

**OVERVIEW**

# Practical and theoretical aspects of mixture-of-experts modeling: An overview

Hien D. Nguyen<sup>1</sup>  | Faicel Chamroukhi<sup>2</sup><sup>1</sup>Department of Mathematics and Statistics, La Trobe University, 3086, Bundoora, Victoria, Australia<sup>2</sup>Normandie Universite, UNICAEN, CNRS, Laboratoire de Mathematiques Nicolas Oresme – LMNO, 14000, Caen, France**Correspondence**Hien D. Nguyen, Department of Mathematics and Statistics, La Trobe University, 3086 Bundoora, Victoria, Australia.  
Email: h.nguyen5@latrobe.edu.au

Mixture-of-experts (MoE) models are a powerful paradigm for modeling data arising from complex data generating processes (DGPs). In this article, we demonstrate how different MoE models can be constructed to approximate the underlying DGPs of arbitrary types of data. Due to the probabilistic nature of MoE models, we propose the maximum quasi-likelihood (MQL) approach as a method for estimating MoE model parameters from data, and we provide conditions under which MQL estimators are consistent and asymptotically normal. The blockwise minorization–maximization (blockwise-MM) algorithm framework is proposed as an all-purpose method for constructing algorithms for obtaining MQL estimators. An example derivation of a blockwise-MM algorithm is provided. We then present a method for constructing information criteria for estimating the number of components in MoE models and provide justification for the classic Bayesian information criterion (BIC). We explain how MoE models can be used to conduct classification, clustering, and regression and illustrate these applications via two worked examples.

This article is categorized under:

Algorithmic Development &gt; Statistics

Technologies &gt; Structure Discovery and Clustering

Technologies &gt; Statistical Fundamentals

**KEYWORDS**

classification, clustering, mixture models, mixture of experts, neural networks

**1 | INTRODUCTION**

Let  $\Xi^T = (\mathbf{X}^T, \mathbf{Y}^T) \in \mathbb{X} \times \mathbb{Y}$  be an observed random pair from some data generating process (DGP), where  $\mathbb{X} \subset \mathbb{R}^p$  and  $\mathbb{Y} \subset \mathbb{R}^q$  for  $p, q \in \mathbb{N}$ . Here,  $(\cdot)^T$  is the matrix transposition operator. We call  $\mathbf{X}$  the input variable and  $\mathbf{Y}$  the output (or response) variable. Suppose that the DGP of interest can be approximated as follows.

Firstly, suppose that there is an unobserved random variable  $Z \in [g] = \{1, \dots, g\}$  ( $g \in \mathbb{N}$ ), where the conditional relationship between  $Z$  and the input can be characterized by

$$\mathbb{P}(Z = z | \mathbf{X} = \mathbf{x}) = \text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}), \quad (1)$$

where  $\boldsymbol{\gamma} \in \mathbb{R}^{d_\gamma}$  ( $d_\gamma \in \mathbb{N}$ ) is some parameter vector,  $\text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}) > 0$ , and  $\sum_{z=1}^g \text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}) = 1$ . Secondly, let the conditional relationship between the response and the input, given  $Z = z$ , be characterized by

$$f(\mathbf{y} | \mathbf{X} = \mathbf{x}, Z = z) = \text{Expert}_z(\mathbf{y} | \mathbf{x}; \boldsymbol{\eta}_z), \quad (2)$$

where  $\boldsymbol{\eta}_z \in \mathbb{R}^{d_\eta}$  ( $d_\eta \in \mathbb{N}$ ) is some parameter vector and  $\text{Expert}_z(\mathbf{y} | \mathbf{x}; \boldsymbol{\eta}_z)$  is a probability density function or probability mass function (PDF or PMF; see e.g., DasGupta, 2011, Chs. 2 and 3). Via characterizations (1) and (2), and by using the law of

total probability, we can characterize the marginal relationship between the response and the input, unconditional on  $Z$ , via the expression

$$\begin{aligned} \text{MoE}(\mathbf{y}|\mathbf{x};\boldsymbol{\theta}) &= \sum_{z=1}^g \text{Gate}_z(\mathbf{x};\boldsymbol{\gamma}) \text{Expert}_z(\mathbf{y}|\mathbf{x};\boldsymbol{\eta}_z) \\ &= f(\mathbf{y}|\mathbf{X}=\mathbf{x}), \end{aligned} \tag{3}$$

where  $f(\mathbf{y}|\mathbf{X}=\mathbf{x})$  is a PDF of the response  $\mathbf{y}$  given the input  $\mathbf{X}=\mathbf{x}$ . Here,  $\boldsymbol{\theta}^\top = (\boldsymbol{\gamma}^\top, \boldsymbol{\eta}_1^\top, \dots, \boldsymbol{\eta}_g^\top)$  is the vector of all parameter elements that are required in characterizing (3). We refer to the approximation of the DGP of  $\Xi$  of form (3) as a  $g$ -component mixture-of-experts (MoE) model. The functions  $\text{Gate}_z$  and  $\text{Expert}_z$  are referred to as gating and expert functions, respectively.

MoE models were first studied as neural networks (NNs) by Jacobs, Jordan, Nowlan, and Hinton (1991), where they were used to model complex and heterogeneous DGPs. A schematic diagram of an MoE model as a NN is provided in Figure 1. Some recent reviews of the MoE literature are provided by Yuksel, Wilson, and Gader (2012) and Masoudnia and Ebrahimpour (2014).

MoE models have been broadly applied to numerous areas of business, science, and technology for the tasks of classification, clustering, and regression. A sample of recent applications that were not covered by Yuksel et al. (2012) and Masoudnia and Ebrahimpour (2014) includes: modeling neural connectivity (Bock & Fine, 2014), fusion and segmentation of images (Camplani, del Blanco, Salgado, Jaureguizar, & Garcia, 2014), segmentation of spectral images (Cohen & Le Pennec, 2014), phone activity recognition (Lee & Cho, 2014), climatic change modeling (Nguyen & McLachlan, 2014), parallel mapping of threads in dynamic runtime environments (Emani & O’Boyle, 2015), cardiac stress monitoring via heart sounds (Herzig, Bickel, Eitan, & Intrator, 2015), aerodynamic performance predictions (Liem, Mader, & Martins, 2015), functional magnetic resonance image analysis (Shoenmakers, Guclu, van Gerven, & Heskes, 2015), heterogeneity modeling in neural connectivity data (Eavani et al., 2016), reinforcement learning (He, Boyd-Graber, Kwon, & Daume III, 2016), landmine detection (Yuksel & Gader, 2016), and attention deficit hyperactivity disorder diagnosis (ADHD) (Yaghoobi Karimu & Azadi, 2017).

Since the reviews by Yuksel et al. (2012) and Masoudnia and Ebrahimpour (2014), there have also been numerous theoretical developments regarding the approximation capacity of MoE models (Mendes & Jiang, 2012; Nguyen, Lloyd-Jones, & McLachlan, 2016; Norets & Pelenis, 2014), the performance of maximum quasi-likelihood (MQL) estimation algorithms and the properties of MQL estimators (Nguyen & McLachlan, 2014, 2016), and the modes by which model selection can be conducted within the MoE framework (Baudry, 2015; Montuelle & Le Pennec, 2014).

Regarding the approximation capacity, Mendes and Jiang (2012) demonstrated that MoE models with soft-max gating functions polynomial regression experts are able approximate arbitrarily complex conditional relationships. A similar result is obtain by Norets and Pelenis (2014) who investigated the use of Gaussian gating functions with Gaussian regression experts. Via the Stone-Weierstrass theorem, Nguyen et al. (2016) demonstrated that the mean function an MoE model with soft-max gating function and linear regression expert was able to approximate any arbitrary continuous function on a compact set.

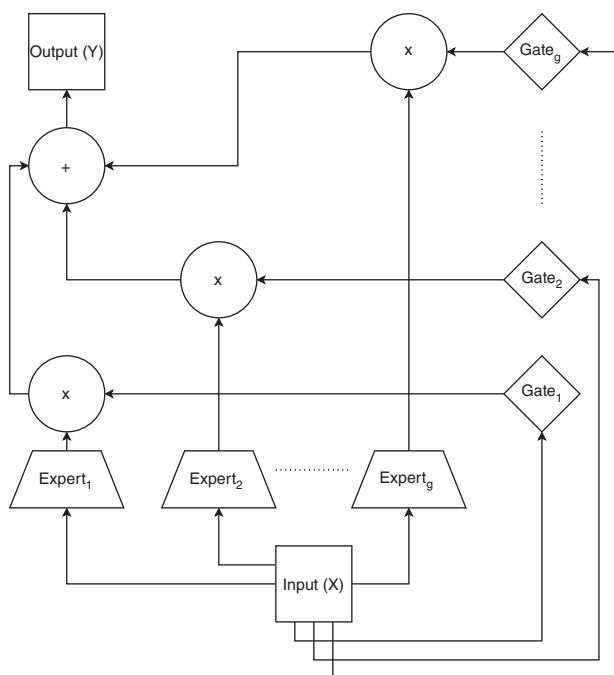


FIGURE 1 Schematic diagram of the NN architecture of a  $g$ -component MoE model as defined by characterizations (1), (2), and (3)

We have concentrated our efforts on the estimation of non-standard MoEs, the theory regarding their approximation capacity, as well as their applications in functional data analysis and signal processing. For example, estimation methodology for non-standard MoE models with Laplace, Student-*t*, and skew-*t* experts were obtained in Nguyen and McLachlan (2016), Chamroukhi (2016), and Chamroukhi (2017), respectively. As mentioned earlier, a power approximation theoretic result is obtained in Nguyen et al. (2016), and numerous applications of MoEs to functional data analysis and signal processing can be found in articles such as Chamroukhi, Same, Govaert, and Aknin (2009), Same, Chamroukhi, Govaert, and Aknin (2011), and Chamroukhi, Glotin, and Same (2013). Throughout our research, we have found that there is a lack of aggregated knowledge regarding the more theoretical and methodological aspects of MoE modeling. As such, we have endeavored to fill the perceived void with this article.

The goal of this article is to provide a concise treatment regarding the practice of constructing MoE models as well as the theoretical justification of such models. As such, the remainder of the article progresses as follows. In Section 2, we discuss the construction of MoE models via the choice of gating and expert functions. In Section 3, we present some of the aforementioned recent theoretical results in a digestible manner and demonstrate their use, where possible. In Section 4, we discuss the problems of classification, clustering, and regression, and show how MoE models can be applied to each of these tasks. In Section 5, we provide some examples of each task. Conclusions are finally drawn in Section 6.

## 2 | MIXTURE-OF-EXPERTS MODELING

We begin by considering the original MoE model of Jacobs et al. (1991) that was designed for the task of multi-speaker vowel discrimination. Here, Jacobs et al. (1991) paired the popular and ubiquitous soft-max gating function

$$\text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}) = \frac{\exp(\alpha_{z0} + \boldsymbol{\alpha}_z^\top \mathbf{x})}{\sum_{\zeta=1}^g \exp(\alpha_{\zeta 0} + \boldsymbol{\alpha}_\zeta^\top \mathbf{x})}, \tag{4}$$

with the multivariate Gaussian distribution expert

$$\text{Expert}_z(\mathbf{y}|\mathbf{x}; \boldsymbol{\eta}_z) = \phi(\mathbf{y}; \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z), \tag{5}$$

where

$$\phi(\mathbf{y}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = |\mathbf{2}\pi\boldsymbol{\Sigma}|^{-1/2} \exp\left[-\frac{1}{2}(\mathbf{y}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{y}-\boldsymbol{\mu})\right],$$

is the multivariate normal density function with mean vector  $\boldsymbol{\mu} \in \mathbb{R}^q$  and positive-definite covariance matrix  $\boldsymbol{\Sigma} \in \mathbb{R}^{q \times q}$ . Here  $\alpha_{z0} \in \mathbb{R}$  and  $\boldsymbol{\alpha}_z^\top = (\alpha_{z1}, \dots, \alpha_{zp}) \in \mathbb{R}^p$  for each  $z \in [g - 1]$ , and  $\alpha_{g0} = 0$  and  $\boldsymbol{\alpha}_g = \mathbf{0}$ , where  $\mathbf{0}$  is the zero vector. We set  $\boldsymbol{\gamma}^\top = (\alpha_{10}, \boldsymbol{\alpha}_1^\top, \dots, \alpha_{g-1,0}, \boldsymbol{\alpha}_{g-1}^\top)$  and  $\boldsymbol{\eta}_z^\top = (\boldsymbol{\mu}_z^\top, \text{vech}^\top \boldsymbol{\Sigma}_z)$ , where  $\text{vech}(\cdot)$  extracts the unique elements of a symmetric matrix (cf. Henderson & Searle, 1979). The MoE of Jacobs et al. (1991) is a direct extension of the usual Gaussian mixture model (GMM; see e.g., McLachlan & Peel, 2000, Ch. 3) that allows for mixing proportions (i.e., the probabilities of  $Z = z$ , for each  $z \in [g]$ ) to depend on the input variable.

In Jordan and Jacobs (1994), the expert of form (5) was extended upon via the multivariate Gaussian regression expert

$$\text{Expert}_z(\mathbf{y}|\mathbf{x}; \boldsymbol{\eta}_z) = \phi(\mathbf{y}; \mathbf{b}_z + \mathbf{B}_z \mathbf{x}, \boldsymbol{\Sigma}_z), \tag{6}$$

where  $\mathbf{b}_z \in \mathbb{R}^q$  and  $\mathbf{B}_z \in \mathbb{R}^{q \times p}$ , for each  $z \in [g]$ . Here,  $\boldsymbol{\eta}_z^\top = (\mathbf{b}_z^\top, \text{vec}^\top \mathbf{B}_z)$ , where  $\text{vec}(\cdot)$  puts all elements of the matrix input into a vector (cf. Henderson & Searle, 1979). Whereas the MoE model with expert (5) can be seen as an extension of the GMM, the MoE with expert (6) is analogously an extension of the multivariate Gaussian mixture regression model of Jones and McLachlan (1992). Robust experts for heterogeneous linear regression models have also been considered by Nguyen and McLachlan (2016), Chamroukhi (2016), and Chamroukhi (2017), where Laplace, Student-*t*, and skew Student-*t* experts are used in place of (6), respectively.

