Neurocomputing 266 (2017) 390–408

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Skew *t* mixture of experts

F. Chamroukhi

Normandie Université, UNICAEN, CNRS, Laboratoire de Mathématiques Nicolas Oresme - LMNO, Département de Mathématiques et d'Informatique, Caen 14000, France

ARTICLE INFO

Article history:
Received 5 November 2015
Revised 30 November 2016
Accepted 19 May 2017
Available online 5 July 2017

Communicated by Zidong Wang

Keywords:
Mixture of experts
Skew *t* distribution
EM algorithm
ECM algorithm
Non-linear regression
Model-based clustering

ABSTRACT

Mixture of experts (MoE) is a popular framework in the fields of statistics and machine learning for modeling heterogeneity in data for regression, classification and clustering. MoE for continuous data are usually based on the normal distribution. However, it is known that for data with asymmetric behavior, heavy tails and atypical observations, the use of the normal distribution is unsuitable. We introduce a new robust non-normal mixture of experts modeling using the skew *t* distribution. The proposed skew *t* mixture of experts, named STMoE, handles these issues of the normal mixtures experts regarding possibly skewed, heavy-tailed and noisy data. We develop a dedicated expectation conditional maximization (ECM) algorithm to estimate the model parameters by monotonically maximizing the observed data log-likelihood. We describe how the presented model can be used in prediction and in model-based clustering of regression data. Numerical experiments carried out on simulated data show the effectiveness and the robustness of the proposed model in fitting non-linear regression functions as well as in model-based clustering. Then, the proposed model is applied to the real-world data of tone perception for musical data analysis, and the one of temperature anomalies for the analysis of climate change data. The obtained results confirm the usefulness of the model for practical data analysis applications.

1. Introduction

Mixture of experts (MoE) [31] is a popular framework in the statistics and machine learning fields for modeling heterogeneity in data for regression, classification and clustering. They consist in a fully conditional mixture model where both the mixing proportions, known as the gating functions, and the component densities, known as the experts, are conditional on some input covariates. MoE have been investigated, in their simple form, as well as in their hierarchical form [34] (e.g., Section 5.12 of [44]) for regression and model-based cluster and discriminant analyses and in different application domains. MoE Have also been investigated for rank data [20] and network data [21] with social science applications. A survey on the topic can be found in [22]. A complete review of the MoE models can be found in [65]. MoE for continuous data are usually based on the normal distribution. Along this paper, we will call the MoE using the normal distribution the normal mixture of experts, abbreviated as NMoE. However, it is well-known that the normal distribution is sensitive to outliers. Moreover, for a set of data containing a group or groups of observations with heavy tails or asymmetric behavior, the use of normal experts may be unsuitable and can unduly affect the fit of the MoE model. In this paper, we attempt to overcome these limitations in MoE by proposing a more adapted and robust mixture of experts model that can deal with possibly skewed, heavy-tailed data and with outliers.

Recently, the problem of sensitivity of NMoE to outliers have been considered by [48] where the authors proposed a Laplace mixture of linear experts (LMoLE) for a robust modeling of non-linear regression data. The model parameters are estimated by maximizing the observed-data likelihood via a minorization-maximization (MM) algorithm. Here, we propose an alternative MoE model, by relaying on other non-normal distribution that generalizes the normal distribution, that is, the skew-*t* distribution introduced quite recently by [4]. We call the proposed MoE model the skew-*t* mixture of experts (STMoE). One may use the *t* distribution, as in the *t* mixture of experts (TMoE) proposed by [9,10] which provides a natural robust extension of the normal distribution to model data with more heavy tails and to deal with possible outliers. The robustness of the *t* distribution may however be not sufficient in the presence of asymmetric observations. In mixture modeling, to deal with this issue regarding skewed data, [41] proposed the univariate skew-*t* mixture model that allows for accommodation of both skewness and thick tails in the data, by relying on the skew-*t* distribution [4]. For the general multivariate case using skew-*t* mixtures, one can refer to [36–38,40,51], and recently, the unifying framework for previous restricted and unrestricted skew-*t* mixtures, using the CFUST distribution [39]. We

E-mail addresses: faicel.chamroukhi@unicaen.fr, faicel.chamroukhi@gmail.com

http://dx.doi.org/10.1016/j.neucom.2017.05.044
0925-2312/

note that the STMoE model presented in this paper is more general and more robust compared to the TMoE model presented in Chamroukhi-TMoE. As discussed in Section 3.2, the TMoE model can be seen as a particular case of the STMoE model when the skewness parameter goes to zero. The presented model is then more robust as it is able to accommodate more complex data distribution where the data or a group of data are skewed and affected by atypical observations. The TMoE model may fail in this context.

The inference in the previously described approaches is performed by maximum likelihood estimation via the expectation-maximization (EM) algorithm or its extensions [15,43], in particular the expectation conditional maximization (ECM) algorithm [46]. [18] have also considered the Bayesian inference framework for namely the skew-*t* mixtures.

For the regression context, the robust modeling of regression data has been studied namely by [5,29,63] who considered a mixture of linear regressions using the *t* distribution. In the same context of regression, [58] proposed the mixture of Laplace regressions, which has been then extended by [48] to the case of mixture of experts, by introducing the Laplace mixture of linear experts (LMoLE). Recently, [66] introduced the scale mixtures of skew-normal distributions for robust mixture regressions. However, unlike our proposed STMoE model, the regression mixture models of [5,29,58,63,66] do not consider conditional mixing proportions, that is, mixing proportions depending on some input variables, as in the case of mixture of experts, which we investigate here. In addition, the approaches of [5,29,58,63] do not consider both the problem of robustness to outliers together with the one of dealing with possibly asymmetric data.

Here we consider the mixture of experts framework for non-linear regression problems and model-based clustering of regression data, and we attempt to overcome the limitations of the NMoE model for dealing with asymmetric, heavy-tailed data and which may contain outliers. We investigate the use of the skew *t* distribution for the experts, rather than the commonly used normal distribution. We propose the skew-*t* mixture of experts (STMoE) model that allows for accommodation of both skewness and heavy tails in the data and which is robust to outliers. This model corresponds to an extension of the unconditional skew *t* mixture model [41], to the mixture of experts (MoE) framework, where the mixture means are regression functions and the mixing proportions are also covariate-varying.

For the model inference, we develop a dedicated expectation conditional maximization (ECM) algorithm to estimate the model parameters by monotonically maximizing the observed data log-likelihood. The expectation-maximization algorithm and its extensions [15,43] are indeed very popular and successful estimation algorithms for mixture models in general and for mixture of experts in particular. Moreover, the EM algorithm for MoE has been shown by [47] to be monotonically maximizing the MoE likelihood. The authors have showed that the EM (with Iteratively Reweighted Least Squares (IRLS) in this case) algorithm has stable convergence and the log-likelihood is monotonically increasing when a learning rate smaller than one is adopted for the IRLS procedure within the M-step of the EM algorithm. They have further proposed an expectation conditional maximization (ECM) algorithm to train MoE, which also has desirable numerical properties. The MoE has also been considered in the Bayesian framework, for example one can cite the Bayesian MoE [61,62] and the Bayesian hierarchical MoE [7]. Related MoE considering the asymmetric *t* distribution in a Bayesian framework is proposed by [67]. Beyond the Bayesian parametric framework, the MoE models have also been investigated within the Bayesian non-parametric framework. We cite for example the Bayesian non-parametric MoE model [54] and the Bayesian non-parametric hierarchical MoE approach of [30] using Gaussian Processes experts for regression. For further models on mixture of experts for regression, the reader can be referred to for example the book of [57]. In this paper, we investigate semi-parametric models under the maximum likelihood estimation framework.

The remainder of this paper is organized as follows. In Section 2 we briefly recall the normal MoE framework. In Section Then, in Section 3, we present the STMoE model and in Section 4 the parameter estimation technique using the ECM algorithm. We then investigate in Section 5 the use of the proposed model for non-linear regression and for prediction. We also show in Section 6 how the model can be used in a model-based clustering prospective. In Section 7, we discuss the model selection. Section 8 is dedicated to the experimental study to assess the proposed model. Finally, in Section 9, conclusions are drawn and we open a future work.

2. Mixture of experts for continuous data

Mixtures of experts [31,34] are used in a variety of contexts including regression, classification and clustering. Here, we consider the MoE framework for fitting (non-linear) regression functions and clustering of univariate continuous data . The aim of regression is to explore the relationship of an observed random variable Y given a covariate vector $\boldsymbol{X} \in \mathbb{R}^p$ via conditional density functions for $Y|\boldsymbol{X} = \boldsymbol{x}$ of the form $f(y|\boldsymbol{x})$, rather than only exploring the unconditional distribution of Y. Thanks to their great flexibility, mixture models [44] has took much attention for non-linear regression problems and we distinguish in particular the classical mixture of regressions model [16,19,27,33,52,53,59,60] and mixture of experts for regression analysis [31,34,64]. The univariate mixture of regressions model assumes that the observed pairs of data $(\boldsymbol{x}, y)$ where $y \in \mathbb{R}$ is the response for some covariate $\boldsymbol{x} \in \mathbb{R}^p$, are generated from K regression functions and are governed by a hidden categorical random variable Z indicating from which component each observation is generated. Thus, the mixture of regressions model decomposes the nonlinear regression model density $f(y|\boldsymbol{x})$ into a convex weighted sum of K regression component models $f_k(y|\boldsymbol{x})$ and can be defined as follows:

$$f(y|\boldsymbol{x}; \boldsymbol{\Psi}) = \sum_{k=1}^{K} \pi_k f_k\left(y|\boldsymbol{x}; \boldsymbol{\Psi}_k\right) \tag{1}$$

where the π_k's defined by $\pi_k = \mathbb{P}(Z = k)$ and represent the non-negative mixing proportions that sum to 1. The model parameter vector is given by $\boldsymbol{\Psi} = (\pi_1, \ldots, \pi_{K-1}, \boldsymbol{\Psi}_1^T, \ldots, \boldsymbol{\Psi}_K^T)^T$, $\boldsymbol{\Psi}_k$ being the parameter vector of the kth component density.

2.1. The mixture of experts (MoE) model

Although similar, the mixture of experts [31] differ from regression mixture models in many aspects. One of the main differences is that the MoE model consists in a fully conditional mixture while in the regression mixture, only the component densities are conditional. Indeed, the mixing proportions are constant for the regression mixture, while in the MoE, they are modeled as a function of the inputs, generally modeled by logistic or a softmax function. Mixture of experts (MoE) for regression analysis [31,34] extend the model (1) by modeling the mixing proportions as function of some covariates $\boldsymbol{r} \in \mathbb{R}^q$. The mixing proportions, known as the gating functions in the context of MoE, are modeled by the multinomial logistic model and are defined by:

$$\pi_k(\boldsymbol{r}; \boldsymbol{\alpha}) = \mathbb{P}(Z = k|\boldsymbol{r}; \boldsymbol{\alpha}) = \frac{\exp\left(\boldsymbol{\alpha}_k^T \boldsymbol{r}\right)}{\sum_{\ell=1}^{K} \exp\left(\boldsymbol{\alpha}_\ell^T \boldsymbol{r}\right)} \tag{2}$$

where $\boldsymbol{r} \in \mathbb{R}^q$ is a covariate vector, $\boldsymbol{\alpha}_k$ is the q-dimensional coefficients vector associated with $\boldsymbol{r}$ and $\boldsymbol{\alpha} = (\boldsymbol{\alpha}_1^T, \ldots, \boldsymbol{\alpha}_{K-1}^T)^T$ is the pa-

rameter vector of the logistic model, with $\boldsymbol{\alpha}_K$ being the null vector. Thus, the MoE model consists in a fully conditional mixture model where both the mixing proportions (the gating functions) and the component densities (the experts) are conditional on predictors (respectively $\boldsymbol{r}$ and $\boldsymbol{x}$).

2.2. The normal mixture of experts (NMoE) model and maximum likelihood estimation

In the case of mixture of experts for regression, it is usually assumed that the experts are normal, that is, follow a normal distribution. A K-component normal mixture of experts (NMoE) ($K>1$) has the following formulation:

$$f(y|\boldsymbol{r},\boldsymbol{x};\boldsymbol{\Psi}) = \sum_{k=1}^{K} \pi_k(\boldsymbol{r};\boldsymbol{\alpha})\mathrm{N}\left(y;\mu(\boldsymbol{x};\boldsymbol{\beta}_k),\sigma_k^2\right) \tag{3}$$

which involves, in the semi-parametric case, component means defined as parametric (non-)linear regression functions $\mu(\boldsymbol{x};\boldsymbol{\beta}_k)$.

The NMoE model parameters are estimated by maximizing the observed data log-likelihood by using the EM algorithm [15,31,34,35,43,47]. Suppose we observe an i.i.d sample of n individuals $(y_1,\ldots,y_n)$ with their respective associated covariates $(\boldsymbol{x}_1,\ldots,\boldsymbol{x}_n)$ and $(\boldsymbol{r}_1,\ldots,\boldsymbol{x}_r)$. Then, under the NMoE model, the observed data log-likelihood for the parameter vector $\boldsymbol{\Psi}$ is given by:

$$\log L(\boldsymbol{\Psi}) = \sum_{i=1}^{n} \log \sum_{k=1}^{K} \pi_k(\boldsymbol{r}_i;\boldsymbol{\alpha})\mathrm{N}\left(y_i;\mu(\boldsymbol{x};\boldsymbol{\beta}_k),\sigma_k^2\right). \tag{4}$$

The E-Step at the mth iteration of the EM algorithm for the NMoE model requires the calculation of the following posterior probability that the individual $(y_i,\ \boldsymbol{x}_i,\ \boldsymbol{r}_i)$ belongs to expert k, given a parameter estimation $\boldsymbol{\Psi}^{(m)}$:

$$\begin{aligned}\tau_{ik}^{(m)} &= \mathbb{P}(Z_i=k|y_i,\boldsymbol{x}_i,\boldsymbol{r}_i;\boldsymbol{\Psi}^{(m)}) \\ &= \frac{\pi_k(\boldsymbol{r};\boldsymbol{\alpha}^{(m)})\mathrm{N}\left(y_i;\mu_k(\boldsymbol{x}_i;\beta_k^{(m)}),\sigma_k^{2\,(m)}\right)}{f(y_i|\boldsymbol{r}_i,\boldsymbol{x}_i;\boldsymbol{\Psi}^{(m)})}.\end{aligned} \tag{5}$$

Then, the M-step calculates the parameter update $\boldsymbol{\Psi}^{(m+1)}$ by maximizing the well-known Q-function (the expected complete-data log-likelihood), that is:

$$\boldsymbol{\Psi}^{(m+1)} = \arg\max_{\boldsymbol{\Psi}\in\boldsymbol{\Omega}} Q(\boldsymbol{\Psi};\boldsymbol{\Psi}^{(m)}) \tag{6}$$

where $\boldsymbol{\Omega}$ is the parameter space. For example, in the case of normal mixture of linear experts (NMoLE) where each expert's mean has the flowing linear form:

$$\mu(\boldsymbol{x};\boldsymbol{\beta}_k) = \boldsymbol{\beta}_k^T\boldsymbol{x}, \tag{7}$$

where $\boldsymbol{\beta}_k \in \mathbb{R}^p$ is the vector of regression coefficients of component k, the updates for each of the expert component parameters consist in analytically solving a weighted Gaussian linear regression problem and are given by:

$$\boldsymbol{\beta}_k^{(m+1)} = \left[\sum_{i=1}^{n} \tau_{ik}^{(m)}\boldsymbol{x}_i\boldsymbol{x}_i^T\right]^{-1} \sum_{i=1}^{n} \tau_{ik}^{(q)} y_i\boldsymbol{x}_i, \tag{8}$$

$$\sigma_k^{2\,(m+1)} = \frac{\sum_{i=1}^{n} \tau_{ik}^{(m)}\left(y_i - \boldsymbol{\beta}_k^{T\,(m+1)}\boldsymbol{x}_i\right)^2}{\sum_{i=1}^{n}\tau_{ik}^{(m)}}. \tag{9}$$

For the mixing proportions, the parameter update $\boldsymbol{\alpha}^{(m+1)}$ cannot however be obtained in a closed form. It is calculated by IRLS [11,13,23,31,34].

However, the normal distribution is not adapted to deal with asymmetric and heavy tailed data. It is also known that the normal distribution is sensitive to outliers. In the proposal, we address these issues regarding the skewness, heavy tails and atypical observations in the data, by proposing a robust MoE modeling by using the skew-t distribution, recently introduced by [4], for the expert components rather than the usually used normal one. The proposed skew t mixture of experts (STMoE) allows for simultaneously accommodating asymmetry and heavy tails in the data and is also robust to outliers.

3. The skew t mixture of experts (STMoE) model

The proposed skew t mixture of experts (STMoE) model is a MoE model in which the expert components have a skew-t density, rather than the standard normal one as in the NMoE model. The skew-t distribution [4], which is a robust generalization the skew-normal distribution [2,3], as well as its stochastic and hierarchical representations, which will be used to define the proposed STMoE model, are recalled in the following section.

