



HAL
open science

Functional Mixtures-of-Experts

Faïcel Chamroukhi, Nhat Thien Pham, van Hà Hoang, Geoffrey Mclachlan

► **To cite this version:**

Faïcel Chamroukhi, Nhat Thien Pham, van Hà Hoang, Geoffrey Mclachlan. Functional Mixtures-of-Experts. *Statistics and Computing*, 2024, 34 (3), pp.98. 10.1007/s11222-023-10379-0. hal-04360919v2

HAL Id: hal-04360919

<https://hal.science/hal-04360919v2>

Submitted on 22 Mar 2024

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

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

Public Domain



Functional mixtures-of-experts

Faïcel Chamroukhi^{1,2} · Nhat Thien Pham² · Van Hà Hoang³ · Geoffrey J. McLachlan⁴

Received: 2 February 2022 / Accepted: 22 December 2023
© The Author(s) 2024

Abstract

We consider the statistical analysis of heterogeneous data for prediction, in situations where the observations include functions, typically time series. We extend the modeling with mixtures-of-experts (ME), as a framework of choice in modeling heterogeneity in data for prediction with vectorial observations, to this functional data analysis context. We first present a new family of ME models, named functional ME (FME), in which the predictors are potentially noisy observations, from entire functions. Furthermore, the data generating process of the predictor and the real response, is governed by a hidden discrete variable representing an unknown partition. Second, by imposing sparsity on derivatives of the underlying functional parameters via Lasso-like regularizations, we provide sparse and interpretable functional representations of the FME models called iFME. We develop dedicated expectation–maximization algorithms for Lasso-like regularized maximum-likelihood parameter estimation strategies to fit the models. The proposed models and algorithms are studied in simulated scenarios and in applications to two real data sets, and the obtained results demonstrate their performance in accurately capturing complex nonlinear relationships and in clustering the heterogeneous regression data.

Keywords Mixtures-of-experts · Functional data analysis · EM algorithms · Maximum likelihood estimation · Lasso regularization · Clustering

Contents

1	Introduction	...	3.4	Non-linear regression and clustering with FME models	...
2	Functional mixtures-of-experts (FME)	...	3.5	Tuning parameters and model selection	...
2.1	ME with functional predictor and scalar response	...	4	Experimental study	...
2.2	Smoothing representation of the functional experts	...	4.1	Evaluation criteria	...
2.3	Smoothing representation of the functional gating network	...	4.2	Simulation studies	...
2.4	The FME model conditional density	...	4.2.1	Simulation parameters and experimental protocol	...
2.5	Maximum likelihood estimation via the EM algorithm	...	4.2.2	Some implementation details	...
3	Regularized maximum likelihood estimation	...	4.2.3	Simulation results	...
3.1	ℓ_1 -regularization and the EM-Lasso algorithm	...	4.3	Application to real data	...
3.2	Interpretable Functional Mixture of Experts (iFME)	...	4.3.1	Canadian weather data	...
3.2.1	Interpretable sparse regularization	...	4.3.2	Diffusion tensor imaging data for multiple sclerosis subjects	...
3.2.2	The iFME model	...	5	Conclusion and discussion	...
3.3	Regularized MLE via the EM-iFME algorithm	...	A	EM for the FME model	...
			B	EM-Lasso for ℓ_1 -regularized MLE of the FME model	...
			B.1	Updating the gating network parameters	...
			B.2	Updating the experts network parameters	...
			C	EM-iFME for updating iFME model parameters	...
			C.1	Updating the gating network parameters	...
			C.2	Updating the expert network parameters	...
			D	Detailed data generating protocol, simulation parameters and experimental protocol	...
			D.1	Data generating protocol	...
			D.2	Simulation parameters and experimental protocol	...
			References

✉ Faïcel Chamroukhi
faïcel.chamroukhi@irt-systemx.fr

¹ IRT SystemX, 2 boulevard Thomas Gobert, 91120 Palaiseau, France

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

³ University of Science, Vietnam National University, Ho Chi Minh City, Vietnam

⁴ School of Mathematics and Physics, University of Queensland, Brisbane, Australia

1 Introduction

Mixture-of-experts (ME), introduced by Jacobs et al. (1991), is a successful and flexible supervised learning architecture that allows one to efficiently represent complex non-linear relationships in observed pairs of heterogeneous data (\mathbf{X}, Y) . The ME model relies on the divide and conquer principle, so that the response Y is gathered from soft-association of several expert responses, each targeted to a homogeneous sub-population of the heterogeneous population, given the input covariates (predictors or features) \mathbf{X} . From the statistical modeling point of view, a ME model is an extension of the finite mixture model (McLachlan and Peel 2000) which explores the unconditional (mixture) distribution of a given set of the features \mathbf{X} . It is thus more tailored to unsupervised learning than to supervised learning and it has the fully conditional mixture model of the form

$$\text{ME}(y|\mathbf{x}) = \sum_{k=1}^K G_k(\mathbf{x})E_k(y|\mathbf{x}). \quad (1)$$

In model (1) the ME distribution of the response y given the predictors \mathbf{x} is a conditional mixture distribution with predictor-dependent mixing weights, referred to as gating functions, $G_k(\mathbf{x})$, and conditional mixture components, referred to as experts $E_k(y|\mathbf{x})$, K being the number of experts.

Mixture of experts (ME) models thus allow one to better capture more complex relationships between y and \mathbf{x} in heterogeneous situations in non-linear regression $y \in \mathbb{R}$, classification $y \in \{1, \dots, G\}$, and in clustering the data by associating each expert component to a cluster. The richness of the class of ME models in terms of conditional density approximation capabilities has been recently demonstrated by proving denseness results (Nguyen et al. 2021a, 2019).

They have been investigated in their simple form, as well as in their hierarchical form (Jordan and Jacobs 1994), for non-linear regression and model-based cluster and discriminant analyses and in different application domains. The inference in this case can be performed by maximum likelihood estimation (MLE) via the expectation–maximization (EM) algorithm (Jordan and Jacobs 1994; McLachlan and Krishnan 2008; Dempster et al. 1977) or, when p is possibly larger than the sample size n , by regularized MLE via dedicated EM-Lasso algorithms as in Khalili (2010), Montuelle et al. (2014), Chamroukhi et al. (2019), Chamroukhi and Huynh (2019), Huynh and Chamroukhi (2019) and Nguyen et al. (2020) which include Lasso-like penalties (Tibshirani 1996). For a more complete review of ME models, the reader is referred to Yuksel et al. (2012) and Nguyen and Chamroukhi (2018). More recent theoretical results about the ME estimation and model selection for different families of ME models, can be found in Nguyen et al. (2020).

To the best of our knowledge, ME models have been exclusively studied in multivariate analysis when the inputs are vectors, i.e. $\mathbf{X} \in \mathcal{X} = \mathbb{R}^p$. However, in many problems, the predictors and/or the responses are observed from smooth functions. Indeed, in many situations, unlike in predictive and cluster analyses of multivariate and potentially high-dimensional heterogeneous data, which have been studied with the ME modeling in (1), the observed data may arise from continuously observed processes, e.g. time series. Thus, a multivariate (vectorial) analysis does not allow one to enough capture the inherent functional structure of the data. In such situations, classical multivariate models are not adapted as they ignore the underlying intrinsic nature and structure of the data. Functional Data Analysis (FDA) (Ramsay and Silverman 2005; Ferraty and Vieu 2006) in which the individual data units are assumed to be functions, rather than vectors, offers an adapted framework to deal with continuously observed data, including in regression, classification and clustering. FDA considers the observed data as (discretized) values of smooth functions, rather than multivariate observations represented in the form of “simple” vectors.

The study of functional data has been considered in most of the statistical modeling and inference problems including regression, classification, clustering, functional graphical models (Qiao et al. 2019), among others. In regression, functional linear models have been introduced including penalized functional regression (Élodie Brunel et al. 2016; Goldsmith et al. 2011) and in particular the FLiRTI approach, a functional linear regression constructed upon interpretable regularization (James et al. 2009), and more generally generalized linear models with functional predictors (Müller et al. 2005; James 2002), which cover functional logistic regression for classification. In classification, we can also cite functional linear discriminant analysis (James and Hastie 2001), and, as a penalized model, Lasso-regularized functional logistic regression (Mousavi and Sørensen 2017). To deal with heterogeneous functional data, the construction of mixture models with functional data analytic aspects have been introduced for model-based clustering (Liu and Yang (2009), Jacques and Preda (2014a)) including Lasso-regularized mixtures for functional data (Devijver 2017; James and Sugar 2003; Jacques and Preda 2014b; Chamroukhi and Nguyen 2019). The resulting functional mixture models are better able to handle functional data structures compared to standard multivariate mixtures.

The problem of clustering and prediction in presence of functional observations from heterogeneous populations, leading to complex distributions, is still however less investigated. In this paper, we investigate the framework of mixtures-of-experts (ME) models, as models of choice in modeling heterogeneity in data for prediction and clustering with vectorial observations, and extend it to the functional data framework. The main novelty of our paper is to get

interpretable results for functional ME (FME). First, the ME framework is extended to a functional data setting to learn from functional predictors, and the statistical inference in the resulting setting (ie. of the FME model) looks for estimating a sparse and interpretable FME model parameters. The key technical challenges brought by this setting are addressed by considering sparsity on the derivatives of the underlying functional parameters of the FME model, thanks to Lasso-like regularizations, solved by a dedicated EM-Lasso algorithm. To the best of our knowledge, this is we first time the ME model is constructed upon functional predictors, and provides interpretable sparse estimates.

Firstly, we introduce in Sect. 2 a new family of ME models, referred to as FME, to relate a functional predictor to a scalar response, and develop a dedicated EM algorithm for the maximum-likelihood parameter estimation. Secondly, to deal with potential high-dimensional setting of the introduced FME model, we develop in Sect. 3 a Lasso-regularized approach, which consists of a penalized MLE estimation via an hybrid EM-Lasso algorithm, which integrates an optimized coordinate ascent procedure to efficiently implement the M-Step. Thirdly, we in particular present in Sect. 3.2 and extended FME model, which is constructed upon a sparse and highly-interpretable regularization of the functional expert and gating parameters. The resulting model, abbreviated as iFME, is fitted by regularized MLE via an dedicated EM algorithm. The developed algorithms for the two introduced ME models are applied and evaluated in Sect. 4 on several simulated scenarios and on real data sets, in both clustering and non-linear regression.

2 Functional mixtures-of-experts (FME)

We wish to derive and fit new mixture-of-experts (ME) models in presence of functional predictors, and potentially, functional responses. In this paper, we first consider ME models with a functional predictor $X(\cdot)$ and a real response Y where the pair arises from heterogeneous population composed of unknown K homogeneous sub-populations. To the best of our knowledge, this is the first time ME models are considered for functional data.

2.1 ME with functional predictor and scalar response

Let $\{X_i(t), Y_i\}_{i=1}^n$, be a sample of n independently and identically distributed (i.i.d.) data pairs where Y_i is a real-valued response and $X_i(t)$ is a functional predictor with $t \in \mathcal{T} \subset \mathbb{R}$, for example the time in time series. First, to model the conditional relationships between the continuous response Y and the functional predictor $X(\cdot)$, given an expert z , we formulate each expert component $E_z(y|x)$ in (1) as a functional regression model (cf. Müller et al. (2005), James et al. (2009)).

The resulting functional expert regression model for the i th observation takes the following stochastic representation

$$Y_i = \beta_{z_i,0} + \int_{\mathcal{T}} X_i(t)\beta_{z_i}(t)dt + \varepsilon_i, \quad i \in [n], \tag{2}$$

where $\beta_{z_i,0}$ is an unknown constant intercept, $\beta_{z_i}(t)$, $t \in \mathcal{T}$ is the function of unknown coefficients of functional expert z_i , and $\varepsilon_i \sim \mathcal{N}(0, \sigma_{z_i}^2)$ are independent Gaussian errors, $z_i \in [K]$ being the unknown label of the expert generating the i th observation. The notation $z_i \in [K]$ means $z_i = 1, \dots, K$ which is used throughout this paper. In this context, the response Y is related to the entire trajectory of $X(\cdot)$. Let $\beta = \{\beta_{z,0}, \beta_z(t), t \in \mathcal{T}\}_{z=1}^K$ represents the set of unknown functional parameters for the experts network.

Now consider the modeling of the gating network in the proposed FME model. As in the context of ME for vectorial data, different choices are possible to model the gating network function, typically softmax-gated or Gaussian-gated ME (e.g. see Nguyen and Chamroukhi (2018), Xu et al. (1994), Chamroukhi et al. (2019)). A standard choice as in Jacobs et al. (1991) to model the gating network $G_z(x)$ in (1) is to use the multinomial logistic (softmax) function as a distribution of the latent variable Z . In this functional data modeling context with $K \geq 2$ experts, we use a multinomial logistic function as an extension of the functional logistic regression presented in Mousavi and Sørensen (2018) for linear classification. The resulting functional softmax gating network then takes the following form

$$\begin{aligned} \pi_z(X(t), t \in \mathcal{T}; \alpha) &= \mathbb{P}(Z = z|X(t), t \in \mathcal{T}; \alpha) \\ &= \frac{\exp\{\alpha_{z,0} + \int_{\mathcal{T}} X(t)\alpha_z(t)dt\}}{1 + \sum_{z'=1}^{K-1} \exp\{\alpha_{z',0} + \int_{\mathcal{T}} X(t)\alpha_{z'}(t)dt\}}, \end{aligned} \tag{3}$$

where $\alpha = \{\alpha_{z,0}, \alpha_z(t), t \in \mathcal{T}\}_{z=1}^K$ is the set of unknown constant intercept coefficients $\alpha_{z,0}$ and functional parameters $\alpha_z(t)$, $t \in \mathcal{T}$ for each expert $z \in [K]$. Note that model (3) is equivalent to assuming that each expert z is related to the entire trajectory $X(\cdot)$ via the following functional linear predictor for the gating network

$$\begin{aligned} h_z(X(t), t \in \mathcal{T}; \alpha) &= \log \left\{ \frac{\pi_z(X(t), t \in \mathcal{T}; \alpha)}{\pi_K(X(t), t \in \mathcal{T}; \alpha)} \right\} \\ &= \alpha_{z,0} + \int_{\mathcal{T}} X(t)\alpha_z(t)dt. \end{aligned} \tag{4}$$

The objective is to estimate the functional parameters α and β of the FME model defined by (2)–(3), from an observed sample. In this setting with functional predictors, this requires estimating a possibly infinite number of coefficients (as many as the number of temporal observations for the predictor). In order to reduce the complexity of the problem, the observed functional predictor can be projected onto a fixed number

of basis functions so that we sufficiently capture enough the functional structure of the data, and sufficiently reduce enough the number of coefficients to estimate.

2.2 Smoothing representation of the functional experts

Here we consider the case of fixed design, that is, the covariates $X_i(t)$ are non-random functions. We suppose that the $X_i(\cdot)$'s are measured with error at any given time t . Hence, instead of observing directly $X_i(t)$, one has a noisy version of it $U_i(t)$, defined as

$$U_i(t) = X_i(t) + \delta_i(t), \quad i \in [n],$$

where $\delta_i(\cdot) \sim \mathcal{N}(0, \sigma_\delta^2)$ are measurement errors assumed to be independent of the $X_i(\cdot)$'s and the Y_i 's. Since the functional predictors $X_i(t)$ are not directly observed, we first construct an approximation of $X_i(t)$ from the noisy predictors $U_i(t)$ by projecting the latter onto a set of continuous basis functions. Let $\mathbf{b}_r(t) = [b_1(t), \dots, b_r(t)]^\top$ be a r -dimensional (B-spline, Fourier, Wavelet) basis, then with r sufficiently large $X_i(t)$ can be represented as

$$X_i(t) = \sum_{j=1}^r x_{ij} b_j(t) = \mathbf{x}_i^\top \mathbf{b}_r(t), \tag{5}$$

where $x_{ij} = \int_{\mathcal{T}} X_i(t) b_j(t) dt$ for $j \in [r]$ and $\mathbf{x}_i = (x_{i1}, \dots, x_{ir})^\top$. Since $X_i(t)$ is not observed, the representation coefficients x_{ij} 's are unknown. Hence we propose an unbiased estimator of x_{ij} defined as

$$\hat{x}_{ij} := \int_{\mathcal{T}} U_i(t) b_j(t) dt.$$

Thus, an estimate $\hat{X}_i(t)$ of $X_i(t)$ is given by

$$\hat{X}_i(t) = \hat{\mathbf{x}}_i^\top \mathbf{b}_r(t), \quad i \in [n], \tag{6}$$

with $\hat{\mathbf{x}}_i = (\hat{x}_{i1}, \dots, \hat{x}_{ir})^\top$.

Similarly, to represent the regression coefficient functions $\beta_z(\cdot)$, consider a p -dimensional basis $\mathbf{b}_p(t) = [b_1(t), b_2(t), \dots, b_p(t)]^\top$. Then the function $\beta_z(t)$ can be represented as

$$\beta_z(t) = \boldsymbol{\eta}_z^\top \mathbf{b}_p(t) \tag{7}$$

where $\boldsymbol{\eta}_z = (\eta_{z,1}, \eta_{z,2}, \dots, \eta_{z,p})^\top$ is the vector of unknown coefficients and the choice of p should ensure the tradeoff between smoothness of the functional predictor and complexity of the estimation problem. We select $r \geq p$ to satisfy the identifiability constraint (see for instance Goldsmith et al. (2011), Ramsay and Silverman (2002)). Furthermore, rather

than assuming a perfect fit of $\beta_z(t)$ by $\mathbf{b}_p(t)$ as in (7), we use for each Gaussian expert regressor z , the following error model as proposed by James et al. (2009) for functional linear regression

$$\beta_z(t) = \boldsymbol{\eta}_z^\top \mathbf{b}_p(t) + e(t)$$

where $e(t)$ represents the approximation error of $\beta_z(t)$ by the linear projection (7). As we choose $p \gg n$, $|e(t)|$ can be assumed to be small.

2.3 Smoothing representation of the functional gating network

Since here we are examining functional predictors, an appropriate representation has also to be given for the gating network (3) with functional parameters $\{\alpha_z(t), t \in \mathcal{T}\}_{z=1}^K$. Due to the infinite number of these parameters, we also represent the gating network by a finite set of basis functions similarly as for the experts network. For the representation of the functional predictors $X_i(t)$, $i \in [n]$, we use $\hat{X}_i(t)$ established in (6). The coefficients function $\alpha_z(t)$ is represented similarly as for the β coefficients function of the experts network, by using a q -dimensional basis $\mathbf{b}_q(t) = [b_1(t), b_2(t), \dots, b_q(t)]^\top$, $q \leq r$, via the projection

$$\alpha_z(t) = \boldsymbol{\zeta}_z^\top \mathbf{b}_q(t), \tag{8}$$

where $\boldsymbol{\zeta}_z = (\zeta_{z1}, \zeta_{z2}, \dots, \zeta_{zq})^\top$ is the vector of softmax coefficients function. Note that here we use the same type of basis functions for both representations of β_z and α_z , but one can use different types of bases if needed. Then, by using the representations (6) and (8) of $X(t)$ and $\alpha_z(t)$, respectively, in the linear predictor $h_z(\cdot)$ defined in (4) for $i \in [n]$, the latter is thus approximated as

$$h_z(U_i(t), t \in \mathcal{T}; \boldsymbol{\alpha}) = \alpha_{z_i,0} + \boldsymbol{\zeta}_{z_i}^\top \mathbf{r}_i, \tag{9}$$

where $\mathbf{r}_i = [\int_{\mathcal{T}} \mathbf{b}_r(t) \mathbf{b}_q(t)^\top dt]^\top \hat{\mathbf{x}}_i$. Thus, following its definition in (3), the functional softmax gating network is approximated as

$$\pi_k(\mathbf{r}_i; \boldsymbol{\xi}) = \frac{\exp\{\alpha_{k,0} + \boldsymbol{\zeta}_k^\top \mathbf{r}_i\}}{1 + \sum_{k'=1}^{K-1} \exp\{\alpha_{k',0} + \boldsymbol{\zeta}_{k'}^\top \mathbf{r}_i\}}, \tag{10}$$

where $\boldsymbol{\xi} = ((\alpha_{1,0}, \boldsymbol{\zeta}_1^\top), \dots, (\alpha_{K-1,0}, \boldsymbol{\zeta}_{K-1}^\top))^\top \in \mathbb{R}^{(q+1)(K-1)}$ is the unknown parameter vector of the functional gating network to be estimated.

2.4 The FME model conditional density

We now have appropriate representations for the functional predictors, as well as for both the functional gating network and the functional experts network, involved in the construction of the functional ME (FME) model (2)-(3). Gathering (6) and (7), the stochastic representation (2) of the FME model can thus be defined as follows,

$$Y_i|u_i(\cdot) = \beta_{z_i,0} + \eta_{z_i}^\top \mathbf{x}_i + \varepsilon_i^*, \quad i \in [n], \tag{11}$$

where $\mathbf{x}_i = \left[\int_{\mathcal{T}} \mathbf{b}_r(t) \mathbf{b}_p(t)^\top dt \right]^\top \widehat{\mathbf{x}}_i$ and $\varepsilon_i^* = \varepsilon_i + \widehat{\mathbf{x}}_i^\top \int_{\mathcal{T}} \mathbf{b}_r(t) e(t) dt$. We can see that the problem is now reduced at the standard version of the ME model. From this stochastic representation under the Gaussian assumption for the error variable ε_i , the conditional density of each approximated functional expert $z_i = k$ is thus given by

$$f(y_i|u_i(\cdot), z_i = k; \theta_k) = \phi(y_i; \beta_{k,0} + \eta_k^\top \mathbf{x}_i, \sigma_k^2), \tag{12}$$

where $\phi(\cdot; \mu, v)$ is the Gaussian probability density function with mean μ and variance v , $\beta_{k,0} + \eta_k^\top \mathbf{x}_i$ is the mean of the approximated functional regression expert, σ_k^2 its variance, and $\theta_k = (\beta_{k,0}, \eta_k^\top, \sigma_k^2)^\top \in \mathbb{R}^{p+2}$ the unknown parameter vector of expert density k , $k \in [K]$ to be estimated. Finally, combining (12) and (10) in the ME model (1), leads to the the following conditional density defining the FME model,

$$f(y_i|u_i(\cdot); \Psi) = \sum_{k=1}^K \pi_k(\mathbf{r}_i; \xi) \phi(y_i; \beta_{k,0} + \eta_k^\top \mathbf{x}_i, \sigma_k^2), \tag{13}$$

where $\Psi = (\xi^\top, \theta_1^\top, \dots, \theta_K^\top)^\top$ is the parameter vector of the model to be estimated.

2.5 Maximum likelihood estimation via the EM algorithm

The FME model (13) is now defined upon an adapted finite representation of the functional predictors, and its parameter estimation can then be performed given an observed data sample. We first consider the maximum likelihood estimation framework via the EM algorithm (Dempster et al. 1977; Jacobs et al. 1991) which has many desirable properties including stability and convergence guarantees (e.g. see McLachlan and Krishnan (2008) for more details). Note that here we use the term maximum likelihood estimation to not unduly clutter the clarity of the text, while as it will be specified later, we refer to the conditional maximum likelihood estimator.

