scMET: Bayesian modelling of DNA methylation heterogeneity at single-cell resolution

Chantiriolnt-Andreas Kapourani^{1,2,*} Ricard Argelaguet^{3,*}
Guido Sanguinetti^{2,4,†} Catalina A. Vallejos^{1,5,†}

¹MRC Human Genetics Unit, University of Edinburgh, UK

²School of Informatics, University of Edinburgh, UK

³European Bioinformatics Institute (EMBL-EBI), Hinxton, UK

⁴SISSA, International School of Advanced Studies, Trieste, Italy

⁵ The Alan Turing Institute, London, UK

*These authors contributed equally †Corresponding author

Email: catalina.vallejos@igmm.ed.ac.uk (C.A.V.); gsanguin@sisa.it (G.S.)

Abstract

High throughput measurements of DNA methylomes at single-cell resolution are a promising resource to quantify the heterogeneity of DNA methylation and uncover its role in gene regulation. However, limitations of the technology result in sparse CpG coverage, effectively posing challenges to robustly quantify genuine DNA methylation heterogeneity. Here we tackle these issues by introducing scMET, a hierarchical Bayesian model which overcomes data sparsity by sharing information across cells and genomic features, resulting in a robust and biologically interpretable quantification of variability. scMET can be used to both identify highly variable features that drive epigenetic heterogeneity and perform differential methylation and differential variability analysis between pre-specified groups of cells. We demonstrate scMET's effectiveness on some recent large scale single cell methylation datasets, showing that the scMET feature selection approach facilitates the characterisation of epigenetically distinct cell populations. Moreover, we illustrate how scMET variability estimates enable the formulation of novel biological hypotheses on the epigenetic regulation of gene expression in early development. An R package implementation of scMET is publicly available at <https://github.com/andreaskapou/scMET>.

Introduction

DNA methylation (DNAm) at cytosine residues plays an important role in the regulation of gene expression (Jaenisch and Bird, 2003). It is also critical for a broad range of biological processes, including X-chromosome inactivation, genomic imprinting and cancer (Avner and Heard, 2001; Baylin and Jones, 2011; Reik and Walter, 2001). The gold standard approach to profile DNAm at single-base resolution is to treat DNA with sodium bisulphite, which efficiently converts unmethylated cytosines to uracils, while leaving methylated cytosines unmodified (Krueger *et al.*, 2012). Although bulk bisulphite sequencing (BS-seq) experiments have paved the way for mapping the methylome landscape across different tissues, they fall

short of explaining the inter-cellular methylation heterogeneity and quantifying its dynamics in a variety of biological contexts (Schwartzman and Tanay, 2015).

More recently, advances in sequencing technologies have enabled the development of protocols that profile DNAm with single-cell resolution (e.g. scBS-seq, Guo *et al.*, 2013; Smallwood *et al.*, 2014) and multiplexing protocols offer scalability to thousands of cells in a single experiment (Luo *et al.*, 2017; Mulqueen *et al.*, 2018). In contrast to gene expression signatures from scRNA-seq experiments, which are influenced by the environment, DNA methylation profiles are highly distinct between cell types and stable across individuals and over the life span (Lister *et al.*, 2013; Mo *et al.*, 2015). Moreover, whilst scRNA-seq assays might fail to capture information about genes with moderate expression levels, cell-level measurements of DNAm offer a more complete coverage across genomic regions (Luo *et al.*, 2017). However, due to the small amounts of initial genomic DNA and the destructive nature of bisulphite on nucleic acids, the output data are often noisy and extremely sparse; that is, a large proportion of CpG dinucleotides is not observed (ranging from 80% to 95%). While tailored computational imputation methods such as Melissa (Kapourani and Sanguinetti, 2019) and DeepCpG (Angermueller *et al.*, 2017) might ameliorate the sparsity problem, disentangling genuine epigenetic variability from technical biases remains a formidable problem.

Here we present scMET, a Bayesian framework that addresses the statistical challenges associated with sparse scBS-seq data and provides novel functionality that is tailored to single-cell level datasets. To overcome sparsity, scMET aggregates the input data within regions (hereafter also referred to as genomic *features*): either by combining CpG information in a sliding window approach or using pre-annotated contexts, such as promoter regions or enhancers (Gravina *et al.*, 2016; Smallwood *et al.*, 2014). To dissect genuine epigenetic variability from the many confounding technical biases, scMET adopts a hierarchical model specification which shares information across cells and genomic features, whilst incorporating feature-level characteristics (e.g. CpG density). Critically, scMET introduces residual *overdispersion* estimates as a measure of DNAm variability that is not confounded by differences in mean methylation. These estimates can be used to perform differential DNAm variability testing among groups of cells, embracing the cellular resolution of the data to provide novel insights which are not possible using traditional differential mean tests on bulk data (Lähnemann *et al.*, 2020). scMET can also identify highly variable features (HVF) which, among others, can be used as input for unsupervised clustering analyses.

scMET scales readily to thousands of cells and features, making it a powerful tool for large scale single-cell epigenetic studies. Our results both on simulated and real datasets demonstrate that it can accurately and robustly quantify DNAm heterogeneity. Results on two recent large scale datasets show that scMET detects biologically relevant highly variable features which result in improved clustering performance. In addition, we show that scMET can facilitate the interrogation of single-cell multi-omics assays, yielding novel biological hypotheses on the role of epigenetic variability in gene regulation in early development.

Results

Quantifying cell-to-cell DNAm heterogeneity with scMET

To disentangle technical from biological variability and overcome data sparsity, scMET couples a hierarchical beta-binomial (BB) model with a generalised linear model (GLM) framework (Fig. 1a-b). For each cell i and feature j , the input for scMET is the number of CpG sites that are observed to be methylated (Y_{ij}) and the total number of sites for which methylation status was recorded (n_{ij}). The BB model uses feature-specific mean parameters μ_j to quantify overall DNAm across all cells and biological *overdispersion* parameters γ_j as a proxy for cell-to-cell DNAm heterogeneity. The latter capture the amount of variability that is not explained by binomial sampling noise, which would only account for technical variation. Hence, γ_j is akin to the overdispersion term used in negative binomial models for RNA-seq data (e.g. Love *et al.*, 2014). Although BB models have been developed for bulk DNAm data (e.g. Dolzhenko and Smith, 2014; Feng *et al.*, 2014), they typically use data from individual CpG sites as input; a strategy prone to fail for the highly sparse scBS-seq data.

The GLM framework is incorporated at two levels. Firstly, to introduce feature-specific covariates $\mathbf{x}_j$ (e.g. CpG density) that may explain differences in mean methylation μ_j across features. Secondly, similar to Eling *et al.* (2018), we use a non-linear regression framework to capture the mean-overdispersion trend that is typically observed in high throughput sequencing data, such as scBS-seq (Fig. 1c). Critically, this trend is used to derive *residual overdispersion* parameters ϵ_j — a measure of cell-to-cell variability that is not confounded by mean methylation. Feature-specific parameters are subsequently used for: (i) feature selection, to identify highly variable features (HVF) that drive cell-to-cell epigenetic heterogeneity (Fig. 1d) and (ii) differential methylation testing, to highlight features that show differences in DNAm mean or variability between specified groups of cells (Fig. 1e).

