

A non-asymptotic approach for model selection via penalization in high-dimensional mixture of experts models*

TrungTin Nguyen[†]

*Normandie Univ, UNICAEN, CNRS, LMNO, 14000 Caen, France.
Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK,
Inria Grenoble Rhone-Alpes, 655 av. de l'Europe, 38335 Montbonnot, France
e-mail: trung-tin.nguyen@inria.fr
url: <http://trung-tinnguyen.github.io/>*

Hien Duy Nguyen

*Department of Mathematical and Physical Sciences,
La Trobe University, Bundoora Melbourne 3066, Victoria Australia.
School of Mathematics and Physics,
University of Queensland, St. Lucia, Queensland, Australia
e-mail: h.nguyen7@uq.edu.au
url: <https://hiendn.github.io/>*

Faïcel Chamroukhi

*Normandie Univ, UNICAEN, CNRS, LMNO, 14000 Caen, France.
IRT SystemX, Palaiseau, France
e-mail: Faïcel.chamroukhi@irt-systemx.fr
url: <https://chamroukhi.com/>*

Florence Forbes

*Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK,
Inria Grenoble Rhone-Alpes, 655 av. de l'Europe, 38335 Montbonnot, France
e-mail: florence.forbes@inria.fr
url: <http://mistis.inrialpes.fr/people/forbes/>*

Abstract: Mixture of experts (MoE) are a popular class of statistical and machine learning models that have gained attention over the years due to their flexibility and efficiency. In this work, we consider Gaussian-gated localized MoE (GLoME) and block-diagonal covariance localized MoE (BLoME) regression models to present nonlinear relationships in heterogeneous data with potential hidden graph-structured interactions between high-dimensional predictors. These models pose difficult statistical estimation and model selection questions, both from a computational and theoretical perspective. This paper is devoted to the study of the problem

*This work is partially supported by the French Ministry of Higher Education and Research (MESRI), French National Research Agency (ANR) grant SMILES ANR-18-CE40-0014, Australian Research Council grant number DP180101192, and the Inria LANDER project.

[†]Corresponding author.

of model selection among a collection of GLoME or BLoME models characterized by the number of mixture components, the complexity of Gaussian mean experts, and the hidden block-diagonal structures of the covariance matrices, in a penalized maximum likelihood estimation framework. In particular, we establish non-asymptotic risk bounds that take the form of weak oracle inequalities, provided that lower bounds for the penalties hold. The good empirical behavior of our models is then demonstrated on synthetic and real datasets.

MSC2020 subject classifications: Primary 62H30, 62E17; secondary 62H12.

Keywords and phrases: Mixture of experts, linear cluster-weighted models, mixture of regressions, Gaussian locally-linear mapping models, clustering, oracle inequality, model selection, penalized maximum likelihood, block-diagonal covariance matrix, graphical Lasso.

Received May 2021.

Contents

1	Introduction	4744
1.1	Mixture of experts models	4744
1.2	Goals	4745
1.3	Related literature	4746
1.4	Main contributions	4746
1.5	Main challenges	4748
1.6	Organization	4749
2	Notation and framework	4749
2.1	GLoME and BLoME models	4749
2.1.1	Gaussian gating networks	4752
2.1.2	Gaussian experts	4752
2.2	High-dimensional regression via GLLiM and BLLiM models	4754
2.3	Collection of GLoME and BLoME models	4756
2.4	Loss functions and penalized maximum likelihood estimator	4757
3	Main results of finite-sample oracle inequalities	4758
3.1	Deterministic collection of GLoME models	4758
3.2	Random subcollection of BLoME models	4760
3.3	Practical application of finite-sample oracle inequalities	4761
4	Proofs of finite-sample oracle inequalities and main mathematical challenges	4763
4.1	Proof for deterministic collections of GLoME models	4763
4.1.1	A general conditional density model selection theorem for deterministic collection of models	4763
4.1.2	Sketch of the proof for LinBoSGaME models	4765
4.1.3	Our contributions on the proof for GLoME models	4768
4.2	Proof of random subcollection of BLoME models	4770
4.2.1	A model selection theorem for MLE among a random subcollection	4770
4.2.2	Sketch of the proof for BLoME models	4771

4.2.3	Our contributions on BLoME models	4771
5	Numerical experiments	4773
5.1	The procedure	4773
5.2	Simulated data sets	4774
5.3	Ethanol data set	4789
6	Conclusion and perspectives	4789
A	Our contributions on lemma proofs	4796
A.1	Proof of Lemma 2.1: Inverse regression trick	4796
A.1.1	Elliptically symmetric distributions	4797
A.1.2	Relation between forward and inverse regression	4799
A.2	Proof of Lemma 4.5: Bracketing entropy of Gaussian gating networks	4801
A.2.1	Proof of Lemma A.5	4802
A.2.2	Proof of Lemma A.6	4803
A.3	Proof of Lemma 4.7: Bracketing entropy of Gaussian experts with 1-block covariance matrices	4806
A.4	Proof of Lemma 4.11	4807
A.5	Proof of Lemma 4.12: Bracketing entropy of Gaussian experts with multi-block-diagonal covariance matrices	4808
A.5.1	Proof of Lemma A.8	4809
A.5.2	Proof of Lemma A.9	4810
	Acknowledgments	4815
	References	4816

1. Introduction

1.1. Mixture of experts models

Mixture of experts (MoE) models, originally introduced as neural network architectures in [48, 50], are flexible models that generalize the classical finite mixture models as well as finite mixtures of regression models [66, Sec. 5.13]. Their flexibility comes from allowing the mixture weights (or the gating functions) to depend on the explanatory variables, along with the component densities (or experts). In the context of regression, MoE models with Gaussian experts and softmax or normalized Gaussian gating functions are the most popular choices and are powerful tools for modeling more complex nonlinear relationships between outputs (responses) and inputs (predictors) that arise from different subpopulations. The popularity of these conditional mixture density models arise largely due to their universal approximation properties, which have been studied in [49, 81, 74, 45, 73, 75], and which improve upon approximation capabilities of unconditional finite mixture models, as studied in [41, 85, 80, 43, 44, 78, 77]. Detailed reviews on the practical and theoretical aspects of MoE models can be found in [101, 59, 72, 76].

In this paper, we wish to investigate MoE models with Gaussian experts and normalized Gaussian gating functions for clustering and regression, first introduced by [99], which extended the original MoE models of [48]. From hereon

in, we refer to these models as *Gaussian-gated localized MoE* (GLoME) models and *block-diagonal covariance for localized MoE* (BLoME) models, and we provide their precise definitions in Section 2.1. Furthermore, the original MoE models with softmax gating functions will be referred to as *linear-combination-of-bounded-functions softmax-gated MoE* (LinBoSGaME) regression models. In particular, we simply refer to affine instances of LinBoSGaME models as *softmax-gated MoE* (SGaME) regression models. It is worth pointing out that the BLoME models generalize GLoME models by utilizing a parsimonious covariance structure, via block-diagonal structures for covariance matrices in the Gaussian experts. It is also interesting to point out that supervised *Gaussian locally-linear mapping* (GLLiM) and *block-diagonal covariance for Gaussian locally-linear mapping* (BLLiM) models in [26] and [34] are affine instances of GLoME and BLoME models, respectively, where the latter pair consider general linear combination of bounded functions instead of affine for mean functions of Gaussian experts, in the prior pair.

One of the main disadvantages of LinBoSGaME models is the difficulty of applying an expectation–maximization (EM) algorithm [27, 65], which requires an internal iterative numerical optimization procedure (*e.g.*, iteratively-reweighted least squares, Newton-Raphson algorithm) to update the softmax parameters. To overcome this problem, we instead use the Gaussian gating network that enables us to link BLoME with hierarchical Gaussian mixture model. Then, the maximization with respect to the parameters of the gating network can be solved analytically with the EM algorithm framework, which decreases the computational complexity of the estimation routine. Furthermore, we can then also make use of well-established theoretical results for finite mixture models.

We note here that both GLoME and BLoME models have been thoroughly studied in the statistics and machine learning literatures in many different guises, including localized MoE [86, 87, 69, 15], normalized Gaussian networks [90], MoE modeling of priors in Bayesian nonparametric regression [83, 82], cluster-weighted modeling [47], deep mixture of linear inverse regressions [55], hierarchical Gaussian locally linear mapping structured mixture (HGLLiM) model [95], multiple-output Gaussian gated mixture of linear experts [73], and approximate Bayesian computation with surrogate posteriors using GLLiM [39].

1.2. Goals

Our main goal is to learn potentially nonlinear relationships between a multivariate output and a high-dimensional input issued from a heterogeneous population. This involves performing regression, clustering and model selection, simultaneously. While estimation can be performed using standard EM algorithms, it crucially depends and requires data-driven hyperparameter choices, including the number of mixture components (or clusters), the degree of complexity of each Gaussian expert’s mean function, and the hidden block-diagonal structures of the covariance matrices. Recall that hyperparameter choices of data-driven learning algorithms belong to the class of model selection problems,

which have attracted a lot of attention in statistics and machine learning over the past 50 years [1, 58, 2, 60, 3]. Specifically, given a set of models, how do we select the one with the lowest possible risk from the data? Note that penalization is one of the main strategies proposed for model selection. It suggests choosing the estimator that minimizes the sum of its empirical risk and some penalty terms corresponding to the fit of the model to the data while avoiding overfitting.

1.3. Related literature

In general, model selection for MoE models is often performed using the Akaike information criterion (AIC) [1], the Bayesian information criterion (BIC) [91] or the BIC-like approximation of integrated classification likelihood (ICL-BIC) [10]. An important limitation of these criteria, however, is that they are only valid asymptotically. This implies that there is no finite-sample guarantee when using AIC, BIC or ICL-BIC, to choose between different levels of complexity. Their use in small samples is, therefore, ad hoc. To overcome such difficulties, [11] proposed a novel approach, called the slope heuristics, supported by a non-asymptotic oracle inequality. This method leads to an optimal data-driven choice of multiplicative constants for the penalties. Recent reviews and practical issues concerning the slope heuristic can be found in [8, 3] and the references given therein.

It should be emphasized that a general model selection result, originally established by [60, Theorem 7.11], ensures a penalized criterion leads to a good model selection and that the penalty is known only up to multiplicative constants and is proportional to the dimensions of models. In particular, these multiplicative constants can be calibrated by the slope heuristic approach in a finite-sample setting. Then, in the spirit of the methods based on concentration inequalities developed in [60, 61, 22, 23], a number of finite-sample oracle type inequalities have been established for the least absolute shrinkage and selection operator (LASSO) [94] and general penalized maximum likelihood estimators (PMLE). These results include the works for high-dimensional Gaussian graphical models [33], Gaussian mixture model selection [62, 63, 64], finite mixture regression models [68, 28, 29, 31, 32], and LinBoSGaME models, outside the high-dimensional setting [70].

1.4. Main contributions

To the best of our knowledge, no attempt has been made in the literature to develop a finite-sample oracle inequality for the framework of MoE regression models for high-dimensional data. In this paper, our first original contribution is to provide finite-sample oracle inequalities, Theorems 3.1 and 3.2 for high-dimensional GLoME and BLoME models, respectively. More specifically, we establish non-asymptotic risk bounds that take the form of weak oracle

inequalities, provided that the lower bounds on the penalties are true, in high-dimensional GLoME and BLoME models, based on an inverse regression strategy and block-diagonal structures of the Gaussian covariance matrices experts. Unlike traditional criteria such as AIC, BIC, or ICL-BIC, which are based on asymptotic theory or a Bayesian approach, our contributed non-asymptotic risk bounds allow the number n of observations to be fixed while the dimensionality and cardinality of the models, characterized by the number of covariates and the size of the response, are allowed to grow with respect to n and can be much larger than n .

In particular, our oracle inequalities show that the Jensen–Kullback–Leibler loss performance of our PMLEs is comparable to that of oracle models if we take sufficiently large constants in front of the penalties, whose shapes are only known up to multiplicative constants and proportional to the dimensions of the models. These theoretical justifications for the shapes of the penalties motivate us to make use of the slope heuristic criterion to select several hyperparameters, including the number of mixture components, the degree of polynomial mean functions, and the potential hidden block-diagonal structures of the covariance matrices of the multivariate predictors.

Moreover, in Theorem 3.1, we extend a corollary of [70, Theorem 1], which can be verified via Lemma 1 from [75], which makes explicit the relationship between softmax and Gaussian gating classes.

Another significant contribution is the numerical experiments for simulated and real data sets in Section 5 that support our theoretical results, and the statistical study of non-asymptotic model selection in GLoME models. To the best of our knowledge, instead of using classical asymptotic approaches for model selection such as AIC, BIC and ICL-BIC, we are the first to illustrate that the slope heuristic works well for MoEs with Gaussian experts and normalized Gaussian gating networks such as GLLiM, BLLiM, GLoME, and BLoME models. Note that our main objective here is to investigate how well the empirical tensorized Kullback–Leibler divergence between the true model (s_0^*) and the selected model $\hat{s}_{\mathbf{m}}^*$ follows the finite-sample oracle inequality of Theorem 3.1, as well as the convergence rate $\mathcal{O}(n^{-1})$ of the error upper bound. Therefore, we focus on 1-dimensional data sets. Beyond the statistical estimation and model selection purposes considered here, the dimensionality reduction capability of GLLiM and BLLiM models in high-dimensional regression data, can be found in [26, Section 6] and [34, Sections 3 and 4], respectively.

Specifically, besides the important theoretical issues regarding the tightness of upper bounds, how to integrate a priori information, and a minimax analysis of our proposed PMLEs, our finite-sample oracle inequalities and the corresponding illustrative numerical experiments help to partially answer the following two important questions raised in the field of MoE regression models: (1) What is the number of mixture components K to choose, given the sample size n ; and (2) Whether it is better to use a few complex experts or a combination of many simple experts, given the total number of parameters. Note that such problems are also considered in the work of [67, Proposition 1], where the authors provided

some qualitative guidance and only suggested a practical method for choosing K and d , involving a complexity penalty or cross-validation. Furthermore, their model is only non-regularized maximum-likelihood estimation, and thus is not suitable for the high-dimensional setting.

Lastly, we emphasize that although the finite-sample oracle inequalities compare performances of our estimators with the best model in the collection, they also allow us to approximate a rich class of conditional densities if we take enough clusters and/or a high enough degree d of polynomials in the mean of the functions of Gaussian experts, in the context of mixture of Gaussian experts [49, 67, 74, 45, 75]. This leads to the upper bounds of the risks being small.

1.5. Main challenges

For the Gaussian gating parameters, the technique for handling the logistic weights in the SGaME models of [70] is not directly applicable to the GLoME or BLoME framework, due to the quadratic form of the canonical link. Therefore, we have to propose a *reparameterization trick*¹ to bound the metric entropy of the Gaussian gating parameters space; see (36) and Lemma 4.5 for more details.

To work with conditional density estimation in the BLLiM and BLoME models, it is natural to make use of a general conditional density model selection theorem from [22, 23], see also Theorem 4.1 for its adaption to our context. However, it is worth mentioning that since the collection of models constructed by the BLLiM model [34] or by some appropriate procedures for BLoME models in practice is usually random, we have to utilize a model selection theorem for MLE among a random subcollection (cf. [29, Theorem 5.1] and [33, Theorem 7.3]).

In particular, our collections of BLoME models must satisfy certain regularity assumptions, see Section 4.2.3 for details, which are not easy to verify due to the complexity of BLoME models and technical reasons. For BLoME models, the main difficulty in proving our oracle inequality lies in bounding the bracketing entropy of the Gaussian gating functions and Gaussian experts with block-diagonal covariance matrices. To solve the first problem, we need to use our previous trick of reparametrizing the Gaussian gating parameter space. For the second one, based on some ideas of Gaussian mixture models from [41, 62], we provide a novel extension for standard MoE models with Gaussian gating networks. Note that our important contributions also extend the recent result on block-diagonal covariance matrices in [33], which is only developed for Gaussian graphical models.

In addition, there are two possible ways to decompose the bracketing entropy of collection models based on different distances. As mentioned in Appendix B.2.1 from [70], their decomposition boils down to assuming that the predictor

¹Note that we only use this nomenclature to perform a change of variables of the Gaussian gating parameters space of GLoME models via the logistic weights of SGaME models. This reparameterization trick does not stand for the well-known one of variational autoencoders (VAEs) in deep learning literature (see [52], for more details).

domain is bounded. And this limiting assumption can then be relaxed by using a smaller distance, namely a tensorized extension of the Hellinger distance $d^{\otimes n}$ in (22), but bounding the corresponding bracketing entropy becomes much more difficult. Nevertheless, it is important to point out that we are able to weaken this limiting assumption by using the smaller distance: $d^{\otimes n}$, for the bracketing entropy of our collection of BLoME models, see Lemma 4.10 for more details. This reinforces our original contributions concerning the control of bracketing entropy of BLoME models.

1.6. Organization

The rest of this paper is organized as follows. In Section 2, we present the notation and framework for GLoME, and BLoME models and their special cases, GLLiM, and BLLiM models. In Section 3, we state the main results of this paper: finite-sample oracle inequalities satisfied by the PMLEs. Section 4 is devoted to the proofs and main mathematical challenges in establishing our new finite-sample oracle inequalities. We experimentally evaluate our new results in simulated and real datasets in Section 5. Some conclusions and perspectives are provided in Section 6. The proofs of technical lemmas can be found in Appendix A.

2. Notation and framework

From now on, for any $L \in \mathbb{N}^*$, we write $[L]$ to denote the set $\{1, \dots, L\}$. We are interested in estimating the law of the random variable $\mathbf{Y} = (\mathbf{Y}_j)_{j \in [L]}$, conditionally on $\mathbf{X} = (\mathbf{X}_j)_{j \in [D]}$. Here and subsequently, given any $n \in \mathbb{N}^*$, $(\mathbf{X}_{[n]}, \mathbf{Y}_{[n]}) := (\mathbf{X}_i, \mathbf{Y}_i)_{i \in [n]}$ denotes a random sample, and \mathbf{x} and \mathbf{y} stands for the observed values of the random variables \mathbf{X} and \mathbf{Y} , respectively. The following assumptions will be needed throughout the paper. 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 $\mathbf{X}_{[n]}$, $\mathbf{Y}_{[n]}$ are independent, and each \mathbf{Y} follows a law with true (but unknown) PDF $s_0(\cdot | \mathbf{X} = \mathbf{x})$, which is approximated via MoE models. For a matrix \mathbf{A} , let $m(\mathbf{A})$ and $M(\mathbf{A})$ be, respectively, the modulus of the smallest and largest eigenvalues of \mathbf{A} .

2.1. GLoME and BLoME models

Motivated by an inverse regression framework, where the role of predictor and response variables will be exchanged such that \mathbf{Y} becomes the input and \mathbf{X} plays the role of a multivariate output, we consider GLoME and BLoME models with inverse conditional probability density functions (PDFs) of the form (1). This construction goes back to the work of [57, 99, 26, 84], and is very useful in a high-dimensional regression context, where typically $D \gg L$. In this way, we

define PDF of \mathbf{X} conditional on $\mathbf{Y} = \mathbf{y}$ as follows:

$$s_{\psi_{K,d}}(\mathbf{x} | \mathbf{y}) = \sum_{k=1}^K g_k(\mathbf{y}; \boldsymbol{\omega}) \phi_D(\mathbf{x}; \mathbf{v}_{k,d}(\mathbf{y}), \boldsymbol{\Sigma}_k), \quad (1)$$

$$g_k(\mathbf{y}; \boldsymbol{\omega}) = \frac{\pi_k \phi_L(\mathbf{y}; \mathbf{c}_k, \boldsymbol{\Gamma}_k)}{\sum_{j=1}^K \pi_j \phi_L(\mathbf{y}; \mathbf{c}_j, \boldsymbol{\Gamma}_j)}. \quad (2)$$

Here, $g_k(\cdot; \boldsymbol{\omega})$ and $\phi_D(\cdot; \mathbf{v}_{k,d}(\cdot), \boldsymbol{\Sigma}_k)$, $k \in [K]$, $K \in \mathbb{N}^*$, $d \in \mathbb{N}^*$, are called normalized Gaussian gating functions (or Gaussian gating networks) and Gaussian experts, respectively. Furthermore, we decompose the parameters of the model as follows: $\boldsymbol{\psi}_{K,d} = (\boldsymbol{\omega}, \mathbf{v}_d, \boldsymbol{\Sigma}) \in \boldsymbol{\Omega}_K \times \boldsymbol{\Upsilon}_{K,d} \times \mathbf{V}_K =: \boldsymbol{\Psi}_{K,d}$, $\boldsymbol{\omega} = (\boldsymbol{\pi}, \mathbf{c}, \boldsymbol{\Gamma}) \in (\boldsymbol{\Pi}_{K-1} \times \mathbf{C}_K \times \mathbf{V}'_K) =: \boldsymbol{\Omega}_K$, $\boldsymbol{\pi} = (\pi_k)_{k \in [K]}$, $\mathbf{c} = (\mathbf{c}_k)_{k \in [K]}$, $\boldsymbol{\Gamma} = (\boldsymbol{\Gamma}_k)_{k \in [K]}$, $\mathbf{v}_d = (\mathbf{v}_{k,d})_{k \in [K]} \in \boldsymbol{\Upsilon}_{K,d}$, and $\boldsymbol{\Sigma} = (\boldsymbol{\Sigma}_k)_{k \in [K]} \in \mathbf{V}_K$. Note that $\boldsymbol{\Pi}_{K-1} = \left\{ (\pi_k)_{k \in [K]} \in (\mathbb{R}^+)^K, \sum_{k=1}^K \pi_k = 1 \right\}$ is a $K-1$ dimensional probability simplex, \mathbf{C}_K is a set of K -tuples of mean vectors of size $L \times 1$, \mathbf{V}'_K is a set of K -tuples of elements in \mathcal{S}_L^{++} , where \mathcal{S}_L^{++} denotes the collection of symmetric positive definite matrices on \mathbb{R}^L , $\boldsymbol{\Upsilon}_{K,d}$ is a set of K -tuples of mean functions from \mathbb{R}^L to \mathbb{R}^D depending on a degree d (e.g., a degree of polynomials), and \mathbf{V}_K is a set containing K -tuples from \mathcal{S}_D^{++} .

Recall that GLLiM and BLLiM models are affine instances of GLoME and BLoME models, and are especially useful for high-dimensional regression data, since there exist link functions between the inverse and forward conditional density; see Figure 1 for comprehensive classification and nomenclature of standard MoE regression models with Gaussian gating networks. Note that the principle of inverse regression is only useful when the functions $\mathbf{v}_{k,d}(\mathbf{y})$ are linear, because there is no explicit way to express the law of $\mathbf{Y} | \mathbf{X}$ from that of $\mathbf{X} | \mathbf{Y}$ for higher degree of polynomials. However, to have more consistent notations with the previous affine results of GLLiM, BLLiM models from [26, 34], we decide to use the inverse regression frameworks instead of the forward one.

In the BLoME model, we wish to make use of block-diagonal covariance structures by replacing $\boldsymbol{\Sigma}_k$ and \mathbf{V}_K with $\boldsymbol{\Sigma}_k(\mathbf{B}_k)$ and $\mathbf{V}_K(\mathbf{B})$, defined in (3), respectively (see, e.g., [34, 33]). This block-diagonal structures for covariance matrices are not only used to trade-off between complexity and sparsity but is also motivated by some real applications, where we want to perform prediction on data sets with heterogeneous observations and hidden graph-structured interactions between covariates; for instance, for gene expression data sets in which conditionally on the phenotypic response, genes interact with few other genes only, i.e., there are small modules of correlated genes (see [34, 33] for more details). To be more precise, for $k \in [K]$, we decompose $\boldsymbol{\Sigma}_k(\mathbf{B}_k)$ into G_k blocks, $G_k \in \mathbb{N}^*$, and we denote by $d_k^{[g]}$ the set of variables into the g th group, for $g \in [G_k]$, and by $\text{card}\left(d_k^{[g]}\right)$ the number of variables in the corresponding set. Then, we define $\mathbf{B}_k = \left(d_k^{[g]}\right)_{g \in [G_k]}$ to be a block structure for the cluster k , and $\mathbf{B} = (\mathbf{B}_k)_{k \in [K]}$ to be the covariate indexes into each group for each clus-

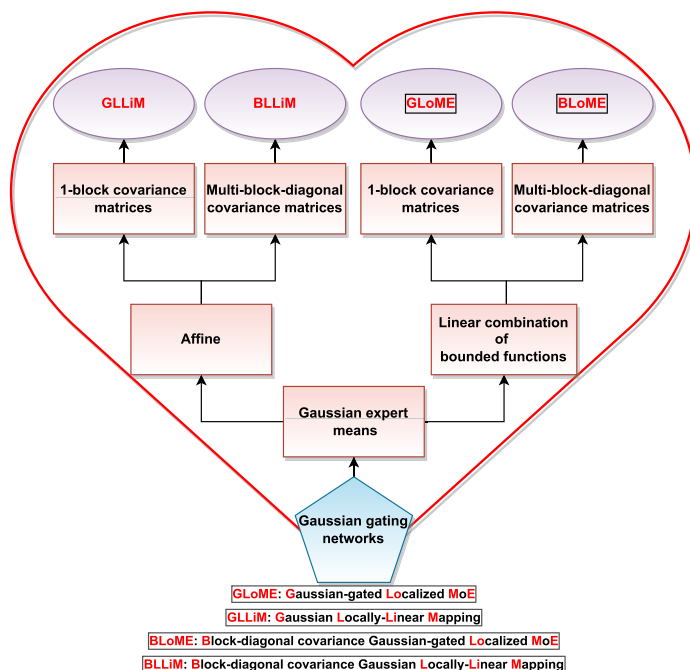


FIG 1. A comprehensive classification and nomenclature of standard MoE regression models with Gaussian gating networks.

ter. In this way, to construct the block-diagonal covariance matrices, up to a permutation, we make the following definition: $\mathbf{V}_K(\mathbf{B}) = (\mathbf{V}_k(\mathbf{B}_k))_{k \in [K]}$, for every $k \in [K]$,

$$\mathbf{V}_k(\mathbf{B}_k) = \left\{ \begin{array}{l} \Sigma_k(\mathbf{B}_k) \in \mathcal{S}_D^{++} \\ \Sigma_k(\mathbf{B}_k) = \mathbf{P}_k \begin{pmatrix} \Sigma_k^{[1]} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \Sigma_k^{[2]} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \dots & \Sigma_k^{[G_k]} \end{pmatrix} \mathbf{P}_k^{-1}, \\ \Sigma_k^{[g]} \in \mathcal{S}_{\text{card}(d_k^{[g]})}^{++}, \forall g \in [G_k] \end{array} \right\}. \quad (3)$$

Here, \mathbf{P}_k corresponds to the permutation leading to a block-diagonal matrix in cluster k . It is worth pointing out that outside the blocks, all coefficients of the matrix are zeros and we also authorize reordering of the blocks: e.g., $\{(1, 3); (2, 4)\}$ is identical to $\{(2, 4); (1, 3)\}$, and the permutation inside blocks: e.g., the partition of 4 variables into blocks $\{(1, 3); (2, 4)\}$ is the same as the partition $\{(3, 1); (4, 2)\}$.

It is interesting to point out that the BLLiM framework aims to model a sample of high-dimensional regression data issued from a heterogeneous population

with hidden graph-structured interaction between covariates. In particular, the BLLiM model is considered as a good candidate for performing model-based clustering and for predicting the response in situations affected by the “curse of dimensionality” phenomenon, where the number of parameters could be larger than the sample size. Indeed, to deal with high-dimensional regression problems, the BLLiM model is based on an inverse regression strategy, which inverts the role of the high-dimensional predictor and the multivariate response. Therefore, the number of parameters to estimate is drastically reduced. More precisely, BLLiM utilizes GLLiM, described in [25, 26], in conjunction with a block-diagonal structure hypothesis on the residual covariance matrices, to make a trade-off between complexity and sparsity. This prediction model is fully parametric and highly interpretable. For instance, it can be used for the analysis of transcriptomic data in molecular biology to classify observations or predict phenotypic states; see *e.g.*, [42, 71, 56]. Indeed, if the predictor variables are gene expression data measured by microarrays or by the RNA-seq technologies, and the response is a phenotypic variable, the BLLiM model not only provides clusters of individuals based on the relation between gene expression data and the phenotype, but also implies a gene regulatory network specific to each cluster of individuals (see [34] for more details).

In order to establish our finite-sample oracle inequalities, Theorems 3.1 and 3.2, we need to explicitly impose some classical boundedness conditions on the parameter space.

2.1.1. Gaussian gating networks

We shall restrict our study to bounded Gaussian gating parameter vectors. Specifically, we assume that there exist deterministic positive constants a_π , A_c , a_Γ , A_Γ , such that $\boldsymbol{\omega}$ belongs to $\tilde{\Omega}_K$, where

$$\tilde{\Omega}_K = \left\{ \boldsymbol{\omega} \in \Omega_K : \forall k \in [K], \|\mathbf{c}_k\|_\infty \leq A_c, \right. \\ \left. a_\Gamma \leq m(\Gamma_k) \leq M(\Gamma_k) \leq A_\Gamma, a_\pi \leq \pi_k \right\}. \quad (4)$$

2.1.2. Gaussian experts

Following the same structure for the Gaussian mean experts from [70], the set $\Upsilon_{K,d}$ will be chosen as a tensor product of compact sets of moderate dimension (*e.g.*, a set of polynomials of degree smaller than d , whose coefficients are smaller in absolute values than T_Υ). Then, $\Upsilon_{K,d}$ is defined as a linear combination of a finite set of bounded functions whose coefficients belong to a compact set. This general setting includes polynomial bases when the covariates are bounded, Fourier bases on an interval, as well as suitably renormalized wavelet dictionaries. More specifically, $\Upsilon_{K,d} = \otimes_{k \in [K]} \Upsilon_{k,d} =: \Upsilon_{k,d}^K$, where $\Upsilon_{k,d} = \Upsilon_{b,d}$,

$\forall k \in [K]$, and

$$\mathbf{\Upsilon}_{b,d} = \left\{ \mathbf{y} \mapsto \left(\sum_{i=1}^d \alpha_i^{(j)} \varphi_{\mathbf{r},i}(\mathbf{y}) \right)_{j \in [D]} =: (\mathbf{v}_{d,j}(\mathbf{y}))_{j \in [D]} : \|\alpha\|_\infty \leq T_{\mathbf{r}} \right\}. \quad (5)$$

Here, the subscript b stands for “bounded functions”, $d \in \mathbb{N}^*$, $T_{\mathbf{r}} \in \mathbb{R}^+$, and $(\varphi_{\mathbf{r},i})_{i \in [d]}$ is a collection of bounded functions on \mathcal{Y} . In particular, we focus on the bounded \mathcal{Y} case and assume that $\mathcal{Y} = [0, 1]^L$, without loss of generality. In this case, $\varphi_{\mathbf{r},i}$ can be chosen as monomials with maximum (non-negative) degree d : $\mathbf{y}^{\mathbf{r}} = \prod_{l=1}^L \mathbf{y}_l^{\mathbf{r}_l}$. Recall that a multi-index $\mathbf{r} = (\mathbf{r}_l)_{l \in [L]}$, $\mathbf{r}_l \in \mathbb{N}^* \cup \{0\}$, $\forall l \in [L]$, is an L -tuple of nonnegative integers. We define $|\mathbf{r}| = \sum_{l=1}^L \mathbf{r}_l$ and the number $|\mathbf{r}|$ is called the order or degree of $\mathbf{y}^{\mathbf{r}}$. Then, $\mathbf{\Upsilon}_{K,d} = \mathbf{\Upsilon}_{p,d}^K$, where

$$\mathbf{\Upsilon}_{p,d} = \left\{ \mathbf{y} \mapsto \left(\sum_{|\mathbf{r}|=0}^d \alpha_{\mathbf{r}}^{(j)} \mathbf{y}^{\mathbf{r}} \right)_{j \in [D]} =: (\mathbf{v}_{d,j}(\mathbf{y}))_{j \in [D]} : \|\alpha\|_\infty \leq T_{\mathbf{r}} \right\}. \quad (6)$$

Here, the subscript p stands for “polynomial”.