In practice, the data are available in the form of discretized values of functions. The noisy functional predictors $U_i(t)$ are usually observed at discrete sampling points $t_{i1} < \dots < t_{im_i}$

with $t_{ij} \in \mathcal{T}$ for $j \in [m_i]$. We suppose that $U_i(t)$ is scaled such that $0 \leq t \leq 1$ and divide the time period $[0, 1]$ up into a fine grid of m_i points t_{i1}, \dots, t_{im_i} . Thus, in (11) we have $\mathbf{x}_i = \left[\sum_{j=1}^{m_i} \mathbf{b}_r(t_{ij}) \mathbf{b}_p(t_{ij})^\top \right]^\top \widehat{\mathbf{x}}_i$, $\mathbf{r}_i = \left[\sum_{j=1}^{m_i} \mathbf{b}_r(t_{ij}) \mathbf{b}_q(t_{ij})^\top \right]^\top \widehat{\mathbf{x}}_i$, where $\widehat{\mathbf{x}}_{ij} = \sum_{j=1}^{m_i} U_i(t_{ij}) \mathbf{b}_j(t_{ij})$. Note that if we choose $p = q = r$, then $\mathbf{x}_i = \mathbf{r}_i = \widehat{\mathbf{x}}_i$. Let $\mathcal{D} = \{(\mathbf{u}_1, y_1), \dots, (\mathbf{u}_n, y_n)\}$ be an i.i.d. sample of n observed data pairs where $\mathbf{u}_i = (u_{i,1}, \dots, u_{i,m_i})$ is the observed functional predictor for the i th response y_i .

We use \mathcal{D} to estimate the parameter vector Ψ by iteratively maximizing the observed data log-likelihood,

$$\log L(\Psi) = \sum_{i=1}^n \log \sum_{k=1}^K \pi_k(\mathbf{r}_i; \xi) \phi(y_i; \beta_{k,0} + \eta_k^\top \mathbf{x}_i, \sigma_k^2), \tag{14}$$

via the EM algorithm. As detailed in ‘‘Appendix’’, the EM algorithm for the FME model is implemented as follows. After starting with an initial solution $\Psi^{(0)}$, it alternates, at each iteration s , between the two following steps, until convergence (when there is no longer a significant change in the values of the log-likelihood (14)).

E-step Calculate the following conditional probability memberships $\tau_{ik}^{(s)}$ (for all $i \in [n]$), that the observed pair (u_i, y_i) originates from the k th expert, given the observed data and the current parameter estimate $\Psi^{(s)}$,

$$\begin{aligned} \tau_{ik}^{(s)} &= \mathbb{P}(Z_i = k | y_i, u_i(\cdot); \Psi^{(s)}) \\ &= \frac{\pi_k(\mathbf{r}_i; \xi^{(s)}) \phi(y_i; \beta_{k,0}^{(s)} + \eta_k^\top \mathbf{x}_i, \sigma_k^{2(s)})}{f(y_i|u_i(\cdot); \Psi^{(s)})}. \end{aligned} \tag{15}$$

M-step Update the value of the parameter vector Ψ by maximizing the Q -function (43) with respect to Ψ . The maximization is performed by separate maximizations with respect to (w.r.t.) the gating network parameters ξ and, for each expert k , w.r.t. the expert network parameters θ_k , for each of the K experts.

Updating the gating network’s parameters ξ consists of maximizing w.r.t. ξ the part of (43) that is a function of ξ . Since we use a softmax-gated expert network in (10), this maximization problem consists of a weighted multinomial logistic problem for which there is no a closed-form solution. We then use a Newton–Raphson (NR) procedure, which iteratively maximizes (44) after starting from an initial parameter vector $\xi^{(0)}$, by updating, at each NR iteration t , the values of the parameter vector ξ according to the following updating formula:

$$\xi^{(t+1)} = \xi^{(t)} - \left[H(\xi; \Psi^{(s)}) \right]_{\xi=\xi^{(t)}}^{-1} g(\xi; \Psi^{(s)}) \Big|_{\xi=\xi^{(t)}} \tag{16}$$

where $H(\xi; \Psi^{(s)})$ and $g(\xi; \Psi^{(s)})$ are, respectively, the Hessian matrix and the gradient vector of $Q(\xi; \Psi^{(s)})$, and are

provided in ‘‘Appendix A’’. At each NR iteration, the Hessian matrix and gradient vector are evaluated at the current value of ξ . We keep updating the gating network parameter ξ according to (16) until there is no significant change in $Q(\xi; \Psi)$. The maximization then provides $\xi^{(s+1)}$ for the next EM iteration. Alternatively, instead of using the NR procedure, one can employ the Minorize-Maximization (MM) algorithm to update the parameter ξ of the gating network. In this approach, the Hessian matrix in (16) is approximated by a square matrix that is independent of ξ . This can offer an improved computational efficiency, although it comes at the cost of requiring a greater number of iterations. For further details, refer to Gormley and Murphy (2008).

Updating the experts network parameters θ_k consists of solving K independent weighted regression problems where the weights are the conditional expert memberships $\tau_{ik}^{(s)}$ given by (15). The updating formulas for the regression parameters $(\beta_{k,0}, \eta_k)$ and the noise variances σ_k^2 for each expert k are straightforward and correspond to weighted versions of those of standard Gaussian linear regression, i.e., weighted ordinary least squares. The updating rules for the experts network parameters are given by the following formulas:

$$\begin{aligned} \beta_{k,0}^{(s+1)} &= \frac{1}{n_k^{(s)}} \sum_{i=1}^n \tau_{ik}^{(s)} (y_i - \mathbf{x}_i^\top \eta_k^{(s)}), \\ \eta_k^{(s+1)} &= \frac{1}{\sum_{i=1}^n \tau_{ik}^{(s)} \mathbf{x}_i \mathbf{x}_i^\top} \sum_{i=1}^n \tau_{ik}^{(s)} (y_i - \beta_{k,0}^{(s+1)}) \mathbf{x}_i, \\ \sigma_k^2{}^{(s+1)} &= \frac{1}{n_k^{(s)}} \sum_{i=1}^n \tau_{ik}^{(s)} \left[y_i - (\beta_{k,0}^{(s+1)} + \mathbf{x}_i^\top \eta_k^{(s+1)}) \right]^2, \end{aligned} \tag{17}$$

where $n_k^{(s)} = \sum_{i=1}^n \tau_{ik}^{(s)}$ represents the expected cardinal number of component k .

This EM algorithm, designed here for the FME that is constructed upon smoothing of the functional data, can be seen as a direct extension of the vectorized version the ME model. While it can hence be expected to provide accurate estimations as in the vector predictors setting, the number of parameters to estimate here in the case of the FME can still be high, for example when a big number of basis functions is used to have more accurate approximation of the functional predictors. In that case, it is better to regularize the maximum likelihood estimator in order to establish a compromise between the quality of fit and complexity.

3 Regularized maximum likelihood estimation

We rely on the LASSO (Tibshirani 1996) as a successful procedure to encourage sparse models in high-dimensional linear regression based on an ℓ_1 -penalty, and include it in this

mixture of experts modeling framework for functional data. The ℓ_1 -regularized ME models have demonstrated their performance from a computational point of view (Chamroukhi and Huynh 2019; Huynh and Chamroukhi 2019) and enjoy good theoretical properties of universal approximation and consistent model selection in the high-dimensional setting (Nguyen et al. 2020, 2021b).

3.1 ℓ_1 -regularization and the EM-Lasso algorithm

We propose an ℓ_1 -regularization of the observed-data log-likelihood (18) to be maximized, along with coordinate ascent algorithms to deal with the high-dimensional setting when updating the parameters within the resulting EM-Lasso algorithm. The objective function in this case is given by the following ℓ_1 -regularized observed-data log-likelihood,

$$\mathcal{L}(\Psi) = \log L(\Psi) - \text{Pen}_{\lambda,\chi}(\Psi), \tag{18}$$

where $\log L(\Psi)$ is the observed-data log-likelihood of Ψ defined by (14), and $\text{Pen}_{\lambda,\chi}(\Psi)$ is a LASSO regularization term encouraging sparsity for the expert and the gating network parameters, defined by

$$\text{Pen}_{\lambda,\chi}(\Psi) = \lambda \sum_{k=1}^K \sum_{j=1}^p |\eta_{k,j}| + \chi \sum_{k=1}^{K-1} \sum_{j=1}^q |\zeta_{k,j}|, \tag{19}$$

where λ and χ are positive real values representing tuning parameters. The maximization of (18) cannot be performed in a closed form but again the EM algorithm can be adapted to iteratively solve the maximization problem. The resulting algorithm for the FME model, called EM-Lasso, takes the following form, as detailed in ‘‘Appendix B’’. After starting with an initial solution $\Psi^{(0)}$, it alternates between the two following steps, until convergence, i.e., when there is no longer a significant change in the values of the ℓ_1 -penalized log-likelihood (18).

E-step The E-Step in this EM-Lasso algorithm is unchanged compared to the previously presented EM algorithm, and only requires the computation of the conditional expert memberships $\tau_{ik}^{(s)}$ according to (15).

M-step In this regularized MLE context, the parameter vector Ψ is now updated by maximizing the regularized Q -function (46), i.e., $\Psi^{(s+1)} = \arg \max_{\Psi} \{Q(\Psi; \Psi^{(s)}) - \text{Pen}_{\lambda,\chi}(\Psi)\}$. This is performed by separate maximizations w.r.t. the gating network parameters ξ and, for each expert k , w.r.t. the expert network parameters $\theta_k, k \in [K]$.

Updating the gating network parameters at iteration s of the EM-Lasso algorithm consists of maximizing the following regularized Q -function w.r.t. ξ ,

$$Q_{\chi}(\xi; \Psi^{(s)}) = Q(\xi; \Psi^{(s)}) - \chi \sum_{k=1}^{K-1} \|\xi_k\|_1, \tag{20}$$

where $Q(\xi; \Psi^{(s)}) = \sum_{i=1}^n \left[\sum_{k=1}^{K-1} \tau_{ik}^{(s)} (\alpha_{k,0} + \xi_k^{\top} \mathbf{r}_i) - \log \left(1 + \sum_{k'=1}^{K-1} \exp\{\alpha_{k',0} + \xi_{k'}^{\top} \mathbf{r}_i\} \right) \right]$. One can see this is equivalent to solving a weighted regularized multinomial logistic regression problem for which $Q_{\chi}(\xi; \Psi^{(s)})$ is its penalized log-likelihood, the weights being the conditional probabilities $\tau_{ik}^{(s)}$. There is no closed-form solution for this kind of problem. We then use an iterative optimization algorithm to seek for a maximizer of $Q_{\chi}(\xi; \Psi^{(s)})$, i.e., an update for the parameters of the gating network. To be effective when the number of parameters to estimate is high, we propose a coordinate ascent algorithm to update the softmax gating network coefficients in this regularized context.

Coordinate ascent for updating the gating network. The idea of the coordinate ascent algorithm (e.g. see Hastie et al. (2015), Huynh and Chamroukhi (2019)), implemented in our context at the s th EM-Lasso iteration to maximize $Q_{\chi}(\xi; \Psi^{(s)})$ at the M-Step, is as follows. The gating function parameter vectors $\xi_k = (\alpha_{k,0}, \xi_k^{\top})^{\top}$ as components of the whole gating network parameters $\xi = (\xi_1^{\top}, \dots, \xi_{K-1}^{\top})^{\top}$, are updated one at a time, while fixing the other gate's parameters to their previous estimates. Furthermore, to update a single gating parameter vector ξ_k , we only update its coefficients ξ_{kj} one at a time, while fixing the other coefficients to their previous values. More precisely, for each single gating function k , we partially approximate the smooth part of $Q_{\chi}(\xi; \Psi^{(s)})$ with respect to ξ_k at the current EM-Lasso estimate, say $\xi^{(t)}$, then optimize the resulting objective function (with respect to ξ_k). This corresponds to solving penalized weighted least squares problems using coordinate ascent. Thus, this results into an inner loop, indexed by m , that cycles over the components of ξ_k and updates them one by one, while the others are kept fixed to their previous values, i.e., $\xi_{kh}^{(m+1)} = \xi_{kh}^{(m)}$ for all $h \neq j$, until the objective function (48) is not significantly increased.

The obtained closed form updates for each coefficient ξ_{kj} , $j \in [q]$, and for the intercept $\alpha_{k,0}$, are as follows

$$\xi_{kj}^{(m+1)} = \frac{\mathcal{S}_{\chi} \left(\sum_{i=1}^n w_{ik} r_{ij} (c_{ik} - \tilde{c}_{ikj}^{(m)}) \right)}{\sum_{i=1}^n w_{ik} r_{ij}^2} \text{ for } j \in [q],$$

$$\alpha_{k,0}^{(m+1)} = \frac{\sum_{i=1}^n w_{ik} (c_{ik} - \mathbf{r}_i^{\top} \xi_k^{(m+1)})}{\sum_{i=1}^n w_{ik}},$$

where $\tilde{c}_{ikj}^{(m)} = \alpha_{k,0}^{(m)} + \mathbf{r}_i^{\top} \xi_k^{(m)} - \xi_{kj}^{(m)} r_{ij}$ is the fitted value excluding the contribution from the j th component of the i th vector \mathbf{r}_{ij} in the design matrix of the gating network and $\mathcal{S}_{\chi}(\cdot)$ is a soft-thresholding operator defined by $\mathcal{S}_{\chi}(u) = \text{sign}(u)(|u| - \chi)_+$ with $(v)_+$ is a shorthand for $\max\{v, 0\}$.

The values $(\alpha_{k,0}^{(m+1)}, \xi_k^{(m+1)})$ obtained at convergence of the coordinate ascent inner loop for the k th gating function are taken as the fixed values of that gating function, in the procedure of updating the next parameter vector ξ_{k+1} . Finally, when all the gating functions have their values updated, i.e., after $K - 1$ inner coordinate ascent loops, to avoid numerical instability, we perform a backtracking line search, before actually updating the gating network's parameters for the next EM-Lasso iteration. More precisely, the update is $\xi^{(t+1)} = (1 - \nu)\xi^{(t)} + \nu\bar{\xi}^{(t)}$, where $\bar{\xi}^{(t)}$ is the output after $K - 1$ inner loops and ν is backtrackingly determined to ensure $Q_{\chi}(\xi^{(t+1)}; \Psi^{(s)}) \geq Q_{\chi}(\xi^{(t)}; \Psi^{(s)})$.

We keep cycling these updated iterates for the parameter vectors ξ_k , until convergence of the whole coordinate ascent (CA) procedure inside the M-Step, i.e., when the relative increase in the Lasso-regularized objective $Q_{\chi}(\xi; \Psi^{(s)})$ related to the softmax gating network is not significant, e.g., less than a small tolerance. Then, the next EM-Lasso iteration is calculated with the updated gating network's parameters $\xi^{(s+1)} = (\tilde{\alpha}_{1,0}, \tilde{\xi}_1^{\top}, \dots, \tilde{\alpha}_{K-1,0}, \tilde{\xi}_{K-1}^{\top})^{\top}$ where the values $\tilde{\alpha}_{k,0}$ and $\tilde{\xi}_{kj}$ for all $k \in [K-1]$, $j \in [q]$ are those obtained for the $\alpha_{k,0}$'s and the ξ_{kj} 's at convergence of the CA algorithm. *Updating the experts network parameters* The maximization step for updating the expert parameters θ_k consists of maximizing the function $Q_{\lambda}(\theta_k; \Psi^{(s)})$ given by

$$Q_{\lambda}(\theta_k; \Psi^{(s)}) = Q(\theta_k; \Psi^{(s)}) - \lambda \|\eta_k\|_1, \tag{21}$$

where $Q(\theta_k; \Psi^{(s)}) = -\frac{1}{2\sigma_k^2} \sum_{i=1}^n \tau_{ik}^{(s)} (y_i - (\beta_{k,0} + \eta_k^{\top} \mathbf{x}_i))^2 - \frac{n}{2} \log(2\pi\sigma_k^2)$. This corresponds to solving a weighted LASSO problem where the weights are the conditional experts memberships $\tau_{ik}^{(s)}$. We then solve it by an iterative optimization algorithm similarly to the previous case of updating the gating network parameters. As it can be seen in "Appendix B.2", updating $(\beta_{k,0}, \eta_k)$ according to (21) is obtained by coordinate ascent as follows. For each $j \in [p]$, the closed-form update for η_{kj} at the m th iteration of the coordinate ascent algorithm within the M-Step of EM-Lasso, is given by

$$\eta_{kj}^{(m+1)} = \frac{\mathcal{S}_{\lambda\sigma_k^2(s)} \left(\sum_{i=1}^n \tau_{ik}^{(s)} x_{ij} (y_i - \tilde{y}_{ij}^{(m)}) \right)}{\sum_{i=1}^n \tau_{ik}^{(s)} x_{ij}^2},$$

$$\beta_{k,0}^{(m+1)} = \frac{\sum_{i=1}^n \tau_{ik}^{(s)} (y_i - \mathbf{x}_i^{\top} \eta_k^{(m+1)})}{\sum_{i=1}^n \tau_{ik}^{(s)}},$$

in which $\tilde{y}_{ij}^{(m)} = \beta_{k,0}^{(m)} + \mathbf{x}_i^{\top} \eta_k^{(m)} - \eta_{kj}^{(m)} x_{ij}$ is the fitted value excluding the contribution from x_{ij} and $\mathcal{S}_{\lambda\sigma_k^2(s)}(\cdot)$ is the soft-thresholding operator. We keep updating the components of $(\beta_{k,0}, \eta_k)$ cyclically until no enough increase in objective function (21). Then, once $(\beta_{k,0}, \eta_k)$ are updated while fixing the variance σ_k^2 , the latter is then updated straightforwardly

as in the case of standard weighted Gaussian regression, and its update is given by

$$\sigma_k^{2(s+1)} = \frac{\sum_{i=1}^n \tau_{ik}^{(s)} \left(y_i - \beta_{k,0}^{(s+1)} - \mathbf{x}_i^\top \boldsymbol{\eta}_k^{(s+1)} \right)^2}{\sum_{i=1}^n \tau_{ik}^{(s)}}$$

where $(\beta_{k,0}^{(s+1)}, \boldsymbol{\eta}_k^{(s+1)}) = (\tilde{\beta}_{k,0}, \tilde{\boldsymbol{\eta}}_k)$ is the solution obtained at convergence of the CA algorithm, which is taken as the update in the next EM-Lasso iteration.

This completes the parameter vector update $\boldsymbol{\Psi}^{(s+1)} = (\boldsymbol{\xi}^{(s+1)}, \boldsymbol{\theta}_1^{(s+1)}, \dots, \boldsymbol{\theta}_K^{(s+1)})$ of the regularized FME model, where $\boldsymbol{\xi}^{(s+1)}$ and $\boldsymbol{\theta}_k^{(s+1)}, k \in [K]$, are, respectively, the updates of the gating network parameters and the experts network parameters, calculated by the coordinate ascent algorithms.

The EM-Lasso algorithm provides an estimate of the FME parameters with sparsity constraints on the values of the parameter vectors $\boldsymbol{\xi}$ and $\boldsymbol{\theta}_k, k \in [K]$. Actually, since here these parameter vectors do not relate directly the original functional inputs, to the output, assuming some of their values is zero is not necessarily justified, as there is no indeed any reason that these values are zero, nor easily interpretable, compared to the sparsity in vectorial generalized linear models, mixture of regressions and ME models.

From now on, we refer to FME and FME-Lasso, respectively, the FME model fitted by EM algorithm in Sect. 2.5 and the regularized FME model fitted by EM-Lasso algorithm, in Sect. 3.1. In the following Section, we introduce a regularization that is interpretable and encourages sparsity in the FME model.

3.2 Interpretable Functional Mixture of Experts (iFME)

In this section, we provide a sparse and, especially highly-interpretable fit, for the coefficient functions $\{\beta_k(t), t \in \mathcal{T}\}$ and $\{\alpha_k(t), t \in \mathcal{T}\}$ representing each of the K functional experts and gating network. We call our approach Interpretable Functional Mixture of Experts (iFME). The presented iFME allows us to control the expert and gating parameter functions while still providing performance as with the standard FME model presented previously.

3.2.1 Interpretable sparse regularization

We rely on the methodology of Functional Linear Regression That’s Interpretable (FLiRTI) presented in James et al. (2009). The idea of the FLiRTI methodology is as follows. We use variable selection with sparsity assumption on appropriate chosen derivatives of the coefficient function, say $\beta_{z_i}(t)$ here, in the case of the functional expert network, to produce

a highly interpretable estimate for the coefficient functions $\beta_{z_i}(t)$. For instance, $\beta_{z_i}^{(0)}(t) = 0$ implies that the predictor $X_i(t)$ has no effect on the response Y_i at t , $\beta_{z_i}^{(1)}(t) = 0$ means that $\beta_{z_i}(t)$ is constant in t , $\beta_{z_i}^{(2)}(t) = 0$ shows that $\beta_{z_i}(t)$ is linear in t , etc. Assuming sparsity in higher-order derivatives of $\beta_{z_i}(t)$, for instance when $d = 3$ or $d = 4$, will however give us a less easily interpretable fit. Hence, for example, if one believes that the expert parameter function $\beta_{z_i}(t)$ is exactly zero over a certain region and exactly linear over other region of t , then it is necessary to estimate $\beta_{z_i}(t)$ such that $\beta_{z_i}^{(0)}(t) = 0$ and $\beta_{z_i}^{(2)}(t) = 0$ over those regions, respectively. In this situation, we need to model $\beta_{z_i}(t)$ assuming that its zeroth and second derivatives are sparse. However, with the EM-Lasso method derived above via the Lasso regularization, there is no actually any reason that we could obtain those desired properties, which may result in an estimate for $\beta_{z_i}(t)$ that is rarely exactly zeros (and/or linear), and making the sparsity and coefficient curves hard to interpret. The same situation may occur with the gating parameter functions. To handle this, we describe in what follows the construction of our iFME model that produces flexible-shape and highly-interpretable estimates for the expert and gating coefficient functions, by simultaneously constraining any two of their derivatives to be sparse.

We start by selecting a p -dimensional basis $\mathbf{b}_p(t)$ and a q -dimensional basis $\mathbf{b}_q(t)$ for approximating the experts and gating networks, respectively. For the expert network, if we divide the time domain into a grid of p evenly spaced points, and let D^d be the d th finite difference operator defined recursively as

$$\begin{aligned} D\mathbf{b}_p(t_j) &= p [\mathbf{b}_p(t_j) - \mathbf{b}_p(t_{j-1})], \\ D^2\mathbf{b}_p(t_j) &= D [D\mathbf{b}_p(t_j)] \\ &= p^2 [\mathbf{b}_p(t_j) - 2\mathbf{b}_p(t_{j-1}) + \mathbf{b}_p(t_{j-2})], \\ &\vdots \\ D^d\mathbf{b}_p(t_j) &= D [D^{d-1}\mathbf{b}_p(t_j)], \end{aligned}$$

then $D^d\mathbf{b}_p(t_j)$ provides an approximation for $\mathbf{b}_p^{(d)}(t_j) = [b_1^{(d)}(t_j), \dots, b_p^{(d)}(t_j)]^\top, j \in [p]$. Let

$$\mathbf{A}_p = \left[\underbrace{D^{d_1}\mathbf{b}_p(t_1), D^{d_1}\mathbf{b}_p(t_2), \dots, D^{d_1}\mathbf{b}_p(t_p)}_{\mathbf{A}_p^{[d_1]}}, \underbrace{D^{d_2}\mathbf{b}_p(t_1), D^{d_2}\mathbf{b}_p(t_2), \dots, D^{d_2}\mathbf{b}_p(t_p)}_{\mathbf{A}_p^{[d_2]}} \right]^\top$$

be the matrix of approximate d_1 th and d_2 th derivative of the basis $\mathbf{b}_p(t)$. We denote $\mathbf{A}_p^{[d_1]}$ the first p rows of \mathbf{A}_p and $\mathbf{A}_p^{[d_2]}$

the remainder, i.e., $\mathbf{A}_p = [\mathbf{A}_p^{[d_1]}, \mathbf{A}_p^{[d_2]}]^\top$. One can see such matrix \mathbf{A}_p is in $\mathbb{R}^{2p \times p}$ and $\mathbf{A}_p^{[d_1]}$ is a square invertible matrix.