3.1. The skew t distribution

Let us denote by $t_\nu(.)$ and $T_\nu(.)$ respectively the pdf and cdf of the standard t distribution with degrees of freedom ν. The skew t distribution, introduced by [4], can be characterized as follows. Let U be an univariate random variable with a standard skew-normal distribution $U\sim\mathrm{SN}(0,\ 1,\ \lambda)$ (which can be shortened as $U\sim\mathrm{SN}(\lambda)$) with pdf given by (A.1). The skew-normal distribution is recalled in Appendix A. Then, let W be an univariate random variable independent of U and following the Gamma distribution, that is, $W\sim\mathrm{Gamma}(\frac{\nu}{2},\frac{\nu}{2})$. A random variable Y having the following representation:

$$Y = \mu + \sigma\frac{U}{\sqrt{W}} \tag{10}$$

follows the skew t distribution $\mathrm{ST}(\mu,\ \sigma^2,\ \lambda,\ \nu)$ with location parameter μ, scale parameter σ, skewness parameter λ and degrees of freedom ν, whose density is defined by:

$$f(y;\mu,\sigma^2,\lambda,\nu) = \frac{2}{\sigma}\ t_\nu(d_y)\ T_{\nu+1}\left(\lambda\ d_y\sqrt{\frac{\nu+1}{\nu+d_y^2}}\right) \tag{11}$$

where $d_y = \frac{y-\mu}{\sigma}$. From the hierarchical distribution of the skew-normal (A.3), a further hierarchical representation of the stochastic representation (10) of the skew t distribution is given by:

$$\begin{aligned} Y_i|u_i,w_i &\sim \mathrm{N}\left(\mu+\delta|u_i|,\frac{1-\delta^2}{w_i}\sigma^2\right), \\ U_i|w_i &\sim \mathrm{N}\left(0,\frac{\sigma^2}{w_i}\right), \\ W_i &\sim \mathrm{Gamma}\left(\frac{\nu}{2},\frac{\nu}{2}\right). \end{aligned} \tag{12}$$

3.2. The skew t mixture of experts (STMoE) model

The skew proposed t mixture of experts (STMoE) model extends the skew t mixture model, which was first introduced by [41], to the MoE framework. In the skew-t mixture model of [41], the mixing proportions and the components means are constant, that is, they are not predictor-depending. In the proposed STMoE, however, we consider skew-t expert components in which both the mixing proportions and the mixture component means are

predictor-depending. More specifically, we use polynomial regressors for the components, as well as multinomial logistic regressors for the mixing proportions. A K-component mixture of skew t experts (STMoE) is therefore defined by:

$$f(y|\boldsymbol{r},\boldsymbol{x};\boldsymbol{\Psi}) = \sum_{k=1}^{K} \pi_k(\boldsymbol{r};\boldsymbol{\alpha})\, \mathrm{ST}\big(y;\mu(\boldsymbol{x};\boldsymbol{\beta}_k),\sigma_k^2,\lambda_k,\nu_k\big). \quad (13)$$

The parameter vector of the STMoE model is $\boldsymbol{\Psi} = (\boldsymbol{\alpha}_1^T,\ldots,\boldsymbol{\alpha}_{K-1}^T,\boldsymbol{\Psi}_1^T,\ldots,\boldsymbol{\Psi}_K^T)^T$ where $\boldsymbol{\Psi}_k = (\boldsymbol{\beta}_k^T,\sigma_k^2,\lambda_k,\nu_k)^T$ is the parameter vector for the kth skew t expert component whose density is defined by

$$f\big(y|\boldsymbol{x};\mu(\boldsymbol{x};\boldsymbol{\beta}_k),\sigma^2,\lambda,\nu\big) = \frac{2}{\sigma}\, t_\nu(d_y(\boldsymbol{x}))\; T_{\nu+1}\left(\lambda\; d_y(\boldsymbol{x})\sqrt{\frac{\nu+1}{\nu+d_y^2(\boldsymbol{x})}}\right) \quad (14)$$

where $d_y(\boldsymbol{x}) = \frac{y-\mu(\boldsymbol{x};\boldsymbol{\beta}_k)}{\sigma}$.

It can be seen that, when the robustness parameter $\nu_k \to \infty$ for each k, the STMoE model (13) reduces to a skew-normal mixture of experts model (SNMoE) (see [9]). On the other hand, if the skewness parameter $\lambda_k = 0$ for each k, the STMoE model reduces to the t mixture of experts model (TMoE) (e.g., see [9,10]). Moreover, when $\nu_k \to \infty$ and $\lambda_k = 0$ for each k, the STMoE approaches the standard NMoE model (3). This therefore makes the STMoE very flexible as it generalizes the previously described MoE models to accommodate situations with asymmetry, heavy tails, and outliers.

3.3. Hierarchical representation of the STMoE model

By introducing the binary latent component-indicators Z_{ik} such that $Z_{ik} = 1$ iff $Z_i = k$, Z_i being the hidden class label of the ith observation, a hierarchical model for the STMoE model can be derived as follows. From the hierarchical representation (12) of the skew t distribution, a hierarchical model for the proposed STMoE model (13) can be derived from its stochastic representation (B.2) given in Appendix B, and is as follows:

$$\begin{aligned} Y_i|u_i,w_i,Z_{ik}=1,\boldsymbol{x}_i &\sim \mathrm{N}\left(\mu(\boldsymbol{x}_i;\boldsymbol{\beta}_k)+\delta_k|u_i|,\frac{1-\delta_k^2}{w_i}\sigma_k^2\right),\\ U_i|w_i,Z_{ik}=1 &\sim \mathrm{N}\left(0,\frac{\sigma_k^2}{w_i}\right), \quad (15)\\ W_i|Z_{ik}=1 &\sim \mathrm{Gamma}\left(\frac{\nu_k}{2},\frac{\nu_k}{2}\right)\\ \boldsymbol{Z}_i|\boldsymbol{r}_i &\sim \mathrm{Mult}\big(1;\pi_1(\boldsymbol{r}_i;\boldsymbol{\alpha}),\ldots,\pi_K(\boldsymbol{r}_i;\boldsymbol{\alpha})\big). \end{aligned}$$

The variables U_i and W_i are treated as hidden in this hierarchical representation, which facilitates the inference scheme and will be used to derive the maximum likelihood estimation of the STMoE model parameters $\boldsymbol{\Psi}$ by using the ECM algorithm.

3.4. Identifiability of the STMoE model

[32] have established that ordered, initialized, and irreducible MoEs are identifiable. Ordered implies that there exist a certain ordering relationship on the experts parameters $\boldsymbol{\Psi}_k$ such that $(\boldsymbol{\alpha}_1^T,\boldsymbol{\Psi}_1^T)^T \prec \ldots \prec (\boldsymbol{\alpha}_K^T,\boldsymbol{\Psi}_K^T)^T$; initialized implies that $\mathbf{w}_K$, the parameter vector of the Kth logistic proportion, is the null vector, and irreducible implies that $\boldsymbol{\Psi}_k \neq \boldsymbol{\Psi}_{k'}$ for any $k \neq k'$. For the proposed STMoE, ordered implies that there exist a certain ordering relationship such that $(\boldsymbol{\beta}_1^T,\sigma_1^2,\lambda_1,\nu_1)^T \prec \ldots \prec (\boldsymbol{\beta}_K^T,\sigma_K^2,\lambda_K,\nu_K)^T$; initialized implies that $\mathbf{w}_K$ is the null vector, as assumed in the model, and finally, irreducible implies that if $k \neq k'$, then one of the following conditions holds: $\boldsymbol{\beta}_k \neq \boldsymbol{\beta}_{k'}$, $\sigma_k \neq \sigma_{k'}$, $\lambda_k \neq \lambda_{k'}$ or $\nu_k \neq \nu_{k'}$. Then, we can establish the identifiability of ordered and initialized irreducible STMoE models by applying Lemma 2 of [32], which requires the validation of the following nondegeneracy condition. The set $\{\mathrm{ST}(y;\mu(\boldsymbol{x};\boldsymbol{\beta}_1),\sigma_1^2,\lambda_1,\nu_1),\ldots,\mathrm{ST}(y;\mu(\boldsymbol{x};\boldsymbol{\beta}_{4K}),\sigma_{4K}^2,\lambda_{4K},\nu_{4K})\}$ contains $4K$ linearly independent functions of y, for any $4K$ distinct quadruplet $(\mu(\boldsymbol{x};\boldsymbol{\beta}_k),\sigma_k^2,\lambda_k,\nu_k)$ for $k = 1,\ldots,4K$. Thus, via Lemma 2 of [32] we have any ordered and initialized irreducible STMoE is identifiable.

4. Maximum likelihood estimation of the STMoE model

The unknown parameter vector $\boldsymbol{\Psi}$ of the STMoE model is estimated by maximizing the following observed-data log-likelihood given an observed i.i.d sample of n observations, that is, the responses $(y_1,\ldots,y_n)$ and the corresponding predictors $(\boldsymbol{x}_1,\ldots,\boldsymbol{x}_n)$ and $(\boldsymbol{r}_1,\ldots,\boldsymbol{r}_n)$:

$$\log L(\boldsymbol{\Psi}) = \sum_{i=1}^{n} \log \sum_{k=1}^{K} \pi_k(\boldsymbol{r}_i;\boldsymbol{\alpha})\mathrm{ST}\big(y;\mu(\boldsymbol{x}_i;\boldsymbol{\beta}_k),\sigma_k^2,\lambda_k,\nu_k\big). \quad (16)$$

We perform this iteratively by a dedicated ECM algorithm. The complete data consist of the observations as well as the latent variables $(u_1,\ldots,u_n)$ and $(w_1,\ldots,w_n)$, and the latent component labels $(z_1,\ldots,z_n)$. Then, from the hierarchical representation of the STMoE (15), the complete-data log-likelihood of $\boldsymbol{\Psi}$ is given by:

$$\begin{aligned} \log L_c(\boldsymbol{\Psi}) &= \sum_{i=1}^{n}\sum_{k=1}^{K} Z_{ik}\Big[\log\big(\mathbb{P}(Z_i=k|\boldsymbol{r}_i)\big)+\log\big(f(w_i|Z_{ik}=1)\big)+\\ &\qquad \log\big(f(u_i|w_i,Z_{ik}=1)\big)+\log\big(f(y_i|u_i,Z_{ik}=1,\boldsymbol{x}_i)\big)\Big]\\ &= \log L_{1c}(\boldsymbol{\alpha}) + \sum_{k=1}^{K}\Big[\log L_{2c}(\boldsymbol{\theta}_k)+\log L_{3c}(\nu_k)\Big] \quad (17) \end{aligned}$$

where $\boldsymbol{\theta}_k = (\boldsymbol{\beta}_k^T,\sigma_k^2,\lambda_k)^T$ and

$$\begin{aligned} \log L_{1c}(\boldsymbol{\alpha}) &= \sum_{i=1}^{n}\sum_{k=1}^{K} Z_{ik}\log\pi_k(\boldsymbol{r}_i;\boldsymbol{\alpha}),\\ \log L_{2c}(\boldsymbol{\theta}_k) &= \sum_{i=1}^{n} Z_{ik}\Big[-\log(2\pi)-\log(\sigma_k^2)-\frac{1}{2}\log(1-\delta_k^2)\\ &\qquad -\frac{w_i\, d_{ik}^2}{2(1-\delta_k^2)}+\frac{w_i\, u_i\, \delta_k\, d_{ik}}{(1-\delta_k^2)\sigma_k}-\frac{w_i\, u_i^2}{2(1-\delta_k^2)\sigma_k^2}\Big],\\ \log L_{3c}(\nu_k) &= \sum_{i=1}^{n} Z_{ik}\Big[-\log\Gamma\Big(\frac{\nu_k}{2}\Big)+\Big(\frac{\nu_k}{2}\Big)\log\Big(\frac{\nu_k}{2}\Big)\\ &\qquad +\Big(\frac{\nu_k}{2}\Big)\log(w_i)-\Big(\frac{\nu_k}{2}\Big)w_i\Big]. \end{aligned}$$

4.1. The ECM algorithm for the STMoE model

The ECM algorithm for the STMoE model starts with an initial parameter vector $\boldsymbol{\Psi}^{(0)}$ and alternates between the E- and CM-steps until convergence.

4.2. E-Step

The E-Step of the CEM algorithm for the STMoE calculates the Q-function, that is the conditional expectation of the complete-data log-likelihood (17), given the observed data $\{y_i,\boldsymbol{x}_i,\boldsymbol{r}_i\}_{i=1}^n$ and a current parameter estimation $\boldsymbol{\Psi}^{(m)}$, m being the current iteration. From (17), the Q-function is given by:

$$Q(\boldsymbol{\Psi};\boldsymbol{\Psi}^{(m)}) = Q_1(\boldsymbol{\alpha};\boldsymbol{\Psi}^{(m)}) + \sum_{k=1}^{K}\Big[Q_2(\boldsymbol{\theta}_k,\boldsymbol{\Psi}^{(m)}) + Q_3(\nu_k,\boldsymbol{\Psi}^{(m)})\Big], \tag{18}$$

where

$$Q_1(\boldsymbol{\alpha};\boldsymbol{\Psi}^{(m)}) = \sum_{i=1}^{n}\sum_{k=1}^{K}\tau_{ik}^{(m)}\log\pi_k(\boldsymbol{r}_i;\boldsymbol{\alpha}),$$

$$Q_2(\boldsymbol{\theta}_k;\boldsymbol{\Psi}^{(m)}) = \sum_{i=1}^{n}\tau_{ik}^{(m)}\Big[-\log(2\pi)-\log(\sigma_k^2)-\frac{1}{2}\log(1-\delta_k^2) - \frac{w_{ik}^{(m)}\,d_{ik}^2}{2(1-\delta_k^2)} + \frac{\delta_k\,d_{ik}\,e_{1,ik}^{(m)}}{(1-\delta_k^2)\sigma_k} - \frac{e_{2,ik}^{(m)}}{2(1-\delta_k^2)\sigma_k^2}\Big],$$

$$Q_3(\nu_k;\boldsymbol{\Psi}^{(m)}) = \sum_{i=1}^{n}\tau_{ik}^{(m)}\Big[-\log\Gamma\Big(\frac{\nu_k}{2}\Big)+\Big(\frac{\nu_k}{2}\Big)\log\Big(\frac{\nu_k}{2}\Big) - \Big(\frac{\nu_k}{2}\Big)\,w_{ik}^{(m)}+\Big(\frac{\nu_k}{2}\Big)e_{3,ik}^{(m)}\Big].$$

It can be seen that computing the Q-function requires the following conditional expectations:

$$\begin{aligned}
\tau_{ik}^{(m)} &= \mathbb{E}_{\boldsymbol{\Psi}^{(m)}}[Z_{ik}|y_i,\boldsymbol{x}_i,\boldsymbol{r}_i],\\
w_{ik}^{(m)} &= \mathbb{E}_{\boldsymbol{\Psi}^{(m)}}[W_i|y_i,Z_{ik}=1,\boldsymbol{x}_i,\boldsymbol{r}_i],\\
e_{1,ik}^{(m)} &= \mathbb{E}_{\boldsymbol{\Psi}^{(m)}}[W_iU_i|y_i,Z_{ik}=1,\boldsymbol{x}_i,\boldsymbol{r}_i],\\
e_{2,ik}^{(m)} &= \mathbb{E}_{\boldsymbol{\Psi}^{(m)}}\big[W_iU_i^2|y_i,Z_{ik}=1,\boldsymbol{x}_i,\boldsymbol{r}_i\big],\\
e_{3,ik}^{(m)} &= \mathbb{E}_{\boldsymbol{\Psi}^{(m)}}[\log(W_i)|y_i,Z_{ik}=1,\boldsymbol{x}_i,\boldsymbol{r}_i].
\end{aligned}$$

Following the expressions of these conditional expectations given namely in the case of the standard skew t mixture model [41], the conditional expectations for the case of the proposed STMoE model can be expressed similarly as:

$$\tau_{ik}^{(m)} = \frac{\pi_k(\boldsymbol{r};\boldsymbol{\alpha}^{(m)})\;\mathrm{ST}\Big(y_i;\mu(\boldsymbol{x}_i;\boldsymbol{\beta}_k^{(m)}),\sigma_k^{2(m)},\lambda_k^{(m)},\nu_k^{(m)}\Big)}{f(y_i|\boldsymbol{r}_i,\boldsymbol{x}_i;\boldsymbol{\Psi}^{(m)})}, \tag{19}$$

$$w_{ik}^{(m)} = \left(\frac{\nu_k^{(m)}+1}{\nu_k^{(m)}+d_{ik}^{2\,(m)}}\right)\times\frac{T_{\nu_k^{(m)}+3}\left(M_{ik}^{(m)}\sqrt{\frac{\nu_k^{(m)}+3}{\nu_k^{(m)}+1}}\right)}{T_{\nu_k^{(m)}+1}\big(M_{ik}^{(m)}\big)}, \tag{20}$$

where $M_{ik}^{(m)} = \lambda_k^{(m)}\;d_{ik}^{(m)}\sqrt{\frac{\nu_k^{(m)}+1}{\nu_k^{(m)}+d_{ik}^{2\,(m)}}}$,

$$e_{1,ik}^{(m)} = \delta_k^{(m)}\Big(y_i-\mu_k(\boldsymbol{x}_i;\boldsymbol{\beta}^{(m)})\Big)w_{ik}^{(m)} + \left[\frac{\sqrt{1-\delta_k^{2\,(m)}}}{\pi\,f(y_i|\boldsymbol{r}_i,\boldsymbol{x}_i;\boldsymbol{\Psi}^{(m)})}\left(\frac{d_{ik}^{2\,(m)}}{\nu_k^{(m)}(1-\delta_k^{2\,(m)})}+1\right)^{-(\frac{\nu_k^{(m)}}{2}+1)}\right], \tag{21}$$

$$e_{2,ik}^{(m)} = \delta_k^{2(m)}\Big(y_i-\mu_k(\boldsymbol{x}_i;\boldsymbol{\beta}^{(m)})\Big)^2 w_{ik}^{(m)} + \left[\Big(1-\delta_k^{2(m)}\Big)\sigma_k^{2(m)} + \frac{\delta_k^{(m)}\Big(y_i-\mu_k(\boldsymbol{x}_i;\boldsymbol{\beta}^{(m)})\Big)\sqrt{1-\delta_k^{2\,(m)}}}{\pi\,f(y_i|\boldsymbol{r}_i,\boldsymbol{x}_i;\boldsymbol{\Psi}^{(m)})} \times\left(\frac{d_{ik}^{2\,(m)}}{\nu_k^{(m)}(1-\delta_k^{2\,(m)})}+1\right)^{-(\frac{\nu_k^{(m)}}{2}+1)}\right], \tag{22}$$

$$e_{3,ik}^{(m)} = w_{ik}^{(m)} - \log\left(\frac{\nu_k^{(m)}+d_{ik}^{2\,(m)}}{2}\right) - \left(\frac{\nu_k^{(m)}+1}{\nu_k^{(m)}+d_{ik}^{2\,(m)}}\right) + \psi\left(\frac{\nu_k^{(m)}+1}{2}\right) + \frac{\lambda_k^{(m)}d_{ik}^{(m)}\Big(d_{ik}^{2\,(m)}-1\Big)}{\sqrt{\Big(\nu_k^{(m)}+1\Big)\Big(\nu_k^{(m)}+d_{ik}^{2\,(m)}\Big)^3}}\times\frac{t_{\nu_k^{(m)}+1}\big(M_{ik}^{(m)}\big)}{T_{\nu_k^{(m)}+1}\big(M_{ik}^{(m)}\big)}. \tag{23}$$

We note that, for (23), we adopted a one-step-late (OSL) approach to compute the conditional expectation $e_{3,ik}^{(m)}$ as described in [38], by setting the integral part in the expression of the corresponding conditional expectation given in [41] to zero, rather than using a Monte Carlo approximation. We also mention that, for the multivariate skew t mixture models, recently [39] presented a series-based truncation approach, which exploits an exact representation of this conditional expectation and which can also be used in place of (23).