For GLoME models, note that any covariance matrix Σ_k can be decomposed into the form $B_k \mathbf{P}_k \mathbf{A}_k \mathbf{P}_k^\top$, such that $B_k = |\Sigma_k|^{1/D}$ is a positive scalar corresponding to the volume, \mathbf{P}_k is the matrix of eigenvectors of Σ_k and \mathbf{A}_k the diagonal matrix of normalized eigenvalues of Σ_k ; $B_- \in \mathbb{R}^+$, $B_+ \in \mathbb{R}^+$, $\mathcal{A}(\lambda_-, \lambda_+)$ is a set of diagonal matrices \mathbf{A}_k , such that $|\mathbf{A}_k| = 1$ and $\forall i \in [D]$, $\lambda_- \leq (\mathbf{A}_k)_{i,i} \leq \lambda_+$, where $\lambda_-, \lambda_+ \in \mathbb{R}$; and $SO(D)$ is the special orthogonal group of dimension D . In this way, we obtain what is known as the classical covariance matrix sets described by [20] for Gaussian parsimonious clustering models, defined by

$$\mathbf{V}_K = \left\{ (B_k \mathbf{P}_k \mathbf{A}_k \mathbf{P}_k^\top)_{k \in [K]} : \forall k \in [K], B_- \leq B_k \leq B_+, \right. \\ \left. \mathbf{P}_k \in SO(D), \mathbf{A}_k \in \mathcal{A}(\lambda_-, \lambda_+) \right\}. \quad (7)$$

For the block-diagonal covariances of Gaussian experts from BLoME models, we assume that there exist some positive constants λ_m and λ_M such that, for every $k \in [K]$,

$$\tilde{\mathbf{V}}_K(\mathbf{B}) = \left\{ (\Sigma_k(\mathbf{B}_k))_{k \in [K]} \in \mathbf{V}_K(\mathbf{B}) : \right. \\ \left. 0 < \lambda_m \leq m(\Sigma_k(\mathbf{B}_k)) \leq M(\Sigma_k(\mathbf{B}_k)) \leq \lambda_M \right\}. \quad (8)$$

Note that this is a typical assumption and is also used in the block-diagonal covariance selection for Gaussian graphical models of [33].

Next, characterizations of GLLiM and BLLiM models are provided in Section 2.2.

2.2. High-dimensional regression via GLLiM and BLLiM models

The GLLiM and BLLiM models, as originally introduced in [26, 34], are used to capture the nonlinear relationship between the response and the set of covariates from a high-dimensional regression data, imposed by a potential hidden graph structured interaction, typically in the case when $D \gg L$, by the K locally affine mappings:

$$\mathbf{Y} = \sum_{k=1}^K \mathbb{I}(Z = k) (\mathbf{A}_k^* \mathbf{X} + \mathbf{b}_k^* + \mathbf{E}_k^*). \quad (9)$$

Here, \mathbb{I} is an indicator function and Z is a latent variable capturing a cluster relationship, such that $Z = k$ if \mathbf{Y} originates from cluster $k \in [K]$. Cluster specific affine transformations are defined by matrices $\mathbf{A}_k^* \in \mathbb{R}^{L \times D}$ and vectors $\mathbf{b}_k^* \in \mathbb{R}^L$. Furthermore, \mathbf{E}_k^* are error terms capturing both the reconstruction error due to the local affine approximations and the observation noise in \mathbb{R}^L .

Following the common assumption that \mathbf{E}_k^* is a zero-mean Gaussian variable with covariance matrix $\mathbf{\Sigma}_k^* \in \mathbb{R}^{L \times L}$, it holds that

$$p(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}, Z = k; \boldsymbol{\psi}_K^*) = \Phi_L(\mathbf{y}; \mathbf{A}_k^* \mathbf{x} + \mathbf{b}_k^*, \mathbf{\Sigma}_k^*), \quad (10)$$

where we denote by $\boldsymbol{\psi}_K^*$ the vector of model parameters and Φ_L is the PDF of a Gaussian distribution of dimension L . In order to allow the affine transformations to be local, \mathbf{X} is defined as a mixture of K Gaussian components as follows:

$$p(\mathbf{X} = \mathbf{x} \mid Z = k; \boldsymbol{\psi}_K^*) = \Phi_D(\mathbf{x}; \mathbf{c}_k^*, \mathbf{\Gamma}_k^*), p(Z = k; \boldsymbol{\psi}_K^*) = \pi_k^*, \quad (11)$$

where $\mathbf{c}_k^* \in \mathbb{R}^D$, $\mathbf{\Gamma}_k^* \in \mathbb{R}^{D \times D}$, $\boldsymbol{\pi}^* = (\pi_k^*)_{k \in [K]} \in \Pi_{K-1}^*$, and Π_{K-1}^* is the $K-1$ dimensional probability simplex. Then, according to formulas for conditional multivariate Gaussian variables and the following hierarchical decomposition

$$p(\mathbf{Y} = \mathbf{y}, \mathbf{X} = \mathbf{x}; \boldsymbol{\psi}_K^*) = \sum_{k=1}^K \pi_k^* \Phi_D(\mathbf{x}; \mathbf{c}_k^*, \mathbf{\Gamma}_k^*) \Phi_L(\mathbf{y}; \mathbf{A}_k^* \mathbf{x} + \mathbf{b}_k^*, \mathbf{\Sigma}_k^*),$$

we obtain the following *forward conditional density* [26]:

$$p(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}; \boldsymbol{\psi}_K^*) = \sum_{k=1}^K \frac{\pi_k^* \Phi_D(\mathbf{x}; \mathbf{c}_k^*, \mathbf{\Gamma}_k^*)}{\sum_{j=1}^K \pi_j^* \Phi_D(\mathbf{x}; \mathbf{c}_j^*, \mathbf{\Gamma}_j^*)} \Phi_L(\mathbf{y}; \mathbf{A}_k^* \mathbf{x} + \mathbf{b}_k^*, \mathbf{\Sigma}_k^*), \quad (12)$$

where $\boldsymbol{\psi}_K^* = (\boldsymbol{\pi}^*, \boldsymbol{\theta}_K^*) \in \Pi_{K-1} \times \Theta_K^* =: \Psi_K^*$. Here, $\boldsymbol{\theta}_K^* = (\mathbf{c}_k^*, \mathbf{\Gamma}_k^*, \mathbf{A}_k^*, \mathbf{b}_k^*, \mathbf{\Sigma}_k^*)_{k \in [K]}$ and

$$\Theta_K^* = (\mathbb{R}^D \times \mathcal{S}_D^{++}(\mathbb{R}) \times \mathbb{R}^{L \times D} \times \mathbb{R}^L \times \mathcal{S}_L^{++}(\mathbb{R}))^K.$$

Without assuming anything on the structure of parameters, the dimension of the model, denoted by $\dim(\cdot)$, is defined as the total number of parameters that has to be estimated, as follows:

$$\dim(\Psi_K^*) = K \left(1 + D(L + 1) + \frac{D(D + 1)}{2} + \frac{L(L + 1)}{2} + L \right) - 1.$$

It is worth mentioning that $\dim(\Psi_K)$ is typically very large compared to the sample size (see, e.g., [26, 34, 84] for more details in their real data sets) whenever D is large and $D \gg L$. Furthermore, it is more realistic to make assumption on the residual covariance matrices Σ_k^* of error vectors \mathbf{E}_k^* rather than on Γ_k^* (cf. [26, Section 3]). This justifies the use of the inverse regression trick from [26], which leads a drastic reduction in the number of parameters to be estimated.

More specifically, in (12), the roles of input and response variables were exchanged such that \mathbf{Y} becomes the covariates and \mathbf{X} plays the role of the multivariate response. Therefore, its corresponding *inverse conditional density* is defined as a *Gaussian locally-linear mapping* (GLLiM) model, based on the previous hierarchical Gaussian mixture model, as follows:

$$p(Z = k; \psi_k) = \pi_k, \tag{13}$$

$$p(\mathbf{Y} = \mathbf{y} \mid Z = k; \psi_K) = \Phi_L(\mathbf{y}; \mathbf{c}_k, \Gamma_k), \tag{14}$$

$$p(\mathbf{X} = \mathbf{x} \mid \mathbf{Y} = \mathbf{y}, Z = k; \psi_K) = \Phi_D(\mathbf{x}; \mathbf{A}_k \mathbf{y} + \mathbf{b}_k, \Sigma_k), \tag{15}$$

$$p(\mathbf{X} = \mathbf{x} \mid \mathbf{Y} = \mathbf{y}; \psi_K) = \sum_{k=1}^K \frac{\pi_k \Phi_L(\mathbf{y}; \mathbf{c}_k, \Gamma_k)}{\sum_{j=1}^K \pi_j \Phi_L(\mathbf{y}; \mathbf{c}_j, \Gamma_j)} \Phi_D(\mathbf{x}; \mathbf{A}_k \mathbf{y} + \mathbf{b}_k, \Sigma_k), \tag{16}$$

where Σ_k is a $D \times D$ covariance structure (usually diagonal in GLLiM models, chosen to reduce the number of parameters) automatically learned from data and ψ_K is the set of parameters, denoted by $\psi_K = (\boldsymbol{\pi}, \boldsymbol{\theta}_K) \in \Pi_{K-1} \times \Theta_K =: \Psi_K$. It is important to note that the BLLiM model imposes block-diagonal structures on $(\Sigma_k)_{k \in [K]}$ to make a trade-off between complexity and sparsity.

The following interesting feature of both GLLiM and BLLiM models is proved in Appendix A.1.

Lemma 2.1. *The parameter ψ_K^* in the forward conditional PDF, defined in (12), can then be deduced from ψ_K in (16) via the following one-to-one correspondence:*

$$\boldsymbol{\theta}_K = \left(\begin{array}{c} \mathbf{c}_k \\ \Gamma_k \\ \mathbf{A}_k \\ \mathbf{b}_k \\ \Sigma_k \end{array} \right)_{k \in [K]} \mapsto \left(\begin{array}{c} \mathbf{C}_k^* \\ \Gamma_k^* \\ \mathbf{A}_k^* \\ \mathbf{b}_k^* \\ \Sigma_k^* \end{array} \right)_{k \in [K]} = \left(\begin{array}{c} \mathbf{A}_k \mathbf{c}_k + \mathbf{b}_k \\ \Sigma_k + \mathbf{A}_k \Gamma_k \mathbf{A}_k^\top \\ \Sigma_k^* \mathbf{A}_k^\top \Sigma_k^{-1} \\ \Sigma_k^* (\Gamma_k^{-1} \mathbf{c}_k - \mathbf{A}_k^\top \Sigma_k^{-1} \mathbf{b}_k) \\ (\Gamma_k^{-1} + \mathbf{A}_k^\top \Sigma_k^{-1} \mathbf{A}_k)^{-1} \end{array} \right)_{k \in [K]} \in \Theta_K^*, \tag{17}$$

with $\boldsymbol{\pi}^* \equiv \boldsymbol{\pi}$.

2.3. Collection of GLoME and BLoME models

For GLoME, we only need to choose the degree of polynomials d and the number of components K among finite sets $\mathcal{D}_{\Upsilon} = [d_{\max}]$ and $\mathcal{K} = [K_{\max}]$, respectively, where $d_{\max} \in \mathbb{N}^*$ and $K_{\max} \in \mathbb{N}^*$ may depend on the sample size n . We wish to estimate the unknown inverse conditional density s_0 by inverse conditional densities belonging to the following collection of models $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$. Here, we define $\mathcal{M} = \{(K, d) : K \in \mathcal{K}, d \in \mathcal{D}_{\Upsilon}\}$, and

$$S_{\mathbf{m}} = \left\{ (\mathbf{x}, \mathbf{y}) \mapsto s_{\psi_{K,d}}(\mathbf{x} | \mathbf{y}) := s_{\mathbf{m}}(\mathbf{x} | \mathbf{y}) : \right. \\ \left. \psi_{K,d} = (\boldsymbol{\omega}, \mathbf{v}_d, \boldsymbol{\Sigma}) \in \tilde{\boldsymbol{\Omega}}_K \times \boldsymbol{\Upsilon}_{K,d} \times \mathbf{V}_K \right\}, \quad (18)$$

where $s_{\psi_{K,d}}$, $\tilde{\boldsymbol{\Omega}}_K$, $\boldsymbol{\Upsilon}_{K,d}$ and \mathbf{V}_K are defined previously in (1), (4), (6) (or more general (5)) and (7), respectively.

While for the BLoME model, we also need to select \mathbf{B} from a list of candidate structures $(\mathcal{B}_k)_{k \in [K]} \equiv (\mathcal{B})_{k \in [K]}$, where \mathcal{B} denotes the set of all possible partitions of the covariables indexed by $[D]$, for each cluster of individuals. The collection of BLoME models is defined as follows: $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$, $\mathcal{M} = \{(K, d, \mathbf{B}) : K \in \mathcal{K}, d \in \mathcal{D}_{\Upsilon}, \mathbf{B} \in (\mathcal{B})_{k \in [K]}\}$,

$$S_{\mathbf{m}} = \left\{ (\mathbf{x}, \mathbf{y}) \mapsto s_{\psi_{K,d,\mathbf{B}}}(\mathbf{x} | \mathbf{y}) := s_{\mathbf{m}}(\mathbf{x} | \mathbf{y}) : \right. \\ \left. \psi_{K,d,\mathbf{B}} = (\boldsymbol{\omega}, \mathbf{v}_d, \boldsymbol{\Sigma}(\mathbf{B})) \in \tilde{\boldsymbol{\Omega}}_K \times \boldsymbol{\Upsilon}_{K,d} \times \tilde{\mathbf{V}}_K(\mathbf{B}) \right\}, \quad (19)$$

where $s_{\psi_{K,d,\mathbf{B}}}$, $\tilde{\boldsymbol{\Omega}}_K$, $\boldsymbol{\Upsilon}_{K,d}$, and $\tilde{\mathbf{V}}_K(\mathbf{B})$ are defined in (1), (4), (6) (or more generally (5)), and (8), respectively. In theory, we would like to consider the whole collection of models $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$. However, the cardinality of \mathcal{B} is large; its size is a Bell number. Even for a moderate number of variables D , it is not possible to explore the set \mathcal{B} , exhaustively. We restrict our attention to a random subcollection \mathcal{B}^R of moderate size. For example, we can consider the BLLiM procedure from [34, Section 2.2].

Remark 2.2. It is worth noting that we can also define the collection of the forward GLoME and BLoME models in the same framework as in (18) and (19), respectively. For example, in GLoME models, the unknown forward conditional density s_0^* is estimated via the following collection of forward models $(S_{\mathbf{m}}^*)_{\mathbf{m} \in \mathcal{M}}$, with $\mathcal{M} = \mathcal{K} \times \mathcal{D}_{\Upsilon}$, and

$$S_{\mathbf{m}}^* = \left\{ (\mathbf{x}, \mathbf{y}) \mapsto s_{\psi_K^*}(\mathbf{y} | \mathbf{x}) := s_{\mathbf{m}}^*(\mathbf{y} | \mathbf{x}) : \right. \\ \left. \psi_K^* = (\boldsymbol{\omega}^*, \mathbf{v}_d^*, \boldsymbol{\Sigma}^*) \in \tilde{\boldsymbol{\Omega}}_K^* \times \boldsymbol{\Upsilon}_{K,d}^* \times \mathbf{V}_K^* =: \tilde{\boldsymbol{\Psi}}_{K,d}^* \right\}, \quad (20)$$

where $\tilde{\boldsymbol{\Omega}}_K^*$, $\boldsymbol{\Upsilon}_{K,d}^*$ and \mathbf{V}_K^* are defined similar to (4), (6) (or more general (5)) and (7), respectively.

Motivated by the inverse conditional densities (16) and the one-to-one correspondence in Lemma 2.1, for the sake of simplicity of notation, we only state our main Theorems 3.1 and 3.2 for the collection of inverse models $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$ in (18) and (19), respectively. However, our finite-sample oracle inequalities, Theorems 3.1 and 3.2 can be applied to the forward model $(S_{\mathbf{m}}^*)_{\mathbf{m} \in \mathcal{M}}$, established in (20), if we consider \mathbf{y} and \mathbf{x} as realizations of predictors and response variables, respectively.

2.4. Loss functions and penalized maximum likelihood estimator

In the maximum likelihood approach, the Kullback–Leibler divergence is the most natural loss function, which is defined for two densities s and t by

$$\text{KL}(s, t) = \begin{cases} \int_{\mathbb{R}^D} \ln \left(\frac{s(y)}{t(y)} \right) s(y) dy & \text{if } sdy \text{ is absolutely continuous w.r.t. } tdy, \\ +\infty & \text{otherwise.} \end{cases}$$

However, to take into account the structure of conditional densities and the random covariates $(\mathbf{Y}_{[n]})$, we consider the *tensorized Kullback–Leibler divergence* $\text{KL}^{\otimes n}$, defined as:

$$\text{KL}^{\otimes n}(s, t) = \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n \text{KL}(s(\cdot | \mathbf{Y}_i), t(\cdot | \mathbf{Y}_i)) \right], \quad (21)$$

if sdy is absolutely continuous w.r.t. tdy , and $+\infty$ otherwise. Note that if the predictors are fixed, this divergence is the classical fixed design type divergence in which there is no expectation. We refer to our result as a *weak oracle inequality*, because its statement is based on a smaller divergence, when compared to $\text{KL}^{\otimes n}$, namely the *tensorized Jensen–Kullback–Leibler divergence*:

$$\text{JKL}_{\rho}^{\otimes n}(s, t) = \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n \frac{1}{\rho} \text{KL}(s(\cdot | \mathbf{Y}_i), (1 - \rho)s(\cdot | \mathbf{Y}_i) + \rho t(\cdot | \mathbf{Y}_i)) \right],$$

with $\rho \in (0, 1)$. We note that $\text{JKL}_{\rho}^{\otimes n}$ was first used in [22, 23]. However, a version of this divergence appears explicitly with $\rho = \frac{1}{2}$ in [60], and it is also found implicitly in [12]. This loss is always bounded by $\frac{1}{\rho} \ln \frac{1}{1-\rho}$ but behaves like $\text{KL}^{\otimes n}$, when t is close to s . The main tools in the proof of such a weak oracle inequality are deviation inequalities for sums of random variables and their suprema. These tools require a boundedness assumption on the controlled functions which is not satisfied by $-\ln \frac{s_{\mathbf{m}}}{s_0}$, and thus also not satisfied by $\text{KL}^{\otimes n}$. Therefore, we consider instead the use of $\text{JKL}_{\rho}^{\otimes n}$. In particular, in general, it holds that $C_{\rho} d^{2\otimes n} \leq \text{JKL}_{\rho}^{\otimes n} \leq \text{KL}^{\otimes n}$, where $C_{\rho} = \frac{1}{\rho} \min \left(\frac{1-\rho}{\rho}, 1 \right) \left(\ln \left(1 + \frac{\rho}{1-\rho} \right) - \rho \right)$ (see [22, Prop. 1]) and $d^{2\otimes n}$ is a tensorized extension of the squared Hellinger distance $d^{2\otimes n}$, defined by

$$d^{2\otimes n}(s, t) = \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n d^2(s(\cdot | \mathbf{Y}_i), t(\cdot | \mathbf{Y}_i)) \right]. \quad (22)$$

Moreover, if we assume that, for any $\mathbf{m} \in \mathcal{M}$ and any $s_{\mathbf{m}} \in S_{\mathbf{m}}$, $s_0 d\lambda \ll s_{\mathbf{m}} d\lambda$, then (see [70, 22])

$$\frac{C_\rho}{2 + \ln \|s_0/s_{\mathbf{m}}\|_\infty} \text{KL}^{\otimes n}(s_0, s_{\mathbf{m}}) \leq \text{JKL}_\rho^{\otimes n}(s_0, s_{\mathbf{m}}). \tag{23}$$

In the context of MLE, given the collection of conditional densities $S_{\mathbf{m}}$, we aim to estimate s_0 by the conditional density $\widehat{s}_{\mathbf{m}}$ that maximizes the likelihood (conditionally to $(\mathbf{y}_i)_{i \in [n]}$) or equivalently that minimizes the negative log-likelihood (NLL), which we can write as:

$$\widehat{s}_{\mathbf{m}} = \operatorname{argmin}_{s_{\mathbf{m}} \in S_{\mathbf{m}}} \sum_{i=1}^n -\ln [s_{\mathbf{m}}(\mathbf{x}_i | \mathbf{y}_i)].$$

We should work with almost minimizer of this quantity and define a η -log-likelihood minimizer as any $\widehat{s}_{\mathbf{m}}$ that satisfies:

$$\sum_{i=1}^n -\ln [\widehat{s}_{\mathbf{m}}(\mathbf{x}_i | \mathbf{y}_i)] \leq \inf_{s_{\mathbf{m}} \in S_{\mathbf{m}}} \sum_{i=1}^n -\ln [s_{\mathbf{m}}(\mathbf{x}_i | \mathbf{y}_i)] + \eta, \tag{24}$$

where the error term η is necessary to avoid any existence issue, *e.g.*, the infimum may not be reached. Roughly speaking, the Ekeland variational principle asserts that, for any extended-valued lower semicontinuous function which is bounded below, one can add a small perturbation to ensure the existence of the minimum, see *e.g.*, [14, Chapter 2]. This framework is also used in [22, 23, 70].

As always, using the NLL of the estimate in each model as a criterion is not sufficient. It is an underestimation of the risk of the estimate and this leads to choosing models that are too complex. In the context of PMLE, by adding a suitable penalty $\text{pen}(\mathbf{m})$, one hopes to create a trade-off between good data fit and model complexity. For a given choice of $\text{pen}(\mathbf{m})$, the *selected model* $S_{\widehat{\mathbf{m}}}$ is chosen as the one whose index is an η' -almost minimizer of the sum of the NLL and this penalty:

$$\sum_{i=1}^n -\ln [\widehat{s}_{\widehat{\mathbf{m}}}(\mathbf{x}_i | \mathbf{y}_i)] + \text{pen}(\widehat{\mathbf{m}}) \leq \inf_{\mathbf{m} \in \mathcal{M}} \left\{ \sum_{i=1}^n -\ln [\widehat{s}_{\mathbf{m}}(\mathbf{x}_i | \mathbf{y}_i)] + \text{pen}(\mathbf{m}) \right\} + \eta'. \tag{25}$$

Note that $\widehat{s}_{\widehat{\mathbf{m}}}$ is then called the η' -penalized likelihood estimate and depends on both the error terms η and η' . From hereon in, the term *selected model (estimate) or best data-driven model (estimate)* is used to indicate that it satisfies the definition in (25).

3. Main results of finite-sample oracle inequalities

3.1. Deterministic collection of GLoME models

Theorem 3.1 provides a lower bound on the penalty function, $\text{pen}(\mathbf{m})$, which guarantees that the PMLE selects a best data-driven model that performs almost as well as the best model. We highlight our main contribution and briefly

establish the proof of Theorem 3.1 in Section 4.1.3. A more detail technical proof can be found in Appendices A.2 and A.3.

Theorem 3.1 (Finite-sample oracle inequality for GLoME models). *Assume that we observe $(\mathbf{x}_{[n]}, \mathbf{y}_{[n]})$, arising from an unknown conditional density s_0 . Given a collection of GLoME models, $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$, there is a constant C such that for any $\rho \in (0, 1)$, for any $\mathbf{m} \in \mathcal{M}$, $z_{\mathbf{m}} \in \mathbb{R}^+$, $\Xi = \sum_{\mathbf{m} \in \mathcal{M}} e^{-z_{\mathbf{m}}} < \infty$ and any $C_1 > 1$, there is a constant κ depending only on ρ and C_1 , such that if for every index $\mathbf{m} \in \mathcal{M}$,*

$$\text{pen}(\mathbf{m}) > \kappa [(C + \ln n) \dim(S_{\mathbf{m}}) + z_{\mathbf{m}}], \tag{26}$$

then the η' -penalized likelihood estimate $\widehat{s}_{\mathbf{m}}$, defined in (24) and (25), satisfies

$$\begin{aligned} \mathbb{E}_{\mathbf{X}_{[n]}, \mathbf{Y}_{[n]}} [\text{JKL}_{\rho}^{\otimes n}(s_0, \widehat{s}_{\mathbf{m}})] &\leq C_1 \inf_{\mathbf{m} \in \mathcal{M}} \left(\inf_{s_{\mathbf{m}} \in S_{\mathbf{m}}} \text{KL}^{\otimes n}(s_0, s_{\mathbf{m}}) + \frac{\text{pen}(\mathbf{m})}{n} \right) \\ &\quad + \frac{\kappa C_1 \Xi}{n} + \frac{\eta + \eta'}{n}. \end{aligned} \tag{27}$$

It is worth noting that Theorem 3.1 extends a corollary of [70, Theorem 1], which can be verified via Lemma 1 from [75], which makes explicit the relationship between softmax and Gaussian gating classes.

Note that we only aim to construct upper bounds and do not focus on the important question of the existence of the corresponding lower bounds or minimax adaptive estimation. However, as a special case of GLoME models, [62, 64] consider a collection of univariate Gaussian mixture models having the same and known component variance. They show that the PMLE $\widehat{s}_{\mathbf{m}}$ is minimax adaptive to the regularity β , $\beta > 0$, of univariate density classes \mathcal{H}_{β} whose logarithm of the elements is locally β -Hölder, with convergence rate $n^{-\frac{2\beta}{2\beta+1}}$ up to a power of $\ln(n)$. Although this result is stated for the Hellinger risk, it remains valid for the Kullback–Leibler divergence if we further assume that $\ln(\|s/t\|_{\infty})$ is uniformly bounded on $\cup_{\mathbf{m} \in \mathcal{M}} S_{\mathbf{m}}$, see *e.g.*, Lemma 7.23 in [60]. Note that this assumption guarantees that the Kullback–Leibler divergence and the Hellinger distance are equivalent.

A special case of GLoME models, namely model-selection (nearly-D-sparse) aggregation in mixture models, is considered by [17, 9, 89, 18, 24] and is related to our model selection results. More precisely, the authors of [89] did not consider PMLEs with Kullback–Leibler loss but only with L_2 -loss or aggregation of a finite number of densities. Similarly, the results from [17, 9] dealt with the L_2 -loss and investigate the Lasso and the Dantzig estimators, respectively, suitably adapted to the problem of density estimation. With respect to the Kullback–Leibler loss, [18] established model selection type oracle inequalities with high probability rather than in expectation. In particular, a bound in expectation with Kullback–Leibler loss in [24] is perhaps the most relevant reference to our result. They established exact oracle inequalities with a rate-optimal remainder term $((\ln K)/n)^{1/2}$, up to a possible logarithm correction, in the problem of convex aggregation when the number K of components is larger than $n^{1/2}$. It

is worth noting that they did not consider PMLEs as we do or as [62, 64], see *e.g.*, (25). Instead, the constraint that the weight vector in the mixture model belongs to the probabilistic simplex acts as a sparsity-inducing penalty. However, in their collection of mixture models, the mixture components are deterministic and chosen from a dictionary obtained on the basis of previous experiments or expert knowledge rather estimated from data. The adjective exact refers to the fact that the “bias term” $\text{KL}^{\otimes n}(s_0, s_{\mathbf{m}})$ is not multiplied by factor strictly larger than one as in our Theorem 3.1.

To the best of our knowledge, providing a minimax analysis for the PMLEs or the problem of model-selection (nearly-D-sparse) aggregation in the context of standard MoE models is still an open question. In particular, it is not trivial to extend such optimal risk bounds regarding Gaussian mixtures, see *e.g.*, [64, Theorem 2.8] or [24, Theorems 2.3 and 3.1], to standard MoE models. However, we wish to provide such a minimax analysis in future research.

In contrast to Theorem 3.1, in Theorem 3.2, we consider a random subcollection of models $\widetilde{\mathcal{M}}$, included in the whole collection \mathcal{M} . This is particularly useful in a high-dimensional context where we cannot test all the models. Therefore, we have to restrict ourselves to a smaller subcollection of models, which is then potentially random.

3.2. Random subcollection of BLoME models

Note that the constructed collection of models with block-diagonal structures for each cluster of individuals is designed, for example, by the BLLiM procedure from [34], where each collection of partitions is sorted by sparsity level. Nevertheless, our finite-sample oracle inequality, Theorem 3.2, still holds for any random subcollection of \mathcal{M} , which is constructed by some suitable tools in the framework of BLoME regression models. We highlight our main contribution and briefly establish the proof of Theorem 3.2 in Section 4.2. A more detailed technical proof can be found in Appendices A.4 and A.5 and Section 4.2.3.

Theorem 3.2 (Finite-sample oracle inequality for BLoME models). *Let $(\mathbf{x}_{[n]}, \mathbf{y}_{[n]})$ be observations coming from an unknown conditional density s_0 . For each $\mathbf{m} = (K, d, \mathbf{B}) \in (\mathcal{K} \times \mathcal{D}_{\mathbf{r}} \times \mathcal{B}) \equiv \mathcal{M}$, let $S_{\mathbf{m}}$ be define by (19). Assume that there exists $\tau > 0$ and $\epsilon_{KL} > 0$ such that, for all $\mathbf{m} \in \mathcal{M}$, one can find $\bar{s}_{\mathbf{m}} \in S_{\mathbf{m}}$, such that*

$$\text{KL}^{\otimes n}(s_0, \bar{s}_{\mathbf{m}}) \leq \inf_{t \in S_{\mathbf{m}}} \text{KL}^{\otimes n}(s_0, t) + \frac{\epsilon_{KL}}{n}, \text{ and } \bar{s}_{\mathbf{m}} \geq e^{-\tau} s_0. \quad (28)$$

Next, we construct some random subcollections $(S_{\mathbf{m}})_{\mathbf{m} \in \widetilde{\mathcal{M}}}$ of $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$ by letting $\widetilde{\mathcal{M}} \equiv (\mathcal{K} \times \mathcal{D}_{\mathbf{r}} \times \mathcal{B}^R) \subset \mathcal{M}$, where \mathcal{B}^R is a random subcollection \mathcal{B} , of moderate size. Consider the collection $(\widehat{s}_{\mathbf{m}})_{\mathbf{m} \in \widetilde{\mathcal{M}}}$ of η -log likelihood minimizers satisfying (24) for all $\mathbf{m} \in \widetilde{\mathcal{M}}$. Then, there is a constant C such that for any $\rho \in (0, 1)$, and any $C_1 > 1$, there are two constants κ and C_2 depending only on

ρ and C_1 such that, for every index, $\mathbf{m} \in \mathcal{M}$, $z_{\mathbf{m}} \in \mathbb{R}^+$, $\Xi = \sum_{\mathbf{m} \in \mathcal{M}} e^{-z_{\mathbf{m}}} < \infty$ and

$$\text{pen}(\mathbf{m}) \geq \kappa [(C + \ln n) \dim(S_{\mathbf{m}}) + (1 \vee \tau)z_{\mathbf{m}}], \tag{29}$$

the η' -penalized likelihood estimator $\widehat{s}_{\mathbf{m}}$, defined as in (25) on the subset $\widetilde{\mathcal{M}} \subset \mathcal{M}$, satisfies

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}_{[n]}, \mathbf{Y}_{[n]}} [\text{JKL}_{\rho}^{\otimes n}(s_0, \widehat{s}_{\mathbf{m}})] \\ & \leq C_1 \mathbb{E}_{\mathbf{X}_{[n]}, \mathbf{Y}_{[n]}} \left[\inf_{\mathbf{m} \in \widetilde{\mathcal{M}}} \left(\inf_{t \in S_{\mathbf{m}}} \text{KL}^{\otimes n}(s_0, t) + 2 \frac{\text{pen}(\mathbf{m})}{n} \right) \right] \\ & \quad + C_2 (1 \vee \tau) \frac{\Xi^2}{n} + \frac{\eta' + \eta}{n}. \end{aligned} \tag{30}$$

It is worth mentioning that the comparison of the two inequalities in Theorems 3.1 and 3.2 involves the following major differences. First, to control the random subcollection, in Theorem 3.2, we further assume condition (28). However, this is not a strong assumption and is satisfied, for example, if s_0 is bounded, with a compact support. Assumption (28) is needed because we consider a random subcollection from the whole collection. Thanks to this assumption, we can utilize Bernstein’s inequality to control the additional randomness. It is important to stress that the parameter τ depends on the true unknown density s_0 and cannot be explicitly determined for this reason. It appears not only in the oracle type inequality (30), but also in the penalty term (29). However, in some special cases of BLoME models, for instance finite mixture regression models, we can construct a larger penalty independent of τ to obtain an oracle type inequality but the price to pay for getting rid of τ in the risk bound is the increased value of error upper bound, see Appendix C in [32] for more details. Second, the constant Ξ associated to the Kraft-type inequality for the collection appears squared in the upper bound of the oracle inequality in Theorem 3.2. This is because of the random subcollection $\widetilde{\mathcal{M}}$ of \mathcal{M} , if the model collection is fixed, we get a linear bound as in Theorem 3.1.