Now, if we consider the following representation for the expert network coefficient function

$$\boldsymbol{\gamma}_{z_i} = \mathbf{A}_p \boldsymbol{\eta}_{z_i} \tag{22}$$

with $\boldsymbol{\gamma}_{z_i} = (\boldsymbol{\gamma}_{z_i}^{[d_1]\top}, \boldsymbol{\gamma}_{z_i}^{[d_2]\top})^\top$, where $\boldsymbol{\gamma}_{z_i}^{[d_1]} = (\boldsymbol{\gamma}_{1,z_i}^{[d_1]}, \dots, \boldsymbol{\gamma}_{p,z_i}^{[d_1]})^\top$, $\boldsymbol{\gamma}_{z_i}^{[d_2]} = (\boldsymbol{\gamma}_{1,z_i}^{[d_2]}, \dots, \boldsymbol{\gamma}_{p,z_i}^{[d_2]})^\top$, and $\boldsymbol{\eta}_{z_i}$ defined as in relation to $\beta_{z_i}(t)$ as in (7), then one can see that $\boldsymbol{\gamma}_{z_i}^{[d_1]}$ provides an approximation to $\beta_{z_i}^{(d_1)}(t)$, i.e. the d_1 th derivative of $\beta_{z_i}(t)$. Respectively, $\boldsymbol{\gamma}_{z_i}^{[d_2]}$ provides an approximation to $\beta_{z_i}^{(d_2)}(t)$, the d_2 th derivative of $\beta_{z_i}(t)$. Therefore, enforcing sparsity in $\boldsymbol{\gamma}_{z_i}$ will constrain $\beta_{z_i}^{(d_1)}(t)$ and $\beta_{z_i}^{(d_2)}(t)$ to be zero at most time points.

Similarly, let $\mathbf{A}_q = [\mathbf{A}_q^{[d_1]}, \mathbf{A}_q^{[d_2]}]^\top \in \mathbb{R}^{2q \times q}$ be the corresponding matrix defined for the gating network. If we consider the following representation for the gating network coefficient function

$$\boldsymbol{\omega}_{z_i} = \mathbf{A}_q \boldsymbol{\zeta}_{z_i} \tag{23}$$

with $\boldsymbol{\omega}_{z_i} = (\boldsymbol{\omega}_{z_i}^{[d_1]\top}, \boldsymbol{\omega}_{z_i}^{[d_2]\top})^\top$, where $\boldsymbol{\omega}_{z_i}^{[d_1]} = (\omega_{1,z_i}^{[d_1]}, \dots, \omega_{q,z_i}^{[d_1]})^\top$, $\boldsymbol{\omega}_{z_i}^{[d_2]} = (\omega_{1,z_i}^{[d_2]}, \dots, \omega_{q,z_i}^{[d_2]})^\top$, and $\boldsymbol{\zeta}_{z_i}$ defined as in relation to $\alpha_{z_i}(t)$ as in (8), then we can derive the same interpretation as above for the coefficient vector $\boldsymbol{\omega}_{z_i}$ and the gating parameter function $\alpha_{z_i}(t)$.

Using simple calculations, the representations (22) and (23) imply the following relations that subsequently used in the iFME model:

$$\boldsymbol{\eta}_{z_i} = (\mathbf{A}_p^{[d_1]})^{-1} \boldsymbol{\gamma}_{z_i}^{[d_1]}, \quad \boldsymbol{\gamma}_{z_i}^{[d_2]} = \mathbf{A}_p^{[d_2]} (\mathbf{A}_p^{[d_1]})^{-1} \boldsymbol{\gamma}_{z_i}^{[d_1]}, \tag{24}$$

and

$$\boldsymbol{\zeta}_{z_i} = (\mathbf{A}_q^{[d_1]})^{-1} \boldsymbol{\omega}_{z_i}^{[d_1]}, \quad \boldsymbol{\omega}_{z_i}^{[d_2]} = \mathbf{A}_q^{[d_2]} (\mathbf{A}_q^{[d_1]})^{-1} \boldsymbol{\omega}_{z_i}^{[d_1]}, \tag{25}$$

respectively.

In fact, one can construct \mathbf{A}_p and \mathbf{A}_q with only one derivative. Then the constraints involved to the d_2 th derivative will be eliminated making the estimation easier, but also limiting the flexibility in the shapes of the functions. That is why in this construction and in our experimental studies, \mathbf{A}_p and \mathbf{A}_q are constructed with multiple derivatives in order to produce curves of $\beta_{z_i}(\cdot)$ and $\alpha_{z_i}(\cdot)$ with such many desired properties.

3.2.2 The iFME model

Combining the stochastic representation of the FME model in (11) for the experts model, the linear predictor definition

in (9), and the relations (24)-(25), we obtain the following iFME model construction

$$Y_i | u_i(\cdot) = \beta_{z_i,0} + \boldsymbol{\gamma}_{z_i}^{[d_1]\top} \mathbf{v}_i + \varepsilon_i^*, \tag{26}$$

$$h_{z_i}(u_i(\cdot); \boldsymbol{\alpha}) = \alpha_{z_i,0} + \boldsymbol{\omega}_{z_i}^{[d_1]\top} \mathbf{s}_i, \tag{27}$$

subject to

$$\boldsymbol{\gamma}_{z_i}^{[d_2]} = \mathbf{A}_p^{[d_2]} \mathbf{A}_p^{[d_1]-1} \boldsymbol{\gamma}_{z_i}^{[d_1]} \quad \text{and} \quad \boldsymbol{\omega}_{z_i}^{[d_2]} = \mathbf{A}_q^{[d_2]} \mathbf{A}_q^{[d_1]-1} \boldsymbol{\omega}_{z_i}^{[d_1]}, \tag{28}$$

where $\mathbf{v}_i = (\mathbf{A}_p^{[d_1]-1})^\top \mathbf{x}_i$ is the new design vector for the experts and $\mathbf{s}_i = (\mathbf{A}_q^{[d_1]-1})^\top \mathbf{r}_i$ the new one for the gating network. Hence, from (26) and (27), the conditional density of each expert and the gating network are now written as

$$f(y_i | u_i(\cdot); \boldsymbol{\psi}_k) = \phi(y_i; \beta_{k,0} + \boldsymbol{\gamma}_k^{[d_1]\top} \mathbf{v}_i, \sigma_k^2) \tag{29}$$

and

$$\pi_k(\mathbf{s}_i; \mathbf{w}) = \frac{\exp\{\alpha_{k,0} + \boldsymbol{\omega}_k^{[d_1]\top} \mathbf{s}_i\}}{1 + \sum_{k'=1}^{K-1} \exp\{\alpha_{k',0} + \boldsymbol{\omega}_{k'}^{[d_1]\top} \mathbf{s}_i\}}, \tag{30}$$

where $\boldsymbol{\psi}_k = (\beta_{k,0}, \boldsymbol{\gamma}_k^{[d_1]\top}, \sigma_k^2)^\top$ is the unknown parameter vector of expert component density k and $\mathbf{w} = (\alpha_{1,0}, \boldsymbol{\omega}_1^{[d_1]\top}, \dots, \alpha_{K-1,0}, \boldsymbol{\omega}_{K-1}^{[d_1]\top})^\top$, with $(\alpha_{K,0}, \boldsymbol{\omega}_K^{[d_1]\top})^\top$ a null vector, is the unknown parameter vector of the gating network. Finally, gathering (29) and (30) as for (13), the iFME model density is now given by

$$f(y_i | u_i(\cdot); \boldsymbol{\Psi}) = \sum_{k=1}^K \pi_k(\mathbf{s}_i; \mathbf{w}) \phi(y_i; \beta_{k,0} + \boldsymbol{\gamma}_k^{[d_1]\top} \mathbf{v}_i, \sigma_k^2), \tag{31}$$

where $\boldsymbol{\Psi} = (\mathbf{w}^\top, \boldsymbol{\psi}_1^\top, \dots, \boldsymbol{\psi}_K^\top)^\top$ is the parameter vector of the model to be estimated. Thus, the iFME model constructed upon the parameter vectors $\boldsymbol{\gamma}_k$'s and $\boldsymbol{\omega}_k$'s, for which the sparsity is assumed to obtain interpretable estimates.

3.3 Regularized MLE via the EM-iFME algorithm

In order to fit the iFME model and to maintain the sparsity in $\boldsymbol{\gamma}_k$ and $\boldsymbol{\omega}_k$, the following EM-iFME algorithm is then developed to maximize the penalized log-likelihood function

$$\mathcal{L}(\boldsymbol{\Psi}) = \sum_{i=1}^n \log f(y_i | u_i(\cdot); \boldsymbol{\Psi}) + \text{Pen}_{\lambda, \chi}(\boldsymbol{\Psi}) \tag{32}$$

with the conditional iFME density $f(y_i|u_i(\cdot); \Psi)$ is defined in (31) and the new sparse and interpretable regularization term is given by

$$\text{Pen}_{\lambda, \chi}(\Psi) = \lambda \sum_{k=1}^K \left(\|\mathbf{y}_k^{[d_1]}\|_1 + \rho \|\mathbf{y}_k^{[d_2]}\|_1 \right) + \chi \sum_{k=1}^{K-1} \left(\|\boldsymbol{\omega}_k^{[d_1]}\|_1 + \varrho \|\boldsymbol{\omega}_k^{[d_2]}\|_1 \right), \tag{33}$$

where ρ and ϱ are, respectively, the weights associated to the d_2 th derivative of the expert and the gating parameter function. The appearance of the weighting parameters ρ and ϱ , besides the usual regularization parameters λ and χ , is motivated by the fact that one may wish to place a greater emphasis on sparsity in the d_2 th derivative than in the d_1 th derivative of the parameter functions, or vice versa (we will see the usage of them in the subsequent section of experimental study). In practice, the selection of ρ and ϱ is more about whether they are greater than or less than one (i.e., the emphasis on d_2 th) rather than select an exact value.

Note that, firstly, unlike the previous FME-Lasso, in iFME model the regularization operates on the functional derivative \mathbf{y}_k 's rather than the functional coefficients $\boldsymbol{\eta}_k$'s for the experts, and on the functional derivatives $\boldsymbol{\omega}_k$'s rather than the functional coefficients $\boldsymbol{\zeta}_k$'s for the gating network. Secondly, maximizing the penalized log-likelihood function (32) with penalization in (33) in iFME model must be coupled with the constrains (28). Follows are the two steps of the proposed EM-iFME algorithm.

E-Step The E-Step for the new iFME model calculates for each observation the conditional probability memberships of each expert k

$$\tau_{ik}^{(s)} = \mathbb{P}(Z_i = k | y_i, u_i(\cdot); \Psi^{(s)}) = \frac{\pi_k(\mathbf{s}_i; \mathbf{w}^{(s)}) \phi(y_i; \beta_{k,0}^{(s)} + \mathbf{v}_i^\top \mathbf{y}_k^{[d_1]^{(s)}}, \sigma_k^{2(s)})}{f(y_i | u_i(\cdot); \Psi^{(s)})}, \tag{34}$$

where $f(y_i|u_i(\cdot); \Psi^{(s)})$ is now calculated according to the iFME density given by (31).

M-Step The maximization is performed by separate maximizations w.r.t. the gating network parameters \mathbf{w} and the experts network parameters $\boldsymbol{\psi}_k$'s.

Updating the gating network parameters. The maximization step for updating the gating network parameters $\mathbf{w} = (\alpha_{1,0}, \boldsymbol{\omega}_1^{[d_1]^\top}, \dots, \alpha_{K-1,0}, \boldsymbol{\omega}_{K-1}^{[d_1]^\top})^\top$ consists of maximizing the function $Q_\chi(\mathbf{w}; \Psi^{(s)})$ given by

$$Q_\chi(\mathbf{w}; \Psi^{(s)}) = Q(\mathbf{w}; \Psi^{(s)})$$

$$- \chi \sum_{k=1}^{K-1} \left(\|\boldsymbol{\omega}_k^{[d_1]}\|_1 + \varrho \|\boldsymbol{\omega}_k^{[d_2]}\|_1 \right), \tag{35}$$

subject to

$$\boldsymbol{\omega}_k^{[d_2]} = \mathbf{A}_q^{[d_2]} \mathbf{A}_q^{[d_1]-1} \boldsymbol{\omega}_k^{[d_1]}, \quad \forall k \in [K - 1], \tag{36}$$

where $Q(\mathbf{w}; \Psi^{(s)}) = \sum_{i=1}^n \left[\sum_{k=1}^{K-1} \tau_{ik}^{(s)} \left(\alpha_{k,0} + \boldsymbol{\omega}_k^{[d_1]^\top} \mathbf{s}_i \right) - \log \left(1 + \sum_{k'=1}^{K-1} \exp \{ \alpha_{k',0} + \boldsymbol{\omega}_{k'}^{[d_1]^\top} \mathbf{s}_i \} \right) \right]$. This is a constrained version of the weighted regularized multinomial logistic regression problem, where the weights are the conditional probabilities $\tau_{ik}^{(s)}$.

To solve it, in the same spirit as when updating the gating network in the previous EM-Lasso algorithm, we use an outer loop that cycles over the gating function parameters to update them one by one. However, to update a single gating function parameter $\mathbf{w}_k = (\alpha_{k,0}, \boldsymbol{\omega}_k^{[d_1]^\top})^\top, k \in [K - 1]$, since the maximization problem (35) is now coupled with an additional constraint (36), rather than using a coordinate ascent algorithm as in EM-Lasso, we use the following alternative approach. For each single gating network k , using a Taylor series expansion, we partially approximate the smooth part of $Q_\chi(\mathbf{w}; \Psi^{(s)})$ defined in (35) w.r.t. \mathbf{w}_k at the current estimate $\mathbf{w}^{(t)}$, then maximize the resulting objective function (w.r.t. \mathbf{w}_k), subject to the corresponding constraint (w.r.t. k) in (36). It corresponds to solving the following penalized weighted least squares problem with constraints,

$$\begin{aligned} \max_{(\alpha_{k,0}, \boldsymbol{\omega}_k^{[d_1]}, \boldsymbol{\omega}_k^{[d_2]})} & - \frac{1}{2} \sum_{i=1}^n w_{ik} \left(c_{ik} - \alpha_{k,0} - \mathbf{s}_i^\top \boldsymbol{\omega}_k^{[d_1]} \right)^2 \\ & - \chi \left(\|\boldsymbol{\omega}_k^{[d_1]}\|_1 + \varrho \|\boldsymbol{\omega}_k^{[d_2]}\|_1 \right) \end{aligned} \tag{37}$$

subject to $\boldsymbol{\omega}_k^{[d_2]} = \mathbf{A}_q^{[d_2]} \mathbf{A}_q^{[d_1]-1} \boldsymbol{\omega}_k^{[d_1]}$,

where $w_{ik} = \pi_k(\mathbf{w}^{(t)}; \mathbf{s}_i) (1 - \pi_k(\mathbf{w}^{(t)}; \mathbf{s}_i))$ are the weights and $c_{ik} = \alpha_{k,0}^{(t)} + \mathbf{s}_i^\top \boldsymbol{\omega}_k^{[d_1]^{(t)}} + \frac{\tau_{ik}^{(s)} - \pi_k(\mathbf{w}^{(t)}; \mathbf{s}_i)}{w_{ik}}$ are the working responses computed given the current estimate $\mathbf{w}^{(t)}$. This problem can be solved by quadratic programming (see, e.g., Gaines et al. (2018)) or by using the Dantzig selector (Candes et al. 2007), which we opt to use in our experimental studies. The details of using Dantzig selector to solve problem (37) are given in ‘‘Appendix C.1’’.

Therefore, if $(\tilde{\alpha}_{k,0}, \tilde{\boldsymbol{\omega}}_k^{[d_1]}, \tilde{\boldsymbol{\omega}}_k^{[d_2]})$ is an optimal solution to problem (37), then $\tilde{\mathbf{w}}_k = (\tilde{\alpha}_{k,0}, \tilde{\boldsymbol{\omega}}_k^{[d_1]^\top})^\top$ is taken as an update for the gating parameter vector \mathbf{w}_k . We keep cycling over $k \in [K - 1]$ until there is no significant increase in the regularized Q -function (35).

Updating the experts network parameters. The maximization step for updating the expert parameter vector $\boldsymbol{\psi}_k =$

$(\beta_{k,0}, \boldsymbol{\gamma}_k^{[d_1]^\top}, \sigma_k^2)^\top$ consists of solving the following problem:

$$\begin{aligned} & \max_{(\beta_{k,0}, \boldsymbol{\gamma}_k^{[d_1]}, \boldsymbol{\gamma}_k^{[d_2]}, \sigma_k^2)} \sum_{i=1}^n \tau_{ik}^{(s)} \log \phi \left(y_i; \beta_{k,0} + \mathbf{v}_i^\top \boldsymbol{\gamma}_k^{[d_1]}, \sigma_k^2 \right) \\ & - \lambda \left(\|\boldsymbol{\gamma}_k^{[d_1]}\|_1 + \rho \|\boldsymbol{\gamma}_k^{[d_2]}\|_1 \right) \tag{38} \\ & \text{subject to } \boldsymbol{\gamma}_k^{[d_2]} = \mathbf{A}_p^{[d_2]} \mathbf{A}_p^{[d_1]-1} \boldsymbol{\gamma}_k^{[d_1]}. \end{aligned}$$

As in the previous EM-Lasso algorithm, we first fix σ_k^2 to its previous estimate and perform the update for $(\beta_{k,0}, \boldsymbol{\gamma}_k^{[d_1]})$, which corresponds to solving a penalized weighted least squares problem with constraints. This is performed by using the Dantzig selector, in the same manner as previously for solving problem (37). The corresponding technical details can be found in ‘‘Appendix C.2’’.

Once the $(\beta_{k,0}, \boldsymbol{\gamma}_k^{[d_1]})$ are updated, the straightforward update for the variance σ_k^2 is given by the standard estimate of a weighted Gaussian regression. More specifically, let $(\tilde{\beta}_{k,0}, \tilde{\boldsymbol{\gamma}}_k^{[d_1]}, \tilde{\boldsymbol{\gamma}}_k^{[d_2]})$ be the solution to the problem (38) (with σ_k^2 fixed to $\sigma_k^{2(s)}$), the updates for expert parameter vector $\boldsymbol{\psi}_k$ are given by

$$\begin{aligned} (\beta_{k,0}^{(s+1)}, \boldsymbol{\gamma}_k^{[d_1]^{(s+1)}}) &= (\tilde{\beta}_{k,0}, \tilde{\boldsymbol{\gamma}}_k^{[d_1]}), \\ \sigma_k^{2(s+1)} &= \frac{\sum_{i=1}^n \tau_{ik}^{(s)} \left(y_i - \beta_{k,0}^{(s+1)} - \mathbf{v}_i^\top \boldsymbol{\gamma}_k^{[d_1]^{(s+1)}} \right)^2}{\sum_{i=1}^n \tau_{ik}^{(s)}}. \end{aligned}$$

Thus, at the end of the M-Step, we obtain a parameter vector update $\boldsymbol{\Psi}^{(s+1)} = (\mathbf{w}^{(s+1)}, \boldsymbol{\psi}_1^{(s+1)}, \dots, \boldsymbol{\psi}_K^{(s+1)})$ for the next EM iteration, where $\mathbf{w}^{(s+1)}$ and $\boldsymbol{\psi}_k^{(s+1)}$, $k \in [K]$, are, respectively, the updates of the gating network parameters and the experts network parameters, calculated by the two procedures described above. Alternating the E-Step and M-Step until convergence, i.e., when there is no longer a significant change in the values of the penalized log-likelihood (32), leads to a penalized maximum likelihood estimate $\hat{\boldsymbol{\Psi}}$ for $\boldsymbol{\Psi}$.

Estimating the coefficient functions Finally, once the parameter vector of iFME model has been estimated, the coefficient functions of the gating network $\alpha_k(t)$, $k \in [K - 1]$ and the ones of the experts network $\beta_k(t)$, $k \in [K]$, can be reconstructed by the following formulas

$$\begin{aligned} \hat{\alpha}_k(t) &= \mathbf{b}_q(t)^\top \mathbf{A}_q^{-1} \hat{\boldsymbol{\omega}}_k^{[d_1]}, \\ \hat{\beta}_k(t) &= \mathbf{b}_p(t)^\top \mathbf{A}_p^{-1} \hat{\boldsymbol{\gamma}}_k^{[d_1]}, \end{aligned} \tag{39}$$

where $\hat{\boldsymbol{\omega}}_k^{[d_1]}$ and $\hat{\boldsymbol{\gamma}}_k^{[d_1]}$ are respectively the regularized maximum likelihood estimates for $\boldsymbol{\omega}_k^{[d_1]}$ and $\boldsymbol{\gamma}_k^{[d_1]}$.

3.4 Non-linear regression and clustering with FME models

Once the model parameters have been estimated, one can then construct an estimate of the unknown functional non-linear regression function, as well as a clustering of the data into groups of similar pairs of observations.

For the aim of functional non-linear regression, the unknown non-linear regression function with functional predictors is given by the following conditional expectation $\hat{y}|u(\cdot) = \mathbb{E}[Y|U(\cdot); \hat{\boldsymbol{\Psi}}]$, which is defined by $\hat{y}_i|u_i(\cdot) = \sum_{k=1}^K \pi_k(\mathbf{r}_i; \hat{\boldsymbol{\xi}})(\hat{\beta}_{k,0} + \hat{\boldsymbol{\eta}}_k^\top \mathbf{x}_i)$ for the FME model (13), and by $\hat{y}_i|u_i(\cdot) = \sum_{k=1}^K \pi_k(\mathbf{s}_i; \hat{\mathbf{w}})(\hat{\beta}_{k,0} + \boldsymbol{\gamma}_k^{[d_1]^\top} \mathbf{v}_i)$ for the iFME model (31).

For clustering, a soft partition of the data into K clusters, represented by the estimated posterior probabilities $\hat{\tau}_{ik} = \mathbb{P}(Z_i = k|u_i, y_i; \hat{\boldsymbol{\Psi}})$, is obtained. A hard partition can also be computed according to the Bayes’ optimal allocation rule. That is, by assigning each curve to the component having the highest estimated posterior probability τ_{ik} , defined by (15) for FME or by (34) for the iFME model, using the MLE $\hat{\boldsymbol{\Psi}}$ of $\boldsymbol{\Psi}$:

$$\hat{z}_i = \arg \max_{1 \leq k \leq K} \tau_{ik}(\hat{\boldsymbol{\Psi}}), \quad i \in [n],$$

where \hat{z}_i denotes the estimated cluster label for the i th curve.

