



HAL
open science

Distributed Learning of Mixtures of Experts

Faïcel Chamroukhi, Nhat Thien Pham

► **To cite this version:**

Faïcel Chamroukhi, Nhat Thien Pham. Distributed Learning of Mixtures of Experts. 2023. hal-04349770

HAL Id: hal-04349770

<https://hal.science/hal-04349770>

Preprint submitted on 18 Dec 2023

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

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

Distributed Learning of Mixtures of Experts

Faïcel Chamroukhi¹, Nhat Thien Pham²

¹IRT SystemX, 91120 Palaiseau, France

²Normandie Univ, UNICAEN, CNRS, LMNO, 14000 Caen, France

December 18, 2023

Abstract

In modern machine learning problems we deal with datasets that are either distributed by nature or potentially large for which distributing the computations is usually a standard way to proceed, since centralized algorithms are in general ineffective. We propose a distributed learning approach for mixtures of experts (MoE) models with an aggregation strategy to construct a reduction estimator from local estimators fitted parallelly to distributed subsets of the data. The aggregation is based on an optimal minimization of an expected transportation divergence between the large MoE composed of local estimators and the unknown desired MoE model. We show that the provided reduction estimator is consistent as soon as the local estimators to be aggregated are consistent, and its construction is performed by a proposed majorization-minimization (MM) algorithm that is computationally effective. We study the statistical and numerical properties for the proposed reduction estimator on experiments that demonstrate its performance compared to namely the global estimator constructed in a centralized way from the full dataset. For some situations, the computation time is more than ten times faster, for a comparable performance. Our source codes are publicly available on Github.

Keywords: Distributed Learning; Mixtures of Experts; Estimator Aggregation; Optimal Transport; Majorization-Minimization Algorithm

Contents

1	Introduction	3
1.1	Motivation	4
1.2	Mixtures of Experts context	4
1.3	Main contributions	5
2	Aggregating distributed MoE models	6
2.1	Problem setting	6
2.2	The core technical challenges brought by our distributed MoE approach	7
2.3	The proposed aggregation strategy	7
3	Estimator reduction via optimal transport	8
3.1	Expected transportation divergence for reduction estimator	10
3.2	Well-posedness and consistency of the reduction estimator	13
4	An MM algorithm for the reduction estimator	14
5	Experimental study	16
6	Conclusion and perspectives	19
7	Acknowledgement	20
	References	20
	Appendix	23
A	Proofs	23
B	Details of updating the gating and expert networks of the DMoE model	26
C	Pseudocode of the main DMoE algorithm	31
D	Details of implementation of numerical experiments and additional results	32

1 Introduction

In modern machine learning problems one has to deal with datasets that are not centralized. This may be related to the application context in which the data can be by nature available at different locations and not accessible in a centralized mode, or distributed for computational issues in case of a large amount of data. Indeed, even if the dataset is fully available in a centralized mode, implementing reasonable learning algorithms may be computationally demanding in case of a large number of examples. The construction of distributed techniques in a Federated Learning setting [Yang et al. \(2019\)](#) in which the model is trained collaboratively under the orchestration of a central server, while keeping the data decentralized, is an increasing area of research. The most attractive strategy is to perform standard inference on local machines to obtain local estimators, then transmits them to a central machine where they are aggregated to produce an overall estimator, while attempting to satisfy some statistical guarantees criteria.

There are many successful attempts in this direction of parallelizing the existing learning algorithms and statistical methods. Those that may be mentioned here include, among others, parallelizing stochastic gradient descent ([Zinkevich et al., 2010](#)), multiple linear regression ([Mingxian et al., 1991](#)), parallel K -means in clustering based on MapReduce ([Zhao et al., 2009](#)), distributed learning for heterogeneous data via model integration ([Merugu and Ghosh, 2005](#)), split-and-conquer approach for penalized regressions ([Chen and ge Xie, 2014](#)), for logistic regression ([Shofiyah and Sofro, 2018](#)), for k -clustering with heavy noise [Li and Guo \(2018\)](#). It is only very recently that a distributed learning approach has been proposed for mixture distributions, specifically for finite Gaussian mixtures ([Zhang and Chen, 2022a](#)). In this paper we focus on mixtures of experts (MoE) models ([Jacobs et al., 1991](#); [Jordan and Xu, 1995](#)) which extend the standard unconditional mixture distributions that are typically used for clustering purposes, to model complex non-linear relationships of a response Y conditionally on some predictors \mathbf{X} , for prediction purposes, while enjoying denseness results, *e.g.* see ([Nguyen et al., 2021](#)), in accommodating the heterogeneous nature of the data pairs (\mathbf{X}, Y) . For an overview of theoretical and practical aspects of MoE modeling, the reader is referred to [Yuksel et al. \(2012\)](#); [Nguyen and Chamroukhi \(2018\)](#).

1.1 Motivation

The class of MoE has attracted a lot of research in the statistics and machine learning community. Training MoE is performed iteratively by maximum likelihood estimation via the well-known EM algorithm (Jacobs et al., 1991; Jordan and Xu, 1995; Dempster et al., 1977). For the high-dimensional setting (when d is large, including for $d \gg N$), Nguyen et al. (2022) recently proved a non-asymptotic oracle inequality result for model selection via Lasso-penalization in MoE. In this paper we consider the context of distributed data with large N which may rend the computations costly or infeasible, degrade the quality and interpretability of the reduced estimator at the central server when aggregated from the local machines, typically when using simple averaging. Our approach also concerns preserving privacy and reducing the communication between the local machines and the central server in the construction of a principled aggregation estimator.

1.2 Mixtures of Experts context

Let $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ be a sample of N independently and identically distributed (i.i.d.) observations from the pair (\mathbf{X}, Y) where $\mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^d$ represents the vector of predictors and $y_i \in \mathcal{Y}$ represents a real response in the regression context or a categorical response in the classification one. For simpler notation, from now on, we assume that the vector \mathbf{x} is augmented by the scalar 1 so that $\mathbf{x} \leftarrow (1, \mathbf{x}^\top)^\top$. A MoE model defines the conditional distribution of the response y given the covariates \mathbf{x} as a mixture distribution with covariate-dependent mixing proportions $\pi_k(\mathbf{x}; \boldsymbol{\alpha})$ and conditional mixture components $\varphi(y|\mathbf{x}; \boldsymbol{\beta}_k)$, via the following conditional probability density function (p.d.f)

$$f(y|\mathbf{x}; \boldsymbol{\theta}) := \sum_{k=1}^K \pi_k(\mathbf{x}; \boldsymbol{\alpha}) \varphi(y|\mathbf{x}; \boldsymbol{\beta}_k), \quad (1)$$

where $\boldsymbol{\theta} = (\boldsymbol{\alpha}^\top, \boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_K^\top)^\top$ is the parameter vector of the model and K is the number of experts. Here, $\pi_k(\mathbf{x}; \boldsymbol{\alpha})$ is the gating network, usually modelled by a softmax function given by $\pi_k(\mathbf{x}; \boldsymbol{\alpha}) = \frac{\exp(\mathbf{x}^\top \boldsymbol{\alpha}_k)}{1 + \sum_{k'=1}^{K-1} \exp(\mathbf{x}^\top \boldsymbol{\alpha}_{k'})}$ with parameter vector $\boldsymbol{\alpha} = (\boldsymbol{\alpha}_1^\top, \dots, \boldsymbol{\alpha}_{K-1}^\top)^\top$, $\boldsymbol{\alpha}_k \in \mathbb{R}^{d+1}$ for $k \in [K-1]$, whereas the conditional density $\varphi(y|\mathbf{x}; \boldsymbol{\beta}_k)$ for $k \in [K]$ is usually a Gaussian p.d.f in regression or the softmax function in multinomial classification. Here we focus on

the MoE model for regression with Gaussian experts with mean $\mathbf{x}^\top \boldsymbol{\beta}_k$ and variance σ_k^2 , with $\boldsymbol{\beta}_k \in \mathbb{R}^{d+1}$, being the regression coefficients vector, and derive it for the direct classification problem as provided in Section B.4 of the supplementary material.

1.3 Main contributions

Our main contributions in the paper can be summarized as follows:

- i) We propose a principled distributed learning approach of mixture of experts models (DMoE) with an aggregation strategy to construct a reduction estimator from local estimators fitted parallelly on decentralized data. To the best of our knowledge, we are the first to study this distributed learning of MoE models from decentralized data or for a large amount of data;
- ii) The collaborative learning of the MoE model is performed well via an aggregation based on an optimal transport minimization (Proposition 3.3) between the large MoE composed of local estimators and the unknown desired MoE model;
- iii) We show that the aggregated reduction estimator is well-posed and consistent as soon as the local estimators are consistent (Proposition 3.4);
- iv) This work extends the work of Zhang and Chen (2022a) on density estimation via unconditional mixture of Gaussians, to the non-trivial case of conditional density estimation via mixture of experts for prediction, with a key result of an MM algorithm for aggregating local models and theoretical results on the statistical consistency of the proposed approach.
- v) The core technical challenges brought by our DMoE are presented in Sections 2.2 and 2.3.

2 Aggregating distributed MoE models

2.1 Problem setting

We consider the aggregation at a central server of MoE models trained separately on decentralized data at local machines to estimate a reduced MoE model.

Suppose that we have decentralized M subsets $\mathcal{D}_1, \dots, \mathcal{D}_M$ of a data set \mathcal{D} of N examples, stored on M local machines. Let N_m be the sample size of the subset \mathcal{D}_m , *i.e.*, $\sum_{m=1}^M N_m = N$. Note that in case of distributing the data for computational reasons, many strategies can be adopted, including by randomly dividing them into enough small disjoint subsets or by constructing small bootstrap sub-samples like in (Kleiner et al., 2014). We propose to approximate the true unknown data distribution via a MoE model of the form $f(\cdot|\mathbf{x}, \boldsymbol{\theta}^*)$ as defined in (1), where $\boldsymbol{\theta}^*$ is the true unknown parameter vector. The class of MoE models is dense as recently demonstrated in (Nguyen et al., 2021) and enjoys attractive approximation capabilities. For simpler notation, we will denote by f^* the true conditional p.d.f $f(\cdot|\mathbf{x}, \boldsymbol{\theta}^*)$.

Consider a local estimation of the MoE model on a data subset \mathcal{D}_m , $m \in [M]$, by maximizing the locally observed-data likelihood (*e.g.*, using EM algorithm). Let \hat{f}_m be the local estimator of the MoE conditional density $f(\cdot|\mathbf{x}, \hat{\boldsymbol{\theta}}_m)$, defined by

$$\hat{f}_m := f(\cdot|\mathbf{x}, \hat{\boldsymbol{\theta}}_m) = \sum_{k=1}^K \pi_k(\mathbf{x}; \hat{\boldsymbol{\alpha}}^{(m)}) \varphi(\cdot|\mathbf{x}; \hat{\boldsymbol{\beta}}_k^{(m)}), \quad (2)$$

where $\hat{\boldsymbol{\theta}}_m = (\hat{\boldsymbol{\alpha}}^{(m)}, \hat{\boldsymbol{\beta}}_1^{(m)}, \dots, \hat{\boldsymbol{\beta}}_K^{(m)})$ is the locally estimated parameter vector on the sub-sample \mathcal{D}_m . The question is how to approximate the true global density f^* from the local densities \hat{f}_m , and more importantly, how to aggregate the local estimators $\hat{\boldsymbol{\theta}}_m$ to produce a single aggregated estimator for the true parameter vector $\boldsymbol{\theta}^*$. Hereafter, we use “local density” and “local estimator” to respectively refer to the conditional density function \hat{f}_m and the parameter vector $\hat{\boldsymbol{\theta}}_m$, of the MoE model estimated at local machine m based on \mathcal{D}_m , $m \in [M]$.