By using a Bayesian formulation, scMET infers the posterior distribution for all model parameters (*Methods*). Moreover, a variational Bayes scheme (Blei *et al.*, 2017) permits scalable analysis to thousands of cells and features (Supplementary Fig. S1), while having comparable posterior inference performance when compared to a Markov Chain Monte Carlo implementation (Supplementary Fig. S2 and S3). As in Bochkina and Richardson (2007), the output generated by scMET is used to implement a probabilistic decision rule to enable HVF selection and differential methylation testing. The decision rule is calibrated to control the expected false discovery rate (EFDR, Newton *et al.*, 2004). A more detailed description of the model specification and its implementation is provided in the *Methods* section. scMET is implemented as an R package and is available at <https://github.com/andreaskapou/scMET>.

Figure 1: Graphical outline for scMET. (a) Overview of the scMET probabilistic graphical model. The random variables and data that form the model, along with the distributional assumptions, are shown. Input values are denoted by gray circles. Model parameters are denoted by white circles. (b) scMET uses single-cell DNAm data as input. The data could consist of measurements obtained from different groups of cells, such as experimental conditions or cell types (represented by green and orange colours in the diagram). For each region of interest (e.g. promoters), the input data is recorded in terms of the number of CpG sites for which a valid measurement was recorded and, among those, the number of methylated CpG sites. (c) By combining a hierarchical beta-binomial specification with a generalized linear model framework, scMET captures the mean-overdispersion relationship (left) that is typically observed in bisulphite sequencing readouts and derives residual overdispersion estimates that are not confounded by mean methylation (right). (d) scMET can be used to identify HVFs that drive epigenetic heterogeneity within a cell population. For example, these could be used as the input of dimensionality reduction techniques or clustering analyses. (e) scMET uses a probabilistic decision rule to perform differential methylation analysis: to identify features that show differences in mean methylation (left) and/or methylation variability (right) between pre-specified groups of cells.

Benchmarking scMET on synthetic data

First, we benchmark the performance of scMET using synthetic data. To mimic the properties observed in real scBS-seq data, we simulated features with rich and poor CpG density (see *Methods* for details about the simulation settings). We compared mean and overdispersion estimates obtained by scMET with respect to BB maximum likelihood estimates (BB MLE), which were obtained separately for each feature. As expected, mean parameters μ_j are easier to infer and estimates were comparable across both methods (Supplementary Fig. S4). However, scMET outperformed BB MLE when inferring overdispersion parameters γ_j , particularly for small numbers of cells (Supplementary Fig. S5).

To assess whether the shrinkage introduced by scMET improves overdispersion estimates in real data, we performed down-sampling experiments based on a subset of the dataset introduced by Luo *et al.* (2017). For this analysis, we focused on 424 inhibitory neurons (a more detailed description is provided in *Methods*). We compared estimates obtained using the full and down-sampled datasets (Supplementary Fig. S6). We observed scMET posterior estimates to be more stable than BB MLE as the sample size decreased, suggesting that scMET leads to more robust inference. This is particularly important for rare cell populations or during early development, where large numbers of cells are difficult to obtain. In combination with the simulation study described above, this showcases the benefits of using a Bayesian hierarchical framework to share information across cells and genomic features.

Finally, we evaluated the performance of scMET as a tool to identify differentially methylated (DM) and differentially variable (DV) features. For this purpose, we generated synthetic data representing two cell types, for varying number of cells and different effect sizes in terms of log-odds ratio (Supplementary Note S2.3). For the DM analysis, we compared scMET against the Fisher’s exact test, which has been previously used for BS-seq data (e.g. Challen *et al.*, 2012). We observed better performance for scMET in terms of the F1-measure, especially for CpG rich features (Supplementary Fig. S7). scMET led to better type I error control, but to more conservative results (Supplementary Fig. S8). In terms of DV testing, the simulation study showed that for small effect sizes we would need more than 200 cells to achieve 50% to 80% power, whereas for features with larger effect sizes we would need around 50 cells per group to achieve 80% power (Supplementary Fig. S9).

scMET improves feature selection for unsupervised analysis of single-cell methylomes from the mouse frontal cortex

To demonstrate the performance of scMET in real data, we considered a dataset where DNA methylation was profiled in 3,069 cells isolated from the frontal cortex of young adult mice (Luo *et al.*, 2017). To date, this is one of the largest and most heterogeneous publicly available scBS-seq datasets. The main source of heterogeneity in this dataset is due to two broad classes of neurons: excitatory ($I = 2,645$) and inhibitory ($I = 424$). Within each class, a hierarchy of sub-populations can be identified according to the cortical depth (Fig. 2a), where excitatory neurons progress from deep layers (mDL-1, mDL-2, mDL-3, mL6-1, mL6-2) to middle (mL5-1, mL5-2, mL4) and superficial layers (mL2/3). These groups were validated in the original study, and thus can be used as a benchmark for clustering analyses.

We applied scMET to genomic features from three different putative regulatory elements: gene promoters within 2kb around transcription start site ($J = 12,774$), distal H3K27ac peaks ($J = 17,284$) and H3K4me1 peaks ($J = 30,374$). As expected, scMET captured the mean-overdispersion relationship within each genomic context, and estimates for residual overdispersion parameters ϵ_j were not confounded by mean DNAm (Supplementary Fig. S10).

Here, we illustrate scMET as a feature selection tool, using residual overdispersion estimates to identify HVFs that can be used as input for unsupervised analyses, such as clustering. For each genomic context, we selected HVFs (Supplementary Fig. S11a) and performed a clustering analysis with varying numbers of HVFs (ranked by decreasing values of their associated tail posterior probabilities) as input. More concretely, we performed dimensionality reduction followed by k -means clustering (*Methods*) and used the adjusted Rand index (ARI, Hubert and Arabie, 1985) to quantify agreement with respect to the sub-populations validated by Luo *et al.* As a comparison, we also evaluated two alternative HVF selection strategies based on Gaussian and binomial models (*Methods*). As expected, the clustering performance improved steadily with increasing number of HVFs for all three methods. However, scMET consistently led to better clustering performance (Fig. 2b and Supplementary Fig. S11b, as well as Supplementary Fig. S12 and S13 for visual inspection in a low dimensional space). We could already separate inhibitory from excitatory neurons using only the top 100 HVFs obtained by scMET, and generally resulted in more distinct cell sub-populations. In all cases, promoters were less able to disentangle the neuronal sub-populations. This is consistent with the lower overdispersion levels observed in this genomic context (Supplementary Fig. S11c).

To facilitate interpretation for the HVFs highlighted by scMET, we linked genomic features to genes by overlapping the genomic coordinates allowing for a maximum distance of 20kb from the transcription start site in the case of distal elements. We explored whether features identified as HVF (red points in Fig. 2c and Supplementary Fig. S14a) were enriched for neuronal markers identified in the Luo *et al.* (2017) study (Supplementary Table S1). This enrichment was observed for distal H3K27ac and H3K4me1 marks, but not for promoters (Supplementary Fig. S14b). As representative examples, we display three distal H3K27ac elements among the HVFs that are located proximal to known gene markers of each neuron class: *Gad1* for inhibitory, *Neurod6* and *Sv2b* for excitatory (Fig. 2d).

