

UC San Diego

UC San Diego Previously Published Works

Title

Spectral Learning of Binomial HMMs for DNA Methylation Data

Permalink

<https://escholarship.org/uc/item/359566qz>

Authors

Zhang, Chicheng
Mukamel, Eran A
Chaudhuri, Kamalika

Publication Date

2018-02-07

Peer reviewed

Spectral Learning of Binomial HMMs for DNA Methylation Data

Chicheng Zhang¹, Eran A. Mukamel², Kamalika Chaudhuri³

¹ Microsoft Research

² Dept. of Computer Science and Engineering, University of California San Diego

³ Dept. of Cognitive Science, University of California San Diego

chicheng.zhang@microsoft.com, kamalika@cs.ucsd.edu, emukamel@ucsd.edu

February 8, 2018

Abstract

We consider learning parameters of Binomial Hidden Markov Models, which may be used to model DNA methylation data. The standard algorithm for the problem is EM, which is computationally expensive for sequences of the scale of the mammalian genome. Recently developed spectral algorithms can learn parameters of latent variable models via tensor decomposition, and are highly efficient for large data. However, these methods have only been applied to categorical HMMs, and the main challenge is how to extend them to Binomial HMMs while still retaining computational efficiency. We address this challenge by introducing a new feature-map based approach that exploits specific properties of Binomial HMMs. We provide theoretical performance guarantees for our algorithm and evaluate it on real DNA methylation data.

1 Introduction

Epigenetic modifications of DNA are critical for a wide range of biological processes. DNA methylation, i.e. the addition of a methyl group to genomic cytosine, is a well known and much studied epigenetic modification of mammalian genomes. However, the locations of methyl marks throughout the genome in particular cell types are largely unknown. Recent advances in genome sequencing technology now make it possible to measure DNA methylation with single-base resolution throughout the genome using whole-genome bisulfite sequencing (BS-Seq) [10, 13]. However, this data are limited by incomplete coverage and stochastic sampling of cells. Probabilistic methods for modeling DNA methylation data sets are therefore essential.

A DNA methylation profile, also called a methylome, is a binary sequence, where a location corresponds to the position of a cytosine base within the genome. A location t has some underlying functional property, modeled by a hidden state h_t , and may be methylated with some probability p_{h_t} that depends on the underlying state. For example, locations within the promoter regions of genes generally have very low probability of DNA methylation, while gene bodies and intergenic regions in general are highly methylated. Measuring methylation in a collection of cells of the same type gives a pair of counts (c_t, μ_t) , where the coverage c_t represents how many measurements were made (i.e. how many cytosines sequenced), and μ_t represents how many of these were methylated. Thus, μ_t is distributed as $\text{Bin}(c_t, p_{h_t})$. A natural framework for modeling this sequence data is therefore a hidden Markov model with binomial distributed observations (Binomial HMM). Given a DNA methylation sequence, our goal is to find the parameters of the underlying model, which describe the states, their methylation probabilities, and also their transitions. These parameters may be further used to segment the genome into compartments, corresponding to hidden states, with potential biologically relevant function [5, 12].

The main challenge in addressing this problem is that mammalian genomes are long sequences, with $\sim 10^9$ nucleotides. This means that algorithms such as EM, which are generally used to learn latent variable models, are extremely expensive as each iteration makes a pass over the entire methylation sequence. Recent

work by [1] has led to the emergence of spectral algorithms that can learn parameters of latent variable models via tensor decomposition. These algorithms are highly efficient and can operate on large volumes of data and provide theoretical guarantees for convergence to a global optimum. However, the challenge is that they only apply to categorical HMMs, and unlike likelihood-based methods, extending them to more complex models while still preserving computational efficiency is non-trivial. One plausible approach is to convert Binomial HMMs to categorical HMMs by converting the observations (c_t, μ_t) to categorical observations. However, this leads to a very high dimensional observation space, which in turn makes the algorithm expensive. A second plausible approach is kernelization [15]; however, the kernelized algorithm requires a running time quadratic in the sequence length, which again is prohibitively large for real data.

To address this, we observe that none of the proposed approaches (direct conversion or kernelization) exploit specific properties of Binomial HMMs. In particular, an important property is that estimating the joint distribution of c_t and μ_t given h_t is not necessary, and a direct estimation of $\{p_h\}_{h=1}^m$ may result in better statistical efficiency. We exploit this key property to propose a novel feature map-based modification of the tensor decomposition algorithm of [1], which we call FEATURE-TENSOR-DECOMP. We equip this algorithm with a novel feature map, called the Beta Map, that is tailored to Binomial HMMs. We then provide a novel parameter recovery procedure, which works with the Beta Map to recover the final binomial probability parameters. Finally, we make the entire algorithm more robust against model mismatch by providing a novel stabilization procedure for recovering the transition parameters.

We evaluate our algorithm both analytically and empirically. Theoretically, we provide performance guarantees which show that the proposed algorithm recovers the parameters correctly provided the methylation sequence is long enough, and the granularity of the Beta Map is fine enough with respect to the difference in the methylation probabilities across states. Empirically, we evaluate our algorithm on synthetic and real DNA methylation data. Our experiments show that in all cases our proposed algorithm is an order of magnitude faster than EM, which allows us to run it on much longer sequences. Our algorithm has lower estimation error as well as lower variance on synthetic data, and manages to recover a known pattern of methylation probabilities on real data.

2 Preliminaries

2.1 The Generative Model

DNA methylation in mammals occurs mainly at CG dinucleotides, which occur roughly once for every 100 bases. Therefore, in this work we model the genome using non-overlapping 100 base-pair bins tiling the genome. A DNA methylation data set from whole-genome bisulfite sequencing is represented by a sequence of pairs of integers $\{x_t = (c_t, \mu_t)\}_{t=1}^l$. At the genomic bin labeled by position t , $c_t \in \{0, 1, \dots, N\}$ is the *coverage*, i.e. number of sequenced DNA fragments that map to CG sites in that bin. The integer $\mu_t \in \{0, 1, \dots, c_t\}$ is the *methylation count*, which represents the number of DNA fragments that were found to be methylated at those CG sites. On average, c_t is typically about 30 for high-quality methylome data sets.

We model this data by a binomial hidden Markov model. At position t , there is an underlying hidden state h_t in $\{1, 2, \dots, m\}$ generating the observation x_t , the dynamics of which are modeled by a Markov chain. A hidden state $h \in \{1, \dots, m\}$ is associated with a methylation probability $p_h \in [0, 1]$. The coverage c_t 's are observed. Given c_t and h_t , the methylation count μ_t is drawn from a binomial distribution, with the mean parameter $p = p_{h_t}$. Thus, the binomial HMM model is represented by parameters (π, T, p) , where $\pi \in \mathbb{R}^m$ is the initial distribution over states, $T \in \mathbb{R}^{m \times m}$ is the transition matrix of the Markov chain over the states, and $p \in \mathbb{R}^m$ is the methylation probability vector.

2.2 Notation

For a matrix M , we denote by M^l its l -th column, and denote by $M_{i,j}$ its (i, j) -th entry. Given a third order tensor $\mathcal{T} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, we use $\mathcal{T}_{i_1, i_2, i_3}$ to denote its (i_1, i_2, i_3) -th entry. The tensor product of vectors v_1, v_2

and v_3 , denoted by $v_1 \otimes v_2 \otimes v_3$, whose (i_1, i_2, i_3) -th entry is $(v_1)_{i_1} (v_2)_{i_2} (v_3)_{i_3}$. A tensor T is called symmetric if $\mathcal{T}_{i_1, i_2, i_3} = \mathcal{T}_{i_{\pi(1)}, i_{\pi(2)}, i_{\pi(3)}}$ for any permutation $\pi : \{1, 2, 3\} \rightarrow \{1, 2, 3\}$. Given a tensor $\mathcal{T} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ and matrices $V_i \in \mathbb{R}^{n_i \times m_i}$, $i = 1, 2, 3$, $\mathcal{T}(V_1, V_2, V_3)$ is a tensor of size $m_1 \times m_2 \times m_3$, whose (i_1, i_2, i_3) -th entry is given by: $\mathcal{T}(V_1, V_2, V_3)_{i_1, i_2, i_3} = \sum_{j_1, j_2, j_3} \mathcal{T}_{j_1, j_2, j_3} (V_1)_{j_1, i_1} (V_2)_{j_2, i_2} (V_3)_{j_3, i_3}$.

