

Mixture of Coupled HMMs for Robust Modeling of Multivariate Healthcare Time Series

Onur Poyraz ONUR.POYRAZ@AALTO.FI and Pekka Marttinen
Department of Computer Science, Aalto University, Finland

PEKKA.MARTTINEN@AALTO.FI

Abstract

Analysis of multivariate healthcare time series data is inherently challenging: irregular sampling, noisy and missing values, and heterogeneous patient groups with different dynamics violating exchangeability. In addition, interpretability and quantification of uncertainty are critically important. Here, we propose a novel class of models, a mixture of coupled hidden Markov models (M-CHMM), and demonstrate how it elegantly overcomes these challenges. To make the model learning feasible, we derive two algorithms to sample the sequences of the latent variables in the CHMM: samplers based on (i) particle filtering and (ii) factorized approximation. Compared to existing inference methods, our algorithms are computationally tractable, improve mixing, and allow for likelihood estimation, which is necessary to learn the mixture model. Experiments on challenging real-world epidemiological and semi-synthetic data demonstrate the advantages of the M-CHMM: improved data fit, capacity to efficiently handle missing and noisy measurements, improved prediction accuracy, and ability to identify interpretable subsets in the data.

Keywords: Markov Chain Monte Carlo (MCMC); Multivariate Time Series; Probabilistic Graphical Models; Robustness.

1. Introduction

Hidden Markov models (HMMs) (Rabiner, 1989) are widely used probabilistic frameworks for modeling longitudinal structures when some underlying hidden process happens in time. They have been widely studied and have efficient state and parameter estimation algorithms (Bishop, 2006; Murphy, 2012; Barber and Cemgil, 2010). HMMs have a hidden state at each time point, representing the system’s status at that particular time. However, many systems in the real world result from multiple interacting processes (i.e., coupled latent Markov chains), causing the state-space

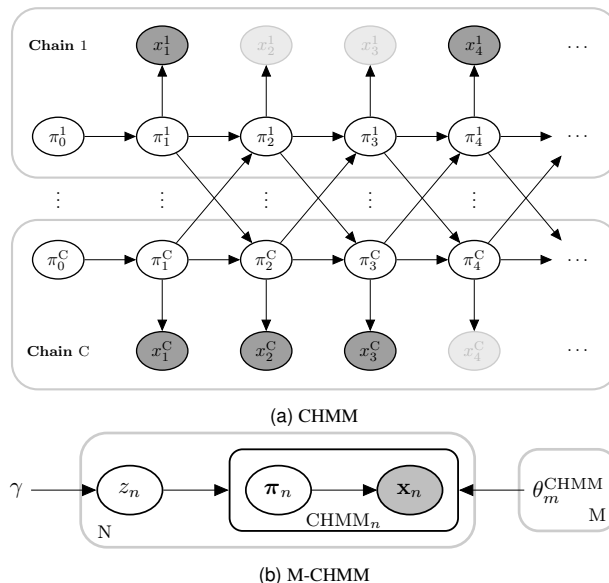


Figure 1: An illustration of (a) CHMM and (b) mixture of CHMMs (M-CHMM), in which each individual n has a cluster label z_n which specifies one of M possible CHMMs that describes the dynamics of the individual.

to grow exponentially with respect to the number of chains; therefore, learning the HMM becomes intractable (Touloupou et al., 2020). Coupled hidden Markov models (CHMMs, Brand et al. (1997)), an extension of HMM, model interactions among multiple chains within individuals by allowing the current latent state of a chain to depend on the previous states of all the chains, as illustrated in Figure 1a. However, the system is still computationally intractable when the number of chains grows.

Real-world problems in epidemiology often involve longitudinal data on multiple interacting processes (Brand et al., 1997; Brand, 1997). The data can be complicated by missing, irregularly collected, and

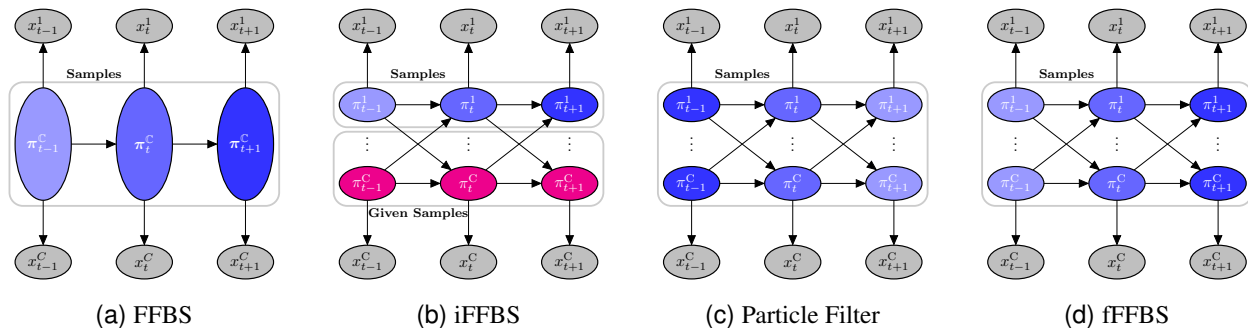


Figure 2: Strategies for sampling the latent variables in CHMM. The shade of blue shows the sampling order (darker colors are sampled first). (a) A CHMM is converted into a joint-HMM, and latent variables of chains are sampled jointly using the FFBS algorithm. (b) Latent variables of each chain are sampled conditionally on the latent variables of other chains. (c, d) Latent variables are sampled jointly at each time step using (c) particle filtering, which results in factorized joint sampling, and (d) factorized FFBS (fFFBS), which makes the joint sampling factorizable by assumption.

noisy observations. Hence, modeling the data is non-trivial, and application of recent neural network methods for longitudinal health data (Alaa et al., 2022; Kodialam et al., 2021; Choi et al., 2016) is infeasible, especially when data are sparse and limited, as in the present work. On the other hand, probabilistic CHMMs provide two advantages for dealing with these challenges: 1) When the collected observations (e.g., medical samples or sensory information) are noisy or missing, a CHMM can model the true state as a latent chain, in which the sensitivity and specificity are available as emission parameters, and 2) the dynamic relations between the variables in the dataset are intuitively modeled by a transition matrix that describes how different latent chains affect each other. Hence, the CHMM provides a flexible, interpretable, and data-efficient alternative for modeling complex longitudinal epidemiological data. However, another challenge in real-world data is that individuals might follow different dynamics; for example, patients may respond to treatments differently, and the standard CHMM fails to capture this.

Multiple methods for inference with CHMMs have been published, which can be broadly categorized into three main branches: 1) Maximum likelihood, 2) variational approximation, and 3) Markov chain Monte Carlo (MCMC). Earlier works (Brand et al., 1997; Saul and Jordan, 1999; Rezek et al., 2000; Zhong and Ghosh, 2001) focused primarily on the first solution; however, exact inference for CHMMs using expectation-maximization (EM) results in computa-

tional issues. Wang et al. (2019) proposed a variational EM to alleviate this. However, our focus here is on MCMC as it is the only method that fully estimates the uncertainty of hidden states and parameters, which is important in many applications, e.g., healthcare. The main bottleneck in MCMC is the sampling of the latent chains. One straightforward solution is to convert the CHMM into a single HMM where the single latent state jointly represents the states of all the chains, and apply the standard forward filtering backward sampling (FFBS) to sample the latent variables (Figure 2a). However, this simple solution is infeasible with a large number of chains. Therefore, several alternative methods have been proposed. Conditional single-site (Dong et al., 2012) or block updates (Spencer et al., 2015) for latent sequences provide computationally less demanding methods but produce highly correlated samples. Individual FFBS (iFFBS, Touloupou et al. (2020)) achieves efficiency by updating each latent chain conditionally on the other chains (Figure 2b); however, this conditioning increases correlation between posterior samples and does not allow for an easy way to estimate the likelihood, necessary with more complex models.

Our first contribution is to extend the class of CHMM to mixtures of CHMMs (M-CHMM), which resolves the challenge of distinctly behaving subgroups in healthcare time series data by learning the groups and the corresponding component models (Figure 1b). *Our second contribution* is to investigate the suitability of two novel algorithms to sample the latent variables

Algorithm 1: Coupled Hidden Markov Model

Input: Prior of η, τ, ϵ ; Initialize π using priors

Output: Posterior parameters

```

1 for  $j$  in  $\mathbb{J}$  do
2   for  $c$  in  $\mathbb{C}$  do
3      $\eta^{c*} \sim p(\eta^c | \pi^c)$ 
4      $\tau^{c*} \sim p(\tau^c | \pi^c)$ 
5      $\epsilon^{c*} \sim p(\epsilon^c | \pi^c, \mathbf{x}^c)$ 
6   end
7    $\pi^* \sim p(\pi | \mathbf{x}, \eta, \tau, \epsilon)$ 
8 end

```

in CHMM: particle filtering (PF, Figure 2c) and factorized FFBS (fFFBS, Figure 2d). They (i) are efficient and scale to a large number of states and chains, (ii) avoid the conditioning on the latent variables of other chains, which reduces auto-correlation and improves mixing, and (iii) allow accurate likelihood estimation by integrating out the latent variables, which is crucial in M-CHMM. We demonstrate the usefulness of our contributions by analyzing a longitudinal multivariate dataset from a clinical trial concerning the efficiency of decolonization of Methicillin-resistant *Staphylococcus aureus* (MRSA) (Huang et al., 2019).

2. Background and Notation

In this section, we provide the necessary background information about CHMMs. We first show the model formulation and a generic MCMC algorithm for the CHMM. Then, we show how to convert a CHMM into an HMM and the algorithm for the HMM. Finally, we outline the current s-o-t-a algorithm iFFBS (Touloupou et al., 2020). Throughout, we use $c \in \mathbb{C} = \{1, \dots, C\}$ to denote a chain (shown as superscripts), where C is the number of chains. $k \in \mathbb{K} = \{1, \dots, K\}$ and $l \in \mathbb{L} = \{1, \dots, L\}$ denote sets of latent and observed states, where K, L are the numbers of values that the latent and observed states can take for each chain (assumed equal for different chains although this could be relaxed). Time steps are represented by $t \in \mathbb{T} = \{1, \dots, T\}$, where T is the number of time steps. Individuals are denoted by $n \in \mathbb{N} = \{1, \dots, N\}$, where N denotes the number of individuals. Subscripts represent the time step(s), except in the description of M-CHMM where they represent the individual(s). Additionally, η, τ , and ϵ denote initial state, transition, and emission probabilities, respectively. Joint set $\theta = \{\eta, \tau, \epsilon\}$ denotes all the parameters of the

Algorithm 2: Hidden Markov Model

Input: Prior of η, τ, ϵ ; Initialize π using priors

Output: Posterior parameters

```

1 for  $j$  in  $\mathbb{J}$  do
2    $\eta^* \sim p(\eta | \pi)$ 
3    $\tau^* \sim p(\tau | \pi)$ 
4    $\epsilon^* \sim p(\epsilon | \pi, \mathbf{x})$ 
5    $\pi^* \sim p(\pi | \mathbf{x}, \eta, \tau, \epsilon)$ 
6 end

```

system. Finally, $j \in \mathbb{J} = \{1, \dots, J\}$ is the iterations in the MCMC, where J is the number of iterations. We emphasize that except for the mixture model, we present the models and algorithms for one individual for simplicity.

2.1. Coupled Hidden Markov Model

A CHMM with a latent sequence $\pi, \pi_t^c \in \{1, \dots, K\}$ and observations $\mathbf{x}, x_t^c \in \{1, \dots, L\}$ is defined as

$$\begin{aligned}
p(\pi_{1:T}^{\mathbb{C}}, \mathbf{x}_{1:T}^{\mathbb{C}}) &= \left[\prod_{c \in \mathbb{C}} \underbrace{p(\pi_1^c)}_{\eta^c} \right] \left[\prod_{t=2}^T \prod_{c \in \mathbb{C}} \underbrace{p(\pi_t^c | \pi_{t-1}^c)}_{\tau^c} \right] \\
&\quad \times \left[\prod_{t=1}^T \prod_{c \in \mathbb{C}} \underbrace{p(x_t^c | \pi_t^c)}_{\epsilon^c} \right]. \tag{1}
\end{aligned}$$

Algorithm 1 shows a generic Gibbs sampler to learn CHMMs. The sampling of the latent sequences (step 7 in Algorithm 1) is important because (i) it is the key to the efficiency, and (ii) marginal likelihood can be estimated as a byproduct. Therefore, it is the primary focus of this paper.