Figure 2: Feature selection using scMET characterises known heterogeneity on the mouse frontal cortex. (a) UMAP (Becht *et al.*, 2019) representation of neuron sub-populations present in the mouse frontal cortex dataset by combining the top 4,000 HVFs identified by scMET across distal H3K27ac and H3K4me1 genomic contexts. Cells are coloured according to the original cell sub-population assignments obtained by Luo *et al.* (b) Clustering performance in terms of adjusted Rand index (ARI), for varying number of HVFs. HVF selection was based on scMET’s residual overdispersion parameters ϵ_j (yellow), binomial variance (grey) and Gaussian variance (blue). A finite grid of HVFs was used for ARI evaluation and non-parametric regression was used to obtain a smoothed interpolation across all values (Methods). (c) Identifying HVFs for the distal H3K27ac genomic context. Red points correspond to features being called as HVF (EFDR = 10% and percentile threshold $\delta_E = 90\%$). To ease interpretation each element is linked to its nearest gene. Labels highlight known neuron marker genes that were used in the Luo *et al.* (2017) study to define the different cell populations. (d) Example distal H3K27ac HVFs whose methylation patterns distinguish the two broad neuronal populations. Panel titles correspond to the nearest genes.

scMET enables differential methylation testing between cellular sub-populations

To showcase scMET as a differential methylation tool, we applied it on the same mouse frontal cortex dataset (Luo *et al.*, 2017), after separating the cells in excitatory and inhibitory groups. Initially, we applied scMET to characterise differential methylation (DM), i.e. changes in mean methylation. Across all genomic contexts, we observed a substantially larger fraction of features being hyper-methylated in inhibitory compared to excitatory neurons (Fig. 3a and Supplementary Fig. S15a). Within distal H3K27ac peaks, for instance, scMET identified 5,242 features to have higher methylation levels in inhibitory neurons, compared to only 935 features showing higher methylation in the excitatory group (Fig. 3a-b). After mapping features to their nearest gene, we observed that DM hits were enriched for known marker genes that differentiate inhibitory and excitatory neurons (Supplementary Fig. S15b).

Besides DM testing, the primary focus of the scMET differential test is to identify changes in cell-to-cell methylation variability. In principle, differential variability (DV) testing could be based on feature-specific overdispersion parameters γ_j , but these results would be confounded by the mean-overdispersion trend (Fig. 1c). Hence, meaningful DV analysis based on γ_j would need to be restricted to non-DM features. Instead, we propose to perform DV analysis based on residual overdispersion parameters ϵ_j . For the mouse frontal cortex dataset, we identified a large number of DV features across genomic contexts, except from promoter regions which showed tighter methylation patterns across inhibitory and excitatory neurons (Fig. 3c and Supplementary Fig. S16a). Critically, the procedure for calling DV features was not confounded by mean methylation levels (Fig. 3c and Supplementary Fig. S16b).

As representative examples, we show two distal H3K27ac peaks that are located proximal to neuronal markers and exhibit higher variability in excitatory neurons (Fig. 3d). Both features show substantial variation across the different sub-types of excitatory neurons: the first is mostly unmethylated in the mL6 cortical layer, the second is mostly unmethylated in the superficial cortical layer. These patterns are consistent with previously reported spatial expression for *Tle4* (mostly expressed in the deep cortical layer, see Sorensen *et al.*, 2015; Zahr *et al.*, 2018) and *Cux1* (which shows expression specificity for the superficial layer, see Georgala *et al.*, 2011; Zahr *et al.*, 2018). It should be noted that these features would be excluded from DV analysis based on γ_j , since they are also DM between the two broad classes (Fig. 3a). In summary, these findings demonstrate the ability of scMET to identify potential markers that drive between and within cell population heterogeneity.

Figure 3: **Summary of changes in methylation patterns (mean and variability) for inhibitory and excitatory neurons.** (a) Identifying differentially methylated (DM) features for the distal H3K27ac genomic context. Green and pink points correspond to features showing higher mean methylation in inhibitory (Inh+) and excitatory (Exc+) neurons, respectively. Labels highlight neuron marker genes that were used in [Luo *et al.* \(2017\)](#) to define the different cell populations. Blue vertical dashed lines correspond to log-odds ratio threshold $\psi_M = \pm \log(2)$. Yellow horizontal dashed line is located at posterior evidence probability cut-offs defined by EFDR = 5%. (b) Representative heatmap of methylation rates (Y_{ij}/n_{ij}) across cells (columns) and features (rows) for the distal H3K27ac genomic context. Cells are grouped in excitatory and inhibitory classes. A set of randomly selected features is displayed. These are grouped according to the DM analysis output as: Exc+, Inh+ and no mean methylation difference. The colour code represents features with low (0, green) and high (1, yellow) mean methylation level. Features with no CpG coverage are denoted with white colour. (c) Identifying differentially variable (DV) features for the distal H3K27ac genomic context. Green and pink points correspond to features showing higher methylation variability in inhibitory and excitatory neurons, respectively. Blue dashed lines correspond to log-odds ratio threshold $\psi_E = \pm \log(1.5)$. Yellow dashed line is located at posterior evidence probability cut-offs defined by EFDR = 5%. For each feature, posterior estimates for the change in residual overdispersion parameter ϵ_j between excitatory and inhibitory neurons is plotted against the posterior tail probability of calling a feature as DV (left). For each feature, posterior estimates for mean methylation parameter μ_j is plotted against posterior estimates for the change in residual overdispersion parameter ϵ_j between excitatory and inhibitory neurons (right). (d) Example features that are called as being more variable in excitatory neurons. Left subplots show broad differences in methylation patterns. Right subplots show methylation patterns separately within each broad neuronal class. Each data point represents the methylation rate for a cell.

Exploring the relationship between transcriptional and DNAm variability using single-cell multi-omics data

As a second use case, we considered a single-cell multi-omics dataset where scNMT-seq (Clark *et al.*, 2018) was employed to profile RNA expression, DNAm and nucleosome occupancy at single-cell resolution, spanning multiple time points from the exit from pluripotency to primary germ layer specification (Argelaguet *et al.*, 2019). The multi-modal nature of this dataset provides a unique opportunity to link cell-to-cell variation between DNAm and transcription across individual cells. Here we used scMET to quantify DNAm variability at promoter elements, which we subsequently contrasted to RNA expression heterogeneity for the corresponding genes. For this analysis we exclusively used promoter elements as, unlike distal regulatory elements, they can be unambiguously matched to their respective genes.

For each gene, we quantified transcriptional heterogeneity using the residual overdispersion estimates generated by BASiCS (Supplementary Fig. S17a, Eling *et al.*, 2018). Promoter DNAm variability was calculated using the residual overdispersion estimates inferred using scMET (Supplementary Fig. S17b). More details about these analyses and the associated data pre-processing steps are described in the *Methods* section.

