



HAL
open science

Bayesian nonparametric mixture of experts for high-dimensional inverse problems

Trungtin Nguyen, Florence Forbes, Julyan Arbel, Hien Duy Nguyen

► **To cite this version:**

Trungtin Nguyen, Florence Forbes, Julyan Arbel, Hien Duy Nguyen. Bayesian nonparametric mixture of experts for high-dimensional inverse problems. 2023. hal-04015203v1

HAL Id: hal-04015203

<https://hal.science/hal-04015203v1>

Preprint submitted on 5 Mar 2023 (v1), last revised 17 Jun 2024 (v3)

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

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

Bayesian nonparametric mixture of experts for high-dimensional inverse problems

TrungTin Nguyen^{*1}, Florence Forbes¹, Julyan Arbel¹, and Hien Duy Nguyen²

¹Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, Inria Grenoble
Rhone-Alpes, 655 av. de l'Europe, 38335 Montbonnot, France.

²School of Mathematics and Physics, University of Queensland, St. Lucia,
Brisbane, Australia.

Abstract

A large class of problems can be formulated as inverse problems, where the goal is to find parameter values that best explain some observed measures. Typical constraints in practice are that the relationships between parameters and observations are highly non-linear, with high-dimensional observations and multi-dimensionally correlated parameters. To deal with these constraints, we consider probabilistic mixtures of locally linear models using inverse regression strategies, namely the Gaussian locally linear mapping (GLLiM) models. These can be seen as special instances of a mixture of experts (MoE) models. The popularity of MoE is largely due to their universal approximation properties, provided that the number of mixture components is large enough. In this paper, we propose a general scheme to design a tractable Bayesian nonparametric GLLiM (BNP-GLLiM) model to avoid any commitment to an arbitrary number of components. A tractable estimation algorithm is designed using a variational Bayesian expectation-maximization. In particular, we establish posterior consistency for the number of components in BNP-GLLiM after the merge-truncate-merge algorithm post-processing. Illustrations on simulated data show good results in terms of recovering the true number of clusters and the mean regression function.

Keywords. Bayesian nonparametrics, Mixture of experts, High-dimensional statistics, Inverse problems, Gaussian locally-linear mapping models, Linear cluster-weighted models, Variational inference, Clustering, Regression, Model selection.

Contents

1	Introduction	3
1.1	Mixture of experts models	3
1.2	Model selection in MoE	3
1.3	Contributions	4
2	High-dimensional BNP-GLLiM model	5
2.1	Inverse regression framework	5
2.2	High-dimensional regression via GLLiM model	5
2.3	Construction of BNP-GLLiM model	7
2.3.1	BNP prior	7
2.3.2	Hierarchical representation of BNP-GLLiM model	9

^{*}Corresponding author, email: trung-tin.nguyen@inria.fr

3	Variational inference	9
3.1	Variational Bayesian expectation-maximization principle	9
3.2	VBEM for a simple BNP-GLLiM model	11
3.2.1	VB-E- τ step	12
3.2.2	VB-E- (α, σ) step	12
3.2.3	VB-E- \mathcal{Z} step	12
3.2.4	VB-M- (s_1, s_2, a) step	13
3.2.5	VB-M- $(\mathbf{c}, \mathbf{\Gamma})$ step	13
3.2.6	VB-M- $(\mathbf{A}, \mathbf{b}, \mathbf{\Sigma})$ step	14
3.3	VBEM for a general and tractable BNP-GLLiM model	14
3.3.1	VB-E- \mathcal{Z} step	15
3.3.2	VB-E- θ^* step	15
3.3.3	VB-M- ρ step	16
4	Predictive conditional density	16
4.1	BNP-GLLiM1	17
4.1.1	Joint density	17
4.1.2	Inverse conditional density	17
4.1.3	Forward conditional density	18
4.2	BNP-GLLiM2	18
4.2.1	Joint density	18
4.2.2	Inverse conditional density	19
4.2.3	Forward conditional density	19
5	ELBO	20
5.1	BNP-GLLiM1	20
5.2	BNP-GLLiM2	21
6	Nonparametric Bayesian model selection	21
6.1	Posterior contraction rate in Bayesian infinite mixtures	21
6.1.1	Problem setup	21
6.1.2	Wasserstein distance for MM	22
6.1.3	Posterior contraction rate in infinite mixtures	22
6.2	MTM post-processing consistency for BNP-GLLiM	22
6.2.1	Link between GLLiM and joint GMM	23
7	Numerical experiments	25
7.1	Data generation	25
7.2	Model selection, clustering and regression tasks via MTM-BNP-GLLiM1	25
8	Perspectives	28
A	Details of VBEM for BNP-GLLiM model	29
A.1	VB-E- τ step from Section 3.2.1	29
A.2	VB-E- (α, σ) step from Section 3.2.2	30
A.3	VB-E- \mathcal{Z} step from Sections 3.2.3 and 3.3.1	32
A.3.1	Updating $\mathbf{\Sigma}_k$	33
A.4	VB-M- $(\mathbf{c}, \mathbf{\Gamma})$ step from Section 3.2.5	34
A.5	VB-M- $(\mathbf{A}, \mathbf{b}, \mathbf{\Sigma})$ step from Section 3.2.6	35
A.5.1	Updating \mathbf{b}_k	35
A.5.2	Updating \mathbf{A}_k	36

B Technical proofs	37
B.1 Proof of Lemma 6.2	37
B.2 Proof of Theorem 4.1	38
B.3 Proof of Theorem 4.2	40
B.4 Proof of Theorem 4.3	41
B.5 Proof of Theorem 4.4	43
B.5.1 Proof of Theorem B.1	44
B.6 Proof of Theorem 4.5	47
B.7 Proof of Theorem 4.6	47
B.8 Proof of Theorem 5.1	49
B.9 Proof of Theorem 5.2	53

1 Introduction

1.1 Mixture of experts models

Many problems can be formulated as inverse problems, where the goal is to find parameter values that best explain observations. Typical constraints in practice are that the relationships between parameters and observations are highly non-linear, with high-dimensional observations and multi-dimensionally correlated parameters. To handle these constraints, we consider probabilistic mixtures of locally linear models using inverse regression strategies, namely the Gaussian locally linear mapping (GLLiM) model in [1], which can be seen as special affine instances of the mixture of experts (MoE) models [2].

Note that these MoEs are alternative neural network architectures to the original one of [3] and [4]. These flexible models are generalizations from the classical mixture models (MM) and mixture of regression (MoR) models [5]. Their flexibility comes from the fact that they allow the mixture weights (or the gating functions) to depend on the explanatory variables, together with the component densities (or the experts). In the context of regression, MoE models with Gaussian experts and softmax or normalized Gaussian gating functions (including GLLiM) are the most popular choices, see *e.g.*, [6, 7, 8] for a comprehensive classification and nomenclature. They are powerful tools for modelling more complex nonlinear relationships between outputs (responses) and inputs (predictors) that arise from different subpopulations.

The popularity of these conditional mixture density models is largely due to their universal approximation properties as well as their good convergence rate, as studied in [9, 10, 11, 12, 13, for MoE], and in [14, 15, for MoR]. It is worth noting that these results improve the approximation capabilities of unconditional MMs, as discussed in [16, 17, 18, 19, 20, 21, 22]. At a high level for universal approximation theorems, given a large enough number of components, MM and MoE models can approximate a large class of unconditional and conditional probability density functions (CPDF), respectively, to any degree of accuracy. See, *e.g.*, [23, 24, 25, 6], for further detailed reviews of practical and theoretical aspects of MoE models.

1.2 Model selection in MoE

Although universal approximation allows us to conclude that, given a sufficient number of components, a finite MoE can approximate any other CPDF to an arbitrary degree of accuracy, it is not clear how to choose a large enough number of components for realistic problems. This motivates a careful study of interesting and important model selection problems for MM and MoE models, which have attracted much attention in statistics and machine learning over the last 50 years, see *e.g.*, [26] for a recent comprehensive review.

When selecting the best data-driven number of components for MoE models, there are several approaches to controlling and accounting for model complexity. Typically, model selection is performed using the Akaike information criterion (AIC; [27]), the Bayesian information criterion (BIC; [28], BIC-GLLiM; [29, 30]) or the BIC-like approximation of integrated classification likelihood (ICL-BIC; [31]).

But an important limitation of these criteria is that they are only asymptotically valid. This means that there are no finite sample guarantees when using AIC, ICL-BIC or BIC to choose between different levels of complexity. Therefore, their use in small sample settings is an ad hoc one.

To overcome such difficulties, and to partially support a recent approach called the slope heuristic of [32], also see [33] for a recent review, the authors of [34, 7, 35, 36, 8] established non-asymptotic risk bounds in the form of weak oracle inequalities, provided that lower bounds on the penalties hold, in high-dimensional regression scenarios for a variety of MoE models, including GLLiM. Another approach is based on the closed testing principle leads to a sequential testing procedure that allows for confidence statements to be made regarding the order of a finite mixture model from [37]. These works lead to an optimal data-driven choice of number of components in finite-sample setting. However, all previous approaches require that a range of models with different values be trained and compared, which can be a computational bottleneck in a high-dimensional framework.

Recently, [38] proposed computationally efficient variational inference approaches to architecture selection in high-dimensional deep Gaussian mixture models using overfitted mixtures, see *e.g.*, [39, 40], where unnecessary components are dropped in the estimation. However, in this work, we are more interested in the more general context of the Bayesian nonparametric (BNP) approach, see *e.g.*, [41, 42], where it is not necessary to know the upper bound of the true number of components as in Bayesian overfitted MM. This is one motivation for the BNP priors that are under consideration here for GLLiM.

Dirichlet process mixture models (DP-MM) and Pitman-Yor process mixture models (PYP-MM) are among the most popular BNP models, particularly suitable for density estimation and probabilistic clustering. However, the posterior of DP-MM or PYP-MM is inconsistent on the number of components if the true number of components is finite and the concentration parameter is known and fixed, see *e.g.*, [43]. This is because a BNP prior such as DP or PYP places zero probability on mixing measures with a finite number of supporting atoms. An interesting recent result in [44] is that consistency for the number of components can be achieved if we impose some assumptions on the prior for the concentration parameter of the DP-MM.

It appears that BNP-GLLiM tends to produce many small extraneous components around the true clusters in our numerical experiment Section 7. This makes it difficult to use them to infer the true number of components when this becomes a quantity of interest [45, 46]. This encourages the use of a novel, simple MTM post-processing algorithm described in [47]. This post-processing consistently estimates the number of components for any general Bayesian prior, even without knowing its exact structure, as long as the posterior for that prior contracts to the true mixing distribution at a known rate.

1.3 Contributions

To deal with high-dimensional and large-scale data, our novel contribution is to propose a general scheme for designing a tractable BNP-GLLiM model. This scheme avoids any commitment to an arbitrary number of components. In particular, it allows a single training run to automatically select the optimal data-driven number of components with respect to the trade-off between model complexity and fit to the data. Its special case, called the BNP mixture model (BNP-MM), has been well studied in [48, 49, 50, 51] but very few results for BNP-GLLiM. This motivates us to design a tractable, theoretically guaranteed estimation algorithm using variational Bayesian expectation maximization (VBEM), see *e.g.*, [52] for a recent review of variational inference.

In particular, to the best of our knowledge, we are the first to establish the consistency theoretical properties for recovering the true number of components in BNP-GLLiM using the post-processing merge-truncate-merge (MTM) algorithm. Finally, our illustrations on simulated data show good results in terms of recovering the true number of clusters and mean regression functions. It is worth emphasising that, for the first time, we provide evidence that MTM consistency holds not only for the MM results of [47], but also in the more general context of MoE models for CPDFs.

Notations. Throughout this paper, $\{1, \dots, D\}$ is abbreviated as $[D]$ for $D \in \mathbb{N}^*$, where \mathbb{N}^* denotes the positive natural numbers. Next, the notation \equiv refers to a definition. It is used to simplify the notation or expression. For a parametric model S , $\dim(S)$ refers to its dimension, *i.e.*, the total number of parameters to be estimated. Furthermore, $[\cdot; \cdot]$ denotes vertical vector concatenation. A coupling between $\boldsymbol{\pi} \equiv (\pi_k)_{k \in [K]}$ and $\boldsymbol{\pi}^0 \equiv (\pi_l^0)_{l \in [K_0]}$ is a joint distribution \mathbf{Q} on $[K] \times [K_0]$, which is expressed as a matrix $\mathbf{Q} = (q_{kl})_{k \in [K], l \in [K_0]} \in [0, 1]^{K \times K_0}$ with marginal probabilities $\sum_{k=1}^K q_{kl} = \pi_l^0$ and $\sum_{l=1}^{K_0} q_{kl} = \pi_k$ for any $k \in [K]$ and $l \in [K_0]$. We use $\mathcal{Q}(\boldsymbol{\pi}, \boldsymbol{\pi}^0)$ to denote the space of all such couplings. Regarding the space of mixing measures, let $\mathcal{E}_K \equiv \mathcal{E}_K(\boldsymbol{\Theta})$ and $\mathcal{O}_K \equiv \mathcal{O}_K(\boldsymbol{\Theta})$ respectively denote the space of all mixing measures with exactly and at most K support points, all in some parameter space $\boldsymbol{\Theta}$. Additionally, denote $\mathcal{G} \equiv \mathcal{G}(\boldsymbol{\Theta}) = \bigcup_{K \in \mathbb{N}_+} \mathcal{E}_K$ the set of all discrete measures with finite supports on $\boldsymbol{\Theta}$. Moreover, $\overline{\mathcal{G}}(\boldsymbol{\Theta})$ denotes the space of all discrete measures (including those with countably infinite supports) on $\boldsymbol{\Theta}$. Finally, $\mathcal{P}(\boldsymbol{\Theta})$ stands for the space of all probability measures on $\boldsymbol{\Theta}$.

Paper organization. The paper is organized as follows. In [Section 2](#), we first discuss how to construct high-dimensional BNP-GLLiM models. VBEM algorithms, predictive CPDFs and ELBO for two useful BNP-GLLiM models are then described correspondingly in [Sections 3 to 5](#). Next, [Section 6](#) shows how we can integrate a post-processing and prove MTM consistency for BNP-GLLiM. This is useful to perform regression, clustering and model selection simultaneously. We experimentally evaluate our new results in simulated datasets in [Section 7](#). Some perspectives are provided in [Section 8](#). All details of VBEM for the BNP-GLLiM model and technical proofs not included in the main paper are relegated to [Appendix A](#) and [Appendix B](#) respectively.

2 High-dimensional BNP-GLLiM model

2.1 Inverse regression framework

We are interested in estimating the law of the low-dimensional random variable $\mathbf{X} = (\mathbf{X}_l)_{l \in [L]}$ conditionally on high-dimensional $\mathbf{Y} = (\mathbf{Y}_d)_{d \in [D]}$, where typically $D \gg L$. We follow an inverse regression framework from [\[53, 1\]](#). Therefore, in training, the low-dimensional variable \mathbf{X} plays the role of the regressor, while the response \mathbf{Y} is a function of \mathbf{X} , possibly corrupted by noise through inverse CPDF $p(\mathbf{Y} | \mathbf{X}; \boldsymbol{\psi})$, where $\boldsymbol{\psi}$ is an inverse parameter. The low dimension of the regressor \mathbf{X} drastically reduces the number of parameters to be estimated. In particular, the forward parameter $\boldsymbol{\psi}^*$ and CPDF $p(\mathbf{X} | \mathbf{Y}; \boldsymbol{\psi}^*)$ are tractable after estimating the inverse parameter $\boldsymbol{\psi}$. Therefore, this density can be used to predict the low-dimensional response \mathbf{x} of a high-dimensional test point \mathbf{y} . This inverse-then-forward regression strategy justifies the unconventional notation: \mathbf{Y} for the high-dimensional input and \mathbf{X} for the low-dimensional response.

Here and subsequently, we refer to the low-dimensional data sample as $\mathcal{X} \equiv \{\mathbf{x}_n\}_{n \in [N]} \subset (\mathbb{R}^L)^N$, the high-dimensional data sample as $\mathcal{Y} \equiv \{\mathbf{y}_n\}_{n \in [N]} \subset (\mathbb{R}^D)^N$. We further denote the realisations by (\mathbf{x}, \mathbf{y}) , which are generated from two random variables \mathbf{X} and \mathbf{Y} .

Throughout the paper, the following assumptions will be required for the training phase. We assume that the covariates \mathbf{X} are independent but not necessarily identically distributed. The assumptions on the responses \mathbf{Y} are stronger: conditional on \mathcal{X} , the \mathcal{Y} , are independent, and each \mathbf{Y} follows a law with true (but unknown) inverse CPDF $s_0(\mathbf{Y} | \mathbf{X} = \mathbf{x})$, which is approximated via GLLiM. In [Section 2.2](#) we first have a more detailed recall of GLLiM model. Then in [Section 2.3](#) we construct the BNP-GLLiM model.

2.2 High-dimensional regression via GLLiM model

The GLLiM models, as originally introduced in [\[1\]](#), are used to capture the non-linear relationship between the response and the set of covariates from a high-dimensional heterogeneous data. More

specifically, the authors of [1] overcame the difficulty of high-to-low regression by tackling the problem the other way round, *i.e.*, low-to-high. This means that the roles of input and response variables are swapped so that the low-dimensional variable \mathbf{X} becomes the regressor as in Section 2.1. GLLiM then relies on a piecewise linear model in the following way. The high-dimensional response \mathbf{Y} is approximated by the local affine mappings K :

$$\mathbf{Y} = \sum_{k=1}^K \mathbb{I}(Z = k) (\mathbf{A}_k \mathbf{X} + \mathbf{b}_k + \mathbf{E}_k). \quad (2.1)$$

Here \mathbb{I} is an indicator function and Z is a latent variable that captures a cluster relationship, such that $Z = k$ if \mathbf{Y} comes from cluster $k \in [K]$. Matrices $\mathbf{A}_k \in \mathbb{R}^{D \times L}$ and vectors $\mathbf{b}_k \in \mathbb{R}^D$ define cluster-specific affine transformations. In addition, \mathbf{E}_k are error terms that capture both the reconstruction error due to the local affine approximations as well as the observation noise in \mathbb{R}^D .

Following the usual assumption that \mathbf{E}_k is a zero-mean Gaussian variable with a covariance matrix $\boldsymbol{\Sigma}_k \in \mathbb{R}^{D \times D}$, it follows that

$$p(\mathbf{y} \mid \mathbf{x}, Z = k; \boldsymbol{\psi}) = \mathcal{N}_D(\mathbf{y} \mid \mathbf{A}_k \mathbf{x} + \mathbf{b}_k, \boldsymbol{\Sigma}_k), \quad (2.2)$$

where $\boldsymbol{\psi}$ is the vector of model parameters, $\mathcal{N}_D(\mathbf{y}; \mathbf{A}_k \mathbf{x} + \mathbf{b}_k, \boldsymbol{\Sigma}_k)$ is the Gaussian CPDF of dimension D . In order to enforce the affine transformations to be local, \mathbf{X} is defined as a mixture of K Gaussian components as follows:

$$p(\mathbf{x} \mid Z = k; \boldsymbol{\psi}) = \mathcal{N}_L(\mathbf{x} \mid \mathbf{c}_k, \boldsymbol{\Gamma}_k), \quad p(Z = k; \boldsymbol{\psi}) = \pi_k, \quad (2.3)$$

where $\mathbf{c}_k \in \mathbb{R}^L$, $\boldsymbol{\Gamma}_k \in \mathbb{R}^{L \times L}$, and $\boldsymbol{\pi} = (\pi_k)_{k \in [K]} \in \left\{ (\pi_k)_{k \in [K]} \in (\mathbb{R}^+)^K, \sum_{k=1}^K \pi_k = 1 \right\} \equiv \boldsymbol{\Pi}_{K-1}$, which is the $K - 1$ dimensional probability simplex. Then, according to the formulas for conditional multivariate Gaussian variables and the hierarchical decomposition

$$\begin{aligned} p(\mathbf{y}, \mathbf{x}; \boldsymbol{\psi}) &= \sum_{k=1}^K p(\mathbf{y} \mid \mathbf{x}, Z = k; \boldsymbol{\psi}) p(\mathbf{x} \mid Z = k; \boldsymbol{\psi}) p(Z = k; \boldsymbol{\psi}), \\ &= \sum_{k=1}^K \pi_k \mathcal{N}_L(\mathbf{x} \mid \mathbf{c}_k, \boldsymbol{\Gamma}_k) \mathcal{N}_D(\mathbf{y} \mid \mathbf{A}_k \mathbf{x} + \mathbf{b}_k, \boldsymbol{\Sigma}_k), \end{aligned}$$

we obtain the following inverse conditional density, defined as a GLLiM model [1]:

$$p(\mathbf{y} \mid \mathbf{x}; \boldsymbol{\psi}) = \sum_{k=1}^K \frac{\pi_k \mathcal{N}_L(\mathbf{x} \mid \mathbf{c}_k, \boldsymbol{\Gamma}_k)}{\sum_{j=1}^K \pi_j \mathcal{N}_L(\mathbf{x} \mid \mathbf{c}_j, \boldsymbol{\Gamma}_j)} \mathcal{N}_D(\mathbf{y} \mid \mathbf{A}_k \mathbf{x} + \mathbf{b}_k, \boldsymbol{\Sigma}_k), \quad (2.4)$$

where $\boldsymbol{\psi} = (\pi_k, \mathbf{c}_k, \boldsymbol{\Gamma}_k, \mathbf{A}_k, \mathbf{b}_k, \boldsymbol{\Sigma}_k)_{k \in [K]} \in \boldsymbol{\Psi}$.

Without assuming anything about the structure of the parameters, the dimension of GLLiM is given by

$$\dim(\boldsymbol{\Psi}) = K \left(1 + D(L + 1) + \frac{D(D + 1)}{2} + \frac{L(L + 1)}{2} + L \right) - 1.$$

It is worth noting that $\dim(\boldsymbol{\Psi})$ is very large compared to the sample size (see *e.g.*, [1] for more details in their real data sets) whenever D is large and $D \gg L$. Furthermore, under the assumption that the K transformations are affine, it is more realistic to make the assumption on the residual covariance matrices $\boldsymbol{\Sigma}_k$ of the error vectors \mathbf{E}_k rather than on $\boldsymbol{\Gamma}_k$ (*cf.*, [1, Section 3]). This justifies the use of the inverse regression trick from [1], which leads to a drastic reduction in the number of parameters to be estimated. For instance, \mathbf{E}_k can be modelled with equal isotropic Gaussian noise, so we have $\boldsymbol{\Sigma}_k = \sigma^2 \mathbf{I}_D, \forall k \in [K]$. The number of parameters to be estimated, *i.e.* the size of $\boldsymbol{\psi}$, is $K - 1 + K(L + L(L + 1)/2 + DL + D)$. For example, it is 30,060 if $K = 10, L = 2, D = 1000$.

However, if a high-to-low regression is estimated directly instead, the size of the parameter vector will be $K - 1 + K(D + LD + D(D + 1)/2 + L)$, which is 5,035,030 in the previous example.

More specifically, in (2.2), the roles of input and response variables should be reversed so that \mathbf{Y} becomes the covariate and \mathbf{X} plays the role of the multivariate response. Therefore, based on a similar argument to the previous hierarchical Gaussian MM (GMM), its corresponding forward conditional density from \mathbb{R}^D to \mathbb{R}^L is defined as follows

$$p(\mathbf{x} | \mathbf{y}; \boldsymbol{\psi}^*) = \sum_{k=1}^K \frac{\pi_k^* \mathcal{N}_D(\mathbf{y} | \mathbf{c}_k^*, \boldsymbol{\Gamma}_k^*)}{\sum_{j=1}^K \pi_j^* \mathcal{N}_D(\mathbf{y} | \mathbf{c}_j^*, \boldsymbol{\Gamma}_j^*)} \mathcal{N}_L(\mathbf{x} | \mathbf{A}_k^* \mathbf{y} + \mathbf{b}_k^*, \boldsymbol{\Sigma}_k^*). \quad (2.5)$$

A useful feature of GLLiM models is described in Lemma 2.1, which is established in [7, Lemma 2.1] and [1], elliptical distribution and multivariate extensions of Gaussian in [54], respectively.

Lemma 2.1. *The parameter $\boldsymbol{\psi}^*$ in the inverse CPDF, defined in (2.5), can then be deduced from $\boldsymbol{\psi}$ in (2.4) via the following one-to-one correspondence:*

$$\begin{pmatrix} \pi_k \\ \mathbf{c}_k \\ \boldsymbol{\Gamma}_k \\ \mathbf{A}_k \\ \mathbf{b}_k \\ \boldsymbol{\Sigma}_k \end{pmatrix}_{k \in [K]} \mapsto \begin{pmatrix} \pi_k^* \\ \mathbf{c}_k^* \\ \boldsymbol{\Gamma}_k^* \\ \mathbf{A}_k^* \\ \mathbf{b}_k^* \\ \boldsymbol{\Sigma}_k^* \end{pmatrix}_{k \in [K]} = \begin{pmatrix} \pi_k \\ \mathbf{A}_k \mathbf{c}_k + \mathbf{b}_k \\ \boldsymbol{\Sigma}_k + \mathbf{A}_k \boldsymbol{\Gamma}_k \mathbf{A}_k^\top \\ \boldsymbol{\Sigma}_k^* \mathbf{A}_k^\top \boldsymbol{\Sigma}_k^{-1} \\ \boldsymbol{\Sigma}_k^* (\boldsymbol{\Gamma}_k^{-1} \mathbf{c}_k - \mathbf{A}_k^\top \boldsymbol{\Sigma}_k^{-1} \mathbf{b}_k) \\ (\boldsymbol{\Gamma}_k^{-1} + \mathbf{A}_k^\top \boldsymbol{\Sigma}_k^{-1} \mathbf{A}_k)^{-1} \end{pmatrix}_{k \in [K]}. \quad (2.6)$$

2.3 Construction of BNP-GLLiM model

2.3.1 BNP prior

Stick-breaking construction of Dirichlet process. Note that the Dirichlet process (DP) [55] is a central BNP prior and is the infinite-dimensional generalization of the Dirichlet distribution. Therefore, for the sake of completeness, let us first recall the definition of the DP. A DP on the space \mathcal{G} is defined as a random process characterized by a concentration parameter α and a base distribution G_0 , denoted by $G \sim \text{DP}(\alpha, G_0)$, such that for any finite partition $\{A_1, \dots, A_p\}$ of \mathcal{G} , the random vector $(G(A_1), \dots, G(A_p))$ is Dirichlet distributed:

$$(G(A_1), \dots, G(A_p)) \sim \text{Dir}(\alpha G_0(A_1), \dots, \alpha G_0(A_p)). \quad (2.7)$$

A DP prior G can be constructed by three methods: the Blackwell-MacQueen urn scheme, the Chinese restaurant method, and the stick-breaking construction. We will use the stick-breaking construction of DP (SBDP):

$$\boldsymbol{\theta}_k^0 | G_0 \stackrel{\text{i.i.d.}}{\sim} G_0, \quad k \in \mathbb{N}^*, \quad (2.8)$$

$$\tau_k | \alpha \stackrel{\text{i.i.d.}}{\sim} \text{Beta}(\tau_k | 1, \alpha), \quad k \in \mathbb{N}^*, \quad (2.9)$$

$$\pi_k(\boldsymbol{\tau}) = \tau_k \prod_{l=1}^{k-1} (1 - \tau_l), \quad k \in \mathbb{N}^*, \quad (2.10)$$

$$G = \sum_{k=1}^{\infty} \pi_k(\boldsymbol{\tau}) \delta_{\boldsymbol{\theta}_k^0} \sim \text{DP}(\alpha, G_0). \quad (2.11)$$

Pitman–Yor process. As a generalized version of SBDP, in the Pitman–Yor process (PYP) [56], the τ_k 's are independent ($\overset{\text{ind}}{\sim}$) but not identically distributed. More specially,

$$\theta_k^0 \mid G_0 \overset{\text{i.i.d.}}{\sim} G_0, \quad k \in \mathbb{N}^*, \quad (2.12)$$

$$\tau_k \mid \alpha, \sigma \overset{\text{ind}}{\sim} \text{Beta}(1 - \sigma, \alpha + k\sigma), \quad k \in \mathbb{N}^*, \quad (2.13)$$

$$\pi_k(\boldsymbol{\tau}) = \tau_k \prod_{l=1}^{k-1} (1 - \tau_l), \quad k \in \mathbb{N}^*, \quad (2.14)$$

$$G = \sum_{k=1}^{\infty} \pi_k(\boldsymbol{\tau}) \delta_{\theta_k^0} \sim \text{PYP}(\alpha, \sigma, G_0). \quad (2.15)$$

Here $\sigma \in (0, 1)$ is a discount parameter and α is a concentration parameter $\alpha > \sigma$. The PYP is a two-parameter generalisation of the DP that allows one to control the tail behaviour when modelling data with either exponential or power-law tails [57, 56]. The PYP is reduced to a DP when $\sigma = 0$. More general stick-breaking representations are possible, *e.g.*, Gibbs-type priors [58, 42] or homogeneous normalised random measures with independent increments [59]). However, in terms of the number of clusters, the PYP provides a clear interpretation. The rich-gets-richer property of the DP is preserved, meaning that there are a small number of large clusters, but there are also a large number of small clusters. The parameter σ reduces the probability that observations join small clusters. The PYP has a power-law behaviour. This can make it more suitable for a number of applications. In other words, the number of clusters grows as $\mathcal{O}(N^\sigma)$ for PYP, while growing more slowly as $\mathcal{O}(\log N)$ for DP.

Since the hyperparameters α and σ can have a significant effect on the growth of the number of clusters with data sample size, it is possible to specify priors for them. For the DP case obtained with $\sigma = 0$, it is suggested in [60] to use a gamma prior, $\alpha \sim \text{Gam}(s_1, s_2)$, where the hyperparameters s_1 and s_2 can be estimated or fixed. A natural question is whether one can also find a tractable prior for the discount parameter σ . Following the work of [48], we use the following prior that satisfies the constraints $\sigma \in (0, 1)$ and $\alpha > -\sigma$,

$$p(\alpha, \sigma \mid s_1, s_2, a) = p(\alpha \mid \sigma; s_1, s_2) p(\sigma \mid a) \quad (2.16)$$

where $p(\alpha \mid \sigma; s_1, s_2)$ is a shifted gamma distribution $\mathcal{SG}(\alpha \mid \sigma; s_1, s_2)$ and $p(\sigma, a)$ is a distribution depending on some parameter a which is not specified at the moment but which can typically be assumed to be a uniform distribution on the interval $(0, 1)$. Such a shifted gamma distribution is the distribution of a variable $U - \sigma$, where σ is considered fixed and U follows a gamma distribution $\text{Gam}(s_1, s_2)$. The PDF of this shifted gamma distribution is obtained from the standard gamma distribution as $\mathcal{SG}(\alpha \mid \sigma; s_1, s_2) = \text{Gam}(\alpha + \sigma \mid s_1, s_2)$.

2.3.2 Hierarchical representation of BNP-GLLiM model

We propose the following hierarchical representation to generate a data point $(\mathbf{y}_n, \mathbf{x}_n)$ within our BNP-GLLiM model:

1. Sample from BNP prior:

$$(\alpha, \sigma) \mid s_1, s_2, a \sim \mathcal{SG}(\alpha \mid \sigma; s_1, s_2) p(\sigma \mid a) \equiv \text{Gam}(\alpha + \sigma \mid s_1, s_2) p(\sigma \mid a), \quad (2.17)$$

$$\tau_k \mid \alpha, \sigma \stackrel{\text{i.i.d.}}{\sim} \text{Beta}(\tau_k \mid 1 - \sigma, \alpha + k\sigma), \quad k \in \mathbb{N}^*, \quad (2.18)$$

$$\pi_k(\boldsymbol{\tau}) = \tau_k \prod_{l=1}^{k-1} (1 - \tau_l), \quad k \in \mathbb{N}^*, \quad (2.19)$$

$$\boldsymbol{\theta}_k^* \mid G_0 \stackrel{\text{i.i.d.}}{\sim} G_0, \quad \boldsymbol{\theta}_k^* \equiv (\mathbf{c}_k, \boldsymbol{\Gamma}_k, \mathbf{A}_k, \mathbf{b}_k, \boldsymbol{\Sigma}_k), \quad k \in \mathbb{N}^*, \quad (2.20)$$

$$G = \sum_{k=1}^{\infty} \pi_k(\boldsymbol{\tau}) \delta_{\boldsymbol{\theta}_k^*} \sim \text{BNP}(\alpha, \sigma, G_0), \quad (2.21)$$

Sample from BNP-GLLiM model: for each $n \in [N]$,

$$\boldsymbol{\theta}_n \mid G \stackrel{\text{i.i.d.}}{\sim} G, \quad (2.22)$$

$$z_n = k \text{ if } \boldsymbol{\theta}_n = \boldsymbol{\theta}_k^* \text{ with probability } \pi_k(\boldsymbol{\tau}), \quad (2.23)$$

$$\mathbf{x}_n \mid z_n, \mathbf{c}, \boldsymbol{\Gamma} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}_L(\mathbf{x}_n \mid \mathbf{c}_{z_n}, \boldsymbol{\Gamma}_{z_n}), \quad (\mathbf{c}, \boldsymbol{\Gamma}) \equiv (\mathbf{c}_k, \boldsymbol{\Gamma}_k)_{k \in \mathbb{N}^*}, \quad (2.24)$$

$$\mathbf{y}_n \mid \mathbf{x}_n, z_n, \mathbf{A}, \mathbf{b}, \boldsymbol{\Sigma} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}_D(\mathbf{y}_n \mid \mathbf{A}_{z_n} \mathbf{x}_n + \mathbf{b}_{z_n}, \boldsymbol{\Sigma}_{z_n}), \quad (\mathbf{A}, \mathbf{b}, \boldsymbol{\Sigma}) \equiv (\mathbf{A}_k, \mathbf{b}_k, \boldsymbol{\Sigma}_k)_{k \in \mathbb{N}^*}. \quad (2.25)$$

Hierarchical representation of BNP-MM. BNP-MM, including DP-MM and PYP-MM, see *e.g.*, [48, 51], has the following hierarchical representation to generate a data point \mathbf{x}_n as a special case of BNP-GLLiM:

1. Sample from BNP prior: (2.26)

$$G = \sum_{k=1}^{\infty} \pi_k(\boldsymbol{\tau}) \delta_{\boldsymbol{\theta}_k^*} \sim \text{BNP}(\alpha, \sigma, G_0), \quad \boldsymbol{\theta}_k^* = (\boldsymbol{\mu}_k, \mathbf{V}_k), \quad (2.27)$$

2. Sample from BNP-MM: for each $n \in [N]$,

$$\boldsymbol{\theta}_n \mid G \stackrel{\text{i.i.d.}}{\sim} G, \quad (2.28)$$

$$z_n = k \text{ if } \boldsymbol{\theta}_n = \boldsymbol{\theta}_k^* \text{ with probability } \pi_k(\boldsymbol{\tau}), \quad (2.29)$$

$$\mathbf{x}_n \mid z_n, \boldsymbol{\theta}^* \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}_L(\mathbf{x}_n \mid \boldsymbol{\theta}_{z_n}^*), \quad \boldsymbol{\theta}^* \equiv (\boldsymbol{\theta}_k^*)_{k \in [K]}. \quad (2.30)$$

It is worth noting that the explicit use of the labels $\mathcal{Z} = (z_1, \dots, z_n)$ in the above BNP-GLLiM construction brings it closer to clustering generative models and opens the way to an HMRF extension, see *e.g.*, [48, 51].

3 Variational inference

We will briefly recall the variational principle in [Section 3.1](#).

3.1 Variational Bayesian expectation-maximization principle

The clustering task consists mainly of estimating the unknown labels $\mathcal{Z} = (z_n)_{n \in [N]}$ from the observed data $(\mathcal{Y}, \mathcal{X}) = (\mathbf{y}_n, \mathbf{x}_n)_{n \in [N]}$, whose joint distribution $p(\mathcal{Y}, \mathcal{X}, \mathcal{Z} \mid \boldsymbol{\Theta}; \phi)$ is determined by a set of parameters denoted by $\boldsymbol{\Theta}$ and often by additional hyperparameters ϕ .

