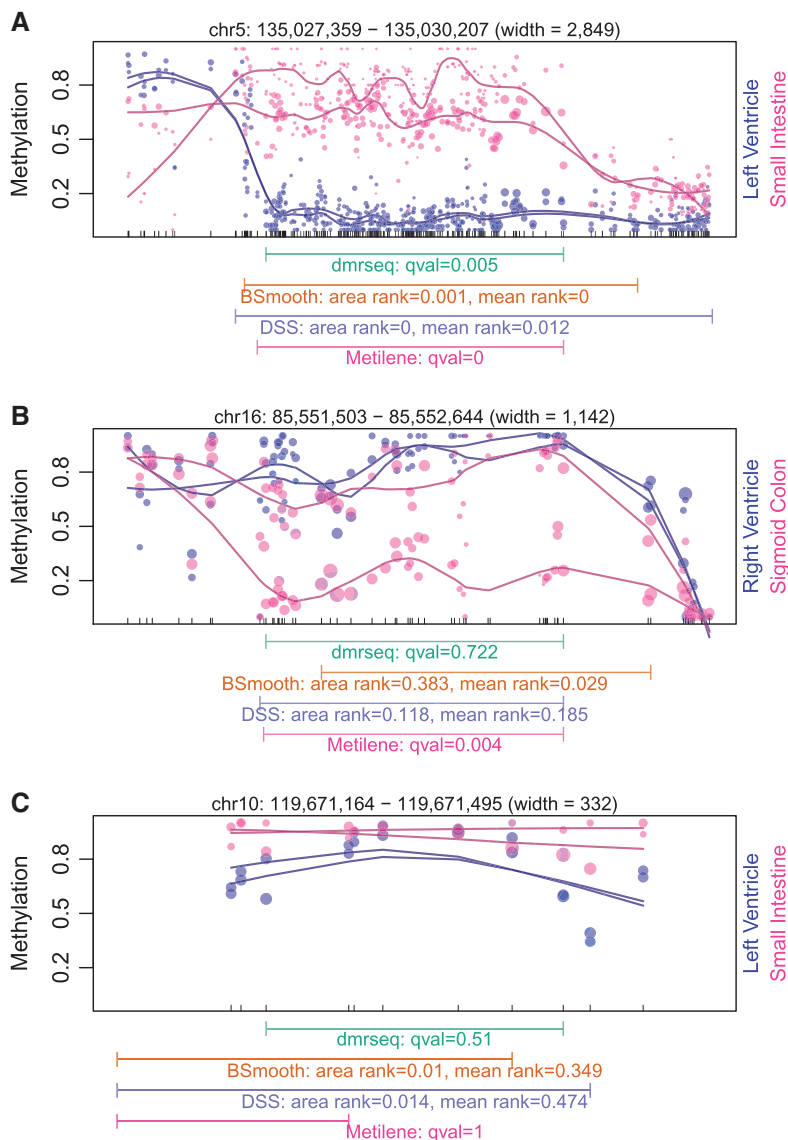


Fig. 2. dmrseq ranks regions by statistical significance. Example regions from the human tissue case studies are displayed for three cases that illustrate the increased variability of regions that are highly ranked by area or mean difference statistics of BSmooth and DSS but not dmrseq. For each case, the q-value is shown for dmrseq and metilene, and the rank percentile by the area statistic and mean difference statistics are both shown for BSmooth and DSS (see Section 2.8 of [supplementary material](#) available at *Biostatistics* online for details). (A) All methods assign a consistently high rank. (B) dmrseq assigns a low rank, but the mean difference statistic of BSmooth and DSS assign a high rank. (C) dmrseq assigns a low rank, but the area statistic of BSmooth and DSS assign a high rank. The condition comparison is indicated by the labels to the right of each plot.

3. ANALYSIS FRAMEWORK

A two-step procedure is carried out to (i) construct de novo candidate regions and (ii) quantify and evaluate the statistical significance of the effect of the covariate of interest on methylation level. Here, we detail each stage of the approach.

3.1. Construction of candidate regions

In Step 1, we detect candidate regions that contain multiple loci showing evidence of a difference in the smoothed pooled methylation proportion between biological conditions. For simplicity of presentation, we assume there are two biological conditions $s \in \{1, 2\}$ with sample indices $j \in C_s$ (see Section 2.7 of [supplementary material](#) available at *Biostatistics* online for the case of more than two conditions). Let M_{ij} be the number of methylated reads and U_{ij} the number of unmethylated reads for locus i of sample j from condition s . The coverage is denoted N_{ij} , where $N_{ij} = M_{ij} + U_{ij}$. The estimate of the mean methylation proportion $\hat{\pi}_{is}$ for loci i in condition s is taken to be the sum of methylated reads from all samples in that condition divided by the sum of all reads (i.e. the coverage) from all samples in condition s : $\hat{\pi}_{is} = \sum_{j \in C_s} M_{ij} / \sum_{j \in C_s} N_{ij}$. This leads to the following estimate of methylation proportion difference $\hat{\beta}_i$ between condition s and s' at loci i : $\hat{\beta}_i = \hat{\pi}_{is} - \hat{\pi}_{is'}$.