When comparing residual overdispersion estimates for RNA expression and promoter DNAm, there was no clear genome-wide association (Fig. 4a). However, when restricting to genes that display high levels of transcriptional variability, two main groups can be identified. The first category corresponds to genes with low levels of promoter DNAm residual overdispersion, and it includes differentiation and germ layer markers such as *Mesp1*, *Lefty2*, *Id3* (mesoderm) and *Cldn6*, *Cer1* and *Krt8* (endoderm). The second category is characterised by genes with high promoter DNAm residual overdispersion and includes known pluripotency markers such as *Dppa5a*, *Zfp42*, *Spp1* and *Peg3*. Representative examples for these genes are displayed in Fig. 4b.

These results suggest the presence of two modes of regulation. On one hand, down-regulation of pluripotency genes is associated with high promoter DNAm heterogeneity, linked to a pronounced increase in promoter DNA methylation throughout the embryonic stages. On the other hand, up-regulation of differentiation genes is not linked to high levels of promoter DNAm variability. This suggests that other genomic contexts or molecular layers might be responsible for their activation (Argelaguet *et al.*, 2019). Finally, we also find genes with low RNA expression variability that display high levels of promoter DNAm heterogeneity (Supplementary Fig. S17c), suggesting that the coupling between promoter DNAm and transcriptional activity is more complex than previously acknowledged during embryonic development stages (Anastasiadi *et al.*, 2018).

Figure 4: scMET applied to the multi-omics scNMT-seq gastrulation dataset reveals a complex linkage between promoter DNAm and RNA expression during embryonic development. (a) Scatter plot displays posterior median estimates for residual DNAm overdispersion parameters ϵ_j in gene promoters (x-axis) versus RNA residual overdispersion of the corresponding genes (y-axis). Among the genes with high levels of RNA heterogeneity, green and pink points correspond to promoters showing high and low levels of DNAm variability, respectively. (b) Representative examples of DNAm and RNA expression patterns across developmental stages for genes with high transcriptional heterogeneity and low (left, pink) or high (right, green) DNAm heterogeneity. Y-axis shows BASiCS log-normalised gene expression (in a $\log(x + 1)$ scale) (top) and promoter DNAm rate (bottom). Cells are stratified by embryonic stage (x-axis).

Discussion

Single-cell DNAm assays can currently profile hundreds to thousands of DNA methylomes, with increasingly complex experimental designs. The high resolution of these measurements enables us to measure cell-to-cell epigenetic variability, as well as uncover the regulatory features that modulate it (Eling *et al.*, 2019). However, the noise and biases intrinsic to such technologies create a need for computational frameworks that can systematically interrogate the data generated, dissecting genuine variability and quantifying uncertainties.

In this study we introduced scMET, a statistical framework for modelling DNA methylation heterogeneity from scBS-seq data. Using a hierarchical Bayesian framework to borrow information across cells and features, scMET robustly quantifies genuine cell-to-cell variability. Our results demonstrated the ability of scMET in highlighting genomic features that drive cell-to-cell heterogeneity across neuronal sub-populations in a large dataset of single-cell methylomes from the mouse frontal cortex. Furthermore, scMET can be used as a quantitative tool to interrogate changes in DNAm patterns between pre-specified cell populations. Unlike common approaches that only detect changes in mean methylation levels (Feng *et al.*, 2014; Hansen *et al.*, 2012), scMET can also identify features with differences in DNAm variability between populations. Importantly, the differential variability estimates are quantified through residual overdispersion parameters, thus accounting for the known confounding relationship between mean and overdispersion in scBS-seq datasets.

scMET uses a GLM framework to explicitly model known biases in the data in the form of additional covariates, such as CpG content. The flexibility of the GLM approach enables it to easily incorporate additional features, such as DNA motifs, which could be important to elucidate the role of sequence or chromatin state in modulating DNA methylation. Additionally, the framework could readily be extended to model joint variability in multiple molecular layers (such as transcriptome and methylome), opening a path to new methodologies in integrative, single-cell multi-omics analyses. Given the increasing prominence of such studies, we expect scMET to become an important tool in the extraction of biological signals from DNAm datasets of increasing complexity.

Methods

Standard statistical models for count data, such as the Poisson and binomial distributions, do not always capture the properties of data generated by high-throughput sequencing assays (e.g. RNA sequencing, bisulphite sequencing). In such cases, the data typically exhibit higher variance than what is predicted by these models — this is often referred to as *overdispersion* (Cox, 1983; Hinde *et al.*, 1998). This overdispersion may relate to *technical variation* (e.g. due to differences in sequencing depth) or to *biological variation* between the units of interest (e.g. cells or subjects) that is linked to genetic, environmental or other factors. Disentangling these sources of variation is a major challenge in computational biology.

The scMET model

Let Y_{ij} represent the number of methylated CpGs out of the n_{ij} CpGs for which DNAm was measured for genomic feature $j \in \{1, \dots, J\}$ in cell $i \in \{1, \dots, I\}$. These genomic features could be defined by pre-annotated regions (e.g. enhancers) or other regions of interest. To capture data overdispersion, scMET assumes a beta-binomial (BB) hierarchical formulation:

$$Y_{ij} | \theta_{ij} \sim \text{Binomial}(n_{ij}, \theta_{ij}), \quad \theta_{ij} | \mu_j, \gamma_j \sim \text{Beta}(\mu_j, \gamma_j). \quad (1)$$

In Eq. (1), the beta distribution is parameterised such that $\mathbb{E}[\theta_{ij}] = \mu_j$ and $\text{Var}[\theta_{ij}] = \gamma_j^2 \mu_j (1 - \mu_j (1 - \gamma_j) - \gamma_j) (1 - \gamma_j)^{-1}$, with $\mu_j \in (0, 1)$ and $\gamma_j \in (0, 1)$. If $\gamma_j = 0$, the model in Eq. (1) reduces to a binomial model with parameters n_{ij} and μ_j . After integrating out the random effects θ_{ij} , it can be seen that μ_j corresponds to the mean methylation across all cells for feature j and that γ_j controls the *overdispersion* that is not captured by binomial sampling. In fact, the BB variance can be written as:

$$\text{Var}[Y_{ij} | \mu_j, \gamma_j] = \underbrace{n_{ij} \mu_j (1 - \mu_j)}_{\text{technical variation}} + \underbrace{n_{ij} \mu_j (1 - \mu_j) (n_{ij} - 1) \gamma_j}_{\text{additional (biological) variation}}. \quad (2)$$

Parameters μ_j and γ_j can be inferred via maximum likelihood estimation. However, due to the high sparsity and noise present in single-cell DNAm data, these estimates can be unstable, especially for overdispersion parameters γ_j (Supplementary Fig. S5 and S6). To overcome this, we use a Bayesian framework with a hierarchical prior specification for μ_j and γ_j , sharing information across sets of similar types of genomic features (e.g. enhancers). Our approach is flexible and can incorporate feature-specific covariates $\mathbf{x}_j$ that explain differences in mean methylation across features. For instance, features with high CpG density tend to have lower methylation levels. These covariates are introduced within a generalized linear model (GLM) framework through the prior on the mean methylation parameters μ_j :

$$\mu_j | \mathbf{w}_\mu, s_\mu, \mathbf{x}_j \sim \text{Logit}\mathcal{N}(f_\mu(\mathbf{x}_j; \mathbf{w}_\mu), s_\mu), \quad \text{where } f_\mu(\mathbf{x}_j; \mathbf{w}_\mu) = \mathbf{w}_\mu^\top \mathbf{x}_j. \quad (3)$$

In Eq. (3), $\text{Logit}\mathcal{N}$ denotes a logit-normal distribution, $\mathbf{w}_\mu$ is a vector of regression coefficients and s_μ is the standard deviation for $\text{logit}(\mu_j)$. Throughout our analyses we assume $\mathbf{x}_j = (1, C_j)$, where C_j denotes the CpG density for feature j . However, scMET is flexible and users can introduce other feature-specific covariates.