The expectation-maximization (EM) algorithm [61, 62] is a generative technique for maximum likelihood estimation (MLE) in the presence of unobserved latent variables or missing data. An EM iteration consists of two steps usually referred to as the E-step in which the expectation of the so-called complete log-likelihood is computed and the M-step in which this expectation is maximized

over Θ . An equivalent way to define EM is the following. As discussed in [63], EM can be viewed as an alternating maximization procedure of a function \mathcal{F}_0 defined, for any probability distribution $q_{\mathcal{Z}}$ over labels \mathcal{Z} , by

$$\mathcal{F}_0(q_{\mathcal{Z}}, \Theta, \phi) = \sum_{\mathcal{Z}} q_{\mathcal{Z}}(\mathcal{Z}) \log p(\mathcal{Y}, \mathcal{X}, \mathcal{Z} | \Theta, \phi) - \mathbb{E}_{q_{\mathcal{Z}}} [\log q_{\mathcal{Z}}(\mathcal{Z})] \quad (3.1)$$

$$= \mathbb{E}_{q_{\mathcal{Z}}} \left[\log \frac{p(\mathcal{Y}, \mathcal{X}, \mathcal{Z} | \Theta, \phi)}{q_{\mathcal{Z}}(\mathcal{Z})} \right] \quad (3.2)$$

where $-\mathbb{E}_{q_{\mathcal{Z}}} [\log q_{\mathcal{Z}}(\mathcal{Z})]$ is the entropy of $q_{\mathcal{Z}}$ and $\mathbb{E}_q[\cdot]$ is the expectation with respect to q . The function \mathcal{F}_0 depends on the observations $(\mathcal{Y}, \mathcal{X})$, which are fixed throughout and are therefore omitted from the notation.

When prior knowledge on the parameters is available, an alternative approach consists of replacing the MLE by a maximum a posteriori (MAP) estimation of Θ using the prior knowledge encoded in a distribution $p(\Theta)$. More precisely, the MLE of Θ is then replaced by a point estimation $\hat{\Theta} = \operatorname{argmax}_{\Theta \in \Theta} p(\Theta | \mathcal{Y}, \mathcal{X})$. In this paper, instead of considering only point estimation of Θ , we carry out a fully Bayesian approach. In this case, we have to compute

$$p(\mathcal{Z} | \mathcal{Y}, \mathcal{X}) = \int_{\Theta} p(\mathcal{Z} | \mathcal{Y}, \mathcal{X}, \Theta) p(\Theta | \mathcal{Y}, \mathcal{X}) d\Theta. \quad (3.3)$$

Integrating out Θ in this way requires the computation of the marginal point distribution $p(\Theta | \mathcal{Y}, \mathcal{X})$ which is usually not available in closed-form. As an alternative to costly simulation-based methods (*e.g.*, Markov Chain Monte Carlo (MCMC)), an EM-like procedure using variational approximation (VA) can provide approximations of the marginal posterior distributions $p(\Theta | \mathcal{Y}, \mathcal{X})$ and $p(\mathcal{Z} | \mathcal{Y}, \mathcal{X})$. This approach is referred to as VBEM for variational Bayesian EM, was introduced by [64].

To deal with the BNP-GLLiM model, we need to use the VBEM with hyperparameter optimization of [65, Figure 2.5 and Algorithm 5.3]. Let $q_{\mathcal{Z}}$ and q_{Θ} denote the distributions over \mathcal{Z} and Θ respectively, which will serve as approximations to the true posteriors. More specifically, in the Bayesian setting, the intractable posterior $p(\mathcal{Z}, \Theta | \mathcal{Y}, \mathcal{X}; \phi)$ is approximated by the variational posterior $q(\mathcal{Z}, \Theta) = q_{\mathcal{Z}}(\mathcal{Z})q_{\Theta}(\Theta)$.

Similar to standard EM, VBEM maximizes the following evidence lower bound (often abbreviated ELBO, sometimes called the variational lower bound or negative variational free energy), defined for arbitrary $q_{\mathcal{Z}}$ and q_{Θ} distributions by

$$\begin{aligned} \mathcal{F}(q_{\mathcal{Z}}, q_{\Theta}, \phi) &= \mathbb{E}_{q_{\mathcal{Z}}q_{\Theta}} \left[\log \frac{p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \phi)}{q_{\mathcal{Z}}(\mathcal{Z})q_{\Theta}(\Theta)} \right] \\ &= \log p(\mathcal{Y}, \mathcal{X} | \phi) - \text{KL}(q_{\mathcal{Z}}q_{\Theta} \| p(\mathcal{Z}, \Theta | \mathcal{Y}, \mathcal{X}, \phi)) \leq \log p(\mathcal{Y}, \mathcal{X} | \phi), \end{aligned} \quad (3.4)$$

alternatively over $q_{\mathcal{Z}}, q_{\Theta}$ and ϕ . Here, KL stands for Kullback-Leibler divergence. It is worth noting that adding a prior on Θ is formally equivalent to considering Θ as missing variables, while the hyperparameters ϕ play the role of the parameters of interest in MLE.

The alternate maximization over \mathcal{F} leads to the VBEM algorithm, which can be decomposed into three steps. It is easy to show, using the KL divergence properties, that the maximization over $q_{\mathcal{Z}}$ and q_{Θ} leads to the following E-steps, see *e.g.*, [66, Appendix A], [65, Theorem 2.1] and [67, Section 10.1.1], which is essentially coordinate ascent in the function space of variational distributions. Furthermore, the following update rules for E-steps converge to a local maximum of $\mathcal{F}(q_{\mathcal{Z}}, q_{\Theta}, \phi)$. At the r th iteration, using current values $\phi^{(r-1)}$ and $q_{\Theta}^{(r-1)}$, we get the following updating,

$$\text{VB-E-}\mathcal{Z}: q_{\mathcal{Z}}^{(r)}(\mathcal{Z}) \propto \exp \mathbb{E}_{q_{\Theta}^{(r-1)}} \left[\log p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \phi^{(r-1)}) \right], \quad (3.5)$$

$$\text{VB-E-}\Theta: q_{\Theta}^{(r)}(\Theta) \propto \exp \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} \left[\log p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \phi^{(r-1)}) \right], \quad (3.6)$$

$$\text{VB-M-}\phi: \phi^{(r)} \propto \operatorname{argmax}_{\phi} \mathbb{E}_{q_{\mathcal{Z}}^{(r)}q_{\Theta}^{(r)}} [\log p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \phi)]. \quad (3.7)$$

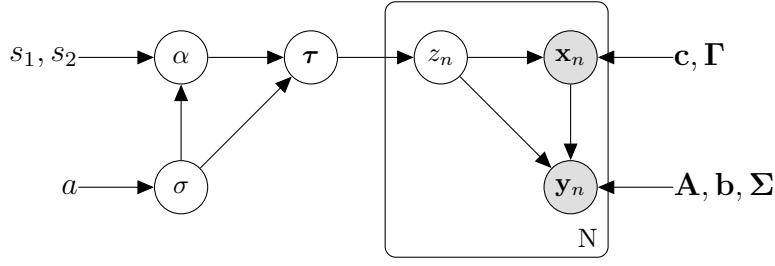


Figure 1: Graphical representation of BNP-GLLiM1: the plate denotes N i.i.d. observations, white filled circles correspond to unobserved (latent) variables while grey filled circles correspond to observed variables.

In practice, we can decide which parameters to treat as genuine parameters Θ or as hyperparameters ϕ , depending on whether some prior knowledge is available for only a subset of the parameters, or whether the model has hyperparameters ϕ for which no prior information is available. Furthermore, for complex models, q_{Θ} and $q_{\mathcal{Z}}$ may need to be further restricted to simpler forms, such as factorised forms, to ensure tractable VB-E steps. This is illustrated in the next [Section 3.3](#) for the BNP-GLLiM inference.

3.2 VBEM for a simple BNP-GLLiM model

In a simple BNP-GLLiM model, or just BNP-GLLiM1 for short, we focus only on the tractable BNP-GLLiM model without any commitment to an arbitrary number of experts. That is why we decide to consider as hyperparameters the set $\phi = (s_1, s_2, a, \theta^*)$ while $\Theta = (\tau, \alpha, \sigma)$ are specific to the BNP part, namely stick-breaking construction of PYP. Recall that $\theta^* = (\mathbf{c}, \mathbf{\Gamma}, \mathbf{A}, \mathbf{b}, \mathbf{\Sigma}) \equiv (\mathbf{c}_k, \mathbf{\Gamma}_k, \mathbf{A}_k, \mathbf{b}_k, \mathbf{\Sigma}_k)_{k \in \mathbb{N}^*}$.

In [Section 2.3.2](#), the base distribution is given in this case as

$$G_0 = \sum_{k=1}^{\infty} \pi_k^0(\tau) \delta_{\theta_k^0} \equiv \sum_{k=1}^{\infty} \pi_k^0(\tau) \delta_{(\mathbf{c}_k^0, \mathbf{\Gamma}_k^0, \mathbf{A}_k^0, \mathbf{b}_k^0, \mathbf{\Sigma}_k^0)}. \quad (3.8)$$

As in [\[68, 69\]](#), BNP-GLLiM1 can be viewed as a hierarchical PYP. See [\[70\]](#) for a review on hierarchical nonparametric priors. For better illustration, we present the following graphical representation for BNP-GLLiM1 in [Figure 1](#).

It then follows that the joint distribution of the observed data \mathcal{X}, \mathcal{Y} and of all the latent variables can be expressed in a hierarchical manner as follows

$$\begin{aligned} p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \phi) &= \prod_{n=1}^N p(\mathbf{y}_n | \mathbf{x}_n, z_n, \Theta; \phi) p(\mathbf{x}_n | z_n, \Theta; \phi) p(\mathcal{Z} | \Theta; \phi) p(\Theta; \phi) \\ &= \prod_{n=1}^N p(\mathbf{y}_n | \mathbf{x}_n, z_n; \mathbf{A}, \mathbf{b}, \mathbf{\Sigma}) p(\mathbf{x}_n | z_n, \mathbf{c}, \mathbf{\Gamma}) p(\mathcal{Z} | \tau; \beta) \\ &\quad \prod_{k \in \mathbb{N}^*} p(\tau_k | \alpha, \sigma) p(\alpha, \sigma | s_1, s_2, a). \end{aligned} \quad (3.9)$$

Here the notation $\prod_{k \in \mathbb{N}^*}$ is a distributional notation.

In most VA, the posterior for the stick-breaking variables is approximated in a factorized form (mean-field approximation). Following the same approach, by factorizing between the latent variables and the parameters, we obtain the following variational distribution

$$q(\mathcal{Z}, \Theta; \phi) = q_{\mathcal{Z}}(\mathcal{Z}; \phi) q_{\Theta}(\Theta; \phi) \equiv q_{\mathcal{Z}}(\mathcal{Z}) q_{\Theta}(\Theta).$$

In particular, the intractable posterior on \mathcal{Z} is approximated as $q_{\mathcal{Z}}(\mathcal{Z})$ that factorizes so as to handle intractability, namely

$$q_{\mathcal{Z}}(\mathcal{Z}) = \prod_{n=1}^N q_{z_n}(z_n). \quad (3.10)$$

Then the infinite state space for each z_j is dealt with by choosing a truncation of the state space to a maximum label $K \in \mathbb{N}^*$, see *e.g.*, [60, 71]. In practice, this consists of assuming that the variational distributions q_{z_n} for $n \in [N]$, satisfy $q_{z_n}(k) = 0$ for $k > K$ and that the variational distribution on τ also factorizes as $q_{\tau}(\tau) = \prod_{k=1}^{K-1} q_{\tau_k}(\tau_k)$ with the additional condition that $\tau_K = 1$. Thus the truncated variational posterior of parameters Θ is given by

$$q_{\Theta}(\Theta) = q_{\alpha, \sigma}(\alpha, \sigma) \prod_{k=1}^{K-1} q_{\tau_k}(\tau_k). \quad (3.11)$$

These forms of $q_{\mathcal{Z}}$ and q_{Θ} lead to our three VB-E steps and four VB-M steps, summarized below and with more detail in [Appendix A](#). Set the initial value of ϕ to $\phi^{(0)}$. Then repeat the following steps iteratively. The iteration index is omitted in the update formulas for simplicity.

VB-E steps

3.2.1 VB-E- τ step

The VB-E- τ step corresponds to a VA in the exponential family case and results in a posterior from the same family as the prior, see *e.g.*, [67, Section 10.4].

More precisely, to achieve this, we use (2.18), (2.23), (3.26), and are only interested in the functional dependence of the right-hand side of (3.6) on the variable τ_k . Thus, any terms that do not depend on τ_k can be included in the additive normalization constant. Then, given for $k \in [K-1]$ that $N_k = \sum_{n=1}^N q_{z_n}(k)$ corresponds to the weight of the cluster k , see more details in [Appendix A.1](#), it holds that

$$q_{\tau_k}(\tau_k) = \text{Beta}(\tau_k \mid \hat{\gamma}_{k,1}, \hat{\gamma}_{k,2}), \text{ where,} \quad (3.12)$$

$$\hat{\gamma}_{k,1} = 1 - \mathbb{E}_{q_{\alpha, \sigma}}[\sigma] + N_k, \quad \hat{\gamma}_{k,2} = \mathbb{E}_{q_{\alpha, \sigma}}[\alpha] + k \mathbb{E}_{q_{\alpha, \sigma}}[\sigma] + \sum_{l=k+1}^K N_l. \quad (3.13)$$

3.2.2 VB-E- (α, σ) step

The (α, σ) variational posterior is more complex, but has a simple form in the DP case ($\sigma = 0$). Specifically, we have to compute

$$\hat{s}_1 = s_1 + K - 1, \quad \hat{s}_2 = s_2 - \sum_{k=1}^{K-1} \psi(\hat{\gamma}_{k,2}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}), \quad (3.14)$$

where $\psi(\cdot)$ is the digamma function defined by $\psi(\gamma) = \frac{d}{d\gamma} \log \Gamma(\gamma) = \frac{\Gamma'(\gamma)}{\Gamma(\gamma)}$.

When $\sigma = 0$ then $q_{\alpha, \sigma} \equiv q_{\alpha, 0}$ is a gamma distribution $\text{Gam}(\hat{s}_1, \hat{s}_2)$ with $\mathbb{E}_{q_{\alpha, \sigma}}[\alpha] = \frac{\hat{s}_1}{\hat{s}_2}$. Otherwise (PYP case), $q_{\alpha, \sigma}$ is only identified up to a normalizing constant but the required $\mathbb{E}_{q_{\alpha, \sigma}}[\alpha]$ and $\mathbb{E}_{q_{\alpha, \sigma}}[\sigma]$ can be computed by importance sampling, see [Appendix A.2](#) for more details.

We next consider the derivation of the update equation for the factor $q_{\mathcal{Z}}(\mathcal{Z})$.

3.2.3 VB-E- \mathcal{Z} step

By using the mean-field approximation (3.10) and the truncation, for more details refer to [Appendix A.3](#), for all $n \in [N]$ and for all $k \in [K]$, this step consists in computing,

$$q_{z_n}(k) = \frac{\rho_{nk}}{\sum_{l=1}^K \rho_{nl}}. \quad (3.15)$$

Here, we define $\log \rho_{nk}$ by

$$\begin{aligned} & -\frac{1}{2} \left\{ D \log(2\pi) + \log \left| \widehat{\Sigma}_k \right| + (\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k)^\top \widehat{\Sigma}_k^{-1} (\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k) \right. \\ & \quad \left. + L \log(2\pi) + \log \left| \widehat{\Gamma}_k \right| + (\mathbf{x}_n - \widehat{\mathbf{c}}_k)^\top \widehat{\Gamma}_k^{-1} (\mathbf{x}_n - \widehat{\mathbf{c}}_k) \right\} \\ & \quad + \psi(\widehat{\gamma}_{k,1}) - \psi(\widehat{\gamma}_{k,1} + \widehat{\gamma}_{k,2}) + \sum_{l=1}^{k-1} [\psi(\widehat{\gamma}_{l,2}) - \psi(\widehat{\gamma}_{l,1} + \widehat{\gamma}_{l,2})]. \end{aligned} \quad (3.16)$$

Note that in the above formula, symbols $(\widehat{\mathbf{c}}_k, \widehat{\Gamma}_k, \widehat{\mathbf{A}}_k, \widehat{\mathbf{b}}_k, \widehat{\Sigma}_k)$ are the hyperparameters more specifically defined in the following [Sections 3.2.5](#) and [3.2.6](#).

It is important to note that (3.15) provides assignment probabilities $q_{z_n}(k)$ rather than intermediate commitments to hard assignments of z_n . However, the hard assignments can be postponed to the end if desired to obtain a segmentation by the following MAP estimation:

$$\widehat{z}_n = \operatorname{argmax}_{k \in [K]} q_{z_n}(k). \quad (3.17)$$

VB-M steps

The maximization step consists of updating the hyperparameters $\phi = (s_1, s_2, a, \boldsymbol{\theta}_{[K]}^*)$, where $\boldsymbol{\theta}_{[K]}^* = (\mathbf{c}_k, \Gamma_k, \mathbf{A}_k, \mathbf{b}_k, \Sigma_k)_{k \in [K]}$, by maximizing the free energy, if they are not set heuristically:

$$\phi^{(r)} = \operatorname{argmax}_{\phi} \mathbb{E}_{q_{\mathcal{Z}}^{(r)} q_{\boldsymbol{\tau}}^{(r)} q_{\alpha, \sigma}^{(r)}} [\log p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \boldsymbol{\tau}, \alpha, \sigma; \phi)]. \quad (3.18)$$

The VB-M-step can therefore be divided into 4 independent sub-steps as listed below. From the conditional independence of (s_1, s_2, a) and $(\mathcal{Y}, \mathcal{X}, \mathcal{Z})$ given $(\boldsymbol{\tau}, \alpha, \sigma)$, the solution for the VB-M- (s_1, s_2) (in the DP case) step is straightforward. Only the M- β , M- (s_1, s_2, a) (in the PYP case) and M- $\boldsymbol{\theta}_{[K]}^*$ steps are more involved.

3.2.4 VB-M- (s_1, s_2, a) step

This step is straightforward in the DP case ($\sigma = 0$). It can be expressed easily using the fact that both the prior and the variational posterior are Gamma distributions, and using the cross-entropy properties,

$$(s_1, s_2)^{(r)} = \operatorname{argmax}_{(s_1, s_2)} \mathbb{E}_{q_{\alpha, 0}^{(r)}} [\log p(\alpha \mid s_1, s_2)] = (\widehat{s}_1^{(r)}, \widehat{s}_2^{(r)}), \quad (3.19)$$

where $(\widehat{s}_1^{(r)}, \widehat{s}_2^{(r)})$ is given in (3.14). We can also solve this step numerically using importance sampling in the more general case of PYP. For more details, see [48, Appendix A.7].

3.2.5 VB-M- (\mathbf{c}, Γ) step

This step divides into K sub-steps that involve the following optimizations

$$\left(\widehat{\mathbf{c}}_k, \widehat{\Gamma}_k \right) \equiv \left(\widehat{\mathbf{c}}_k^{(r)}, \widehat{\Gamma}_k^{(r)} \right) = \operatorname{argmax}_{(\mathbf{c}_k, \Gamma_k)} \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [\log p(\mathcal{X} \mid \mathcal{Z}; \mathbf{c}_k, \Gamma_k)]. \quad (3.20)$$

We can then update the Gaussian gating parameters as follows:

$$\widehat{\mathbf{c}}_k = \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n, \quad (3.21)$$

$$\widehat{\Gamma}_k = \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) (\mathbf{x}_n - \widehat{\mathbf{c}}_k) (\mathbf{x}_n - \widehat{\mathbf{c}}_k)^\top. \quad (3.22)$$

The technical details will be left to the [Appendix A.4](#).

3.2.6 VB-M-($\mathbf{A}, \mathbf{b}, \Sigma$) step

Using the same idea, this step is divided into K sub-steps, which include the following optimization problems

$$\left(\widehat{\mathbf{A}}_k, \widehat{\mathbf{b}}_k, \widehat{\Sigma}_k\right) \equiv \left(\widehat{\mathbf{A}}_k^{(r)}, \widehat{\mathbf{b}}_k^{(r)}, \widehat{\Sigma}_k^{(r)}\right) = \underset{(\mathbf{A}_k, \mathbf{b}_k, \Sigma_k)}{\operatorname{argmax}} \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}; \mathbf{A}_k, \mathbf{b}_k, \Sigma_k)]. \quad (3.23)$$

It requires to define the following quantities,

$$\begin{aligned} \bar{\mathbf{x}}_k &= \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n, \\ \bar{\mathbf{y}}_k &= \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n, \\ \mathbf{X}_k &= \frac{1}{\sqrt{N_k}} \left(\sqrt{q_{z_1}(k)}(\mathbf{x}_1 - \bar{\mathbf{x}}_k), \dots, \sqrt{q_{z_N}(k)}(\mathbf{x}_N - \bar{\mathbf{x}}_k) \right), \\ \mathbf{Y}_k &= \frac{1}{\sqrt{N_k}} \left(\sqrt{q_{z_1}(k)}(\mathbf{y}_1 - \bar{\mathbf{y}}_k), \dots, \sqrt{q_{z_N}(k)}(\mathbf{y}_N - \bar{\mathbf{y}}_k) \right). \end{aligned}$$

Then we can update the parameters for the Gaussian experts, see [Appendix A.5](#), for more details, as follows:

$$\begin{aligned} \widehat{\mathbf{A}}_k &= \mathbf{Y}_k \mathbf{X}_k^\top (\mathbf{X}_k \mathbf{X}_k^\top)^{-1}, \\ \widehat{\mathbf{b}}_k &= \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n), \\ \widehat{\Sigma}_k &= \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k)(\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k)^\top. \end{aligned}$$

3.3 VBEM for a general and tractable BNP-GLLiM model

To obtain a general and tractable BNP-GLLiM model, or simply BNP-GLLiM2 for simplicity, for variational inverse conditional estimation, we specify the prior on the gating parameters (\mathbf{c}_k, Γ_k) as a normal-inverse-Wishart (NIW) distribution parameterized by $\boldsymbol{\rho}_k = (\mathbf{m}_k, \lambda_k, \Psi_k, \nu_k)$ with a PDF

$$p(\mathbf{c}_k, \Gamma_k \mid \boldsymbol{\rho}_k) \equiv \mathcal{NIW}(\mathbf{c}_k, \Gamma_k \mid \boldsymbol{\rho}_k) = \mathcal{N}(\mathbf{c}_k \mid \mathbf{m}_k, \lambda_k^{-1} \Gamma_k) \mathcal{IW}(\Gamma_k \mid \Psi_k, \nu_k). \quad (3.24)$$

Then we decide to take into account as hyperparameters and parameters, respectively, as follows

$$\boldsymbol{\phi} = (s_1, s_2, a, \beta, (\boldsymbol{\rho}_k, \mathbf{A}_k, \mathbf{b}_k, \Sigma_k)_{k \in \mathbb{N}^*}), \text{ while } \boldsymbol{\Theta} = (\boldsymbol{\tau}, \alpha, \sigma, \boldsymbol{\theta}^*), \boldsymbol{\theta}^* = (\boldsymbol{\theta}_k^*)_{k \in \mathbb{N}^*} \equiv (\mathbf{c}_k, \Gamma_k)_{k \in \mathbb{N}^*}.$$

In [Section 2.3.2](#), the base distribution is given in this case as

$$G_0 = \sum_{k=1}^{\infty} \pi_k^0(\boldsymbol{\tau}) \left(\delta_{(\mathbf{A}_k^0, \mathbf{b}_k^0, \Sigma_k^0)}, \mathcal{NIW}(\cdot, \cdot \mid \boldsymbol{\rho}_k^0) \right). \quad (3.25)$$

As in [\[68, 69\]](#), BNP-GLLiM2 can also be viewed as a hierarchical PYP. For better illustration, we present the following graphical representation for BNP-GLLiM2 in [Figure 2](#).

In BNP-GLLiM2, it follows that the joint distribution of the observed data \mathcal{X}, \mathcal{Y} and all latent variable can be expressed hierarchically as

$$\begin{aligned} p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \boldsymbol{\Theta}; \boldsymbol{\phi}) &= \prod_{n=1}^N p(\mathbf{y}_n \mid \mathbf{x}_n, z_n, \boldsymbol{\Theta}; \boldsymbol{\phi}) p(\mathbf{x}_n \mid z_n, \boldsymbol{\Theta}; \boldsymbol{\phi}) p(\mathcal{Z} \mid \boldsymbol{\Theta}; \boldsymbol{\phi}) p(\boldsymbol{\Theta}; \boldsymbol{\phi}) \\ &= \prod_{n=1}^N p(\mathbf{y}_n \mid \mathbf{x}_n, z_n; \mathbf{A}, \mathbf{b}, \Sigma) p(\mathbf{x}_n \mid z_n, \mathbf{c}, \Gamma) p(\mathcal{Z} \mid \boldsymbol{\tau}) \\ &\quad \prod_{k \in \mathbb{N}^*} p(\boldsymbol{\tau}_k \mid \alpha, \sigma) p(\alpha, \sigma \mid s_1, s_2, a) \prod_{k \in \mathbb{N}^*} p(\mathbf{c}_k, \Gamma_k; \boldsymbol{\rho}_k). \end{aligned} \quad (3.26)$$

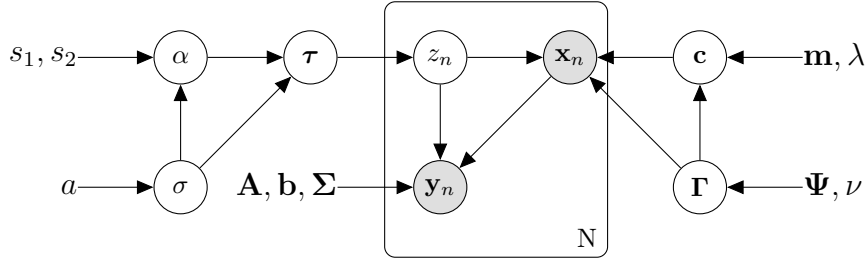


Figure 2: Graphical representation of BNP-GLLiM2: the plate denotes N i.i.d. observations, white filled circles correspond to unobserved (latent) variables while grey filled circles correspond to observed variables.

Following the same idea as in [Section 3.2](#), we only consider the truncated variational posterior of parameters Θ as follows

$$q_{\Theta}(\Theta) = q_{\alpha, \sigma}(\alpha, \sigma) \prod_{k=1}^{K-1} q_{\tau_k}(\tau_k) \prod_{k=1}^K q_{\theta_k^*}(\theta_k^*). \quad (3.27)$$

These forms of $q_{\mathcal{Z}}$ and q_{Θ} lead to our four VB-E steps and three VB-M steps, summarized below with details in [Appendix A](#). Set the initial value of ϕ to $\phi^{(0)}$. Then, repeat iteratively the following steps. The iteration index is omitted in the update formulas for simplicity.

VB-E steps

Note that the VB-E- τ , VB-E- (α, σ) steps are the same as in [Section 3.2](#). We only highlight the modified steps as follows.

We first consider the derivation of the update equation for the factor $q_{\mathcal{Z}}(\mathcal{Z})$.

3.3.1 VB-E- \mathcal{Z} step

By using the mean-field approximation [\(3.10\)](#) and the truncation, see [Appendix A.3](#) for more details, for all $n \in [N]$ and for $k \in [K]$, this step consists in computing,

$$q_{z_n}(k) = \frac{\rho_{nk}}{\sum_{l=1}^K \rho_{nl}}. \quad (3.28)$$

Here, given \mathcal{N}_n represents the neighbors of n , we define $\log \rho_{nk}$ by

$$\begin{aligned} & -\frac{1}{2} \left\{ \log |\widehat{\Sigma}_k| + (\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k)^\top \widehat{\Sigma}_k^{-1} (\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k) \right. \\ & \left. + \log \left| \frac{\widehat{\Psi}_k}{2} \right| - \sum_{l=1}^L \psi \left(\frac{\widehat{\nu}_k + (1-l)}{2} \right) + \widehat{\nu}_k (\mathbf{x}_n - \widehat{\mathbf{m}}_k)^\top \widehat{\Psi}_k^{-1} (\mathbf{x}_n - \widehat{\mathbf{m}}_k) + \frac{L}{\widehat{\lambda}_k} \right\} \\ & + \psi(\widehat{\gamma}_{k,1}) - \psi(\widehat{\gamma}_{k,1} + \widehat{\gamma}_{k,2}) + \sum_{l=1}^{k-1} [\psi(\widehat{\gamma}_{l,2}) - \psi(\widehat{\gamma}_{l,1} + \widehat{\gamma}_{l,2})]. \end{aligned}$$

Note that in the above formula, symbols $(\widehat{\mathbf{m}}_k, \widehat{\lambda}_k, \widehat{\Psi}_k, \widehat{\nu}_k)$ and $(\widehat{\mathbf{A}}_k, \widehat{\mathbf{b}}_k, \widehat{\Sigma}_k)$ are the hyperparameters more specifically defined in the following [Sections 3.2.6](#) and [3.3.2](#).

3.3.2 VB-E- θ^* step

This step is divided into K parts where the computation is similar to that in standard Bayesian GMM with a choice of conjugate prior. Hence, for each $k \leq K$, the variational posterior is a Normal-inverse-Wishart distribution defined as

$$q_{\theta_k^*}(\mathbf{c}_k, \Gamma_k) = \mathcal{NIW}(\mathbf{c}_k, \Gamma_k \mid \widehat{\mathbf{m}}_k, \widehat{\lambda}_k, \widehat{\Psi}_k, \widehat{\nu}_k). \quad (3.29)$$

Here, the hyperparameters are updated as follows (see, *e.g.*, [67, Section 10.2.1]):

$$\widehat{\lambda}_k = \lambda_k + N_k, \quad \widehat{\nu}_k = \nu_k + N_k, \quad N_k = \sum_{n=1}^N q_{z_n}(k) \quad (3.30)$$

$$\widehat{\Psi}_k = \Psi_k + N_k \mathbf{S}_k + \frac{\lambda_k N_k}{\lambda_k + N_k} (\mathbf{m}_k - \bar{\mathbf{c}}_k)(\mathbf{m}_k - \bar{\mathbf{c}}_k)^\top, \quad (3.31)$$

$$\widehat{\mathbf{m}}_k = \frac{\lambda_k \mathbf{m}_k + N_k \bar{\mathbf{c}}_k}{\lambda_k + N_k} = \frac{\lambda_k \mathbf{m}_k + N_k \bar{\mathbf{c}}_k}{\widehat{\lambda}_k}, \quad (3.32)$$

$$\bar{\mathbf{c}}_k = \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n, \quad (3.33)$$

$$\mathbf{S}_k = \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) (\mathbf{x}_n - \bar{\mathbf{c}}_k)(\mathbf{x}_n - \bar{\mathbf{c}}_k)^\top. \quad (3.34)$$

VB-M steps

The maximization step consists of updating the hyperparameters $\phi = (s_1, s_2, a, (\boldsymbol{\rho}_k, \mathbf{A}_k, \mathbf{b}_k, \boldsymbol{\Sigma}_k)_{k \in [K]})$, where $\boldsymbol{\rho}_k = (\mathbf{m}_k, \lambda_k, \Psi_k, \nu_k)$, $k \in [K]$, by maximizing the free energy, if they are not set heuristically:

$$\phi^{(r)} = \operatorname{argmax}_{\phi} \mathbb{E}_{q_{\mathcal{Z}}^{(r)} q_{\boldsymbol{\tau}}^{(r)} q_{\alpha, \sigma}^{(r)} q_{\boldsymbol{\theta}^*}^{(r)}} [\log p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \boldsymbol{\tau}, \alpha, \sigma, \boldsymbol{\theta}^*; \phi)]. \quad (3.35)$$

The VB-M-step can therefore be divided into four independent sub-steps as listed below. From the conditional independence of $(s_1, s_2, a, \boldsymbol{\rho})$ and $(\mathcal{Y}, \mathcal{X}, \mathcal{Z})$ given $(\boldsymbol{\tau}, \alpha, \sigma, \boldsymbol{\theta}^*)$, the solutions for the VB-M- (s_1, s_2) (in the DP case) and VB-M- $\boldsymbol{\rho}$ steps are straightforward. Only the M- (s_1, s_2, a) step (in the PYP case) and $(\mathbf{A}_k, \mathbf{b}_k, \boldsymbol{\Sigma}_k)_{k \in [K]}$ are more involved.

Note that the VB-M- (s_1, s_2, a) , VB-M- $(\mathbf{A}, \mathbf{b}, \boldsymbol{\Sigma})$ steps are the same as in Section 3.2. We only highlight the modified step below.

3.3.3 VB-M- $\boldsymbol{\rho}$ step

This step divides into K sub-steps that involve again cross-entropies,

$$\rho_k^{(r)} = \operatorname{argmax}_{\rho} \mathbb{E}_{q_{\boldsymbol{\theta}^*}^{(r)}} [\log p(\mathbf{c}_k, \boldsymbol{\Gamma}_k; \rho_k)] = \widehat{\rho}_k^{(r)} \quad (3.36)$$

where $\widehat{\rho}_k^{(r)} = (\widehat{\lambda}_k^{(r)}, \widehat{\nu}_k^{(r)}, \widehat{\Psi}_k^{(r)}, \widehat{\mathbf{m}}_k^{(r)})$ is given in (3.30).

4 Predictive conditional density

The most popular uses of BNP-GLLiM with discrete random probability measures, such as the one displayed in (2.25), relate to conditional density estimation and data clustering. More precisely, we will be interested in the predictive conditional density for a new value $(\widehat{\mathbf{y}}, \widehat{\mathbf{x}})$ of the observed variables. Note that associated with these observations will be a corresponding latent variable $\widehat{\mathbf{z}}$.

If $\sigma \neq 0$, we can use the following remark. Suppose there is enough training data. Then the posterior distribution of all parameters is usually highly peaked. This leads to the second approximation, where the integral reduces to a point evaluation at the posterior mean of each parameter, see *e.g.*, [72, 73, 74, 75, 76].

When $\sigma = 0$, we can analytically approximate the such densities via several following lemmas. In the following Theorems 4.1 to 4.6, the notation “ \approx ” means that we approximate the desired densities of the BNP-GLLiM1 by a mixture of Gaussians using factorized VA posteriors and a truncation of the number of components.

4.1 BNP-GLLiM1

4.1.1 Joint density

We first show how to compute the joint density $p(\widehat{\mathbf{y}}, \widehat{\mathbf{x}}, \mathcal{X}, \mathcal{Y})$ via [Theorem 4.1](#), which is proved in [Appendix B.2](#).

Theorem 4.1. *Via defining $\widehat{\mathbf{w}} \equiv [\widehat{\mathbf{x}}; \widehat{\mathbf{y}}]$, we obtain*

$$\boldsymbol{\mu}_k \equiv \mathbb{E}[\widehat{\mathbf{w}}] = \begin{pmatrix} \widehat{\mathbf{c}}_k \\ \widehat{\mathbf{A}}_k \widehat{\mathbf{c}}_k + \widehat{\mathbf{b}}_k \end{pmatrix}, \quad \mathbf{V}_k = \text{cov}[\widehat{\mathbf{w}}] = \begin{pmatrix} \widehat{\boldsymbol{\Gamma}}_k & \widehat{\boldsymbol{\Gamma}}_k \widehat{\mathbf{A}}_k^\top \\ \widehat{\mathbf{A}}_k \widehat{\boldsymbol{\Gamma}}_k & \widehat{\boldsymbol{\Sigma}}_k + \widehat{\mathbf{A}}_k \widehat{\boldsymbol{\Gamma}}_k \widehat{\mathbf{A}}_k^\top \end{pmatrix}, \quad (4.1)$$

and

$$\begin{aligned} p(\widehat{\mathbf{y}}, \widehat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) &\approx \sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{c}}_k, \widehat{\boldsymbol{\Gamma}}_k) \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) \\ &\equiv \sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_{L+D}(\widehat{\mathbf{w}} \mid \boldsymbol{\mu}_k, \mathbf{V}_k). \end{aligned} \quad (4.2)$$

Furthermore, it holds that

$$\mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] = \mathbb{E}_{q_{\tau_k}} [\tau_k] \prod_{l=1}^{k-1} \mathbb{E}_{q_{\tau_l}} [1 - \tau_l]. \quad (4.3)$$

Here, we have for any $k \in [K]$,

$$\begin{aligned} \mathbb{E}_{q_{\tau_k}} [\tau_k] &= \frac{\widehat{\gamma}_{k,1}}{\widehat{\gamma}_{k,1} + \widehat{\gamma}_{k,2}}, \quad \mathbb{E}_{q_{\tau_k}} [1 - \tau_k] = 1 - \mathbb{E}_{q_{\tau_k}} [\tau_k] = \frac{\widehat{\gamma}_{k,2}}{\widehat{\gamma}_{k,1} + \widehat{\gamma}_{k,2}}, \\ \widehat{\gamma}_{k,1} &= 1 - \mathbb{E}_{q_{\alpha,\sigma}} [\sigma] + N_k, \quad \widehat{\gamma}_{k,2} = \mathbb{E}_{q_{\alpha,\sigma}} [\alpha] + k \mathbb{E}_{q_{\alpha,\sigma}} [\sigma] + \sum_{l=k+1}^K N_l, \quad N_k = \sum_{n=1}^N q_{z_n}(k). \end{aligned}$$

4.1.2 Inverse conditional density

We then show how to approximate the inverse conditional density $p(\widehat{\mathbf{y}} \mid \widehat{\mathbf{x}}, \mathcal{X}, \mathcal{Y})$. This predictive density in BNP-GLLiM1 is approximated by a GLLiM via [Theorem 4.2](#) with the proof in [Appendix B.3](#).

