



HAL
open science

Parameter estimation of hidden Markov models: comparison of EM and quasi-Newton methods with a new hybrid algorithm.

Sidonie Foulon, Thérèse Truong, Anne-Louise Leutenegger, Hervé Perdry

► To cite this version:

Sidonie Foulon, Thérèse Truong, Anne-Louise Leutenegger, Hervé Perdry. Parameter estimation of hidden Markov models: comparison of EM and quasi-Newton methods with a new hybrid algorithm.. 2024. hal-04685772

HAL Id: hal-04685772

<https://hal.science/hal-04685772v1>

Preprint submitted on 3 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Parameter estimation of hidden Markov models: comparison of EM and quasi-Newton methods with a new hybrid algorithm.

Sidonie Foulon^{1,2,3,*}, Thérèse Truong^{1,2}, Anne-Louise Leutenegger³, and Hervé Perdry¹

¹Université Paris-Saclay, UVSQ, Inserm, CESP, 16 av PV Couturier, 94807, Villejuif, France

²Institut Gustave Roussy, 114 rue Edouard-Vaillant, 94805, Villejuif, France

³NeuroDiderot, Inserm, Université Paris Cité, UMR1141, 48 bd Sérurier, 75019, Paris, France

*Corresponding author: Sidonie Foulon, sidonie.foulon@inserm.fr

2024

Abstract

Hidden Markov Models (HMM) model a sequence of observations that are dependent on a hidden (or latent) state that follow a Markov chain. These models are widely used in diverse fields including ecology, speech recognition, and genetics.

Parameter estimation in HMM is typically performed using the Baum-Welch algorithm, a special case of the Expectation-Maximisation (EM) algorithm. While this method guarantee the convergence to a local maximum, its convergence rates is usually slow.

Alternative methods, such as the direct maximisation of the likelihood using quasi-Newton methods (such as L-BFGS-B) can offer faster convergence but can be more complicated to implement due to challenges to deal with the presence of bounds on the space of parameters.

We propose a novel hybrid algorithm, QNEM, that combines the Baum-Welch and the quasi-Newton algorithms. QNEM aims to leverage the strength of both algorithms by switching from one method to the other based on the convexity of the likelihood function.

We conducted a comparative analysis between QNEM, the Baum-Welch algorithm, an EM acceleration algorithm called SQUAREM (Varadhan, 2008, Scand J Statist), and the L-BFGS-B quasi-Newton method by applying these algorithms to four examples built on different models. We estimated the parameters of each model using the different algorithms and evaluated their performances.

Our results show that the best-performing algorithm depends on the model considered. QNEM performs well overall, always being faster or equivalent to L-BFGS-B. The Baum-Welch and SQUAREM algorithms are faster than the quasi-Newton and QNEM algorithms in certain scenarios with multiple optimum. In conclusion, QNEM offers a promising alternative to existing algorithms.

Keywords

Hidden Markov models, Baum-Welch, quasi-Newton, L-BFGS-B, SQUAREM, optimisation, computational statistics

1 Introduction

Hidden Markov Models (HMMs) are widely used in time series analysis, with application across diverse fields, including ecology, speech recognition and many others. A particular interest is found in biology, especially in genetics: HMM are often used for genome annotation. For example, CpG islands identification (regions of the genome with a high frequency of cytosine-guanine dinucleotides) (Wu et al., 2010), can be performed through the use of HMM. It can also be used to differentiate exons (protein-coding sequences) and introns (non-coding sequences) within genes (Arimura et al., 2005).

To compute the maximum likelihood estimators of the parameters of HMM, two approaches are naturally employed: the Baum-Welch algorithm (Baum and Petrie, 1966; Welch, 2003), which is a special case of the Expectation-Maximisation (EM) algorithm (Dempster et al., 1977) (cf section 2.2 below) and direct maximisation of the likelihood using a quasi-Newton method (section 2.1), especially BFGS (Broyden-Fletcher-Goldfarb-Shanno) algorithm (Broyden, 1970; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970). Both methods have advantages and drawbacks. In particular, the Baum-Welch algorithm accommodates naturally the presence of bounds on the space of parameters, which can be a hurdle for the quasi-Newton, in particular when some parameters live in an open interval. On the other hand, the Baum-Welch algorithm usually converges more slowly than the quasi-Newton method.

In this paper we propose a new algorithm, QNEM, for HMM parameter estimation. This algorithm starts with one or more EM steps, and switches to a quasi-Newton algorithm, the BFGS, as soon as a criterion on local convexity of the log-likelihood is met. Quasi-Newton iterations are then performed, until convergence or until the aforementioned criterion on local convexity is no longer satisfied, in which case it switches back to the EM algorithm, etc.

We compare its performances with several algorithms: the Baum-Welch algorithm, an accelerated EM algorithm called SQUAREM (section 3.3) and the L-BFGS-B (Limited-memory BFGS with Bound constraints) algorithm, the “off-the-shelf” quasi-Newton method optimized for problems with bounded parameters.

We first review the theory of HMMs, and then expose the four algorithms of interest as well as the examples on which the comparisons are performed. Finally, we present the results obtained and discuss them.

2 Hidden Markov Models

A sequence of random variables measured at successive moments $(X_i)_{i \in \mathbb{N}}$ is a Markov chain if it satisfies the Markov property: for any time i , the distribution of X_{i+1} conditional to $(X_j)_{j \leq i}$ is equal to the distribution of X_{i+1} conditional to the single value X_i . The model is classically illustrated by a directed acyclic graph (DAG) in Figure 1.




Figure 1: Representation of a Markov chain. The X_i are the observations.

One of the property of a Markov chain is the homogeneity of the sequence. A sequence is homogeneous if the transition probabilities $\mathbb{P}(X_{i+1}|X_i)$ are identical in every point of the sequence. This property is sometimes contradicted by a given sequence of observations, making the Markov model inappropriate. A more complex model, in which there is an additional hidden layer $(S_i)_{i \in \mathbb{N}}$ which is a Markov chain, and the distribution of X_i is assumed to depend only on the “hidden state” S_i , can be more appropriate. This model is called Hidden Markov Model (HMM) and is illustrated by the DAG in Figure 2.




Figure 2: Representation of an Hidden Markov chain. The X_i represents the observations and the S_i , the hidden states.