3.3. Practical application of finite-sample oracle inequalities

It is important to mention that Theorems 3.1 and 3.2 provide some theoretical justification regarding the penalty shapes when using the slope heuristic for our collection of MoE models, including GLLiM, GLoME, BLLiM, BLoME models.

More precisely, our oracle inequalities show that the performance in $\text{JKL}_{\rho}^{\otimes n}$ loss of our PMLEs are roughly comparable to that of oracle models if we take large enough constants in front of the penalties, whose forms are only known up to multiplicative constants and proportional to the dimensions of models. This motivate us to make use of the slope heuristic criterion to select our data-driven hyperparameters, including the number of mixture components, the degree of

polynomial Gaussian expert's mean functions, and the potential hidden block-diagonal structures of the covariance matrices of the multivariate predictor, see more in Section 5.

To be more precise, we consider the condition (31) for GLoME models. As shown in the proof of Theorem 3.1, in fact we can replace the assumption on $\text{pen}(\mathbf{m})$ by a milder one. More precisely, given a constant \mathfrak{C} , which we specify later, there is a constant κ depending only on ρ and C_1 , such that for every index $\mathbf{m} \in \mathcal{M}$, $\text{pen}(\mathbf{m})$ is bounded from below by

$$\kappa \left(\dim(S_{\mathbf{m}}) \left(2 \left(\sqrt{\mathfrak{C}} + \sqrt{\pi} \right)^2 + \left(\ln \frac{n}{(\sqrt{\mathfrak{C}} + \sqrt{\pi})^2 \dim(S_{\mathbf{m}})} \right)_+ \right) + z_{\mathbf{m}} \right). \quad (31)$$

In our numerical experiment, namely, Figures 3 and 4, we confirm the validity of the linear penalty shape assumption, which supports the use of penalties proportional to the dimension. This implies that the logarithm terms are not detected in practice as shown in Figures 3 and 4 and thus only the preponderant term in $\dim(S_{\mathbf{m}})$ is retained in the penalty form.

In particular, based on Appendix B.4 from [22], we can make explicit the dependence of the theoretical constant κ , with respect to ρ and C_1 as follows. For instance, in Theorem 3.1, given any $\rho \in (0, 1)$ and $C_1 > 1$, define $\epsilon_{\text{pen}} = 1 - \frac{1}{C_1}$. Then κ is determined by

$$\kappa = \frac{\kappa'_0 (\kappa'_1 + \kappa'_2)^2 \left(\sqrt{1 + \frac{72C_\rho \epsilon_{\text{pen}}}{\rho \kappa'_0 (\kappa'_1 + \kappa'_2)^2}} + 1 \right)}{2C_\rho \epsilon_{\text{pen}}} + \frac{18}{\rho},$$

where ϵ_d is a given positive constant and

$$\begin{aligned} \kappa'_0 &= \frac{2(2 + \epsilon_d)}{1 + \epsilon_d}, \kappa'_1 = \frac{1}{\sqrt{\rho(1-\rho)}} \left(3\kappa'_3 \sqrt{2} + 12 + 16\sqrt{\frac{1-\rho}{\rho}} \right), \\ \kappa'_3 &\leq 27, \kappa'_2 = \frac{1}{\sqrt{\rho(1-\rho)}} \left(42 + \frac{3}{4\sqrt{\kappa'_0}} \right). \end{aligned}$$

For example, if $\rho = \frac{1}{2}$, $C_1 = 2$, $\epsilon_d = 1$, $\kappa'_3 = 27$, then $\epsilon_{\text{pen}} = 1 - \frac{1}{C_1} = \frac{1}{2}$, $\kappa'_0 = 3$, $\kappa'_1 = 56 + 162\sqrt{2}$, $\kappa'_2 = 84 + \frac{\sqrt{3}}{2}$, and

$$\begin{aligned} C_\rho &= \frac{1}{\rho} \min \left(\frac{1-\rho}{\rho}, 1 \right) \left(\ln \left(1 + \frac{\rho}{1-\rho} \right) - \rho \right) = 2 \ln 2 - 1, \\ \kappa &\approx 2126069. \end{aligned} \quad (32)$$

According to the previous example, we can see that the theoretical penalty is lower bounded by κ , which can be large in practice. This result is not surprising since, according to [22, Appendix B.4, page 40, line 7], if we choose ϵ_d small enough then κ scales proportionally to

$$\frac{1}{C_\rho \rho (1-\rho) \epsilon_{\text{pen}}} = \frac{\rho}{(1-\rho)^2 \left(\ln \left(1 + \frac{\rho}{1-\rho} \right) - \rho \right)}$$

and thus explodes to $+\infty$ when ρ goes to 1 and C_1 goes to 1. Therefore, it is important to study a natural question as to whether the constant κ appearing in the penalty can be estimated from the data without losing theoretical guaranties on the performance? No definite answer exists so far for MoE regression models, however at least our numerical experiment in Section 5 shows that the slope heuristic proposed by [11] may lead to a good practical solution. In particular, we seek to mathematically and fully justify the slope heuristic in MoE regression models as in least-squares regression on a random (or fixed) design with regression (projection) estimators, respectively, see [11, 5, 4, 3] for greater details. Furthermore, as is often the case in empirical processes theory, the constant κ appearing in the bound is pessimistic as shown in (32). Numerical experiments in Section 5 show that there is a hope that this is only a technical issue. For instance, Figures 5 and 6 show that the practical values of κ , selecting by slope heuristic, are generally very small compared to the theoretical values.

4. Proofs of finite-sample oracle inequalities and main mathematical challenges

4.1. Proof for deterministic collections of GLoME models

The deterministic collection of MoE models include GLLiM, GLoME, SGaME and LinBoSGaME models, where finite-sample oracle inequalities have only been well studied for the latter two models [22, Theorem 2], see also [23, Theorem 2.2.] and [70, Theorem 2]. We first summarize this general model selection theorem and the techniques that [70] used to control the bracketing entropy of LinBoSGaME models with softmax gating networks in Sections 4.1.1 and 4.1.2, respectively. Then, we explain why such techniques can not be directly applied to our collection of GLoME models via Section 4.1.3 and propose our approach to highlight the main issues and our contributions.

4.1.1. A general conditional density model selection theorem for deterministic collection of models

Before stating a general model selection theorem for conditional densities, we have to present some regularity assumptions.

First, we need an information theory type assumption to control the complexity of our collection. We assume the existence of a Kraft-type inequality for the collection [60, 6].

Assumption 4.1 (K). *There is a family $(z_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$ of non-negative numbers and a real number Ξ such that*

$$\Xi = \sum_{\mathbf{m} \in \mathcal{M}} e^{-z_{\mathbf{m}}} < +\infty.$$

For technical reasons, a separability assumption, always satisfied in the setting of this paper, is also required. It is a mild condition, which is classical

in empirical process theory [97, 96]. Assumption 4.2 allows us to work with a countable subset.

Assumption 4.2 (Sep). For every model $S_{\mathbf{m}}$ in the collection $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$, there exists some countable subset $S'_{\mathbf{m}}$ of $S_{\mathbf{m}}$ and a set $\mathcal{X}'_{\mathbf{m}}$ with $\iota(\mathcal{X} \setminus \mathcal{X}'_{\mathbf{m}}) = 0$, where ι denotes Lebesgue measure, such that for every $t \in S_{\mathbf{m}}$, there exist some sequences $(t_k)_{k \geq 1}$ of elements of $S'_{\mathbf{m}}$, such that for every $\mathbf{y} \in \mathcal{Y}$ and every $\mathbf{x} \in \mathcal{X}'_{\mathbf{m}}$, $\ln(t_k(\mathbf{x} | \mathbf{y})) \xrightarrow{k \rightarrow +\infty} \ln(t(\mathbf{x} | \mathbf{y}))$.

Next, recall that the bracketing entropy of a set S with respect to any distance d , denoted by $\mathcal{H}_{[\cdot],d}(\delta, S)$, is defined as the logarithm of the minimal number $\mathcal{N}_{[\cdot],d}(\delta, S)$ of brackets $[t^-, t^+]$ covering S , such that $d(t^-, t^+) \leq \delta$. That is,

$$N_{[\cdot],d}(\delta, S) := \min \left\{ n \in \mathbb{N}^* : \exists [t_k^-, t_k^+]_{k \in [n]} \text{ s.t. } d(t_k^-, t_k^+) \leq \delta, S \subset \bigcup_{k=1}^n [t_k^-, t_k^+] \right\}. \tag{33}$$

Here, $s \in [t_k^-, t_k^+]$ means that $t_k^-(\mathbf{x}, \mathbf{y}) \leq s(\mathbf{x}, \mathbf{y}) \leq t_k^+(\mathbf{x}, \mathbf{y}), \forall (\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$.

We also need the following important Assumption 4.3 on Dudley-type integral of these bracketing entropies, which is often utilized in empirical process theory [97, 96, 53].

Assumption 4.3 (H). For every model $S_{\mathbf{m}}$ in the collection \mathcal{S} , there is a non-decreasing function $\phi_{\mathbf{m}}$ such that $\delta \mapsto \frac{1}{\delta} \phi_{\mathbf{m}}(\delta)$ is non-increasing on $(0, \infty)$ and for every $\delta \in \mathbb{R}^+$,

$$\int_0^\delta \sqrt{\mathcal{H}_{[\cdot],d^{\otimes n}}(\delta, S_{\mathbf{m}}(\tilde{s}, \delta))} d\delta \leq \phi_{\mathbf{m}}(\delta),$$

where $S_{\mathbf{m}}(\tilde{s}, \delta) = \{s_{\mathbf{m}} \in S_{\mathbf{m}} : d^{\otimes n}(\tilde{s}, s_{\mathbf{m}}) \leq \delta\}$. The model complexity of $S_{\mathbf{m}}$ is then defined as $\mathcal{D}_{\mathbf{m}} = n\delta_{\mathbf{m}}^2$, where $\delta_{\mathbf{m}}$ is the unique root of $\frac{1}{\delta} \phi_{\mathbf{m}}(\delta) = \sqrt{n}\delta$.

Observe that the model complexity does not depend on the bracketing entropies of the global models $S_{\mathbf{m}}$, but rather on those of smaller localized sets $S_{\mathbf{m}}(\tilde{s}, \delta)$. We are now able to state an important weak oracle inequality, Theorem 4.1, originally from [22, 23, Theorem 2].

Theorem 4.1. Assume that we observe $(\mathbf{x}_{[n]}, \mathbf{y}_{[n]})$, arising from an unknown conditional density s_0 . Let $\mathcal{S} = (S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$ be an at most countable conditional density model collection. Assume that Assumption 4.1 (K), Assumption 4.2 (Sep), and Assumption 4.3 (H) hold for every model $S_{\mathbf{m}} \in \mathcal{S}$. Then, for any $\rho \in (0, 1)$ and any $C_1 > 1$, there is a constant κ depending only on ρ and C_1 , such that for every index $\mathbf{m} \in \mathcal{M}$,

$$\text{pen}(\mathbf{m}) \geq \kappa (n\delta_{\mathbf{m}}^2 + z_{\mathbf{m}})$$

with $\delta_{\mathbf{m}}$ is the unique root of $\frac{1}{\delta} \phi_{\mathbf{m}}(\delta) = \sqrt{n}\delta$, such that the η' -penalized likelihood estimator $\widehat{s}_{\mathbf{m}}$ satisfies

$$\mathbb{E}_{\mathbf{X}_{[n]}, \mathbf{Y}_{[n]}} [\text{JKL}_{\rho}^{\otimes n}(s_0, \widehat{s}_{\mathbf{m}})] \leq C_1 \inf_{\mathbf{m} \in \mathcal{M}} \left(\inf_{s_{\mathbf{m}} \in S_{\mathbf{m}}} \text{KL}^{\otimes n}(s_0, s_{\mathbf{m}}) + \frac{\text{pen}(\mathbf{m})}{n} \right)$$

$$+ \frac{\kappa C_1 \Xi}{n} + \frac{\eta + \eta'}{n}. \tag{34}$$

For the sake of generality, Theorem 4.1 is relatively abstract. Since the assumptions of Theorem 4.1 are as general as possible. But from the practical point of view, a natural question is the existence of interesting model collections that satisfy these assumptions. We will sketch the proof for collection of LinBoSGaME models from [70, Theorem 1] and show that their result is not directly applicable to the GLoME setting. The main reason is that the technique for handling the linear combination of bounded functions for the weight functions of logistic schemes of [70] is not valid for the Gaussian gating parameters in GLoME models. Therefore, we propose a *reparameterization trick* to bound the metric entropy of the Gaussian gating parameters space; see Section 4.1.3 for more details.

4.1.2. Sketch of the proof for LinBoSGaME models

To prove the main conditional density model selection theorem for LinBoSGaME models, [70, Theorem 1], the authors have to make use of Theorem 4.1. Then, they need to prove that their collection of LinBoSGaME models have to satisfy Assumption 4.1 (K), Assumption 4.2 (Sep), and Assumption 4.3 (H). However, they did not verify Assumption 4.1 (K) and Assumption 4.2 (Sep) and considered them as primitive assumptions on their LinBoSGaME models because of the complexity of LinBoSGaME models and technical reasons. Therefore, the main difficulty remains on verifying Assumption 4.3 (H) via bracketing entropy controls of the linear combination of bounded functions for the weight functions of logistic schemes.

Firstly, they define the following distance over conditional densities:

$$\sup_{\mathbf{y}} d_{\mathbf{x}}(s, t) = \sup_{\mathbf{y} \in \mathcal{Y}} \left(\int_{\mathcal{X}} \left(\sqrt{s(\mathbf{x} | \mathbf{y})} - \sqrt{t(\mathbf{x} | \mathbf{y})} \right)^2 d\mathbf{x} \right)^{1/2}.$$

This leads straightforwardly to $d^{2 \otimes n}(s, t) \leq \sup_{\mathbf{y}} d_{\mathbf{x}}^2(s, t)$. Then, they also define

$$\sup_{\mathbf{y}} d_k(g, g') = \sup_{\mathbf{y} \in \mathcal{Y}} \left(\sum_{k=1}^K \left(\sqrt{g_k(\mathbf{y})} - \sqrt{g'_k(\mathbf{y})} \right)^2 \right)^{1/2},$$

for any gating functions g and g' . To this end, given any densities s and t over \mathcal{X} , the following distance, depending on \mathbf{y} , is constructed as follows:

$$\begin{aligned} \sup_{\mathbf{y}} \max_k d_{\mathbf{x}}(s, t) &= \sup_{\mathbf{y} \in \mathcal{Y}} \max_{k \in [K]} d_{\mathbf{x}}(s_k(\cdot, \mathbf{y}), t_k(\cdot, \mathbf{y})) \\ &= \sup_{\mathbf{y} \in \mathcal{Y}} \max_{k \in [K]} \left(\int_{\mathcal{X}} \left(\sqrt{s_k(\mathbf{x}, \mathbf{y})} - \sqrt{t_k(\mathbf{x}, \mathbf{y})} \right)^2 d\mathbf{x} \right)^{1/2}. \end{aligned}$$

Then, they prove that definition of complexity of model $S_{\mathbf{m}}$ in Assumption 4.3 (H) is related to an classical entropy dimension with respect to a Hellinger type divergence $d^{\otimes n}$, due to Proposition 4.2.

Proposition 4.2 (Proposition 2 from [22]). *For any $\delta \in (0, \sqrt{2}]$, such that $\mathcal{H}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \dim(S_{\mathbf{m}}) (C_{\mathbf{m}} + \ln(\frac{1}{\delta}))$, the function*

$$\phi_{\mathbf{m}}(\delta) = \delta \sqrt{\dim(S_{\mathbf{m}})} \left(\sqrt{C_{\mathbf{m}}} + \sqrt{\pi} + \sqrt{\ln\left(\frac{1}{\min(\delta, 1)}\right)} \right)$$

satisfies Assumption 4.3 (H). Furthermore, the unique solution $\delta_{\mathbf{m}}$ of $\frac{1}{\delta} \phi_{\mathbf{m}}(\delta) = \sqrt{n\delta}$, satisfies

$$n\delta_{\mathbf{m}}^2 \leq \dim(S_{\mathbf{m}}) \left(2 \left(\sqrt{C_{\mathbf{m}}} + \sqrt{\pi} \right)^2 + \left(\ln \frac{n}{(\sqrt{C_{\mathbf{m}}} + \sqrt{\pi})^2 \dim(S_{\mathbf{m}})} \right)_+ \right).$$

Therefore, Proposition 4.2 implies that Assumption 4.3 (H) can be proved via Lemma 4.3.

Lemma 4.3. *For any $\delta \in (0, \sqrt{2}]$, the collection of LinBoSGaME models, $\mathcal{S} = (S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$, satisfies*

$$\mathcal{H}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \dim(S_{\mathbf{m}}) \left(C_{\mathbf{m}} + \ln\left(\frac{1}{\delta}\right) \right).$$

Lemma 4.3 is then obtained by decomposing the entropy terms between the softmax gating functions and the Gaussian experts. Note that both LinBoSGaME and GLoME models share the same structures of Gaussian experts mean, see Figure 1 again for more details. Therefore, in Section 4.1.3, we only highlight our contributions regarding the control of bracketing entropy for the parameter of gating network compared to [70].

More precisely, the authors of [70] rewrite the softmax gating parameters' space as follows:

$$\begin{aligned} \mathbf{W}_{K, d_{\mathbf{w}}} &= \{0\} \otimes \mathbf{W}^{K-1}, \mathbf{W} = \left\{ \mathbf{y} \mapsto \sum_{d=1}^{d_{\mathbf{w}}} \omega_d \theta_{\mathbf{w}, d}(\mathbf{y}) \in \mathbb{R} : \max_{d \in [d_{\mathbf{w}}]} |\omega_d| \leq T_{\mathbf{w}} \right\}, \\ \mathcal{P}_K &= \left\{ \mathbf{y} \mapsto \left(\frac{e^{\mathbf{w}_k(\mathbf{y})}}{\sum_{l=1}^K e^{\mathbf{w}_l(\mathbf{y})}} \right)_{k \in [K]} =: (g_{\mathbf{w}, k}(\mathbf{y}))_{k \in [K]}, \mathbf{w} \in \mathbf{W}_{K, d_{\mathbf{w}}} \right\}. \end{aligned}$$

Then, they also require the definition of metric entropy of the set \mathcal{W}_K : $\mathcal{H}_{d_{\|\sup\|_{\infty}}}(\delta, \mathcal{W}_K)$, which measures the logarithm of the minimal number of balls of radius at most δ , according to a distance $d_{\|\sup\|_{\infty}}$, needed to cover \mathcal{W}_K where

$$d_{\|\sup\|_{\infty}} \left((s_k)_{k \in [K]}, (t_k)_{k \in [K]} \right) = \max_{k \in [K]} \sup_{\mathbf{y} \in \mathcal{Y}} \|s_k(\mathbf{y}) - t_k(\mathbf{y})\|_2, \tag{35}$$

for any K -tuples of functions $(s_k)_{k \in [K]}$, $(t_k)_{k \in [K]}$ and $\|s_k(\mathbf{y}) - t_k(\mathbf{y})\|_2$ is the Euclidean distance in \mathbb{R}^L . By using Lemma 5 and Proposition 2 from [70], Lemma 4.3 holds true if we can prove Lemma 4.4. Note that the first inequality of Lemma 4.4 comes from [70, Lemma 4] and describes relationship between the bracketing entropy of \mathcal{P}_K and the entropy of \mathcal{W}_K .

Lemma 4.4. For all $\delta \in (0, \sqrt{2}]$, there exists a constant $C_{\mathcal{W}_K}$ such that

$$\begin{aligned} \mathcal{H}_{[\cdot], \sup_{\mathbf{y}} d_k} \left(\frac{\delta}{5}, \mathcal{P}_K \right) &\leq \mathcal{H}_{d_{\|\sup\|_\infty}} \left(\frac{3\sqrt{3}\delta}{20\sqrt{K-1}}, \mathcal{W}_K \right) \\ &\leq \dim(\mathcal{W}_K) \left(C_{\mathcal{W}_K} + \ln \left(\frac{20\sqrt{K-1}}{3\sqrt{3}\delta} \right) \right). \end{aligned}$$

By the linearity from the construction of linear combination of a finite set of bounded functions whose coefficients belong to a compact set, in the argument from [70, Proof of Part 1 of Lemma 1, Page 1689], the second inequality of Lemma 4.4 is then easily established as follows.

Proof of the second inequality of Lemma 4.4. Note that for all

$$\mathbf{w} = (0, w_k)_{k \in [K-1]} \in \mathbf{W}_{K, d_{\mathbf{w}}}, \quad \mathbf{v} = (0, v_k)_{k \in [K-1]} \in \mathbf{W}_{K, d_{\mathbf{w}}},$$

it holds that

$$\begin{aligned} d_{\|\sup\|_\infty}(\mathbf{w} - \mathbf{v}) &= \max_{k \in [K-1]} \|\mathcal{W}_K - \mathbf{v}_k\|_\infty \\ &= \max_{k \in [K-1]} \sup_{\mathbf{y} \in \mathbf{Y}} \left| \sum_{i=1}^{d_{\mathbf{w}}} \omega_{k,i}^{\mathbf{w}} \theta_{\mathbf{w},i}(\mathbf{y}) - \sum_{i=1}^{d_{\mathbf{w}}} \omega_{k,i}^{\mathbf{v}} \theta_{\mathbf{w},i}(\mathbf{y}) \right| \\ &\leq \max_{k \in [K-1]} \sum_{i=1}^{d_{\mathbf{w}}} |\omega_{k,i}^{\mathbf{w}} - \omega_{k,i}^{\mathbf{v}}| \underbrace{\sup_{\mathbf{y} \in \mathbf{Y}} |\theta_{\mathbf{w},i}(\mathbf{y})|}_{\leq 1} \\ &\leq d_{\mathbf{w}} \max_{k \in [K-1], i \in [d_{\mathbf{w}}]} |\omega_{k,i}^{\mathbf{w}} - \omega_{k,i}^{\mathbf{v}}|. \end{aligned}$$

Therefore, we obtain

$$\begin{aligned} &\mathcal{H}_{d_{\|\sup\|_\infty}} \left(\frac{3\sqrt{3}\delta}{20\sqrt{K-1}}, \mathcal{W}_K \right) \\ &\leq \mathcal{H}_{d_{\|\sup\|_\infty}} \left(\frac{3\sqrt{3}\delta}{20\sqrt{K-1}d_{\mathbf{w}}}, \left\{ \boldsymbol{\omega} \in \mathbb{R}^{(K-1)d_{\mathbf{w}}} : \|\boldsymbol{\omega}\|_\infty \leq T_{\mathbf{w}} \right\} \right) \\ &\leq (K-1)d_{\mathbf{w}} \ln \left(1 + \frac{20\sqrt{K-1}d_{\mathbf{w}}T_{\mathbf{w}}}{3\sqrt{3}\delta} \right) \\ &\leq (K-1)d_{\mathbf{w}} \ln \left(\sqrt{2} + \frac{20\sqrt{K-1}d_{\mathbf{w}}T_{\mathbf{w}}}{3\sqrt{3}} + \ln \left(\frac{1}{\delta} \right) \right). \quad \square \end{aligned}$$

However, proving Lemma 4.5, our equivalent of Lemma 4.4, will be much harder and could be used for controlling the bracketing entropy for many standard MoE regression models with Gaussian gating networks, see Section 4.1.3 for more details.

4.1.3. Our contributions on the proof for GLoME models

To prove Theorem 3.1, we also need to make use of Theorem 4.1. Then, our model collection has to satisfy Assumption 4.1 (K), Assumption 4.2 (Sep), and Assumption 4.3 (H).

Note that in the proof of LinBoSGaME models, the authors of [70] did not prove Assumption 4.1 (K) and Assumption 4.2 (Sep) and considered them as assumptions on their LinBoSGaME models because of the complexity of LinBoSGaME models and technical reasons. However, in our proof for GLoME models, we consider an explicit example where the model is defined by $\mathcal{M} = \mathcal{K} \times \mathcal{D}_{\mathbf{r}} = [K_{\max}] \times [d_{\max}]$, $K_{\max}, d_{\max} \in \mathbb{N}^*$, leads to the Assumption 4.1 (K) is always satisfied. It is interesting to find the optimal family $(z_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$ satisfying Assumption 4.1 (K). To the best of our knowledge, this question is only partially answered in some special cases of MoE regression models, *e.g.*, Gaussian graphical model [33], Gaussian finite mixture models [62], finite mixture of Gaussian regression models [29, 32]. However, for the standard MoE regression models such as LinBoSGaME and GLoME models, such question still remains open due to the complexity of models. We hope to resolve this important and interesting problem in our future work. Furthermore, note that the Assumption 4.2 (Sep) is true when we consider Gaussian densities [60].

Therefore, our model has only to satisfy the remaining Assumption 4.3 (H). Following the same strategy as in the proof of LinBoSGaME models, the main task for GLoME models is to establish Lemma 4.5, which is much more difficult compared to Lemma 4.4 and is proved in Appendix A.2.

For the Gaussian gating parameters, to make use of the first inequality from Lemma 4.4 of [70], we propose the following *reparameterization trick* of the Gaussian gating space, which is defined in (2), via the logistics scheme \mathcal{P}_K and the nonlinear space \mathcal{W}_K as follows:

$$\begin{aligned} \mathcal{W}_K &= \left\{ \mathbf{y} \mapsto (\ln(\pi_k \phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k)))_{k \in [K]} =: (\mathbf{w}_k(\mathbf{y}; \boldsymbol{\omega}))_{k \in [K]} = \mathbf{w}(\mathbf{y}; \boldsymbol{\omega}) : \boldsymbol{\omega} \in \tilde{\Omega}_K \right\}, \\ \mathcal{P}_K &= \left\{ \mathbf{y} \mapsto \left(\frac{e^{\mathbf{w}_k(\mathbf{y})}}{\sum_{l=1}^K e^{\mathbf{w}_l(\mathbf{y})}} \right)_{k \in [K]} =: (g_{\mathbf{w},k}(\mathbf{y}))_{k \in [K]}, \mathbf{w} \in \mathcal{W}_K \right\}. \end{aligned} \quad (36)$$

We aim to provide the following important upper bound for metric entropy of nonlinear space Lemma 4.5, which play a key step for controlling the bracketing entropy not only for GLoME models but also for any standard MoE regression models with Gaussian gating networks, *e.g.*, BLLiM and BLoME models, see again Figure 1 for comprehensive descriptions of this general class.

Lemma 4.5. *For all $\delta \in (0, \sqrt{2}]$, there exists a constant $C_{\mathcal{W}_K}$ such that*

$$\begin{aligned} \mathcal{H}_{[\cdot], \sup_{\mathbf{y}} d_k} \left(\frac{\delta}{5}, \mathcal{P}_K \right) &\leq \mathcal{H}_{d_{\|\cdot\|_{\infty}}} \left(\frac{3\sqrt{3}\delta}{20\sqrt{K-1}}, \mathcal{W}_K \right) \\ &\leq \dim(\mathcal{W}_K) \left(C_{\mathcal{W}_K} + \ln \left(\frac{20\sqrt{K-1}}{3\sqrt{3}\delta} \right) \right). \end{aligned}$$

More precisely, Assumption 4.3 (H) is proved due to Lemma 4.3 and the following Lemma 4.6 by using the fact that $\sum_{k=1}^K g_{\mathbf{w},k}(\mathbf{y}) = 1, \forall \mathbf{y} \in \mathcal{Y}, \forall \mathbf{w} \in \mathcal{W}_K$ and $\mathcal{H}_{[\cdot],d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \mathcal{H}_{[\cdot],\sup_{\mathbf{y}} d_{\mathbf{x}}}(\delta, S_{\mathbf{m}})$, which is obtained by definition of bracketing entropy and $d^{\otimes n}(s, t) \leq \sup_{\mathbf{y}} d_{\mathbf{x}}(s, t)$.

Lemma 4.6 (Lemma 5 from [70]). *Let*

$$\mathcal{G}_{K,d} = \left\{ \mathcal{X} \times \mathcal{Y} \ni (\mathbf{x}, \mathbf{y}) \mapsto (\Phi_D(\mathbf{x}; \mathbf{v}_{k,d}(\mathbf{y}), \boldsymbol{\Sigma}_k))_{k \in [K]} : \mathbf{v}_d \in \mathfrak{V}_{K,d}, \boldsymbol{\Sigma} \in \mathbf{V}_K \right\}.$$

For all $\delta \in (0, \sqrt{2}]$ and $\mathbf{m} \in \mathcal{M}$,

$$\mathcal{H}_{[\cdot],\sup_{\mathbf{y}} d_{\mathbf{x}}}(\delta, S_{\mathbf{m}}) \leq \mathcal{H}_{[\cdot],\sup_{\mathbf{y}} d_k} \left(\frac{\delta}{5}, \mathcal{P}_K \right) + \mathcal{H}_{[\cdot],\sup_{\mathbf{y}} \max_k d_{\mathbf{x}}} \left(\frac{\delta}{5}, \mathcal{G}_{K,d} \right).$$

By making use of Lemma 4.6, the remaining task is to control the bracketing entropy of the Gaussian gating functions and experts separately via Lemmas 4.5 and 4.7, which are proved in Appendices A.2 and A.3, respectively.