Given a sample S of size M drawn from some distribution, we use $\hat{\mathbb{E}}[\cdot]$ to denote the empirical expectation over S , that is, given a function f , $\hat{\mathbb{E}}[f(x)] = \frac{1}{M} \sum_{x \in S} f(x)$.

2.3 Background: Spectral Learning for Categorical HMMs

Our proposed algorithm builds on [1], who provide a spectral algorithm for parameter estimation in categorical Hidden Markov Models. [1] consider categorical HMMs – where observations x_t are drawn from a categorical distribution on $\{1, \dots, n\}$. The distribution of x_t when the hidden state is h is given by O_h , where O is a parameter matrix in $\mathbb{R}^{n \times m}$. As the model is different from the binomial HMM, the results of [1] do not directly apply to our setting; we show in Section 3 how they can be adapted with a few novel modifications.

We begin with a brief overview of the algorithm of [1], abbreviated as TENSOR-DECOMP. The key idea is that if certain conditions on the parameters of the HMM holds, then, decomposing a tensor of third moments of the observations can give a transformed version of the HMM parameters. The TENSOR-DECOMP algorithm has three steps – first, it constructs certain cooccurrence matrices and tensors; second, it decomposes the cooccurrence tensor after transforming it to ensure tractable decomposition, and finally, the parameters are recovered from the decomposition results.

- **Step 1: Construct Cooccurrence Matrices and Tensors.** First, compute the empirical matrices $\hat{P}_{i,j} := \hat{\mathbb{E}}[x_i \otimes x_j]$ (where i, j are distinct elements from $\{1, 2, 3\}$ ¹) and tensor $\hat{\mathcal{T}} := \hat{\mathbb{E}}[x_1 \otimes x_2 \otimes x_3]$.

- **Step 2: Transform and Decompose.** The tensor $\hat{\mathcal{T}}$ is related to the parameter O , which can in theory be recovered by decomposing $\hat{\mathcal{T}}$. However, directly decomposing $\hat{\mathcal{T}}$ is computationally intractable [7], and so the tensor is converted to symmetric orthogonal form.

To symmetrize $\hat{\mathcal{T}}$, we compute the matrices $S_1 := \hat{P}_{2,3} \hat{P}_{1,3}^\dagger$ and $S_3 := \hat{P}_{2,1} \hat{P}_{3,1}^\dagger$. Then, compute $\mathcal{G} := \hat{\mathcal{T}}(S_1, I, S_3)$, which is symmetric. To orthogonalize \mathcal{G} , compute a matrix $J := S_3 \hat{P}_{3,2}$; take an SVD of J and take the top m singular vectors U_m , and singular values in diagonal matrix S_m , getting the orthogonalization matrix $W = U_m S_m^{-1/2}$. Perform a linear transformation over tensor \mathcal{G} using W , getting the symmetric orthogonal tensor $\mathcal{H} = \mathcal{G}(W, W, W)$.

\mathcal{H} is then decomposed via the Tensor Power Method [1] to recover the eigenvectors v_l 's and the eigenvalues λ_l .

- **Step 3: Recover Parameters.** The columns of O can now be recovered as $\hat{O}^l = (W^T)^\dagger \hat{\lambda}_l \hat{v}_l$, for $l = 1, 2, \dots, m$. Compute an estimate of the joint probability of h_2 and h_1 : $\hat{H}_{21} := \hat{O}^T \hat{P}_{21} \hat{O}^\dagger$. Then estimate the initial probability vector and the transition matrix by $\hat{\pi} := 1^T \hat{H}_{21}^T$ and $\hat{T} := \hat{H}_{21} \text{diag}(\hat{\pi})^{-1}$.

Observe that TENSOR-DECOMP has two advantages over popular parameter estimation algorithms such as EM. First, it only needs to make one pass over the data, and is thus computationally efficient if the size of the observation space is not too large. In contrast, EM proceeds iteratively and needs to make one pass over the data per iteration. Second, it achieves statistical consistency if the data is generated from an HMM (see Theorem 5.1 of [1]), in contrast with methods such as EM, which do not have statistical consistency guarantees.

¹Although only the first three observations are used, the algorithm can be generalized to use all three consecutive observations in the sequences.

3 Algorithm

The main limitation of spectral algorithms is that unlike likelihood-based methods, they cannot be readily extended to more general models. In particular, the TENSOR-DECOMP algorithm for categorical HMMs does not directly apply to our setting. One plausible approach is to convert our model to a categorical HMM by converting the observations (c_t, μ_t) to categorical observations. However, this leads to a very high dimensional observation space, leading to high computational cost. A second plausible approach is kernelization [15]; however, the kernelized algorithm requires a running time quadratic in the sequence length, which again is prohibitively large for real data.

To address this challenge, we observe that neither direct conversion nor kernelization exploits specific properties of Binomial HMMs. In particular, an important property is that estimating the joint distribution of c_t and μ_t given h_t is not necessary, and a direct estimation of $\{p_h\}_{h=1}^m$ may result in better statistical efficiency.

We exploit this key property to propose a novel feature map-based modification of TENSOR-DECOMP, which we call FEATURE-TENSOR-DECOMP. We equip FEATURE-TENSOR-DECOMP with a novel feature map, called the Beta Map, that is tailored to Binomial HMMs. We then provide a novel parameter recovery procedure, which works with the Beta Map to recover the final binomial probability parameters. Finally, we make the entire algorithm more robust against model mismatch by providing a novel stabilization procedure for recovering the transition parameters.

3.1 Key Components

We next describe the key modifications that we make to TENSOR-DECOMP to adapt it to binomial HMMs.

Modification 1: Feature Map. Our first contribution is to use a feature map ϕ to map our discrete observations $x = (c, \mu)$ to a D -dimensional vector $\phi(x)$. Thus, instead of computing the cooccurrence matrices and tensors in Step 1 of TENSOR-DECOMP, we now compute the empirical *feature co-occurrence matrices and tensors*:

$$\hat{P}_{i,j} := \hat{\mathbb{E}}[\phi(x_i) \otimes \phi(x_j)], \quad \hat{\mathcal{T}} := \hat{\mathbb{E}}[\phi(x_1) \otimes \phi(x_2) \otimes \phi(x_3)],$$

and apply the remaining steps of TENSOR-DECOMP on these matrices and tensors. We call the resulting algorithm FEATURE-TENSOR-DECOMP.

What does FEATURE-TENSOR-DECOMP recover? Define the matrix C as a $D \times m$ matrix whose j -th row is: $C^j = \mathbb{E}[\phi(x)|h = j]$; this is analogous to the observation matrix O in [1]. Provided certain conditions hold, we show in Section 4 that FEATURE-TENSOR-DECOMP can, given sufficiently many sequences, provably recover a high quality estimate of C .

Modification 2: Beta Mapping. We next propose a novel feature map, called the Beta Map, that is tailored to Binomial HMMs; the map is inspired by the popular Beta-Binomial conjugate prior-posterior system in Bayesian inference.