3.5 Tuning parameters and model selection

In practice, appropriate values of the tuning parameters should be chosen. In using FME, this cover the selection of K , the number of experts, and r , p , and q , the dimensions of B-spline bases used to approximate, respectively, the predictors, the experts, and the gating network functions, although they can be chosen to be equal. For the FME-Lasso, additionally the ℓ_1 penalty constants χ and λ in (19) should be chosen. For the iFME model, the tuning parameters include also d_1 , d_2 , the two derivatives related to the sparsity constraints, and ρ, ϱ , the weights associated to the d_2 th derivative of the expert and gating functions (see (33)).

The selection of the tuning parameters can be performed by a cross-validation procedure with a grid search scheme to select the best combination. An alternative is to use the well-known BIC criterion (Schwarz 1978) or, in our paper, its extension, called modified BIC (Städler et al. 2010) defined as

$$\text{mBIC} = L(\hat{\boldsymbol{\Psi}}) - \text{df}(\hat{\boldsymbol{\Psi}}) \frac{\log n}{2}, \tag{40}$$

where $\hat{\boldsymbol{\Psi}}$ is the obtained log-likelihood estimator (for the FME model) or penalized log-likelihood estimator (for the FME-Lasso and iFME models), and the number of degrees

of freedom $df(\widehat{\Psi})$ is the effective number of parameters of the model, given by

$$df(\widehat{\Psi}) = \begin{cases} df(\zeta) + (K - 1) + df(\eta) + K + K & \text{for the FME and FME-Lasso models,} \\ df(\omega) + (K - 1) + df(\gamma) + K + K & \text{for the iFME model,} \end{cases}$$

in which the quantities $df(\zeta)$, $df(\eta)$, $df(\omega)$ and $df(\gamma)$ are, respectively, the counts for non-zero free coefficients in the vectors ζ , η , ω , and γ . Note that, because of the constraints (28) for the iFME model, free coefficients in ω and γ consist of only the part related to the d_1 derivative. That is, $df(\zeta) = \sum_{k=1}^{K-1} \sum_{j=1}^q \mathbf{1}_{\{\zeta_{kj} \neq 0\}}$, $df(\eta) = \sum_{k=1}^K \sum_{j=1}^p \mathbf{1}_{\{\eta_{kj} \neq 0\}}$, $df(\omega) = \sum_{k=1}^{K-1} \sum_{j=1}^q \mathbf{1}_{\{\omega_{kj}^{[d_1]} \neq 0\}}$ and $df(\gamma) = \sum_{k=1}^K \sum_{j=1}^p \mathbf{1}_{\{\gamma_{kj}^{[d_1]} \neq 0\}}$. We apply both the BIC and the modified BIC in our experimental study.

4 Experimental study

We study the performances of the FME, FME-Lasso, and iFME models in regression and clustering problems by considering simulated scenarios and real data with functional predictors and scalar responses. The interests of this study consist of the prediction performance as well as the estimation of the functional parameters, i.e., the expert and gating functions in the ME model, along with the clustering partition of the considered heterogeneous data.

4.1 Evaluation criteria

We will use the following criteria, for where applicable, to assess and compare the performances of the models and the related algorithms. For regression evaluation, we use the *relative prediction error* (RPE) and the *correlation* (Corr) index as standard criteria to evaluate the prediction performance, i.e., to quantify the relationship between the true and the predicted values of the scalar outputs. The RPE is defined by $RPE = \sum_{i=1}^n (y_i - \widehat{y}_i)^2 / \sum_{i=1}^n y_i^2$, where y_i and \widehat{y}_i are, respectively, the true and the predicted response of the i th observation in the testing set. For clustering evaluation, we use the *Rand index* (RI), the *adjusted Rand index* (ARI), and the *clustering error* (ClusErr), as standard criteria to evaluate the clustering performance, i.e., to quantify how similar the testing observations are presented in the true partition compared to the predicted partition. To evaluate the parameters estimation performance, we compute the *mean squared error* (MSE) between the true and the estimated functional parameters. The MSE between a true function $g(\cdot)$ and its estimate $\widehat{g}(\cdot)$ is defined by

$$MSE(\widehat{g}(\cdot)) = \frac{1}{m} \sum_{j=1}^m (g(t_j) - \widehat{g}(t_j))^2, \tag{41}$$

with m being the number of time points taken into account. The function $g(\cdot)$ here corresponds to an expert function $\beta(\cdot)$, or a gating function $\alpha(\cdot)$. The parameter functions are reconstructed from the model parameters using (7) and (8) for both FME and FME-Lasso models, and using (39) for iFME model.

The values of these criteria are averaged over N Monte Carlo runs ($N = 100$ for simulation, for the real data, see the corresponding section). Note that, the average over N sample replicates of $MSE(\widehat{g}(\cdot))$ is equivalent to the usual Mean Integrated Squared Error (MISE) defined by $MISE(\widehat{g}(\cdot)) = \mathbb{E} \left[\int_{\mathcal{T}} (\widehat{g}(t) - g(t))^2 dt \right]$, where the integral here is calculated numerically by a Riemann sum over the grid t_1, \dots, t_m .

4.2 Simulation studies

4.2.1 Simulation parameters and experimental protocol

We generate data from a K -component functional mixture of Gaussian experts model that relates a scalar response $y \in \mathbb{R}$ to a univariate functional predictor $X(t)$, $t \in \mathcal{T}$ defined on a domain $\mathcal{T} \subset \mathbb{R}$. The data generation protocol is detailed in ‘‘Appendix D.1’’. The parameters that were used in this data generating process (54) are described and provided in the simulation parameters and experimental protocol section in ‘‘Appendix D.2’’. We study different simulation scenarios (sample sizes n , number of observations per time-series input m , noise levels σ_δ^2, \dots) summarised in Sect. D.2 and Table. 9.

Figure 1 displays, for each scenario, 10 randomly taken predictors colored according to their corresponding clusters.

For each run, the concerned dataset is randomly split into a training set and testing set of equal size, the model parameters are estimated using training set, with the tuning parameters selected by maximizing the modified BIC (40). The evaluation criteria are computed on testing set and reported for each model accordingly.

4.2.2 Some implementation details

For all scenarios, for all datasets, we implemented the three proposed models with 10 EM runs and with a tolerance of 10^{-6} . For the iFME model, in principle, for each parameter function the two derivatives d_1 and d_2 to be penalized, and the weights for the latter (i.e., ρ and ϱ) can be seen as tun-

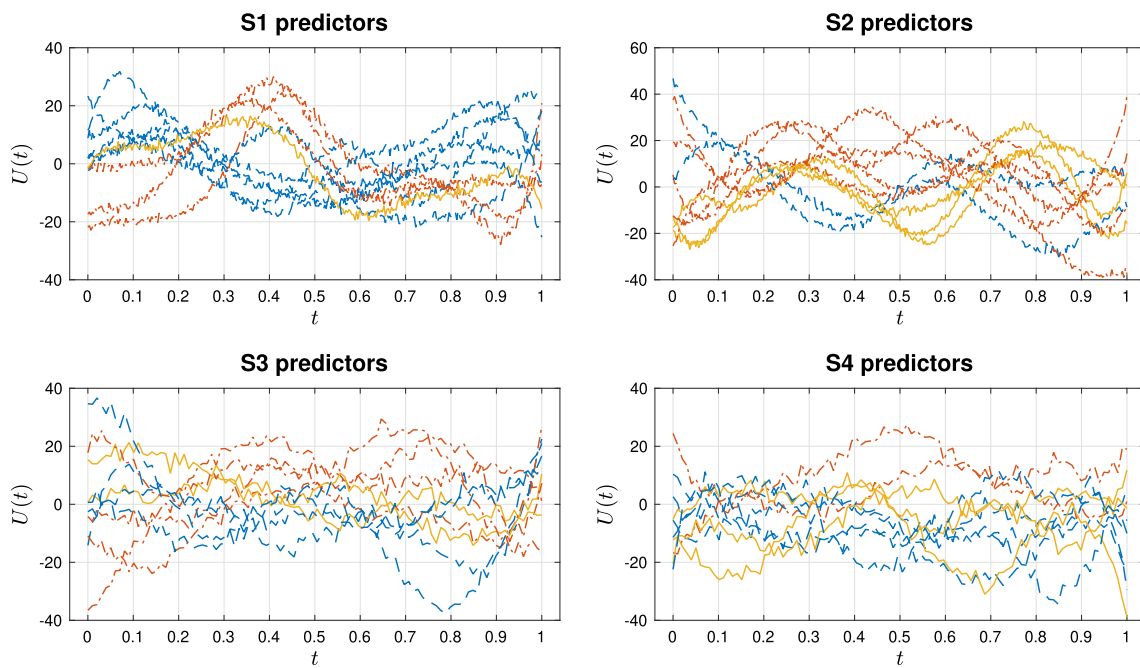


Fig. 1 10 randomly taken predictors in scenarios *S1* (large m and small σ_δ), *S2* (small m and small σ_δ), *S3* (large m and large σ_δ), and *S4* (small m and large σ_δ). Here, the noisy predictors $U_i(\cdot)$ are colored (blue, red, yellow) according to their true cluster labels $Z_i \in \{1, 2, 3\}$

ing parameters. However, such an implementation could be computationally expensive in this simulation study with 400 datasets in total and 10 EM runs for each dataset. Therefore, we opted to fix d_1 and d_2 for all implementations ($d_1 = 0$ and $d_2 = 3$ for both expert and gating networks.) and left ρ and ϱ to be selected in some sets of targeted values. The choices of the targeted values were made by the following straightforward arguments. Since $\beta_1(\cdot)$ and $\beta_2(\cdot)$ have zero-valued regions, the weight for their zeroth derivative in penalization term should be large, equivalently, the weight for the third derivative should be small, so ρ is selected in a set of small values: $\rho \in \{10^{-2}, 10^{-3}, 10^{-4}\}$. On the other hand, for $\alpha_1(t)$ and $\alpha_2(t)$, we select $\varrho \in \{10, 10^2, 10^3\}$ as we should emphasize sparsity in their third derivative.

4.2.3 Simulation results

Clustering and prediction performances We report in Table 1 the results of regression and clustering tasks on simulated datasets in the four considered scenarios.

The mean and standard error of the relative predictions error (RPE) and correlation (Corr) summarize the regression performance, while the mean and standard error of the Rand index (RI), adjusted Rand index (ARI) and clustering error (ClusErr) summarize the clustering performance.

As we can see from Table 1, all the models have very good performances on both regression and clustering tasks, with high correlation, RI, ARI, and small RPE and clustering error. The iFME appears to slightly have a better performance than

the others in all scenarios. The low standard errors confirm the stability of the algorithms.

Figure 2 shows the clustering results obtained by the models with highest values of the modified BIC criterion, on a dataset selected in scenario *S1*. Here, we plotted the responses against the predictors at two specific time points: $t_1 = 0$ and $t_{50} = 0.5$. The highly accurate predictions (in both regression and clustering) can be seen visually through Fig. 2. This figure also shows that it is difficult to cluster these data according to a few number of time observations, for example in \mathbb{R}^2 , according to $\{(U_i(t_0), y_i)\}_{i=1}^n$ or $\{(U_i(t_{50}), y_i)\}_{i=1}^n$, which suggests using functional approaches.

Comparison with functional regression mixtures (FMR): Finally, we compare our proposed models with the functional mixture regression (FMR) model proposed in Yao et al. (2010). In their approach, the functional predictors are first projected onto its eigenspace, then the obtained new coordinates are fed to the standard mixture regression model to estimate the weights and the coefficients of the $\hat{\beta}(\cdot)$ in that eigenspace. They performed functional principal component analysis (FPCA) to obtain estimates for the eigenfunctions and the principal component scores (which serve as predictors). The number of relevant FPCA components are chosen automatically for each dataset by selecting the minimum number of components that explain 90% of the total variation of the predictors. It is noticed that, in their simulation studies, the authors computed the scalar responses by using conditional prediction, i.e., the true y_i were used to determine

Table 1 Evaluation criteria of FME, FME-Lasso and iFME models for test data in scenarios $S1, \dots, S4$

	RPE	Corr	ARI	ClusErr
$S1 (m = 300, \sigma_\delta^2 = 1)$				
FME	.1552 _(.1282)	.9188 _(.0602)	.7852 _(.0934)	.0899 _(.0434)
FME-Lasso	.1390 _(.1074)	.9224 _(.0603)	.7670 _(.1436)	.1112 _(.0891)
iFME	.1334_(.0977)	.9287_(.0504)	.7997_(.0796)	.0792_(.0340)
$S2 (m = 100, \sigma_\delta^2 = 1)$				
FME	.1600 _(.1698)	.9164 _(.0712)	.7966_(.0881)	.0852 _(.0427)
FME-Lasso	.1656 _(.1316)	.9071 _(.0733)	.7455 _(.1520)	.1236 _(.0961)
iFME	.1540_(.0950)	.9192_(.0476)	.7955 _(.0822)	.0820_(.0341)
$S3 (m = 300, \sigma_\delta^2 = 4)$				
FME	.1724 _(.1807)	.9110 _(.0778)	.7798 _(.0838)	.0918 _(.0415)
FME-Lasso	.1457 _(.1504)	.9197 _(.0785)	.7629 _(.1524)	.1115 _(.0928)
iFME	.1432_(.0860)	.9228_(.0448)	.7987_(.0899)	.0783_(.0361)
$S4 (m = 100, \sigma_\delta^2 = 4)$				
FME	.2251 _(.4462)	.9048 _(.0822)	.7816_(.1040)	.0927 _(.0497)
FME-Lasso	.1432_(.1229)	.9188_(.0683)	.7526 _(.1654)	.1215 _(.1055)
iFME	.1639 _(.0978)	.9125 _(.0521)	.7798 _(.0848)	.0864_(.0325)

The reported values are the averages of 100 Monte Carlo runs with standard errors in parentheses. The bold values correspond to the best solution




Fig. 2 Scatter plots of \hat{y}_i against $U_i(t_1)$ (top panels) and $U_i(t_{150})$ (bottom panels) on a randomly selected dataset. Here, the clustering errors are 5.5%, 4.75% and 5.0% for FME, FME-Lasso and iFME models, respectively

which cluster the observation belongs to. Then the predicted \hat{y}_i is calculated as the conditional mean of the density of the corresponding cluster. For comparison with that approach, we also used this strategy to make predictions in our models. We further employed the FMR model with the B-spline functional representation, instead of the functional PCA, the number of B-spline functions is set to be the same as the number of basis functions used in our models. Table 2 shows the evaluation criteria corresponding to the considered models evaluated on 100 datasets in scenario $S3$. Here, FMR-PC

is the original model of Yao et al. (2010) and FMR-B is the modified one with B-spline bases. As expected, FME, FME-Lasso, and iFME, which are more flexible compared to the the FMR model, allows to capture more complexity in the data, thanks to the predictor-depending mixture weights, and provides clearly better results than the FMR alternatives.

Parameter estimation performance To evaluate the parameter estimation performance, we consider the functional parameter functions estimation error as defined by (41). This error between the true function and the estimated one,

Table 2 Performance comparison of the models for datasets in scenarios $S3$

	RPE	Corr	ARI	ClusErr
FME	.0614 _(.0896)	.9681 _(.0469)	.8523 _(.0934)	.0618 _(.0490)
FME-Lasso	.0145 _(.0254)	.9924 _(.0133)	.9582 _(.0641)	.0174 _(.0325)
iFME	.0139_(.0086)	.9928_(.0046)	.9594_(.0430)	.0151_(.0166)
FMR-B	.0460 _(.0399)	.9768 _(.0201)	.7064 _(.1044)	.1124 _(.0434)
FMR-PC	.0191 _(.0331)	.9902 _(.0164)	.8345 _(.0794)	.0612 _(.0361)

The reported values are the averages of 100 Monte Carlo samples with standard errors in parentheses

provides an evaluation of how well the proposed models reconstruct the hidden functional gating and expert networks. In this evaluation, we considered scenario $S1$, i.e., $m = 300$, $\sigma_\delta^2 = 1$. Moreover, in order to provide an idea of the impact of training size on parameter estimation, we run the models with different training sizes (share the same scenario $S1$) and report the MSE for each function, for each model in Table 3. It shows that there are significant improvements, even with small training size, when using iFME model in estimating the gating network, compared to the FME and FME-Lasso model.

Now, in order to evaluate the designed sparsity of zeroth and third derivatives of the reconstructed functions, we compute the MSEs versus their true values, i.e., zeros, on each designed intervals. In particular, we divide the domain $\mathcal{T} = [0, 1]$ into three parts: $\mathcal{T}_1 = [0, 0.3)$, $\mathcal{T}_2 = [0.3, 0.7)$, and $\mathcal{T}_3 = [0.7, 1]$, then for each model the MSEs on the corresponding intervals are reported in Table 4.

The reported values show that, as expected, iFME model is better than both FME and FME-Lasso in providing sparse solutions with respect to the derivatives of the parameter functions.

Table 5 shows the means and the standard errors of the estimated intercepts and variances. Note that these coefficients are not considered in the penalization. For the intercepts $\beta_{k,0}$, all the models estimated them very well, while for the intercepts $\alpha_{k,0}$, iFME is slightly better than the others. For the variances, FME-Lasso gave the estimated values closest with the true values on average.

The initialization is crucial for the EM algorithm. In all of our experimental studies, the model parameter vector Ψ of the FME model was initialized as follows (similar for the FME-Lasso and iFME models). Firstly, we perform a K -means algorithm on the predictors $\{\mathbf{x}_i\}_{i=1}^n$, then for each estimated group k , $(\beta_{k,0}^{(0)}, \boldsymbol{\eta}_k^{(0)})$ is initialized as the solution to the linear regression problem $y_i = \beta_{k,0} + \mathbf{x}_i^\top \boldsymbol{\eta}_k$, for i belongs to group k th. The parameter $\sigma_k^{(0)}$ is then initialized as the estimated variance within group k th. For the gating parameter, we simply drawn each component of $\boldsymbol{\xi}^{(0)}$ randomly from the uniform distribution $\mathcal{U}(0, 1)$.

However, to see the impact of different initialization strategies, we performed a side experiment with data taken from $S1$

scenario. In particular, the initialization for expert parameters is fixed as described before, while that for gating parameter varies in “rand.,” “zeros”, and “LR”. Here, “rand.” is the random strategy described above, “zeros” is the strategy where all coefficients of $\boldsymbol{\xi}$ are initialized as zeros, i.e., we are putting equal weights for the experts; and “LR” is the strategy where we perform a logistic regression with the predictors are \mathbf{r}_i (for FME, FME-Lasso) or \mathbf{s}_i (for iFME), and the responses are the labels resulted by K -means on them. Table 6 shows the time to convergence, and the log-likelihood, RPE on testing sets corresponding with the three strategies. We can see that, in terms of log-likelihood and RPE, simple strategies as “rand.” and “zeros” perform quite well, while in terms of time to convergence, LR is a good choice for the iFME model.

Finally, to illustrate the selection of the number of expert components using BIC and/or modified BIC in this simulation study, we provide, in Fig. 3, the plots of these criteria against the number of experts for each model. Here, we implemented the models with all fixed tuning parameters, except K which varies in the set $\{1, \dots, 6\}$. We can observe that BIC selects the correct number $K = 3$ for both FME-Lasso and iFME, while it selects $K = 4$ for FME. However, the modified BIC selects $K = 4$ for both FME and FME-Lasso, and selects the true number of components $K = 3$ for iFME.

4.3 Application to real data

In this section, we apply the FME, FME-Lasso and iFME models to two well-known real datasets, Canadian weather and Diffusion Tensor Imaging (DTI). For each dataset, we perform clustering and investigate the prediction performance, estimate the functional mixture of experts models with different number of experts K and perform the selection of K using modified BIC, and discuss the obtained results.