Our prior specification is also designed to capture the mean-overdispersion relationship that is typically observed in the data generated by high-throughput sequencing assays, such as scBS-seq (Fig. 1c). Here, we follow the approach in Eling *et al.* (2018), introducing a non-linear regression model through an informative prior for γ_j :

$$\gamma_j | \mu_j, \mathbf{w}_\gamma, s_\gamma \sim \text{Logit}\mathcal{N}(f_\gamma(\mu_j; \mathbf{w}_\gamma), s_\gamma), \quad \text{where } f_\gamma(\mu_j; \mathbf{w}_\gamma) = w_{\gamma 1} + \sum_{l=2}^L w_{\gamma l} g_l(\mu_j). \quad (4)$$

Here, $f_\gamma(\mu_j; \mathbf{w})$ can be interpreted as the overdispersion (logit scale) that is predicted by mean methylation levels μ_j (fitted black line in Fig. 1c), $g_l(\mu_j)$ represent radial basis function kernels (defined as in [Kapourani and Sanguinetti, 2016](#)), and $w_{\gamma 1}, \dots, w_{\gamma L}$ are regression coefficients. Unless otherwise stated, we use $L = 4$ throughout our analyses. The remaining elements of the prior are described in Supplementary Note 2.1.

The prior distribution in Eq. (4) can be rewritten as a non-linear regression model

$$\text{logit}(\gamma_j) = f_\gamma(\mu_j; \mathbf{w}_\gamma) + \epsilon_j, \quad \epsilon_j \sim \mathcal{N}(0, s_\gamma), \quad (5)$$

where ϵ_j corresponds to a feature-specific *residual overdispersion* parameter that captures deviations from the overall trend. Hence, a feature that exhibits positive ϵ_j values has more variation than expected for features with similar mean methylation. Accordingly, negative ϵ_j values indicate less variation than expected for features with similar mean methylation.

Implementation

The posterior distribution for the model parameters in scMET is not amenable to analytic solutions. Hence, we resort to variational Bayes (VB, [Blei et al., 2017](#)) and Markov Chain Monte Carlo (MCMC, [Liang et al., 2010](#)) implementations using the Stan probabilistic programming language ([Carpenter et al., 2017](#)). scMET is publicly available as an R package at <https://github.com/andreaskapou/scMET> and will be shortly submitted to Bioconductor.

Identifying highly variable features

Residual overdispersion parameters ϵ_j can be used to label highly variable features (HVF) within a population of cells. Our decision rule is based on tail posterior probabilities ([Bochkina and Richardson, 2007](#)) associated to whether ϵ_j exceed a pre-specified threshold ϵ_0 :

$$\pi_j^E(\epsilon_0) \equiv P(\epsilon_j > \epsilon_0 \mid \text{data}), \quad (6)$$

As a default choice, we define ϵ_0 based on the distribution of posterior estimates for residual overdispersion parameters ϵ_j across all features. In particular, we define ϵ_0 to match the δ_E -th percentile of the distribution. Unless otherwise stated, we set as default $\delta_E = 0.9$.

The probabilities in Eq. (6) can be estimated by counting the proportion of posterior draws (obtained by VB or MCMC) for which the chosen criteria are met ([Lewin et al., 2006](#)). scMET labels as HVFs those for which their associated tail posterior probabilities are above a given posterior evidence threshold α_H ($0.6 < \alpha_H < 1$), where α_H is calibrated via the expected false discovery rate (EFDR; [Newton et al., 2004](#)), see also Supplementary Note S2.2.

Differential testing

scMET provides a similar probabilistic rule to label differentially methylated (DM) and differentially variable (DV) features across experimental conditions or cell types (Fig 1e). Here, we define DM features as those for which mean methylation varies across the groups of cells under study. More concretely, let μ_j^A and μ_j^B be the mean methylation parameters associated with feature j in groups A and B. We quantify differences in mean methylation as the log-odds ratio (LOR):

$$\text{LOR}(\mu_j^A, \mu_j^B) = \text{logit}(\mu_j^A) - \text{logit}(\mu_j^B). \quad (7)$$

Similar to the HVF analysis, our decision rule for DM testing is defined as:

$$\pi_{jAB}^M(\psi_M) \equiv P(|\text{LOR}(\mu_j^A, \mu_j^B)| > \psi_M \mid \text{data}) > \alpha_M, \quad (8)$$

where α_M ($0.6 < \alpha_M < 1$) is a posterior evidence threshold chosen to match a desired EFDR level and ψ_M is a LOR threshold which can be interpreted as a minimum effect size to be detected by the test. As default, we use $\psi_M = \log(2)$, i.e. a two fold change in odds ratio.

Beyond highlighting DM features, scMET embraces the cellular resolution of scBS-seq data to perform differential variability (DV) analyses, identifying changes in cell-to-cell DNAm variability across groups. Although overdispersion parameters γ_j could be used as the input for the DV test, the results would be confounded by the mean-overdispersion relationship that is typically observed within each genomic context (Fig 1c). Instead, we propose to perform DV analysis based on ϵ_j — a measure of cell-to-cell DNAm variability that is not confounded by differences in mean methylation. Let γ_j^A and γ_j^B denote the overdispersion parameters linked to feature j in groups A and B. To label DV features based on residual overdispersion, we make use of Eq. (5), and decompose the LOR between γ_j^A and γ_j^B parameters as:

$$\text{LOR}(\gamma_j^A, \gamma_j^B) = \underbrace{f_\mu^A(\mathbf{x}_j; \mathbf{w}_\mu^A) - f_\mu^B(\mathbf{x}_j; \mathbf{w}_\mu^B)}_{\text{mean contribution}} + \underbrace{\epsilon_j^A - \epsilon_j^B}_{\text{residual change}}. \quad (9)$$

In Eq. (9), the first term captures the changes in overdispersion that are explained by mean methylation and the second term captures residual overdispersion changes after accounting for the mean methylation. This residual change is used to identify features with statistically significant differences in residual overdispersion. For a given posterior evidence threshold α_E ($0.6 < \alpha_E < 1$) and tolerance threshold ψ_E , the following rule is used to identify DV features:

$$\pi_{jAB}^E(\psi_E) \equiv P(|\epsilon_j^A - \epsilon_j^B| > \psi_E \mid \text{data}) > \alpha_E. \quad (10)$$

As default we set $\psi_E = \log(1.5)$, i.e. 50% change in overdispersion LOR between the groups. As above, the posterior evidence threshold α_E is calibrated via the EFDR, see Supplementary Note 2.2.

