

A Bayesian Binomial Regression Model with Latent Gaussian Processes for Modelling DNA Methylation

Aliaksandr Hubin	Geir O. Storvik	Paul E. Grini	Melinka A. Butenko
Norwegian Computing Center Dep. Mathematics University of Oslo	Dep. Mathematics University of Oslo Norwegian Computing Center	Dep. Biosciences University of Oslo	Dep. Biosciences University of Oslo

Abstract

Epigenetic observations are represented by the total number of reads from a given pool of cells and the number of methylated reads, making it reasonable to model this data by a binomial distribution. There are numerous factors that can influence the probability of success in a particular region. Moreover, there is a strong spatial (alongside the genome) dependence of these probabilities. We incorporate dependence on the covariates and the spatial dependence of the methylation probability for observations from a pool of cells by means of a binomial regression model with a latent Gaussian field and a logit link function. We apply a Bayesian approach including prior specifications on model configurations. We run a mode jumping Markov chain Monte Carlo algorithm (MJMCMC) across different choices of covariates in order to obtain the joint posterior distribution of parameters and models. This also allows finding the best set of covariates to model methylation probability within the genomic region of interest and individual marginal inclusion probabilities of the covariates.

Keywords: Bayesian binomial regression, latent Gaussian field, mode jumping Markov chain Monte Carlo, integrated nested Laplace approximations, Bayesian variable selection, Bayesian model averaging, epigenetic data, genetic patterns.

1. Introduction

Epigenetic modifications contribute to the generation of phenotypic plasticity, but the understanding of its contribution to phenotypic alterations and how the genome influences epigenetic variants requires further investigation (Schmitz, Schultz, Urich, Nery, Pelizzola, Libiger, Alix, McCosh, Chen, Schork *et al.* 2013). Epigenetic changes are crucial for the development and differentiation of various cell types in an organism, as well as for normal cellular processes. Epigenetic modifications modulate gene expression and modifications found in the promoter or regulatory elements play a prominent role in activating or suppressing transcript levels. This creates interesting research possibilities, which are often challenging from the statistical

point of view. For example, Li, Cassese, Guindani, and Vannucci (2019) suggested a Bayesian negative binomial regression model to study the influence of methylation (used as covariates) on RNA-Seq gene expression counts (used as observations). Also, Tang, Zhou, Wang, Huang, and Jin (2017) developed a Gaussian Bayesian regression model to link the differential gene expression (measured as \log_2 fold change) to various exogenous variables including tumour suppressor genes categories, mean methylation values and genomic segment distributions. In turn, Ma, Liu, Zhang, Huang, and Tang (2017) suggested a multiple network for epigenetic studies and implemented the Cox proportional hazard model to analyze the association of methylation profile of each epigenetic module with the patient survival. Recently, high-throughput epigenetics experiments have enabled researchers to measure genome-wide epigenetic profiles. This allows performing Epigenome-wide association studies (EWAS), which also hold promise for the detection of new regulatory mechanisms that may be susceptible to modification by environmental and lifestyle factors (Michels, Binder, Dedeurwaerder, Epstein, Greally, Gut, Houseman, Izzi, Kelsey, Meissner *et al.* 2013).

A major task today is the development of models and statistical methods for linking epigenetic patterns to genomic and/or environmental variables and interpreting them. Unlike the papers mentioned above (Ma *et al.* 2017; Tang *et al.* 2017; Li *et al.* 2019), we use methylation data as responses and link them to genomic and phenotypic variables (used as covariates). Moreover, by means of performing careful statistical modelling, our model takes into account that epigenetic data are spatially correlated (along locations in the genome) with high noise levels. Due to the availability of the data, our focus will be on the model plant *Arabidopsis thaliana*. For instance, Becker, Haggmann, Müller, Koenig, Stegle, Borgwardt, and Weigel (2011) previously analysed Arabidopsis data consisting of epigenetic observations on a set of 10 lines, which were separately propagated in a common environment for 30 generations. These were compared with two independent lines propagated for only three generations. Their analysis aimed at global summaries of structures but was based on individual and (site-wise) hypothesis testing methods combined with false discovery rate control methodology. In this paper, however, we limit ourselves to finding a pattern of signals appearing along the single genome that significantly influences the methylation probability of the corresponding organism. This is done by means of applying a binomial regression model with latent Gaussian variables, which take into account both spatial dependence and variability that can not be explained by the exogenous variables alone. The chosen latent Gaussian variable is a sum of a random walk $RW(1)$ component and an independent IG component. Model selection and parameter estimation within models are performed simultaneously in a Bayesian framework, applying the mode jumping Markov chain Monte Carlo (MJMCMC) algorithm developed by Hubin and Storvik (2018) to perform the computations involved. MJMCMC outputs posterior model probabilities allowing to find the best combination of explanatory genomic variables and compute marginal inclusion probabilities for the importance of individual variables. Our approach also allows to generate a model-averaged classification of the methylation status at different locations and make imputations for those locations that do not have enough observations, whilst the currently used approach is to simply ignore these locations.