4.3.1 Canadian weather data

Canadian weather data has been introduced in Ramsay and Silverman (2005) and is also available in the R package `fda`. This dataset consists of $m = 365$ daily temperature measurements (averaged over the year 1961 to 1994) at $n = 35$

Table 3 Average of 100 Monte Carlo runs of MSE between the estimated functions resulted by FME, FME-Lasso and iFME models in S1 scenario

	$\widehat{\beta}_1(\cdot)$	$\widehat{\beta}_2(\cdot)$	$\widehat{\beta}_3(\cdot)$	$\widehat{\alpha}_1(\cdot)$	$\widehat{\alpha}_2(\cdot)$
Training size: 100					
FME	11.68 _(56.63)	6.21 _(24.34)	8.57 _(27.60)	5.2e+4 _(7.2e+4)	5.7e+4 _(5.9e+6)
FME-Lasso	0.53 _(0.19)	0.66 _(0.55)	0.62 _(0.79)	17.06 _(28.68)	19.50 _(37.07)
iFME	1.09 _(0.78)	1.05 _(0.82)	2.44 _(3.06)	5.76 _(7.78)	4.25 _(3.58)
Training size: 200					
FME	0.77 _(0.48)	0.61 _(0.30)	0.19 _(0.20)	8.55 _(11.99)	9.57 _(8.29)
FME-Lasso	0.54 _(0.18)	0.54 _(0.20)	0.31 _(0.81)	10.71 _(16.63)	12.64 _(16.33)
iFME	0.55 _(0.50)	0.48 _(0.38)	1.26 _(2.91)	5.34 _(6.21)	2.81 _(2.82)
Training size: 300					
FME	0.62 _(0.56)	0.59 _(0.50)	0.16 _(0.18)	12.58 _(19.83)	16.49 _(30.97)
FME-Lasso	0.60 _(0.52)	0.57 _(0.44)	0.17 _(0.19)	11.26 _(16.56)	13.15 _(24.80)
iFME	0.56 _(0.43)	0.59 _(0.50)	0.49 _(0.47)	2.98 _(3.12)	3.16 _(3.02)
Training size: 400					
FME	0.54 _(0.23)	0.53 _(0.19)	0.11 _(0.12)	8.27 _(11.73)	7.08 _(8.53)
FME-Lasso	0.52 _(0.18)	0.51 _(0.16)	0.15 _(0.18)	6.09 _(8.98)	5.14 _(6.51)
iFME	0.34 _(0.19)	0.38 _(0.25)	0.74 _(0.62)	3.06 _(4.21)	2.77 _(3.21)
Training size: 500					
FME	0.49 _(0.22)	0.50 _(0.28)	0.09 _(0.11)	5.17 _(7.52)	3.90 _(5.69)
FME-Lasso	0.49 _(0.18)	0.49 _(0.21)	0.09 _(0.11)	4.62 _(7.28)	7.63 _(16.27)
iFME	0.35 _(0.23)	0.40 _(0.25)	0.55 _(0.63)	2.11 _(2.82)	2.68 _(4.70)

Table 4 MSE of the derivatives of reconstructed functions on the corresponding interested intervals, in which, $\mathcal{T}_1 = [0, 0.3]$, $\mathcal{T}_2 = [0.3, 0.7]$, $\mathcal{T}_3 = [0.7, 1]$ and $\mathcal{T} = [0, 1]$

	on \mathcal{T}_2		on $\mathcal{T}_1 \cup \mathcal{T}_3$		on \mathcal{T}		
	$\widehat{\beta}_1^{(0)}(\cdot)$	$\widehat{\beta}_2^{(0)}(\cdot)$	$\widehat{\beta}_1^{(3)}(\cdot)$	$\widehat{\beta}_2^{(3)}(\cdot)$	$\widehat{\beta}_3^{(3)}(\cdot)$	$\widehat{\alpha}_1^{(3)}(\cdot)$	$\widehat{\alpha}_2^{(3)}(\cdot)$
Training size: 100							
FME	0.26 (0.15)	0.48 (0.91)	0.15 (0.07)	0.12 (0.08)	3.96e-08 (1.37e-07)	9.24e-06 (3.32e-05)	1.08e-05 (2.62e-05)
FME-Lasso	0.27 (0.19)	0.26 (0.16)	0.14 (0.07)	0.15 (0.09)	1.64e-09 (7.64e-09)	7.37e-09 (2.64e-08)	5.12e-09 (1.70e-08)
iFME	0.40 (0.26)	0.37 (0.30)	0.20 (0.10)	0.18 (0.11)	6.48e-09 (2.44e-08)	4.05e-09 (1.33e-08)	2.46e-09 (4.51e-09)
Training size: 200							
FME	0.43 (0.56)	0.22 (0.09)	0.16 (0.10)	0.14 (0.05)	2.23e-11 (2.69e-11)	1.64e-10 (2.33e-10)	3.83e-10 (8.90e-10)
FME-Lasso	0.21 (0.11)	0.20 (0.09)	0.15 (0.05)	0.15 (0.05)	3.62e-11 (6.04e-11)	2.36e-10 (3.38e-10)	2.96e-10 (6.41e-10)
iFME	0.18 (0.08)	0.17 (0.09)	0.12 (0.05)	0.12 (0.05)	9.95e-10 (2.00e-09)	1.44e-09 (2.72e-09)	8.52e-10 (1.22e-09)
Training size: 300							
FME	0.25 (0.30)	0.22 (0.25)	0.16 (0.10)	0.15 (0.09)	1.72e-11 (2.30e-11)	2.55e-10 (4.61e-10)	1.01e-10 (1.42e-10)
FME-Lasso	0.25 (0.28)	0.22 (0.23)	0.16 (0.09)	0.15 (0.08)	1.73e-11 (2.52e-11)	1.41e-10 (2.00e-10)	1.64e-10 (3.77e-10)
iFME	0.25 (0.19)	0.24 (0.19)	0.16 (0.08)	0.15 (0.08)	5.94e-10 (1.26e-09)	5.41e-10 (7.26e-10)	7.15e-10 (1.05e-09)

Table 4 continued

	on \mathcal{T}_2		on $\mathcal{T}_1 \cup \mathcal{T}_3$		on \mathcal{T}		
	$\hat{\beta}_1^{(0)}(\cdot)$	$\hat{\beta}_2^{(0)}(\cdot)$	$\hat{\beta}_1^{(3)}(\cdot)$	$\hat{\beta}_2^{(3)}(\cdot)$	$\hat{\beta}_3^{(3)}(\cdot)$	$\hat{\alpha}_1^{(3)}(\cdot)$	$\hat{\alpha}_2^{(3)}(\cdot)$
Training size: 400							
FME	0.23 (0.12)	0.23 (0.10)	0.16 (0.05)	0.16 (0.06)	1.62e-11 (3.53e-11)	1.29e-10 (3.59e-10)	8.14e-11 (1.02e-10)
FME-Lasso	0.22 (0.11)	0.21 (0.10)	0.16 (0.05)	0.16 (0.05)	1.90e-11 (3.55e-11)	9.71e-11 (1.63e-10)	1.03e-10 (1.54e-10)
iFME	0.14 (0.08)	0.16 (0.09)	0.08 (0.03)	0.08 (0.04)	1.87e-08 (1.42e-08)	1.13e-08 (1.77e-08)	9.91e-09 (1.05e-08)
Training size: 500							
FME	0.22 (0.11)	0.21 (0.10)	0.17 (0.05)	0.17 (0.05)	2.12e-11 (4.66e-11)	8.98e-11 (2.70e-10)	1.03e-10 (2.83e-10)
FME-Lasso	0.21 (0.09)	0.20 (0.09)	0.17 (0.05)	0.17 (0.05)	1.90e-11 (3.64e-11)	7.02e-11 (1.44e-10)	8.91e-11 (1.69e-10)
iFME	0.14 (0.07)	0.16 (0.08)	0.08 (0.04)	0.08 (0.04)	1.55e-08 (1.34e-08)	1.03e-08 (1.58e-08)	8.34e-09 (1.20e-08)

The reported values are the averaged of 100 Monte Carlo samples with standard errors in parentheses

Table 5 Intercepts and variances obtained by FME, FME-Lasso and iFME models in scenario S1

True value	$\hat{\beta}_{1,0}$	$\hat{\beta}_{2,0}$	$\hat{\beta}_{3,0}$	$\hat{\alpha}_{1,0}$	$\hat{\alpha}_{2,0}$	$\hat{\sigma}_1$	$\hat{\sigma}_2$	$\hat{\sigma}_3$
	-5	0	5	-10	-10	5	5	5
Training size: 100								
FME	-16.90 (18.91)	-3.18 (3.25)	4.17 (8.87)	-31.37 (30.74)	-25.20 (41.27)	32.38 (40.05)	15.19 (25.15)	21.47 (24.71)
FME-Lasso	-6.48 (9.37)	-0.07 (1.95)	5.51 (4.63)	-12.53 (8.99)	-12.69 (10.99)	6.86 (10.98)	6.73 (8.08)	10.01 (20.60)
iFME	-4.95 (1.14)	0.21 (1.08)	5.05 (0.90)	-9.44 (3.85)	-9.39 (3.31)	5.70 (1.46)	5.75 (1.49)	4.70 (1.51)
Training size: 200								
FME	-4.83 (0.78)	0.14 (0.54)	4.89 (0.38)	-23.77 (26.36)	-16.10 (6.65)	5.77 (1.26)	5.63 (0.95)	5.24 (0.79)
FME-Lasso	-4.97 (0.66)	0.07 (0.60)	4.99 (0.45)	-12.70 (4.45)	-13.72 (4.38)	5.94 (1.23)	5.99 (1.12)	4.92 (1.30)
iFME	-4.85 (0.72)	0.19 (0.68)	4.99 (0.42)	-9.30 (3.15)	-9.21 (2.83)	5.77 (0.90)	5.72 (1.02)	5.28 (1.27)
Training size: 300								
FME	-4.99 (0.55)	-0.06 (0.54)	4.98 (0.26)	-16.82 (13.09)	-17.18 (14.89)	5.92 (1.14)	5.87 (0.94)	5.34 (0.84)
FME-Lasso	-5.03 (0.54)	-0.03 (0.53)	4.98 (0.27)	-13.06 (4.54)	-12.96 (4.96)	6.08 (1.10)	5.99 (1.07)	5.37 (0.89)
iFME	-4.93 (0.61)	0.04 (0.56)	5.00 (0.29)	-8.96 (2.48)	-8.85 (2.57)	5.54 (0.75)	5.53 (0.72)	5.53 (0.87)
Training size: 400								
FME	-5.02 (0.43)	0.01 (0.40)	5.04 (0.28)	-16.37 (9.93)	-15.39 (7.83)	5.99 (1.01)	6.20 (1.09)	5.47 (0.77)
FME-Lasso	-5.03	-0.02	5.02	-13.18	-12.65	5.99	6.13	5.45

Table 5 continued

True value	$\hat{\beta}_{1,0}$	$\hat{\beta}_{2,0}$	$\hat{\beta}_{3,0}$	$\hat{\alpha}_{1,0}$	$\hat{\alpha}_{2,0}$	$\hat{\sigma}_1$	$\hat{\sigma}_2$	$\hat{\sigma}_3$
	-5	0	5	-10	-10	5	5	5
iFME	(0.43)	(0.38)	(0.27)	(4.61)	(4.57)	(0.94)	(1.01)	(0.84)
	-4.89	0.05	5.01	-10.17	-9.82	5.08	5.22	5.57
	(0.45)	(0.35)	(0.27)	(2.79)	(2.61)	(0.60)	(0.58)	(0.85)
Training size: 500								
FME	-5.00	-0.01	5.01	-14.09	-15.42	6.09	6.06	5.44
	(0.41)	(0.38)	(0.25)	(7.26)	(9.64)	(0.98)	(0.93)	(0.69)
FME-Lasso	-4.99	-0.01	5.01	-11.35	-12.38	6.13	6.06	5.44
	(0.40)	(0.37)	(0.26)	(3.27)	(4.15)	(0.98)	(0.90)	(0.68)
iFME	-4.92	0.06	5.02	-9.23	-9.61	5.23	5.13	5.60
	(0.44)	(0.34)	(0.27)	(1.93)	(2.68)	(0.55)	(0.55)	(0.75)

The reported values are the averages of 100 Monte Carlo samples with standard errors in parentheses

Table 6 Comparison of different initialization strategies. The reported values are the mean and standard error (in parentheses) over 100 Monte Carlo runs

	Log likelihood			RPE			Time (s)		
	Rand	Zeros	LR	Rand	Zeros	LR	Rand	Zeros	LR
FME	-215.52	-221.03	-285.23	0.19	0.20	0.38	0.21	0.16	0.52
	(24.77)	(18.51)	(69.30)	(0.16)	(0.13)	(0.29)	(0.50)	(0.12)	(1.99)
FME-Lasso	-235.35	-239.47	-244.99	0.14	0.13	0.20	1.13	1.08	4.63
	(21.98)	(18.54)	(63.09)	(0.11)	(0.10)	(0.16)	(1.23)	(2.48)	(15.51)
iFME	-254.65	-256.26	-336.87	0.13	0.13	0.35	6.58	4.17	3.91
	(7.83)	(13.76)	(80.38)	(0.09)	(0.09)	(0.30)	(4.45)	(3.65)	(2.08)

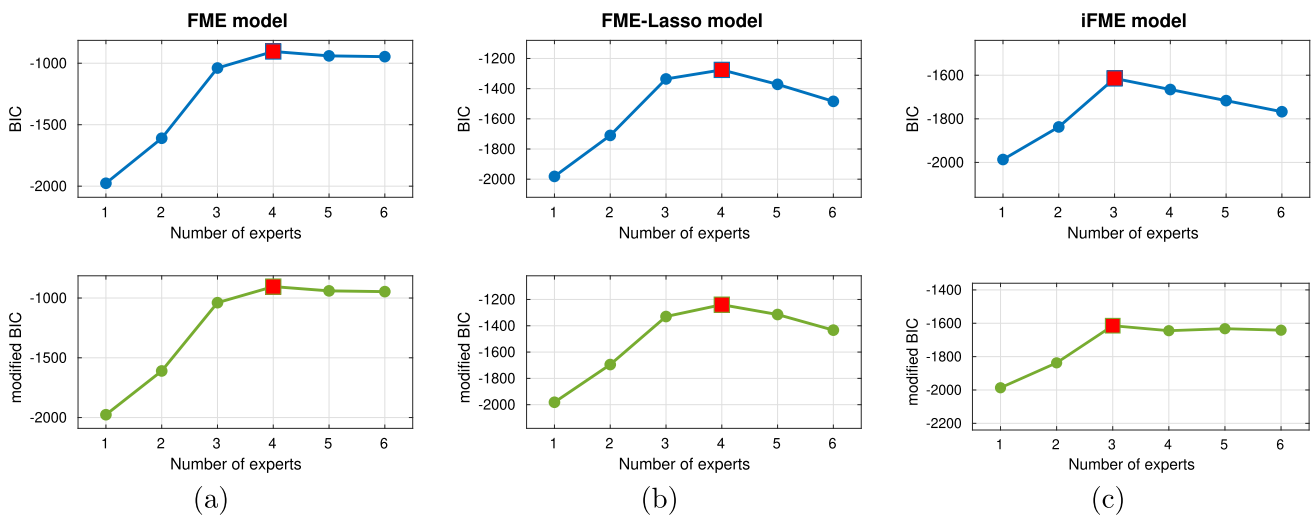


Fig. 3 Values of BIC (top) and modified BIC (bottom) for **a** FME, **b** FME-Lasso and **c** iFME model versus the number of experts K , fitted on a randomly taken dataset in scenario $S1$. Here, the square points correspond to the highest values

weather stations in Canada, and their corresponding average annual precipitation (in log scale). The weather stations are located in 4 climate zones: Atlantic, Pacific, Continental and Arctic (Fig. 4a). In this dataset, presented in Fig. 4b, the noisy functional predictors $U_i(\cdot)$ are the curves of 365 averaged daily temperature measurements, the scalar responses Y_i are the corresponding total precipitation at each station

i , during the year, for 35 stations. Its station climate zone is taken as a cluster label (the cluster label $Z_i \in \{1, \dots, 4\}$).

The aim here is to use the daily temperature curves (functional predictors) to predict the precipitations (scalar responses) at each station. Moreover, in addition to predicting the precipitation values, we are interested in clustering the temperature curves (therefore the stations), as well as

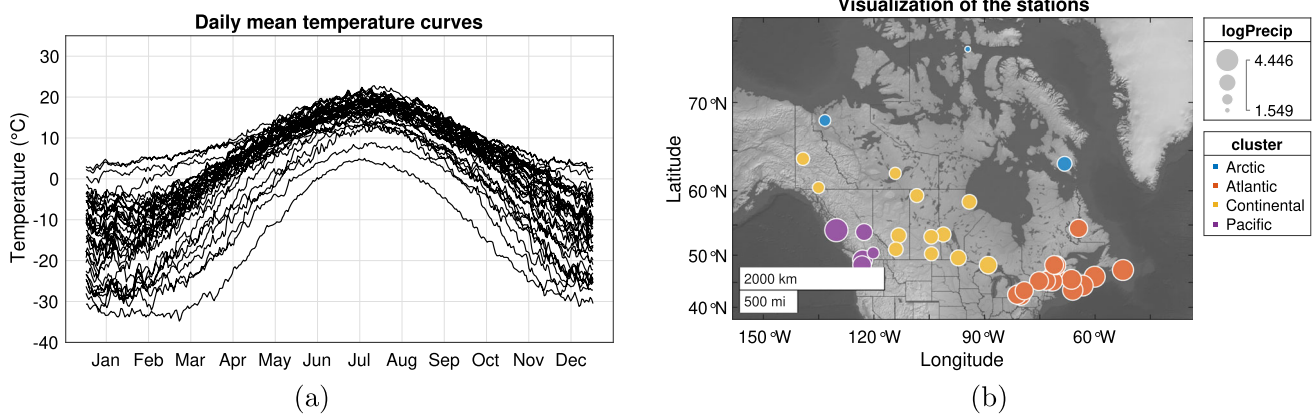


Fig. 4 **a** 35 daily mean temperature measurement curves; **b** geographical visualization of the stations, in which the sizes of the bubbles corresponds to their log of precipitation values and the colors correspond to their climate regions

Table 7 7-fold cross-validated correlation (Corr), sum of squared errors (SSE) and relative prediction error (RPE) of predictions on Canadian weather data

	Corr	SSE	RPE
FME	0.690	14.410	0.290
FME-Lasso	0.632	2.011	0.045
iFME	0.944	0.582	0.012

identifying the periods of time of the year that have effect on prediction for each group of curves.

Firstly, in order to assess the prediction performance of the FME, FME-Lasso and iFME models on this dataset, we implement the models by selecting the tuning parameters, including the number of expert components K in the set $\{1, 2, 3, 4, 5, 6\}$, by maximizing the modified BIC criterion, given its performance as shown in the simulation study. We report in Table 7 the results in terms of correlations, *sum of squared errors* (SSE) and relative prediction errors.

According to the obtained results, iFME provides the best results w.r.t. all the criteria. The cross-validated RPE provided by iFME is only of 1.2%, the next is FME-Lasso with 4.5%, while the FME model has the worst RPE value. Note that in James et al. (2009), the authors applied their proposed model to Canadian weather data and obtained a 10-fold cross-validated SSE of 4.77. Clearly, with the smaller cross-validated SSEs, the FME-Lasso and iFME models significantly improve the prediction.

Figure 5 shows the obtained results with the FME, FME-Lasso and iFME models, with $K = 4$, and with the derivatives $d_1 = 0, d_2 = 3$ for the iFME model. The estimated experts functions and gating functions are presented in the two top panels of the curve, while the clustering for the temperature curves and the stations are shown in the two bottom panels.

As we can see, all models provide reasonable clustering for the curves which may be corresponding to different complicated underlying meteorological forecasting mechanisms. Particularly, although not using any spatial information, merely temperature information, the obtained clustering for the stations is also comparable with the original labels of the stations. For example, the FME and FME-Lasso models identify exactly the Arctic stations, while iFME identifies exactly the Pacific stations, and all of the models provide reasonable spatially organized clusters. However, what is interesting here is the shape of the expert and gating functions $\hat{\alpha}(\cdot)$'s and $\hat{\beta}(\cdot)$'s obtained by the models. While FME and FME-Lasso gave less interpretable estimations, iFME appears to give, as it can be seen in the two top-right panels, piece-wise zero-valued and possibly quadratic estimated functions, which have a wide range of flat relationship from January to February and from June to September.

Motivated by the above results, on direction of identifying the periods of time of the year that truly have an effect on prediction, we implement the iFME model with $K = 2$, and the derivative levels d_1 and d_2 are set to be the zeroth and the third derivatives. The reason for the choices of d_1 and d_2 is that the penalization on the zeroth derivative would take into account zero ranges in the expert and gating functions, while the penalization on the third derivative, would take into account the smoothness for the changes between the periods of times in the functions. The obtained results are shown in Fig. 6.

As we can see, there are differences in the prediction mechanisms of the models between the northern stations and the southern stations. At southern stations, the obtained $\hat{\beta}_2(t)$ shows that there is a negative relationship in the spring and a positive relationship in the late fall, but no relationship in the remaining period of the year. This phenomenon is concordant with the result obtained in James et al. (2009), where the authors obtained the same relationships in the same periods of time. However, our iFME model additionally suggests that,

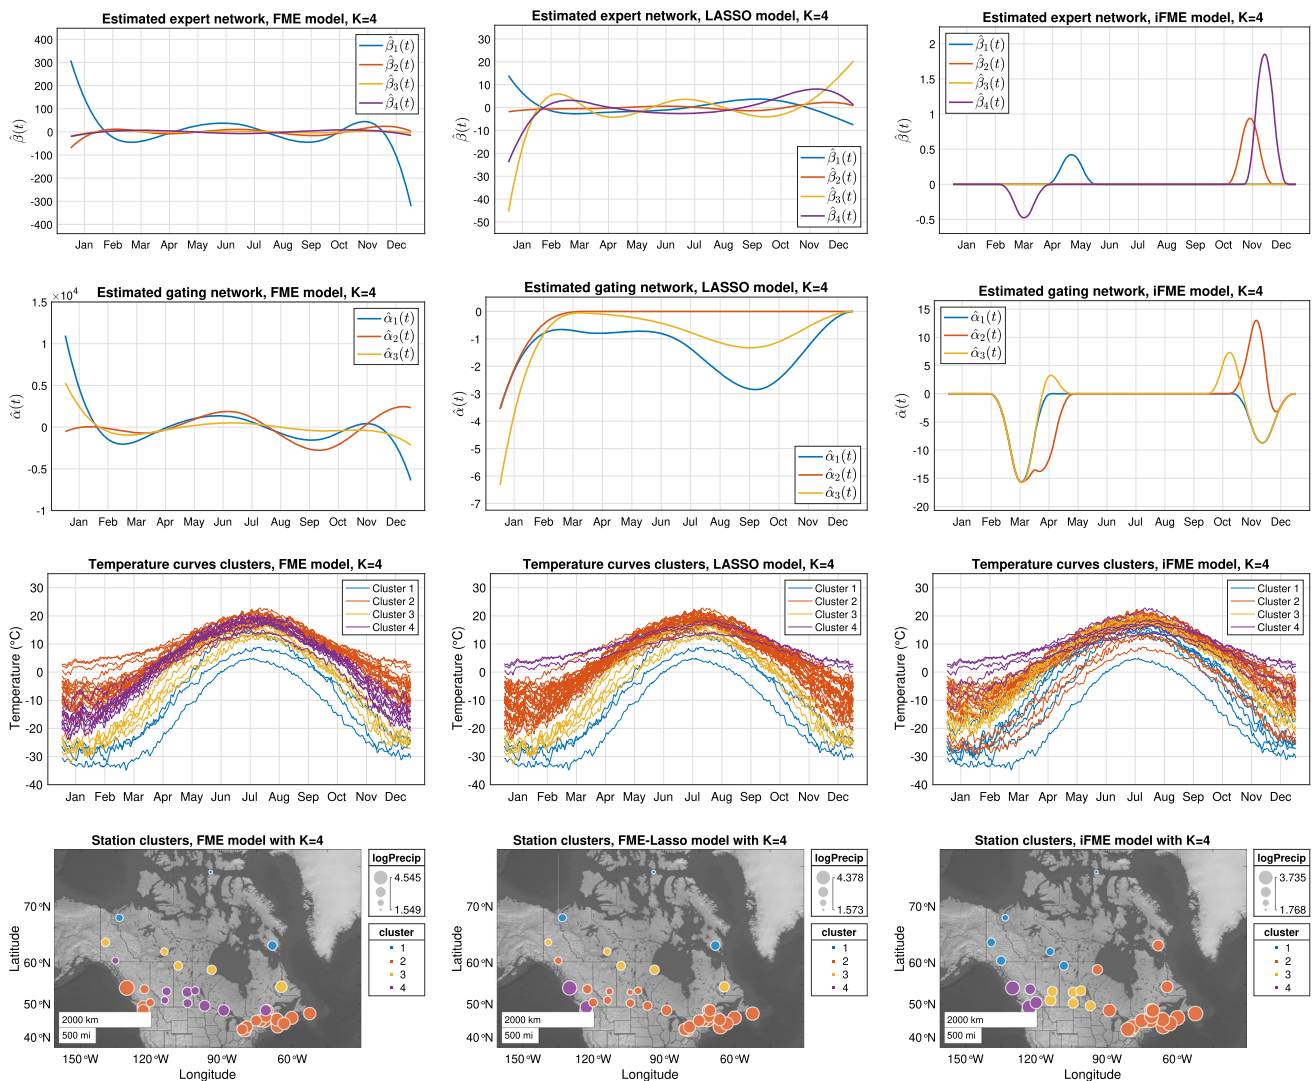


Fig. 5 Results obtained by FME (left panels), FME-Lasso (middle panels) and iFME (right panels) on Canadian weather data with $K = 4$. For each column, the panels are respectively the estimated functional

experts network, estimated functional gating network, estimated clusters of the temperature curves and estimated clusters of the stations

at the northern stations, the relationship between temperature and precipitation may differ from that of southern stations. This may be explained by the differences in mean temperatures and climatic characteristics between the two regions.

Finally, Fig. 7 displays the values of modified BIC for varying number of expert component for the proposed models on Canadian weather data. According to these values, FME-Lasso and iFME select $K = 2$, while FME selects $K = 3$.