Define the function $\varphi_{\text{beta},D}((c, \mu), t) = \frac{1}{\text{B}(\mu+1, c-\mu+1)} t^\mu (1-t)^{c-\mu}$ as the density of the Beta distribution with shape parameters $\mu + 1$ and $c - \mu + 1$, where $\text{B}(\cdot, \cdot)$ is the Beta function. Now, given an observation (c, μ) , the Beta Map $\phi_{\text{beta},D}(c, \mu)$ is an n -dimensional vector whose i -th entry is:

$$(\phi_{\text{beta},D}(c, \mu))_i := \int_{(i-1)/D}^{i/D} \varphi_{\text{beta},D}((c, \mu), t) dt \tag{1}$$

We apply Algorithm FEATURE-TENSOR-DECOMP with the feature map $\phi_{\text{beta},D}$.

The Beta Map has two highly desirable properties. First, it maps the observation (μ_t, c_t) to a probability mass function with mean close to $\frac{\mu_t}{c_t}$, which is p_{h_t} in expectation. Second, if c_t is large, then the feature map is highly concentrated around $\approx \frac{\mu_t}{c_t}$, reflecting higher confidence in the estimated p_{h_t} .

Modification 3: Recovery of the Binomial Probabilities. As we observe earlier, running FEATURE-TENSOR-DECOMP with the Beta Map will recover the expected feature map matrix C , and not the binomial probabilities. We now provide a novel recovery procedure to estimate p_h from C .

Suppose $\phi = \phi_{\text{beta}, D}$ is the Beta mapping, and \hat{C} is the estimated feature map recovered by FEATURE-TENSOR-DECOMP. Observe that for any h , the expected Beta map C^h is a mixture of Beta distributions, where each mixture component corresponds to a pair (c, μ) and has shape parameter $(\mu + 1, c - \mu + 1)$ (and hence mean $\approx \frac{\mu+1}{c+2}$) and mixing weight $p(c_t = c, \mu_t = \mu | h_t = h)$. As the mean of the mixture distribution is equal to the weighted average over the component means, this gives us: $\frac{1}{D} \sum_{i=1}^D \frac{i}{D} (C)_{i,h} \approx \mathbb{E}[\frac{\mu+1}{c+2} | h]$. Assuming that h and c are independent, the right hand side simplifies to $a + (1 - 2a)p_h$, where $a := \mathbb{E}[\frac{1}{c+2}]$. Using $\hat{a} := \hat{\mathbb{E}}[\frac{1}{c+2}]$ instead of a , this gives the following recovery equation (2) for p_h :

$$\hat{p}_h := \frac{\frac{1}{D} \sum_{i=1}^D \frac{i}{D} (\hat{C})_{i,h} - \hat{a}}{1 - 2\hat{a}} \quad (2)$$

Due to the estimation error in C and the discretization of the Beta mapping, we cannot hope to recover p_h exactly. However, if C is accurately recovered, and the granularity of the Beta map is fine, then our estimate of p_h is accurate. Our estimation grows accurate with increasing granularity of discretization, at the expense of a higher running time.

Modification 4: Stabilization. Model mismatch in real data often leads to unstable solutions in [1], especially in recovering the transition matrix T . To prevent instability, we propose an alternative approach: a least squares formulation to recover π and T . Given \hat{C} as an estimator of $\mathbb{E}[\phi(x)|h]$, we propose to solve the following optimization problem:

$$\min_{H_{2,1}: \forall i,j (H_{2,1})_{i,j} \geq 0, \sum_{i,j} (H_{2,1})_{i,j} = 1} \|P_{2,1} - \hat{C}H_{2,1}\hat{C}^T\|_F^2$$

Here, $(H_{2,1})_{i,j}$ is our proposed estimator of $\mathbb{P}(x_2 = i, x_1 = j)$, $i, j \in \{1, \dots, m\}$. Next, we recover the transition matrix and initial probability by applying the formulae $\hat{\pi} = 1^T H_{2,1}$ and $\hat{T} = H_{2,1} \text{diag}(\hat{\pi})^{-1}$. The key difference between this procedure and the Step 3 of TENSOR-DECOMP is that, our optimization problem ensures that our estimators of π and T are entrywise positive, and thus no postprocessing are needed for subsequent usage of the parameters.

3.2 Extension: Multiple Cell Types

Finally, an additional goal for us is to study multiple aligned methylation sequences in order to identify differential methylation states, where the expected methylation probabilities are different across different cell types.

Here, we observe two coverage methylation pairs per location, one for each cell type, so an observation $x = ((c^1, \mu^1), (c^2, \mu^2))$. Our goal is to estimate a pair of methylation probabilities (p_h^1, p_h^2) for each state h that is shared across cells. To this end, we construct a concatenated feature map:

$$\Phi(x) = \begin{bmatrix} \phi(c^1, \mu^1) \\ \phi(c^2, \mu^2) \end{bmatrix}$$

Following the tensor decomposition algorithm in Section 2, we can recover the expected feature map given hidden states:

$$C = \mathbb{E}[\Phi(x)|h] = \begin{bmatrix} \mathbb{E}[\phi(c^1, \mu^1)|h] \\ \mathbb{E}[\phi(c^2, \mu^2)|h] \end{bmatrix}$$

Applying the recovery procedure of p_h to each block gives a pair (p_h^1, p_h^2) for each hidden state h . If we see a large difference between p_h^1 (methylation probability in cell type 1) and p_h^2 (methylation probability in cell type 2), then we identify state h as a differential methylation state.

4 Performance Guarantees

The TENSOR-DECOMP algorithm has provable guarantees when an underlying condition, called the Rank Condition, on the parameters of the HMM holds. For FEATURE-TENSOR-DECOMP, the analogous condition is the Feature Rank Condition below.

Assumption 1 (Feature Rank Condition). *The expected feature map matrix $C \in \mathbb{R}^{D \times m}$ and the transition matrix $T \in \mathbb{R}^{m \times m}$ are of full column rank.*

The Feature Rank Condition is satisfied for Beta Maps when the p_h values are well separated and the granularity of the Beta Map is high. Formally, if q is the minimum gap $\min_{i \neq j} |p_i - p_j|$, then, Theorem 1 shows that as long as the discretization parameter D and the coverage c are above some function of q , the Feature Rank Condition is satisfied. For simplicity we assume here that the coverage c is fixed.

Theorem 1. *Suppose $D \geq \frac{4}{q}$ and $c \geq \frac{512}{q^2}$. Consider the Beta feature map $\phi_{\text{beta},D}(x)$ defined as in Equation (1). Then C , the expected feature map matrix, has minimum singular value at least $\frac{1}{2\sqrt{D}}$, and is thus of full column rank.*

Provided the conditions of Theorem 1 hold, we can show statistical consistency of FEATURE-TENSOR-DECOMP; the proof is given in the appendix.

Theorem 2 (Statistical consistency of FEATURE-TENSOR-DECOMP). *Suppose FEATURE-TENSOR-DECOMP receives M iid samples drawn from a binomial hidden Markov model represented by parameters (π, T, p_c, p) . In addition, suppose T is of full rank, π is positive entrywise, and the coverage c is large enough. Then with high probability, the ℓ_2 distances between the outputs \hat{p} , \hat{T} and $\hat{\pi}$ and the respective underlying parameters π, T, p converge to zero, with increasing sample size M and Beta Map dimension D .*

5 Experiments

Our evaluation of the empirical performance of FEATURE-TENSOR-DECOMP has two major goals. First, we aim to validate our theoretical results by examining how it performs against EM when data is truly generated from a Binomial HMM. Real data typically has model mismatch, and our second goal is to investigate how FEATURE-TENSOR-DECOMP performs under model mismatch by comparing it with EM on real DNA methylation data.

5.1 Validation on Synthetic Data

We begin with simulations on synthetic data with known underlying generative parameters.