4.3. M-Step

The M-step maximizes the Q-function (18) with respect to $\boldsymbol{\Psi}$ and provides the parameter vector update $\boldsymbol{\Psi}^{(m+1)}$. From (18), it can be seen that the maximization of Q can be performed by separately maximizing Q_1 with respect to the parameters $\boldsymbol{\alpha}$ of the mixing proportions, and for each expert k ($k=1,\ldots,K$), Q_2 with respect to $(\boldsymbol{\beta}_k^T,\sigma_k^2)^T$ and λ_k, and Q_3 with respect to ν_k. The maximization of Q_2 and Q_3 is carried out by conditional maximization (CM) steps by updating $(\boldsymbol{\beta}_k,\sigma_k^2)$ and then updating $(\lambda,\ \nu_k)$ with the given updated parameters. This leads to the following CM steps. On the $(m+1)$th iteration of the M-step, the STMoE model parameters are updated as follows.

CM-Step 1. Calculate $\boldsymbol{\alpha}^{(m+1)}$ maximizing the function $Q_1(\boldsymbol{\alpha};\boldsymbol{\Psi}^{(m)})$:

$$\boldsymbol{\alpha}^{(m+1)} = \arg\max_{\boldsymbol{\alpha}} Q_1(\boldsymbol{\alpha};\boldsymbol{\Psi}^{(m)}). \tag{24}$$

Contrarily to the case of the standard mixture model and mixture of regression models, this maximization in the case of the proposed STMoE does not exist in closed form. It is performed iteratively by Iteratively Reweighted Least Squares (IRLS).

The Iteratively Reweighted Least Squares (IRLS) algorithm. The IRLS algorithm is used to maximize $Q_1(\boldsymbol{\alpha},\boldsymbol{\Psi}^{(m)})$ with respect to the parameter vector $\boldsymbol{\alpha}$ of the softmax function in the M step at each iteration m of the ECM algorithm. The IRLS is a Newton–Raphson algorithm, which consists in starting with a vector $\boldsymbol{\alpha}^{(0)}$, and, at the $l+1$ iteration, updating the estimation of $\boldsymbol{\alpha}$ as follows:

$$\boldsymbol{\alpha}^{(l+1)} = \boldsymbol{\alpha}^{(l)} - \left[\frac{\partial^2 Q_1(\boldsymbol{\alpha},\boldsymbol{\Psi}^{(m)})}{\partial\boldsymbol{\alpha}\partial\boldsymbol{\alpha}^T}\right]^{-1}_{\boldsymbol{\alpha}=\boldsymbol{\alpha}^{(l)}}\frac{\partial Q_1(\boldsymbol{\alpha},\boldsymbol{\Psi}^{(m)})}{\partial\boldsymbol{\alpha}}\bigg|_{\boldsymbol{\alpha}=\boldsymbol{\alpha}^{(l)}} \tag{25}$$

where $\frac{\partial^2 Q_1(\boldsymbol{\alpha},\boldsymbol{\Psi}^{(m)})}{\partial\boldsymbol{\alpha}\partial\boldsymbol{\alpha}^T}$ and $\frac{\partial Q_1(\boldsymbol{\alpha},\boldsymbol{\Psi}^{(m)})}{\partial\boldsymbol{\alpha}}$ are respectively the Hessian matrix and the gradient vector of $Q_1(\boldsymbol{\alpha},\boldsymbol{\Psi}^{(m)})$. At each IRLS iteration the Hessian and the gradient are evaluated at $\boldsymbol{\alpha}=\boldsymbol{\alpha}^{(l)}$ and are computed similarly as in [12]. The parameter update $\boldsymbol{\alpha}^{(m+1)}$ is taken at convergence of the IRLS algorithm (25). Then, for $k=1\ldots,K$,

CM-Step 2. Calculate $(\boldsymbol{\beta}_k^{T(m+1)}, \sigma_k^{2(m+1)})^T$ by maximizing $Q_2(\boldsymbol{\theta}_k; \boldsymbol{\Psi}^{(m)})$ w.r.t $(\boldsymbol{\beta}_k^T, \sigma_k^2)^T$. For the skew t mixture of linear experts (STMoLE) case, where the expert means are linear regressors, that is, of the form (7), this maximization can be performed in a closed form and provides the following updates:

$$\boldsymbol{\beta}_k^{(m+1)} = \left[\sum_{i=1}^{n} \tau_{ik}^{(q)} w_{ik}^{(m)} \boldsymbol{x}_i \boldsymbol{x}_i^T\right]^{-1} \sum_{i=1}^{n} \tau_{ik}^{(q)} \left(w_{ik}^{(m)} y_i - \boldsymbol{e}_{1,ik}^{(m)} \delta_k^{(m+1)}\right) \boldsymbol{x}_i, \tag{26}$$

$$\sigma_k^{2(m+1)} = \frac{\sum_{i=1}^{n} \tau_{ik}^{(m)} \left[w_{ik}^{(m)} \left(\boldsymbol{y}_i - \boldsymbol{\beta}_k^{T(m+1)} \boldsymbol{x}_i\right)^2 - 2\delta_k^{(m+1)} \boldsymbol{e}_{1,ik}^{(m)} (y_i - \boldsymbol{\beta}_k^{T(m+1)} \boldsymbol{x}_i) + \boldsymbol{e}_{2,ik}^{(m)}\right]}{2\left(1 - \delta_k^{2(m)}\right) \sum_{i=1}^{n} \tau_{ik}^{(m)}}. \tag{27}$$

CM-Step 3. The skewness parameters λ_k are updated by maximizing $Q_2(\boldsymbol{\theta}_k; \boldsymbol{\Psi}^{(m)})$ w.r.t λ_k, with $\boldsymbol{\beta}_k$ and σ_k^2 fixed at the update $\boldsymbol{\beta}_k^{(m+1)}$ and $\sigma_k^{2(m+1)}$, respectively. It can be easily shown that the maximization to obtain $\delta_k^{(m+1)}$ $(k = 1, \ldots, K)$ consists in solving the following equation in δ_k:

$$\delta_k(1 - \delta_k^2) \sum_{i=1}^{n} \tau_{ik}^{(m)} + (1 + \delta_k^2) \sum_{i=1}^{n} \tau_{ik}^{(m)} \frac{d_{ik}^{(m+1)} e_{1,ik}^{(m)}}{\sigma_k^{(m+1)}} - \delta_k \sum_{i=1}^{n} \tau_{ik}^{(m)} \left[w_{ik}^{(m)} d_{ik}^{2\,(m+1)} + \frac{e_{2,ik}^{(m)}}{\sigma_k^{2(m+1)}}\right] = 0. \tag{28}$$

CM-Step 4. Similarly, the degrees of freedom ν_k are updated by maximizing $Q_3(\nu_k; \boldsymbol{\Psi}^{(m)})$ w.r.t ν_k with $\boldsymbol{\beta}_k$ and σ_k^2 fixed at $\boldsymbol{\beta}_k^{(m+1)}$ and $\sigma_k^{2(m+1)}$, respectively. An update $\nu_k^{(m+1)}$ is calculated as solution of the following equation in ν_k:

$$-\psi\left(\frac{\nu_k}{2}\right) + \log\left(\frac{\nu_k}{2}\right) + 1 + \frac{\sum_{i=1}^{n} \tau_{ik}^{(m)} \left(\boldsymbol{e}_{3,ik}^{(m)} - w_{ik}^{(m)}\right)}{\sum_{i=1}^{n} \tau_{ik}^{(m)}} = 0. \tag{29}$$

The two scalar non-linear Eqs. (28) and (29) can be solved similarly as in the TMoE model, that is with a root finding algorithm, such as Brent's method [8].

As mentioned before, one can see that, when the robustness parameter $\nu_k \to \infty$ for all the components, the parameter updates for the STMoE model correspond to those of the SNMoE model (see [9]). On the other hand, when the skewness parameters $\lambda_k = 0$, the STMoE parameter updates correspond to those of the TMoE model ([9]). Finally, when both the degrees of freedom $\nu_k \to \infty$ and the skewness $\lambda_k = 0$, we obtain the parameter updates of the standard NMoE model. The STMoE therefore provides a more general framework for inferring flexible MoE models and attempts to simultaneously accommodate data with asymmetric distribution heavy tails and outliers.

Here the ECM algorithm is used to infer the STMoE model parameters. We note that there is a good generalization of the EM algorithm, that is the Minorization-Maximization (MM) algorithm [28]. The MM algorihm, used in the MoE framewrok namely by [20,48] can also be a good alternative to the ECM algorithm used here. On the other hand, the ECM algorithm divides the space of model-parameters to perform sequentially the optimization in each sub-space. It may also be convenient to divide the space of the hidden variables and to alternate the optimisation, cyclically within each sub-space. This scheme is known as the Alternating ECM (AECM) algorithm [45].

5. Prediction using the STMoE

The goal in regression is to be able to make predictions for the response variable(s) given some new value of the predictor variable(s) on the basis of a model trained on a set of training data. In regression analysis using mixture of experts, the aim is therefore to predict the response y given new values of the predictors $(\boldsymbol{x}, \boldsymbol{r})$, on the basis of a MoE model characterized by a parameter vector $\hat{\boldsymbol{\Psi}}$ inferred from a set of training data, here, by maximum likelihood via EM. These predictions can be expressed in terms of the predictive distribution of y, which is obtained by substituting the maximum likelihood parameter $\hat{\boldsymbol{\Psi}}$ into (1) and (2) to give:

$$f(y|\boldsymbol{x}, \boldsymbol{r}; \hat{\boldsymbol{\Psi}}) = \sum_{k=1}^{K} \pi_k(\boldsymbol{r}; \hat{\boldsymbol{\alpha}}) f_k(y|\boldsymbol{x}; \hat{\boldsymbol{\Psi}}_k).$$

Using f, we might then predict y for a given set of $\boldsymbol{x}$'s and $\boldsymbol{r}$'s as the expected value under f, that is by calculating the prediction $\hat{y} = \mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|\boldsymbol{r}, \boldsymbol{x})$. We thus need to compute the expectation of the mixture of experts model. It is easy to show (see for example Section 1.2.4 in [17]) that the mean and the variance of a mixture of experts distribution of the form (5) are respectively given by

$$\mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|\boldsymbol{r}, \boldsymbol{x}) = \sum_{k=1}^{K} \pi_k(\boldsymbol{r}; \hat{\boldsymbol{\alpha}}_n) \mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x}), \tag{30}$$

$$\mathbb{V}_{\hat{\boldsymbol{\Psi}}}(Y|\boldsymbol{r}, \boldsymbol{x}) = \sum_{k=1}^{K} \pi_k(\boldsymbol{r}; \hat{\boldsymbol{\alpha}}_n) \left[\left(\mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x})\right)^2 + \mathbb{V}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x})\right] - \left[\mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|\boldsymbol{r}, \boldsymbol{x})\right]^2, \tag{31}$$

where $\mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x})$ and $\mathbb{V}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x})$ are respectively the component-specific (expert) means and variances. The mean and the variance for the two MoE models described here are given as follows.

NMoE. For the NMoE model, the normal expert means and variances are respectively given by $\mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x}) = \hat{\boldsymbol{\beta}}_k^T \boldsymbol{x}$ and $\mathbb{V}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x}) = \hat{\sigma}_k^2$. Then, from (30) it follows that the mean of the NMoE is given by

$$\mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|\boldsymbol{r}, \boldsymbol{x}) = \sum_{k=1}^{K} \pi_k(\boldsymbol{r}; \hat{\boldsymbol{\alpha}}_n) \hat{\boldsymbol{\beta}}_k^T \boldsymbol{x}. \tag{32}$$

STMoE. The mean and the variance for a skew t random variable, for this scalar case, can be easily computed as in Section 4.2 in [4] for a non-zero location parameter. Thus, for the STMoE model, the expert means for $\hat{\nu}_k > 1$, are given by

$$\mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x}) = \hat{\boldsymbol{\beta}}_k^T \boldsymbol{x} + \hat{\sigma}_k \; \hat{\delta}_k \; \xi(\hat{\nu}_k)$$

and the expert variances for $\hat{\nu}_k > 2$ are given by

$$\mathbb{V}_{\hat{\boldsymbol{\Psi}}}(Y|Z = k, \boldsymbol{x}) = \left(\frac{\hat{\nu}_k}{\hat{\nu}_k - 2} - \hat{\delta}_k^2 \; \xi^2(\hat{\nu}_k)\right) \hat{\sigma}_k^2,$$

where $\xi(\hat{\nu}_k) = \sqrt{\frac{\hat{\nu}_k}{\pi}} \frac{\Gamma\left(\frac{\hat{\nu}_k}{2} - \frac{1}{2}\right)}{\Gamma\left(\frac{\hat{\nu}_k}{2}\right)}$. Then, following (30), the mean of the proposed STMoE is thus given by:

$$\mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y|\boldsymbol{r}, \boldsymbol{x}) = \sum_{k=1}^{K} \pi_k(\boldsymbol{r}; \hat{\boldsymbol{\alpha}}) \left(\hat{\boldsymbol{\beta}}_k^T \boldsymbol{x} + \hat{\sigma}_k \, \hat{\delta}_k \, \xi(\hat{\nu}_k) \right). \tag{33}$$

Finally, the variance for each MoE model is obtained by using (31) with the specified corresponding means and variances calculated in the above.

6. Model-based clustering using the STMoE

The MoE models can also be used for a model-based clustering perspective to provide a partition of the regression data into K clusters. Model-based clustering using the proposed STMoE consists in assuming that the observed data $\{\boldsymbol{x}_i, \boldsymbol{r}_i, y_i\}_{i=1}^n$ are generated from a K component mixture of skew t experts, with parameter vector $\boldsymbol{\Psi}$ where the STMoE components are interpreted as clusters and hence associated to clusters. The problem of clustering therefore becomes the one of estimating the MoE parameters $\boldsymbol{\Psi}$, which is performed here by using the dedicated ECM algorithm presented in Section 4.1. Once the parameters are estimated, the provided posterior component memberships τ_{ik} given by (19) represent a fuzzy partition of the data. A hard partition of the data can then be obtained from the posterior memberships by applying the MAP rule, that is:

$$\hat{z}_i = \arg\max_{k=1}^{K} \hat{\tau}_{ik} \tag{34}$$

where $\hat{z}_i$ represents the estimated cluster label for the ith individual.

7. Model selection for the STMoE

One of the issues in mixture model-based clustering is model selection. The problem of model selection for the STMoE models presented here in its general form is equivalent to the one of choosing the optimal number of experts K, the degree p of the polynomial regression and the degree q for the logistic regression. The optimal value of the triplet (K, p, q) can be computed by using some model selection criteria such as the Akaike Information Criterion (AIC) [1], the Bayesian Information Criterion (BIC) [56] or the Integrated Classification Likelihood criterion (ICL) [6], etc. The AIC and BIC are are penalized observed log-likelihood criteria which can be defined as functions to be maximized and are respectively given by:

$$\text{AIC}(K, p, q) = \log L(\hat{\boldsymbol{\Psi}}) - \eta_{\boldsymbol{\Psi}},$$

$$\text{BIC}(K, p, q) = \log L(\hat{\boldsymbol{\Psi}}) - \frac{\eta_{\boldsymbol{\Psi}} \log(n)}{2}.$$

The ICL criterion consists in a penalized complete-data log-likelihood and can be expressed as follows:

$$\text{ICL}(K, p, q) = \log L_c(\hat{\boldsymbol{\Psi}}) - \frac{\eta_{\boldsymbol{\Psi}} \log(n)}{2}.$$

In the above, $\log L(\hat{\boldsymbol{\Psi}})$ and $\log L_c(\hat{\boldsymbol{\Psi}})$ are respectively the incomplete (observed) data log-likelihood and the complete data log-likelihood, obtained at convergence of the ECM algorithm for the corresponding MoE model and $\eta_{\boldsymbol{\Psi}}$ is the number of free model parameters. The number of free parameters $\eta_{\boldsymbol{\Psi}}$ is given by $\eta_{\boldsymbol{\Psi}} = K(p+q+3) - q - 1$ for the NMoE model and $\eta_{\boldsymbol{\Psi}} = K(p+q+5) - q - 1$ for the proposed STMoE model. Indeed, for each component, the STMoE have two additional parameters to be estimated, which are the robustness and the skewness parameters.

However, note that in MoE it is common to use mixing proportions modeled as logistic transformation of linear functions of the covariates, that is the covariate vector in (2) is given by $\boldsymbol{r}_i = (1, r_i)^T$ (corresponding to $q = 2$), r_i being an univariate covariate variable. This is also adopted in this work. Moreover, for the case of linear experts, that is when the experts are linear regressors with parameter vector $\boldsymbol{\beta}_k$ for which the corresponding covariate vector $\boldsymbol{x}_i$ in (7) is given by $\boldsymbol{x}_i = (1, x_i)^T$ (corresponding to $p = 2$), r_i being an univariate covariate variable, the model selection reduces to choosing the number of experts K. Here we mainly consider this linear case. However, for a general use of the proposed STMoE model, even though the model selection criteria such as AIC, BIC, ICL can be easily computed, the direct model selection is difficult due to the large model space dimension $\nu_{\boldsymbol{\Psi}}$. A searching strategy is then required to optimise the way of exploring the model space.

8. Experimental study

This section is dedicated to the evaluation of the proposed approach on simulated data and real-world data . We evaluated the performance of proposed ECM algorithm[1] for the STMoE model in terms of modeling, robustness to outliers and clustering.

8.1. Initialization and stopping rules

The parameters $\boldsymbol{\alpha}_k$ $(k = 1, \ldots, K-1)$ of the mixing proportions are initialized randomly, including an initialization at the null vector for one run (corresponding to equal mixing proportions). Then, the common parameters $(\boldsymbol{\beta}_k, \sigma_k^2)$ $(k = 1, \ldots, K)$ are initialized from a random partition of the data into K clusters. This corresponds to fitting a normal mixture of experts where the initial values of the parameters are respectively given by (8) and (9) with the posterior memberships τ_{ik} replaced by the hard assignments Z_{ik} issued from the random partition. For the STMoE model, the robustness parameters ν_k $(k = 1, \ldots, K)$ is initialized randomly in the range [1, 200] and the skewness parameters λ_k $(k = 1, \ldots, K)$ is initialized by randomly initializing the parameter δ_k in $(-1, 1)$ from the relation $\lambda_k = \frac{\delta_k}{\sqrt{1-\delta_k^2}}$. Then, the proposed ECM algorithm for each model is stopped when the relative variation of the observed-data log-likelihood $\frac{\log L(\boldsymbol{\Psi}^{(m+1)}) - \log L(\boldsymbol{\Psi}^{(m)})}{|\log L(\boldsymbol{\Psi}^{(m)})|}$ reaches a prefixed threshold (for example $\epsilon = 10^{-6}$). For each model, this process is repeated 10 times and the solution corresponding the highest log-likelihood is finally selected.