2.2 The core technical challenges brought by our distributed MoE approach

A natural and intuitive way to approximate the true density f^* is to consider the weighted average of the densities \hat{f}_m 's with the weights λ_m 's being the sample proportions N_m/N that sum to one. This weighted average density, that can be defined by

$$\bar{f}^W = \sum_{m=1}^M \lambda_m \hat{f}_m, \quad (3)$$

as an average model approximates very well the true density f^* . However, this aggregation strategy has two issues. Firstly, it only provides an approximation to the conditional density value, *i.e.*, for each $\mathbf{x} \in \mathcal{X}$ we have $\bar{f}^W(y|\mathbf{x}) \approx f^*(y|\mathbf{x})$, but does not approximate directly the true parameter $\boldsymbol{\theta}^*$, which is desirable for many reasons. Secondly, since each \hat{f}_m is a mixture of K expert components, we can express the density \bar{f}^W in (3) as

$$\bar{f}^W = \sum_{m=1}^M \sum_{k=1}^K \lambda_m \pi_k(\mathbf{x}; \hat{\boldsymbol{\alpha}}^{(m)}) \varphi(\cdot|\mathbf{x}; \hat{\boldsymbol{\beta}}_k^{(m)}).$$

One can see that $\sum_{m=1}^M \sum_{k=1}^K \lambda_m \pi_k(\mathbf{x}; \hat{\boldsymbol{\alpha}}^{(m)}) = 1$ for all $\mathbf{x} \in \mathcal{X}$, so \bar{f}^W can be viewed as a MoE model with MK components in which the component densities are $\varphi(\cdot|\mathbf{x}; \hat{\boldsymbol{\beta}}_k^{(m)})$ and the gating functions are $\lambda_m \pi_k(\mathbf{x}; \hat{\boldsymbol{\alpha}}^{(m)})$. Therefore, although this MK -component MoE approximates well the true density f^* , it has a wrong number of components, *i.e.*, MK instead of K , which makes the approximation results difficult to interpret. Hence, this aggregation is not desirable.

2.3 The proposed aggregation strategy

Let \mathcal{M}_K denotes the space of all K -component MoE models as in (1). Two other common strategies to aggregate the local densities $\hat{f}_1, \dots, \hat{f}_M$ are via

$$\bar{f}^B := f(y|\mathbf{x}; \bar{\boldsymbol{\theta}}^B) = \arg \inf_{g \in \mathcal{M}_K} \sum_{m=1}^M \lambda_m \rho(\hat{f}_m, g), \quad (4)$$

and

$$\bar{f}^R := f(y|\mathbf{x}; \bar{\boldsymbol{\theta}}^R) = \arg \inf_{g \in \mathcal{M}_K} \rho(\bar{f}^W, g), \quad (5)$$

where $\rho(\cdot, \cdot)$ is some divergence defined on the space of finite mixture of experts distributions. The solutions \bar{f}^B and \bar{f}^R are often known as barycenter solution and reduction solution, respectively, with their interpretations are given as follows.

In the case of \bar{f}^B , we are finding a K -component MoE model g that can be viewed as a barycenter of the local models $\hat{f}_1, \dots, \hat{f}_M$ with respect to (w.r.t) the weights $\lambda_1, \dots, \lambda_M$ and the divergence $\rho(\cdot, \cdot)$. Whereas in the case of \bar{f}^R , we are finding a K -component MoE model g that is closest to the weighted average density \bar{f}^W defined in (3), w.r.t the divergence $\rho(\cdot, \cdot)$. Because we are searching for solutions in \mathcal{M}_K , both \bar{f}^B and \bar{f}^R solve the problem of wrong number of components in \bar{f}^W . Both solutions are desirable, and in fact, can be shown to be connected under specific choices of $\rho(\cdot, \cdot)$. However, we prefer the reduction solution \bar{f}^R for addressing the following core technical reasons. First, as we have already remarked, the MK -component MoE \bar{f}^W is a good approximation to the true density f^* , so it makes more sense to find a K -component MoE model that approximates \bar{f}^W , which we already know is very good. Second, the barycenter approach may lead to a counter-intuitive solution under some specific $\rho(\cdot, \cdot)$, *e.g.*, as already shown in Zhang and Chen (2022a) for the case of univariate Gaussian mixture models and the 2-Wasserstein divergence with Euclidean ground distance. Finally, as we can see, given the same divergence $\rho(\cdot, \cdot)$, the computation in (4) will be more expensive than that of (5), especially in our large data context.

The proposed reduction strategy addresses these key technical challenges mentioned above and provides a desirable reduced MoE \bar{f}^R that has the correct number of components and is constructed upon a well-posed and consistent reduction estimator $\bar{\theta}^R$ of θ^* as soon as the local ones are.

3 Estimator reduction via optimal transport

Thus, one objective is to develop an efficient algorithm to find the reduction estimator $\bar{\theta}^R$. We shall denote the components of the reduction estimator by $\bar{\theta}^R = (\bar{\alpha}^R, \bar{\beta}_1^R, \dots, \bar{\beta}_K^R)$. A description of our considered problem can be seen in Fig. 1, in which, the gating and expert parameters are explicitly indicated for each component of the local and the reduced models.




Figure 1: Description of the proposed aggregation framework. Left panels are the components of the local ME models with their corresponding estimated parameters. Right panels are the components of the aggregated K -component ME model, *i.e.*, the desired solution. The unknown parameters are therefore $\bar{\alpha}^R, \bar{\beta}_1^R, \dots, \bar{\beta}_K^R$.

For mathematical convenience, we will incorporate the weights λ_m into the local gating functions $\pi_k(\mathbf{x}; \hat{\alpha}^{(m)})$ and write them simply by $\{\hat{\pi}_\ell(\mathbf{x})\}_{\ell \in [MK]}$, which will be also referred to as gating functions. The corresponding experts in \bar{f}^W will be abbreviated by $\{\hat{\varphi}_\ell(y|\mathbf{x})\}_{\ell \in [MK]}$. Similarly, for $g \in \mathcal{M}_K$, we will abbreviate its K gating functions by $\{\pi_k(\mathbf{x})\}_{k \in [K]}$, and its K experts by $\{\varphi_k(y|\mathbf{x})\}_{k \in [K]}$. We then have, for $g \in \mathcal{M}_K$,

$$\bar{f}^W = \sum_{\ell=1}^{MK} \hat{\pi}_\ell(\mathbf{x}) \hat{\varphi}_\ell(\cdot|\mathbf{x}) \text{ and } g = \sum_{k=1}^K \pi_k(\mathbf{x}) \varphi_k(\cdot|\mathbf{x}). \quad (6)$$

By substituting (3) into (5) we can express explicitly the reduced density as

$$\bar{f}^R = \arg \inf_{g \in \mathcal{M}_K} \rho \left(\sum_{m=1}^M \lambda_m \hat{f}_m, g \right), \quad (7)$$

that is, we seek for a K -component MoE model g of form as in (1) that is closest to the MK -component MoE $\bar{f}^W = \sum_{m=1}^M \lambda_m \hat{f}_m$ w.r.t some divergence $\rho(\cdot, \cdot)$. In general, an analytic solution is difficult to obtain for such reduction problem, especially in the context of MoE models where the conditional density is constructed upon several different gating and expert functions. Therefore, one has to find a numerical solution for the the reduction model \bar{f}^R .

Firstly, the solution much depends on the choice of the divergence $\rho(\cdot, \cdot)$, which measures the goodness of a candidate model $g \in \mathcal{M}_K$. The best candidate g should minimize its divergence from the large MoE model \bar{f}^W . Secondly, computing the divergence between two MoE models is not explicit and difficult. Therefore, borrowing the idea in [Zhang and Chen \(2022a\)](#) that finds a reduction estimator for finite Gaussian mixtures based on minimizing transportation divergence between the large starting mixture and the desired mixture, we establish a framework for solving a reduction estimator for MoE models formalized in (7).

3.1 Expected transportation divergence for reduction estimator

Let $h = \sum_{\ell=1}^L \hat{\pi}_\ell(\mathbf{x}) \hat{\varphi}_\ell(y|\mathbf{x})$, $L \in \mathbb{N}^*$, and $g = \sum_{k=1}^K \pi_k(\mathbf{x}) \varphi_k(y|\mathbf{x})$ be two MoE models of L and K components, respectively. Here, h will play the role as \bar{f}^W in problem (5). We wish to define a divergence that measures the dissimilarity between the two MoE models h and g . Let Φ denotes the family of the component conditional densities $\varphi(\cdot|\mathbf{x}; \boldsymbol{\beta}) =: \varphi(\cdot|\mathbf{x})$. Let $\hat{\boldsymbol{\pi}}$ and $\boldsymbol{\pi}$ be, respectively, the column vectors of gating functions $\hat{\pi}_\ell(\cdot)$'s and $\pi_k(\cdot)$'s. For $\mathbf{x} \in \mathcal{X}$, we denote

$$\begin{aligned} \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi}) &:= \mathbf{\Pi}(\hat{\boldsymbol{\pi}}(\mathbf{x}), \boldsymbol{\pi}(\mathbf{x})) \\ &= \{ \mathbf{P} \in \mathbb{U}^{L \times K} : \mathbf{P} \mathbf{1}_K = \hat{\boldsymbol{\pi}}(\mathbf{x}), \mathbf{P}^\top \mathbf{1}_L = \boldsymbol{\pi}(\mathbf{x}) \}, \end{aligned}$$

where \mathbb{U} denotes the interval $[0, 1]$, $\mathbf{1}_K$ and $\mathbf{1}_L$ are vectors of all ones. In other words, for $\mathbf{x} \in \mathcal{X}$, $\mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi})$ denotes the set of all matrices \mathbf{P} of size $L \times K$, with entries $P_{\ell k} \in [0, 1]$, satisfying the marginal constraints

$$\sum_{k=1}^K P_{\ell k} = \hat{\pi}_\ell(\mathbf{x}), \quad \forall \ell \in [L], \quad \text{and} \quad \sum_{\ell=1}^L P_{\ell k} = \pi_k(\mathbf{x}), \quad \forall k \in [K]. \quad (8)$$

Given $\mathbf{x} \in \mathcal{X}$, one can think of such matrix $\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi})$ as a plan that transports an amount of material distributed according to $\hat{\boldsymbol{\pi}}(\mathbf{x})$ to distributed according to $\boldsymbol{\pi}(\mathbf{x})$, for short, we will say “transports $\hat{\boldsymbol{\pi}}(\mathbf{x})$ to $\boldsymbol{\pi}(\mathbf{x})$ ”. Suppose that the transportation is costing for each unit of material. We are then interested in the optimal way, *i.e.*, with a minimum cost, to transports $\hat{\boldsymbol{\pi}}(\mathbf{x})$ to $\boldsymbol{\pi}(\mathbf{x})$. The optimal plan is clearly depends on \mathbf{x} . We will write $\mathbf{P}(\mathbf{x})$ to indicate such optimal plan.