4.3.2 Diffusion tensor imaging data for multiple sclerosis subjects

We now apply our proposed models to the diffusion tensor imaging (DTI) data for subjects with multiple sclerosis (MS),

discussed in Goldsmith et al. (2012). The data come from a longitudinal study investigating the cerebral white matter tracts of subjects with multiple sclerosis, recruited from an outpatient neurology clinic and healthy controls. We are interested in the underlying relationship between the fractional anisotropy profile (FAP) from the corpus callosum and the paced auditory serial addition test (PASAT) score, which is a commonly used examination of cognitive function affected by MS. The FAP curves are derived from DTI data, which are obtained by a Magnetic Resonance Imaging (MRI) scanner. Each curve is recorded at 93 locations along the corpus callosum. The PASAT score is the number of correct answers out of 60 questions, and thus ranges from 0 to 60. In our context, the FAP curves serve as the noisy functional predictors $U_i(\cdot)$ and the PASAT scores serve as

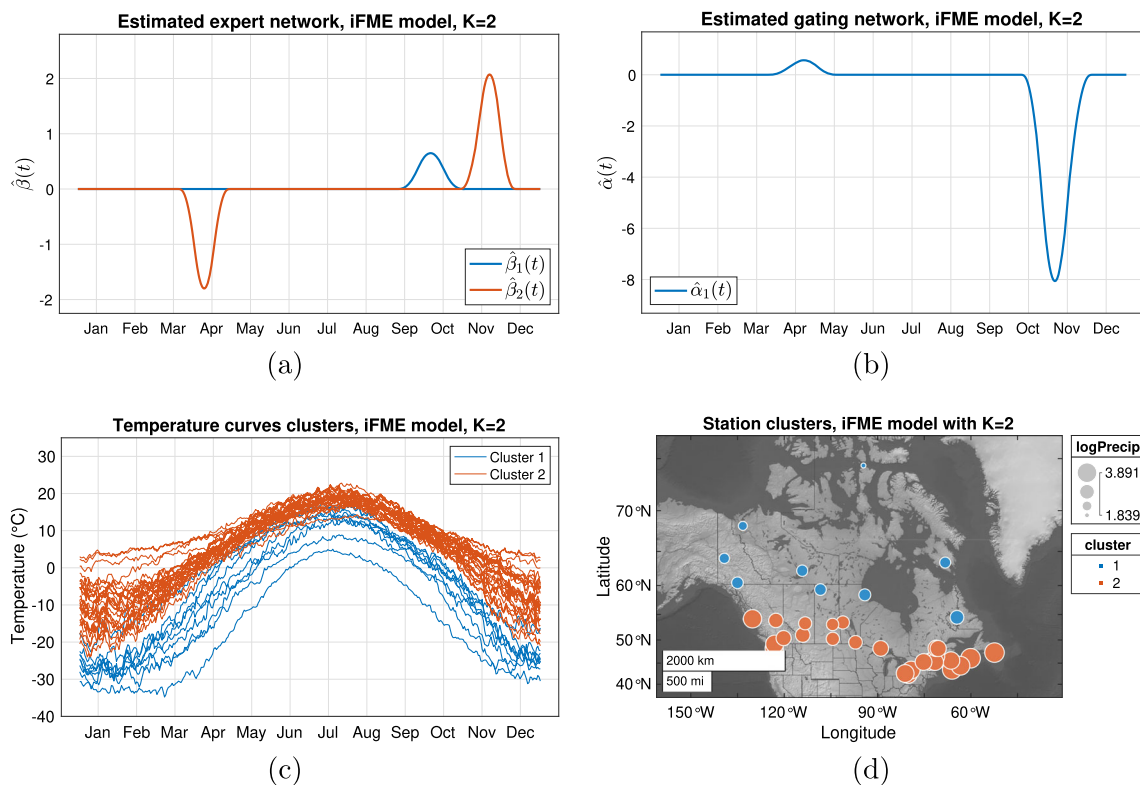


Fig. 6 Results obtained by iFME model with $K = 2, d_1 = 0, d_2 = 3$: **a** estimated functional expert network, **b** estimated functional gating network, **c** estimated clusters of the temperature curves, and **d** estimated clusters of the stations




Fig. 7 Values of modified BIC for **a** FME, **b** FME-Lasso and **c** iFME model, versus the number of experts K , fitted on Canadian weather data. The square points correspond to highest values. Here, iFME is implemented with $d_1 = 0, d_2 = 3$ and $\rho = \varrho = 100$

the scalar responses Y_i . So this dataset consists of $n = 99$ pairs $(U_i(\cdot), Y_i), i \in [n]$, with each $U_i(\cdot)$ contains $m = 93$ fractional anisotropy values. Figure 9a shows all the predictors (FAP curves), and Fig. 8 (most-right panels) shows the responses (PASAT Scores).

In Ciarleglio and Ogden (2016), the authors applied their Wavelet-based functional mixture regression (WBFMR) model with two components to this dataset, and observed that there is one group in which there is no association between the FAP and the PASAT score for those subjects belonging to it. We accordingly fix $K = 2$ in our models. Figure 8 displays the obtained results for each of the three models, and Fig. 9a shows the functional predictors FAP curves clus-

tered with the iFME model. In this implementation, we tried iFME model with two different combinations of d_1 and d_2 : $(d_1, d_2) = (0, 2)$ and $(d_1, d_2) = (0, 3)$. As expected, when d_2 is the second derivative, the reconstructed parameter functions are piecewise zero and linear, while when d_2 is the third derivative, the reconstructed functions have smooth changes along the tract location.

In Fig. 8, we have the three following observations. First, as it can be observed in Fig. 8 right-panel, all models give a threshold of 50 that clusters the PASAT scores, this is the same as the threshold observed in Ciarleglio and Ogden (2016). Second, the absolute values of the coefficient functions $\hat{\beta}_2(t)$'s are significantly smaller than those of $\hat{\beta}_1(t)$'s,

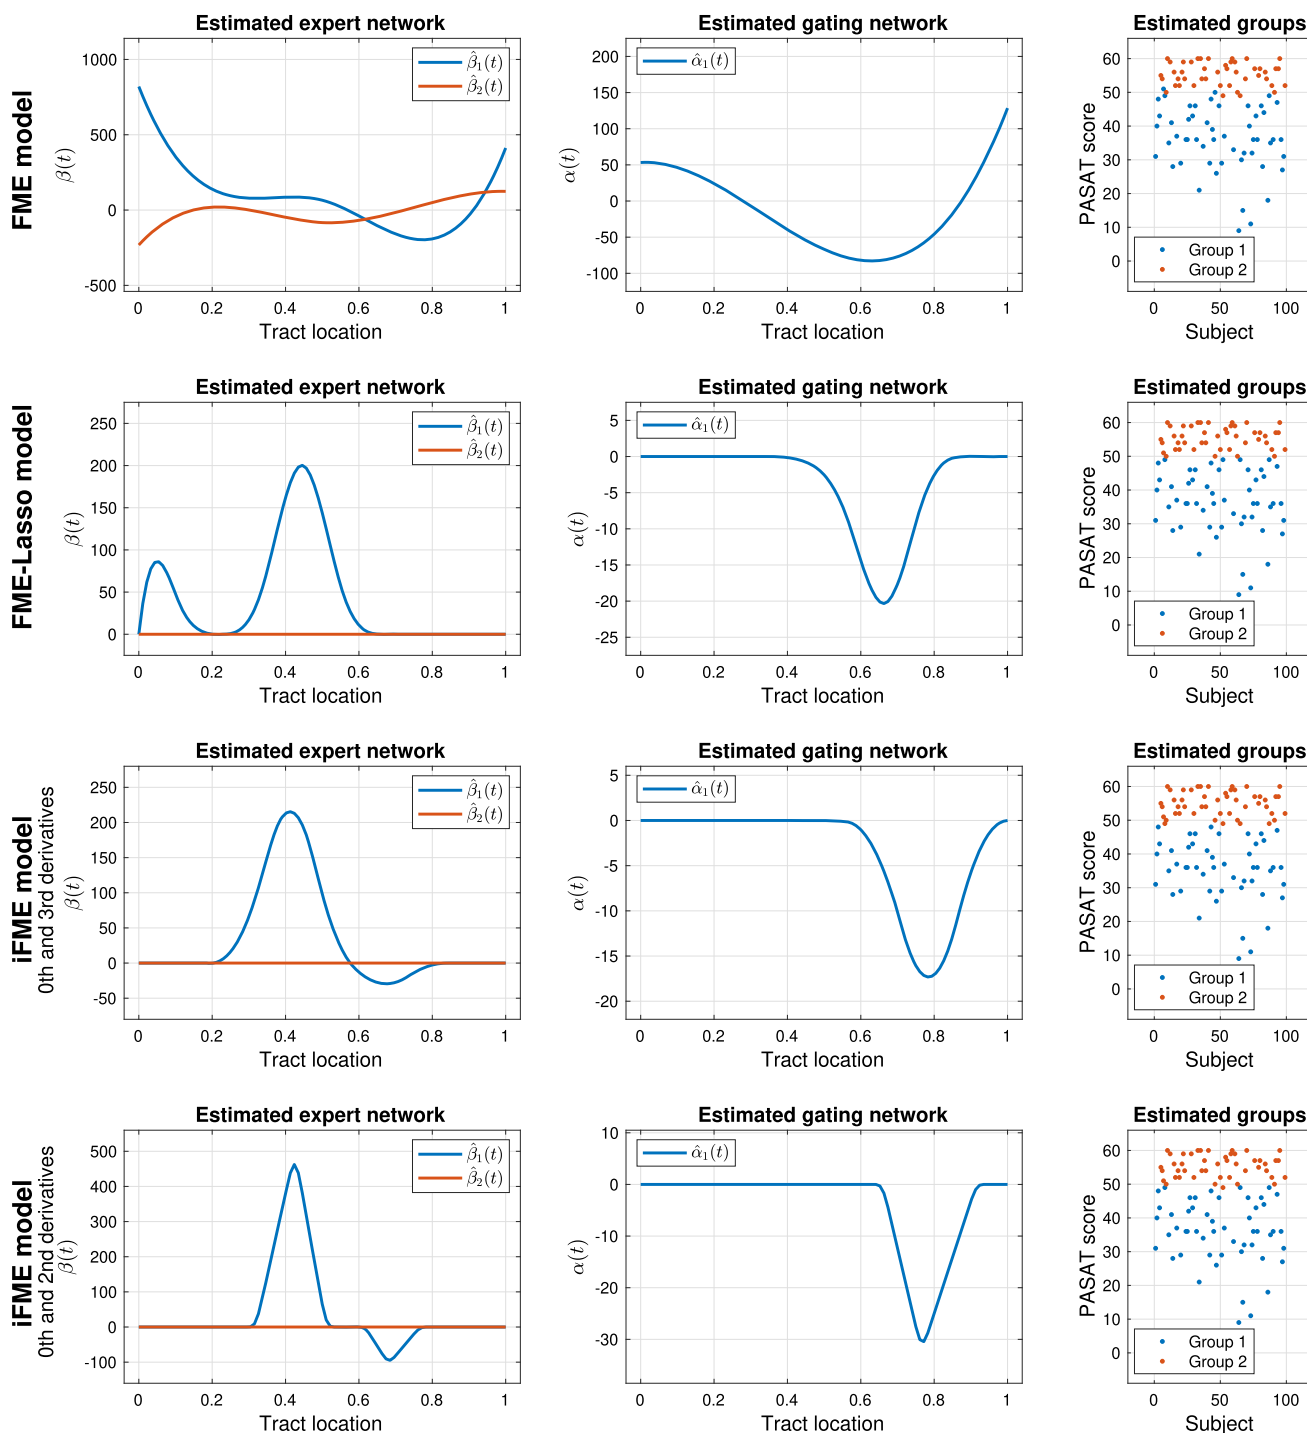
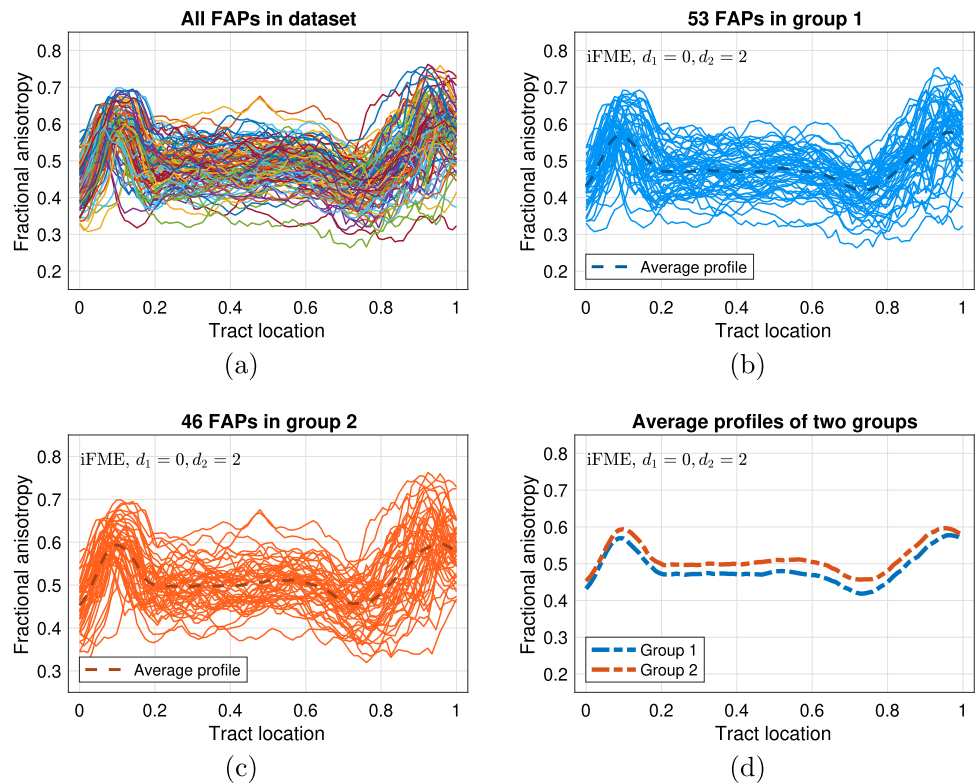


Fig. 8 The estimated expert and gating coefficient functions, the estimated groups of the PASAT scores, resulted by FME, FME-Lasso and iFME models with $K = 2$ for the DTI dataset. For iFME model, the

upper is implemented with penalization on the zeroth and third derivatives, while the lower is with penalized zeroth and second derivatives

Fig. 9 DTI data with **a** FAP curves of all subjects, **b** Cluster 1 and **c** Cluster 2, obtained by iFME model with $d_1 = 0$, $d_2 = 2$, and **d** the point-wise average of the curves in each of the two clusters



this is again the same with the result obtained by the WBFMR model. Third, when $\hat{\beta}_2(t)$ is estimated as zero in the FME-Lasso and iFME models, the shape of $\hat{\beta}_1(t)$ is almost the same as the shape obtained in Ciarleglio and Ogden (2016), particularly, the peak at around the tract location of 0.42. These confirm the underlying relationship between the fractional anisotropy and the cognitive function: higher fractional anisotropy values between the locations about 0.2 to 0.7 results in higher PASAT scores for subjects in Group 1. The clustering of the FAP curves, resulted by the iFME model with $d_1 = 0$, $d_2 = 2$, is shown in Fig. 9 (b)-(d).

Next, to compare with Ciarleglio and Ogden (2016), we investigate the prediction performance of the proposed models with respect to the leave-one-out cross validated relative prediction errors defined by $CVRPE = \frac{\sum_{i=1}^n (y_i - \hat{y}_i^{(-i)})^2}{\sum_{i=1}^n y_i^2}$, where y_i is the true score for subject i and $\hat{y}_i^{(-i)}$ is the score predicted by the model fit on data without subject i . In this implementation, we keep fixing $K = 2$ and select the other tuning parameters by maximizing the modified BIC. The CVRPEs corresponding to the models are provided in Table 8. Note that, for comparison, in Ciarleglio and Ogden (2016), the CVRPE of their WBFMR model is 0.0315 and of the wavelet based functional linear model (FLM) is 0.0723.

Finally, we present in Fig. 10 the selection of the number of experts K with modified BIC. In this case, FME and FME-Lasso select $K = 2$, and iFME selects $K = 4$.

Table 8 CVRPEs of the models on the DTI data

	CVRPE
FME	0.0273
FME-Lasso	0.0280
iFME (with 0th and 3rd derivatives)	0.0271
iFME (with 0th and 2nd derivatives)	0.0267

5 Conclusion and discussion

The first algorithm for mixtures-of-experts constructed upon functional predictors is presented in this paper. Beside the classic maximum likelihood parameter estimation, we proposed two other regularized versions that allow sparse and notably interpretable solutions, by regularizing in particular the derivatives of the underlying functional parameters of the model, after projecting onto a set of continuous basis functions. The performances of the proposed approaches are evaluated in data prediction and clustering via experiments involving simulated and two real-world functional data.

The presented FME models can be extended in different ways. First direct extensions of the modeling framework presented here can be considered with categorical response, to perform supervised classification with functional predictors, or with vector response, to perform multivariate functional regression. Then, it may be interesting to consider the exten-

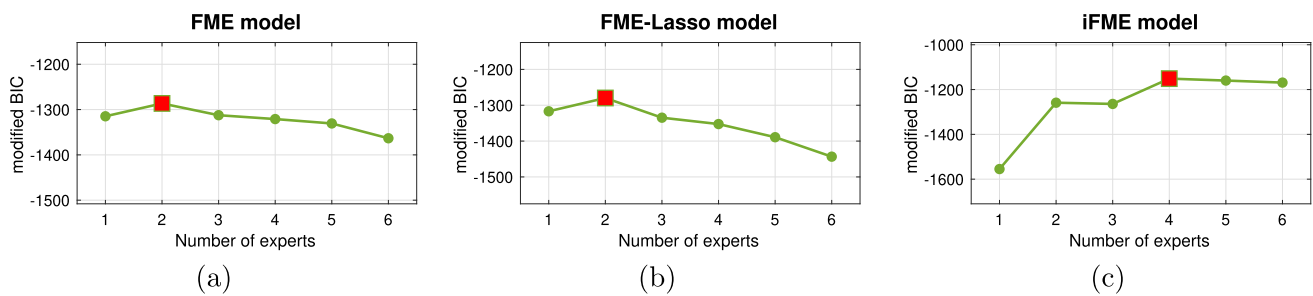


Fig. 10 Values of modified BIC for **a** FME, **b** FME-Lasso and **c** iFME, versus the number of experts K , fitted on DTI data. The square points correspond to highest values

sion of the FME model to setting involving vector (or scalar) predictors and functional responses (Chiou et al. 2004).

Another extension, which we intend also to investigate in the future, concerns the case when we observe pairs of functional data, i.e., a sample of n functional data pairs $\{X_i(t), Y_i(u)\}_{i=1}^n$, $t \in \mathcal{T} \subset \mathbb{R}$, $u \in \mathcal{U} \subset \mathbb{R}$, where $Y_i(\cdot)$ is a functional response, explained by a functional predictor $X_i(\cdot)$. The modeling with such FME extension then takes the form $Y_i(u) = \beta_{z_i,0}(u) + \int_{\mathcal{T}} X_i(t) \beta_{z_i}(t, u) dt + \varepsilon_i(u)$, to explain the functional response Y by the functional predictor X via the unknown discrete variable z . The particularity with this model is that, for the clustering, as well as for the prediction, we model the relation between Y at any time u and the entire curve of X , or the entire curve of each variable X_{ij} in the case of multivariate functional predictor X_i .

Since we look for sparse estimates in the space of the derivatives of the functional parameters, and provided the resulting estimates with highly structured shapes, we believe that using the fused LASSO (Tibshirani et al. 2005) rather than the LASSO could be expected to take accommodate more the targeted sparsity. The choice of the regularization constants for the derivatives is more about the targeted penalization magnitude (i.e., very large or very small) rather than the exact values. Typically, in the case one has no idea about the shape of the functional experts, but still want them to be interpretable, then it is practical to consider the zeroth and third derivatives (since they produce smooth changes), with the grid for ρ containing both very small and very large values. The same for the functional gating network. Furthermore, if we believe that the relation of the response on the predictor is very sparse, except over some small regions where they are linearly related, then we will penalize the zeroth and second derivatives of the functional experts, and we will put very small regularization on the second derivative. This could produce curves such as those in bottom panels of Fig. 8. For a more general tuning of the regularization constant values, we recommend performing a complete cross-validation study to select the best values of the regularization constants.

Acknowledgements This research was partly supported by an Ethel Raybould Fellowship, University of Queensland (FC), the Australian Research Council Grant Number DP180101192 (GJM), the French National Research Agency ANR Grant SMILES ANR-18-CE40-0014 (FC and NTP), Région Normandie grant RIN ASterICs (FC and VHH), Vietnam National University - Ho Chi Minh City (VNU- HCM) Grant Number C2022-18-01 (VHH), and the exploratory research facility “EXPLO” at SystemX (FC).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A EM for the FME model

The FME model can be fitted by iteratively maximizing the observed-data log-likelihood (14) iteratively via the EM algorithm. For FME, the EM takes the following form. The complete-data log-likelihood upon which the EM principle is constructed is defined by

$$\log L_c(\Psi) = \sum_{i=1}^n \sum_{k=1}^K Z_{ik} \log \left[\pi_k(\mathbf{r}_i; \xi) \phi(y_i; \beta_{k,0} + \boldsymbol{\eta}_k^\top \mathbf{x}_i, \sigma_k^2) \right], \quad (42)$$

Z_{ik} being an indicator binary-valued variable such that $Z_{ik} = 1$ if $Z_i = k$ (i.e., if the i th pair (x_i, y_i) is generated from the k th expert component) and $Z_{ik} = 0$, otherwise.

E-step This step computes at each EM iteration s the expectation of the complete-data log-likelihood (42), given the observed data \mathcal{D} , and the current parameter vector $\Psi^{(s)}$:

$$\begin{aligned}
 Q(\Psi; \Psi^{(s)}) &= \mathbb{E} \left[\log L_c(\Psi) | \mathcal{D}; \Psi^{(s)} \right] \\
 &= \sum_{i=1}^n \sum_{k=1}^K \tau_{ik}^{(s)} \log \left[\pi_k(\mathbf{r}_i; \xi) \phi(y_i; \beta_{k,0} + \eta_k^\top \mathbf{x}_i, \sigma_k^2) \right],
 \end{aligned}
 \tag{43}$$

where $\tau_{ik}^{(s)} = \mathbb{P}(Z_i = k | y_i, u_i(\cdot); \Psi^{(s)})$ is the conditional probability that the observed pair $(u_i(\cdot), y_i)$ is generated by the k th expert. This step therefore only requires the computation of the conditional probabilities $\tau_{ik}^{(s)}$ as defined by (15).

M-step This step updates the value of the parameter vector Ψ by maximizing the Q -function (43) with respect to Ψ , that is $\Psi^{(s+1)} = \arg \max_{\Psi} Q(\Psi; \Psi^{(s)})$, via separate maximizations w.r.t. the gating network parameters, and the experts network parameters as follows.