In order to give more weight to measurements with higher coverage, this estimate pools together samples within the same condition. To account for biological variability between samples and further reduce influence of observations with low coverage, smoothed individual loci estimates $\hat{\beta}_i^{\text{Smooth}}$ are obtained using a local-likelihood smoother ([Loader, 1999](#)) with smoothing weights w_i equal to the median coverage at loci i scaled by the average median absolute deviation (MAD) within the sample groups $\bar{\delta}_i$: $w_i = \text{med}_j(N_{ij}) / \bar{\delta}_i$, where $\bar{\delta}_i = \frac{1}{2} \sum_s \text{med}_{j \in C_s} \left| \frac{M_{ij}}{N_{ij}} - \text{med}_{k \in C_s} \left(\frac{M_{ik}}{N_{ik}} \right) \right|$ and $\text{med}_j(x_j)$ is the median of x_j over all j . This places more emphasis on observations with high coverage and low variability within sample group (see Section 2.1 of the [supplementary material](#) available at *Biostatistics* online for more details).

Candidate regions are defined by segmenting the genome into groups of loci with a smoothed pooled proportion difference $\hat{\beta}_i^{\text{Smooth}}$ in the same direction that is greater than some threshold in absolute value (refer to Section 2.2 of [supplementary material](#) available at *Biostatistics* online for more details). Maximum spacing between loci within a candidate region is controlled by a predetermined value, and loci at the start and end of the region with low difference values are trimmed (refer to Section 2.3 of [supplementary material](#) available at *Biostatistics* online for more details). The threshold can be thought of as the minimum difference in methylation proportion that is considered biologically significant without regard to FP, as significance of the candidate regions is assessed in the next step. We find that a value of 0.10 works well in practice and choose this for the default value, but this choice can be informed by the specific application at hand. For example, if interest lies primarily in detecting large magnitude differences, the threshold may need to be raised. Additionally, the threshold may need to be lowered when large magnitude differences are not expected.

3.2. Assessing significance of regions

In the second step, we assess the significance of candidate regions. This task is complicated by the fact that the null statistics are calculated on an enriched set of regions. In general, the null distribution generated by the type of selection procedure described in the previous section is not known. A natural approach would be to carry out a permutation test to control family-wise error rate (FWER), which is done by [Jaffe and others \(2012\)](#) to infer DMRs from array data. However, this is not feasible when we have only a few samples per population as is most often the case with WGBS. Thus, we set out to construct a statistic that can be comparable across the genome so that the signal can be compared among regions. Such an

exchangeable statistic allows us to generate an approximate null distribution by pooling genome-wide candidate regions detected from permutations.

To generate an approximately exchangeable region statistic that measures the strength of methylation difference, we need to account for sources of variation that are known to vary across the genome, including biological variability from sample to sample (Hansen and others, 2012), as well as covariance of nearby loci (Benjamini and others, 2016). Failing to do so may result in large test statistics just by chance for regions with high variability, leading to increased FDR or decreased power. For example, if we use an area-based statistic (Hansen and others, 2012) or a mean difference statistic averaged across loci, power to detect DMRs is greatly reduced in simulation studies (Figure S5 and Section 4.1 of [supplementary material](#) available at *Biostatistics* online).

Since we need to compute the statistic over potentially hundreds of thousands of candidate regions, we also favor an approach that provides efficient and stable estimation procedures. For these reasons, we make use of generalized least squares (GLS) regression model with a nested autoregressive correlated error structure for the effect of interest on transformed methylation proportions, the advantages of which are described in detail in the next subsections.

3.2.1. Estimation of region statistics with generalized least squares models To account for sampling variability, we assume that methylation counts for region r are Binomially distributed with probability p_{ijr} , where $M_{ijr}|N_{ijr}, p_{ijr} \sim \text{Binomial}(N_{ijr}, p_{ijr})$. To model biological variability, we allow the binomial proportion for samples in condition $s \in \{1, 2\}$ to vary according to a beta distribution with shape parameters α_{irs} and β_{irs} , where $p_{ijr} \sim \text{Beta}(\alpha_{irs}, \beta_{irs})$. Let $\pi_{irs} = \alpha_{irs}/(\alpha_{irs} + \beta_{irs})$ denote the mean of this Beta distribution. We are interested in estimating and assessing the significance of the difference in mean methylation levels across a region r for two biological conditions.

Our approach models transformed methylation proportions using GLS to obtain an approximation of the effect of interest. While directly modeling counts with either a Beta-Binomial Generalized Linear Model (GLM) or a Generalized Linear Mixed Model (GLMM) would allow us to accommodate complex covariance structures across samples and loci, it also results in complex likelihoods that require iterative maximization for each candidate region. Further, these procedures are subject to instability of estimation for methylation levels near the boundaries (zero and one) or non-identifiability in the case of separation as they occur in GLM (Gelman and others, 2008) and GLMM (Abrahantes and Aerts, 2012) estimation. GLS models, in contrast, are efficient and stable to estimate due to the availability of approximate closed-form parameter estimates. Though GLS does not model counts directly, we incorporate information lost after transformation of methylation proportions through specification of a variance estimate that depends on coverage.

We choose the arcsine link function $Z_{ijr} = \arcsin(2M_{ijr}/N_{ijr} - 1)$ to obtain transformed methylation proportions, as proposed by Park and Wu (2016) for DML analysis, for its desirable ability to stabilize the dependence of the variance on the mean methylation level. While the variance of methylation proportions M_{ijr}/N_{ijr} depends on the mean parameter π_{ijr} , the variance of Z_{ijr} only depends on coverage N_{ijr} and the dispersion of the Beta-Binomial distribution (refer to Section 2.6 of [supplementary material](#) available at *Biostatistics* online for more details). This helps us to form a statistic involving the transformed proportions that is exchangeable across regions that have different mean methylation values.

