



HAL
open science

Approximation results regarding the multiple-output Gaussian gated mixture of linear experts model

Hien Duy Nguyen, Faicel Chamroukhi, Florence Forbes

► **To cite this version:**

Hien Duy Nguyen, Faicel Chamroukhi, Florence Forbes. Approximation results regarding the multiple-output Gaussian gated mixture of linear experts model. *Neurocomputing*, 2019, 366, pp.208-214. 10.1016/j.neucom.2019.08.014 . hal-02265793

HAL Id: hal-02265793

<https://hal.science/hal-02265793>

Submitted on 12 Aug 2019

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

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

Approximation results regarding the multiple-output Gaussian gated mixture of linear experts model

Hien D. Nguyen^{1*}, Faicel Chamroukhi², and Florence Forbes³

May 28, 2019

Abstract

Mixture of experts (MoE) models are a class of artificial neural networks that can be used for functional approximation and probabilistic modeling. An important class of MoE models is the class of mixture of linear experts (MoLE) models, where the expert functions map to real topological output spaces. Recently, Gaussian gated MoLE models have become popular in applied research. There are a number of powerful approximation results regarding Gaussian gated MoLE models, when the output space is univariate. These results guarantee the ability of Gaussian gated MoLE mean functions to approximate arbitrary continuous functions, and Gaussian gated MoLE models themselves to approximate arbitrary conditional probability density functions. We utilize

^{*}(Corresponding author email: h.nguyen5@latrobe.edu.au). ¹Department of Mathematics and Statistics, La Trobe University, Melbourne Victoria 3086, Australia. ²Laboratoire de Mathématiques Nicolas Oresme, Université de Caen Normandie, 14000 Caen-Cedex, France. ³Université de Grenoble Alpes, Inria, CNRS, Grenoble INP[†], LJK, 38000 Grenoble, France. [†]Institute of Engineering Université de Grenoble Alpes.

and extend upon the univariate approximation results in order to prove a pair of useful results for situations where the output spaces are multivariate. We do this by proving a pair of lemmas regarding the combination of univariate MoLE models, which are interesting in their own rights.

Keywords: artificial neural network; conditional model; Gaussian distribution; mean function; multiple-output; multivariate analysis

1 Introduction

Mixture of experts (MoE) models are a class of probabilistic artificial neural networks that were first introduced by Jacobs et al. (1991), and further developed in Jordan & Jacobs (1994) and Jordan & Xu (1995). In the contemporary setting, MoE models have become highly popular and successful in a range of applications including audio classification, bioinformatics, climate prediction, face recognition, financial forecasting, handwriting recognition, and text classification, among many others; see Yuksel et al. (2012), Masoudnia & Ebrahimpour (2014), and Nguyen & Chamroukhi (2018) and the references therein.

Let $\mathbb{X} \subseteq \mathbb{R}^p$ and \mathbb{Y} be input and output spaces (the specific nature of the output space will be discussed in the sequel), respectively, where $p \in \mathbb{N}$ (the zero-exclusive natural numbers). Let $\mathbf{X} \in \mathbb{X}$ and $\mathbf{Y} \in \mathbb{Y}$ be observable random variables, where \mathbf{X} may also be taken to be non-stochastic (i.e. $\mathbf{X} = \mathbf{x}$ with probability one, for some fixed $\mathbf{x} \in \mathbb{X}$). In addition to \mathbf{X} and \mathbf{Y} , define a third latent random variable $Z \in [n] = \{1, \dots, n\}$, such that

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

where $\text{Gate}_z(\mathbf{x}; \boldsymbol{\alpha})$ are parametric functions (known as gating functions), which depend on some

vector $\boldsymbol{\alpha}$ in a real space of fixed dimension. We call n the number of experts in the MoE. Here, the gating functions are required to satisfy the conditions $\text{Gate}_z(\boldsymbol{x}; \boldsymbol{\alpha}) > 0$, and $\sum_{z=1}^n \text{Gate}_z(\boldsymbol{x}; \boldsymbol{\alpha}) = 1$, for each $z \in [n]$, \boldsymbol{x} , and $\boldsymbol{\alpha}$.

The probability density functions (PDFs) of \mathbf{Y} , given $\mathbf{X} = \boldsymbol{x}$ and $Z = z$, are referred to as expert functions, which are parametric and can be written as

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

where $\boldsymbol{\beta}_z$ is a parameter vector in a real space of fixed dimensionality, for each z . For brevity, we write $f(\mathbf{y}|\boldsymbol{x}, z) = f(\mathbf{y}|\mathbf{X} = \boldsymbol{x}, Z = z; \boldsymbol{\beta}_z)$. We combine the gating functions (1) and expert functions (2), via the law of total probability, to produce the conditional PDF of \mathbf{Y} given $\mathbf{X} = \boldsymbol{x}$:

$$f(\mathbf{y}|\boldsymbol{x}; \boldsymbol{\theta}) = \sum_{z=1}^n \text{Gate}_z(\boldsymbol{x}; \boldsymbol{\alpha}) \text{Expert}_z(\mathbf{y}; \boldsymbol{x}, \boldsymbol{\beta}_z),$$

where, $\boldsymbol{\theta}$ is a vector that contains the elements of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}_z$ ($z \in [n]$). We refer to $f(\mathbf{y}|\boldsymbol{x}; \boldsymbol{\theta}) = f(\mathbf{y}|\mathbf{X} = \boldsymbol{x}; \boldsymbol{\theta})$ as the MoE model.

Depending on the choices of gating and expert functions, numerous classes of MoE models can be specified. For example, if \mathbb{Y} is a binary or categorical output space, then one can consider a logistic or multinomial logistic form (see, e.g. Jordan & Jacobs, 1994; Chen et al., 1999). If $\mathbb{Y} \subset \mathbb{N}$, then one may follow Grun & Leisch (2008) and utilize Poisson experts. When $\mathbb{Y} \subseteq (0, \infty)$ or $\mathbb{Y} \subseteq [0, 1]$, the mixture of gamma or beta experts are most appropriate (see, e.g. Jiang & Tanner, 1999a; Grun et al., 2012).

In this article, we are only concerned with the case where $\mathbb{Y} \subseteq \mathbb{R}^q$ ($q \in \mathbb{N}$), and when the mean

of the expert functions are linear in \mathbf{x} , so that

$$\mathbb{E}(\mathbf{y}|\mathbf{X} = \mathbf{x}, Z = z) = \mathbf{a}_z + \mathbf{B}_z^\top \mathbf{x} = \begin{bmatrix} a_{z1} + \mathbf{b}_{z1}^\top \mathbf{x} \\ \vdots \\ a_{zq} + \mathbf{b}_{zq}^\top \mathbf{x} \end{bmatrix}, \quad (3)$$

where we put the elements of $a_{zj} \in \mathbb{R}$ and $\mathbf{b}_{zj} \in \mathbb{R}^p$ ($j \in [q]$) into $\boldsymbol{\beta}_z$, for each z . Here $(\cdot)^\top$ is the transposition operator, $\mathbf{a}_z^\top = (a_{z,1}, \dots, a_{z,q}) \in \mathbb{R}^q$, and $\mathbf{B}_z \in \mathbb{R}^{p \times q}$ is a matrix with j th column $\mathbf{b}_{z,j}$. Following the nomenclature of Nguyen & McLachlan (2016), we refer to MoE models with the characteristic above as mixture of linear experts (MoLE) models.

Define the q -dimensional multivariate normal distribution by its PDF

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

where $\boldsymbol{\mu} \in \mathbb{R}^q$ is a mean vector and $\boldsymbol{\Sigma} \in \mathbb{R}^{q \times q}$ is a symmetric positive-definite covariance matrix. The multivariate normal linear experts were used to specify MoLE models in the foundational works of Jacobs et al. (1991) and Jordan & Jacobs (1994). Alternative MoLE models using Laplace, student- t , and skew student- t linear experts have also been considered in Nguyen & McLachlan (2016), Chamroukhi (2016), and Chamroukhi (2017), respectively.