Data Generation. We generate a single sequence binomial HMM with four hidden states. For each state, we generate a methylation probability p_h . To ensure there is a gap in the probabilities, p_1 and p_2 are uniformly drawn from $[0, 0.3]$, and p_3 and p_4 from $[0.7, 1]$. We generate a transition matrix T as: $T = 0.2I_4 + 0.8U$, where all elements of matrix U are drawn independently and uniformly from $[0, 1]$, and then each column is normalized to sum to 1. We generate the initial probability vector π by normalizing a random vector u , with entries drawn uniformly and independently from $[0, 1]$. We consider three coverage settings – the coverage c_t drawn from a Poisson with means 25 (low coverage), 50 (medium coverage) and 100 (high coverage). For each set of parameters, we draw 8 sequences of size from $\{2^7, \dots, 2^{13}\}$.

Methodology. We compare FEATURE-TENSOR-DECOMP with EM. Both algorithms have a hyperparameter – the number of states m that we set to the correct value 4. For FEATURE-TENSOR-DECOMP the number of tensor power iterations per component is 2 and the granularity for Beta Map is 30. EM is stopped at iteration t if the fractional decrease in the log-likelihood is below 0.001.

Recall that neither FEATURE-TENSOR-DECOMP nor EM produces the hidden states in correct order, and hence to evaluate the estimation error, we need to match the estimated states with the true ones. We




Figure 1: Estimation error vs sample size for FEATURE-TENSOR-DECOMP (red) and EM (blue). From top to bottom: three settings with varying mean coverage parameters: 25, 50 and 100. All results are averaged over 20 repeated trials; the error bars represent the standard deviation of the estimation error value over these trials.

do so by using the Hungarian algorithm to find the minimum cost matching between states where the cost of matching a state p to \hat{p} is the difference $|p - \hat{p}|$. The estimation error is defined as the cost of the best cost matching. We repeat the experiment for 20 trials, and plot the mean and the standard error of the estimation error in Figure 1. The running times are reported in Table 1.

Results. From Figure 1, we observe that for small training sample size, EM and FEATURE-TENSOR-DECOMP have comparable estimation error, whereas for large sample size (> 1000), FEATURE-TENSOR-DECOMP performs substantially better. In general, we find that FEATURE-TENSOR-DECOMP is robust as our theory predicts, while the results of EM depend strongly on initialization. Table 1 illustrates that FEATURE-TENSOR-DECOMP is an order of magnitude faster than EM for all sample sizes, which means that it can be indeed be run on larger datasets. We remark that the average running time is not strictly monotonic increasing with sample size, which may be attributed to the fact that we use caching of the feature maps to speed up the implementation. These observations indicate that FEATURE-TENSOR-DECOMP should be preferred when large amounts of data are available, while EM may be used for smaller data problems. An additional observation is that the estimation error under EM typically has higher variance than FEATURE-TENSOR-DECOMP for larger sample sizes. Finally, the degree of coverage does not appear to make a big difference to the results.

Sample Size \ Algorithm	FTD	EM
128	0.6758	5.838
256	0.9948	5.718
512	1.4427	7.713
1024	1.8173	17.468
2048	1.4253	28.105
4096	2.7181	62.430
8192	4.2342	82.218

Table 1: Average running time (in seconds) for EM and FEATURE-TENSOR-DECOMP with varying training sample sizes. All results are averaged over 20 repeated trials. We only present the result when the average coverage is 25; the cases when the average coverage are 50, 100 are similar.

5.2 Real Data Experiments

We next investigate how FEATURE-TENSOR-DECOMP performs on real DNA methylation data.

Data. We use the DNA methylation dataset from [10], and use chromosome 1 of the mouse genome. The data has a pair of (coverage, methylation) counts for every bin consisting of 100 DNA base pairs. We select two cell types – excitatory neurons (E) and VIP cells (V) – which are known to have differential methylation. The data has two replicate sequences for each cell type; we merge them by adding the coverage and methylation counts at each location. The data also has contextual information for each position that is derived from the underlying DNA sequence. We are specifically interested in the CG context that is of biological relevance. We extract the subsequence that is restricted to the CG context, which gives us a sequence of length 1923719.

Methodology. For real data, we compare three algorithms – FEATURE-TENSOR-DECOMP (FTD), EM and FEATURE-TENSOR-DECOMP followed by 3 rounds of EM (FTD+EM). The motivation for studying the last algorithm is that it has been previously reported to have good performance on real data [17, 3]. We set the number of hidden states to $m = 6$. As running EM until convergence on this long sequence is slow, we run it for 10 iterations and report the results.

Evaluation on real data is more challenging since we do not have the ground truth parameters. We get around this by using the following two measures. First, we compute the log-likelihood of the estimated parameters over a separate test set. Second, we look at the cell types E and V which are known to have differential methylation regions in the CG context, and verify that this is reflected in our results. Finally, for aligned sequences from multiple cell types, the observation probabilities given the hidden states are computed under a conditional independence assumption.

5.3 Results

Figure 2 plots the log-likelihood on a test set under EM, FTD and FTD+EM. The running times are reported in Table 2, and the estimated methylation probability matrices are plotted in Figure 3.

From Figure 2, we see that EM achieves slightly better log-likelihood than FTD and FTD+EM. This is not surprising since EM directly maximizes the likelihood unlike FTD. We see from Table 2 that for the sample size of 4×10^4 , FTD is two orders of magnitude faster than that of EM. The bottleneck in FTD+EM’s running time is the 3 rounds of EM. In addition, FTD processes almost the whole sequence (of length 1.6×10^6) in a relatively short amount of time (<280 s), which is faster than EM on a sample of size 4×10^4 (>1175 s).

In Figure 3, we permute the columns of each estimated methylation probability matrix to make them as aligned as possible. We mark the aligned states using vertical lines between the columns. We can draw a few conclusions from the estimated matrices: first, all algorithms identify a hidden state with extremely low methylation level on both cells (column 1, <0.025) and a hidden state with extremely high methylation level on both cells (column 2, >0.91). Second, all algorithms identify a hidden state with low methylation on cell

Algorithm (sample size)	Running Time (s)
FTD(4×10^4)	11.608894
FTD+ EM(4×10^4)	357.817675
EM(4×10^4)	1125.509846
FTD(1.6×10^6)	274.1687

Table 2: Running Time for a set of algorithm and sample size pairs.




Figure 2: Test log-likelihood vs training sample Size for algorithms FEATURE-TENSOR-DECOMP (red), EM (blue), and EM with FEATURE-TENSOR-DECOMP initialization (green).

E and high methylation on cell V (column 6), although the difference is most stark for FEATURE-TENSOR-DECOMP run on a sample size of 1.6×10^6 . Third, the EM based algorithms (FTD and FTD+EM) find a state that has relatively high methylation on cell E and low methylation on cell V (column 5), whereas FTD does not identify this state. FTD+EM appears to achieve the best of both worlds – relatively low running time with correctly identified states.

6 Related Work

There has been a large body of work on hidden Markov models with structured emission distributions [11]. The binomial hidden Markov model has been proposed in [4], and is used in [6] for financial applications.

