

Tensor-variate Mixture of Experts

Noémie Jaquier¹, Robert Haschke² and Sylvain Calinon¹

Abstract—When data are organized in matrices or arrays of higher dimensions (tensors), classical regression methods first transform these data into vectors, therefore ignoring the underlying structure of the data and increasing the dimensionality of the problem. This flattening operation typically leads to overfitting when only few training data is available. In this paper, we present a mixture of experts model that exploits tensorial representations for regression of tensor-valued data. The proposed formulation takes into account the underlying structure of the data and remains efficient when few training data are available. Evaluation on artificially generated data, as well as offline and real-time experiments recognizing hand movements from tactile myography prove the effectiveness of the proposed approach.

Index Terms—Tensor methods, mixture of experts, generalized linear model, tactile myography

I. INTRODUCTION

The data collected by a robot are often naturally represented as matrices or tensors, i.e., generalization of matrices to arrays of higher dimensions [1]. Examples include video streams [2], movements in multiple coordinate systems [3], electroencephalography (EEG) [4], [5] or tactile myography (TMG) data [6]. Many approaches described in the literature consist of reorganizing the elements of these tensors into vectors before applying learning algorithms based on linear algebra, operating on vector spaces. This flattening operation ignores the underlying structure of the original data. Moreover, the dimensionality of the resulting problem dramatically increases, creating high computational and memory requirements. Finally, the number of model parameters to estimate in the learning method may become high, which constitutes an important issue in the cases where only few training data are available.

With the burst of multidimensional data available in various fields of research, important efforts were turned toward extending standard dimensionality reduction and learning techniques to tensor data. In this context, Ozdemir *et al.* [7] proposed a review of tensor decomposition methods by dividing them into three categories of problems usually targeted by principal component analysis (PCA), namely low-rank tensor approximation, low-rank tensor recovery, and feature extraction. In particular, multilinear PCA (MPCA) [8] and weighted MPCA [5] were proposed to extract features from tensor objects as a preprocessing step for classification. Similarly, linear discriminant analysis was extended to multilinear discriminant

analysis in [9] and factor analysis was adapted to tensor data in [10]. In the context of regression, Guo *et al.* [11] proposed generalizations of ridge regression (RR) and support-vector regression (SVR) methods to tensor data, where they showed the superiority of tensor-based algorithms over the vector-based algorithms in various applications. A similar extension of RR to tensor data was proposed in [12] with an application in magnetic resonance imaging (MRI). Following a similar process, tensor-variate logistic regression (LR) was proposed in [13], [14] for the classification of multidimensional data. Moreover, kernel-based frameworks such as Gaussian processes (GPs) [2] were also adapted to tensors [15].

In this paper, we introduce a tensor-variate mixture of experts (TME) model for regression. Mixture of experts (ME) models, first introduced by Jacobs *et al.* [16], combine the predictions of several experts based on their probability of being active in a given region of the input space. Each expert is a regression function and a gate determines the regions of the input space where each expert is trustworthy. The output of the model is a weighted sum of the experts predictions. Over the years, ME models were widely improved with different gates, regression and classification models for the experts, and used in many applications (see [17] for a review).

In order to handle tensor data in a ME model, we propose to use tensor-variate models for the experts and for the gate. As explained in Section III, the experts are defined using tensor linear models and the gate is set as a tensor-variate softmax function. Both elements are based on the inner product between the input tensor data and model parameters (see Section II). The resulting TME model is trained with an EM algorithm using the CANDECOMP/PARAFAC (CP) decomposition [18], [19].

The functionality of the proposed approach is first evaluated and compared to the corresponding vector-based approach using artificially generated data (Section IV-A). The TME model is then exploited in the context of prosthetic hands to recognize hand movements from tactile myography (TMG).

Our TMG sensor is made of 320 cells organized in a 8×40 array forming a bracelet [6], therefore providing intrinsically matrix-valued data. We show the effectiveness of our approach in an offline experiment with the aim of detecting finger and wrist movements from TMG data (Section IV-B). The TME model outperforms the standard ME model and achieves similar performances as a nonlinear model, namely a GP, at a lower computational cost. Note that a ME model with linear experts was shown to achieve similar performance as more complex nonlinear methods at a lower computational cost in wrist movements recognition based on electromyographic (EMG) signals [20]. We finally validate the use of the proposed approach in a real-time teleoperation experiment, where participants controlled a robotic arm and hand by moving their

This work was supported by the Swiss National Science Foundation and the German Research Foundation (SNSF/DFG project TACT-HAND).