2.2. CHMM as a joint-HMM

The CHMM is equivalent to an HMM with a latent sequence $\pi, \pi_t^{\mathbb{C}} \in \{1, \dots, K^{\mathbb{C}}\}$ and observations $\mathbf{x}, \mathbf{x}_t^{\mathbb{C}} \in \{1, \dots, L^{\mathbb{C}}\}$, which is defined as

$$\begin{aligned}
p(\pi_{1:T}^{\mathbb{C}}, \mathbf{x}_{1:T}^{\mathbb{C}}) &= p(\underbrace{\pi_1^{\mathbb{C}}}_{\eta}) \left[\prod_{t=2}^T \underbrace{p(\pi_t^{\mathbb{C}} | \pi_{t-1}^{\mathbb{C}})}_{\tau} \right] \\
&\quad \times \left[\prod_{t=1}^T \underbrace{p(\mathbf{x}_t^{\mathbb{C}} | \pi_t^{\mathbb{C}})}_{\epsilon} \right]. \tag{2}
\end{aligned}$$

Algorithm 2 shows the Gibbs sampler for the single HMM. An efficient way of updating the latent sequences (step 5 on Algorithm 2) is to use the *forward-filtering backward-sampling (FFBS)* algorithm.

2.3. FFBS

Forward-filtering is the estimation of the current hidden state by using all observations so far, $p(\boldsymbol{\pi}_t^c | \mathbf{x}_{1:t}^c)$ (Barber and Cemgil, 2010) given model parameters. By defining $\alpha(\boldsymbol{\pi}_t^c) = p(\boldsymbol{\pi}_t^c | \mathbf{x}_{1:t}^c)$ it yields the following recursive equation known as α -recursion:

$$\alpha(\boldsymbol{\pi}_t^c) \propto \underbrace{p(\mathbf{x}_t^c | \boldsymbol{\pi}_t^c)}_{\text{corrector}} \underbrace{\sum_{\boldsymbol{\pi}_{t-1}^c} p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t-1}^c) \alpha(\boldsymbol{\pi}_{t-1}^c)}_{\text{predictor}} \quad (3)$$

which is derived in Appendix A. The recursion starts with $\alpha(\boldsymbol{\pi}_1^c)$ and the filtered distribution $\alpha(\cdot)$ is propagated through to the next time-step, where it acts like a ‘prior’ (Barber and Cemgil, 2010; Barber, 2012). In Equation 3, the normalization constant of $\alpha(\boldsymbol{\pi}_t^c)$ is $Z_t = p(\mathbf{x}_t^c | \mathbf{x}_{1:t-1}^c)$, and the marginal likelihood is available as a byproduct of the forward-filtering:

$$\mathcal{L}(\mathbf{x}_{1:T}^c) \triangleq p(\mathbf{x}_{1:T}^c) = \prod_{t=1}^T Z_t. \quad (4)$$

The latent variables are sampled from:

$$p(\boldsymbol{\pi}_{1:T}^c | \mathbf{x}_{1:T}^c) = p(\boldsymbol{\pi}_T^c | \mathbf{x}_{1:T}^c) \prod_{t=1}^{T-1} p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t+1}^c, \mathbf{x}_{1:t}^c). \quad (5)$$

The posterior sampling in Equation 5 requires ‘time-reversed’ transitions $p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t+1}^c, \mathbf{x}_{1:t}^c)$, where

$$p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t+1}^c, \mathbf{x}_{1:t}^c) \propto p(\boldsymbol{\pi}_{t+1}^c | \boldsymbol{\pi}_t^c) \alpha(\boldsymbol{\pi}_t^c), \quad (6)$$

which are obtained using α -recursions from the forward filtering (Barber, 2012). The backward sampling starts at the end of the sequence and proceeds backwards recursively using the previously sampled latent state $\boldsymbol{\pi}_{t+1}^c = \hat{\boldsymbol{\pi}}_{t+1}^c$:

$$\hat{\boldsymbol{\pi}}_T^c \sim \alpha(\boldsymbol{\pi}_T^c), \quad (7)$$

$$\hat{\boldsymbol{\pi}}_t^c \sim p(\boldsymbol{\pi}_{t+1}^c = \hat{\boldsymbol{\pi}}_{t+1}^c | \boldsymbol{\pi}_t^c) \alpha(\boldsymbol{\pi}_t^c). \quad (8)$$

This procedure is known as forward-filtering backward-sampling (Barber, 2012), and Figure 2a illustrates this algorithm. The computational complexity for one time-step is of the order $\mathcal{O}((K^c)^2)$.

2.4. Individual FFBS (iFFBS)

Since the chains are interconnected in a CHMM, applying the FFBS separately (not jointly) on each chain will not be valid for the CHMM. It is still usable (Sherlock et al., 2013), but it will underestimate the inter-chain effects and fail if those are strong. Touloupou et al. (2020) reformulated the FFBS algorithm for a

single chain by conditioning the updates on the states of the other chains, resulting in a proper posterior sampling. The algorithm (iFFBS) is illustrated in Figure 2b. By defining $\alpha(\boldsymbol{\pi}_t^c) = p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{1:t+1}^{-c}, \mathbf{x}_{1:t}^c)$, where $\neg c = \mathbb{C} \setminus c$ refers to all the chains except c itself, the modified α -recursion is as follows:

$$\begin{aligned} \alpha(\boldsymbol{\pi}_t^c) \propto & \underbrace{p(\mathbf{x}_t^c | \boldsymbol{\pi}_t^c)}_{\text{corrector}} \underbrace{\sum_{\boldsymbol{\pi}_{t-1}^c} p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t-1}^c) \alpha(\boldsymbol{\pi}_{t-1}^c)}_{\text{predictor}} \\ & \times \underbrace{\prod_{\hat{c} \in \mathbb{C} \setminus c} p(\boldsymbol{\pi}_{t+1}^{\hat{c}} | \boldsymbol{\pi}_t^c)}_{\text{modifying mass}}. \end{aligned} \quad (9)$$

Here, the *modifying mass* is the only additional term compared to the FFBS for a single chain. Intuitively, it corresponds to the effect of the updated chain on the other chains in the next time step (i.e., outgoing edges from the variable of interest). The marginal likelihood of the whole system is not directly available as in the FFBS algorithm (Touloupou et al., 2020); here we used a heuristic way to approximate it for each chain separately using forward filtering conditionally on the other chains, and take the sum. Details of this approximation is given in Appendix B.

Respectively, the posterior of the latent sequences for each chain conditionally on other chains is:

$$\begin{aligned} p(\boldsymbol{\pi}_{1:T}^c | \boldsymbol{\pi}_{1:T}^{-c}, \mathbf{x}_{1:T}^c) &= p(\boldsymbol{\pi}_T^c | \boldsymbol{\pi}_{1:T}^{-c}, \mathbf{x}_{1:T}^c) \\ &\times \prod_{t=1}^{T-1} p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t+1}^c, \boldsymbol{\pi}_{1:t+1}^{-c}, \mathbf{x}_{1:t}^c), \end{aligned} \quad (10)$$

where the ‘time-reversed’ transitions satisfy:

$$p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t+1}^c, \boldsymbol{\pi}_{1:t+1}^{-c}, \mathbf{x}_{1:t}^c) \propto p(\boldsymbol{\pi}_{t+1}^c | \boldsymbol{\pi}_t^c, \boldsymbol{\pi}_t^{-c}) \alpha(\boldsymbol{\pi}_t^c). \quad (11)$$

Therefore, this algorithm allows using the backward sampling procedure described in Equation 7 and Equation 8 for each chain given the others (for details and derivations, see Appendix B and Touloupou et al. (2020)). The computational complexity for one time-step is $\mathcal{O}(K^2C)$. This approach is useful yet creates correlated latent sequence samples because of the dependency on the other chains during sampling, which decreases sampling efficiency.

3. Methods

In this section, we first formally define the Mixture of CHMMs (M-CHMM) model. Then we explain how we parametrize the transition matrix in the CHMM. Finally, we propose two novel latent sequence samplers which allow us to use M-CHMM more accurately, and one of which exploits the properties of the design of the transition matrix.

Algorithm 3: Mixture of CHMMs

Input: Prior of $\gamma, \eta, \tau, \epsilon$; Initialize π using priors

Output: Posterior parameters

```

1 for  $j$  in  $\mathbb{J}$  do
2   for  $m$  in  $\mathbb{M}$  do
3     for  $c$  in  $\mathbb{C}$  do
4        $\eta_m^{c*} \sim p(\eta_m^c \mid \pi^c, z_n = m)$ 
5        $\tau_m^{c*} \sim p(\tau_m^c \mid \pi, z_n = m)$ 
6        $\epsilon_m^{c*} \sim p(\epsilon_m^c \mid \pi^c, \mathbf{x}^c, z_n = m)$ 
7     end
8   end
9    $\mathbf{z}^* \sim p(\mathbf{z} \mid \mathbf{x}, \eta, \tau, \epsilon, \gamma)$ 
10   $\gamma^* \sim p(\gamma \mid \mathbf{z})$ 
11   $\pi^* \sim p(\pi \mid \mathbf{x}, \eta, \tau, \epsilon, \mathbf{z})$ 
12 end

```

3.1. Mixture of CHMMs (M-CHMM)

We propose the mixture of CHMMs as a flexible model for multivariate time-series data:

$$p(\mathbf{x}_n) = \sum_{m=1}^M \gamma_m \mathcal{L}_{\text{CHMM}}(\mathbf{x}_n \mid \theta_m^{\text{CHMM}}), \quad (12)$$

where the mixing coefficients γ satisfy $\sum_{m=1}^M \gamma_m = 1$ and $0 \leq \gamma_m \leq 1$. $\mathcal{L}_{\text{CHMM}}(\mathbf{x}_n \mid \theta_m^{\text{CHMM}})$ denotes the marginal likelihood for the m th component distribution (π_n is integrated out). For inference, we augment the model with cluster labels $z_n \in \mathbb{M} = \{1, \dots, M\}$, yielding an equivalent formulation:

$$p(\mathbf{x}_n \mid z_n) = \mathcal{L}_{\text{CHMM}}(\mathbf{x}_n \mid \theta_{z_n}^{\text{CHMM}}), \text{ and} \quad (13)$$

$$z_n \sim \text{Categorical}(\gamma), \quad \forall n \in \mathbb{N}. \quad (14)$$

The graphical model of the M-CHMM is shown in Figure 1b, and Algorithm 3 shows a generic Gibbs sampler. Note that updating the cluster labels z_n requires evaluating the component distributions in Equation 13 for each alternative cluster.

3.2. Design of Transition Matrix in CHMM

Defining the transition matrix is a critical part of the CHMM since all the heavy computation relies on it. There is a trade-off between sparsity, which can lead to an efficient implementation (Sherlock et al., 2013), and the wish to carry as much information about the other chains as possible. We follow the formulation of Poyraz et al. (2022) and model the dependencies between the chains with parameters β , where $\beta^c \in \beta$ denotes all parameters related to a target chain c . Further, each $\beta_k^{c \leftarrow \hat{c}} \in \beta^c$ is a matrix of the same size

as the transition matrix describing the impact the chain \hat{c} has on the target chain c when \hat{c} is in state k . We assume that the rows of $\beta_k^{c \leftarrow \hat{c}}$ sum to zero, which removes redundancy and allows a one-to-one mapping between β and τ . We assume that there is a baseline state k_0 for each covariate chain \hat{c} , such that if the chain \hat{c} is in that state it does not affect other chains, i.e., $\beta_{k_0}^{c \leftarrow \hat{c}} = \mathbf{0}$. The transition matrix for the chain c , at time t , denoted by τ_t^c , is defined as

$$\mu_t^c = \beta_0^{c \leftarrow c} + \sum_{\hat{c} \in \mathbb{C} \setminus c} \sum_{k \in \mathbb{K}} \beta_k^{c \leftarrow \hat{c}} \mathbb{I}[\pi_{t-1}^{\hat{c}} = k], \quad (15)$$

$$\tau_t^c = \sigma_{\text{row}}(\mu_t^c). \quad (16)$$

Hence, τ_t^c is obtained in Equation 16 by applying a row-wise softmax operator σ_{row} to the unnormalized transition matrix μ_t^c . Parameter $\beta_0^{c \leftarrow c}$ corresponds to an intercept, and it specifies the transition matrix of the target chain c when all other chains are in the baseline state k_0 , i.e., the probabilities $p(\pi_t^c \mid \pi_{t-1}^c, \pi_{t-1}^{\hat{c}} = k_0 \forall \hat{c} \neq c)$. Parameters $\beta_k^{c \leftarrow \hat{c}}$ that represent the impact of the other chains on the target chain c are added to $\beta_0^{c \leftarrow c}$. The design of the transition matrix in Equation 15 and Equation 16 yields the following factorization over chains:

$$p(\pi_t^c \mid \pi_{t-1}^c) \propto f_0(\pi_t^c, \pi_{t-1}^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} f_{\pi_{t-1}^{\hat{c}}}(\pi_t^c, \pi_{t-1}^c), \quad (17)$$

where each factor f uses the previous and next time step of target chain (π_t^c, π_{t-1}^c) as inputs and is a function of the latent states of other chains $\pi_{t-1}^{\hat{c}}$, except for f_0 which corresponds to an intercept. The derivation is in Appendix C. β parameters are updated with a Metropolis-Hasting step within the Gibbs sampler as proposed by Sherlock et al. (2013), see details in Appendix E. All methods in this paper are implemented using this trick for a fair comparison. Detailed prior distributions are given in the Appendix F.