### 2.1 | Mixture of generalized linear experts models

In the same way that the MoE model with experts in the form of (6) can be seen as a heterogeneous linear regression model, heterogeneous versions of other generalized linear models (GLMs; cf. Nelder & Wedderburn, 1972 and McCullagh & Nelder, 1989) can be constructed via MoE modeling. The first of such models was considered by Jordan and Jacobs (1994), where the logistic regression expert

$$\text{Expert}_z(y|\mathbf{x};\boldsymbol{\eta}_z) = \left[ \frac{\exp(\beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x})}{1 + \exp(\beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x})} \right]^y \left[ \frac{1}{1 + \exp(\beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x})} \right]^{1-y}, \quad (7)$$

was proposed for the modeling of binary response data  $Y \in \{0, 1\}$ . Here,  $\beta_{0z} \in \mathbb{R}$ ,  $\boldsymbol{\beta}_z \in \mathbb{R}^p$  and  $\boldsymbol{\eta}_z^\top = (\beta_{0z}, \boldsymbol{\beta}_z^\top)$  for each  $z \in [g]$ .

Another MoE model in this class is the MoE with Poisson regression experts of Grun and Leisch (2008), where

$$\text{Expert}_z(y|\mathbf{x};\boldsymbol{\eta}_z) = \frac{\exp[y(\beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x})]}{y!} \exp[-\exp(\beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x})],$$

which was proposed for modeling of count response data  $Y \in \{0\} \cup \mathbb{N}$ . Here,  $\boldsymbol{\eta}_z$  is the same as that of (7). Further models considered in the literature include the MoE models with gamma experts of Jiang and Tanner (1999a) for modeling positive responses  $Y \in [0, \infty)$  as well as the MoE models with multinomial logistic experts of Chen, Xu, and Chi (1999) for modeling categorical responses  $Y \in [K]$  for some  $K \in \mathbb{N}$ , where  $K > 2$ .

## 2.2 | Gating functions

The majority of MoE models that are applied in practice tend to utilize soft-max gating functions in the form given by (4). In Xu, Jordan, and Hinton (1995), the Gaussian gating functions of form

$$\text{Gate}_z(\mathbf{x};\boldsymbol{\gamma}) = \frac{\pi_z \phi(\mathbf{x}; \mathbf{m}_z, \mathbf{S}_z)}{\sum_{\zeta=1}^g \pi_\zeta \phi(\mathbf{x}; \mathbf{m}_\zeta, \mathbf{S}_\zeta)}, \quad (8)$$

was proposed, where  $\pi_z > 0$  for each  $z \in [g]$ ,  $\sum_{z=1}^g \pi_z = 1$ , and

$$\boldsymbol{\gamma}^\top = \left( \pi_1, \mathbf{m}_1^\top, \text{vech}^\top \mathbf{S}_1, \dots, \pi_g, \mathbf{m}_g^\top, \text{vech}^\top \mathbf{S}_g \right).$$

The Gating functions of form (8) have seen interest use in the literature under the cluster-weighted modeling framework of Ingrassia, Minotti, and Vittadini (2012) and the Gaussian locally-linear mapping framework of Deleforge, Forbes, and Horaud (2015). It can be shown that under some restrictions, there is an equivalence between the class of gating functions of form (4) and (8) (cf. Ingrassia et al., 2012; Norets & Pelenis, 2014).

Although it is possible to utilize any set of functions that meet the restrictions  $\text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}) > 0$ , and  $\sum_{z=1}^g \text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}) = 1$ , there are few alternatives to (4) and (8) that are considered in the literature. Some of these considered alternatives include the exponential family gating functions of Xu et al. (1995) and the Student- $t$  gating functions of Perthame, Forbes, Olivier, and Deleforge (2016).

**Remark 1:** Aside from the simple MoE models that can be characterized via the simple architecture of Figure 1, there are more intricate constructions that are possible for the modeling of complex data. Examples of extensions to the MoE modeling framework include the Mixed-effects MoE models of Ng and McLachlan (2007) and Ng and McLachlan (2014), and the hierarchical MoE models of Jordan and Jacobs (1992) and Jordan and Jacobs (1994) that can be used to fit highly heterogeneous and nonlinear data. We find the hierarchical MoE models to be particularly interesting moving forward as they present a direction for construction of deep generative NNs. Works in this direction include van den Oord and Schrauwen (2014), Theis and Bethge (2015), and Variani, McDermott, and Heigold (2015).

## 3 | THEORETICAL RESULTS

### 3.1 | Approximation theorems

We begin by considering some approximation theory results regarding the most popular class of MoE models: the mixture of linear experts with gates of form (4), and

$$\text{Expert}_z(y|\mathbf{x};\boldsymbol{\eta}_z) = h_z(y; \beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x}),$$

where  $h_z(\cdot; \mu)$  is a PDF with support  $\mathbb{R}$  and mean value  $\mu \in \mathbb{R}$ . Here,  $\boldsymbol{\eta}_z$  is the same as that of (7). Examples of such MoE models include the  $q = 1$  case of the linear experts of Jordan and Jacobs (1994), the Laplace experts of Nguyen and

McLachlan (2016), and the Student- $t$  experts of Chamroukhi (2016). Under characterization (3), the expectation of the response given the input of such MoE models can be written as

$$\begin{aligned} \mathbb{E}(Y|X = \mathbf{x}) &= \sum_{z=1}^g \frac{\exp(\alpha_{z0} + \boldsymbol{\alpha}_z^\top \mathbf{x})}{\sum_{\zeta=1}^g \exp(\alpha_{\zeta 0} + \boldsymbol{\alpha}_\zeta^\top \mathbf{x})} (\beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x}) \\ &= m(\mathbf{x}). \end{aligned} \tag{9}$$

Let  $\mathbb{C}(\mathbb{X})$  be the class of continuous functions and let

$$\mathbb{M}(\mathbb{X}) = \left\{ m(\mathbf{x}) : m \text{ has form (9)}, \boldsymbol{\theta} \in \mathbb{R}^{2g(p+1)-p-1}, g \in \mathbb{N} \right\},$$

be the class of mean functions obtained from the mixture of linear experts models described above, over the domain  $\mathbb{X}$ . The following result from Nguyen et al. (2016) was obtained as a direct consequence of the Stone-Weierstrass theorem (cf. Cotter, 1990).

**Theorem 1:** If  $\mathbb{X} \subset \mathbb{R}^q$  is compact, then the class  $\mathbb{M}(\mathbb{X})$  is dense within the class  $\mathbb{C}(\mathbb{X})$ . That is, for any  $c \in \mathbb{C}(\mathbb{X})$  and  $\epsilon > 0$ , there exists an  $m \in \mathbb{M}(\mathbb{X})$  such that  $\sup_{\mathbf{x} \in \mathbb{X}} |c(\mathbf{x}) - m(\mathbf{x})| < \epsilon$ .

Theorem 1 can be viewed as a universal approximation theorem in the style of the famous result by Cybenko (1989). The theorem states that any continuous function over a compact subset of the Euclidean space can be modeled arbitrarily closely by a mixture of linear experts mean function of form (9). Unfortunately, the theorem does not provide an approximation rate.

Let  $\|\cdot\|_{\mathbb{X},s}$  denote the  $\mathcal{L}_s$  norm over the support  $\mathbb{X}$ , for  $s \in (1, \infty]$ . Define  $\mathbb{W}_s^k(\mathbb{X})$  to be the Sobolev class of continuously differentiable functions in  $\mathcal{L}^s(\mathbb{X})$  (i.e., functions with finite  $\mathcal{L}_s$  norm over the support  $\mathbb{X}$ ) with  $k \in \mathbb{N}$  derivatives, where the sum of the  $\mathcal{L}_s$  norms of the derivatives is bounded. The following result regarding the estimation of functions in Sobolev classes was obtained by Zeevi, Meir, and Maiorov (1998).

**Theorem 2:** Assume that  $\mathbb{X} \subset \mathbb{R}^q$  is compact and define  $\mathbb{M}_g(\mathbb{X})$  to be the subset of  $\mathbb{M}(\mathbb{X})$  where  $g$  is fixed. There exists an absolute positive constant  $c$  such that

$$\sup_{w \in \mathbb{W}_s^k(\mathbb{X})} \inf_{m_g \in \mathbb{M}_g(\mathbb{X})} \|w(\mathbf{x}) - m_g(\mathbf{x})\|_{\mathbb{X},s} \leq \frac{c}{g^{k/q}}.$$

Theorem 2 sacrifices the generality of approximating over all continuous functions as a tradeoff for a uniform approximation rate result. The theorem states that an increase in the number of components in the MoE model increases the accuracy of approximation. However, the rate of increase is itself accelerated by greater differentiability of the target class and decelerated by increasing dimensionality of the support  $\mathbb{X}$ .

We now suppose that  $\boldsymbol{\Xi} \in \mathbb{R}^p \times \mathbb{R}$  is generated from some unknown DGP that can be characterized by a joint density function  $f_0(\boldsymbol{\xi})$  and where the conditional relationship between the response and the input can be characterized by a conditional density function  $f_0(y|X = \mathbf{x})$ . Suppose that we wish to approximate  $f_0(y|X = \mathbf{x})$  by an MoE with Gaussian gating functions of form

$$\text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}) = \frac{\pi_z \phi(\mathbf{x}; \mathbf{m}_z, s_z^2 \mathbf{I})}{\sum_{\zeta=1}^g \pi_\zeta \phi(\mathbf{x}; \mathbf{m}_\zeta, s_\zeta^2 \mathbf{I})}, \tag{10}$$

where  $s_z^2 > 0$  and  $\mathbf{I}$  is the identity function, and experts of form

$$\text{Expert}_z(y|\mathbf{x}; \boldsymbol{\eta}_z) = h(y; \mu_z, \sigma_z), \tag{11}$$

with  $\mu_z \in \mathbb{R}$  and  $\sigma_z$ , for each  $z \in [g]$ . Here, the function  $h(\cdot, \mu, \sigma^2)$  is taken to be of the form

$$h(\cdot; \mu, \sigma^2) = \sigma^{-1} \psi\left(\frac{\cdot - \mu}{\sigma}\right),$$

where  $\psi(y)$  is a probability density function that is a bounded, continuous and symmetric function (about zero), and monotonically decreasing in  $|y|$ . Furthermore, we assume that  $\log h(y; \mu, \sigma)$  is integrable with respect to the  $f_0(\boldsymbol{\xi}) d\boldsymbol{\xi}$ .

Let the class of  $g$ -component MoE models of form (3) with gating functions of form (10) and experts of form (11) over the support  $\mathbb{X}$  be denoted  $\mathbb{M}\mathbb{E}_g(\mathbb{X})$ . Further, denote the Euclidean norm by  $\|\cdot\|$ . The following result is available from Norets and Pelenis (2014).

**Theorem 3:** Assume that (A1)  $\mathbb{X} \subset \mathbb{R}^q$  is compact, (A2)  $f_0(y|\mathbf{X} = \mathbf{x})$  is continuous in  $\boldsymbol{\xi}$ , and (A3) there exists an  $s > 0$  such that

$$\int \log \frac{f_0(y|\mathbf{X} = \mathbf{x})}{\inf_{\|a-y\| \leq s} \inf_{\|b-x\| \leq s} f_0(a|\mathbf{X} = \mathbf{b})} f_0(\boldsymbol{\xi}) d\boldsymbol{\xi} < \infty.$$

If Assumptions (A1)–(A3) are fulfilled then for any  $\epsilon > 0$ , there exists a  $g \in \mathbb{N}$  and an  $\text{MoE}(y|\mathbf{x}; \boldsymbol{\theta}) \in \mathbb{M}\mathbb{E}_g(\mathbb{X})$ , such that

$$\int \log \frac{f_0(y|\mathbf{X} = \mathbf{x})}{\text{MoE}(y|\mathbf{x}; \boldsymbol{\theta})} f_0(\boldsymbol{\xi}) d\boldsymbol{\xi} < \epsilon.$$

Theorem 3 states that the class of MoE models with Gaussian gating and a suitable location-scale experts can densely approximate arbitrary continuous PDFs over compact supports with respect to the conditional Kullback–Leibler divergence (Kullback & Leibler, 1951). This is a powerful result and extends upon well-known denseness theorems regarding approximations of marginal distributions by mixtures of location-scale PDFs (e.g., DasGupta, 2008, Thm. 33.2).

Although Theorem 3 uses Gaussian gating functions, because of the mapping between Gaussian gating and soft-max gating functions, it also applies when the gating functions are of form (4). An alternative denseness result to Theorem 3 is that of Jiang and Tanner (1999a), which is difficult to state but can also be applied to MoE models with GLM experts.