¹N. Jaquier and S. Calinon are with the Idiap Research Institute, Rue Marconi 19, CH-1920 Martigny, Switzerland. {noemie.jaquier, sylvain.calinon}@idiap.ch

²R. Haschke is with the Center of Excellence - Cognitive Interaction Technology (CITEC), Bielefeld University, D-33619 Bielefeld, Germany. rhaschke@techfak.uni-bielefeld.de

wrist and closing/opening their hand (Section IV-C).

The contributions of this paper are two-fold: (i) we propose a tensor-variate mixture of experts model that exploits tensorial representations to take into account the structure of the data in the regression process; and (ii) we demonstrate the efficiency of tensor-based approaches in the context of prosthetic hands to recognize hand movements from tactile myography. In particular, we present a teleoperation experiment based on tactile myography to control a robotic arm and hand.

II. BACKGROUND

Tensors are generalization of matrices to arrays of higher dimensions [1], where vectors and matrices correspond to 1st and 2nd-order tensors. Tensor representation permits to represent and exploit the intrinsic structure of multidimensional arrays. We introduce here the tensor operations necessary for the proposed TME, as well as the extensions of two forms of the generalized linear model, namely ridge regression and logistic regression, to tensor-variate data.

A. Tensor operations

The inner product of two tensors $\mathcal{X}, \mathcal{Y} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_M}$ is defined as the sum of the products of their entries, so that

$$\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{i_1=1}^{I_1} \sum_{i_2=1}^{I_2} \dots \sum_{i_M=1}^{I_M} x_{i_1, i_2, \dots, i_M} y_{i_1, i_2, \dots, i_M}. \quad (1)$$

Note that the inner product of two tensors is equivalent to the Frobenius inner product of their m -mode matricization or unfolding $\mathbf{X}_{(m)}, \mathbf{Y}_{(m)} \in \mathbb{R}^{I_m \times (I_1 \dots I_{m-1} I_{m+1} \dots I_M)}$ and to the inner product of their vectorization $\text{vec}(\mathcal{X}), \text{vec}(\mathcal{Y}) \in \mathbb{R}^{I_1 \dots I_M}$

$$\langle \mathcal{X}, \mathcal{Y} \rangle = \langle \mathbf{X}_{(m)}, \mathbf{Y}_{(m)} \rangle = \langle \text{vec}(\mathcal{X}), \text{vec}(\mathcal{Y}) \rangle. \quad (2)$$

A rank-one tensor \mathcal{Y} of order M is a tensor that can be written as the outer product of M vectors, i.e.,

$$\mathcal{Y} = \mathbf{u}^{(1)} \circ \mathbf{u}^{(2)} \circ \dots \circ \mathbf{u}^{(M)}, \quad (3)$$

with \circ the outer product between vectors, so that each element of \mathcal{Y} is equal to $x_{i_1, i_2, \dots, i_M} = u_{i_1}^{(1)} u_{i_2}^{(2)} \dots u_{i_M}^{(M)}$ and M the number of dimensions or modes of the tensor.

The CANDECOM/PARAFAC (CP) decomposition factorizes a tensor $\mathcal{Y} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_M}$ as a sum of R rank-one tensors, i.e.,

$$\mathcal{Y} \approx \sum_{r=1}^R \mathbf{u}_r^{(1)} \circ \mathbf{u}_r^{(2)} \circ \dots \circ \mathbf{u}_r^{(M)}. \quad (4)$$

The smallest number of rank-one tensors that generate \mathcal{Y} as their sum is defined as the rank of the tensor \mathcal{Y} . It corresponds to the smallest number of components $R = \text{rank}(\mathcal{Y})$ in the CP decomposition. The CP decomposition can also be expressed in terms of the m -mode matricization and vectorization of the tensor $\mathbf{Y}_{(m)}$ and $\text{vec}(\mathcal{Y})$ as

$$\mathbf{Y}_{(m)} \approx \mathbf{U}^{(m)} \mathbf{U}^{(-m)\top}, \quad (5)$$

$$\text{vec}(\mathcal{Y}) \approx (\mathbf{U}^{(1)} \odot \dots \odot \mathbf{U}^{(M)}) \mathbf{1}_R, \quad (6)$$