In some situations, the Markov model is known to be adequate, but it can be impossible to directly access the Markov chain, but only variables depending on the current state of the Markov chain. In this case, the HMM provides a natural framework for analysis. HMMs allow to retrieve the states S_i of the latent Markov chain, using the observations X_i whose distribution depends on the hidden states.

2.1 Direct maximisation of the likelihood with a quasi-Newton type method

Direct maximisation of the likelihood (DML) with quasi-Newton methods is a possible way to estimate the HMM parameters.

Quasi-Newton methods are particularly suited for maximisation of differentiable scalar func-

tions, such as likelihood functions. They have the advantage of avoiding the calculation of the Hessian matrix by approximating the Hessian matrix (or its inverse) at each iteration. The off-the-shelf methods are BFGS or L-BFGS-B when some parameters are constrained, standing respectively for Broyden-Fletcher-Goldfarb-Shanno (Broyden, 1970; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970), and for Limited-memory BFGS with Bound constraints.

This method relies on a forward step, computing probabilities of the hidden states, either conditionally to the X_i , or jointly with the X_i . These probabilities allow to compute the likelihood. The most common choice for the DML is using joint probabilities going forward (Thompson, 2000; Welch, 2003; Leutenegger et al., 2003). A log-transformation is often employed to get rid of the numerical underflow implied by the joint probabilities (Zucchini and MacDonald, 2009). In this paper, we chose to use the conditional version of the algorithm for the forward step (cf below section 3.2.1), which avoids the numerical underflow and therefore going through the log-transformation.

2.2 Baum-Welch algorithm, an EM algorithm

The Baum-Welch algorithm (Baum and Petrie, 1966; Welch, 2003) was developed to estimate the HMM parameters.

The respective contributions of Baum and Welch in the development of the Baum-Welch algorithm is explained in Welch’s Shannon lecture, in 2003 (Welch, 2003). After working separately in the calculation of a posteriori probabilities of the hidden states, Baum and Welch worked together on the re-estimation of the parameters from these a posteriori probabilities.

The Baum-Welch algorithm was later recognised as a particular case of the more general Expectation-Maximisation (EM) algorithm formalised by Dempster, Laird and Rubin in 1977 (Dempster et al., 1977).

The Baum-Welch algorithm relies on a forward-backward algorithm for calculating the a posteriori probabilities of the hidden states (E-step). This algorithm is constructed with a forward step as sketched before. We chose to use the forward algorithm based on conditional probabilities, which is more commonly used for the Baum-Welch algorithm. It is also possible to use the forward algorithm based on the joint probabilities, together with an adequate version of the backward step (cf (Thompson, 2000), chapter 6). Then the parameters are re-estimated from the a posteriori probabilities (M-step). This step depends on the model considered.

3 Material and methods

For the following, let $\mathbf{X} = (X_1, \dots, X_L) = X_1^L$ the observations and $\mathbf{S} = (S_1, \dots, S_L) = S_1^L$ the hidden states along the chain of length L .

Let $T_\theta(s, t) = \mathbb{P}_\theta(S_{i+1} = t | S_i = s)$ be the transition probability from state s to state t and $E_\theta(x, s) = \mathbb{P}_\theta(X_i = x | S_i = s)$ the emission probability of observation x from state s .

We consider four maximum likelihood methods to estimate the parameter θ .

3.1 Direct maximisation of the likelihood

As seen in for example Zucchini (Zucchini and MacDonald, 2009) and Turner (Turner, 2008), a variant of the forward algorithm based on joint probabilities can be used to compute the likelihood. If $a_i(s) = \mathbb{P}(S_i = s, X_1^{i-1} = x_1^{i-1})$ and $b_i(s) = \mathbb{P}_\theta(S_i = s, X_1^i = x_1^i)$, the quantities $a_i(s)$ and $b_i(s)$ can be computed recursively as follows:

1. Initialisation :

$$a_1(s) = \pi(s) \tag{1}$$

$$b_1(s) = a_1(s)E_\theta(x_1, s) \tag{2}$$

with $\pi(s)$ the initial distribution of state s .

2. Recursion :

$$a_i(s) = \sum_t b_{i-1}(t)T_\theta(t, s) \tag{3}$$

$$b_i(s) = a_i(s)E_\theta(x_i, s) \tag{4}$$

for $i \in 2, \dots, L$

Finally the likelihood can be computed as

$$L(\theta; \mathbf{X}) = \mathbb{P}_\theta(X_1^L = x_1^L) = \sum_s a_L(s) \tag{5}$$

As the values of the $a_i(s)$ and $b_i(s)$ become very small when i increases, this algorithm is prone to numerical underflow. The usual solution is to compute the values on the logarithmic scale (Zucchini and MacDonald, 2009), which involves rewriting the above equations accordingly. In this paper, we prefer to use the version of the forward algorithm based on conditional probabilities.

This forward algorithm computes classically two conditional probabilities alternatively:

$$\alpha_i(s) = \mathbb{P}_\theta(S_i = s | X_1^{i-1} = x_1^{i-1}) \tag{6}$$

the “forecast” probability and

$$\beta_i(s) = \mathbb{P}_\theta(S_i = s | X_1^i = x_1^i) \tag{7}$$

the “filtering” probability.

To compute the log-likelihood we need to compute also for each i $\gamma_i = \log(\mathbb{P}_\theta(X_1^i = x_1^i))$. These quantities can be computed recursively as follows:

1. Initialisation :

$$\alpha_1(s) = \pi_s \tag{8}$$

$$\beta_1(s) = \frac{E_\theta(x_1, s) \times \alpha_1(s)}{\sum_i E_\theta(x_1, i) \times \alpha_1(i)} \tag{9}$$

$$\gamma_1 = \log \left(\sum_s \alpha_1(s) E_\theta(x_1, s) \right). \tag{10}$$

2. Recursion :

$$\alpha_i(s) = \sum_t T_\theta(t, s) \times \beta_{i-1}(t) \tag{11}$$

$$\beta_i(s) = \frac{E_\theta(x_i, s) \times \alpha_i(s)}{\sum_t E_\theta(x_i, t) \times \alpha_i(t)} \tag{12}$$

$$\gamma_i = \gamma_{i-1} + \log \left(\sum_s E_\theta(x_i, s) \alpha_i(s) \right) \tag{13}$$

for $i \in 2, \dots, L$

Finally, the log-likelihood $\ell(\theta; \mathbf{X})$ is equal to γ_L . The likelihood is still computed on the logarithmic scale, but the $\alpha_i(s)$ and $\beta_i(s)$ can be computed without rewriting all the equations on the logarithmic scale, which make the computation simpler.