In the MoE literature, there are two dominant choices for gating functions. The first, and by far the most popular, is the soft-max gate:

$$\text{Gate}_z(\mathbf{x}; \boldsymbol{\alpha}) = \frac{\exp(c_z + \mathbf{d}_z^\top \mathbf{x})}{\sum_{\zeta=1}^n \exp(c_\zeta + \mathbf{d}_\zeta^\top \mathbf{x})}, \quad (4)$$

where $c_z \in \mathbb{R}$ and $\mathbf{d}_z \in \mathbb{R}^p$ ($z \in [n]$) are put in the parameter vector $\boldsymbol{\alpha}$. This choice of gating was

originally considered in Jacobs et al. (1991).

The second of the dominant gating functions is the Gaussian gating function, or normalized-Gaussian radial basis gate (cf. Wang & Mendel, 1992), of the form

$$\text{Gate}_z(\mathbf{x}; \boldsymbol{\alpha}) = \frac{\pi_z \phi_p(\mathbf{x}; \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z)}{\sum_{\zeta=1}^n \pi_{\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{\zeta}, \boldsymbol{\Sigma}_{\zeta})}, \quad (5)$$

where $\pi_z > 0$ and $\sum_{z=1}^n \pi_z = 1$, and the unique elements of π_z , $\boldsymbol{\mu}_z$, and $\boldsymbol{\Sigma}_z$ ($z \in [n]$) are put in the parameter vector $\boldsymbol{\alpha}$. This gating choice was originally considered, in the MoE context, by Xu et al. (1995), although it had been used in the radial basis functions context by Wang & Mendel (1992). The Gaussian gating function has recently gained some popularity in the literature. For example, Ingrassia et al. (2012) used the Gaussian gating function within the framework of cluster-weighted modeling, and Deleforge et al. (2015a) used the Gaussian gates within the locally-linear mapping framework. Here, both cluster-weighted models and locally-linear mappings are types of MoE models. The Gaussian gates have also been used by Norets & Pelenis (2014) and Norets & Pati (2017) for MoE modeling of priors in Bayesian nonparametric regression. Under some restrictions, one can show that the class of soft-max gates is a subset of the Gaussian gates (cf. Ingrassia et al., 2012, Cor. 5).

A class of related gating functions to (5) are the student- t gates. This type of gating has been explored in Ingrassia et al. (2012), Ingrassia et al. (2014), and Perthame et al. (2018). Multivariate probit gates have also been considered in Geweke & Keane (2007).

Given any particular choice of gating, we can write the MoLE mean function as

$$\mathbf{m}(\mathbf{x}; \boldsymbol{\theta}) = \mathbb{E}(\mathbf{y} | \mathbf{X} = \mathbf{x}) = \sum_{z=1}^n \text{Gate}_z(\mathbf{x}; \boldsymbol{\alpha}) [\mathbf{a}_z + \mathbf{B}_z^{\top} \mathbf{x}].$$

An important property of MoLE models is their richness of representation capability. This representational richness has been characterized in a number of ways via various theoretical results. In Zeevi et al. (1998) and Jiang & Tanner (1999b), the single-output ($q = 1$) soft-max gated MoLE mean function was proved to be dense in an appropriate Sobolev space, under assumptions on differentiability and measurability. We define the notion of denseness in the manner of Cheney & Light (2000, Ch. 22), in the sequel. In Wang & Mendel (1992), the single-output Gaussian gated MoLE mean function was proved to be dense in the class of continuous functions, using the Stone-Weierstrass theorem (cf. Stone, 1948). Also via the Stone-Weierstrass theorem, Nguyen et al. (2016) proved that the single-output soft-max gated MoLE mean function is dense in the class of continuous functions.

Distributional approximation theorems have also been obtained. For example, Jiang & Tanner (1999a) proved that the class of single-output soft-max gated MoLE models can approximate any conditional density with mean function characterized via a ridge-type relationship with the input vector (cf. Pinkus, 2015) to an arbitrary degree of accuracy, with respect to the Hellinger distance and Kullback-Leibler divergence (see Pollard, 2002, Ch. 3). Replacing the linear mean functions (3) by polynomials, Mendes & Jiang (2012) obtained an approximation result regarding conditional densities with Sobolev class mean functions, instead of ridge-type mean functions.

We note that the results of Jiang & Tanner (1999a), Jiang & Tanner (1999b), and Mendes & Jiang (2012) are more general than what has been discussed here. That is, the results from the aforementioned papers extend to various generalized linear MoE models, and are not restricted to the MoLE context.

In a similar manner to Jiang & Tanner (1999a) and Mendes & Jiang (2012), Norets (2010) and Pelenis (2014) showed that the single-output soft-max gated MoLE models can approximate any

conditional density, regardless of mean function (under some regularity conditions), to an arbitrary degree of accuracy, with respect to a Kullback-Leibler type divergence. Extending upon the results of Norets (2010) and Pelenis (2014), Norets & Pelenis (2014) proved that the same approximation result holds for Gaussian gated MoLE models.

In recent years, numerous articles have described practical applications of multi-output MoLE models ($q > 1$; MO). For example, Chamroukhi et al. (2013) utilized such models for time series segmentation of human activity data. An application of MO-MoLE models to the analyze genomics data appears in Montuelle & Le Pennec (2014). Such models have also been used in image reconstruction and spectroscopic remote sensing applications (Deleforge et al., 2015b), as well as in sound source separation applications (Deleforge et al., 2015a). Time series applications of MO-MoLE models have been considered by Prado et al. (2006) and Kalliovirta et al. (2016).

Unfortunately, the single-output approximation theorems that have been previously cited no longer apply directly to MO-MoLE models. This is because there is no currently available results that allow for the pooling of marginal univariate effects, to best of our knowledge. That is, one cannot simply assume that the individual modeling of each output variable as an MoLE results in an MO-MoLE model when viewed across all of the variables, simultaneously, to the best of our knowledge, this current work is the first article to establish such a result via Lemmas 2 and 3.

In this paper, we utilize the previous results of Wang & Mendel (1992) and Norets & Pelenis (2014) in order to state useful approximation theorems to justify the use of MO-MoLE models for the analysis of functionally complex data and those data that arise from complex distributions. The approximation theorems regarding MO-MoLE models are presented as Theorems 3 and 4.

Theorem 3 states that we can arbitrarily well approximate the marginal conditional densities of any multivariate regression data generating process (DGP) in the conditional Kullback-Leibler

divergence, provided we utilize an MO-MoLE model with a sufficiently large but finite number of experts. Similarly, Theorem 4 states that we can utilize the mean function of an MO-MoLE model to arbitrarily well estimate all output variables of a continuous multivariate function over a compact support, simultaneously, given a sufficiently large but finite number of experts in our model. Recently there has been some interest in the use of deep variants of MO-MoLE models for multivariate density estimation and functional approximation (see, e.g., Shazeer et al., 2017, Fu et al., 2018, and Zhao et al., 2018). Our results provide empirical justification for the empirical effectiveness in modeling complex multivariate data of these deep variants and the already noted shallow counterparts that have been cited earlier.

In order to prove Theorems 3 and 4, we also proved a pair of technical lemmas regarding the combination of univariate MoLE models. These lemmas are interesting in their own rights, and are presented as Lemmas 2 and 3.

To the best of our knowledge, the approximation capabilities of the MO-MoLE models have not considered in previous articles on the topic. This is because past works have primarily focused on the derivation of estimation algorithms for MO-MoLE algorithms and the probabilistic properties of the estimators of such models under various DGPs. The assumption that MO-MoLE models would provide good approximations extrapolated from works regarding the related class of finite mixture models (cf. DasGupta, 2008, Sec. 33.1 and Norets & Pelenis, 2012). Our results are the first available theorems that explain the empirical effectiveness of MO-MoLE models in practice.