### 3.2 | Maximum quasi-likelihood estimation

Let  $\{\boldsymbol{\Xi}_i\}_{i=1}^n$  be an independent and identically distributed (IID) random sample of  $n \in \mathbb{N}$  observations from some DGP that can be characterized by the PDF  $f_0(\boldsymbol{\xi})$  and where the conditional relationship between each response  $\mathbf{Y}_i$  given input  $\mathbf{X}_i$  ( $i \in [n]$ ) can be characterized by the conditional PDF  $f_0(y|\mathbf{X} = \mathbf{x})$ . Further, let  $\{\boldsymbol{\xi}_i\}_{i=1}^n$  be some fixed observation of  $\{\boldsymbol{\Xi}_i\}_{i=1}^n$ . For some fixed  $g \in \mathbb{N}$ , and without knowledge of either  $f_0(\boldsymbol{\xi})$  and  $f_0(y|\mathbf{X} = \mathbf{x})$ , suppose that we wish to estimate the MoE model of form (3), for some class of gating and expert functions, that best approximates  $f_0(y|\mathbf{X} = \mathbf{x})$ . As proposed by Zeevi et al. (1998) (see also White, 1982), we can do so by obtaining the MQL estimator  $\hat{\boldsymbol{\theta}}_n$ : a local maximizer of the log-quasi-likelihood function

$$Q_n(\boldsymbol{\theta}) = \sum_{i=1}^n \log \text{MoE}(y_i|\mathbf{x}_i; \boldsymbol{\theta}). \quad (12)$$

Often, the task of obtaining a local maximizer of (12) can be difficult. For example, we cannot obtain closed form solutions to the usual first-order condition (FOC) for differentiable functions when the gating functions are of form (4) and the experts are of form (6). As such, iterative or numerical schemes are often employed to conduct maximization. In Nguyen and McLachlan (2014) and Nguyen and McLachlan (2016), the authors considered the blockwise-MM (minorization–maximization) algorithm framework of Lange (2016); see Nguyen, 2017 for a concise tutorial on MM algorithms.

### 3.3 | Minorization–maximization algorithms

The blockwise-MM algorithm framework can be described as follows. Suppose that we have some objective function  $\Omega(\mathbf{u})$ , where

$$\mathbf{u}^\top = (\mathbf{u}_1^\top, \dots, \mathbf{u}_k^\top) \in \mathbb{U} = \prod_{j=1}^k \mathbb{U}_j \subset \prod_{j=1}^k \mathbb{R}^{d_k},$$

for some  $k \in \mathbb{N}$ , where  $\Omega(\mathbf{u})$  is difficult to maximize (e.g., due to lack of closed form FOC or lack of differentiability). Here,  $d_k \in \mathbb{N}$  for each  $j \in [k]$ . Suppose that in each coordinate  $j$ , there exists a function  $M_j(\mathbf{u}_j; \mathbf{v})$  that is easy to manipulate, such that (B1)  $M_j(\mathbf{v}_j; \mathbf{v}) = \Omega(\mathbf{v})$  and (B2)  $M_j(\mathbf{u}_j; \mathbf{v}) \leq \Omega(\mathbf{w})$ , where  $\mathbf{w}^\top = (\mathbf{v}_1^\top, \dots, \mathbf{v}_{j-1}^\top, \mathbf{u}_j^\top, \mathbf{v}_{j+1}^\top, \dots, \mathbf{v}_k^\top)$ , for all  $\mathbf{v}^\top = (\mathbf{v}_1^\top, \dots, \mathbf{v}_k^\top) \in \mathbb{U}$ . We say that  $M_j(\mathbf{u}_j; \mathbf{v})$  is the  $j$ th blockwise minorizer of  $\Omega(\mathbf{v})$  at  $\mathbf{v}$ , or that  $M_j(\mathbf{u}_j; \mathbf{v})$  minorizes  $\Omega(\mathbf{v})$  at  $\mathbf{v}$  with respect to the  $j$ th coordinate.



Construct a blockwise-MM algorithm by firstly initializing it with some value  $\mathbf{u}^{(0)}$ . Next, at the  $r$ th iteration of the algorithm ( $r \in \mathbb{N}$ ), set  $\mathbf{u}_j^{(r)}$  to

$$\mathbf{u}_j^{(r)} = \begin{cases} \arg \max_{\mathbf{u}_j \in \mathbb{U}_j} M_j(\mathbf{u}_j; \mathbf{u}^{(r-1)}) & \text{if } j = (r \bmod k) + 1, \\ \mathbf{u}_j^{(r-1)} & \text{otherwise,} \end{cases} \quad (13)$$

for each  $j \in [k]$ , and then set  $\mathbf{u}^{(r)\top} = (\mathbf{u}_1^{(r)\top}, \dots, \mathbf{u}_k^{(r)\top})$ . From (B1), (B2), and rule (13), we obtain

$$\begin{aligned} \Omega(\mathbf{u}^{(r-1)}) &= M_{(r \bmod k) + 1} \left( \mathbf{u}_{(r \bmod k) + 1}^{(r-1)}; \mathbf{u}^{(r-1)} \right) \\ &\leq M_{(r \bmod k) + 1} \left( \mathbf{u}_{(r \bmod k) + 1}^{(r)}; \mathbf{u}^{(r-1)} \right) \\ &\leq \Omega(\mathbf{u}^{(r)}) \end{aligned} \quad (14)$$

The sequence of inequalities (14) indicates that the sequence of blockwise-MM iterates  $\{\mathbf{u}^{(r)}\}$  generates a sequence of objective evaluates  $\{\Omega(\mathbf{u}^{(r)})\}$  that is monotonically increasing in  $r$ .

Denote directional derivative of  $F(\mathbf{u})$  in the direction of  $\boldsymbol{\delta}$  by

$$\Omega'_{\boldsymbol{\delta}}(\mathbf{u}) = \liminf_{\lambda \downarrow 0} \frac{\Omega(\mathbf{u} + \lambda \boldsymbol{\delta}) - \Omega(\mathbf{u})}{\lambda},$$

and define a stationary point of  $\Omega(\mathbf{u})$  as any point  $\mathbf{u}^*$  such that  $\Omega'_{\boldsymbol{\delta}}(\mathbf{u}^*) \geq 0$  for all  $\boldsymbol{\delta}$  such that  $\mathbf{u}^* + \boldsymbol{\delta} \in \mathbb{U}$ . For all  $j \in [k]$ , make the following assumptions: (C1)  $M_j(\mathbf{u}_j; \mathbf{u}) = \Omega(\mathbf{u})$ , for all  $\mathbf{u} \in \mathbb{U}$ ; (C2)  $M_j(\mathbf{u}_j; \mathbf{v}) \leq \Omega(\mathbf{v})$ , for all  $\mathbf{u}_j \in \mathbb{U}_j$  and  $\mathbf{v} \in \mathbb{U}$ ; (C3)  $M'_{j, \boldsymbol{\delta}_j}(\mathbf{u}_j, \mathbf{v}) \Big|_{\mathbf{u}_j = \mathbf{v}_j} = \Omega'_{\boldsymbol{\delta}}(\mathbf{v})$ , for all  $\boldsymbol{\delta}^\top = (\boldsymbol{\delta}_1^\top, \dots, \boldsymbol{\delta}_j^\top, \dots, \mathbf{0}^\top)$  such that  $\mathbf{u}_j + \boldsymbol{\delta}_j \in \mathbb{U}_j$ ; and (C4)  $M_j(\mathbf{u}_j; \mathbf{v})$  is continuous in  $(\mathbf{u}_j^\top, \mathbf{v}^\top)$ , for each  $j \in [k]$ . Assumptions (C1)–(C4) are validated if each  $M_j(\mathbf{u}_j; \mathbf{v})$  is continuous and differentiable in  $(\mathbf{u}_j^\top, \mathbf{v}^\top)$ , and are blockwise minorizers that fulfill assumptions (B1) and (B2).

Say that a function  $\Omega(\mathbf{u})$  is regular at a point  $\mathbf{v} \in \mathbb{U}$  if we have  $\Omega'_{\boldsymbol{\delta}}(\mathbf{v}) \geq 0$  for all  $\boldsymbol{\delta}^\top = (\boldsymbol{\delta}_1^\top, \boldsymbol{\delta}_2^\top, \dots, \boldsymbol{\delta}_k^\top)$  with  $\Omega'_{\boldsymbol{\delta}_k}(\mathbf{v}) \geq 0$ , where  $\boldsymbol{\delta}_j^{0\top} = (\mathbf{0}^\top, \dots, \boldsymbol{\delta}_j^\top, \dots, \mathbf{0}^\top)$  and  $\mathbf{v}_j + \boldsymbol{\delta}_j \in \mathbb{U}_j$ , for all  $j \in [k]$ . Further say that  $\mathbf{u}^*$  is a coordinate-wise maximum of  $\Omega(\mathbf{u})$  if  $\mathbf{u}^* \in \mathbb{U}$  satisfies  $\Omega(\mathbf{u}^* + \boldsymbol{\delta}_j^0) \leq \Omega(\mathbf{u}^*)$ , for every  $j \in [k]$  and  $\boldsymbol{\delta}_j^0$  such that  $\mathbf{u}^* + \boldsymbol{\delta}_j^0 \in \mathbb{U}$ . Let  $\mathbf{u}^{(\infty)} = \lim_{r \rightarrow \infty} \mathbf{u}^{(r)}$  be the limit point of a blockwise-MM algorithm defined via the update rule (13). The following theorem is available from Razaviyayn, Hong, and Luo (2013).

**Theorem 4:** For all  $j \in [k]$ , assume that  $M_j(\mathbf{u}_j; \mathbf{v})$  is quasi-concave in  $\mathbf{u}_j \in \mathbb{U}_j$ , for fixed  $\mathbf{v} \in \mathbb{U}$ , and fulfills Assumptions (C1)–(C4). Further assume that there is a unique solution to the problem:

$$\arg \max_{\mathbf{u}_j \in \mathbb{U}_j} M_j(\mathbf{u}_j; \mathbf{v}),$$

for each  $j$ , for any  $\mathbf{v} \in \mathbb{U}$ . If  $\mathbf{u}^{(\infty)}$  is a limit point of a blockwise-MM algorithm defined by rule (13), then  $\mathbf{u}^{(\infty)}$  is a coordinate-wise maximum of  $\Omega(\mathbf{u})$ . Furthermore, if  $\mathbf{u}^{(\infty)}$  is regular, then  $\mathbf{u}^{(\infty)}$  is a stationary point of  $\Omega(\mathbf{u})$ .

Theorem 4 states that under some generous conditions, limit points of the blockwise-MM algorithm that are defined by rule (13) are stationary points of the objective. Furthermore, the convergence towards these stationary points is monotonically increasing in nature, with respect to the sequence of objective evaluations.

We furthermore note that global convergence results are also available for generalized blockwise-MM algorithms, where the iterates satisfy the relationships  $M_j(\mathbf{u}^{(r)}; \mathbf{u}^{(r-1)}) \geq M_j(\mathbf{u}^{(r-1)}; \mathbf{u}^{(r-1)})$  (for  $j \in [k]$ ), but where  $\mathbf{u}^{(r)}$  does not satisfy rule (13). In such cases, the algorithm retains its monotonicity with respect to the sequence of objective evaluations. It can further be shown that the limit points globally converge to a fixed-point via results such as that of Meyer (1976). However, the nature of the fixed-points cannot be stated in general and must be established on a case-to-case basis. As Theorem 4 is generally sufficient, we will not engage in further discussions of such results.

### 3.4 | An MM algorithm for the MQL estimation of an example MoE model

We consider the MQL estimation of an MoE model with soft-max gating functions and Gaussian regression experts from

data  $\{\xi_i\}_{i=1}^n$ , where  $\xi_i^\top = (\mathbf{x}_i^\top, y_i) \in \mathbb{R}^p \times \mathbb{R}$ . Using characterization (3), the MoE model can be written as

$$\begin{aligned} \text{MoE}(y|\mathbf{x};\boldsymbol{\theta}) &= \sum_{z=1}^g \text{Gate}_z(\mathbf{x};\boldsymbol{\gamma}) \text{Expert}_z(y|\mathbf{x};\boldsymbol{\eta}_z) \\ &= \sum_{z=1}^g \frac{\exp(\alpha_{z0} + \boldsymbol{\alpha}_z^\top \mathbf{x})}{\sum_{\zeta=1}^g \exp(\alpha_{\zeta 0} + \boldsymbol{\alpha}_\zeta^\top \mathbf{x})} \phi(y; \beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x}, \sigma_z^2), \end{aligned}$$

where  $\beta_{0z} \in \mathbb{R}$ ,  $\boldsymbol{\beta}_z \in \mathbb{R}^p$ ,  $\sigma_z^2 > 0$ , and  $\boldsymbol{\eta}_z^\top = (\beta_{0z}, \boldsymbol{\beta}_z^\top, \sigma_z^2)$ , for each  $z \in [g]$ . The log-quasi-likelihood function can then be written as

$$Q_n(\boldsymbol{\theta}) = \sum_{i=1}^n \log \sum_{z=1}^g \frac{\exp(\alpha_{z0} + \boldsymbol{\alpha}_z^\top \mathbf{x}_i)}{\sum_{\zeta=1}^g \exp(\alpha_{\zeta 0} + \boldsymbol{\alpha}_\zeta^\top \mathbf{x}_i)} \phi(y_i; \beta_{0z} + \boldsymbol{\beta}_z^\top \mathbf{x}_i, \sigma_z^2). \quad (15)$$

We observe that although (15) is smooth in all coordinates of  $\boldsymbol{\theta}$ , it is in the log-sum-exp form and thus a closed form solution to the usual FOC cannot be obtained. We thus turn to constructing a blockwise-MM algorithm for obtaining suitable roots of (15).