Competing methods

The BB MLE method corresponds to estimating the parameters of the beta-binomial model in Eq. (1) independently per feature using maximum likelihood. The VGAM package was used for parameter estimation (Yee, 2015).

For HVF selection, two additional strategies were considered. The binomial model, where features are ranked according to binomial variance given by $1/I \sum_i \theta_{ij}(1 - \theta_{ij})$, where $\theta_{ij} = Y_{ij}/n_{ij}$ is the methylation rate for feature j in cell i . The Gaussian model on methylation rates $\theta_{ij} \sim \mathcal{N}(\mu_j, \sigma_j)$, where features are ranked according to σ_j .

For differential mean methylation testing on the synthetic datasets, we compared scMET with the Fisher’s exact test. Features with log-odds ratio $> \log(1.5)$ between the two groups and FDR $< 10\%$ (based on Benjamini-Hochberg procedure) were called as differentially methylated. Whilst methods have been proposed for DV testing using bulk methylation data (e.g. Phipson and Oshlack, 2014), to the best of our knowledge, scMET is the first DV approach tailored to single cell methylation data.

Simulation study

We simulated $J = 300$ features for varying number of cells ranging from $I = 20$ up to $I = 1000$. To mimic the properties observed in real scBS-seq data, we assume that for each feature we have coverage for a subset of cells given by $I_j \sim \text{Binomial}(I, p_j)$, where $p_j \sim \text{Uniform}(0.4, 0.8)$ to generate diverse I_j across features. We also simulate two alternative regions that have rich

($N = 50$) and poor ($N = 15$) CpG density. That is, the number of CpGs (n_{ij}) are simulated from $\text{Binomial}(N, q_j)$, where $q_j \sim \text{Uniform}(0.4, 0.8)$ to generate a broad range of CpG coverage across features. Next, for each feature we generate mean methylation parameters $\mu_j \sim \text{Logit}\mathcal{N}(\mathbf{w}_\mu^\top \mathbf{x}_j, 1)$, where $\mathbf{w}_\mu = (-0.5, -1.5)$ and $\mathbf{x}_j = (1, C_j)$ are feature-specific covariates, where C_j denotes the CpG density. The negative weight on $\mathbf{w}_\mu$ is used to simulate the known negative association between mean methylation and CpG density. Next, we simulated feature-specific overdispersion parameters $\gamma_j \sim \text{Logit}\mathcal{N}(\mathbf{w}_\gamma^\top \mathbf{g}_j(\mu_j), 0.25)$ to mimic the mean-overdispersion relationship. We set $\mathbf{w}_\gamma = (-1.2, -0.3, 1.1, -0.9)$ and $\mathbf{g}_j(\mu_j)$ is a vector of basis function values with methylation level μ_j . Finally, we simulated the number of methylated CpGs from BB distribution, using the VGAM package, that is, $Y_{ij} \sim \text{BB}(n_{ij}, \mu_j, \gamma_j)$.

For differential testing analysis, we used the above approach to generate cells from the first group (group A). For DM analysis, 15% of features were randomly selected and their corresponding μ_j were randomly increased or decreased by three different LOR thresholds: 2, 3, and 5, to generate cells from the second group (group B). Similarly, for DV analysis 15% of features were randomly selected from the first group and their corresponding γ_j were randomly increased or decreased by three different LOR thresholds: 2, 3, and 5, to generate cells from second group.

Mouse frontal cortex dataset

Data processing

The dataset is available from the Gene Expression Omnibus repository under accession number GSE97179. Details on quality control and data pre-processing can be found in [Luo et al. \(2017\)](#). Supplementary Table S2 contains metadata for the 3,069 cells, such as cell type annotations and cortical layer information. We aggregated closely related cellular sub-populations with less than 25 cells, following the hierarchy established in [Luo et al. \(2017\)](#). DNA methylation was quantified using mCG dinucleotides over three genomic contexts: (1) gene promoters ($\pm 2\text{kb}$ windows around the transcription start sites of genes extracted from ENSEMBL version 87 [Yates et al., 2016](#)), (2) Distal H3K27ac ChIP-seq peaks and (3) H3K4me1 ChIP-seq peaks. The latter was based on two ChIP-seq datasets that were profiled in adult (8 weeks) mouse cortex as part of the ENCODE project (see Supplementary Table S3).

For each genomic feature $j \in \{1, \dots, J\}$ and cell $i \in \{1, \dots, I\}$, the following censoring procedure was applied: we recorded Y_{ij} as a missing value if methylation coverage was available in less than 3 CpGs (i.e. $n_{ij} < 3$). The purpose of this censoring step was to exclude observations with very low coverage for which DNAm quantification is less reliable. Subsequently, we removed features that did not have CpG coverage in at least 15 cells. In addition, we excluded features that had mean methylation across cells lower than 0.1 or higher than 0.9; the rationale being that fully (un)methylated features do not drive methylation heterogeneity and will not provide information for identifying cell sub-populations.

Down-sampling experiment

Using the characterised sub-populations from [Luo et al. \(2017\)](#) we performed down-sampling experiments on 424 inhibitory neurons. Both scMET and BB MLE methods were run once on the full dataset (424 cells) to generate pseudo-ground truth parameter estimates. Subsequently, 20, 50, 100 and 200 cells were randomly down-sampled from the full population prior to parameter estimation. This procedure was repeated 5 times for each sample size. The same censoring step as described above was applied. Moreover, due to smaller sample sizes, we filtered genomic features that did not have CpG coverage in at least 5 cells.

HVF analysis

HVF analysis was applied on 12,774 gene promoters, 17,284 distal H3K27ac peaks, and 30,374 H3K4me1 peaks. To model the mean-overdispersion relationship we used $L = 4$ radial basis function kernels and kept default hyper-parameter values (Supplementary Note S2.1). The total number of iterations was set to 50,000 and convergence was attained when the evidence lower bound difference between two consecutive iterations was less than $1e-04$.

Differential analysis

For differential analysis between excitatory and inhibitory neurons, we only included features with CpG coverage in at least 15 cells, in both sub-populations. This resulted in 12,611 gene promoters, 13,075 distal H3K27ac peaks, and 20,212 H3K4me1 peaks. To model the mean-overdispersion relationship we used $L = 4$ radial basis function kernels and kept default hyper-parameter values (Supplementary Note S2.1). The total number of iterations was set to 50,000 and convergence was attained when the evidence lower bound difference between two consecutive iterations was less than $1e-04$.

Dimensionality reduction

Dimensionality reduction was applied using a Bayesian Factor Analysis algorithm, as implemented in the MOFA2 package (Argelaguet *et al.*, 2020). The motivation for this method, as opposed to the conventional Principal Component Analysis, is to handle the large presence of missing values without need for imputation. A second (non-linear) dimensionality reduction step was applied using UMAP (as implemented in the `umap` package) to project the data into a two-dimensional space (Supplementary Fig. S12 and S13).