The rest of the paper is organized as follows. The univariate results of Wang & Mendel (1992) and Norets & Pelenis (2014) are presented in Section 2. The main results of the paper are stated in Section 3. Proofs of the main theorems are provided in Section 4. Discussions and conclusions are presented in Section 5. Supporting results are reported in the Appendix.

2 Preliminary results

The approximation result of Norets & Pelenis (2014) requires the following setup. Suppose that we observe the data pair $(\mathbf{X}, Y) \in \mathbb{X} \times \mathbb{Y}$, where $\mathbb{Y} \subset \mathbb{R}$, is generated from a DGP that can be characterized by a marginal PDF $g_{\mathbf{X}}(\mathbf{x})$ and conditional PDF $g_{Y|\mathbf{X}}(y|\mathbf{x})$. Let G denote the joint probability measure that is implied by the joint PDF $g_{Y|\mathbf{X}}(y|\mathbf{x})g_{\mathbf{X}}(\mathbf{x})$.

Make the assumptions that:

[A1] $g_{Y|\mathbf{X}}(y|\mathbf{x})$ is a continuous function in both $\mathbf{x} \in \mathbb{X}$ and $y \in \mathbb{Y}$, almost surely with respect to G , and

[A2] there exists some $\rho > 0$ such that

$$\int_{\mathbb{X} \times \mathbb{Y}} \log \frac{g_{Y|\mathbf{X}}(y|\mathbf{x})}{\inf_{\{(s,t): \|\mathbf{x}-s\| \leq \rho, \|y-t\| \leq \rho\}} g_{Y|\mathbf{X}}(t|s)} dG(\mathbf{x}, y) < \infty,$$

where $\|\cdot\|$ is the Euclidean norm.

As stated by Norets & Pelenis (2014), condition [A2] is a technical requirement that the log relative changes in $g_{Y|\mathbf{X}}(y|\mathbf{x})$ are finite, on average, and that $g_{Y|\mathbf{X}}(y|\mathbf{x})$ is positive for all pairs of \mathbf{x} and y .

Write the class of MO-MoLE models with Gaussian gates and Gaussian linear experts over \mathbb{X} as

$$\mathcal{L}_q(\mathbb{X}) = \left\{ f : f(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \frac{\sum_{z=1}^n \pi_z \phi_p(\mathbf{x}; \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z) \phi_q(\mathbf{y}; \mathbf{a}_z + \mathbf{B}_z^\top \mathbf{x}, \mathbf{C}_z)}{\sum_{\zeta=1}^n \pi_\zeta \phi_p(\mathbf{x}; \boldsymbol{\mu}_\zeta, \boldsymbol{\Sigma}_\zeta)}, n \in \mathbb{N} \right\}.$$

Here, $\mathbf{C}_z \in \mathbb{R}^{q \times q}$ is a symmetric positive-definite covariance matrix, for each $z \in [n]$. Furthermore, define the subclass $\mathcal{L}_q^*(\mathbb{X}) \subset \mathcal{L}_q(\mathbb{X})$, where

$$\mathcal{L}_q^*(\mathbb{X}) = \left\{ f : f(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \frac{\sum_{z=1}^n \pi_z \phi_p(\mathbf{x}; \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z) \phi_q(\mathbf{y}; \mathbf{a}_z, \mathbf{C}_z)}{\sum_{\zeta=1}^n \pi_\zeta \phi_p(\mathbf{x}; \boldsymbol{\mu}_\zeta, \boldsymbol{\Sigma}_\zeta)}, n \in \mathbb{N} \right\}.$$

The following result is a direct consequence of Norets & Pelenis (2014, Thm. 3.1).

Theorem 1. *Let \mathbb{X} be compact and $\mathbb{Y} \subset \mathbb{R}$. If the data pair (\mathbf{X}, Y) arises from a DGP that is characterized by the joint probability measure G , and if $g_{Y|\mathbf{X}}$ is a conditional PDF that satisfies [A1] and [A2], then for every $\epsilon > 0$, there exist n and $\boldsymbol{\theta}$ that characterize an MoLE model $f \in \mathcal{L}_1^*(\mathbb{X})$, such that*

$$\int_{\mathbb{X} \times \mathbb{Y}} \log \frac{g_{Y|\mathbf{X}}(y|\mathbf{x})}{f(y|\mathbf{x}; \boldsymbol{\theta})} dG(\mathbf{x}, y) < \epsilon.$$

We now consider the approximation theorem of Wang & Mendel (1992). Let $\mathcal{C}(\mathbb{X})$ denote the class of all continuous functions with support $\mathbb{X} \subset \mathbb{R}^p$. For a pair of single-output functions u and v on \mathbb{X} , we can define the uniform distance between u and v as $d_\infty(u, v) = \|u - v\|_\infty$, where $\|u(\mathbf{x})\|_\infty = \sup_{\mathbf{x} \in \mathbb{X}} |u(\mathbf{x})|$ is the uniform norm over the support \mathbb{X} .

The following definition is taken from Cheney & Light (2000, Ch. 22). Suppose that $\mathcal{U}(\mathbb{X})$ and $\mathcal{V}(\mathbb{X})$ are two classes of functions on \mathbb{X} . If \mathcal{U} and \mathcal{V} are normed vector spaces (with respect to an appropriate norm), then we say that \mathcal{U} is dense in \mathcal{V} , if the closure of \mathcal{U} is \mathcal{V} . That is, we say that \mathcal{U} is dense in \mathcal{V} with respect to the uniform norm, if for each $v \in \mathcal{V}$ and $\epsilon > 0$, there exists a linear combination $u \in \mathcal{U}$, such that $d_\infty(u, v) < \epsilon$.

For $\mathbb{Y} = \mathbb{R}^q$, denote the class of Gaussian gated MoLE mean functions over the support \mathbb{X} by

$$\mathcal{M}_q(\mathbb{X}) = \left\{ \mathbf{m} : \mathbf{m}(\mathbf{x}; \boldsymbol{\theta}) = \sum_{z=1}^n \frac{\pi_z \phi_p(\mathbf{x}; \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z)}{\sum_{\zeta=1}^n \pi_\zeta \phi_p(\mathbf{x}; \boldsymbol{\mu}_\zeta, \boldsymbol{\Sigma}_\zeta)} [\mathbf{a}_z + \mathbf{B}_z^\top \mathbf{x}], \mathbf{x} \in \mathbb{X}, n \in \mathbb{N} \right\}.$$

Further define the subclass $\mathcal{M}_q^*(\mathbb{X}) \subset \mathcal{M}_q(\mathbb{X})$, where

$$\mathcal{M}_q^*(\mathbb{X}) = \{ \mathbf{m} \in \mathcal{M}_q(\mathbb{X}) : \mathbf{B}_z = \mathbf{0}, \text{ for each } z \in [n] \},$$

and $\mathbf{0}$ is a matrix containing only zeros of appropriate dimensionality. In the $q = 1$ case, the following result was proved by Wang & Mendel (1992), using the Stone-Weierstrass theorem.

Theorem 2. *If $\mathbb{X} \subset \mathbb{R}^p$ is a compact set, then the set $\mathcal{M}_1^*(\mathbb{X})$ is dense in $\mathcal{C}(\mathbb{X})$, with respect to the uniform norm. Subsequently, since $\mathcal{M}_1^*(\mathbb{X}) \subset \mathcal{M}_1(\mathbb{X})$, it follows that $\mathcal{M}_1(\mathbb{X})$ is also dense in $\mathcal{C}(\mathbb{X})$, with respect to the uniform norm.*

3 Main results

Extending from the work of Norets & Pelenis (2014), we now consider the approximation capabilities of MO-MoLE models. To do so, we require the following definitions.