Let  $\mathbf{u}$  and  $\mathbf{v}$  be such that  $u_j > 0$  and  $v_j > 0$  for each  $j \in [k]$ . We can minorize  $F(\mathbf{u}) = \log \sum_{j=1}^k u_j$  in all coordinates, simultaneously, by the minorizer

$$M(\mathbf{u}; \mathbf{v}) = \sum_{j=1}^k \frac{v_j}{\sum_{z=1}^k v_z} \log u_j - \sum_{j=1}^k \frac{v_j}{\sum_{z=1}^k v_z} \log \frac{v_j}{\sum_{z=1}^k v_z}, \quad (16)$$

from Zhou and Lange (2010). Applying (16), we obtain the minorizer

$$\begin{aligned} R(\boldsymbol{\theta}; \boldsymbol{\theta}^{(r-1)}) &= \sum_{i=1}^n \sum_{z=1}^g \tau_z(\xi_i; \boldsymbol{\theta}^{(r-1)}) (\alpha_{z0} + \boldsymbol{\alpha}_z^\top \mathbf{x}_i) \\ &\quad - \sum_{i=1}^n \log \sum_{\zeta=1}^g \exp(\alpha_{\zeta 0} + \boldsymbol{\alpha}_\zeta^\top \mathbf{x}_i) \\ &\quad - \frac{1}{2} \sum_{i=1}^n \sum_{z=1}^g \tau_z(\xi_i; \boldsymbol{\theta}^{(r-1)}) \log \sigma_z^2 \\ &\quad - \frac{1}{2} \sum_{i=1}^n \sum_{z=1}^g \tau_z(\xi_i; \boldsymbol{\theta}^{(r-1)}) \frac{(y - \beta_{0z} - \boldsymbol{\beta}_z^\top \mathbf{x}_i)^2}{\sigma_z^2}, \end{aligned} \quad (17)$$

at  $\boldsymbol{\theta}^{(r-1)}$  (in all coordinates), where

$$\tau_z(\xi; \boldsymbol{\theta}) = \frac{\text{Gate}_z(\mathbf{x}_i; \boldsymbol{\gamma}) \text{Expert}_z(y_i | \mathbf{x}_i; \boldsymbol{\eta}_z)}{\text{MoE}(y_i | \mathbf{x}_i; \boldsymbol{\theta})},$$

and

$$C(\boldsymbol{\theta}^{(r-1)}) = -\frac{n}{2} \log 2\pi - \sum_{i=1}^n \sum_{z=1}^g \tau_z(\xi_i; \boldsymbol{\theta}^{(r-1)}) \log \tau_z(\xi_i; \boldsymbol{\theta}^{(r-1)}),$$

is a constant that does not depend on  $\boldsymbol{\theta}$ .

Define  $\tilde{\mathbf{x}}_i^\top = (1, \mathbf{x}_i^\top)$ , for each  $i \in [n]$ , and  $\tilde{\boldsymbol{\alpha}}_z^\top = (\alpha_{z0}, \boldsymbol{\alpha}_z^\top)$  and  $\tilde{\boldsymbol{\beta}}_z^\top = (\beta_{0z}, \boldsymbol{\beta}_z^\top)$ , for each  $z \in [g]$ . We can rewrite (17) as

$$R(\boldsymbol{\theta}; \boldsymbol{\theta}^{(r-1)}) = R_0(\boldsymbol{\theta}; \boldsymbol{\theta}^{(r-1)}) + R_g(\boldsymbol{\theta}; \boldsymbol{\theta}^{(r-1)}) + C(\boldsymbol{\theta}^{(r-1)}),$$

where

$$\begin{aligned} R_0(\boldsymbol{\theta}; \boldsymbol{\theta}^{(r-1)}) &= \sum_{i=1}^n \sum_{z=1}^g \tau_z(\xi_i; \boldsymbol{\theta}^{(r-1)}) \tilde{\boldsymbol{\alpha}}_z^\top \tilde{\mathbf{x}}_i \\ &\quad - \sum_{i=1}^n \log \sum_{\zeta=1}^g \exp(\tilde{\boldsymbol{\alpha}}_\zeta^\top \tilde{\mathbf{x}}_i), \end{aligned}$$



and

$$R_g(\boldsymbol{\theta}; \boldsymbol{\theta}^{(r-1)}) = -\frac{1}{2} \sum_{i=1}^n \sum_{z=1}^g \tau_z(\boldsymbol{\xi}_i; \boldsymbol{\theta}^{(r-1)}) \log \sigma_z^2 - \frac{1}{2} \sum_{i=1}^n \sum_{z=1}^g \tau_z(\boldsymbol{\xi}_i; \boldsymbol{\theta}^{(r-1)}) \frac{(y - \tilde{\boldsymbol{\beta}}_z^\top \tilde{\mathbf{x}}_i)^2}{\sigma_z^2}.$$

Let  $\boldsymbol{\theta}^\top = (\boldsymbol{\theta}_1^\top, \dots, \boldsymbol{\theta}_g^\top)$  be a partitioning of the coordinates of  $\boldsymbol{\theta}$ , where  $\boldsymbol{\theta}_z = \tilde{\boldsymbol{\alpha}}_z^\top$ , for  $z \in [g - 1]$ , and  $\boldsymbol{\theta}_g^\top = (\tilde{\boldsymbol{\beta}}_1^\top, \dots, \tilde{\boldsymbol{\beta}}_g^\top, \sigma_1^2, \dots, \sigma_g^2)$ . For twice differentiable and concave functions  $\Omega(\mathbf{u})$ , Bohning and Lindsay (1988) proposed the minorizer at  $\mathbf{v}$  (in all coordinates)

$$M(\mathbf{u}; \mathbf{v}) = \Omega(\mathbf{v}) + (\mathbf{u} - \mathbf{v})^\top \nabla \Omega(\mathbf{v}) + \frac{1}{2} (\mathbf{u} - \mathbf{v})^\top \mathbf{H} (\mathbf{u} - \mathbf{v}), \tag{18}$$

where  $\mathbf{H} = \text{Hess}(F)(\mathbf{u})$  is negative semi-definite for all  $\mathbf{u} \in \mathbb{U}$  and  $\mathbf{H}$  is positive definite. Here,  $\nabla(\cdot)$  is the gradient operator and  $\text{Hess}(\cdot)$  is the Hessian operator. Let

$$R_z(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r)}) = R_0(\boldsymbol{\vartheta}_z^{(r-1)}; \boldsymbol{\theta}^{(r-1)}),$$

where  $\boldsymbol{\vartheta}^{(r-1)\top} = (\boldsymbol{\theta}_1^{(r-1)\top}, \dots, \boldsymbol{\theta}_{z-1}^{(r-1)\top}, \boldsymbol{\theta}_z^\top, \boldsymbol{\theta}_{z+1}^{(r-1)\top}, \dots, \boldsymbol{\theta}_{g-1}^{(r-1)\top}, \boldsymbol{\theta}_g^{(r-1)\top})$ . Applying (18) to  $R_z(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r-1)})$ , we obtain the coordinate-wise minorizer of (15):

$$S_z(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r-1)}) = R_0(\boldsymbol{\theta}^{(r-1)}; \boldsymbol{\theta}^{(r-1)}) + \nabla R_z(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r-1)}) \Big|_{\boldsymbol{\theta}_z = \boldsymbol{\theta}_z^{(r-1)}} - \frac{1}{8} (\boldsymbol{\theta}_z - \boldsymbol{\theta}_z^{(r-1)})^\top \mathbf{H} (\boldsymbol{\theta}_z - \boldsymbol{\theta}_z^{(r-1)}) + R_g(\boldsymbol{\theta}^{(r-1)}; \boldsymbol{\theta}^{(r-1)}) + C(\boldsymbol{\theta}^{(r-1)}) \tag{19}$$

for each  $z \in [g - 1]$ , where  $\mathbf{H} = \sum_{i=1}^n \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top$  and

$$\nabla R_z(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r-1)}) = \sum_{i=1}^n [\tau_z(\boldsymbol{\xi}_i; \boldsymbol{\theta}^{(r-1)}) - \text{Gate}_z(\mathbf{x}_i; \tilde{\boldsymbol{\gamma}}_z^{(r-1)})] \tilde{\mathbf{x}}_i.$$

This is obtained by noting that

$$\text{Hess}(R_z)(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r-1)}) = - \sum_{i=1}^n \text{Gate}_z(\mathbf{x}_i; \tilde{\boldsymbol{\gamma}}_z^{(r-1)}) [1 - \text{Gate}_z(\mathbf{x}_i; \tilde{\boldsymbol{\gamma}}_z^{(r-1)})] \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top,$$

where  $\tilde{\boldsymbol{\gamma}}_z^{(r-1)\top} = (\boldsymbol{\theta}_1^{(r-1)\top}, \dots, \boldsymbol{\theta}_{z-1}^{(r-1)\top}, \boldsymbol{\theta}_z^\top, \boldsymbol{\theta}_{z+1}^{(r-1)\top}, \dots, \boldsymbol{\theta}_{g-1}^{(r-1)\top})$ , and that  $a(1 - a) \leq 1/4$  for any  $a \in (0, 1)$ .

Notice that (19) is a quadratic and thus is concave and has a unique maximizer with respect to  $\boldsymbol{\theta}_z$ . We can obtain the maximizer by solving the FOC  $\nabla S_z(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r-1)}) = \mathbf{0}$ , which yields the solution

$$\boldsymbol{\theta}_z^{(r)} = 4 \times \mathbf{H}^{-1} \nabla R_z(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r-1)}) \Big|_{\boldsymbol{\theta}_z = \boldsymbol{\theta}_z^{(r-1)}} + \boldsymbol{\theta}_z^{(r-1)}, \tag{20}$$

to the problem

$$\arg \max_{\boldsymbol{\theta}_z} S_z(\boldsymbol{\theta}_z; \boldsymbol{\theta}^{(r-1)}),$$

for each  $z \in [g - 1]$ . Recall that  $\tilde{\boldsymbol{\alpha}}_g = \mathbf{0}$  and thus does not require updating.

Next, a minorizer of (15) in  $\boldsymbol{\theta}_g$  can be obtained by simply holding all other coordinates constant. That is,

$$S_g(\boldsymbol{\theta}_g; \boldsymbol{\theta}^{(r-1)}) = R_0(\boldsymbol{\theta}^{(r-1)}; \boldsymbol{\theta}^{(r-1)}) + R_g(\boldsymbol{\vartheta}_g^{(r-1)}; \boldsymbol{\theta}^{(r-1)}) + C(\boldsymbol{\theta}^{(r-1)}),$$

is a coordinate-wise minorizer of (15) in  $\boldsymbol{\theta}_g$ . The quasi-concavity of  $S_g(\boldsymbol{\theta}_g; \boldsymbol{\theta}^{(r-1)})$  and the solution to its FOC  $\nabla S_g(\boldsymbol{\theta}_g; \boldsymbol{\theta}^{(r-1)}) = \mathbf{0}$  can be obtained via slight modifications to the results of Nguyen and McLachlan (2015). For completeness, the solution  $\boldsymbol{\theta}_z^{(r)}$  containing

$$\tilde{\beta}_z^{(r)} = \left[ \sum_{i=1}^n \tau_z(\xi_i; \theta^{(r-1)}) \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top \right]^{-1} \sum_{i=1}^n \tau_z(\xi_i; \theta^{(r-1)}) y_i \tilde{\mathbf{x}}_i, \quad (21)$$

and

$$\sigma_z^{2(r)} = \frac{\sum_{i=1}^n \tau_z(\xi_i; \theta^{(r-1)}) (y_i - \tilde{\beta}_z^{(r)\top} \tilde{\mathbf{x}}_i)^2}{\sum_{i=1}^n \tau_z(\xi_i; \theta^{(r-1)})}, \quad (22)$$

for each  $z \in [g]$ , uniquely solves the problem

$$\arg \max_{\theta_g} S_g(\theta_z; \theta^{(r-1)}).$$

Together updates (20)–(22) can be applied within rule (13) in order to generate a blockwise-MM algorithm for obtaining the MQL estimator of (15). Since all blockwise solutions are unique and each blockwise minorizers is quasi-concave, we obtain the full conclusion of Theorem 4, as the objective function is smooth and thus regularity is not an issue.

For each initialization  $\theta^{(0)}$ , the blockwise-MM algorithm tends towards a single solution. Unfortunately, like many mixture-type models, the log-quasi-likelihood of MoE models tend to be highly multimodal. As such, numerous initializations should be considered in order to locate a good local maximizer, which can then be considered as candidates for the MQL estimator. One technique for choosing good initializations is that of McLachlan (1988).

We note that although we have derived an algorithm that is entirely within the MM framework, it is possible to replace some of the updates with numerical or alternative optimization schemes that are outside of the MM paradigm. For example, in Ng and McLachlan (2004), a Newton procedure was utilized to update  $\theta_z$  for each  $z \in [g - 1]$ , upon firstly minorizing (15) by (16). As long as the alternative schemes yields solutions that satisfy rule (13) for some notion of blockwise minorization functions, the resulting hybrid blockwise-MM algorithms that are produced will retain the desirable properties bestowed by Theorem 4.

### 3.5 | Asymptotic properties of the maximum quasi-likelihood estimator

We now consider the asymptotic properties of the MQL estimator. The consistency and asymptotic normality of the MQL estimator for the MoE model with Gaussian experts was proved in Zeevi et al. (1998). Further results for GLM experts were obtained in Jiang and Tanner (2000). We shall provide a general scheme for deriving such results for arbitrary MoE models, using the extremum estimation concept of Amemiya (1985).