A line of work has developed spectral learning algorithms for latent variable models, including HMMs. [1] provides an elegant algorithm for categorical HMMs based on moment matching and tensor decomposition, which provably recovers the parameters when the training data is generated from an HMM. The algorithm provides no guarantees under model mismatch, and works only for categorical HMMs, which has a finite output space. [14] applies this algorithm to Chip-Seq data to recover underlying states. [15] proposes kernelized

$$\begin{array}{l}
 \left[\begin{array}{cccccc} 0.023 & 0.953 & 0.229 & 0.859 & 0.709 & 0.583 \\ 0.023 & 0.954 & 0.476 & 0.781 & 0.457 & 0.903 \end{array} \right] \text{EM}(4 \times 10^4) \\
 \left[\begin{array}{cccccc} 0.023 & 0.955 & 0.749 & 0.472 & 0.858 & 0.168 \\ 0.026 & 0.948 & 0.827 & 0.379 & 0.419 & 0.811 \end{array} \right] \text{FTD + EM}(4 \times 10^4) \\
 \left[\begin{array}{cccccc} 0.01 & 0.916 & 0.298 & 0.345 & 0.649 & 0.206 \\ 0.01 & 0.925 & 0.257 & 0.450 & 0.880 & 0.535 \end{array} \right] \text{FTD}(4 \times 10^4) \\
 \left[\begin{array}{cccccc} 0.01 & 0.99 & 0.598 & 0.838 & 0.950 & 0.132 \\ 0.01 & 0.99 & 0.519 & 0.990 & 0.913 & 0.752 \end{array} \right] \text{FTD}(1.6 \times 10^6)
 \end{array}$$

Figure 3: The methylation probability matrices (of size 2×6) for each algorithm on datasets of different sizes. Row 1: Cell E, Row 2: Cell V.

versions of spectral methods that generalize the algorithm of [1] to handle data with rich observation spaces; however, a direct application of [15]’s algorithm has running time quadratic in the sample size, thus making it prohibitive for methylation data. [16] studies spectral learning for a structured HMM with multiple cell types, namely HMM with tree hidden states (HMM-THS) [2]. Their model is different from ours, in that our HMM has only one state controlling the observations at each time, whereas HMM-THS has one state per sequence. [18] studies the setting where there are two sequences, foreground and background, and the goal is to find states that appear exclusively in the foreground sequence. Their problem setting is different from ours, in that the roles of the two sequences are asymmetric, unlike ours. Finally, it has been observed both empirically and theoretically in the spectral learning literature that EM initialized with the output of the spectral methods achieve better accuracy in many tasks, such as crowdsourcing [17] and in learning a mixture of linear regressors [3].

Acknowledgments

We thank NSF under IIS-1617157 for research support. EAM is supported by NIH/NINDS (NS080911).

References

- [1] Anima Anandkumar, Rong Ge, Daniel Hsu, Sham M. Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *CoRR*, abs/1210.7559, 2012.
- [2] Jacob Biesinger, Yuanfeng Wang, and Xiaohui Xie. Discovering and mapping chromatin states using a tree hidden Markov model. *BMC Bioinformatics*, 14(Suppl 5):S4, 2013.
- [3] A. Chaganty and P. Liang. Spectral experts for estimating mixtures of linear regressions. In *ICML*, 2013.
- [4] Christophe Couvreur. Hidden markov models and their mixtures. Technical report, Quantitative Research Centre, NatWestGroup, 1996.
- [5] Jason Ernst and Manolis Kellis. ChromHMM: automating chromatin-state discovery and characterization. *Nature Publishing Group*, 9(3):215–216, March 2012.
- [6] Giacomo Giampieri, Mark Davis, and Martin Crowder. A hidden markov model of default interaction. *Quantitative Finance*, 2005.
- [7] Christopher J Hillar and Lek-Heng Lim. Most tensor problems are np-hard. *Journal of the ACM (JACM)*, 60(6):45, 2013.
- [8] D. Hsu, S.M. Kakade, and T Zhang. A spectral algorithm for learning hidden Markov models. *Journal of Computer and System Sciences*, 78:1460–1480, 2012.
- [9] Olivier Marchal and Julyan Arbel. On the sub-gaussianity of the beta and dirichlet distributions. *arXiv preprint arXiv:1705.00048*, 2017.
- [10] Alisa Mo, Eran A Mukamel, Fred P Davis, Chongyuan Luo, Gilbert L Henry, Serge Picard, Mark A Urich, Joseph R Nery, Terrence J Sejnowski, Ryan Lister, Sean R Eddy, Joseph R Ecker, and Jeremy Nathans. Epigenomic signatures of neuronal diversity in the mammalian brain. *Neuron*, 86(6):1369–1384, June 2015.
- [11] Lawrence R Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

- [12] Roadmap Epigenomics Consortium, Wouter Meuleman, Angela Yen, Jianrong Wang, Michael J Ziller, Viren Amin, Lucas D Ward, Abhishek Sarkar, Gerald Quon, Richard S Sandstrom, Matthew L Eaton, Yi-Chieh Wu, Andreas R Pfenning, Xinchun Wang, Melina Claussnitzer, Yaping Liu, Cristian Coarfa, R Alan Harris, Charles B Epstein, Danny Leung, Wei Xie, R David Hawkins, Ryan Lister, Chibo Hong, Philippe Gascard, Andrew J Mungall, Richard Moore, Angela Tam, Theresa K Canfield, R Scott Hansen, Rajinder Kaul, Peter J Sabo, Mukul S Bansal, Jesse R Dixon, Kai-How Farh, Soheil Feizi, Rosa Karlic, Ah-Ram Kim, Daofeng Li, Tim R Mercer, Vitor Onuchic, Paz Polak, Pradipta Ray, Richard C Sallari, Kyle T Siebenthall, Nicholas A Sinnott-Armstrong, Robert E Thurman, Jie Wu, Bo Zhang, Xin Zhou, Arthur E Beaudet, Laurie A Boyer, Philip L De Jager, Susan J Fisher, Steven J M Jones, Wei Li, Marco A Marra, James A Thomson, Thea D Tlsty, Li-Huei Tsai, Wei Wang, Robert A Waterland, Michael Q Zhang, Lisa H Chadwick, Bradley E Bernstein, Joseph F Costello, Joseph R Ecker, Martin Hirst, Aleksandar Milosavljevic, Bing Ren, John A Stamatoyannopoulos, Ting Wang, Manolis Kellis, Anshul Kundaje, Jason Ernst, Misha Bilenky, Alireza Heravi-Moussavi, Pouya Kheradpour, Zhizhuo Zhang, John W Whitaker, Matthew D Schultz, Noam Shores, Elizabeta Gjoneska, Eric Chuah, Annaick Carles, Ashwinikumar Kulkarni, Rebecca Lowdon, Ginell Elliott, Shane J Neph, Nisha Rajagopal, Michael Stevens, Peggy J Farnham, David Haussler, Michael T McManus, Shamil Sunyaev, and Alexander Meissner. Integrative analysis of 111 reference human epigenomes. *Nature*, 518(7539):317–330, February 2015.
- [13] Matthew D Schultz, Yupeng He, John W Whitaker, Manoj Hariharan, Eran A Mukamel, Danny Leung, Nisha Rajagopal, Joseph R Nery, Mark A Urich, Huaming Chen, Shin Lin, Yiing Lin, Inkyung Jung, Anthony D Schmitt, Siddarth Selvaraj, Bing Ren, Terrence J Sejnowski, Wei Wang, and Joseph R Ecker. Human body epigenome maps reveal noncanonical DNA methylation variation. *Nature*, June 2015.
- [14] J. Song and K. C. Chen. Spectacle: fast chromatin state annotation using spectral learning. *Genome Biology*, 16:33, 2015.
- [15] Le Song, Animashree Anandkumar, Bo Dai, and Bo Xie. Nonparametric estimation of multi-view latent variable models. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 640–648, 2014.
- [16] Chicheng Zhang, Jimin Song, Kamalika Chaudhuri, and Kevin Chen. Spectral learning of large structured hmms for comparative epigenomics. In *Advances in Neural Information Processing Systems*, pages 469–477, 2015.
- [17] Yuchen Zhang, Xi Chen, Dengyong Zhou, and Michael I. Jordan. Spectral methods meet EM: A provably optimal algorithm for crowdsourcing. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [18] J. Zou, D. Hsu, D. Parkes, and R. Adams. Contrastive learning using spectral methods. In *NIPS*, 2013.