The L-BFGS-B algorithm is then performed on the likelihood computed in the last step using the R `optim` function. The gradient of the likelihood along θ can be computed using automatic differentiation (our implementation relies on the R package `salad` (Perdry, 2024)).

3.2 Baum-Welch algorithm

The Baum-Welch EM algorithm is composed by two steps:

- E step: knowing θ , compute the a posteriori probabilities of the hidden states along the chain (forward-backward algorithm)
- M step: re-estimate θ using these probabilities

The Baum-Welch algorithm alternates these 2 steps until convergence. As mentioned in section 2.2, we will use here the forward-backward algorithm presented in (Zucchini and MacDonald, 2009) or (Delmas and Jourdain, 2006).

3.2.1 E step : Forward-Backward algorithm

The forward-backward algorithm runs through the observations twice: the first time forward, ie. from first to last observations ; the second backward, ie. from last to first observation.

After computing the forward quantities shown in equations (6) and (7), the following “backward quantities” will be computed:

$$\delta_i(s, t) = \mathbb{P}_\theta(S_{i-1} = s, S_i = t | X_1^L = x_1^L) \quad (14)$$

the “smoothing” probability and

$$\varphi_i(s) = \mathbb{P}_\theta(S_i = s | X_1^L = x_1^L) \quad (15)$$

the “marginal” probability.

1. Initialisation:

$$\varphi_L(s) = \beta_L(s) \quad (16)$$

2. Recursion:

$$\delta_i(s, t) = T_\theta(s, t) \times \frac{\beta_{i-1}(s)}{\alpha_i(t)} \times \varphi_i(t) \quad (17)$$

$$\varphi_{i-1}(s) = \sum_t \delta_i(s, t) \quad (18)$$

for $i \in L, \dots, 2$.

At the end of the forward-backward, we have computed the probabilities $\varphi_i(s)$ of the hidden states, and the probability $\delta_i(s, t)$ of two consecutive hidden states, given all the observations. These probabilities will be used in the M step.

It is worth to note that the backward algorithm has a higher computational cost than the forward algorithm, due to the necessity to calculate the $\delta_i(s, t)$ for all pairs of states s and t , at each time point i . As mentioned previously, there exists a variant of the backward algorithm based the forward quantities a_i and b_i involving joint probabilities, as seen in (Thompson, 2000).

3.2.2 M step

The aim in M step is to re-estimate the parameter θ that will be used in the following E step.

Note $\theta^{(k)}$ the parameter at the k -th iteration. For each iteration w of the Baum-Welch algorithm, in M step we are going to estimate the parameter $\theta^{(k+1)}$ for the next iteration.

The complete likelihood of θ , ie the likelihood assuming the hidden states are observed, is:

$$L(\theta; \mathbf{S} = \mathbf{s}, \mathbf{X} = \mathbf{x}) = \mathbb{P}_\theta(S_1 = s_1) \cdot \prod_{i=2}^L \mathbb{P}_\theta(S_i = s_i | S_{i-1} = s_{i-1}) \cdot \prod_{i=1}^L \mathbb{P}_\theta(X_i = x_i | S_i = s_i) \quad (19)$$

Now the log-likelihood of θ is:

$$\begin{aligned} \ell(\theta; \mathbf{S} = \mathbf{s}, \mathbf{X} = \mathbf{x}) &= \log \mathbb{P}_\theta(S_1 = s_1) + \sum_{i=2}^L \log \mathbb{P}_\theta(S_i = s_i | S_{i-1} = s_{i-1}) \\ &+ \sum_{i=1}^L \log \mathbb{P}_\theta(X_i = x_i | S_i = s_i) \end{aligned} \quad (20)$$

The M step relies on the expected value of the log-likelihood:

$$\begin{aligned} Q(\theta; \theta^{(k)}) &= \mathbb{E}(\ell(\theta; \mathbf{S} = \mathbf{s}, \mathbf{X} = \mathbf{x}) | \theta^{(k)}) \\ &= \sum_s (\log \mathbb{P}_\theta(S_1 = s) \cdot \varphi_1(s)) + \sum_{i=2}^L \left[\sum_{s,t} (\log \mathbb{P}_\theta(S_i = t | S_{i-1} = s) \cdot \delta_i(s, t)) \right] \\ &+ \sum_{i=1}^L \left[\sum_s (\log \mathbb{P}_\theta(X_i = x_i | S_i = s) \cdot \varphi_i(s)) \right] \end{aligned} \quad (21)$$

The M step ends with finding the parameters that maximise $Q(\theta; \theta^{(k)})$:

$$\theta^{(k+1)} = \operatorname{argmax}_\theta Q(\theta; \theta^{(k)}) \quad (22)$$

In concrete examples, the $\delta_i(s, t)$ will be used to compute the transition parameters of the Markov chain, and the $\phi_i(s)$ the parameters relatives to the emission probabilities.

3.3 EM acceleration

The SQUAREM method was presented in Varadhan, 2008 (Varadhan and Roland, 2008). It is a method created to accelerate the EM algorithm. This method seems interesting when we aim to apply EM to high-dimensional data or data based on complex statistical models. In a nutshell, the SQUAREM makes two EM iterations, and uses these values to extrapolate the trajectory of the algorithm, and estimate a future point several iterations ahead. It then iterates again from this extrapolated point. The constraints on the parameters, if any, are simply dealt with by constraining the extrapolated point to be in the feasible space. We applied this method to our examples to test its performances as well.

3.4 QNEM: a mix of the two algorithms

We propose to mix the Baum-Welch algorithm and the quasi-Newton algorithm BFGS as presented in the algorithm 1.

Algorithm 1 QNEM algorithm

0. Input : $\theta^{(0)}$ = starting point.
Set $k = 0$ and $H_0 = \mathbb{I}$.

1. Compute

$$\theta^{(k+1)} = BW(\theta^{(k)})$$

2. Let $s_k = \theta^{(k+1)} - \theta^{(k)}$ and $y_k = \nabla f_{k+1} - \nabla f_k$.
If the curvature condition $s'_k \cdot y_k > 0$ is not satisfied, go to 1.