8.2. Experiments on simulation data sets

In this section we perform an experimental study on simulated data sets to apply and assess the proposed model. Two sets of experiments have been performed. The first experiment aims at observing the effect of the sample size on the estimation quality and the second one aims at observing the impact of the presence of outliers in the data on the estimation quality, that is the robustness of the models.

8.2.1. Experiment 1

For this first experiment on simulated data, each simulated sample consisted of n observations with increasing values of the sample size n: 50, 100, 200, 500, 1000. The simulated data are generated from a two component mixture of linear experts, that is $K = 2$, $p = q = 1$. The covariate variables $(\boldsymbol{x}_i, \boldsymbol{r}_i)$ are simulated such

[1] The codes have been implemented in Matlab and are available upon request from the author.

Table 1
Parameter values used in simulation.

Parameters					
Component 1	$\boldsymbol{\alpha}_1 = (0, 10)^T$	$\boldsymbol{\beta}_1 = (0, 1)^T$	$\sigma_1 = 0.1$	$\lambda_1 = 3$	$\nu_1 = 5$
Component 2	$\boldsymbol{\alpha}_2 = (0, 0)^T$	$\boldsymbol{\beta}_2 = (0, -1)^T$	$\sigma_2 = 0.1$	$\lambda_2 = -10$	$\nu_2 = 7$

Table 2
MSE $\times 10^3$ between each component of the estimated parameter vector of the STMoE model and the actual one for a varying sample size n.

Param. n	α_{10}	α_{11}	β_{10}	β_{11}	β_{20}	β_{21}	σ_1	σ_2	λ_1	λ_2	ν_1	ν_2
50	525	5737	0.965	2.440	4.388	0.667	0.954	0.608	3115	16095	15096	4643
100	457	1815	0.847	0.852	0.742	0.660	0.844	0.303	2013	7844	5360	263
200	247	785	0.816	0.348	0.473	0.556	0.362	0.297	700	3847	3135	167
500	31	565	0.363	0.091	0.314	0.398	0.091	0.061	7.8	1078	223	8.6
1000	8.5	68	0.261	0.076	0.233	0.116	0.026	0.002	2.8	554	49.4	0.79

that $\boldsymbol{x}_i = \boldsymbol{r}_i = (1, x_i)^T$ where x_i is simulated uniformly over the interval $(-1, 1)$. We consider each of the two models for data generation (NMoE and STMoE), that is, given the covariates, the response $y_i|\{\boldsymbol{x}_i, \boldsymbol{r}_i; \boldsymbol{\Psi}\}$ is simulated according to the generative process of the models (3) and (15). For each generated sample, we fit each of the two models. Thus, the results are reported for the two models with data generated from each of them. We consider the mean square error (MSE) between each component of the true parameter vector and the estimated one, which is given by $\|\boldsymbol{\Psi}_j - \hat{\boldsymbol{\Psi}}_j\|^2$. The squared errors are averaged on 100 trials. The used simulation parameters $\boldsymbol{\Psi}$ for each model are given in Table 1.

8.2.2. Obtained results

Table 2 shows the obtained results in terms of the MSE for the STMoE. One can observe that, for the proposed model, the parameter estimation error is decreasing as n increases, which is related the convergence property of the maximum likelihood estimator. One can also observe that the error decreases significantly for $n \geq 500$, especially for the regression coefficients and the scale parameters.

In addition to the previously showed results, we plotted in Figs. 1 and 2 the estimated quantities provided by applying respectively the NMoE model and the proposed STMoE model, and their true counterparts for $n = 500$ for the same the data set which was generated according the NMoE model. The upper-left plots show the estimated mean function, the estimated expert component mean functions, and the corresponding true ones. The upper-right plots show the estimated mean function with the estimated confidence region computed as plus and minus twice the estimated (pointwise) standard deviation of the model as presented in Section 5, and their true counterparts. The bottom-left plots show the true expert component mean functions and the true partition, and the bottom-right plots show their estimated counterparts.

One can clearly see that the estimations provided by the proposed model are very close to the true ones which correspond to those of the NMoE model in this case. This shows that the proposed algorithm performs well and provides an additional support to the fact that the corresponding proposed STMoE model is good generalization of the normal mixture of experts (NMoE), as it clearly approaches the NMoE as shown in this simulated examples. Fig. 3 shows the true and estimated MoE mean functions and component mean functions by fitting the proposed STMoE model to a simulated data set of $n = 500$ observations. Each model was considered for data generation. The upper plot corresponds to the NMoE model and the bottom plot corresponds to the STMoE model. Finally, Fig. 4 shows the corresponding true and estimated partitions. Again, one can clearly see that both the estimated models are precise. The fitted functions are close to the true ones. In addition, one can also see that the partitions estimated by the STMoE model are close the actual partitions. The proposed STMoE model can therefore be used as alternative to the NMoE model for both regression and model-based clustering.

Table 3
MSE $\times 10^3$ between the estimated mean function and the true one for each of the two models for a varying probability c of outliers for each simulation. The first column indicates the model used for generating the data and the second one indicates the model used for inference.

Model	c	0%	1%	2%	3%	4%	5%
NMoE	NMoE	**0.178**	1.057	1.241	3.631	13.25	28.96
	STMoE	0.258	**0.741**	**0.794**	**0.696**	**0.697**	**0.626**
STMoE	NMoE	0.710	0.7238	1.048	6.066	12.45	31.64
	STMoE	**0.280**	**0.186**	**0.447**	**0.600**	**0.509**	**0.602**

8.2.3. Experiment 2

In this experiment we examine the robustness of the proposed STMoE model to outliers versus the standard NMoE one. For that, we considered each of the two models (NMoE and STMoE) for data generation. For each generated sample, each of the two models in considered for the inference. The data were generated exactly in the same way as in Experiment 1, except for some observations which were generated with a probability c from a class of outliers. We considered the same class of outliers as in [48], that is the predictor x is generated uniformly over the interval $(-1, 1)$ and the response y is set the value -2. We apply the MoE models by setting the covariate vectors as before, that is, $\boldsymbol{x} = \boldsymbol{r} = (1, x)^T$. We considered varying probability of outliers $c = 0\%, 1\%, 2\%, 3\%, 4\%, 5\%$ and the sample size of the generated data is $n = 500$. An example of simulated sample containing 5% outliers is shown in Fig. 5. As a criterion of evaluation of the impact of the outliers on the quality of the results, we considered the MSE between the true regression mean function and the estimated one. This MSE is calculated as $\frac{1}{n}\sum_{i=1}^{n}\|\mathbb{E}_{\boldsymbol{\Psi}}(Y_i|\boldsymbol{r}_i, \boldsymbol{x}_i) - \mathbb{E}_{\hat{\boldsymbol{\Psi}}}(Y_i|\boldsymbol{r}_i, \boldsymbol{x}_i)\|^2$ where the expectations are computed as in Section 5.

8.2.4. Obtained results

Table 3 shows, for each of the two models, the results in terms of mean squared error (MSE) between the true mean function and the estimated one, for an increasing number of outliers in the data. First, one can see that, when there is no outliers ($c = 0\%$) and when the data follow a NMoE distribution, the error of fitting a NMoE is very slightly less than the one of fitting the proposed STMoE model. The STMoE then is still competitive. However, when the data do not contain outliers and follow a STMoE distribution, fitting a NMoE is restrictive since the error of this one is high compared to the one obtained by fitting a STMoE model. More importantly, it can be seen that, as expected, when there is outliers

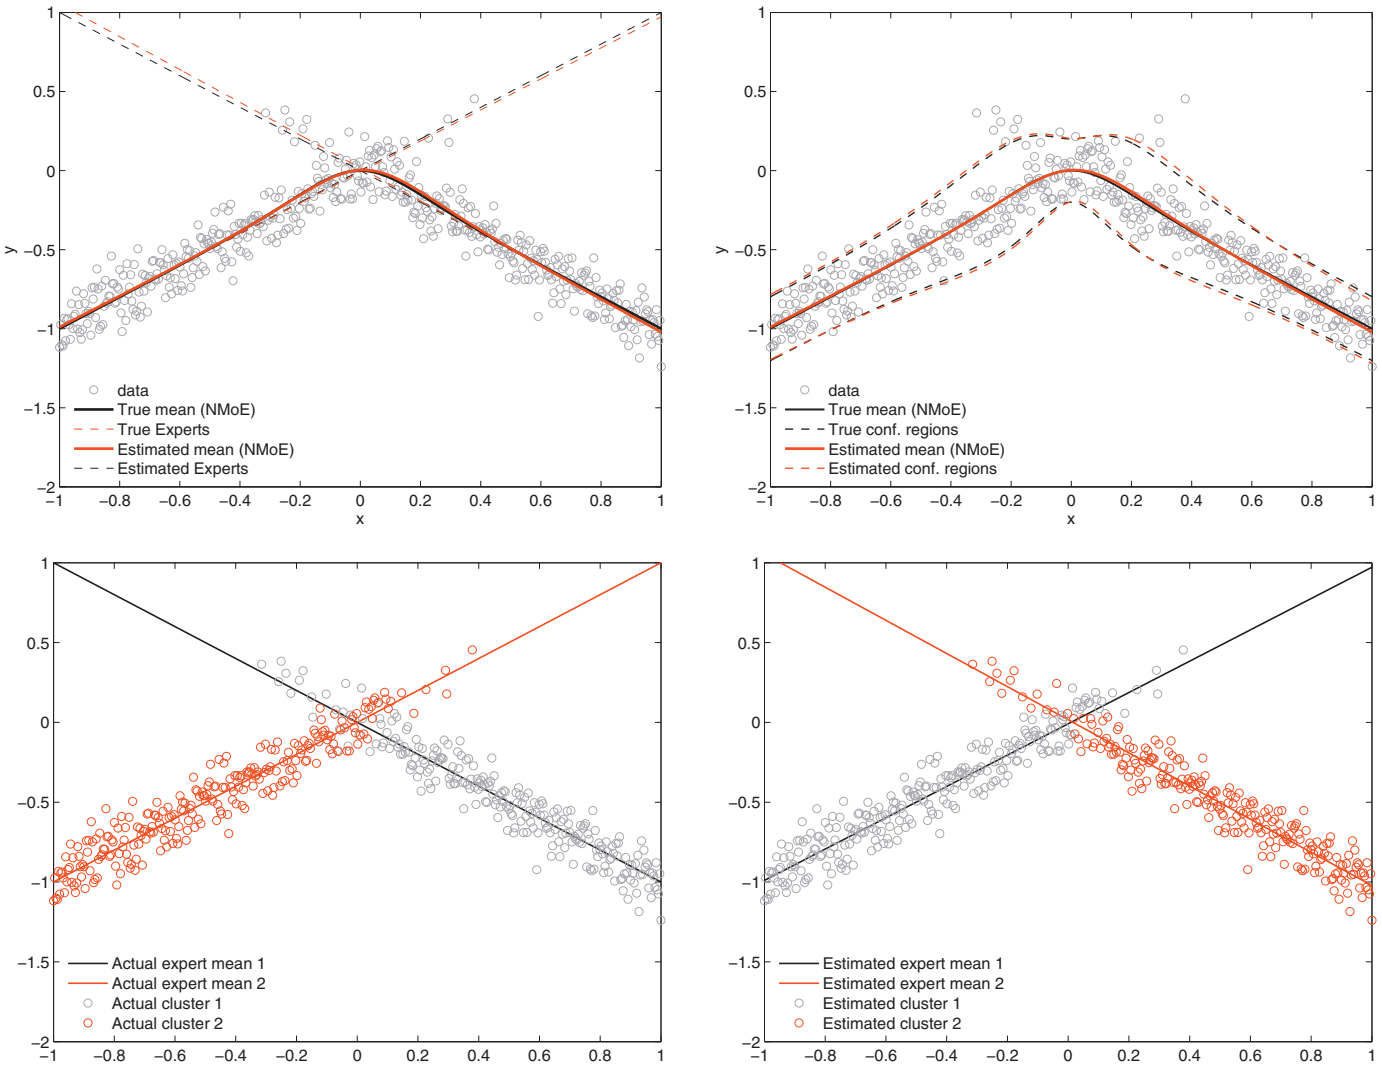

Fig. 1. Fitted NMoE model to a data set generated according to the NMoE model.

in the data, including the situations with only few atypical data points (1% and 2 %), the NMoE does not provide an adapted fit and is clearly outperformed by the proposed STMoE model. This includes the two situations, that is, including when the data are not generated according to the STMoE model. The errors of the NMoE model are high compared to those of the STMoE. This confirms that te STMoE is much more robust to outliers compared to the normal mixture of experts because the expert components in the STMoE model follow the robust skew t distribution. The NMoE is sensitive to outliers. On the other hand it can be seen that, when the number of outliers is increasing, increase in the error of the NMoE is more pronounced compared to the one of STMoE model. The error for the STMoE may indeed slightly increase, remain stable or even decrease in some situations. This provides an additional support to the expected robustness of the STMoE compared to the NMoE.

Then, in order to highlight the robustness to noise of the TMoE and STMoE models, in addition to the previously shown numerical results, Figs 5 and 6 show an example of results obtained on the same data set by, respectively, the NMoE and the STMoE. The data are generated by the NMoE model and contain $c = 5\%$ of outliers.

In this example, we clearly see that the NMoE model is severely affected by the outliers. It provides a rough fit especially for the second component whose estimation is corresponds to a rough approximation due to the atypical data. However, one can see that the STMoE model clearly provides a precise fit; the estimated mean function and expert components are very close to the true ones. The STMoE is robust to outliers, in terms of estimating the true model as well as in terms of estimating the true partition of the data (as shown in the middle plots of the data). Notice that for the STMoE, the confidence regions are not shown because for this situation the estimated degrees of freedom are less than 2 (1.6097 and 1.5311) for the STMoE); Hence, the variance for this model in that case is not defined (see Section 5). The STMoE model provides indeed components with small degrees of freedom corresponding to highly heavy tails, which allow to handle outliers in this noisy case. While the variance is not estimable here, resampling techniques can be used to evaluate it, such as the techniques of [49] for producing standard errors and confidence intervals for mixture parameters in model-based clustering.

8.3. Application to two real-world data sets

In this section, we consider an application to two real-world data sets: the tone perception data set and the temperature anomalies data set shown in Fig. 7.

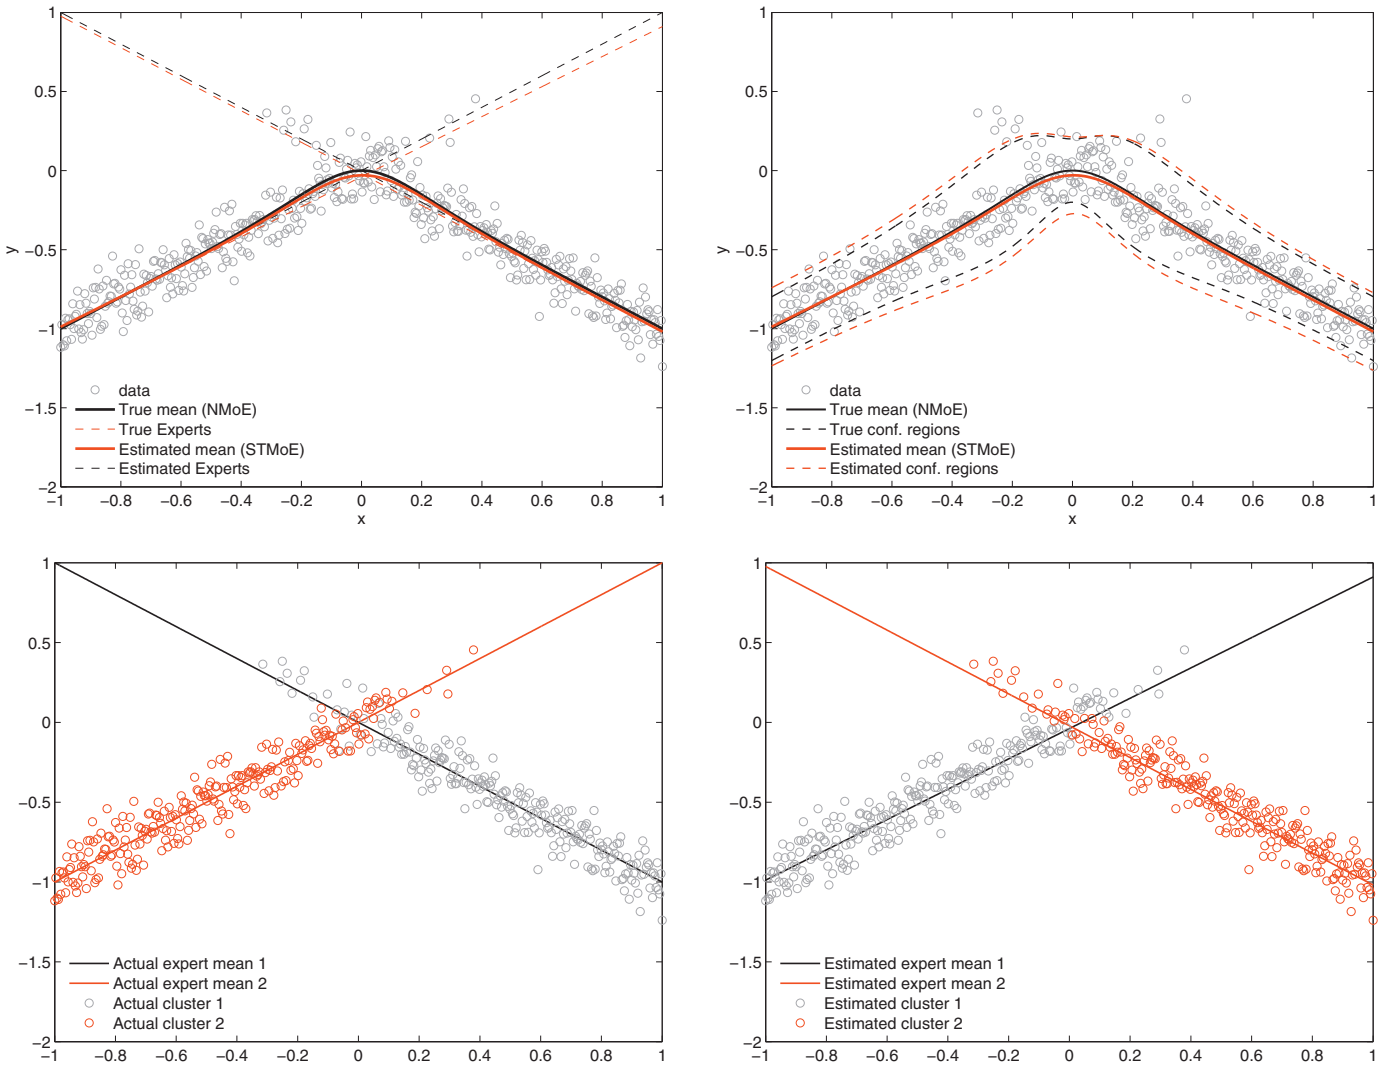

Fig. 2. Fitted STMoE model to a data set generated according to the NMoE model.