3.3. Latent Sequence Samplers

Sampling the latent sequences is a crucial part of the MCMC inference for the CHMM, as it is the bottleneck in terms of scalability. Here, we propose two novel latent sequence samplers, which are scalable and provides an unbiased estimation of the component distribution required in M-CHMM.

3.3.1. PARTICLE FILTER FOR CHMM (PF)

Here we formulate particle filter to approximate the posterior distribution of the latent sequence, $p(\pi_{1:T}^c \mid \mathbf{x}_{1:T}^c)$, by sequentially approximating distributions, $p(\pi_{1:t}^c \mid \mathbf{x}_{1:t}^c)$, for each t . To approximate

the posterior, particle filtering places a weighted point mass at the locations of P samples (Cemgil, 2014), or particles, $\pi_{1:t}^{c,p}$, $p \in \{1, \dots, P\}$:

$$p\left(\pi_{1:t}^c \mid \mathbf{x}_{1:t}^c\right) \approx \hat{p}\left(\pi_{1:t}^c \mid \mathbf{x}_{1:t}^c\right) \triangleq \sum_{p=1}^P \frac{w_t^p}{\sum_l w_t^l} \delta_{\pi_{1:t}^{c,p}}. \quad (18)$$

The weighted set of particles is constructed sequentially for $t = 1, \dots, T$. At time $t = 1$, the standard importance sampling is applied and samples are drawn from a proposal distribution, $\pi_1^{c,p} \sim q(\pi_1^c) = \prod_{c \in \mathbb{C}} q(\pi_1^c)$. For $t > 1$, the optimal proposal distribution $q(\pi_t^c) \propto p(\pi_t^c \mid \pi_{t-1}^c) p(\mathbf{x}_t^c \mid \pi_t^c)$ is used. Also, for $t > 1$, the systematic resampling step is applied with a probability proportional to the importance weights w_{t-1}^p if the effective sample size (ESS) is below a given threshold (here 0.5, see Doucet et al. (2009)). The whole procedure is as follows:

$$\text{resample} \quad a_{t-1}^p \sim \text{Cat}\left(w_{t-1}^p / \sum_k w_{t-1}^k\right) \quad (19)$$

$$\text{propose} \quad \pi_t^{c,p} \sim q\left(\pi_t^c \mid \pi_{t-1}^{c,a_{t-1}^p}\right) \quad (20)$$

$$\text{append} \quad \pi_{1:t}^{c,p} = \left(\pi_{1:t-1}^{c,a_{t-1}^p}, \pi_t^{c,p}\right) \quad (21)$$

$$\text{reweight} \quad w_t^p = \frac{p\left(\pi_t^c \mid \pi_{t-1}^{c,a_{t-1}^p}\right) p\left(\mathbf{x}_t^c \mid \pi_t^{c,p}\right)}{q\left(\pi_t^c \mid \pi_{t-1}^{c,a_{t-1}^p}\right)} \quad (22)$$

The final particles (samples) $\pi_{1:T}^{c,p}$ and weights w_T^p define the particle filter approximation to the posterior. An unbiased estimate of the marginal likelihood can be calculated as a side-product (Naesseth et al., 2018);

$$\mathcal{L}\left(\mathbf{x}_{1:T}^c\right) \triangleq \hat{p}\left(\mathbf{x}_{1:T}^c\right) = \prod_{t=1}^T \frac{1}{P} \sum_{p=1}^P w_t^p \quad (23)$$

The algorithm is illustrated in Figure 2c and the computational complexity for one time-step is $\mathcal{O}(KCP)$.

3.3.2. FACTORIZED FFBS (FFFBS)

The complexity of the FFBS for the single HMM comes from the fact that $\alpha(\pi_t^c)$ in Equation 3 depends on all the latent variables, requiring a summation over all combinations. To overcome this, we assume the following factorization:

$$p\left(\pi_t^c \mid \mathbf{x}_{1:t}^c\right) \approx \prod_{c \in \mathbb{C}} p\left(\pi_t^c \mid x_t^c, \mathbf{x}_{1:t-1}^c\right), \quad (24)$$

for all t . Equation 24 does not hold exactly because previous π_{t-1}^c make π_t^c dependent even conditionally on observations $\mathbf{x}_{1:t-1}^c$ from all chains. However, it is accurate when within-chain dependencies outweigh

between-chain dependencies, or when the emission probabilities reflect the underlying state accurately in which case the information in π_{t-1}^c can be extracted from observations $\mathbf{x}_{1:t-1}^c$. If we define $\alpha(\pi_t^c) = p(\pi_t^c \mid x_t^c, \mathbf{x}_{1:t-1}^c)$, then $\alpha(\pi_t^c) = \prod_{c \in \mathbb{C}} \alpha(\pi_t^c)$. The approximated α -recursion for each chain is given by:

$$\begin{aligned} \alpha(\pi_t^c) &\propto \underbrace{p\left(x_t^c \mid \pi_t^c\right)}_{\epsilon^c} \sum_{\pi_{t-1}^c} \alpha(\pi_{t-1}^c) \\ &\times \underbrace{\sum_{\pi_{t-1}^{c \setminus c}} p\left(\pi_t^c \mid \pi_{t-1}^c\right)}_{\tau_t^c} \prod_{\hat{c} \in \mathbb{C} \setminus c} \alpha(\pi_{t-1}^{\hat{c}}), \end{aligned} \quad (25)$$

as derived in Appendix D. Here, by using Equation 24, the summation over all latent state combinations in Equation 3 is re-arranged, such that for each chain there is a summation over the latent states of the other chains inside a dynamic transition matrix τ_t^c . So, τ_t^c requires the summation over $C-1$ latent states of the other chains, which equals to K^{C-1} summations. However, we can express any τ_t^c as following by exploiting the proportionality in Equation 17:

$$\tau_t^c = f_0(\pi_t^c, \pi_{t-1}^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} \sum_{\pi_{t-1}^{\hat{c}}} f_{\pi_{t-1}^{\hat{c}}}(\pi_t^c, \pi_{t-1}^{\hat{c}}) \alpha(\pi_{t-1}^{\hat{c}}), \quad (26)$$

which is derived in Appendix D. By plugging Equation 26 into Equation 25, the big summation over all latent state combinations in Equation 3 is separated into independent summations for each chain. Therefore, we only need KC summations instead of K^C , which makes fFFBS scalable to a high number of states and chains. In Equation 25, the normalization constant of $\alpha(\pi_t^c)$ is $Z_t^c = p(\mathbf{x}_t^c \mid \mathbf{x}_{1:t-1}^c)$, and the marginal likelihood is a byproduct of the forward-filtering under the factorization assumption:

$$\mathcal{L}\left(\mathbf{x}_{1:T}^c\right) \triangleq p\left(\mathbf{x}_{1:T}^c\right) = \prod_{t=1}^T \prod_{c=1}^C Z_t^c. \quad (27)$$

The latent variables are sampled from:

$$p\left(\pi_{1:T}^c \mid \mathbf{x}_{1:T}^c\right) = p\left(\pi_T^c \mid \mathbf{x}_{1:T}^c\right) \prod_{t=1}^{T-1} p\left(\pi_t^c \mid \pi_{t+1}^c, \mathbf{x}_{1:t}^c\right), \quad (28)$$

where the ‘time-reversed’ transitions satisfy:

$$\begin{aligned} p\left(\pi_t^c \mid \pi_{t+1}^c, \mathbf{x}_{1:t}^c\right) \\ \propto \prod_{c \in \mathbb{C}} \alpha(\pi_t^c) f_0(\pi_{t+1}^c, \pi_t^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} f_{\pi_{t+1}^{\hat{c}}}(\pi_{t+1}^{\hat{c}}, \pi_t^{\hat{c}}), \end{aligned} \quad (29)$$

which is derived in Appendix D. Similar to the FFBS, the backward sampling starts at the end $\hat{\pi}_T^c \sim \alpha(\pi_T^c)$, $\forall c \in \mathbb{C}$, and proceeds recursively backwards using Equation 29 with previously sampled $\hat{\pi}_{t+1}^c$, as illustrated in Figure 2d. The computational complexity for one time-step is $\mathcal{O}(K^3C)$.

4. Experiments and Results

In this section, we first introduce datasets used in experiments and our experimental setup. Then, we show that our proposed sampling algorithms for CHMM outperform existing alternatives in terms of stability, computational efficiency, and ability to detect known clusters using the M-CHMM. Finally, we show how using the M-CHMM instead of the standard CHMM improves model fit and prediction accuracy with complex real-world data.

4.1. Datasets and Experimental Setup

We used a randomized controlled trial (RCT) dataset from the CLEAR (Changing Lives by Eradicating Antibiotic Resistance) (Huang et al., 2019). The trial subjects were randomized into two groups (*control*: $n = 1063$ and *treatment*: $n = 1058$) to test the efficiency of a treatment protocol to clear the bacterium which colonizes the patient in different body parts (nares, throat, and skin, i.e., chains). Measurements indicating the presence or absence of a bacterium in the body parts were collected from the subjects at hospital discharge (t_0) and on three follow-up visits 1, 3, and 6 months after the discharge (t_1, t_3 , and t_6). Measurements were noisy, and some were missing. Overall, 19671 measurements (10039 and 9632 from the control and treatment) were received. We treated skipped months (t_2, t_4 , and t_5) as additional missing observations. The dimensions of latent and observed variables were $K = L = 2$, where the latent chains corresponded to whether the bacterium was truly present and observations were the noisy measurements. Hence, emission parameters give us the sensitivity and specificity of the measurements.

We created two sets of semi-synthetic datasets using the CLEAR data. For the first set (SS-1), we trained a single HMM on each trial group using the FFBS algorithm and simulated synthetic datasets with different numbers of time points (5, 7, 10, 20, 50, and 100) using the trained models. The SS-1 is the most realistic regeneration of the original CLEAR dataset. The second set (SS-2) was simulated from two CHMMs trained on control and treatment groups, respectively, using the iFFBS algorithm (not to create any advantage to our proposed models), setting the sequence length to 7 to mimic the original data. Since this dataset is created using the CHMM formulation, we can check whether CHMM training with different samplers can learn the true underlying parameters (β -parameters described in the section 3). The posterior

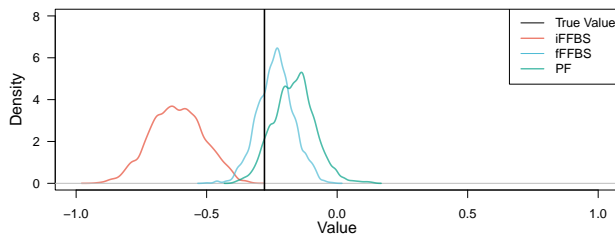


Figure 3: Estimated posterior distributions for one of the β parameters with different latent sequence samplers.

predictive checkings for the data-generator models for SS-1 and SS-2 are given in the Appendix G. The semi-synthetic datasets always had 1000 individuals in each group. The single HMM formulation of the CHMM and the state-of-the-art iFFBS were used as baseline samplers.