We assume a linear effect on the arcsine link-transformed methylation proportion parameters:

$$\arcsin(2\pi_{ijr} - 1) = \sum_{l=1}^{L_r} \beta_{0lr} 1_{[i=l]} + \beta_{1r} X_j = \mathbf{X}\boldsymbol{\beta}_r. \quad (3.1)$$

Here β_{0lr} are loci-specific intercept terms that account for variation on overall methylation levels across the region, where $l = 1, \dots, L_r$ and L_r denotes the number of loci in region r . The coefficient for the effect

of interest (e.g. biological group) is β_{1r} . We denote the design matrix as \mathbf{X} and the $(L_r + 1)$ -length vector of all coefficients $(\beta_{01r}, \beta_{02r}, \dots, \beta_{0L_r r}, \beta_{1r})$ as $\boldsymbol{\beta}_r$. This leads to the following model for the transformed response $\mathbf{Z}_r = (Z_{11r}, \dots, Z_{L_r r})$ in region r

$$\mathbf{Z}_r = \mathbf{X}\boldsymbol{\beta}_r + \boldsymbol{\epsilon}_r, \tag{3.2}$$

where we assume that $E[\boldsymbol{\epsilon}_r] = 0$ and $\text{Var}[\boldsymbol{\epsilon}_r] = \mathbf{V}_r$, which can be fit by GLS given an estimate of the covariance matrix \mathbf{V}_r . Since GLS allows arbitrary covariance structures, we use an autoregressive correlation structure to account for the correlation of methylation levels among nearby loci. To account for the dependence of the variance on coverage as mentioned above, we use variance weights. More details on the specific structure and estimation of \mathbf{V}_r are given in the next section.

With the above model, we assess the strength of the effect of the covariate of interest on methylation level within region r using the t -statistic t_r from the Wald test of the null hypothesis that $\beta_{1r} = 0$. Parameter estimates and their standard errors are obtained with the `gls` function in the `nlme` package (Pinheiro and others, 2017). Significance is evaluated by permutation using a pooled null distribution as described in detail in Section 3.2.3.

3.2.2. Covariance of methylation levels within regions In the estimation of the covariance matrix \mathbf{V}_r , we take into account biological variability through variance weighting, and correlation of nearby loci through an autocorrelation structure. The variance weighting is done to account for the dependence of the variance of transformed values Z_{ijr} on coverage. This variance depends non-linearly on N_{ijr} (Section 2.6 of supplementary material available at *Biostatistics* online), but in order to enable efficient closed-form estimation with GLS, we further approximate it by $\text{Var}(Z_{ijr}) \approx \sigma_r^2/N_{ijr}$. In addition, in order to construct a valid permutation test where the variance conditional on the effect of interest is invariant to permutation, we assume this variance is identical for all samples at a given loci by approximating N_{ijr} by $\text{med}_j(N_{ijr}) = N_{i,r}$.

To model correlation of nearby loci, we use the flexible continuous autoregressive correlation structure of order 1, abbreviated CAR(1). Under CAR(1), the correlation parameter depends on the length of the interval between the two observations considered through the equation $\rho_r(\tau) = e^{-\phi_r|\tau|}$, where τ is the length of the interval between two observations and ϕ_r is the positive positive continuous-time autoregressive coefficient (following the notation of Jones and Boadi-Boateng 1991) for region r . Thus, for subject j , the predicted methylation value for loci i at location t_{ijr} in region r given the methylation value at loci $i - 1$ is $\hat{Z}_{ijr} = \hat{Z}_{i-1,jr} e^{-\phi_r|t_{ijr} - t_{i-1,jr}|}$. If the error variance of the CAR(1) process is $\sigma_r^2 = \sigma_r^2/N_{i,r}$, and we let the correlation structure be nested within subject (i.e. such that observations from two subjects are independent), it follows that the covariance matrix for a given sample can be written

$$\text{Cov}(\mathbf{Z}_r) = \mathbf{V}_{jr} = \sigma_r^2 \mathbf{R}_{jr}, \text{ where the } mn^{\text{th}} \text{ element of } \mathbf{R}_{jr} \text{ is } \frac{e^{-\phi_r|t_{mjr} - t_{njr}|}}{\sqrt{N_{m,r}N_{n,r}}} \tag{3.3}$$

and for two subjects j and j' , $\text{Cov}(Z_{j'r}, Z_{ij'r}) = 0$.

The estimation of ϕ_r is computationally efficient to carry out on small to moderately sized regions. However, for larger regions with more than 40 loci we use the slightly simpler AR(1) correlation structure since it is many times faster to compute. This discrete formulation assumes that observations are equally spaced, and that observations that are separated by lag 1 are correlated with region-specific correlation parameter ρ_r . In addition, observations that are separated by m positions are correlated by ρ_r^m . Under this discrete formulation, the mn^{th} element of \mathbf{R}_{jr} from 3.3 becomes $\rho_r^{m-n} / \sqrt{N_{m,r}N_{n,r}}$.

The CAR(1) structure simplifies to the AR(1) process under certain conditions when observations are equally spaced (Jones and Boadi-Boateng, 1991). Thus the discrete AR(1) can be viewed as an

approximation of the CAR(1) when correlations are positive and the two provide increasingly more similar estimates as observations approach constant spacing. Indeed, when comparing model fits under both correlation structures in simulated data, the t -statistics for the coefficient of interest under CAR(1) generally converge to the estimates under AR(1) as the number of loci increases (Figure S1 and Section 2.5 of [supplementary material](#) available at *Biostatistics* online).