Definition 3.1 (Expected transportation divergence). *The expected transportation divergence between two MoE models h and g is defined by*

$$\begin{aligned} \mathcal{T}_c(h, g) &= \mathbb{E} \left[\inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi})} \sum_{\ell=1}^L \sum_{k=1}^K P_{\ell k} c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x})) \right] \\ &=: \mathbb{E} \left[\mathcal{T}_c(h, g, \mathbf{x}) \right], \end{aligned} \quad (9)$$

where $c(\cdot, \cdot)$ is a real-valued bivariate function defined on $\Phi \times \Phi$ satisfies $c(\varphi_1, \varphi_2) \geq 0$ for all $\varphi_1, \varphi_2 \in \Phi$, and the equality holds if and only if $\varphi_1 = \varphi_2$.

Here, the expectation is taken w.r.t \mathbf{x} and the function c is often referred to as the cost function. Assume that, for a given $\mathbf{x} \in \mathcal{X}$, $\mathbf{P}(\mathbf{x})$ is an optimal solution to the optimization problem inside the expectation operator in (9). Then the value of $\mathcal{T}_c(h, g, \mathbf{x})$ can be interpreted as the optimal total cost of transporting $\hat{\boldsymbol{\pi}}(\mathbf{x})$ to $\boldsymbol{\pi}(\mathbf{x})$ with the unit cost of transportation is proportional to the values $c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x}))$'s. This interpretation is common in optimal transport and known as Kantorovich formulation [Oberman and Ruan \(2015\)](#). The value $\mathcal{T}_c(h, g)$ is therefore defined as the expectation of these optimal transportation costs. The cost function c must be chosen such that it reflects the dissimilarity between the conditional densities $\hat{\varphi}_{\ell}(\cdot|\mathbf{x})$ and $\varphi_k(\cdot|\mathbf{x})$, and also has to be easy in calculation. We show later that the Kullback-Leibler (KL) divergence as choice for c is suitable for MoE models.

Thus, by considering $\rho(\cdot, \cdot)$ being the expected transportation divergence $\mathcal{T}_c(\cdot, \cdot)$, problem (5) becomes

$$\begin{aligned} \bar{f}^R &= \arg \inf_{g \in \mathcal{M}_K} \mathcal{T}_c(\bar{f}^W, g) \\ &= \arg \inf_{g \in \mathcal{M}_K} \left\{ \mathbb{E} \left[\inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi})} \sum_{\ell=1}^{MK} \sum_{k=1}^K P_{\ell k} c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x})) \right] \right\}. \end{aligned} \quad (10)$$

We will now use $\mathcal{T}_c(g)$ to refer to $\mathcal{T}_c(\bar{f}^W, g)$, since our objective now is to minimize the transportation divergence $\mathcal{T}_c(\bar{f}^W, g)$ as a function of $g \in \mathcal{M}_K$. Problem (10) involves two optimization problems: one over $\mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi})$ and one over \mathcal{M}_K . However, we show that the constraint $\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi})$ in fact can be relaxed to $\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)$. This means we only need \mathbf{P} to satisfy the constraints in the first part of (8), those in the second part (8) are redundant. In other words, instead of searching, for each \mathbf{x} , a plan \mathbf{P} that satisfies $\boldsymbol{\pi}(\mathbf{x})$, we can move

$\boldsymbol{\pi}$, the gating functions constructing g , to match \mathbf{P} . Indeed, for \bar{f}^W and g as in (6), let us define $\mathcal{R}_c(\bar{f}^W, g)$ as a function of g as follows

$$\mathcal{R}_c(\bar{f}^W, g) = \mathbb{E} \left[\inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)} \sum_{\ell=1}^{MK} \sum_{k=1}^K P_{\ell k} c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x})) \right] \quad (11)$$

$$=: \mathbb{E} \left[\mathcal{R}_c(\bar{f}^W, g, \mathbf{x}) \right]. \quad (12)$$

Here, the $MK \times K$ matrix \mathbf{P} is now free of the constraints involving the gating network $\boldsymbol{\pi}$ of g , namely, for all $\mathbf{x} \in \mathcal{X}$, the equality $\mathbf{P}^\top \mathbf{1}_{MK} = \boldsymbol{\pi}(\mathbf{x})$ is unnecessary. Then, the following proposition shows that solving $\bar{f}^R = \arg \inf_{g \in \mathcal{M}_K} \mathcal{T}_c(\bar{f}^W, g)$ can actually be reduced to solving $\bar{f}^R = \arg \inf_{g \in \mathcal{M}_K} \mathcal{R}_c(\bar{f}^W, g)$, which is much easier. Before presenting the proposition, we define $\mathcal{P}(\bar{f}^W, g, \mathbf{x})$ a function of $g \in \mathcal{M}_K$ and $\mathbf{x} \in \mathcal{X}$ as follows

$$\mathcal{P}(\bar{f}^W, g, \mathbf{x}) = \arg \inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)} \sum_{\ell=1}^{MK} \sum_{k=1}^K [P_{\ell k} c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x}))]. \quad (13)$$

This function returns the optimal plan for the optimization problem inside the expectation operator in the definition of $\mathcal{R}_c(\bar{f}^W, g)$. Similarly to $\mathcal{T}_c(g)$, let the dependence of the functions on \bar{f}^W in the background, so that $\mathcal{R}_c(\bar{f}^W, g)$, $\mathcal{R}_c(\bar{f}^W, g, \mathbf{x})$ and $\mathcal{P}(\bar{f}^W, g, \mathbf{x})$ will be written as $\mathcal{R}_c(g)$, $\mathcal{R}_c(g, \mathbf{x})$ and $\mathcal{P}(g, \mathbf{x})$, respectively.

Proposition 3.2. *Let $\mathcal{T}_c(g)$, $\mathcal{R}_c(g)$ and $\mathcal{P}(g, \mathbf{x})$ defined as above. Then $\inf_{g \in \mathcal{M}_K} \mathcal{T}_c(g) = \inf_{g \in \mathcal{M}_K} \mathcal{R}_c(g)$. The reduction solution is hence given by $\bar{f}^R = \arg \inf_{g \in \mathcal{M}_K} \mathcal{R}_c(g)$ and the gating functions of \bar{f}^R can be calculated by*

$$\pi_k(\mathbf{x}) = \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}(\bar{f}^R, \mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}, \quad (14)$$

where $\mathcal{P}_{\ell k}(\bar{f}^R, \mathbf{x})$ denotes the entry (ℓ, k) of $\mathcal{P}(\bar{f}^R, \mathbf{x})$.

The proof of Proposition 3.2 is given in Appendix A.1 of the supplementary material. Thus, thanks to Proposition 3.2, our objective now is to minimize $\mathcal{R}_c(g)$ w.r.t g . By replacing $\mathcal{T}_c(g)$ by $\mathcal{R}_c(g)$, problem (10) becomes

$$\begin{aligned} \bar{f}^R &= \arg \inf_{g \in \mathcal{M}_K} \mathcal{R}_c(g) \\ &= \arg \inf_{g \in \mathcal{M}_K} \left\{ \mathbb{E} \left[\inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)} \sum_{\ell=1}^{MK} \sum_{k=1}^K P_{\ell k} c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x})) \right] \right\}. \end{aligned} \quad (15)$$

This optimization can be done with the help of majorization-minimization (MM) algorithm (*e.g.*, [Lange \(2004\)](#)). We defer the numerical approach for solving the problem (15) in Section 4, now we will show that the problem is well-posed and state the conditions for the consistency of the reduction estimator.

3.2 Well-posedness and consistency of the reduction estimator

We make the following standard assumptions before the consistency proposition.

- A1. The dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ is an i.i.d. sample from the K -component MoE model $f(y|\mathbf{x}, \boldsymbol{\theta}^*)$ that is ordered and initialized ([Jiang and Tanner, 1999b](#)).
- A2. The cost function $c(\cdot, \cdot)$ is continuous in both arguments, and $c(\varphi_1, \varphi_2) \rightarrow 0$ if and only if $\varphi_1 \rightarrow \varphi_2$ in distribution.
- A3. $c(\cdot, \cdot)$ is convex in the second argument.

■ **Well-posedness.** First, we observe that for any $\mathbf{x} \in \mathcal{X}$, the minimization problem inside the expectation operator in (15) is in fact a linear programming (LP) problem. Moreover, for $\mathbf{x} \in \mathcal{X}$, we see that $\mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)$ is a nonempty set, and the sum $\sum_{\ell=1}^{MK} \sum_{k=1}^K P_{\ell k} c(\bar{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x}))$ is bounded below by zero for all $\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)$. Therefore, this LP problem has a global minimizer (*e.g.*, see [Boyd and Vandenberghe \(2004\)](#)). Hence, for all $g \in \mathcal{M}_K$, there exists non-negative finite expectation for the random quantity $\mathcal{R}_c(g, \mathbf{x})$. Let $\mathcal{C}(\mathbf{P}, \boldsymbol{\varphi})$ be a function of the transport plan w.r.t the component densities of g , given by

$$\mathcal{C}(\mathbf{P}, \boldsymbol{\varphi}) = \sum_{\ell=1}^{MK} \sum_{k=1}^K P_{\ell k} c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x})),$$

where $\mathbf{P} \in \mathbb{U}^{MK \times K}$, $\mathbb{U} = [0, 1] \subset \mathbb{R}$, and $\boldsymbol{\varphi} = (\varphi_1(\cdot|\mathbf{x}), \dots, \varphi_K(\cdot|\mathbf{x})) \in \boldsymbol{\Phi}^K$. Let $\mathcal{I}(g, \mathbf{x})$ be a function of g and \mathbf{x} defined by $\mathcal{I}(g, \mathbf{x}) = \inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)} \mathcal{C}(\mathbf{P}, \boldsymbol{\varphi})$, *i.e.*, the optimal cost given g and \mathbf{x} inside the expectation operator in (15). Then, the objective function of the problem (15) can then be written as

$$\mathcal{R}_c(g) = \mathbb{E}[\mathcal{I}(g, \mathbf{x})].$$

Proposition 3.3. *Assume there exists $\Delta \in \mathbb{R}_+$ such that $\mathcal{I}(g, \mathbf{x}) \leq \Delta$ for all $g \in \mathcal{M}_K$, $\mathbf{x} \in \mathcal{X}$. Then $\mathcal{R}_c(g)$ is continuous and convex as a function of $g \in \mathcal{M}_K$. It follows that the problem (15) has a global solution.*

The proof of the Proposition 3.3 is provided in Appendix A.2 of the supplementary material. Note that the condition on the boundedness of the optimal transportation costs is common, and the solution to the problem (15) is not necessarily unique. The assumption A3. is necessary to show that problem (15) is well-posed, and it holds for most of divergences including the KL one used in this paper, see *e.g.*, Dragomir (2013).

■ **Consistency.** The reduction estimator $\bar{\theta}^R$ has a desired property that it is a consistent estimator of the true parameter θ^* as soon as the local estimators are consistent estimators of θ^* .

Proposition 3.4. *Let $\bar{\theta}^R$ be the parameter of the reduction density \bar{f}^R defined in (5) with ρ being the expected transportation divergence \mathcal{T}_c . Suppose assumptions A1.-A2. are satisfied. Then $\bar{\theta}^R$ is a consistent estimator of θ^* .*

The assumption that the parameters are ordered and initialized in Assumption A1. is common and necessary because under it the MoE models are identifiable (Jiang and Tanner, 1999a). Assumption A2. ensures the divergence between \bar{f}^R and f^* tends to 0 as $N \rightarrow \infty$. The proof of Proposition 3.4 is given in the supplementary material Sec. A.3.