2. Mathematical model

We model the number of methylated reads $Y_t \in \{1, \dots, n_t\}$ per position in the genome (nucleotide base position) to be binomially distributed with the number of trials equal to the number of reads, $n_t \in \mathbb{N}$, for position $t \in \{t_1, \dots, t_T\}$ (where T is the total number of genomic positions in the addressed genomic region) and corresponding probability of success $p_t \in \mathbb{R}_{[0,1]}$. The probability p_t is modeled via the logit link to the covariates $\mathbf{X}_t = \{X_{t1}, \dots, X_{td}\}$. These covariates might be a position within a gene (e.g promoter or coding region), an indicator of the underlying genetic structure, or other types of features (our choice of the covariates is given in Section 3). A latent Gaussian $RW(1)$ process $\delta_t \in \mathbb{R}$ and a latent independent Gaussian process (IG) ζ_t are included into the model in order to take into account spatial

dependence of methylation probabilities along the genome and the variance which is not explained by the covariates. Other explored latent Gaussian variables were also tested prior to variable selection on a full model before the selection of this structure was done (see the detail in the Appendix A of the paper). This gives the following model formulation:

$$\Pr(Y_t = y | n_t, p_t) = \binom{n_t}{y} p_t^y (1 - p_t)^{n_t - y}, \quad (1)$$

$$\text{logit}(p_t) = \beta_0 + \sum_{j=1}^d \gamma_j \beta_j X_{tj} + \delta_t + \zeta_t, \quad (2)$$

$$\delta_t = \delta_{t-1} + \epsilon_t, \quad \epsilon_t \stackrel{iid}{\sim} N(0, \tau_\epsilon^{-1}), \quad (3)$$

$$\zeta_t \stackrel{iid}{\sim} N(0, \tau_\zeta^{-1}), \quad (4)$$

where $\beta_j \in \mathbb{R}, j \in \{0, \dots, d\}$ are regression coefficients of the model showing whether and in which way the corresponding covariate influences the probability of methylation on average, $\gamma_j \in \{0, 1\}, i \in \{1, \dots, d\}$ are latent indicators, defining if covariate j is included into the model ($\gamma_j = 1$) or not ($\gamma_j = 0$), $\{\epsilon_t\}$ are the error terms of $RW(1)$ process $\{\delta_t\}$, which are normally distributed with zero mean and precision τ_ϵ^{-1} . Finally, τ_ζ^{-1} is the precision term of the IG process $\{\zeta_t\}$. We then put the following priors for the parameters of the model:

$$\gamma_i \sim \text{Bernoulli}(q), \beta_i | \gamma_i \sim I(\gamma_i = 1)N(0, \tau_\beta^{-1}), \tau_\beta, \tau_\epsilon \text{ and } \tau_\zeta \sim \text{Gamma}(1, 5 \cdot 10^{-5}). \quad (5)$$

Here, $q = 0.5$ is the prior Bernoulli probability of including a covariate into the model, and $I(\cdot)$ is the indicator function.

We perform analysis for the model defined by Equations (1)-(5) by means of the MJMCMC algorithm (Hubin and Storvik 2018). The algorithm is capable of efficiently moving in the defined model space by means of both accurately exploring the modes of the probability mass and switching between these modes using large jumps combined with local optimization and randomization.

3. Bayesian inference

Let $\gamma = (\gamma_1, \dots, \gamma_p)$, which uniquely defines a specific model. Assuming the constant term β_0 is always included, there are $L = 2^d$ different models to consider. Define $\theta = \{\beta, \psi, \tau_\beta, \tau_\epsilon, \tau_\zeta\} \in \Theta$, which describes the parameters of the models. Also let $\mathbf{Y} = (y_1, \dots, y_n)$ and $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$. We want to make inference jointly on models and their parameters $p(\gamma, \theta | \mathbf{Y}, \mathbf{X})$. We also want to find a set of the best models with respect to posterior marginal model probabilities $p(\gamma | \mathbf{Y}, \mathbf{X})$. Finally, we want to obtain the marginal inclusion probabilities $p(\gamma_j = 1 | \mathbf{Y}, \mathbf{X}), j \in \{1, \dots, d\}$ for individual covariates.