3.2.3. Permutation to generate a null set of regions The values of the covariate of interest (e.g. biological group) are permuted and the previous steps repeated in order to generate a set of statistics under the null hypothesis. Since the statistics account for known sources of variation that would otherwise prevent to comparison of regions across the genome, we can pool them together to form an approximate null distribution with as few as two samples per population. The empirical p -value is calculated by comparing the observed test statistics to the entire null set of statistics from all permutations. Control of FDR is carried out by adjusting the p -values using the procedure of [Benjamini and Hochberg \(1995\)](#).

4. RESULTS

For each of the data sets described in Section 2, we applied dmrseq, as well as three widely used methods for DMR detection: BSmooth ([Hansen and others, 2012](#)), DSS ([Park and Wu, 2016](#)), and metilene ([Juhling and others, 2016](#)). Each approach was evaluated based on the criteria detailed in the next subsections. For specific details on software implementation, refer to the [supplementary material](#) available at *Biostatistics* online (Section 3).

4.1. Simulation using dendritic cell data

Specificity was evaluated by identifying DMRs in null comparisons of two (N2) and three (N3) samples per group. Sensitivity was evaluated by identifying simulated DMRs in comparisons of two (D2) and three (D3) samples per group. Performance of each method is assessed by its ability to identify as many of the simulated DMRs as possible, while identifying as few DMRs as possible in the null comparison.

dmrseq did not identify any DMRs at the 0.05 level for the null comparisons N2 or N3 (Table 1). This remains true even when increasing the FDR threshold to 0.5 in both settings. In contrast, metilene identified a small number of DMRs, DSS identified many hundreds, and BSmooth tens of thousands using default settings (specific parameter specifications provided in Section 3 of [supplementary material](#) available at *Biostatistics* online). When applied to the data sets with simulated DMRs (D2 and D3), dmrseq is able to accurately control the FDR (Figure 1B), whereas metilene cannot (Figure S3 of [supplementary material](#) available at *Biostatistics* online). Note that analogous results cannot be obtained from DSS or BSmooth, as there is no way to specify FDR level.

BSmooth and DSS identify similar numbers of FP regions in D2 and D3 compared to the null setting of N2 and N3, and far more than dmrseq and metilene (Table 1). Although both BSmooth and DSS have favorable numbers of TPs, it is clear that this comes at the expense of lack of control of FDR (Figure 3). Similarly, metilene has favorable numbers of FPs, but this comes at the expense of low power. Further, dmrseq achieves higher power than the alternative methods at similar observed FDR levels, regardless of the effect size, CpG density, or coverage levels of the true regions (Figures S12–S14 of [supplementary material](#) available at *Biostatistics* online).

Although FDR thresholds are not available for BSmooth or DSS, we also investigated the sensitivity and specificity of other settings beyond defaults of the thresholds at the single-loci level (the loci t -statistic cutoff for BSmooth, and the loci p -value for DSS). Making these thresholds more conservative generally reduced the numbers of FP, but once again dmrseq was consistently able to identify more true positives at

Table 1. Null comparison and simulated DMR results for two sample size settings. Null comparison results for sample size 2 (N2) and sample size 3 (N3) are shown in the top two rows; simulated DMR results for sample size 2 (D2) and sample size 3 (D3) are shown in the bottom two rows. Numbers of DMRs identified by dmrseq and metilene are shown at the 0.05 FDR level. Default settings were used for BSmooth and DSS. True positives (TPs) is the number of simulated DMRs that are overlapped by at least one identified DMR. False positives (FPs) are DMRs that do not overlap any of the simulated DMRs. Note that all DMRs detected in N2 and N3 are FPs, since there are no true DMRs in the null setting

Comparison	Method	DMRs detected	TPs (unique)	FPs
		Null		
N2	dmrseq	0	—	0
	BSmooth	76 563	—	76 563
	DSS	661	—	661
	metilene	31	—	31
N3	dmrseq	0	—	0
	BSmooth	76 319	—	76 319
	DSS	770	—	770
	metilene	27	—	27
		Simulation		
D2	dmrseq	914	816	42
	BSmooth	73 252	2466	70 688
	DSS	2086	762	655
	metilene	329	210	30
D3	dmrseq	1620	1455	78
	BSmooth	72 764	2646	69 999
	DSS	2858	1257	763
	metilene	652	441	27

similar numbers of FP (see Section 4 and Figure S4 of [supplementary material](#) available at *Biostatistics* online).

We also stress that although lower FP rates could be achieved in this simulation study for BSmooth and DSS, individual loci thresholds do not correspond directly to specific FDRs at the region level. As a result, in practice, one must choose a threshold either by default settings, or by trial and error.

4.2. Human tissue and murine leukemia experimental data

To assess functional relevance of the results, the human tissue and murine leukemia studies were evaluated empirically based on the observed association of DMRs with differential expression by RNA-seq. Differentially expressed (DE) genes were identified using DESeq2 version 1.14.1 ([Love and others, 2014](#)). The association between expression level and methylation level was assessed for DMRs and DE genes using three measures of overlap: the gene body, promoter region, and island shore of a DE gene (see Section 1.5 of [supplementary material](#) available at *Biostatistics* online for more details). Methylation levels in each of these region types has been shown to be associated with the expression level of nearby genes ([Irizarry and others, 2009](#); [Lou and others, 2014](#)). Specifically, a DMR–DE gene pair is expected to have higher methylation values in the sample group with lower expression. The odds that the DMR and DE statistics are in opposing directions are calculated at various FDR cutoffs for dmrseq and metilene to assess whether top-ranked DMRs are more likely to be biologically relevant. The same is done for various cutoffs for the