4 An MM algorithm for the reduction estimator

We see that, the optimization problem (15) is accompanied by the objective function $\mathcal{R}_c(g)$ defined upon another optimization, w.r.t the transportation plan \mathbf{P} , that makes the approach such as gradient descent hard to apply directly. For that we derive an MM algorithm (Lange, 2004), which requires the definition of a so-called majorant function, to solve the problem (15). The MM algorithm in our problem consists of starting from an initial model $g^{(0)} \in \mathcal{M}_K$, then at each iteration t , find a majorant function for $\mathcal{R}_c(g)$ at $g^{(t)}$, and minimize it to obtain the next model $g^{(t+1)}$. The generated sequence $(g^{(t)})_{t \geq 1}$ satisfies $\mathcal{R}_c(g^{(0)}) \geq \mathcal{R}_c(g^{(1)}) \geq \dots$

We keep running until convergence is reached, *i.e.*, there is no significant change in the value of $\mathcal{R}_c(g)$. The local minimum obtained at convergence is then taken as the desired model \bar{f}^R . This descent property is attractive and makes it very stable, as compared to other (*e.g.* gradient descent) algorithms for which the monotonicity property often does not hold. Moreover, in gradient descent, we need to calculate the gradient of the objective function in (15), which is not obvious, as it is defined upon another optimization (w.r.t the transportation plan \mathbf{P}).

The following proposition provides a majorizer for $\mathcal{R}_c(g)$.

Proposition 4.1. *Let $g^{(t)}$ be the model obtained at the t -th MM iteration and let*

$$\mathcal{S}_c(g, g^{(t)}) = \mathbb{E} \left[\sum_{\ell=1}^{MK} \sum_{k=1}^K \mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x}) c(\widehat{\varphi}_\ell(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x})) \right], \quad (16)$$

where $\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})$ is given by

$$\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x}) = \begin{cases} \widehat{\pi}_\ell(\mathbf{x}) & \text{if } k = \arg \inf_{k' \in [K]} c(\widehat{\varphi}_\ell(\cdot|\mathbf{x}), \varphi_{k'}^{(t)}(\cdot|\mathbf{x})) \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

Then $\mathcal{S}_c(g, g^{(t)})$ is a majorant function of $\mathcal{R}_c(g)$ at $g^{(t)}$.

The proof of Proposition 4.1 is provided in Appendix A.4 of the supplementary material. Here we see that, the plan with entries $\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})$ given in (17) is an optimal plan for the optimization problem (13) when $g = g^{(t)}$ (this is clarified more in the proof of the proposition). Then, the next step is to minimize $\mathcal{S}_c(g, g^{(t)})$, w.r.t g , to obtain a next point $g^{(t+1)}$ for the MM algorithm. As we can see, the optimization for $\mathcal{S}_c(g, g^{(t)})$ can be performed separately. At each iteration t -th, the parameter vector of the expert k can be optimized via

$$\varphi_k^{(t+1)}(\cdot|\mathbf{x}) = \arg \inf_{\varphi \in \Phi} \left\{ \mathbb{E} \left[\sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x}) c(\widehat{\varphi}_\ell(\cdot|\mathbf{x}), \varphi(\cdot|\mathbf{x})) \right] \right\}. \quad (18)$$

According to each specification of the experts and the cost function $c(\cdot, \cdot)$, problem (18) has a certain form that is supposed to be easy to solve, as for the cases of the Gaussian and logistic regression experts, with c is the KL-divergence. Hence, our algorithm alternates between the two following steps until convergence:

- calculating the majorant function $\mathcal{S}_c(g, g^{(t)})$ as in (16), which in fact only involves updating $\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})$ according to (17);

- updating the experts parameters by solving (18).

The gating function parameters are updated based on equation (14), in which the theoretical solution \bar{f}^R is replaced by model g^* obtained by the MM algorithm. For that, we need a supporting sample \mathcal{D}_S available at the central server, details in Section B of the supplementary material

Communication and computation costs The communication cost in our case involves sending to the central machine the local parameters with a cost of $\mathcal{O}(M.K.2d)$ and the sample \mathcal{D}_S with a cost of $\mathcal{O}(N/M)$. Moreover, the communication in our approach occurs solely in one direction (local to central) and takes place only once, whereas in many other approaches, maintaining communication (in both directions and/or between nodes) is crucial and potentially impose high costs. The Algorithm 1 provided in Section C of the supplementary material gives a pseudo-code that summarizes our main algorithm.

5 Experimental study

We compared, on simulated and real datasets, the performances of our estimator $\bar{\theta}^R$ to the following ones:

- Global estimator θ^G : the MLE of the MoE based on the full dataset calculated in a centralized way, *i.e.*, $\theta^G \in \arg \inf_{\theta} \sum_{i=1}^N \log f(y|\mathbf{x}; \theta)$ with f defined in (1). Note that here the generic term global does not refer to a global solution.
- Middle estimator $\bar{\theta}^{Mid}$: the parameter of the local density that gives the smallest sum of transportation divergences weighted by the λ_m 's, defined as, $\bar{f}^{Mid} = f(y|\mathbf{x}; \bar{\theta}^{Mid}) = \arg \inf_{g \in \{\hat{f}_1, \dots, \hat{f}_M\}} \sum_{m=1}^M \lambda_m \mathcal{T}_c(\hat{f}_m, g)$.
- Weighted average estimator $\bar{\theta}^W$: an ad-hoc estimator defined as the pointwise weighted average of the local estimators, *i.e.*, $\bar{\theta}^W = \sum_{m=1}^M \lambda_m \hat{\theta}_m$.

The metrics used to compare the estimators are:

- The expected transportation divergence, defined in (9) with c being KL divergence, between the true MoE model and the estimated MoE model.
- Log likelihood value of the estimated parameter evaluated on the testing set.
- Mean squared error (MSE) between the true and estimated parameters.
- Relative prediction error on the testing set, defined by $\text{RPE} = \sum_{i=1}^n (y_i - \hat{y}_i)^2 / \sum_{i=1}^n y_i^2$, where y_i and \hat{y}_i are the true and the predicted responses.
- Adjusted Rand Index (ARI) between the true clustering in the testing set and the clustering predicted by the estimated MoE model.
- Learning time, *i.e.*, the total running time to obtain the estimator.

We conducted simulations in which we fixed $K = 4$ and $d = 20$ and considered datasets of different sizes, from moderate to very large ($N \in \{10^5, 5 \times 10^5, 10^6\}$), to illustrate the statistical performance of the estimators. The testing sets (20% of the data) has 50k, 125k and 500k observations, respectively. It is worth mentioning that each of the datasets with one million observations occupies around 1.6 GB of memory. The estimators are computed using training sets and the metrics (except the learning time) are evaluated on testing sets. Additional details of the data generating process, the implementation, as well as the convergence behavior of the MM algorithm in a typical run can be found in Appendix D of the supplementary material.

Figure 2 shows the box plots of 100 Monte Carlo runs of the evaluation metrics when $N = 10^6$. As we can see, the performance of our reduction estimator is as good as that of the global estimator when M is 4 or 16 machines. When M is 64 or 128 machines, the errors of the reduction estimator are slightly higher than the global estimator, but still better than the middle and weighted average estimators. In terms of transportation distance, we can see that with the reduction estimator, the obtained MoE model is as close as the MoE model formed by the global estimator, even when $M = 128$ machines. Finally, the distributed approach requires much less learning time than the global one, for example it is three to ten times faster when using 4 to 64 machines.




Figure 2: Performance of the Global MoE (G), Reduction (R), Middle (M) and Weighted average (W) estimator for sample size $N = 10^6$ and M machines.

In Figure 3, we compare the evaluation metrics of the estimators across sample sizes when 128 machines were used. With this large number of machines, the learning time is significantly reduced when using the distributed approaches, especially for datasets with one million observations. The performance metrics such as MSE, RPE and ARI behave as expected across estimators and across sample sizes.

We investigate the prediction performance of the proposed distributed approach and the global approach on the Multilevel Monitoring of Activity and Sleep in Healthy people (MMASH) dataset. This dataset consists of 24 hours of continuous inter-beat interval data (IBI), triaxial accelerometer data, sleep quality, physical activity and psychological characteristics of 22 healthy young males. More details about the experiment setup of MMASH are provided in Rossi et al. (2020). The size of our training set is $N = 1162400$, and the size of testing set is 290600. The RPE, RMSE, Scattered Index (SI) and the learning time of each approach are given in Table 1. We can see, the distributed approach significantly reduced the learning time while the prediction metrics are not too different.




Figure 3: Performance of the Global (G), vs Reduction (R), Middle (M) and Weighted average (W) estimator using 128 machines for different sample sizes N .

Table 1: Comparison of Global (G) and Reduction (R) estimators on MMASH data.

	RPE	RMSE	SI	Time (m)
G	1.03%	8.58	9.01%	57.32
R, $M = 64$	1.45%	9.46	9.54%	7.84
R, $M = 128$	1.45%	9.49	9.58%	6.32

6 Conclusion and perspectives

In this paper, a principled distributed learning approach of MoE models was proposed. The approach is based on searching for a desired K -component MoE model that globally minimizes its expected transportation divergence from a large MoE composed of local MoE models fitted on local machines. An MM algorithm, with a well-designed majorizer, was proposed to solve the resulting minimization problem. The numerical studies demonstrate the effectiveness of the proposed approach through extensive experiments conducted on both simulated and real-world datasets. The demonstrated well-posedness and consistency of the

aggregated estimator, coupled with the promising performances in regression, underscores the potential practical applicability of the proposed approach in real-world scenarios. The numerical experiments of the proved updating formulas of the algorithm for the classification case to evaluate it in practice is deferred to future work.

While MoE are highly effective in capturing nonlinearity (e.g., see (Nguyen et al., 2016)) and are of broad interest in the statistics and machine learning communities, a potential use of our approach with deep neural networks is to consider for example multi-layer perception or more complex architectures for the experts, to make the impact broader. Moreover, we theoretically provided guarantees of consistency of the reduction estimator in the distributed setting regardless of the type of experts (provided they fulfill the assumptions as discussed before) which is potentially applicable to DNNs. Finally, because we focused on the non-trivial problem of models aggregation to construct a consistent MoE model with a fixed number of experts, our approach requires the local models have the same number of experts, which may not occur in practice if standard model selection approaches for MoE are applied on local machines. For its use in its current state, one may for instance use a sufficiently large number of experts given the denseness property of MoE. Choosing an aggregated estimator with the optimized number of experts is an open perspective kept for a future work.

7 Acknowledgement

This research was supported by the exploratory research facility “EXPLO” at SystemX.

References

- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.
- Chen, X. and ge Xie, M. (2014). A split-and-conquer approach for analysis of extraordinarily large data. *Statistica Sinica*, 24:1655–1684.
- 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, B*, 39(1):1–38.
- Dragomir, S. (2013). A generalization of f -divergence measure to convex functions defined on linear spaces. *Communications in Mathematical Analysis*, 15.