Theorem 4.2. *We approximate the inverse conditional density of the BNP-GLLiM1 as follows:*

$$p(\widehat{\mathbf{y}} \mid \widehat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) \approx \sum_{k=1}^K g_k(\widehat{\mathbf{x}} \mid \widehat{\boldsymbol{\Theta}}, \widehat{\boldsymbol{\phi}}, \mathcal{X}, \mathcal{Y}) \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k). \quad (4.4)$$

Here, the gating posteriors belong to a $K - 1$ dimensional probability simplex and are given by

$$g_k(\widehat{\mathbf{x}} \mid \widehat{\boldsymbol{\Theta}}, \widehat{\boldsymbol{\phi}}, \mathcal{X}, \mathcal{Y}) = \frac{\mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{c}}_k, \widehat{\boldsymbol{\Gamma}}_k)}{\sum_{l=1}^K \mathbb{E}_{q_\tau} [\pi_l(\boldsymbol{\tau})] \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{c}}_l, \widehat{\boldsymbol{\Gamma}}_l)}, \quad k \in [K], \quad (4.5)$$

For prediction, we also have

$$\mathbb{E}[\widehat{\mathbf{y}} \mid \widehat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}] \approx \sum_{k=1}^K g_k(\widehat{\mathbf{x}} \mid \widehat{\boldsymbol{\Theta}}, \widehat{\boldsymbol{\phi}}, \mathcal{X}, \mathcal{Y}) [\widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k]. \quad (4.6)$$

4.1.3 Forward conditional density

To deal with high-dimensional regression data, namely high-to-low regression, given the inverse conditional density $p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y})$, we want to approximate the following forward conditional density via [Theorem 4.3](#), whose proof is provided in [Appendix B.4](#).

Theorem 4.3. *We approximate the forward conditional density of the BNP-GLLiM1 by*

$$p(\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) \approx \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*, \hat{\Sigma}_k^*), \quad (4.7)$$

which is again a GLLiM. Here, we defined

$$\begin{aligned} g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) &= \frac{\mathbb{E}_{q_\tau}[\pi_k(\tau)] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_k^*, \hat{\Gamma}_k^*)}{\sum_{l=1}^K \mathbb{E}_{q_\tau}[\pi_l(\tau)] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_l^*, \hat{\Gamma}_l^*)}, \\ \hat{\Sigma}_k^* &= \left(\hat{\Gamma}_k^{-1} + \hat{\mathbf{A}}_k^\top \hat{\Sigma}_k^{-1} \hat{\mathbf{A}}_k \right)^{-1}, \\ \hat{\mathbf{A}}_k^* &= \hat{\Sigma}_k^* \hat{\mathbf{A}}_k^\top \hat{\Sigma}_k^{-1}, \\ \hat{\mathbf{b}}_k^* &= \hat{\Sigma}_k^* \left[\hat{\Gamma}_k^{-1} \hat{\mathbf{c}}_k - \hat{\mathbf{A}}_k^\top \hat{\Sigma}_k^{-1} \hat{\mathbf{b}}_k \right], \\ \hat{\mathbf{c}}_k^* &= \hat{\mathbf{A}}_k \hat{\mathbf{c}}_k + \hat{\mathbf{b}}_k, \\ \hat{\Gamma}_k^* &= \hat{\Sigma}_k + \hat{\mathbf{A}}_k \hat{\Gamma}_k \hat{\mathbf{A}}_k^\top. \end{aligned} \quad (4.8)$$

Furthermore, we approximate the expectation and covariance matrix of $\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}$ as follows:

$$\begin{aligned} \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] &\approx \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \left(\hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^* \right), \\ \text{var}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] &\approx \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \left[\hat{\Sigma}_k^* + \left(\hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^* \right) \left(\hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^* \right)^\top \right] \\ &\quad - \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}]^\top. \end{aligned} \quad (4.9)$$

4.2 BNP-GLLiM2

4.2.1 Joint density

We first show how to compute the joint density $p(\hat{\mathbf{y}}, \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y})$ via [Theorem 4.4](#), which is proved in [Appendix B.5](#).

Theorem 4.4. *We approximate the joint density of BNP-GLLiM2 by a mixture of product between Gaussian and Student's t -distributions as follows:*

$$p(\hat{\mathbf{y}}, \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) \approx \sum_{k=1}^K \mathbb{E}_{q_\tau}[\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} | \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\Sigma}_k). \quad (4.10)$$

Here, the positive semidefinite shape matrices of Student's t -distributions are given by

$$\mathbf{L}_k = \frac{(\hat{\nu}_k + 1 - L) \hat{\lambda}_k}{1 + \hat{\lambda}_k} \hat{\Psi}_k. \quad (4.11)$$

4.2.2 Inverse conditional density

We then show how to approximate the inverse conditional density $p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y})$. This predictive density in BNP-GLLiM2 is approximated by a MoE via [Theorem 4.5](#) with the proof in [Appendix B.6](#).

Theorem 4.5. *We approximate the inverse conditional density of BNP-GLLiM2 by a MoE as follows:*

$$p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) \approx \sum_{k=1}^K g_k(\hat{\mathbf{x}} | \hat{\Theta}, \hat{\phi}, \mathcal{X}, \mathcal{Y}) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\Sigma}_k). \quad (4.12)$$

Here, the gating posteriors are defined as

$$g_k(\hat{\mathbf{x}} | \hat{\Theta}, \hat{\phi}, \mathcal{X}, \mathcal{Y}) = \frac{\mathbb{E}_{q_\tau}[\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} | \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L)}{\sum_{l=1}^K \mathbb{E}_{q_\tau}[\pi_l(\tau)] \text{St}(\hat{\mathbf{x}} | \hat{\mathbf{m}}_l, \mathbf{L}_l, \hat{\nu}_l + 1 - L)}, \quad k \in [K]. \quad (4.13)$$

Furthermore, for any $k \in [K]$, it holds that

$$\begin{aligned} \mathbb{E}_{q_\tau}[\pi_k(\tau)] &= \mathbb{E}_{q_{\tau_k}}[\tau_k] \prod_{l=1}^{k-1} \mathbb{E}_{q_{\tau_l}}[1 - \tau_l], \\ \mathbb{E}_{q_{\tau_k}}[\tau_k] &= \frac{\hat{\gamma}_{k,1}}{\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}}, \quad \mathbb{E}_{q_{\tau_k}}[1 - \tau_k] = 1 - \mathbb{E}_{q_{\tau_k}}[\tau_k] = \frac{\hat{\gamma}_{k,2}}{\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}}, \\ \hat{\gamma}_{k,1} &= 1 - \mathbb{E}_{q_{\alpha,\sigma}}[\sigma] + N_k, \quad \hat{\gamma}_{k,2} = \mathbb{E}_{q_{\alpha,\sigma}}[\alpha] + k \mathbb{E}_{q_{\alpha,\sigma}}[\sigma] + \sum_{l=k+1}^K N_l, \quad N_k = \sum_{n=1}^n q_{z_n}(k). \end{aligned}$$

The prediction task is carried out via the following approximation

$$\mathbb{E}[\hat{\mathbf{y}} | \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}] \approx \sum_{k=1}^K g_k(\hat{\mathbf{x}} | \hat{\Theta}, \hat{\phi}, \mathcal{X}, \mathcal{Y}) [\hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k]. \quad (4.14)$$

4.2.3 Forward conditional density

To deal with high-dimensional regression data, namely high-to-low regression, given the inverse conditional density $p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y})$, we want to approximate the following forward conditional density via [Theorem 4.6](#), whose proof is provided in [Appendix B.7](#).

Theorem 4.6. *It holds that*

$$p(\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) \approx \sum_{i=1}^I \sum_{k=1}^K g_{ki}(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^*(\eta_i) \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*(\eta_i), \hat{\Sigma}_k^*(\eta_i)),$$

which is a mixture of Gaussian experts, where, for all $k \in [K], i \in [I]$,

$$\begin{aligned} g_{ki}(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) &= \frac{\mathbb{E}_{q_\tau}[\pi_k(\tau)] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_k^*, \hat{\Gamma}_k^*(\eta_i)) \text{Gam}(\eta_i | \frac{\hat{\nu}_k+1-L}{2}, \frac{\hat{\nu}_k+1-L}{2})}{\sum_{i=1}^I \sum_{l=1}^K \mathbb{E}_{q_\tau}[\pi_l(\tau)] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_l^*, \hat{\Gamma}_l^*(\eta_i)) \text{Gam}(\eta_i | \frac{\hat{\nu}_l+1-L}{2}, \frac{\hat{\nu}_l+1-L}{2})}, \\ \hat{\Sigma}_k^*(\eta_i) &= (\eta_i \mathbf{L}_k + \hat{\mathbf{A}}_k^\top \hat{\Sigma}_k^{-1} \hat{\mathbf{A}}_k)^{-1}, \\ \hat{\mathbf{A}}_k^*(\eta_i) &= \hat{\Sigma}_k^*(\eta_i) \hat{\mathbf{A}}_k^\top \hat{\Sigma}_k^{-1}, \\ \hat{\mathbf{b}}_k^*(\eta_i) &= \hat{\Sigma}_k^*(\eta_i) [\eta_i \mathbf{L}_k \hat{\mathbf{m}}_k - \hat{\mathbf{A}}_k^\top \hat{\Sigma}_k^{-1} \hat{\mathbf{b}}_k], \\ \hat{\mathbf{c}}_k^* &= \hat{\mathbf{A}}_k \hat{\mathbf{m}}_k + \hat{\mathbf{b}}_k, \\ \hat{\Gamma}_k^*(\eta_i) &= \hat{\Sigma}_k + \hat{\mathbf{A}}_k (\eta_i \mathbf{L}_k)^{-1} \hat{\mathbf{A}}_k^\top. \end{aligned} \quad (4.15)$$

Here, $\eta_i, i \in [I]$, are chosen via discretizing η -space, $[0, U_\eta]$, into a grid, e.g., uniform. Note that for simplicity, we evaluate the integrand as a Riemann integral with a truncated value $0 < U_\eta < \infty$ and a number of point $I \in \mathbb{N}^*$ for approximating the integration but we can use any scheme to approximate such 1-dimensional integration.

5 ELBO

Evaluating the ELBO in (3.4) allows us not only to monitor the bound during the re-estimation to test for convergence, but also to check both the mathematical expressions for the solutions and their software implementation. Indeed, the value of this bound (3.4) at each step of the iterative re-estimation procedure should not decrease [77], in particular recent results for Bayesian nonparametric mixture models in [51, Appendix A].

Recall that $\hat{\phi} = \left(\hat{s}_1, \hat{s}_2, \hat{a}, \left(\hat{\rho}_k, \hat{\mathbf{A}}_k, \hat{\mathbf{b}}_k, \hat{\Sigma}_k \right)_{k \in \mathbb{N}^*} \right)$. Here, in order to keep the notation uncluttered, we will sometimes omit the subscripts on the expectation operators because each expectation is taken with respect to all of the random variables in its argument, and the hat superscript $\hat{\cdot}$ on the hyperparameters $\hat{\phi}$ of q distribution.

When $\sigma = 0$, we can analytically compute the ELBO in the BNP-GLLiM via [Theorems 5.1](#) and [5.2](#), which are proved in [Appendices B.8](#) and [B.9](#), respectively.

5.1 BNP-GLLiM1

Theorem 5.1. *When $\sigma = 0$, the ELBO in the BNP-GLLiM1 is analytically derived as follows:*

$$\begin{aligned} \mathcal{F} \left(q_{\mathcal{Z}}, q_{\Theta}, \hat{\phi} \right) &= \mathbb{E} \left[\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\mathcal{X} \mid \mathcal{Z}, \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\mathcal{Z} \mid \Theta; \hat{\phi}) \right] \\ &\quad + \mathbb{E} \left[\log p(\Theta; \hat{\phi}) \right] - \mathbb{E} [\log q(\mathcal{Z})] - \mathbb{E} [\log q(\Theta)]. \end{aligned} \quad (5.1)$$

Here, we have the following update formulas:

$$\mathbb{E} \left[\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi}) \right] = \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \log \mathcal{N}_D \left(\mathbf{y}_n \mid \hat{\mathbf{A}}_k \mathbf{x}_n + \hat{\mathbf{b}}_k, \hat{\Sigma}_k \right) \quad (5.2)$$

$$\mathbb{E} \left[\log p(\mathcal{X} \mid \mathcal{Z}, \Theta; \hat{\phi}) \right] = \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \log \mathcal{N}_L \left(\mathbf{x}_n \mid \hat{\mathbf{c}}_k, \hat{\Gamma}_k \right), \quad (5.3)$$

$$\mathbb{E} \left[\log p(\mathcal{Z} \mid \Theta; \hat{\phi}) \right] = \sum_{k=1}^K N_k \left[\psi(\hat{\gamma}_{k,1}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}) + \sum_{l=1}^{k-1} [\psi(\hat{\gamma}_{l,2}) - \psi(\hat{\gamma}_{l,1} + \hat{\gamma}_{l,2})] \right], \quad (5.4)$$

$$\mathbb{E} \left[\log p(\Theta; \hat{\phi}) \right] = \sum_{k=1}^{K-1} \mathbb{E} [\log p(\tau_k \mid \alpha)] + \mathbb{E} [\log p(\alpha \mid \hat{s}_1, \hat{s}_2)], \quad (5.5)$$

$$\mathbb{E} [\log p(\tau_k \mid \alpha)] = \frac{\hat{s}_1 - \hat{s}_2}{\hat{s}_2} [\psi(\hat{\gamma}_{k,2}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2})] + \psi(\hat{s}_1) - \log \hat{s}_2,$$

$$\mathbb{E} [\log p(\alpha \mid \hat{s}_1, \hat{s}_2)] = -\log \Gamma(\hat{s}_1) + (\hat{s}_1 - 1) \psi(\hat{s}_1) + \log(\hat{s}_2) - \hat{s}_1,$$

$$\mathbb{E} [\log q(\mathcal{Z})] = \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \log q_{z_n}(k), \quad (5.6)$$

$$\mathbb{E} [\log q(\Theta)] = \mathbb{E} [\log q_{\alpha,0}(\alpha)] + \sum_{k=1}^{K-1} \mathbb{E} [\log q_{\tau_k}(\tau_k)], \quad (5.7)$$

$$\mathbb{E} [\log q_{\alpha,0}(\alpha)] = \log \Gamma(\hat{s}_1) + (\hat{s}_1 - 1) \psi(\hat{s}_1) + \log \hat{s}_2 - \hat{s}_1, \quad (5.8)$$

$$\mathbb{E} [\log q_{\tau_k}(\tau_k)] = \sum_{l=1}^2 (\hat{\gamma}_{k,l} - 1) \{ \psi(\hat{\gamma}_{k,l}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}) \} + \log \frac{\Gamma(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2})}{\Gamma(\hat{\gamma}_{k,1}) \Gamma(\hat{\gamma}_{k,2})}. \quad (5.9)$$

Note that if the free energy is computed at the end of each VBEM iteration, as in [Section 3.2.4](#), we have $\mathbb{E} [\log q_{\alpha,0}(\alpha)] = \mathbb{E} [\log p(\alpha \mid \hat{s}_1, \hat{s}_2)]$.

5.2 BNP-GLLiM2

Theorem 5.2. *When $\sigma = 0$, the ELBO in BNP-GLLiM2 is determined analytically as follows:*

$$\begin{aligned} \mathcal{F} \left(q_{\mathcal{Z}}, q_{\Theta}, \hat{\phi} \right) &= \mathbb{E} \left[\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\mathcal{X} \mid \mathcal{Z}, \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\mathcal{Z} \mid \Theta; \hat{\phi}) \right] \\ &\quad + \mathbb{E} \left[\log p(\Theta; \hat{\phi}) \right] - \mathbb{E} \left[\log q(\mathcal{Z}) \right] - \mathbb{E} \left[\log q(\Theta) \right]. \end{aligned} \quad (5.10)$$

Here, we have the following update formulas:

$$\mathbb{E} \left[\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi}) \right] = \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \log \mathcal{N}_D \left(\mathbf{y}_n \mid \hat{\mathbf{A}}_k \mathbf{x}_n + \hat{\mathbf{b}}_k, \hat{\Sigma}_k \right), \quad (5.11)$$

$$\begin{aligned} &\mathbb{E} \left[\log p(\mathcal{X} \mid \mathcal{Z}, \Theta; \hat{\phi}) \right] \\ &= \frac{1}{2} \sum_{k=1}^K N_k \left[\log \tilde{\Gamma}_k - L \log(2\pi) - L \hat{\lambda}_k^{-1} - \hat{\nu}_k \text{Tr} \left(\mathbf{S}_k \hat{\Psi}_k^{-1} \right) - \hat{\nu}_k (\bar{\mathbf{x}}_k - \hat{\mathbf{m}}_k) \hat{\Psi}_k^{-1} (\bar{\mathbf{x}}_k - \hat{\mathbf{m}}_k) \right], \end{aligned} \quad (5.12)$$

$$\mathbb{E} \left[\log p(\mathcal{Z} \mid \Theta; \hat{\phi}) \right] = \sum_{k=1}^K N_k \left[\psi(\hat{\gamma}_{k,1}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}) + \sum_{l=1}^{k-1} [\psi(\hat{\gamma}_{l,2}) - \psi(\hat{\gamma}_{l,1} + \hat{\gamma}_{l,2})] \right], \quad (5.13)$$

$$\mathbb{E} \left[\log p(\Theta; \hat{\phi}) \right] = \sum_{k=1}^{K-1} \mathbb{E} \left[\log p(\tau_k \mid \alpha) \right] + \mathbb{E} \left[\log p(\alpha \mid \hat{s}_1, \hat{s}_2) \right] + \sum_{k=1}^K \mathbb{E} \left[\log p(\mathbf{c}_k, \mathbf{\Gamma}_k; \hat{\rho}_k) \right], \quad (5.14)$$

$$\mathbb{E} \left[\log p(\tau_k \mid \alpha) \right] = \frac{\hat{s}_1 - \hat{s}_2}{\hat{s}_2} [\psi(\hat{\gamma}_{k,2}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2})] + \psi(\hat{s}_1) - \log \hat{s}_2,$$

$$\mathbb{E} \left[\log p(\alpha \mid \hat{s}_1, \hat{s}_2) \right] = -\log \Gamma(\hat{s}_1) + (\hat{s}_1 - 1) \psi(\hat{s}_1) + \log \hat{s}_2 - \hat{s}_1,$$

$$\mathbb{E} \left[\log p(\mathbf{c}_k, \mathbf{\Gamma}_k; \hat{\rho}_k) \right] = \frac{1}{2} L \log \left(\frac{\hat{\lambda}_k}{2\pi} \right) - \frac{L}{2} - \frac{L}{2} \hat{\nu}_k + \log B \left(\hat{\Psi}_k^{-1}, \hat{\nu}_k \right) - \frac{\hat{\nu}_k - L}{2} \log \tilde{\Gamma}_k,$$

$$\log \tilde{\Gamma}_k = \sum_{l=1}^L \psi \left(\frac{\hat{\nu}_k + 1 - l}{2} \right) + L \log 2 + \log \left| \hat{\Psi}_k \right|,$$

$$\mathbb{E} \left[\log q(\mathcal{Z}) \right] = \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \log q_{z_n}(k), \quad (5.15)$$

$$\mathbb{E} \left[\log q(\Theta) \right] = \mathbb{E} \left[\log q_{\alpha,0}(\alpha) \right] + \sum_{k=1}^{K-1} \mathbb{E} \left[\log q_{\tau_k}(\tau_k) \right] + \sum_{k=1}^K \mathbb{E} \left[\log q_{\mathbf{c}_k, \mathbf{\Gamma}_k}(\mathbf{c}_k, \mathbf{\Gamma}_k) \right], \quad (5.16)$$

$$\mathbb{E} \left[\log q_{\alpha,0}(\alpha) \right] = -\log \Gamma(\hat{s}_1) + (\hat{s}_1 - 1) \psi(\hat{s}_1) + \log \hat{s}_2 - \hat{s}_1,$$

$$\mathbb{E} \left[\log q_{\tau_k}(\tau_k) \right] = \sum_{l=1}^2 (\hat{\gamma}_{k,l} - 1) \{ \psi(\hat{\gamma}_{k,l}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}) \} + \log \frac{\Gamma(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2})}{\Gamma(\hat{\gamma}_{k,1}) \Gamma(\hat{\gamma}_{k,2})},$$

$$\mathbb{E} \left[\log q_{\mathbf{c}_k, \mathbf{\Gamma}_k}(\mathbf{c}_k, \mathbf{\Gamma}_k) \right] = \frac{L}{2} \log \frac{\hat{\lambda}_k}{2\pi} - \frac{L}{2} + \log B \left(\hat{\Psi}_k^{-1}, \hat{\nu}_k \right) - \frac{\hat{\nu}_k - L}{2} \log \tilde{\Gamma}_k - \frac{\hat{\nu}_k L}{2}. \quad (5.17)$$

6 Nonparametric Bayesian model selection

6.1 Posterior contraction rate in Bayesian infinite mixtures

6.1.1 Problem setup

We first recall the GMM where we have i.i.d. samples $(\mathbf{W}_n)_{n \in [N]} \equiv \mathcal{W}$ coming from a true but unknown distribution P_{G_0} with given PDF

$$p_{G_0} \equiv \int \mathcal{N}(\mathbf{w} \mid \boldsymbol{\theta}) dG_0(\boldsymbol{\theta}) = \sum_{k=1}^{K_0} \pi_k^0 \mathcal{N}(\mathbf{w} \mid \boldsymbol{\theta}_k^0), \quad \boldsymbol{\theta}_k^0 \equiv (\boldsymbol{\mu}_k^0, \mathbf{V}_k^0), \quad (6.1)$$

where $G_0 = \sum_{k=1}^{K_0} \pi_k^0 \delta_{\theta_k^0}$ is a true but unknown mixing distribution with exactly K_0 number of support points, for some unknown K_0 . Furthermore, Θ is a chosen parameter space, where we empirically believe that the true parameters belong to. In a well-specified setting, all support points of G_0 reside in Θ , but this may not be the case in a misspecified setting. In this section, we assume that the GMM is well-specified, *i.e.*, the data are i.i.d. samples from the mixture density p_{G_0} , where mixing measure G_0 has K_0 support atoms in compact parameter space Θ .

A Bayesian modeler places a prior distribution Π on a suitable subspace of $\bar{\mathcal{G}}(\Theta)$. Then, the posterior distribution over G is given by:

$$\Pi(G \in B \mid \mathcal{W}) \equiv \frac{\int_B \prod_{n=1}^N p_G(\mathbf{W}_n) d\Pi(G)}{\int_{\bar{\mathcal{G}}(\Theta)} \prod_{n=1}^N p_G(\mathbf{W}_n) d\Pi(G)} \quad (6.2)$$

Here, the GMM p_G is defined in (6.1) with $K \leq \infty$ unknown number of support points. We are interested in the posterior contraction behavior of G toward G_0 , in addition to recovering the true number of components K_0 .

We next recall the notion of Wasserstein distance for mixing measures that prove useful in the next sections.

6.1.2 Wasserstein distance for MM

It is useful to analyze the identifiability and convergence of parameter estimation in mixture models using the notion of Wasserstein distance, as in [18, 20]. This distance can be defined as the optimal cost of moving masses in the transformation from one probability measure to another [78, 79].

Definition 6.1. Suppose Θ is equipped with a metric d . The W_r Wasserstein distance for two discrete measures $G = \sum_{k=1}^K \pi_k \delta_{\theta_k}$ and $G_0 = \sum_{k=1}^{K_0} \pi_k^0 \delta_{\theta_k^0}$ is

$$W_r(G, G_0) = \inf_{\mathbf{Q} \in \mathcal{Q}(\pi, \pi^0)} \left[\sum_{i,j} q_{ij} d^r(\theta_i, \theta_i^0) \right]^{1/r}. \quad (6.3)$$

It should be emphasized that if a sequence of probability measures $G_N \in \mathcal{O}_{K_0}$ converges to $G_0 \in \mathcal{E}_{K_0}$ under the W_r metric at a rate $\omega_N = o(1)$ for some $r \geq 1$, then there exists a subsequence of G_N such that the set of atoms of G_N converges to the K_0 atoms of G_0 , up to a permutation of the atoms, at the same rate ω_N .

6.1.3 Posterior contraction rate in infinite mixtures

With the similar idea as in [47], our starting point is the availability of a mixing measure sample G that is drawn from the posterior distribution $\Pi(G \mid \mathcal{W})$, where \mathcal{W} are i.i.d. samples of the mixing density p_{G_0} . Under certain conditions on the kernel density, it can be established that for some Wasserstein metric W_r ,

$$\Pi(G \in \bar{\mathcal{G}}(\Theta) : W_r(G, G_0) \leq \delta \omega_N \mid \mathcal{W}) \xrightarrow{N \rightarrow \infty} 1 \text{ in } P_{G_0}\text{-probability.} \quad (6.4)$$

for *all* constant $\delta > 0$, while $\omega_N = o(1)$ is a vanishing rate. Thus ω_N can be assumed to be (slightly) slower than the actual rate of posterior contraction of the mixture measure. We can also write that ω_N is a rate such that, under the posterior distribution $\Pi(G \mid \mathcal{W})$, $W_r(G, G_0) = o_{P_{G_0}}(\omega_N)$. See [18, 80, 20] for concrete examples of posterior contraction rates in infinite and (overfitted) finite mixtures.

6.2 MTM post-processing consistency for BNP-GLLiM

It is worth noting that we can prove a GLLiM model on (\mathbf{X}, \mathbf{Y}) , see (2.4), with unconstrained parameters $\psi = (\pi_k, \mathbf{c}_k, \Gamma_k, \mathbf{A}_k, \mathbf{b}_k, \Sigma_k)_{k \in [K]}$, is equivalent to a GMM on the joint variable $[\mathbf{X}; \mathbf{Y}]$ with unrestricted parameters, via Lemma 6.2.

6.2.1 Link between GLLiM and joint GMM

We establish a link between GLLiM and joint GMM in [Lemma 6.2](#), which is briefly proved in [Appendix B.1](#).

Lemma 6.2. *A GLLiM model on (\mathbf{X}, \mathbf{Y}) with unconstrained parameters $\boldsymbol{\psi} = (\pi_k, \mathbf{c}_k, \boldsymbol{\Gamma}_k, \mathbf{A}_k, \mathbf{b}_k, \boldsymbol{\Sigma}_k)_{k \in [K]}$, defined in [\(2.4\)](#), is equivalent to a GMM on the joint variable $[\mathbf{X}; \mathbf{Y}] \equiv \mathbf{W}$ with unconstrained parameters $\boldsymbol{\nu} = \{\boldsymbol{\mu}_k, \mathbf{V}_k, \rho_k\}_{k=1}^K$, i.e.,*

$$p(\mathbf{w} \mid \boldsymbol{\psi}) = \sum_{k=1}^K \rho_k \mathcal{N}_{L+D}(\mathbf{w} \mid \boldsymbol{\mu}_k, \mathbf{V}_k). \quad (6.5)$$

The parameter $\boldsymbol{\psi}$ can be expressed as a function of $\boldsymbol{\nu}$ by:

$$\begin{aligned} \pi_k &= \rho_k, \\ \mathbf{c}_k &= \boldsymbol{\mu}_k^{\mathbf{x}}, \\ \boldsymbol{\Gamma}_k &= \mathbf{V}_k^{\mathbf{xx}}, \\ \mathbf{A}_k &= \mathbf{V}_k^{\mathbf{xy}\top} (\mathbf{V}_k^{\mathbf{xx}})^{-1}, \\ \mathbf{b}_k &= \boldsymbol{\mu}_k^{\mathbf{y}} - \mathbf{V}_k^{\mathbf{xy}\top} (\mathbf{V}_k^{\mathbf{xx}})^{-1} \boldsymbol{\mu}_k^{\mathbf{x}}, \\ \boldsymbol{\Sigma}_k &= \mathbf{V}_k^{\mathbf{yy}} - \mathbf{V}_k^{\mathbf{xy}\top} (\mathbf{V}_k^{\mathbf{xx}})^{-1} \mathbf{V}_k^{\mathbf{xy}}. \end{aligned} \quad (6.6)$$

Here, we have defined

$$\boldsymbol{\mu}_k = \begin{bmatrix} \boldsymbol{\mu}_k^{\mathbf{x}} \\ \boldsymbol{\mu}_k^{\mathbf{y}} \end{bmatrix}, \quad \mathbf{V}_k = \begin{bmatrix} \mathbf{V}_k^{\mathbf{xx}} & \mathbf{V}_k^{\mathbf{xy}} \\ \mathbf{V}_k^{\mathbf{yx}} & \mathbf{V}_k^{\mathbf{yy}} \end{bmatrix}.$$

Note that the symmetry $\mathbf{V}_k^\top = \mathbf{V}_k$ of the covariance matrix implies that $\mathbf{V}_k^{\mathbf{xx}}$ and $\mathbf{V}_k^{\mathbf{yy}}$ are symmetric, while $\mathbf{V}_k^{\mathbf{xy}\top} = \mathbf{V}_k^{\mathbf{yx}}$. The parameter vector $\boldsymbol{\nu}$ can be expressed as a function of $\boldsymbol{\psi}$ by:

$$\begin{aligned} \rho_k &= \pi_k, \\ \boldsymbol{\mu}_k &= \begin{bmatrix} \mathbf{c}_k \\ \mathbf{A}_k \mathbf{c}_k + \mathbf{b}_k \end{bmatrix}, \\ \mathbf{V}_k &= \begin{bmatrix} \boldsymbol{\Gamma}_k & (\mathbf{A}_k \boldsymbol{\Gamma}_k)^\top \\ \mathbf{A}_k \boldsymbol{\Gamma}_k & \boldsymbol{\Sigma}_k + \mathbf{A}_k \boldsymbol{\Gamma}_k \mathbf{A}_k^\top \end{bmatrix}. \end{aligned} \quad (6.7)$$

[Lemma 6.2](#) implies that BNP joint GMM and BNP-GLLiM are considered equivalent with respect to number of components in the model selection problem. Therefore, using the results from [\[47, Theorem 3.2\]](#) and [Lemma 6.2](#), it holds that the result of the MTM [Algorithm 1](#) for BNP joint GMM is a consistent estimate of both the number of components and the mixing measure. The latter also admits the upper bound of the posterior contraction rate ω_N . The detailed and main idea of the pseudocode of the MTM [Algorithm 1](#) for BNP joint GMM follows the novel MTM algorithm for BNP-MM of [\[47, Algorithm 1\]](#). Broadly speaking, there are two main phases. The first stage is a probabilistic procedure for the merging of atoms that may be in close proximity. The second stage consists of a deterministic procedure for the truncation of extraneous atoms and their appropriate merging with the remaining atoms in a systematic way. More details can be found in [\[47\]](#).

We aim to provide the theoretical guarantee, [Theorem 6.3](#), for the outcome of [Algorithm 1](#).

Theorem 6.3 (MTM consistency for BNP joint GMM). *Let G be a posterior sample from posterior distribution of any Bayesian procedure, namely, $\Pi(G \mid \mathcal{W})$ according to which the upper bound [\(6.4\)](#) holds for all $\delta > 0$. Let \tilde{G} and \tilde{K} be the outcome of [Algorithm 1](#) applied to G , for an arbitrary constant $c > 0$. Then the following hold*

Algorithm 1 MTM Algorithm for BNP joint GMM

Require: Posterior sample $G = \sum_{k=1}^K \pi_k \delta_{\theta_k}$, posterior contraction rate ω_N from (6.4), and a tuning parameter c .

Ensure: Discrete measure \tilde{G} and its number of supporting atoms \tilde{K} .

Stage 1: Merge procedure:

1: Reorder atoms $\{\theta_k\}_{k \in [K]}$ by simple random sampling without replacement with corresponding weights $\{\pi_1, \pi_2, \dots\}$.

let τ_1, τ_2, \dots denote the new indices, and set $\mathcal{E} = \{\tau_j\}_j$ as the existing set of atoms.

2: Sequentially for each index $\tau_j \in \mathcal{E}$, if there exists an index $\tau_i < \tau_j$ such that $d(\theta_{\tau_i}, \theta_{\tau_j}) \leq \omega_N$, then:

update $\pi_{\tau_i} = \pi_{\tau_i} + \pi_{\tau_j}$, and remove τ_j from \mathcal{E} .

3: Collect $G' = \sum_{j: \tau_j \in \mathcal{E}} \pi_{\tau_j} \delta_{\theta_{\tau_j}}$.

write G' as $\sum_{k=1}^K q_k \delta_{\phi_k}$ so that $q_1 \geq q_2 \geq \dots$

Stage 2: Truncate-Merge procedure:

4: Set $\mathcal{A} = \{i : q_i > (c\omega_N)^r\}$, $\mathcal{N} = \{i : q_i \leq (c\omega_N)^r\}$.

5: For each index $i \in \mathcal{A}$, if there is $j \in \mathcal{A}$ such that $j < i$ and $q_i \|\phi_i - \phi_j\|^r \leq (c\omega_N)^r$, then remove i from \mathcal{A} and add it to \mathcal{N} .

6: For each $i \in \mathcal{N}$, find atom ϕ_j among $j \in \mathcal{A}$ that is nearest to ϕ_i

update $q_j = q_j + q_i$.

7: Return $\tilde{G} = \sum_{j \in \mathcal{A}} q_j \delta_{\phi_j}$ and $\tilde{K} = |\mathcal{A}|$.

(a) $\Pi(\tilde{K} = K_0 \mid \mathcal{W}) \xrightarrow{N \rightarrow \infty} 1$ in P_{G_0} -probability.

(b) For all $\delta > 0$, $\Pi(G \in \bar{\mathcal{G}}(\Theta) : W_r(\tilde{G}, G_0) \leq \delta \omega_N \mid \mathcal{W}) \xrightarrow{N \rightarrow \infty} 1$ in P_{G_0} -probability.

Remark 6.4. To this theorem we add some comments on posterior consistency for the number of components in BNP-GLLiM after the MTM algorithm post-processing.

(i) As a complementary result to [47], the aim of this paper is to study the practical viability of MTM Algorithm 1 and Theorem 6.3 in the context of high-dimensional BNP-GLLiM. In order to do this, we first need to specify the metric d in Θ , e.g.,

$$d(\theta_{\tau_i}, \theta_{\tau_j}) = \|\mu_{\tau_i} - \mu_{\tau_j}\| + \|\mathbf{V}_{\tau_i} - \mathbf{V}_{\tau_j}\|. \quad (6.8)$$

Here $\|\cdot\|$ denotes either the l_2 norm elements in \mathbb{R}^{L+D} or the entrywise l_2 norm for matrices in $\mathbb{R}^{(L+D) \times (L+D)}$.

(ii) In practice, one may not have a mixing measure G sampled from the posterior $\Pi(\cdot \mid \mathcal{W})$, but rather a sample of G itself. In particular, to deal with large data sets, we need to use VBEM. Therefore, for example in BNP-GLLiM1, instead, we only obtain a sample F_N from the variational posterior $G_V = \sum_{k=1}^K \mathbb{E}_{q_{\mu}}[\pi_k(\mu)] \delta_{\theta_k}$. Here $\mathbb{E}_{q_{\tau}}[\pi_k(\tau)]$ and $\theta_k = (\mu_k, \mathbf{V}_k)$ are defined in (4.3) and (4.1) respectively. However, as long as F_N is sufficiently close to G in the sense that $W_r(F_N, G) \lesssim W_r(G, G_0)$, we can still apply the MTM algorithm to F_N instead. This requires an extension of the above Theorem 6.3 to cover this scenario and verify this approximation condition, which we leave for future work.