Lemma 4.7. *For all $\delta \in (0, \sqrt{2}]$, there exists a constant $C_{\mathcal{G}_{K,d}}$ such that*

$$\mathcal{H}_{[\cdot],\sup_{\mathbf{y}} \max_k d_{\mathbf{x}}} \left(\frac{\delta}{5}, \mathcal{G}_{K,d} \right) \leq \dim(\mathcal{G}_{K,d}) \left(C_{\mathcal{G}_{K,d}} + \ln \left(\frac{1}{\delta} \right) \right). \quad (37)$$

To this end, Lemma 4.6 allows us to conclude that given $\mathfrak{C} = C_{\mathcal{W}_K} + \ln \left(\frac{5K_{\max} \sqrt{K_{\max}}}{a_W} \right) + C_{\mathcal{G}_{K,d}}$,

$$\begin{aligned} \mathcal{H}_{[\cdot],\sup_{\mathbf{y}} d_{\mathbf{x}}}(\delta, S_{\mathbf{m}}) &\leq \mathcal{H}_{[\cdot],\sup_{\mathbf{y}} d_k} \left(\frac{\delta}{5}, \mathcal{P}_K \right) + \mathcal{H}_{[\cdot],\sup_{\mathbf{y}} \max_k d_{\mathbf{x}}} \left(\frac{\delta}{5}, \mathcal{G}_{K,d} \right) \\ &\leq \dim(S_{\mathbf{m}}) \left(\mathfrak{C} + \ln \left(\frac{1}{\delta} \right) \right). \end{aligned}$$

Then, Proposition 4.2 leads to

$$n\delta_{\mathbf{m}}^2 \leq \dim(S_{\mathbf{m}}) \left(2 \left(\sqrt{\mathfrak{C}} + \sqrt{\pi} \right)^2 + \left(\ln \frac{n}{\left(\sqrt{\mathfrak{C}} + \sqrt{\pi} \right)^2 \dim(S_{\mathbf{m}})} \right)_+ \right).$$

Finally, Theorem 4.1 implies that for any given collection of GLoME models $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$, the oracle inequality of Theorem 3.1 is satisfied if $\text{pen}(\mathbf{m})$ is bounded from below by

$$\kappa \left(\dim(S_{\mathbf{m}}) \left(2 \left(\sqrt{\mathfrak{C}} + \sqrt{\pi} \right)^2 + \left(\ln \frac{n}{\left(\sqrt{\mathfrak{C}} + \sqrt{\pi} \right)^2 \dim(S_{\mathbf{m}})} \right)_+ \right) + z_{\mathbf{m}} \right).$$

4.2. Proof of random subcollection of BLoME models

The random subcollection of MoE models include BLLiM and BLoME models where finite-sample oracle inequalities were only well studied for finite mixture of Gaussian regression models via a model selection theorem for MLE among a random subcollection of models in regression framework of [29, Theorem 5.1], see also [33, Theorem 7.3]. This is an extension of a whole collection of conditional densities from [22, Theorem 2], and of [60, Theorem 7.11], working only for density estimation. In Section 4.2.1, we first summarize this theorem and the techniques that [32] used to control the bracketing entropy of finite mixture of Gaussian regression models with joint rank and variable selection for parsimonious estimation in a high-dimensional framework. Then, we explain why such techniques can not be applied to our collection of BLLiM and BLoME models to highlight the main challenges and our contributions.

4.2.1. A model selection theorem for MLE among a random subcollection

We can now state the main result of [29, Theorem 5.1] for the model selection theorem for MLE among a random subcollection.

Theorem 4.8 (Theorem 5.1 from [29]). *Let $(\mathbf{x}_{[n]}, \mathbf{y}_{[n]})$ be observations coming from an unknown conditional density s_0 . Let the model collection $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$ be an at most countable collection of conditional density sets. Assume that Assumption 4.1 (K), Assumption 4.2 (Sep), and Assumption 4.3 (H) hold for every $\mathbf{m} \in \mathcal{M}$. Let $\epsilon_{KL} > 0$, and $\bar{s}_{\mathbf{m}} \in S_{\mathbf{m}}$, such that*

$$\text{KL}^{\otimes n}(s_0, \bar{s}_{\mathbf{m}}) \leq \inf_{t \in S_{\mathbf{m}}} \text{KL}^{\otimes n}(s_0, t) + \frac{\epsilon_{KL}}{n};$$

and let $\tau > 0$, such that

$$\bar{s}_{\mathbf{m}} \geq e^{-\tau} s_0. \quad (38)$$

Introduce $(S_{\mathbf{m}})_{\mathbf{m} \in \tilde{\mathcal{M}}}$, a random subcollection of $(S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$. Consider the collection $(\hat{s}_{\mathbf{m}})_{\mathbf{m} \in \tilde{\mathcal{M}}}$ of η -log likelihood minimizer satisfying (24) for all $\mathbf{m} \in \tilde{\mathcal{M}}$. Then, for any $\rho \in (0, 1)$, and any $C_1 > 1$, there are two constants κ and C_2 depending only on ρ and C_1 , such that, for every index $\mathbf{m} \in \mathcal{M}$,

$$\text{pen}(\mathbf{m}) \geq \kappa (\mathcal{D}_{\mathbf{m}} + (1 \vee \tau) z_{\mathbf{m}}),$$

and where the model complexity $\mathcal{D}_{\mathbf{m}}$ is defined in Assumption 4.3 (H), the η' -penalized likelihood estimator $\hat{s}_{\mathbf{m}}^{\wedge}$, defined as in (25) on the subset $\tilde{\mathcal{M}} \subset \mathcal{M}$, satisfies

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}_{[n]}, \mathbf{Y}_{[n]}} [\text{JKL}_{\rho}^{\otimes n}(s_0, \hat{s}_{\mathbf{m}}^{\wedge})] \\ & \leq C_1 \mathbb{E}_{\mathbf{X}_{[n]}, \mathbf{Y}_{[n]}} \left[\inf_{\mathbf{m} \in \tilde{\mathcal{M}}} \left(\inf_{t \in S_{\mathbf{m}}} \text{KL}^{\otimes n}(s_0, t) + 2 \frac{\text{pen}(\mathbf{m})}{n} \right) \right] \\ & \quad + C_2 (1 \vee \tau) \frac{\Xi^2}{n} + \frac{\eta' + \eta}{n}. \end{aligned}$$

4.2.2. Sketch of the proof for BLoME models

To work with conditional density estimation in the BLLiM and BLoME models, it is natural to make use of Theorem 4.1. However, it is worth mentioning that because the model collection constructed by the BLLiM [34] or by some suitable procedures for BLoME models in practice is usually random, we have to use a model selection theorem for MLE among a random subcollection (cf. [29, Theorem 5.1] and [33, Theorem 7.3]).

More precisely, we explain how Theorem 4.8 implies the finite-sample oracle inequality, Theorem 3.2. To this end, our collections of BLoME models has to satisfy some regularity assumptions, see Section 4.2.3 for more details. For BLoME models, the main difficulty in proving our oracle inequality lies in bounding the bracketing entropy of the Gaussian gating functions and Gaussian experts with block-diagonal covariance matrices. To overcome the former issue, we follow a reparameterization trick of the Gaussian gating parameters space in (36) and Lemma 4.5. For the second one, based on some ideas of Gaussian mixture models from [41, 62], we contribute a novel extension for standard MoE models with Gaussian gating networks. Note that our contributions extend the recent novel result on block-diagonal covariance matrices in [33], which is only developed for Gaussian graphical models.

4.2.3. Our contributions on BLoME models

It should be stressed that all we need is to verify that Assumption 4.3 (H), Assumption 4.2 (Sep) and Assumption 4.1 (K) hold for every $\mathbf{m} \in \mathcal{M}$. According to the result from [29, Section 5.3], Assumption 4.2 (Sep) holds when we consider Gaussian densities and the assumption defined by (38) is true if we assume further that the true conditional density s_0 is bounded and compactly supported. Furthermore, since we restricted d and K to $\mathcal{D}_{\mathbf{r}} = [d_{\max}]$ and $\mathcal{K} = [K_{\max}]$, respectively, it is true that there exists a family $(z_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$ and $\Xi > 0$ such that, Assumption 4.1 (K) is satisfied. Therefore, the proof for the remaining Assumption 4.3 (H) is our novel contribution. In particular, to the best of our knowledge, there are no results that can be directly applied to Assumption 4.3 (H) for BLoME models due to their complexity with the Gaussian gating functions and Gaussian experts with block-diagonal covariance matrices. This highlights our contributions to this challenging problem concerning the work of [41, 62, 29, 32, 33, 70].

By using Lemma 4.3, Assumption 4.3 (H) holds true if we can prove that for any $\delta \in (0, \sqrt{2}]$, the collection of LinBoSGaME models, $\mathcal{S} = (S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$, satisfies

$$\mathcal{H}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \dim(S_{\mathbf{m}}) \left(C_{\mathbf{m}} + \ln \left(\frac{1}{\delta} \right) \right). \quad (39)$$

Proof of (39). Note that (39) can be established by first decomposing the entropy term between the Gaussian gating functions and the Gaussian experts.

Motivated by our *reparameterization trick* of the Gaussian gating space in Section 4.1.3 and (36), we define Gaussian experts $\mathcal{G}_{K,d,\mathbf{B}}$ as follows.

$$\mathcal{G}_{K,d,\mathbf{B}} = \left\{ (\mathbf{x}, \mathbf{y}) \mapsto (\phi(\mathbf{x}; \mathbf{v}_{k,d}(\mathbf{y}), \boldsymbol{\Sigma}_k(\mathbf{B}_k)))_{k \in [K]} : \mathbf{v}_d \in \Upsilon_{K,d}, \boldsymbol{\Sigma}(\mathbf{B}) \in \mathbf{V}_K(\mathbf{B}) \right\}.$$

There are two possible ways to decompose the bracketing entropy of $S_{\mathbf{m}}$ based on different distances. For the first approach, we can use Lemma 4.9 [70, Lemma 5]:

Lemma 4.9. *For all $\delta \in (0, \sqrt{2}]$ and $\mathbf{m} \in \mathcal{M}$,*

$$\mathcal{H}_{[\cdot], \sup_{\mathbf{y}} d_{\mathbf{x}}}(\delta, S_{\mathbf{m}}) \leq \mathcal{H}_{[\cdot], \sup_{\mathbf{y}} d_k} \left(\frac{\delta}{5}, \mathcal{P}_K \right) + \mathcal{H}_{[\cdot], \sup_{\mathbf{y}} \max_k d_{\mathbf{x}}} \left(\frac{\delta}{5}, \mathcal{G}_{K,d,\mathbf{B}} \right).$$

As mentioning in Appendix B.2.1 from [70], Lemma 4.9 boils down to assuming that \mathbf{Y} is bounded. Furthermore, they also claim that this boundedness assumption can be relaxed when using smaller distance $d^{\otimes n}$ but bounding the corresponding bracketing entropy becomes much more challenging. We successfully weaken such boundedness assumption via utilizing the smaller distance: $d^{\otimes n}$, for the bracketing entropy of $S_{\mathbf{m}}$ although bounding such bracketing entropy for \mathcal{W}_K and $\mathcal{G}_{K,\mathbf{B}}$ becomes much more challenging. This reinforces our new contributions concerning the control of bracketing entropy of BLoME models. Consequently, this leads to the second approach via Lemma 4.10 [70, Lemma 6].

Lemma 4.10. *For all $\delta \in (0, \sqrt{2}]$,*

$$\mathcal{H}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \mathcal{H}_{[\cdot], d_{\mathcal{P}_K}} \left(\frac{\delta}{2}, \mathcal{P}_K \right) + \mathcal{H}_{[\cdot], d_{\mathcal{G}_{K,d,\mathbf{B}}}} \left(\frac{\delta}{2}, \mathcal{G}_{K,d,\mathbf{B}} \right),$$

where

$$\begin{aligned} d_{\mathcal{P}_K}^2(g^+, g^-) &= \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n d_k^2(g^+(\mathbf{Y}_i), g^-(\mathbf{Y}_i)) \right] \\ &= \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^K \left(\sqrt{g_k^+(\mathbf{Y}_i)} - \sqrt{g_k^-(\mathbf{Y}_i)} \right)^2 \right], \\ d_{\mathcal{G}_{K,d,\mathbf{B}}}^2(\phi^+, \phi^-) &= \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^K d_{\mathbf{x}}^2(\phi_k^+(\cdot, \mathbf{Y}_i), \phi_k^-(\cdot, \mathbf{Y}_i)) \right] \\ &= \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^K \int_{\mathcal{X}} \left(\sqrt{\phi_k^+(\mathbf{x}, \mathbf{Y}_i)} - \sqrt{\phi_k^-(\mathbf{x}, \mathbf{Y}_i)} \right)^2 d\mathbf{x} \right]. \end{aligned}$$

Next, we make use of Lemma 4.11, which is proved in Appendix A.4, to provide an upper bound on the bracketing entropy of $S_{\mathbf{m}}$ and \mathcal{P}_K on the corresponding distances $d^{\otimes n}$ and $d_{\mathcal{P}_K}$, respectively.

Lemma 4.11. *It holds that*

$$d^{\otimes n}(s, t) \leq \sup_{\mathbf{y}} d_{\mathbf{x}}(s, t), \text{ and } \mathcal{H}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \mathcal{H}_{[\cdot], \sup_{\mathbf{y}} d_{\mathbf{x}}}(\delta, S_{\mathbf{m}}), \quad (40)$$

$$d_{\mathcal{P}_K}(g^+, g^-) \leq \sup_{\mathbf{y}} d_k(g^+, g^-), \text{ and } \mathcal{H}_{[\cdot], d_{\mathcal{P}_K}}\left(\frac{\delta}{2}, \mathcal{P}_K\right) \leq \mathcal{H}_{[\cdot], \sup_{\mathbf{y}} d_k}\left(\frac{\delta}{2}, \mathcal{P}_K\right). \tag{41}$$

Lemmas 4.10 and 4.11 imply that

$$\mathcal{H}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \mathcal{H}_{[\cdot], \sup_{\mathbf{y}} d_k}\left(\frac{\delta}{2}, \mathcal{P}_K\right) + \mathcal{H}_{[\cdot], d_{\mathcal{G}_{K,d,\mathbf{B}}}}\left(\frac{\delta}{2}, \mathcal{G}_{K,d,\mathbf{B}}\right).$$

Based on this metric, one can first relate the bracketing entropy of \mathcal{P}_K to $\mathcal{H}_{d_{\|\sup\|\infty}}(\delta, \mathcal{W}_K)$, and then obtain the upper bound for its entropy via Lemma 4.5. Then, we present our main contribution for BLoME models via Lemma 4.12. This lemma allows us to construct the Gaussian brackets to handle the metric entropy for Gaussian experts, which is established in Appendix A.5.

Lemma 4.12.

$$\mathcal{H}_{[\cdot], d_{\mathcal{G}_{K,d,\mathbf{B}}}}\left(\frac{\delta}{2}, \mathcal{G}_{K,d,\mathbf{B}}\right) \leq \dim(\mathcal{G}_{K,d,\mathbf{B}}) \left(C_{\mathcal{G}_{K,d,\mathbf{B}}} + \ln\left(\frac{1}{\delta}\right) \right). \tag{42}$$

Finally, (39) can easily be proved via Lemmas 4.5 and 4.12. □

5. Numerical experiments

Note that our numerical experiments in Section 5, and the codes written in the R programming language [93] are available in the following link:

<https://github.com/Trung-TinGUYEN/NamsGLOME-Simulation>.

5.1. The procedure

We illustrate our theoretical results in settings similar to those considered by [21] and [70], including simulated as well as real data sets. We first observe n random samples $(\mathbf{x}_i, \mathbf{y}_i)_{i \in [n]}$ from an forward conditional density s_0^* , and look for the best estimate among $s_{\mathbf{m}}^* \in S_{\mathbf{m}}^*$, $\mathbf{m} \in \mathcal{M}$, defined in (20). We considered the simple case where the mean experts are linear functions, which leads to GLoME and supervised GLLiM are identical models. Our aim is to estimate the best number of components K , as well as the model parameters. As described in more detail in [26], we use a GLLiM-EM algorithm to estimate the model parameters for each K , and select the optimal model using the penalized approach that was described earlier. More precisely, in the following numerical experiments, the GLoME model is learned using functions from the package `xLLiM`, available on CRAN. It solves the inverse regression problem, defined in (16), and obtain the inverse maximum likelihood estimates (MLE) $(\hat{s}_{\mathbf{m}}(\mathbf{x}_i | \mathbf{y}_i))_{i \in [N]}$, $\mathbf{m} \in \mathcal{M}$, then via (17), we obtain the forward MLE $(\hat{s}_{\mathbf{m}}^*(\mathbf{y}_i | \mathbf{x}_i))_{i \in [N]}$, $\mathbf{m} \in \mathcal{M}$.

According to the general procedure for model selection, we first compute the forward MLE for each model $\mathbf{m} \in \mathcal{M}$, where $\mathcal{M} = \mathcal{K}$. Then, we select the model

that satisfies the definition (25) with $\text{pen}(\mathbf{m}) = \kappa \dim(S_{\mathbf{m}}^*)$, where κ is a positive hyperparameter. The point we want to stress here is that the theoretical result stated in Theorem 3.1 gives the general form of penalty functions but it does not provide explicit penalties since the lower bounds on penalty functions in (26) are defined up to an unknown multiplicative constant κ and mixture parameters are not bounded in practice. Nevertheless, Theorem 3.1 is required to justify the shape of penalties using for the slope heuristic and allows the number of variables D as well as responses L to be large compared to the fixed sample size n . In particular, our Theorem 3.1 and Remark 2.2 guarantee that there exists a κ large enough for which the estimate has the desired properties. Therefore, we need a data-driven method to choose κ . According to the AIC or the BIC, we can select $\kappa = 1$ or $\kappa = \frac{\ln n}{2}$. An important limitation of these criteria, however, is that they are only valid asymptotically. This implies that there are no finite-sample guarantees when using AIC or BIC for choosing between different levels of complexity. Their use in small sample settings is thus ad hoc. To overcome such difficulties, Birgé and Massart [11] proposed a novel approach, called slope heuristics, supported by a non-asymptotic oracle inequality. This method leads to an optimal data-driven choice of multiplicative constants for penalties, *e.g.*, κ in our framework. Thus, we shall concentrate our attention on the slope heuristic for choosing the number of mixture components in our numerical experiments.

5.2. Simulated data sets

Note that our main objective here is to investigate how well the empirical tensorized Kullback–Leibler divergence between the true model (s_0^*) and the selected model $\hat{s}_{\mathbf{m}}^*$ follows the finite-sample oracle inequality of Theorem 3.1, as well as the rate of convergence of the error term. Therefore, we focus on 1-dimensional data sets, that is, with $L = D = 1$. Beyond the statistical estimation and model selection objectives considered here, the dimensionality reduction capability of GLLiM in high-dimensional regression data, typically $D \gg L$, can be found in [26, Section 6].

We construct simulated data sets following two scenarios: a *well-specified* (WS) case in which the true forward conditional density belongs to the class of proposed models:

$$s_0^*(y | x) = \frac{\Phi(x; 0.2, 0.1)}{\Phi(x; 0.2, 0.1) + \Phi(x; 0.8, 0.15)} \Phi(y; -5x + 2, 0.09) \\ + \frac{\Phi(x; 0.8, 0.15)}{\Phi(x; 0.2, 0.1) + \Phi(x; 0.8, 0.15)} \Phi(y; 0.1x, 0.09),$$

and a *misspecified* (MS) case, whereupon such an assumption is not true:

$$s_0^*(y | x) = \frac{\Phi(x; 0.2, 0.1)}{\Phi(x; 0.2, 0.1) + \Phi(x; 0.8, 0.15)} \Phi(y; x^2 - 6x + 1, 0.09) \\ + \frac{\Phi(x; 0.8, 0.15)}{\Phi(x; 0.2, 0.1) + \Phi(x; 0.8, 0.15)} \Phi(y; -0.4x^2, 0.09).$$


Figures 2a and 2e show some typical realizations of 2000 data points arising from the WS and MS scenarios. Note that by using GLoME, our estimator performs well in the WS setting (Figures 2b to 2d). In the MS case, we expect our algorithm to automatically balance the model bias and its variance (Figures 2f to 2h), which leads to the choice of a complex model, with 4 mixture components. This observation will be elaborated upon in the subsequent experiments.

Firstly, by using the `capushe` (CALibrating Penalties Using Slope HEuristics) package in **R** [8], we can select the penalty coefficient, along with the number of mixture components K . This heuristic comprises two possible criteria: the slope criterion and the jump criterion. The first criterion consists of computing the *asymptotic* slope of the log-likelihood (Figures 3 and 4), drawn according to the model dimension, and then penalizing the log-likelihood by twice the slope times the model dimension. Regarding the second criterion, one aims to represent the dimension of the selected model according to κ (Figures 5 and 6), and find $\hat{\kappa}$, such that if $\kappa < \hat{\kappa}$, then the dimension of the selected model is large, and of reasonable size, otherwise. The slope heuristic prescribes then the use of $\kappa = 2\hat{\kappa}$. In our simulated data sets, Figures 7 to 10 show that the jump criterion appears to work better. The slope criterion sometimes chooses highly complex models in the WS case, with the problem exacerbated in the MS case.


Next, a close inspection shows that the bias-variance trade-off differs between the two examples. We run our experiment over 100 trials with $K \in \mathcal{K} = [20]$, using both the jump and slope criteria. The first remark is that the best choice of $K = 2$ appears to be selected with very high probability, even for large samples ($n = 10000$) in the WS case. This can be observed in Figures 7a, 7b, 9a and 9b. In the MS case, the best choice for K should balance between the model approximation error term and the variance one, which is observed in Figures 8a, 8b, 10a and 10b. Here, the larger the number of samples n , the larger the value of K that is selected as optimal.

From hereon in, we only focus on the jump criterion, due to its stability regarding the choice of K . We wish to measure the performances of our chosen GLoME models in term of the tensorized Kullback–Leibler divergence, $\text{KL}^{\otimes n}$, which can not be calculated exactly in the case of Gaussian mixtures. Therefore, we evaluate the divergence using a Monte Carlo simulation, since we know the true density. We note that the variability of this randomized approximation has been demonstrated to be negligible in practice, which is also supported in the numerical experiments by [70]. More precisely, we compute the Monte Carlo approximation for the tensorized Kullback–Leibler divergence as follows. First, note that the Monte Carlo approximation for tensorized Kullback–Leibler divergence between the true model, s_0^* , and the selected model $\hat{s}_{\mathbf{m}}^*$ can be approximated as


$$\begin{aligned} \text{KL}^{\otimes n} \left(s_0^*, \hat{s}_{\mathbf{m}}^* \right) &= \mathbb{E}_{\mathbf{X}} \left[\frac{1}{n} \sum_{i=1}^n \text{KL} \left(s_0^* (\cdot \mid \mathbf{x}_i), \hat{s}_{\mathbf{m}}^* (\cdot \mid \mathbf{x}_i) \right) \right] \\ &\approx \frac{1}{n} \sum_{i=1}^n \text{KL} \left(s_0^* (\cdot \mid \mathbf{x}_i), \hat{s}_{\mathbf{m}}^* (\cdot \mid \mathbf{x}_i) \right) \end{aligned}$$




(a) Typical realization of example WS




(b) Clustering by GLoME




(c) 2D view of the resulting conditional density with the 2 regression components




(d) Gating network probabilities




(e) Typical realization of example MS



(f) Clustering by GLoME



(g) 2D view of the resulting conditional density with the 4 regression components



(h) Gating network probabilities

FIG 2. Clustering deduced from the estimated conditional density of GLoME by a MAP principle with 2000 data points of example WS and MS. The dash and solid black curves present the true and estimated mean functions.




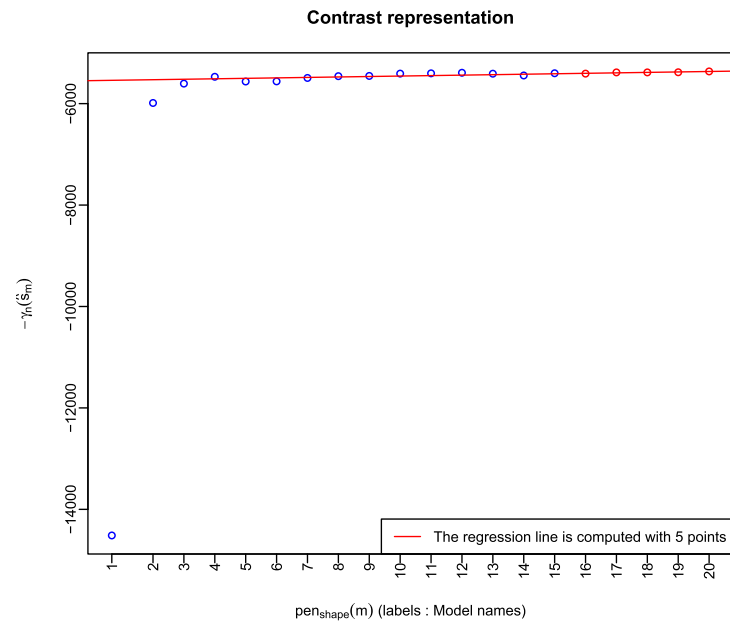


FIG 3. Plot of the selected model dimension using the slope criterion with 2000 data points.




(a) Example WS.




(b) Example MS.

FIG 4. Plot of the selected model dimension using the slope criterion with 10000 data points.




(a) Example WS.




(b) Example MS.

FIG 5. Plot of the selected model dimension using the jump criterion with 2000 data points.




(a) Example WS.




(b) Example MS.

FIG 6. Plot of the selected model dimension using the jump criterion with 10000 data points.




(a) 2000 data points.




(b) 10000 data points.

FIG 7. Comparison histograms of selected K in WS case using jump criterion over 100 trials.




(a) 2000 data points.




(b) 10000 data points.

FIG 8. Comparison histograms of selected K in MS case using jump criterion over 100 trials.




(a) 2000 data points.




(b) 10000 data points.

FIG 9. Comparison histograms of selected K in WS case using slope criterion over 100 trials.



(a) 2000 data points



(b) 10000 data points.

FIG 10. Comparison histograms of selected K in MS case using slope criterion over 100 trials.

$$= \frac{1}{n} \sum_{i=1}^n \frac{1}{n_y} \sum_{j=1}^{n_y} \ln \left(\frac{s_0^*(\mathbf{y}_{i,j} | \mathbf{x}_i)}{\widehat{s}_{\mathbf{m}}^*(\mathbf{y}_{i,j} | \mathbf{x}_i)} \right),$$

where the data $\mathbf{x}_i, i \in [n]$, and $(\mathbf{y}_{i,j})_{j \in [n_y]}$ are drawn from $s_0^*(\cdot | \mathbf{x}_i)$. Then, $\mathbb{E} \left[\text{KL}^{\otimes n} \left(s_0^*, \widehat{s}_{\mathbf{m}}^* \right) \right]$ is approximated again by averaging over $N_t = 100$ Monte Carlo trials. Therefore, the simulated data used for approximation can be written as $(\mathbf{x}_i, \mathbf{y}_{i,j})_t$ with $i \in [n], j \in [n_y], t \in [N_t]$.

Based on the approximation, Figures 11 and 12 show the box plots and the mean of the tensorized Kullback–Leibler divergence over 100 trials, based on the jump criterion. Our boxplots confirm that the mean tensorized Kullback–Leibler divergence between \widehat{s}_K^* and s_0^* , over $K \in \{1, \dots, 20\}$ number of mixture components, is always larger than the mean of tensorized Kullback–Leibler divergence between the penalized estimator \widehat{s}_K^* and s_0^* , which is consistent with Theorem 3.1. In particular, if the true model belongs to our nested collection, the mean tensorized Kullback–Leibler divergence seems to behave like $\frac{\dim(S_{\mathbf{m}}^*)}{2n}$ (shown by a dotted line), which can be explained by the AIC heuristic. More precisely, we firstly assume that

$$S_{\mathbf{m}}^* = \left\{ \mathcal{X} \times \mathcal{Y} \ni (\mathbf{x}, \mathbf{y}) \mapsto s_{\mathbf{m}}^* := s_{\psi_{\mathbf{m}}^*}(\mathbf{y} | \mathbf{x}) : \psi_{\mathbf{m}}^* \in \Psi_{\mathbf{m}}^* \subset \mathbb{R}^{\dim(S_{\mathbf{m}}^*)} \right\}$$

is identifiable and make some strong regularity assumptions on $\psi_{\mathbf{m}}^* \mapsto s_{\psi_{\mathbf{m}}^*}$. Further, we assume the existence of $\dim(S_{\mathbf{m}}^*) \times \dim(S_{\mathbf{m}}^*)$ matrices $\mathbf{A}(\psi_{\mathbf{m}}^*)$ and $\mathbf{B}(\psi_{\mathbf{m}}^*)$, which are defined as follows:

$$\begin{aligned} [\mathbf{A}(\psi_{\mathbf{m}}^*)]_{k,l} &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \int \frac{-\partial^2 \ln s_{\psi_{\mathbf{m}}^*}(\mathbf{y} | \mathbf{x}_i)}{\partial \psi_{\mathbf{m},k}^* \partial \psi_{\mathbf{m},l}^*}(\mathbf{y} | \mathbf{x}_i) s_0^*(\mathbf{y} | \mathbf{x}_i) d\mathbf{y} \right], \\ [\mathbf{B}(\psi_{\mathbf{m}}^*)]_{k,l} &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \int \frac{\partial \ln s_{\psi_{\mathbf{m}}^*}(\mathbf{y} | \mathbf{x}_i)}{\partial \psi_{\mathbf{m},k}^*}(\mathbf{y} | \mathbf{x}_i) \frac{\partial \ln s_{\psi_{\mathbf{m}}^*}(\mathbf{y} | \mathbf{x}_i)}{\partial \psi_{\mathbf{m},l}^*}(\mathbf{y} | \mathbf{x}_i) s_0^*(\mathbf{y} | \mathbf{x}_i) d\mathbf{y} \right]. \end{aligned}$$


Then, the results from [98] and [22] imply that $\mathbb{E}[\text{KL}^{\otimes n}(s_0^*, \widehat{s}_{\mathbf{m}}^*)]$ is asymptotically equivalent to

$$\text{KL}^{\otimes n}(s_0^*, s_{\psi_{\mathbf{m}}^{**}}) + \frac{1}{2n} \text{tr} \left(\mathbf{B}(\psi_{\mathbf{m}}^{**}) \mathbf{A}(\psi_{\mathbf{m}}^{**})^{-1} \right),$$


where we defined $\psi_{\mathbf{m}}^{**} = \text{argmin}_{s_{\psi_{\mathbf{m}}^*} \in S_{\mathbf{m}}^*} \text{KL}^{\otimes n}(s_0^*, s_{\psi_{\mathbf{m}}^*})$.

In particular, $\mathbb{E}[\text{KL}^{\otimes n}(s_0^*, \widehat{s}_{\mathbf{m}}^*)]$ is asymptotically equivalent to $\frac{1}{2n} \dim(S_{\mathbf{m}}^*)$, whenever s_0^* belongs to the model collection $S_{\mathbf{m}}^*$. Furthermore, even though there is no theoretical guarantee, the slope of the mean error in the misspecified case seems also to grow at the same rate as $\frac{\dim(S_{\mathbf{m}}^*)}{2n}$, for large enough number of mixture components ($K \geq 6$ in the WS case and $K \geq 9$ in the MS case).

Figure 13 shows that the error decays when the sample size n grows, when using the penalty based on the jump criterion. The first remark is that we observed the error decay is of order t/n , as predicted in by the theory, where t is some constant, as expected in the well-specified case. The rate of convergence for the misspecified case seems to be slower.




(a) 2000 data points.




(b) 10000 data points.

FIG 11. Box-plot of the tensorized Kullback–Leibler divergence according to the number of mixture components in WS case using the jump criterion over 100 trials. The tensorized Kullback–Leibler divergence of the penalized estimator \hat{s}_K is shown in the right-most box-plot of each graph.




(a) 2000 data points.




(b) 10000 data points.

FIG 12. Box-plot of the tensorized Kullback-Leibler divergence according to the number of mixture components in MS case using the jump criterion over 100 trials. The tensorized Kullback-Leibler divergence of the penalized estimator \hat{s}_K is shown in the right-most box-plot of each graph.



(a) WS Example: The slope of the free regression line is ≈ -1.287 and $t = 3$.



(b) MS Example: The slope of the free regression line is ≈ -0.6120 and $t = 20$.

FIG 13. Tensorized Kullback–Leibler divergence between the true and selected densities based on the jump criterion, represented in a log-log scale, using 30 trials. A free least-square regression with standard error and a regression with slope -1 were added to stress the two different behavior for each graph.