3. Update H_k using the BFGS formula

$$H_{k+1} = H_k - \frac{H_k \cdot y_k \cdot y'_k \cdot H_k}{y'_k \cdot H_k \cdot y_k} + \frac{s_k \cdot s'_k}{y'_k \cdot s_k}$$

Let $k = k + 1$.

4. Let $p_k = -H_k \cdot \nabla f_k$ the search direction.
Perform a backtracking search for α such that Armijo condition is satisfied

$$f(\theta_k + \alpha \cdot p_k) \leq f(\theta_k + c \cdot \alpha (\nabla f_k)' p_k)$$

5. Let $\theta^{(k+1)} = \theta^{(k)} + \alpha \cdot p_k$.
If the curvature condition is not satisfied, reinitialise $H_k = \mathbb{I}$ and go back to 1.
Else go to 3.
-

If the current value of the parameter is $\theta^{(k)}$, the BFGS algorithm finds $\theta^{(k+1)}$ by performing a linear search in a search direction $p_k = -H_k \nabla f_k$ where H_k is the current approximation of the inverse of the Hessian of the objective function, and ∇f_k is its gradient in $\theta^{(k)}$. The H_k is updated using the formula

$$H_{k+1} = H_k - \frac{H_k \cdot y_k \cdot y'_k \cdot H_k}{y'_k \cdot H_k \cdot y_k} + \frac{s_k \cdot s'_k}{y'_k \cdot s_k}$$

where $s_k = \theta^{(k+1)} - \theta^{(k)}$ and $y_k = \nabla f_{k+1} - \nabla f_k$. If H_k is positive definite, H_{k+1} will be positive definite as soon as the scalar product of s_k and y_k is positive:

$$s'_k \cdot y_k > 0 \tag{23}$$

This is the curvature condition. If the function is strongly convex it always holds after a linear search which successfully decreases the value of the objective function. In the general case, it may not hold, leading to the use of complex linear search methods to avoid this situation, such as the Nocedal line search, see (Nocedal and Wright, 2006) chapter 3.

The QNEM starts by performing Baum-Welch algorithm steps, until the curvature condition (equation 23) is met; it then switches to BFGS algorithm. Our BFGS iterations use a simple backtracking search, until the Armijo condition is met. This condition ensures that the objective function decreased.

As long as the curvature condition holds, BFGS iterations are continued. If at some point it does not hold, the QNEM falls back to the Baum-Welch algorithm, until the curvature condition is met and the quasi-Newton iterations are performed again, until convergence. The box constraint are simply dealt with by projecting the gradient on the box when θ is on the border.

3.5 Stopping criterion

For the four algorithms, we choose to use the same stopping criterion to decide when the convergence point is reached. The criterion is the classic relative convergence tolerance (*reltol*), which is based on the likelihood of the current iteration and the previous one. At the k -th iteration, the algorithm stops if

$$\frac{|\ell^{(k-1)} - \ell^{(k)}|}{|\ell^{(k-1)}| + \textit{reltol}} < \textit{reltol}$$

where $\ell^{(k)}$ is the likelihood of the k -th iteration and *reltol* is a small constant, depending on the machine precision; in our experiment it is $1.49 \cdot 10^{-8}$.

3.6 Examples

3.6.1 Umbrella example

We will first use a simple example based on a fictitious situation. Every day of the year, in a secret underground installation, a security guard observes whether the director comes in with an umbrella or not (Gao, 2021). During 56 days, the security guard keep track of his observations on his note pad as follows: U if the director carries an umbrella; N otherwise. The guards aims to predict the daily weather conditions (hidden states) based on the umbrella observations: rainy (noted R) or dry (noted D) day. Here the hidden state is the weather. Since the security guard is in a underground installation, he can not observe directly the weather: he has to rely only on his observations of the umbrella status to determine the weather. The parameters to estimate are a the probability of weather state transition and b the probability of error in the umbrella status (not carrying an umbrella on rainy day or carrying an umbrella on dry day). The transition and emission matrix are presented respectively in Tables 1 and 2.

To apply the quasi-Newton algorithm (L-BGFS-B) to this example, it was necessary to bound the values of a and b to the intervals $[0.01, 0.99]$. Using the interval $[0, 1]$ would be preferable, but the likelihood is not always defined on the boundary, which leads to a failure of the algorithm. The algorithm can not accommodate open intervals such as $(0, 1)$.

Transition matrix	$S_i = D$	$S_i = R$
$S_{i-1} = D$	1-a	a
$S_{i-1} = R$	a	1-a

Table 1: Transition matrix for the umbrella example

Emission matrix	$S_i = D$	$S_i = R$
$X_i = N$	1-b	b
$X_i = U$	b	1-b

Table 2: Emission matrix for the umbrella example

3.6.2 Old Faithful geyser example

Following an example presented in Zucchini (Zucchini and MacDonald, 2009), we work with the Old Faithful geyser data (Azzalini and Bowman, 1990), available in the R package MASS, considering the 272 duration of the eruptions. We make two uses of the data. First, we dichotomise at three minutes of eruption to divide the observations in two groups "Duration < 3min" and "Duration ≥ 3 min" (noted as "Dinf3" and "Dsup3" respectively) to work in a discrete observations case. This is similar to what is done in (Zucchini and MacDonald, 2009). Second, we work with the continuous values, assuming they are drawn in Gaussian distributions whose parameters depend on the hidden states. In the sequence of dichotomised values, it appears that each short eruption was followed by a long eruption, but there are sequences of long eruptions. We decided to assume that the hidden Markov chain has three states: "short", "long" and "steady long", noted respectively "S", "L" and "Sl". The parameters to estimate in the dichotomised model are a the probability to change from state "short" to "steady long", b the probability to stay in "steady long" state, c , d and e the probability to observe a "Duration ≥ 3 min" in respectively "short", "long" and "steady long" states. This emission matrix is presented in Table 4. For the continuous model, the parameters to estimate are the same a and b and six other parameters μ_s , μ_l and μ_{sl} , the mean of the Gaussian distributions in which duration of eruptions are drawn, and σ_s , σ_l and σ_{sl} , their standard deviation. The transition matrix, common to both the dichotomised and continuous models is presented in Table 3

Transition matrix	$S_i = S$	$S_i = L$	$S_i = Sl$
$S_{i-1} = S$	0	1-a	a
$S_{i-1} = L$	1	0	0
$S_{i-1} = Sl$	1-b	0	b

Table 3: Transition matrix for the geyser example

Emission matrix	$S_i = S$	$S_i = L$	$S_i = Sl$
$X_i = Dinf3$	1-c	1-d	1-e
$X_i = Dsup3$	c	d	e

Table 4: Emission matrix for the dichotomised version of the geyser example

Again, for the quasi-Newton L-BFGS-B algorithm, we had to constraint the probabilities a , b , c , d , and e to the interval $[0.01, 0.99]$, and to set a lower bound of 0.01 for σ_s , σ_l , and σ_{sl} .