Let  $\Omega_n(\mathbf{u}) = \Omega(\mathbf{u}; \Xi_1, \dots, \Xi_n)$  be an arbitrary objective function that takes random inputs  $\{\Xi_i\}_{i=1}^n$  and is parameterized by  $\theta$ . Suppose that we wish to obtain the properties of the extremum estimator

$$\hat{\mathbf{u}}_n = \arg \max_{\mathbf{u} \in \mathbb{U}} \Omega_n(\mathbf{u}), \quad (23)$$

for some Euclidean subset  $\mathbb{U} \subset \mathbb{R}^d$  ( $d \in \mathbb{N}$ ). Suppose that there is some  $\mathbf{u}_0$  that naturally connects  $\Omega_n(\mathbf{u})$  to the DGP of  $\{\Xi_i\}_{i=1}^n$ . We say that  $\hat{\mathbf{u}}_n$  is consistent if it converges to  $\mathbf{u}_0$  in probability. The following theorem of Amemiya (1985) provides a simple set of assumptions that can be used to establish the consistency of (23).

**Theorem 5:** Make the following assumptions: (D1) let  $\mathbb{U}$  be open; (D2) let  $\Omega_n(\mathbf{u})$  be measurable in  $\{\Xi_i\}_{i=1}^n$  for all  $\mathbf{u} \in \mathbb{U}$ , and let  $\nabla \Omega_n(\mathbf{u})$  exist and be continuous in an open neighborhood of  $\mathbf{u}_0 \in \mathbb{U}$ ; and (D3) let  $n^{-1} \Omega_n(\mathbf{u})$  converge to a non-stochastic function  $\Omega(\mathbf{u})$  in probability uniformly in  $\mathbf{u}$ , in an open neighborhood of  $\mathbf{u}_0$ , and let  $\Omega(\mathbf{u})$  attain a strict local maximum at a root  $\mathbf{u}_0$ . If (D1)–(D3) are fulfilled and

$$\mathbb{U}_n = \{\mathbf{u} : \nabla \Omega_n(\mathbf{u}) = \mathbf{0} \text{ and } \mathbf{u} \text{ is a strict local maximum}\}, \quad (24)$$

then

$$\lim_{n \rightarrow \infty} \mathbb{P} \left( \inf_{\mathbf{u} \in \mathbb{U}_n} \|\mathbf{u} - \mathbf{u}_0\| > \epsilon \right) = 0,$$

for any  $\epsilon > 0$ . Here,  $\mathbb{U}_n = \{\text{arbitrary element of } \mathbb{U}\}$  when definition (24) results in the empty set.

The conclusion of Theorem 5 is that there exists a consistent root,  $\mathbf{u}_0$ , that is a local maximizer of the objective function  $\Omega_n(\mathbf{u})$ . The result is useful in MoE modeling due to the general lack of universal identifiability of MoE models (cf. Jiang &

Tanner, 1999b), which leads to multiple roots corresponding to the same underlying DGP. Furthermore, the result is also useful due to the highly multimodal nature of MoE log-quasi-likelihood functions. We note that in general, it is not obvious which of many roots is the consistent one that leads to the best approximation of the underlying DGP. Following a suggestion of Amemiya (1985), we can gain some confidence regarding a particular root based on its reasonability from a scientific or contextual perspective, or if it is the limit point of an algorithm when initialized from numerous starting points. A theoretical method for choosing between multiple roots in MQL estimation was proposed by Gan and Jiang (1999). Upon establishing the consistency of a root  $\hat{\mathbf{u}}_n$ , we can then deduce its asymptotic normality via the following theorem of Amemiya (1985).

**Theorem 6:** Make the following assumptions: (E1) let  $\text{Hess}(\Omega_n)(\mathbf{u})$  exist and be continuous in a convex neighborhood of  $\mathbf{u}_0$ ; (E2) let  $n^{-1} \text{Hess}(\Omega_n)(\mathbf{u}_n^*)$  converge to a finite and non-singular matrix

$$\mathbf{I}_1(\mathbf{u}_0) = \lim_{n \rightarrow \infty} \mathbb{E} n^{-1} \text{Hess}(\Omega_n)(\mathbf{u}_0),$$

in probability for any sequence  $\mathbf{u}_n^*$  that converges to  $\mathbf{u}_0$  in probability; (E3) let  $n^{-1/2} \nabla \Omega_n(\mathbf{u})|_{\mathbf{u}=\mathbf{u}_0}$  be asymptotically normal with mean  $\mathbf{0}$  and covariance matrix  $\mathbf{I}_2(\mathbf{u}_0)$ , where

$$\mathbf{I}_2(\mathbf{u}_0) = \lim_{n \rightarrow \infty} \mathbb{E} n^{-1} \nabla \Omega_n(\mathbf{u})|_{\mathbf{u}=\mathbf{u}_0} \nabla^\top \Omega_n(\mathbf{u})|_{\mathbf{u}=\mathbf{u}_0}.$$

If  $\{\hat{\mathbf{u}}_n\}_{n=1}^\infty$  is a sequence that is obtained by choosing an element from  $\mathbb{U}_n$ , for each  $n \in \mathbb{N}$ , such that  $\hat{\mathbf{u}}_n$  converges to  $\mathbf{u}_0$  in probability, then  $n^{-1/2}(\hat{\mathbf{u}}_n - \mathbf{u}_0)$  is asymptotically normal with mean  $\mathbf{0}$  and covariance matrix

$$\mathbf{I}(\mathbf{u}_0) = \mathbf{I}_1^{-1}(\mathbf{u}_0) \mathbf{I}_2(\mathbf{u}_0) \mathbf{I}_1^{-1}(\mathbf{u}_0).$$

Theorem 6 allows for the construction of asymptotic hypothesis tests and confidence intervals regarding the obtained consistent root  $\hat{\mathbf{u}}_n$ . Such tests and intervals can be constructed via results such as those from Hayashi (2000, Sec. 7.4). For the purpose of hypothesis testing, knowledge of  $\mathbf{u}_0$  is assumed. However, when constructing confidence intervals, the DGP is generally unknown. Thus, one must estimate  $\mathbf{I}(\mathbf{u}_0)$  in such constructions. When  $\Omega_n(\mathbf{u}) = \sum_{i=1}^n \omega(\Xi_i; \mathbf{u})$ , a natural estimator for  $\mathbf{I}(\mathbf{u}_0)$  is

$$\hat{\mathbf{I}}_n(\hat{\mathbf{u}}_n) = \hat{\mathbf{I}}_{1,n}^{-1}(\hat{\mathbf{u}}_n) \hat{\mathbf{I}}_{2,n}(\hat{\mathbf{u}}_n) \hat{\mathbf{I}}_{1,n}^{-1}(\hat{\mathbf{u}}_n), \tag{25}$$

where

$$\hat{\mathbf{I}}_{1,n}(\hat{\mathbf{u}}_n) = n^{-1} \sum_{i=1}^n \text{Hess}(\omega)(\Xi_i; \hat{\mathbf{u}}_n),$$

and

$$\hat{\mathbf{I}}_{2,n}(\hat{\mathbf{u}}_n) = n^{-1} \sum_{i=1}^n \nabla \omega(\Xi_i; \mathbf{u})|_{\mathbf{u}=\hat{\mathbf{u}}_n} \nabla^\top \omega(\Xi_i; \mathbf{u})|_{\mathbf{u}=\hat{\mathbf{u}}_n}.$$

Results such as the one of Boos and Stefanski (2013, Thm. 7.3) can be used to establish the validity of (25).

Fix  $g \in \mathbb{N}$  and let  $\mathbb{U}$  be the space of valid values that  $\theta$  can take in the log-quasi-likelihood function 15. Set  $\Omega_n(\theta) = Q_n(\theta)$  and let  $Q_n(\theta) = \sum_{i=1}^n q(\theta; \Xi_i)$ , where  $q(\theta; \xi_i) = \log \text{MoE}(\mathbf{y}_i | \mathbf{x}_i; \theta)$  with gating function and experts as per Section 3.4 (the MoE model has gating functions of form (4) and the experts are of form (6)). For convenience, suppose that  $\{\Xi_i\}_{i=1}^n$  is an IID sample from a DGP with continuous PDF over a compact support. By definition of the MoE model and its parameter vector  $\theta$ , the space of valid values  $\mathbb{U}$  is an open subset of a Euclidean space and hence validates assumption (D1). Since the MoE is constructed from gating functions of form (4) and experts of form (6), it is continuously differentiable, and thus its logarithm is also continuously differentiable. Since the PDF of  $\Xi_i$  is continuous and the support is compact,  $\Omega_n(\theta)$  is also measurable, thus validating (D2). Since the PDF of  $\Xi_i$  is continuous and the support is compact, and since  $\{\Xi_i\}_{i=1}^n$  is an IID sample, it is procedural to validate (D3) for  $\Omega(\theta) = \mathbb{E} q(\Xi_i; \theta)$  via a uniform law of large numbers such as that of Jennrich (1969). We therefore obtain the conclusion of Theorem 5 for the MQL estimator  $\hat{\theta}_n$  of the MoE model from Section 3.4.

Further, we note that the MoE model above also has continuous Hessian for all valid inputs from  $\mathbb{U}$  and thus (E1) is valid. Assumption (E3) is valid because (D3) implies that  $\mathbf{u}_0$  solves  $\nabla \mathbb{E} q(\Xi_i; \theta) = \mathbf{0}$  and because we can swap the gradient and expectation operator since the PDF of  $\Xi_i$  is continuous and the support is compact. An application of the multivariate central limit theorem yields the desired result. Finally, we must make assumption that  $\mathbb{E} \text{Hess}(q)(\Xi_i; \theta_0)$  is non-singular in

order to validate (E2). Such an assumption is standard in the literature and cannot be done away with in general. We therefore have the conclusion of Theorem 6 for the MQL estimator  $\hat{\theta}_n$  of the MoE model from Section 3.4.

### 3.6 | Choice of number of components

Thus far, we have assumed that the number of experts (components)  $g \in \mathbb{N}$  is some known constant. However, in reality, its value in the best MoE approximation of the DGP is unknown, as well as the value of the parameter vector  $\theta$ . In the literature, a popular method for choosing the value of  $g$  among many candidates is to use an information criterion such as Akaike information criterion (AIC; Akaike, 1974), the Bayesian information criterion (BIC; Schwarz, 1978), or the Integrated complete-likelihood information criterion (ICL; Biernacki, Celeux, & Govaert, 2000).

The aforementioned criteria have been implemented in articles such as Grun and Leisch (2007), Grun and Leisch (2008), Chamroukhi et al. (2009), and Nguyen and McLachlan (2016). It is notable that until recently, the only theoretical justification for any of these criteria is for the BIC, which was demonstrated to be consistent in Olteanu and Rynkiewicz (2011). The MoE model results from Olteanu and Rynkiewicz (2011) can be viewed as extensions of the marginal mixture model results from Keribin (2000). Unfortunately, the result of Olteanu and Rynkiewicz (2011) is difficult to state concisely. A more recent approach by Baudry (2015) allows for a much simpler statement of an information criterion consistency theorem. Using the notation from Section 3.5, we paraphrase Baudry (2015, Thm. 8.1) below.

**Theorem 7:** Let  $\{\mathbb{U}_g\}_{g=1}^G$  be a set of parameter spaces, for any  $g \in [G]$ , such that  $\mathbb{U}_g \subset \mathbb{R}^{d_g}$ , where  $d_g \in \mathbb{N}$  for each  $g$ , and  $d_1 \leq \dots \leq d_G$ . Further let

$$\mathbb{U}_0^{[g]} = \{\mathbf{u} : \mathbf{u} = \arg \max_{\mathbf{u} \in \mathbb{U}_g} \mathbb{E} \Omega(\mathbf{u})\},$$

for some  $O_n(\mathbf{u}) = O(\mathbf{u}; \Xi_1, \dots, \Xi_n)$  that is a function of both  $\mathbf{u} \in \mathbb{U}_g$  and  $\{\Xi_i\}_{i=1}^n$ . Make the following assumptions: (F1) let

$$\mathbb{G}_0 = \left\{ g : g = \arg \max_{g \in [G]} \mathbb{E} \Omega(\mathbf{u}_0^{[g]}), \mathbf{u}_0^{[g]} \in \mathbb{U}_0^{[g]} \right\},$$

and assume that  $g = \min \mathbb{G}_0$ ; (F2) for all  $g \in [G]$ ,  $\hat{\mathbf{u}}_n^{[g]} \in \mathbb{U}_n^{[g]}$ , where

$$\mathbb{U}_n^{[g]} = \left\{ \mathbf{u} : \Omega_n(\mathbf{u}) \geq \Omega_n(\mathbf{u}_0^{[g]}), \Omega_n(\mathbf{u}) \rightarrow \mathbb{E} \Omega(\mathbf{u}_0^{[g]}) \text{ in probability} \right\};$$

(F3) for all  $g \in [g]$ , define  $\text{pen}_n(g)$  to be such that  $\text{pen}_n(g) > 0$ ,  $n^{-1} \text{pen}_n(g) \rightarrow 0$  in probability, as  $n \rightarrow \infty$ , and

$$\text{pen}_n(g) - \text{pen}_n(g^*) \rightarrow \infty,$$

in probability, as  $n \rightarrow \infty$ , when  $g > g^*$ ; (F4)

$$\Omega_n(\hat{\mathbf{u}}_n^{[g_0]}) - \Omega_n(\hat{\mathbf{u}}_n^{[g]}) \rightarrow C,$$

in probability, where  $C$  is a constant, for any  $g \in \mathbb{G}_0$ . If (F1)–(F4) are fulfilled and if selection of  $g$  is based upon the generic information criterion:

$$\hat{g}_n = \min \left\{ g : g = \arg \max_{g \in [G]} \left[ \Omega_n(\hat{\mathbf{u}}_n^{[g]}) - \text{pen}_n(g) \right] \right\}, \quad (26)$$

then  $\mathbb{P}(\hat{g}_n \neq g_0) \rightarrow 0$  as  $n \rightarrow \infty$ .