5.3. Ethanol data set

We now consider the use of GLoME models for performing clustering and regression tasks on a real data set. Following the numerical experiments from [100] and [70], we demonstrate our model on the ethanol data set of [16]. The data comprises of 88 observations, which represent the relationship between the engine's concentration of nitrogen oxide (NO) emissions and the equivalence ratio (ER), a measure of the air-ethanol mix, used as a spark-ignition engine fuel in a single-cylinder automobile test (Figures 18a and 18e). Our goal is then to estimate the parameters of a GLoME model, as well as the number of mixture components.


More precisely, we first use the EM algorithm from the `xLLiM` package to compute the forward PMLE of (1), for each $K \in [12]$, on the Ethanol data set. Then, based on the slope heuristic (Figures 14 to 17), we select the best model. Given the estimators of the model chosen, we obtain the estimated conditional density and clustering by applying the maximum a posteriori probability (MAP) rule (Figures 18 and 19).

Because we only have 88 data points and 6 parameters per class, the EM algorithm is strongly dependent on the random initialization of the parameters. One solution is that we can modify slightly that procedure in order to guarantee that at least 10 points are assigned to each class so that the estimated parameters are more stable (cf. [70]). In this work, we wish to investigate how well our proposed PMLE performs for detecting the best number of mixture components for the GLoME model. Thus, we run our experiment over 100 trials with different initializations for the EM algorithm. Histograms of selected values of K are presented in Figures 14 and 15. Notice that it is quite unlikely that the true conditional PDF of Ethanol data set belongs to our hypothesised collection of GLoME models. In fact, this phenomenon has been observed in the MS case, Figures 7 to 10, on the simulated data set. We believe that this is due to the simplistic affine models used in our experiments. Furthermore, it seems that the jump criterion outperformed the slope criterion in the stability of order selection for GLoME models, as previously observed.


Based on the highest empirical probabilities in all situations, our procedure selects $K = 4$ components, which is consistent with the results from [70]. It is worth noting that if we consider the regression of NO with respect to ER, our proposed PMLE of GLoME performs very well for both the clustering and regression tasks (Figure 19). Here, instead of considering the variable NO as the covariate, we use it as the response variable. Then, the resulting clustering, the estimated mean function (black curve) and mixing probabilities are more easily interpretable. This is very similar to the results obtained in [70].

6. Conclusion and perspectives

We have studied the PMLEs for GLoME and BLoME regression models. Our main contributions are to establish non-asymptotic risk bounds that take the




(a) Jump criterion.




(b) Slope criterion.

FIG 14. Histogram of selected K of GLoME on Ethanol data set based on NO using slope heuristic.

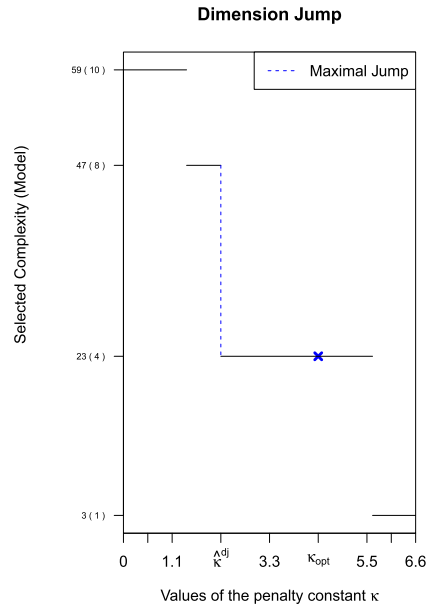


(a) Jump criterion.




(b) Slope criterion.

FIG 15. Histogram of selected K of GLoME on Ethanol data set based on ER using slope heuristic.




(a) Jump criterion based on NO




(b) Jump criterion based on ER

FIG 16. The jump criterion corresponding to the models chosen with highest empirical probabilities.




(a) Slope criterion based on NO




(b) Slope criterion based on ER


FIG 17. The slope criterion corresponding to the models chosen with highest empirical probabilities.




(a) Raw Ethanol data set based on NO.




(b) Clustering by GLoME based on NO.




(c) 3D view of the resulting conditional density based on NO with the 4 clusters.




(d) 2D view of the same conditional density on NO.




(e) Raw Ethanol data set based on ER.



(f) Clustering by GLoME based on ER.



(g) 3D view of the estimated conditional density with the 4 clusters.



(h) 2D view of the same conditional density.

FIG 18. Estimated conditional density with 4 components based upon on the covariate NO or ER from the Ethanol data set.

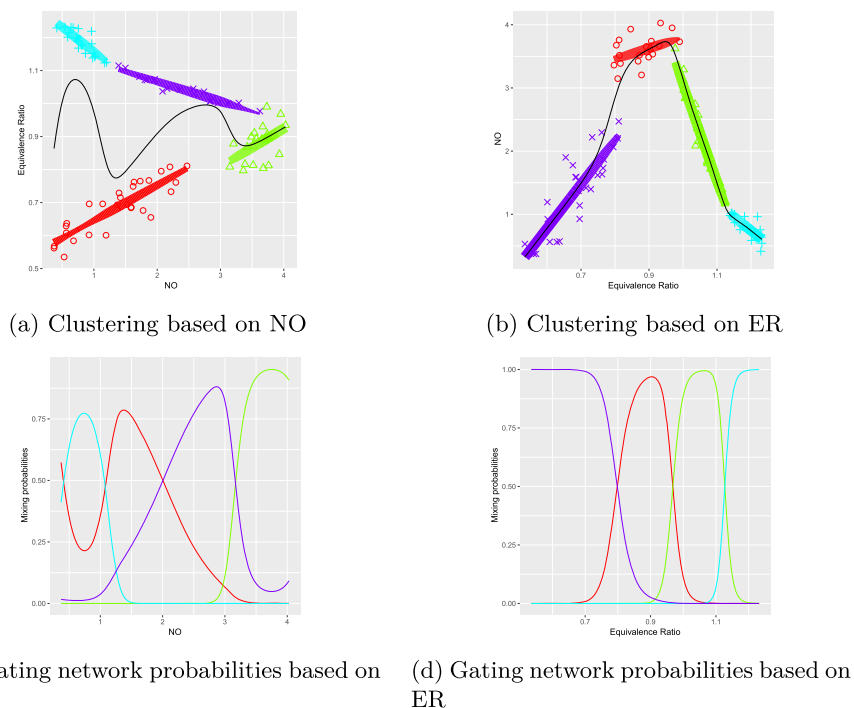


FIG 19. Clustering of Ethanol data set. The black curves present the estimated mean functions. The size of the component mean functions corresponds to the posterior mixture proportions.

form of finite-sample weak oracle inequalities, provided that lower bounds on the penalties hold true. We believe that our contributions help to popularize GLoME and BLoME models as well as slope heuristics, by giving some non-asymptotic theoretical foundations for model selection technique in this area and demonstrating some interesting numerical schemes and experiments. In particular, we aim to provide extensions of the current finite-sample oracle inequalities, Theorems 3.1 and 3.2, to more general frameworks where Gaussian experts are replaced by the elliptical distributions in the future work.

Some important theoretical issues regarding the tightness of the bounds of the PMLEs are still open. As per the SGaME models from [70], we also have to deal with three potential issues: the differences of divergence on the left ($\text{JKL}_\rho^{\otimes n}$) and the right ($\text{KL}^{\otimes n}$) hand side, $C_1 > 1$, and the relationship between $\frac{\text{pen}(\mathbf{m})}{n}$ and the variance. The first issue is important as in general we have $\text{JKL}_\rho^{\otimes n}(s_0, s_{\mathbf{m}}) \leq \text{KL}^{\otimes n}(s_0, s_{\mathbf{m}})$. However, (23) ensures that the two divergences are *equivalent* under regularity conditions. Namely, when

$$\sup_{\mathbf{m} \in \mathcal{M}} \sup_{s_{\mathbf{m}} \in S_{\mathbf{m}}} \|s_0/s_{\mathbf{m}}\|_\infty < \infty.$$

Such a strong assumption is satisfied as long as \mathcal{X} is compact, s_0 is com-

pactly supported, and the regression functions are uniformly bounded, and under the condition that there is a uniform lower bound on the eigenvalues of the covariance matrices. When $s_0 \notin (S_{\mathbf{m}})_{\mathbf{m} \in \mathcal{M}}$, the error bound converges to $C_1 \inf_{s_{\mathbf{m}} \in S_{\mathbf{m}}} \text{KL}^{\otimes n}(s_0, s_{\mathbf{m}})$, which may be large. Nevertheless, as we consider GLoME models, some recent results from [73, 75] imply that if we take a sufficiently large number of mixture components, we can approximate a broad class of densities, and thus the term on the right hand side is small for K sufficiently large. This improves the error bound even when s_0 does not belong to $S_{\mathbf{m}}$ for any $\mathbf{m} \in \mathcal{M}$. For the last issue, it holds that $\frac{\text{pen}(\mathbf{m})}{n}$ is approximately proportional to the asymptotic variance, $\frac{\text{dim}(S_{\mathbf{m}})}{n}$, in the parametric case, see its justification in Section 3.3.

Finally, it is also interesting to point out that the weak inequalities in Theorems 3.1 and 3.2 make use of different divergences and require some regularity assumptions to be considered proper oracle inequalities. To illustrate the strictness of the compactness assumption for s_0 , we only need to consider s_0 as a univariate Gaussian PDF, which obviously does not satisfy such a hypothesis. This motivates [79, Theorem 3.1] to investigate an l_1 -ball model selection procedure and an l_1 -oracle inequality of SGaME with the LASSO estimator. Such an l_1 -oracle inequality can be considered as a complementary to Theorem 3.1. In particular, the most intriguing property of the l_1 -oracle inequality is that it requires only the boundedness assumption of the parameters of the model, which is also required in Theorem 3.1, as well as in [92, 68, 30, 79]. Note that such a mild assumption is quite common when working with MLE (cf. [7, 62]), to tackle the problem of the unboundedness of the likelihood at the boundary of the parameter space [66, 88], and to prevent it from diverging. Nevertheless, by using the smaller divergence: $\text{JKL}_{\rho}^{\otimes n}$ (or more strict assumptions on s_0 and $s_{\mathbf{m}}$, with the same divergence $\text{KL}^{\otimes n}$ as ours), Theorem 3.1 obtain a faster convergence rate of order $1/n$ compared to $1/\sqrt{n}$ in the l_1 -oracle inequality. Therefore, in the case where there is no guarantee of a compact support of s_0 or uniformly bounded regression functions, the l_1 -oracle inequality gives a theoretical foundation for the l_1 -ball model selection procedure, with the order of rate convergence of $1/\sqrt{n}$, with only the boundedness assumption on the parameter space.

Appendix A: Our contributions on lemma proofs

A.1. Proof of Lemma 2.1: Inverse regression trick

We firstly aim to provide the proof of one-to-one correspondence defines the link between the inverse and forward conditional distributions not only for the special case of Gaussian distribution in (17) but also for elliptical distributions (cf. [19, 38]). It is worth mentioning that the multivariate normal distribution, multivariate t -distribution and multivariate Laplace distribution are some instances of elliptical distributions (cf. [40, Chapter 1], [46]). Note that a statement similar to the following has been proved in the linear regression setting in

[35, Section 2.2]. We include a proof for mixture of regression models, which is an extension to the aforementioned result.

A.1.1. Elliptically symmetric distributions

We will provide the proof of Lemma 2.1 by using some general results regarding elliptical distributions.

Definition A.1. Let \mathbf{X} be a D -dimensional random vector. Then \mathbf{X} is said to be *elliptically distributed* (or simply *elliptical*) if and only if there exist a vector $\boldsymbol{\mu} \in \mathbb{R}^D$, a positive semidefinite matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{D \times D}$ and a function $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}$ such that the characteristic function $\mathbf{t} \mapsto \varphi_{\mathbf{X}-\boldsymbol{\mu}}(\mathbf{t})$ of $\mathbf{X} - \boldsymbol{\mu}$ corresponds to $\mathbf{t} \mapsto \phi(\mathbf{t}^\top \boldsymbol{\Sigma} \mathbf{t})$, $\mathbf{t} \in \mathbb{R}^D$. We write $\mathbf{X} \sim \mathcal{E}_D(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \phi)$ to denote that \mathbf{X} is elliptical.

The function ϕ is referred to as the *characteristic generator* of \mathbf{X} . When $D = 1$ the class of elliptical distributions coincides with the class of univariate symmetric distributions. Thanks to Proposition 1 from [40], it holds that every affinely transformed elliptical random vector is elliptically distributed. Moreover, the following stochastic representation theorem, Theorem A.2, shows that the converse is true if the transformation matrix has full rank.

Theorem A.2 (Theorem 1 from [19]). *Let \mathbf{X} be a D -dimensional random vector. Then $\mathbf{X} \sim \mathcal{E}_D(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \phi)$ with $\text{rank}(\boldsymbol{\Sigma}) = k$ if and only if*

$$\mathbf{X} = \boldsymbol{\mu} + \mathcal{R}\boldsymbol{\Lambda}\mathbf{U}^{(k)},$$

where $\mathbf{U}^{(k)}$ is a k -dimensional random vector uniformly distributed on the unit hypersphere with $k - 1$ dimensions $\mathcal{S}^{k-1} = \{\mathbf{x} \in \mathbb{R}^D : \|\mathbf{x}\|_2 = 1\}$, \mathcal{R} is a non-negative random variable with distribution function F related to ϕ being stochastically independent of $\mathbf{U}^{(k)}$, $\boldsymbol{\mu} \in \mathbb{R}^D$ and $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}\boldsymbol{\Lambda}^\top$ is a rank factorization of $\boldsymbol{\Sigma}$ where $\boldsymbol{\Lambda} \in \mathbb{R}^{D \times k}$ with $\text{rank}(\boldsymbol{\Lambda}) = k$.

Note that via the transformation matrix $\boldsymbol{\Lambda}$, the spherical random vector $\mathbf{U}^{(k)}$ produces elliptically contoured density surfaces, whereas the generating random variable \mathcal{R} determines the distribution’s shape, in particular the heaviness of the distribution’s tails. Further, $\boldsymbol{\mu}$ determines the location of the random vector \mathbf{X} . The matrix $\boldsymbol{\Sigma}$ is called the *dispersion matrix* or *scatter matrix* of \mathbf{X} . Therefore, it holds that every elliptical distribution belongs to a location-scale-family [51] defined by an underlying spherical standard distribution.

Example A.3 (Multivariate normal distribution). Let $\boldsymbol{\mu} \in \mathbb{R}^D$ and $\boldsymbol{\Lambda} \in \mathbb{R}^{D \times k}$ such that $\boldsymbol{\Sigma} := \boldsymbol{\Lambda}\boldsymbol{\Lambda}^\top \in \mathbb{R}^{D \times D}$ is positive definite. The random vector $\mathbf{X} \sim \Phi_D(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is elliptically distributed since \mathbf{X} is representable as $\mathbf{X} = \boldsymbol{\mu} + \sqrt{\chi_k^2}\boldsymbol{\Lambda}\mathbf{U}^{(k)}$. The underlying spherical standard distribution is the standard normal distribution. Further, since $s \mapsto \exp(-s/2)$ is the characteristic generator for the class of normal distributions, the characteristic function of $\mathbf{X} - \boldsymbol{\mu}$ corresponds to $\mathbf{t} \mapsto \phi_{\mathbf{X}-\boldsymbol{\mu}}(\mathbf{t}) = \exp(-\mathbf{t}^\top \boldsymbol{\Sigma} \mathbf{t})$, $\mathbf{t} \in \mathbb{R}^d$.

We next describe some important results on the conditional distributions of elliptical random vectors (cf. [19, Corollary 5], [40, Chapter 1]).

Theorem A.4. *Let $\mathbf{X} \sim \mathcal{E}_D(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \phi)$ with $\text{rank}(\boldsymbol{\Sigma}) = k$. It holds that:*

- (a) $\mathbb{E}(\mathbf{X}) = \boldsymbol{\mu}$.
- (b) $\text{var}(\mathbf{X}) = \frac{\mathbb{E}(\mathcal{R}^2)}{k} \boldsymbol{\Sigma} = -2\phi'(0)\boldsymbol{\Sigma}$, if ϕ is differentiable at 0.
- (c) *The sum of independent elliptically distributed random vector with the same dispersion matrix $\boldsymbol{\Sigma}$ is elliptically too. Furthermore, the sum of two dependent elliptical random vectors with the same dispersion matrix, which are dependent only through their radial parts \mathcal{R} , is also elliptical [46, Theorem 4.1]. More precisely, let \mathcal{R} and $\tilde{\mathcal{R}}$ be nonnegative random variables and let $\mathbf{X} := \boldsymbol{\mu} + \mathcal{R}\mathbf{Z} \sim \mathcal{E}_D(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \phi)$ and $\tilde{\mathbf{X}} := \tilde{\boldsymbol{\mu}} + \tilde{\mathcal{R}}\tilde{\mathbf{Z}} \sim \mathcal{E}_D(\tilde{\boldsymbol{\mu}}, \boldsymbol{\Sigma}, \tilde{\phi})$, where $(\mathcal{R}, \tilde{\mathcal{R}})$, \mathbf{Z} , and $\tilde{\mathbf{Z}}$ are independent. Then $\mathbf{X} + \tilde{\mathbf{X}} \sim \mathcal{E}_D(\boldsymbol{\mu} + \tilde{\boldsymbol{\mu}}, \boldsymbol{\Sigma}, \phi^*)$.*
- (d) *Affine transformation: every affinely transformed and particularly every linearly combined elliptical random vector is elliptical, too. More formally, for any $\mathbf{b} \in \mathbb{R}^L$, $\mathbf{A} \in \mathbb{R}^{L \times D}$, and $\mathbf{Y} = \mathbf{b} + \mathbf{A}\mathbf{X}$ where $\mathbf{X} = \boldsymbol{\mu} + \mathcal{R}\boldsymbol{\Lambda}\mathbf{U}^{(k)}$ with $\boldsymbol{\Lambda} \in \mathbb{R}^{D \times k}$, it follows that $\mathbf{Y} = \mathcal{E}_L(\mathbf{b} + \mathbf{A}\boldsymbol{\mu}, \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top, \phi)$ since*

$$\mathbf{Y} = \mathbf{b} + \mathbf{A}(\boldsymbol{\mu} + \mathcal{R}\boldsymbol{\Lambda}\mathbf{U}^{(k)}) = (\mathbf{A}\boldsymbol{\mu} + \mathbf{b}) + \mathcal{R}\mathbf{A}\boldsymbol{\Lambda}\mathbf{U}^{(k)}. \tag{43}$$

- (e) *Marginal distribution: let $\mathcal{P}_m \in \{0, 1\}^{m \times D}$ ($m \leq D$) be a permutation and deletion matrix, i.e., \mathcal{P}_m has only binary entries of 0's and 1's and $\mathcal{P}_m\mathcal{P}_m^\top = \mathbf{I}_m$. So the transformation $\mathcal{P}_m\mathbf{X} =: \mathbf{Y}$ permutes and deletes certain components of \mathbf{X} such that \mathbf{Y} is a k -dimensional random vector containing the remaining components of \mathbf{X} and having a (multivariate) marginal distribution with respect to the joint distribution of \mathbf{X} . Then by (43), we obtain $\mathbf{Y} \sim \mathcal{E}_m(\mathcal{P}_m\boldsymbol{\mu}, \mathcal{P}_m\boldsymbol{\Sigma}\mathcal{P}_m^\top, \phi)$ since*

$$\mathbf{Y} = \mathcal{P}_m(\boldsymbol{\mu} + \mathcal{R}\boldsymbol{\Lambda}\mathbf{U}^{(k)}) = \mathcal{P}_m\boldsymbol{\mu} + \mathcal{R}\mathcal{P}_m\boldsymbol{\Lambda}\mathbf{U}^{(k)}. \tag{44}$$

- (f) *Conditional distribution: let $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2)$, where \mathbf{X}_1 is a k -dimensional sub-vector of \mathbf{X} , and let $\boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{pmatrix} \in \mathbb{R}^{D \times D}$. Provided the conditional random vector $\mathbf{X}_2 \mid \mathbf{X}_1 = \mathbf{x}_1$ exists, it is also elliptically distributed: $\mathbf{X}_2 \mid (\mathbf{X}_1 = \mathbf{x}_1) \sim \mathcal{E}_{D-k}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*, \phi^*)$. Moreover, it can be presented stochastically by*

$$\mathbf{X}_2 \mid (\mathbf{X}_1 = \mathbf{x}_1) = \boldsymbol{\mu}^* + \mathcal{R}^*\mathbf{U}^{(D-k)}\boldsymbol{\Gamma}^*$$

and $\mathbf{U}^{(D-k)}$ is uniformly distributed on $\mathcal{S}^{(D-k-1)}$, and

$$\begin{aligned} \mathcal{R}^* &= \mathcal{R}\sqrt{1 - \beta} \mid \left(\mathcal{R}\sqrt{\beta}\mathbf{U}^{(k)} = \boldsymbol{\Sigma}_{11}^{-1}(\mathbf{x}_1 - \boldsymbol{\mu}_1) \right), \\ \boldsymbol{\mu}^* &= \boldsymbol{\mu}_2 + \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}(\mathbf{x}_1 - \boldsymbol{\mu}_1), \end{aligned}$$

$$\Sigma^* = \Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12}.$$

where $\beta \sim \text{Beta}(k/2, (D - k)/2)$ and $\mathcal{R}, \beta, \mathbf{U}^{(k)}$ and $\mathbf{U}^{(D-k)}$ are mutually independent, and $\Sigma^* = (\Gamma^*)^\top \Gamma^*$.

A.1.2. Relation between forward and inverse regression

Proposition 1 from [26], a multivariate extension of [47], leads to a link between GLLiM, defined in (16) models, and a Gaussian mixture model on the joint variable $[\mathbf{X}; \mathbf{Y}]$. This result motivates us to establish the general proof for the relationship between forward and inverse mixture of elliptical regression models. More precisely, we consider the following generative model, conditionally on the cluster label:

$$[\mathbf{X}; \mathbf{Y}] \mid (Z = k) \sim \mathcal{E}_{L+D}(\mathbf{m}_k, \mathbf{V}_k, \phi), \tag{45}$$

where \mathcal{E}_{L+D} denotes an elliptical distribution of dimension $D + L$, and are \mathbf{m}_k and \mathbf{V}_k its location and scale parameters, respectively.

When applying the inverse regression strategy in the context of mixture of elliptical locally-linear mapping, the key point is to account for (46):

$$\mathbf{X} = \sum_{k=1}^K \mathbb{I}(Z = k) (\mathbf{A}_k \mathbf{Y} + \mathbf{b}_k + \mathbf{E}_k), \tag{46}$$

where $\mathbf{A}_k \in \mathbb{R}^{D \times L}$ and vector $\mathbf{b}_k \in \mathbb{R}^D$, and \mathbf{E}_k is an error term capturing both the reconstruction error due to the local affine approximation and the observation noise in \mathbb{R}^D , into the parameterization of \mathbf{m}_k and \mathbf{V}_k . Given $Z = k$, it follows from (45) that \mathbf{Y} is also elliptical distribution by using Theorem A.4 (e) and \mathbf{Y} can be assumed to have a location $\mathbf{c}_k \in \mathbb{R}^L$ and a scale matrix $\Gamma_k \in \mathbb{R}^{L \times L}$. We assume further that the error term $\mathbf{E}_k \sim \mathcal{E}(\mathbf{0}, \phi_{e_k}, \Sigma_k)$ is an unobserved centered elliptical random noise with residual covariance matrix Σ_k of type ϕ_{e_k} , and is independent of \mathbf{Y} . Then, using (46) and Theorem A.4, we have

$$\begin{aligned} \mathbb{E}(\mathbf{X} \mid (Z = k)) &= \mathbb{E}(\mathbf{A}_k \mathbf{Y} + \mathbf{b}_k + \mathbf{E}_k) = \mathbf{A}_k \mathbf{c}_k + \mathbf{b}_k, \\ \text{var}(\mathbf{X} \mid (Z = k)) &= \text{var}(\mathbf{A}_k \mathbf{Y}) + \text{var}(\mathbf{E}_k) = \mathbf{A}_k \Gamma_k \mathbf{A}_k^\top + \Sigma_k, \\ \text{cov}(\mathbf{X}, \mathbf{Y} \mid (Z = k)) &= \mathbf{A}_k \Gamma_k, \text{cov}(\mathbf{Y}, \mathbf{X}) = \text{cov}(\mathbf{Y}, \mathbf{A}_k \mathbf{Y}) = \Gamma_k^\top \mathbf{A}_k^\top, \\ \mathbf{m}_k &= \begin{bmatrix} \mathbf{A}_k \mathbf{c}_k + \mathbf{b}_k \\ \mathbf{c}_k \end{bmatrix}, \\ \mathbf{V}_k &= \begin{bmatrix} \Sigma_k + \mathbf{A}_k \Gamma_k \mathbf{A}_k^\top & \mathbf{A}_k \Gamma_k \\ (\mathbf{A}_k \Gamma_k)^\top & \Gamma_k \end{bmatrix}. \end{aligned} \tag{47}$$

Note that in the forward and inverse regression problems of elliptical locally-linear mapping (containing the Gaussian case (10), (11), (15) and (14)), by using Theorem A.4, the joint distribution defined in (45) allows us to consider

a mixture of linear regression problem, characterized by the following marginal and conditional distributions:

$$\mathbf{X} \mid (Z = k) \sim \mathcal{E}_D(\mathbf{c}_k^*, \mathbf{\Gamma}_k^*, \phi), \quad (48)$$

$$\mathbf{Y} \mid (\mathbf{X}, Z = k) = \mathbf{A}_k^* \mathbf{X} + \mathbf{b}_k^* + \mathbf{E}_k^*, \quad (49)$$

$$\mathbf{Y} \mid (Z = k) \sim \mathcal{E}_D(\mathbf{c}_k, \mathbf{\Gamma}_k, \phi), \quad (50)$$

$$\mathbf{X} \mid (\mathbf{Y}, Z = k) = \mathbf{A}_k \mathbf{Y} + \mathbf{b}_k + \mathbf{E}_k, \quad (51)$$

where $\mathbf{E}_k \sim \mathcal{E}(\mathbf{0}, \phi_{e_k}, \mathbf{\Sigma}_k)$ and $\mathbf{E}_k^* \sim \mathcal{E}(\mathbf{0}, \phi_{e_k^*}, \mathbf{\Sigma}_k^*)$.

Then, we claim that the joint distribution, defined in (45) and (47), leads to the marginal and the conditional distributions of Equations (48) to (51) and to a mapping between their mean and variance parameters. Indeed, by using conditioning properties of elliptical distributions, see more in Theorem A.4, implies the following marginal for \mathbf{X} and conditional distribution for \mathbf{Y} given \mathbf{X} as follows:

$$\mathbf{X} \mid (Z = k) \sim \mathcal{E}_D(\mathbf{A}_k \mathbf{c}_k + \mathbf{b}_k, \mathbf{\Sigma}_k + \mathbf{A}_k \mathbf{\Gamma}_k \mathbf{A}_k^\top, \phi), \quad (52)$$

$$\mathbf{Y} \mid (\mathbf{X}, Z = k) \sim \mathcal{E}_L(\mathbf{m}_k^{yx}, \mathbf{\Sigma}_k^{yx}, \tilde{\phi}), \quad (53)$$

where the explicit expression of the characteristic function $\tilde{\phi}$ can be found in [19, Corollary 5], and

$$\begin{aligned} \mathbf{m}_k^{yx} &= \mathbf{c}_k + \mathbf{\Gamma}_k^\top \mathbf{A}_k^\top (\mathbf{\Sigma}_k + \mathbf{A}_k \mathbf{\Gamma}_k \mathbf{A}_k^\top)^{-1} (\mathbf{X} - \mathbf{A}_k \mathbf{c}_k - \mathbf{b}_k), \\ \mathbf{\Sigma}_k^{yx} &= \mathbf{\Gamma}_k - \mathbf{\Gamma}_k^\top \mathbf{A}_k^\top (\mathbf{\Sigma}_k + \mathbf{A}_k \mathbf{\Gamma}_k \mathbf{A}_k^\top)^{-1} \mathbf{A}_k \mathbf{\Gamma}_k = (\mathbf{\Gamma}_k^{-1} + \mathbf{A}_k^\top \mathbf{\Sigma}_k^{-1} \mathbf{A}_k)^{-1}, \end{aligned}$$

with the fact that the last equality is the Woodbury matrix identity. Note that the locations and scale matrices of the conditional distribution do not depend upon the third parameter of the joint distribution, and consequently, we do not describe the explicit expression for $\tilde{\phi}$. We then utilize again the Woodbury matrix identity and the symmetric property of $\mathbf{\Gamma}$ to identify (48) and (49) with (52) and (53), respectively, which implies the following important connections:

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

$$= \Sigma_k^* (\Gamma_k^{-1} \mathbf{c}_k - \mathbf{A}_k^\top \Sigma_k^{-1} \mathbf{b}_k).$$

Therefore, the desired results then are obtained by using the fact that the multivariate normal distribution not only has the property of having Gaussian marginal and conditional distributions [13, Sections 2.3.1 and 2.3.2] but also belongs to elliptical distributions (detailed in Example A.3). Furthermore, it should be stressed that several versions of multivariate t -distributions (e.g., Section 5.5, page 94 of [54], [36]) have the previous property. This leads to the inverse regression model based on the multivariate t -distributions [84]. It will be interesting to find other sub-classes of elliptically contoured distributions that have the closedness property on marginal and conditional distribution so that the previous inverse regression trick can be applied.

A.2. Proof of Lemma 4.5: Bracketing entropy of Gaussian gating networks

Note that the first inequality of Lemma 4.5 comes from [70, Lemma 4] and describes relationship between the bracketing entropy of \mathcal{P}_K and the entropy of \mathcal{W}_K . Therefore, Lemma 4.5 is obtained by proving that there exists a constant $C_{\mathcal{W}_K}$ such that $\forall \delta \in (0, 2]$,

$$\mathcal{H}_{d_{\|\cdot\|_\infty}}(\delta, \mathcal{W}_K) \leq \dim(\mathcal{W}_K) \left(C_{\mathcal{W}_K} + \ln \left(\frac{1}{\delta} \right) \right), \quad (54)$$

where $\dim(\mathcal{W}_K) = K - 1 + KL + K \frac{L(L+1)}{2}$.

In order to establish the proof for (54), we have to construct firstly the δ_π -covering $\Pi_{K-1, \omega}$ of Π_{K-1} via Lemma A.5, which is proved in Appendix A.2.1.

Lemma A.5 (Covering number of probability simplex with maximum norm). *Given any $\delta_\pi > 0$, any $\pi \in \Pi_{K-1}$, we can choose $\hat{\pi} \in \Pi_{K-1, \omega}$, an δ_π -covering of Π_{K-1} , so that $\max_{k \in [K]} |\pi_k - \hat{\pi}_k| \leq \delta_\pi$. Furthermore, it holds that*

$$\mathcal{N}(\delta_\pi, \Pi_{K-1}, \|\cdot\|_\infty) \leq \frac{K (2\pi e)^{K/2}}{\delta_\pi^{K-1}}. \quad (55)$$

Then, by definition of the covering number, (54) is obtained immediately via Lemma A.6, which controls the covering number of \mathcal{W}_K and is proved in Appendix A.2.2.