3.6.3 HBD segments example

Consanguineous individuals are offspring of related parents. In the absence of pedigree information, consanguineous individuals can be identified through their genome which carries Homozygous-By-Descent (HBD) segments. HBD segments are regions of the genome where the two copies of the genome are identical and come from a common ancestor of the parents. These segments are characteristic of consanguineous individuals. Reconstructing these regions allow to determine whether an individual is consanguineous and his/her degree of consanguinity. The coefficient of consanguinity f can be defined as the expected proportion of HBD genome. HMM are used to compute for each variant of the genome, its probability to be in a HBD segment. These probabilities can be used to reconstruct the HBD segments, but this requires to first estimate the parameters of the HMM. The parameters to estimate are the coefficient of consanguinity f and another parameter a related to the mean length of HBD segments (Leutenegger et al., 2003; Gazal et al., 2014). cf. Figure 3

To test the algorithms, we will generate the genome of an individual from a given set of parameters $\theta = (f = 0.0625, a = 0.064)$, which corresponds to the expected values for an offsprings of first cousins. The data were created from an inhouse R package named **Mozza** (available at <https://github.com/genostats/Mozza>, still in development). We then rounded the distances between positions to 0.1 centiMorgan (cM). For each of these new positions, we randomly selected only one observation, to create submaps of positions. This corresponds to a total of 1050 observations.

The transition and emission probabilities are presented in Tables 5 and 6. We note d the distance between positions (fixed here at 0.1 cM), pA and pa the allele frequency of the reference and alternative allele respectively and ϵ the genotyping or mutation error.

Transition matrix	$S_i = \text{nHBD}$	$S_i = \text{HBD}$
$S_{i-1} = \text{nHBD}$	$(1 - \exp^{-ad})(1 - f) + \exp^{-ad}$	$(1 - \exp^{-ad})f$
$S_{i-1} = \text{HBD}$	$(1 - \exp^{-ad})(1 - f)$	$(1 - \exp^{-ad})f + \exp^{-ad}$

Table 5: Transition matrix for the HBD example

Emission matrix	$S_i = \text{nHBD}$	$S_i = \text{HBD}$
$X_i = \text{AA}$	pA^2	$(1 - \epsilon)pA + \epsilon \cdot pA^2$
$X_i = \text{Aa}$	$2 \cdot pA \cdot pa$	$2\epsilon pA \cdot pa$
$X_i = \text{aa}$	pa^2	$(1 - \epsilon)pa + \epsilon \cdot pa^2$
$X_i = \text{NA}$	1	1

Table 6: Emission matrix for the HBD example

In this example, the probability f was constrained to the interval $[0.01, 0.99]$ and the lower-bound was 0.01 for the parameter a when applying the quasi-Newton algorithm.




Figure 3: This figure shows on the upper left a representation of the genealogy of a consanguineous individual C and on the upper right his/her genome. In red are represented the Homozygous By Descent (HBD) segments. On the lower part is represented a chromosome with the information about the hidden states (2 possible values : HBD or non-HBD) and the observations (genotypes with B the reference allele and b the alternative allele).

3.7 Comparison of the algorithms

We compare the likelihood maximisation methods by comparing

- the number of iteration to convergence
- the number of forwards step runs
- the number of backwards step runs
- the running time
- the percentage of runs that failed to converge
- the running time of forward step
- the running time of backward step

We also studied the final estimations of the parameters and calculated the percentage of appearance of each corresponding likelihoods.

For all four examples, the tests were made on 1000 uniform draws within a wide possibility of values for the initial θ , some of them being unlikely.

4 Results

4.1 Umbrella example

The results are presented in Table 7.

Umbrella example	Nb of iterations						Mean nb of steps		Mean time (s)	Percent of convergence
	Min	Q1	Q2	Mean	Q3	Max	Fw	Bw		
qN	4	13	16	15.77	18	37	15.77	0	0.03	100
BW	3	18	24	27.95	31	338	27.95	27.95	0.06	100
SQUAREM	3	20	24	26.76	29	179	17.70	17.67	0.04	100
QNEM	2	8	11	11.32	13	118	12.63*	2.80	0.03	100

Table 7: This table recapitulates for each algorithm the number of iterations, the mean number of steps, the mean running time and the percentage of convergence in the umbrella example. The mean running times are in seconds. For all algorithms, each backward step takes ~ 1.7 ms, but the forward step for the Baum-Welch (BW) and the SQUAREM algorithms is ~ 0.5 ms only while it is ~ 1.1 ms for the quasi-Newton (qN) and the QNEM algorithms, due to the computation of the derivatives. * = The number of forward steps is higher than the number of iterations because the backtracking is not counted in the number of iterations.

We can see that, overall, the QNEM algorithm achieves convergence with the lowest mean number of iterations, averaging 11.32 across all runs. However, the quasi-Newton algorithm

has the lowest maximum number of iterations (37 iterations vs. 118-338 iterations for the other three algorithms). The number of iterations of the Baum-Welch and the SQUAREM algorithms are comparable and systematically higher than those of the QNEM and the quasi-Newton algorithms.