By Bayes formula, $p(\gamma, \theta | \mathbf{Y}, \mathbf{X}) = p(\theta | \gamma, \mathbf{Y}, \mathbf{X})p(\gamma | \mathbf{Y}, \mathbf{X})$. In Section 3.1 we describe how to compute $p(\theta | \gamma, \mathbf{Y}, \mathbf{X})$ and $p(\mathbf{Y} | \mathbf{X}, \gamma)$. For the time being, we assume that marginal likelihoods $p(\mathbf{Y} | \mathbf{X}, \gamma)$ are available for a given γ . Then by Bayes formula:

$$p(\gamma | \mathbf{Y}, \mathbf{X}) = \frac{p(\mathbf{Y} | \mathbf{X}, \gamma)p(\gamma)}{\sum_{\gamma' \in \Omega} p(\mathbf{Y} | \mathbf{X}, \gamma')p(\gamma')}. \quad (6)$$

In order to calculate $p(\gamma | \mathbf{Y}, \mathbf{X})$ we have to iterate through the whole model space Ω , which becomes computationally infeasible for large d . The ordinary MCMC estimate is based on a number of MCMC samples $\gamma^{(i)}, i = 1, \dots, W$:

$$\tilde{p}(\gamma | \mathbf{Y}, \mathbf{X}) = W^{-1} \sum_{i=1}^W I(\gamma^{(i)} = \gamma) \xrightarrow{W \rightarrow \infty} p(\gamma | \mathbf{Y}, \mathbf{X}). \quad (7)$$

An alternative, named the renormalized model (RM) estimates by [Clyde, Ghosh, and Littman \(2011\)](#), is

$$\hat{p}(\gamma|\mathbf{Y}, \mathbf{X}) = \frac{p(\mathbf{Y}|\mathbf{X}, \gamma)p(\gamma)}{\sum_{\gamma' \in \mathbb{V}} p(\mathbf{Y}|\mathbf{X}, \gamma')p(\gamma')} \mathbf{I}(\gamma \in \mathbb{V}) \xrightarrow[\mathbb{V} \rightarrow \Omega]{d} p(\gamma|\mathbf{Y}, \mathbf{X}), \quad (8)$$

where now \mathbb{V} is the set of visited models during the MCMC (or any other model space exploration algorithm) run. Although both (8) and (7) are asymptotically consistent, (8) will often be a preferable estimator since convergence of the MCMC based approximation (7) is much slower, see [Clyde *et al.* \(2011\)](#); [Hubin and Storvik \(2018\)](#).

We aim at approximating $p(\gamma|\mathbf{Y}, \mathbf{X})$ by means of searching for some subspace \mathbb{V} of Ω making the approximation (8) as precise as possible. Models with high values of $p(\mathbf{Y}|\mathbf{X}, \gamma)p(\gamma)$ and regions of relatively high posterior mass are important to be included into \mathbb{V} . Missing them in \mathbb{V} can introduce significant biases in our estimates. Note that these aspects are just as important for the standard MCMC estimate (7). The difference is that while the number of times a specific model is visited is important when using (7), for (8) it is enough that the model is visited at least once. In this context, the denominator of (8), which we would like to be as high as possible, becomes an extremely relevant measure for the quality of the search in terms of being able to capture whether the algorithm visits all of the modes, whilst the size of \mathbb{V} should be low in order to save computational time.

The marginal inclusion probability $p(\gamma_j = 1|\mathbf{Y}, \mathbf{X})$ is defined as:

$$p(\gamma_j = 1|\mathbf{Y}, \mathbf{X}) = \sum_{\gamma' \in \Omega} \mathbf{I}(\gamma'_j = 1)p(\gamma'|\mathbf{Y}, \mathbf{X}). \quad (9)$$

It can be approximated either by the MCMC estimator:

$$\tilde{p}(\gamma_j = 1|\mathbf{Y}, \mathbf{X}) = W^{-1} \sum_{i=1}^W \mathbf{I}(\gamma_j^{(i)} = 1) \xrightarrow[W \rightarrow \infty]{d} p(\gamma_j = 1|\mathbf{Y}, \mathbf{X}), \quad (10)$$

or using the renormalized approach:

$$\hat{p}(\gamma_j = 1|\mathbf{Y}, \mathbf{X}) = \sum_{\gamma' \in \mathbb{V}} \mathbf{I}(\gamma'_j = 1)p(\gamma'|\mathbf{Y}, \mathbf{X}) \xrightarrow[\mathbb{V} \rightarrow \Omega]{d} p(\gamma_j = 1|\mathbf{Y}, \mathbf{X}), \quad (11)$$

giving a measure of importance for the covariates of the model.