7 Numerical experiments

7.1 Data generation

We illustrate our theoretical results on simulated datasets in a more general setting for the BNP approach compared to those considered by [81, 82, 7]. More specifically, we consider the following true inverse CPDF from GLLiM model as follows

$$s_0(\mathbf{y} | \mathbf{x}; \boldsymbol{\psi}^0) = \sum_{k=1}^{K_0} \frac{\pi_k^0 \mathcal{N}_L(\mathbf{x} | \mathbf{c}_k^0, \boldsymbol{\Gamma}_k^0)}{\sum_{j=1}^{K_0} \pi_j^0 \mathcal{N}_L(\mathbf{x} | \mathbf{c}_j^0, \boldsymbol{\Gamma}_j^0)} \mathcal{N}_D(\mathbf{y} | \mathbf{A}_k^0 \mathbf{x} + \mathbf{b}_k^0, \boldsymbol{\Sigma}_k^0).$$

Here $K_0 = 3$, $L = D = 1$, and $\boldsymbol{\psi}^0 = (\boldsymbol{\pi}^0, \mathbf{c}^0, \boldsymbol{\Gamma}^0, \mathbf{A}^0, \mathbf{b}^0, \boldsymbol{\Sigma}^0)$, where

$$\begin{aligned} \boldsymbol{\pi}^0 &= (0.3, 0.4, 0.3), & \mathbf{c}^0 &= (1, 0.05, -1), & \boldsymbol{\Gamma}^0 &= (0.1, 0.2, 0.1), \\ \mathbf{A}^0 &= (-15, 3, 15), & \mathbf{b}^0 &= (-2, 1, -2), & \boldsymbol{\Sigma}^0 &= (0.5, 0.3, 0.5). \end{aligned}$$

Figure 3a shows typical $N = 1000$ realizations of the true inverse CPDF from GLLiM, representing a π -shape simulation with three clusters without labels.

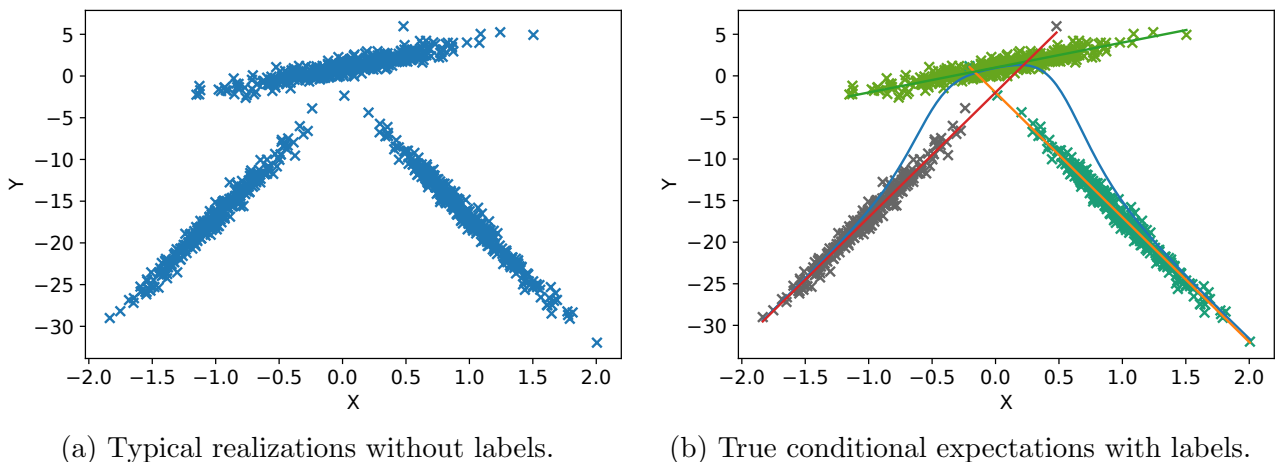


Figure 3: Typical 1000 realizations of GLLiM's true inverse CPDF with its true conditional expectations.

7.2 Model selection, clustering and regression tasks via MTM-BNP-GLLiM1

Our goal is to evaluate the inverse and forward CPDF, as well as the conditional means, to investigate the empirical performance of our MTM-BNP-GLLiM1 in the previous simulation. In Figure 4 it is clear that with the help of MTM Algorithm 1, MTM-BNP-GLLiM1 can perform regression, clustering and model selection well simultaneously. Without the MTM procedure, BNP-GLLiM1 performs poorly in model selection, clustering and CPDF estimation, except for conditional expectations as shown in Figure 5.

Next, we illustrate the performance of the MTM algorithm when applied to the variational posterior from BNP-GLLiM1. More specifically, the samples in our 100 trials are drawn from $G_V = \sum_{k=1}^K \mathbb{E}_{q_{\boldsymbol{\mu}}} [\pi_k(\boldsymbol{\mu})] \delta_{\boldsymbol{\theta}_k}$, where $\mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})]$ and $\boldsymbol{\theta}_k = (\boldsymbol{\mu}_k, \mathbf{V}_k)$ are defined in (4.3) and (4.1) respectively.

We know that for some constant \tilde{C} , which depends on the covariance matrix \mathbf{V}_k^0 , the location parameters $\boldsymbol{\mu}_k^0$ and the weights π_k^0 , the contraction rate of mixing measures under the location Gaussian DP-MM is $\tilde{C} (\log(N))^{-1/2}$ with respect to the W_2 -norm. Similar to [47], our first attempt to choose

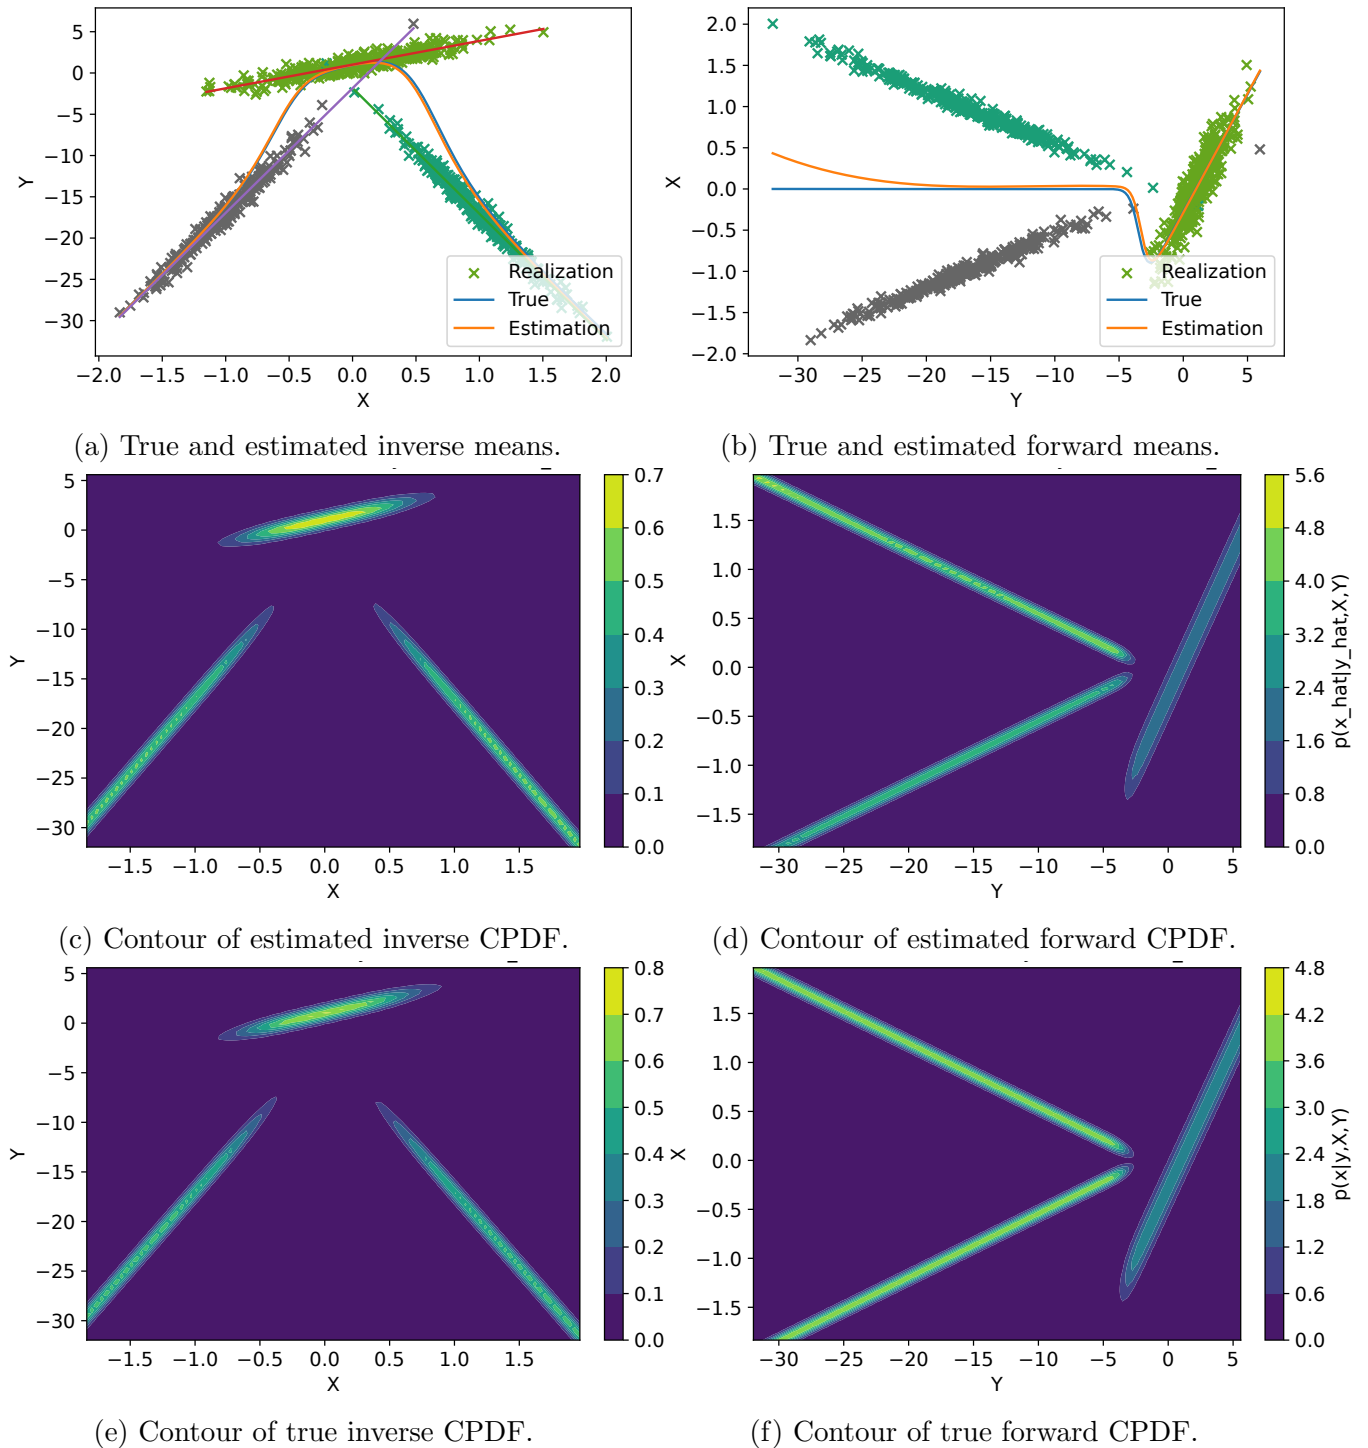


Figure 4: True and estimated inverse and forward CPDF of GLLiM with a good number of components ($K_{\text{MTM}} = 3$) using MTM algorithm for post-processing in MTM-BNP-GLLiM1.

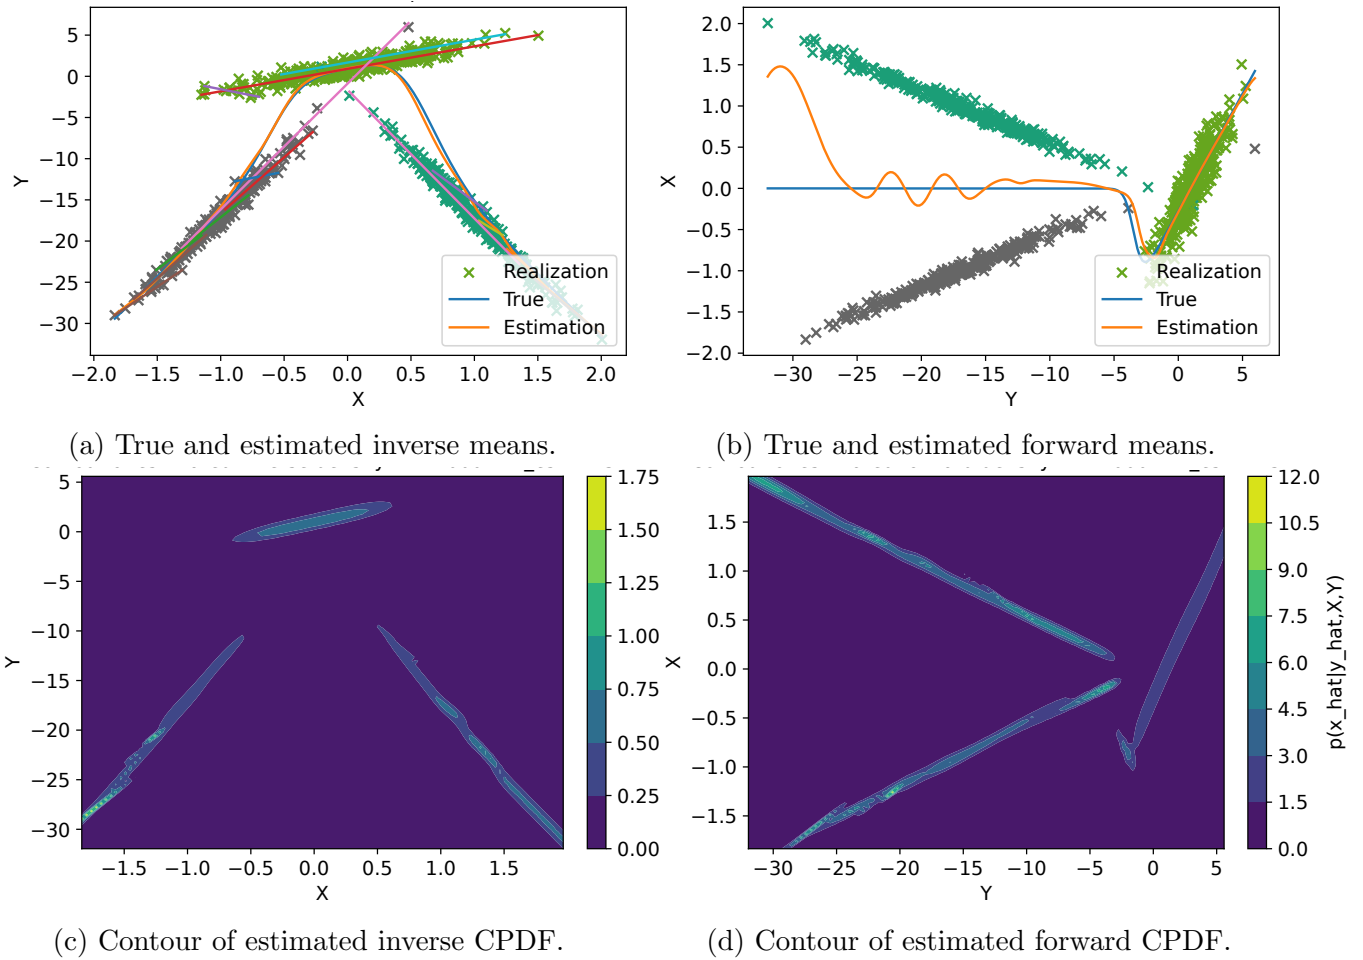


Figure 5: True and estimated inverse and forward means and CPDFs of GLLiM without MTM algorithm for post-processing in BNP-GLLiM with truncated number of components $K = 20$.

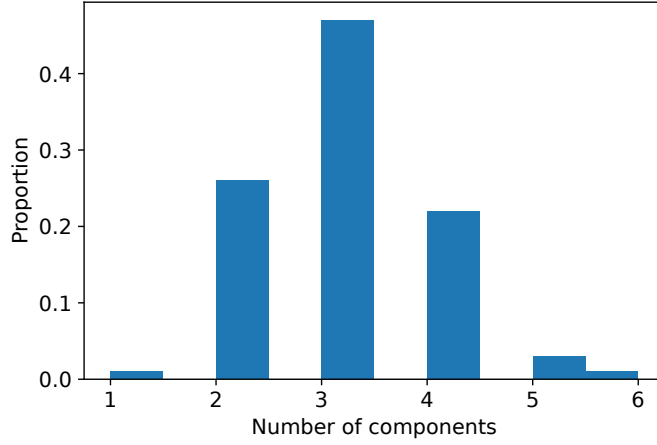


Figure 6: Histogram of K_{MTM} using variational posterior sample with 100 trials and $c = 0.45$.

w_N to satisfy (6.4) is $w_N = \left(\frac{\log(\log(N))}{\log(N)}\right)^{1/2}$. In fact, we can choose any w_N , as long as $\frac{w_N}{\log(N)^{-1/2}} \rightarrow \infty$, in order for w_N to satisfy (6.4).

Since we only work with finite sample N , it is not expected that the posterior probability for $K_{\text{MTM}} = K_0$ is close to 1 and the input c to Algorithm 1 should be chosen so that $\frac{\tilde{C}}{(\log(\log(N)))^{1/2}} \leq c$. Furthermore, based on [47, Equation (26)] with a useful lower bound on the posterior mass the mode, for any $1 > \epsilon > 0$, $(1 - \epsilon) \left(1 - \sum_{k=1}^3 \frac{c_0^{r/2}}{\pi_k^0}\right)$, we hope to identify K_0 via the posterior mode with a reasonable estimate. To guarantee $K = K_0$ consistently using the posterior mode safety, we have to choose $c < c_0$, with c_0 satisfying

$$(1 - \epsilon) \left(1 - \sum_{k=1}^3 \frac{c_0^{r/2}}{\pi_k^0}\right) > \frac{1}{2} \Leftrightarrow \frac{1 - 2\epsilon}{2(1 - \epsilon)} \left(\sum_{k=1}^3 \frac{1}{\pi_k^0}\right)^{-1} > c_0^{r/2} > c^{r/2}. \quad (7.1)$$

Therefore, we can choose

$$\left[\frac{1 - 2\epsilon}{4(1 - \epsilon)} \left(\sum_{k=1}^3 \frac{1}{\pi_k^0}\right)^{-1}\right]^{2/r} = c_0 > c, \text{ for all } \frac{1}{2} > \epsilon > 0. \quad (7.2)$$

In particular, it is unrealistic to obtain the exact computation of the upper bound c_0 and the lower bound $\frac{\tilde{C}}{(\log(\log(N)))^{1/2}}$. However, a reasonable estimate may be possible. To this end, we followed the same setting as in [47] to simplify such a complicated procedure by considering a large range of c , and show that there is a range where we can robustly identify the true number of components via the posterior mode. Here, $c = 0.45$ leads to a quite good posterior mode in our numerical experiments, see more in Figure 6.

Although we do not have a theoretical result for the convergence rate of the variational posterior of BNP-GLLiM to the true data generating process, Figure 6 seems to suggest that MTM-BNP-GLLiM1 gave a comparable good result to the location Gaussian DP-MM in the simulation studies in [47].

8 Perspectives

As indicated in Remark 6.4, there is a crucial need to formally establish general conditions on the prior, the likelihood and the variational class to characterize the convergence rate of the variational posterior of BNP-GLLiM to the true data generating process. Using the very similar ‘‘prior mass and testing’’ conditions as in [83], we believe that an interesting but challenging extension from the work of [84, variational posterior unconditional distributions for MM models] and [85, 86, 87, an adaptive Bayesian estimation for MM and MoE but for true posterior distribution] can help shed some light and answer this important question.

Furthermore, it is important to establish an extensional convergence property of our VBEM algorithm for BNP-GLLiM. This property is only known for GMM from [88]. A potential improvement of the VBEM algorithm developed for BNP-GLLiM can be achieved by combining it with MCMC, taking advantage of both inference approaches as in [89]. Then, it is also of interest to investigate the performance of BNP-GLLiM on high-dimensional real-world datasets. Finally, as mentioned in Section 7.2, the selection of a good data-driven tuning parameter c as the same idea from the slope heuristic of [32] is crucial for the success of the MTM procedure for any BNP model. We leave these interesting but challenging questions for future studies.

Appendix A Details of VBEM for BNP-GLLiM model

A.1 VB-E- τ step from Section 3.2.1

To achieve results from Section 3.2.1, we make use of (2.18), (2.23), (3.26), and are only interested in the functional dependence of the right-hand side of (3.6) on the variable τ_k . Thus any terms that do not depend on τ_k can be absorbed into the additive normalization constant, giving

$$\begin{aligned}
q_{\tau_k}(\tau_k) &= \exp \left\{ \mathbb{E}_{q_{\alpha,\sigma}} [\log p(\tau_k | \alpha, \sigma)] + \sum_{n=1}^N \mathbb{E}_{q_{z_n} q_{\tau_{\setminus\{k\}}}} [\log \pi_{z_n}(\boldsymbol{\tau})] \right\} + \text{constant} \\
&\propto \exp \left\{ -\mathbb{E}_{q_{\alpha,\sigma}} [\sigma] \log \tau_k + [\mathbb{E}_{q_{\alpha,\sigma}} [\alpha] + k\mathbb{E}_{q_{\alpha,\sigma}} [\sigma] - 1] \log(1 - \tau_k) \right. \\
&\quad \left. + \sum_{n=1}^N q_{z_n}(k) \log \tau_k + \sum_{n=1}^N \sum_{j=k+1}^K q_{z_n}(j) \log(1 - \tau_k) \right\} \\
&= \text{Beta}(\tau_k | \hat{\gamma}_{k,1}, \hat{\gamma}_{k,2}). \tag{A.1}
\end{aligned}$$

Here,

$$\hat{\gamma}_{k,1} = 1 - \mathbb{E}_{q_{\alpha,\sigma}} [\sigma] + \sum_{n=1}^N q_{z_n}(k) = 1 - \mathbb{E}_{q_{\sigma}} [\sigma] + \sum_{n=1}^N q_{z_n}(k), \tag{A.2}$$

$$\hat{\gamma}_{k,2} = \mathbb{E}_{q_{\alpha,\sigma}} [\alpha] + k\mathbb{E}_{q_{\alpha,\sigma}} [\sigma] + \sum_{n=1}^N \sum_{j=k+1}^K q_{z_n}(j) = \mathbb{E}_{q_{\alpha,\sigma}} [\alpha] + k\mathbb{E}_{q_{\alpha,\sigma}} [\sigma] + \sum_{j=k+1}^K \sum_{n=1}^N q_{z_n}(j). \tag{A.3}$$

Furthermore, we used the fact that

$$\begin{aligned}
\log p(\tau_k | \alpha, \sigma) &= \log [\text{Beta}(\tau_k | 1 - \sigma, \alpha + k\sigma)] = \log \left[\frac{\Gamma(1 - \sigma + k\sigma)}{\Gamma(1 - \sigma)\Gamma(\alpha + k\sigma)} \tau_k^{1-\sigma-1} (1 - \tau_k)^{\alpha+k\sigma-1} \right], \\
\log \pi_{z_n}(\boldsymbol{\tau}) &= \log \left[\tau_{z_n} \prod_{l=1}^{z_n-1} (1 - \tau_l) \right] = \log \tau_{z_n} + \sum_{l=1}^{z_n-1} \log(1 - \tau_l), \\
q_{\tau_{\setminus\{k\}}}(\boldsymbol{\tau}_l) &= \prod_{i=1, i \neq k}^{K-1} q_{\tau_i}(\tau_i), \quad d\tau_{\setminus\{k\}} = \prod_{i=1, i \neq k}^{K-1} d\tau_i. \tag{A.4}
\end{aligned}$$

Finally, we have for $k \in [K]$, let $N_k = \sum_{n=1}^N q_{z_n}(k)$ correspond to the weight of cluster k , then

$$q_{\tau_k}(\tau_k) = \text{Beta}(\tau_k | \hat{\gamma}_{k,1}, \hat{\gamma}_{k,2}), \tag{A.5}$$

$$\hat{\gamma}_{k,1} = 1 - \mathbb{E}_{q_{\alpha,\sigma}} [\sigma] + N_k, \quad \hat{\gamma}_{k,2} = \mathbb{E}_{q_{\alpha,\sigma}} [\alpha] + k\mathbb{E}_{q_{\alpha,\sigma}} [\sigma] + \sum_{l=k+1}^K N_l. \tag{A.6}$$

A.2 VB-E- (α, σ) step from Section 3.2.2

In the PY case, to achieve results from Section 3.2.2, we make use of (2.18), (2.17), (3.26), (3.27), and are only interested in the functional dependence of the right-hand side of (3.6) to the variables (α, σ) . Thus, any terms that do not depend on (α, σ) can be included in the additive normalization constant. This results in $q_{\alpha, \sigma}(\alpha, \sigma)$ being proportional to

$$\begin{aligned}
\tilde{q}_{\alpha, \sigma}(\alpha, \sigma) &= p(\alpha, \sigma \mid s_1, s_2, a) \exp \left\{ \mathbb{E}_{\prod_{k=1}^{K-1} q_{\tau_k}} \left[\log \prod_{k=1}^{K-1} p(\tau_k \mid \alpha, \sigma) \right] \right\} \\
&= p(\alpha, \sigma \mid s_1, s_2, a) \prod_{k=1}^{K-1} \frac{\Gamma(1 - \sigma + \alpha + k\sigma)}{\Gamma(1 - \sigma)\Gamma(\alpha + k\sigma)} \exp \left\{ \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [-\sigma \log \tau_k + (\alpha - 1 + k\sigma) \log(1 - \tau_k)] \right\} \\
&= p(\alpha, \sigma \mid s_1, s_2, a) \frac{1}{\Gamma(1 - \sigma)^{K-1}} \prod_{k=1}^{K-1} \frac{[\alpha + (k-1)\sigma] \Gamma(\alpha + (k-1)\sigma)}{\Gamma(\alpha + k\sigma)} \\
&\quad \times \exp \left\{ \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [-\sigma (\log \tau_k - k \log(1 - \tau_k)) + (\alpha - 1) \log(1 - \tau_k)] \right\} \\
&= p(\alpha, \sigma \mid s_1, s_2, a) \frac{1}{\Gamma(1 - \sigma)^{K-1}} \prod_{k=1}^{K-1} [\alpha + (k-1)\sigma] \frac{\prod_{l=0}^{K-2} \Gamma(\alpha + l\sigma)}{\prod_{k=1}^{K-1} \Gamma(\alpha + k\sigma)} \\
&\quad \times \exp \left\{ \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [-\sigma (\log \tau_k - k \log(1 - \tau_k)) + (\alpha - 1) \log(1 - \tau_k)] \right\} \\
&= p(\alpha, \sigma \mid s_1, s_2, a) \frac{\Gamma(\alpha)}{\Gamma(1 - \sigma)^{K-1} \Gamma(\alpha + (K-1)\sigma)} \prod_{k=1}^{K-1} [\alpha + (k-1)\sigma] \\
&\quad \times \exp \left\{ -\sigma \left[\sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log \tau_k] - \sum_{k=1}^{K-1} k \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)] \right] + (\alpha - 1) \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)] \right\}, \tag{A.7}
\end{aligned}$$

where we used the fact that $\Gamma(x+1) = x\Gamma(x)$. Except in the DP-GLLiM case, *i.e.*, $\sigma = 0$, the normalizing constant, $(\int \tilde{q}_{\alpha, \sigma}(\alpha, \sigma) d(\alpha, \sigma))^{-1}$, for $\tilde{q}_{\alpha, \sigma}$ is not tractable. However, to perform VBEM in Section 3.3, we do not need the full $q_{\alpha, \sigma}$ distribution, but only the means $\mathbb{E}_{q_{\alpha, \sigma}}[\alpha]$ and $\mathbb{E}_{q_{\alpha, \sigma}}[\sigma]$. One solution, therefore, is to use importance sampling or MCMC to compute these expectations by

means of Monte Carlo sums. Via utilizing the prior on (α, σ) given in (2.17), it holds that

$$\begin{aligned}
\tilde{q}_{\alpha, \sigma}(\alpha, \sigma) &= \frac{1}{\Gamma(s_1)} s_2^{s_1} (\alpha + \sigma)^{s_1 - 1} \exp\{-s_2(\alpha + \sigma)\} \frac{\Gamma(\alpha)p(\sigma | a)}{\Gamma(1 - \sigma)^{K-1}\Gamma(\alpha + (K-1)\sigma)} \prod_{k=1}^{K-1} [\alpha + (k-1)\sigma] \\
&\quad \times \exp\left\{-\sigma \left[\sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log \tau_k] - \sum_{k=1}^{K-1} k \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)] \right] + (\alpha - 1) \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)] \right\} \\
&= \frac{1}{\Gamma(s_1)} s_2^{s_1} (\alpha + \sigma)^{s_1 - 1} \exp\left\{-\left[s_2 - \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)]\right] (\alpha + \sigma)\right\} \\
&\quad \times e^{-\sigma\xi} \exp\left\{-\sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)]\right\} \frac{\Gamma(\alpha)p(\sigma | a)}{\Gamma(1 - \sigma)^{K-1}\Gamma(\alpha + (K-1)\sigma)} \prod_{k=1}^{K-1} [\alpha + (k-1)\sigma] \\
&= \frac{1}{\Gamma(s_1)} (\alpha + \sigma)^{s_1 - 1} \exp\left\{-\left[s_2 - \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)]\right] (\alpha + \sigma)\right\} \left[s_2 - \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)]\right]^{s_1 - s_1} \\
&\quad \times e^{-\sigma\xi} \exp\left\{-\sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)]\right\} \frac{\Gamma(\alpha)p(\sigma | a)}{\Gamma(1 - \sigma)^{K-1}\Gamma(\alpha + (K-1)\sigma)} \prod_{k=1}^{K-1} [\alpha + (k-1)\sigma] s_2^{s_1} \\
&\propto \text{Gam}(\alpha + \sigma | \hat{s}_1, \hat{s}_2) e^{-\sigma\xi} \frac{\Gamma(\alpha)p(\sigma | a)}{\Gamma(1 - \sigma)^{K-1}\Gamma(\alpha + (K-1)\sigma)} \prod_{k=1}^{K-1} \frac{\alpha + (k-1)\sigma}{\alpha + \sigma}. \tag{A.8}
\end{aligned}$$

Here, given that $\psi(\cdot)$ is the digamma function defined by $\psi(z) = \frac{d}{dz} \log \Gamma(z) = \frac{\Gamma'(z)}{\Gamma(z)}$, we have

$$\xi = \sum_{k=1}^{K-1} \mathbb{E}_{q_{\tau_k}} [\log \tau_k] - \sum_{k=1}^{K-1} (k-1) \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)], \tag{A.9}$$

$$\mathbb{E}_{q_{\tau_k}} [\log \tau_k] = \psi(\hat{\gamma}_{k,1}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}), \tag{A.10}$$

$$\mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)] = \psi(\hat{\gamma}_{k,2}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}), \tag{A.11}$$

$$\hat{s}_1 = s_1 + K - 1, \quad \hat{s}_2 = s_2 - \sum_{k=1}^{K-1} [\psi(\hat{\gamma}_{k,2}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2})]. \tag{A.12}$$

We propose to use the following important distribution $\nu(\alpha, \sigma) = \text{Gam}(\alpha + \sigma | \hat{s}_1, \hat{s}_2) p(\sigma | a)$ where $p(\sigma | a)$ is the uniform distribution on $[0, 1]$, denoted as $\mathcal{U}_{[0,1]}(\sigma)$. Then we obtain an expression for the importance weights,

$$W(\alpha, \sigma) = \frac{\tilde{q}_{\alpha, \sigma}(\alpha, \sigma)}{\nu(\alpha, \sigma)} = e^{-\sigma\xi} \frac{\Gamma(\alpha)}{\Gamma(1 - \sigma)^{K-1}\Gamma(\alpha + (K-1)\sigma)} \prod_{k=1}^{K-1} \frac{\alpha + (k-1)\sigma}{\alpha + \sigma}. \tag{A.13}$$

The importance sampling scheme then consists of

- For $m \in [M]$, first simulate independently σ_m from $\mathcal{U}_{[0,1]}(\sigma)$ and then simulate conditionally α_m with the σ_m -shifted gamma $\mathcal{G}(\sigma_m, \hat{s}_1, \hat{s}_2)$. This later simulation is easily obtained by simulating a standard $\gamma(\alpha'_m | \hat{s}_1, \hat{s}_2)$ and then subtracting σ_m from the result.
- Compute the importance weights $w_m = W(\alpha_m, \sigma_m)$.
- Approximate the means

$$\mathbb{E}_{q_{\alpha, \sigma}} [\alpha] \simeq \sum_{m=1}^M \frac{w_m}{\sum_{i=1}^M w_i} \alpha_m, \quad \mathbb{E}_{q_{\alpha, \sigma}} [\sigma] \simeq \sum_{m=1}^M \frac{w_m}{\sum_{i=1}^M w_i} \sigma_m. \tag{A.14}$$

Note that this complication is due to the PY. In the DP-GLLiM case, by substituting $\sigma = 0$ in (A.8), the E step is much simpler, as it reduces to computing the approximate posterior expectation of α , namely

$$\mathbb{E}_{q_{\alpha,0}}[\alpha] = \frac{\widehat{s}_1}{\widehat{s}_2}, \quad q_{\alpha,0} = \text{Gam}(\alpha \mid \widehat{s}_1, \widehat{s}_2). \quad (\text{A.15})$$

A.3 VB-E-Z step from Sections 3.2.3 and 3.3.1

In some situations, it is useful to use a 1-of- K binary vector \mathbf{z}_n to represent the latent variable z_n for each observation $(\mathbf{y}_n, \mathbf{x}_n)$. To be more precise, we introduce a K -dimensional binary random variable $\mathbf{z}_n = (z_{nk})_{k \in [K]}$, $K \leq \infty$, with a 1-of- K representation in which a particular element z_{nk} is equal to 1, *i.e.*, $z_n = k$, and all other elements are equal to 0. The values of z_{nk} thus satisfy $z_{nk} \in \{0, 1\}$ and $\sum_{k \in [K]} z_{nk} = 1, \forall n \in \mathbb{N}^*$. If there is no confusion, we also denote \mathcal{Z} as the latent matrix $\mathcal{Z} = (z_{nk})_{n \in [N], k \in [K]}$. It is worth to mentioning that when using a 1-of- K representation of \mathbf{z}_n , we can also write down marginal the conditional distributions of \mathcal{X} and $\mathcal{Y} \mid \mathcal{X}$, corresponding to (2.24) and (2.25), respectively, in the form

$$p(\mathcal{X} \mid \mathcal{Z}, \mathbf{c}, \mathbf{\Gamma}) = \prod_{n=1}^N \prod_{k=1}^K \mathcal{N}_L(\mathbf{x}_n \mid \mathbf{c}_k, \mathbf{\Gamma}_k)^{z_{nk}}, \quad (\text{A.16})$$

$$p(\mathbf{x}_n \mid \mathbf{c}, \mathbf{\Gamma}) = \sum_{k=1}^K p_{z_n}(k) \mathcal{N}_L(\mathbf{x}_n \mid \mathbf{c}_k, \mathbf{\Gamma}_k), \quad p_{z_n}(k) \equiv p(z_n = k) = \pi_k(\boldsymbol{\tau}), \quad (\text{A.17})$$

$$p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}; \mathbf{A}, \mathbf{b}, \mathbf{\Sigma}) = \prod_{n=1}^N \prod_{k=1}^K \mathcal{N}_D(\mathbf{y}_n \mid \mathbf{A}_k \mathbf{x}_n + \mathbf{b}_k, \mathbf{\Sigma}_k)^{z_{nk}}, \quad (\text{A.18})$$

$$p(\mathbf{y}_n \mid \mathbf{x}_n; \mathbf{A}, \mathbf{b}, \mathbf{\Sigma}) = \sum_{k=1}^K \frac{p_{z_n}(k) \mathcal{N}_L(\mathbf{x}_n \mid \mathbf{c}_k, \mathbf{\Gamma}_k)}{\sum_{l=1}^K p_{z_n}(l) \mathcal{N}_L(\mathbf{x}_n \mid \mathbf{c}_l, \mathbf{\Gamma}_l)} \mathcal{N}_D(\mathbf{y}_n \mid \mathbf{A}_k \mathbf{x}_n + \mathbf{b}_k, \mathbf{\Sigma}_k). \quad (\text{A.19})$$