Let $\mathbb{Y} = \prod_{j=1}^q \mathbb{Y}_j$, such that $\mathbb{Y}_j \subset \mathbb{R}$. Suppose that the data pair $(\mathbf{X}, \mathbf{Y}) \in \mathbb{X} \times \mathbb{Y}$ is generated from a DGP that can be characterized by a marginal PDF $g_{\mathbf{X}}(\mathbf{x})$ and that admits the univariate conditional PDFs $g_{Y_j|\mathbf{X}}(y_j|\mathbf{x})$, for each $j \in [q]$, where $\mathbf{Y}^\top = (Y_1, \dots, Y_q)$, and subsequently $\mathbf{y}^\top = (y_1, \dots, y_q)$. Let the probability measure that is implied by the PDF $g_{Y_j|\mathbf{X}}(y_j|\mathbf{x}) g_{\mathbf{X}}(\mathbf{x})$ be written as G_j , for each $j \in [q]$.

Make Assumptions [A1] and [A2] regarding each of the conditional PDFs $g_{Y_j|\mathbf{X}}(y_j|\mathbf{x})$. That is, assume that:

[B1] for each $j \in [q]$, $g_{Y_j|\mathbf{X}}(y_j|\mathbf{x})$ is a continuous function in both $\mathbf{x} \in \mathbb{X}$ and $y_j \in \mathbb{Y}_j$, almost surely with respect to G_j , and

[B2] for each $j \in [q]$, there exists some $\rho_j > 0$ such that

$$\int_{\mathbb{X} \times \mathbb{Y}_j} \log \frac{g_{Y_j|\mathbf{X}}(y|\mathbf{x})}{\inf_{\{(s,t): \|\mathbf{x}-\mathbf{s}\| \leq \rho, \|y-t\| \leq \rho\}} g_{Y_j|\mathbf{X}}(t|\mathbf{s})} dG(\mathbf{x}, y) < \infty.$$

Using Theorem 1, we obtain the following generalization regarding MO-MoLE models from the class $\mathcal{L}_q^*(\mathbb{X})$, and subsequently, the class $\mathcal{L}_q(\mathbb{X})$. The proof appears in Section 4.

Theorem 3. *Let \mathbb{X} be compact and $\mathbb{Y} = \prod_{j=1}^q \mathbb{Y}_j$, where $\mathbb{Y}_j \subset \mathbb{R}$. Assume that the DGP of (\mathbf{X}, \mathbf{Y}) is compatible with each of the joint probability measures G_j ($j \in [q]$). If the conditional PDFs $g_{Y_j|\mathbf{X}}$ ($j \in [q]$) are such that Assumptions [B1] and [B2] are satisfied, then there exist n and $\boldsymbol{\theta}$ that characterize an MoLE model $f \in \mathcal{L}_q(\mathbb{X})$, such that for some $\epsilon > 0$,*

$$\int_{\mathbb{X} \times \mathbb{Y}_j} \log \frac{g_{Y_j|\mathbf{X}}(y_j|\mathbf{x})}{f(y_j|\mathbf{x}; \boldsymbol{\theta})} dG_j(\mathbf{x}, y_j) < \epsilon$$

is satisfied simultaneously for all $j \in [q]$.

We now extend upon the result of Wang & Mendel (1992) in order to state a theorem regarding the approximation capabilities of MO-MoLE mean functions. Define the space of MO continuous functions over \mathbb{X} as

$$\mathcal{C}_q(\mathbb{X}) = \{ \mathbf{m}^\top(\mathbf{x}) = (m_1(\mathbf{x}), \dots, m_q(\mathbf{x})) : m_j \in \mathcal{C}(\mathbb{X}), j \in [q] \}.$$

We wish to determine the relationship between the class $\mathcal{M}_q(\mathbb{X})$ and $\mathcal{C}_q(\mathbb{X})$, for $q > 1$. In order to state such a relationship, we require an appropriate distance function. Following the approach of Chiou et al. (2014), we utilize summation to induce a multivariate norm and distance function as follows. Let $\mathbf{u}^\top = (u_1, \dots, u_q)$ and $\mathbf{v}^\top = (v_1, \dots, v_q)$ be a pair of MO functions on \mathbb{X} . Denote the induced distance between \mathbf{u} and \mathbf{v} by $d_{q,\infty}(\mathbf{u}, \mathbf{v}) = \|\mathbf{u} - \mathbf{v}\|_{q,\infty}$, where $\|\mathbf{u}(\mathbf{x})\|_{q,\infty} = \sum_{j=1}^q \|u_j(\mathbf{x})\|_\infty$.

We prove that the operator $\|\cdot\|_{q,\infty}$ satisfies the definition of a norm in the Appendix. Our following result generalizes Theorem 2. The proof appears in Section 4.

Theorem 4. *If $\mathbb{X} \subset \mathbb{R}^p$ is a compact set and $q \in \mathbb{N}$, then the sets of MO-MoLE mean functions $\mathcal{M}_q^*(\mathbb{X})$ and $\mathcal{M}_q(\mathbb{X})$ are dense in $\mathcal{C}_q(\mathbb{X})$, with respect to the induced norm.*

We note that both Theorems 4 and 3 require that the gating functions are of the Gaussian form, given by (5). We note that Nguyen et al. (2016, Thm. 1) provides a version of Theorem 2 that utilizes the soft-max gating function instead of the Gaussian gating function, under the same compactness assumption on \mathbb{X} . Similarly, Pelenis (2014, Thm. 1) provides a substitute for Theorem 1, under almost identical assumptions, for soft-max gated MoLEs with Gaussian linear experts. An additional assumption that $\int_{\mathbb{Y}} y^2 g_{Y|\mathbf{X}}(y|\mathbf{x}) d\mathbf{x} < \infty$ for all $\mathbf{x} \in \mathbb{X}$ is required, in order to apply the result of Pelenis (2014). Thus, one can largely replace the Gaussian gating functions in Theorems 4 and 3 by the soft-max gating functions of form (4), and still obtain the conclusions of the two results.

Theorems 3 and 4 are directly applicable to the MO-MoLE models that are considered in Prado et al. (2006), Chamroukhi et al. (2013), Montuelle & Le Pennec (2014), (Deleforge et al., 2015b), (Deleforge et al., 2015a), and Kalliovirta et al. (2016). For example, the MO-MoLE models of Chamroukhi et al. (2013) and Deleforge et al., 2015b take the forms:

$$f(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \sum_{z=1}^n \frac{\exp(c_z + \mathbf{d}_z^\top \mathbf{x})}{\sum_{\zeta=1}^n \exp(c_\zeta + \mathbf{d}_\zeta^\top \mathbf{x})} \phi_q(\mathbf{y}; \mathbf{a}_z + \mathbf{B}_z^\top \mathbf{x}, \boldsymbol{\Omega}_z)$$

and

$$f(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \sum_{z=1}^n \frac{\pi_z \phi_p(\mathbf{x}; \boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z)}{\sum_{\zeta=1}^n \pi_\zeta \phi_p(\mathbf{x}; \boldsymbol{\mu}_\zeta, \boldsymbol{\Sigma}_\zeta)} \phi_q(\mathbf{y}; \mathbf{a}_z + \mathbf{B}_z^\top \mathbf{x}, \boldsymbol{\Omega}_z),$$

where $\boldsymbol{\Omega}_z$ is a positive definite and symmetric matrix, for each $z \in [n]$. Thus both MO-MoLE models satisfy the assumptions of 3 and 4. We can therefore conclude that with sufficiently many

experts n , both models are able to arbitrarily well approximate mean functions and conditional marginal density functions of the underlying DGPs. This therefore explains why these models, and the other cited MO-MoLE models are able to well approximate their target functions in the respective articles.

The distributional approximation and denseness results provide some theoretical justification for the flexibility and goodness-of-fit of such models in the simulation studies and applications that are presented in the listed references. Furthermore, we note that the results are directly applicable to any class of MoLE models with gating functions that includes Gaussian as a subclass. For example, it is hypothetically possible to construct a family of skew normal gated MoLE models using the skew normal distributions of Azzalini & Dalla Valle (1996), where the skew normal density function replaces the Gaussian density function in (5). Since the skew normal distribution includes the Gaussian distribution as a special case, the results of our theorems would immediately apply to such a construction.