4.2. Comparison of Latent Sequence Samplers

Stability. We conducted this analysis on the SS-2 dataset. First, we compared the error in forward filtering probabilities for different algorithms, considering the single HMM as the ground truth (Appendix G). Here, we run FFBS and fFFBS just once since they can calculate forward probabilities deterministically. However, PF and iFFBS are based on samples, so they were averaged over 1000 iterations. iFFBS results in fluctuating probabilities, while the fFFBS algorithm has a slight bias due to the factorization assumption. On the other hand, PF converges to the true forward filtering probabilities. Next, we compared the estimated posterior distributions for the transition parameters (β -parameters) on SS-2 data; one of the parameters is shown in Figure 3 and the rest in Appendix G. We see that fFFBS and PF are more accurate than iFFBS and that some posteriors are shrunk towards 0 because of the prior, as expected. We note that the FFBS algorithm is excluded since it is inapplicable for this test.

Computational Efficiency. We also compared all algorithms in terms of runtime (Table 1) and effective sample size (in Appendix G). This analysis shows that the single HMM is intractable when the number of chains or latent state dimension increases. The fFFBS and iFFBS algorithms scale approximately similarly when the dimensions increase. On the other hand, PF is much more efficient if the number of particles stays

Table 1: Time comparison with different latent state and chain numbers. For each setup, sequence length $T=20$, a number of iterations $J=100$, and particle number $P=10$ are selected.

| Setup | FFBS | iFFBS | fFFBS | PF |
|-------------|------|-------|-------|------|
| $K=2, C=4$ | 1.2s | 5.7s | 8.6s | 8.6s |
| $K=4, C=4$ | 90s | 13s | 25s | 10s |
| $K=2, C=8$ | 92s | 20s | 42s | 21s |
| $K=4, C=8$ | NA | 52s | 132s | 25s |
| $K=8, C=4$ | NA | 40s | 120s | 13s |
| $K=2, C=16$ | NA | 109s | 239s | 63s |
| $K=8, C=8$ | NA | 200s | 650s | 35s |
| $K=4, C=16$ | NA | 337s | 775s | 83s |
| $K=8, C=16$ | NA | 1313s | 3826s | 154s |

the same. However, a larger number of particles may be helpful with a higher dimension.

Ability to detect known clusters with M-CHMM. Here, we compared the inference algorithms with the M-CHMM in an unsupervised task by measuring how well the model learns the two clusters in the semi-synthetic SS-1 data corresponding to the control and treatment groups. The results in Figure 4 show that increasing the length of the observed sequence enhances accuracy in general, and that our proposed samplers (fFFBS and PF) achieve better and more consistent results than the alternatives. This is expected because (i) compared to the iFFBS, they can accurately estimate the likelihoods of the component models required when updating the cluster labels, and (ii) compared to the FFBS, they benefitted from the structure within the latent sequence, which made the inference more efficient and stable especially when the sequence length is insufficient.

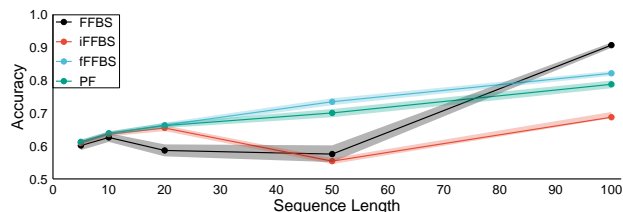


Figure 4: Clustering accuracy with different sequence lengths. Solid lines represents the mean and shaded areas the 95% confidence interval.

Table 2: Comparison of models with negative log-likelihood (NLL) in the test set for control and treatment groups using 5-fold cross-validation (CV). A lower score is better.

| Model | Control | Treatment |
|----------------|--------------------------------|--------------------------------|
| CHMM | 999 ± 44 | 812 ± 42 |
| M-CHMM (M = 2) | 855 ± 50 | 663 ± 31 |
| M-CHMM (M = 3) | 808 ± 48 | 585 ± 32 |
| M-CHMM (M = 4) | 779 ± 48 | 599 ± 28 |
| M-CHMM (M = 5) | 772 ± 46 | 588 ± 19 |
| M-CHMM (M = 6) | 758 ± 50 | 618 ± 34 |
| M-CHMM (M = 7) | 765 ± 55 | 609 ± 36 |
| M-CHMM (M = 8) | 755 ± 47 | 623 ± 47 |
| M-CHMM (M = 9) | 779 ± 38 | 626 ± 36 |

4.3. Modeling the real-world data with the M-CHMM

Here, we modeled the control and treatment groups separately with the M-CHMM to understand what kind of different patient subgroups they included. From this point onwards, we only used PF for inference since it was the most accurate and scalable in the previous experiments. Table 2 compares different numbers of clusters using negative log-likelihood (NLL) scores in the test set. The best predictive accuracy was obtained with 8 clusters in the control group and 3 clusters in the treatment group. Figure 5 shows the results of posterior predictive checking for CHMM and M-CHMM with the optimal number of clusters. We see that the CHMM fails to model the decrease in the probability of bacterial colonization in the different body parts over time, but the M-CHMM accurately captures it and successfully regenerates the observations. This demonstrates the flexibility of the M-CHMM to capture complex patterns in real-world time-series data by identifying groups of individuals with different dynamics.

Next, we investigated the learned clusters within the treatment group (the control group is analyzed in Appendix G). Figure 6 clearly shows one cluster of patients who responded well to the treatment (probability of bacterial colonization dropped rapidly) and another cluster where the bacterium was persistent (probability stayed constantly high), which reflects the medical understanding of Huang et al. (2019). The third cluster included patients colonized with a low probability already in the initial phase. Figure 7 show

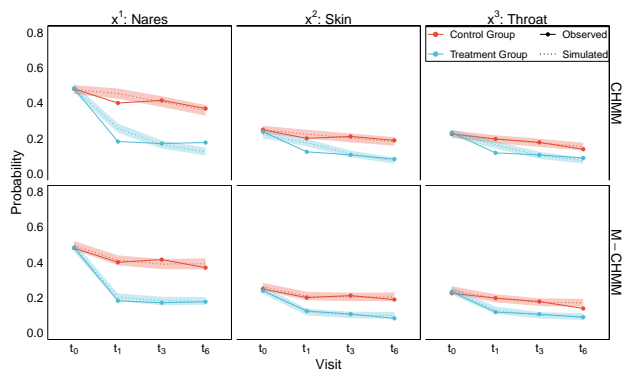


Figure 5: Posterior predictive checking of bacterial colonization for CHMM, and M-CHMM.

the sensitivity and specificity of the measurements. Overall, similar values compatible with literature are estimated for different clusters (Carr et al., 2018), with a slight shift of cluster m_2 , for which a possible explanation is that the treatment may affect the measurement accuracy (Carr et al., 2018).

5. Discussion

We proposed the M-CHMM as a flexible extension of the CHMM for analyses of nonexchangeable multivariate data from longitudinal clinical trials by identifying clusters of patients with different temporal dynamics. The CHMMs are ideal for systems consisting of multiple interacting latent processes where the aim is to learn about the dynamics of the system and interactions between its components. Many systems in epidemiology and healthcare match this description and, consequently, CHMMs have previously been found useful in multiple diverse biomedical applications: for disease co-morbidities (Maag et al., 2021), for hourly measurements of multiple vital signs

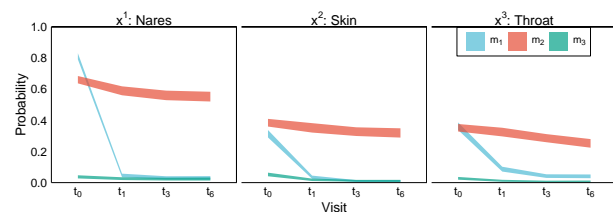


Figure 6: The dynamics of the clusters detected by M-CHMM with 3 clusters.

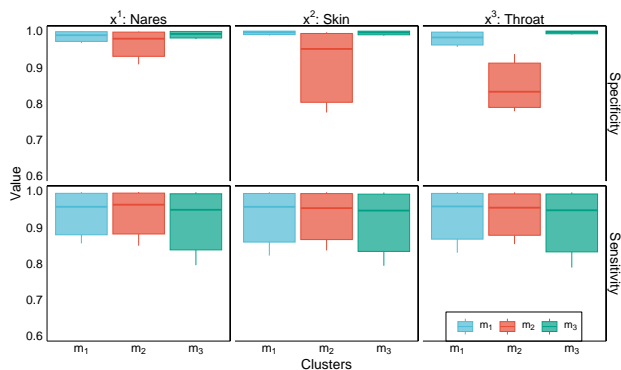


Figure 7: The specificity and sensitivity analysis of measurements for 3 clusters.

of ICU patients (Pohle et al., 2021), for accounting for genetic relatedness between individuals when detecting genetic copy-number variants jointly across individuals (Wang et al., 2019), for disease transmission between individuals (Touloupou et al., 2020) or between body sites of an individual (Poyraz et al., 2022), and for interactions between parasite species in a host (Sherlock et al., 2013). Our extension to mixtures of CHMMs is expected to be useful in all situations where the CHMMs are applicable, by providing increased flexibility. Furthermore, different systems, e.g., time-series from different patients, may not behave in a similar way, e.g., due to differences in treatment response. In such situations, the M-CHMM can identify differently behaving patients and thereby learn patterns that may provide insight into the data, as well as improve model fit and prediction accuracy, which can be beneficial in downstream use cases.

During our investigation, we found existing algorithms to sample the latent sequence of the CHMM either computationally demanding or not providing a way to estimate the likelihood, which prevented their robust application with M-CHMMs. To resolve this, we formulated two latent sequence samplers for CHMMs, and showed that they are more efficient and scalable than existing methods, and allow accurate likelihood estimation, facilitating inference with more complex models. By using them, we demonstrated the ability of the M-CHMM to model complex real-world multivariate time-series data with noisy and missing measurements. To facilitate a wide adoption, we provide the method as an R-package with C-code optimization of computationally heavy parts.¹

1. Available at <https://github.com/onurpoyraz/M-CHMM>.

Acknowledgments

The authors gratefully acknowledge Susan S. Huang for providing the CLEAR (Changing Lives by Eradicating Antibiotic Resistance) dataset to be used in this study (available at doi:[10.6084/m9.figshare.19786678.v4](https://doi.org/10.6084/m9.figshare.19786678.v4)). This work was supported by the Academy of Finland (Flagship programme: Finnish Center for Artificial Intelligence FCAI, and grants 336033, 352986) and EU (H2020 grant 101016775 and NextGenerationEU).