The observations \mathcal{X} and \mathcal{Y} are therefore i.i.d. and generated from the same GMM (A.17) and infinite GLLiM (A.18), respectively. Similarly, (2.23) can be written down in the form

$$p(\mathcal{Z} \mid \boldsymbol{\tau}) = \prod_{n=1}^N \prod_{k=1}^K \pi_k(\boldsymbol{\tau})^{z_{nk}}. \quad (\text{A.20})$$

By using the decomposition (3.9), the representation (A.16), (A.18) and absorbing any terms that are independent on \mathcal{Z} into the additive normalization constant, we obtain

$$\begin{aligned} \log q_{\mathcal{Z}}(\mathcal{Z}) &\equiv \log q_{\mathcal{Z}}^{(r)}(\mathcal{Z}) = \mathbb{E}_{q_{\Theta}} \left[\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}; \widehat{\mathbf{A}}, \widehat{\mathbf{b}}, \widehat{\mathbf{\Sigma}}) p(\mathcal{X} \mid \mathcal{Z}; \widehat{\mathbf{c}}, \widehat{\mathbf{\Gamma}}) p(\mathcal{Z} \mid \boldsymbol{\tau}) \right] + \text{constant} \\ &\propto \mathbb{E}_{q_{\Theta}} \left[\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}; \widehat{\mathbf{A}}, \widehat{\mathbf{b}}, \widehat{\mathbf{\Sigma}}) \right] + \mathbb{E}_{q_{\Theta}} \left[\log p(\mathcal{X} \mid \mathcal{Z}; \widehat{\mathbf{c}}, \widehat{\mathbf{\Gamma}}) \right] + \mathbb{E}_{q_{\Theta}} \left[\log p(\mathcal{Z} \mid \boldsymbol{\tau}) \right] \\ &= \sum_{n=1}^N \sum_{k=1}^K z_{nk} \mathbb{E}_{q_{\Theta}} \left[\log \mathcal{N}_D(\mathbf{y}_n \mid \widehat{\mathbf{A}}_k \mathbf{x}_n + \widehat{\mathbf{b}}_k, \widehat{\mathbf{\Sigma}}_k) \right] + \sum_{n=1}^N \sum_{k=1}^K z_{nk} \mathbb{E}_{q_{\Theta}} \left[\log \mathcal{N}_L(\mathbf{x}_n \mid \widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k) \right] \\ &\quad + \mathbb{E}_{q_{\Theta}} \left[\sum_{n=1}^N \log(\pi_{z_n}(\boldsymbol{\tau})) \right] \\ &= \sum_{n=1}^N \sum_{k=1}^K z_{nk} \log \mathcal{N}_D(\mathbf{y}_n \mid \widehat{\mathbf{A}}_k \mathbf{x}_n + \widehat{\mathbf{b}}_k, \widehat{\mathbf{\Sigma}}_k) + \sum_{n=1}^N \sum_{k=1}^K z_{nk} \log \mathcal{N}_L(\mathbf{x}_n \mid \widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k) \\ &\quad + \sum_{n=1}^N \sum_{k=1}^K z_{nk} \mathbb{E}_{q_{\tau}} \left[\log(\pi_k(\boldsymbol{\tau})) \right] = \sum_{n=1}^N \sum_{k=1}^K z_{nk} \log \rho_{nk}. \end{aligned} \quad (\text{A.21})$$

Here, we used the fact that

$$\begin{aligned}
\log \rho_{nk} &= \log \mathcal{N}_D \left(\mathbf{y}_n \mid \widehat{\mathbf{A}}_k \mathbf{x}_n + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k \right) + \log \mathcal{N}_L \left(\mathbf{x}_n \mid \widehat{\mathbf{c}}_k, \widehat{\boldsymbol{\Gamma}}_k \right) + \mathbb{E}_{q_{\boldsymbol{\tau}}} [\log (\pi_k (\boldsymbol{\tau}))] \\
&= -\frac{D}{2} \log(2\pi) - \frac{1}{2} \log \left| \widehat{\boldsymbol{\Sigma}}_k \right| - \frac{1}{2} \left(\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k \right)^\top \widehat{\boldsymbol{\Sigma}}_k^{-1} \left(\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k \right) \\
&\quad - \frac{L}{2} \log(2\pi) - \frac{1}{2} \log \left| \widehat{\boldsymbol{\Gamma}}_k \right| - \frac{1}{2} \left(\mathbf{x}_n - \widehat{\mathbf{c}}_k \right)^\top \widehat{\boldsymbol{\Gamma}}_k^{-1} \left(\mathbf{x}_n - \widehat{\mathbf{c}}_k \right) \\
&\quad + \psi(\widehat{\gamma}_{k,1}) - \psi(\widehat{\gamma}_{k,1} + \widehat{\gamma}_{k,2}) + \sum_{l=1}^{k-1} \psi(\widehat{\gamma}_{l,2}) - \psi(\widehat{\gamma}_{l,1} + \widehat{\gamma}_{l,2}).
\end{aligned} \tag{A.22}$$

By taking exponential of both sides and taking into account the normalized constant, it holds that

$$q_{\mathcal{Z}}(\mathcal{Z}) = \frac{1}{\sum_{l=1}^K \rho_{nl}} \prod_{n=1}^N \prod_{k=1}^K \rho_{nk}^{z_{nk}} = \prod_{n=1}^N \prod_{k=1}^K r_{nk}^{z_{nk}}, \quad r_{nk} = \frac{\rho_{nk}}{\sum_{l=1}^K \rho_{nl}}. \tag{A.23}$$

Note also that $z_n = k$ if and only if the latent matrix \mathcal{Z} reduces to a sparse matrix \mathcal{Z}_{nk} which has only one position different from 0, namely $z_{nk} = 1$. This leads to the following simplified notation:

$$\log q_{z_n}(k) \equiv \log q_{z_n}(z_n = k) \equiv \log q_{\mathcal{Z}}(\mathcal{Z}_{nk}) = r_{nk}. \tag{A.24}$$

With respect to the VBEM for GT-BNP-GLLiM model from Section 3.3.1, it is almost similar to the previous step, except that we have to take into account the randomness of \mathbf{c} and $\boldsymbol{\Gamma}$. Namely, we have

$$\begin{aligned}
q_{z_n}(z_n) &\propto \exp \mathbb{E}_{q_{\boldsymbol{\theta}}} \left[\log \left(p \left(\mathbf{y}_n \mid \mathbf{x}_n, z_n; \widehat{\mathbf{A}}_{z_n}, \widehat{\mathbf{b}}_{z_n}, \widehat{\boldsymbol{\Sigma}}_{z_n} \right) p \left(\mathbf{x}_n \mid z_n, \mathbf{c}_{z_n}, \boldsymbol{\Gamma}_{z_n} \right) p(\mathbf{z} \mid \boldsymbol{\tau}) \right) \right] \\
&= \exp \left\{ \log p \left(\mathbf{y}_n \mid \mathbf{x}_n, z_n; \widehat{\mathbf{A}}_{z_n}, \widehat{\mathbf{b}}_{z_n}, \widehat{\boldsymbol{\Sigma}}_{z_n} \right) + \mathbb{E}_{q_{\boldsymbol{\theta}_{z_n}^*}} [\log p \left(\mathbf{x}_n \mid z_n, \mathbf{c}_{z_n}, \boldsymbol{\Gamma}_{z_n} \right)] + \mathbb{E}_{q_{\boldsymbol{\tau}}} [\log \pi_{z_n}(\boldsymbol{\tau})] \right\}.
\end{aligned} \tag{A.25}$$

Here, for $z_n = k$, it holds that

$$\begin{aligned}
\mathbb{E}_{q_{\boldsymbol{\theta}_{z_n}^*}} \left[\log p \left(\mathbf{x}_n \mid z_n, \widehat{\mathbf{c}}_{z_n}, \widehat{\boldsymbol{\Gamma}}_{z_n} \right) \right] &= \mathbb{E}_{q_{\boldsymbol{\theta}_{z_n}^*}} \left[\log \mathcal{N}_L \left(\mathbf{x}_n \mid \widehat{\mathbf{c}}_k, \widehat{\boldsymbol{\Gamma}}_k \right) \right] \\
&= -\frac{L}{2} \log(2\pi) - \frac{1}{2} \mathbb{E}_{q_{\boldsymbol{\Gamma}_k}} \left[\log \left| \widehat{\boldsymbol{\Gamma}}_k \right| \right] - \frac{1}{2} \mathbb{E}_{q_{\boldsymbol{\theta}_{z_n}^*}} \left[\left(\mathbf{x}_n - \widehat{\mathbf{c}}_k \right)^\top \widehat{\boldsymbol{\Gamma}}_k^{-1} \left(\mathbf{x}_n - \widehat{\mathbf{c}}_k \right) \right],
\end{aligned}$$

where we used the fact that

$$\begin{aligned}
\mathbb{E}_{q_{\boldsymbol{\Gamma}_k}} \left[\log \left| \widehat{\boldsymbol{\Gamma}}_k \right| \right] &= \log \left| \frac{\widehat{\boldsymbol{\Psi}}_k}{2} \right| - \sum_{l=1}^L \psi \left(\frac{\widehat{\nu}_k + (1-l)}{2} \right), \\
\mathbb{E}_{q_{\boldsymbol{\theta}_{z_n}^*}} \left[\left(\mathbf{x}_n - \widehat{\mathbf{c}}_k \right)^\top \widehat{\boldsymbol{\Gamma}}_k^{-1} \left(\mathbf{x}_n - \widehat{\mathbf{c}}_k \right) \right] &= \widehat{\nu}_k \left(\mathbf{x}_n - \widehat{\mathbf{m}}_k \right)^\top \widehat{\boldsymbol{\Psi}}_k^{-1} \left(\mathbf{x}_n - \widehat{\mathbf{m}}_k \right) + \frac{L}{\widehat{\lambda}_k}.
\end{aligned}$$

Plugging in all of the above expression back into (A.25) yields the desired results in Section 3.3.1.

A.3.1 Updating $\boldsymbol{\Sigma}_k$

By using matrix derivatives, the derivative of the log likelihood with respect to $\boldsymbol{\Sigma}_k^{-1}$ is given by

$$\begin{aligned}
&\frac{\partial}{\partial \boldsymbol{\Sigma}_k^{-1}} f_1 \left(\widehat{\mathbf{A}}_k, \mathbf{b}_k, \boldsymbol{\Sigma}_k^{-1} \right) \\
&= -\frac{1}{2} \sum_{n=1}^N q_{z_n}(k) \frac{\partial}{\partial \boldsymbol{\Sigma}_k^{-1}} \left[-\log \left| \boldsymbol{\Sigma}_k^{-1} \right| + \text{Tr} \left[\boldsymbol{\Sigma}_k^{-1 \top} \left(\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k \right) \left(\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k \right)^\top \right] \right] \\
&= -\frac{1}{2} \sum_{n=1}^N q_{z_n}(k) \left[-\boldsymbol{\Sigma}_k + \left(\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k \right) \left(\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k \right)^\top \right] \\
&= \frac{N_k}{2} \boldsymbol{\Sigma}_k - \frac{1}{2} \sum_{n=1}^N q_{z_n}(k) \left(\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k \right) \left(\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k \right)^\top.
\end{aligned} \tag{A.26}$$

Finally, setting to zero yields

$$\widehat{\boldsymbol{\Sigma}}_k = \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \left(\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k \right) \left(\mathbf{y}_n - \widehat{\mathbf{A}}_k \mathbf{x}_n - \widehat{\mathbf{b}}_k \right)^\top. \quad (\text{A.27})$$

A.4 VB-M-($\mathbf{c}, \mathbf{\Gamma}$) step from Section 3.2.5

This step divides into K sub-steps that involve the following optimizations

$$\left(\widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k \right) \equiv \left(\widehat{\mathbf{c}}_k^{(r)}, \widehat{\mathbf{\Gamma}}_k^{(r)} \right) = \underset{(\mathbf{c}_k, \mathbf{\Gamma}_k)}{\operatorname{argmax}} \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [\log p(\mathcal{X} | \mathcal{Z}; \mathbf{c}_k, \mathbf{\Gamma}_k)]. \quad (\text{A.28})$$

By definition, we have

$$\begin{aligned} & \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [\log (p(\mathcal{X} | \mathcal{Z}; \mathbf{c}_k, \mathbf{\Gamma}_k))] \\ &= \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} \left[\log \prod_{n=1}^N \mathcal{N}_L(\mathbf{x}_n | \mathbf{c}_k, \mathbf{\Gamma}_k)^{z_{nk}} \right] \\ &= \sum_{n=1}^N \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [z_{nk} \log \mathcal{N}_L(\mathbf{x}_n | \mathbf{c}_k, \mathbf{\Gamma}_k)] \\ &= \sum_{n=1}^N \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [z_{nk}] \log \mathcal{N}_L(\mathbf{x}_n | \mathbf{c}_k, \mathbf{\Gamma}_k) \\ &= \sum_{n=1}^N q_{z_n}(k) \log \mathcal{N}_L(\mathbf{x}_n | \mathbf{c}_k, \mathbf{\Gamma}_k) \\ &= \sum_{n=1}^N q_{z_n}(k) \left[-\frac{L}{2} \log(2\pi) - \frac{1}{2} \log |\mathbf{\Gamma}_k| - \frac{1}{2} (\mathbf{x}_n - \mathbf{c}_k)^\top \mathbf{\Gamma}_k^{-1} (\mathbf{x}_n - \mathbf{c}_k) \right] \\ &\equiv f_2(\mathbf{c}_k, \mathbf{\Gamma}_k). \end{aligned} \quad (\text{A.29})$$

We aim to solve the following optimization

$$\left(\widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k \right) = \underset{(\mathbf{c}_k, \mathbf{\Gamma}_k)}{\operatorname{argmax}} f_2(\mathbf{c}_k, \mathbf{\Gamma}_k). \quad (\text{A.30})$$

Similarly with Appendices A.3.1 and A.5.1, we obtain the following update:

$$\begin{aligned} \widehat{\mathbf{c}}_k &= \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n, \\ \widehat{\mathbf{\Gamma}}_k &= \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) (\mathbf{x}_n - \widehat{\mathbf{c}}_k) (\mathbf{x}_n - \widehat{\mathbf{c}}_k)^\top. \end{aligned} \quad (\text{A.31})$$

A.5 VB-M-($\mathbf{A}, \mathbf{b}, \Sigma$) step from Section 3.2.6

By definition, we have

$$\begin{aligned}
& \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}; \mathbf{A}_k, \mathbf{b}_k, \Sigma_k)] \\
&= \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} \left[\log \prod_{n=1}^N \mathcal{N}_D(\mathbf{y}_n \mid \mathbf{A}_k \mathbf{x}_n + \mathbf{b}_k, \Sigma_k)^{z_{nk}} \right] \\
&= \sum_{n=1}^N \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [z_{nk} \log \mathcal{N}_D(\mathbf{y}_n \mid \mathbf{A}_k \mathbf{x}_n + \mathbf{b}_k, \Sigma_k)] \\
&= \sum_{n=1}^N \mathbb{E}_{q_{\mathcal{Z}}^{(r)}} [z_{nk}] \log \mathcal{N}_D(\mathbf{y}_n \mid \mathbf{A}_k \mathbf{x}_n + \mathbf{b}_k, \Sigma_k) \\
&= \sum_{n=1}^N q_{z_n}(k) \log \mathcal{N}_D(\mathbf{y}_n \mid \mathbf{A}_k \mathbf{x}_n + \mathbf{b}_k, \Sigma_k) \\
&= \sum_{n=1}^N q_{z_n}(k) \left[-\frac{D}{2} \log(2\pi) - \frac{1}{2} \log |\Sigma_k| - \frac{1}{2} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k)^\top \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \right] \\
&\equiv f_1(\mathbf{A}_k, \mathbf{b}_k, \Sigma_k). \tag{A.32}
\end{aligned}$$

We aim to solve the following optimization

$$(\hat{\mathbf{A}}_k, \hat{\mathbf{b}}_k, \hat{\Sigma}_k) = \underset{(\mathbf{A}_k, \mathbf{b}_k, \Sigma_k)}{\operatorname{argmax}} f_1(\mathbf{A}_k, \mathbf{b}_k, \Sigma_k). \tag{A.33}$$

A.5.1 Updating \mathbf{b}_k

The derivative of the log likelihood with respect to \mathbf{b}_k is given by

$$\begin{aligned}
\frac{\partial}{\partial \mathbf{b}_k} f_1(\mathbf{A}_k, \mathbf{b}_k, \Sigma_k) &= -\frac{1}{2} \sum_{n=1}^N q_{z_n}(k) \frac{\partial}{\partial \mathbf{b}_k} \left[(\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k)^\top \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \right] \\
&= -\sum_{n=1}^N q_{z_n}(k) \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \frac{\partial}{\partial \mathbf{b}_k} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \\
&= \sum_{n=1}^N q_{z_n}(k) \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k).
\end{aligned}$$

Setting this derivative to zero, we obtain the solution for VB-M- \mathbf{b} step given by

$$\begin{aligned}
& \sum_{n=1}^N q_{z_n}(k) \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) = 0 \\
&\Leftrightarrow \sum_{n=1}^N q_{z_n}(k) \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n) - \sum_{n=1}^N q_{z_n}(k) \Sigma_k^{-1} \mathbf{b}_k = 0 \\
&\Leftrightarrow \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n) - \sum_{n=1}^N q_{z_n}(k) \mathbf{b}_k = 0 \text{ (left multiplying by } \Sigma_k) \\
&\Leftrightarrow \mathbf{b}_k = \frac{1}{\sum_{n=1}^N q_{z_n}(k)} \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n) \equiv \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n). \tag{A.34}
\end{aligned}$$

A.5.2 Updating \mathbf{A}_k

The derivative of the log likelihood with respect to \mathbf{A}_k is given by

$$\begin{aligned} \frac{\partial}{\partial \mathbf{A}_k} f_1(\mathbf{A}_k, \mathbf{b}_k, \Sigma_k) &= -\frac{1}{2} \sum_{n=1}^N q_{z_n}(k) \frac{\partial}{\partial \mathbf{A}_k} \left[(\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k)^\top \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \right] \\ &= -\sum_{n=1}^N q_{z_n}(k) \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \frac{\partial}{\partial \mathbf{A}_k} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \\ &= \sum_{n=1}^N q_{z_n}(k) \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \mathbf{x}_n^\top. \end{aligned}$$

Then, we set this derivative w.r.t. \mathbf{A}_k equal to zero, giving

$$\begin{aligned} \sum_{n=1}^N q_{z_n}(k) \Sigma_k^{-1} (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \mathbf{x}_n^\top &= 0 \\ \Leftrightarrow \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n - \mathbf{b}_k) \mathbf{x}_n^\top &= 0 \text{ (left multiplying by } \Sigma_k) \\ \Leftrightarrow \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \mathbf{x}_n^\top - \sum_{n=1}^N q_{z_n}(k) \mathbf{A}_k \mathbf{x}_n \mathbf{x}_n^\top - \sum_{n=1}^N q_{z_n}(k) \mathbf{b}_k \mathbf{x}_n^\top &= 0 \\ \Leftrightarrow \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \mathbf{x}_n^\top - \mathbf{A}_k \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n \mathbf{x}_n^\top - \mathbf{b}_k \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n^\top &= 0 \\ \Leftrightarrow \mathbf{A}_k \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n \mathbf{x}_n^\top &= \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \mathbf{x}_n^\top - \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \mathbf{A}_k \mathbf{x}_n) \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n^\top \text{ (using (A.34) for } \mathbf{b}_k) \\ \Leftrightarrow \mathbf{A}_k \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n \mathbf{x}_n^\top &= \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \mathbf{x}_n^\top - \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n^\top + \mathbf{A}_k \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n^\top \\ \Leftrightarrow \mathbf{A}_k \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n \left(\mathbf{x}_n^\top - \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n^\top \right) &= \sum_{n=1}^N q_{z_n}(k) \left(\mathbf{y}_n - \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \right) \mathbf{x}_n^\top \\ \Leftrightarrow N_k \mathbf{A}_k \mathbf{X}_k \mathbf{X}_k^\top &= N_k \mathbf{Y}_k \mathbf{X}_k^\top \Leftrightarrow \mathbf{A}_k = \mathbf{Y}_k \mathbf{X}_k^\top \left(\mathbf{X}_k \mathbf{X}_k^\top \right)^{-1}. \end{aligned}$$

Here, the last equality is obtained by firstly define the following quantities,

$$\begin{aligned} \bar{\mathbf{x}}_k &= \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n, \\ \bar{\mathbf{y}}_k &= \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n, \\ \mathbf{X}_k &= \frac{1}{\sqrt{N_k}} \left(\sqrt{q_{z_1}(k)} (\mathbf{x}_1 - \bar{\mathbf{x}}_k), \dots, \sqrt{q_{z_N}(k)} (\mathbf{x}_N - \bar{\mathbf{x}}_k) \right)_{L \times N}, \\ \mathbf{Y}_k &= \frac{1}{\sqrt{N_k}} \left(\sqrt{q_{z_1}(k)} (\mathbf{y}_1 - \bar{\mathbf{y}}_k), \dots, \sqrt{q_{z_N}(k)} (\mathbf{y}_N - \bar{\mathbf{y}}_k) \right)_{D \times N}. \end{aligned}$$

Then, we used the fact that

$$\begin{aligned}
\sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n \left(\mathbf{x}_n^\top - \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n^\top \right) &= \sum_{n=1}^N q_{z_n}(k) \mathbf{x}_n (\mathbf{x}_n^\top - \bar{\mathbf{x}}_k^\top) \\
&= \sum_{n=1}^N q_{z_n}(k) (\mathbf{x}_n - \bar{\mathbf{x}}_k) (\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top + \sum_{n=1}^N q_{z_n}(k) \bar{\mathbf{x}}_k (\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top \\
&= \sum_{n=1}^N q_{z_n}(k) (\mathbf{x}_n - \bar{\mathbf{x}}_k) (\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top \\
&= N_k \mathbf{X}_k \mathbf{X}_k^\top, \tag{A.35}
\end{aligned}$$

and

$$\begin{aligned}
\sum_{n=1}^N q_{z_n}(k) \left(\mathbf{y}_n - \frac{1}{N_k} \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \right) \mathbf{x}_n^\top &= \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \bar{\mathbf{y}}_k) \mathbf{x}_n^\top \\
&= \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \bar{\mathbf{y}}_k) (\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top + \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \bar{\mathbf{y}}_k) \bar{\mathbf{x}}_k^\top \\
&= \sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \bar{\mathbf{y}}_k) (\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top \\
&= N_k \mathbf{Y}_k \mathbf{X}_k^\top. \tag{A.36}
\end{aligned}$$

Here, we used the fact that

$$\begin{aligned}
\sum_{n=1}^N q_{z_n}(k) \bar{\mathbf{x}}_k (\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top &= \sum_{n=1}^N q_{z_n}(k) \bar{\mathbf{x}}_k \mathbf{x}_n^\top - \bar{\mathbf{x}}_k N_k \bar{\mathbf{x}}_k^\top = \sum_{n=1}^N q_{z_n}(k) \bar{\mathbf{x}}_k \mathbf{x}_n^\top - \sum_{n=1}^N q_{z_n}(k) \bar{\mathbf{x}}_k \mathbf{x}_n^\top = 0, \\
\sum_{n=1}^N q_{z_n}(k) (\mathbf{y}_n - \bar{\mathbf{y}}_k) \bar{\mathbf{x}}_k^\top &= \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \bar{\mathbf{x}}_k^\top - N_k \bar{\mathbf{y}}_k \bar{\mathbf{x}}_k^\top = \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \bar{\mathbf{x}}_k^\top - \sum_{n=1}^N q_{z_n}(k) \mathbf{y}_n \bar{\mathbf{x}}_k^\top = 0, \tag{A.37}
\end{aligned}$$

and for each $i, j \in [L]$, it holds that

$$\begin{aligned}
\left[\sum_{n=1}^N q_{z_n}(k) (\mathbf{x}_n - \bar{\mathbf{x}}_k) (\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top \right]_{ij} &= \sum_{n=1}^N q_{z_n}(k) \left[(\mathbf{x}_n - \bar{\mathbf{x}}_k) (\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top \right]_{ij} \\
&= \sum_{n=1}^N q_{z_n}(k) [(\mathbf{x}_n - \bar{\mathbf{x}}_k)]_{i1} [(\mathbf{x}_n - \bar{\mathbf{x}}_k)^\top]_{1j} \\
&\equiv \sum_{n=1}^N q_{z_n}(k) [(\mathbf{x}_n - \bar{\mathbf{x}}_k)]_i [(\mathbf{x}_n - \bar{\mathbf{x}}_k)]_j \\
&= \left[\left(\sqrt{q_{z_1}(k)} (\mathbf{x}_1 - \bar{\mathbf{x}}_k), \dots, \sqrt{q_{z_N}(k)} (\mathbf{x}_N - \bar{\mathbf{x}}_k) \right) \right]_i \cdot \left[\left(\sqrt{q_{z_1}(k)} (\mathbf{x}_1 - \bar{\mathbf{x}}_k), \dots, \sqrt{q_{z_N}(k)} (\mathbf{x}_N - \bar{\mathbf{x}}_k) \right) \right]_j^\top \\
&= N_k \left[\mathbf{X}_k \mathbf{X}_k^\top \right]_{ij}. \tag{A.38}
\end{aligned}$$

Appendix B Technical proofs

B.1 Proof of Lemma 6.2

For the sake of completeness, we highlighted the main techniques using for the proof of Lemma 6.2 in this paper, see [1, Proposition 1] for more details. We first want to prove (6.6). Using the partition of

a joint Gaussian with $\mathbf{x}_b \equiv \mathbf{x}$, $\boldsymbol{\mu}_b = \boldsymbol{\mu}_k^{\mathbf{x}}$, $\boldsymbol{\Sigma}_{bb} \equiv \mathbf{V}_k^{\mathbf{xx}}$, $\mathbf{x}_a \equiv \mathbf{y}$, $\boldsymbol{\mu}_a \equiv \boldsymbol{\mu}_k^{\mathbf{y}}$, $\boldsymbol{\Sigma}_{aa} \equiv \mathbf{V}_k^{\mathbf{yy}}$, we obtain

$$\begin{aligned} p(\mathbf{x}_a | \mathbf{x}_b) &= \mathcal{N}(\mathbf{x}_a | \boldsymbol{\mu}_{a|b}, \boldsymbol{\Gamma}_{aa}^{-1}), \quad \boldsymbol{\mu}_{a|b} = \boldsymbol{\mu}_a - \boldsymbol{\Gamma}_{aa}^{-1} \boldsymbol{\Gamma}_{ab} (\mathbf{x}_b - \boldsymbol{\mu}_b) = \boldsymbol{\mu}_a - \boldsymbol{\Sigma}_{ab} \boldsymbol{\Sigma}_{bb}^{-1} (\mathbf{x}_b - \boldsymbol{\mu}_b), \\ p(\mathbf{x}_b) &= \mathcal{N}(\mathbf{x}_b | \boldsymbol{\mu}_b, \boldsymbol{\Sigma}_{bb}). \end{aligned} \quad (\text{B.1})$$

Recall that

$$\begin{aligned} p(\mathbf{y} | \mathbf{x}, Z = k; \boldsymbol{\psi}) &= \mathcal{N}_D(\mathbf{y} | \mathbf{A}_k \mathbf{x} + \mathbf{b}_k, \boldsymbol{\Sigma}_k), \\ p(\mathbf{x} | Z = k; \boldsymbol{\psi}) &= \mathcal{N}_L(\mathbf{x} | \mathbf{c}_k, \boldsymbol{\Gamma}_k), \quad p(Z = k; \boldsymbol{\psi}) = \pi_k. \end{aligned} \quad (\text{B.2})$$

By identifying the parameters of (B.1) and (B.2), it holds that

$$\begin{aligned} \pi_k &= \rho_k \\ \mathbf{c}_k &= \boldsymbol{\mu}_k^{\mathbf{x}}, \\ \boldsymbol{\Gamma}_k &= \mathbf{V}_k^{\mathbf{xx}}, \\ \mathbf{A}_k &= -\boldsymbol{\Gamma}_{aa}^{-1} \boldsymbol{\Gamma}_{ab} = \mathbf{V}_k^{\mathbf{xy}^\top} (\mathbf{V}_k^{\mathbf{xx}})^{-1}, \\ \mathbf{b}_k &= \boldsymbol{\mu}_a + \boldsymbol{\Gamma}_{aa}^{-1} \boldsymbol{\Gamma}_{ab} \boldsymbol{\mu}_b = \boldsymbol{\mu}_k^{\mathbf{y}} - \mathbf{V}_k^{\mathbf{xy}^\top} (\mathbf{V}_k^{\mathbf{xx}})^{-1} \boldsymbol{\mu}_k^{\mathbf{x}}, \\ \boldsymbol{\Sigma}_k &= \boldsymbol{\Gamma}_{aa}^{-1} = \boldsymbol{\Sigma}_{aa} - \boldsymbol{\Sigma}_{ab} \boldsymbol{\Sigma}_{bb}^{-1} \boldsymbol{\Sigma}_{ba} = \mathbf{V}_k^{\mathbf{yy}} - \mathbf{V}_k^{\mathbf{xy}^\top} (\mathbf{V}_k^{\mathbf{xx}})^{-1} \mathbf{V}_k^{\mathbf{xy}}. \end{aligned} \quad (\text{B.3})$$

$$\quad (\text{B.4})$$

The following decomposition of the joint probability distribution will be used:

$$\begin{aligned} p(\mathbf{w} | \boldsymbol{\psi}) &= \sum_{k=1}^K p(\mathbf{y} | \mathbf{x}, Z = k; \boldsymbol{\psi}) p(\mathbf{X} = \mathbf{x} | Z = k; \boldsymbol{\psi}) p(Z = k; \boldsymbol{\psi}) \\ &= \sum_{k=1}^K \pi_k \mathcal{N}_D(\mathbf{y} | \mathbf{A}_k \mathbf{x} + \mathbf{b}_k, \boldsymbol{\Sigma}_k) \mathcal{N}_L(\mathbf{x} | \mathbf{c}_k, \boldsymbol{\Gamma}_k) \\ &\equiv \sum_{k=1}^K \rho_k \mathcal{N}_{L+D}(\mathbf{w} | \boldsymbol{\mu}_k, \mathbf{V}_k). \end{aligned} \quad (\text{B.5})$$

By using result for the joint Gaussian, see *e.g.*, (B.11), we obtain the desired result (6.7).

Finally, Lemma 6.2 is proved via using the following two statements [1, Lemmas 1 and 2]:

- (i) For any $\rho_k \in \mathbb{R}$, $\boldsymbol{\mu}_k \in \mathbb{R}^{L+D}$, and $\mathbf{V}_k \in \mathcal{S}_+^{L+D}$, there is a set of parameters $\mathbf{c}_k \in \mathbb{R}^L$, $\boldsymbol{\Lambda}_k \in \mathcal{S}_+^L$, $\pi_k \in \mathbb{R}$, $\mathbf{A}_k \in \mathbb{R}^{D \times L}$, $\mathbf{b}_k \in \mathbb{R}^D$, $\boldsymbol{\Sigma}_k \in \mathcal{S}_+^D$ such that (6.6) holds.
- (ii) Reciprocally, for any $\mathbf{c}_k \in \mathbb{R}^L$, $\boldsymbol{\Lambda}_k \in \mathcal{S}_+^L$, $\pi_k \in \mathbb{R}$, $\mathbf{A}_k \in \mathbb{R}^{D \times L}$, $\mathbf{b}_k \in \mathbb{R}^D$, $\boldsymbol{\Sigma}_k \in \mathcal{S}_+^D$, there is a set of parameters $\rho_k \in \mathbb{R}$, $\boldsymbol{\mu}_k \in \mathbb{R}^{L+D}$ and $\mathbf{V}_k \in \mathcal{S}_+^{L+D}$ such that (6.7) holds.

B.2 Proof of Theorem 4.1

Recall that we defined $\boldsymbol{\Theta} = (\boldsymbol{\tau}, \alpha, \sigma)$. Then,

$$\begin{aligned} p(\hat{\mathbf{y}}, \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) &= \sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \hat{\mathbf{z}}, \boldsymbol{\Theta}, \mathcal{X}, \mathcal{Y}) p(\hat{\mathbf{x}} | \hat{\mathbf{z}}, \boldsymbol{\Theta}, \mathcal{X}, \mathcal{Y}) p(\hat{\mathbf{z}} | \boldsymbol{\Theta}, \mathcal{X}, \mathcal{Y}) p(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\Theta} \\ &= \sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \hat{\mathbf{z}}; \mathbf{A}, \mathbf{b}, \boldsymbol{\Sigma}) p(\hat{\mathbf{x}} | \hat{\mathbf{z}}, \mathbf{c}, \boldsymbol{\Gamma}) p(\hat{\mathbf{z}} | \boldsymbol{\tau}; \beta) p(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\Theta} \equiv D_1. \end{aligned} \quad (\text{B.6})$$

Note that in (B.6), $p(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y})$ is in fact the (unknown) true posterior distribution of the parameters given a sample $(\mathcal{X}, \mathcal{Y})$. Because the integrations w.r.t. true posterior distribution are intractable, we approximate the predictive conditional density by replacing the true posterior distribution $p(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y})$ with its truncated variational posterior of parameters $\boldsymbol{\Theta}$ given by

$$q_{\boldsymbol{\Theta}}(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) = q_{\alpha, \sigma}(\alpha, \sigma | \mathcal{X}, \mathcal{Y}) \prod_{k=1}^{K-1} q_{\tau_k}(\tau_k | \mathcal{X}, \mathcal{Y}). \quad (\text{B.7})$$

Recall that the infinite state space for each z_j is dealt with by choosing a truncation of the state space to a maximum label K [60]. In practice, this consists of assuming that the variational distributions q_{z_n} for $n \in [N]$, satisfy $q_{z_n}(k) = 0$ for $k > K$ and that the variational distribution on τ also factorizes as $q_\tau(\boldsymbol{\tau}) = \prod_{k=1}^{K-1} q_{\tau_k}(\tau_k)$ with the additional condition that $\tau_K = 1$. In particular, here we choose $K = \hat{K}$ where \hat{K} is estimated from some suitable procedures.