4 Proofs of main results

The following lemmas streamline the proofs of Theorems 4 and 3. The first lemma is well known and characterizes the functional form of the product of two Gaussian PDFs. A proof of the lemma can be found in Bromiley (2014). The proofs of Lemmas 2 and 3 appear in the Appendix.

Lemma 1. *If $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \in \mathbb{R}^p$ and $\boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2 \in \mathbb{R}^{p \times p}$ are symmetric positive-definite covariance matrices, then*

$$\phi_p(\mathbf{x}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) \phi_p(\mathbf{x}; \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2) = c \phi_p(\mathbf{x}; \boldsymbol{\mu}_{12}, \boldsymbol{\Sigma}_{12}),$$

where $c > 0$, $\boldsymbol{\Sigma}_{12}^{-1} = \boldsymbol{\Sigma}_1^{-1} + \boldsymbol{\Sigma}_2^{-1}$, and $\boldsymbol{\mu}_{12} = \boldsymbol{\Sigma}_{12} (\boldsymbol{\Sigma}_1^{-1} \boldsymbol{\mu}_1 + \boldsymbol{\Sigma}_2^{-1} \boldsymbol{\mu}_2)$.

Lemma 2. If $\mathbf{m}_{[1]}, \mathbf{m}_{[2]} \in \mathcal{M}_q^*(\mathbb{X})$, for some \mathbb{X} , then $\mathbf{m}_{[12]} \in \mathcal{M}_q^*(\mathbb{X})$, where $\mathbf{m}_{[12]} = \mathbf{m}_{[1]} + \mathbf{m}_{[2]}$.

Lemma 3. If $f_{[1]} \in \mathcal{L}_q^*(\mathbb{X})$ and $f_{[2]} \in \mathcal{L}_r^*(\mathbb{X})$, for some \mathbb{X} ($q, r \in \mathbb{N}$), then $f_{[12]} \in \mathcal{L}_{q+r}^*(\mathbb{X})$, where $f_{[12]} = f_{[1]}f_{[2]}$.

4.1 Proof of Theorem 3

By Theorem 1, under Assumptions [B1] and [B2], for each $j \in [q]$ and $\epsilon > 0$, there exists an n_j and $\boldsymbol{\theta}_j$ that specifies a function

$$f(y_j|\mathbf{x}; \boldsymbol{\theta}_j) = \sum_{z=1}^{n_j} \frac{\pi_{jz} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{jz}, \boldsymbol{\Sigma}_{jz}) \phi_1(y_j; a_{jz}, \sigma_{jz}^2)}{\sum_{\zeta=1}^{n_j} \pi_{j\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{j\zeta}, \boldsymbol{\Sigma}_{j\zeta})},$$

in $\mathcal{L}_1^*(\mathbb{X})$, such that

$$\int_{\mathbb{X} \times \mathbb{Y}_j} \log \frac{g_{Y_j|\mathbf{X}}(y_j|\mathbf{x})}{f(y_j|\mathbf{x}; \boldsymbol{\theta}_j)} dG(\mathbf{x}, y_j) < \epsilon$$

is satisfied.

We complete the proof constructively. That is, we can show that the product of the marginal PDFs $f(y_j|\mathbf{x}; \boldsymbol{\theta}_j)$ yields a joint PDF $f(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})$, which is in the class $\mathcal{L}_q^*(\mathbb{X})$. This is achieved via repeated applications of Lemma 3. We obtain the desired conclusion by noting that $\mathcal{L}_q^*(\mathbb{X}) \subset \mathcal{L}_q(\mathbb{X})$.

4.2 Proof of Theorem 4

Let \mathbb{X} be a compact set. Define \mathbf{e}_j to be a column vector with 1 in the j th position and 0, elsewhere. Let $\mathbf{u}^\top(\mathbf{x}) = (u_1(\mathbf{x}), \dots, u_q(\mathbf{x})) \in \mathcal{C}_q(\mathbb{X})$ be an arbitrary continuous MO function over \mathbb{X} . By

Theorem 2, there exists an MO mean function

$$m_j(\mathbf{x}; \boldsymbol{\theta}_j) = \sum_{z=1}^{n_j} \frac{\pi_{jz} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{jz}, \boldsymbol{\Sigma}_{jz})}{\sum_{\zeta=1}^{n_j} \pi_{j\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{j\zeta}, \boldsymbol{\Sigma}_{j\zeta})} a_{jz},$$

for each $j \in [q]$, such that $d_\infty(m_j(\mathbf{x}; \boldsymbol{\theta}_j), u_j(\mathbf{x})) < \epsilon/q$, for any $\epsilon > 0$. Here, $\boldsymbol{\theta}_j$ is a parameter vector that contains the unique elements of $\boldsymbol{\mu}_{jz}$, $\boldsymbol{\Sigma}_{jz}$, and $a_{jz} \in \mathbb{R}$, for each $z \in n_j$, where $n_j \in \mathbb{N}$, for each $j \in [q]$. Now, write

$$\mathbf{m}_j(\mathbf{x}; \boldsymbol{\theta}_j) = m_j(\mathbf{x}; \boldsymbol{\theta}_j) \times \mathbf{e}_j$$

and note that for any $k \neq j$, $\mathbf{m}_{jk}(\mathbf{x}; \boldsymbol{\theta}_j) = 0$, for all $\mathbf{x} \in \mathbb{X}$.

Consider the fact that the j th coordinate of the function

$$\mathbf{m}(\mathbf{x}; \boldsymbol{\theta}) = \sum_{j=1}^q \mathbf{m}_j(\mathbf{x}; \boldsymbol{\theta}_j) \tag{6}$$

is only influenced by the j th functional $\mathbf{m}_j(\mathbf{x}; \boldsymbol{\theta}_j)$, by construction. Thus, at each coordinate j , we have

$$d_\infty(m_j(\mathbf{x}), u_j(\mathbf{x})) = d_\infty(m_j(\mathbf{x}; \boldsymbol{\theta}_j), u_j(\mathbf{x})) < \epsilon/q.$$

By definition of the induced distance, we therefore obtain the result that

$$\begin{aligned} d_{q,\infty}(\mathbf{m}, \mathbf{u}) &= \sum_{j=1}^q d_\infty(m_j(\mathbf{x}; \boldsymbol{\theta}_j), u_j(\mathbf{x})) \\ &< q \times (\epsilon/q) = \epsilon. \end{aligned}$$

It suffices to show that (6) is a function in the class $\mathcal{M}_q^*(\mathbb{X})$. We obtain such a result by

repeated application of Lemma 2.

5 Discussion and conclusions

Theorem 3 implies that all q univariate $g_{Y_j|\mathbf{X}}(y_j|\mathbf{x})$ conditional PDFs ($j \in [q]$) of a q -variate target conditional PDF $g_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$ can be approximated to an arbitrary degree of accuracy via a Gaussian gated MoLE model with Gaussian linear experts, with respect to a Kullback-Leibler like divergence, assuming the fulfillment of Assumptions [B1] and [B2]. Unfortunately, the statement of the theorem provides no guarantees regarding the approximation accuracy of the dependence structures between each of the q univariate variables y_j , conditioned on the observation $\mathbf{X} = \mathbf{x}$. Using Theorem 1, we cannot prove such a result using algebraic manipulations alone, in the manner that has been used to prove Theorem 3. Proving that dependence structures can also be approximated to an arbitrary degree of accuracy is a topic of ongoing research in the literature. Such results may be sought via adaptations and extensions of the joint density approximation results of DasGupta (2008, Sec. 33.1) or Norets & Pelenis (2012), to the problem of multivariate conditional density approximation.