A Sample Complexity - Proof of Theorem 2

We present Theorem 3 below, which immediately implies Theorem 2. To see this, suppose the coverage c is $\geq \frac{512}{q^2}$, and we are given given learning parameters $\epsilon, \delta \in (0, 1)$. Now, set $D_0 = \max(\frac{4}{q}, \frac{16}{\epsilon})$. By Theorem 1, for all $D \geq D_0$, $\sigma_{\min}(C) \geq \frac{1}{2\sqrt{D}}$. Given such D , by Theorem 3, we can find a value of $N_0 = \text{poly}(\frac{1}{\min_i \pi_i}, D, \frac{1}{\sigma_{\min}(T)}, \frac{1}{\epsilon}, \ln \frac{1}{\delta})$, such that for every $N \geq N_0$, with probability $1 - \delta$, the distances between \hat{p} , \hat{T} and $\hat{\pi}$ and the respective underlying parameters π, T, p are all at most ϵ in terms of the respective error metrics.

Theorem 3 (Statistical Consistency of FEATURE-TENSOR-DECOMP). *Suppose FEATURE-TENSOR-DECOMP receives M iid samples (x_1, x_2, x_3) as input, which are drawn from a binomial hidden Markov model represented by parameters (π, T, p_c, p) . In addition, suppose the distribution on coverage c is such that $a := \mathbb{E}[\frac{1}{c+2}] \leq \frac{3}{8}$. Then, given learning parameters ϵ and δ in $(0, \frac{1}{16})$, if the Beta feature map dimension D is $\frac{16}{\epsilon}$, and the number of samples M is at least $\text{poly}(\frac{1}{\min_i \pi_i}, \frac{1}{\sigma_{\min}(C)}, \frac{1}{\sigma_{\min}(T)}, \frac{1}{\epsilon}, \ln \frac{1}{\delta}, D)$, then with probability $1 - \delta$, the output \hat{p} , \hat{T} and $\hat{\pi}$ satisfies that*

$$\begin{aligned} \|p - \hat{p}\Pi\|_2 &\leq \epsilon \\ \|T - \Pi^\dagger \hat{T}\Pi\|_F &\leq \epsilon \\ \|\pi - \Pi^\dagger \hat{\pi}\|_2 &\leq \epsilon \end{aligned}$$

for some permutation matrix Π . Here $\sigma_{\min}(M)$ denotes the minimum singular value of matrix M .

A.1 Recovering Initial Probability, Transition Matrix and Expected Feature Map

To prove Theorem 2, we will apply the sample complexity bounds of TENSOR-DECOMP for HMM [1]. Specifically, we will apply a result implicit in [1], which appears explicitly in [16].

Theorem 4 (Initial Probability, Transition Matrix and Expected Feature Map Consistency). *Suppose FEATURE-TENSOR-DECOMP receives m iid samples (x_1, x_2, x_3) as input, which are drawn from a categorical hidden Markov model represented by parameters (π, T, O) . Then, given parameters ϵ and δ in $(0, 1)$, if the number of samples M is at least $\text{poly}(\frac{1}{\min_i \pi_i}, \frac{1}{\sigma_{\min}(C)}, \frac{1}{\sigma_{\min}(T)}, \frac{1}{\epsilon}, \ln \frac{1}{\delta})$, then with probability $1 - \delta$, FEATURE-TENSOR-DECOMP produces estimated expected feature map \hat{C} , transition matrix \hat{T} and initial probability $\hat{\pi}$ such that*

$$\begin{aligned} \|C - \hat{C}\Pi\|_F &\leq \epsilon \\ \|T - \Pi^\dagger \hat{T}\Pi\|_F &\leq \epsilon \\ \|\pi - \Pi^\dagger \hat{\pi}\|_2 &\leq \epsilon \end{aligned}$$

for some permutation matrix Π . Here $\sigma_{\min}(M)$ denotes the minimum singular value of a matrix M .

Proof. The proof is almost the same as the proof of ([16], Theorem 1), by taking V as a set of size $D = 1$. The only difference between our proof and theirs is the argument for the concentration of the raw moments, which we address in Lemma 1 below. \square

Lemma 1. *Suppose we are given M iid triples $(x_{i,1}, x_{i,2}, x_{i,3})$, $i \in \{1, 2, \dots, M\}$. Then with probability $1 - \delta$, the following concentration inequalities hold simultaneously:*

$$\begin{aligned} \|P_{12} - \hat{P}_{12}\|_F &\leq \epsilon(M, \delta) \\ \|P_{23} - \hat{P}_{23}\|_F &\leq \epsilon(M, \delta) \\ \|P_{13} - \hat{P}_{13}\|_F &\leq \epsilon(M, \delta) \\ \|\mathcal{T} - \hat{\mathcal{T}}\|_F &\leq \epsilon(M, \delta) \end{aligned}$$

where $\epsilon(M, \delta) = \sqrt{\frac{4+4\ln(8/\delta)}{M}}$.

Proof. The lemma is a direct consequence of Lemma 2 below by taking ξ_i as the vectorization of $\phi(x_{i,1}) \otimes \phi(x_{i,2})$, $\phi(x_{i,2}) \otimes \phi(x_{i,3})$, $\phi(x_{i,1}) \otimes \phi(x_{i,3})$ and $\phi(x_{i,1}) \otimes \phi(x_{i,2}) \otimes \phi(x_{i,3})$ respectively, along with a union bound. \square

Lemma 2. *Suppose we are given a sequence of M iid d -dimensional vectors ξ_i , $i = 1, 2, \dots, M$, and $\|\xi_i\|_1 \leq 1$ almost surely. In addition, denote by Ξ the expectation of the ξ_i 's, and denote by $\hat{\xi} := \frac{1}{M} \sum_{i=1}^M \xi_i$ the empirical mean of the ξ_i 's. Then, with probability $1 - \delta$,*

$$\|\hat{\xi} - \Xi\|_2 \leq \sqrt{\frac{4 + 4 \ln \frac{2}{\delta}}{M}}$$

Proof. The lemma follows from the ideas in ([8], Appendix A). We first show that $\mathbb{E}\|\hat{\xi} - \Xi\|_2 \leq \frac{2}{M}$. We justify it as follows:

$$\mathbb{E}\|\hat{\xi} - \Xi\|_2^2 = \mathbb{E}\left\| \sum_{i=1}^M \frac{1}{M} (\xi_i - \Xi) \right\|_2^2 = \frac{1}{M^2} \sum_{i=1}^M \mathbb{E}\|\xi_i - \Xi\|_2^2 \leq \frac{4}{M}$$

where the last inequality is due to that $\|\xi_i\|_2 \leq \|\xi_i\|_1 \leq 1$, and $\|\Xi\|_2 \leq \mathbb{E}\|\xi_i\|_2 \leq 1$. Consequently, $\mathbb{E}\|\hat{\xi} - \Xi\|_2 \leq \sqrt{\mathbb{E}\|\hat{\xi} - \Xi\|_2^2} \leq \sqrt{\frac{4}{M}}$.