References

- Asem Alaa, Erik Mayer, and Mauricio Barahona. Ice-node: Integration of clinical embeddings with neural ordinary differential equations. In *Machine Learning for Healthcare Conference*, pages 537–564. PMLR, 2022. doi:[10.48550/arXiv.2207.01873](https://doi.org/10.48550/arXiv.2207.01873).
- David Barber. *Bayesian reasoning and machine learning*. Cambridge University Press, 2012.
- David Barber and A Taylan Cemgil. Graphical models for time-series. *IEEE Signal Processing Magazine*, 27(6):18–28, 2010. doi:[10.1109/MSP.2010.938028](https://doi.org/10.1109/MSP.2010.938028).
- Christopher M Bishop. *Pattern recognition and machine learning*. Springer Science+ Business Media, 2006.
- Matthew Brand. Coupled hidden Markov models for modeling interacting processes, 1997.
- Matthew Brand, Nuria Oliver, and Alex Pentland. Coupled hidden Markov models for complex action recognition. In *Proceedings of IEEE computer society conference on computer vision and pattern recognition*, pages 994–999. IEEE, 1997. doi:[10.1109/CVPR.1997.609450](https://doi.org/10.1109/CVPR.1997.609450).
- Amy L Carr, Mitchell J Daley, Kathryn Givens Merkel, and Dusten T Rose. Clinical utility of methicillin-resistant staphylococcus aureus nasal screening for antimicrobial stewardship: a review of current literature. *Pharmacotherapy: The Journal of Human Pharmacology and Drug Therapy*, 38(12):1216–1228, 2018. doi:[10.1002/phar.2188](https://doi.org/10.1002/phar.2188).
- A Taylan Cemgil. A tutorial introduction to monte carlo methods, markov chain monte carlo and particle filtering. *Academic Press Library in Signal Processing*, 1:1065–1114, 2014. doi:[10.1016/B978-0-12-396502-8.00019-X](https://doi.org/10.1016/B978-0-12-396502-8.00019-X).
- Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F Stewart, and Jimeng Sun. Doctor ai: Predicting clinical events via recurrent neural networks. In *Machine learning for healthcare conference*, pages 301–318. PMLR, 2016. doi:[10.48550/arXiv.1511.05942](https://doi.org/10.48550/arXiv.1511.05942).
- Sanjoy Dasgupta and Daniel Hsu. On-line estimation with the multivariate gaussian distribution. In *International Conference on Computational Learning Theory*, pages 278–292. Springer, 2007. doi:[10.1007/978-3-540-72927-3_21](https://doi.org/10.1007/978-3-540-72927-3_21).
- Wen Dong, Alex Pentland, and Katherine A Heller. Graph-coupled hmms for modeling the spread of infection. *arXiv preprint arXiv:1210.4864*, 2012. doi:[10.48550/arXiv.1210.4864](https://doi.org/10.48550/arXiv.1210.4864).
- Arnaud Doucet, Adam M Johansen, et al. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of nonlinear filtering*, 12(656-704): 3, 2009.
- Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*. CRC press, 2013.
- Susan S Huang, Raveena Singh, James A McKinnell, Steven Park, Adrijana Gombosy, Samantha J Eells, Daniel L Gillen, Diane Kim, Syma Rashid, Raul Macias-Gil, et al. Decolonization to reduce postdischarge infection risk among MRSA carriers. *New England Journal of Medicine*, 380(7):638–650, 2019. doi:[10.1056/NEJMoa1716771](https://doi.org/10.1056/NEJMoa1716771).
- Rohan Kodialam, Rebecca Boiarsky, Justin Lim, Neil Dixit, Aditya Sai, and David Sontag. Deep contextual clinical prediction with reverse distillation. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*, 2021. doi:[10.1609/aaai.v35i1.16099](https://doi.org/10.1609/aaai.v35i1.16099).
- Basil Maag, Stefan Feuerriegel, Mathias Kraus, Maytal Saar-Tsechansky, and Thomas Züger. Modeling longitudinal dynamics of comorbidities. In *Proceedings of the Conference on Health, Inference, and Learning*, pages 222–235, 2021. doi:[10.1145/3450439.3451871](https://doi.org/10.1145/3450439.3451871).
- Kevin P Murphy. *Machine learning: A probabilistic perspective*. MIT press, 2012.
- Christian Naesseth, Scott Linderman, Rajesh Ranganath, and David Blei. Variational sequential

- monte carlo. In Amos Storkey and Fernando Perez-Cruz, editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 968–977. PMLR, 09–11 Apr 2018. doi:[10.48550/arXiv.1705.11140](https://doi.org/10.48550/arXiv.1705.11140).
- Jennifer Pohle, Roland Langrock, Mihaela van der Schaar, Ruth King, and Frants Havmand Jensen. A primer on coupled state-switching models for multiple interacting time series. *Statistical Modelling*, 21(3):264–285, 2021. doi:[10.1177/1471082X20956423](https://doi.org/10.1177/1471082X20956423).
- Onur Poyraz, Mohamad RA Sater, Loren G Miller, James A McKinnell, Susan S Huang, Yonatan H Grad, and Pekka Marttinen. Modelling methicillin-resistant staphylococcus aureus decolonization: interactions between body sites and the impact of site-specific clearance. *Journal of the Royal Society Interface*, 19(191):20210916, 2022. doi:[10.1098/rsif.2021.0916](https://doi.org/10.1098/rsif.2021.0916).
- Lawrence R Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989. doi:[10.1109/5.18626](https://doi.org/10.1109/5.18626).
- Ilead Rezek, Peter Sykacek, and Stephen J Roberts. Learning interaction dynamics with coupled hidden Markov models. *IEE Proceedings-Science, Measurement and Technology*, 147(6):345–350, 2000. doi:[10.1049/ip-smt:20000851](https://doi.org/10.1049/ip-smt:20000851).
- Gareth O Roberts, Jeffrey S Rosenthal, et al. Optimal scaling for various Metropolis-Hastings algorithms. *Statistical science*, 16(4):351–367, 2001. doi:[10.1214/ss/1015346320](https://doi.org/10.1214/ss/1015346320).
- Lawrence K Saul and Michael I Jordan. Mixed memory markov models: Decomposing complex stochastic processes as mixtures of simpler ones. *Machine learning*, 37(1):75–87, 1999. doi:[10.1023/A:1007649326333](https://doi.org/10.1023/A:1007649326333).
- Chris Sherlock, Paul Fearnhead, Gareth O Roberts, et al. The random walk Metropolis: Linking theory and practice through a case study. *Statistical Science*, 25(2):172–190, 2010. doi:[10.1214/10-STS327](https://doi.org/10.1214/10-STS327).
- Chris Sherlock, Tatiana Xifara, Sandra Telfer, and Mike Begon. A coupled hidden Markov model for disease interactions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 62(4):609–627, 2013. doi:[10.1111/rssc.12015](https://doi.org/10.1111/rssc.12015).
- Simon EF Spencer, Thomas E Besser, Rowland N Cobbold, and Nigel P French. ‘super’ or just ‘above average’? supershedders and the transmission of escherichia coli o157: H7 among feedlot cattle. *Journal of the Royal Society Interface*, 12(110):20150446, 2015. doi:[10.1098/rsif.2015.0446](https://doi.org/10.1098/rsif.2015.0446).
- Panayiota Touloupou, Bärbel Finkenstädt, and Simon EF Spencer. Scalable Bayesian inference for coupled hidden Markov and semi-Markov models. *Journal of Computational and Graphical Statistics*, 29(2):238–249, 2020. doi:[10.1080/10618600.2019.1654880](https://doi.org/10.1080/10618600.2019.1654880).
- Xiaoqiang Wang, Emilie Lebarbier, Julie Aubert, and Stéphane Robin. Variational inference for coupled hidden markov models applied to the joint detection of copy number variations. *The international journal of biostatistics*, 15(1), 2019. doi:[10.1515/ijb-2018-0023](https://doi.org/10.1515/ijb-2018-0023).
- Shi Zhong and Joydeep Ghosh. A new formulation of coupled hidden Markov models. *Dept. Elect. Comput. Eng., Univ. Austin, Austin, TX, USA*, 2001.

Appendix A. Derivation of FFBS on CHMM

If we define $\alpha(\boldsymbol{\pi}_t^{\mathbb{C}}) = p(\boldsymbol{\pi}_t^{\mathbb{C}} | \mathbf{x}_{1:t}^{\mathbb{C}})$, the full α -recursion or *forward recursion* of the single HMM in which all chains are handled jointly is derived as follows;

$$\alpha(\boldsymbol{\pi}_t^{\mathbb{C}}) = \sum_{\boldsymbol{\pi}_{t-1}^{\mathbb{C}}} p(\boldsymbol{\pi}_{t-1}^{\mathbb{C}}, \boldsymbol{\pi}_t^{\mathbb{C}} | \mathbf{x}_{1:t}^{\mathbb{C}}) \quad (30)$$

$$\propto \sum_{\boldsymbol{\pi}_{t-1}^{\mathbb{C}}} p(\mathbf{x}_t^{\mathbb{C}} | \boldsymbol{\pi}_t^{\mathbb{C}}) p(\boldsymbol{\pi}_t^{\mathbb{C}} | \boldsymbol{\pi}_{t-1}^{\mathbb{C}}) p(\boldsymbol{\pi}_{t-1}^{\mathbb{C}} | \mathbf{x}_{1:t-1}^{\mathbb{C}}) \quad (31)$$

$$= p(\mathbf{x}_t^{\mathbb{C}} | \boldsymbol{\pi}_t^{\mathbb{C}}) \sum_{\boldsymbol{\pi}_{t-1}^{\mathbb{C}}} p(\boldsymbol{\pi}_t^{\mathbb{C}} | \boldsymbol{\pi}_{t-1}^{\mathbb{C}}) p(\boldsymbol{\pi}_{t-1}^{\mathbb{C}} | \mathbf{x}_{1:t-1}^{\mathbb{C}}) \quad (32)$$

$$= \underbrace{p(\mathbf{x}_t^{\mathbb{C}} | \boldsymbol{\pi}_t^{\mathbb{C}})}_{\text{corrector}} \underbrace{\sum_{\boldsymbol{\pi}_{t-1}^{\mathbb{C}}} p(\boldsymbol{\pi}_t^{\mathbb{C}} | \boldsymbol{\pi}_{t-1}^{\mathbb{C}}) \alpha(\boldsymbol{\pi}_{t-1}^{\mathbb{C}})}_{\text{predictor}}. \quad (33)$$

Appendix B. Derivation of iFFBS

If we define $\alpha(\pi_t^c) = p(\pi_t^c | \boldsymbol{\pi}_{1:t+1}^{-c}, \mathbf{x}_{1:t}^c)$, the modified α -recursion of CHMM using iFFBS algorithm can be derived for a single chain c as following:

$$\alpha(\pi_t^c) = \sum_{\pi_{t-1}^c} p(\pi_t^c, \pi_{t-1}^c | \mathbf{x}_{1:t-1}^c, x_t^c, \boldsymbol{\pi}_{1:t}^{-c}, \boldsymbol{\pi}_{t+1}^{-c}) \quad (34)$$

$$\propto \sum_{\pi_{t-1}^c} p(x_t^c | \pi_t^c) p(\pi_t^c | \boldsymbol{\pi}_{t-1}^c) p(\pi_{t-1}^c | \boldsymbol{\pi}_{1:t}^{-c}, \mathbf{x}_{1:t-1}^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} p(\pi_{t+1}^{\hat{c}} | \boldsymbol{\pi}_t^{\mathbb{C}}) \quad (35)$$

$$= \underbrace{p(x_t^c | \pi_t^c)}_{\text{corrector}} \underbrace{\sum_{\pi_{t-1}^c} p(\pi_t^c | \boldsymbol{\pi}_{t-1}^c) \alpha(\pi_{t-1}^c)}_{\text{predictor}} \underbrace{\prod_{\hat{c} \in \mathbb{C} \setminus c} p(\pi_{t+1}^{\hat{c}} | \boldsymbol{\pi}_t^{\mathbb{C}})}_{\text{modifying mass}} \quad (36)$$

If we normalize $\alpha(\pi_t^c) = p(\pi_t^c | \boldsymbol{\pi}_{1:t+1}^{-c}, \mathbf{x}_{1:t}^c)$ in two step in which, first before multiplying the *modifying mass* and after multiplication, the first normalization constant will be $Z_t^c = p(x_t^c | \mathbf{x}_{1:t-1}^c, \boldsymbol{\pi}_{1:t}^{-c})$. In iFFBS algorithm, marginal likelihood is not available but we heuristically approximate it as follows:

$$\mathcal{L}(\mathbf{x}_{1:T}^{\mathbb{C}}) \triangleq p(\mathbf{x}_{1:T}^{\mathbb{C}}) \approx \prod_{t=1}^T \prod_{c=1}^{\mathbb{C}} Z_t^c. \quad (37)$$

Appendix C. Design of the Transition Probabilities

We notice that the unnormalized transition log-probabilities μ_t^c are an additive function of the latent states of the other chains, and transition probabilities τ_t^c are obtained by a normalization of the exponential function. Hence we can write:

$$p(\pi_t^c | \boldsymbol{\pi}_{t-1}^{\mathbb{C}}) \propto e^{\beta_0^{c \leftarrow c} + \sum_{\hat{c} \in \mathbb{C} \setminus c} \sum_{k \in \mathbb{K}} \beta_k^{c \leftarrow \hat{c}} \mathbb{I}[\pi_{t-1}^{\hat{c}} = k]} \quad (38)$$

$$= e^{\beta_0^{c \leftarrow c}} \prod_{\hat{c} \in \mathbb{C} \setminus c} \prod_{k \in \mathbb{K}} e^{\beta_k^{c \leftarrow \hat{c}} \mathbb{I}[\pi_{t-1}^{\hat{c}} = k]} \quad (39)$$

$$=: f_0(\pi_t^c, \boldsymbol{\pi}_{t-1}^{\mathbb{C}}) \prod_{\hat{c} \in \mathbb{C} \setminus c} f_{\pi_{t-1}^{\hat{c}}}(\pi_t^c, \boldsymbol{\pi}_{t-1}^{\mathbb{C}}). \quad (40)$$

[Equation 40](#) means that the transition probability from π_{t-1}^c to π_t^c in chain c factorizes into (an element-wise) product of matrices over chains, where the matrix corresponding to each chain $\hat{c} \neq c$ depends on the state $\pi_{t-1}^{\hat{c}}$ of the chain \hat{c} in the previous time-step.