- Jacobs, R. A., Jordan, M. I., Nowlan, S. J., and Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87.
- Jiang, W. and Tanner, M. (1999a). On the identifiability of mixtures-of-experts. *Neural Networks*, 12(9):1253–1258.
- Jiang, W. and Tanner, M. A. (1999b). On the Identifiability of Mixtures-of-Experts. *Neural Networks*, 12:197–220.
- Jordan, M. I. and Xu, L. (1995). Convergence results for the EM approach to mixtures of experts architectures. *Neural Networks*, 8(9):1409–1431.
- Kleiner, A., Talwalkar, A., Sarkar, P., and Jordan, M. I. (2014). A scalable bootstrap for massive data. *Journal of the Royal Statistical Society: Series B: Statistical Methodology*, pages 795–816.
- Lange, K. (2004). The MM Algorithm. In *Optimization*, pages 119–136. Springer.
- Li, S. and Guo, X. (2018). Distributed k-clustering for data with heavy noise. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Merugu, S. and Ghosh, J. (2005). A distributed learning framework for heterogeneous data sources. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 208–217.
- Mingxian, X., Miller, J. J., and Wegman, E. J. (1991). Parallelizing multiple linear regression for speed and redundancy: an empirical study. *Journal of Statistical Computation and Simulation*, 39(4):205–214.
- 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*, pages e1246–n/a.
- Nguyen, H. D., Lloyd-Jones, L. R., and McLachlan, G. J. (2016). A universal approximation theorem for mixture-of-experts models. *Neural Computation*, 28(12):2585–2593.
- 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(1):13.
- Nguyen, T., Nguyen, H. D., Chamroukhi, F., and Forbes, F. (2022). A non-asymptotic approach for model selection via penalization in high-dimensional mixture of experts models. *Electronic Journal of Statistics*, 16(2):4742 – 4822.
- Oberman, A. M. and Ruan, Y. (2015). An efficient linear programming method for optimal

transportation. *arXiv: Numerical Analysis*.

- Rossi, A., Da Pozzo, E., Menicagli, D., Tremolanti, C., Priami, C., Sîrbu, A., Clifton, D. A., Martini, C., and Morelli, D. (2020). A public dataset of 24-h multi-levels psychophysiological responses in young healthy adults. *Data*, 5(4).
- Shofiyah, F. and Sofro, A. (2018). Split and conquer method in penalized logistic regression with lasso (application on credit scoring data). *Journal of Physics: Conference Series*, 1108:012107.
- Yang, Q., Liu, Y., Chen, T., and Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.*, 10(2).
- Yuksel, S. E., Wilson, J. N., and Gader, P. D. (2012). Twenty years of mixture of experts. *IEEE Trans. Neural Netw. Learning Syst.*, 23(8):1177–1193.
- Zhang, Q. and Chen, J. (2022a). Distributed Learning of Finite Gaussian Mixtures. *Journal of Machine Learning Research*, 23(99):1–40.
- Zhang, Q. and Chen, J. (2022b). Distributed learning of finite gaussian mixtures. *Journal of Machine Learning Research*, 23(99):1–40.
- Zhao, W., Ma, H., and He, Q. (2009). Parallel k-means clustering based on mapreduce. In Jaatun, M. G., Zhao, G., and Rong, C., editors, *Cloud Computing*, pages 674–679. Springer Berlin Heidelberg.
- Zinkevich, M., Weimer, M., Li, L., and Smola, A. (2010). Parallelized stochastic gradient descent. In Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., and Culotta, A., editors, *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc.

Appendix

This supplementary material provides proofs for the Propositions 3.2, 3.3, 3.4, and 4.1, in Appendices A.1, A.2, A.3, and A.4, respectively. In Appendices B.1 and B.2 we then provide details for the updating formulas for the gating network and the experts network, respectively. The formulations of these updating formulas rely on the Propositions B.1 for regression and B.2 for classification, both being proved in Appendices B.3 and B.5, respectively. Appendix C provides the summarizing pseudocode of the main algorithm. Further implementation details and experimental results are provided in Appendix D.

A Proofs

A.1 Proof of Proposition 3.2

First, by definition, $\mathcal{R}_c(g)$ is obtained by relaxing the constraint in $\mathcal{T}_c(g)$ from $\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi})$ to $\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)$, therefore the inequality

$$\inf_{g \in \mathcal{M}_K} \mathcal{T}_c(g) \geq \inf_{g \in \mathcal{M}_K} \mathcal{R}_c(g)$$

is trivial.

The inequality in the opposite direction is also true. Indeed, let $g^* = \arg \inf_{g \in \mathcal{M}_K} \mathcal{R}_c(g)$ be a minimizer of $\mathcal{R}_c(g)$, *i.e.*, we have $\mathcal{R}_c(g^*) = \inf_{g \in \mathcal{M}_K} \mathcal{R}_c(g)$. Then, in the remainder we will show that $\inf_{g \in \mathcal{M}_K} \mathcal{T}_c(g) \leq \mathcal{R}_c(g^*)$. We denote the gating and expert functions of g^* by $\{\pi^*(\mathbf{x}), \varphi_k^*(\cdot|\mathbf{x})\}$, $k \in [K]$.

According to (14), the gating functions are calculated by

$$\pi_k^*(\mathbf{x}) = \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}(g^*, \mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}. \quad (19)$$

where $\mathcal{P}_{\ell k}(g^*, \mathbf{x})$ denotes the entry (ℓ, k) of $\mathcal{P}(g^*, \mathbf{x})$ as defined in (13). The relation (19) means that the matrix $\mathcal{P}(g^*, \mathbf{x})$ is belonging to $\mathbf{\Pi}_{\mathbf{x}}(\cdot, \boldsymbol{\pi}^*)$. On the other hand, by its definition $\mathcal{P}(g^*, \mathbf{x})$ is obviously belonging to $\mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)$. Hence, $\mathcal{P}(g^*, \mathbf{x}) \in \mathbf{\Pi}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi}^*)$. Note that, this holds for all $\mathbf{x} \in \mathcal{X}$. Therefore, taking expectation we have

$$\begin{aligned} \inf_{g \in \mathcal{M}_K} \mathcal{T}_c(g) &= \inf_{g \in \mathcal{M}_K} \left\{ \mathbb{E} \left[\inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi})} \sum_{\ell=1}^{KM} \sum_{k=1}^K \mathcal{P}_{\ell k} c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k(\cdot|\mathbf{x})) \right] \right\} \\ &\leq \mathbb{E} \left[\inf_{\mathbf{P} \in \mathbf{\Pi}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \boldsymbol{\pi}^*)} \sum_{\ell=1}^L \sum_{k=1}^K \mathcal{P}_{\ell k}(g^*, \mathbf{x}) c(\hat{\varphi}_{\ell}(\cdot|\mathbf{x}), \varphi_k^*(\cdot|\mathbf{x})) \right] \\ &= \mathcal{R}_c(g^*), \end{aligned}$$

which completes the proof. \square

A.2 Proof of Proposition 3.3

Continuity. Firstly, we show that, conditional on \mathbf{x} , $\mathcal{I}(g, \mathbf{x})$ is continuous in g . The assumptions A2. and A3. on the continuity and convexity of the cost function c immediately lead to the continuity and convexity in φ of the function $\mathcal{C}(\mathbf{P}, \varphi)$. Moreover, $\mathcal{C}(\mathbf{P}, \varphi)$ as an affine function is obviously continuous and convex in \mathbf{P} . Indeed, because $\mathcal{C}(\mathbf{P}, \varphi)$ is a continuous function, and $\mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)$ is a compact set for all $\mathbf{x} \in \mathcal{X}$, the function defined by $\varphi \mapsto \inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)} \mathcal{C}(\mathbf{P}, \varphi)$ is continuous, moreover, the infimum is attained. By definition, $\mathcal{I}(g, \mathbf{x})$ depends on g only through φ , it follows that $\mathcal{I}(g, \mathbf{x})$ is also continuous in g .

Then, since $\mathcal{I}(g, \mathbf{x})$ is bounded for all g and \mathbf{x} by the assumption, from the Lebesgue's dominated convergence theorem we have that: for all sequence $(g^{(t)})_{t \geq 1}$ such that $g^{(t)} \rightarrow g$,

$$= \lim_{t \rightarrow \infty} \int_{\mathcal{X}} \mathcal{I}(g^{(t)}, \mathbf{x}) d\mu(\mathbf{x}) = \int_{\mathcal{X}} \mathcal{I}(g, \mathbf{x}) d\mu(\mathbf{x}) = \mathcal{R}_c(g),$$

that is $\mathcal{R}_c(g)$ is continuous.

Convexity. We observe that $\mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)$ is a convex set for all $\mathbf{x} \in \mathcal{X}$ (it can be verified easily by definition). It follows that $\mathcal{I}(g, \mathbf{x})$ is convex in g because of the facts that $\mathcal{C}(\mathbf{P}, \varphi)$ is convex and taking infimum over a convex set preserves convexity. Then, let $g_1, g_2 \in \mathcal{M}_K$ and $\lambda_1, \lambda_2 \in [0, 1]$ such that $\lambda_1 + \lambda_2 = 1$, one have

$$\begin{aligned} \mathcal{R}_c(\lambda_1 g_1 + \lambda_2 g_2) &= \int_{\mathcal{X}} \mathcal{I}(\lambda_1 g_1 + \lambda_2 g_2, \mathbf{x}) d\mu(x) \\ &\leq \int_{\mathcal{X}} (\lambda_1 \mathcal{I}(g_1, \mathbf{x}) + \lambda_2 \mathcal{I}(g_2, \mathbf{x})) d\mu(x) \\ &= \lambda_1 \mathcal{R}_c(g_1) + \lambda_2 \mathcal{R}_c(g_2), \end{aligned}$$

which is the definition for the convexity of $\mathcal{R}_c(g)$. □

A.3 Proof of Proposition 3.4

For each $m \in [M]$, the assumption that the local estimator $\hat{\boldsymbol{\theta}}_m$ is consistent implies that its expert parameters $\hat{\boldsymbol{\beta}}_1^{(m)}, \dots, \hat{\boldsymbol{\beta}}_K^{(m)}$ must be ordered and its gating parameter $\hat{\boldsymbol{\alpha}}^{(m)}$ must be initialized as those of $\boldsymbol{\theta}^*$. Therefore, their weighted average $\bar{\boldsymbol{\theta}}^W = \sum_{m=1}^M \lambda_m \hat{\boldsymbol{\theta}}_m$ is also a consistent estimator of $\boldsymbol{\theta}^*$, *i.e.*, we have $\bar{\boldsymbol{\theta}}^W \rightarrow \boldsymbol{\theta}^*$ in probability. It follows that the expected transportation divergence between the two corresponding MoE models, \bar{f}^W and f^* , will tend to zero almost surely as $N \rightarrow \infty$, *i.e.*, we have $\mathcal{T}(\bar{f}^W, f^*) \rightarrow 0$.

On the other hand, by definition, the MoE model corresponding to the reduction estimator $\bar{\boldsymbol{\theta}}^R$ is the one, among all $g \in \mathcal{M}_K$, that minimizes the expected transportation divergence from the weighted average model \bar{f}^W . Therefore, we have

$$\mathcal{T}(\bar{f}^W, \bar{f}^R) \leq \mathcal{T}(\bar{f}^W, f^*) \rightarrow 0 \tag{20}$$

almost surely.

Consequently, we also have that $\mathcal{T}(\bar{f}^R, f^*) \rightarrow 0$ almost surely. Indeed, if it does not hold, the definition (9) implies that there exists a nonzero-measure set $\hat{\mathcal{X}} \subset \mathcal{X}$ on which $\mathcal{T}(\bar{f}^R, f^*, \mathbf{x}) \not\rightarrow 0$. On the other hand, because $\mathcal{T}(\bar{f}^W, f^*) \rightarrow 0$, we must have $\mathcal{T}(\bar{f}^W, f^*, \mathbf{x}) \rightarrow 0$, and so $\mathcal{T}(\bar{f}^W, \bar{f}^R, \mathbf{x}) \rightarrow 0$, almost surely on $\hat{\mathcal{X}}$. Because we are conditional on \mathbf{x} , we can borrow exactly the technique used in the proof of Theorem 9 in Zhang and Chen (2022b) which basically argues that the assumption $\mathcal{T}(\bar{f}^R, f^*, \mathbf{x}) \not\rightarrow 0$ finally leads to a contradiction to the fact that $\mathcal{T}(\bar{f}^W, \bar{f}^R, \mathbf{x}) \rightarrow 0$ almost surely on $\hat{\mathcal{X}}$ (*i.e.*, contradiction to (20)). That means we must have $\mathcal{T}(\bar{f}^R, f^*, \mathbf{x}) \rightarrow 0$ on all nonzero-measure subsets of \mathcal{X} , and hence its expectation, *i.e.*, $\mathcal{T}(\bar{f}^R, f^*) \rightarrow 0$.