Updating the the gating network parameters Updating the the gating network’s parameters ξ consists of maximizing w.r.t. ξ the following function

$$\begin{aligned}
 Q(\xi; \Psi^{(s)}) &= \sum_{i=1}^n \sum_{k=1}^K \tau_{ik}^{(s)} \log \pi_k(\mathbf{r}_i; \xi) \\
 &= \sum_{i=1}^n \left[\sum_{k=1}^{K-1} \tau_{ik}^{(s)} (\alpha_{k,0} + \xi_k^\top \mathbf{r}_i) - \log \left(1 + \sum_{k'=1}^{K-1} \exp\{\alpha_{k',0} + \xi_{k'}^\top \mathbf{r}_i\} \right) \right].
 \end{aligned}
 \tag{44}$$

This consists of a weighted multinomial logistic problem for which there is no closed-form solution. This can be performed by the Newton-Raphson (NR) algorithm which iteratively maximizes (44) according to the procedure (16).

Let us denote by ξ_1, \dots, ξ_{K-1} the parameter vectors $(\alpha_{1,0}, \xi_1^\top)^\top, \dots, (\alpha_{K-1,0}, \xi_{K-1}^\top)^\top$. Since there are $K - 1$ parameter vectors to be estimated, the Hessian matrix $H(\xi; \Psi^{(s)})$ is a block-matrix, consists of $(K - 1) \times (K - 1)$ blocks, in which each block $H_{k\ell}(\xi; \Psi^{(s)})$, for $k, \ell \in [K - 1]$, is given by:

$$\begin{aligned}
 H_{k\ell}(\xi; \Psi^{(s)}) &= \frac{\partial^2 Q(\xi; \Psi^{(s)})}{\partial \xi_k \partial \xi_\ell^\top} \\
 &= - \sum_{i=1}^n \pi_k(\mathbf{r}_i; \xi^{(t)}) \left[\delta_{k\ell} - \pi_\ell(\mathbf{r}_i; \xi^{(t)}) \right] \mathbf{r}_i \mathbf{r}_i^\top,
 \end{aligned}$$

where $\delta_{k\ell}$ is the Kronecker symbol ($\delta_{k\ell} = 1$ if $k = \ell$, 0 otherwise). The gradient vector consists of $K - 1$ gradients corresponding to the vectors ξ_k , for $k \in [K - 1]$, and is given by

$$g(\xi; \Psi^{(s)}) = \frac{\partial Q(\xi; \Psi^{(s)})}{\partial \xi} = \left[g_1(\xi^{(t)}), \dots, g_{K-1}(\xi^{(t)}) \right]^\top,$$

where, for $k \in [K - 1]$, $g_k(\xi^{(t)}) = \frac{\partial Q(\xi; \Psi^{(s)})}{\partial \xi_k} = \sum_{i=1}^n \left[\tau_{ik}^{(s)} - \pi_k(\xi^{(t)}; \mathbf{r}_i) \right] \mathbf{r}_i^\top$.

Updating the experts network parameters Updating the experts network’s parameters θ_k consists of maximizing the function $Q(\theta_k; \Psi^{(s)})$ given by

$$\begin{aligned}
 Q(\theta_k; \Psi^{(s)}) &= \sum_{i=1}^n \tau_{ik}^{(s)} \log \phi(y_i; \beta_{k,0} + \eta_k^\top \mathbf{x}_i, \sigma_k^2) \\
 &= - \frac{1}{2\sigma_k^2} \sum_{i=1}^n \tau_{ik}^{(s)} \left[y_i - (\beta_{k,0} + \eta_k^\top \mathbf{x}_i) \right]^2 - \frac{n}{2} \log(2\pi\sigma_k^2).
 \end{aligned}$$

Thus, updating $\theta_k = (\beta_{k,0}, \eta_k^\top, \sigma_k^2)^\top$, consists of a weighted Gaussian regression problem where the weights are the conditional memberships $\tau_{ik}^{(s)}$, and the updates are given by (17).

B EM-Lasso for ℓ_1 -regularized MLE of the FME model

The EM-Lasso algorithm for the maximization of (18) firstly requires the construction of the penalized complete-data log-likelihood

$$\mathcal{L}_c(\Psi) = \log L_c(\Psi) - \text{Pen}_{\lambda, \chi}(\Psi)
 \tag{45}$$

where $\log L_c(\Psi)$ is the non-regularized complete-data log-likelihood log-likelihood defined by (42). Thus, the EM algorithm for the FME model is implemented as follows. After starting with an initial solution $\Psi^{(0)}$, it alternates between the two following steps, until convergence (when there is no longer a significant change in the values of the penalized log-likelihood (18)).

E-step This step computes the expectation of the complete-data log-likelihood (45), given the observed data \mathcal{D} , using the current parameter vector $\Psi^{(s)}$:

$$\begin{aligned}
 \mathcal{Q}_{\lambda, \chi}(\Psi; \Psi^{(s)}) &= \mathbb{E} \left[\mathcal{L}_c(\Psi) | \mathcal{D}; \Psi^{(s)} \right] \\
 &= Q(\Psi; \Psi^{(s)}) - \text{Pen}_{\lambda, \chi}(\Psi),
 \end{aligned}
 \tag{46}$$

which only requires the computation of the posterior probabilities of component membership $\tau_{ik}^{(s)}$ ($i \in [n]$), for each of the K experts as defined by (15).

M-step This step updates the value of the parameter vector Ψ by maximizing the Q -function (46) with respect to Ψ , that is, by computing the parameter vector update

$$\Psi^{(s+1)} = \arg \max_{\Psi} \mathcal{Q}_{\lambda, \chi}(\Psi; \Psi^{(s)}).
 \tag{47}$$

The maximization is performed by separate maximizations w.r.t. the gating network parameters and the experts network parameters.

B.1 Updating the gating network parameters

Updating the gating network parameters at the s th EM iteration consists of maximizing the following function

$$Q_\chi(\xi; \Psi^{(s)}) = Q(\xi; \Psi^{(s)}) - \chi \sum_{k=1}^{K-1} \|\zeta_k\|_1,$$

with

$$Q(\xi; \Psi^{(s)}) = \sum_{i=1}^n \left[\sum_{k=1}^{K-1} \tau_{ik}^{(s)} (\alpha_{k,0} + \zeta_k^\top \mathbf{r}_i) - \log \left(1 + \sum_{k'=1}^{K-1} \exp\{\alpha_{k',0} + \zeta_{k'}^\top \mathbf{r}_i\} \right) \right]$$

where $\xi = (\alpha_{1,0}, \zeta_1^\top, \dots, \alpha_{K-1,0}, \zeta_{K-1}^\top)^\top \in \mathbb{R}^{(q+1)(K-1)}$ is the gating network parameter vector and $\Psi^{(s)}$ is the current estimation of model's parameters. One can see this is equivalent to solving a weighted regularized multinomial logistic problem for which $Q_\chi(\xi; \Psi^{(s)})$ is its penalized log-likelihood. There is no closed-form solution for this kind of problem. We then use an iterative optimization algorithm to seek for a maximizer of $Q_\chi(\xi; \Psi^{(s)})$, i.e., an update for the parameters of the gating network. The idea is to update only a single gate at a time, while fixing other gate's parameters to their previous estimates. Again, to update that single gate, we only update one component at a time, while fixing the other components to their previous values. This procedure for updating the gating network parameters is supported by the methodology of coordinate ascent algorithm: if the objective function consists of a concave, differentiable function and a sum of concave functions then the maximizer can be achieved by iteratively maximizing with respect to each coordinate direction at a time.

Coordinate ascent for updating the gating network Suppose at the s th EM iteration, we wish to update the gates one by one such that it maximizes $Q_\chi(\xi; \Psi^{(s)})$. To do that, we create an outer loop, indexed by t , which cycles over the gates. For each single gate, say gate k , we partially approximate the smooth part of $Q_\chi(\xi; \Psi^{(s)})$ with respect to $(\alpha_{k,0}, \zeta_k)$ at $\xi^{(t)}$, then optimize the obtained objective function (with respect to $(\alpha_{k,0}, \zeta_k)$) by solving a penalized weighted least square problem using coordinate ascent algorithm. Note that $\xi^{(t)}$ denotes the current value of ξ at the iteration t th of the outer loop, while $\xi^{(s)}$ is the value of ξ before entering the outer loop.

In particular, using Taylor expansion, one has a quadratic approximation for smooth part of $Q_\chi(\xi; \Psi^{(s)})$ with respect

to $(\alpha_{k,0}, \zeta_k)$ at $\xi^{(t)}$ given by

$$l_k(\alpha_{k,0}, \zeta_k) = -\frac{1}{2} \sum_{i=1}^n w_{ik} (c_{ik} - \alpha_{k,0} - \mathbf{r}_i^\top \zeta_k)^2 + C(\xi^{(t)}),$$

where

$$w_{ik} = \pi_k(\xi^{(t)}; \mathbf{r}_i) \left[1 - \pi_k(\xi^{(t)}; \mathbf{r}_i) \right], \quad (\text{weights})$$

$$c_{ik} = \alpha_{k,0}^{(t)} + \mathbf{r}_i^\top \zeta_k^{(t)} + \frac{\tau_{ik}^{(s)} - \pi_k(\xi^{(t)}; \mathbf{r}_i)}{w_{ik}},$$

(working response)

and $C(\xi^{(t)})$ is a function of $\xi^{(t)}$. After calculating the partial quadratic approximation $l_k(\alpha_{k,0}, \zeta_k)$ about the current estimator $\xi^{(t)}$, we then solve the following penalized weighted least square problem

$$\max_{(\alpha_{k,0}, \zeta_k)} l_k(\alpha_{k,0}, \zeta_k) - \chi \|\zeta_k\|_1, \quad \chi > 0, \tag{48}$$

to obtain an update for the parameters of gate k .

As mentioned above, this problem could be solved by coordinate ascent algorithm. This means we will create an inner loop, indexed by m , cycles over the components of $(\alpha_{k,0}, \zeta_k)$ and update them one by one until the objective function of (48) does not gain any significant increase. For each $j \in [q]$, using the soft-thresholding operator (see Hastie et al. (2015), sec. 5.4), one can obtain the closed form update for ζ_{kj} as follows

$$\zeta_{kj}^{(m+1)} = \frac{\mathcal{S}_\chi \left(\sum_{i=1}^n w_{ik} \mathbf{r}_{ij} (c_{ik} - \tilde{c}_{ikj}^{(m)}) \right)}{\sum_{i=1}^n w_{ik} \mathbf{r}_{ij}^2},$$

in which $\tilde{c}_{ikj}^{(m)} = \alpha_{k,0}^{(m)} + \mathbf{r}_i^\top \zeta_k^{(m)} - \zeta_{kj}^{(m)} \mathbf{r}_{ij}$ is the fitted value excluding the contribution from \mathbf{r}_{ij} , $\mathcal{S}_\chi(\cdot)$ is a soft-thresholding operator defined by $\mathcal{S}_\chi(u) = \text{sign}(u)(|u| - \chi)_+$ and $(v)_+$ a shorthand for $\max\{v, 0\}$. Note that at each iteration of the inner loop, only one component is updated while the others are kept to their previous values, that means $\zeta_{kh}^{(m+1)} = \zeta_{kh}^{(m)}$ for all $h \neq j$. For $\alpha_{k,0}$, the closed-form update is given by

$$\alpha_{k,0}^{(m+1)} = \frac{\sum_{i=1}^n w_{ik} (c_{ik} - \mathbf{r}_i^\top \zeta_k^{(m+1)})}{\sum_{i=1}^n w_{ik}}.$$

Once the inner loop converges, the new values of $(\alpha_{k,0}, \zeta_k)$ are used for the updating procedure of the next gate. When all the gates have their new values, i.e., after $K - 1$ inner loops, we perform a backtracking line search before actually updating the gating network's parameters for the next t -indexed iteration. More precisely, the update is $\xi^{(t+1)} = (1 - \nu)\xi^{(t)} + \nu\bar{\xi}^{(t)}$ where $\bar{\xi}^{(t)}$ is the output after $K - 1$

inner loops and ν is backtrackingly determined to ensure $\mathcal{Q}_\chi(\xi^{(t+1)}; \Psi^{(s)}) > \mathcal{Q}_\chi(\xi^{(t)}; \Psi^{(s)})$.

We keep running the t -indexed loop until convergence, i.e., when there is no significant relative variation in $\mathcal{Q}_\chi(\xi; \Psi^{(s)})$. Once $\alpha_{k,0}$ and ζ_{kj} reach their optimal values $\tilde{\alpha}_{k,0}$ and $\tilde{\zeta}_{kj}$ for all $k \in [K - 1], j \in [q]$, the update for gating network's parameters is then $\xi^{(s+1)} = (\tilde{\alpha}_{1,0}, \tilde{\zeta}_1^\top, \dots, \tilde{\alpha}_{K-1,0}, \tilde{\zeta}_{K-1}^\top)^\top$.

B.2 Updating the experts network parameters

The maximization step for updating the expert parameters θ_k consists of maximizing the function $\mathcal{Q}_\lambda(\theta_k; \Psi^{(s)})$ given by

$$\mathcal{Q}_\lambda(\theta_k; \Psi^{(s)}) = Q(\theta_k; \Psi^{(s)}) - \lambda \|\eta_k\|_1,$$

with

$$Q(\theta_k; \Psi^{(s)}) = -\frac{1}{2\sigma_k^2} \sum_{i=1}^n \tau_{ik}^{(s)} \left[y_i - (\beta_{k,0} + \eta_k^\top \mathbf{x}_i) \right]^2 - \frac{n}{2} \log(2\pi\sigma_k^2),$$

where $\theta_k = (\beta_{k,0}, \eta_k^\top, \sigma_k^2)^\top \in \mathbb{R}^{p+2}$ is the unknown vector and $\Psi^{(s)}$ is the current estimation of model's parameters. There is no closed-form solution for this optimization problem, we then solve it by an iterative optimization algorithm similarly to updating the gating network parameters. We first perform the update for $(\beta_{k,0}, \eta_k)$ while fixing σ_k^2 . This corresponds to solving a weighted LASSO problem where the weights are the the posterior experts memberships $\tau_{ik}^{(s)}$. Once $(\beta_{k,0}, \eta_k)$ has new value, the variance σ_k^2 is updated straightforwardly by the standard update of a weighted Gaussian regression.

More specifically, when σ_k^2 , the variance of expert k , is fixed to $\sigma_k^{2(s)}$, updating $(\beta_{k,0}, \eta_k)$ consists of solving the following weighted LASSO problem:

$$\max_{(\beta_{k,0}, \eta_k)} -\frac{1}{2\sigma_k^{2(s)}} \sum_{i=1}^n \tau_{ik}^{(s)} \left[y_i - (\beta_{k,0} + \eta_k^\top \mathbf{x}_i) \right]^2 - \frac{n}{2} \log(2\pi\sigma_k^{2(s)}) - \lambda \sum_{j=1}^q |\eta_{kj}|, \quad \lambda > 0, \tag{49}$$

which can be solved by coordinate ascent algorithm. For each $j \in [p]$, the closed-form update for η_{kj} is given by

$$\eta_{kj}^{(m+1)} = \frac{\mathcal{S}_{\lambda\sigma_k^{2(s)}} \left(\sum_{i=1}^n \tau_{ik}^{(s)} x_{ij} (y_i - \tilde{y}_{ij}^{(m)}) \right)}{\sum_{i=1}^n \tau_{ik}^{(s)} x_{ij}^2},$$

in which $\tilde{y}_{ij}^{(m)} = \beta_{k,0}^{(m)} + \mathbf{x}_i^\top \eta_k^{(m)} - \eta_{kj}^{(m)} x_{ij}$ is the fitted value excluding the contribution from x_{ij} and $\mathcal{S}_\chi(\cdot)$ is the soft-thresholding operator (see Hastie et al. (2015), sec. 5.4). Here m denotes the m th iteration of the coordinate ascent algorithm. The update for $\beta_{k,0}$ is

$$\beta_{k,0}^{(m+1)} = \frac{\sum_{i=1}^n \tau_{ik}^{(s)} (y_i - \mathbf{x}_i^\top \eta_k^{(m+1)})}{\sum_{i=1}^n \tau_{ik}^{(s)}}.$$

We keep updating the components of $(\beta_{k,0}, \eta_k)$ cyclically until the change in objective function of (49) is small enough. So, the update for $(\beta_{k,0}, \eta_k)$ in this EM iteration is then $(\beta_{k,0}^{(s+1)}, \eta_k^{(s+1)}) = (\beta_{k,0}^*, \eta_k^*)$ where the latter is the optimal solution of (49). Finally, the update for σ_k^2 is given by

$$\sigma_k^{2(s+1)} = \frac{\sum_{i=1}^n \tau_{ik}^{(s)} (y_i - \beta_{k,0}^{(s+1)} - \mathbf{x}_i^\top \eta_k^{(s+1)})^2}{\sum_{i=1}^n \tau_{ik}^{(s)}}.$$

Hence, the update for the value of parameters vector Ψ at M-step, i.e., the solution to problem (47), is $\Psi^{(s+1)} = (\xi^{(s+1)}, \theta_1^{(s+1)}, \dots, \theta_K^{(s+1)})$ where $\xi^{(s+1)}$ and $\theta_k^{(s+1)}, k \in [K]$, are solved by maximizing $\mathcal{Q}_\chi(\xi; \Psi^{(s)})$ and $\mathcal{Q}_\lambda(\theta_k; \Psi^{(s)})$, respectively, using the algorithms described above. The EM algorithm monotonically increases (18). Furthermore, the sequence of parameter estimates generated by the EM algorithm converges toward a local maximum of the log-likelihood function (Dempster et al. 1977; McLachlan and Krishnan 2008; Wu 1983).

C EM-iFME for updating iFME model parameters

C.1 Updating the gating network parameters

This section presents the using of Dantzig selector to solve problem (37). Let us simplify the subscript k in the notations and rewrite the problem under matrix form as follows

$$\max_{\tilde{\omega} \in \mathbb{R}^{2q+1}} -\frac{1}{2} \|\mathbf{c}_w - \mathbf{X}_w \tilde{\omega}\|_2^2 - \chi \|\Omega \tilde{\omega}\|_1$$

subject to $\mathbf{A} \tilde{\omega} = \mathbf{0}_q,$ (50)

where $\tilde{\omega} = (\alpha_0, \omega^{[d_1]^\top}, \omega^{[d_2]^\top})^\top$ is the unknown coefficients vector, $\mathbf{c}_w = (\sqrt{w_1}c_1, \dots, \sqrt{w_n}c_n)^\top$ is the weighted working response vector, $\mathbf{X}_w = [\sqrt{w}|\mathbf{S}_w| \mathbf{0}_{n \times q}] \in \mathbb{R}^{2q+1}$ is the weighted design matrix, $\Omega = \text{diag}(0, \mathbf{1}_q^\top, \varrho \mathbf{1}_q^\top)$ is the diagonal weighting matrix and $\mathbf{A} = [\mathbf{0}_q | \mathbf{A}_q^{[d_2]} \mathbf{A}_q^{[d_1]^{-1}} | -\mathbf{I}_q]$ is the constraints matrix. Here, $\sqrt{w} = (\sqrt{w_1}, \dots, \sqrt{w_n})^\top$, $\mathbf{S}_w = [\sqrt{w_1} \mathbf{s}_1, \dots, \sqrt{w_n} \mathbf{s}_n]^\top$, with \mathbf{s}_i are the design vectors (see (28)), $\mathbf{0}_{n \times q} \in \mathbb{R}^{n \times q}$ contains 0's, $\mathbf{0}_q \in \mathbb{R}^q$ contains

0's, $\mathbf{1}_q \in \mathbb{R}^q$ contains 1's and \mathbf{I}_q is the identity matrix in $\mathbb{R}^{q \times q}$. This problem can be viewed as the problem of finding a sparse solution via Lasso estimate for a linear regression model with constraints. Therefore, we can solve it alternatively by Dantzig selector estimate as the solution to the following problem

$$\begin{aligned} & \max_{\tilde{\omega} \in \mathbb{R}^{2q+1}} -\|\Omega\tilde{\omega}\|_1 \\ & \text{subject to } \begin{cases} |\mathbf{X}_w^\top(\mathbf{c}_w - \mathbf{X}_w\tilde{\omega})| \leq \chi \mathbf{1}_{2q+1}, \\ \mathbf{A}\tilde{\omega} = \mathbf{0}_q, \end{cases} \end{aligned}$$

where the absolute value operator is understood component-wise. By decomposing $\tilde{\omega}$ into its positive and negative parts, $\tilde{\omega} = \tilde{\omega}_+ - \tilde{\omega}_-$, the above problem becomes

$$\begin{aligned} & \max_{(\tilde{\omega}_+, \tilde{\omega}_-) \in \mathbb{R}^{4q+2}} -[0, \mathbf{1}_q^\top, \varrho \mathbf{1}_q^\top, 0, \mathbf{1}_q^\top, \varrho \mathbf{1}_q^\top] \begin{bmatrix} \tilde{\omega}_+ \\ \tilde{\omega}_- \end{bmatrix} \\ & \text{subject to } \begin{cases} \begin{bmatrix} \mathbf{X}_w^\top \mathbf{X}_w & -\mathbf{X}_w^\top \mathbf{X}_w \\ -\mathbf{X}_w^\top \mathbf{X}_w & \mathbf{X}_w^\top \mathbf{X}_w \end{bmatrix} \begin{bmatrix} \tilde{\omega}_+ \\ \tilde{\omega}_- \end{bmatrix} \leq \begin{bmatrix} \chi + \mathbf{X}_w^\top \mathbf{c}_w \\ \chi - \mathbf{X}_w^\top \mathbf{c}_w \end{bmatrix}, \\ \begin{bmatrix} \mathbf{A} & -\mathbf{A} \end{bmatrix} \begin{bmatrix} \tilde{\omega}_+ \\ \tilde{\omega}_- \end{bmatrix} = \mathbf{0}_{4q+2}, \\ \tilde{\omega}_+ \geq \mathbf{0}_{2q+1}, \quad \tilde{\omega}_- \geq \mathbf{0}_{2q+1}, \end{cases} \end{aligned} \tag{51}$$

which is a standard linear program with $4q + 2$ variables, therefore can be easily solved by available toolboxes for linear programming, e.g., Matlab's linprog function. Finally, the solution $\tilde{\omega}^* = (\alpha_0^*, \omega^{[d_1]^* \top}, \omega^{[d_2]^* \top})^\top$ to the original problem could be retrieved by the relation $\tilde{\omega}^* = \tilde{\omega}_+^* - \tilde{\omega}_-^*$ where $(\tilde{\omega}_+^{*\top}, \tilde{\omega}_-^{*\top})^\top$ is the solution of (51).