Twice the negative of  $\Omega_n(\hat{\mathbf{u}}_n^{[g]}) - \text{pen}_n(g)$ , in (26), is often referred to as the information criterion. Although tedious, the assumptions of Theorem 7 are generally valid. Assumption (F1) states that we are searching for a parsimonious model, and (F2) is valid if the hypothesis of Theorem 5 are valid for each  $g \in [G]$ . Assumption (F3) states that the constructed information criterion must involve a penalty (at the discretion of the investigator) that becomes smaller as more observations are observed and that is capable of ordering different complexities of models. Assumption (F4) is difficult to rationalize, although it can be validated by application of Baudry (2015, Cor. 8.2).

Consider the penalty

$$\text{pen}_n(g) = \dim(\theta) \log n, \quad (27)$$

where  $\dim(\cdot)$  computes the dimension (i.e., number of elements) of a vector. This is the BIC penalty function. For the soft-max gated mixture of Gaussian regression experts, we can show that

$$\dim(\boldsymbol{\theta}) = (3 + 2p)g - p - 1.$$

We can validate (F3) by noting that  $\lim_{n \rightarrow \infty} n^{-1} \log n = 0$  and that  $\log n$  is strictly increasing. Using Baudry (2015, Cor. 8.2) to validate (F4), and the assumptions made in the example from Section 3.5 to validate (F2), we can show that the BIC (i.e., rule (26) with penalty (27)) consistently selects the most parsimonious MoE model with gating functions of form (4) and experts of form (6), with respect to number of components  $g$ .

## 4 | APPLICATIONS OF MIXTURE-OF-EXPERTS MODELS

The three primary applications of a statistical model are to conduct classification, clustering, and regression. Here, we refer to classification as the conduct of supervised learning using labeled data  $\{\boldsymbol{\Xi}_i\}_{i=1}^n$ , where  $Y \in [K]$ , and then assessing the success of the learnt discriminant model via some additional unlabeled data. The problem of clustering is take some unlabeled data  $\{\mathbf{X}_i\}_{i=1}^n$  set and to sort the observations into similarity groups, where the individual observations in each group are close in some sense, and the individuals between the groups are sufficiently different. Lastly, the problem of regression is to produce a feasible conditional probabilistic relationship between the data  $Y_i$  and  $\mathbf{X}_i$  ( $i \in [n]$ ), given some observed sample  $\{\boldsymbol{\Xi}_i\}_{i=1}^n$ . There are numerous examples of applications of MoEs to all three problems in the literature. We refer the reader to Yuksel et al. (2012) for an bibliographical reference.

### 4.1 | Classification

We can conduct classification via MoE modeling by using an MoE model with multinomial logistic experts. That is, suppose that we observe data  $\{\boldsymbol{\Xi}_i\}_{i=1}^n$ , where  $\boldsymbol{\Xi}_i^\top = (\mathbf{X}_i^\top, Y_i) \in \mathbb{X} \times [K]$ , for some  $K \in \mathbb{N}$ . Conditioned on  $\mathbf{X}_i = \mathbf{x}_i$ , suppose that the PMF of  $Y_i$  can be best approximated by an MoE of form (3) with multinomial logistic expert functions of form

$$\text{Expert}_z(y|\mathbf{x}, \boldsymbol{\eta}_z) = \prod_{l=1}^K \left[ \frac{\exp(\beta_{z l 0} + \boldsymbol{\beta}_{z l}^\top \mathbf{x})}{\sum_{\ell=1}^K \exp(\beta_{z \ell 0} + \boldsymbol{\beta}_{z \ell}^\top \mathbf{x})} \right]^{\mathbb{I}(y=l)}, \quad (28)$$

where  $\beta_{z l 0} \in \mathbb{R}$  and  $\boldsymbol{\beta}_{z l} \in \mathbb{R}^p$  for each  $z \in [g]$  and  $l \in [K - 1]$ , and  $\beta_{z K 0}$  and  $\boldsymbol{\beta}_{z K} = \mathbf{0}$  for each  $z \in [g]$ . We set

$$\boldsymbol{\eta}_z^\top = (\beta_{z 1 0}, \boldsymbol{\beta}_{z 1}^\top, \dots, \beta_{z, K-1, 0}, \boldsymbol{\beta}_{z, K-1}^\top),$$

for each  $z$ . Here,  $\mathbb{I}(A)$  is the indicator function that takes value 1 if proposition  $A$  is true and 0 otherwise.

Respectively, let  $g_0$  and  $\boldsymbol{\theta}_0$  be the number of components and parameter vector that best approximates the DGP of interest. The MoE model of form (3), with experts of form (28), has the probabilistic interpretation

$$\mathbb{P}(Y = y | \mathbf{X} = \mathbf{x}) = \text{MoE}(y | \mathbf{x}; \boldsymbol{\theta}_0). \quad (29)$$

Let  $\mathbf{d}^\top = (\mathbf{x}^\top, y)$  be an arbitrary data point that is generated via the same DGP as that of interest, and suppose that we only have knowledge of  $\mathbf{x}$  and wish to estimate  $y$ . Using interpretation (29), we can obtain the MAP (maximum a posteriori probability) rule for classification:

$$\begin{aligned} \hat{y} &= \arg \max_{y \in [K]} \text{MoE}(y | \mathbf{x}; \boldsymbol{\theta}_0) \\ &= \arg \max_{y \in [K]} \sum_{z=1}^{g_0} \text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}_0) \text{Expert}_z(y | \mathbf{x}; \boldsymbol{\eta}_{0,z}). \end{aligned}$$

If  $g_0$  and  $\boldsymbol{\theta}_0$  are unknown, then we can estimate these quantities by  $\hat{g}_n$  and  $\hat{\boldsymbol{\theta}}_n$ , respectively, in order to obtain the plugin-MAP rule

$$\begin{aligned} \hat{y} &= \arg \max_{y \in [K]} \text{MoE}(y | \mathbf{x}; \hat{\boldsymbol{\theta}}_n) \\ &= \arg \max_{y \in [K]} \sum_{z=1}^{\hat{g}_n} \text{Gate}_z(\mathbf{x}; \hat{\boldsymbol{\gamma}}_n) \text{Expert}_z(y | \mathbf{x}; \hat{\boldsymbol{\eta}}_{n,z}). \end{aligned} \quad (30)$$

Decision theoretic justification for the MAP and plugin-MAP rules for classification are detailed in McLachlan (1992).

## 4.2 | Clustering

When conducting clustering, we assume that our data  $\{\Xi_i\}_{i=1}^n$  arises from a DGP that can be best characterized by an MoE model, which is defined via the hierarchical construction that is characterized by Equations (1)–(3). That is, we assume that each  $\Xi_i$  has a latent label  $Z_i \in [g]$  that determines which of the  $g$  experts of form (2) that it was generated from.

Let  $\mathbf{d}$  be an arbitrary data point that is generated via a DGP that is best characterized by the  $g_0$ -component MoE (3) with some parameter vector  $\theta_0$ . Suppose that we wish to estimate  $Z$ , the expert from which  $\mathbf{d}$  was generated. As in Section 4.1, we can utilize a pair of MAP rule, depending on the nature of the clustering problem. The MAP rules for clustering are to estimate  $Z$  by

$$\hat{z} = \arg \max_{z \in [g]} \frac{\text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}) \text{Expert}_z(\mathbf{y} | \mathbf{x}; \boldsymbol{\eta}_z)}{\text{MoE}(\mathbf{y} | \mathbf{x}; \boldsymbol{\theta})},$$

which is equivalent to estimating  $Z$  by

$$\hat{z} = \arg \max_{z \in [g]} \mathbb{P}(Z = z | \mathbf{X}_i = \mathbf{x}_i, \mathbf{Y}_i = \mathbf{y}_i),$$

or by

$$\hat{z} = \arg \max_{z \in [g]} \text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}),$$

which is equivalent to

$$\hat{z} = \arg \max_{z \in [g]} \mathbb{P}(Z = z | \mathbf{X}_i = \mathbf{x}_i).$$

The estimation of the number of clusters experts, or clusters in such an application, can be conducted via the technique of Section 3.6. Furthermore, in general, we do not know the parameter vector  $\theta_0$  and thus we must estimate by the MQL estimator  $\hat{\theta}_n$  from a realization of  $\{\Xi_i\}_{i=1}^n$ . Using the estimated number of clusters (components)  $\hat{g}_n$  and the MQL estimator for the MoE model with  $\hat{g}_n$  clusters  $\hat{\theta}_n$ , we obtain the plugin-MAP rules for clustering:

$$\hat{z} = \arg \max_{z \in [\hat{g}_n]} \frac{\text{Gate}_z(\mathbf{x}; \hat{\boldsymbol{\gamma}}_n) \text{Expert}_z(\mathbf{y} | \mathbf{x}; \hat{\boldsymbol{\eta}}_{n,z})}{\text{MoE}(\mathbf{y} | \mathbf{x}; \hat{\boldsymbol{\theta}}_n)}, \quad (31)$$

and

$$\hat{z} = \arg \max_{z \in [\hat{g}_n]} \text{Gate}_z(\mathbf{x}; \hat{\boldsymbol{\gamma}}_n). \quad (32)$$

## 4.3 | Regression

Suppose now that we do not care that the DGP can be characterized via the hierarchical construction of Equations (1)–(3), but only that the conditional relationship between the response  $\mathbf{Y}$  given input  $\mathbf{X} = \mathbf{x}$  has form (3), for some parameter vector  $\theta_0$ . We are often interested in using form (3) in order to estimate some functionals of the DGP of  $\mathbf{d}$  such as the mean  $\mathbb{E}(\mathbf{Y} | \mathbf{X} = \mathbf{x})$  or higher moments. For example, in the case where the gating function is of form (4) and the expert arises from a location-scale family of conditional density functions, we can write  $\mathbb{E}(\mathbf{Y} | \mathbf{X} = \mathbf{x})$  in form (9).

Since the form of the MoE that best approximate the DGP of  $\mathbf{d}$  is often unknown, we must estimate it via the realization of some random sample  $\{\Xi_i\}_{i=1}^n$ . Often both the number of components  $g_0$  and the parameter vector  $\theta_0$  require estimation. As in Sections 4.1 and 4.2,  $g_0$  can be estimated via  $\hat{g}_n$ , obtained via the technique from Section 3.6, and  $\theta_0$  can be estimated by the MQL estimator  $\hat{\theta}_n$ .

Upon obtaining the estimators above, conditional moment functions can be easily computed. For example, if  $\mathbb{E}[H(\mathbf{Y}) | \mathbf{X} = \mathbf{x}]$  exists, for some real-valued function  $H(\mathbf{y})$  of  $\mathbf{y} \in \mathbb{Y}$ , then

$$\mathbb{E}[H(\mathbf{Y}) | \mathbf{X} = \mathbf{x}] = \sum_{z=1}^{g_0} \text{Gate}_z(\mathbf{x}; \boldsymbol{\gamma}_0) \int_{\mathbb{Y}} H(\mathbf{y}) \text{Expert}_z(\mathbf{y} | \mathbf{x}; \boldsymbol{\eta}_{0,z}) \mathbf{d}\mathbf{y}, \quad (33)$$

and can be estimated by

$$\hat{\mathbb{E}}[H(\mathbf{Y}) | \mathbf{X} = \mathbf{x}] = \sum_{z=1}^{\hat{g}_n} \text{Gate}_z(\mathbf{x}; \hat{\boldsymbol{\gamma}}_n) \int_{\mathbb{Y}} H(\mathbf{y}) \text{Expert}_z(\mathbf{y} | \mathbf{x}; \hat{\boldsymbol{\eta}}_{n,z}) \mathbf{d}\mathbf{y}.$$



We can obtain Equation (9) by making the substitution  $H(Y) = Y$  into (33) for an MoE with soft-max gating functions and location-scale experts. If the conditional variance function of each expert is constant and equal to  $\sigma_z^2$  for each  $z \in [g]$ , we can write the variance function of the response  $Y$ , given the input  $\mathbf{X} = \mathbf{x}$ , as

$$\text{var}(Y|\mathbf{X} = \mathbf{x}) = \sum_{z=1}^{g_0} \text{Gate}_z(\mathbf{x}; \gamma_0) \left[ (\beta_{z,0} + \beta_z^\top \mathbf{x})^2 + \sigma_z^2 \right] - [\mathbb{E}(Y|\mathbf{X} = \mathbf{x})]^2.$$

Furthermore, we can estimate  $\text{var}(Y|\mathbf{X} = \mathbf{x})$  by

$$\widehat{\text{var}}(Y|\mathbf{X} = \mathbf{x}) = \sum_{z=1}^{\hat{g}_n} \text{Gate}_z(\mathbf{x}; \hat{\gamma}_n) \left[ (\hat{\beta}_{z,0} + \hat{\beta}_z^\top \mathbf{x})^2 + \hat{\sigma}_z^2 \right] - [\hat{\mathbb{E}}(Y|\mathbf{X} = \mathbf{x})]^2.$$

Confidence intervals can also be constructed about any regression function. For example, see Nguyen and McLachlan (2014) regarding the construction of asymptotic confidence intervals around the mean function of an MoE model using the asymptotic normality conclusion of Theorem 6.

## 5 | CASE STUDIES

To demonstrate the applications of MoE models that are described in Section 4, we present the two following examples. We note that all computation for both examples are performed within the **R** programming environment (R Core Team, 2016) via the **flexmix** package of Grun and Leisch (2008), which allows for estimation of generic MoE models with soft-max gating functions. The optimization procedures utilized in **flexmix** are hybrid MM algorithms for MQL estimation and are described in Grun and Leisch (2008).