Lemma A.6. *Given a bounded set \mathcal{Y} in \mathbb{R}^L such that $\mathcal{Y} = \{\mathbf{y} \in \mathbb{R}^L : \|\mathbf{y}\|_\infty \leq C_{\mathcal{Y}}\}$, it holds that \mathcal{W}_K has a covering number satisfied $\mathcal{N}(\delta, \mathcal{W}_K, d_{\|\cdot\|_\infty}) \leq C \delta^{-\dim(\mathcal{W}_K)}$, for some constant C .*

Indeed, Lemma A.6 implies the desired result by noting that

$$\mathcal{H}_{d_{\|\cdot\|_\infty}}(\delta, \mathcal{W}_K) = \ln \mathcal{N}(\delta, \mathcal{W}_K, d_{\|\cdot\|_\infty}) \leq \ln \left[\frac{C}{\delta^{\dim(\mathcal{W}_K)}} \right]$$

$$\begin{aligned}
 &= \dim(\mathcal{W}_K) \left[\frac{1}{\dim(\mathcal{W}_K)} \ln C + \ln \left(\frac{1}{\delta} \right) \right] \\
 &= \dim(\mathcal{W}_K) \left(C_{\mathcal{W}_K} + \ln \left(\frac{1}{\delta} \right) \right).
 \end{aligned}$$

A.2.1. Proof of Lemma A.5

Note that [41, Lemma 2] provide a result for controlling a δ_π -Hellinger bracketing of $\mathbf{\Pi}_{K-1}$. However, such result can not be applied for our Lemma A.5 since they use δ_π -Hellinger bracketing entropy while we use δ_π -covering number for the probability complex with maximum norm.

Given any $\boldsymbol{\pi} = (\pi_k)_{k \in [K]} \in \mathbf{\Pi}_{K-1}$, let $\boldsymbol{\xi} = (\boldsymbol{\xi}_k)_{k \in [K]}$ where $\boldsymbol{\xi}_k = \sqrt{\pi_k}$, $\forall k \in [K]$. Then $\boldsymbol{\pi} \in \mathbf{\Pi}_{K-1}$ if and only if $\boldsymbol{\xi} \in Q^+ \cap U$, where U is the surface of the unit sphere and Q^+ is the positive quadrant of \mathbb{R}^K . Next, we divide the unit cube in \mathbb{R}^K into disjoint cubes with sides parallel to the axes and sides of length δ_π/\sqrt{K} . Let $(\mathbf{C}_j)_{j \in [N]}$ is the subset of these cubes that have non-empty intersection with $Q^+ \cap U$. For any $j \in [N]$, let $\boldsymbol{\nu}_j = (\boldsymbol{\nu}_{j,k})_{k \in [K]}$ be the center of the cube \mathbf{C}_j and $\boldsymbol{\nu}_j^2 = (\boldsymbol{\nu}_{j,k}^2)_{k \in [K]}$.

Then $\{\boldsymbol{\nu}_j\}_{j \in [N]}$ is a $\delta_\pi / (2\sqrt{K})$ -covering of $Q^+ \cap U$, since we have for any $\boldsymbol{\xi} = (\boldsymbol{\xi}_k)_{k \in [K]} \in Q^+ \cap U$, there exists $j_0 \in [N]$ such that $\boldsymbol{\xi} \in \mathbf{C}_{j_0}$, and

$$\|\boldsymbol{\xi} - \boldsymbol{\nu}_{j_0}\|_\infty = \max_{k \in [K]} |\boldsymbol{\xi}_k - \boldsymbol{\nu}_{j_0,k}| \leq \frac{\delta_\pi}{2\sqrt{K}}. \tag{56}$$

Therefore, it follows that $\mathbf{\Pi}_{K-1,\omega} := \{\boldsymbol{\nu}_j^2\}_{j \in [N]}$ is a δ_π -covering of $\mathbf{\Pi}_{K-1}$, since for any $\boldsymbol{\pi} = (\pi_k)_{k \in [K]} \in \mathbf{\Pi}_{K-1}$, (56) leads to the existence of $j_0 \in [N]$, such that

$$\begin{aligned}
 \|\boldsymbol{\pi} - \boldsymbol{\nu}_{j_0}^2\|_\infty &= \max_{k \in [K]} |\boldsymbol{\xi}_k^2 - \boldsymbol{\nu}_{j_0,k}^2| = \max_{k \in [K]} \{|\boldsymbol{\xi}_k - \boldsymbol{\nu}_{j_0,k}| |\boldsymbol{\xi}_k + \boldsymbol{\nu}_{j_0,k}|\} \\
 &\leq \frac{\delta_\pi}{2\sqrt{K}} \max_{k \in [K]} |\boldsymbol{\xi}_k + \boldsymbol{\nu}_{j_0,k}| \leq \frac{\delta_\pi}{\sqrt{K}} \leq \delta_\pi,
 \end{aligned}$$

where we used the fact that $\max_{k \in [K]} |\boldsymbol{\xi}_k + \boldsymbol{\nu}_{j_0,k}| \leq 2$. Now, it remains to count the number of cubes N . Let $\mathcal{T}_a = \{\mathbf{z} \in Q^+ : \|\mathbf{z}\|_2 \leq a\}$ and let $\mathcal{C} = \bigcup_{j \in [N]} \mathbf{C}_j$. Note that $\mathcal{C} \subset \mathcal{T}_{1+\delta_\pi} - \mathcal{T}_{1-\delta_\pi} \equiv \mathcal{T}$, and so

$$\text{Volume}(\mathcal{T}) \geq \text{Volume}(\mathcal{C}) = N \left(\frac{\delta_\pi}{\sqrt{K}} \right)^K.$$

Note that here we use the notation π for the Archimedes' constant, which differs from $\boldsymbol{\pi} = (\pi_k)_{k \in [K]}$ for the mixing proportion of the GLoME model. Then, we define $\mathcal{V}_K(a) = a^K \pi^{K/2}$ as the volume of a sphere of radius a . Since $z! \geq z^z e^{-z}$

and $(1 + \delta_\pi)^K - (1 - \delta_\pi)^K = K \int_{1-\delta_\pi}^{1+\delta_\pi} z^{K-1} dz \leq 2\delta_\pi K (1 + \delta_\pi)^{K-1}$, it follows that

$$\begin{aligned} & \mathcal{N}(\delta_\pi, \mathbf{\Pi}_{K-1}, \|\cdot\|_\infty) \\ & \leq N \leq \frac{\text{Volume}(\mathcal{C})}{(\delta_\pi/\sqrt{K})^K} = \frac{1}{2^K} \frac{\mathcal{V}_K(1 + \delta_\pi) - \mathcal{V}_K(1 - \delta_\pi)}{(\delta_\pi/\sqrt{K})^K} \\ & = \frac{1}{2^K} \frac{[(1 + \delta_\pi)^K - (1 - \delta_\pi)^K]}{(\delta_\pi/\sqrt{K})^K} \frac{\pi^{K/2}}{(K/2)!} \leq \left(\frac{\pi e}{2}\right)^{K/2} \frac{[(1 + \delta_\pi)^K - (1 - \delta_\pi)^K]}{\delta_\pi^K} \\ & \leq \frac{K(2\pi e)^{K/2}}{\delta_\pi^{K-1}}. \end{aligned}$$

A.2.2. Proof of Lemma A.6

In order to find an upper bound for a covering number of \mathcal{W}_K , we wish to construct a finite δ -covering $\mathbf{W}_{K,\omega}$ of \mathcal{W}_K , with respect to the distance $d_{\|\cdot\|_\infty}$. That is, given any $\delta > 0$, $\mathbf{w}(\cdot; \omega) \in \mathcal{W}_K$, we aim to prove that there exists $\mathbf{w}(\cdot; \hat{\omega}) \in \mathbf{W}_{K,\omega}$ such that

$$d_{\|\cdot\|_\infty}(\mathbf{w}(\cdot; \omega), \mathbf{w}(\cdot; \hat{\omega})) = \max_{k \in [K]} \sup_{\mathbf{y} \in \mathcal{Y}} |\mathcal{W}_K(\mathbf{y}; \omega) - \mathcal{W}_K(\mathbf{y}; \hat{\omega})| \leq \delta. \quad (57)$$

In order to accomplish such task, given any positive constants $\delta_c, \delta_\Gamma, \delta_\pi$, and any $k \in [K]$, let us define

$$\begin{aligned} \mathcal{F} &= \{\mathcal{Y} \ni \mathbf{y} \mapsto \ln(\phi_L(\mathbf{y}; \mathbf{c}, \mathbf{\Gamma})) : \|\mathbf{c}\|_\infty \leq A_c, a_\Gamma \leq m(\mathbf{\Gamma}) \leq M(\mathbf{\Gamma}) \leq A_\Gamma\}, \\ \mathcal{F}_{\mathbf{c}_k} &= \left\{ \ln(\phi_L(\cdot; \mathbf{c}_k, \mathbf{\Gamma}_k)) : \ln(\phi_L(\cdot; \mathbf{c}_k, \mathbf{\Gamma}_k)) \in \mathcal{F}, \right. \\ & \quad \left. \mathbf{c}_{k,j} \in \{-C\mathbf{y} + l\delta_c/L : l = 0, \dots, \lceil 2C\mathbf{y}L/\delta_c \rceil\}, j \in [L] \right\}, \end{aligned} \quad (58)$$

$$\begin{aligned} \mathcal{F}_{\mathbf{c}_k, \mathbf{\Gamma}_k} &= \left\{ \ln(\phi_L(\cdot; \mathbf{c}_k, \mathbf{\Gamma}_k)) : \ln(\phi_L(\cdot; \mathbf{c}_k, \mathbf{\Gamma}_k)) \in \mathcal{F}_{\mathbf{c}_k}; [\text{vec}(\mathbf{\Gamma}_k)]_{i,j} = \gamma_{i,j} \frac{\delta_\Gamma}{L^2}, \right. \\ & \quad \left. \gamma_{i,j} = \gamma_{j,i} \in \mathbb{Z} \cap \left[-\left\lfloor \frac{L^2 A_\Gamma}{\delta_\Gamma} \right\rfloor, \left\lfloor \frac{L^2 A_\Gamma}{\delta_\Gamma} \right\rfloor \right], i \in [L], j \in [L] \right\}, \end{aligned} \quad (59)$$

$$\begin{aligned} \mathbf{W}_{K,\omega} &= \left\{ \mathbf{w}(\cdot; \omega) : \mathbf{w}(\cdot; \omega) \in \mathcal{W}_K, \right. \\ & \quad \left. \forall k \in [K], \ln(\phi_L(\cdot; \mathbf{c}_k, \mathbf{\Gamma}_k)) \in \mathcal{F}_{\mathbf{c}_k, \mathbf{\Gamma}_k}, \boldsymbol{\pi} \in \mathbf{\Pi}_{K-1,\omega} \right\}. \end{aligned} \quad (60)$$

Here, $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ are ceiling and floor functions, respectively, and $\text{vec}(\cdot)$ is an operator that stacks matrix columns into a column vector. In particular, we denote $\mathbf{\Pi}_{K-1,\omega}$ as a δ_π -covering of $\mathbf{\Pi}_{K-1}$, which is defined in Lemma A.5.

By the previous definition, it holds that $\forall k \in [K], \mathcal{F}_{\mathbf{c}_k, \mathbf{\Gamma}_k} \subset \mathcal{F}_{\mathbf{c}_k} \subset \mathcal{F}$, and $\mathbf{W}_{K, \omega} \subset \mathcal{W}_K$.

Next, we claim that $\mathbf{W}_{K, \omega}$ is a finite δ -covering of \mathcal{W}_K with respect to the distance $d_{\|\cdot\|_{\infty}}$. To do this, for any $\mathbf{w}(\cdot; \omega) = (\ln(\pi_k \phi_L(\cdot; \mathbf{c}_k, \mathbf{\Gamma}_k)))_{k \in [K]} \in \mathcal{W}_K$, $\ln(\phi_L(\cdot; \mathbf{c}_k, \mathbf{\Gamma}_k)) \in \mathcal{F}$, $\boldsymbol{\pi} \in \mathbf{\Pi}_{K-1}$, and for any $k \in [K]$, by (58), we first choose a function $\ln(\phi_L(\cdot; \hat{\mathbf{c}}_k, \mathbf{\Gamma}_k)) \in \mathcal{F}_{\mathbf{c}_k}$ so that

$$\|\hat{\mathbf{c}}_k - \mathbf{c}_k\|_1 = \sum_{j=1}^L |\hat{\mathbf{c}}_{k,j} - \mathbf{c}_{k,j}| \leq L \frac{\delta_{\mathbf{c}}}{L} = \delta_{\mathbf{c}}.$$

Furthermore, by (59), we can obtain a result to construct the covariance matrix lattice. That is, any $\ln(\phi_L(\cdot; \hat{\mathbf{c}}_k, \mathbf{\Gamma}_k)) \in \mathcal{F}_{\mathbf{c}_k}$ can be approximated by $\ln(\phi_L(\cdot; \hat{\mathbf{c}}_k, \hat{\mathbf{\Gamma}}_k)) \in \mathcal{F}_{\mathbf{c}_k, \mathbf{\Gamma}_k}$ such that

$$\begin{aligned} \left\| \text{vec}(\hat{\mathbf{\Gamma}}_k) - \text{vec}(\mathbf{\Gamma}_k) \right\|_1 &\equiv \left\| \text{vec}(\hat{\mathbf{\Gamma}}_k - \mathbf{\Gamma}_k) \right\|_1 \\ &= \sum_{i=1}^L \sum_{j=1}^L \left| \left[\text{vec}(\hat{\mathbf{\Gamma}}_k - \mathbf{\Gamma}_k) \right]_{i,j} \right| \leq \frac{L^2 \delta_{\mathbf{\Gamma}}}{L^2} = \delta_{\mathbf{\Gamma}}. \end{aligned} \quad (61)$$

Note that since for any $k \in [K]$, $(\mathbf{y}, \mathbf{c}_k, \text{vec}(\mathbf{\Gamma}_k)) \mapsto \ln(\phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k))$ is differentiable, it is also continuous w.r.t. \mathbf{y} and its parameters \mathbf{c}_k and $\mathbf{\Gamma}_k$. Thus, for every fixed $\mathbf{y} \in \mathcal{Y}$, for every $\hat{\mathbf{c}}_k, \mathbf{c}_k \in \mathcal{X}$ with $\hat{\mathbf{c}}_k \leq \mathbf{c}_k$, and for every $\hat{\mathbf{\Gamma}}_k, \mathbf{\Gamma}_k$, where $\text{vec}(\hat{\mathbf{\Gamma}}_k) \leq \text{vec}(\mathbf{\Gamma}_k)$, we can apply the mean value theorem (see [37, Lemma 2.5.1]) to $\ln(\phi_L(\mathbf{y}; \cdot, \mathbf{\Gamma}_k))$ and $\ln(\phi_L(\mathbf{y}; \hat{\mathbf{c}}_k, \cdot))$ on the intervals $[\hat{\mathbf{c}}_k, \mathbf{c}_k]$ and $[\text{vec}(\hat{\mathbf{\Gamma}}_k), \text{vec}(\mathbf{\Gamma}_k)]$ for some $z_{\mathbf{c}_k} \in (\hat{\mathbf{c}}_k, \mathbf{c}_k)$ and $z_{\mathbf{\Gamma}_k} \in (\text{vec}(\hat{\mathbf{\Gamma}}_k), \text{vec}(\mathbf{\Gamma}_k))$, respectively, to get

$$\ln(\phi_L(\mathbf{y}; \hat{\mathbf{c}}_k, \mathbf{\Gamma}_k)) - \ln(\phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k)) = (\hat{\mathbf{c}}_k - \mathbf{c}_k)^\top \nabla_{\mathbf{c}_k} \ln(\phi_L(\mathbf{y}; z_{\mathbf{c}_k}, \mathbf{\Gamma}_k)),$$

$$\begin{aligned} &\ln(\phi_L(\mathbf{y}; \hat{\mathbf{c}}_k, \hat{\mathbf{\Gamma}}_k)) - \ln(\phi_L(\mathbf{y}; \hat{\mathbf{c}}_k, \mathbf{\Gamma}_k)) \\ &= (\text{vec}(\hat{\mathbf{\Gamma}}_k) - \text{vec}(\mathbf{\Gamma}_k))^\top \nabla_{\text{vec}(\mathbf{\Gamma}_k)} \ln(\phi_L(\mathbf{y}; \hat{\mathbf{c}}_k, z_{\mathbf{\Gamma}_k})). \end{aligned}$$

Moreover, $(\mathbf{y}, \mathbf{c}_k, \text{vec}(\mathbf{\Gamma}_k)) \mapsto \nabla_{\mathbf{c}_k} \ln(\phi_L(\mathbf{y}; z_{\mathbf{c}_k}, \mathbf{\Gamma}_k))$ and $(\mathbf{y}, \mathbf{c}_k, \text{vec}(\mathbf{\Gamma}_k)) \mapsto \nabla_{\text{vec}(\mathbf{\Gamma}_k)} \ln(\phi_L(\mathbf{y}; \hat{\mathbf{c}}_k, z_{\mathbf{\Gamma}_k}))$ are continuous functions on the compact set $\mathcal{U} := \mathcal{Y} \times \mathcal{Y} \times [a_{\mathbf{\Gamma}}, A_{\mathbf{\Gamma}}]^{L^2}$ leads to they attain minimum and maximum values (see [37, Theorem 1.8.8]). That is, we can set

$$\begin{aligned} \mathbf{0} < (C_{\mathbf{c}})_{1, \dots, L}^\top &:= \max_{k \in [K]} \sup_{(\mathbf{y}, \mathbf{c}_k, \text{vec}(\mathbf{\Gamma}_k)) \in \mathcal{U}} \|\nabla_{\mathbf{c}_k} \ln(\phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k))\| < \infty, \\ \mathbf{0} < (C_{\mathbf{\Gamma}})_{1, \dots, L^2}^\top &:= \max_{k \in [K]} \sup_{(\mathbf{y}, \mathbf{c}_k, \text{vec}(\mathbf{\Gamma}_k)) \in \mathcal{U}} \|\nabla_{\text{vec}(\mathbf{\Gamma}_k)} \ln(\phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k)(x))\| < \infty. \end{aligned}$$

Therefore, by the Cauchy–Schwarz inequality, we have

$$\begin{aligned} \sup_{\mathbf{y} \in \mathcal{Y}} |\ln(\phi_L(\mathbf{y}; \widehat{\mathbf{c}}_k, \mathbf{\Gamma}_k)) - \ln(\phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k))| &\leq |\widehat{\mathbf{c}}_k - \mathbf{c}_k|^\top (C_{\mathbf{c}})_{1, \dots, L}^\top \\ &= C_{\mathbf{c}} \|\widehat{\mathbf{c}}_k - \mathbf{c}_k\|_1 \leq C_{\mathbf{c}} \delta_{\boldsymbol{\pi}}, \end{aligned}$$

$$\sup_{\mathbf{y} \in \mathcal{Y}} \left| \ln(\phi_L(\mathbf{y}; \widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k)) - \ln(\phi_L(\mathbf{y}; \widehat{\mathbf{c}}_k, \mathbf{\Gamma}_k)) \right| \leq C_{\mathbf{\Gamma}} \left\| \text{vec}(\widehat{\mathbf{\Gamma}}_k - \mathbf{\Gamma}_k) \right\|_1 \leq C_{\mathbf{\Gamma}} \delta_{\mathbf{\Gamma}},$$

and by using the triangle inequality, it follows that

$$\max_{k \in [K]} \sup_{\mathbf{y} \in \mathcal{Y}} \left| \ln(\phi_L(\mathbf{y}; \widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k)) - \ln(\phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k)) \right| \leq C_{\mathbf{c}} \delta_{\mathbf{c}} + C_{\mathbf{\Gamma}} \delta_{\mathbf{\Gamma}}. \quad (62)$$

Moreover, for every $\boldsymbol{\pi} \in \mathbf{\Pi}_{K-1}$, Lemma A.5 implies that we can choose $\widehat{\boldsymbol{\pi}} \in \mathbf{\Pi}_{K-1, \omega}$ so that $\max_{k \in [K]} |\pi_k - \widehat{\pi}_k| \leq \delta_{\boldsymbol{\pi}}$. Notice that $[a_{\boldsymbol{\pi}}, \infty) \ni t \mapsto \ln(t)$, $a_{\boldsymbol{\pi}} > 0$ is a Lipschitz continuous function on $[a_{\boldsymbol{\pi}}, \infty]$. Indeed, by the mean value theorem, it holds that there exists $c \in (t_1, t_2)$, such that

$$|\ln(t_1) - \ln(t_2)| = \ln'(c) |t_1 - t_2| \leq \frac{1}{a_{\boldsymbol{\pi}}} |t_1 - t_2|, \text{ for all } t_1, t_2 \in [a_{\boldsymbol{\pi}}, \infty). \quad (63)$$

Therefore, (57) can be obtained by the following evaluation

$$\begin{aligned} & d_{\|\cdot\|_{\infty}}(\mathbf{w}(\cdot; \boldsymbol{\omega}), \mathbf{w}(\cdot; \widehat{\boldsymbol{\omega}})) \\ &= \max_{k \in [K]} \sup_{\mathbf{y} \in \mathcal{Y}} \left| \ln(\pi_k \phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k)) - \ln(\widehat{\pi}_k \phi_L(\mathbf{y}; \widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k)) \right| \\ &= \max_{k \in [K]} \sup_{\mathbf{y} \in \mathcal{Y}} \left| \ln(\pi_k) - \ln(\widehat{\pi}_k) + \ln(\phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k)) - \ln(\phi_L(\mathbf{y}; \widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k)) \right| \\ &\leq \max_{k \in [K]} |\ln(\pi_k) - \ln(\widehat{\pi}_k)| + \max_{k \in [K]} \sup_{\mathbf{y} \in \mathcal{Y}} \left| \ln(\phi_L(\mathbf{y}; \mathbf{c}_k, \mathbf{\Gamma}_k)) - \ln(\phi_L(\mathbf{y}; \widehat{\mathbf{c}}_k, \widehat{\mathbf{\Gamma}}_k)) \right| \\ &\leq \frac{1}{a_{\boldsymbol{\pi}}} \max_{k \in [K]} |\pi_k - \widehat{\pi}_k| + C_{\mathbf{c}} \delta_{\mathbf{c}} + C_{\mathbf{\Gamma}} \delta_{\mathbf{\Gamma}} \text{ (using (63) and (62))} \\ &\leq \frac{\delta_{\boldsymbol{\pi}}}{a_{\boldsymbol{\pi}}} + C_{\mathbf{c}} \delta_{\mathbf{c}} + C_{\mathbf{\Gamma}} \delta_{\mathbf{\Gamma}} \text{ (using Lemma A.5)} \leq \frac{\delta}{3} + \frac{\delta}{3} + \frac{\delta}{3} = \delta, \end{aligned}$$

where we choose $\delta_{\boldsymbol{\pi}} = \frac{\delta a_{\boldsymbol{\pi}}}{3}$, $\delta_{\mathbf{c}} = \frac{\delta}{3C_{\mathbf{c}}}$, $\delta_{\mathbf{\Gamma}} = \frac{\delta}{3C_{\mathbf{\Gamma}}}$. Finally, we get the covering number

$$\begin{aligned} & \mathcal{N}(\delta, \mathcal{W}_K, d_{\|\cdot\|_{\infty}}) \\ &\leq \text{card}(\mathbf{W}_{K, \omega}) = \left[\frac{2C_{\mathbf{y}}L}{\delta_{\mathbf{c}}} \right]^{KL} \left[\frac{2A_{\mathbf{\Gamma}}L^2}{\delta_{\mathbf{\Gamma}}} \right]^{\frac{L(L+1)}{2}K} \mathcal{N}(\delta_{\boldsymbol{\pi}}, \mathbf{\Pi}_{K-1}, \|\cdot\|_{\infty}) \\ &= \frac{C}{\delta^{KL+K\frac{L(L+1)}{2}+K-1}} = \frac{C}{\delta^{\dim(\mathcal{W}_K)}}, \end{aligned}$$

where

$$C = (6C_{\mathbf{c}}C_{\mathbf{y}}L)^{KL} (6C_{\mathbf{\Gamma}}A_{\mathbf{\Gamma}}L^2)^{\frac{L(L+1)}{2}K} \left(\frac{3}{a_{\boldsymbol{\pi}}} \right)^{K-1} K (2\pi e)^{K/2}.$$

A.3. Proof of Lemma 4.7: Bracketing entropy of Gaussian experts with 1-block covariance matrices

Note that the proof of Lemma 4.7 is adapted from [70] and is given here for the sake of completeness.

Indeed, Lemmas 1 and 2 from [70] imply that there exists a constant $C_{\mathbf{r}_{K,d}} = \ln\left(\sqrt{2} + \sqrt{D}dT_{\mathbf{r}}\right)$ (when $\mathbf{r}_{K,d} = \mathbf{r}_{b,d}^K$) or $C_{\mathbf{r}_{K,d}} = \ln\left(\sqrt{2} + \sqrt{D}\binom{d+L}{L}T_{\mathbf{r}}\right)$ (when $\mathbf{r}_{K,d} = \mathbf{r}_{p,d}^K$) such that, $\forall \delta \in (0, \sqrt{2})$,

$$\mathcal{H}_{d\|\sup\|_{\infty}}(\delta, \mathbf{r}_{K,d}) \leq \dim(\mathbf{r}_{K,d}) \left(C_{\mathbf{r}_{K,d}} + \ln \frac{1}{\delta} \right). \tag{64}$$

Next, we rely on Proposition A.7 for constructing of Gaussian brackets to for the Gaussian experts.

Proposition A.7 (Proposition 2 from [70]). *Let $\kappa \geq \frac{17}{29}$ and $\gamma_{\kappa} = \frac{25(\kappa - \frac{1}{2})}{49(1 + \frac{2\kappa}{5})}$.*

For any $0 < \delta \leq \sqrt{2}$, any $D \geq 1$, and any $\delta_{\Sigma} \leq \frac{1}{5\sqrt{\kappa^2 \cosh(\frac{2\kappa}{5}) + \frac{1}{2}}} \frac{\delta}{D}$, let

$(\mathbf{v}_d, B, \mathbf{A}, \mathbf{P}) \in \mathbf{r}_{K,d} \times [B_-, B_+] \times \mathcal{A}(\lambda_-, \lambda_+) \times SO(D)$ and $(\tilde{\mathbf{v}}_d, \tilde{B}, \tilde{\mathbf{A}}, \tilde{\mathbf{P}}) \in \mathbf{r}_{K,d} \times [B_-, B_+] \times \mathcal{A}(\lambda_-, +\infty) \times SO(D)$, and define $\Sigma = B\mathbf{P}\mathbf{A}\mathbf{P}^{\top}$, $\tilde{\Sigma} = \tilde{B}\tilde{\mathbf{P}}\tilde{\mathbf{A}}\tilde{\mathbf{P}}^{\top}$,

$$t^{-}(\mathbf{x}, \mathbf{y}) = (1 + \kappa\delta_{\Sigma})^{-D} \Phi_D\left(\mathbf{x}; \tilde{\mathbf{v}}_d(\mathbf{y}), (1 + \delta_{\Sigma})^{-1} \tilde{\Sigma}\right), \text{ and}$$

$$t^{+}(\mathbf{x}, \mathbf{y}) = (1 + \kappa\delta_{\Sigma})^D \Phi_D\left(\mathbf{x}; \tilde{\mathbf{v}}_d(\mathbf{y}), (1 + \delta_{\Sigma}) \tilde{\Sigma}\right).$$

If

$$\begin{cases} \forall \mathbf{y} \in \mathcal{Y}, \|\mathbf{v}_d(\mathbf{y}) - \tilde{\mathbf{v}}_d(\mathbf{y})\|^2 \leq D\gamma_{\kappa}\mathcal{L}_- \lambda_- \frac{\lambda_-}{\lambda_+} \delta_{\Sigma}^2 \\ (1 + \frac{2}{25}\delta_{\Sigma})^{-1} \tilde{B} \leq B \leq \tilde{B} \\ \forall i \in [D], \left| \mathbf{A}_{i,i}^{-1} - \tilde{\mathbf{A}}_{i,i}^{-1} \right| \leq \frac{1}{10} \frac{\delta_{\Sigma}}{\lambda_+} \\ \forall \mathbf{x} \in \mathbb{R}^D, \|\mathbf{P}\mathbf{x} - \tilde{\mathbf{P}}\mathbf{x}\| \leq \frac{1}{10} \frac{\lambda_-}{\lambda_+} \delta_{\Sigma} \|\mathbf{x}\| \end{cases}$$

then $[t^{-}, t^{+}]$ is a $\frac{\delta}{5}$ Hellinger bracket such that $t^{-}(\mathbf{x}, \mathbf{y}) \leq \Phi_D(\mathbf{x}; \mathbf{v}_d(\mathbf{y}), \Sigma) \leq t^{+}(\mathbf{x}, \mathbf{y})$.

Then, following the same argument as in [70, Appendix B.2.3], Proposition A.7 allows us to construct nets over the spaces of the means, the volumes, the eigenvector matrices, the normalized eigenvalue matrices and then control the bracketing entropy of $\mathcal{G}_{K,d}$. More precisely, three different contexts are considered for the mean, volume and matrix parameters. They can be all known ($\star = 0$), unknown but common to all classes ($\star = c$), unknown and possibly different for every class ($\star = K$). For example, $[\mathbf{v}_K, B_0, \mathbf{P}_0, \mathbf{A}_0]$ denotes a model

in which only mean vectors are assumed to be free. Then, we obtain

$$\mathcal{H}_{[\cdot], \sup_{\mathbf{y}} \max_k d_{\mathbf{x}}} \left(\frac{\delta}{5}, \mathcal{G}_{K,d} \right) \leq \dim(\mathcal{G}_{K,d}) \left(C_{\mathcal{G}_{K,d}} + \ln \left(\frac{1}{\delta} \right) \right), \quad (65)$$

where $\dim(\mathcal{G}_{K,d}) = Z_{\mathbf{v},*} + Z_{B,*} + \frac{D(D-1)}{2} Z_{\mathbf{P},*} + (D-1)Z_{\mathbf{A},*}$. Here, $Z_{\mathbf{v},K} = \dim(\Upsilon_{K,d})$, $Z_{\mathbf{v},c} = \dim(\Upsilon_1)$, $Z_{\mathbf{v},0} = 0$, $Z_{B,0} = Z_{\mathbf{P},0} = Z_{\mathbf{A},0} = 0$, $Z_{B,c} = Z_{\mathbf{P},c} = Z_{\mathbf{A},c} = 1$, $Z_{B,K} = Z_{\mathbf{P},K} = Z_{\mathbf{A},K} = K$, and given a universal constant c_U ,

$$\begin{aligned} C_{\mathcal{G}_{K,d}} &= \ln \left(5D \sqrt{\kappa^2 \cosh \left(\frac{2\kappa}{5} \right) + \frac{1}{2}} \right) + C_{\Upsilon_{K,d}} + \frac{1}{2} \ln \left(\frac{\lambda_+}{D\gamma_k B_- \lambda_-^2} \right) \\ &+ \ln \left(\frac{4 + 129 \ln \left(\frac{B_+}{B_-} \right)}{10} \right) + \frac{D(D-1)}{2} \ln(c_U) \\ &+ \ln \left(\frac{10\lambda_+}{\lambda_-} \right) + \ln \left(\frac{4}{5} + \frac{52\lambda_+}{5\lambda_-} \ln \left(\frac{\lambda_+}{\lambda_-} \right) \right). \end{aligned}$$

A.4. Proof of Lemma 4.11

We first aim to prove that $d^{2\otimes n}(s, t) \leq \sup_{\mathbf{y}} d_{\mathbf{x}}^2(s, t)$. Indeed, by definition, it follows that

$$\begin{aligned} d^{2\otimes n}(s, t) &= \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n d_{\mathbf{x}}^2(s(\cdot | \mathbf{Y}_i), t(\cdot | \mathbf{Y}_i)) \right] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{Y}_{[n]}} [d_{\mathbf{x}}^2(s(\cdot | \mathbf{Y}_i), t(\cdot | \mathbf{Y}_i))] \\ &= \frac{1}{n} \sum_{i=1}^n \int_{\mathcal{Y}} d_{\mathbf{x}}^2(s(\cdot | \mathbf{y}), t(\cdot | \mathbf{y})) s_{\mathbf{x},0}(\mathbf{y}) d\mathbf{y} \\ &\leq \sup_{\mathbf{y}} d_{\mathbf{x}}^2(s, t) \frac{1}{n} \sum_{i=1}^n \int_{\mathcal{Y}} s_{\mathbf{x},0}(\mathbf{y}) d\mathbf{y} = \sup_{\mathbf{y}} d_{\mathbf{x}}^2(s, t), \end{aligned}$$

where $s_{\mathbf{x},0}$ denotes that marginal PDF of s_0 , w.r.t. \mathbf{x} . Consequently, it holds that $d^{\otimes n}(s, t) = \sqrt{d^{2\otimes n}(s, t)} \leq \sqrt{\sup_{\mathbf{y}} d_{\mathbf{x}}^2(s, t)} = \sup_{\mathbf{y}} d_{\mathbf{x}}(s, t)$. To prove that

$$\mathcal{H}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \mathcal{H}_{[\cdot], \sup_{\mathbf{y}} d_{\mathbf{x}}}(\delta, S_{\mathbf{m}}),$$

it is sufficient to check that

$$\mathcal{N}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}) \leq \mathcal{N}_{[\cdot], \sup_{\mathbf{y}} d_{\mathbf{x}}}(\delta, S_{\mathbf{m}}).$$

By using the definition of bracketing entropy in (33) and $d^{\otimes n}(s, t) \leq \sup_{\mathbf{y}} d_{\mathbf{x}}(s, t)$, given

$$A = \left\{ n \in \mathbb{N}^* : \exists t_1^-, t_1^+, \dots, t_n^-, t_n^+ \text{ s.t. } \sup_{\mathbf{y}} d_{\mathbf{x}}(s, t) (t_k^-, t_k^+) \leq \delta, S_{\mathbf{m}} \subset \bigcup_{k=1}^n [t_k^-, t_k^+] \right\},$$

$$B = \left\{ n \in \mathbb{N}^* : \exists t_1^-, t_1^+, \dots, t_n^-, t_n^+ \text{ s.t. } d^{\otimes n}(t_k^-, t_k^+) \leq \delta, S_{\mathbf{m}} \subset \bigcup_{k=1}^n [t_k^-, t_k^+] \right\},$$

it leads to that $A \subset B$ and then (40) follows, since

$$\mathcal{N}_{[\cdot], \sup_{\mathbf{y}} d_{\mathbf{x}}(s, t)}(\delta, S_{\mathbf{m}}) = \min A \geq \min B = \mathcal{N}_{[\cdot], d^{\otimes n}}(\delta, S_{\mathbf{m}}).$$

With the similar argument as in the proof of (40), it holds that $d_{\mathcal{P}_K}(g^+, g^-) \leq \sup_{\mathbf{y}} d_k(g^+, g^-)$ and (41) is proved.