3.1. Integrated nested Laplace approximations

Within hierarchical models with latent Gaussian structures, integrated nested Laplace approximations (INLA) for efficient inference on the posterior distribution ([Rue, Martino, and Chopin 2009](#)) can be used. Following the INLA terminology, we define \mathbf{z} to be the set of latent Gaussian variables and the regression parameters $\boldsymbol{\beta}$ while $\boldsymbol{\eta}$ contains the remaining parameters (a low-dimensional vector). The INLA approach is based on two steps. First the marginal posterior of the hyperparameters is approximated by

$$p(\boldsymbol{\eta}|\mathbf{Y}, \mathbf{X}, \gamma) \propto \frac{p(\mathbf{z}, \boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \gamma)}{p(\mathbf{z}|\boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \gamma)} = \frac{p(\mathbf{z}, \boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \gamma)}{\tilde{p}_G(\mathbf{z}|\boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \gamma)} \Big|_{\mathbf{z}=\mathbf{z}^*(\boldsymbol{\eta})} + \mathcal{O}(T^{-3/2}). \quad (12)$$

Here, $\tilde{p}_G(\mathbf{z}|\boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \gamma)$ is the Gaussian approximation of $p(\mathbf{z}|\boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \gamma)$, and $\mathbf{z}^*(\boldsymbol{\eta})$ is the mode of the distribution $p(\mathbf{z}|\boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \gamma)$. The posterior mode of the hyperparameters is found by maximizing the corresponding Laplace approximation using some gradient descent method (like for example the Newton-Raphson routine). Then an area with a relatively high posterior density of the hyperparameters is explored with either some grid-based procedure or variational Bayes.

The second step involves the approximation of the latent variables for every set of the explored hyperparameters. Here, the computational complexity of the approximation depends on the likelihood type for the data $\mathbf{Y}|\mathbf{X}$. If it is Gaussian, the posterior of the latent variables is Gaussian, and the approximation is exact and fully tractable. In the case the likelihood is skewed or heavy tails are present, a Gaussian approximation of the latent variables tends to become inaccurate and another Laplace approximation should be used:

$$\tilde{p}_{\text{LA}}(z_i|\boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\gamma}) \propto \frac{p(\mathbf{z}, \boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\gamma})}{\tilde{p}_{\text{GG}}(\mathbf{z}_{-i}|z_i, \boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\gamma})} \Big|_{\mathbf{z}_{-i}=\mathbf{z}_{-i}^*(z_i, \boldsymbol{\eta})}. \quad (13)$$

Here, \tilde{p}_{GG} is the Gaussian approximation to $p(\mathbf{z}_{-i}|z_i, \boldsymbol{\eta}, \mathbf{Y}, \mathbf{X}, \boldsymbol{\gamma})$ and $\mathbf{z}_{-i}^*(z_i, \boldsymbol{\eta})$ is the corresponding posterior mode. The error rate of (13) is $\mathcal{O}(T^{-3/2})$. The full Laplace approximation of the latent fields defined in equation (13) is rather time-consuming, hence more crude lower-order Laplace approximations are often used instead (typically increasing the error rate to $\mathcal{O}(T^{-1})$, Tierney and Kadane 1986). Once the posterior distribution of the latent variables given the hyperparameters is approximated, the uncertainty in the hyperparameters can be marginalized out using the law of total probability (Rue *et al.* 2009):

$$\tilde{p}(z_i|\mathbf{Y}, \mathbf{X}, \boldsymbol{\gamma}) = \sum_k \tilde{p}_{\text{LA}}(z_i|\boldsymbol{\eta}_k, \mathbf{Y}, \mathbf{X}, \boldsymbol{\gamma}) \tilde{p}(\boldsymbol{\eta}_k|\mathbf{Y}, \mathbf{X}, \boldsymbol{\gamma}) \Delta_k, \quad (14)$$

where Δ_k is the area weight corresponding to the grid exploration of the posterior distribution of the hyperparameters.

Computing the marginal likelihood

The marginal likelihood is defined as follows: For data $\{\mathbf{Y}, \mathbf{X}\}$ and model $\boldsymbol{\gamma}$, which includes some unknown parameters $\boldsymbol{\theta}$, the marginal likelihood is given by

$$p(\mathbf{Y}|\mathbf{X}, \boldsymbol{\gamma}) = \int_{\Theta} p(\mathbf{Y}|\mathbf{X}, \boldsymbol{\theta}, \boldsymbol{\gamma}) p(\boldsymbol{\theta}|\boldsymbol{\gamma}) d\boldsymbol{\theta} \quad (15)$$

where $p(\boldsymbol{\theta}|\boldsymbol{\gamma})$ is the prior for $\boldsymbol{\theta}$ under model $\boldsymbol{\gamma}$ while $p(\mathbf{Y}|\mathbf{X}, \boldsymbol{\gamma}, \boldsymbol{\theta})$ is the likelihood function conditional on $\boldsymbol{\theta}$. Again, consider $\boldsymbol{\theta} = (\boldsymbol{\eta}, \mathbf{z})$.