### 5.1 | Three-class problem

We generate data  $\{\xi_i\}_{i=1}^n$  from the three-class problem of Chen et al. (1999) and Ng and McLachlan (2004). In the three-class problem, each data point  $\mathbf{d}_i^\top = (\mathbf{x}_i^\top, y_i)$  consists of the input  $\mathbf{x}_i = (x_{i1}, x_{i2})$ , which is a realization of a random variable  $\mathbf{X}_i^\top = (X_{i1}, X_{i2})$ , where  $X_{ij}$  is uniformly distributed over the interval  $[-5, 5]$ , for each  $i \in [n]$  and  $j \in \{1, 2\}$ . Depending on  $\mathbf{X}_i = \mathbf{x}_i$ , the value of the response is a categorical variable  $Y_i \in [3]$ , such that  $Y_i = 1$  by default, unless  $\mathbf{X}_i$  is within a ball of radius two around the origin, in which case  $Y_i = 2$ , or if  $\mathbf{X}_i$  is within the square with corners  $(-4, 4)$  and  $(-2, 2)$  or the square with corners  $(2, 2)$  and  $(4, 4)$ , in which case  $Y_i = 3$ . A visualization of a realization of an  $n = 1000$  observations sample from the three-class problem is provided in Figure 2. From the sample, we obtain 759, 156, and 85 observations, with category labels  $y_i = 1, 2, 3$ , respectively.

In order to construct a classifier, based on the sample of  $n = 1000$  observations from Figure 2, we estimate soft-max gated MoE models with multinomial logistic experts for  $K = 3$  classes, as per Section 4.1, with varying numbers of components  $g \in [9]$ . The BIC values (i.e., twice the negative of the log-quasi-likelihood subtract the penalty of form (27)) for each





FIGURE 2 A realization of an  $n = 1000$  observations sample from the three-class problem. The plot symbols indicate the class label



**FIGURE 3** BIC values for  $g \in [9]$  obtained from the MQL estimators on the  $n = 1000$  observations sample from Figure 2. The filled marker indicates the best model obtained via rule (26)


$g$ , obtained via MQL estimation, are plotted in Figure 3. From the figure, we observe that the best MoE model is that with  $\hat{g}_n = 4$  components. Applying rule (30), we obtain a classification accuracy of 91.4% using the fitted MoE model classifier.

In order to better assess the performance of the MoE model classifier, defined by rule 30, we generate a new sample of  $n = 2500$  observations from the sample process as described above. Using the fitted classifier, we obtain a test set accuracy rate (on the new sample) of 90.1%. A plot of the new sample along with the classifications via the MoE model classifier is displayed in Figure 4.

From Figure 4, we can obtain a visualization of the decision boundaries upon which the classification rule (30) assigns new point. We can see that the decision boundaries are not perfect fits to the rigid shapes of the true label boundaries of the DGP. However, upon inspection of Figure 2, we see that the classifier models the dense regions of each class very well. This is especially apparent when inspecting the dense regions of observations with  $y_i = 3$  from the original sample. Also, where the classifier decision boundary exceeds the circle that determines when  $y_i = 2$ , we note that there are fewer observations and thus the lack of fit in that region of the domain is explainable. Considering that the percentage of observations in the training are 79.4%, 11.7%, and 8.9%, for  $y_i = \{1, 2, 3\}$ , respectively, the classification rate of 90.1% is a good result.

## 5.2 | Switch operation power signals

In this example, we analyze a time series data set arising from electrical signals at a switching point on the French railway, under a switching operation. The data were originally studied in Chamroukhi et al. (2009), Chamroukhi, Same, Govaert, and Aknin (2010), and Same et al. (2011). An instance of such a signal, over a period of approximately 6 seconds, is presented in Figure 5. The signals are measured at  $n = 550$  equally-spaced time points  $x_i$  ( $i \in [n]$ ) that are normalized to be in the unit interval. We let  $y_i$  be the measurement of the power at each time point  $x_i$ , in Watts. Together  $\{\xi_i\}_{i=1}^n$  forms our sample of interest, where  $\xi_i^\top = (x_i, y_i)$ . We wish to model the power signals as a function of time.



**FIGURE 4** Plot of the additional  $n = 2500$  observations sample from the three-class problem. The plot symbols indicate the true class labels  $y_i$ ,  $i \in [n]$ . The color indicates the classification via the fitted MoE classifier,  $\hat{y}_i$ . Here, blue, green, and red correspond to  $\hat{y}_i = 1, 2, 3$ , respectively


Upon observation, it is clear that the time series in Figure 5 is highly nonlinear. Following the analysis by Chamroukhi et al. (2009), we model the DGP for the data via an MoE model with soft-max gating functions and quadratic-mean Gaussian regression experts of form

$$\text{Expert}_z(y|x; \boldsymbol{\eta}_z) = \phi(y; \beta_{0z} + \beta_{1z}x + \beta_{2z}x^2, \sigma_z^2),$$


where  $\boldsymbol{\eta}_z^T = (\beta_{0z}, \beta_{1z}, \beta_{2z}, \sigma_z^2)$ , for  $z \in [g]$ . Figure 6 displays the BIC for each  $g \in [10]$ , obtained via MQL estimation. Using the optimal model with  $\hat{g}_n = 8$ , we can obtain the MoE model that best approximates the DGP for the data, and the expectation curve for the MoE model, of form (9), is plotted in Figure 7. We observe that that the curve is a good fit for the data and models its primary features, without being so specific as to model its idiosyncrasies.

As the data arises from an electrical control of a Railway switching point, it undergoes multiple stages of control. Each of the  $\hat{g}_n = 8$  components from the obtained MoE model can be seen as one of these stages or sub-stages. We can utilize the clustering rule (32) in order to assign each time point to one of these stages. Figure 8 displays the segments of the time series that are assigned to each of the  $\hat{g}_n = 8$  components, along with the mean curve corresponding to the respective component. We utilize clustering rule (32) instead of rule (31) as the different stages of control under a signal switching are entirely


**FIGURE 5** Instance of an electrical signal at a switching point, undergoing a switching operation. The abscissa displays the time at which the signal is measured (normalized to the unit interval) and the ordinate displays the value of the signal, in Watts




**FIGURE 6** BIC values for  $g \in [9]$  obtained from the MQL estimators on the  $n = 550$  observations sample from the time series that is displayed in Figure 5. The filled marker indicates the best model obtained via rule (26)



**FIGURE 7** The original signal is plotted as a solid curve and the fitted mean function for the  $\hat{g}_n = 8$  component MoE model, of form (9), is plotted as a dotted curve. The abscissa displays the time at which the signal is measured (normalized to the unit interval) and the ordinate displays the value of the signal, in Watts





**FIGURE 8** Each of the segments is colored in one of  $\hat{g}_n = 8$  colors, and are plotted as solid curves. The dotted curves visualize the mean function of the corresponding MoE model component that the segment is clustered to. The abscissa displays the time at which the signal is measured (normalized to the unit interval) and the ordinate displays the value of the signal, in Watts

time-based. Here, the modeling of the power curves is incidental in identifying the different stages of control, in time. Rule (31) is more appropriate when clustering data that arise from multiple functions, with respect to time, that are pooled together.

## 6 | CONCLUSIONS

MoE modeling is a powerful paradigm for approximating unknown DGPs, and for conducting classification, clustering, and regression. We have demonstrated how MoE models can be constructed for different data types, and we have provided theoretical results regarding the accuracy by which MoE models can approximate arbitrary DGPs and their mean functions.

When faced with data from an unknown DGP, MQL estimation can be used to estimate MoE models that are best fitted to the data in question. We have demonstrated that MQL estimation for MoE models can be conducted via blockwise-MM algorithms, and we have provided conditions under which the MQL estimators are consistent and asymptotically normal. A generic theorem is also provided for the construction of arbitrary information criteria for selecting the number of components of an MoE model.

To demonstrate their usefulness, we provide details regarding the conduct of classification, clustering, and regression via MoE models. Two examples have been provided to illustrate how these modes of application can be conducted in practice.

As with any review, summary, or tutorial article, we have omitted some details and topics for the sake of brevity and flow. A set of interest topics that we have omitted are works regarding the application of MoE models to the modeling of stationary time series processes, variable selection in MoE models via regularization, and finite sample model selection in MoE models. For the interested reader, we provide some details regarding these topics and the relevant literature, below.

The use of MoE for stationary time series modeling was first explored by Zeevi, Meir, and Adler (1999) who considered MoE models with soft-max gating functions and Gaussian autoregressive experts. The model of Zeevi et al. (1999) was further investigated in Carvalho and Tanner (2005a) alongside generic autoregressive expert functions. The family of MoE models with autoregressive GLM experts is explored in Carvalho and Tanner (2005b). A detailed investigation of the MoE model with autoregressive Poisson experts appears in Carvalho and Tanner (2007). A robust model using autoregressive Student- $t$  experts is considered in Carvalho and Skoulakis (2010). The use of hierarchical MoE models for modeling of univariate and multivariate time series process was investigated in Huerta, Jiang, and Tanner (2003) and Prado, Molina, and Huerta (2006), respectively. Recent applications of Gaussian-gated MoE models for univariate and multivariate time series process appear in Kalliovirta, Meitz, and Saikkonen (2015) and Kalliovirta, Meitz, and Saikkonen (2016), respectively.

In our exposition, we have left out details regarding variable selection in MoE models with regression experts due to the topic being overly specific and because we cannot do it justice within the confines of this article. There has been a lot of recent interest in the topic of sparse variable selection via model regularization, that extend upon the pioneering work of Tibshirani (1996) and Fan and Li (2001). In the context of mixture modeling, studies regarding the performance of such estimators under various assumptions on underlying DGPs can be found in Khalili and Chen (2007), Stadler, Buhlmann, and van de Geer (2010), and Khalili and Lin (2013). Extensions of these regularization results to MoE models remain a recent area of interest and appear in Khalili (2010), Peralta and Soto (2014), and Shohoudi, Khalili, Wolfson, and Asgharian (2016).

We note that the information criteria approach from Section 3.6 is not the only available paradigm for choosing the number of components in an MoE model. Recent works by Cohen and Le Pennec (2014) and Montuelle and Le Pennec (2014) have demonstrated that the finite-sample variable selection approach of Massart (2007) can be adapted for use in the MoE context. Unfortunately, both works are limited to model selection for MoE models with Gaussian regression experts, only. An interesting future direction is to extend these works to construct model selection rules for general MoE models.

Lastly, although we focus only on ML estimation of MoE models, there are other frameworks under which estimation can be conducted. Some examples of alternative frameworks include least-squares estimation (Bradshaw, Duchateau, & Bersini, 1997), Bayesian inference (Bishop & Svensen, 2002), regularized minimum cross-entropy estimation (Lu, 2006), and stochastic configuration algorithms (Wang & Li, 2017). The use of these alternative paradigms may allow for feasible estimation of MoE models in various settings where ML estimation may not be computationally possible, or may allow for precise inference to be drawn in situations where ML estimation leads to inaccurate conclusions.

## CONFLICT OF INTEREST

The authors have declared no conflicts of interest for this article.

## RELATED WIRES ARTICLES

[Laplace mixture of linear experts](#)

## ORCID

Hien D. Nguyen  <http://orcid.org/0000-0002-9958-432X>

## REFERENCES

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*, 716–723.
- Amemiya, T. (1985). *Advanced econometrics*. Cambridge, MA: Harvard University Press.
- Baudry, J.-P. (2015). Estimation and model selection for model-based clustering with the conditional classification likelihood. *Electronic Journal of Statistics*, *9*, 1041–1077.
- Biernacki, C., Celeux, G., & 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.
- Bishop, C. M., & Svensen, M. (2002). Bayesian hierarchical mixture of experts. In *Proceedings of the Nineteenth Conference on Uncertainty in Artificial Intelligence*.
- Bock, A. S., & Fine, I. (2014). Anatomical and functional plasticity in early blind individuals and the mixture of experts architecture. *Frontiers in Human Neuroscience*, *8*, 971.
- Bohning, D., & Lindsay, B. R. (1988). Monotonicity of quadratic-approximation algorithms. *Annals of the Institute of Mathematical Statistics*, *40*, 641–663.
- Boos, D. D., & Stefanski, L. A. (2013). *Essential statistical inference: Theory and methods*. New York, NY: Springer.
- Bradshaw, N. P., Duchateau, A., & Bersini, H. (1997). Global least-squares vs. EM training for the Gaussian mixture of experts. In *Proceedings of the International Conference on Artificial Neural Networks 1997*.
- Camplani, M., del Blanco, C. R., Salgado, L., Jaureguizar, F., & Garcia, N. (2014). Multi-sensor background subtraction by fusing multiple region-based probabilistic classifiers. *Pattern Recognition Letters*, *50*, 23–33.
- Carvalho, A. X., & Skoulakis, G. (2010). Time series mixtures of generalized t experts: ML estimation and an application to stock return density forecasting. *Econometric Reviews*, *29*, 642–687.
- Carvalho, A. X., & Tanner, M. A. (2005a). Mixtures-of-experts of autoregressive time series: Asymptotic normality and model specification. *IEEE Transactions on Neural Networks*, *16*, 39–56.
- Carvalho, A. X., & Tanner, M. A. (2005b). Modeling nonlinear time series with local mixtures of generalized linear models. *The Canadian Journal of Statistics*, *33*, 97–113.
- Carvalho, A. X., & Tanner, M. A. (2007). Modelling nonlinear count time series with local mixtures of Poisson autoregressions. *Computational Statistics and Data Analysis*, *51*, 5266–5294.
- Chamroukhi, F. (2016). Robust mixture of experts modeling using the t distribution. *Neural Networks*, *79*, 20–36.
- Chamroukhi, F. (2017). Robust mixture of experts modeling using the skew t distribution. *Neurocomputing*, *266*, 390–408.
- Chamroukhi, F., Glotin, H., & Same, A. (2013). Model-based functional mixture discriminant analysis with hidden process regression for curve classification. *Neurocomputing*, *112*, 153–163.
- Chamroukhi, F., Same, A., Govaert, G., & Akin, P. (2009). Time series modeling by a regression approach based on a latent process. *Neural Networks*, *22*, 593–602.
- Chamroukhi, F., Same, A., Govaert, G., & Akin, P. (2010). A hidden process regression model for functional data description. Application to curve discrimination. *Neurocomputing*, *73*, 1210–1221.
- Chen, K., Xu, L., & Chi, H. (1999). Improved learning algorithms for mixture of experts in multiclass classification. *Neural Networks*, *12*, 1229–1252.
- Cohen, S. X., & Le Pennec, E. (2014). Unsupervised segmentation of spectral images with a spatialized Gaussian mixture model and model selection. *Oil and Gas Science and Technology – Revue d'IFP Energies nouvelles*, *69*, 245–259.
- Cotter, N. E. (1990). The Stone-Weierstrass theorem and its application to neural networks. *IEEE Transactions on Neural Networks*, *1*, 290–295.
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, *2*, 303–314.
- DasGupta, A. (2008). *Asymptotic theory of statistics and probability*. New York, NY: Springer.
- DasGupta, A. (2011). *Probability for statistics and machine learning*. New York, NY: Springer.
- Deleforge, A., Forbes, F., & Horaud, R. (2015). High-dimensional regression with Gaussian mixtures and partially-latent response variables. *Statistics and Computing*, *25*, 893–911.
- Eavani, H., Hsieh, M. K., An, Y., Erus, G., Beason-Held, L., Resnick, S., & Davatzikos, C. (2016). Capturing heterogeneous group differences using mixture-of-experts: Application to a study of aging. *NeuroImage*, *125*, 498–514.
- Emani, M. K., & O'Boyle, M. (2015). Celebrating diversity: A mixture of experts approach for runtime mapping in dynamic environments. *ACM SIGPLAN Notices*, *50*, 499–508.
- Fan, J., & Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, *96*, 1348–1360.