Finally, we note that Theorems 1–4 do not provide rates, regarding the reduction of approximation error as functions of q and n . Rate results would require stronger assumptions on the space of approximands. For example, we may utilize the results of Zeevi et al. (1998) in order to obtain an approximation rate for functional approximations from the class $\mathcal{M}_q(\mathbb{X})$, under the additional assumption that the MO approximand is a member of some appropriate Sobolev space. Similarly, using the results of Jiang & Tanner (1999a), we may obtain approximation rates for conditional approximations from the class $\mathcal{L}_q(\mathbb{X})$, under the additional assumptions that the approximand

univariate conditional PDFs satisfy are restricted to affine-dependence structures, with respect to the input vector.

In this paper, we sought to prove the most general results that were available, regarding the approximation capability of the Gaussian gated MoLE model. As such, we do not wish to impose more assumptions than is strictly necessary in order to establish meaningful theorems. We leave the establishment of further interesting results that may require more stringent assumptions to the future.

Acknowledgments

Hien Nguyen is funded by Australian Research Council (ARC) grants DE170101134 and DP180101192, and a La Trobe University startup grant. This research is funded directly by the Inria LANDeR project. The authors thank the anonymous Reviewers for their useful comments that have improved the exposition of the article.

Appendix

The induced norm

Let \mathcal{U} be a normed vector space, and let \mathbf{u} and \mathbf{v} be arbitrary elements of \mathcal{U} . We say that the operator $\|\cdot\|$ is a norm on \mathcal{U} if it satisfies the following assumptions: (i) $\|\mathbf{u}\| \geq 0$ and $\|\mathbf{u}\| = 0$ if and only if $\mathbf{u} = \mathbf{0}$, (ii) For every $c \in \mathbb{R}$, $\|c\mathbf{u}\| = |c| \|\mathbf{u}\|$, and (iii) $\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|$ (cf. Oden & Demkowicz, 2010, Sec. 4.6).

Proposition 1. *For any vector space $\mathcal{U}_q(\mathbb{X})$ of MO functions on \mathbb{X} , the operator $\|\cdot\|_{\Sigma}$ satisfies the*

definition of a norm.

Proof. Let $\mathbf{u}^\top = (u_1, \dots, u_q)$ and $\mathbf{v}^\top = (v_1, \dots, v_q)$ be two arbitrary elements in \mathcal{U}_q . Recall that the operator $\|\cdot\|_\infty$ is a norm over any vector space of single-output functions. This implies that $\|u_j\|_\infty \geq 0$ for each $j \in [q]$ and thus $\|\mathbf{u}\|_{q,\infty} = \sum_{j=1}^q \|u_j\|_\infty \geq 0$.

Suppose that $\|\mathbf{u}\|_{q,\infty} = 0$. This implies that each component of $\sum_{j=1}^q \|u_j\|_\infty$ must equal to zero since no component may take a negative value. However, since $\|\cdot\|_\infty$ is a norm, this implies that $\mathbf{u} = \mathbf{0}$. Now suppose that $\mathbf{u} = \mathbf{0}$. The direct definition of $\|\cdot\|_{q,\infty}$ leads to the result that $\|\mathbf{u}\|_{q,\infty} = 0$. Thus, together, $\|\cdot\|_{q,\infty}$ fulfills Assumption (i).

Assumption (ii) is shown to be fulfilled by observing the direct chain of equalities:

$$\begin{aligned} \|\mathbf{c}\mathbf{u}\|_{q,\infty} &= \sum_{j=1}^q \|cu_j\|_\infty \\ &= \sum_{j=1}^q |c| \|u_j\|_\infty \\ &= |c| \sum_{j=1}^q \|u_j\|_\infty \\ &= |c| \|\mathbf{u}\|_{q,\infty}, \end{aligned}$$

where the second line is due to the fact that $\|\cdot\|_\infty$ is a norm.

Assumption (iii) is also shown to be fulfilled by observing the chain of arguments:

$$\begin{aligned}
\|\mathbf{u} + \mathbf{v}\|_{q,\infty} &= \sum_{j=1}^q \|u_j + v_j\|_\infty \\
&\leq \sum_{j=1}^q [\|u_j\|_\infty + \|v_j\|_\infty] \\
&= \sum_{j=1}^q \|u_j\|_\infty + \sum_{j=1}^q \|v_j\|_\infty \\
&= \|\mathbf{u}\|_{q,\infty} + \|\mathbf{v}\|_{q,\infty},
\end{aligned}$$

where the second line is again due to the fact that $\|\cdot\|_\infty$ is a norm. The proof is thus complete. \square

Proof of Lemma 2

Since $\mathbf{m}_{[1]}, \mathbf{m}_{[2]} \in \mathcal{M}_q^*(\mathbb{X})$, we can write $\mathbf{y}_{[k]}$ as

$$\mathbf{m}_{[k]}(\mathbf{x}; \boldsymbol{\theta}_k) = \sum_{z=1}^{n_k} \frac{\pi_{kz} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{kz}, \boldsymbol{\Sigma}_{kz})}{\sum_{\zeta=1}^{n_k} \pi_{k\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{k\zeta}, \boldsymbol{\Sigma}_{k\zeta})} \mathbf{a}_{kz},$$

where $\boldsymbol{\theta}_k$ contains the unique elements of $\boldsymbol{\mu}_{kz}$, $\boldsymbol{\Sigma}_{kz}$, and \mathbf{a}_{kz} ($z \in [n_k]$; $n_k \in \mathbb{N}$), for each $k \in \{1, 2\}$.

Next, we write

$$\begin{aligned}
\mathbf{m}_{[12]}(\mathbf{x}) &= \sum_{k=1}^2 \sum_{z=1}^{n_k} \frac{\pi_{kz} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{kz}, \boldsymbol{\Sigma}_{kz})}{\sum_{\zeta=1}^{n_k} \pi_{k\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{k\zeta}, \boldsymbol{\Sigma}_{k\zeta})} \mathbf{a}_{kz} \\
&= \sum_{z=1}^{n_1} \frac{\pi_{1z} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{1z}, \boldsymbol{\Sigma}_{1z}) \sum_{\zeta=1}^{n_2} \pi_{2\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{2\zeta}, \boldsymbol{\Sigma}_{2\zeta})}{\prod_{k=1}^2 \sum_{\zeta=1}^{n_k} \pi_{k\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{k\zeta}, \boldsymbol{\Sigma}_{k\zeta})} \mathbf{a}_{1z} \\
&\quad + \sum_{z=1}^{n_2} \frac{\pi_{2z} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{2z}, \boldsymbol{\Sigma}_{2z}) \sum_{\zeta=1}^{n_1} \pi_{1\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{1\zeta}, \boldsymbol{\Sigma}_{1\zeta})}{\prod_{k=1}^2 \sum_{\zeta=1}^{n_k} \pi_{k\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{k\zeta}, \boldsymbol{\Sigma}_{k\zeta})} \mathbf{a}_{2z}.
\end{aligned}$$

For each $s \in [n_1]$ and $t \in [n_2]$, we can perform the following mappings: $\mathbf{a}_{(st)} = \mathbf{a}_{1s} + \mathbf{a}_{2t}$, $\bar{\pi}_{(st)} = \pi_{1s}\pi_{2t}$, $\Sigma_{(st)}^{-1} = \Sigma_{1s}^{-1} + \Sigma_{2t}^{-1}$, and $\boldsymbol{\mu}_{(st)} = \Sigma_{(st)} \left(\Sigma_{1s}^{-1} \boldsymbol{\mu}_{1s} + \Sigma_{2t}^{-1} \boldsymbol{\mu}_{2t} \right)$.