Next, note that replacing a ξ_i with a ξ'_i changes the function $\|\hat{\xi} - \Xi\|_2$ by at most $\frac{2}{M}$. Applying McDiarmid's inequality, we get that with probability $1 - \delta$,

$$\left| \|\hat{\xi} - \Xi\|_2 - \mathbb{E}\|\hat{\xi} - \Xi\|_2 \right| \leq \sqrt{\frac{4 \ln \frac{2}{\delta}}{M}}$$

This implies that

$$\|\hat{\xi} - \Xi\|_2 \leq \mathbb{E}\|\hat{\xi} - \Xi\|_2 + \sqrt{\frac{4 \ln \frac{2}{\delta}}{M}} \leq \sqrt{\frac{4 + 4 \ln \frac{2}{\delta}}{M}}.$$

\square

A.2 From Expected Feature Map to Binomial Probability

We have shown in the above section that the expected feature map, transition matrix and initial probability vector can be accurately recovered. However, the recovery accuracy of the binomial probability p still remains unaddressed. In this section, we address this issue. Specifically, we establish Lemma 3, which shows the following implication: if we get an accurate enough recovery of C_2 and we have a large enough D (the discretization granularity of our Beta feature map), then we can get an accurate recovery of p . Recall that $\hat{a} := \hat{\mathbb{E}}[\frac{1}{c+2}]$ is the empirical mean of $a := \mathbb{E}[\frac{1}{c+2}]$.

Lemma 3. *Suppose that the hidden state h and the coverage c are independent. For any $\epsilon \in (0, \frac{1}{16})$, if $D \geq \frac{16}{\epsilon}$, and for two columns h and h' , $\|\hat{C}^{h'} - C^h\|_2 \leq \frac{\epsilon}{16\sqrt{D}}$, $|\hat{a} - a| \leq \frac{\epsilon}{256}$ and $a \leq \frac{3}{8}$, then*

$$|\hat{p}_{h'} - p_h| \leq \epsilon.$$

Before going into the proof of Lemma 3, we need a lemma that characterize the recovery of p in the setting of infinite D and infinite sample size.

Lemma 4. *The binomial probability p_h can be written in terms of the Beta density as follows:*

$$p_h = \frac{\int_0^1 t \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt - a}{1 - 2a}. \quad (3)$$

Proof. Recall that from standard calculus,

$$\begin{aligned} & \int_0^1 \frac{t \cdot t^\mu (1-t)^{c-\mu}}{\text{B}(\mu+1, c-\mu+1)} dt \\ &= \int_0^1 \frac{t \cdot t^{(\mu+1)-1} (1-t)^{(c-\mu+1)-1}}{\text{B}(\mu+1, c-\mu+1)} dt = \frac{\mu+1}{c+2}. \end{aligned}$$

Therefore,

$$\begin{aligned} & \int_0^1 t \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt = \mathbb{E} \left[\frac{\mu+1}{c+2} \middle| h \right] \\ &= \mathbb{E} \left[\frac{cp_h + 1}{c+2} \middle| h \right] = \mathbb{E} \left[\frac{c}{c+2} \right] p_h + \mathbb{E} \left[\frac{1}{c+2} \right], \end{aligned}$$

where the last step uses the independence assumption between h and c . Recall that $a = \mathbb{E}[\frac{1}{c+2}]$, the above can be rewritten as:

$$\int_0^1 t \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt = (1-2a)p_h + a.$$

The equality in the lemma statement follows by algebra. \square

Proof of Lemma 3. Without loss of generality, suppose h and h' are the same.

Denote $\hat{\gamma}_{D,h}$ (resp. $\gamma_{D,h}$, γ_h) by $\frac{1}{D} \sum_{i=1}^D \frac{i}{D} \hat{C}_{h,i}$ (resp. $\frac{1}{D} \sum_{i=1}^D \frac{i}{D} C_{h,i}$, $\int_0^1 t \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt$). Observe that $\hat{\gamma}_{D,h}$, $\gamma_{D,h}$, and γ_h are all in $[0, 1]$. Using this notation, the recovery formula for p_h in FEATURE-TENSOR-DECOMP (Equation (2)) can be written as

$$\hat{p}_h = \frac{\hat{\gamma}_{D,h} - \hat{a}}{1 - 2\hat{a}}.$$

In addition, Lemma 4 implies that p_h can be written as

$$p_h = \frac{\gamma_h - a}{1 - 2a}.$$

Note that by triangle inequality, $|\hat{\gamma}_{D,h} - \gamma_{D,h}| \leq \|C^h - \hat{C}^h\|_1 \leq \frac{\epsilon}{16\sqrt{D}} \sqrt{D} = \frac{\epsilon}{16}$. In addition, we can bound $|\gamma_h - \gamma_{D,h}|$ as follows:

$$\begin{aligned} & |\gamma_h - \gamma_{D,h}| \\ &= \left| \sum_{i=1}^D \frac{i}{D} \mathbb{E}[(\phi_{\text{beta}}(x))_i | h] - \int_0^1 t \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt \right| \\ &= \left| \sum_{i=1}^D \frac{i}{D} \int_{\frac{i-1}{D}}^{\frac{i}{D}} \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt - \sum_{i=1}^D \int_{\frac{i-1}{D}}^{\frac{i}{D}} t \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt \right| \\ &\leq \sum_{i=1}^D \left| \int_{\frac{i-1}{D}}^{\frac{i}{D}} \left(\frac{i}{D} - t \right) \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt \right| \\ &\leq \sum_{i=1}^D \frac{1}{D} \int_{\frac{i-1}{D}}^{\frac{i}{D}} \mathbb{E}[\varphi_{\text{beta}}(x, t) | h] dt \\ &\leq \frac{1}{D} \leq \frac{\epsilon}{16} \end{aligned}$$

Combining the above two facts, we have $|\hat{\gamma}_h - \gamma_h| \leq |\hat{\gamma}_h - \gamma_{D,h}| + |\gamma_h - \gamma_{D,h}| \leq \frac{\epsilon}{8}$.

On the other hand, denote ξ (resp. $\hat{\xi}$) by $\frac{1}{1-2a}$ (resp. $\frac{1}{1-2\hat{a}}$). Note that since $a \leq \frac{3}{8}$, $\xi \leq 4$. Thus,

$$|\xi - \hat{\xi}| = \frac{2|a - \hat{a}|}{(1-2\hat{a})(1-2a)} \leq 64|a - \hat{a}| \leq \frac{\epsilon}{4}$$

Therefore,

$$\begin{aligned} & |p_h - \hat{p}_h| \\ &= |(\hat{\xi} - \xi)(\hat{\gamma}_h - \gamma_h) + \xi(\hat{\gamma}_h - \gamma_h) + \gamma_h(\hat{\xi} - \xi)| \\ &\leq |(\hat{\xi} - \xi)(\hat{\gamma}_h - \gamma_h)| + |\xi(\hat{\gamma}_h - \gamma_h)| + |\gamma_h(\hat{\xi} - \xi)| \\ &\leq \frac{\epsilon}{4} \frac{\epsilon}{8} + 4 \frac{\epsilon}{8} + \frac{\epsilon}{4} \\ &\leq \epsilon. \end{aligned}$$

This completes the proof. \square

A.3 Putting It Together

Built on Theorem 4 and Lemma 3, we are ready to prove Theorem 3.

Proof of Theorem 3. First, given the choice of $D \geq \frac{16}{\epsilon}$, if we have a sample of size $\text{poly}(\frac{1}{\min_i \pi_i}, \frac{1}{\sigma_{\min}(C)}, \frac{1}{\sigma_{\min}(T)}, \frac{1}{\epsilon}, \ln \frac{1}{\delta}, D)$, by Theorem 4, with probability $1 - \delta/2$, produces estimated expected feature map \hat{C} , transition matrix \hat{T} and initial probability $\hat{\pi}$ such that

$$\begin{aligned} \|C - \hat{C}\Pi\|_2 &\leq \frac{\epsilon}{16\sqrt{D}}, \\ \|T - \Pi^\dagger \hat{T}\Pi\|_F &\leq \frac{\epsilon}{16\sqrt{D}} \leq \epsilon, \\ \|\pi - \Pi^\dagger \hat{\pi}\|_2 &\leq \frac{\epsilon}{16\sqrt{D}} \leq \epsilon. \end{aligned} \tag{4}$$

for some permutation matrix Π . In addition, if the sample size M is at least $O(\frac{\ln \frac{1}{\delta}}{\epsilon^2})$, by Hoeffding's inequality, with probability $1 - \delta/2$,

$$|\hat{a} - a| \leq \frac{\epsilon}{256} \tag{5}$$

Denote by E the intersection of the above two events. By union bound, the probability of E happening is at least $1 - \delta$. Conditioned on event E , the recovery accuracy of the transition matrix and the initial probability are satisfied. We now argue that the recovery accuracy of the binomial probability is also satisfied.