Appendix D. Derivation of the Factorized FFBS (fFFBS)

D.1. Forward Filtering

If we define $\alpha(\pi_t^c) = p(\pi_t^c | x_t^c, \mathbf{x}_{1:t-1}^c)$ given factorization assumption defined in [Equation 24](#), the approximated α -recursion for each chain can be derived as follows;

$$\alpha(\pi_t^c) = \sum_{\pi_{t-1}^c} p(\pi_t^c, \boldsymbol{\pi}_{t-1}^c | x_t^c, \mathbf{x}_{1:t-1}^c) \quad (41)$$

$$\propto \sum_{\pi_{t-1}^c} p(x_t^c | \pi_t^c) p(\pi_t^c | \boldsymbol{\pi}_{t-1}^c) p(\boldsymbol{\pi}_{t-1}^c | \mathbf{x}_{1:t-1}^c) \quad (42)$$

$$= p(x_t^c | \pi_t^c) \sum_{\pi_{t-1}^c} p(\pi_t^c | \boldsymbol{\pi}_{t-1}^c) p(\boldsymbol{\pi}_{t-1}^c | \mathbf{x}_{1:t-1}^c) \quad (43)$$

$$\approx p(x_t^c | \pi_t^c) \sum_{\pi_{t-1}^c} p(\pi_t^c | \boldsymbol{\pi}_{t-1}^c) \prod_{\hat{c} \in \mathbb{C}} \underbrace{p(\pi_{t-1}^{\hat{c}} | x_{t-1}^{\hat{c}}, \mathbf{x}_{1:t-2}^{\hat{c}})}_{\alpha(\pi_{t-1}^{\hat{c}})} \quad (44)$$

$$= \underbrace{p(x_t^c | \pi_t^c)}_{\epsilon^c} \sum_{\pi_{t-1}^c} \alpha(\pi_{t-1}^c) \underbrace{\sum_{\pi_{t-1}^{c \setminus c}} p(\pi_t^c | \boldsymbol{\pi}_{t-1}^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} \alpha(\pi_{t-1}^{\hat{c}})}_{\tau_t^c} \quad (45)$$

$$= \epsilon^c \sum_{\pi_{t-1}^c} \alpha(\pi_{t-1}^c) \tau_t^c. \quad (46)$$

Here, we see that τ_t^c has a role analogous to the transition matrix in the regular forward-filtering algorithm, i.e., the probability $p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t-1}^c)$ in [Equation 33](#). In detail, it is a matrix specifying the probabilities of moving from π_{t-1}^c to π_t^c , weighing the transitions by the probabilities of the states $\pi_{t-1}^{\hat{c}}$ in the other chains \hat{c} .

D.2. Reinterpretation of Transition Matrix

By using the proportionality in [Equation 17](#) we can define dynamic transition matrices τ_t^c in [Equation 25](#) as follows:

$$\tau_t^c = \sum_{\pi_{t-1}^{c \setminus c}} p(\pi_t^c | \boldsymbol{\pi}_{t-1}^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} \alpha(\pi_{t-1}^{\hat{c}}) \quad (47)$$

$$\propto \sum_{\pi_{t-1}^{c \setminus c}} f_0(\pi_t^c, \pi_{t-1}^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} f_{\pi_{t-1}^{\hat{c}}}(\pi_t^c, \pi_{t-1}^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} \alpha(\pi_{t-1}^{\hat{c}}) \quad (48)$$

$$= f_0(\pi_t^c, \pi_{t-1}^c) \sum_{\pi_{t-1}^{c \setminus c}} \prod_{\hat{c} \in \mathbb{C} \setminus c} f_{\pi_{t-1}^{\hat{c}}}(\pi_t^c, \pi_{t-1}^c) \alpha(\pi_{t-1}^{\hat{c}}) \quad (49)$$

$$= f_0(\pi_t^c, \pi_{t-1}^c) \prod_{\hat{c} \in \mathbb{C} \setminus c} \sum_{\pi_{t-1}^{\hat{c}}} f_{\pi_{t-1}^{\hat{c}}}(\pi_t^c, \pi_{t-1}^c) \alpha(\pi_{t-1}^{\hat{c}}) \quad (50)$$

where rows of τ_t^c must sum to one by definition. So, we can normalize it at the final step, and τ_t^c will have an equality, not proportionality.

D.3. Time Reversed Transition Probability

Given Equation 24 at each t and Equation 17, we can expand the time reversed transition probabilities as follows:

$$p(\boldsymbol{\pi}_t^c | \boldsymbol{\pi}_{t+1}^c, \mathbf{x}_{1:t}^c) \propto p(\boldsymbol{\pi}_{t+1}^c | \boldsymbol{\pi}_t^c) p(\boldsymbol{\pi}_t^c | \mathbf{x}_{1:t}^c) \quad (51)$$

$$= p(\boldsymbol{\pi}_{t+1}^c | \boldsymbol{\pi}_t^c) \alpha(\boldsymbol{\pi}_t^c) \quad (52)$$

$$\approx \prod_{c \in \mathcal{C}} p(\pi_{t+1}^c | \pi_t^c) \alpha(\pi_t^c) \quad (53)$$

$$\propto \prod_{c \in \mathcal{C}} \alpha(\pi_t^c) \left[f_0(\pi_{t+1}^c, \pi_t^c) \prod_{\hat{c} \in \mathcal{C} \setminus c} f_{\pi_t^{\hat{c}}}(\pi_{t+1}^c, \pi_t^c) \right] \quad (54)$$

$$= \prod_{c \in \mathcal{C}} \alpha(\pi_t^c) \left[f_0(\pi_{t+1}^c, \pi_t^c) \prod_{\hat{c} \in \mathcal{C} \setminus c} f_{\pi_t^{\hat{c}}}(\pi_{t+1}^{\hat{c}}, \pi_t^{\hat{c}}) \right] \quad (55)$$

Here, in final step, we rearranged the factors, where the factors are grouped by the outgoing factors. We see that the factorization assumption leads to efficient computation by allowing us to operate on matrices whose size equals the size of a transition matrix for one chain only.

Appendix E. Adaptive Metropolis-Hastings within Gibbs Sampling

We have described the design specifications transition matrix of the CHMM in Section 3.2, in which $\boldsymbol{\beta}^c$ is defined as a set of matrices. Here, by abusing the notation, we consider all the parameters related to a chain c as a single vector, denoted as $\boldsymbol{\beta}^c$. Since transitions are changing at each time step, it is not feasible to calculate their sufficient statistics as in HMM. To resolve this, we sample the $\boldsymbol{\beta}^{c*}$ parameters using a Metropolis-Hastings(MH) step within the Gibbs sampler conditional on the latent states of all the chains. The transition parameters for each chain $\boldsymbol{\tau}^c$ are fully determined by the $\boldsymbol{\beta}^c$ parameters. To get a draw from $p(\boldsymbol{\beta}^c | \boldsymbol{\pi}_{1:T}^c, \boldsymbol{\pi}_{1:T}^{-c})$ we use the following steps;

1. Make a proposal $\boldsymbol{\beta}^{c*} \sim q(\boldsymbol{\beta}^c)$
2. Transform $\boldsymbol{\beta}^{c*}$ samples to transition probabilities $\boldsymbol{\tau}_{1:T}^{c*}$ according to latent states of the other chains, $\boldsymbol{\pi}^{-c}$
3. Accept the proposal $\boldsymbol{\beta}^{c*}$ with probability

$$\alpha(\boldsymbol{\tau}_{1:T}^{c*}) = \min \left\{ 1, \frac{p(\boldsymbol{\beta}^{c*} | \boldsymbol{\pi}_{1:T}^c, \boldsymbol{\pi}_{1:T}^{-c}) q(\boldsymbol{\beta}^c | \boldsymbol{\beta}^{c*})}{p(\boldsymbol{\beta}^c | \boldsymbol{\pi}_{1:T}^c, \boldsymbol{\pi}_{1:T}^{-c}) q(\boldsymbol{\beta}^{c*} | \boldsymbol{\beta}^c)} \right\} \quad (56)$$

where

$$p(\boldsymbol{\beta}^c | \boldsymbol{\pi}_{1:T}^c, \boldsymbol{\pi}_{1:T}^{-c}) \propto p(\boldsymbol{\pi}_{1:T}^c | \boldsymbol{\beta}^c, \boldsymbol{\pi}_{1:T}^{-c}) p(\boldsymbol{\beta}^c) \quad (57)$$

The first quantity on the right-hand side of the Equation 57 can be directly calculated with the transition parameters $\boldsymbol{\tau}_{1:T}^c$ for the corresponding chain. We used the Gaussian proposal in which the previously sampled $\boldsymbol{\beta}^c$ is the mean such that $q(\boldsymbol{\beta}^{c*} | \boldsymbol{\beta}^c) = \mathcal{N}(\boldsymbol{\beta}^{c*} | \boldsymbol{\beta}^c, \psi^2 \Sigma)$ where ψ is scaling factor and Σ is covariance matrix. We initially set $\psi \approx 2.38/\sqrt{d}$ as Metropolis jumping scaling factor since it is theoretically the most efficient scaling factor (Gelman et al., 2013), where d is the dimension of the sampling. As a variance parameter, we used a fixed diagonal covariance matrix $\Sigma = 0.01 \times I$, where I is the identity matrix, (which corresponds to a step size giving the optimal acceptance rate of $\approx 23\%$ (Roberts et al., 2001)) during the warm-up period and then used the online covariance matrix estimated from the previous samples Dasgupta and Hsu (2007). Then we adaptively scaled the scaling factor ψ as described in Sherlock et al. (2010). Since the proposal distribution is symmetric (i.e., Normal), $q(\boldsymbol{\beta}^c | \boldsymbol{\beta}^{c*})$ and $q(\boldsymbol{\beta}^{c*} | \boldsymbol{\beta}^c)$ in Equation 56, cancel out each other.