As a consequence, $\bar{\beta}_k^R \rightarrow \beta_k^*$ in probability, because any possibility that it does not hold will lead to a contradiction to (20). Indeed, suppose that there exists $k \in [K]$ such that the expert parameter $\bar{\beta}_k^R$ does not tend to β_k^* at some event. Then due to the property of the cost function $c(\cdot, \cdot)$, the value $\mathcal{T}(\bar{f}^R, f^*)$ must be bounded from below by some positive constant, which obviously contradicts the fact (20).

We are left to show that $\bar{\alpha}^R$ also converge to α^* in probability. The estimation of $\bar{\alpha}^R$ via MLE in (25) makes itself a consistent estimator for the parameter of the logistic regression model with predictors $\mathbf{x}_s \in \mathbf{X}_S$ and responses a_{sk} given in (24). Since $\hat{\alpha}^{(m)} \xrightarrow{p} \alpha^*$ and the fact that the gating function is continuous w.r.t the parameter, for all $\mathbf{x}_s \in \mathbf{X}_S$ we have $\pi_k(\mathbf{x}_s; \hat{\alpha}^{(m)}) \xrightarrow{p} \pi_k(\mathbf{x}_s; \alpha^*)$ by the Slutsky theorem. It follows

$$\sum_{m=1}^M \lambda_m \pi_k(\mathbf{x}_s; \hat{\alpha}^{(m)}) \xrightarrow{p} \pi_k(\mathbf{x}_s; \alpha^*), \quad \forall \mathbf{x}_s \in \mathbf{X}_S. \quad (21)$$

Moreover, because at the convergence of the MM algorithm, the obtained experts parameters are supposed to be the theoretical solution $\bar{\beta}_k^R$, the matrix $\mathcal{P}_{\ell k}(g^*, \mathbf{x})$ in (24) can be written by

$$\mathcal{P}_{\ell k}(g^*, \mathbf{x}_s) = \begin{cases} \hat{\pi}_{\ell}(\mathbf{x}_s) & \text{if } k = \arg \inf_{k' \in [K]} c(\hat{\varphi}_{\ell}(\cdot | \mathbf{x}_s), \varphi_{k'}^R(\cdot | \mathbf{x}_s)) \\ 0 & \text{otherwise,} \end{cases} \quad (22)$$

where $\varphi_{k'}^R(\cdot | \mathbf{x}_s)$ denotes $\varphi_{k'}(\cdot | \mathbf{x}_s; \bar{\beta}_k^R)$, the k th expert density of the reduced MoE model. Therefore, by noting the convention about the connection between $\lambda_m \pi_k(\mathbf{x}_s; \hat{\alpha}^{(m)})$ and $\hat{\pi}_{\ell}(\mathbf{x}_s)$, we can see that the responses a_{sk} are in fact the approximations of $\pi_k(\mathbf{x}_s; \alpha^*)$. This means $\bar{\alpha}^R$ is also a consistent estimation of α^* . \square

A.4 Proof of Proposition 4.1

We will use the definition of majorant function to prove that $\mathcal{S}_c(g, g^{(t)})$ given in (16) is a majorant function of $\mathcal{R}_c(g)$ given in (11) at $g^{(t)}$. That is, we will show that $\mathcal{S}_c(g, g^{(t)}) \geq \mathcal{R}_c(g)$

and the equality holds when $g = g^{(t)}$.

First, for all $g \in \mathcal{M}_K$ and $\mathbf{x} \in \mathcal{X}$, from the definition of $\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})$ we see that

$$\sum_{k=1}^K \mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x}) = \hat{\pi}_{\ell}(\mathbf{x}), \quad \forall \ell \in [MK].$$

This means the transportation plan $\mathbf{P} = [\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})]_{\ell k}$ is a member of $\mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)$. Therefore, by the definition of $\mathcal{S}_c(g, g^{(t)})$ and $\mathcal{R}_c(g)$ we have $\mathcal{S}_c(g, g^{(t)}) \geq \mathcal{R}_c(g)$.

Finally, at $g = g^{(t)}$, by definition the functions $\mathcal{S}_c(g, g^{(t)})$ and $\mathcal{R}_c(g)$ are respectively given by

$$\begin{aligned} \mathcal{R}_c(g^{(t)}) &= \mathbb{E} \left[\inf_{\mathbf{P} \in \mathcal{P}_{\mathbf{x}}(\hat{\boldsymbol{\pi}}, \cdot)} \sum_{\ell=1}^{MK} \sum_{k=1}^K P_{\ell k} c \left(\hat{\varphi}_{\ell}(\cdot | \mathbf{x}), \varphi_k^{(t)}(\cdot | \mathbf{x}) \right) \right], \\ \mathcal{S}_c(g^{(t)}, g^{(t)}) &= \mathbb{E} \left[\sum_{\ell=1}^{MK} \sum_{k=1}^K \mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x}) c \left(\hat{\varphi}_{\ell}(\cdot | \mathbf{x}), \varphi_k^{(t)}(\cdot | \mathbf{x}) \right) \right]. \end{aligned} \quad (23)$$

Obviously, the transportation $\mathbf{P} = [\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})]_{\ell k}$ with $\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})$ given in (17) is a minimizer of the minimization in (23). This means that the two quantities above are equal. \square

B Details of updating the gating and expert networks of the DMoE model

B.1 Updating the gating network of the distributed MoE model

As we can see, up to this point all the optimizations carried on the central machine are involving the expectations w.r.t \mathbf{x} . In practice, the distribution of \mathbf{x} is unknown. Therefore, our approach is to create a small sample $\mathcal{D}_S = \{(\mathbf{x}_s, y_s)\}_{s=1}^S$ of size S by randomly drawing from the full dataset \mathcal{D} . This \mathcal{D}_S acts as a supporting sample for the calculations on the central machine. Then, the expectation $\mathbb{E}_{\mu(\mathbf{x})}[\cdot]$ will be understood as the empirical expectation calculated based on the empirical distribution $\mu(\cdot) = S^{-1} \sum_{s=1}^S \delta_{\mathbf{x}_s}(\cdot)$. The size of \mathcal{D}_S is also an issue for further consideration. However, let us assume that \mathcal{D}_S has a sufficiently large size to characterize the distribution of the variable \mathbf{x} . We will denote by \mathbf{X}_S the set of all observations \mathbf{x}_s in the supporting sample \mathcal{D}_S , and where there is no ambiguity we also use \mathbf{X}_S to denote the matrix with each row corresponds to one row vector \mathbf{x}_s^\top .

Although the equation (14) gives a formula to compute the mixing weights for every covariate \mathbf{x} , to obtain the final desired MoE model it is necessary to find the parameter $\bar{\boldsymbol{\alpha}}^R$ for the gating functions $\pi_k(\cdot, \bar{\boldsymbol{\alpha}}^R)$. Therefore, we propose to estimate $\bar{\boldsymbol{\alpha}}^R$ via solving a softmax regression problem in which the predictors are \mathbf{X}_S and the responses are computed

according to (14). In particular, let $\mathbf{A} = [a_{sk}]$ be a matrix in $\mathbb{R}^{S \times K}$ defined by

$$a_{sk} = \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}(g^*, \mathbf{x}_s), \quad (24)$$

where g^* is the model obtained at convergence of the MM algorithm, $\mathbf{x}_s \in \mathcal{D}_S$, and $\mathcal{P}_{\ell k}(g^*, \mathbf{x}_s)$ is computed according to (17). Then $\bar{\boldsymbol{\alpha}}^R$ can be estimated via the maximum likelihood approach

$$\begin{aligned} \bar{\boldsymbol{\alpha}}^R &= \arg \max_{\boldsymbol{\alpha}} \sum_{s=1}^S \sum_{k=1}^K a_{sk} \log \pi_k(\mathbf{x}_s; \boldsymbol{\alpha}) \\ &= \arg \max_{\boldsymbol{\alpha}} \sum_{s=1}^S \left[\sum_{k=1}^{K-1} a_{sk} \boldsymbol{\alpha}_k^\top \mathbf{x}_s - \log \left(1 + \sum_{k'=1}^{K-1} \exp\{\boldsymbol{\alpha}_{k'}^\top \mathbf{x}_s\} \right) \right]. \end{aligned} \quad (25)$$

This problem can be solved efficiently, for example, by the Newton-Raphson procedure. Here, we note that the parameter vector is also initialized, *i.e.*, the K th gate's parameter is constrained to be zero, this is necessary for proving the consistency of $\bar{\boldsymbol{\theta}}^R$.

Now, we are left with the question of how to obtain the update for the expert parameter $\boldsymbol{\beta}_k^{(t+1)}$ as the solution of (18). In the next subsections, we will present the update formulas for the two cases of the Gaussian regression expert and logistic regression expert when the cost function $c(\cdot, \cdot)$ is KL-divergence.

B.2 Updating the Gaussian regression experts parameter

Now, we consider the case of $\varphi(\cdot|\mathbf{x})$ being Gaussian regression experts, *i.e.*, $\varphi(\cdot|\mathbf{x}) = \varphi(\cdot|\mathbf{x}^\top \boldsymbol{\beta}; \sigma^2)$, where $\boldsymbol{\beta} \in \mathbb{R}^{d+1}$ and $\sigma^2 \in \mathbb{R}_+$ are the parameters of the expert. The KL-divergence between two Gaussian regression experts takes the following form

$$\text{KL}(\varphi_1(\cdot|\mathbf{x}) \parallel \varphi_2(\cdot|\mathbf{x})) = \frac{1}{2} \left(\log \frac{\sigma_2^2}{\sigma_1^2} + \frac{\sigma_1^2}{\sigma_2^2} + \frac{(\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)^\top \mathbf{x} \mathbf{x}^\top (\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1)}{\sigma_2^2} - 1 \right).$$

Therefore, with the experts being Gaussian and the cost function $c(\cdot, \cdot)$ being $\text{KL}(\cdot \parallel \cdot)$, the objective function of the problem (18) can be written as a function of $\boldsymbol{\beta}$ and σ^2 as

$$\mathcal{K}(\boldsymbol{\beta}, \sigma^2) := \mathbb{E} \left[\sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}) \frac{1}{2} \left(\log \frac{\sigma^2}{\hat{\sigma}_\ell^2} + \frac{\hat{\sigma}_\ell^2}{\sigma^2} + \frac{(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_\ell)^\top \mathbf{x} \mathbf{x}^\top (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}_\ell)}{\sigma^2} - 1 \right) \right], \quad (26)$$

where $\mathcal{P}_{\ell k}^{(t)}(\mathbf{x})$ denotes $\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})$, which is calculated by (17) with $c(\hat{\varphi}_\ell(\cdot|\mathbf{x}), \varphi_{k'}^{(t)}(\cdot|\mathbf{x}))$ is now replaced by $\text{KL}(\hat{\varphi}_\ell(\cdot|\mathbf{x}) \parallel \varphi_{k'}^{(t)}(\cdot|\mathbf{x}))$. Then, (18) can be simply rewritten by

$$(\boldsymbol{\beta}_k^{(t+1)}, \sigma_k^{2(t+1)}) = \arg \inf_{(\boldsymbol{\beta}, \sigma^2) \in \mathbb{R}^{d+1} \times \mathbb{R}_+} \mathcal{K}(\boldsymbol{\beta}, \sigma^2).$$

The following proposition gives a solution to this problem, *i.e.*, the updating formula for the parameters of the Gaussian regression experts.