8.3.1. Tone perception data set

The first analyzed data set is the real tone perception data set[2] which goes back to [14]. It was recently studied by [5,29] and [58] by using robust regression mixture models based on, respectively, the t distribution and the Laplace distribution. In the tone perception experiment, a pure fundamental tone was played to a trained musician. Electronically generated overtones were added, determined by a stretching ratio ("stretch ratio" $= 2$) which corresponds to the harmonic pattern usually heard in traditional definite pitched instruments. The musician was asked to tune an adjustable tone to the octave above the fundamental tone and a "tuned" measurement gives the ratio of the adjusted tone to the fundamental. The obtained data consists of $n = 150$ pairs of "tuned" variables, considered here as predictors (x), and their corresponding "strech ratio" variables considered as responses (y). To apply the proposed MoE models, we set the response $y_i (i = 1, \ldots, 150)$ as the "strech ratio" variables and the covariates $\boldsymbol{x}_i = \boldsymbol{r}_i = (1, x_i)^T$ where x_i is the "tuned" variable of the ith observation. We also follow the study in [5] and [58] by using two mixture components. Model selection results are given later in Table 5.

Fig. 8 shows the scatter plots of the tone perception data and the linear expert components of the fitted NMoE model and the proposed STMoE model. One can observe that we obtain a good fit with the two models. The NMoE fit differs very slightly from the one of the STMoE. The two regression lines may correspond to correct tuning and tuning to the first overtone, respectively, as analyzed in [5]. Fig. 9 shows the log-likelihood profiles for each of the two models. It can namely be seen that training the skew t mixture of experts may take more iterations than the normal MoE model. The STMoE has indeed more parameters to estimate (additional skewness and robustness parameters). However, in terms of computing time, the algorithm is fast and converges in only few seconds (around 10 seconds for this example) on a personal laptop with 2,9 GHz processor and 8 GB memory. The values of estimated parameters for the tone perception data set are given in Table 4. One can see that the regression coefficients are very similar for the two models. One can also see that the STMoE model retrieves a skewed component and with high degrees of freedom compared to the other component. This one may be seen as approaching the one of a skew-normal MoE model, while the second one in approaching a t distribution, that is the one of a t-MoE model.

We also performed a model selection procedure on this data set to choose the best number of MoE components for a number of components between 1 and 5. We used BIC, AIC, and ICL. Table 5 gives the obtained values of the model selection criteria. One can see that for the NMoE model, the three criteria overestimate the number of components, but for both BIC and ICL, the

[2] Source: http://artax.karlin.mff.cuni.cz/r-help/library/fpc/html/tonedata.html.

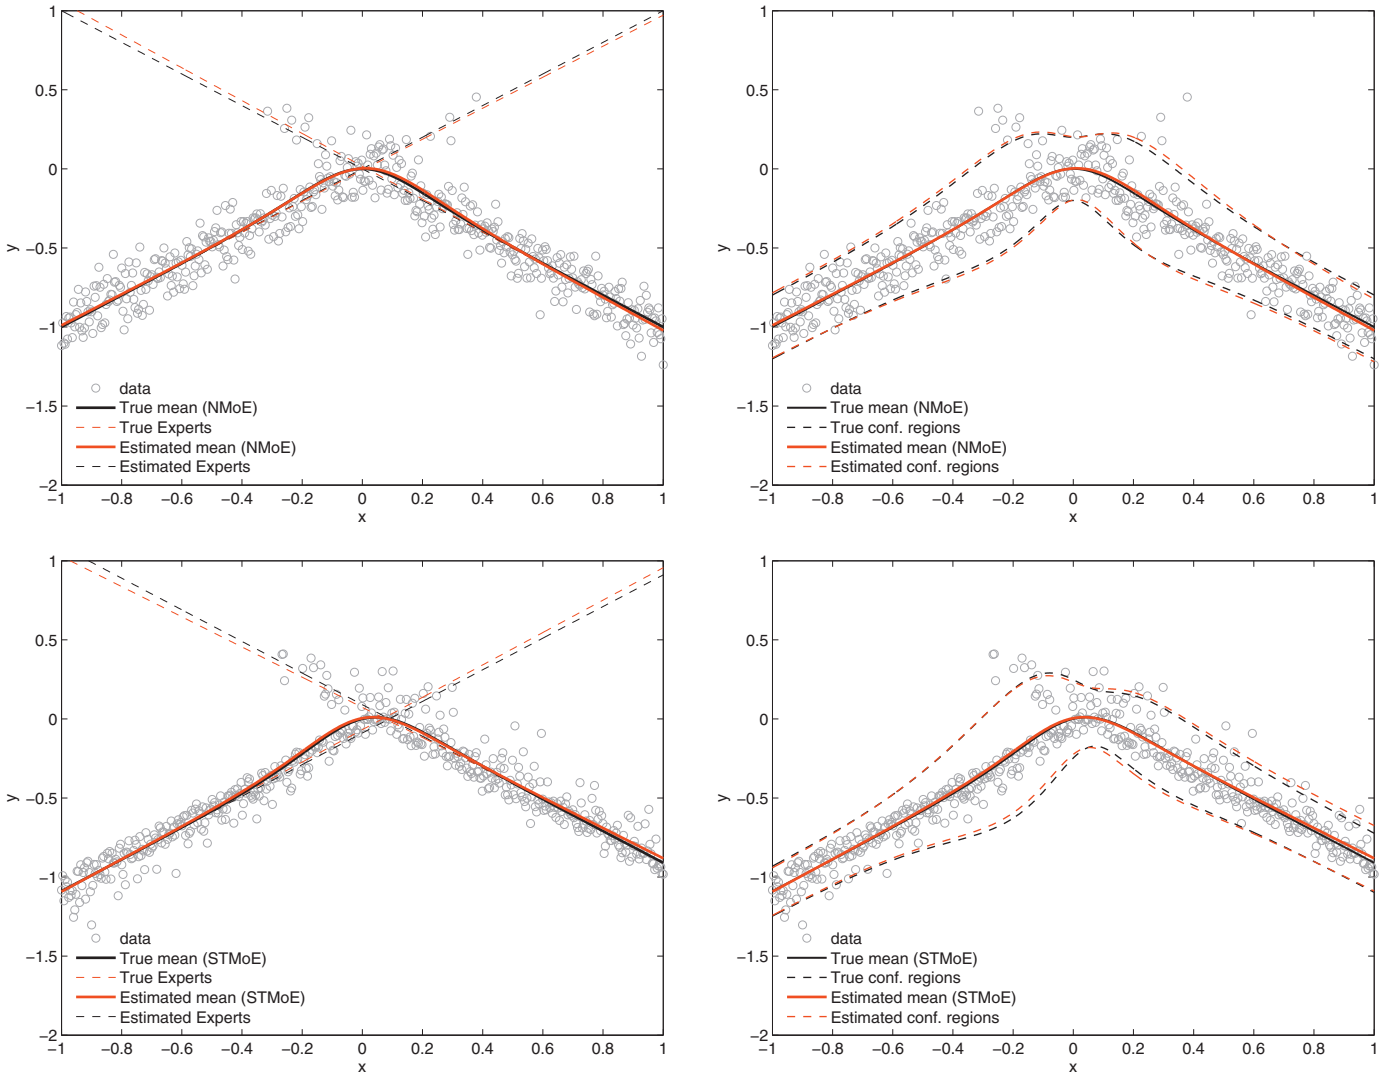

Fig. 3. The true and estimated mean function and expert mean functions by fitting the standard NMoE model (up) and the proposed STMoE model (bottom) to a simulated data set of $n = 500$ observations generating according to the corresponding model.

Table 4
Values of the estimated MoE parameters for the original Tone perception data set.

Param. Model	α_{10}	α_{11}	β_{10}	β_{11}	β_{20}	β_{21}	σ_1	σ_2	λ_1	λ_2	ν_1	ν_2
NMoE	−2.690	0.796	−0.029	0.995	1.913	0.043	0.137	0.047	–	–	–	–
STMoE	−3.044	0.824	−0.058	0.944	1.944	0.032	0.200	0.032	93.386	−0.011	19.070	1.461

Table 5
Choosing the number of expert components K for the original tone perception data by using the information criteria BIC, AIC, and ICL. Bold value indicates the highest value for each criterion.

	NMoE			STMoE		
K	BIC	AIC	ICL	BIC	AIC	ICL
1	1.866	6.382	1.866	69.532	77.059	69.532
2	122.805	134.847	107.384	**92.435**	110.499	**82.455**
3	118.193	137.763	76.524	77.9753	106.576	52.564
4	121.703	148.798	94.460	77.7092	116.847	56.365
5	**141.696**	**176.318**	**123.655**	79.043	**128.719**	67.748

solution with two components is also likely and is the most competitive to the selected one with 5 components. In deed, it can be seen that, if the number mixture of components is fixed at 4 rather than 5, both BIC and ICL would select the right number of components in that case. AIC performs poorly for the two models and overestimates the number of components. On the other hand, for the proposed STMoE model, both BIC and ICL retrieve the correct number of components. Then, one can conclude that the BIC and the ICL are the criteria that one would suggest for the analysis of this data with the proposed model.

Now we examine the sensitivity of the MoE models to outliers based on this real data set. For this, we adopt the same scenario used in [5] and [58] (the last and more difficult scenario) by adding 10 identical pairs (0, 4) to the original data set as outliers in the y-direction, considered as high leverage outliers. We apply the MoE models in the same way as before.

The upper plots in Fig. 10 clearly show that the normal mixture of experts fit is sensitive to outliers. However, note that for this situation, compared to the normal regression mixture result in [5], and the Laplace regression mixture and the t regression mixture results in [58], the fitted NMoE model is affected less severely by the outliers. This may be attributed to the fact that the mixing

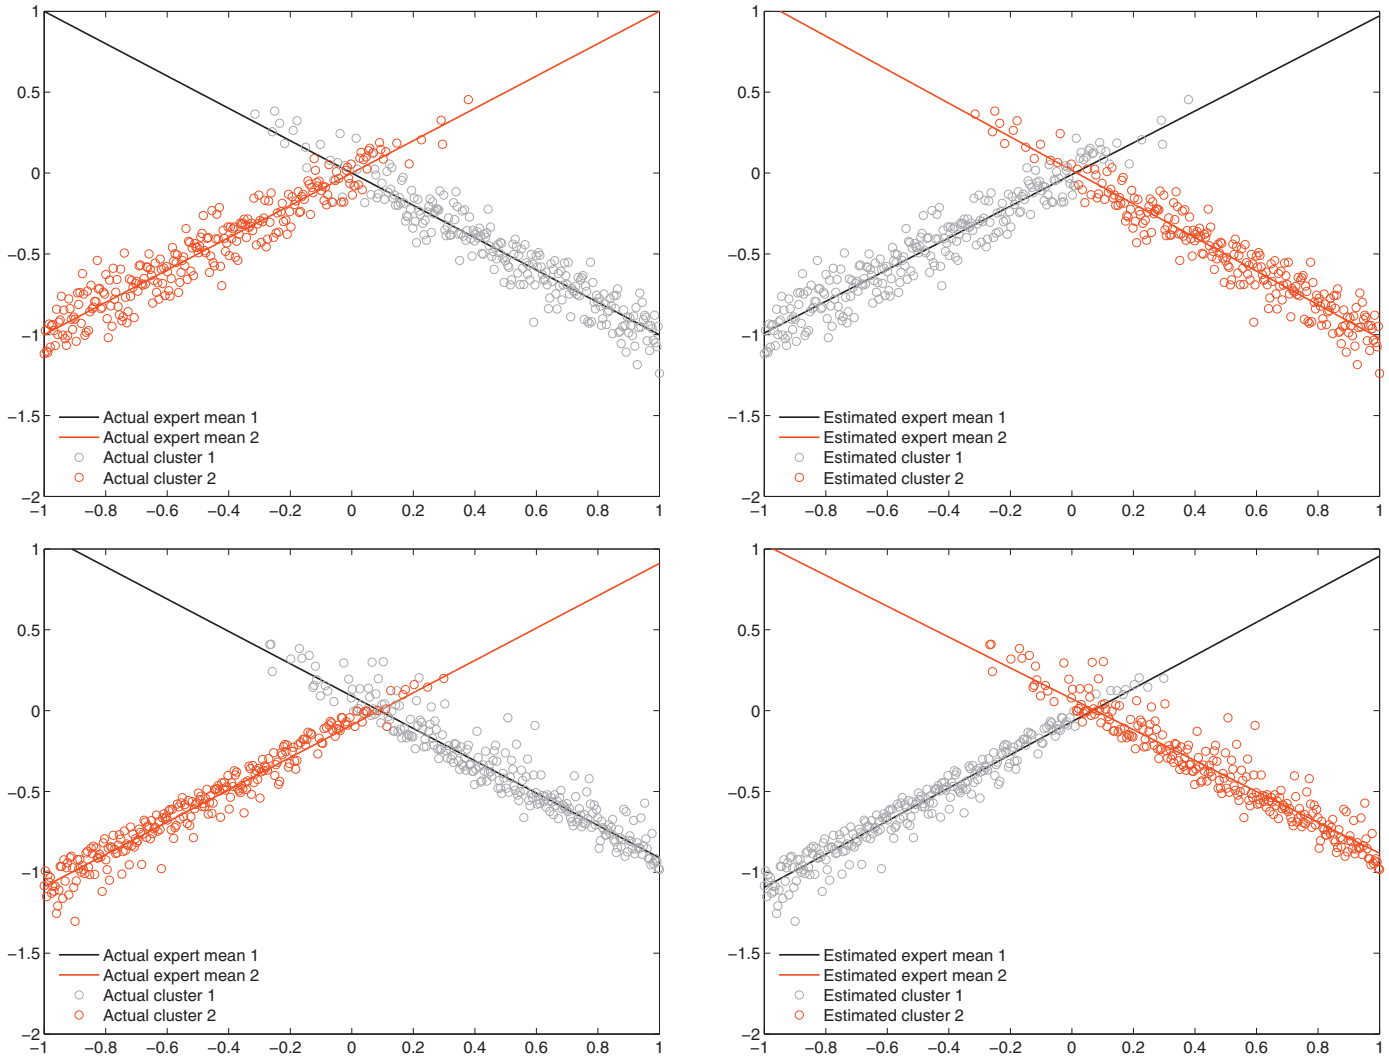

Fig. 4. The true and estimated partitions by fitting the standard NMoE model (up) and the proposed STMoE model (bottom) to the simulated data sets shown in Fig. 3.

Table 6
Values of the estimated MoE parameters for the tone perception data set with added outliers.

Param. Model	α_{10}	α_{11}	β_{10}	β_{11}	β_{20}	β_{21}	σ_1	σ_2	λ_1	λ_2	ν_1	ν_2
NMoE	0.811	0.150	3.117	−0.285	1.907	0.046	0.700	0.050	–	–	–	–
STMoE	−3.004	0.732	−0.246	1.016	1.808	0.060	0.212	0.088	156.240	1.757	81.355	1.630

proportions here are depending on the predictors, which is not the case in these regression mixture models, namely those of [5] and [58]. One can also see that, even the regression mean functions are affected severely by the outliers, the provided partitions are still reasonable and similar to those provided in the previous non-noisy case. Then, the plots in Fig. 10 also clearly show that the STMoE model provides a precise robust fit. For the STMoE, even if the fit differs very slightly compared to the case with outliers, the obtained fits for both situations (with and without outliers) are very reasonable. Moreover, we notice that, as showed in [58], for this situation with outliers, the t mixture of regressions fails; The fit is affected severely by the outliers. However, for the proposed STMoE model, the ten high leverage outliers have no significant impact on the fitted experts. This is because here the mixing proportions depend on the inputs, which is not the case for the regression mixture model described in [58]. Fig. 11 shows the log-likelihood profiles for each of the two models, which show a similar behavior than the one in the case without outliers. The values of estimated MoE parameters in this case with outliers are given in Table 6. One can see that the SNMoE model parameters are identical to those of the NMoE, with a skewness close to zero. The regression coefficients for the second expert component are very similar for the two models. For the STMoE model, it retrieves a skewed normal component while the second component is approaching a t distribution with a small degrees of freedom.

8.3.2. Temperature anomalies data set

In this experiment, we examine another real-world data set related to climate change analysis. The NASA GISS Surface Temperature (GISTEMP) analysis provides a measure of the changing global surface temperature with monthly resolution for the period since 1880, when a reasonably global distribution of meteorological stations was established. The GISS analysis is updated monthly, however the data presented here[3] are updated annually as issued from

[3] Source: [55], http://cdiac.ornl.gov/ftp/trends/temp/hansen/gl_land.txt.