Clustering

A finite grid of HVFs (from 50 to 1,000 with step size of 50) was selected by each of the competing methods. Subsequently, clustering analysis was performed using the k -means algorithm on the latent space defined by the MOFA factors (fixed to 15). The number of clusters was set to the number of cell types as characterised by Luo *et al.* (2017) (see also *Data processing* section). We assessed clustering performance using the ARI and cluster purity. A non-parametric regression (implemented by the `loess` function) was used to obtain a smoothed interpolation across all HVF values.

scNMT-seq gastrulation dataset

Data processing

The parsed scNMT-seq gastrulation dataset was downloaded from ftp://ftp.ebi.ac.uk/pub/databases/scnmt_gastrulation. Raw files are available from the Gene Expression Omnibus repository under accession number GSE121708. Details on the quality control and data processing can be found in Argelaguet *et al.* (2019). We selected all cells from E4.5 to E7.5 days after excluding the extra-embryonic visceral endoderm cells, as they display distinct DNA methylation profiles. Supplementary Table S4 contains sample metadata for the 848 cells retained for analysis. DNA methylation was quantified over gene promoters (± 2 kb windows around the transcription start sites of genes extracted from ENSEMBL version 87 Yates *et al.*, 2016).

Calculation of DNA methylation and RNA expression heterogeneity

For the DNAm data, we applied the same censoring procedure and feature exclusion criteria as described in the pre-processing of the [Luo *et al.* \(2017\)](#) dataset. This resulted in 13,785 gene promoters for downstream analysis. Residual overdispersion estimates were calculated by scMET with default parameter values using the same number of iterations and convergence criteria described above.

For the RNA expression data, we removed lowly expressed genes (no counts in less than 10 cells and average count across expressed cells less than 5). This resulted in 14,076 genes for downstream analysis. Residual overdispersion estimates were obtained using BASiCS ([Eling *et al.*, 2018](#)). The algorithm was run using 20,000 iterations, applying a burn in of 10,000 and thinning of 10. An empirical Bayes approach was used to derive the prior hyperparameters associated to gene-specific mean expression parameters within BASiCS.

In the comparison displayed in Fig 4, we focused on the 10,192 genes contained in the intersection of the lists obtained above.

Data availability

All datasets analysed in this article are publicly available in the cited references.

Code availability

scMET is implemented in an open-source and publicly available R package at <https://github.com/andreaskapou/scMET>. The code used to process and analyse the data is available at <https://github.com/andreaskapou/scMET-analysis>. The following software versions were used throughout our analyses: scMET (0.99.1), BASiCS (2.0.1), coda (0.19.3), MOFA2 (0.99.7), rstan (2.19.3) and VGAM (1.1.3).

Acknowledgements

We thank Oliver Stegle and John C. Marioni for constructive feedback during early stages of this project. C.A.K. is a cross-disciplinary post-doctoral fellow supported by funding from the University of Edinburgh and Medical Research Council (core grant to the MRC Institute of Genetics and Molecular Medicine). C.A.V. is a Chancellor's Fellow funded by the University of Edinburgh. R.A. was funded by the European Molecular Biology Laboratory (EMBL) international PhD programme.

Authors' contributions

All authors conceived and planned the methods and experiments. C.A.K. developed the statistical model and software implementation. C.A.K. and R.A. analysed the data and wrote the manuscript, with input from all authors. G.S. and C.A.V. supervised the study. All authors contributed to the interpretation of results and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Ethics approval and consent to participate

Ethical approval was not needed for this study.

References

- Anastasiadi, D., Esteve-Codina, A., and Piferrer, F. (2018). Consistent inverse correlation between DNA methylation of the first intron and gene expression across tissues and species. *Epigenetics & Chromatin*, **11**(1), 37.
- Angermueller, C., Lee, H. J., Reik, W., and Stegle, O. (2017). DeepCpG: accurate prediction of single-cell DNA methylation states using deep learning. *Genome Biology*, **18**(1), 67.
- Argelaguet, R., Clark, S. J., Mohammed, H., Stapel, L. C., Krueger, C., Kapourani, C. A., Imaz-Rosshandler, I., Lohoff, T., Xiang, Y., Hanna, C. W., Smallwood, S., Ibarra-Soria, X., Buettner, F., Sanguinetti, G., Xie, W., Krueger, F., Götting, B., Rugg-Gunn, P. J., Kelsey, G., Dean, W., Nichols, J., Stegle, O., Marioni, J. C., and Reik, W. (2019). Multi-omics profiling of mouse gastrulation at single-cell resolution. *Nature*, **576**(7787), 487–491.
- Argelaguet, R., Arnol, D., Bredikhin, D., Deloro, Y., Velten, B., Marioni, J. C., and Stegle, O. (2020). MOFA+: a statistical framework for comprehensive integration of multi-modal single-cell data. *Genome Biology*, **21**(1), 1–17.
- Avner, P. and Heard, E. (2001). X-chromosome inactivation: counting, choice and initiation. *Nature Reviews Genetics*, **2**(1), 59.
- Baylin, S. B. and Jones, P. a. (2011). A decade of exploring the cancer epigenome - biological and translational implications. *Nature Reviews Cancer*, **11**(10), 726–734.
- Becht, E., McInnes, L., Healy, J., Dutertre, C.-A., Kwok, I. W. H., Ng, L. G., Ginhoux, F., and Newell, E. W. (2019). Dimensionality reduction for visualizing single-cell data using UMAP. *Nature Biotechnology*, **37**(1), 38–44.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, **112**(518), 859–877.
- Bochkina, N. and Richardson, S. (2007). Tail posterior probability for inference in pairwise and multiclass gene expression data. *Biometrics*, **63**(4), 1117–1125.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, **76**(1).
- Challen, G. A., Sun, D., Jeong, M., Luo, M., Jelinek, J., Berg, J. S., Bock, C., Vasanthakumar, A., Gu, H., Xi, Y., Liang, S., Lu, Y., Darlington, G. J., Meissner, A., Issa, J.-P. J., Godley, L. A., Li, W., and Goodell, M. A. (2012). Dnmt3a is essential for hematopoietic stem cell differentiation. *Nature Genetics*, **44**(1), 23–31.
- Clark, S. J., Argelaguet, R., Kapourani, C. A., Stubbs, T. M., Lee, H. J., Alda-Catalinas, C., Krueger, F., Sanguinetti, G., Kelsey, G., Marioni, J. C., Stegle, O., and Reik, W. (2018). ScNMT-seq enables joint profiling of chromatin accessibility DNA methylation and transcription in single cells. *Nature Communications*, **9**(1), 1–9.
- Cox, D. R. (1983). Some remarks on overdispersion. *Biometrika*, **70**(1), 269–274.
- 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), 1–8.
- Eling, N., Richard, A. C., Richardson, S., Marioni, J. C., and Vallejos, C. A. (2018). Correcting the Mean-Variance Dependency for Differential Variability Testing Using Single-Cell RNA Sequencing Data. *Cell Systems*, **7**(3), 284–294.