Using Lemma 1, we can write

$$\begin{aligned} \mathbf{m}_{[12]}(\mathbf{x}) &= \sum_{s=1}^{n_1} \sum_{t=1}^{n_2} \frac{c_{st} \bar{\pi}_{(st)} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{(st)}, \Sigma_{(st)})}{\sum_{\xi=1}^{n_1} \sum_{\zeta=1}^{n_2} c_{\xi\zeta} \bar{\pi}_{(\xi\zeta)} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{(\xi\zeta)}, \Sigma_{(\xi\zeta)})} \mathbf{a}_{(st)} \\ &= \sum_{s=1}^{n_1} \sum_{t=1}^{n_2} \frac{\pi_{(st)} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{(st)}, \Sigma_{(st)})}{\sum_{\xi=1}^{n_1} \sum_{\zeta=1}^{n_2} \pi_{(\xi\zeta)} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{(\xi\zeta)}, \Sigma_{(\xi\zeta)})} \mathbf{a}_{(st)}, \end{aligned}$$

where $\pi_{(st)} = c_{st} \bar{\pi}_{(st)} / \sum_{\xi=1}^{n_1} \sum_{\zeta=1}^{n_2} c_{\xi\zeta} \bar{\pi}_{(\xi\zeta)}$, for each s and t . Note that this implies that $\pi_{(st)} > 0$ ($s \in [n_1]$, $t \in [n_2]$) and $\sum_{s=1}^{n_1} \sum_{t=1}^{n_2} \pi_{(st)} = 1$, as required, since $c_{st} > 0$.

Finally, utilizing some pairing function (see e.g., Smorynski, 1991, Sec. 1.3), we may map every pair $(s, t) \in [n_1] \times [n_2]$ uniquely to a $z \in [n_{[12]}]$, where $n_{[12]} = n_1 n_2$. Using this mapping, we can then write

$$\begin{aligned} \mathbf{m}_{[12]}(\mathbf{x}) &= \sum_{z=1}^{n_{[12]}} \frac{\pi_{[12]z} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{[12]z}, \Sigma_{[12]z})}{\sum_{\zeta=1}^{n_{[12]}} \pi_{[12]\zeta} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{[12]\zeta}, \Sigma_{[12]\zeta})} \mathbf{a}_{[12]z} \\ &= \mathbf{m}_{[12]}(\mathbf{x}; \boldsymbol{\theta}_{[12]}), \end{aligned}$$

where $\boldsymbol{\theta}_{[12]}$ is a parameter vector that contains the unique elements of $\pi_{[12]z}$, $\boldsymbol{\mu}_{[12]z}$, $\Sigma_{[12]z}$, and $\mathbf{a}_{[12]z}$ for each $z \in [n_{[12]}]$. Thus, we have shown that $\mathbf{m}_{[12]} = \mathbf{m}_{[1]} + \mathbf{m}_{[2]}$ is in the class of functions $\mathcal{M}_q^*(\mathbb{X})$.

Proof of Lemma 3

Since $f_{[1]} \in \mathcal{L}_q^*(\mathbb{X})$ and $f_{[2]} \in \mathcal{L}_r^*(\mathbb{X})$, we can write

$$f_{[1]}(\mathbf{y}_{[1]}|\mathbf{x}; \boldsymbol{\theta}_1) = \sum_{z=1}^{n_1} \frac{\pi_{1z} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{1z}, \boldsymbol{\Sigma}_{1z}) \phi_q(\mathbf{y}_{[1]}; \mathbf{a}_{1z}, \mathbf{C}_{1z})}{\sum_{\zeta=1}^{n_1} \pi_{1\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{1\zeta}, \boldsymbol{\Sigma}_{1\zeta})},$$

and

$$f_{[2]}(\mathbf{y}_{[2]}|\mathbf{x}; \boldsymbol{\theta}_2) = \sum_{z=1}^{n_2} \frac{\pi_{2z} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{2z}, \boldsymbol{\Sigma}_{2z}) \phi_r(\mathbf{y}_{[2]}; \mathbf{a}_{2z}, \mathbf{C}_{2z})}{\sum_{\zeta=1}^{n_2} \pi_{2\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{2\zeta}, \boldsymbol{\Sigma}_{2\zeta})},$$

where $\boldsymbol{\theta}_k$ contains the unique elements of $\boldsymbol{\mu}_{kz}$, $\boldsymbol{\Sigma}_{kz}$, \mathbf{a}_{kz} , and \mathbf{C}_{kz} ($z \in [n_k]$; $n_k \in \mathbb{N}$), for each $k \in \{1, 2\}$. Here, $\mathbf{y}^\top = (\mathbf{y}_{[1]}^\top, \mathbf{y}_{[2]}^\top)$, where $\mathbf{y}_{[1]} \in \mathbb{R}^q$ and $\mathbf{y}_{[2]} \in \mathbb{R}^r$.

Next, write

$$f_{[12]}(\mathbf{y}|\mathbf{x}) = \sum_{z=1}^{n_1} \frac{\pi_{1z} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{1z}, \boldsymbol{\Sigma}_{1z}) \phi_q(\mathbf{y}_{[1]}; \mathbf{a}_{1z}, \mathbf{C}_{1z})}{\sum_{\zeta=1}^{n_1} \pi_{1\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{1\zeta}, \boldsymbol{\Sigma}_{1\zeta})} \times \sum_{z=1}^{n_2} \frac{\pi_{2z} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{2z}, \boldsymbol{\Sigma}_{2z}) \phi_r(\mathbf{y}_{[2]}; \mathbf{a}_{2z}, \mathbf{C}_{2z})}{\sum_{\zeta=1}^{n_2} \pi_{2\zeta} \phi_p(\mathbf{x}; \boldsymbol{\mu}_{2\zeta}, \boldsymbol{\Sigma}_{2\zeta})},$$

and make the following mapping for each $s \in [n_1]$ and $t \in [n_2]$: $\bar{\pi}_{(st)} = \pi_{1s} \pi_{2t}$, $\boldsymbol{\Sigma}_{(st)}^{-1} = \boldsymbol{\Sigma}_{1s}^{-1} + \boldsymbol{\Sigma}_{2t}^{-1}$, and $\boldsymbol{\mu}_{(st)} = \boldsymbol{\Sigma}_{(st)} (\boldsymbol{\Sigma}_{1s}^{-1} \boldsymbol{\mu}_{1s} + \boldsymbol{\Sigma}_{2t}^{-1} \boldsymbol{\mu}_{2t})$. Furthermore, for each s and t ,

$$\begin{aligned} \phi_q(\mathbf{y}_{[1]}; \mathbf{a}_{1s}, \mathbf{C}_{1s}) \phi_r(\mathbf{y}_{[2]}; \mathbf{a}_{2t}, \mathbf{C}_{2t}) &= \phi_{q+r} \left(\begin{bmatrix} \mathbf{y}_{[1]} \\ \mathbf{y}_{[2]} \end{bmatrix}; \begin{bmatrix} \mathbf{a}_{1s} \\ \mathbf{a}_{2t} \end{bmatrix}, \begin{bmatrix} \mathbf{C}_{1s} & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_{2t} \end{bmatrix} \right) \\ &= \phi_{q+r}(\mathbf{y}; \mathbf{a}_{(st)}, \mathbf{C}_{(st)}), \end{aligned}$$

specifies a $(q+r)$ -dimensional multivariate normal PDF.

Using Lemma 1, we can write

$$\begin{aligned}
f_{[12]} &= \sum_{s=1}^{n_1} \sum_{t=1}^{n_2} \frac{c_{st} \bar{\pi}_{(st)} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{(st)}, \boldsymbol{\Sigma}_{(st)}) \phi_{q+r}(\mathbf{y}; \mathbf{a}_{(st)}, \mathbf{C}_{(st)})}{\sum_{\xi=1}^{n_1} \sum_{\zeta=1}^{n_2} c_{\xi\zeta} \bar{\pi}_{(\xi\zeta)} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{(\xi\zeta)}, \boldsymbol{\Sigma}_{(\xi\zeta)})} \\
&= \sum_{s=1}^{n_1} \sum_{t=1}^{n_2} \frac{\pi_{(st)} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{(st)}, \boldsymbol{\Sigma}_{(st)}) \phi_{q+r}(\mathbf{y}; \mathbf{a}_{(st)}, \mathbf{C}_{(st)})}{\sum_{\xi=1}^{n_1} \sum_{\zeta=1}^{n_2} \pi_{(\xi\zeta)} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{(\xi\zeta)}, \boldsymbol{\Sigma}_{(\xi\zeta)})},
\end{aligned}$$

where $\pi_{(st)} = c_{st} \bar{\pi}_{(st)} / \sum_{\xi=1}^{n_1} \sum_{\zeta=1}^{n_2} c_{\xi\zeta} \bar{\pi}_{(\xi\zeta)}$, for each s and t .