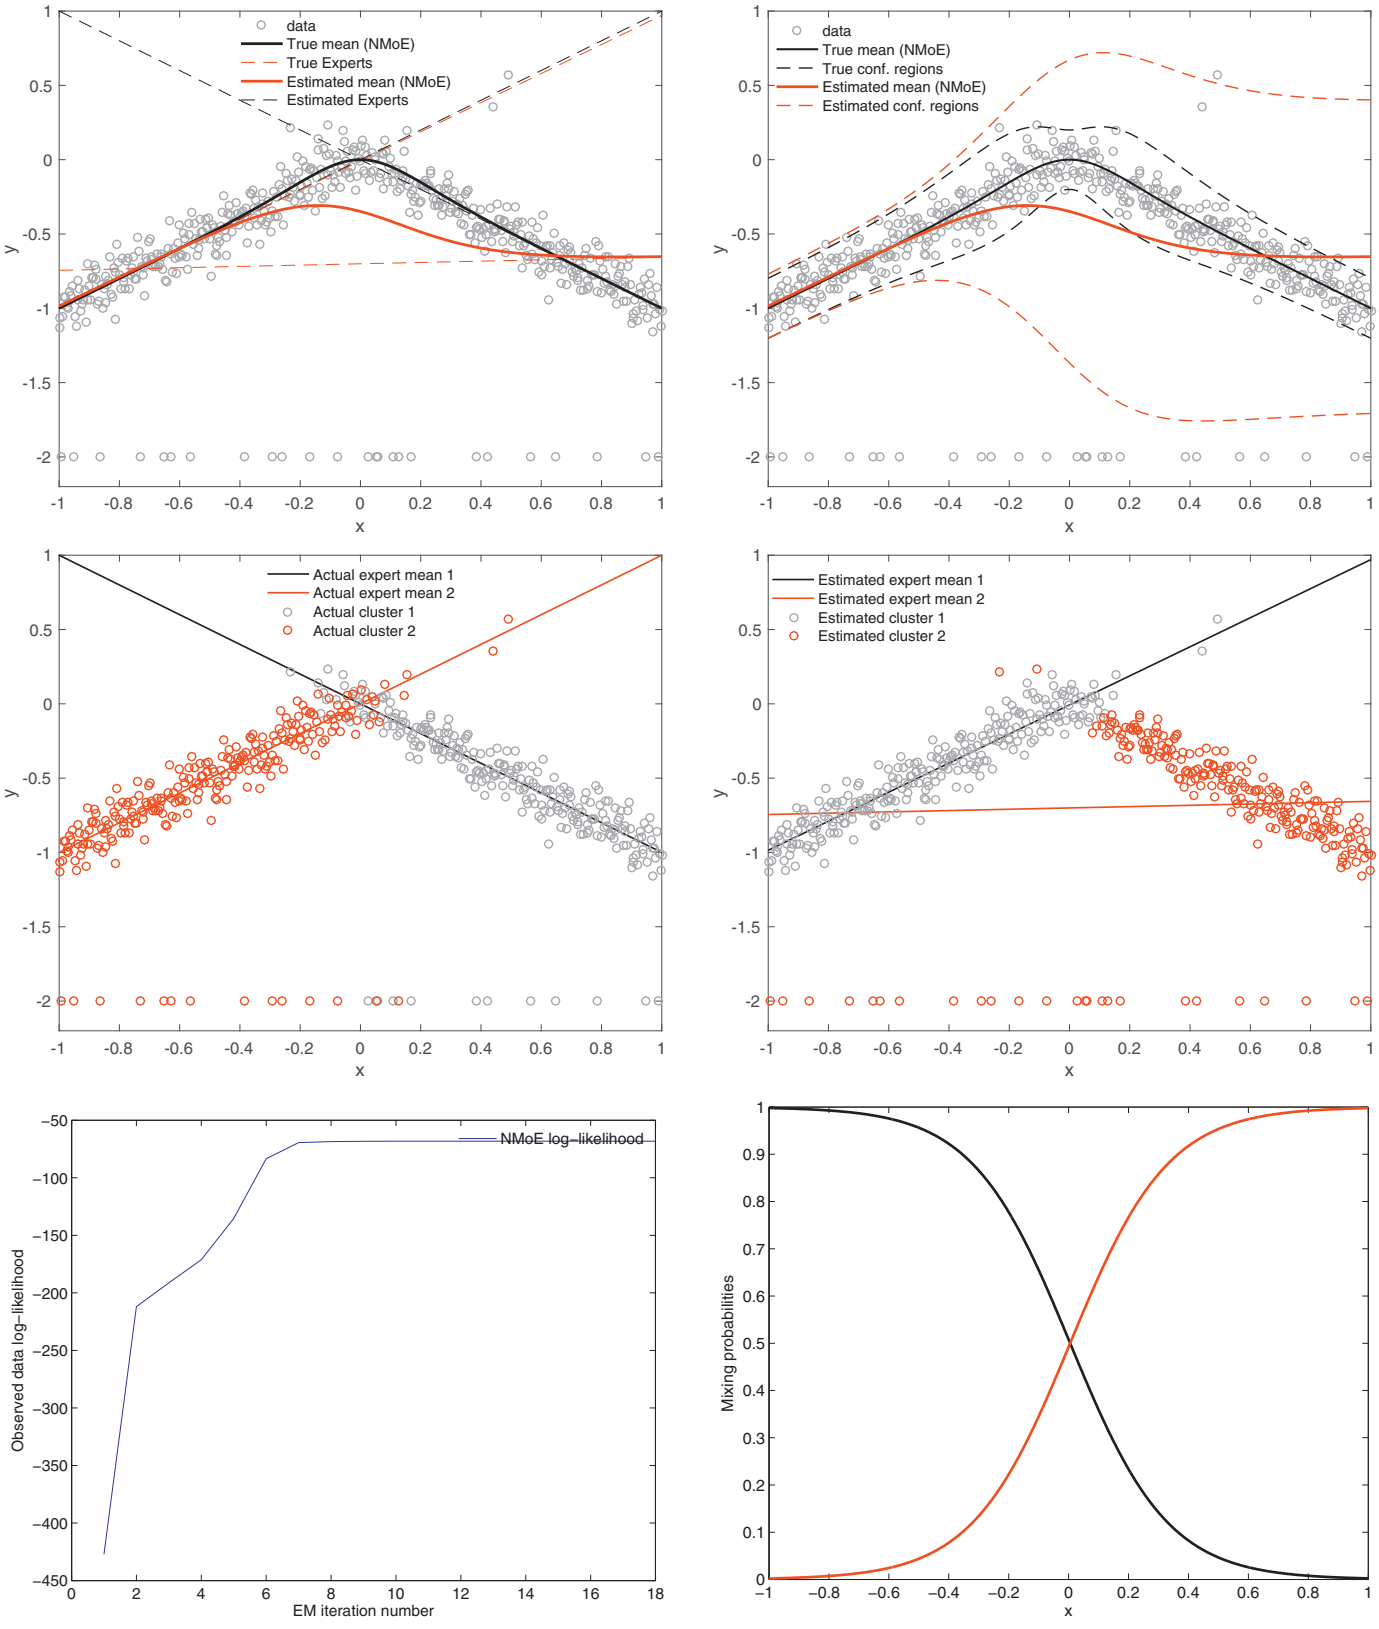

Fig. 5. Fitted NMoE model to a data set of $n = 500$ observations generated according to the NMoE model and including 5% of outliers.

the Carbon Dioxide Information Analysis Center (CDIAC), which has served as the primary climate-change data and information analysis center of the U.S. Department of Energy since 1982. The data consist of $n = 135$ yearly measurements of the global annual temperature anomalies (in degrees C) computed using data from land meteorological stations for the period of 1882 – 2012. These data have been analyzed earlier by [24,25] and recently by [48] by using the Laplace mixture of linear experts (LMoLE).

To apply the proposed non-normal mixture of expert model, we consider mixtures of two experts as in [48]. This number of com-

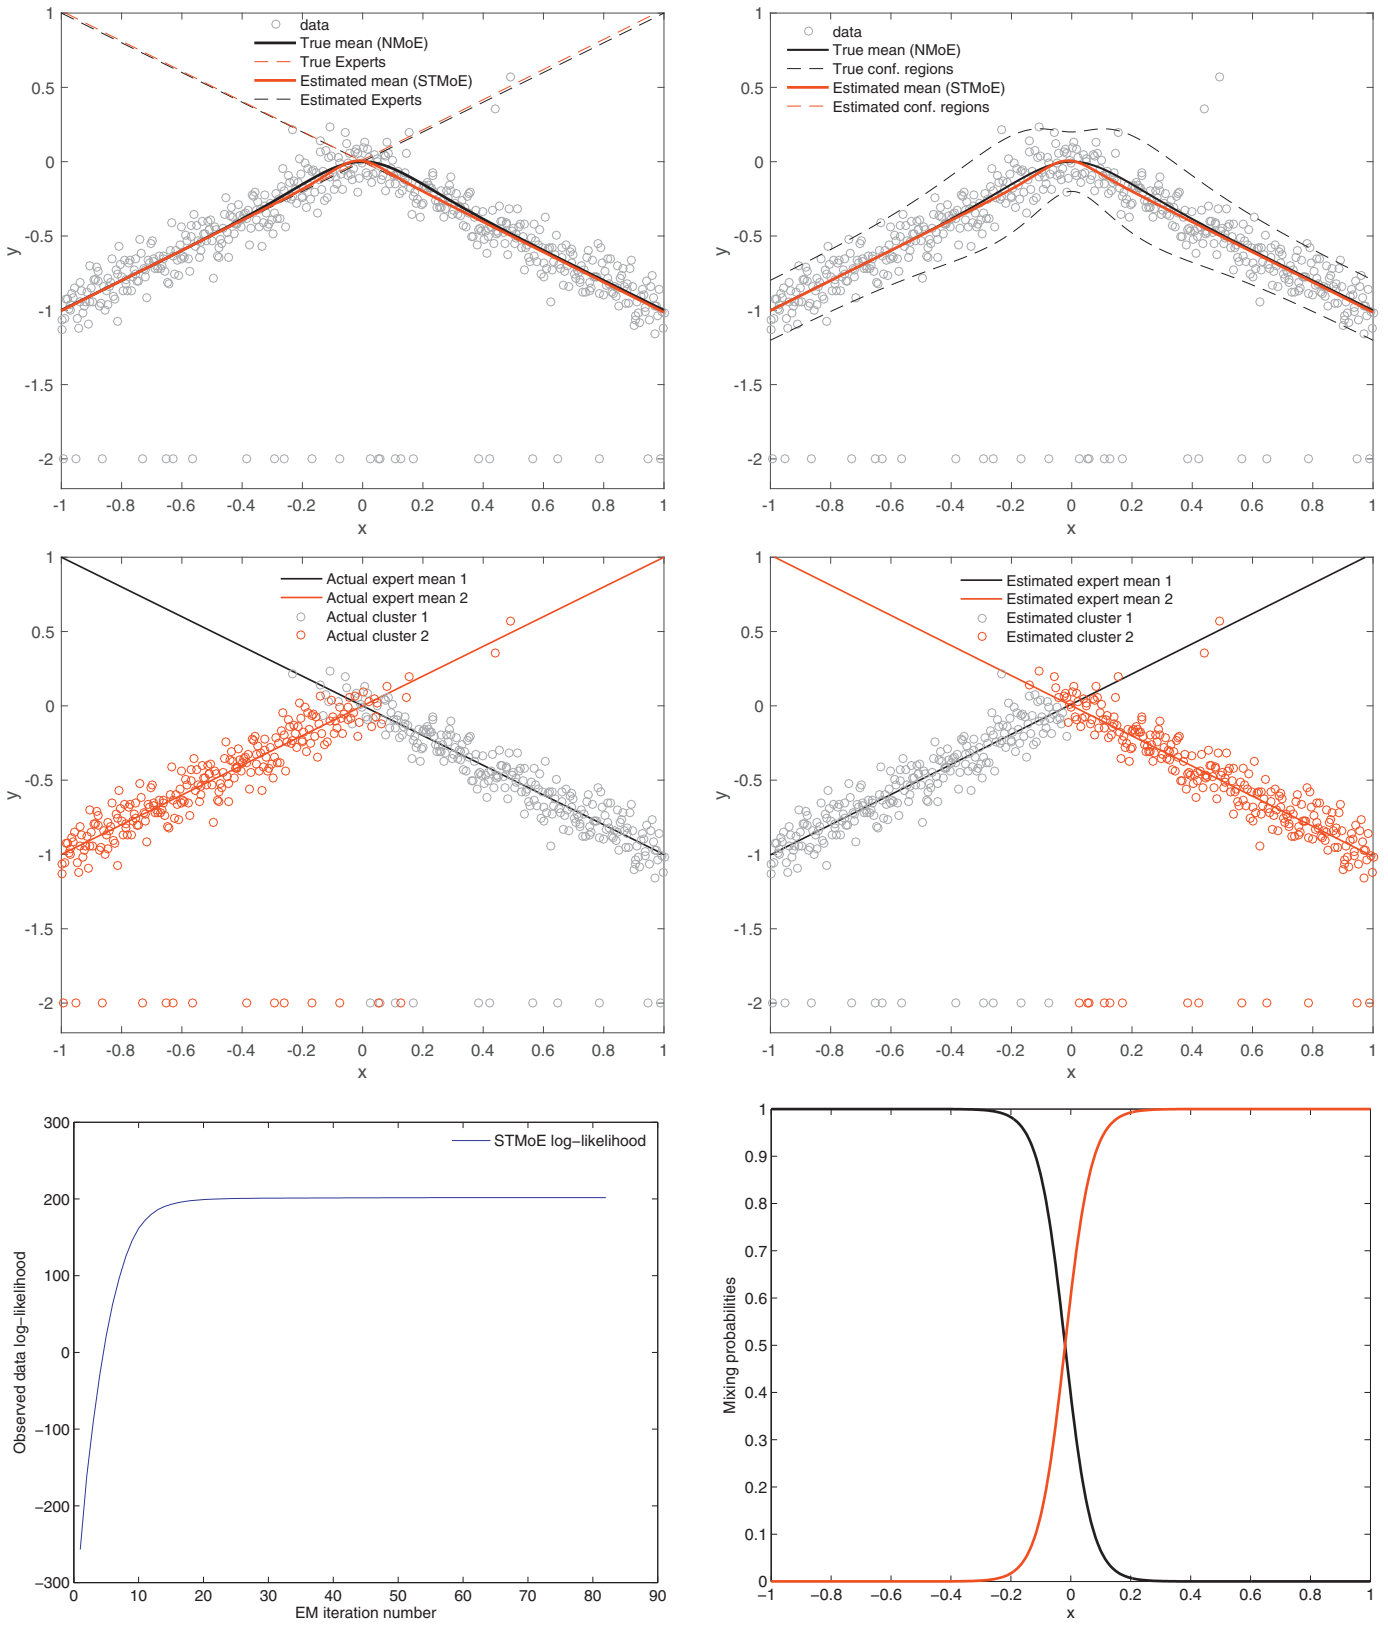

Fig. 6. Fitted STMoE model to a data set of $n = 500$ observations generated according to the NMoE model and including 5% of outliers.

ponents is also the one provided by the model selection criteria as shown later in Table 8. Indeed, as mentioned by [48], [25] found that the data could be segmented into two periods of global warming (before 1940 and after 1965), separated by a transition period where there was a slight global cooling (i.e. 1940 to 1965). Documentation of the basic analysis method is provided by [24,25]. We set the response $y_i (i = 1, \ldots, 135)$ as the temperature anomalies and the covariates $\boldsymbol{x}_i = \boldsymbol{r}_i = (1, x_i)^T$ where x_i is the year of the ith observation.

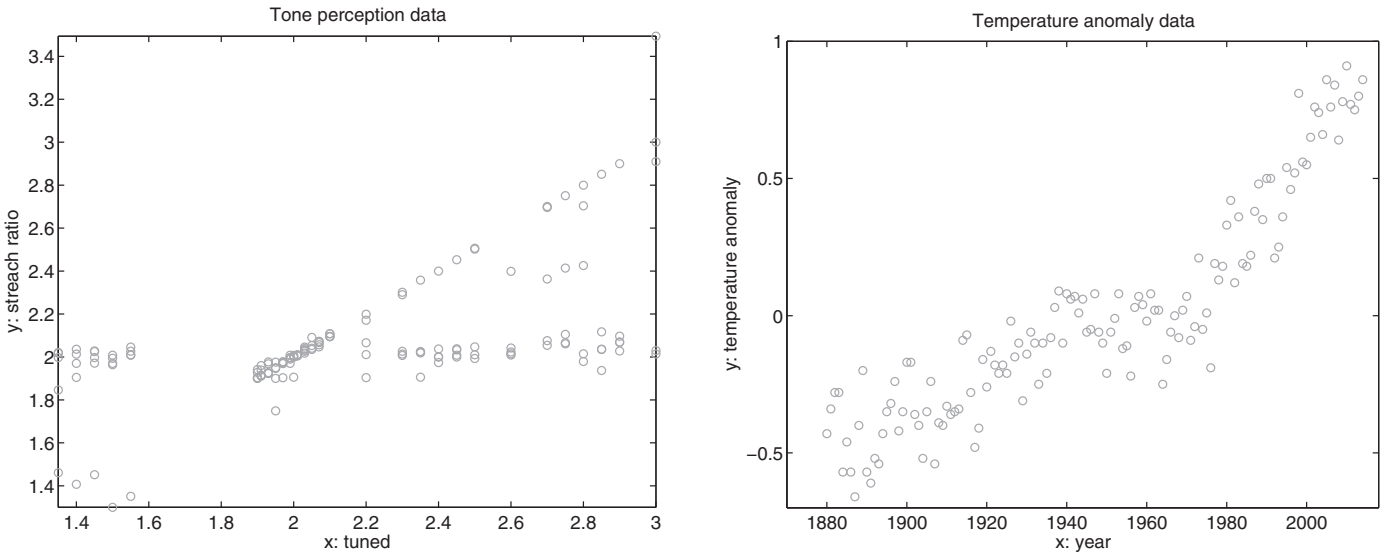

Fig. 7. Scatter plots of the tone perception data and the temperature anomalies data.