Apart from the total number of iterations, it is interesting to look at the mean number of forward and backward steps used by the four algorithms. The quasi-Newton algorithm does not perform any backward steps. For the quasi-Newton and Baum-Welch algorithms, the number of forward steps corresponds to the mean number of iterations on the 1000 draws of the initial starting point. The Baum-Welch algorithm compute the same number of forward and backward steps. For the SQUAREM algorithm, the mean number of forward and backward steps is slightly different, with 17.70 forward steps and 17.67 backward steps. Finally, the QNEM algorithm executes an average of 12.63 forward steps and 2.80 backward steps, where the backward steps correspond to the calls to the Baum-Welch algorithm. We note that a forward step takes about 1.5ms for the Baum-Welch and the SQUAREM algorithms and around 11.5ms for the other two, which is longer because of the additional computation of the derivatives. A backward step takes about 8ms for all algorithms. The total mean time spent in forward and backward steps for a run is almost the same for the quasi-Newton algorithm and the QNEM (~ 18 ms). It is higher for the SQUAREM, and even longer for the Baum-Welch algorithm (~ 39 ms and 61.5ms respectively). This trend is in accordance with the total mean execution time, where the Baum-Welch algorithm is the slowest one and the quasi-Newton and the QNEM algorithms are the fastest.

Likelihood	Parameters		Percent of runs			
	a	b	qN	BW	SQUAREM	QNEM
33.3	0.06	0.18 0.82	89.9	70.1	69.9	70.1
33.8	ND	0.5	10.1	29.9	30.1	29.3
Other			0	0	0	0.6

Table 8: This table shows the different convergence points reached with the umbrella example and their percent. ND = Non determined.

We also examined the final values obtained for θ . They can be divided into three subsets, corresponding to three different regions of the likelihood as shown in Figure 4. Two of these regions are points with $a = 0.06$ and $b = 0.18$ or $b = 0.82$, corresponding to the highest likelihood value (negative log-likelihood equal to 33.3). The third region is the segment $b = 0.5$ and $a \in [0.5, 1]$ which corresponds to a crest line (negative log-likelihood equal to 33.8). The parameters and the corresponding proportion of runs converging to these parameters for each algorithm are presented in Table 8. The θ with the best likelihood is reached almost 90% of the time with the quasi-Newton algorithm but only 70% of the time for the other three algorithms.




Figure 4: The likelihood with 100 convergence points.

4.2 Old Faithful geyser example

We will first present the results obtained with the dichotomised data, then with the continuous values.

4.2.1 Dichotomised values

As previously done, we tested the algorithms on 1000 random initial sets of parameters. The results are presented in Table 9.

The quasi-Newton and the QNEM algorithms had a lower number of iterations compared to the SQUAREM and the Baum-Welch algorithms. The latter showed a mean number of iterations almost 6 times higher. The lowest number of iterations was always reached by the QNEM.

Moreover, the quickest algorithm was the QNEM as well, with a mean of 0.26 seconds by run, compared to the quasi-Newton, the Baum-Welch and the SQUAREM algorithms, which had respectively 0.33 seconds, 1.21 seconds and 0.71 seconds. The mean time spent in the forward and backward step is the lower for the QNEM as well, with 215ms compared to 280ms for the quasi-Newton algorithm, 1197ms for the Baum-Welch algorithm and 693ms for the SQUAREM.

We note that the QNEM did not converge for two runs.

We then present the table of the convergence points in Table 10. We see that the optimum

Geyser example (dichotomised)	Nb of iterations						Mean nb of steps		Mean time (s)	Percent of convergence
	Min	Q1	Q2	Mean	Q3	Max	Fw	Bw		
qN	12	18	21	24.23	26	66	24.23	0	0.33	100
BW	12	82	139	126	161	324	126	126	1.21	100
SQUAREM	15	71	113	108	140	290	78.59	71.89	0.71	100
QNEM	6	14	16	16.27	19	64	17.55*	1.62	0.26	99.8

Table 9: This table recapitulates for each algorithm the number of iterations, the mean number of steps, the mean running time and the percentage of convergence in the dichotomised geyser example. The mean running times are in seconds. For all algorithms, each backward step takes ~ 8 ms, but the forward step for the Baum-Welch (BW) and the SQUAREM algorithms is ~ 1.5 ms only while it is ~ 11.5 ms for the quasi-Newton (qN) and the QNEM algorithms, due to the computation of the derivatives. * = The number of forward steps is higher than the number of iterations because the backtracking is not counted in the number of iterations.

Likelihood	Parameters					Percent of runs			
	a	b	c	d	e	qN	BW	SQUAREM	QNEM
144.5	0.79	0.57	0*	1*	0.95	52.6	50.6	50.8	42.6
149.5	0.58	0.50	1*	0*	0.57	41.2	42.4	42.4	36.2
172.8	0* ND	ND 0*	0.55 / 0.74	0.74 / 0.55 ND	ND	6.2	7.0	0.9 5.9	4.4 4.5
Other						0	0	0	11.1

Table 10: This table shows the different convergence points reached with the dichotomised geyser example and their percent. ND = Non determined. * = These values were rounded for the quasi-Newton results because of the bound-constraints.

is obtained most of the time by all algorithms but it represents only $\sim 50\%$ of the runs for all algorithms except the QNEM, for which it's closer to 40%. There are two other local optimum reached by the algorithms. Moreover, we observe that the QNEM converges to other local optima not reached by the other algorithms.

4.2.2 Continuous values

For the continuous case, the summary of the results is presented in Table 11.

Geyser example (continuous)	Nb of iterations						Mean nb of steps		Mean time (s)	Percent of convergence
	Min	Q1	Q2	Mean	Q3	Max	Fw	Bw		
qN	21	60	78	79.41	98	159	79.41	0	1.47	93.2
BW	8	21	23	25.32	26	129	25.32	25.32	0.24	100
SQUAREM	12	30	32	34.04	35	118	23.23	22.59	0.22	100
QNEM	4	19	23	24.01	27	87	31.61*	1.89	0.66	99.5

Table 11: This table recapitulates for each algorithm the number of iterations, the mean number of steps, the mean running time and the percentage of convergence in the continuous geyser example. The mean running times are in seconds. For all algorithms, each backward step takes ~ 8 ms, but the forward step for the Baum-Welch (BW) and the SQUAREM algorithms is ~ 1.5 ms only while it is ~ 17.5 ms for the quasi-Newton (qN) and the QNEM algorithms, due to the computation of the derivatives. * = The number of forward steps is higher than the number of iterations because the backtracking is not counted in the number of iterations.

The quasi-Newton algorithm showed the biggest amount of iterations, with a maximum of 159 iterations. The QNEM keeps showing the less number of iterations globally.