In a similar manner to the approach from Lemma 2, we may map every pair $(s, t) \in [n_1] \times [n_2]$ uniquely to a $z \in [n_{[12]}]$, where $n_{[12]} = n_1 n_2$. Using this mapping, we can then write

$$\begin{aligned}
f_{[12]}(\mathbf{y}|\mathbf{x}) &= \sum_{z=1}^{n_{[12]}} \frac{\pi_{[12]z} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{[12]z}, \boldsymbol{\Sigma}_{[12]z}) \phi_{q+r}(\mathbf{y}; \mathbf{a}_{(st)}, \mathbf{C}_{(st)})}{\sum_{\zeta=1}^{n_{[12]}} \pi_{[12]\zeta} \phi_d(\mathbf{x}; \boldsymbol{\mu}_{[12]\zeta}, \boldsymbol{\Sigma}_{[12]\zeta})} \\
&= f_{[12]}(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}_{[12]}),
\end{aligned}$$

where $\boldsymbol{\theta}_{[12]}$ is a parameter vector that contains the unique elements of $\pi_{[12]z}$, $\boldsymbol{\mu}_{[12]z}$, $\boldsymbol{\Sigma}_{[12]z}$, $\mathbf{a}_{[12]z}$, and $\mathbf{C}_{[12]z}$, for each $z \in [n_{[12]}]$. Thus, we have shown that $f_{[12]} = f_{[1]} f_{[2]}$ is in the class of functions $\mathcal{L}_{q+r}^*(\mathbb{X})$.

References

- Azzalini, A. & Dalla Valle, A. (1996). the multivariate skew-normal distribution. *Biometrika*, 83, 715–726.
- Bromiley, P. A. (2014). *Products and convolutions of Gaussian probability density functions*. Technical Report 2003-003, TINA-VISION, Manchester.

- Chamroukhi, F. (2016). Robust mixture of experts modeling using the t distribution. *Neural Networks*, 79, 20–36.
- Chamroukhi, F. (2017). Skew t mixture of experts. *Neurocomputing*, 266, 390–408.
- Chamroukhi, F., Mohammed, S., Trabelsi, D., Oukhellou, L., & Amirat, Y. (2013). Joint segmentation of multivariate time series with hidden process regression for human activity recognition. *Neurocomputing*, 120, 633–644.
- Chen, K., Xu, L., & Chi, H. (1999). Improved learning algorithms for mixture of experts in multiclass classification. *Neural Networks*, 12, 1229–1252.
- Cheney, W. & Light, W. (2000). *A Course in Approximation Theory*. Pacific Grove: Brooks/Cole.
- Chiou, J.-M., Chen, Y.-T., & Yang, Y.-F. (2014). Multivariate functional principal component analysis: a normalization approach. *Statistica Sinica*, 24, 1571–1596.
- DasGupta, A. (2008). *Asymptotic Theory Of Statistics And Probability*. New York: Springer.
- Deleforge, A., Forbes, F., & Horaud, R. (2015a). Acoustic space learning for sound-source separation and localization on binaural manifolds. *International Journal of Neural Systems*, 25, 1440003.
- Deleforge, A., Forbes, F., & Horaud, R. (2015b). High-dimensional regression with Gaussian mixtures and partially-latent response variables. *Statistics and Computing*, 25, 893–911.
- Fu, H., Gong, M., Wang, C., & Tao, D. (2018). MoE-SPNet: a mixture of experts scene parsing network. *Pattern Recognition*, 84, 226–236.

- Geweke, J. & Keane, M. (2007). Smoothly mixing regressions. *Journal of Econometrics*, 138, 252–290.
- Grun, B., Kosmidis, I., & Zeileis, A. (2012). Extended beta regression in R: shaken, stirred, mixed, and partitioned. *Journal of Statistical Software*, 48, 1–25.
- 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.
- Ingrassia, S., Minotti, S. C., & Punzo, A. (2014). Model-based clustering via linear cluster-weighted models. *Computational Statistics and Data Analysis*, 71, 159–182.
- 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.
- 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 approximation rate of hierarchical mixtures-of-experts for generalized linear models. *Neural Computation*, 11, 1183–1198.
- Jordan, M. I. & Jacobs, R. A. (1994). Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6, 181–214.

- Jordan, M. I. & Xu, L. (1995). Convergence results for the EM approach to mixtures of experts architectures. *Neural Networks*, 8, 1409–1431.
- Kalliovirta, L., Meitz, M., & Saikkonen, P. (2016). Gaussian mixture vector autoregression. *Journal of Econometrics*, 192, 485–498.
- Masoudnia, S. & Ebrahimpour, R. (2014). Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42, 275–293.
- Mendes, E. F. & Jiang, W. (2012). On convergence rates of mixture of polynomial experts. *Neural Computation*, 24, 3025–3051.
- 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.
- Nguyen, H. D. & Chamroukhi, F. (2018). Practical and theoretical aspects of mixture-of-experts modeling: an overview. *WIREs Data Mining and Knowledge Discovery*, (pp. e1246).
- 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. (2016). Laplace mixture of linear experts. *Computational Statistics and Data Analysis*, 93, 177–191.
- Norets, A. (2010). Approximation of conditional densities by smooth mixtures of regressions. *Annals of Statistics*, 38, 1733–1766.
- Norets, A. & Pati, D. (2017). Adaptive Bayesian estimation of conditional densities. *Econometric Theory*, 33, 980–1012.

- Norets, A. & Pelenis, J. (2012). Bayesian modeling of joint and conditional distributions. *Journal of Econometrics*, 168, 332–346.
- Norets, A. & Pelenis, J. (2014). Posterior consistency in conditional density estimation by covariate dependent mixtures. *Econometric Theory*, 30, 606–646.
- Oden, J. T. & Demkowicz, L. F. (2010). *Applied Functional Analysis*. Boca Raton: CRC Press.
- Pelenis, J. (2014). Bayesian regression with heteroscedastic error density and parametric mean function. *Journal of Econometrics*, 178, 624–638.
- Perthame, E., Forbes, F., & Deleforge, A. (2018). Inverse regression approach to robust nonlinear high-to-low dimensional mapping. *Journal of Multivariate Analysis*, 163, 1–14.
- Pinkus, A. (2015). *Ridge Functions*. Cambridge: Cambridge University Press.
- Pollard, D. (2002). *A User's Guide to Measure Theoretic Probability*. Cambridge: Cambridge University Press.
- Prado, R., Molina, F., & Huerta, G. (2006). Multivariate time series modeling and classification via hierarchical VAR mixture. *Computational Statistics and Data Analysis*, 51, 1445–1462.
- Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., & Dean, J. (2017). Outrageously large neural networks: the sparsely-gated mixture-of-experts layer. In *Proceedings of the International Conference on Learning Representation*.
- Smorynski, C. (1991). *Logical Number Theory I: An Introduction*. Berlin: Springer.
- Stone, M. H. (1948). The generalized Weierstrass approximation theorem. *Mathematical Magazine*, 21, 237–254.

- Wang, L.-X. & Mendel, J. M. (1992). Fuzzy basis functions, universal approximation, and orthogonal least-squares learning. *IEEE Transactions on Neural Networks*, 3, 807–814.
- 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., 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. 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.
- Zhao, T., Chen, Q., Kuang, Z., Yu, J., Zhang, W., & Fan, J. (2018). Deep mixture of diverse experts for large-scale visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41, 1072–1087.