INLA approximates marginal likelihoods by

$$\tilde{p}(\mathbf{Y}|\mathbf{X}, \boldsymbol{\gamma}) = \int_{\mathbf{Z}} \frac{p(\mathbf{Y}, \mathbf{z}, \boldsymbol{\eta}|\mathbf{X}, \boldsymbol{\gamma})}{\tilde{\pi}_G(\boldsymbol{\eta}|\mathbf{Y}, \mathbf{X}, \mathbf{z}, \boldsymbol{\gamma})} \Big|_{\boldsymbol{\eta}=\boldsymbol{\eta}^*(\mathbf{z}|\boldsymbol{\gamma})} d\mathbf{z}, \quad (16)$$

where $\boldsymbol{\eta}^*(\mathbf{z}|\boldsymbol{\gamma})$ is some chosen value of $\boldsymbol{\eta}$, typically the posterior mode, while $\tilde{\pi}_G(\boldsymbol{\eta}|\mathbf{Y}, \mathbf{X}, \mathbf{z}, \boldsymbol{\gamma})$ is a Gaussian approximation to $\pi(\boldsymbol{\eta}|\mathbf{Y}, \mathbf{X}, \mathbf{z}, \boldsymbol{\gamma})$. The integration of \mathbf{z} over the support \mathbf{Z} can be performed by an empirical Bayes (EB) approximation or using numerical integration based on a central composite design (CCD) or a grid (see Rue *et al.* 2009, for details).

A toy example on computing the marginal likelihood. Consider an example from Neal (2008), in which we assume the following model $\boldsymbol{\gamma}$:

$$Y|z, \boldsymbol{\gamma} \sim N(z, \tau_1^{-1}); \quad z|\boldsymbol{\gamma} \sim N(0, \tau_0^{-1}). \quad (17)$$

Then obviously the marginal likelihood is available analytically as

$$Y|\boldsymbol{\gamma} \sim N(0, \tau_0^{-1} + \tau_1^{-1}), \quad (18)$$

and we have a benchmark to compare approximations to. The harmonic mean estimator (Raftery, Newton, Satagopan, and Krivitsky 2006) is given by

$$\tilde{p}(Y|\boldsymbol{\gamma}) = \frac{W}{\sum_{i=1}^W \frac{1}{p(Y|z_i, \boldsymbol{\gamma})}}$$

where $z_i \sim p(z|Y, \gamma)$. This estimator is consistent, however, often requires unreasonably many iterations to converge. We performed the experiments with $\tau_1 = 1$ and τ_0 being either 0.001, 0.1 or 10. The harmonic mean is obtained based on $W = 10^7$ simulations. 5 runs of the harmonic mean procedure are performed for each scenario. For INLA we used the default tuning parameters from the package (Rue *et al.* 2009). As one can see from Table 1, INLA gives

Table 1: Comparison of INLA, harmonic mean and exact marginal likelihood

τ_0	τ_1	T	Exact	INLA	H.mean				
0.001	1	2	-7.8267	-7.8267	-2.4442	-2.4302	-2.5365	-2.4154	-2.4365
0.1	1	2	-3.2463	-3.2463	-2.3130	-2.3248	-2.5177	-2.4193	-2.3960
10	1	2	-2.9041	-2.9041	-2.9041	-2.9041	-2.9042	-2.9041	-2.9042

extremely precise results even for a huge variance of the latent variable, whilst the harmonic mean can often become extremely crude even for 10^7 iterations. More examples showing the accuracy of INLA are summarized in Hubin and Storvik (2016) and Friel and Wyse (2012), which perform a comparison of a number of approaches to computing the marginal likelihood, including Laplace approximations, harmonic mean approximations, Chib’s method, Chib and Jelizkov’s method and INLA. The studies show that INLA is a fast method that enjoys giving precise approximations to the marginal likelihood even if the number of samples T is limited.

3.2. Mode jumping Markov chain Monte Carlo

The main problem with the standard Metropolis-Hastings algorithms is the trade-off between possibilities of large jumps (by which we understand proposals with a large neighbourhood) and high acceptance probabilities. Large jumps will typically result in proposals with low probabilities. In a continuous setting, Tjelmeland and Hegstad (1999) solved this by introducing local optimization after large jumps, which results in proposals with higher acceptance probabilities. Hubin and Storvik (2018) adopted this approach to the discrete model selection setting and suggested the following algorithm:

Algorithm 1 Mode jumping MCMC

- 1: Generate a large jump $\boldsymbol{\chi}_0^*$ according to a proposal distribution $q_l(\boldsymbol{\chi}_0^*|\boldsymbol{\gamma})$.
- 2: Perform a local optimization, defined through $\boldsymbol{\chi}_k^* \sim q_o(\boldsymbol{\chi}_k^*|\boldsymbol{\chi}_0^*)$.
- 3: Perform small randomization to generate the proposal $\boldsymbol{\gamma}^* \sim q_r(\boldsymbol{\gamma}^*|\boldsymbol{\chi}_k^*)$.
- 4: Generate backwards auxiliary variables $\boldsymbol{\chi}_0 \sim q_l(\boldsymbol{\chi}_0|\boldsymbol{\gamma}^*)$, $\boldsymbol{\chi}_k \sim q_o(\boldsymbol{\chi}_k|\boldsymbol{\chi}_0)$.
- 5: Put

$$\boldsymbol{\gamma}' = \begin{cases} \boldsymbol{\gamma}^* & \text{with probability } r_{mh}(\boldsymbol{\gamma}, \boldsymbol{\gamma}^*; \boldsymbol{\chi}_k, \boldsymbol{\chi}_k^*); \\ \boldsymbol{\gamma} & \text{otherwise,} \end{cases}$$

where

$$r_{mh}^*(\boldsymbol{\gamma}, \boldsymbol{\gamma}^*; \boldsymbol{\chi}_k, \boldsymbol{\chi}_k^*) = \min \left\{ 1, \frac{\pi(\boldsymbol{\gamma}^*)q_r(\boldsymbol{\gamma}|\boldsymbol{\chi}_k)}{\pi(\boldsymbol{\gamma})q_r(\boldsymbol{\gamma}^*|\boldsymbol{\chi}_k^*)} \right\}. \quad (19)$$

Here, a large jump corresponds to changing a large number of γ_j ’s while the local optimization will be some iterative procedure based on, at each iteration, changing a small number of components until a local mode is reached. For this algorithm, three proposals need to be specified; $q_l(\cdot|\cdot)$ specifying the first large jump, $q_o(\cdot|\cdot)$ specifying the local optimizer, and $q_r(\cdot|\cdot)$ specifying the last randomization. All of them are described in detail in Hubin and Storvik (2018). The convergence of the MJMCMC procedures is shown in Theorem 1 in Hubin and Storvik (2018).

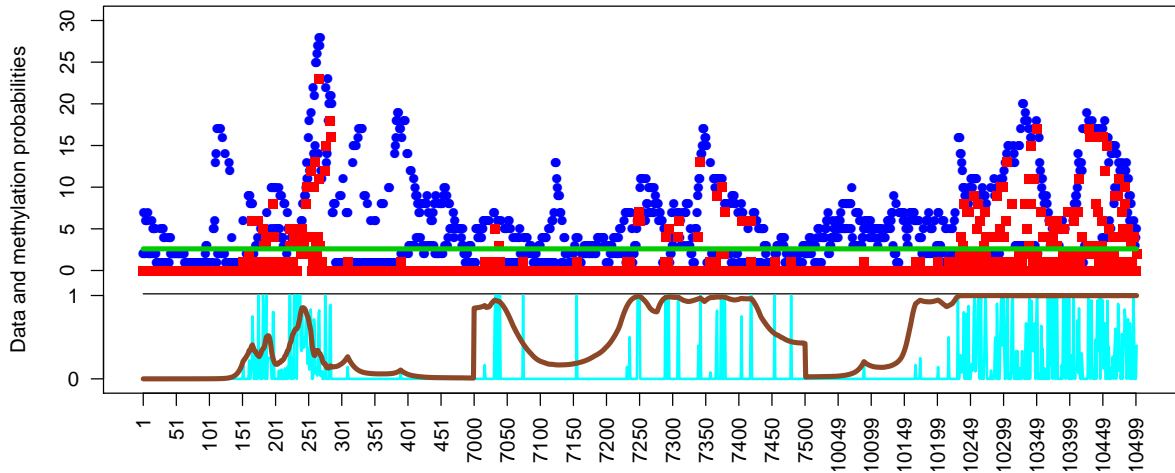


Figure 1: Epigenetic observations, where blue dots are total number of reads, red dots - number of methylated reads, the green line corresponds to 2 total reads distinguishing the inference and the identification data, light blue line gives naïve probabilities as rates, brown line - probabilities as the posterior mean of the probability of success parameter from the posterior mode model.

4. Data description

The addressed data set consists of 1502 observations from chromosome one of the Arabidopsis plant belonging to five predefined groups of genes. This data set was divided into 950 observations (with more than 2 reads, see Figure 1) for inference and 552 observations (with less than 3 reads) for model-based identification of methylation probabilities for the positions with the lack of data.