In term of running time, the lowest are reached by the EM-based algorithms: the Baum-Welch algorithm and the SQUAREM, with a mean running time near 0.23s. Therefore, the quasi-Newton was slow on this continuous case, with a mean running time of 1.47s. The hybrid QNEM showed an intermediate time of 0.66s. This difference of time can be explained by the time spent in the forward step, which is 1.5 millisecond for the forward used in the SQUAREM and the Baum-Welch (conditional probabilities forward) against 17.5 milliseconds for the forward used in the other two algorithms (conditional probabilities forward with derivatives) for a comparable mean number of forward steps.

The quasi-Newton algorithm did not converge for 68 sets of initial parameters and the QNEM for 5 sets of parameters.

In Table 12, the best likelihood is reached most of the time by the Baum-Welch, the SQUAREM and the QNEM algorithms but not by the quasi-Newton algorithm. The latter converges more often to local optima. Most of the time, the quasi-Newton converge to the point with the third best likelihood. This behavior can explain its slow convergence in this example.

Likelihood	Parameters								Percent of runs			
	a	b	μ_s	μ_l	μ_{sl}	σ_s	σ_l	σ_{sl}	qN	BW	SQUAREM	QNEM
265.7	0.61	0.65	2.0	4.58	4.09	0.22	0.24	0.64	19.1	40.7	40.5	35.7
275.1	0.54	0.46	4.44	1.98	3.23	0.3	0.19	1.03	29.2	38.6	38.6	30.6
303.2	0.93	0.34	2.37	1.91	4.34	0.77	0.20	0.37	21.9	9.0	9.6	9.9
316.2	1*	0.28	4.45	ND	2.79	0.30	ND	1	0.1	0	0	6.4
335	0.44 0.56	0*	3.72	4.24 2.0	2.0 4.24	1.05	0.45 0.22	0.22 0.45	7.5	5.5	4.9	5.2
Other									22.2	6.2	6.4	12.2

Table 12: This table shows the different convergence points reached with the continuous geyser example and their percent. ND = Non determined. * = These values were rounded for the quasi-Newton results because of the bound-constraints.

4.3 HBD segments example

For this last example, the results are presented in Table 13.

HBD segments example	Nb of iterations						Mean nb of steps		Mean time (s)	Percent of convergence
	Min	Q1	Q2	Mean	Q3	Max	Fw	Bw		
qN	7	11	14	15.36	18	38	15.36	0	0.32	100
BW	5	63	73	68.53	78	86	68.53	68.53	2.40	100
SQUAREM	5	50	57	56.19	63	88	38.66	38.66	1.37	100
QNEM	5	9	11	11.99	14	27	12.83*	1.40	0.32	100

Table 13: This table recapitulates for each algorithm the number of iterations, the mean number of steps, the mean running time and the percentage of convergence in the HBD segments example. The mean running times are in seconds. For all algorithms, each backward step takes ~ 30.3 ms, but the forward step for the Baum-Welch (BW) and the SQUAREM algorithms is ~ 5.3 ms only while it is ~ 20.2 ms for the quasi-Newton (qN) and the QNEM algorithms, due to the computation of the derivatives. * = The number of forward steps is higher than the number of iterations because the backtracking is not counted in the number of iterations.

The smallest number of iterations is always hit by the QNEM, with a maximum of 27 iterations. It is closely followed by the quasi-Newton algorithm. On the opposite, the Baum-Welch algorithm shows less good performances in term of number of iterations with a mean of 68.53 iterations. The SQUAREM algorithm, is a bit better than the Baum-Welch algorithm with a mean of 56.19 iterations.

In term of running time, the QNEM and the quasi-Newton shows a similar mean running time of 0.32s, whereas the SQUAREM and the Baum-Welch algorithms are way longer with respectively 1.37s and 2.40s of running time. These running times are linked to the time spent in forward and backward steps. Even if the time spent in the forward step by the two EM-based algorithms (5.3ms) is lower than the time spent in the QNEM and quasi-Newton's forward step

(around 20ms), the two latter need less forward steps (respectively a mean of 12.83 and 15.36 iterations) than the Baum-Welch and the SQUAREM algorithms (respectively a mean of 68.53 and 38.66 iterations).

For all algorithms, only one likelihood optimum is reached with all 1000 draws of the starting value θ . In fact, this likelihood admits only one optimum.

5 Discussion

The Baum-Welch (or EM) algorithm is a natural choice when it comes to HMM parameter estimation, offering several advantages: its conceptual simplicity as well as the ease of implementation, and the fact it naturally accommodates the possible constraints on the parameters. However, while the first iterations may move quickly in the space of parameters, once it is close to the solution its convergence is usually slow.

Another possibility is to use a direct maximisation of the likelihood, particularly via a quasi-Newton method. It is expected to be more efficient than the Baum-Welch algorithm. However, the computation of the likelihood, and of its derivatives, on the logarithmic scale, can be tedious. In the literature, it is usually done using a joint probability version of the forward algorithm (equations 1 to 5), in which all quantities need to be computed in the logarithmic scale. We showed that it is possible to use the conditional probability version of the forward algorithm (equations 8 to 13) to compute the log-likelihood in a simpler manner.

A practical problem with the use of the quasi-Newton algorithm is the presence of bounds on the parameters. Typically, in our examples, many parameters take their values in $[0, 1]$, and others in $[0, +\infty)$, with possibly an infinite value of the log-likelihood on the border of the space of parameters. The off-the-shelf quasi-Newton method for this case is L-BFGS-B, which tends to favour in its first iterations points on the border of the space, leading to a failure of the algorithm. A common solution, but not totally satisfying, is to use constraints such as $[0.01, 0.99]$ instead.

An appealing alternative is to use an EM acceleration, such as the SQUAREM algorithm. It is simple to implement from an existing Baum-Welch implementation. Bounds on the parameters are easily dealt with, including when the likelihood is infinite on the border.

We propose a new algorithm, QNEM, which alternates between Baum-Welch algorithm iterations and BFGS iterations. Our intuition was that the Baum-Welch algorithm can help to explore efficiently the space of parameters, and the BFGS can allow to achieve a fast convergence once the algorithm is near to a local optimum. The algorithm switches between the two kind of iterations using a convexity criterion. The box-constraints are simply managed with a projection of the gradient on the box when the parameter θ gets on the border.