A.5. Proof of Lemma 4.12: Bracketing entropy of Gaussian experts with multi-block-diagonal covariance matrices

It is worth mentioning that without any structures on covariance matrices of Gaussian experts from the collection \mathcal{M} , Lemma 4.12 can be proved using [70, Proposition 2 and Appendix B.2.3], for constructing of Gaussian brackets to deal with the Gaussian experts. However, dealing with block-diagonal covariance matrices with random subcollection is not a trivial extension. We have to establish more constructive bracketing entropies in the spirits of [62, 29, 33].

Given any $k \in [K]$, by defining

$$\mathcal{G}_{d, \mathbf{B}_k} = \left\{ (\mathbf{x}, \mathbf{y}) \mapsto \phi(\mathbf{x}; \mathbf{v}_{k,d}(\mathbf{y}), \boldsymbol{\Sigma}_k(\mathbf{B}_k)) =: \phi_k : \mathbf{v}_{k,d} \in \mathbf{V}_{k,d}, \boldsymbol{\Sigma}_k(\mathbf{B}_k) \in \mathbf{V}_k(\mathbf{B}_k) \right\}, \quad (66)$$

it follows that $\mathcal{G}_{K,d,\mathbf{B}} = \prod_{k=1}^K \mathcal{G}_{d,\mathbf{B}_k}$, where \prod stands for the cartesian product. By using Lemma A.8, which is proved in Appendix A.5.1, it follows that

$$\mathcal{H}_{[\cdot], d_{\mathcal{G}_{K,d,\mathbf{B}}}}\left(\frac{\delta}{2}, \mathcal{G}_{K,d,\mathbf{B}}\right) \leq \sum_{k=1}^K \mathcal{H}_{[\cdot], d_{\mathcal{G}_{d,\mathbf{B}_k}}}\left(\frac{\delta}{2\sqrt{K}}, \mathcal{G}_{d,\mathbf{B}_k}\right). \quad (67)$$

Lemma A.8. Given $\mathcal{G}_{K,d,\mathbf{B}} = \prod_{k=1}^K \mathcal{G}_{d,\mathbf{B}_k}$, where $\mathcal{G}_{d,\mathbf{B}_k}$ is defined in (66), it holds that

$$\mathcal{N}_{[\cdot], d_{\mathcal{G}_{K,d,\mathbf{B}}}}\left(\frac{\delta}{2}, \mathcal{G}_{K,d,\mathbf{B}}\right) \leq \prod_{k=1}^K \mathcal{N}_{[\cdot], d_{\mathcal{G}_{d,\mathbf{B}_k}}}\left(\frac{\delta}{2\sqrt{K}}, \mathcal{G}_{d,\mathbf{B}_k}\right),$$

where for any $\phi^+, \phi^- \in \mathcal{G}_{K,d,\mathbf{B}}$ and any $\phi_k^+, \phi_k^- \in \mathcal{G}_{d,\mathbf{B}_k}, k \in [K]$,

$$d_{\mathcal{G}_{K,d,\mathbf{B}}}^2(\phi^+, \phi^-) = \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^K d^2(\phi_k^+(\cdot, \mathbf{Y}_i), \phi_k^-(\cdot, \mathbf{Y}_i)) \right],$$

$$d_{\mathcal{G}_{d,\mathbf{B}_k}}^2(\phi_k^+, \phi_k^-) = \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n d^2(\phi_k^+(\cdot, \mathbf{Y}_i), \phi_k^-(\cdot, \mathbf{Y}_i)) \right].$$

Lemma 4.12 is proved via (67) and Lemma A.9, which is proved in Appendix A.5.2.

Lemma A.9. *By defining $\mathcal{G}_{d, \mathbf{B}_k}$ as in (66), for all $\delta \in (0, \sqrt{2}]$, it holds that*

$$\begin{aligned} \mathcal{H}_{[\cdot], d_{\mathcal{G}_{d, \mathbf{B}_k}}} \left(\frac{\delta}{2}, \mathcal{G}_{d, \mathbf{B}_k} \right) &\leq \dim(\mathcal{G}_{d, \mathbf{B}_k}) \left(C_{\mathcal{G}_{d, \mathbf{B}_k}} + \ln \left(\frac{1}{\delta} \right) \right), \text{ where} \quad (68) \\ D_{\mathbf{B}_k} &= \sum_{g=1}^{G_k} \frac{\text{card} \left(d_k^{[g]} \right) \left(\text{card} \left(d_k^{[g]} \right) - 1 \right)}{2}, \\ C_{\mathcal{G}_{d, \mathbf{B}_k}} &= \frac{D_{\mathbf{B}_k} \ln \left(\frac{6\sqrt{6}\lambda_M D^2 (D-1)}{\lambda_m D_{\mathbf{B}_k}} \right) + \dim(\mathbf{Y}_{k,d}) \ln \left(\frac{6\sqrt{2D} \exp(C_{\mathbf{Y}_{k,d}})}{\sqrt{\lambda_m}} \right)}{\dim(\mathcal{G}_{d, \mathbf{B}_k})}. \end{aligned}$$

Indeed, (67) and (68) lead to

$$\begin{aligned} \mathcal{H}_{[\cdot], d_{\mathcal{G}_{K,d, \mathbf{B}}}} \left(\frac{\delta}{2}, \mathcal{G}_{K,d, \mathbf{B}} \right) &\leq \sum_{k=1}^K \mathcal{H}_{[\cdot], d_{\mathcal{G}_{d, \mathbf{B}_k}}} \left(\frac{\delta}{2\sqrt{K}}, \mathcal{G}_{d, \mathbf{B}_k} \right) \\ &\leq \sum_{k=1}^K \dim(\mathcal{G}_{d, \mathbf{B}_k}) \left(C_{\mathcal{G}_{d, \mathbf{B}_k}} + \ln(\sqrt{K}) + \ln \left(\frac{1}{\delta} \right) \right) \\ &\leq \dim(\mathcal{G}_{K,d, \mathbf{B}}) \left(C_{\mathcal{G}_{K,d, \mathbf{B}}} + \ln \left(\frac{1}{\delta} \right) \right). \end{aligned}$$

Here, $C_{\mathcal{G}_{K,d, \mathbf{B}}} = \sum_{k=1}^K C_{\mathcal{G}_{d, \mathbf{B}_k}} + \ln(\sqrt{K})$ and note that $\dim(\mathcal{G}_{K,d, \mathbf{B}}) = \sum_{k=1}^K \dim(\mathcal{G}_{d, \mathbf{B}_k})$, $\dim(\mathcal{G}_{d, \mathbf{B}_k}) = D_{\mathbf{B}_k} + \dim(\mathbf{Y}_{k,d})$, $\dim(\mathbf{Y}_{k,d}) = D d_{\mathbf{Y}_{k,d}}$, $C_{\mathbf{Y}_{k,d}} = \sqrt{D} d_{\mathbf{Y}_{k,d}} T_{\mathbf{Y}_{k,d}}$ (in cases where linear combination of bounded functions are used for means, *i.e.*, $\mathbf{Y}_{k,d} = \mathbf{Y}_b$) or $\dim(\mathbf{Y}_{k,d}) = D \binom{d_{\mathbf{Y}_{k,d}} + L}{L}$, $C_{\mathbf{Y}_{k,d}} = \sqrt{D} \binom{d_{\mathbf{Y}_{k,d}} + L}{L} T_{\mathbf{Y}_{k,d}}$ (in cases where we use polynomial means, *i.e.*, $\mathbf{Y}_{k,d} = \mathbf{Y}_p$).

A.5.1. Proof of Lemma A.8

By the definition of the bracketing entropy in (33), for each $k \in [K]$, let $\left\{ \left[\phi_k^{l,-}, \phi_k^{l,+} \right] \right\}_{1 \leq l \leq \mathcal{N}_{\mathcal{G}_{d, \mathbf{B}_k}}}$ be a minimal covering of δ_k brackets for $d_{\mathcal{G}_{d, \mathbf{B}_k}}$ of $\mathcal{G}_{d, \mathbf{B}_k}$, with cardinality $\mathcal{N}_{\mathcal{G}_{d, \mathbf{B}_k}}$. This leads to

$$\forall l \in \left[\mathcal{N}_{\mathcal{G}_{d, \mathbf{B}_k}} \right], d_{\mathcal{G}_{d, \mathbf{B}_k}} \left(\phi_k^{l,-}, \phi_k^{l,+} \right) \leq \delta_k.$$

Therefore, we claim that the set $\left\{ \prod_{k=1}^K \left[\phi_k^{l,-}, \phi_k^{l,+} \right] \right\}_{1 \leq l \leq \mathcal{N}_{\mathcal{G}_{d, \mathbf{B}_k}}}$ is a covering of $\frac{\delta}{2}$ -bracket for $d_{\mathcal{G}_{K,d, \mathbf{B}}}$ of $\mathcal{G}_{K,d, \mathbf{B}}$ with cardinality $\prod_{k=1}^K \mathcal{N}_{[\cdot], d_{\mathcal{G}_{d, \mathbf{B}_k}}}(\delta_k, \mathcal{G}_{d, \mathbf{B}_k})$. Indeed, let any $\phi = (\phi_k)_{k \in [K]} \in \mathcal{G}_{K,d, \mathbf{B}}$. Consequently, for each $k \in [K]$, $\phi_k \in \mathcal{G}_{d, \mathbf{B}_k}$, there exists $l(k) \in \left[\mathcal{N}_{\mathcal{G}_{d, \mathbf{B}_k}} \right]$, such that

$$\phi_k^{l(k),-} \leq \phi_k \leq \phi_k^{l(k),+}, d_{\mathcal{G}_{d, \mathbf{B}_k}}^2 \left(\phi_k^{l(k),+}, \phi_k^{l(k),-} \right) \leq (\delta_k)^2.$$

Then, it follows that $\phi \in [\phi^-, \phi^+] \in \{\prod_{k=1}^K [\phi_k^{l,-}, \phi_k^{l,+}]\}_{1 \leq l \leq \mathcal{N}_{\mathcal{G}_{d, \mathbf{B}_k}}}$, with $\phi^- = (\phi_k^{l(k),-})_{k \in [K]}$, $\phi^+ = (\phi_k^{l(k),+})_{k \in [K]}$, which implies that $\{\prod_{k=1}^K [\phi_k^{l,-}, \phi_k^{l,+}]\}_{1 \leq l \leq \mathcal{N}_{\mathcal{G}_{d, \mathbf{B}_k}}}$ is a bracket covering of $\mathcal{G}_{K,d,\mathbf{B}}$.

Now, we want to verify that the size of this bracket is $\delta/2$ by choosing $\delta_k = \frac{\delta}{2\sqrt{K}}, \forall k \in [K]$. It follows that

$$\begin{aligned} d_{\mathcal{G}_{K,d,\mathbf{B}}}^2(\phi^-, \phi^+) &= \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^K d^2 \left(\phi_k^{l(k),-}(\cdot, \mathbf{Y}_i), \phi_k^{l(k),+}(\cdot, \mathbf{Y}_i) \right) \right] \\ &= \sum_{k=1}^K \mathbb{E}_{\mathbf{Y}_{[n]}} \left[\frac{1}{n} \sum_{i=1}^n d^2 \left(\phi_k^{l(k),-}(\cdot, \mathbf{Y}_i), \phi_k^{l(k),+}(\cdot, \mathbf{Y}_i) \right) \right] \\ &= \sum_{k=1}^K d_{\mathcal{G}_{d,\mathbf{B}_k}}^2 \left(\phi_k^{l(k),-}, \phi_k^{l(k),+} \right) \leq K \left(\frac{\delta}{2\sqrt{K}} \right)^2 = \left(\frac{\delta}{2} \right)^2. \end{aligned}$$

To this end, by definition of a minimal $\frac{\delta}{2}$ -bracket covering number for $\mathcal{G}_{K,d,\mathbf{B}}$, Lemma A.8 is proved.

A.5.2. Proof of Lemma A.9

To provide the upper bound of the bracketing entropy in (68), our technique is adapted from the work of [41] for unidimensional Gaussian mixture families, which is recently generalized to multidimensional case by [62] for Gaussian mixture models. Furthermore, we make use of the results from [33] to deal with block-diagonal covariance matrices, $\mathbf{V}_k(\mathbf{B}_k)$, $k \in [K]$, and from [70] to handle the means of Gaussian experts $\boldsymbol{\Upsilon}_{k,d}$, $k \in [K]$.

The main idea in our approach is to define firstly a net over the parameter spaces of Gaussian experts, $\boldsymbol{\Upsilon}_{k,d} \times \mathbf{V}_k(\mathbf{B}_k)$, $k \in [K]$, and to construct a bracket covering of $\mathcal{G}_{d,\mathbf{B}_k}$ according to the tensorized Hellinger distance. Note that $\dim(\mathcal{G}_{d,\mathbf{B}_k}) = \dim(\boldsymbol{\Upsilon}_{k,d}) + \dim(\mathbf{V}_k(\mathbf{B}_k))$.

Step 1: Construction of a net for the block-diagonal covariance matrices. Firstly, for $k \in [K]$, we denote by $\text{Adj}(\boldsymbol{\Sigma}_k(\mathbf{B}_k))$ the adjacency matrix associated to the covariance matrix $\boldsymbol{\Sigma}_k(\mathbf{B}_k)$. Note that this matrix of size D^2 can be defined by a vector of concatenated upper triangular vectors. We are going to make use of the result from [33] to handle the block-diagonal covariance matrices $\boldsymbol{\Sigma}_k(\mathbf{B}_k)$, via its corresponding adjacency matrix. To do this, we need to construct a discrete space for $\{0,1\}^{D(D-1)/2}$, which is a one-to-one correspondence (bijection) with

$$\mathcal{A}_{\mathbf{B}_k} = \{\mathbf{A}_{\mathbf{B}_k} \in \mathcal{S}_D(\{0,1\}) : \exists \boldsymbol{\Sigma}_k(\mathbf{B}_k) \in \mathbf{V}_k(\mathbf{B}_k) \text{ s.t. } \text{Adj}(\boldsymbol{\Sigma}_k(\mathbf{B}_k)) = \mathbf{A}_{\mathbf{B}_k}\},$$

where $\mathcal{S}_D(\{0,1\})$ is the set of symmetric matrices of size D taking values on $\{0,1\}$.

Then, we want to deduce a discretization of the set of covariance matrices. Let h denotes Hamming distance on $\{0, 1\}^{D(D-1)/2}$ defined by

$$d(z, z') = \sum_{i=1}^n \mathbb{I}\{z \neq z'\}, \text{ for all } z, z' \in \{0, 1\}^{D(D-1)/2}.$$

Let $\{0, 1\}_{\mathbf{B}_k}^{D(D-1)/2}$ be the subset of $\{0, 1\}^{D(D-1)/2}$ of vectors for which the corresponding graph has structure $\mathbf{B}_k = \left(d_k^{[g]}\right)_{g \in [G_k]}$. Corollary 1 and Proposition 2 from Supplementary Material A of [33] imply that there exists some subset \mathcal{R} of $\{0, 1\}^{D(D-1)/2}$, as well as its equivalent $\mathcal{A}_{\mathbf{B}_k}^{\text{disc}}$ for adjacency matrices such that, given $\epsilon > 0$, and

$$\begin{aligned} \tilde{\mathcal{S}}_{\mathbf{B}_k}^{\text{disc}}(\epsilon) = \left\{ \Sigma_k(\mathbf{B}_k) \in \mathcal{S}_D^{++}(\mathbb{R}) : \text{Adj}(\Sigma_k(\mathbf{B}_k)) \in \mathcal{A}_{\mathbf{B}_k}^{\text{disc}}, \right. \\ \left. [\Sigma_k(\mathbf{B}_k)]_{i,j} = \sigma_{i,j}\epsilon, \sigma_{i,j} \in \left[\frac{-\lambda_M}{\epsilon}, \frac{\lambda_M}{\epsilon} \right] \cap \mathbb{Z} \right\}, \end{aligned}$$

it holds that for all $(\Sigma_k(\mathbf{B}_k), \tilde{\Sigma}_k(\mathbf{B}_k)) \in (\tilde{\mathcal{S}}_{\mathbf{B}_k}^{\text{disc}}(\epsilon))^2$ s.t. $\Sigma_k(\mathbf{B}_k) \neq \tilde{\Sigma}_k(\mathbf{B}_k)$,

$$\begin{aligned} \left\| \Sigma_k(\mathbf{B}_k) - \tilde{\Sigma}_k(\mathbf{B}_k) \right\|_2^2 &\leq \frac{D_{\mathbf{B}_k}}{2} \wedge \epsilon^2, \\ \text{card}(\tilde{\mathcal{S}}_{\mathbf{B}_k}^{\text{disc}}(\epsilon)) &\leq \left(\left\lfloor \frac{2\lambda_M}{\epsilon} \right\rfloor \frac{D(D-1)}{2D_{\mathbf{B}_k}} \right)^{D_{\mathbf{B}_k}}, \end{aligned} \quad (69)$$

$$D_{\mathbf{B}_k} = \dim(\mathbf{V}_k(\mathbf{B}_k)) = \sum_{g=1}^{G_k} \frac{\text{card}(d_k^{[g]}) (\text{card}(d_k^{[g]}) - 1)}{2}. \quad (70)$$

By choosing $\epsilon^2 \leq \frac{D_{\mathbf{B}_k}}{2}$, given $\Sigma_k(\mathbf{B}_k) \in \mathbf{V}_k(\mathbf{B}_k)$, then there exists $\tilde{\Sigma}_k(\mathbf{B}_k) \in \tilde{\mathcal{S}}_{\mathbf{B}_k}^{\text{disc}}(\epsilon)$, such that

$$\left\| \Sigma_k(\mathbf{B}_k) - \tilde{\Sigma}_k(\mathbf{B}_k) \right\|_2^2 \leq \epsilon^2. \quad (71)$$

Step 2: Construction of a net for the mean functions. Based on $\tilde{\Sigma}_k(\mathbf{B}_k)$, we can construct the following bracket covering of $\mathcal{G}_{d, \mathbf{B}_k}$ by defining the nets for the means of Gaussian experts. The proof of Lemma 1, page 1693, from [70] implies that

$$\mathcal{N}_{[\cdot, \sup_{\mathbf{y}} \|\cdot\|_2]}(\delta_{\mathbf{r}_{k,d}}, \mathbf{r}_{k,d}) \leq \left(\frac{\exp(C_{\mathbf{r}_{k,d}})}{\delta_{\mathbf{r}_{k,d}}} \right)^{\dim(\mathbf{r}_{k,d})}.$$

Here $\dim(\mathbf{r}_{k,d}) = Dd_{\mathbf{r}_{k,d}}$, and $C_{\mathbf{r}_{k,d}} = \sqrt{D}d_{\mathbf{r}_{k,d}}T_{\mathbf{r}_{k,d}}$ in the general case or $\dim(\mathbf{r}_{k,d}) = D\binom{d_{\mathbf{r}_{k,d}}+L}{L}$, and $C_{\mathbf{r}_{k,d}} = \sqrt{D}\binom{d_{\mathbf{r}_{k,d}}+L}{L}T_{\mathbf{r}_{k,d}}$ in the special case of

polynomial means. Then, by the definition of bracketing entropy in (33), for any minimal $\delta_{\mathbf{r}_{k,d}}$ -bracketing covering of the means from Gaussian experts, denoted by $G_{\mathbf{r}_{k,d}}(\delta_{\mathbf{r}_{k,d}})$, it is true that

$$\text{card}(G_{\mathbf{r}_{k,d}}(\delta_{\mathbf{r}_{k,d}})) \leq \left(\frac{\exp(C_{\mathbf{r}_{k,d}})}{\delta_{\mathbf{r}_{k,d}}} \right)^{\dim(\mathbf{r}_{k,d})}. \quad (72)$$

Therefore, given $\alpha > 0$, which is specified later, we claim that the set

$$\left\{ [l, u] \left| \begin{array}{l} l(\mathbf{x}, \mathbf{y}) = (1 + 2\alpha)^{-D} \phi(\mathbf{x}; \tilde{\mathbf{v}}_{k,d}(\mathbf{y}), (1 + \alpha)^{-1} \tilde{\Sigma}_k(\mathbf{B}_k)), \\ u(\mathbf{x}, \mathbf{y}) = (1 + 2\alpha)^D \phi(\mathbf{x}; \tilde{\mathbf{v}}_{k,d}(\mathbf{y}), (1 + \alpha) \tilde{\Sigma}_k(\mathbf{B}_k)), \\ \tilde{\mathbf{v}}_{k,d} \in G_{\mathbf{r}_{k,d}}(\delta_{\mathbf{r}_{k,d}}), \tilde{\Sigma}_k(\mathbf{B}_k) \in \tilde{S}_{\mathbf{B}_k}^{\text{disc}}(\epsilon) \end{array} \right. \right\},$$

is a $\delta_{\mathbf{r}_{k,d}}$ -brackets set over $\mathcal{G}_{d, \mathbf{B}_k}$. Indeed, let $\mathcal{X} \times \mathcal{Y} \ni (\mathbf{x}, \mathbf{y}) \mapsto f(\mathbf{x}, \mathbf{y}) = \phi(\mathbf{x}; \mathbf{v}_{k,d}(\mathbf{y}), \Sigma_k(\mathbf{B}_k))$ be a function of $\mathcal{G}_{d, \mathbf{B}_k}$, where $\mathbf{v}_{k,d} \in \mathbf{r}_{k,d}$ and $\Sigma_k(\mathbf{B}_k) \in \mathbf{V}_k(\mathbf{B}_k)$. According to (71), there exists $\tilde{\Sigma}_k(\mathbf{B}_k) \in \tilde{S}_{\mathbf{B}_k}^{\text{disc}}(\epsilon)$, such that

$$\left\| \Sigma_k(\mathbf{B}_k) - \tilde{\Sigma}_k(\mathbf{B}_k) \right\|_2^2 \leq \epsilon^2.$$

By definition of $G_{\mathbf{r}_{k,d}}(\delta_{\mathbf{r}_{k,d}})$, there exists $\tilde{\mathbf{v}}_{k,d} \in G_{\mathbf{r}_{k,d}}(\delta_{\mathbf{r}_{k,d}})$, such that

$$\sup_{\mathbf{y} \in \mathcal{Y}} \|\tilde{\mathbf{v}}_{k,d}(\mathbf{y}) - \mathbf{v}_{k,d}(\mathbf{y})\|_2^2 \leq \delta_{\mathbf{r}_{k,d}}^2. \quad (73)$$

Step 3: Upper bound of the number of the bracketing entropy. Next, we wish to make use of Lemma A.10 to evaluate the ratio of two Gaussian densities.

Lemma A.10 (Proposition C.1 from [62]). *Let $\phi(\cdot; \boldsymbol{\mu}_1, \Sigma_1)$ and $\phi(\cdot; \boldsymbol{\mu}_2, \Sigma_2)$ be two Gaussian densities. If $\Sigma_2 - \Sigma_1$ is a positive definite matrix then for all $\mathbf{x} \in \mathbb{R}^D$,*

$$\frac{\phi(\mathbf{x}; \boldsymbol{\mu}_1, \Sigma_1)}{\phi(\mathbf{x}; \boldsymbol{\mu}_2, \Sigma_2)} \leq \sqrt{\frac{|\Sigma_2|}{|\Sigma_1|}} \exp \left[\frac{1}{2} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^\top (\Sigma_2 - \Sigma_1)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \right].$$

The following Lemma A.11 allows us to fulfill the assumptions of Lemma A.10.

Lemma A.11. *Assume that $0 < \epsilon < \lambda_m^2/9$, and set $\alpha = 3\sqrt{\epsilon}/\lambda_m$. Then, for every $k \in [K]$, $(1 + \alpha) \tilde{\Sigma}_k(\mathbf{B}_k) - \Sigma_k(\mathbf{B}_k)$ and $\Sigma_k(\mathbf{B}_k) - (1 + \alpha)^{-1} \tilde{\Sigma}_k(\mathbf{B}_k)$ are both positive definite matrices. Moreover, for all $\mathbf{x} \in \mathbb{R}^D$,*

$$\begin{aligned} \mathbf{x}^\top \left[(1 + \alpha) \tilde{\Sigma}_k(\mathbf{B}_k) - \Sigma_k(\mathbf{B}_k) \right] \mathbf{x} &\geq \epsilon \|\mathbf{x}\|_2^2, \\ \mathbf{x}^\top \left[\Sigma_k(\mathbf{B}_k) - (1 + \alpha)^{-1} \tilde{\Sigma}_k(\mathbf{B}_k) \right] \mathbf{x} &\geq \epsilon \|\mathbf{x}\|_2^2. \end{aligned}$$

Proof of Lemma A.11. For all $\mathbf{x} \neq \mathbf{0}$, since

$$\sup_{\lambda \in \text{vp}(\boldsymbol{\Sigma}_k(\mathbf{B}_k) - \tilde{\boldsymbol{\Sigma}}_k(\mathbf{B}_k))} |\lambda| = \left\| \boldsymbol{\Sigma}_k(\mathbf{B}_k) - \tilde{\boldsymbol{\Sigma}}_k(\mathbf{B}_k) \right\|_2 \leq \epsilon,$$

where vp denotes the spectrum of matrix, $-\epsilon \geq -\lambda_m/3$, and $\alpha = 3\epsilon/\lambda_m$, it follow that

$$\begin{aligned} & \mathbf{x}^\top \left[(1 + \alpha) \tilde{\boldsymbol{\Sigma}}_k(\mathbf{B}_k) - \boldsymbol{\Sigma}_k(\mathbf{B}_k) \right] \mathbf{x} \\ &= (1 + \alpha) \mathbf{x}^\top \left[\tilde{\boldsymbol{\Sigma}}_k(\mathbf{B}_k) - \boldsymbol{\Sigma}_k(\mathbf{B}_k) \right] \mathbf{x} + \alpha \mathbf{x}^\top \boldsymbol{\Sigma}_k(\mathbf{B}_k) \mathbf{x} \\ &\geq -(1 + \alpha) \left\| \tilde{\boldsymbol{\Sigma}}_k(\mathbf{B}_k) - \boldsymbol{\Sigma}_k(\mathbf{B}_k) \right\|_2 \|\mathbf{x}\|_2^2 + \alpha \lambda_m \|\mathbf{x}\|_2^2 \\ &\geq (\alpha \lambda_m - (1 + \alpha) \epsilon) \|\mathbf{x}\|_2^2 = (\alpha \lambda_m - \alpha \epsilon - \epsilon) \|\mathbf{x}\|_2^2 \\ &\geq \left(\frac{2}{3} \alpha \lambda_m - \epsilon \right) \|\mathbf{x}\|_2^2 = \epsilon \|\mathbf{x}\|_2^2 > 0, \end{aligned}$$

and

$$\begin{aligned} & \mathbf{x}^\top \left[\boldsymbol{\Sigma}_k(\mathbf{B}_k) - (1 + \alpha)^{-1} \tilde{\boldsymbol{\Sigma}}_k(\mathbf{B}_k) \right] \mathbf{x} \\ &= (1 + \alpha)^{-1} \mathbf{x}^\top \left[\boldsymbol{\Sigma}_k(\mathbf{B}_k) - \tilde{\boldsymbol{\Sigma}}_k(\mathbf{B}_k) \right] \mathbf{x} + \left(1 - (1 + \alpha)^{-1} \right) \mathbf{x}^\top \boldsymbol{\Sigma}_k(\mathbf{B}_k) \mathbf{x} \\ &\geq \left(\frac{\alpha \lambda_m - \epsilon}{1 + \alpha} \right) \|\mathbf{x}\|_2^2 = \frac{2\epsilon}{1 + \alpha} \|\mathbf{x}\|_2^2 \geq \epsilon \|\mathbf{x}\|_2^2 > 0 \quad (\text{since } 0 < \alpha < 1). \quad \square \end{aligned}$$

By Lemma A.10 and the same argument as in the proof of Lemma B.9 from [62], given $0 < \epsilon < \lambda_m/3$, where ϵ is chosen later, and $\alpha = 3\epsilon/\lambda_m$, we obtain

$$\max \left\{ \frac{l(\mathbf{x}, \mathbf{y})}{f(\mathbf{x}, \mathbf{y})}, \frac{f(\mathbf{x}, \mathbf{y})}{u(\mathbf{x}, \mathbf{y})} \right\} \leq (1 + 2\alpha)^{-\frac{D}{2}} \exp \left(\frac{\|\mathbf{v}_{k,d}(\mathbf{y}) - \tilde{\mathbf{v}}_{k,d}(\mathbf{y})\|_2^2}{2\epsilon} \right). \quad (74)$$

Because $\ln(\cdot)$ is a non-decreasing function, $\ln(1 + 2\alpha) \geq \alpha, \forall \alpha \in [0, 1]$. Combined with (73) where $\delta_{\mathbf{r}_{k,d}}^2 = D\alpha\epsilon$, we conclude that

$$\begin{aligned} & \max \left\{ \ln \left(\frac{l(\mathbf{x}, \mathbf{y})}{f(\mathbf{x}, \mathbf{y})} \right), \ln \left(\frac{f(\mathbf{x}, \mathbf{y})}{u(\mathbf{x}, \mathbf{y})} \right) \right\} \\ &\leq -\frac{D}{2} \ln(1 + 2\alpha) + \frac{\delta_{\mathbf{r}_{k,d}}^2}{2\epsilon} \leq -\frac{D}{2} \alpha + \frac{\delta_{\mathbf{r}_{k,d}}^2}{2\epsilon} = 0. \end{aligned}$$

This means that $l(\mathbf{x}, \mathbf{y}) \leq f(\mathbf{x}, \mathbf{y}) \leq u(\mathbf{x}, \mathbf{y}), \forall (\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}$. Hence, it remains to bound the size of bracket $[l, u]$ w.r.t. $d_{\mathcal{G}_d, \mathbf{B}_k}$. To this end, we aim to verify that $d_{\mathcal{G}_d, \mathbf{B}_k}^2(l, u) \leq \frac{\delta}{2}$. To do that, we make use of the following Lemma A.12.