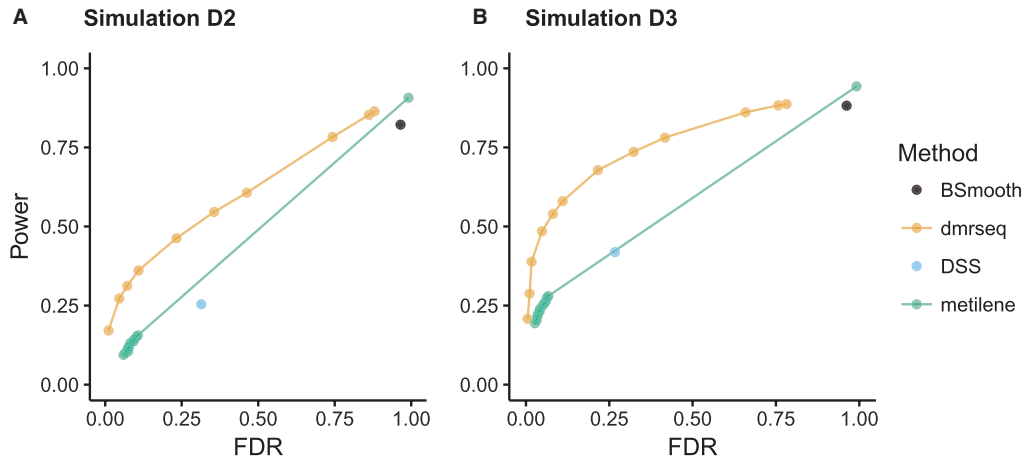


Fig. 3. dmrseq is more powerful than other methods. FDR and power results for (A) Simulation D2 and (B) Simulation D3, with method denoted by shading. dmrseq and metilene results are displayed for several different FDR cutoffs. Since region level FDR control is not possible for BSmooth and DSS, results using default settings are displayed. Power is calculated as the proportion of simulated DMRs overlapped by at least one identified DMR. FDR is calculated as the proportion of DMRs identified that do not overlap with any of the simulated DMRs.

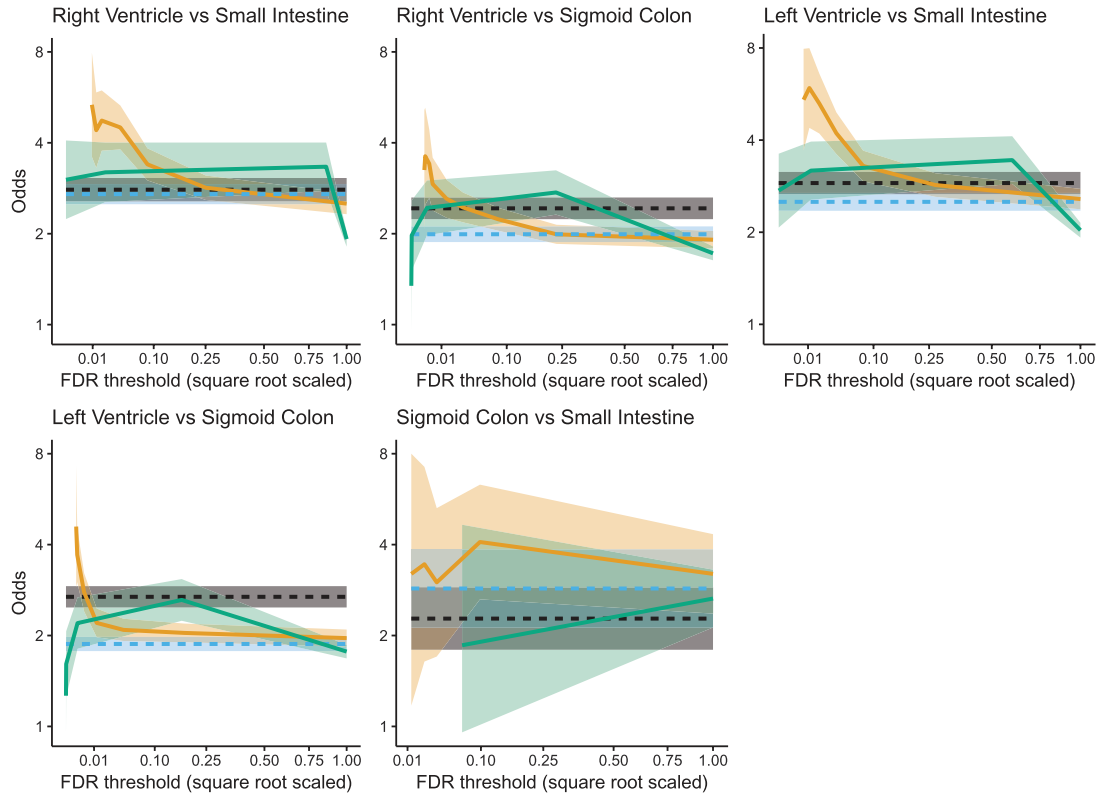
numbers of top-ranking regions by effect size. Additionally, for each cutoff we calculate the number of CpGs covered and the proportion of detected DMRs that overlap the island shore of a DE gene.

To qualitatively assess the ability of the dmrseq region-level summary statistic to rank DMRs as compared to other methods, we display example regions from the case studies. These examples illustrate the increased variability of regions that are highly ranked by naive statistics but not dmrseq (Figures 2 and S15 for the human and murine case studies, respectively). We include DMRs with concordant rankings that exhibit clear differences between two sample groups (Figures 2A and S15A). In contrast, the regions with discordant rankings between dmrseq q-value and mean difference (Figures 2B and S15B) and area statistics (Figure 2C and S15C) exhibit considerable variability between samples or loci (See Section 2.8 of [supplementary material](#) available at *Biostatistics* online for more details).