We compared the four algorithms on four different examples. The umbrella example (cf. sections 3.6.1 and 4.1) is a toy example with two hidden states and two possible observed states. As expected, the quasi-Newton was much faster than the Baum-Welch algorithm. The SQUAREM succeeds in accelerating the Baum-Welch algorithm, but is still much slower than the quasi-Newton. Our algorithm, QNEM, is equivalent to the quasi-Newton in term of computational time.

Our second example was based on the Old Faithful data, with three hidden states and two possible observed states (cf. sections 3.6.2 and 4.2.1). In that case, the QNEM is the faster of all algorithms, followed by the quasi-Newton, then by the SQUAREM, and the Baum-Welch is the slowest.

In our third example, the Old Faithful data were used again, with the same model for the hidden Markov chain, but now the observed states are continuous (cf. sections 3.6.2 and 4.2.2). Surprisingly, for this model, the quasi-Newton is by far the slowest method. The Baum-Welch algorithm is faster, and the SQUAREM is even slightly faster. The QNEM ranks third, with intermediate performances.

Our last example, which was our motivating example, addresses a concrete question in genetics: making inference about HBD segments on the genome. It has two hidden states, and the observations are discrete (cf. sections 3.6.3 and 4.3). In that case, just as in the first example, the QNEM and the quasi-Newton were equivalent, and performed better than the SQUAREM and the Baum-Welch in term of computational time. Our results are in line with those of (Bertrand et al., 2019), who also compared the L-BFGS-B algorithm to the EM algorithm in term of number of iterations and running time on a more complex HMM model aimed to identify HBD segments. In their work, the quasi-Newton method is several time faster than the EM algorithm.

In all examples except the HBD segments example, there are multiple local likelihood optimum, leading to multiple final estimations of the parameters. In a concrete situation, to deal with this issue, one would need to run the algorithms multiple times from different initialisation of the parameters, to retain the highest likelihood. In our examples, the quasi-Newton and the QNEM seems more prone to converge to a local optimum which is not a global optimum. It is worth to note that in our numeric experiments, the starting points are drawn in very large intervals, which may exacerbate this behaviour.

There was no algorithm uniformly more efficient. In particular, it was surprising to us to find that on our third example, the Baum-Welch and the SQUAREM are the more efficient ones. The presence of multiple optima in this example seems to play a role in this behaviour. The QNEM was the most efficient on our second example, and equivalent to the quasi-Newton on the first and fourth example. This new optimization algorithm for HMM parameter estimation seems to be an interesting alternative to the other methods, offering a good compromise in most situations.

Code availability

Related to this paper, we created a R package available at <https://github.com/SidonieFoulon/steveHMM>.

Acknowledgment

This work is part of the Inserm cross-cutting program GOLD (GenOmics variability in health & Disease)

Funding sources

Sidonie Foulon is funded for this work by a doctoral grant from the French Ministry of Research.

References

- Kazuhiko Arimura, Ikuo Yoshihara, and Kunihiro Yamamori. Identification of exon-intron boundary by hidden markov model and evaluation with the human genome. *34*:269–274, 2005.
- Adelchi Azzalini and Adrian W Bowman. A look at some data on the old faithful geyser. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 39(3):357–365, 1990.
- Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.
- Amandine R Bertrand, Naveen K Kadri, Laurence Flori, Mathieu Gautier, and Tom Druet. Rzooroh: an r package to characterize individual genomic autozygosity and identify homozygous-by-descent segments. *Methods in Ecology and Evolution*, 10(6):860–866, 2019.
- Charles G Broyden. The convergence of a class of double-rank minimization algorithms: 2. the new algorithm. *IMA journal of applied mathematics*, 6(3):222–231, 1970.
- Jean-François Delmas and Benjamin Jourdain. *Modèles aléatoires: applications aux sciences de l'ingénieur et du vivant*. Springer, 2006.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1):1–22, 1977.
- Roger Fletcher. A new approach to variable metric algorithms. *The computer journal*, 13(3):317–322, 1970.
- Alice Gao. Inference in Hidden Markov Models, November 2021. URL https://www.cs.toronto.edu/~axgao/cs486686_f21/lecture_notes/Lecture_14_on_Hidden_Markov_Models_1.pdf.
- Steven Gazal, Mourad Sahbatou, Marie-Claude Babron, Emmanuelle Génin, and Anne-Louise Leutenegger. Fsuite: exploiting inbreeding in dense snp chip and exome data. *Bioinformatics*, 30(13):1940–1941, 2014.
- Donald Goldfarb. A family of variable-metric methods derived by variational means. *Mathematics of computation*, 24(109):23–26, 1970.
- Anne-Louise Leutenegger, Bernard Prum, Emmanuelle Génin, Christophe Verny, Arnaud Lemainque, Françoise Clerget-Darpoux, and Elizabeth A Thompson. Estimation of the inbreeding coefficient through use of genomic data. *The American Journal of Human Genetics*, 73(3):516–523, 2003.
- Jorge Nocedal and Stephen J Wright. *Numerical optimization*. Springer, 2006.

- Hervé Perdry. *salad: Simple Automatic Differentiation*, 2024. URL <https://CRAN.R-project.org/package=salad>.
- David F Shanno. Conditioning of quasi-newton methods for function minimization. *Mathematics of computation*, 24(111):647–656, 1970.
- Elizabeth Alison Thompson. Statistical inference from genetic data on pedigrees. IMS, 2000.
- Rolf Turner. Direct maximization of the likelihood of a hidden markov model. *Computational Statistics & Data Analysis*, 52(9):4147–4160, 2008.
- Ravi Varadhan and Christophe Roland. Simple and globally convergent methods for accelerating the convergence of any em algorithm. *Scandinavian Journal of Statistics*, 35(2):335–353, 2008.
- Lloyd R Welch. Hidden markov models and the baum-welch algorithm. *IEEE Information Theory Society Newsletter*, 53(4):10–13, 2003.
- Hao Wu, Brian Caffo, Harris A Jaffee, Rafael A Irizarry, and Andrew P Feinberg. Redefining cpg islands using hidden markov models. *Biostatistics*, 11(3):499–514, 2010.
- Walter Zucchini and Iain L MacDonald. *Hidden Markov models for time series: an introduction using R*. Chapman and Hall/CRC, 2009.