Apart from the observations represented by the methylated versus the total number of reads we have data on various exogenous variables (covariates). Among these covariates, we address a factor with 3 levels corresponding to whether the location belongs to a CGH, CHH or CHG genetic region, where H is either A, C or T and thus generating two covariates X_{CGH} and X_{CHH} . The second group of factors indicates whether the distance to the previous cytosine nucleobase (C) in DNA is 1, 2, 3, 4, 5, from 6 to 20 or greater than 20 inducing six binary covariates $X_{DT1}, X_{DT2}, X_{DT3}, X_{DT4}, X_{DT5}$, and $X_{DT6:20}$. We also include such 1D distance as a continuous covariate X_{DIST} . The third addressed group of factors corresponds to whether the location belongs to a gene from a particular group of genes of biological interest. These groups are indicated as M_a, M_g and M_d , yielding two additional covariates X_{M_a}, X_{M_g} . Additionally, we have a covariate X_{CODE} indicating if the corresponding nucleobase is in the coding region of a gene and a covariate X_{STRD} indicating if the nucleobase is on a ” + ” or a ” - ” strand. Finally, we have a continuous covariate $X_{EXPR} \in \mathbb{R}^+$ representing expression level for the corresponding gene and interactions between expression levels and gene groups $X_{EXPR,a}, X_{EXPR,g}, X_{EXPR,d} \in \mathbb{R}^+$. Thus multiple predictors with respect to a strict choice of the reference levels for categorical variables in our example induced $d = 17$ potentially important covariates. The correlation structure for the addressed variables is depicted in Figure 2, where one clearly sees multiple correlations, which in turn are likely to induce multiple modes of the marginal likelihood.

5. Results

The MJMCMC algorithm was run until around 10 000 unique models (7.6% of the model space) were explored. We parallelized the search on 10 CPUs. Default tuning parameters from Hubin and Storvik (2018) were used.




Figure 2: Heat-map correlation plot between the addressed covariates.

According to the marginal inclusion probabilities reported in Figure 3, only three factors X_{CHG} , X_{CGH} and X_{CODE} are clearly significant for inference on the methylation patterns for the addressed epigenetic region, factors X_{M_a} and X_{M_g} also have some significance. Table 2 gives the marginal posterior model probability and posterior means of the parameters for the best model in the explored subset of models from the model space. This model is both the posterior mode model in the set of explored models and the median probability model (Barbieri and Berger 2004). We have also compared the selected model with alternative models based on the optimal sets of covariates but with other latent Gaussian structures and found our model to be the best in terms of the marginal log likelihood (see the Appendix A of the paper).




Figure 3: Barplots of RM estimates (Hubin and Storvik 2018) of marginal inclusion probabilities of the covariates.

Table 2: Posterior means for the best model in terms of marginal posterior probability (PMP).

PMP	β_0	β_{CHG}	β_{CGH}	β_{CODE}	τ_ϵ	τ_ζ
0.4276	-8.8255	2.4717	5.2122	6.4240	7.5075	1.2109

Based on the best model, we carried out computations of methylation probabilities for the locations in both the inference set and the identification set. Highly methylated regions are located between observations 7000-7050, 7250-7400, and 10150-10500, see Figure 1. Note that the model is quite sceptical to the methylation status of locations 7051 - 7249, despite a number of observations with a high proportion of methylated reads in this region. Further-

more, we compared the results with the naïve approach based on computing the proportion of methylated reads (light blue line in Figure 1), which is currently addressed in the biological literature as a standard way to evaluate methylation probability of a given nucleobase. The results show that the naïve approach should not be trusted in the presence of spatially correlated data and the probabilities corresponding to it can be strongly biased.

6. Discussion

During cancer development, the changes in DNA methylation patterns occur within the gene promoter, CpG islands and their shores (Tang *et al.* 2017). Hence in future, it would be of interest to obtain additional covariates such as whether the corresponding nucleotide base position belongs to a particular part of the non-coding gene region like a promoter, an intron or remnants of transposable elements, and whether the nucleobase is within a CpG island, and see how these covariates are influencing the underlying methylation patterns. At the same time, in this paper we looked only at a subset of the genomic locations associated with the groups of genes of biological interest, however, in the future, it would be of interest to address the whole genome. That would induce working with extremely large data, which in turn creates new methodological challenges. In particular, in order to make the efficient inference, it will be important to allow sub-sampling for INLA within a given model and MJMCMC in the discrete marginal space of models. It will also be of interest to allow logical expressions for the binary covariates to be included in the model following Hubin, Storvik, and Frommlet (2018). Finally, joint inference on the covariates and various latent Gaussian variables by means of MJMCMC can be of interest.

Acknowledgements

The authors would like to thank the Center for Evolutionary Life Sciences (CELS) at the University of Oslo and BBPUBL19 project at the Norwegian Computing Center for fully financially supporting this work. Additionally, we thank Dr Olav Nikolai Breivik (Norwegian Computing Center) for thoroughly proofreading the article.