Proposition B.1. *Let the notations be as defined above. If $c(\cdot, \cdot)$ is chosen to be $\text{KL}(\cdot \|\cdot)$, and $\varphi_k^{(t+1)}(\cdot | \mathbf{x})$ in (18) are Gaussian regression experts, then the update formulas for their parameters are given by*

$$\boldsymbol{\beta}_k^{(t+1)} = (\mathbf{X}_S^\top \mathbf{D}_k^{(t)} \mathbf{X}_S)^{-1} \mathbf{X}_S^\top \mathbf{X}_S \odot (\mathbf{W}_k^{(t)\top} \widehat{\mathbf{B}}^\top) \mathbf{1}_p, \quad (27a)$$

$$\sigma_k^{2(t+1)} = \frac{1}{\text{trace}(\mathbf{D}_k^{(t)})} \left(\widehat{\boldsymbol{\Sigma}}^\top \mathbf{W}_k^{(t)} \mathbf{1}_S + \mathbf{1}_{MK}^\top \mathbf{W}_k^{(t)} \odot (\boldsymbol{\beta}^{(t+1)} - \widehat{\mathbf{B}})^\top \mathbf{X}_S^\top \mathbf{X}_S (\boldsymbol{\beta}^{(t+1)} - \widehat{\mathbf{B}}) \mathbf{1}_{MK} \right), \quad (27b)$$

in which, the matrices $\mathbf{D}_k^{(t)}$, $\mathbf{W}_k^{(t)}$ and $\widehat{\mathbf{B}}$ are given in (28)-(30).

The detailed calculations are given in Appendix B.3.

B.3 Proof of Proposition B.1

To see this, it is necessary to minimize the function $\mathcal{K}(\boldsymbol{\beta}, \sigma^2)$ given in (26). Firstly, we observe that $\mathcal{K}(\boldsymbol{\beta}, \sigma^2)$ is continuously differentiable and convex w.r.t $\boldsymbol{\beta}$, so any stationary point is a global minimizer of $\mathcal{K}(\boldsymbol{\beta}, \sigma^2)$ as a function of $\boldsymbol{\beta}$. On the other hand, $\mathcal{K}(\boldsymbol{\beta}, \sigma^2)$ is coercive w.r.t σ^2 , then it has a global minimizer that satisfies the first-order optimality condition.

The partial derivatives of $\mathcal{K}(\boldsymbol{\beta}, \sigma^2)$ w.r.t $\boldsymbol{\beta}$ and σ^2 are given by

$$\begin{aligned} \frac{\partial \mathcal{K}}{\partial \boldsymbol{\beta}} &= \mathbb{E} \left[\sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}) \frac{\mathbf{x} \mathbf{x}^\top (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_\ell)}{\sigma^2} \right], \\ \frac{\partial \mathcal{K}}{\partial \sigma^2} &= \mathbb{E} \left[\sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}) \left(\frac{\sigma^2 - \widehat{\sigma}_\ell^2 - (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_\ell)^\top \mathbf{x} \mathbf{x}^\top (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_\ell)}{2\sigma^4} \right) \right]. \end{aligned}$$

Replacing the expectation operator by the empirical expectation and setting $\frac{\partial \mathcal{K}}{\partial \boldsymbol{\beta}}$ to $\mathbf{0}$ one gets

$$\begin{aligned} \frac{1}{S} \sum_{s=1}^S \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_s) \mathbf{x}_s \mathbf{x}_s^\top \boldsymbol{\beta} &= \frac{1}{S} \sum_{s=1}^S \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_s) \mathbf{x}_s \mathbf{x}_s^\top \widehat{\boldsymbol{\beta}}_\ell \\ \Leftrightarrow \sum_{s=1}^S \mathbf{x}_s \left(\sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_s) \right) \mathbf{x}_s^\top \boldsymbol{\beta} &= \sum_{s=1}^S \mathbf{x}_s \mathbf{x}_s^\top \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_s) \widehat{\boldsymbol{\beta}}_\ell. \end{aligned}$$

Let us denote

$$\widehat{\mathbf{B}} = [\widehat{\boldsymbol{\beta}}_1, \dots, \widehat{\boldsymbol{\beta}}_{MK}] \in \mathbb{R}^{p \times MK}, \quad (28)$$

$$\mathbf{D}_k^{(t)} = \text{diag} \left(\sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_1), \dots, \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_S) \right) \in \mathbb{U}^{S \times S}, \quad (29)$$

$$\mathbf{W}_k^{(t)} = \begin{bmatrix} \mathcal{P}_{1k}^{(t)}(\mathbf{x}_1) & \mathcal{P}_{1k}^{(t)}(\mathbf{x}_2) & \dots & \mathcal{P}_{1k}^{(t)}(\mathbf{x}_S) \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{P}_{MK,k}^{(t)}(\mathbf{x}_1) & \mathcal{P}_{MK,k}^{(t)}(\mathbf{x}_2) & \dots & \mathcal{P}_{MK,k}^{(t)}(\mathbf{x}_S) \end{bmatrix} \in \mathbb{U}^{MK \times S}, \quad (30)$$

where $\mathbb{U} = [0, 1] \subset \mathbb{R}$. Then the above equation can be written under matrix form as

$$\mathbf{X}_S^\top \mathbf{D}_k^{(t)} \mathbf{X}_S \boldsymbol{\beta} = \mathbf{X}_S^\top \mathbf{X}_S \odot (\mathbf{W}_k^\top \widehat{\mathbf{B}}^\top) \mathbf{1}_p,$$

where \odot denotes the Hadamard product. This follows the root for $\boldsymbol{\beta}$ is given by

$$\boldsymbol{\beta} = (\mathbf{X}_S^\top \mathbf{D}_k^{(t)} \mathbf{X}_S)^{-1} \mathbf{X}_S^\top \mathbf{X}_S \odot (\mathbf{W}_k^{(t)\top} \widehat{\mathbf{B}}^\top) \mathbf{1}_p.$$

Similarly, setting $\frac{\partial \mathcal{K}}{\partial \sigma^2}$ to 0 one gets

$$\sum_{s=1}^S \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_s) \sigma^2 = \sum_{s=1}^S \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_s) \left(\widehat{\sigma}_\ell^2 + (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_\ell)^\top \mathbf{x} \mathbf{x}^\top (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_\ell) \right),$$

which can be written under matrix form as

$$\text{trace}(\mathbf{D}_k^{(t)}) \sigma^2 = \widehat{\boldsymbol{\Sigma}}^\top \mathbf{W}_k^{(t)} \mathbf{1}_S + \mathbf{1}_{MK}^\top \mathbf{W}_k^{(t)} \odot (\boldsymbol{\beta} - \widehat{\mathbf{B}})^\top \mathbf{X}_S^\top \mathbf{X}_S (\boldsymbol{\beta} - \widehat{\mathbf{B}}) \mathbf{1}_{MK},$$

where $\widehat{\boldsymbol{\Sigma}} = (\widehat{\sigma}_1^2, \dots, \widehat{\sigma}_{MK}^2)^\top \in \mathbb{R}_+^{MK}$. The conclusion is followed immediately. \square

B.4 Updating the logistic regression experts parameters

When the problem under consideration is classification, it is common to use logistic regression model to model the component densities. For simplicity, we consider here the case of binary responses, *i.e.*, $y \in \{0, 1\}$. The conditional density in this case takes the following form

$$\varphi(y|\mathbf{x}; \boldsymbol{\beta}) = \left[\frac{\exp(\mathbf{x}^\top \boldsymbol{\beta})}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta})} \right]^y \left[\frac{1}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta})} \right]^{1-y},$$

where $\boldsymbol{\beta} \in \mathbb{R}^{d+1}$ is the parameter vector. Let $\varphi_1(y|\mathbf{x}; \boldsymbol{\beta}_1)$ and $\varphi_2(y|\mathbf{x}; \boldsymbol{\beta}_2)$ be the conditional densities of two binary logistic regression experts, then the KL-divergence between them is given by

$$\begin{aligned} \text{KL}(\varphi_1(\cdot|\mathbf{x}) \parallel \varphi_2(\cdot|\mathbf{x})) &= \frac{\exp(\mathbf{x}^\top \boldsymbol{\beta}_1)}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta}_1)} \left(\log \frac{\exp(\mathbf{x}^\top \boldsymbol{\beta}_1)}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta}_1)} - \log \frac{\exp(\mathbf{x}^\top \boldsymbol{\beta}_2)}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta}_2)} \right) \\ &\quad + \frac{1}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta}_1)} \left(\log \frac{1}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta}_1)} - \log \frac{1}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta}_2)} \right). \end{aligned} \quad (31)$$

Analogously to the case of Gaussian regression experts, substituting the expression of KL-divergence between two logistic experts into (18), the objective function can be rewritten as a function of $\boldsymbol{\beta}$ by

$$\mathcal{K}_{Bi}(\boldsymbol{\beta}) := \mathbb{E} \left[\sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}) \left(\log(1 + \exp(\mathbf{x}^\top \boldsymbol{\beta})) - \mathbf{x}^\top \boldsymbol{\beta} \frac{\exp(\mathbf{x}^\top \widehat{\boldsymbol{\beta}}_\ell)}{1 + \exp(\mathbf{x}^\top \widehat{\boldsymbol{\beta}}_\ell)} + \text{const} \right) \right], \quad (32)$$

where $\mathcal{P}_{\ell k}^{(t)}(\mathbf{x})$, similarly to the case of Gaussian experts, is calculated by (17) with $c(\widehat{\varphi}_\ell(\cdot|\mathbf{x}), \varphi_{k'}^{(t)}(\cdot|\mathbf{x}))$ is now replaced by $\text{KL}(\widehat{\varphi}_\ell(\cdot|\mathbf{x}) \parallel \varphi_{k'}^{(t)}(\cdot|\mathbf{x}))$ given in (31).

Proposition B.2. *Let the notations be as defined above. If $c(\cdot, \cdot)$ is chosen to be $\text{KL}(\cdot \parallel \cdot)$, and $\varphi_k^{(t+1)}(\cdot|\mathbf{x})$ in (18) are logistic regression experts, then the update formulas for their parameters are given by*

$$\boldsymbol{\beta}_k^{(t+1)} = (\mathbf{X}_S^\top \mathbf{X}_S)^{-1} \mathbf{X}_S^\top (\log(\mathbf{V}_k^{(t)}) - \log(1 - \mathbf{V}_k^{(t)})), \quad \mathbf{V}_k = \mathbf{D}_k^{(t)-1} \mathbf{W}_k^{(t)\top} \mathbf{U}^\top, \quad (33)$$

in which, the matrices $\mathbf{D}_k^{(t)}$ and $\mathbf{W}_k^{(t)}$ the same with those in Proposition B.1, and \mathbf{U} is given in (36).

The detailed calculations can be found in Appendix B.5.