4.2.1. Tissue specificity in human samples For DSS, metilene, and dmrseq, the number of DMRs found (Table S2 of [supplementary material](#) available at *Biostatistics* online) parallels the numbers of DE genes found by DESeq2 (Table S4 of [supplementary material](#) available at *Biostatistics* online), but DSS generally found far more DMRs and metilene far fewer. For BSmooth, however, the number of DMRs identified was similar for all comparisons. This happens because the cutoff for the individual loci statistics is set by default at a quantile of the observed statistics, resulting in a similar number of loci being deemed significant.

The tissue-specific DMRs found by dmrseq are enriched for inverse associations with DE genes, and this enrichment is stronger for DMRs with lower FDRs (Figure 4, Figures S10–S11 of [supplementary material](#) available at *Biostatistics* online). Additionally, enrichment of dmrseq DMRs is generally stronger than that of alternative methods. While metilene also provides an FDR estimate, there is no consistent association between the FDR ranking and strength of association with expression. DMRs identified by BSmooth and DSS cannot be ranked by FDR and the default settings may not be ideal, so we also rank DMRs by effect size (raw methylation difference) with optimized parameter settings (see Section 3.2 of [supplementary material](#) available at *Biostatistics* online). The BSmooth and DSS DMRs with highest effect sizes exhibit comparable enrichment to dmrseq, with metilene considerably lower

A Roadmap Tissue Comparisons



B Murine Leukemia Models

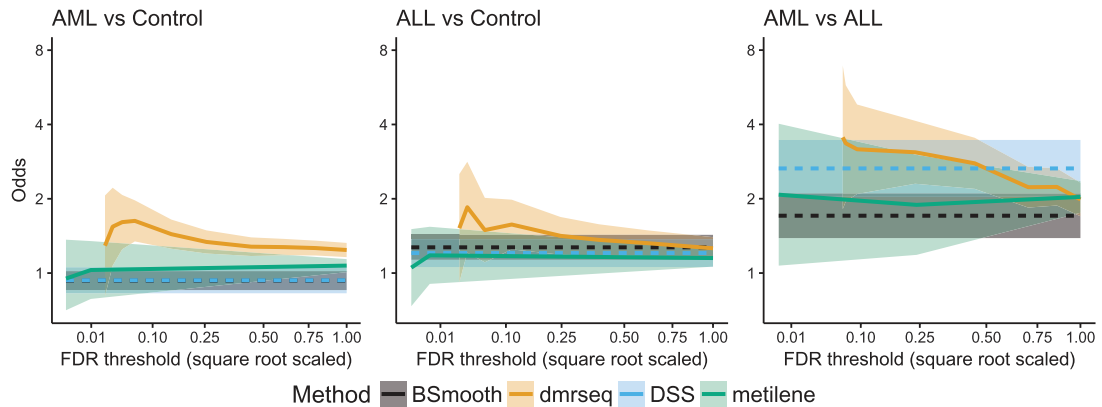


Fig. 4. dmrseq achieves stronger inverse association of methylation and differential expression at lower FDR thresholds. Odds of inverse association between methylation difference of (A) tissue-specific DMRs and (B) murine leukemia DMRs with differential expression of nearby DE genes (log2 scale) is displayed on the y-axis. For dmrseq and metilene, the x-axis represents the FDR threshold (square-root scaled) for which the odds calculation (cumulative) is performed. Since FDR cannot be specified for BSmooth or DSS, the odds are calculated over all DMRs identified and displayed as a horizontal line. Note that the comparison between left and right ventricles is not shown, since no DE genes were identified. A DMR is considered to overlap a gene if it includes any part of the 2 kb region upstream of the TSS.

(Figures S6 and S7 of [supplementary material](#) available at *Biostatistics* online). However, arbitrary cutoffs of effect size do not directly correspond to significance level.

4.2.2. DNMT3a loss in murine leukemia models In the murine leukemia models, dmrseq finds the most DMRs in the comparison of AML and the control (Table S5 of [supplementary material](#) available at *Biostatistics* online), which is also the comparison for which the most DE genes were identified (see Table S7 of [supplementary material](#) available at *Biostatistics* online). In contrast, DSS and metline both find the most DMRs in the comparison with the fewest DE genes identified, and BSmooth identified similar numbers of DMRs in each comparison, each with far more DMRs than the other methods.

The murine leukemia DMRs found by dmrseq are enriched for inverse associations with DE genes, and this enrichment is stronger for DMRs with lower FDRs (Figure 4, Figures S10–S11 of [supplementary material](#) available at *Biostatistics* online). Additionally, enrichment is generally stronger than that of BSmooth, DSS, and metline. While metline also provides an FDR estimate, there is no consistent association between the FDR ranking and strength of association with expression. Similar to the tissue specificity analysis, BSmooth and DSS DMRs with highest effect sizes exhibit comparable enrichment to dmrseq, with metilene considerably lower, and the enrichment when including all DMRs often drops lower for BSmooth, DSS, or metline than for dmrseq (Figures S8 and S9 of [supplementary material](#) available at *Biostatistics* online).

5. DISCUSSION

We have described dmrseq, a method useful for discovering and prioritizing DMRs from WGBS data. The approach is based on rigorous statistical reasoning and is the first method that permits accurate inference on DMRs that are found by scanning the genome. By developing a transformation that results in summary statistics from candidate regions being exchangeable, we are able to borrow strength across the genome to build a null distribution that permits inference with a sample size as small as two. We have demonstrated how the method clearly outperforms currently used tools with several experimental data examples and Monte Carlo simulation. The method is implemented as open source software in the form of an R package.