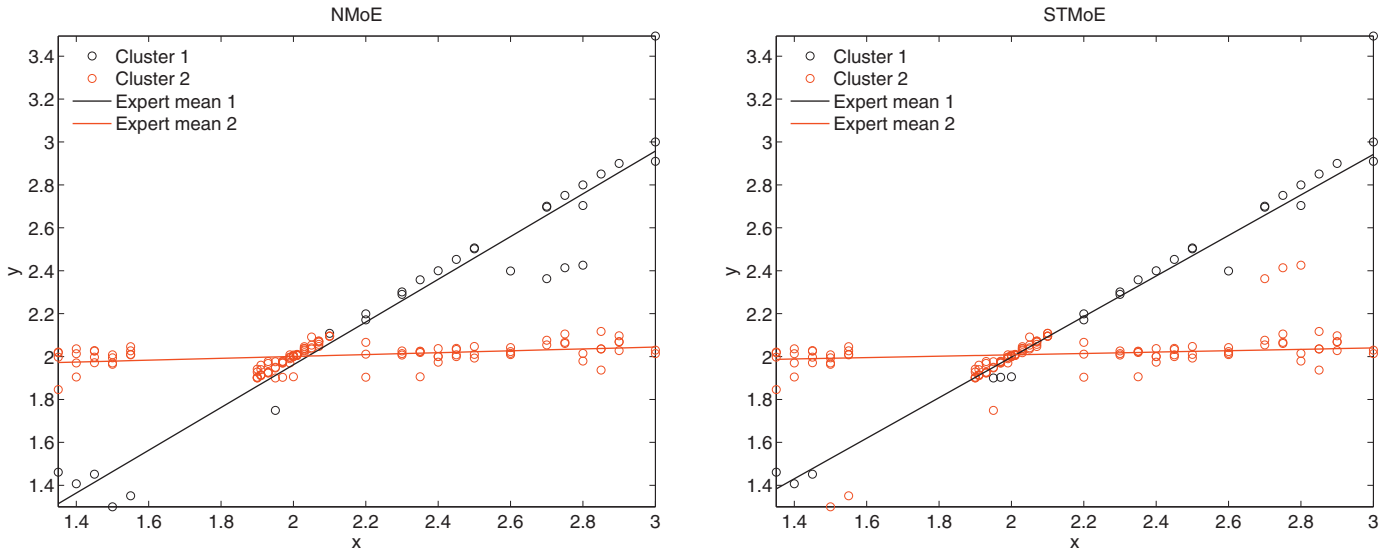

Fig. 8. The fitted MoLE to the original tone data set with the NMoE model (left) and the STMoE model (right). The predictor x is the actual tone ratio and the response y is the perceived tone ratio.

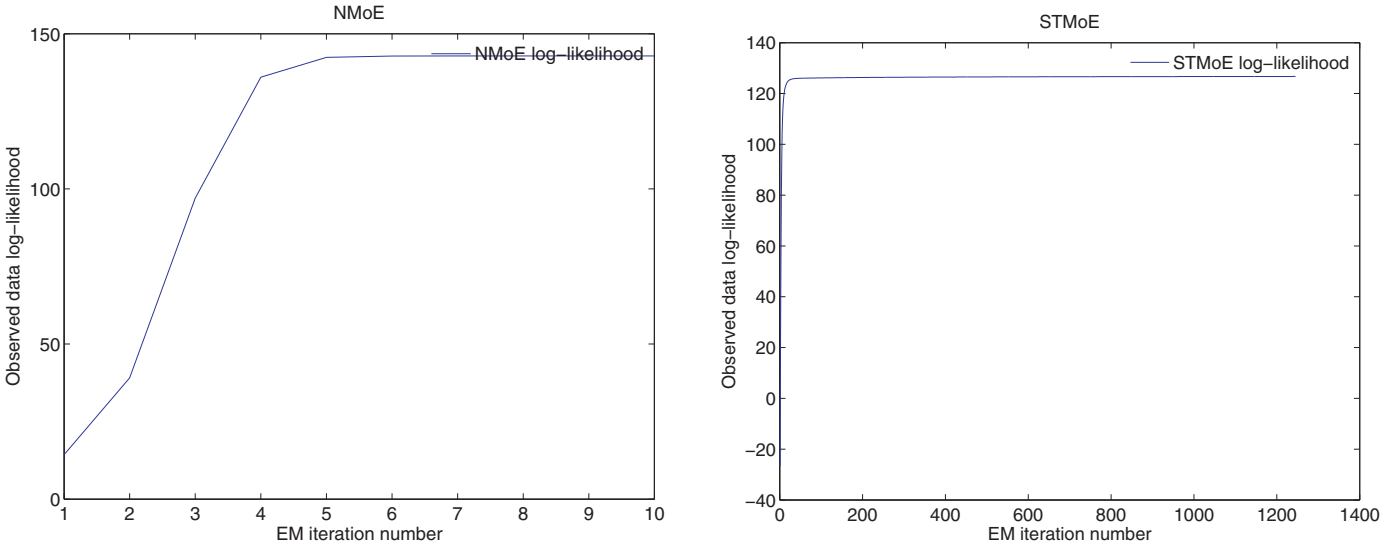

Fig. 9. The log-likelihood during the EM iterations when fitting the MoLE models to the original tone data set with the NMoE model (left) and the STMoE model (right).

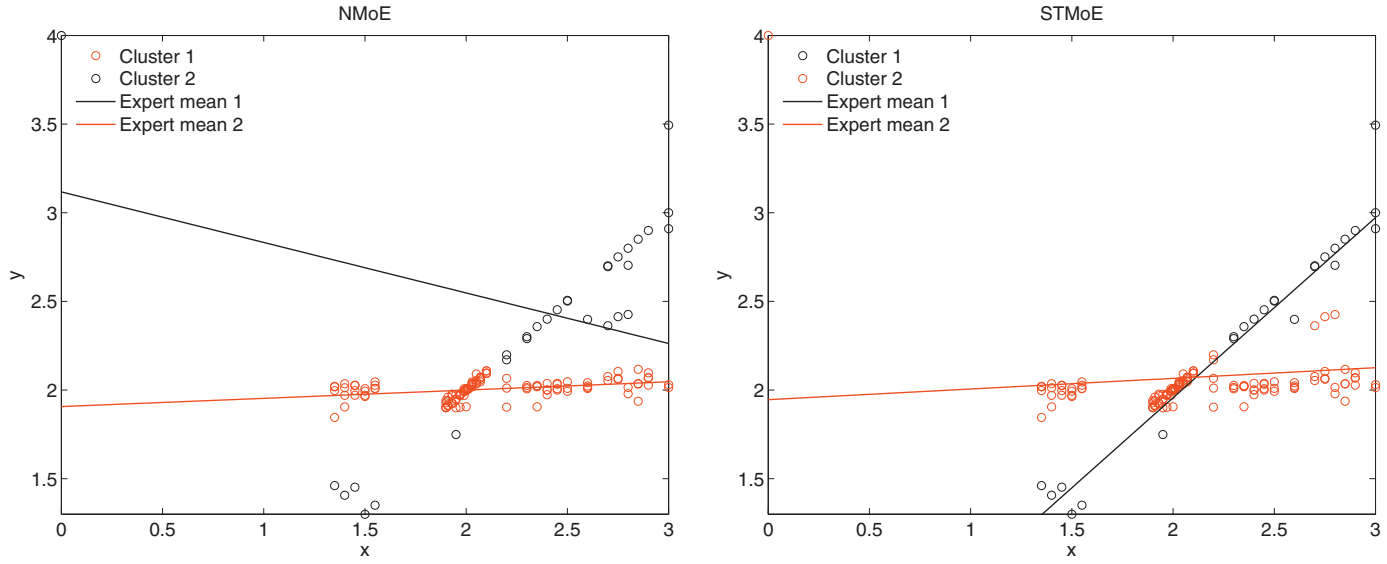

Fig. 10. Fitting MoLE to the tone data set with ten added outliers (0, 4) with the NMoE model fit (left) and the STMoE model fit (right). The predictor x is the actual tone ratio and the response y is the perceived tone ratio.

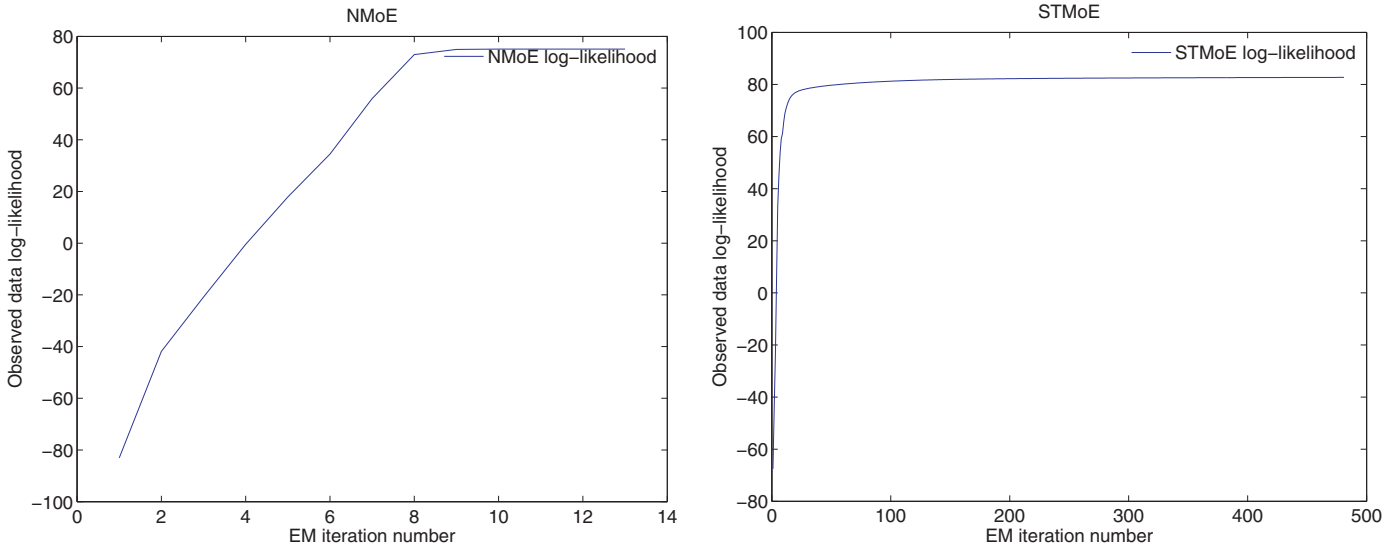

Fig. 11. The log-likelihood during the EM iterations when fitting the MoLE models to the tone data set with ten added outliers (0, 4) with the NMoE model (left) and the STMoE model (right).

Table 7
Values of the estimated MoE parameters for the temperature anomalies data set.

Param. model	α_{10}	α_{11}	β_{10}	β_{11}	β_{20}	β_{21}	σ_1	σ_2	λ_1	λ_2	ν_1	ν_2
NMoE	946.483	−0.481	−12.805	0.006	−41.073	0.020	0.115	0.110	–	–	–	–
STMoE	931.966	−0.474	−12.848	0.006	−40.876	0.020	0.113	0.105	0.024	-0.015	41.048	17.589

Table 8
Choosing the number of expert components K for the temperature anomalies data by using the information criteria BIC, AIC, and ICL. Underlined value indicates the highest value for each criterion.

	NMoE			STMoE		
K	BIC	AIC	ICL	BIC	AIC	ICL
1	46.062	50.420	46.062	40.971	48.234	40.971
2	**79.916**	91.537	**79.624**	**69.638**	**87.069**	**69.341**
3	71.396	90.280	58.487	54.126	81.726	30.655
4	66.727	92.875	54.752	42.308	80.0773	20.494
5	59.510	**92.920**	51.242	28.037	75.974	-8.881

Figs. 12–14 respectively show, for each of the two MoE models, the two fitted linear expert components, the corresponding means and confidence regions computed as plus and minus twice the estimated (pointwise) standard deviation as presented in Section 5, and the log-likelihood profiles. One can observe that the two models are successfully applied on the data set and provide very similar results. These results are also similar to those found by [48] who used a Laplace mixture of linear experts. The values of estimated MoE parameters for the temperature anomalies data set are given in Table 7. One can see that the parameters common for the two models are quasi-identical. It can also be seen the STMoE

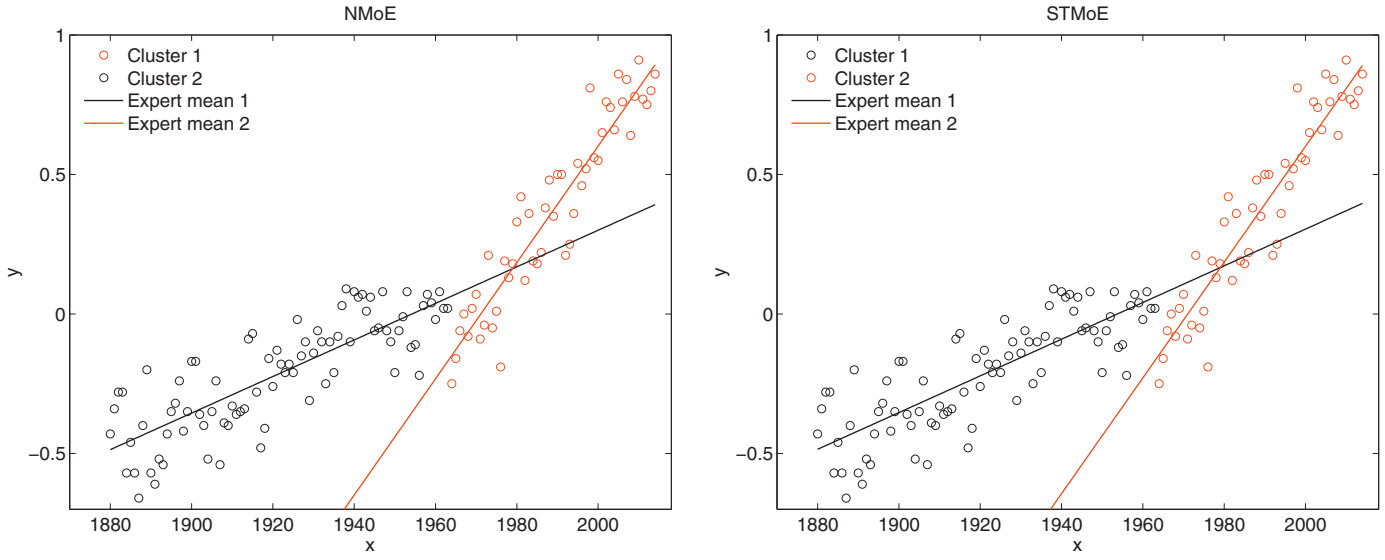

Fig. 12. Fitting the MoLE models to the temperature anomalies data set with the NMoE model (left) and the STMoE model (right). The predictor x is the year and the response y is the temperature anomaly.

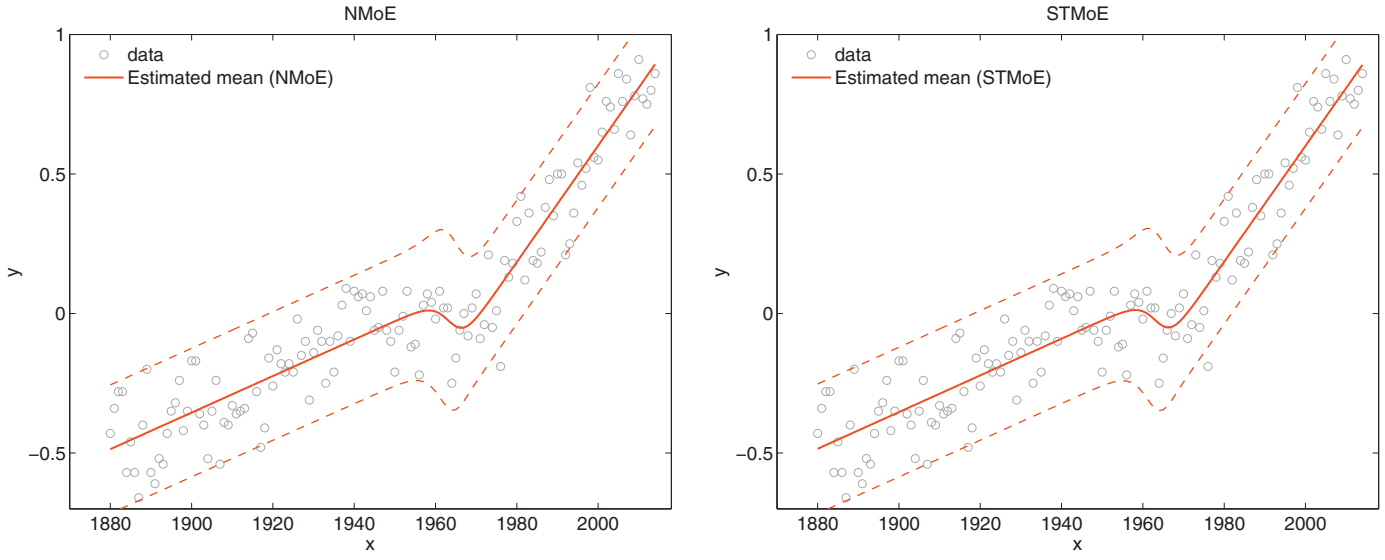

Fig. 13. The fitted MoLE models to the temperature anomalies data set with the NMoE model fit (left) and the STMoE model fit (right). The predictor x is the year and the response y is the temperature anomaly. The shaded region represents plus and minus twice the estimated (pointwise) standard deviation as presented in Section 5.

model provides a solution with a skewness close to zero. This may support the hypothesis of non-asymmetry for this data set. Then, the STMoE solution provides a degrees of freedom more than 17, which tends to approach a normal distribution. On the other hand, the regression coefficients are also similar to those found by [48] who used a Laplace mixture of linear experts.

We also performed a model selection procedure on the temperature anomalies data set to choose the best number of MoE components from values between 1 and 5. Table 8 gives the obtained values of the used model selection criteria, that is BIC, AIC, and ICL. One can see that, except the result provided by the AIC for the NMoE model which provide a high number of components, all the others results provide evidence for two components in the data.

9. Concluding remarks and future work

In this paper we proposed a new non-normal MoE model, which generalizes the normal MoE model and attempts to simultaneously accommodate heavy tailed data with possible outliers and asymmetric distribution. The proposed STMoE is based on the flexible skew t distribution that is suggested for possibly non-symmetric, heavy tailed and noisy data. We developed a CEM algorithm for model inference and described the use of the model in non-linear regression and prediction as well as in model-based clustering. The developed model was successfully applied and validated on simulation studies and two real data sets. The results obtained on simulated data confirm the good performance of the model in terms of density estimation, non-linear regression function approximation and clustering. In addition, the simulation results provide evidence of the robustness of the STMoE model to outliers, compared to the standard alternative NMoE model. The proposed model was also successfully applied to two different real data sets, including situations with outliers. The model selection using information criteria tends to promote using BIC and ICL against AIC which may perform poorly in the analyzed data. The obtained results support the potential benefit of the proposed approach for practical applications.