B.5 Proof of Proposition B.2

To see the update formula for the parameter vector of the logistic regression expert, it is necessary to minimize the function $\mathcal{K}_{Bi}(\boldsymbol{\beta})$ given in (32). The first order derivative of $\mathcal{K}_{Bi}(\boldsymbol{\beta})$ is given by

$$\frac{\partial \mathcal{K}}{\partial \boldsymbol{\beta}} = \mathbb{E} \left[\sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}) \left(\frac{\mathbf{x} \exp(\mathbf{x}^\top \boldsymbol{\beta})}{1 + \exp(\mathbf{x}^\top \boldsymbol{\beta})} - \frac{\mathbf{x} \exp(\mathbf{x}^\top \widehat{\boldsymbol{\beta}}_\ell)}{1 + \exp(\mathbf{x}^\top \widehat{\boldsymbol{\beta}}_\ell)} \right) \right].$$

Replacing the expectation operator by the empirical expectation and setting $\frac{\partial \mathcal{K}}{\partial \boldsymbol{\beta}}$ to $\mathbf{0}$ we obtain the following equation

$$\sum_{s=1}^S \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_s) \frac{\mathbf{x}_s \exp(\mathbf{x}_s^\top \boldsymbol{\beta})}{1 + \exp(\mathbf{x}_s^\top \boldsymbol{\beta})} = \sum_{s=1}^S \sum_{\ell=1}^{MK} \mathcal{P}_{\ell k}^{(t)}(\mathbf{x}_s) \frac{\mathbf{x}_s \exp(\mathbf{x}_s^\top \widehat{\boldsymbol{\beta}}_\ell)}{1 + \exp(\mathbf{x}_s^\top \widehat{\boldsymbol{\beta}}_\ell)}. \quad (34)$$

Let us denote

$$\mathbf{E} = \left[\frac{\exp(\mathbf{x}_1^\top \boldsymbol{\beta})}{1 + \exp(\mathbf{x}_1^\top \boldsymbol{\beta})}, \dots, \frac{\exp(\mathbf{x}_S^\top \boldsymbol{\beta})}{1 + \exp(\mathbf{x}_S^\top \boldsymbol{\beta})} \right]^\top \in \mathbb{U}^S, \quad (35)$$

$$\mathbf{U} = \begin{bmatrix} \frac{\exp(\mathbf{x}_1^\top \widehat{\boldsymbol{\beta}}_1)}{1 + \exp(\mathbf{x}_1^\top \widehat{\boldsymbol{\beta}}_1)} & \frac{\exp(\mathbf{x}_1^\top \widehat{\boldsymbol{\beta}}_2)}{1 + \exp(\mathbf{x}_1^\top \widehat{\boldsymbol{\beta}}_2)} & \cdots & \frac{\exp(\mathbf{x}_1^\top \widehat{\boldsymbol{\beta}}_{MK})}{1 + \exp(\mathbf{x}_1^\top \widehat{\boldsymbol{\beta}}_{MK})} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\exp(\mathbf{x}_S^\top \widehat{\boldsymbol{\beta}}_1)}{1 + \exp(\mathbf{x}_S^\top \widehat{\boldsymbol{\beta}}_1)} & \frac{\exp(\mathbf{x}_S^\top \widehat{\boldsymbol{\beta}}_2)}{1 + \exp(\mathbf{x}_S^\top \widehat{\boldsymbol{\beta}}_2)} & \cdots & \frac{\exp(\mathbf{x}_S^\top \widehat{\boldsymbol{\beta}}_{MK})}{1 + \exp(\mathbf{x}_S^\top \widehat{\boldsymbol{\beta}}_{MK})} \end{bmatrix} \in \mathbb{U}^{S \times MK}. \quad (36)$$

Then the equation (34) can be rewritten under matrix form as

$$\mathbf{X}_S^\top \mathbf{D}_k^{(t)} \mathbf{E} \mathbf{1}_S = \mathbf{X}_S^\top \mathbf{W}_k^{(t)\top} \mathbf{U}^\top \mathbf{1}_S.$$

We see that, the above equation holds if $\mathbf{E} = \mathbf{D}_k^{(t)-1} \mathbf{W}_k^{(t)\top} \mathbf{U}^\top$, or equivalently, $\mathbf{X}_S \boldsymbol{\beta} = \log(\mathbf{V}_k^{(t)}) - \log(1 - \mathbf{V}_k^{(t)})$ with $\mathbf{V}_k := \mathbf{D}_k^{(t)-1} \mathbf{W}_k^{(t)\top} \mathbf{U}^\top$. Therefore, we obtain

$$\boldsymbol{\beta} = (\mathbf{X}_S^\top \mathbf{X}_S)^{-1} \mathbf{X}_S^\top (\log(\mathbf{V}_k^{(t)}) - \log(1 - \mathbf{V}_k^{(t)})),$$

which completes the proof. \square

C Pseudocode of the main DMoE algorithm

The Algorithm 1 gives a pseudo-code that summarizes our main algorithm of aggregating MoE models.

Algorithm 1 Aggregating MoE local estimators

Input:

- local gating parameters $\hat{\boldsymbol{\alpha}}^{(m)}$, $m \in [M]$;
- local expert parameters $\hat{\boldsymbol{\beta}}_k^{(m)}$, $m \in [M]$, $k \in [K]$;
- sample proportions λ_m , $m \in [M]$.

Output: Reduction estimator $\bar{\boldsymbol{\theta}}^R = (\bar{\boldsymbol{\alpha}}^R, \bar{\boldsymbol{\beta}}_1^R, \dots, \bar{\boldsymbol{\beta}}_K^R)$.

- 1: Initialize a MoE model $g^{(0)}$, *i.e.*, initialize $\boldsymbol{\theta}^{(0)} = (\boldsymbol{\alpha}^{(0)}, \boldsymbol{\beta}_1^{(0)}, \dots, \boldsymbol{\beta}_K^{(0)})$. Assign $t \leftarrow 0$
 - 2: Sampling a supporting sample \mathcal{D}_S from \mathcal{D}
 - 3: **repeat**
 - 4: **for** $k \in [K]$ **do**
 - 5: **for** $\ell \in [MK]$ **do**
 - 6: For all $\mathbf{x} \in \mathcal{D}_S$, compute $\mathcal{P}_{\ell k}(g^{(t)}, \mathbf{x})$ according to (17)
 - 7: **end for**
 - 8: Update expert parameters by solving (18). Depending on whether the expert is Gaussian
 or logistic regression model, update formula takes the form (27) or (33), respectively
 - 9: **end for**
 - 10: Assign $t \leftarrow t + 1$
 - 11: **until** the objective function (16) converges.
 - 12: Estimate the gating parameter according to (25).
-

D Details of implementation of numerical experiments and additional results

D.1 Data generating process

In this simulation, we fixed $K = 4$ and $d = 20$. We considered datasets of different sizes, from moderate to very large, to illustrate the statistical performance of the estimators. The datasets were generated as follows. Firstly we specified a *true mixture* of K components, which to be estimated, and the centers $\mathbf{x}^{(k)}$ of the clusters $k \in [K]$. Then for each $N \in \{100000, 500000, 1000000\}$, we generated 100 random dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ by drawing, for each k , N/K points from a multivariate Gaussian distribution with mean $\mathbf{x}^{(k)}$ and covariance matrix defined by $\Sigma_{uv} = \frac{1}{4}^{|u-v|}$, for $u, v \in [d]$. Finally, the responses y_i are generated by the standard generating process for MoE model, *i.e.*,

$$\begin{aligned} y_i | Z_i = k, \mathbf{x}_i &\sim \mathcal{N}\left(\mathbf{x}_i^\top \boldsymbol{\beta}_k^*; \sigma_k^2\right) \\ Z_i | \mathbf{x}_i &\sim \mathcal{M}\left(1, (\pi_1(\mathbf{x}_i; \boldsymbol{\alpha}^*), \dots, \pi_K(\mathbf{x}_i; \boldsymbol{\alpha}^*))\right), \end{aligned}$$

where $\mathcal{M}(1, \boldsymbol{\pi})$ is the multinomial distribution with parameter vector $\boldsymbol{\pi}$. Here, the true parameter vector of the model to be estimated is $\boldsymbol{\theta}^* := (\boldsymbol{\alpha}^{*\top}, \boldsymbol{\beta}_1^{*\top}, \dots, \boldsymbol{\beta}_K^{*\top}, \sigma_1^{*2}, \dots, \sigma_K^{*2})^\top$. To have more fairly data generating process, these true parameter vectors and the cluster centers are also selected randomly from the set of integers in the interval $[-5, 5]$. For each of the N 's considered above, we also generated 100 corresponding testing sets such that the ratio of training-testing split is 80% – 20%, more precisely, the testing sets will have 50k, 125k and 500k observations, respectively. It is worth mentioning that each of the datasets with one million observations consumes roundly 1.6 GB of memory. The estimators are computed using training sets and the metrics (except the learning time) are evaluated on testing sets.

D.2 Implementation details and additional results

For the distributed approaches, *i.e.*, the reduction, middle and weighted average estimators, we consider four settings of M , respectively, 4, 16, 64 and 128 machines. The data are distributed equally to the local machines, and the size of the supporting set is taken to be equal to the size of local data, *i.e.*, $S = N_m = N/M$ for all $m \in [M]$.

The Global estimator is obtained by running the EM algorithm, on a single computer. The learning time for each of them is therefore the total time of running the corresponding algorithm. On the other hand, for the reduction estimator, since the local inferences can be performed parallelly on multiple machines, the learning time is recorded as the sum of the

maximum local time and the aggregating time, *i.e.*, the time consumed by the [Algorithm 1](#). Similarly, the learning time of Middle and Weighted average estimator is the sum of the maximum local time and the time of the corresponding calculations to obtain them. All EM algorithms involving the global and local inferences are run with five random initializations. Finally, for each of the obtained estimators, we compute the transportation distance from the true model, the log-likelihood, MSE, RPE and ARI on the corresponding testing sets.




Figure 4: Performance of the Global MoE (G), Reduction (R), Middle (M) and Weighted average (W) estimator for sample size $N = 10^5$ and M machines.

Figure 6 shows the convergence behavior of the MM algorithm in a typical run (picked up randomly from our experiments). The objective function (16) can be seen to converge in about 35 MM iterations.




Figure 5: Performance of the Global MoE (G), Reduction (R), Middle (M) and Weighted average (W) estimator for sample size $N = 500000$ and M machines.

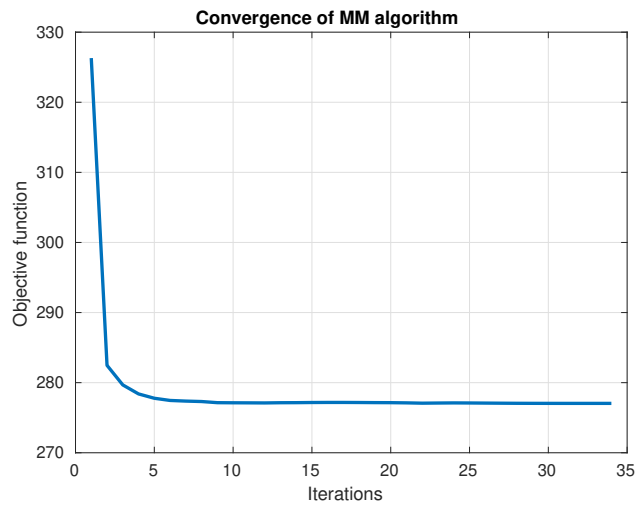


Figure 6: The objective function against iterations to illustrate the decreasing and convergence behavior for a randomly taken dataset.