6. SOFTWARE

The R package dmrseq is available on GitHub at <https://github.com/kdkorthauer/dmrseq>.

SUPPLEMENTARY MATERIAL

Supplementary material is available online at <http://biostatistics.oxfordjournals.org>. The simulated benchmark data is freely available for download via FigShare (https://figshare.com/projects/Whole_Genome_Bisulfite_Sequencing_WGBS_Benchmark_Data/27532). Annotated scripts for the simulation and case study analyses are available in the GitHub repository <https://github.com/kdkorthauer/dmrseqPaper>.

ACKNOWLEDGMENTS

The authors thank two anonymous reviewers for providing comments and suggestions that helped us to improve quality and clarity of the manuscript. *Conflict of Interest*: None declared.

FUNDING

National Institutes of Health (NIH) (R01 grants R01HG005220, R01GM083084, and U41HG007000), in part.

REFERENCES

- ABRAHANTES, J. C. AND AERTS, M. (2012). A solution to separation for clustered binary data. *Statistical Modelling* **12**, 3–27.
- AKALIN, A., KORMAKSSON, M., LI, S., GARRETT-BAKELMAN, F. E., FIGUEROA, M. E., MELNICK, A. AND MASON, C. E. (2012). methylKit: a comprehensive R package for the analysis of genome-wide DNA methylation profiles. *Genome Biology* **13**, R87.
- ARYEE, M. J., JAFFE, A. E., CORRADA-BRAVO, H., LADD-ACOSTA, C., FEINBERG, A. P., HANSEN, K. D. AND IRIZARRY, R. A. (2014). Minfi: a flexible and comprehensive Bioconductor package for the analysis of infinium DNA methylation microarrays. *Bioinformatics* **30**, 1363–1369.
- BENJAMINI, Y. AND HOCHBERG, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B* **57**, 289–300.
- BENJAMINI, Y., TAYLOR, J. AND IRIZARRY, R. A. (2016). Selection corrected statistical inference for region detection with high-throughput assays. *bioRxiv*, doi:10.1101/082321.
- BIRD, A. (2002). DNA methylation patterns and epigenetic memory. *Genes & Development* **16**, 6–21.
- CLEVELAND, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association* **74**, 829–836.
- DOLZHENKO, E. AND SMITH, A. D. (2014). Using beta-binomial regression for high-precision differential methylation analysis in multifactor whole-genome bisulfite sequencing experiments. *BMC Bioinformatics* **15**, 215.
- GELMAN, A., JAKULIN, A., PITTAU, M. G. AND SU, Y.-S. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics* **2**, 1360–1383.
- HANSEN, K. D., LANGMEAD, B. AND IRIZARRY, R. A. (2012). BSmooth: from whole genome bisulfite sequencing reads to differentially methylated regions. *Genome Biology* **13**, R83.
- HE, Y., HARIHARAN, M., GORKIN, D. U., DICKEL, D. E., LUO, C., CASTANON, R. G., NERY, J. R., LEE, A. Y., WILLIAMS, B. A., TROUT, D. and others. (2017). Spatiotemporal DNA methylome dynamics of the developing mammalian fetus. *bioRxiv*, doi:10.1101/166744.
- HEBESTREIT, K., DUGAS, M. AND KLEIN, H. U. (2013). Detection of significantly differentially methylated regions in targeted bisulfite sequencing data. *Bioinformatics* **29**, 1647–1653.
- IRIZARRY, R. A., LADD-ACOSTA, C., WEN, B., WU, Z., MONTANO, C., ONYANGO, P., CUI, H., GABO, K., RONGIONE, M., WEBSTER, M. and others. (2009). The human colon cancer methylome shows similar hypo- and hypermethylation at conserved tissue-specific CpG island shores. *Nature Genetics* **41**, 178–186.
- JAFFE, A. E., MURAKAMI, P., LEE, H., LEEK, J. T., FALLIN, M. D., FEINBERG, A. P. AND IRIZARRY, R. A. (2012). Bump hunting to identify differentially methylated regions in epigenetic epidemiology studies. *International Journal of Epidemiology* **41**, 200–209.
- JONES, R. H. AND BOADI-BOATENG, F. (1991). Unequally Spaced Longitudinal Data with AR(1) Serial Correlation. *Biometrics* **47**, 161–175.
- JUHLING, F., KRETZMER, H., BERNHART, S. H., OTTO, C., STADLER, P. F. AND HOFFMANN, S. (2016). metilene: fast and sensitive calling of differentially methylated regions from bisulfite sequencing data. *Genome Research* **26**, 256–262.
- KHAMIS, A. M., LIOZNOVA, A. V., ARTEMOV, A. V., RAMENSKY, V., BAJIC, V. B. AND MEDVEDEVA, Y. A. (2017). CpG traffic lights are markers of regulatory regions in humans. *bioRxiv*, doi:10.1101/095968.