For simplicity, we consider the case when $\beta = 0, \sigma = 0$. Then we have

$$\begin{aligned}
 D_1 &\approx \sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \hat{\mathbf{z}}; \hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{\boldsymbol{\Sigma}}) p(\hat{\mathbf{x}} | \hat{\mathbf{z}}, \hat{\mathbf{c}}, \hat{\boldsymbol{\Gamma}}) p(\hat{\mathbf{z}} | \boldsymbol{\tau}) q_\Theta(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\Theta} \\
 &= \sum_{k=1}^{\infty} \int \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k) \pi_k(\boldsymbol{\tau}) q_\Theta(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\Theta} \\
 &\approx \sum_{k=1}^K \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k) \int \pi_k(\boldsymbol{\tau}) q_\Theta(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\Theta} \\
 &= \sum_{k=1}^K \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k) \int \pi_k(\boldsymbol{\tau}) q_\tau(\boldsymbol{\tau} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\tau} \underbrace{\int q_{\alpha,0}(\alpha | \mathcal{X}, \mathcal{Y}) d\alpha}_{=1} \\
 &= \sum_{k=1}^K \mathbb{E}_{q_\tau}[\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k) \\
 &\equiv \sum_{k=1}^K \mathbb{E}_{q_\tau}[\pi_k(\boldsymbol{\tau})] \mathcal{N}_{L+D}(\hat{\mathbf{w}} | \mathbb{E}[\mathbf{w}], \text{cov}[\mathbf{w}]) \\
 &= \sum_{k=1}^K \mathbb{E}_{q_{\tau_k}}[\tau_k] \prod_{l=1}^{k-1} \mathbb{E}_{q_{\tau_l}}[1 - \tau_l] \mathcal{N}_{L+D}(\hat{\mathbf{w}} | \mathbb{E}[\mathbf{w}], \text{cov}[\mathbf{w}]).
 \end{aligned}$$

Here, by defining $\hat{\mathbf{w}} \equiv [\hat{\mathbf{x}}; \hat{\mathbf{y}}]$, we used the fact that,

$$\mathbb{E}[\mathbf{w}] = \begin{pmatrix} \hat{\mathbf{c}}_k \\ \hat{\mathbf{A}}_k \hat{\mathbf{c}}_k + \hat{\mathbf{b}}_k \end{pmatrix}, \quad \text{cov}[\mathbf{w}] = \begin{pmatrix} \hat{\boldsymbol{\Gamma}}_k & \hat{\boldsymbol{\Gamma}}_k \hat{\mathbf{A}}_k^\top \\ \hat{\mathbf{A}}_k \hat{\boldsymbol{\Gamma}}_k & \hat{\boldsymbol{\Sigma}}_k + \hat{\mathbf{A}}_k \hat{\boldsymbol{\Gamma}}_k \hat{\mathbf{A}}_k^\top \end{pmatrix}. \quad (\text{B.8})$$

Indeed, we made use of the following result for the joint Gaussian, see, *e.g.*, [67, Eq. (2.115), page 93]. Given a marginal Gaussian distribution for \mathbf{x} and a conditional Gaussian distribution for \mathbf{y} given \mathbf{x} in the form

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \boldsymbol{\Gamma}^{-1}), \quad (\text{B.9})$$

$$p(\mathbf{y} | \mathbf{x}) = \mathcal{N}(\mathbf{y} | \mathbf{A}\mathbf{x} + \mathbf{b}, \mathbf{L}^{-1}), \quad (\text{B.10})$$

then the joint distribution of $\mathbf{w} \equiv [\mathbf{x}; \mathbf{y}]$ is given by

$$\begin{aligned}
 p(\mathbf{w}) &= \mathcal{N}(\mathbf{w} | \mathbb{E}[\mathbf{w}], \text{cov}[\mathbf{w}]), \text{ where} \\
 \text{cov}[\mathbf{w}] &= \begin{pmatrix} \boldsymbol{\Gamma}^{-1} & \boldsymbol{\Gamma}^{-1} \mathbf{A}^\top \\ \mathbf{A} \boldsymbol{\Gamma}^{-1} & \mathbf{L}^{-1} + \mathbf{A} \boldsymbol{\Gamma}^{-1} \mathbf{A}^\top \end{pmatrix}, \quad \mathbb{E}[\mathbf{w}] = \begin{pmatrix} \boldsymbol{\mu} \\ \mathbf{A}\boldsymbol{\mu} + \mathbf{b} \end{pmatrix}. \quad (\text{B.11})
 \end{aligned}$$

In our situation, the desired result is obtained via using $\mathbf{y} \equiv \hat{\mathbf{y}}, \mathbf{A} \equiv \hat{\mathbf{A}}_k, \mathbf{b} \equiv \hat{\mathbf{b}}_k, \mathbf{L}^{-1} = \hat{\boldsymbol{\Sigma}}_k, \mathbf{x} \equiv \hat{\mathbf{x}}, \boldsymbol{\mu} \equiv \hat{\mathbf{c}}_k, \boldsymbol{\Gamma}^{-1} \equiv \hat{\boldsymbol{\Gamma}}_k$.

Furthermore, we also used the fact that

$$\begin{aligned}
 \mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})] &= \int \tau_k q_{\tau_k}(\tau_k | \mathcal{X}, \mathcal{Y}) d\tau_k \int \prod_{l=1}^{k-1} (1 - \tau_l) \prod_{j=1, j \neq k}^{K-1} q_{\tau_j}(\tau_j | \mathcal{X}, \mathcal{Y}) \prod_{j=1, j \neq k}^K d\tau_j \\
 &= \mathbb{E}_{q_{\tau_k}} [\tau_k] \int \prod_{l=1}^{k-1} (1 - \tau_l) \prod_{j=1}^{k-1} q_{\tau_j}(\tau_j | \mathcal{X}, \mathcal{Y}) \underbrace{\int \prod_{j=k+1}^{K-1} q_{\tau_j}(\tau_j | \mathcal{X}, \mathcal{Y}) \prod_{j=k+1}^{K-1} d\tau_j \prod_{j=1}^{k-1} d\tau_j}_{=1} \\
 &= \mathbb{E}_{q_{\tau_k}} [\tau_k] \prod_{l=1}^{k-1} \int (1 - \tau_l) q_{\tau_l}(\tau_l | \mathcal{X}, \mathcal{Y}) d\tau_l \\
 &= \mathbb{E}_{q_{\tau_k}} [\tau_k] \prod_{l=1}^{k-1} \mathbb{E}_{q_{\tau_l}} [1 - \tau_l]. \tag{B.12}
 \end{aligned}$$

Next, we aim to prove that

$$\sum_{k=1}^K \mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})] = 1. \tag{B.13}$$

Indeed, recall that we have defined

$$\tau_k | \alpha, \sigma \stackrel{\text{ind}}{\sim} \text{Beta}(\tau_k | 1 - \sigma, \alpha + k\sigma), \quad k \in \mathbb{N}^*, \tag{B.14}$$

$$\pi_k(\boldsymbol{\tau}) = \tau_k \prod_{l=1}^{k-1} (1 - \tau_l), \quad k \in \mathbb{N}^*, \tag{B.15}$$

$$p(\mathcal{Z} | \boldsymbol{\tau}; \beta) \propto \left(\prod_{n=1}^N \pi_{z_n}(\boldsymbol{\tau}) \right) \exp\left(\beta \sum_{m \sim n} \delta_{(z_m = z_n)} \right), \tag{B.16}$$

and to deal with the infinite state space for each z_j , we considered a truncation of the state space to a maximum label $K \equiv K_{\max}$, $K_{\max} \in \mathbb{N}^*$ [60]. In practice, this consists of assuming that the variational distributions q_{z_n} for $n \in [N]$, satisfy $q_{z_n}(k) = 0$ for $k > K$ and that the variational distribution on $\boldsymbol{\tau}$ also factorizes as $q_{\boldsymbol{\tau}}(\boldsymbol{\tau}) = \prod_{k=1}^{K-1} q_{\tau_k}(\tau_k)$ with the additional condition that $\tau_K = 1$. Based on the proof from [42, Lemma 3.4], it holds that a necessary and sufficient condition to guarantee that these π_k 's sum to 1 almost surely, *i.e.*,

$$\sum_{k=1}^{\infty} \pi_k(\boldsymbol{\tau}) = \sum_{k=1}^{\infty} \tau_k \prod_{l=1}^{k-1} (1 - \tau_l) = 1, \tag{B.17}$$

is that the expectation $\mathbb{E} \left[\prod_{l=1}^{k-1} (1 - \tau_l) \right]$ tends to 0 as k tends to ∞ . In particular, if τ_1, τ_2, \dots are i.i.d., *e.g.*, when $\sigma = 0$, it suffices that $p(\tau_1 > 0) > 0$. Then

$$1 = \mathbb{E}_{q_{\boldsymbol{\tau}}} \left[\sum_{k=1}^{\infty} \pi_k(\boldsymbol{\tau}) \right] = \sum_{k=1}^{\infty} \mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})] = \sum_{k=1}^K \mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})].$$

B.3 Proof of Theorem 4.2

From the product rule of probability, we see that this conditional distribution can be evaluated from the joint and marginal distributions. Furthermore, by integrating out $\hat{\mathbf{z}}$ and $\boldsymbol{\Theta}$, the predictive conditional density is then given by

$$p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) = \frac{p(\hat{\mathbf{y}}, \hat{\mathbf{x}} | \mathcal{X}, \mathcal{Y})}{p(\hat{\mathbf{x}} | \mathcal{X}, \mathcal{Y})} = \frac{\sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{y}}, \hat{\mathbf{x}}, \hat{\mathbf{z}}, \boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\hat{\mathbf{z}} d\boldsymbol{\Theta}}{\sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{x}}, \hat{\mathbf{z}}, \boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\hat{\mathbf{z}} d\boldsymbol{\Theta}} \equiv \frac{D_1}{D_2}. \tag{B.18}$$

Next, with a similar step as in the proof of Theorem 4.1, we also obtain

$$D_2 \approx \sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k). \quad (\text{B.19})$$

Therefore, we obtain

$$\begin{aligned} p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) &\approx \frac{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k)}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k)} \\ &= \sum_{k=1}^K \frac{\mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k)}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k)} \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k) \\ &\equiv \sum_{k=1}^K g_k(\hat{\mathbf{x}} | \hat{\boldsymbol{\Theta}}, \hat{\boldsymbol{\phi}}, \mathcal{X}, \mathcal{Y}) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k), \end{aligned} \quad (\text{B.20})$$

which is a mixture of Gaussian experts since we have

$$g_k(\hat{\mathbf{x}} | \hat{\boldsymbol{\Theta}}, \hat{\boldsymbol{\phi}}, \mathcal{X}, \mathcal{Y}) = \frac{\mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k)}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k)}, \quad k \in [K], \quad (\text{B.21})$$

belongs to a $K - 1$ dimensional probability simplex.

B.4 Proof of Theorem 4.3

To deal with high-dimensional regression data, namely high-to-low regression, given the inverse conditional density $p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y})$, we want to compute the following forward conditional density

$$p(\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) = \frac{p(\hat{\mathbf{x}}, \hat{\mathbf{y}} | \mathcal{X}, \mathcal{Y})}{p(\hat{\mathbf{y}} | \mathcal{X}, \mathcal{Y})} = \frac{p(\hat{\mathbf{x}}, \hat{\mathbf{y}} | \mathcal{X}, \mathcal{Y})}{\int_{\hat{\mathbf{x}}} p(\hat{\mathbf{x}}, \hat{\mathbf{y}} | \mathcal{X}, \mathcal{Y}) d\hat{\mathbf{x}}} = \frac{D_1}{\int_{\hat{\mathbf{x}}} D_1(\hat{\mathbf{x}}) d\hat{\mathbf{x}}} \equiv \frac{D_1}{D_3}. \quad (\text{B.22})$$

Then, we have to compute or approximate D_3 . Using Theorem 4.1, we obtain

$$\begin{aligned} D_3 &\approx \sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \int \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k) d\hat{\mathbf{x}} \\ &= \sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{c}}_k + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k + \hat{\mathbf{A}}_k \hat{\boldsymbol{\Gamma}}_k \hat{\mathbf{A}}_k^\top). \end{aligned}$$

Indeed, we made use of the following results for marginal and conditional Gaussians, see, *e.g.*, [67, Eq. (2.115), page 93]. Given a marginal Gaussian distribution for \mathbf{x} and a conditional Gaussian distribution for \mathbf{y} given \mathbf{x} in the form

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \boldsymbol{\Gamma}^{-1}), \quad (\text{B.23})$$

$$p(\mathbf{y} | \mathbf{x}) = \mathcal{N}(\mathbf{y} | \mathbf{A}\mathbf{x} + \mathbf{b}, \mathbf{L}^{-1}), \quad (\text{B.24})$$

then the marginal distribution of \mathbf{y} and the conditional distribution of \mathbf{x} given \mathbf{y} are given by

$$p(\mathbf{y}) = \int p(\mathbf{y} | \mathbf{x}) p(\mathbf{x}) d\mathbf{x} = \mathcal{N}(\mathbf{y} | \mathbf{A}\boldsymbol{\mu} + \mathbf{b}, \mathbf{L}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top), \quad (\text{B.25})$$

$$p(\mathbf{x} | \mathbf{y}) = \mathcal{N}(\mathbf{x} | \boldsymbol{\Sigma} [\mathbf{A}^\top \mathbf{L}(\mathbf{y} - \mathbf{b}) + \boldsymbol{\Gamma}\boldsymbol{\mu}], \boldsymbol{\Sigma}), \quad \boldsymbol{\Sigma} = (\boldsymbol{\Gamma} + \mathbf{A}^\top \mathbf{L} \mathbf{A})^{-1}. \quad (\text{B.26})$$

In our situation, the desired result is obtained via using $\mathbf{y} \equiv \hat{\mathbf{y}}, \mathbf{A} \equiv \hat{\mathbf{A}}_k, \mathbf{b} \equiv \hat{\mathbf{b}}_k, \mathbf{L}^{-1} = \hat{\mathbf{\Sigma}}_k, \mathbf{x} \equiv \hat{\mathbf{x}}, \boldsymbol{\mu} \equiv \hat{\mathbf{c}}_k, \boldsymbol{\Gamma}^{-1} \equiv \hat{\boldsymbol{\Gamma}}_k$.

Finally, we obtain

$$\begin{aligned}
 & p(\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) \\
 & \approx \sum_{k=1}^K \frac{\mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\mathbf{\Sigma}}_k) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k)}{\sum_{k=1}^K \mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{c}}_k + \hat{\mathbf{b}}_k, \hat{\mathbf{\Sigma}}_k + \hat{\mathbf{A}}_k \hat{\boldsymbol{\Gamma}}_k \hat{\mathbf{A}}_k^\top)} \\
 & = \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\boldsymbol{\Theta}}^*, \hat{\boldsymbol{\phi}}^*, \mathcal{X}, \mathcal{Y}) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*, \hat{\mathbf{\Sigma}}_k^*), \tag{B.27}
 \end{aligned}$$

where

$$g_k(\hat{\mathbf{y}} | \hat{\boldsymbol{\Theta}}^*, \hat{\boldsymbol{\phi}}^*, \mathcal{X}, \mathcal{Y}) = \frac{\mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_k^*, \hat{\boldsymbol{\Gamma}}_k^*)}{\sum_{k=1}^K \mathbb{E}_{q_{\boldsymbol{\tau}}} [\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_k^*, \hat{\boldsymbol{\Gamma}}_k^*)}. \tag{B.28}$$

Here, we used the fact that $p(\hat{\mathbf{y}}, \hat{\mathbf{x}} | \hat{z} = k) = p(\hat{\mathbf{x}} | \hat{\mathbf{y}}, \hat{z} = k) p(\hat{\mathbf{y}} | \hat{z} = k)$, namely,

$$\begin{aligned}
 & \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\mathbf{\Sigma}}_k) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{c}}_k, \hat{\boldsymbol{\Gamma}}_k) \\
 & = \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{\Sigma}}_k^* [\hat{\mathbf{A}}_k^\top \hat{\mathbf{\Sigma}}_k^{-1} (\hat{\mathbf{y}} - \hat{\mathbf{b}}_k) + \hat{\boldsymbol{\Gamma}}_k^{-1} \hat{\mathbf{c}}_k], \hat{\mathbf{\Sigma}}_k^*) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{c}}_k + \hat{\mathbf{b}}_k, \hat{\mathbf{\Sigma}}_k + \hat{\mathbf{A}}_k \hat{\boldsymbol{\Gamma}}_k \hat{\mathbf{A}}_k^\top) \\
 & = \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*, \hat{\mathbf{\Sigma}}_k^*) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_k^*, \hat{\boldsymbol{\Gamma}}_k^*),
 \end{aligned}$$

with

$$\begin{aligned}
 \hat{\mathbf{\Sigma}}_k^* &= (\hat{\boldsymbol{\Gamma}}_k^{-1} + \hat{\mathbf{A}}_k^\top \hat{\mathbf{\Sigma}}_k^{-1} \hat{\mathbf{A}}_k)^{-1}, \\
 \hat{\mathbf{A}}_k^* &= \hat{\mathbf{\Sigma}}_k^* \hat{\mathbf{A}}_k^\top \hat{\mathbf{\Sigma}}_k^{-1}, \\
 \hat{\mathbf{b}}_k^* &= \hat{\mathbf{\Sigma}}_k^* [\hat{\boldsymbol{\Gamma}}_k^{-1} \hat{\mathbf{c}}_k - \hat{\mathbf{A}}_k^\top \hat{\mathbf{\Sigma}}_k^{-1} \hat{\mathbf{b}}_k], \\
 \hat{\mathbf{c}}_k^* &= \hat{\mathbf{A}}_k \hat{\mathbf{c}}_k + \hat{\mathbf{b}}_k, \\
 \hat{\boldsymbol{\Gamma}}_k^* &= \hat{\mathbf{\Sigma}}_k + \hat{\mathbf{A}}_k \hat{\boldsymbol{\Gamma}}_k \hat{\mathbf{A}}_k^\top. \tag{B.29}
 \end{aligned}$$

When required, it is straightforward to approximate the expectation and covariance matrix of $\hat{\mathbf{x}} |$

$\hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}$ as follows:

$$\begin{aligned} \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] &\approx \int (\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*, \hat{\Sigma}_k^*) d\hat{\mathbf{x}} \\ &= \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \int (\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*, \hat{\Sigma}_k^*) d\hat{\mathbf{x}} \\ &= \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) (\hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*), \end{aligned}$$

$$\begin{aligned} \text{var}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] &= \mathbb{E}[(\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) (\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y})^\top] - \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}]^\top \\ &\approx \int (\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) (\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y})^\top \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*, \hat{\Sigma}_k^*) d\hat{\mathbf{x}} \\ &\quad - \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}]^\top \\ &\approx \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \int (\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) (\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y})^\top \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*, \hat{\Sigma}_k^*) d\hat{\mathbf{x}} \\ &\quad - \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}]^\top \\ &\approx \sum_{k=1}^K g_k(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) \left[\hat{\Sigma}_k^* + (\hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*) (\hat{\mathbf{A}}_k^* \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*)^\top \right] \\ &\quad - \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}] \mathbb{E}[\hat{\mathbf{x}} | \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}]^\top, \end{aligned}$$

where we used the following definitions

$$\text{cov}(\mathbf{X}, \mathbf{Y}) = \mathbb{E}(\mathbf{X}\mathbf{Y}^\top) - \mathbb{E}(\mathbf{X}) \mathbb{E}(\mathbf{Y})^\top, \quad \text{var}(\mathbf{X}) = \text{cov}(\mathbf{X}, \mathbf{X}).$$

B.5 Proof of Theorem 4.4

Recall that we defined $\Theta = (\tau, \alpha, \sigma, \theta^*)$, $\theta^* = (\theta_k^*)_{k \in \mathbb{N}^*} := (\mathbf{c}_k, \mathbf{\Gamma}_k)_{k \in \mathbb{N}^*}$. Then,

$$\begin{aligned} p(\hat{\mathbf{y}}, \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) &= \sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \hat{\mathbf{z}}, \Theta, \mathcal{X}, \mathcal{Y}) p(\hat{\mathbf{x}} | \hat{\mathbf{z}}, \Theta, \mathcal{X}, \mathcal{Y}) p(\hat{\mathbf{z}} | \Theta, \mathcal{X}, \mathcal{Y}) p(\Theta | \mathcal{X}, \mathcal{Y}) d\Theta \\ &= \sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{y}} | \hat{\mathbf{x}}, \hat{\mathbf{z}}; \mathbf{A}, \mathbf{b}, \Sigma) p(\hat{\mathbf{x}} | \hat{\mathbf{z}}, \mathbf{c}, \mathbf{\Gamma}) p(\hat{\mathbf{z}} | \tau; \beta) p(\Theta | \mathcal{X}, \mathcal{Y}) d\Theta \equiv T_1. \end{aligned} \quad (\text{B.30})$$

Note that in (B.30), $p(\Theta | \mathcal{X}, \mathcal{Y})$ is in fact the (unknown) true posterior distribution of the parameters given a sample $(\mathcal{X}, \mathcal{Y})$. Because the integrations w.r.t. true posterior distribution are intractable, we approximate the predictive conditional density by replacing the true posterior distribution $p(\Theta | \mathcal{X}, \mathcal{Y})$ with its truncated variational posterior of parameters Θ given by

$$q_{\Theta}(\Theta | \mathcal{X}, \mathcal{Y}) = q_{\alpha, \sigma}(\alpha, \sigma | \mathcal{X}, \mathcal{Y}) \prod_{k=1}^{K-1} q_{\tau_k}(\tau_k | \mathcal{X}, \mathcal{Y}) \prod_{k=1}^K q_{\theta_k^*}(\theta_k^* | \mathcal{X}, \mathcal{Y}). \quad (\text{B.31})$$

Recall that the infinite state space for each z_j is dealt with by choosing a truncation of the state space to a maximum label K [60]. In practice, this consists of assuming that the variational distributions q_{z_n} for $n \in [N]$, satisfy $q_{z_n}(k) = 0$ for $k > K$ and that the variational distribution on τ also factorizes as $q_{\tau}(\tau) = \prod_{k=1}^{K-1} q_{\tau_k}(\tau_k)$ with the additional condition that $\tau_K = 1$. In particular, here we choose $K = \hat{K}$ where \hat{K} is estimated from some suitable procedures.

For simplicity, we consider the case when $\beta = 0, \sigma = 0$. Then, we obtain

$$\begin{aligned}
 T_1 &\approx \sum_{\widehat{\mathbf{z}}} \int p(\widehat{\mathbf{y}} | \widehat{\mathbf{x}}, \widehat{\mathbf{z}}; \mathbf{A}, \mathbf{b}, \boldsymbol{\Sigma}) p(\widehat{\mathbf{x}} | \widehat{\mathbf{z}}, \mathbf{c}, \boldsymbol{\Gamma}) p(\widehat{\mathbf{z}} | \boldsymbol{\tau}) q_{\Theta}(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\Theta} \\
 &= \sum_{k=1}^{\infty} \int \mathcal{N}_D(\widehat{\mathbf{y}} | \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k) \pi_k(\boldsymbol{\tau}) q_{\Theta}(\boldsymbol{\Theta} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\Theta} \\
 &\approx \sum_{k=1}^K \int \mathcal{N}_D(\widehat{\mathbf{y}} | \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k) \int \pi_k(\boldsymbol{\tau}) q_{\boldsymbol{\tau}}(\boldsymbol{\tau} | \mathcal{X}, \mathcal{Y}) d\boldsymbol{\tau} \\
 &\quad \times \underbrace{\int q_{\alpha,0}(\alpha | \mathcal{X}, \mathcal{Y}) d\alpha}_{=1} \prod_{k=1}^K q_{\theta_k^*}(\mathbf{c}_k, \boldsymbol{\Gamma}_k | \mathcal{X}, \mathcal{Y}) d\mathbf{c} d\boldsymbol{\Gamma} \\
 &= \sum_{k=1}^K \mathbb{E}_{q_{\boldsymbol{\tau}}}[\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\widehat{\mathbf{y}} | \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) \int \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k) q_{\theta_k^*}(\mathbf{c}_k, \boldsymbol{\Gamma}_k | \mathcal{X}, \mathcal{Y}) d\mathbf{c}_k d\boldsymbol{\Gamma}_k \\
 &\quad \times \underbrace{\int \prod_{j=1, j \neq k}^K q_{\theta_j^*}(\mathbf{c}_j, \boldsymbol{\Gamma}_j | \mathcal{X}, \mathcal{Y}) \prod_{j=1, j \neq k}^K d\mathbf{c}_j d\boldsymbol{\Gamma}_j}_{=1} \\
 &= \sum_{k=1}^K \mathbb{E}_{q_{\boldsymbol{\tau}_k}}[\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\widehat{\mathbf{y}} | \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) \underbrace{\int \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k) q_{\theta_k^*}(\mathbf{c}_k, \boldsymbol{\Gamma}_k | \mathcal{X}, \mathcal{Y}) d\mathbf{c}_k d\boldsymbol{\Gamma}_k}_{=St(\widehat{\mathbf{x}} | \widehat{\mathbf{m}}_k, \mathbf{L}_k, \widehat{\nu}_k + 1 - L) \text{ (Lemma B.1)}} \\
 &= \sum_{k=1}^K \mathbb{E}_{q_{\boldsymbol{\tau}_k}}[\pi_k(\boldsymbol{\tau})] St(\widehat{\mathbf{x}} | \widehat{\mathbf{m}}_k, \mathbf{L}_k, \widehat{\nu}_k + 1 - L) \mathcal{N}_D(\widehat{\mathbf{y}} | \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k). \tag{B.32}
 \end{aligned}$$

Here we used the following [Lemma B.1](#)

Lemma B.1. *For each $k \in [K]$, it holds that*

$$\int \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k) q_{\theta_k^*}(\mathbf{c}_k, \boldsymbol{\Gamma}_k | \mathcal{X}, \mathcal{Y}) d\mathbf{c}_k d\boldsymbol{\Gamma}_k = St(\widehat{\mathbf{x}} | \widehat{\mathbf{m}}_k, \mathbf{L}_k, \widehat{\nu}_k + 1 - L). \tag{B.33}$$

B.5.1 Proof of Theorem B.1

By definition, we obtain

$$\begin{aligned}
 &\int \int \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k^{-1}) q(\boldsymbol{\pi} | \mathcal{X}) q(\mathbf{c}_k, \boldsymbol{\Gamma}_k | \mathcal{X}) d\mathbf{c}_k d\boldsymbol{\Gamma}_k \\
 &= \int \int \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k^{-1}) \mathcal{N}_L(\mathbf{c}_k | \widehat{\mathbf{m}}_k, (\widehat{\lambda}_k \boldsymbol{\Gamma}_k)^{-1}, \mathcal{X}) \mathcal{W}(\boldsymbol{\Gamma}_k | \widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \mathcal{X}) d\mathbf{c}_k d\boldsymbol{\Gamma}_k \\
 &= \int \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k^{-1}) \mathcal{N}_L(\mathbf{c}_k | \widehat{\mathbf{m}}_k, (\widehat{\lambda}_k \boldsymbol{\Gamma}_k)^{-1}, \mathcal{X}) \mathcal{W}(\boldsymbol{\Gamma}_k | \widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \mathcal{X}) d\mathbf{c}_k d\boldsymbol{\Gamma}_k \\
 &= \int \left[\int \mathcal{N}_L(\widehat{\mathbf{x}} | \mathbf{c}_k, \boldsymbol{\Gamma}_k^{-1}) \mathcal{N}_L(\mathbf{c}_k | \widehat{\mathbf{m}}_k, (\widehat{\lambda}_k \boldsymbol{\Gamma}_k)^{-1}, \mathcal{X}) d\mathbf{c}_k \right] \mathcal{W}(\boldsymbol{\Gamma}_k | \widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \mathcal{X}) d\boldsymbol{\Gamma}_k \\
 &= \int \mathcal{N}_L(\widehat{\mathbf{x}} | \widehat{\mathbf{m}}_k, (1 + \widehat{\lambda}_k^{-1}) \boldsymbol{\Gamma}_k^{-1}, \mathcal{X}) \mathcal{W}(\boldsymbol{\Gamma}_k | \widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \mathcal{X}) d\boldsymbol{\Gamma}_k. \tag{B.34}
 \end{aligned}$$

When the size of the data set is large, *i.e.*, $N \rightarrow \infty$, this predictive distribution (B.34) becomes a mixture of Gaussians with component means $\widehat{\mathbf{m}}_k$ and precisions \mathbf{L}_k . In particular, we made use of the following results for marginal and conditional Gaussians, see, *e.g.*, [67, Eq. (2.115), page 93]. Given a marginal Gaussian distribution for \mathbf{x} and a conditional Gaussian distribution for \mathbf{y} given \mathbf{x} in the form

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \boldsymbol{\Gamma}^{-1}), \tag{B.35}$$

$$p(\mathbf{y} | \mathbf{x}) = \mathcal{N}(\mathbf{y} | \mathbf{A}\mathbf{x} + \mathbf{b}, \mathbf{L}^{-1}), \tag{B.36}$$

then the marginal distribution of \mathbf{y} and the conditional distribution of \mathbf{x} given \mathbf{y} are given by

$$p(\mathbf{y}) = \mathcal{N}\left(\mathbf{y} \mid \mathbf{A}\boldsymbol{\mu} + \mathbf{b}, \mathbf{L}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top\right), \quad (\text{B.37})$$

$$p(\mathbf{x} \mid \mathbf{y}) = \mathcal{N}\left(\mathbf{x} \mid \boldsymbol{\Sigma} \left\{ \mathbf{A}^\top \mathbf{L}(\mathbf{y} - \mathbf{b}) + \boldsymbol{\Gamma}\boldsymbol{\mu} \right\}, \boldsymbol{\Sigma}\right), \quad (\text{B.38})$$

where

$$\boldsymbol{\Sigma} = \left(\boldsymbol{\Gamma} + \mathbf{A}^\top \mathbf{L} \mathbf{A}\right)^{-1}. \quad (\text{B.39})$$

In our situation, via using $\mathbf{y} \equiv \widehat{\mathbf{x}}$, $\mathbf{x} \equiv \mathbf{c}_k$, $\mathbf{A} \equiv \mathbf{I}$, $\mathbf{b} \equiv \mathbf{0}$, $\mathbf{L}^{-1} = \boldsymbol{\Gamma}_k^{-1}$, $\boldsymbol{\mu} \equiv \widehat{\mathbf{m}}_k$, $\boldsymbol{\Gamma}^{-1} \equiv \left(\widehat{\lambda}_k \boldsymbol{\Gamma}_k\right)^{-1}$, we obtain

$$\begin{aligned} p(\widehat{\mathbf{x}} \mid \boldsymbol{\Gamma}_k^{-1}, \mathcal{X}) &= \int \mathcal{N}_L(\widehat{\mathbf{x}} \mid \mathbf{c}_k, \boldsymbol{\Gamma}_k^{-1}) \mathcal{N}_L(\mathbf{c}_k \mid \widehat{\mathbf{m}}_k, \left(\widehat{\lambda}_k \boldsymbol{\Gamma}_k\right)^{-1}, \mathcal{X}) d\mathbf{c}_k \\ &= \mathcal{N}_L\left(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, \boldsymbol{\Gamma}_k^{-1} + \left(\widehat{\lambda}_k \boldsymbol{\Gamma}_k\right)^{-1}, \mathcal{X}\right) \\ &= \mathcal{N}_L\left(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, \left(\frac{1 + \widehat{\lambda}_k}{\widehat{\lambda}_k}\right) \boldsymbol{\Gamma}_k^{-1}, \mathcal{X}\right). \end{aligned} \quad (\text{B.40})$$

Notice that the Wishart distribution is a conjugate prior for the Gaussian distribution with known mean and unknown precision. Therefore, it holds that the product of

$$\mathcal{N}_L\left(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, \left(1 + \widehat{\lambda}_k^{-1}\right) \boldsymbol{\Gamma}_k^{-1}, \mathcal{X}\right) \mathcal{W}\left(\boldsymbol{\Gamma}_k \mid \widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \mathcal{X}\right)$$

is again a Wishart distribution without normalized. This can be verified by focusing on the dependency on $\boldsymbol{\Gamma}_k$. More precisely, by using the trace trick of quadratic form, $(\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top \boldsymbol{\Gamma}_k (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) = \text{Tr}\left((\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)(\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top \boldsymbol{\Gamma}_k\right)$, we obtain

$$\begin{aligned} &\mathcal{N}_L\left(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, \left(1 + \widehat{\lambda}_k^{-1}\right) \boldsymbol{\Gamma}_k^{-1}, \mathcal{X}\right) \mathcal{W}\left(\boldsymbol{\Gamma}_k \mid \widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \mathcal{X}\right) \\ &= \underbrace{\frac{B(\widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k)}{\left(2\pi \left(1 + \widehat{\lambda}_k^{-1}\right)\right)^{L/2}}}_{:=C(\widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \widehat{\lambda}_k)} |\boldsymbol{\Gamma}_k|^{1/2 + (\widehat{\nu}_k - L - 1)/2} \exp\left\{-\frac{1}{2 \left(1 + \widehat{\lambda}_k^{-1}\right)} (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top \boldsymbol{\Gamma}_k (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) - \frac{1}{2} \text{Tr}\left(\widehat{\boldsymbol{\Psi}}_k^{-1} \boldsymbol{\Gamma}_k\right)\right\} \\ &= C\left(\widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \widehat{\lambda}_k\right) |\boldsymbol{\Gamma}_k|^{(\widehat{\nu}_k + 1 - L - 1)/2} \exp\left\{-\frac{1}{2} \text{Tr}\left(\left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)(\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top \boldsymbol{\Gamma}_k + \widehat{\boldsymbol{\Psi}}_k^{-1} \boldsymbol{\Gamma}_k\right)\right\} \\ &= C\left(\widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \widehat{\lambda}_k\right) |\boldsymbol{\Gamma}_k|^{(\widehat{\nu}_k + 1 - L - 1)/2} \exp\left\{-\frac{1}{2} \text{Tr}\left\{\left[\left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)(\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top + \widehat{\boldsymbol{\Psi}}_k^{-1}\right] \boldsymbol{\Gamma}_k\right\}\right\} \\ &= \frac{C\left(\widehat{\boldsymbol{\Psi}}_k, \widehat{\nu}_k, \widehat{\lambda}_k\right)}{B\left(\widehat{\boldsymbol{\Psi}}_k^*, \widehat{\nu}_k^*\right)} \mathcal{W}\left(\boldsymbol{\Gamma}_k \mid \boldsymbol{\Psi}_k^*, \widehat{\nu}_k^*\right). \end{aligned} \quad (\text{B.41})$$

Here, $\widehat{\nu}_k^* = \widehat{\nu}_k + 1$, and

$$\begin{aligned}
 \Psi_k^* &= \left[\left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top + \widehat{\Psi}_k^{-1} \right]^{-1}, \\
 \left| \widehat{\Psi}_k^* \right|^{(\widehat{\nu}_k+1)/2} &= \left| \left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top + \widehat{\Psi}_k^{-1} \right|^{-\widehat{(\nu}_k+1)/2} \\
 &= \left| \widehat{\Psi}_k^{-1} \left[\left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} \widehat{\Psi}_k (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top + \mathbf{I} \right] \right|^{-\widehat{(\nu}_k+1)/2} \\
 &= \left| \widehat{\Psi}_k \right|^{(\widehat{\nu}_k+1)/2} \left| \left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} \widehat{\Psi}_k (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top + \mathbf{I} \right|^{-\widehat{(\nu}_k+1)/2} \\
 &= \left| \widehat{\Psi}_k \right|^{(\widehat{\nu}_k+1)/2} \left[1 + \left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top \widehat{\Psi}_k (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) \right]^{-\widehat{(\nu}_k+1)/2}. \tag{B.42}
 \end{aligned}$$

Via using the normalization constant we have

$$\begin{aligned}
 &\int \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, (1 + \widehat{\lambda}_k^{-1}) \Gamma_k^{-1}, \mathcal{X}) \mathcal{W}(\Gamma_k \mid \widehat{\Psi}_k, \widehat{\nu}_k, \mathcal{X}) d\Gamma_k \\
 &= \frac{C(\widehat{\Psi}_k, \widehat{\nu}_k, \widehat{\lambda}_k)}{B(\widehat{\Psi}_k^*, \widehat{\nu}_k^*)} \underbrace{\int \mathcal{W}(\Gamma_k \mid \Psi_k^*, \widehat{\nu}_k^*) d\Gamma_k}_{=1} = \frac{C(\widehat{\Psi}_k, \widehat{\nu}_k, \widehat{\lambda}_k)}{B(\widehat{\Psi}_k^*, \widehat{\nu}_k^*)} = \frac{B(\widehat{\Psi}_k, \widehat{\nu}_k)}{(2\pi (1 + \widehat{\lambda}_k^{-1}))^{L/2}} \frac{1}{B(\widehat{\Psi}_k^*, \widehat{\nu}_k^*)} \\
 &= \frac{1}{(2\pi (1 + \widehat{\lambda}_k^{-1}))^{L/2}} \frac{\left| \widehat{\Psi}_k \right|^{-\widehat{\nu}_k/2} \left(2^{\widehat{\nu}_k L/2} \pi^{L(L-1)/4} \prod_{l=1}^L \Gamma\left(\frac{\widehat{\nu}_k+1-l}{2}\right) \right)^{-1}}{\left| \widehat{\Psi}_k^* \right|^{-\widehat{(\nu}_k+1)/2} \left(2^{(\widehat{\nu}_k+1)L/2} \pi^{L(L-1)/4} \prod_{l=1}^L \Gamma\left(\frac{(\widehat{\nu}_k+1)+1-l}{2}\right) \right)^{-1}} \\
 &= \frac{1}{(\pi (1 + \widehat{\lambda}_k^{-1}))^{L/2}} \frac{\left| \widehat{\Psi}_k \right|^{-\widehat{\nu}_k/2} \Gamma\left(\frac{\widehat{\nu}_k+1}{2}\right) \Gamma\left(\frac{\widehat{\nu}_k}{2}\right) \dots \Gamma\left(\frac{\widehat{\nu}_k+2-L}{2}\right)}{\left| \widehat{\Psi}_k^* \right|^{-\widehat{(\nu}_k+1)/2} \Gamma\left(\frac{\widehat{\nu}_k}{2}\right) \Gamma\left(\frac{\widehat{\nu}_k-1}{2}\right) \dots \Gamma\left(\frac{\widehat{\nu}_k+2-L}{2}\right) \Gamma\left(\frac{\widehat{\nu}_k+1-L}{2}\right)} \\
 &= \frac{\Gamma\left(\frac{\widehat{\nu}_k+1}{2}\right)}{\Gamma\left(\frac{\widehat{\nu}_k+1-L}{2}\right) \pi^{L/2}} \frac{\left| \widehat{\Psi}_k \right|^{-\widehat{\nu}_k/2}}{(1 + \widehat{\lambda}_k^{-1})^{L/2}} \left| \widehat{\Psi}_k \right|^{(\widehat{\nu}_k+1)/2} \left[1 + \left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top \widehat{\Psi}_k (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) \right]^{-\widehat{(\nu}_k+1)/2} \\
 &= \frac{\Gamma\left(\frac{\widehat{\nu}_k+1}{2}\right)}{\Gamma\left(\frac{\widehat{\nu}_k+1-L}{2}\right) \pi^{L/2}} \frac{\left| \widehat{\Psi}_k \right|^{1/2}}{(1 + \widehat{\lambda}_k^{-1})^{L/2}} \left[1 + \left(1 + \widehat{\lambda}_k^{-1}\right)^{-1} (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top \widehat{\Psi}_k (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k) \right]^{-\widehat{(\nu}_k+1)/2} \\
 &= \text{St}(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, \mathbf{L}_k, \widehat{\nu}_k + 1 - L). \tag{B.43}
 \end{aligned}$$