- Gan, L., & Jiang, J. (1999). A test for global maximum. *Journal of the American Statistical Association*, *94*, 847–854.
- Grun, B., & Leisch, F. (2007). Fitting finite mixtures of generalized linear regressions in R. *Computational Statistics and Data Analysis*, *51*, 5247–5252.
- Grun, B., & Leisch, F. (2008). Flexmix version 2: Finite mixtures with concomitant variables and varying and constant parameters. *Journal of Statistical Software*, *28*, 1–35.
- Hayashi, F. (2000). *Econometrics*. Princeton, NJ: Princeton University Press.
- He, H., Boyd-Graber, J., Kwon, K., & Daume III, H. (2016). Opponent modeling in deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning*.
- Henderson, H. V., & Searle, S. R. (1979). Vec and vech operators for matrices with some uses in Jacobians and multivariate statistics. *The Canadian Journal of Statistics*, *7*, 65–81.
- Herzig, J., Bickel, A., Eitan, A., & Intrator, N. (2015). Monitoring cardiac stress using features extracted from  $S_1$  heart sounds. *IEEE Transactions on Biomedical Engineering*, *62*, 1169–1178.
- Huerta, G., Jiang, W., & Tanner, M. A. (2003). Time series modeling via hierarchical mixtures. *Statistica Sinica*, *13*, 1097–1118.
- Ingrassia, S., Minotti, S. C., & Vittadini, G. (2012). Local statistical modeling via a cluster-weighted approach with elliptical distributions. *Journal of Classification*, *29*, 363–401.
- Jacobs, R. A., Jordan, M. I., Nowlan, S. J., & Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural Computation*, *3*, 79–87.
- Jennrich, R. I. (1969). Asymptotic properties of non-linear least squares estimators. *Annals of Mathematical Statistics*, *40*, 633–643.
- Jiang, W., & Tanner, M. A. (1999a). Hierarchical mixtures-of-experts for exponential family regression models: Approximation and maximum likelihood estimation. *Annals of Statistics*, *27*, 987–1011.
- Jiang, W., & Tanner, M. A. (1999b). On the identifiability of mixture-of-experts. *Neural Networks*, *12*, 1253–1258.
- Jiang, W., & Tanner, M. A. (2000). On the asymptotic normality of hierarchical mixtures-of-experts for generalized linear models. *IEEE Transactions on Information Theory*, *46*, 1005–1013.
- Jones, P. N., & McLachlan, G. J. (1992). Fitting finite mixture models in a regression context. *Australian Journal of Statistics*, *34*, 233–240.
- Jordan, M. I., & Jacobs, R. A. (1992). Hierarchies of adaptive experts. In *Advances in Information Processing Systems* (pp. 985–992).
- Jordan, M. I., & Jacobs, R. A. (1994). Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, *6*, 181–214.
- Kalliovirta, L., Meitz, M., & Saikkonen, P. (2015). A Gaussian mixture autoregressive model for univariate time series. *Journal of Time Series Analysis*, *36*, 247–266.
- Kalliovirta, L., Meitz, M., & Saikkonen, P. (2016). Gaussian mixture vector autoregression. *Journal of Econometrics*, *192*, 485–498.
- Karimu, R. Y., & Azadi, S. (2017). Diagnosing the ADHD using a mixture of expert fuzzy models. *International Journal of Fuzzy Systems*, 1–5.
- Keribin, C. (2000). Consistent estimation of the order of mixture models. *Sankhya A*, *62*, 49–65.
- Khalili, A. (2010). New estimation and feature selection methods in mixture-of-experts models. *The Canadian Journal of Statistics*, *38*, 519–539.
- Khalili, A., & Chen, J. (2007). Variable selection in finite mixture of regression models. *Journal of the American Statistical Association*, *102*, 1025–1038.
- Khalili, A., & Lin, S. (2013). Regularization in finite mixture of regression models with diverging number of parameters. *Biometrics*, *69*, 436–446.
- Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *Annals of Mathematical Statistics*, *22*, 79–86.
- Lange, K. (2016). *MM optimization algorithms*. Philadelphia, PA: SIAM.
- Lee, Y.-S., & Cho, S.-B. (2014). Activity recognition with android phone using mixture-of-experts co-trained with labeled and unlabeled data. *Neurocomputing*, *126*, 106–115.
- Liem, R. P., Mader, C. A., & Martins, J. R. R. A. (2015). Surrogate models and mixtures of experts in aerodynamic performance prediction for aircraft mission analysis. *Aerospace Science and Technology*, *43*, 126–151.
- Lu, Z. (2006). A regularized minimum cross-entropy algorithm on mixture of experts for time series prediction and curve detection. *Pattern Recognition Letters*, *27*, 947–955.
- Masoudnia, S., & Ebrahimpour, R. (2014). Mixture of experts: A literature survey. *Artificial Intelligence Review*, *42*, 275–293.
- Massart, P. (2007). *Concentration inequalities and model selection*. New York, NY: Springer.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models*. Boca Raton, FL: CRC Press.
- McLachlan, G. J. (1988). On the choice of starting values for the EM algorithm in fitting mixture models. *The Statistician*, *37*, 417–425.
- McLachlan, G. J. (1992). *Discriminant analysis and statistical pattern recognition*. New York, NY: Wiley.
- McLachlan, G. J., & Peel, D. (2000). *Finite mixture models*. New York, NY: Wiley.
- Mendes, E. F., & Jiang, W. (2012). On convergence rates of mixture of polynomial experts. *Neural Computation*, *24*, 3025–3051.
- Meyer, R. R. (1976). Sufficient conditions for the convergence of monotonic mathematical programming algorithms. *Journal of Computer and System Sciences*, *12*, 108–121.
- Montuelle, L., & Le Pennec, E. (2014). Mixture of Gaussian regressions model with logistic weights, a penalized maximum likelihood approach. *Electronic Journal of Statistics*, *8*, 1661–1695.
- Nelder, J. A., & Wedderburn, R. W. M. (1972). Generalized linear models. *Journal of the Royal Statistical Society, Series A*, *135*, 370–384.
- Ng, S. K., & McLachlan, G. J. (2007). Extension of mixture-of-experts networks for binary classification of hierarchical data. *Artificial Intelligence in Medicine*, *41*, 57–67.
- Ng, S. K., & McLachlan, G. J. (2014). Mixture of random effects models for clustering multilevel growth trajectories. *Computational Statistics and Data Analysis*, *71*, 43–51.
- Ng, S.-K., & McLachlan, G. J. (2004). Using the EM algorithm to train neural networks: Misconceptions and a new algorithm for multiclass classification. *IEEE Transactions on Neural Networks*, *15*, 738–749.
- Nguyen, H. D. (2017). An introduction to MM algorithms for machine learning and statistical estimation. *WIREs Data Mining and Knowledge Discovery*, *7*, e1198.
- Nguyen, H. D., Lloyd-Jones, L. R., & McLachlan, G. J. (2016). A universal approximation theorem for mixture-of-experts models. *Neural Computation*, *28*, 2585–2593.
- Nguyen, H. D., & McLachlan, G. J. (2014). Asymptotic inference for hidden process regression models. In *Proceedings of the 2014 I.E. Statistical Signal Processing Workshop*.
- Nguyen, H. D., & McLachlan, G. J. (2015). Maximum likelihood estimation of Gaussian mixture models without matrix operations. *Advances in Data Analysis and Classification*, *9*, 371–394.
- Nguyen, H. D., & McLachlan, G. J. (2016). Laplace mixture of linear experts. *Computational Statistics and Data Analysis*, *93*, 177–191.
- Norets, A., & Pelenis, J. (2014). Posterior consistency in conditional density estimation by covariate dependent mixtures. *Econometric Theory*, *30*, 606–646.
- Olteanu, M., & Rynkiewicz, J. (2011). Asymptotic properties of mixture-of-experts models. *Neurocomputing*, *74*, 1444–1449.
- Peralta, B., & Soto, A. (2014). Embedded local feature selection within mixture of experts. *Information Sciences*, *269*, 176–187.
- Perthame, E., Forbes, F., Olivier, B., & Deleforge, A. (2016). Non linear robust regression in high dimension. In *The XXVIIIth International Biometric Conference*.
- Prado, R., Molina, F., & Huerta, G. (2006). Multivariate time series modeling and classification via hierarchical VAR mixtures. *Computational Statistics and Data Analysis*, *51*, 1445–1462.



- R Core Team. (2016). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Razaviyayn, M., Hong, M., & Luo, Z.-Q. (2013). A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM Journal of Optimization*, 23, 1126–1153.
- Same, A., Chamroukhi, F., Govaert, G., & Akinin, P. (2011). Model-based clustering and segmentation of time series with change in regime. *Advances in Data Analysis and Classification*, 5, 301–321.
- Schwarz, G. (1978). Estimating the dimensions of a model. *Annals of Statistics*, 6, 461–464.
- Shoenmakers, S., Guclu, U., van Gerven, M., & Heskes, T. (2015). Gaussian mixture models and semantic gating improve reconstructions from human brain activity. *Frontiers in Computational Neuroscience*, 8, 173.
- Shohoudi, A., Khalili, A., Wolfson, D. B., & Asgharian, M. (2016). Simultaneous variable selection and de-coursening in multi-path change-point models. *Journal of Multivariate Analysis*, 147, 202–217.
- Stadler, N., Buhlmann, P., & van de Geer, S. (2010).  $l_1$ -penalization for mixture regression models. *TEST*, 19, 209–256.
- Theis, L., & Bethge, M. (2015). Generative image modeling using spatial LSTMs. In *Advances in Neural Information Processing Systems*.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B*, 58, 267–288.
- van den Oord, A., & Schrauwen, B. (2014). Factoring variations in natural images with deep Gaussian mixture models. In *Advances in Neural Information Processing*.
- Variani, E., McDermott, E., & Heigold, G. (2015). A Gaussian mixture model layer jointly optimized with discriminative features within a deep neural network architecture. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Wang, D., & Li, M. (2017). Stochastic configuration networks: Fundamentals and algorithms. *IEEE Transactions on Cybernetics*, 47, 3466–3479.
- White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, 50, 1–25.
- Xu, L., Jordan, M. I., & Hinton, G. E. (1995). An alternative model for mixtures of experts. In *Advances in Neural Information Processing Systems* (pp. 633–640).
- Yuksel, S. E., & Gader, P. D. (2016). Context-based classification via mixture of hidden Markov model experts with applications in landmine detection. *IET Computer Vision*, 10, 873–883.
- Yuksel, S. E., Wilson, J. N., & Gader, P. D. (2012). Twenty years of mixture of experts. *IEEE Transactions on Neural Networks and Learning Systems*, 23, 1177–1193.
- Zeevi, A., Meir, R., & Adler, R. J. (1999). Non-linear models for time series using mixtures of autoregressive models. *Preprint*.
- Zeevi, A. J., Meir, R., & Maierov, V. (1998). Error bounds for functional approximation and estimation using mixtures of experts. *IEEE Transactions on Information Theory*, 44, 1010–1025.
- Zhou, H., & Lange, K. (2010). MM algorithms for some discrete multivariate distributions. *Journal of Computational and Graphical Statistics*, 19, 645–665.

**How to cite this article:** Nguyen HD, Chamroukhi F. Practical and theoretical aspects of mixture-of-experts modeling: An overview. *WIREs Data Mining Knowl Discov*. 2018;e1246. <https://doi.org/10.1002/widm.1246>