Lemma A.12 (Proposition C.3 from [62]). *Let $\phi(\cdot; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ and $\phi(\cdot; \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$ be two Gaussian densities with full rank covariance. It holds that*

$$\begin{aligned} & d^2(\phi(\cdot; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1), \phi(\cdot; \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)) \\ &= 2 \left\{ 1 - 2^{D/2} |\boldsymbol{\Sigma}_1 \boldsymbol{\Sigma}_2|^{-1/4} |\boldsymbol{\Sigma}_1^{-1} + \boldsymbol{\Sigma}_2^{-1}|^{-1/2} \exp^{-\frac{1}{4}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^\top (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)} \right\}. \end{aligned}$$

Therefore, using the fact that $\cosh(t) = \frac{e^{-t} + e^t}{2}$, Lemma A.12 leads to, for all $\mathbf{y} \in \mathcal{Y}$:

$$\begin{aligned} & d^2(l(\cdot, \mathbf{y}), u(\cdot, \mathbf{y})) \\ &= \int_{\mathcal{X}} \left[l(\mathbf{x}, \mathbf{y}) + u(\mathbf{x}, \mathbf{y}) - 2\sqrt{l(\mathbf{x}, \mathbf{y})u(\mathbf{x}, \mathbf{y})} \right] d\mathbf{x} \\ &= (1 + 2\alpha)^{-D} + (1 + 2\alpha)^D - 2 \\ &\quad + d^2 \left(\phi \left(\cdot; \tilde{\mathbf{v}}_{k,d}(\mathbf{y}), (1 + \alpha)^{-1} \tilde{\Sigma}_k(\mathbf{B}_k) \right), \phi \left(\cdot; \tilde{\mathbf{v}}_{k,d}(\mathbf{y}), (1 + \alpha) \tilde{\Sigma}_k(\mathbf{B}_k) \right) \right) \\ &= 2 \cosh [D \ln (1 + 2\alpha)] - 2 \\ &\quad + 2 \left[1 - 2^{D/2} \left[(1 + \alpha)^{-1} + (1 + \alpha) \right]^{-D/2} \left| \tilde{\Sigma}_k(\mathbf{B}_k) \right|^{-1/2} \left| \tilde{\Sigma}_k(\mathbf{B}_k) \right|^{1/2} \right] \\ &= 2 \cosh [D \ln (1 + 2\alpha)] - 2 + 2 - 2 [\cosh (\ln (1 + \alpha))]^{-D/2} \\ &= 2g (D \ln (1 + 2\alpha)) + 2h (\ln (1 + \alpha)), \end{aligned}$$

where $g(t) = \cosh(t) - 1 = \frac{e^{-t} + e^t}{2} - 1$, and $h(t) = 1 - \cosh(t)^{-D/2}$. The upper bounds of terms g and h separately imply that, for all $\mathbf{y} \in \mathcal{Y}$,

$$d^2(l(\cdot, \mathbf{y}), u(\cdot, \mathbf{y})) \leq 2 \left(2 \cosh \left(\frac{1}{\sqrt{6}} \right) \alpha^2 D^2 + \frac{1}{4} \alpha^2 D^2 \right) \leq 6\alpha^2 D^2 = \frac{\delta^2}{4},$$

where we choose $\alpha = \frac{3\epsilon}{\lambda_m}$, $\epsilon = \frac{\delta \lambda_m}{6\sqrt{6}D}$, $\forall \delta \in (0, 1]$, $D \in \mathbb{N}^*$, $\lambda_m > 0$, which appears in (74) and satisfies $\alpha = \frac{\delta}{2\sqrt{6}D}$ and $0 < \epsilon < \frac{\lambda_m}{3}$. Indeed, studying functions g and h yields

$$\begin{aligned} g'(t) &= \sinh(t), g''(t) = \cosh(t) \leq \cosh(c), \forall t \in [0, c], c \in \mathbb{R}_+, \\ h'(t) &= \frac{D}{2} \cosh(t)^{-D/2-1} \sinh(t), \\ h''(t) &= \frac{D}{2} \left(-\frac{D}{2} - 1 \right) \cosh(t)^{-D/2-2} \sinh^2(t) + \frac{D}{2} \cosh(t)^{-D/2} \\ &= \frac{D}{2} \left(1 - \left(\frac{D}{2} + 1 \right) \left(\frac{\sinh(t)}{\cosh(t)} \right)^2 \right) \cosh(t)^{-D/2} \leq \frac{D}{2}, \end{aligned}$$

where we used the fact that $\cosh(t) \geq 1$. Then, since $g(0) = 0$, $g'(0) = 0$, $h(0) = 0$, $h'(0) = 0$, by applying Taylor's Theorem, it is true that

$$\begin{aligned} g(t) &= g(t) - g(0) - g'(0)t = R_{0,1}(t) \leq \cosh(c) \frac{t^2}{2}, \forall t \in [0, c], \\ h(t) &= h(t) - h(0) - h'(0)t = R_{0,1}(t) \leq \frac{D}{2} \frac{t^2}{2} \leq \frac{D^2}{2} \frac{t^2}{2}, \forall t \geq 0. \end{aligned}$$

We wish to find an upper bound for $t = D \ln(1 + 2\alpha)$, $D \in \mathbb{N}^*$, $\alpha = \frac{\delta}{2\sqrt{6}D}$, $\delta \in (0, 1]$. Since \ln is an increasing function, then we have

$$t = D \ln \left(1 + \frac{\delta}{\sqrt{6}D} \right) \leq D \ln \left(1 + \frac{1}{\sqrt{6}D} \right) \leq D \frac{1}{\sqrt{6}D} = \frac{1}{\sqrt{6}}, \forall \delta \in (0, 1],$$

since $\ln\left(1 + \frac{1}{\sqrt{6D}}\right) \leq \frac{1}{\sqrt{6D}}, \forall D \in \mathbb{N}^*$. Then, since $\ln(1 + 2\alpha) \leq 2\alpha, \forall \alpha \geq 0$,

$$g(D \ln(1 + 2\alpha)) \leq \cosh\left(\frac{1}{\sqrt{6}}\right) \frac{(D \ln(1 + 2\alpha))^2}{2} \leq \cosh\left(\frac{1}{\sqrt{6}}\right) \frac{D^2}{2} 4\alpha^2,$$

$$h(\ln(1 + \alpha)) \leq \frac{D^2}{2} \frac{(\ln(1 + \alpha))^2}{2} \leq \frac{D^2 \alpha^2}{4}.$$

Note that the set of $\delta/2$ -brackets $[l, u]$ over $\mathcal{G}_{d, \mathbf{B}_k}$ is totally defined by the parameter spaces $\tilde{S}_{\mathbf{B}_k}^{\text{disc}}(\epsilon)$ and $G_{\mathbf{r}_{k,d}}(\delta_{\mathbf{r}_{k,d}})$. This leads to an upper bound of the $\delta/2$ -bracketing entropy of $\mathcal{G}_{d, \mathbf{B}_k}$ evaluated from an upper bound of the two set cardinalities. Hence, given any $\delta > 0$, by choosing $\epsilon = \frac{\delta \lambda_m}{6\sqrt{6D}}, \alpha = \frac{3\epsilon}{\lambda_m} = \frac{\delta}{2\sqrt{6D}}$, and $\delta_{\mathbf{r}_{k,d}}^2 = D\alpha\epsilon = D \frac{\delta}{2\sqrt{6D}} \frac{\delta \lambda_m}{6\sqrt{6D}} = \frac{\delta^2 \lambda_m}{72D}$, it holds that

$$\begin{aligned} & \mathcal{N}_{[\cdot], d_{\mathcal{G}_{d, \mathbf{B}_k}}}\left(\frac{\delta}{2}, \mathcal{G}_{d, \mathbf{B}_k}\right) \\ & \leq \text{card}\left(\tilde{S}_{\mathbf{B}_k}^{\text{disc}}(\epsilon)\right) \times \text{card}\left(G_{\mathbf{r}_{k,d}}(\delta_{\mathbf{r}_{k,d}})\right) \\ & \leq \left(\left\lfloor \frac{2\lambda_M}{\epsilon} \right\rfloor \frac{D(D-1)}{2D_{\mathbf{B}_k}}\right)^{D_{\mathbf{B}_k}} \left(\frac{\exp(C_{\mathbf{r}_{k,d}})}{\delta_{\mathbf{r}_{k,d}}}\right)^{\dim(\mathbf{r}_{k,d})} \quad (\text{using (70) and (72)}) \\ & \leq \left(\frac{2\lambda_M 6\sqrt{6D} D(D-1)}{\delta \lambda_m 2D_{\mathbf{B}_k}}\right)^{D_{\mathbf{B}_k}} \left(\frac{6\sqrt{2D} \exp(C_{\mathbf{r}_{k,d}})}{\delta \sqrt{\lambda_m}}\right)^{\dim(\mathbf{r}_{k,d})} \\ & = \left(\frac{6\sqrt{6}\lambda_M D^2(D-1)}{\lambda_m D_{\mathbf{B}_k}}\right)^{D_{\mathbf{B}_k}} \left(\frac{6\sqrt{2D} \exp(C_{\mathbf{r}_{k,d}})}{\sqrt{\lambda_m}}\right)^{\dim(\mathbf{r}_{k,d})} \left(\frac{1}{\delta}\right)^{D_{\mathbf{B}_k} + \dim(\mathbf{r}_{k,d})}. \end{aligned}$$

Finally, by definition of bracketing entropy in (33), we obtain

$$\begin{aligned} & \mathcal{H}_{[\cdot], d_{\mathcal{G}_{d, \mathbf{B}_k}}}\left(\frac{\delta}{2}, \mathcal{G}_{d, \mathbf{B}_k}\right) \\ & \leq D_{\mathbf{B}_k} \ln\left(\frac{6\sqrt{6}\lambda_M D^2(D-1)}{\lambda_m D_{\mathbf{B}_k}}\right) + \dim(\mathbf{r}_{k,d}) \ln\left(\frac{6\sqrt{2D} \exp(C_{\mathbf{r}_{k,d}})}{\sqrt{\lambda_m}}\right) \\ & \quad + (D_{\mathbf{B}_k} + \dim(\mathbf{r}_{k,d})) \ln\left(\frac{1}{\delta}\right) = \dim(\mathcal{G}_{d, \mathbf{B}_k}) \left(C_{\mathcal{G}_{d, \mathbf{B}_k}} + \ln\left(\frac{1}{\delta}\right)\right), \end{aligned}$$

where $\dim(\mathcal{G}_{d, \mathbf{B}_k}) = D_{\mathbf{B}_k} + \dim(\mathbf{r}_{k,d})$ and

$$C_{\mathcal{G}_{d, \mathbf{B}_k}} = \frac{D_{\mathbf{B}_k} \ln\left(\frac{6\sqrt{6}\lambda_M D^2(D-1)}{\lambda_m D_{\mathbf{B}_k}}\right) + \dim(\mathbf{r}_{k,d}) \ln\left(\frac{6\sqrt{2D} \exp(C_{\mathbf{r}_{k,d}})}{\sqrt{\lambda_m}}\right)}{\dim(\mathcal{G}_{d, \mathbf{B}_k})}.$$

Acknowledgments

TrungTin Nguyen is supported by a ‘‘Contrat doctoral’’ from the French Ministry of Higher Education and Research. Faïcel Chamroukhi is granted by the

French National Research Agency (ANR) grant [SMILES ANR-18-CE40-0014](#). Hien Duy Nguyen is funded by Australian Research Council grant number DP180101192. This research is funded directly by the Inria [LANDER](#) project. TrungTin Nguyen also sincerely acknowledges Inria Grenoble-Rhône-Alpes Research Centre for a valuable Visiting PhD Fellowship working with [STATIFY](#) team so that this research can be completed, Erwan LE PENNEC and Lucie Montuelle for providing the simulations for the SGame models. Finally, we thank the Editor-in-Chief, Associate Editor, and Reviewers for their valuable comments, which enabled us to produce a much improved manuscript.

References

- [1] AKAIKE, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control* **19** 716–723. [MR0423716](#)
- [2] ANDERSON, D. and BURNHAM, K. (2002). *Model Selection and Multi-model Inference*, 2 ed. *A Practical Information-Theoretic Approach*. Springer-Verlag, New York. [MR1919620](#)
- [3] ARLOT, S. (2019). Minimal penalties and the slope heuristics: a survey. *Journal de la Société Française de Statistique* **160** 1–106. [MR4021408](#)
- [4] ARLOT, S. and BACH, F. (2009). Data-driven calibration of linear estimators with minimal penalties. *Advances in Neural Information Processing Systems* **22**.
- [5] ARLOT, S. and MASSART, P. (2009). Data-driven calibration of penalties for least-squares regression. *Journal of Machine Learning Research* **10**. ISBN: 1532-4435.
- [6] BARRON, A. R., HUANG, C., LI, J. and LUO, X. (2008). The MDL principle, penalized likelihoods, and statistical risk. *Festschrift in Honor of Jorma Rissanen on the Occasion of his 75th Birthday* 33–63.
- [7] BAUDRY, J.-P. (2009). Sélection de modèle pour la classification non supervisée. Choix du nombre de classes., PhD thesis, Université Paris-Sud XI.
- [8] BAUDRY, J.-P., MAUGIS, C. and MICHEL, B. (2012). Slope heuristics: overview and implementation. *Statistics and Computing* **22** 455–470. [MR2865029](#)
- [9] BERTIN, K., LE PENNEC, E. and RIVOIRARD, V. (2011). Adaptive Dantzig density estimation. *Annales de l’I.H.P. Probabilités et statistiques* **47** 43–74. Publisher: Gauthier-Villars. [MR2779396](#)
- [10] BIERNACKI, C., CELEUX, G. and GOVAERT, G. (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22** 719–725.
- [11] BIRGÉ, L. and MASSART, P. (2007). Minimal penalties for Gaussian model selection. *Probability Theory and Related Fields* **138** 33–73. [MR2288064](#)
- [12] BIRGÉ, L., MASSART, P. et al. (1998). Minimum contrast estimators on sieves: exponential bounds and rates of convergence. *Bernoulli* **4** 329–375. [MR1653272](#)

- [13] BISHOP, C. M. (2006). *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg. [MR2247587](#)
- [14] BORWEIN, J. M. and ZHU, Q. J. (2004). *Techniques of Variational Analysis*. Springer. [MR2144010](#)
- [15] BOUCHARD, G. (2003). Localised Mixtures of Experts for Mixture of Regressions. In *Between Data Science and Applied Data Analysis* (M. SCHADER, W. GAUL and M. VICHI, eds.) 155–164. Springer Berlin Heidelberg, Berlin, Heidelberg.
- [16] BRINKMAN, N. D. (1981). Ethanol fuel—single-cylinder engine study of efficiency and exhaust emissions. *SAE Transactions* 1410–1424.
- [17] BUNEA, F., TSYBAKOV, A. B., WEGKAMP, M. H. and BARBU, A. (2010). SPADES and mixture models. *The Annals of Statistics* **38** 2525–2558. Publisher: Institute of Mathematical Statistics. [MR2676897](#)
- [18] BUTUCEA, C., DELMAS, J.-F., DUTFOY, A. and FISCHER, R. (2017). Optimal exponential bounds for aggregation of estimators for the Kullback-Leibler loss. *Electronic Journal of Statistics* **11** 2258–2294. Publisher: Institute of Mathematical Statistics and Bernoulli Society. [MR3654825](#)
- [19] CAMBANIS, S., HUANG, S. and SIMONS, G. (1981). On the theory of elliptically contoured distributions. *Journal of Multivariate Analysis* **11** 368–385. [MR0629795](#)
- [20] CELEUX, G. and GOVAERT, G. (1995). Gaussian parsimonious clustering models. *Pattern Recognition* **28** 781–793.
- [21] CHAMROUKHI, F., SAMÉ, A., GOVAERT, G. and AKNIN, P. (2010). A hidden process regression model for functional data description. Application to curve discrimination. *Neurocomputing* **73** 1210–1221.
- [22] COHEN, S. and LE PENNEC, E. (2011). Conditional density estimation by penalized likelihood model selection and applications. *Technical report, INRIA*.
- [23] COHEN, S. X. and LE PENNEC, E. (2013). Partition-based conditional density estimation. *ESAIM: Probability and Statistics* **17** 672–697. [MR3126157](#)
- [24] DALALYAN, A. S. and SEBBAR, M. (2018). Optimal Kullback–Leibler aggregation in mixture density estimation by maximum likelihood. *Mathematical Statistics and Learning* **1** 1–35. ISBN: 2520-2316. [MR4049449](#)
- [25] DELEFORGE, A., FORBES, F., BA, S. and HORAUD, R. (2015). Hyper-spectral image analysis with partially latent regression and spatial Markov dependencies. *IEEE Journal of Selected Topics in Signal Processing* **9** 1037–1048.
- [26] DELEFORGE, A., FORBES, F. and HORAUD, R. (2015). High-dimensional regression with Gaussian mixtures and partially-latent response variables. *Statistics and Computing* **25** 893–911. [MR3375624](#)
- [27] DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* **39** 1–38. [MR0501537](#)
- [28] DEVIJVER, E. (2015). An l_1 -oracle inequality for the Lasso in multivariate

- finite mixture of multivariate Gaussian regression models. *ESAIM: PS* **19** 649–670. [MR3433431](#)
- [29] DEVIJVER, E. (2015). Finite mixture regression: a sparse variable selection by model selection for clustering. *Electronic Journal of Statistics* **9** 2642–2674. [MR3432429](#)
- [30] DEVIJVER, E. (2015). An l_1 -oracle inequality for the Lasso in finite mixture of multivariate Gaussian regression. *ESAIM: Probability and Statistics* **19** 649–670. [MR3433431](#)
- [31] DEVIJVER, E. (2017). Model-based regression clustering for high-dimensional data: application to functional data. *Advances in Data Analysis and Classification* **11** 243–279. [MR3656031](#)
- [32] DEVIJVER, E. (2017). Joint rank and variable selection for parsimonious estimation in a high-dimensional finite mixture regression model. *Journal of Multivariate Analysis* **157** 1–13. [MR3641732](#)
- [33] DEVIJVER, E. and GALLOPIN, M. (2018). Block-diagonal covariance selection for high-dimensional Gaussian graphical models. *Journal of the American Statistical Association* **113** 306–314. [MR3803466](#)
- [34] DEVIJVER, E., GALLOPIN, M. and PERTHAME, E. (2017). Nonlinear network-based quantitative trait prediction from transcriptomic data. *arXiv preprint arXiv:1701.07899*.
- [35] DEVIJVER, E. and PERTHAME, E. (2020). Prediction regions through Inverse Regression. *Journal of Machine Learning Research* **21** 1–24. [MR4119181](#)
- [36] DING, P. (2016). On the conditional distribution of the multivariate t distribution. *The American Statistician* **70** 293–295. [MR3535516](#)
- [37] DUISTERMAAT, J. J. and KOLK, J. A. (2004). *Multidimensional real analysis I: differentiation* **86**. Cambridge University Press. [MR2121976](#)
- [38] FANG, K. T., KOTZ, S. and NG, K. W. (1990). *Symmetric Multivariate And Related Distributions*. Chapman and Hall. [MR1071174](#)
- [39] FORBES, F., NGUYEN, H. D., NGUYEN, T. T. and ARBEL, J. (2021). Approximate Bayesian computation with surrogate posteriors. *Preprint hal-03139256*.
- [40] FRAHM, G. (2004). *Generalized Elliptical Distributions: Theory and Applications*. Universität zu Köln.
- [41] GENOVESE, C. R. and WASSERMAN, L. (2000). Rates of convergence for the Gaussian mixture sieve. *The Annals of Statistics* **28** 1105–1127. [MR1810921](#)
- [42] GOLUB, T. R., SLONIM, D. K., TAMAYO, P., HUARD, C., GAASENBEEK, M., MESIROV, J. P., COLLIER, H., LOH, M. L., DOWNING, J. R., CALIGIURI, M. A., BLOOMFIELD, C. D. and LANDER, E. S. (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science (New York, N.Y.)* **286** 531–537.
- [43] HO, N., NGUYEN, X. et al. (2016). Convergence rates of parameter estimation for some weakly identifiable finite mixtures. *The Annals of Statistics* **44** 2726–2755. [MR3576559](#)
- [44] HO, N., NGUYEN, X. et al. (2016). On strong identifiability and conver-

- gence rates of parameter estimation in finite mixtures. *Electronic Journal of Statistics* **10** 271–307. [MR3466183](#)
- [45] HO, N., YANG, C.-Y. and JORDAN, M. I. (2019). Convergence Rates for Gaussian Mixtures of Experts. *arXiv preprint arXiv:1907.04377*.
- [46] HULT, H. and LINDSKOG, F. (2002). Multivariate extremes, aggregation and dependence in elliptical distributions. *Advances in Applied Probability* **34** 587–608. [MR1929599](#)
- [47] INGRASSIA, S., MINOTTI, S. C. and VITTADINI, G. (2012). Local statistical modeling via a cluster-weighted approach with elliptical distributions. *Journal of Classification* **29** 363–401. [MR2972778](#)
- [48] JACOBS, R. A., JORDAN, M. I., NOWLAN, S. J. and HINTON, G. E. (1991). Adaptive mixtures of local experts. *Neural Computation* **3** 79–87.
- [49] JIANG, W. and TANNER, M. A. (1999). Hierarchical mixtures-of-experts for exponential family regression models: approximation and maximum likelihood estimation. *Annals of Statistics* 987–1011. [MR1724038](#)
- [50] JORDAN, M. I. and JACOBS, R. A. (1994). Hierarchical mixtures of experts and the EM algorithm. *Neural Computation* **6** 181–214.
- [51] KELKER, D. (1970). Distribution theory of spherical distributions and a location-scale parameter generalization. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)* **32** 419–430. [MR0287628](#)
- [52] KINGMA, D. P. and WELING, M. (2013). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*. *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*.
- [53] KOSOROK, M. R. (2007). *Introduction to empirical processes and semi-parametric inference*. Springer Science & Business Media. [MR2724368](#)
- [54] KOTZ, S. and NADARAJAH, S. (2004). *Multivariate T-Distributions and Their Applications*. Cambridge University Press, Cambridge. [MR2038227](#)
- [55] LATHUILIÈRE, S., JUGE, R., MESEJO, P., MUÑOZ-SALINAS, R. and HORAUD, R. (2017). Deep mixture of linear inverse regressions applied to head-pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 4817–4825.
- [56] LÊ CAO, K.-A., ROSSOUW, D., ROBERT-GRANIÉ, C. and BESSE, P. (2008). A sparse PLS for variable selection when integrating omics data. *Statistical Applications in Genetics and Molecular Biology* **7** Article 35. [MR2457048](#)
- [57] LI, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association* **86** 316–327. [MR1137117](#)
- [58] MALLOWS, C. L. (1973). Some comments on CP. *Technometrics* **15** 661–675.
- [59] MASOUDNIA, S. and EBRAHIMPOUR, R. (2014). Mixture of experts: a literature survey. *Artificial Intelligence Review* **42** 275–293.
- [60] MASSART, P. (2007). *Concentration Inequalities and Model Selection: Ecole d’Eté de Probabilités de Saint-Flour XXXIII-2003*. Springer. [MR2319879](#)
- [61] MASSART, P. and MEYNET, C. (2011). The Lasso as an l_1 -ball model selection procedure. *Electronic Journal of Statistics* **5** 669–687. [MR2820635](#)

- [62] MAUGIS, C. and MICHEL, B. (2011). A non asymptotic penalized criterion for Gaussian mixture model selection. *ESAIM: Probability and Statistics* **15** 41–68. [MR2870505](#)
- [63] MAUGIS, C. and MICHEL, B. (2011). Data-driven penalty calibration: A case study for Gaussian mixture model selection. *ESAIM: PS* **15** 320–339. [MR2870518](#)
- [64] MAUGIS-RABUSSEAU, C. and MICHEL, B. (2013). Adaptive density estimation for clustering with Gaussian mixtures. *ESAIM: Probability and Statistics* **17** 698–724. [MR3126158](#)
- [65] MCLACHLAN, G. J. and KRISHNAN, T. (1997). *The EM Algorithm and Extensions*. Wiley. [MR1417721](#)
- [66] MCLACHLAN, G. J. and PEEL, D. (2000). *Finite Mixture Models*. John Wiley & Sons. [MR1789474](#)
- [67] MENDES, E. F. and JIANG, W. (2012). On convergence rates of mixtures of polynomial experts. *Neural Computation* **24** 3025–3051. [MR2905767](#)
- [68] MEYNET, C. (2013). An l_1 -oracle inequality for the Lasso in finite mixture Gaussian regression models. *ESAIM: Probability and Statistics* **17** 650–671. [MR3126156](#)
- [69] MOERLAND, P. (1999). Classification using localized mixture of experts. In *Ninth International Conference on Artificial Neural Networks* **2** 838–843.
- [70] MONTUELLE, L., LE PENNEC, E. et al. (2014). Mixture of Gaussian regressions model with logistic weights, a penalized maximum likelihood approach. *Electronic Journal of Statistics* **8** 1661–1695. [MR3263134](#)
- [71] NGUYEN, D. V. and ROCKE, D. M. (2002). Tumor classification by partial least squares using microarray gene expression data. *Bioinformatics* **18** 39–50.
- [72] NGUYEN, H. D. and CHAMROUKHI, F. (2018). Practical and theoretical aspects of mixture-of-experts modeling: An overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **8** e1246.
- [73] NGUYEN, H. D., CHAMROUKHI, F. and FORBES, F. (2019). Approximation results regarding the multiple-output Gaussian gated mixture of linear experts model. *Neurocomputing* **366** 208–214.
- [74] NGUYEN, H. D., LLOYD-JONES, L. R. and MCLACHLAN, G. J. (2016). A universal approximation theorem for mixture-of-experts models. *Neural Computation* **28** 2585–2593. [MR3866414](#)
- [75] NGUYEN, H. D., NGUYEN, T., CHAMROUKHI, F. and MCLACHLAN, G. J. (2021). Approximations of conditional probability density functions in Lebesgue spaces via mixture of experts models. *Journal of Statistical Distributions and Applications* **8** 13.
- [76] NGUYEN, T. (2021). Model Selection and Approximation in High-dimensional Mixtures of Experts Models: from Theory to Practice, PhD Thesis, Normandie Université.
- [77] NGUYEN, T., CHAMROUKHI, F., NGUYEN, H. and MCLACHLAN, G. (2021). Approximation of probability density functions via location-scale finite mixtures in Lebesgue spaces. *Communications in Statistics – Theory*

and Methods.

- [78] NGUYEN, T., NGUYEN, H. D., CHAMROUKHI, F. and McLACHLAN, G. J. (2020). Approximation by finite mixtures of continuous density functions that vanish at infinity. *Cogent Mathematics & Statistics* **7** 1750861. [MR4095609](#)
- [79] NGUYEN, T., NGUYEN, H. D., CHAMROUKHI, F. and McLACHLAN, G. J. (2020). An l_1 -oracle inequality for the Lasso in mixture-of-experts regression models. *arXiv preprint arXiv:2009.10622*.
- [80] NGUYEN, X. (2013). Convergence of latent mixing measures in finite and infinite mixture models. *The Annals of Statistics* **41** 370–400. [MR3059422](#)
- [81] NORETS, A. et al. (2010). Approximation of conditional densities by smooth mixtures of regressions. *The Annals of statistics* **38** 1733–1766. [MR2662358](#)
- [82] NORETS, A. and PATI, D. (2017). Adaptive Bayesian estimation of conditional densities. *Econometric Theory* **33** 980–1012. [MR3667043](#)
- [83] NORETS, A. and PELENIS, J. (2014). Posterior consistency in conditional density estimation by covariate dependent mixtures. *Econometric Theory* **30** 606–646. [MR3205608](#)
- [84] PERTHAME, E., FORBES, F. and DELEFORGE, A. (2018). Inverse regression approach to robust nonlinear high-to-low dimensional mapping. *Journal of Multivariate Analysis* **163** 1–14. [MR3732337](#)
- [85] RAKHLIN, A., PANCHENKO, D. and MUKHERJEE, S. (2005). Risk bounds for mixture density estimation. *ESAIM: PS* **9** 220–229. [MR2148968](#)
- [86] RAMAMURTI, V. and GHOSH, J. (1996). Structural adaptation in mixture of experts. In *Proceedings of 13th International Conference on Pattern Recognition* **4** 704–708 vol.4.
- [87] RAMAMURTI, V. and GHOSH, J. (1998). Use of localized gating in mixture of experts networks. In *Proc.SPIE* **3390**.
- [88] REDNER, R. A. and WALKER, H. F. (1984). Mixture densities, maximum likelihood and the EM algorithm. *SIAM Review* **26** 195–239. [MR0738930](#)
- [89] RIGOLLET, P. (2012). Kullback–Leibler aggregation and misspecified generalized linear models. *The Annals of Statistics* **40** 639–665. Publisher: Institute of Mathematical Statistics. [MR2933661](#)
- [90] SATO, M. and ISHII, S. (2000). On-line EM algorithm for the normalized gaussian network. *Neural Computation* **12** 407–432.
- [91] SCHWARZ, G. et al. (1978). Estimating the dimension of a model. *The Annals of Statistics* **6** 461–464. [MR0468014](#)
- [92] STADLER, N., BUHLMANN, P. and VAN DE GEER, S. (2010). l_1 -penalization for mixture regression models. *TEST* **19** 209–256. [MR2677722](#)
- [93] R CORE TEAM (2020). R: A language and environment for statistical computing. *Vienna, Austria*.
- [94] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* **58** 267–288. [MR1379242](#)
- [95] TU, C.-C., FORBES, F., LEMASSON, B. and WANG, N. (2019). Prediction

- with high dimensional regression via hierarchically structured Gaussian mixtures and latent variables. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **68** 1485–1507. [MR4022823](#)
- [96] VAN DE GEER, S. (2000). *Empirical Processes in M-estimation* **6**. Cambridge university press. [MR1739079](#)
- [97] VAN DER VAART, A. and WELLNER, J. (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series in Statistics. Springer **58** 59. [MR1385671](#)
- [98] WHITE, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica* **50** 1–25. [MR0640163](#)
- [99] XU, L., JORDAN, M. and HINTON, G. E. (1995). An Alternative Model for Mixtures of Experts. In *Advances in Neural Information Processing Systems* (G. TESAURO, D. TOURETZKY and T. LEEN, eds.) **7**. MIT Press.
- [100] YOUNG, D. S. (2014). Mixtures of regressions with changepoints. *Statistics and Computing* **24** 265–281. [MR3165553](#)
- [101] YUKSEL, S. E., WILSON, J. N. and GADER, P. D. (2012). Twenty years of mixture of experts. *IEEE Transactions on Neural Networks and Learning Systems* **23** 1177–1193.