- Eling, N., Morgan, M. D., and Marioni, J. C. (2019). Challenges in measuring and understanding biological noise. *Nature Reviews Genetics*, **20**(9), 536–548.
- Feng, H., Conneely, K. N., and Wu, H. (2014). A Bayesian hierarchical model to detect differentially methylated loci from single nucleotide resolution sequencing data. *Nucleic Acids Research*, **42**(8), e69.
- Georgala, P. A., Manuel, M., and Price, D. J. (2011). The generation of superficial cortical layers is regulated by levels of the transcription factor Pax6. *Cerebral Cortex*, **21**(1), 81–94.
- Gravina, S., Dong, X., Yu, B., and Vijg, J. (2016). Single-cell genome-wide bisulfite sequencing uncovers extensive heterogeneity in the mouse liver methylome. *Genome Biology*, **17**(150), 2–8.
- Guo, H., Zhu, P., Wu, X., Li, X., Wen, L., and Tang, F. (2013). Single-cell methylome landscapes of mouse embryonic stem cells and early embryos analyzed using reduced representation bisulfite sequencing. *Genome Research*, pages 2126–2135.
- Hansen, K. D., Langmead, B., and Irizarry, R. A. (2012). BSmooth: from whole genome bisulfite sequencing reads to differentially methylated regions. *Genome Biology*, **13**(10), R83.
- Hinde, J., Demétrio, C. G. B., and Others (1998). Overdispersion: models and estimation. *Computational Statistics and Data Analysis*, **27**(2), 151–170.
- Hubert, L. and Arabie, P. (1985). Comparing partitions. *Journal of Classification*, **2**(1), 193–218.
- Jaenisch, R. and Bird, A. (2003). Epigenetic regulation of gene expression: how the genome integrates intrinsic and environmental signals. *Nature Genetics*, **33**(March), 245–254.
- Kapourani, C. A. and Sanguinetti, G. (2016). Higher order methylation features for clustering and prediction in epigenomic studies. *Bioinformatics*, **32**(17), i405–i412.
- Kapourani, C.-A. and Sanguinetti, G. (2019). Melissa: Bayesian clustering and imputation of single-cell methylomes. *Genome Biology*, **20**(61), 1–15.
- Krueger, F., Kreck, B., Franke, A., and Andrews, S. R. (2012). DNA methylome analysis using short bisulfite sequencing data. *Nature Methods*, **9**(2), 145–151.
- Lähnemann, D., Köster, J., Szczurek, E., McCarthy, D. J., Hicks, S. C., Robinson, M. D., Valles, C. A., Campbell, K. R., Beerenwinkel, N., Mahfouz, A., and Others (2020). Eleven grand challenges in single-cell data science. *Genome Biology*, **21**(1), 1–35.
- Lewin, A., Richardson, S., Marshall, C., Glazier, A., and Aitman, T. (2006). Bayesian modeling of differential gene expression. *Biometrics*, **62**(1), 10–18.
- Liang, F., Liu, C., and Carroll, R. (2010). *Advanced Markov chain Monte Carlo methods: learning from past samples*. John Wiley & Sons.
- Lister, R., Mukamel, E. A., Nery, J. R., Urich, M., Puddifoot, C. A., Johnson, N. D., Lucero, J., Huang, Y., Dwork, A. J., Schultz, M. D., Yu, M., Tonti-Filippini, J., Heyn, H., Hu, S., Wu, J. C., Rao, A., Esteller, M., He, C., Haghghi, F. G., Sejnowski, T. J., Behrens, M. M., and Ecker, J. R. (2013). Global epigenomic reconfiguration during mammalian brain development. *Science*, **341**(6146), 629.
- 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**(12), 550.
- Luo, C., Keown, C. L., Kurihara, L., Zhou, J., He, Y., Li, J., Castanon, R., Lucero, J., Nery, J. R., Sandoval, J. P., Bui, B., Sejnowski, T. J., Harkins, T. T., Mukamel, E. A., Behrens, M. M., and Ecker, J. R. (2017). Single-cell methylomes identify neuronal subtypes and regulatory elements in mammalian cortex. *Science*, **357**(6351), 600–604.

- Mo, A., Mukamel, E. A., Davis, F. P., Luo, C., Henry, G. L., Picard, S., Urich, M. A., Nery, J. R., Sejnowski, T. J., Lister, R., Eddy, S. R., Ecker, J. R., and Nathans, J. (2015). Epigenomic Signatures of Neuronal Diversity in the Mammalian Brain. *Neuron*, **86**(6), 1369–1384.
- Mulqueen, R. M., Pokholok, D., Norberg, S. J., Torkenczy, K. A., Fields, A. J., Sun, D., Sinnamon, J. R., Shendure, J., Trapnell, C., O’Roak, B. J., Xia, Z., Steemers, F. J., and Adey, A. C. (2018). Highly scalable generation of DNA methylation profiles in single cells. *Nature Biotechnology*, **36**(5), 428—431.
- Newton, M. A., Noueiry, A., Sarkar, D., and Ahlquist, P. (2004). Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics*, **5**(2), 155–176.
- Phipson, B. and Oshlack, A. (2014). DiffVar: a new method for detecting differential variability with application to methylation in cancer and aging. *Genome Biology*, **15**(9), 465.
- Reik, W. and Walter, J. (2001). Genomic imprinting: parental influence on the genome. *Nature Reviews Genetics*, **2**(1), 21.
- Schwartzman, O. and Tanay, A. (2015). Single-cell epigenomics: techniques and emerging applications. *Nature Reviews Genetics*, **16**(12), 716–726.
- Smallwood, S. a., Lee, H. J., Angermueller, C., Krueger, F., Saadeh, H., Peat, J., Andrews, S. R., Stegle, O., Reik, W., and Kelsey, G. (2014). Single-cell genome-wide bisulfite sequencing for assessing epigenetic heterogeneity. *Nature Methods*, **11**(8), 817–20.
- Sorensen, S. A., Bernard, A., Menon, V., Royall, J. J., Glattfelder, K. J., Desta, T., Hirokawa, K., Mortrud, M., Miller, J. A., Zeng, H., and Others (2015). Correlated gene expression and target specificity demonstrate excitatory projection neuron diversity. *Cerebral Cortex*, **25**(2), 433–449.
- Yates, A., Akanni, W., Amode, M. R., Barrell, D., Billis, K., Carvalho-Silva, D., Cummins, C., Clapham, P., Fitzgerald, S., Gil, L., Girón, C. G., Gordon, L., Hourlier, T., Hunt, S. E., Janacek, S. H., Johnson, N., Juettemann, T., Keenan, S., Lavidas, I., Martin, F. J., Maurel, T., McLaren, W., Murphy, D. N., Nag, R., Nuhn, M., Parker, A., Patricio, M., Pignatelli, M., Rahtz, M., Riat, H. S., Sheppard, D., Taylor, K., Thormann, A., Vullo, A., Wilder, S. P., Zadissa, A., Birney, E., Harrow, J., Muffato, M., Perry, E., Ruffier, M., Spudich, G., Trevanion, S. J., Cunningham, F., Aken, B. L., Zerbino, D. R., and Flicek, P. (2016). Ensembl 2016. *Nucleic Acids Research*, **44**(D1), D710—D716.
- Yee, T. W. (2015). *Vector generalized linear and additive models: with an implementation in R*. Springer.
- Zahr, S. K., Yang, G., Kazan, H., Borrett, M. J., Yuzwa, S. A., Voronova, A., Kaplan, D. R., and Miller, F. D. (2018). A translational repression complex in developing mammalian neural stem cells that regulates neuronal specification. *Neuron*, **97**(3), 520–537.