Here,

$$\mathbf{L}_k = \frac{(\widehat{\nu}_k + 1 - L) \widehat{\lambda}_k}{1 + \widehat{\lambda}_k} \widehat{\Psi}_k, \tag{B.44}$$

and Δ^2 is the squared Mahalanobis distance defined by

$$\Delta^2 = (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k)^\top \mathbf{L}_k (\widehat{\mathbf{x}} - \widehat{\mathbf{m}}_k). \tag{B.45}$$

Then, the last equality holds since we have

$$\begin{aligned}
 \text{St}(\hat{\mathbf{x}} \mid \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L) &= \frac{\Gamma\left(\frac{\hat{\nu}_k + 1 - L}{2} + \frac{L}{2}\right)}{\Gamma\left(\frac{\hat{\nu}_k + 1 - L}{2}\right)} \frac{|\mathbf{L}_k|^{1/2}}{\pi^{L/2} (\hat{\nu}_k + 1 - L)^{L/2}} \left[1 + \frac{\Delta^2}{\hat{\nu}_k + 1 - L}\right]^{-(\hat{\nu}_k + 1 - L)/2 - L/2} \\
 &= \frac{\Gamma\left(\frac{\hat{\nu}_k + 1}{2}\right)}{\Gamma\left(\frac{\hat{\nu}_k + 1 - L}{2}\right)} \frac{(\hat{\nu}_k + 1 - L)^{L/2} |\hat{\Psi}_k|^{1/2}}{\pi^{L/2} (\hat{\nu}_k + 1 - L)^{L/2} (1 + \hat{\lambda}_k^{-1})^{L/2}} \left[1 + \frac{\Delta^2}{\hat{\nu}_k + 1 - L}\right]^{-(\hat{\nu}_k + 1 - L)/2 - L/2} \\
 &= \frac{\Gamma\left(\frac{\hat{\nu}_k + 1}{2}\right)}{\Gamma\left(\frac{\hat{\nu}_k + 1 - L}{2}\right)} \frac{|\hat{\Psi}_k|^{1/2}}{\pi^{L/2} (1 + \hat{\lambda}_k^{-1})^{L/2}} \left[1 + \frac{(\hat{\nu}_k + 1 - L) \hat{\lambda}_k (\hat{\mathbf{x}} - \hat{\mathbf{m}}_k)^\top \hat{\Psi}_k (\hat{\mathbf{x}} - \hat{\mathbf{m}}_k)}{(1 + \hat{\lambda}_k) \hat{\nu}_k + 1 - L}\right]^{-(\hat{\nu}_k + 1)/2} \\
 &= \frac{\Gamma\left(\frac{\hat{\nu}_k + 1}{2}\right)}{\Gamma\left(\frac{\hat{\nu}_k + 1 - L}{2}\right)} \frac{|\hat{\Psi}_k|^{1/2}}{\pi^{L/2} (1 + \hat{\lambda}_k^{-1})^{L/2}} \left[1 + \frac{(\hat{\mathbf{x}} - \hat{\mathbf{m}}_k)^\top \hat{\Psi}_k (\hat{\mathbf{x}} - \hat{\mathbf{m}}_k)}{(1 + \hat{\lambda}_k^{-1})}\right]^{-(\hat{\nu}_k + 1)/2}. \tag{B.46}
 \end{aligned}$$

B.6 Proof of Theorem 4.5

From the product rule of probability, we see that this conditional distribution can be evaluated from the joint and marginal distributions. Furthermore, by integrating out $\hat{\mathbf{z}}$ and Θ , the predictive conditional density is then given by

$$p(\hat{\mathbf{y}} \mid \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) = \frac{p(\hat{\mathbf{y}}, \hat{\mathbf{x}} \mid \mathcal{X}, \mathcal{Y})}{p(\hat{\mathbf{x}} \mid \mathcal{X}, \mathcal{Y})} = \frac{\sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{y}}, \hat{\mathbf{x}}, \hat{\mathbf{z}}, \Theta \mid \mathcal{X}, \mathcal{Y}) d\hat{\mathbf{z}} d\Theta}{\sum_{\hat{\mathbf{z}}} \int p(\hat{\mathbf{x}}, \hat{\mathbf{z}}, \Theta \mid \mathcal{X}, \mathcal{Y}) d\hat{\mathbf{z}} d\Theta} \equiv \frac{T_1}{T_2}. \tag{B.47}$$

Next, with a similar step as in the proof of Theorem 4.4, we also obtain

$$T_2 = \sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} \mid \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L). \tag{B.48}$$

Therefore, we obtain

$$\begin{aligned}
 p(\hat{\mathbf{y}} \mid \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y}) &\approx \frac{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} \mid \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L) \mathcal{N}_D(\hat{\mathbf{y}} \mid \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\Sigma}_k)}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} \mid \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L)} \\
 &= \sum_{k=1}^K \frac{\mathbb{E}_{q_\tau} [\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} \mid \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L)}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} \mid \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L)} \mathcal{N}_D(\hat{\mathbf{y}} \mid \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\Sigma}_k) \\
 &\equiv \sum_{k=1}^K g_k(\hat{\mathbf{x}} \mid \hat{\Theta}, \hat{\phi}, \mathcal{X}, \mathcal{Y}) \mathcal{N}_D(\hat{\mathbf{y}} \mid \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\Sigma}_k), \tag{B.49}
 \end{aligned}$$

which is a mixture of Gaussian experts since we have

$$g_k(\hat{\mathbf{x}} \mid \hat{\Theta}, \hat{\phi}, \mathcal{X}, \mathcal{Y}) = \frac{\mathbb{E}_{q_\tau} [\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} \mid \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L)}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\tau)] \text{St}(\hat{\mathbf{x}} \mid \hat{\mathbf{m}}_k, \mathbf{L}_k, \hat{\nu}_k + 1 - L)}, \quad k \in [K], \tag{B.50}$$

belongs to a $K - 1$ dimensional probability simplex.

B.7 Proof of Theorem 4.6

To deal with high-dimensional regression data, namely high-to-low regression, given the inverse conditional density $p(\hat{\mathbf{y}} \mid \hat{\mathbf{x}}, \mathcal{X}, \mathcal{Y})$, we want to compute the following forward conditional density

$$p(\hat{\mathbf{x}} \mid \hat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) = \frac{p(\hat{\mathbf{x}}, \hat{\mathbf{y}} \mid \mathcal{X}, \mathcal{Y})}{p(\hat{\mathbf{y}} \mid \mathcal{X}, \mathcal{Y})} = \frac{p(\hat{\mathbf{x}}, \hat{\mathbf{y}} \mid \mathcal{X}, \mathcal{Y})}{\int_{\hat{\mathbf{x}}} p(\hat{\mathbf{x}}, \hat{\mathbf{y}} \mid \mathcal{X}, \mathcal{Y}) d\hat{\mathbf{x}}} = \frac{T_1}{\int_{\hat{\mathbf{x}}} T_1(\hat{\mathbf{x}}) d\hat{\mathbf{x}}} \equiv \frac{T_1}{T_3}. \tag{B.51}$$

Then, we have to compute or numerically approximate D_3 . Using Theorem 4.4 and definition of Student's t-distribution, we obtain

$$T_3 = \sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] D_k. \quad (\text{B.52})$$

Then, by definition of Student's t-distribution, it holds that

$$\begin{aligned} D_k &= \int \text{St}(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, \mathbf{L}_k, \widehat{\nu}_k + 1 - L) \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) d\widehat{\mathbf{x}} \\ &= \int \int_0^\infty \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, (\eta \mathbf{L}_k)^{-1}) \text{Gam}\left(\eta \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right) \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) d\eta d\widehat{\mathbf{x}} \\ &= \int_0^\infty \int \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, (\eta \mathbf{L}_k)^{-1}) d\widehat{\mathbf{x}} \text{Gam}\left(\eta \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right) d\eta \\ &= \int_0^\infty \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{m}}_k + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k + \eta^{-1} \widehat{\mathbf{A}}_k \mathbf{L}_k^{-1} \widehat{\mathbf{A}}_k^\top) \text{Gam}\left(\eta \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right) d\eta. \end{aligned}$$

Furthermore, we used the fact that

$$\int \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, (\eta \mathbf{L}_k)^{-1}) d\widehat{\mathbf{x}} = \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{m}}_k + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k + \widehat{\mathbf{A}}_k (\eta \mathbf{L}_k)^{-1} \widehat{\mathbf{A}}_k^\top).$$

Indeed, we made use of the following results for marginal and conditional Gaussians, see, *e.g.*, [67, Eq. (2.115), page 93]. Given a marginal Gaussian distribution for \mathbf{x} and a conditional Gaussian distribution for \mathbf{y} given \mathbf{x} in the form

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}^{-1}), \quad (\text{B.53})$$

$$p(\mathbf{y} \mid \mathbf{x}) = \mathcal{N}(\mathbf{y} \mid \mathbf{A}\mathbf{x} + \mathbf{b}, \mathbf{L}^{-1}), \quad (\text{B.54})$$

then the marginal distribution of \mathbf{y} and the conditional distribution of \mathbf{x} given \mathbf{y} are given by

$$p(\mathbf{y}) = \int p(\mathbf{y} \mid \mathbf{x}) p(\mathbf{x}) d\mathbf{x} = \mathcal{N}(\mathbf{y} \mid \mathbf{A}\boldsymbol{\mu} + \mathbf{b}, \mathbf{L}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top), \quad (\text{B.55})$$

$$p(\mathbf{x} \mid \mathbf{y}) = \mathcal{N}(\mathbf{x} \mid \boldsymbol{\Sigma} \{ \mathbf{A}^\top \mathbf{L}(\mathbf{y} - \mathbf{b}) + \boldsymbol{\Gamma}\boldsymbol{\mu} \}, \boldsymbol{\Sigma}), \boldsymbol{\Sigma} = (\boldsymbol{\Gamma} + \mathbf{A}^\top \mathbf{L} \mathbf{A})^{-1}. \quad (\text{B.56})$$

In our situation, the desired result is obtained via using $\mathbf{y} \equiv \widehat{\mathbf{y}}$, $\mathbf{A} \equiv \widehat{\mathbf{A}}_k$, $\mathbf{b} \equiv \widehat{\mathbf{b}}_k$, $\mathbf{L}^{-1} = \widehat{\boldsymbol{\Sigma}}_k$, $\mathbf{x} \equiv \widehat{\mathbf{x}}$, $\boldsymbol{\mu} \equiv \widehat{\mathbf{m}}_k$, $\boldsymbol{\Gamma}^{-1} \equiv (\eta \mathbf{L}_k)^{-1}$.

Therefore, we obtain

$$\begin{aligned} &p(\widehat{\mathbf{x}} \mid \widehat{\mathbf{y}}, \mathcal{X}, \mathcal{Y}) \\ &\approx \frac{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \text{St}(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, \mathbf{L}_k, \widehat{\nu}_k + 1 - L) \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k)}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \int \text{St}(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, \mathbf{L}_k, \widehat{\nu}_k + 1 - L) \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) d\widehat{\mathbf{x}}} \\ &= \frac{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \int_0^\infty \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{x}} + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k) \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{m}}_k, (\eta \mathbf{L}_k)^{-1}) \text{Gam}\left(\eta \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right) d\eta}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \int_0^\infty \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{A}}_k \widehat{\mathbf{m}}_k + \widehat{\mathbf{b}}_k, \widehat{\boldsymbol{\Sigma}}_k + \eta^{-1} \widehat{\mathbf{A}}_k \mathbf{L}_k^{-1} \widehat{\mathbf{A}}_k^\top) \text{Gam}\left(\eta \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right) d\eta} \\ &= \frac{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \int_0^\infty \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{c}}_k^*, \widehat{\boldsymbol{\Gamma}}_k^*(\eta)) \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{A}}_k^*(\eta) \widehat{\mathbf{y}} + \widehat{\mathbf{b}}_k^*(\eta), \widehat{\boldsymbol{\Sigma}}_k^*(\eta)) \text{Gam}\left(\eta \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right) d\eta}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \int_0^\infty \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{c}}_k^*, \widehat{\boldsymbol{\Gamma}}_k^*(\eta)) \text{Gam}\left(\eta \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right) d\eta} \\ &\approx \frac{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \frac{U_\eta}{R} \sum_{i=1}^I \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{c}}_k^*, \widehat{\boldsymbol{\Gamma}}_k^*(\eta_i)) \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{A}}_k^*(\eta_i) \widehat{\mathbf{y}} + \widehat{\mathbf{b}}_k^*(\eta_i), \widehat{\boldsymbol{\Sigma}}_k^*(\eta_i)) \text{Gam}\left(\eta_i \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right)}{\sum_{k=1}^K \mathbb{E}_{q_\tau} [\pi_k(\boldsymbol{\tau})] \frac{U_\eta}{R} \sum_{i=1}^I \mathcal{N}_D(\widehat{\mathbf{y}} \mid \widehat{\mathbf{c}}_k^*, \widehat{\boldsymbol{\Gamma}}_k^*(\eta_i)) \text{Gam}\left(\eta_i \mid \frac{\widehat{\nu}_k + 1 - L}{2}, \frac{\widehat{\nu}_k + 1 - L}{2}\right)} \\ &\approx \sum_{i=1}^I \sum_{k=1}^K g_{ki} (\widehat{\mathbf{y}} \mid \widehat{\boldsymbol{\Theta}}^*, \widehat{\boldsymbol{\Phi}}^*, \mathcal{X}, \mathcal{Y}) \mathcal{N}_L(\widehat{\mathbf{x}} \mid \widehat{\mathbf{A}}_k^*(\eta_i) \widehat{\mathbf{y}} + \widehat{\mathbf{b}}_k^*(\eta_i), \widehat{\boldsymbol{\Sigma}}_k^*(\eta_i)), \end{aligned}$$

where, for all $k \in [K], i \in [I]$,

$$g_{ki}(\hat{\mathbf{y}} | \hat{\Theta}^*, \hat{\phi}^*, \mathcal{X}, \mathcal{Y}) = \frac{\mathbb{E}_{q_\tau}[\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_k^*, \hat{\boldsymbol{\Gamma}}_k^*(\eta_i)) \text{Gam}(\eta_i | \frac{\hat{\nu}_k+1-L}{2}, \frac{\hat{\nu}_k+1-L}{2})}{\sum_{i=1}^I \sum_{k=1}^K \mathbb{E}_{q_\tau}[\pi_k(\boldsymbol{\tau})] \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_k^*, \hat{\boldsymbol{\Gamma}}_k^*(\eta_i)) \text{Gam}(\eta_i | \frac{\hat{\nu}_k+1-L}{2}, \frac{\hat{\nu}_k+1-L}{2})}.$$

Here, we used the fact that $p(\hat{\mathbf{y}}, \hat{\mathbf{x}} | \hat{z} = k) = p(\hat{\mathbf{x}} | \hat{\mathbf{y}}, \hat{z} = k)p(\hat{\mathbf{y}} | \hat{z} = k)$, namely,

$$\begin{aligned} & \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{x}} + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k) \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{m}}_k, (\eta \mathbf{L}_k)^{-1}) \\ &= \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\boldsymbol{\Sigma}}_k^* [\hat{\mathbf{A}}_k^\top \hat{\boldsymbol{\Sigma}}_k^{-1} (\hat{\mathbf{y}} - \hat{\mathbf{b}}_k) + \eta \mathbf{L}_k \hat{\mathbf{m}}_k], \hat{\boldsymbol{\Sigma}}_k^*) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{A}}_k \hat{\mathbf{m}}_k + \hat{\mathbf{b}}_k, \hat{\boldsymbol{\Sigma}}_k + \hat{\mathbf{A}}_k (\eta \mathbf{L}_k)^{-1} \hat{\mathbf{A}}_k^\top) \\ &= \mathcal{N}_L(\hat{\mathbf{x}} | \hat{\mathbf{A}}_k^*(\eta) \hat{\mathbf{y}} + \hat{\mathbf{b}}_k^*(\eta), \hat{\boldsymbol{\Sigma}}_k^*(\eta)) \mathcal{N}_D(\hat{\mathbf{y}} | \hat{\mathbf{c}}_k^*, \hat{\boldsymbol{\Gamma}}_k^*(\eta)), \end{aligned}$$

where

$$\begin{aligned} \hat{\boldsymbol{\Sigma}}_k^*(\eta) &= (\eta \mathbf{L}_k + \hat{\mathbf{A}}_k^\top \hat{\boldsymbol{\Sigma}}_k^{-1} \hat{\mathbf{A}}_k)^{-1}, \\ \hat{\mathbf{A}}_k^*(\eta) &= \hat{\boldsymbol{\Sigma}}_k^*(\eta) \hat{\mathbf{A}}_k^\top \hat{\boldsymbol{\Sigma}}_k^{-1}, \\ \hat{\mathbf{b}}_k^*(\eta) &= \hat{\boldsymbol{\Sigma}}_k^*(\eta) [\eta \mathbf{L}_k \hat{\mathbf{m}}_k - \hat{\mathbf{A}}_k^\top \hat{\boldsymbol{\Sigma}}_k^{-1} \hat{\mathbf{b}}_k], \\ \hat{\mathbf{c}}_k^* &= \hat{\mathbf{A}}_k \hat{\mathbf{m}}_k + \hat{\mathbf{b}}_k, \\ \hat{\boldsymbol{\Gamma}}_k^*(\eta) &= \hat{\boldsymbol{\Sigma}}_k + \hat{\mathbf{A}}_k (\eta \mathbf{L}_k)^{-1} \hat{\mathbf{A}}_k^\top. \end{aligned}$$

The last approximation is deduced by using the fact that one simplistic strategy for evaluating integration would be to discretize η -space (1-dimensional) into a uniform grid and to evaluate the integrand as a Riemann integral with a truncated value $0 < U_\eta < \infty$ and a number of point $I \in \mathbb{N}^*$ for approximating the integration.

B.8 Proof of Theorem 5.1

Using the sum and product rules for both discrete and continuous variables, the ELBO in BNP-GLLiM (3.4) is given by

$$\begin{aligned} \mathcal{F}(q_{\mathcal{Z}}, q_{\Theta}, \hat{\phi}) &= \mathbb{E}_{q_{\mathcal{Z}} q_{\Theta}} \left[\log \frac{p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi})}{q(\mathcal{Z}) q(\Theta)} \right] \equiv \mathbb{E} \left[\log \frac{p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi})}{q(\mathcal{Z}) q(\Theta)} \right] \\ &= \sum_{\mathcal{Z}} \int \int \int q(\mathcal{Z}) q(\Theta) \log \left[\frac{p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi})}{q(\mathcal{Z}) q(\Theta)} \right] d\mathcal{Z} d\Theta \\ &= \mathbb{E} \left[\log p(\mathcal{Y} | \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\mathcal{X} | \mathcal{Z}, \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\mathcal{Z} | \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\Theta; \hat{\phi}) \right] \\ &\quad - \mathbb{E} \left[\log q(\mathcal{Z}) \right] - \mathbb{E} \left[\log q(\Theta) \right]. \end{aligned} \tag{B.57}$$

Next, we evaluate the various terms in the ELBO (B.57).

Proof of (5.2)

Via the mean field approximation and the truncation, we have the following computations.

$$\begin{aligned}
 \mathbb{E} \left[\log p(\mathcal{Y} \mid \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi}) \right] &= \mathbb{E} \left[\log \prod_{n=1}^N p \left(\mathbf{y}_n \mid \mathbf{x}_n, z_n, \Theta; \hat{\phi} \right) \right] \\
 &= \mathbb{E} \left[\log \prod_{n=1}^N \prod_{k=1}^K \mathcal{N}_D \left(\mathbf{y}_n \mid \hat{\mathbf{A}}_k \mathbf{x}_n + \hat{\mathbf{b}}_k, \hat{\Sigma}_k \right)^{z_{nk}} \right] \\
 &= \sum_{n=1}^N \sum_{k=1}^K \mathbb{E} \left[z_{nk} \log \mathcal{N}_D \left(\mathbf{y}_n \mid \hat{\mathbf{A}}_k \mathbf{x}_n + \hat{\mathbf{b}}_k, \hat{\Sigma}_k \right) \right] \\
 &= \sum_{n=1}^N \sum_{k=1}^K \mathbb{E}_{q_{\mathcal{Z}}} [z_{nk}] \log \mathcal{N}_D \left(\mathbf{y}_n \mid \hat{\mathbf{A}}_k \mathbf{x}_n + \hat{\mathbf{b}}_k, \hat{\Sigma}_k \right) \\
 &= \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \log \mathcal{N}_D \left(\mathbf{y}_n \mid \hat{\mathbf{A}}_k \mathbf{x}_n + \hat{\mathbf{b}}_k, \hat{\Sigma}_k \right), \tag{B.58}
 \end{aligned}$$

where

$$\log \mathcal{N}_D \left(\mathbf{y}_n \mid \hat{\mathbf{A}}_k \mathbf{x}_n + \hat{\mathbf{b}}_k, \hat{\Sigma}_k \right) = -\frac{D}{2} \log(2\pi) - \frac{1}{2} \log \left| \hat{\Sigma}_k \right| - (\mathbf{y}_n - \hat{\mathbf{A}}_k \mathbf{x}_n - \hat{\mathbf{b}}_k)^\top \hat{\Sigma}_k^{-1} (\mathbf{y}_n - \hat{\mathbf{A}}_k \mathbf{x}_n - \hat{\mathbf{b}}_k). \tag{B.59}$$

Proof of (5.3)

Similarly, we obtain

$$\begin{aligned}
 \mathbb{E} \left[\log p(\mathcal{X} \mid \mathcal{Z}, \Theta; \hat{\phi}) \right] &= \mathbb{E} \left[\log \prod_{n=1}^N p \left(\mathbf{x}_n \mid z_n, \Theta; \hat{\phi} \right) \right] \\
 &= \mathbb{E} \left[\log \prod_{n=1}^N \prod_{k=1}^K \mathcal{N}_L \left(\mathbf{x}_n \mid \hat{\mathbf{c}}_k, \hat{\Gamma}_k \right)^{z_{nk}} \right] \\
 &= \sum_{n=1}^N \sum_{k=1}^K \mathbb{E} \left[z_{nk} \log \mathcal{N}_L \left(\mathbf{x}_n \mid \hat{\mathbf{c}}_k, \hat{\Gamma}_k \right) \right] \\
 &= \sum_{n=1}^N \sum_{k=1}^K \mathbb{E}_{q_{\mathcal{Z}}} [z_{nk}] \log \mathcal{N}_L \left(\mathbf{x}_n \mid \hat{\mathbf{c}}_k, \hat{\Gamma}_k \right) \\
 &= \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \log \mathcal{N}_L \left(\mathbf{x}_n \mid \hat{\mathbf{c}}_k, \hat{\Gamma}_k \right), \tag{B.60}
 \end{aligned}$$

where

$$\log \mathcal{N}_L \left(\mathbf{x}_n \mid \hat{\mathbf{c}}_k, \hat{\Gamma}_k \right) = -\frac{L}{2} \log(2\pi) - \frac{1}{2} \log \left| \hat{\Gamma}_k \right| - \frac{1}{2} (\mathbf{x}_n - \hat{\mathbf{c}}_k)^\top \hat{\Gamma}_k^{-1} (\mathbf{x}_n - \hat{\mathbf{c}}_k). \tag{B.61}$$

Proof of (5.4)

Via calculation, it follows the expressions of the following quantities,

$$\begin{aligned}
 \mathbb{E}_{q_{\tau_k}} [\log(\tau_k)] &= \psi(\hat{\gamma}_{k,1}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}), \\
 \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)] &= \psi(\hat{\gamma}_{k,2}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}). \tag{B.62}
 \end{aligned}$$

When $\beta = 0$, via using (B.62), it holds that

$$\begin{aligned}
 \mathbb{E} \left[\log p(\mathcal{Z} \mid \Theta; \hat{\phi}) \right] &= \mathbb{E} \left[\log \prod_{n=1}^N \prod_{k=1}^K [\pi_k(\boldsymbol{\tau})]^{z_{nk}} \right] \\
 &= \sum_{n=1}^N \sum_{k=1}^K \mathbb{E}_{q_{\mathcal{Z}}} [z_{nk}] \mathbb{E}_{q_{\Theta}} \left[\log \left[\tau_k \prod_{l=1}^{k-1} (1 - \tau_l) \right] \right] \\
 &= \sum_{k=1}^K \sum_{n=1}^N q_{z_{nk}} \left[\mathbb{E}_{q_{\tau_k}} [\log \tau_k] + \sum_{l=1}^{k-1} \mathbb{E}_{q_{\tau_l}} [\log (1 - \tau_l)] \right] \\
 &= \sum_{k=1}^K N_k \left[\mathbb{E}_{q_{\tau_k}} [\log \tau_k] + \sum_{l=1}^{k-1} \mathbb{E}_{q_{\tau_l}} [\log (1 - \tau_l)] \right] \\
 &= \sum_{k=1}^K N_k \left[\psi(\hat{\gamma}_{k,1}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}) + \sum_{l=1}^{k-1} [\psi(\hat{\gamma}_{l,2}) - \psi(\hat{\gamma}_{l,1} + \hat{\gamma}_{l,2})] \right].
 \end{aligned} \tag{B.63}$$

When $\beta \neq 0$, it holds that

$$p(\mathcal{Z}; \hat{\beta}, \boldsymbol{\pi}) = C(\boldsymbol{\tau}, \hat{\beta}) \left(\prod_{n=1}^N \pi_{z_n} \right) \exp \left(\hat{\beta} \sum_{m \sim n} \delta_{(z_m = z_n)} \right), \tag{B.64}$$

where C denotes the normalizing constant that depends on $\boldsymbol{\tau}$ and $\hat{\beta}$. Unfortunately, due to the intractable normalizing constant C , it is necessary to evaluate the terms involving $p(\mathcal{Z}; \hat{\beta}, \boldsymbol{\pi})$ in approximate manner.

Proof of (5.5)

Given a chosen truncated value $K \in \mathbb{N}^*$, it holds that

$$\mathbb{E}_{q_{\Theta}} \left[\log p(\Theta; \hat{\phi}) \right] = \sum_{k=1}^{K-1} \mathbb{E}_{q_{\Theta}} [\log p(\tau_k \mid \alpha, \sigma)] + \mathbb{E}_{q_{\Theta}} [\log p(\alpha, \sigma \mid \hat{s}_1, \hat{s}_2, \hat{a})].$$

Here, we have

$$\begin{aligned}
 \mathbb{E}_{q_{\Theta}} [\log p(\tau_k \mid \alpha, \sigma)] &= \mathbb{E}_{q_{\Theta}} [\log \text{Beta}(\tau_k \mid 1 - \sigma, \alpha + k\sigma)] \\
 &= \mathbb{E}_{q_{\Theta}} \left[\log \tau_k^{-\sigma} (1 - \tau_k)^{\alpha + k\sigma - 1} + \log C(\alpha, \sigma) \right], \\
 &= -\mathbb{E}_{q_{\alpha, \sigma}} [\sigma] \mathbb{E}_{q_{\tau_k}} [\log \tau_k] + \mathbb{E}_{q_{\alpha, \sigma}} [\alpha + k\sigma - 1] \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)] + \mathbb{E}_{q_{\alpha, \sigma}} [\log C(\alpha, \sigma)].
 \end{aligned}$$

where we have defined

$$C(\alpha, \sigma) = \frac{\Gamma(1 - \sigma + \alpha + k\sigma)}{\Gamma(1 - \sigma)\Gamma(\alpha + k\sigma)}.$$

Next, for the sake of simplicity, we use for σ a uniform prior $\mathcal{U}_{[0,1]}(\sigma)$ so that parameter a does not have to be taken into account. Then it holds that

$$\begin{aligned}
 \mathbb{E}_{q_{\Theta}} [\log p(\alpha, \sigma \mid \hat{s}_1, \hat{s}_2)] &= \mathbb{E}_{q_{\alpha, \sigma}} [\log \text{Gam}(\alpha + \sigma \mid \hat{s}_1, \hat{s}_2)] + \mathbb{E}_{q_{\alpha, \sigma}} [\log \mathcal{U}_{[0,1]}(\sigma)] \\
 &= \log \left[\frac{1}{\Gamma(\hat{s}_1) \hat{s}_2^{\hat{s}_1}} \right] + (\hat{s}_1 - 1) \mathbb{E}_{q_{\alpha, \sigma}} [\log(\alpha + \sigma)] - \hat{s}_2 \mathbb{E}_{q_{\alpha, \sigma}} [\alpha + \sigma] \\
 &= \log \left[\frac{1}{\Gamma(\hat{s}_1) \hat{s}_2^{\hat{s}_1}} \right] + (\hat{s}_1 - 1) \mathbb{E}_{q_{\alpha, \sigma}} [\log(\alpha + \sigma)] - \hat{s}_2 \mathbb{E}_{q_{\alpha, \sigma}} [\alpha + \sigma].
 \end{aligned}$$

When $\sigma \neq 0$, the normalizing constant for $q_{\alpha, \sigma}(\alpha, \sigma)$ is not tractable. Nevertheless, to compute the ELBO, we do not need the full $q_{\alpha, \sigma}$ distribution but only the means $\mathbb{E}_{q_{\alpha, \sigma}} [\sigma]$, $\mathbb{E}_{q_{\alpha, \sigma}} [\alpha + k\sigma - 1]$,

$\mathbb{E}_{q_{\alpha,\sigma}} [\log C(\alpha, \sigma)]$, $\mathbb{E}_{q_{\alpha,\sigma}} [\log(\alpha + \sigma)]$ and $\mathbb{E}_{q_{\alpha,\sigma}} [\alpha + \sigma]$. One solution is therefore to use importance sampling or MCMC to compute these expectations via Monte Carlo sums.

When $\sigma = 0$, using integration by parts, it holds that $\Gamma(\alpha + 1) = \alpha\Gamma(\alpha)$ and hence $C(\alpha, \sigma) \equiv C(\alpha) = \alpha$. Furthermore, the posterior $q_{\alpha,\sigma} \equiv q_\alpha$ is again a gamma distribution $\text{Gam}(\alpha \mid \hat{s}_1, \hat{s}_2)$ with $\mathbb{E}_{q_{\alpha,\sigma}} [\alpha] \equiv \mathbb{E}_{q_\alpha} [\alpha] = \frac{\hat{s}_1}{\hat{s}_2}$ and $\mathbb{E}_{q_{\alpha,\sigma}} [\log \alpha] \equiv \mathbb{E}_{q_\alpha} [\log \alpha] = \psi(\hat{s}_1) - \log(\hat{s}_2)$. Therefore, we have the following tractable formulas:

$$\begin{aligned} \mathbb{E}_{q_\Theta} [\log p(\tau_k \mid \alpha, \sigma)] &\equiv \mathbb{E}_{q_\Theta} [\log p(\tau_k \mid \alpha)] \\ &= [\mathbb{E}_{q_{\alpha,0}} [\alpha] - 1] \mathbb{E}_{q_{\tau_k}} [\log(1 - \tau_k)] + \mathbb{E}_{q_{\alpha,0}} [\log \alpha], \\ &= \frac{\hat{s}_1 - \hat{s}_2}{\hat{s}_2} [\psi(\hat{\gamma}_{k,2}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2})] + \psi(\hat{s}_1) - \log \hat{s}_2, \end{aligned} \quad (\text{B.65})$$

$$\begin{aligned} \mathbb{E}_{q_\Theta} [\log p(\alpha, \sigma \mid \hat{s}_1, \hat{s}_2)] &\equiv \mathbb{E}_{q_\Theta} [\log p(\alpha \mid \hat{s}_1, \hat{s}_2)] \\ &= \log \left[\frac{1}{\Gamma(\hat{s}_1)} \hat{s}_2^{\hat{s}_1} \right] + (\hat{s}_1 - 1) [\psi(\hat{s}_1) - \log(\hat{s}_2)] - \hat{s}_1. \end{aligned} \quad (\text{B.66})$$

Proof of (5.6)

Due to the mean-field approximation (3.10) and truncation, this step is analytically computed as follows:

$$\begin{aligned} \mathbb{E}_{q_Z} [\log q(Z)] &= \mathbb{E}_{q_Z} \left[\log \prod_{n=1}^N q_{z_n}(z_n) \right] = \mathbb{E}_{q_Z} \left[\log \prod_{n=1}^N \prod_{k=1}^K q_{z_n}(k)^{z_{nk}} \right] \\ &= \sum_{n=1}^N \sum_{k=1}^K \log q_{z_n}(k) \mathbb{E}_{q_Z} [z_{nk}] = \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \log q_{z_n}(k). \end{aligned} \quad (\text{B.67})$$

Proof of (5.7)

We have

$$\mathbb{E} [\log q(\Theta)] = \mathbb{E} [\log q_{\alpha,\sigma}(\alpha, \sigma)] + \sum_{k=1}^{K-1} \mathbb{E} [\log q_{\tau_k}(\tau_k)]. \quad (\text{B.68})$$

Note that these terms involving expectations of the logs of the q distributions simply represent the negative entropies of those distributions.

Since $q_{\alpha,\sigma}(\alpha, \sigma)$ is not tractable, when $\sigma \neq 0$, we cannot calculate analytically $\mathbb{E} [\log q_{\alpha,\sigma}(\alpha, \sigma)]$. Furthermore, it is also difficult to approximate it using MCMC or importance sampling.

When $\sigma = 0$, the posterior $q_{\alpha,\sigma} \equiv q_\alpha$ is again a gamma distribution $\text{Gam}(\alpha \mid \hat{s}_1, \hat{s}_2)$ with

$$\begin{aligned} \mathbb{E} [\log q_{\alpha,0}(\alpha)] &\equiv \mathbb{E} [\log \text{Gam}(\alpha \mid \hat{s}_1, \hat{s}_2)] \\ &= -\text{H}[\text{Gam}(\alpha \mid \hat{s}_1, \hat{s}_2)] \\ &= -\log \Gamma(\hat{s}_1) + (\hat{s}_1 - 1) \psi(\hat{s}_1) + \log \hat{s}_2 - \hat{s}_1. \end{aligned} \quad (\text{B.69})$$

Since we had $q_{\tau_k}(\tau_k) = \text{Beta}(\tau_k \mid \hat{\gamma}_{k,1}, \hat{\gamma}_{k,2})$, its differential entropy is given by

$$\begin{aligned} \mathbb{E} [\log q_{\tau_k}(\tau_k)] &= -\text{H}[\text{Beta}(\tau_k \mid \hat{\gamma}_{k,1}, \hat{\gamma}_{k,2})] \\ &= \sum_{l=1}^2 (\hat{\gamma}_{k,l} - 1) \{ \psi(\hat{\gamma}_{k,l}) - \psi(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2}) \} + \log \frac{\Gamma(\hat{\gamma}_{k,1} + \hat{\gamma}_{k,2})}{\Gamma(\hat{\gamma}_{k,1}) \Gamma(\hat{\gamma}_{k,2})}. \end{aligned} \quad (\text{B.70})$$

B.9 Proof of Theorem 5.2

Using the sum and product rules for both discrete and continuous variables, the ELBO in BNP-GLLiM (3.4) is given by

$$\begin{aligned}
 \mathcal{F}(q_{\mathcal{Z}}, q_{\Theta}, \hat{\phi}) &= \mathbb{E}_{q_{\mathcal{Z}}q_{\Theta}} \left[\log \frac{p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi})}{q(\mathcal{Z})q_{\Theta}(\Theta)} \right] \equiv \mathbb{E} \left[\log \frac{p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi})}{q(\mathcal{Z})q(\Theta)} \right] \\
 &= \sum_{\mathcal{Z}} \int \int \int q(\mathcal{Z})q(\Theta) \log \left[\frac{p(\mathcal{Y}, \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi})}{q(\mathcal{Z})q(\Theta)} \right] d\mathcal{Z}d\Theta \\
 &= \mathbb{E} \left[\log p(\mathcal{Y} | \mathcal{X}, \mathcal{Z}, \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\mathcal{X} | \mathcal{Z}, \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\mathcal{Z} | \Theta; \hat{\phi}) \right] + \mathbb{E} \left[\log p(\Theta; \hat{\phi}) \right] \\
 &\quad - \mathbb{E} \left[\log q(\mathcal{Z}) \right] - \mathbb{E} \left[\log q(\Theta) \right]. \tag{B.71}
 \end{aligned}$$

Note that the proof of (5.11), (5.13), (5.15) are the same as in the proof of Theorem 5.1.