- KRUEGER, F. AND ANDREWS, S. R. (2011). Bismark: a flexible aligner and methylation caller for Bisulfite-Seq applications. *Bioinformatics* **27**, 1571–1572.
- LEE, W. AND MORRIS, J. S. (2016). Identification of differentially methylated loci using wavelet-based functional mixed models. *Bioinformatics* **32**, 664–672.
- LEEK, J. T., SCHARPF, R. B., BRAVO, H. C., SIMCHA, D., LANGMEAD, B., JOHNSON, W. E., GEMAN, D., BAGGERLY, K. AND IRIZARRY, R. A. (2010). Tackling the widespread and critical impact of batch effects in high-throughput data. *Nature Reviews Genetics* **11**, 733–739.
- LOADER, C. (1999). *Local Regression and Likelihood*. New York: Springer.
- LOU, S., LEE, H.-M., QIN, H., LI, J.-W., GAO, Z., LIU, X., CHAN, L. L., KL LAM, V., SO, W.-Y., WANG, Y. *and others*. (2014). Whole-genome bisulfite sequencing of multiple individuals reveals complementary roles of promoter and gene body methylation in transcriptional regulation. *Genome Biology* **15**, 408.
- LOVE, M. I., HUBER, W. AND ANDERS, S. (2014). Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biology* **15**, 550.
- LUN, A. T. AND SMYTH, G. K. (2014). De novo detection of differentially bound regions for ChIP-seq data using peaks and windows: controlling error rates correctly. *Nucleic Acids Research* **42**, e95.
- MARX, V. (2016). Genetics: profiling DNA methylation and beyond. *Nature Methods* **13**, 119–122.
- MAYO, T. R., SCHWEIKERT, G. AND SANGUINETTI, G. (2015). M3D: a kernel-based test for spatially correlated changes in methylation profiles. *Bioinformatics* **31**, 809–816.
- PACIS, A., TAILLEUX, L., MORIN, A. M., LAMBOURNE, J., MACISAAC, J. L., YOTOVA, V., DUMAINE, A., DANCKAERT, A., LUCA, F., GRENIER, J. C. *and others*. (2015). Bacterial infection remodels the DNA methylation landscape of human dendritic cells. *Genome Research* **25**, 1801–1811.
- PARK, Y., FIGUEROA, M. E., ROZEK, L. S. AND SARTOR, M. A. (2014). MethylSig: a whole genome DNA methylation analysis pipeline. *Bioinformatics* **30**, 2414–2422.
- PARK, Y. AND WU, H. (2016). Differential methylation analysis for BS-seq data under general experimental design. *Bioinformatics* **32**, 1446–1453.
- PINHEIRO, J., BATES, D., DEBROY, S., SARKAR, D. AND R CORE TEAM. (2017). nlme: linear and nonlinear mixed effects models. R package version 3.1-131. <https://CRAN.R-project.org/package=nlme>.
- ROBINSON, M. D., KAHRAMAN, A., LAW, C. W., LINDSAY, H., NOWICKA, M., WEBER, L. M. AND ZHOU, X. (2014). Statistical methods for detecting differentially methylated loci and regions. *Frontiers in Genetics* **5**, 324.
- SAITO, Y., TSUJI, J. AND MITUYAMA, T. (2014). Bisulfighter: accurate detection of methylated cytosines and differentially methylated regions. *Nucleic Acids Research* **42**, e45.
- SCHULTZ, M. D., HE, Y., WHITAKER, J. W., HARIHARAN, M., MUKAMEL, E. A., LEUNG, D., RAJAGOPAL, N., NERY, J. R., URICH, M. A., CHEN, H. *and others*. (2015). Human body epigenome maps reveal noncanonical DNA methylation variation. *Nature* **523**, 212–216.
- SHAFI, A., MITREA, C., NGUYEN, T. AND DRAGHICI, S. (2018). A survey of the approaches for identifying differential methylation using bisulfite sequencing data. *Briefings in Bioinformatics* **19**, 737–753.
- SIEGMUND, D. O., ZHANG, N. R. AND YAKIR, B. (2011). Miscellanea false discovery rate for scanning statistics. *Biometrika* **98**, 979–985.
- SMITH, Z. D. AND MEISSNER, A. (2013). DNA methylation: roles in mammalian development. *Nature Reviews Genetics* **14**, 204–220.
- SUN, D., XI, Y., RODRIGUEZ, B., PARK, H. J., TONG, P., MEONG, M., GOODELL, M. A. AND LI, W. (2014). MOABS: model based analysis of bisulfite sequencing data. *Genome Biology* **15**, R38.

- WEN, Y., CHEN, F., ZHANG, Q., ZHUANG, Y. AND LI, Z. (2016). Detection of differentially methylated regions in whole genome bisulfite sequencing data using local Getis-Ord statistics. *Bioinformatics* **32**, 3396–3404.
- WU, H., XU, T., FENG, H., CHEN, L., LI, B., YAO, B., QIN, Z., JIN, P. AND CONNEELY, K. N. (2015). Detection of differentially methylated regions from whole-genome bisulfite sequencing data without replicates. *Nucleic Acids Research* **43**, e141.
- YANG, L., RODRIGUEZ, B., MAYLE, A., PARK, H. J., LIN, X., LUO, M., JEONG, M., CURRY, C. V., KIM, S. B., RUAU, D. and others. (2016). DNMT3A loss drives enhancer hypomethylation in FLT3-ITD-associated leukemias. *Cancer Cell* **29**, 922–934.
- YU, X. AND SUN, S. (2016). HMM-DM: identifying differentially methylated regions using a hidden Markov model. *Statistical Applications in Genetics and Molecular Biology* **15**, 69–81.
- ZILLER, M. J., HANSEN, K. D., MEISSNER, A. AND ARYEE, M. J. (2015). Coverage recommendations for methylation analysis by whole-genome bisulfite sequencing. *Nature Methods* **12**, 230–232.

[Received August 31, 2017; revised December 19, 2017; accepted for publication January 21, 2018]