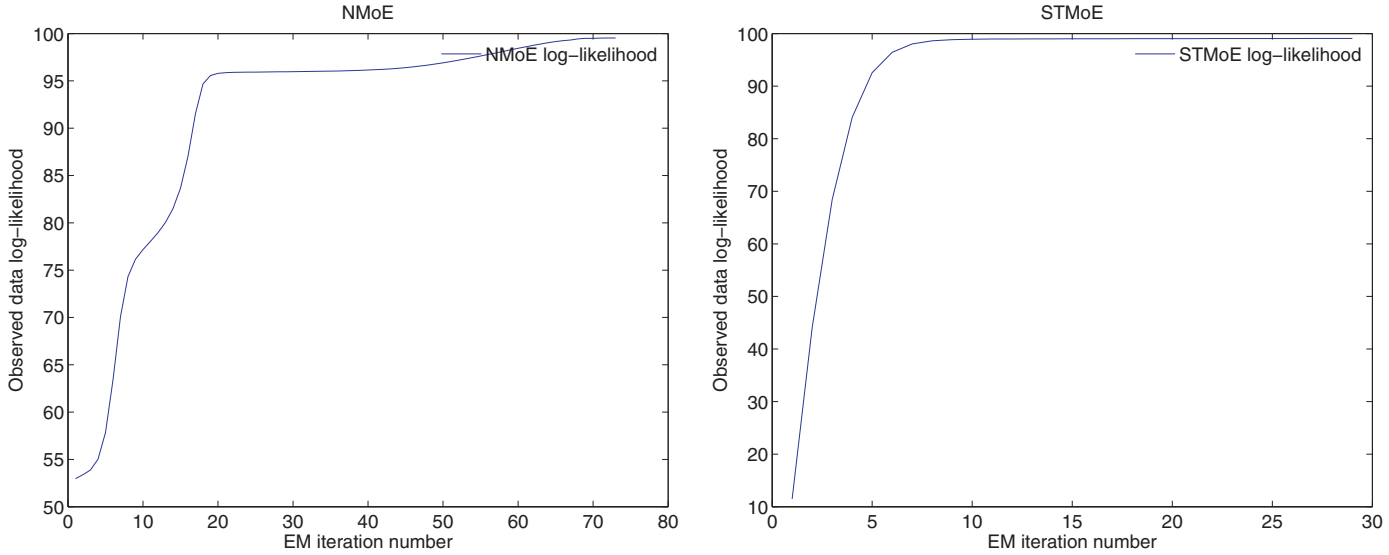

Fig. 14. The log-likelihood during the EM iterations when fitting the MoLE models to the temperature anomalies data set with the NMoE model (left) and the STMoE model (right).

One interesting future direction is to extend the proposed model to the hierarchical mixture of experts framework [34]. Another natural future extension of this work is also to consider the case of MoE for multiple regression on multivariate responses rather than simple univariate regression. In that case, one may consider the multivariate skew-t and the multivariate Normal inverse Gaussian distribution [50] which may be more stable in high-dimensional settings compared to the multivariate skew-t.

Appendix A. The skew-normal distribution

As introduced by [2,3], a random variable Y follows a univariate skew-normal distribution with location parameter $\mu \in \mathbb{R}$, scale parameter $\sigma^2 \in (0, \infty)$ and skewness parameter $\lambda \in \mathbb{R}$ if it has the density

$$f(y; \mu, \sigma^2, \lambda) = \frac{2}{\sigma}\phi\left(\frac{y-\mu}{\sigma}\right)\Phi\left(\lambda\left(\frac{y-\mu}{\sigma}\right)\right) \tag{A.1}$$

where $\phi(.)$ and $\Phi(.)$ denote, respectively, the probability density function (pdf) and the cumulative distribution function (cdf) of the standard normal distribution. It can be seen from (A.1) that when $\lambda = 0$, the skew-normal reduces to the normal distribution. As presented by [3,26], if

$$Y = \mu + \delta|U| + \sqrt{1-\delta^2}E \tag{A.2}$$

where $\delta = \frac{\lambda}{\sqrt{1+\lambda^2}}$, U and E are independent random variables following the normal distribution $\mathrm{N}(0, \sigma^2)$, then Y follows the skew-normal distribution with pdf $\mathrm{SN}(\mu, \sigma^2, \lambda)$ given by (A.1). In the above, $|U|$ denotes the magnitude of U. This stochastic representation of the skew-normal distribution leads to the following hierarchical representation in an incomplete data framework, as presented in [42]:

$$\begin{aligned} &Y|u \sim \mathrm{N}\left(\mu + \delta|u|, (1-\delta^2)\sigma^2\right), \\ &U \sim \mathrm{N}(0, \sigma^2). \end{aligned} \tag{A.3}$$

Appendix B. Stochastic representation of the STMoE model

The skew t mixture of experts model is characterized as follows. Suppose that conditional on a categorical variable $Z_i = z_i \in \{1, \ldots, K\}$ representing the hidden label of the component generating the ith observation and which, conditional on some predictor $\boldsymbol{r}_i$, follows the multinomial distribution (B.1):

$$Z_i|\boldsymbol{r}_i \sim \mathrm{Mult}\left(1; \pi_1(\boldsymbol{r}_i; \boldsymbol{\alpha}), \ldots, \pi_K(\boldsymbol{r}_i; \boldsymbol{\alpha})\right) \tag{B.1}$$

where each of the probabilities $\pi_{z_i}(\boldsymbol{r}_i; \boldsymbol{\alpha}) = \mathbb{P}(Z_i = z_i|\boldsymbol{r}_i)$ is given by the multinomial logistic function (2). Now suppose a random variable Y_i having the following representation:

$$Y_i = \mu(\boldsymbol{x}_i; \boldsymbol{\beta}_{z_i}) + \sigma_{z_i}\frac{E_i}{\sqrt{W_i}} \tag{B.2}$$

where E_i and W_i are independent univariate random variables with, respectively, a standard skew-normal distribution $E_i \sim \mathrm{SN}(\lambda_{z_i})$, and a Gamma distribution $W_i \sim \mathrm{Gamma}(\frac{\nu_{z_i}}{2}, \frac{\nu_{z_i}}{2})$, and $\boldsymbol{x}_i$ and $\boldsymbol{r}_i$ are some given covariate variables. Then, the variable Y_i is said to follow the skew t mixture of experts (STMoE) defined by (13).

References

[1] H. Akaike, A new look at the statistical model identification, IEEE Trans. Autom. Control 19 (6) (1974) 716–723.
[2] A. Azzalini, A class of distributions which includes the normal ones, Scand. J. Stat. (1985) 171–178.
[3] A. Azzalini, Further results on a class of distributions which includes the normal ones, Scand. J. Stat. (1986) 199–208.
[4] A. Azzalini, A. Capitanio, Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t distribution, J. R. Stat. Soc. Ser. B 65 (2003) 367–389.
[5] X. Bai, W. Yao, J.E. Boyer, Robust fitting of mixture regression models, Comput. Stat Data Anal. 56 (7) (2012) 2347–2359.
[6] C. Biernacki, G. Celeux, G. Govaert, Assessing a mixture model for clustering with the integrated completed likelihood, IEEE Trans. Pattern Anal. Mach. Intell. 22 (7) (2000) 719–725.
[7] C. Bishop, M. Svensén, Bayesian hierarchical mixtures of experts., in: Proceeding UAI'03 Proceedings of the Nineteenth conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann Publishers Inc., San Francisco, CA, 2003, pp. 57–64.
[8] R.P. Brent, Algorithms for minimization without derivatives. Chapter 5, Prentice-Hall, Englewood Cliffs, NJ, 1973.
[9] F. Chamroukhi, Non-normal mixtures of experts, arXiv:1506.06707v2 [stat.ME]. (2015).
[10] F. Chamroukhi, Robust mixture of experts modeling using the t distribution, Neural Netw. 79 (2016) 20–36.
[11] F. Chamroukhi, A. Samé, G. Govaert, P. Aknin, A regression model with a hidden logistic process for feature extraction from time series, in: Proceedings of the International Joint Conference on Neural Networks (IJCNN), Atlanta, GA, 2009, pp. 489–496.

[12] F. Chamroukhi, A. Samé, G. Govaert, P. Aknin, Time series modeling by a regression approach based on a latent process, Neural Netw. 22 (5–6) (2009) 593–602.
[13] K. Chen, L. Xu, H. Chi, Improved learning algorithms for mixture of experts in multiclass classification, Neural Netw. 12 (9) (1999) 1229–1252.
[14] E.A. Cohen, Some effects of inharmonic partials on interval perception, Music Percept. 1 (1984).
[15] A.P. Dempster, N.M. Laird, D.B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, J. R. Stat. Soc. B 39(1) (1977) 1–38.
[16] S. Faria, G. Soromenho, Fitting mixtures of linear regressions, J. Stat. Comput. Simul. 80 (2) (2010) 201–225.
[17] S. Frühwirth-Schnatter, Finite Mixture and Markov Switching Models (Springer Series in Statistics), Springer Verlag, New York, 2006.
[18] S. Frühwirth-Schnatter, S. Pyne, Bayesian inference for finite mixtures of univariate and multivariate skew-normal and skew-t distributions, Biostatistics 11 (2) (2010) 317–336.
[19] S. Gaffney, P. Smyth, Trajectory clustering with mixtures of regression models, in: Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining, ACM Press, 1999, pp. 63–72.
[20] I.C. Gormley, T.B. Murphy, A mixture of experts model for rank data with applications in election studies, Ann. Appl. Stat. 2 (4) (2008) 1452–1477.
[21] I.C. Gormley, T.B. Murphy, A mixture of experts latent position cluster model for social network data, Stat. Methodol. 7 (3) (2010) 385–405.
[22] I.C. Gormley, T.B. Murphy, Mixture of experts modelling with social science applications, J. Comput. Gr. Stat. 19 (2) (2011) 332–353.
[23] P. Green, Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives, J. R. Stat. Soc. B 46 (2) (1984) 149–192.
[24] J. Hansen, R. Ruedy, J. Glascoe, M. Sato, Giss analysis of surface temperature change, J. Geophys. Res. 104 (1999) 30997–31022.
[25] J. Hansen, R. Ruedy, S. M., M. Imhoff, W. Lawrence, D. Easterling, T. Peterson, T. Karl, A closer look at united states and global surface temperature change, J. Geophys. Res. 106 (2001) 23947–23963.
[26] N. Henze, A probabilistic representation of the skew-normal distribution, Scand. J. Stat. (1986) 271–275.
[27] D. Hunter, D. Young, Semiparametric mixtures of regressions, J. Nonparametric Stat. 24 (1) (2012) 19–38.
[28] D.R. Hunter, K. Lange, A tutorial on MM algorithms, Am. Stat. 58 (2004) 30–37.
[29] S. Ingrassia, S. Minotti, G. Vittadini, Local statistical modeling via a cluster-weighted approach with elliptical distributions, J. Classif. 29 (3) (2012) 363–401.
[30] R.M.-S. J. Q. Shi, D.M. Titterington, Hierarchical gaussian process mixtures for regression, Stat. Comput. 15 (1) (2005) 31–41.
[31] R.A. Jacobs, M.I. Jordan, S.J. Nowlan, G.E. Hinton, Adaptive mixtures of local experts, Neural Comput. 3 (1) (1991) 79–87.
[32] W. Jiang, M.A. Tanner, On the identifiability of mixtures-of-experts, Neural Netw. 12 (1999) 197–220.
[33] P.N. Jones, G.J. McLachlan, Fitting finite mixture models in a regression context, Aust. J. Stat. 34 (2) (1992) 233–240.
[34] M.I. Jordan, R.A. Jacobs, Hierarchical mixtures of experts and the EM algorithm, Neural Comput. 6 (1994) 181–214.
[35] M.I. Jordan, L. Xu, Convergence results for the EM approach to mixtures of experts architectures, Neural Netw. 8 (9) (1995) 1409–1431.
[36] S.X. Lee, G.J. McLachlan, Model-based clustering and classification with non-normal mixture distributions, Stat. Methods Appl. 22 (4) (2013) 427–454.
[37] S.X. Lee, G.J. McLachlan, On mixtures of skew normal and skew *t*-distributions, Adv. Data Anal. Classif. 7 (3) (2013) 241–266.
[38] S.X. Lee, G.J. McLachlan, Finite mixtures of multivariate skew t-distributions: some recent and new results, Stat. Comput. 24 (2) (2014) 181–202.
[39] S.X. Lee, G.J. McLachlan, Finite mixtures of canonical fundamental skew *t*-distributions, Stat. Comput. 26 (3) (2016) 573–589.
[40] T.I. Lin, Robust mixture modeling using multivariate skew *t* distributions, Stat. Comput. 20 (3) (2010) 343–356.
[41] T.I. Lin, J.C. Lee, W.J. Hsieh, Robust mixture modeling using the skew *t* distribution, Stat. Comput. 17 (2) (2007) 81–92.
[42] T.I. Lin, J.C. Lee, S.Y. Yen, Finite mixture modelling using the skew normal distribution, Stat. Sin. 17 (2007) 909–927.
[43] G.J. McLachlan, T. Krishnan, The EM Algorithm and Extensions, 2nd ed., Wiley, Hoboken, NJ, 2008.
[44] G.J. McLachlan, D. Peel., Finite Mixture Models, Wiley, Hoboken, NJ, 2000.
[45] X.-L. Meng, D. van Dyk, The EM algorithm-an old folk song sung to the fast tune (with discussion), J. R. Stat. Soc. Ser. B 4 (59) (1997) 511–567.
[46] X.L. Meng, D.B. Rubin, Maximum likelihood estimation via the ECM algorithm: a general framework, Biometrika 80 (2) (1993) 267–278.
[47] S.-K. Ng, G.J. McLachlan, Using the em algorithm to train neural networks: misconceptions and a new algorithm for multiclass classification., IEEE Trans. Neural Netw. 15 (3) (2004) 738–749.
[48] H.D. Nguyen, G.J. McLachlan, Laplace mixture of linear experts, Comput. Stat. Data Anal. 93 (2016) 177–191.
[49] A. O'Hagan, T.B. Murphy, I.C. Gormley, On estimation of parameter uncertainty in model-based clustering, arXiv:1510.00551 [stat.CO]. (2015).
[50] A. O'Hagan, T.B. Murphy, I.C. Gormley, P.D. McNicholas, D. Karlis, Clustering with the multivariate normal inverse gaussian distribution, Comput. Stat. Data Anal. 93 (2016) 18–30.
[51] S. Pyne, X. Hu, K. Wang, E. Rossin, T. Lin, L. Maier, C. Baecher-Allan, G. McLachlan, P. Tamayo, D. Hafler, P. De Jager, J. Mesirow, Automated high-dimensional flow cytometric data analysis, Proc. Nat. Acad. Sci. USA (2009).
[52] R.E. Quandt, A new approach to estimating switching regressions, J. Am. Stat. Assoc. 67 (338) (1972) 306–310.
[53] R.E. Quandt, J.B. Ramsey, Estimating mixtures of normal distributions and switching regressions, J. Am. Stat. Assoc. 73 (364) (1978) 730–738.
[54] C.E. Rasmussen, Z. Ghahramani, Infinite mixtures of gaussian process experts, in: Advances in Neural Information Processing Systems 14, MIT Press, 2001, pp. 881–888.
[55] R. Ruedy, M. Sato, K. Lo, NASA GISS surface temperature (GISTEMP) analysis, 2015, Center for Climate Systems Research, NASA Goddard Institute for Space Studies 2880 Broadway, New York, NY 10025 USA, (DOI:10.3334/CDIAC/cli.001).
[56] G. Schwarz, Estimating the dimension of a model, Ann. Stat. 6 (1978) 461–464.
[57] J.Q. Shi, T. Choi, Gaussian Process Regression Analysis for Functional Data, Chapman & Hall/CRC Press, 2011.
[58] W. Song, W. Yao, Y. Xing, Robust mixture regression model fitting by laplace distribution, Comput. Stat. Data Anal. 71 (0) (2014) 128–137.
[59] R.D.D. Veaux, Mixtures of linear regressions, Comput. Stat. Data Anal. 8 (3) (1989) 227–245.
[60] K. Viele, B. Tong, Modeling with mixtures of linear regressions, Stat. Comput. 12 (2002) 315–330.
[61] S. Waterhouse, D. Mackay, T. Robinson, in: Bayesian methods for mixtures of experts, MIT Press, 1996, pp. 351–357.
[62] S.R. Waterhouse, Classification and regression using Mixtures of Experts, Department of Engineering, Cambridge University, 1997 Ph.D. thesis.
[63] Y. Wei, Robust mixture regression models using t-distribution, Technical Report, Master Report, Department of Statistics, Kansas State University, 2012.
[64] D. Young, D. Hunter, Mixtures of regressions with predictor-dependent mixing proportions, Comput. Stat. Data Anal. 55 (10) (2010) 2253–2266.
[65] S.E. Yuksel, J.N. Wilson, P.D. Gader, Twenty years of mixture of experts., IEEE Trans. Neural Netw. Learn Syst. 23 (8) (2012) 1177–1193.
[66] C.B. Zeller, C.R.B. Cabral, V.H. Lachos, Robust mixture regression modelling based on scale mixtures of skew-normal distributions, Test 25 (2) (2016) 375–396.
[67] F. Li, M. Villani, R. Kohn, Flexible modeling of conditional distributions using smooth mixtures of asymmetric student *t* densities, J. Stat. Plan. Inference 140 (12) (2010) 3638–3654.

Faicel Chamroukhi is since 2016 a Full Professor of Statistics and Data Science at University of Caen, department of Mathematics and Computer Science and the Lab of Mathematics Nicolas Oresme (LMNO) UMR CNRS 6139. From 2011 to 2015, he has been an Associate Professor at University of Toulon and the Information Sciences and Systems Lab. (LSIS) UMR CNRS 7296. In 2015–2016 he was awarded a CNRS research leave at the Lab of Mathematics Paul Painlevé, UMR CNRS 8524 in Lille, where he has also been an Invited Scientist at INRIA. He received his Master degree of Engineering Sciences from Pierre & Marie Curie (Paris 6) University in 2007 and his Ph.D. degree in applied mathematics and computer science, in the area of statistical leaning and data analysis from Compiègne University of Technology in 2010. In 2011, he was qualified for the position of Associate Professor in applied mathematics (CNU 26), computer science (CNU 27), and signal processing (CNU 61). In 2015, he received his Accreditation to Supervise Research (HDR) in applied mathematics and computer science, in the area of statistical learning and data science, from Toulon University. In 2016, he was qualified for the position of Professor in the area of applied mathematics, computer science, and signal processing (CNU 26, 27, 61). His multidisciplinary research is in the area of Data Science and includes statistics, machine learning and statistical signal processing, with a particular focus of the statistical methodology and inference of latent data models for complex heterogeneous high dimensional and massive data, temporal data, functional data, and their application to real-world problems including dynamical systems, acoustic/speech processing, life sciences (medicine, biology), information retrieval, social networks.