Appendix F. Implementation details

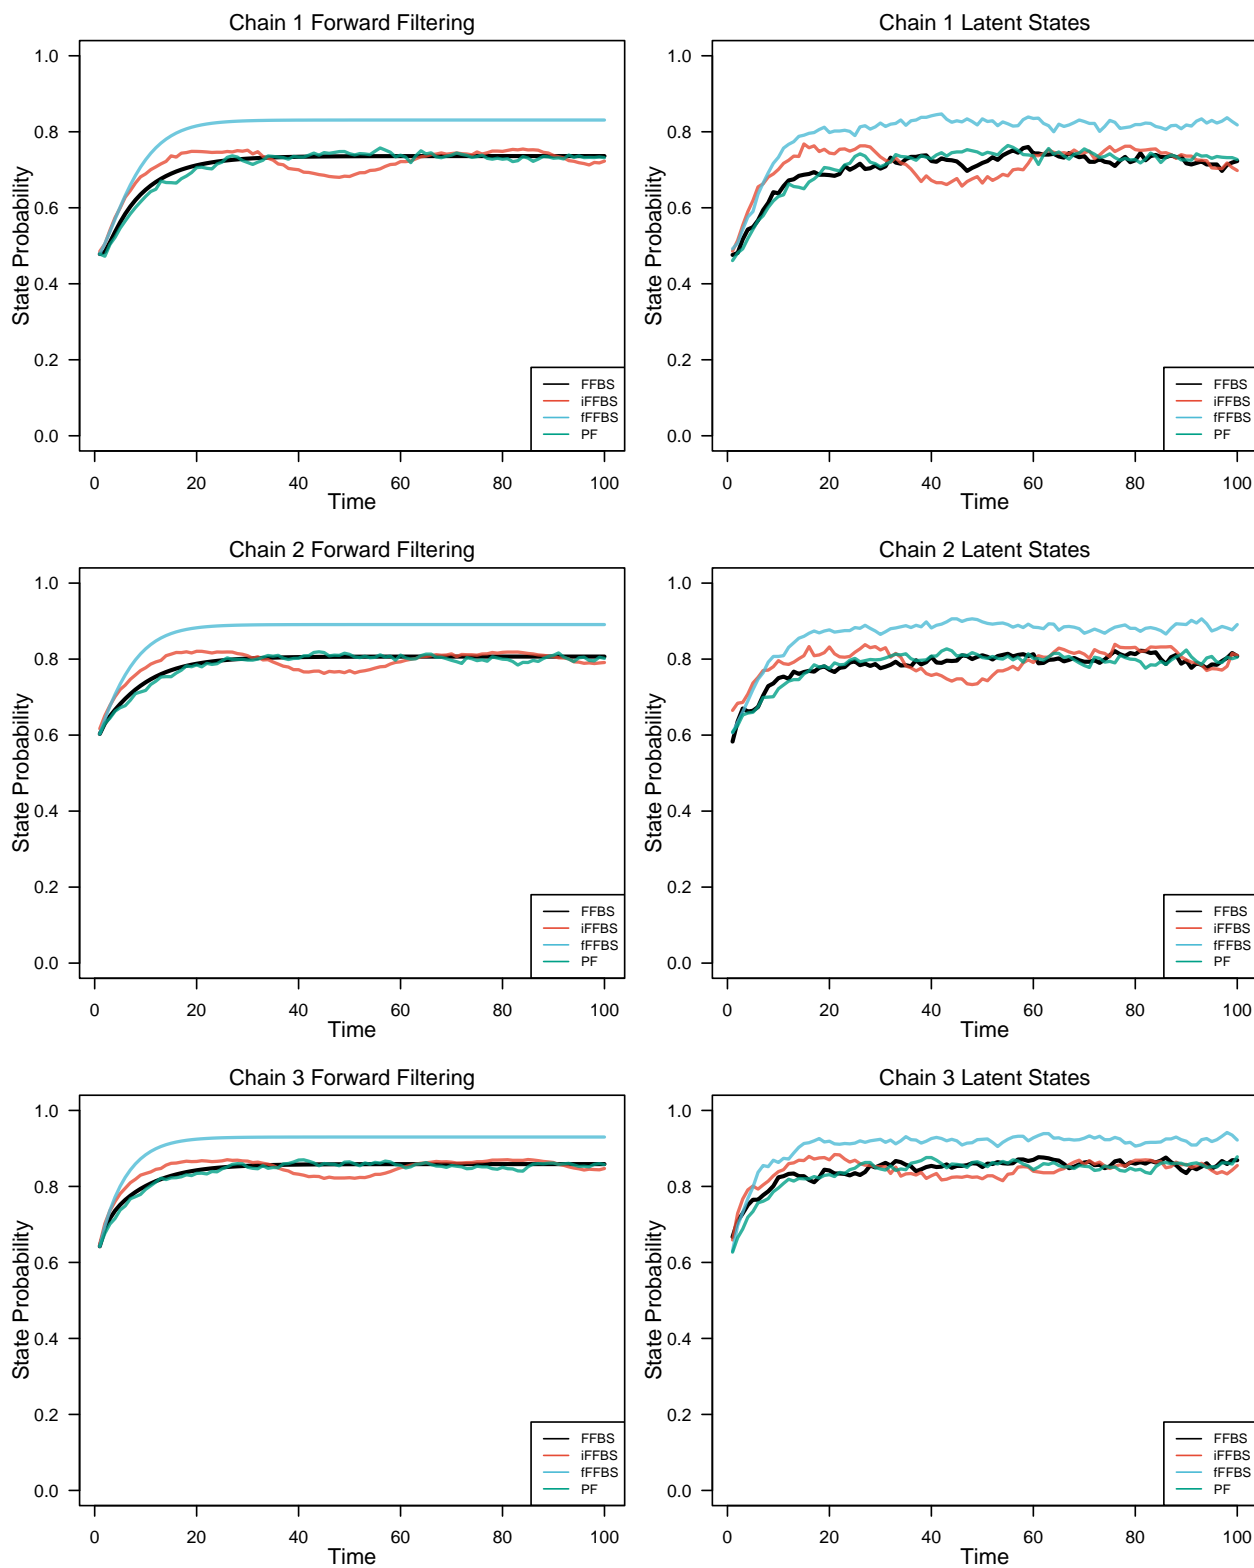
We set the prior of β_0^c as $N(\beta_0^c | 0, 1)$ for any chain, which is almost uninformative so that the estimates are not affected strongly by the prior. For the rest of the β parameters denoted by β_k , we used sparsity encouraging Horseshoe prior with mean and scale parameters of 0 and 0.25, respectively. We used a uniform prior on the initial state probabilities π_0 and weak Dirichlet priors for the rows of emission probabilities E such that we set the value to 30 for specificity and 15 for sensitivity. We set the rest of the emission priors to 1, which corresponds to the uniform prior. In such a formulation, except for the initialization, the prior has a negligible effect because it is summed with observation counts during inference. We drew 20,000 MCMC samples and set the warm-up length as 10,000. Posterior probabilities are calculated using the rest of MCMC samples.

Appendix G. Additional Results

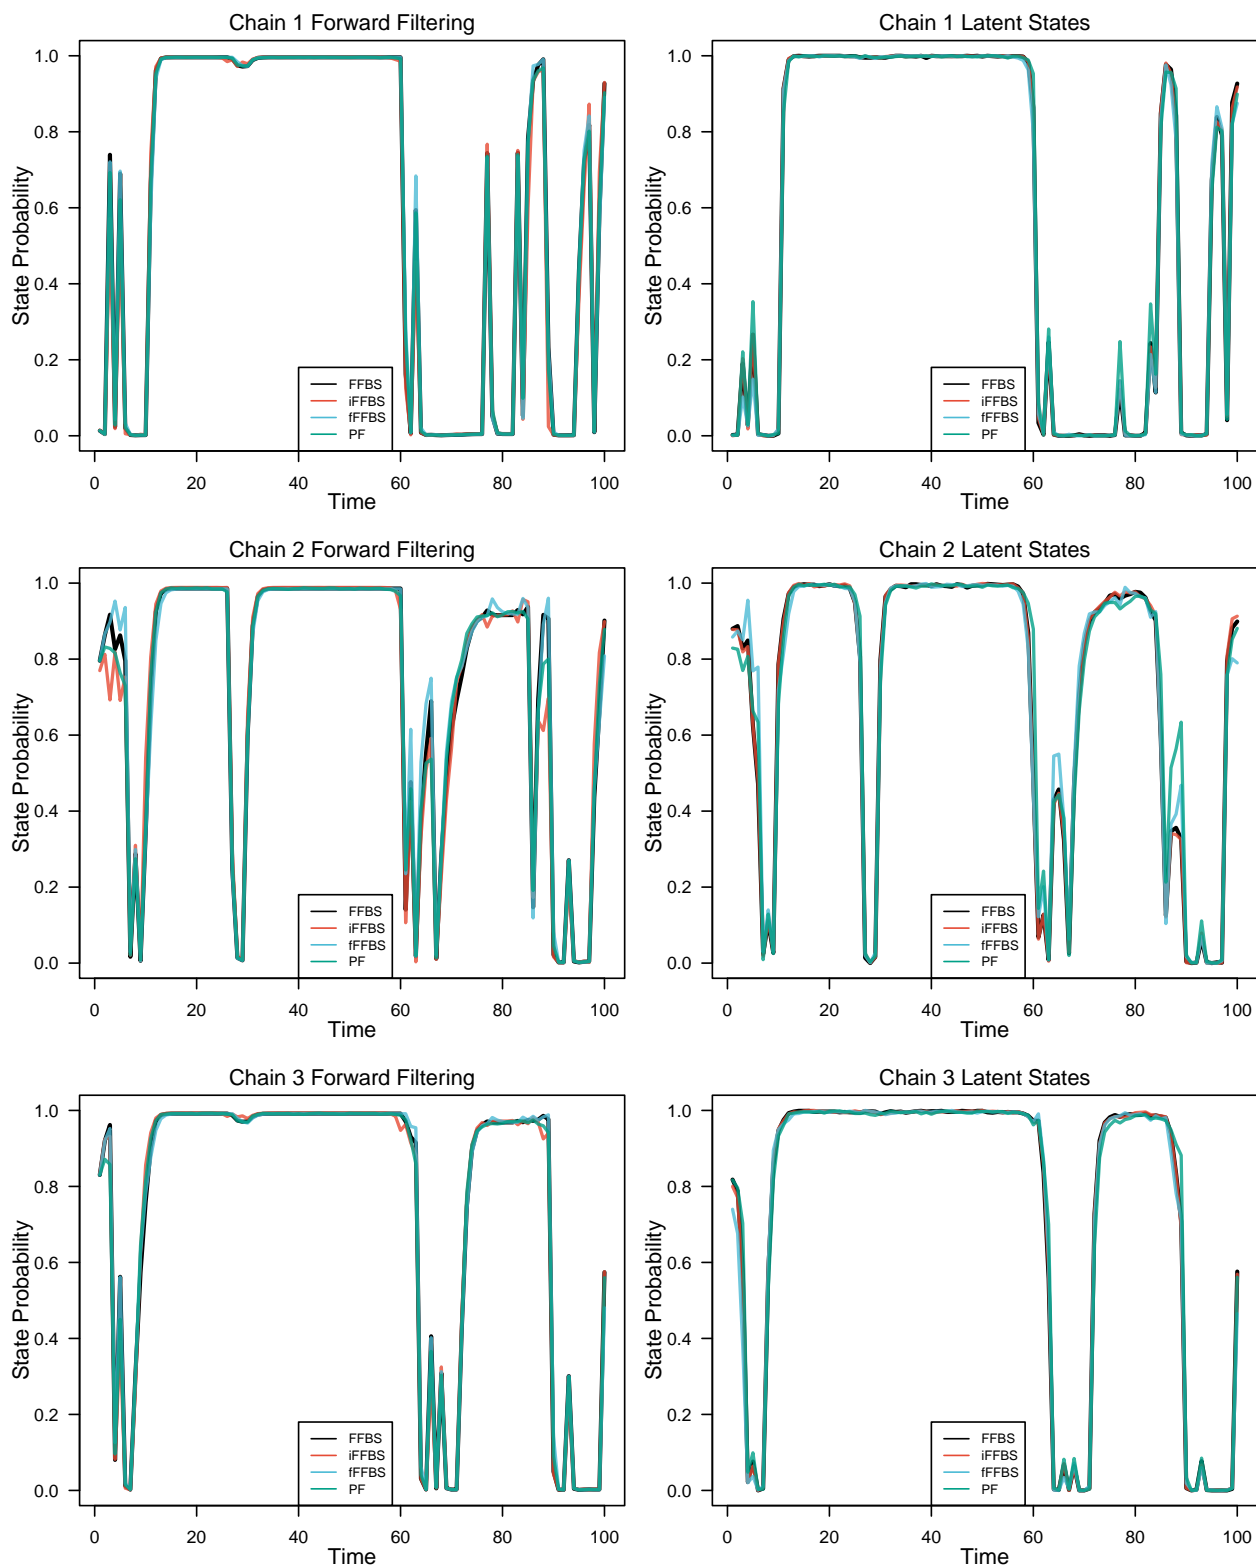
In this section, we provided additional analysis and results which supports the outcomes of the manuscript.

Supplementary Table 1: Effective sample size (ESS) of each model with 10000 MCMC samples.