Denote by π the permutation induced by Π , i.e. for a m -dimensional row vector v , $v\Pi = (v_{\pi(1)}, \dots, v_{\pi(m)})$. Then equation (4) implies that for all $i \in \{1, \dots, m\}$,

$$\|C^i - \hat{C}^{\pi(i)}\|_2 \leq \frac{\epsilon}{16\sqrt{D}}.$$

Now, applying Lemma 3, we get that for all $i \in \{1, \dots, m\}$,

$$|p_i - \hat{p}_{\pi(i)}| \leq \epsilon.$$

written in the matrix form,

$$\|p - \hat{p}\Pi\|_2 \leq \epsilon.$$

\square

B Proof of Theorem 1

In this section, we give the proof of Theorem 1, showing that if the coverage is large enough, the binomial probabilities are sufficiently separated, and the discretization parameter D is large enough, then the feature rank condition is well-satisfied.

Proof of Theorem 1. Define vector C^i as the i th column of the expected feature map matrix C . We note that for each i , C^i is the average of $\phi_{\text{beta},D}(c, \mu)$ for different (c, μ) 's. In addition, for every (c, μ) , $\phi_{\text{beta},D}(c, \mu)$ is a probability vector. This implies that for every i , C^i is a probability vector. To show that C is of full column rank, we show that each column C_i is approximately supported on disjoint coordinates. Formally we have the following claim.

Claim 1. *There exist sets $S_i \in \{1, 2, \dots, D\}$ for each $i \in \{1, 2, \dots, m\}$, such that:*

1. S_i 's are disjoint, that is, for all distinct i, j , $S_i \cap S_j = \emptyset$.
2. Each C^i is well-supported on S_i : $\sum_{j \in S_i} C_{i,j} \geq \frac{3}{4}$.

Proof. Fix a column $j \in \{1, \dots, m\}$. Define a random variable μ drawn from $\text{Bin}(c, p_j)$, and conditioned on μ , a random variable G is drawn from the Beta distribution with shape parameters $\mu + 1$ and $c - \mu + 1$. Recall that $\phi_{\text{beta},D}(c, \mu)_i$ is defined as $\int_{(i-1)/D}^{i/D} \varphi_{\text{beta},D}((c, \mu), t) dt$, and is thus equal to $\mathbb{P}[G \in (\frac{i-1}{D}, \frac{i}{D}) | \mu]$. By the law of total expectation,

$$\begin{aligned} C_{i,j} &= \mathbb{E}[\phi_{\text{beta},D}(x)_i | h = j] \\ &= \mathbb{E}[\mathbb{E}[\phi_{\text{beta},D}(x)_i | \mu] | h = j] \\ &= \mathbb{E}[\mathbb{P}[G \in (\frac{i-1}{D}, \frac{i}{D}) | \mu] | h = j] \\ &= \mathbb{P}[G \in (\frac{i-1}{D}, \frac{i}{D})]. \end{aligned}$$

Now, for each $j \in \{1, 2, \dots, m\}$, define set $S_j = \left\{ i \in \{1, \dots, D\} : i \in [D(p_j - 4\sqrt{\frac{\ln 4}{c}}) - 1, D(p_j + 4\sqrt{\frac{\ln 4}{c}}) + 1] \right\}$.

Under this setting of the S_j 's, we now prove the two items in the lemma statement respectively.

- We now show the first item. Consider two distinct indices $j, l \in \{1, 2, \dots, m\}$. Without loss of generality, suppose $p_l < p_j$. We show that the maximum element of S_l is strictly smaller than the minimum element of S_j , thus establishing the disjointness of the two sets. To this end, we show that

$$D(p_j - 4\sqrt{\frac{\ln 4}{c}}) - 1 > D(p_l + 4\sqrt{\frac{\ln 4}{c}}) + 1.$$

The reason is as follows: since $D \geq \frac{4}{q} \geq \frac{4}{p_j - p_l}$, $\frac{2}{D} \leq \frac{p_j - p_l}{2}$. In addition, since $c \geq \frac{512}{q^2} > \frac{256 \ln 4}{(p_j - p_l)^2}$, $8\sqrt{\frac{\ln 4}{c}} < \frac{p_j - p_l}{2}$. Therefore, the above inequality holds.

- For the second item, we apply concentration inequalities on Beta and binomial distributions. First, by Hoeffding's Inequality, with probability $\frac{7}{8}$,

$$\mu \in [cp_j - \sqrt{c \ln 4}, cp_j + \sqrt{c \ln 4}]$$

Second, given μ , G is Beta distributed with shape parameters μ and $c - \mu$, and is thus $\frac{1}{4c}$ -sub-Gaussian [9]. We have that with probability $\frac{7}{8}$,

$$cG \in [\mu - \sqrt{c \ln 4}, \mu + \sqrt{c \ln 4}]$$

Therefore, by union bound and algebra, we have that with probability $\frac{3}{4}$,

$$G \in [p_j - 2\sqrt{\frac{\ln 4}{c}}, p_j + 2\sqrt{\frac{\ln 4}{c}}]$$

This gives that

$$\begin{aligned} \sum_{i \in S_j} C_{i,j} &= \mathbb{P}[G \in \cup_{i \in S_j} (\frac{i-1}{D}, \frac{i}{D})] \\ &\geq \mathbb{P}[G \in [p_j - 2\sqrt{\frac{\ln 4}{c}}, p_j + 2\sqrt{\frac{\ln 4}{c}}]] \\ &\geq \frac{3}{4}. \end{aligned}$$

where the first inequality is by the fact that $[p_j - 2\sqrt{\frac{\ln 4}{c}}, p_j + 2\sqrt{\frac{\ln 4}{c}}]$ is a subset of $\cup_{i \in S_j} (\frac{i-1}{D}, \frac{i}{D})$. □

Provided the above claim holds, we now lower bound the minimum singular value of C . It suffices to show that, for every vector x in \mathbb{R}^m , $\|Cx\|_1 \geq \frac{1}{2}\|x\|_1$. Indeed, if the above is true, then $\|Cx\|_2 \geq \frac{1}{\sqrt{D}}\|Cx\|_1 \geq \frac{1}{2\sqrt{D}}\|x\|_1 \geq \frac{1}{2\sqrt{D}}\|x\|_2$, implying that the minimum singular value of C is at least $\frac{1}{2\sqrt{D}}$.

Consider vector $x \in \mathbb{R}^m$. Let matrix T be such that

$$T_{i,j} = \begin{cases} C_{i,j} & i \in S_j \\ 0 & i \notin S_j \end{cases}$$

and let matrix S to be $C - T$. Note that as the S_j 's are disjoint,

$$\|Tx\|_1 = \sum_j \sum_{i \in S_j} |C_{i,j}x_j| \geq \frac{3}{4}\|x\|_1.$$

Also,

$$\|Sx\|_1 \leq \sum_j \sum_{i \notin S_j} |C_{i,j}x_j| \leq \sum_j \frac{1}{4}|x_j| \leq \frac{1}{4}\|x\|_1.$$

Thus,

$$\begin{aligned} \|Cx\|_1 &= \|Tx + Sx\|_1 \\ &\geq \|Tx\|_1 - \|Sx\|_1 \geq \frac{1}{2}\|x\|_1. \end{aligned}$$

The lemma follows. □