References

- Barbieri MM, Berger JO (2004). “Optimal Predictive Model Selection.” *The Annals of Statistics*, **32**(3), 870–897.
- Becker C, Hagmann J, Müller J, Koenig D, Stegle O, Borgwardt K, Weigel D (2011). “Spontaneous Epigenetic Variation in the Arabidopsis Thaliana Methylome.” *Nature*, **480**(7376), 245.
- Clyde MA, Ghosh J, Littman ML (2011). “Bayesian Adaptive Sampling for Variable Selection and Model Averaging.” *Journal of Computational and Graphical Statistics*, **20**(1), 80–101.
- Friel N, Wyse J (2012). “Estimating the Evidence - A Review.” *Statistica Neerlandica*, **66**(3), 288–308. ISSN 1467-9574.
- Hubin A, Storvik G (2016). “Estimating the Marginal Likelihood with Integrated Nested Laplace Approximation (INLA).”
- Hubin A, Storvik G (2018). “Mode Jumping MCMC for Bayesian Variable Selection in GLMM.” *Computational Statistics & Data Analysis*, **127**, 281–297.
- Hubin A, Storvik G, Frommlet F (2018). “A Novel Algorithmic Approach to Bayesian Logic Regression.” *Bayesian Analysis*. To be printed as a discussion paper in 2020.

- Li Q, Cassese A, Guindani M, Vannucci M (2019). “Bayesian Negative Binomial Mixture Regression Models for the Analysis of Sequence Count and Methylation Data.” *Biometrics*, **75**(1), 183–192.
- Ma X, Liu Z, Zhang Z, Huang X, Tang W (2017). “Multiple Network Algorithm for Epigenetic Modules via the Integration of Genome-wide DNA Methylation and Gene Expression Data.” *BMC Bioinformatics*, **18**(1), 72.
- Michels KB, Binder AM, Dedeurwaerder S, Epstein CB, Grealley JM, Gut I, Houseman EA, Izzi B, Kelsey KT, Meissner A, *et al.* (2013). “Recommendations for the Design and Analysis of Epigenome-wide Association Studies.” *Nature Methods*, **10**(10), 949.
- Neal R (2008). “The Harmonic Mean of the Likelihood: Worst Monte Carlo Method Ever.”
- Raftery AE, Newton MA, Satagopan JM, Krivitsky PN (2006). “Estimating the Integrated Likelihood via Posterior Simulation Using the Harmonic Mean Identity.” *Technical report*, Memorial Sloan-Kettering Cancer Center, Dept. of Epidemiology & Biostatistics. Working Paper Series.
- Rue H, Martino S, Chopin N (2009). “Approximate Bayesian Inference for Latent Gaussian Models by Using Integrated Nested Laplace Approximations.” *Journal of the Royal Statistical Society*, **71**(2), 319–392.
- Schmitz RJ, Schultz MD, Urich MA, Nery JR, Pelizzola M, Libiger O, Alix A, McCosh RB, Chen H, Schork NJ, *et al.* (2013). “Patterns of Population Epigenomic Diversity.” *Nature*, **495**(7440), 193.
- Tang B, Zhou Y, Wang CM, Huang THM, Jin VX (2017). “Integration of DNA Methylation and Gene Transcription across Nineteen Cell Types Reveals Cell Type-specific and Genomic Region-dependent Regulatory Patterns.” *Scientific Reports*, **7**(1), 3626.
- Tierney L, Kadane JB (1986). “Accurate Approximations for Posterior Moments and Marginal Densities.” *Journal of the American Statistical Association*, **81**(393), 82–86.
- Tjelmeland H, Hegstad BK (1999). “Mode Jumping Proposals in MCMC.” *Scandinavian Journal of Statistics*, **28**, 205–223.

Appendix

Alternative latent models

We also looked at several alternative latent Gaussian variables. These are reported in Table 3 together with the marginal likelihoods (Bayes factors can be computed in a straight-forward fashion). The default priors for the parameters from INLA (Rue *et al.* 2009) were used. Here, we clearly see that the chosen structure of the latent Gaussian field is significantly outperforming all other alternatives for both the null and the full models, which gives strong evidence to our choice. This is also true for the model selected by MJMCMC, which is reported in Table 2.

Table 3: Estimates of the marginal log likelihood for different candidate latent processes, which are Random Walk ($RW(1)$), Ornstein–Uhlenbeck (OU) and auto-regressive processes $AR(p)$ of order p from 1 to 3. Here $FULL$ is the model with all covariates included, $NULL$ is the model without covariates, and $BEST$ is the model with the covariates from Table 2.

$IG +$	$RW(1)$	OU	$AR(1)$	$AR(2)$	$AR(3)$
$FULL$	-771.7173	-1103.8190	-1116.8367	-1103.8188	-811.2963
$NULL$	-875.4917	-876.5565	-876.8325	-876.5565	-875.4917
$BEST$	-591.0837	-706.8185	-713.7015	-706.8185	-697.3347

Affiliation:

Aliaksandr Hubin

Department of Statistical Analysis, Machine Learning and Image Analysis

Norwegian Computing Center

P.O. Box 114 Blindern, NO-0314 Oslo, Norway

E-mail: aliaksandr.hubin@nr.no

URL: <https://www.nr.no/?q=publicationprofile&query=ahu/>