C.2 Updating the expert network parameters

This section presents how to solve problem (38). Firstly, we fix σ_k^2 to its previous estimate and perform an update for $(\beta_{k,0}, \boldsymbol{\gamma}_k^{[d_1]})$, it corresponds to solving a penalized weighted least square problem with constraints. From now on, let us simplify the subscript k in the notations and rewrite the problem under matrix form as follows

$$\begin{aligned} & \max_{\tilde{\boldsymbol{\gamma}} \in \mathbb{R}^{2p+1}} -\frac{1}{2} \|\mathbf{y}_{\sigma\tau} - \mathbf{X}_{\sigma\tau}\tilde{\boldsymbol{\gamma}}\|_2^2 - \lambda \|\Lambda\tilde{\boldsymbol{\gamma}}\|_1 - \frac{n_k}{2} \log(2\pi\sigma^2) \\ & \text{subject to } \mathbf{A}\tilde{\boldsymbol{\gamma}} = \mathbf{0}_p, \end{aligned} \tag{52}$$

where $\tilde{\boldsymbol{\gamma}} = (\beta_0, \boldsymbol{\gamma}^{[d_1]^\top}, \boldsymbol{\gamma}^{[d_2]^\top})^\top$ is the unknown coefficients vector, $\mathbf{y}_{\sigma\tau} = (\sigma\sqrt{\tau_1}y_1, \dots, \sigma\sqrt{\tau_n}y_n)^\top \in \mathbb{R}^n$ is the weighted response vector, $\mathbf{X}_{\sigma\tau} = \boldsymbol{\sigma} \odot [\sqrt{\boldsymbol{\tau}}|\mathbf{V}_\tau|\mathbf{0}_{n \times p}] \in \mathbb{R}^{n \times (2p+1)}$ is the weighted design matrix, $\Lambda = \text{diag}(0, \mathbf{1}_p^\top, \rho \mathbf{1}_p^\top)$ is the diagonal weighting matrix, $n_k = \sum_{i=1}^n \tau_{ik}$ and $\mathbf{A} = [\mathbf{0}_p | \mathbf{A}_p^{[d_2]} \mathbf{A}_p^{[d_1]-1} | -\mathbf{I}_p] \in \mathbb{R}^{p \times (2p+1)}$ is the constraints matrix. Here, $\sqrt{\boldsymbol{\tau}} = (\sqrt{\tau_1}, \dots, \sqrt{\tau_n})^\top \in \mathbb{R}^n$, $\mathbf{V}_\tau = [\sqrt{\tau_1}\mathbf{v}_1, \dots, \sqrt{\tau_n}\mathbf{v}_n]^\top \in \mathbb{R}^{n \times p}$, with \mathbf{v}_i are the design

vectors (see (28)), $\mathbf{0}_{n \times p} \in \mathbb{R}^{n \times p}$ contains 0's, $\mathbf{0}_p \in \mathbb{R}^p$ contains 0's, $\mathbf{1}_p \in \mathbb{R}^p$ contains 1's and \mathbf{I}_p is the identity matrix in $\mathbb{R}^{p \times p}$. As the last term in the objective function is independent of $\tilde{\boldsymbol{\gamma}}$, this problem is similar to the problem (50) and then can be solved by Dantzig selector estimate as the solution to the following problem

$$\begin{aligned} & \max_{\tilde{\boldsymbol{\gamma}} \in \mathbb{R}^{2p+1}} -\|\Lambda\tilde{\boldsymbol{\gamma}}\|_1 \\ & \text{subject to } \begin{cases} |\mathbf{X}_{\sigma\tau}^\top(\mathbf{y}_{\sigma\tau} - \mathbf{X}_{\sigma\tau}\tilde{\boldsymbol{\gamma}})| \leq \lambda \mathbf{1}_{2p+1}, \\ \mathbf{A}\tilde{\boldsymbol{\gamma}} = \mathbf{0}_p. \end{cases} \end{aligned}$$

Similarly to the problem in the gating network update, by decomposing $\tilde{\boldsymbol{\gamma}}$ into its positive and negative parts, $\tilde{\boldsymbol{\gamma}} = \tilde{\boldsymbol{\gamma}}_+ - \tilde{\boldsymbol{\gamma}}_-$, the above problem becomes

$$\begin{aligned} & \max_{(\tilde{\boldsymbol{\gamma}}_+, \tilde{\boldsymbol{\gamma}}_-) \in \mathbb{R}^{4p+2}} -[0, \mathbf{1}_p^\top, \rho \mathbf{1}_p^\top, 0, \mathbf{1}_p^\top, \varrho \mathbf{1}_p^\top] \begin{bmatrix} \tilde{\boldsymbol{\gamma}}_+ \\ \tilde{\boldsymbol{\gamma}}_- \end{bmatrix} \\ & \text{subject to } \begin{cases} \begin{bmatrix} \mathbf{X}_{\sigma\tau}^\top \mathbf{X}_{\sigma\tau} & -\mathbf{X}_{\sigma\tau}^\top \mathbf{X}_{\sigma\tau} \\ -\mathbf{X}_{\sigma\tau}^\top \mathbf{X}_{\sigma\tau} & \mathbf{X}_{\sigma\tau}^\top \mathbf{X}_{\sigma\tau} \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{\gamma}}_+ \\ \tilde{\boldsymbol{\gamma}}_- \end{bmatrix} \leq \begin{bmatrix} \lambda + \mathbf{X}_{\sigma\tau}^\top \mathbf{y}_{\sigma\tau} \\ \lambda - \mathbf{X}_{\sigma\tau}^\top \mathbf{y}_{\sigma\tau} \end{bmatrix}, \\ \begin{bmatrix} \mathbf{A} & -\mathbf{A} \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{\gamma}}_+ \\ \tilde{\boldsymbol{\gamma}}_- \end{bmatrix} = \mathbf{0}_{4p+2}, \\ \tilde{\boldsymbol{\gamma}}_+ \geq \mathbf{0}_{2p+1}, \quad \tilde{\boldsymbol{\gamma}}_- \geq \mathbf{0}_{2p+1}, \end{cases} \end{aligned} \tag{53}$$

which is a standard linear program with $4q + 2$ variables and can be solved similarly to the gating network case. The solution $\tilde{\boldsymbol{\gamma}}^* = (\beta_0^*, \boldsymbol{\gamma}^{[d_1]^* \top}, \boldsymbol{\gamma}^{[d_2]^* \top})^\top$ to the original problem is retrieved by the relation $\tilde{\boldsymbol{\gamma}}^* = \tilde{\boldsymbol{\gamma}}_+^* - \tilde{\boldsymbol{\gamma}}_-^*$ where $(\tilde{\boldsymbol{\gamma}}_+^{*\top}, \tilde{\boldsymbol{\gamma}}_-^{*\top})^\top$ is the solution of (53).

Finally, the update for σ_k^2 is

$$\sigma_k^{2(s+1)} = \frac{\sum_{i=1}^n \tau_{ik}^{(s)} (y_i - \beta_{k,0}^{(s+1)} - \mathbf{v}_i^\top \boldsymbol{\gamma}_k^{[d_1]^{(s+1)}})^2}{\sum_{i=1}^n \tau_{ik}^{(s)}}$$

in which $\beta_{k,0}^{(s+1)}, \boldsymbol{\gamma}_k^{[d_1]^{(s+1)}}$ are the new updates for $\beta_{k,0}$ and $\boldsymbol{\gamma}_k^{[d_1]}$.

D Detailed data generating protocol, simulation parameters and experimental protocol

D.1 Data generating protocol

In the simulated data, the data generation protocol is as follows. We consider a K -component functional mixture of Gaussian experts model that relates a scalar response $y \in \mathbb{R}$ to a univariate functional predictor $X(t), t \in \mathcal{T}$ defined on a domain $\mathcal{T} \subset \mathbb{R}$. Given the model parameters (defined in the next paragraph) $\boldsymbol{\beta} = \{\beta_{k,0}, \beta_k(t), \sigma_k^2\}_{k=1}^K$ and $\boldsymbol{\alpha} = \{\alpha_{k,0}, \alpha_k(t)\}_{k=1}^K, t \in \mathcal{T}$, we first construct the

Table 9 Simulation scenarios with different noise level σ_δ^2 and curve sampling level m

Scenario	S1	S2	S3	S4
σ_δ^2	1	1	4	4
m	300	100	300	100

functional predictors $X_i(\cdot)$ for $i \in [n]$ using the representation defined in (5), with detailed parameterization in (55). Then, for each $i \in [n]$, conditional on the functional predictor $X_i(\cdot)$, a hidden categorical random variable $Z_i \in [K]$ is generated following the multinomial distribution $\mathcal{M}\left(1, (\pi_1(X_i(t); \alpha), \dots, \pi_K(X_i(t); \alpha))\right)$, where $\pi_k(X_i(t); \alpha)$ for $k \in [K]$ is given by (3). Finally, conditional on $Z_i = z_i$ and $X_i(\cdot)$, the scalar response Y_i is obtained by sampling from the Gaussian distribution with mean $\beta_{z_i,0} + \int_{\mathcal{T}} X_i(t)\beta_{z_i}(t)dt$ and variance $\sigma_{z_i}^2$. The value z_i is then the true cluster label of the predictor $X_i(\cdot)$. This hierarchical generative process can be summarized as follows:

$$\begin{aligned}
 Y_i|Z_i = z_i, X_i(t) &\sim \mathcal{N}\left(\beta_{z_i,0} + \int_{\mathcal{T}} X_i(t)\beta_{z_i}(t)dt; \sigma_{z_i}^2\right), \\
 Z_i|X_i(t) &\sim \mathcal{M}\left(1, (\pi_1(X_i(t); \alpha), \dots, \pi_K(X_i(t); \alpha))\right).
 \end{aligned}
 \tag{54}$$

The true generated data, i.e., $(X_i(\cdot), Y_i, Z_i)_{i=1}^n$, are used for evaluating the prediction and clustering performance.

Finally, to mimic real data, in our simulations, the generated predictors $X_i(\cdot)$ are contaminated with measurement errors, that means we will not use $X_i(\cdot)$ for analysis, but

$$U_i(t) = X_i(t) + \delta_i(t),$$

with $\delta_i(t)$ is an independent Gaussian noise with mean zero and constant variance σ_δ^2 for all $t \in \mathcal{T}$. We considered different noise levels σ_δ^2 (see Table 9).

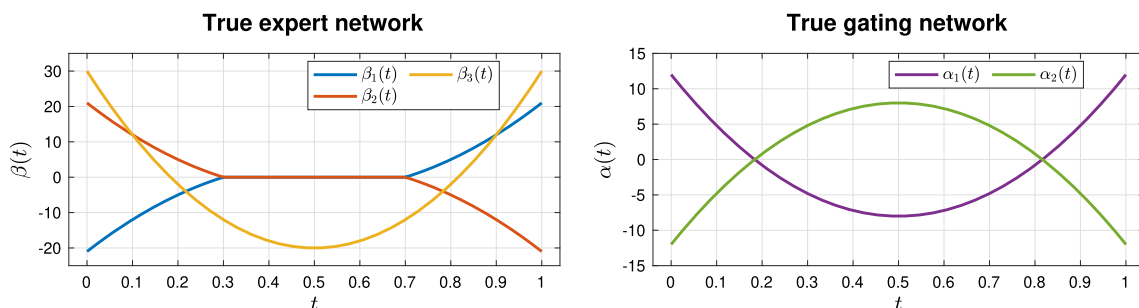


Fig. 11 The true expert and gating functions used in simulations

D.2 Simulation parameters and experimental protocol

The parameters that were used in the data generating process (54) are as follows. We consider $K = 3$ components and a time domain $\mathcal{T} = [0, 1]$.

Functional experts and gating parameters The functional experts parameters are given by

$$\begin{aligned}
 \beta_1(t) &= \begin{cases} -50(t - 0.5)^2 + 4 & \text{if } 0 \leq t < 0.3, \\ 0 & \text{if } 0.3 \leq t < 0.7, \\ 50(t - 0.5)^2 - 4 & \text{if } 0.7 \leq t \leq 1, \end{cases} \\
 \beta_2(t) &= -\beta_1(t), \\
 \beta_3(t) &= 100(t - 0.5)^2 - 10, \quad 0 \leq t \leq 1, \\
 (\beta_{1,0}, \beta_{2,0}, \beta_{3,0})^\top &= (-5, 0, 5)^\top, \\
 (\sigma_1^2, \sigma_2^2, \sigma_3^2)^\top &= (5, 5, 5)^\top,
 \end{aligned}$$

and the functional gating network parameters are given by

$$\begin{aligned}
 \alpha_1(t) &= 80(t - 0.5)^2 - 8, \\
 \alpha_2(t) &= -\alpha_1(t), \quad \alpha_3(t) = \mathbf{0}, \quad 0 \leq t \leq 1, \\
 (\alpha_{1,0}, \alpha_{2,0}, \alpha_{3,0})^\top &= (-10, -10, 0)^\top.
 \end{aligned}$$

Note that to satisfy the identifiability condition (see Jiang and Tanner (1999)), the experts are ordered, for instance $(\beta_{1,0}, \beta_1, \sigma_1^2) < \dots < (\beta_{K,0}, \beta_K, \sigma_K^2)$, and the last gating network parameters, $\alpha_{K,0}, \alpha_K(t)$, are initialized, i.e. fixed to zeros. Here, “ $<$ ” is the lexicographical order on \mathbb{R}^{p+2} .

As it can be seen in Fig. 11, the expert parameter functions $\beta_1(t)$ and $\beta_2(t)$ have a flat region in the interval $0.3 \leq t < 0.7$, outside of which they are quadratic, while $\beta_3(t)$ and the gating parameter functions $\alpha_1(t), \alpha_2(t)$ are all quadratic on the whole domain. By choosing these functional experts and gating parameters, we can later compare the sparsity in the zeroth and third derivatives of the solutions obtained by the proposed models.

Functional predictors parameters In this simulation, the functional predictors $X_i(\cdot)$ are constructed using the following formula,

$$X_i(t) = \mathbf{x}_i^\top \mathbf{b}(t), \quad t \in \mathcal{T}, \tag{55}$$

in which $\mathbf{b}(t) = [b_1(t), \dots, b_{10}(t)]^\top$ is a 10-dimensional B-spline basis $\mathbf{x}_i \in \mathbb{R}^{10}$ is a coefficient vector defined as $\mathbf{x}_i = \mathbf{W}\mathbf{v}_i$, where $\mathbf{W} \in \mathbb{R}^{10 \times 10}$ is a matrix of i.i.d. random values from the uniform distribution $\mathcal{U}(0, 1)$ and $\mathbf{v}_i \in \mathbb{R}^{10}$ is a vector of i.i.d. random values from the normal distribution $\mathcal{N}(0, 10)$. Here, the matrix \mathbf{W} acts as a factor to facilitate the fluctuation of the generated $X_i(\cdot)$.

Noisy functional predictors Since in real practical situations we do not usually directly observe the true functional predictors $X(\cdot)$'s, but only a noisy and discretized version of them, we thus consider several scenarios with different noise and sampling levels. To this end, in the data generating protocol we consider noisy versions $U_i(t_j) = X_i(t_j) + \delta_i(t_j)$ of the functional predictors values $X_i(t_j)$, where $\delta_i(t_j) \sim \mathcal{N}(0, \sigma_\delta^2)$ is a centred Gaussian noise with variance σ_δ^2 , for all $i \in [n]$, $j \in [m]$. We investigate simulated scenarios $S1, \dots, S4$ with curve length m and the noise level σ_δ^2 of the functional predictors, including two levels of sampling $m \in \{300, 100\}$, and two levels of noise $\sigma_\delta^2 \in \{1, 4\}$. The resulting four considered scenarios of simulated data are presented in Table 9.

For each considered scenario, we generate $N = 100$ datasets, each dataset contains n pairs of $(U_i(t), Y_i)$, with $n \in \{200, 400, 600, 800, 1000\}$.

References

- Brunel, Élodie., Mas, A., Roche, A.: Non-asymptotic adaptive prediction in functional linear models. *J. Multivar. Anal.* **143**, 208–232 (2016)
- Candes, E., Tao, T., et al.: The Dantzig selector: statistical estimation when p is much larger than n . *Ann. Stat.* **35**(6), 2313–2351 (2007)
- Chamroukhi, F., Lecocq, F., Nguyen, H.D.: Regularized estimation and feature selection in mixtures of Gaussian-gated experts models. In: *Research School on Statistics and Data Science*, pp. 42–56. Springer (2019)
- Chamroukhi, F., Huynh, B.-T.: Regularized maximum likelihood estimation and feature selection in mixtures-of-experts models. *J. Soc. Fr. Stat.* **160**(1), 57–85 (2019)
- Chamroukhi, F., Nguyen, H.D.: Model-based clustering and classification of functional data. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.* **9**(4), e1298 (2019)
- Chiou, J.-M., Müller, H.-G., Wang, J.-L.: Functional response models. *Stat. Sin.* **14**(3), 675–693 (2004)
- Ciarleglio, A., Ogden, R.T.: Wavelet-based scalar-on-function finite mixture regression models. *Comput. Stat. Data Anal.* **93**, 86–96 (2016)
- Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. B* **39**(1), 1–38 (1977)
- Devijver, E.: Model-based clustering for high-dimensional data. Application to functional data. *Adv. Data Anal. Classif.* **11**, 243–279 (2017)
- Ferraty, F., Vieu, P.: *Nonparametric Functional Data Analysis: Theory and Practice* (Springer Series in Statistics). Springer, Berlin (2006)
- Gaines, B.R., Kim, J., Zhou, H.: Algorithms for fitting the constrained lasso. *J. Comput. Graph. Stat.* **27**(4), 861–871 (2018). (PMID: **30618485**)
- Goldsmith, J., Bobb, J., Crainiceanu, C.M., Caffo, B., Reich, D.: Penalized functional regression. *J. Comput. Graph. Stat.* **20**(4), 830–851 (2011)
- Goldsmith, J., Crainiceanu, C.M., Caffo, B., Reich, D.: Longitudinal penalized functional regression for cognitive outcomes on neuronal tract measurements. *J. R. Stat. Soc. Ser. C (Appl. Stat.)* **61**(3), 453–469 (2012)
- Gormley, I.C., Murphy, T.B.: A mixture of experts model for rank data with applications in election studies. *Ann. Appl. Stat.* **2**(4), 1452–1477 (2008)
- Hastie, T., Tibshirani, R., Wainwright, M.: *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall/CRC, New York (2015)
- Huynh, T., Chamroukhi, F.: Estimation and feature selection in mixtures of generalized linear experts models. *arXiv preprint arXiv:1810.12161* (2019)
- Jacobs, R.A., Jordan, M.I., Nowlan, S.J., Hinton, G.E.: Adaptive mixtures of local experts. *Neural Comput.* **3**(1), 79–87 (1991)
- Jacques, J., Preda, C.: Functional data clustering: a survey. *Adv. Data Anal. Classif.* **8**(3), 231–255 (2014)
- Jacques, J., Preda, C.: Model-based clustering for multivariate functional data. *Comput. Stat. Data Anal.* **71**, 92–106 (2014)
- James, G.M.: Generalized linear models with functional predictor variables. *J. R. Stat. Soc. B* **64**, 411–432 (2002)
- James, G.M., Hastie, T.J.: Functional linear discriminant analysis for irregularly sampled curves. *J. R. Stat. Soc. B* **63**, 533–550 (2001)
- James, G.M., Sugar, C.: Clustering for sparsely sampled functional data. *J. Am. Stat. Assoc.* **98**(462), 397–408 (2003)
- James, G.M., Wang, J., Zhu, J.: Functional linear regression that's interpretable. *Ann. Stat.* **37**(5A), 2083–2108 (2009)
- Jiang, W., Tanner, M.: On the identifiability of mixtures-of-experts. *Neural Netw.* **12**(9), 1253–1258 (1999)
- Jordan, M.I., Jacobs, R.A.: Hierarchical mixtures of experts and the EM algorithm. *Neural Comput.* **6**, 181–214 (1994)
- Khalili, A.: New estimation and feature selection methods in mixture-of-experts models. *Can. J. Stat.* **38**(4), 519–539 (2010)
- Liu, X., Yang, M.: Simultaneous curve registration and clustering for functional data. *Comput. Stat. Data Anal.* **53**(4), 1361–1376 (2009)
- McLachlan, G.J., Krishnan, T.: *The EM Algorithm and Extensions*, 2nd edn. Wiley, New York (2008)
- McLachlan, G.J., Peel, D.: *Finite Mixture Models*. Wiley, New York (2000)
- Montuelle, L., Le Pennec, E., et al.: Mixture of gaussian regressions model with logistic weights, a penalized maximum likelihood approach. *Electron. J. Stat.* **8**(1), 1661–1695 (2014)
- Mousavi, S., Sørensen, H.: Multinomial functional regression with wavelets and lasso penalization. *Econometr. Stat.* **25**, 150–166 (2017)
- Mousavi, S.N., Sørensen, H.: Functional logistic regression: a comparison of three methods. *J. Stat. Comput. Simul.* **88**(2), 250–268 (2018)
- Müller, H.-G., Stadtmüller, U., et al.: Generalized functional linear models. *Ann. Stat.* **33**(2), 774–805 (2005)
- Nguyen, T., Nguyen, H.D., Chamroukhi, F., Forbes, F.: A non-asymptotic penalization criterion for model selection in mixture of experts models. *arXiv preprint* (2021b) [arXiv:2104.02640](https://arxiv.org/abs/2104.02640)
- Nguyen, T., Nguyen, H.D., Chamroukhi, F., McLachlan, G.J.: An l_1 -oracle inequality for the lasso in mixture-of-experts regression models. *arXiv preprint* (2020) [arXiv:2009.10622](https://arxiv.org/abs/2009.10622)
- Nguyen, H.D., Nguyen, T., Chamroukhi, F., McLachlan, G.J.: Approximations of conditional probability density functions in Lebesgue spaces via mixture of experts models. *J. Stat. Distrib. Appl.* **8**(1), 13 (2021a)

- Nguyen, H.D., Chamroukhi, F.: Practical and theoretical aspects of mixture-of-experts modeling: an overview. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.* **8**, e1246 (2018)
- Nguyen, H.D., Chamroukhi, F., Forbes, F.: Approximation results regarding the multiple-output Gaussian gated mixture of linear experts model. *Neurocomputing* **366**, 208–214 (2019)
- Qiao, X., Guo, S., James, G.M.: Functional graphical models. *J. Am. Stat. Assoc.* **114**(525), 211–222 (2019)
- Ramsay, J.O., Silverman, B.W.: *Applied Functional Data Analysis: Methods and Case Studies*. Springer Series in Statistics. Springer, Berlin (2002)
- Ramsay, J.O., Silverman, B.W.: *Functional Data Analysis*. Springer Series in Statistics. Springer, New York (2005)
- Schwarz, G.: Estimating the dimension of a model. *Ann. Stat.* **6**(2), 461–464 (1978)
- Städler, N., Bühlmann, P., Van De Geer, S.: ℓ_1 -penalization for mixture regression models. *TEST* **19**(2), 209–256 (2010)
- Tibshirani, R.: Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B* **58**(1), 267–288 (1996)
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., Knight, K.: Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **67**(1), 91–108 (2005)
- Wu, C.F.J.: On the convergence properties of the EM algorithm. *Ann. Stat.* **11**(1), 95–103 (1983)
- Xu, L., Jordan, M., Hinton, G.E.: An alternative model for mixtures of experts. *Adv. Neural. Inf. Process. Syst.* **7**, 633–640 (1994)
- Yao, F., Fu, Y., Lee, T.C.M.: Functional mixture regression. *Biostatistics* **12**(2), 341–353 (2010)
- Yuksel, S.E., Wilson, J.N., Gader, P.D.: Twenty years of mixture of experts. *IEEE Trans. Neural Netw. Learn. Syst.* **23**(8), 1177–1193 (2012)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.