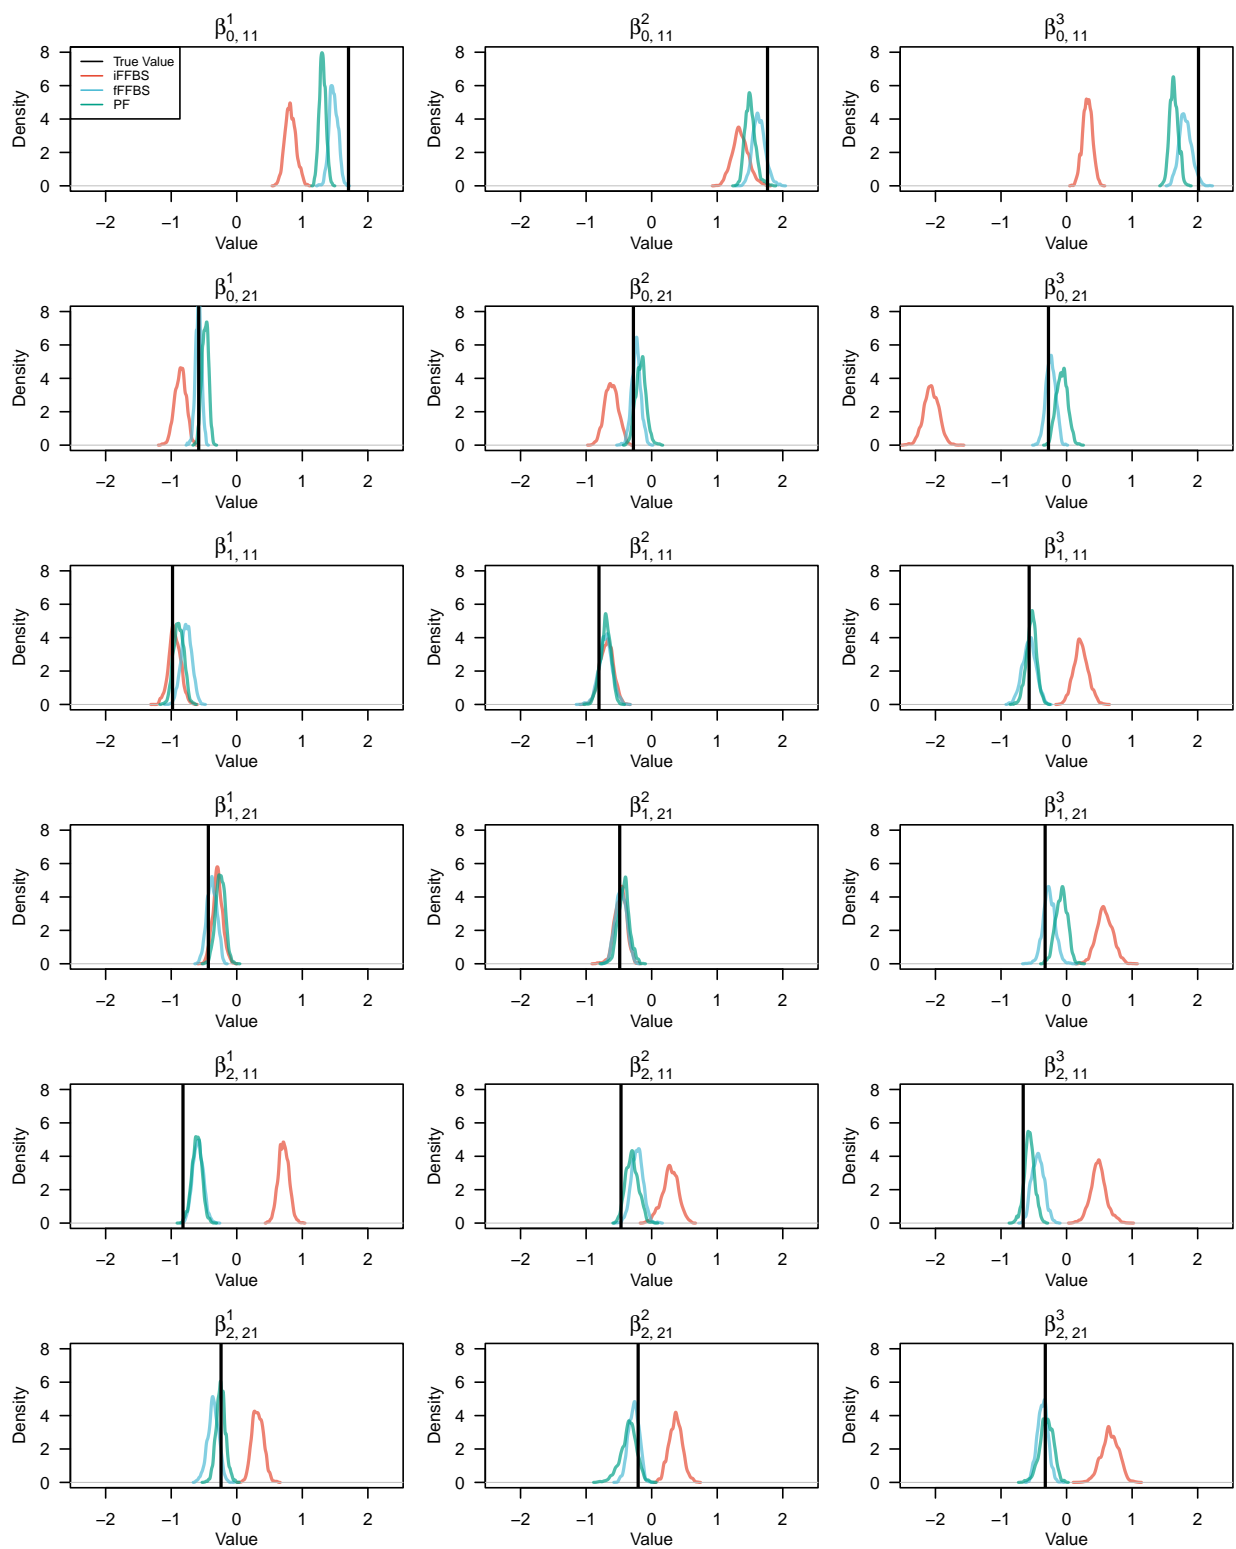
| Model | Control Group | | | | Treatment Group | | | |
|--------------------|---------------|-------|-------|------|-----------------|-------|-------|------|
| | FFBS | iFFBS | fFFBS | PF | FFBS | iFFBS | fFFBS | PF |
| CHMM | 1101 | 2253 | 708 | 3147 | 2115 | 3099 | 1482 | 1756 |
| M-CHMM ($M = 2$) | 966 | 664 | 676 | 958 | 529 | 576 | 787 | 915 |
| M-CHMM ($M = 3$) | 499 | 249 | 503 | 627 | 610 | 549 | 695 | 786 |
| M-CHMM ($M = 4$) | 550 | 203 | 242 | 343 | 494 | 409 | 378 | 456 |
| M-CHMM ($M = 5$) | 366 | 235 | 483 | 374 | 518 | 264 | 263 | 387 |
| M-CHMM ($M = 6$) | 415 | 247 | 366 | 341 | 461 | 307 | 320 | 355 |
| M-CHMM ($M = 7$) | 467 | 247 | 339 | 300 | 509 | 327 | 412 | 421 |
| M-CHMM ($M = 8$) | 439 | 264 | 350 | 335 | 382 | 307 | 321 | 360 |
| M-CHMM ($M = 9$) | 463 | 283 | 317 | 360 | 482 | 343 | 367 | 435 |



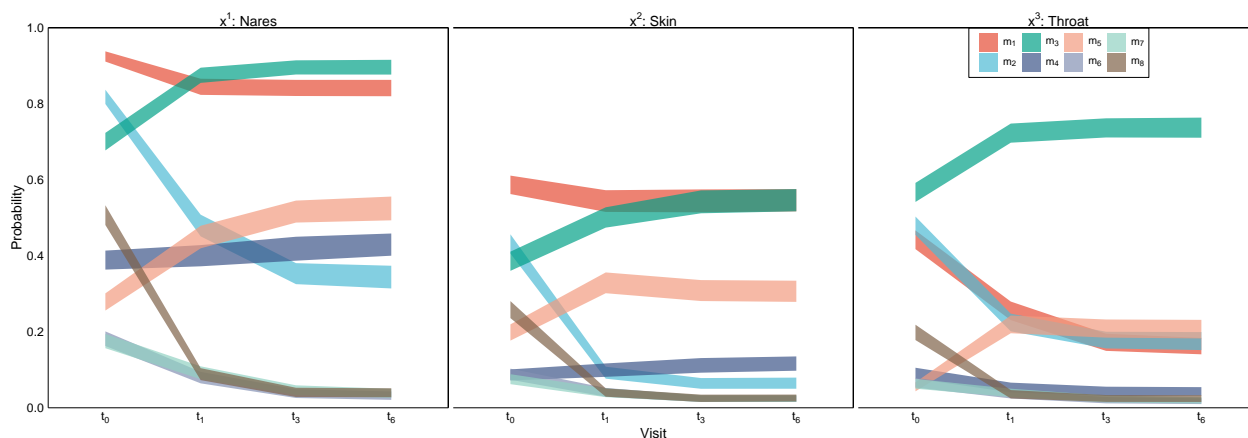
Supplementary Figure 1: Error accumulation without observation.



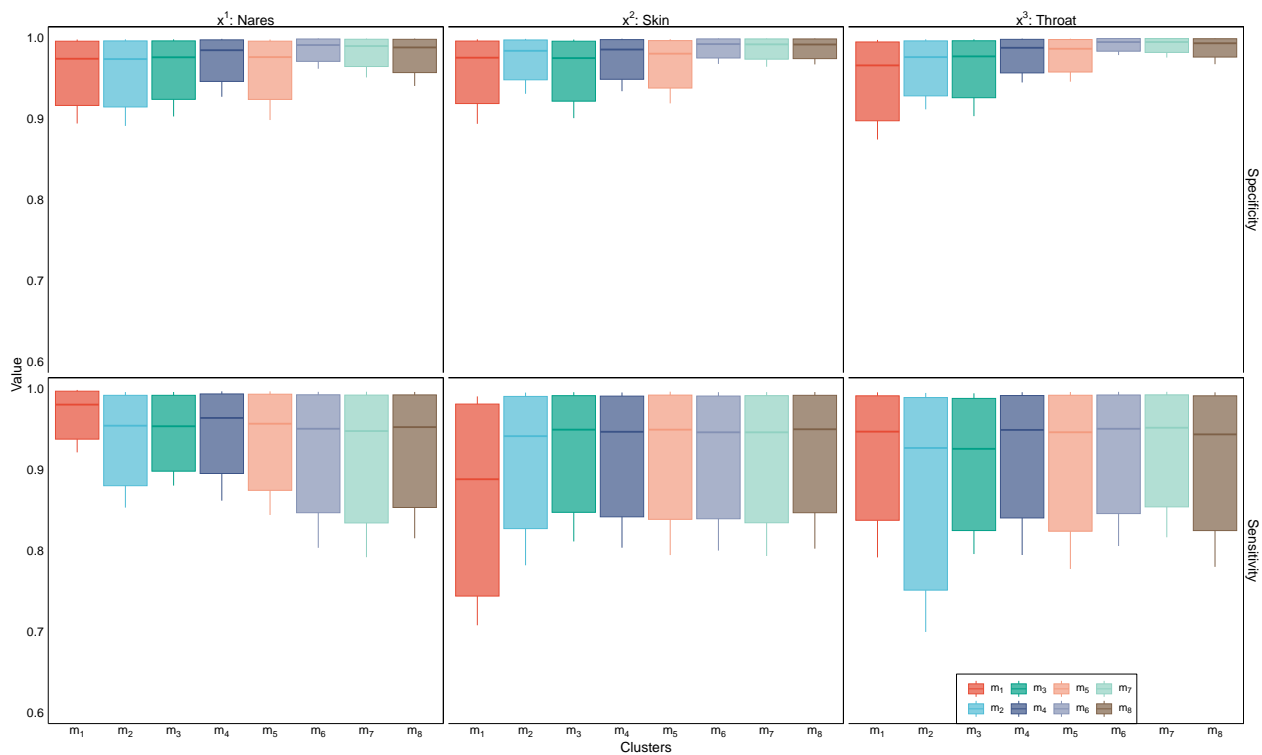
Supplementary Figure 2: Error accumulation with observation.



Supplementary Figure 3: Convergence of β parameters for each chain with different latent sequence samplers.



Supplementary Figure 4: The dynamics of the clusters in the control group detected by M-CHMM with 8 clusters where $n = 130, 117, 121, 144, 135, 154, 122$ and 129 , respectively, for $m_1, m_2, m_3, m_4, m_5, m_6, m_7$ and m_8 . The color scheme is the same with [Supplementary Figure 5](#).



Supplementary Figure 5: The specificity and sensitivity analysis of measurements in control group by cluster in M-CHMM, where $n = 130, 117, 121, 144, 135, 154, 122$ and 129 , respectively, for $m_1, m_2, m_3, m_4, m_5, m_6, m_7$ and m_8 . The color scheme is the same with [Supplementary Figure 4](#).