Proof of (5.12)

$$\begin{aligned}
 \mathbb{E} \left[\log p(\mathcal{X} | \mathcal{Z}, \Theta; \hat{\phi}) \right] &= \mathbb{E} \left[\log \prod_{n=1}^N p(\mathbf{x}_n | z_n, \Theta; \hat{\phi}) \right] = \mathbb{E} \left[\log \prod_{n=1}^N \prod_{k=1}^K \mathcal{N}_L(\mathbf{x}_n | \mathbf{c}_k, \mathbf{\Gamma}_k)^{z_{nk}} \right] \\
 &= \sum_{n=1}^N \sum_{k=1}^K \mathbb{E} [z_{nk} \log \mathcal{N}_L(\mathbf{x}_n | \mathbf{c}_k, \mathbf{\Gamma}_k)] \\
 &= \sum_{n=1}^N \sum_{k=1}^K \mathbb{E}_{q_{\mathcal{Z}}} [z_{nk}] \mathbb{E}_{q_{\mathbf{c}_k, \mathbf{\Gamma}_k}} [\log \mathcal{N}_L(\mathbf{x}_n | \mathbf{c}_k, \mathbf{\Gamma}_k)] \\
 &= \sum_{n=1}^N \sum_{k=1}^K q_{z_n}(k) \left[-\frac{L}{2} \log(2\pi) - \frac{1}{2} \mathbb{E} [\log |\mathbf{\Gamma}_k|] - \frac{1}{2} \mathbb{E} \left[(\mathbf{x}_n - \mathbf{c}_k)^\top \mathbf{\Gamma}_k^{-1} (\mathbf{x}_n - \mathbf{c}_k) \right] \right] \\
 &= \sum_{k=1}^K \sum_{n=1}^N q_{z_n}(k) \left[-\frac{L}{2} \log(2\pi) - \frac{1}{2} \mathbb{E}_{q_{\mathbf{\Gamma}_k}} [\log |\mathbf{\Gamma}_k|] - \frac{1}{2} \mathbb{E}_{q_{\mathbf{c}_k, \mathbf{\Gamma}_k}} \left[(\mathbf{x}_n - \mathbf{c}_k)^\top \mathbf{\Gamma}_k^{-1} (\mathbf{x}_n - \mathbf{c}_k) \right] \right] \tag{Lemma B.2} \\
 &= \frac{1}{2} \sum_{k=1}^K \sum_{n=1}^N q_{z_n}(k) \left[-\log \tilde{\mathbf{\Gamma}}_k - L \log(2\pi) - L \hat{\lambda}_k^{-1} - \hat{\nu}_k (\mathbf{x}_n - \hat{\mathbf{m}}_k) \hat{\Psi}_k^{-1} (\mathbf{x}_n - \hat{\mathbf{m}}_k) \right] \\
 &= \frac{1}{2} \sum_{k=1}^K N_k \left[\log \tilde{\mathbf{\Gamma}}_k - L \log(2\pi) - L \hat{\lambda}_k^{-1} \right] - \frac{1}{2} \sum_{k=1}^K \sum_{n=1}^N q_{z_n}(k) \left[\hat{\nu}_k (\mathbf{x}_n - \hat{\mathbf{m}}_k) \hat{\Psi}_k^{-1} (\mathbf{x}_n - \hat{\mathbf{m}}_k) \right] \\
 &= \frac{1}{2} \sum_{k=1}^K N_k \left[\log \tilde{\mathbf{\Gamma}}_k - L \log(2\pi) - L \hat{\lambda}_k^{-1} - \hat{\nu}_k \text{Tr}(\mathbf{S}_k \hat{\Psi}_k^{-1}) \right] - \hat{\nu}_k (\bar{\mathbf{x}}_k - \hat{\mathbf{m}}_k) \hat{\Psi}_k^{-1} (\bar{\mathbf{x}}_k - \hat{\mathbf{m}}_k) \\
 &\quad \text{(using (B.75) from Lemma B.2)}. \tag{B.72}
 \end{aligned}$$

To obtain (B.72), we have to use the following Lemma B.2.

Lemma B.2. *We can compute the expectations w.r.t. the variational distributions of the parameters as follows:-re*

$$\log \tilde{\mathbf{\Gamma}}_k \equiv \mathbb{E}_{q_{\mathbf{\Gamma}_k}} [\log |\mathbf{\Gamma}_k|] = \sum_{l=1}^L \psi \left(\frac{\hat{\nu}_k + 1 - l}{2} \right) + L \log 2 + \log \left| \hat{\Psi}_k \right|, \tag{B.73}$$

$$\mathbb{E}_{q_{\mathbf{c}_k, \mathbf{\Gamma}_k}} \left[(\mathbf{x}_n - \mathbf{c}_k)^\top \mathbf{\Gamma}_k^{-1} (\mathbf{x}_n - \mathbf{c}_k) \right] = L \hat{\lambda}_k^{-1} + \hat{\nu}_k (\mathbf{x}_n - \hat{\mathbf{m}}_k) \hat{\Psi}_k^{-1} (\mathbf{x}_n - \hat{\mathbf{m}}_k). \tag{B.74}$$

Furthermore, for each $k \in [K]$, it holds that

$$\sum_{n=1}^N q_{z_n}(k) \left[\hat{\nu}_k (\mathbf{x}_n - \hat{\mathbf{m}}_k) \hat{\Psi}_k^{-1} (\mathbf{x}_n - \hat{\mathbf{m}}_k) \right] = N_k \left[\hat{\nu}_k \text{Tr} \left(\mathbf{S}_k \hat{\Psi}_k^{-1} \right) + \hat{\nu}_k (\bar{\mathbf{x}}_k - \hat{\mathbf{m}}_k) \hat{\Psi}_k^{-1} (\bar{\mathbf{x}}_k - \hat{\mathbf{m}}_k) \right]. \quad (\text{B.75})$$

Proof of (5.14)

Given a chosen truncated value $K \in \mathbb{N}^*$, it holds that

$$\mathbb{E}_{q_{\Theta}} \left[\log p(\Theta; \hat{\phi}) \right] = \sum_{k=1}^{K-1} \mathbb{E}_{q_{\Theta}} \left[\log p(\tau_k | \alpha, \sigma) \right] + \mathbb{E}_{q_{\Theta}} \left[\log p(\alpha, \sigma | \hat{s}_1, \hat{s}_2, \hat{a}) \right] + \sum_{k=1}^K \mathbb{E}_{q_{\Theta}} \left[\log p(\mathbf{c}_k, \mathbf{\Gamma}_k; \hat{\rho}_k) \right].$$

Note that $\mathbb{E}_{q_{\Theta}} \left[\log p(\tau_k | \alpha, \sigma) \right]$ and $\mathbb{E}_{q_{\Theta}} \left[\log p(\alpha, \sigma | \hat{s}_1, \hat{s}_2, \hat{a}) \right]$ are calculated in the same way as in [Theorem 5.1](#).

Finally, we have to compute the remaining term

$$\begin{aligned} \mathbb{E}_{q_{\Theta}} \left[\log p(\mathbf{c}_k, \mathbf{\Gamma}_k; \hat{\rho}_k) \right] &= \mathbb{E}_{q_{\mathbf{c}_k, \mathbf{\Gamma}_k}} \left[-\frac{L}{2} \log(2\pi) - \frac{1}{2} \log \hat{\lambda}_k^{-L} |\mathbf{\Gamma}_k| - \frac{1}{2} (\mathbf{c}_k - \hat{\mathbf{m}}_k)^\top (\hat{\lambda}_k^{-1} \mathbf{\Gamma}_k)^{-1} (\mathbf{c}_k - \hat{\mathbf{m}}_k) \right] \\ &\quad + \mathbb{E}_{q_{\mathbf{\Gamma}_k}} \left[\log \mathcal{W} \left(\mathbf{\Gamma}_k^{-1} | \hat{\Psi}_k^{-1}, \hat{\nu}_k \right) \right] \\ &= -\frac{1}{2} L \log(2\pi) + \frac{1}{2} L \log \hat{\lambda}_k - \frac{1}{2} \mathbb{E}_{q_{\mathbf{\Gamma}_k}} \left[\log |\mathbf{\Gamma}_k| \right] - \frac{1}{2} \hat{\lambda}_k \mathbb{E}_{q_{\mathbf{c}_k, \mathbf{\Gamma}_k}} \left[(\mathbf{c}_k - \hat{\mathbf{m}}_k)^\top \mathbf{\Gamma}_k^{-1} (\mathbf{c}_k - \hat{\mathbf{m}}_k) \right] \\ &\quad + \mathbb{E}_{q_{\mathbf{\Gamma}_k}} \left[\log B \left(\hat{\Psi}_k^{-1}, \hat{\nu}_k \right) + \frac{\hat{\nu}_k - L - 1}{2} \log |\mathbf{\Gamma}_k^{-1}| - \frac{1}{2} \text{Tr} \left(\hat{\Psi}_k \mathbf{\Gamma}_k^{-1} \right) \right] \\ &= \frac{1}{2} L \log \left(\frac{\hat{\lambda}_k}{2\pi} \right) - \frac{1}{2} \mathbb{E}_{q_{\mathbf{\Gamma}_k}} \left[\log |\mathbf{\Gamma}_k| \right] - \frac{1}{2} \hat{\lambda}_k \mathbb{E}_{q_{\mathbf{c}_k, \mathbf{\Gamma}_k}} \left[(\mathbf{c}_k - \hat{\mathbf{m}}_k)^\top \mathbf{\Gamma}_k^{-1} (\mathbf{c}_k - \hat{\mathbf{m}}_k) \right] \\ &\quad + \log B \left(\hat{\Psi}_k^{-1}, \hat{\nu}_k \right) - \frac{\hat{\nu}_k - L - 1}{2} \mathbb{E}_{q_{\mathbf{\Gamma}_k}} \left[\log |\mathbf{\Gamma}_k| \right] - \frac{1}{2} \text{Tr} \left(\hat{\Psi}_k \mathbb{E}_{q_{\mathbf{\Gamma}_k}} \left[\mathbf{\Gamma}_k^{-1} \right] \right) \\ &= \frac{1}{2} L \log \left(\frac{\hat{\lambda}_k}{2\pi} \right) - \frac{1}{2} \hat{\lambda}_k \left[L \hat{\lambda}_k^{-1} + \hat{\nu}_k (\hat{\mathbf{m}}_k - \hat{\mathbf{m}}_k)^\top \hat{\Psi}_k^{-1} (\hat{\mathbf{m}}_k - \hat{\mathbf{m}}_k) \right] \\ &\quad + \log B \left(\hat{\Psi}_k^{-1}, \hat{\nu}_k \right) - \frac{\hat{\nu}_k - L}{2} \log \tilde{\mathbf{\Gamma}}_k - \frac{1}{2} \hat{\nu}_k \text{Tr} \left(\hat{\Psi}_k \hat{\Psi}_k^{-1} \right) \text{ (using Lemma B.2)} \\ &= \frac{1}{2} L \log \left(\frac{\hat{\lambda}_k}{2\pi} \right) - \frac{L}{2} - \frac{L}{2} \hat{\nu}_k + \log B \left(\hat{\Psi}_k^{-1}, \hat{\nu}_k \right) - \frac{\hat{\nu}_k - L}{2} \log \tilde{\mathbf{\Gamma}}_k, \end{aligned}$$

where

$$\log \tilde{\mathbf{\Gamma}}_k \equiv \mathbb{E}_{q_{\mathbf{\Gamma}_k}} \left[\log |\mathbf{\Gamma}_k| \right] = \sum_{l=1}^L \psi \left(\frac{\hat{\nu}_k + 1 - l}{2} \right) + L \log 2 + \log \left| \hat{\Psi}_k \right|. \quad (\text{B.76})$$

Proof of (5.16)

We have

$$\mathbb{E} \left[\log q(\Theta) \right] = \mathbb{E} \left[\log q_{\alpha, \sigma}(\alpha, \sigma) \right] + \sum_{k=1}^{K-1} \mathbb{E} \left[\log q_{\tau_k}(\tau_k) \right] + \sum_{k=1}^K \mathbb{E} \left[\log q_{\mathbf{c}_k, \mathbf{\Gamma}_k}(\mathbf{c}_k, \mathbf{\Gamma}_k) \right]. \quad (\text{B.77})$$

Note that these terms involving expectations of the logs of the q distributions simply represent the negative entropies of those distributions. In particular, the first two terms are calculated in the same way as in [Theorem 5.1](#).

Similarly, we obtain

$$\begin{aligned}
& \mathbb{E} [\log q_{\mathbf{c}_k, \mathbf{\Gamma}_k}(\mathbf{c}_k, \mathbf{\Gamma}_k)] \\
&= \mathbb{E} \left[\log \mathcal{N}_L(\mathbf{c}_k \mid \widehat{\mathbf{m}}_k, \widehat{\lambda}_k^{-1} \mathbf{\Gamma}_k) \right] + \mathbb{E} \left[\log \mathcal{W}(\mathbf{\Gamma}_k^{-1} \mid \widehat{\Psi}_k^{-1}, \widehat{\nu}_k) \right] \\
&= -\mathbb{H} \left[\mathcal{N}_L(\mathbf{c}_k \mid \widehat{\mathbf{m}}_k, \widehat{\lambda}_k^{-1} \mathbf{\Gamma}_k) \right] - \mathbb{H} \left[\mathcal{W}(\mathbf{\Gamma}_k^{-1} \mid \widehat{\Psi}_k^{-1}, \widehat{\nu}_k) \right] \\
&= \frac{L}{2} \log \frac{\widehat{\lambda}_k}{2\pi} + \frac{1}{2} \mathbb{E} [\log |\mathbf{\Gamma}_k|] - \frac{L}{2} + \log B(\widehat{\Psi}_k^{-1}, \widehat{\nu}_k) - \frac{\widehat{\nu}_k - L - 1}{2} \mathbb{E} [\log |\mathbf{\Gamma}_k|] - \frac{\widehat{\nu}_k L}{2} \\
&= \frac{L}{2} \log \frac{\widehat{\lambda}_k}{2\pi} - \frac{L}{2} + \log B(\widehat{\Psi}_k^{-1}, \widehat{\nu}_k) - \frac{\widehat{\nu}_k - L}{2} \log \widetilde{\mathbf{\Gamma}}_k - \frac{\widehat{\nu}_k L}{2}. \tag{B.78}
\end{aligned}$$

References

- [1] Antoine Deleforge, Florence Forbes, and Radu Horaud. High-dimensional regression with gaussian mixtures and partially-latent response variables. *Statistics and Computing*, 25(5):893–911, 2015. (Cited on pages 3, 5, 6, 7, 37, and 38.)
- [2] Lei Xu, Michael Jordan, and Geoffrey E Hinton. An Alternative Model for Mixtures of Experts. In G Tesauro, D Touretzky, and T Leen, editors, *Advances in Neural Information Processing Systems*, volume 7. MIT Press, 1995. (Cited on page 3.)
- [3] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991. Publisher: MIT Press. (Cited on page 3.)
- [4] Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural computation*, 6(2):181–214, 1994. Publisher: MIT Press. (Cited on page 3.)
- [5] G J McLachlan and D Peel. *Finite Mixture Models*. John Wiley & Sons, 2000. (Cited on page 3.)
- [6] TrungTin Nguyen. *Model Selection and Approximation in High-dimensional Mixtures of Experts Models: from Theory to Practice*. PhD Thesis, Normandie Université, December 2021. (Cited on page 3.)
- [7] TrungTin Nguyen, Hien Duy Nguyen, Faicel Chamroukhi, and Florence Forbes. A non-asymptotic approach for model selection via penalization in high-dimensional mixture of experts models. *Electronic Journal of Statistics*, 16(2):4742 – 4822, 2022. Publisher: Institute of Mathematical Statistics and Bernoulli Society. (Cited on pages 3, 4, 7, and 25.)
- [8] TrungTin Nguyen, Dung Ngoc Nguyen, Hien Duy Nguyen, and Faicel Chamroukhi. A non-asymptotic theory for model selection in high-dimensional mixture of experts via joint rank and variable selection. *Preprint hal-03984011*, February 2023. (Cited on pages 3 and 4.)
- [9] Wenxin Jiang and Martin A Tanner. Hierarchical mixtures-of-experts for exponential family regression models: approximation and maximum likelihood estimation. *Annals of Statistics*, pages 987–1011, 1999. Publisher: JSTOR. (Cited on page 3.)
- [10] Andriy Norets. Approximation of conditional densities by smooth mixtures of regressions. *The Annals of Statistics*, 38(3):1733 – 1766, 2010. Publisher: Institute of Mathematical Statistics. (Cited on page 3.)
- [11] Hien D Nguyen, Luke R Lloyd-Jones, and Geoffrey J McLachlan. A universal approximation theorem for mixture-of-experts models. *Neural computation*, 28(12):2585–2593, 2016. Publisher: MIT Press. (Cited on page 3.)

-
- [12] Hien D Nguyen, Faicel Chamroukhi, and Florence Forbes. Approximation results regarding the multiple-output Gaussian gated mixture of linear experts model. *Neurocomputing*, 366:208–214, 2019. (Cited on page 3.)
- [13] Hien Duy Nguyen, TrungTin Nguyen, Faicel Chamroukhi, and Geoffrey John McLachlan. Approximations of conditional probability density functions in Lebesgue spaces via mixture of experts models. *Journal of Statistical Distributions and Applications*, 8(1):13, August 2021. (Cited on page 3.)
- [14] Nhat Ho, Chiao-Yu Yang, and Michael I. Jordan. Convergence Rates for Gaussian Mixtures of Experts. *Journal of Machine Learning Research*, 23(323):1–81, 2022. (Cited on page 3.)
- [15] Dat Do, Linh Do, and XuanLong Nguyen. Strong identifiability and parameter learning in regression with heterogeneous response. *arXiv preprint arXiv:2212.04091*, 2022. (Cited on page 3.)
- [16] Christopher R. Genovese and Larry Wasserman. Rates of convergence for the Gaussian mixture sieve. *The Annals of Statistics*, 28(4):1105 – 1127, 2000. Publisher: Institute of Mathematical Statistics. (Cited on page 3.)
- [17] Alexander Rakhlin, Dmitry Panchenko, and Sayan Mukherjee. Risk bounds for mixture density estimation. *ESAIM: PS*, 9:220–229, 2005. (Cited on page 3.)
- [18] XuanLong Nguyen. Convergence of latent mixing measures in finite and infinite mixture models. *The Annals of Statistics*, 41(1):370–400, 2013. Publisher: Institute of Mathematical Statistics. (Cited on pages 3 and 22.)
- [19] Nhat Ho and XuanLong Nguyen. Convergence rates of parameter estimation for some weakly identifiable finite mixtures. *The Annals of Statistics*, 44(6):2726 – 2755, 2016. Publisher: Institute of Mathematical Statistics and Bernoulli Society. (Cited on page 3.)
- [20] Nhat Ho and XuanLong Nguyen. On strong identifiability and convergence rates of parameter estimation in finite mixtures. *Electronic Journal of Statistics*, 10(1):271–307, 2016. Publisher: The Institute of Mathematical Statistics and the Bernoulli Society. (Cited on pages 3 and 22.)
- [21] TrungTin Nguyen, Hien Duy Nguyen, Faicel Chamroukhi, and Geoffrey J. McLachlan. Approximation by finite mixtures of continuous density functions that vanish at infinity. *Cogent Mathematics & Statistics*, 7(1):1750861, January 2020. Publisher: Cogent OA. (Cited on page 3.)
- [22] TrungTin Nguyen, Faicel Chamroukhi, Hien D. Nguyen, and Geoffrey J. McLachlan. Approximation of probability density functions via location-scale finite mixtures in Lebesgue spaces. *Communications in Statistics - Theory and Methods*, pages 1–12, May 2022. Publisher: Taylor & Francis. (Cited on page 3.)
- [23] S E Yuksel, J N Wilson, and P D Gader. Twenty Years of Mixture of Experts. *IEEE Transactions on Neural Networks and Learning Systems*, 23(8):1177–1193, 2012. (Cited on page 3.)
- [24] Saeed Masoudnia and Reza Ebrahimpour. Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42(2):275–293, 2014. (Cited on page 3.)
- [25] Hien D Nguyen and Faicel Chamroukhi. Practical and theoretical aspects of mixture-of-experts modeling: An overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4):e1246, 2018. Publisher: Wiley Online Library. (Cited on page 3.)
- [26] Gilles Celeux, Sylvia Frühwirth-Schnatter, and Christian P Robert. Model selection for mixture models—perspectives and strategies. In *Handbook of mixture analysis*, pages 117–154. Chapman and Hall/CRC, 2019. (Cited on page 3.)

- [27] H Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974. (Cited on page 3.)
- [28] Gideon Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978. Publisher: Institute of Mathematical Statistics. (Cited on page 3.)
- [29] Florence Forbes, Hien Duy Nguyen, TrungTin Nguyen, and Julyan Arbel. Mixture of expert posterior surrogates for approximate Bayesian computation. In *JDS 2022 - 53èmes Journées de Statistique de la Société Française de Statistique (SFdS)*, Lyon, France, June 2022. (Cited on page 3.)
- [30] Florence Forbes, Hien Duy Nguyen, TrungTin Nguyen, and Julyan Arbel. Summary statistics and discrepancy measures for approximate Bayesian computation via surrogate posteriors. *Statistics and Computing*, 32(5):85, October 2022. (Cited on page 3.)
- [31] C. Biernacki, G. Celeux, and G. Govaert. Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(7):719–725, 2000. (Cited on page 3.)
- [32] Lucien Birgé and Pascal Massart. Minimal penalties for Gaussian model selection. *Probability Theory and Related Fields*, 138(1):33–73, 2007. Publisher: Springer. (Cited on pages 4 and 29.)
- [33] Sylvain Arlot. Minimal penalties and the slope heuristics: a survey. *Journal de la Société Française de Statistique*, 160(3):1–106, 2019. (Cited on page 4.)
- [34] TrungTin Nguyen, Faicel Chamroukhi, Hien Duy Nguyen, and Florence Forbes. Non-asymptotic model selection in block-diagonal mixture of polynomial experts models. *Preprint. arXiv:2104.08959*, May 2021. arXiv: 2104.08959. (Cited on page 4.)
- [35] TrungTin Nguyen, Faicel Chamroukhi, Hien Duy Nguyen, and Florence Forbes. Model selection by penalization in mixture of experts models with a non-asymptotic approach. In *JDS 2022 - 53èmes Journées de Statistique de la Société Française de Statistique (SFdS)*, Lyon, France, June 2022. (Cited on page 4.)
- [36] TrungTin Nguyen, Hien D. Nguyen, Faicel Chamroukhi, and Geoffrey J. McLachlan. Non-asymptotic oracle inequalities for the Lasso in high-dimensional mixture of experts. *arXiv:2009.10622*, February 2023. (Cited on page 4.)
- [37] Hien D. Nguyen, Daniel Fryer, and Geoffrey J. McLachlan. Order selection with confidence for finite mixture models. *Journal of the Korean Statistical Society*, October 2022. (Cited on page 4.)
- [38] Lucas Kock, Nadja Klein, and David J. Nott. Variational inference and sparsity in high-dimensional deep Gaussian mixture models. *Statistics and Computing*, 32(5):70, September 2022. (Cited on page 4.)
- [39] Judith Rousseau and Kerrie Mengersen. Asymptotic behaviour of the posterior distribution in overfitted mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(5):689–710, 2011. (Cited on page 4.)
- [40] Florence Forbes, Alexis Arnaud, Benjamin Lemasson, and Emmanuel Barbier. Component elimination strategies to fit mixtures of multiple scale distributions. In *RSSDS 2019 - Research School on Statistics and Data Science*, volume 1150 of *Communications in Computer and Information Science*, pages 81–95, Melbourne, Australia, July 2019. Springer. (Cited on page 4.)
- [41] Nils Lid Hjort, Chris Holmes, Peter Müller, and Stephen G. Walker, editors. *Bayesian Nonparametrics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2010. (Cited on page 4.)

- [42] Subhashis Ghosal and Aad Van der Vaart. *Fundamentals of nonparametric Bayesian inference*, volume 44. Cambridge University Press, 2017. (Cited on pages 4, 8, and 40.)
- [43] Jeffrey W. Miller and Matthew T. Harrison. Inconsistency of Pitman-Yor Process Mixtures for the Number of Components. *Journal of Machine Learning Research*, 15(96):3333–3370, 2014. (Cited on page 4.)
- [44] F Ascolani, A Lijoi, G Rebaudo, and G Zanella. Clustering consistency with Dirichlet process mixtures. *Biometrika*, September 2022. (Cited on page 4.)
- [45] Steven N. MacEachern and Peter Müller. Estimating Mixture of Dirichlet Process Models. *Journal of Computational and Graphical Statistics*, 7(2):223–238, 1998. (Cited on page 4.)
- [46] Peter J. Green and Sylvia Richardson. Modelling Heterogeneity With and Without the Dirichlet Process. *Scandinavian Journal of Statistics*, 28(2):355–375, 2001. (Cited on page 4.)
- [47] Aritra Guha, Nhat Ho, and XuanLong Nguyen. On posterior contraction of parameters and interpretability in Bayesian mixture modeling. *Bernoulli*, 27(4):2159 – 2188, 2021. Publisher: Bernoulli Society for Mathematical Statistics and Probability. (Cited on pages 4, 22, 23, 24, 25, and 28.)
- [48] Hongliang Lü, Julyan Arbel, and Florence Forbes. Bayesian nonparametric priors for hidden Markov random fields. *Statistics and Computing*, 30(4):1015–1035, July 2020. (Cited on pages 4, 8, 9, and 13.)
- [49] J. Arbel, G. Kon Kam King, A. Lijoi, L. Nieto-Barajas, and I. Prünster. BNPdensity: Bayesian nonparametric mixture modelling in R. *Australian & New Zealand Journal of Statistics*, 63(3):542–564, 2021. (Cited on page 4.)
- [50] Riccardo Corradin, Antonio Canale, and Bernardo Nipoti. BNPmix: An R Package for Bayesian Nonparametric Modeling via Pitman-Yor Mixtures. *Journal of Statistical Software*, 100(15):1 – 33, November 2021. (Cited on page 4.)
- [51] J.-B. Durand, F. Forbes, C.D. Phan, L. Truong, H.D. Nguyen, and F. Dama. Bayesian nonparametric spatial prior for traffic crash risk mapping: A case study of Victoria, Australia. *Australian & New Zealand Journal of Statistics*, 64(2):171–204, 2022. (Cited on pages 4, 9, and 20.)
- [52] David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017. Publisher: Taylor & Francis. (Cited on page 4.)
- [53] Ker-Chau Li. Sliced Inverse Regression for Dimension Reduction. *Journal of the American Statistical Association*, 86(414):316–327, February 1991. (Cited on page 5.)
- [54] Salvatore Ingrassia, Simona C Minotti, and Giorgio Vittadini. Local Statistical Modeling via a Cluster-Weighted Approach with Elliptical Distributions. *Journal of Classification*, 29(3):363–401, 2012. (Cited on page 7.)
- [55] Thomas S. Ferguson. A Bayesian Analysis of Some Nonparametric Problems. *The Annals of Statistics*, 1(2):209 – 230, 1973. Publisher: Institute of Mathematical Statistics. (Cited on page 7.)
- [56] Jim Pitman and Marc Yor. The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *The Annals of Probability*, 25(2):855–900, April 1997. ISBN: 0091-1798 Publisher: JSTOR. (Cited on page 8.)
- [57] Hemant Ishwaran and Lancelot F James. Gibbs Sampling Methods for Stick-Breaking Priors. *Journal of the American Statistical Association*, 96(453):161–173, March 2001. Publisher: Taylor & Francis. (Cited on page 8.)

- [58] Pierpaolo De Blasi, Stefano Favaro, Antonio Lijoi, Ramsés H. Mena, Igor Prünster, and Matteo Ruggiero. Are Gibbs-type priors the most natural generalization of the Dirichlet process? *IEEE transactions on pattern analysis and machine intelligence*, 37(2):212–229, 2015. ISBN: 0162-8828 Publisher: IEEE. (Cited on page 8.)
- [59] Stefano Favaro, Antonio Lijoi, Consuelo Nava, Bernardo Nipoti, Igor Pruenster, and Yee Whye Teh. On the stick-breaking representation for homogeneous NRMIs. *Bayesian Analysis*, 11(3):697–724, 2016. ISBN: 1936-0975 Publisher: International Society for Bayesian Analysis. (Cited on page 8.)
- [60] David M. Blei and Michael I. Jordan. Variational inference for Dirichlet process mixtures. *Bayesian analysis*, 1(1):121–143, 2006. ISBN: 1936-0975 Publisher: International Society for Bayesian Analysis. (Cited on pages 8, 12, 39, 40, and 43.)
- [61] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data Via the EM Algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, September 1977. Publisher: John Wiley & Sons, Ltd. (Cited on page 9.)
- [62] Geoffrey J. McLachlan and Thriyambakam Krishnan. *The EM algorithm and extensions*. Wiley, New York, USA, 1997. (Cited on page 9.)
- [63] Radford M. Neal and Geoffrey E. Hinton. A View of the EM Algorithm that Justifies Incremental, Sparse, and other Variants. In Michael I. Jordan, editor, *Learning in Graphical Models*, pages 355–368. Springer Netherlands, Dordrecht, 1998. (Cited on page 10.)
- [64] Matthew J Beal and Zoubin Ghahramani. The variational Bayesian EM algorithm for incomplete data: with application to scoring graphical model structures. *Bayesian statistics*, 7(453-464):210, 2003. (Cited on page 10.)
- [65] Matthew James Beal. *Variational algorithms for approximate Bayesian inference*. PhD thesis, University of London, University College London (United Kingdom), 2003. (Cited on page 10.)
- [66] Lotfi Chaari, Thomas Vincent, Florence Forbes, Michel Dojat, and Philippe Ciuciu. Fast joint detection-estimation of evoked brain activity in event-related fMRI using a variational approach. *IEEE transactions on Medical Imaging*, 32(5):821–837, 2013. ISBN: 0278-0062 Publisher: IEEE. (Cited on page 10.)
- [67] Christopher M Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006. (Cited on pages 10, 12, 16, 39, 41, 44, and 48.)
- [68] Yee Whye Teh, Michael I. Jordan, Matthew J. Beal, and David M. Blei. Hierarchical Dirichlet Processes. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006. Publisher: Taylor & Francis. (Cited on pages 11 and 14.)
- [69] Yee Whye Teh. A Hierarchical Bayesian Language Model Based On Pitman-Yor Processes. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 985–992, Sydney, Australia, July 2006. Association for Computational Linguistics. (Cited on pages 11 and 14.)
- [70] Yee Whye Teh and Michael I. Jordan. Hierarchical Bayesian nonparametric models with applications. In Nils Lid Hjort, Chris Holmes, Peter Müller, and Stephen G.Editors Walker, editors, *Bayesian Nonparametrics*, Cambridge Series in Statistical and Probabilistic Mathematics, pages 158–207. Cambridge University Press, 2010. (Cited on page 11.)
- [71] Chong Wang, John Paisley, and David M. Blei. Online Variational Inference for the Hierarchical Dirichlet Process. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15

- of *Proceedings of Machine Learning Research*, pages 752–760, Fort Lauderdale, FL, USA, April 2011. PMLR. (Cited on page 12.)
- [72] Chao Yuan and Claus Neubauer. Variational Mixture of Gaussian Process Experts. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008. (Cited on page 16.)
- [73] Chen Luo and Shiliang Sun. Variational Mixtures of Gaussian Processes for Classification. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 4603–4609, 2017. (Cited on page 16.)
- [74] Di Wu and Jinwen Ma. A Two-Layer Mixture Model of Gaussian Process Functional Regressions and Its MCMC EM Algorithm. *IEEE Transactions on Neural Networks and Learning Systems*, 29(10):4894–4904, 2018. (Cited on page 16.)
- [75] Trung Nguyen and Edwin Bonilla. Fast Allocation of Gaussian Process Experts. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 145–153, Beijing, China, June 2014. PMLR. Issue: 1. (Cited on page 16.)
- [76] James Ross and Jennifer Dy. Nonparametric Mixture of Gaussian Processes with Constraints. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1346–1354, Atlanta, Georgia, USA, June 2013. PMLR. Issue: 3. (Cited on page 16.)
- [77] Markus Svensén and Christopher M. Bishop. Robust Bayesian mixture modelling. *Neurocomputing*, 64:235–252, 2005. (Cited on page 20.)
- [78] Cédric Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, 2003. (Cited on page 22.)
- [79] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009. (Cited on page 22.)
- [80] Fengnan Gao and Aad van der Vaart. Posterior contraction rates for deconvolution of Dirichlet-Laplace mixtures. *Electronic Journal of Statistics*, 10(1):608 – 627, 2016. Publisher: Institute of Mathematical Statistics and Bernoulli Society. (Cited on page 22.)
- [81] Faicel Chamroukhi, Allou Samé, Gérard Govaert, and Patrice Aknin. A hidden process regression model for functional data description. Application to curve discrimination. *Neurocomputing*, 73(7-9):1210–1221, 2010. Publisher: Elsevier. (Cited on page 25.)
- [82] Lucie Montuelle and Erwan Le Pennec. Mixture of Gaussian regressions model with logistic weights, a penalized maximum likelihood approach. *Electronic Journal of Statistics*, 8(1):1661–1695, 2014. Publisher: The Institute of Mathematical Statistics and the Bernoulli Society. (Cited on page 25.)
- [83] Subhashis Ghosal, Jayanta K. Ghosh, and Aad W. van der Vaart. Convergence rates of posterior distributions. *The Annals of Statistics*, 28(2):500 – 531, 2000. Publisher: Institute of Mathematical Statistics. (Cited on page 28.)
- [84] Fengshuo Zhang and Chao Gao. Convergence rates of variational posterior distributions. *The Annals of Statistics*, 48(4):2180 – 2207, 2020. Publisher: Institute of Mathematical Statistics. (Cited on page 28.)
- [85] Willem Kruijer, Judith Rousseau, and Aad van der Vaart. Adaptive Bayesian density estimation with location-scale mixtures. *Electronic Journal of Statistics*, 4(none):1225 – 1257, 2010. Publisher: Institute of Mathematical Statistics and Bernoulli Society. (Cited on page 28.)

-
- [86] Weining Shen, Surya T Tokdar, and Subhashis Ghosal. Adaptive Bayesian multivariate density estimation with Dirichlet mixtures. *Biometrika*, 100(3):623–640, 2013. Publisher: Oxford University Press. (Cited on page 28.)
- [87] Andriy Norets and Debdeep Pati. Adaptive Bayesian estimation of conditional densities. *Econometric Theory*, 33(4):980–1012, 2017. Publisher: Cambridge University Press. (Cited on page 28.)
- [88] D. M. Titterington and Bo Wang. Convergence properties of a general algorithm for calculating variational Bayesian estimates for a normal mixture model. *Bayesian Analysis*, 1(3):625 – 650, 2006. Publisher: International Society for Bayesian Analysis. (Cited on page 29.)
- [89] Francisco Ruiz and Michalis Titsias. A Contrastive Divergence for Combining Variational Inference and MCMC. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5537–5545. PMLR, June 2019. (Cited on page 29.)