where \odot denotes the Khatri-Rao product, $\mathbf{U}^{(m)} \in \mathbb{R}^{I_m \times R}$ are factor matrices defined as $\mathbf{U}^{(m)} = [\mathbf{u}_1^{(m)} \dots \mathbf{u}_R^{(m)}]$, $\mathbf{U}^{(-m)} = (\mathbf{U}^{(M)} \odot \dots \odot \mathbf{U}^{(m+1)} \odot \mathbf{U}^{(m-1)} \odot \dots \odot \mathbf{U}^{(1)})$, $\mathbf{1}_R \in \mathbb{R}^R$ is a vector of R ones. The Khatri-Rao product of two matrices $\mathbf{A} \in \mathbb{R}^{I \times K}$ and $\mathbf{B} \in \mathbb{R}^{J \times K}$ results in a matrix $\mathbf{A} \odot \mathbf{B} \in \mathbb{R}^{IJ \times K}$ whose columns are equal to the Kronecker product of the columns of \mathbf{A} and \mathbf{B} , i.e., $\mathbf{A} \odot \mathbf{B} = [\mathbf{a}_1 \otimes \mathbf{b}_1 \dots \mathbf{a}_K \otimes \mathbf{b}_K]$.

If \mathcal{Y} follows exactly a CP decomposition (4), the inner product (2) can equivalently be written as

$$\begin{aligned} \langle \mathcal{X}, \mathcal{Y} \rangle &= \langle \mathbf{X}_{(m)}, \mathbf{U}^{(m)} \mathbf{U}^{(-m)\top} \rangle \\ &= \langle \mathbf{X}_{(m)} \mathbf{U}^{(-m)}, \mathbf{U}^{(m)} \rangle, \end{aligned} \quad (7)$$

by exploiting (5) and the properties of the Frobenius norm and matrix trace.

B. Generalized linear model for tensors

Given a vector-valued input data \mathbf{x} , the generalized linear model (GLM) is given by

$$y = f(\mathbf{x}^\top \mathbf{w} + b) = f(\langle \mathbf{x}, \mathbf{w} \rangle + b), \quad (8)$$

where y is the predicted output, \mathbf{w} is a vector of weights, b is the bias and f is a function, see Figure 1-top. This model can be naturally extended to matrix-valued data \mathbf{X} as

$$\begin{aligned} y &= f(\mathbf{w}^{(1)\top} \mathbf{X} \mathbf{w}^{(2)} + b) \\ &= f(\langle \mathbf{X}, \mathbf{w}^{(1)} \circ \mathbf{w}^{(2)} \rangle + b), \end{aligned} \quad (9)$$

where $\mathbf{w}^{(1)}$ and $\mathbf{w}^{(2)}$ are vectors of weights. Following a similar procedure, the model is generalized to M -dimensional tensor-valued data as

$$y = f(\langle \mathcal{X}, \mathbf{w}^{(1)} \circ \dots \circ \mathbf{w}^{(M)} \rangle + b), \quad (10)$$

as shown in Figure 1-middle. The key advantages of this representation compared to vector-valued representation $y = f(\text{vec}(\mathcal{X})^\top \mathbf{w} + b)$ are that the underlying structure of the tensor-valued data is taken into account in the model and that the number of parameters is reduced from $\prod_{m=1}^M I_m$ to $\sum_{m=1}^M I_m$. Moreover, more complex features can be represented by encoding the weight tensor as a sum of R rank-one tensors with

$$y = f(\langle \mathcal{X}, \sum_{r=1}^R \mathbf{w}_r^{(1)} \circ \dots \circ \mathbf{w}_r^{(M)} \rangle + b). \quad (11)$$

This model is represented in Figure 1-bottom.

Similarly to the vector case, we obtain the tensor-valued linear and logistic regression models by defining the function f as identity and as the softmax function, respectively.

C. Tensor Ridge regression

Given a vector-valued input data \mathbf{x} , the classical regression model is of the form

$$y = \mathbf{x}^\top \mathbf{w} + b + \epsilon = \langle \mathbf{x}, \mathbf{w} \rangle + b + \epsilon, \quad (12)$$

where y is the predicted output, \mathbf{w} is a vector of weights, b is the bias and ϵ is a zero-mean Gaussian noise variable.

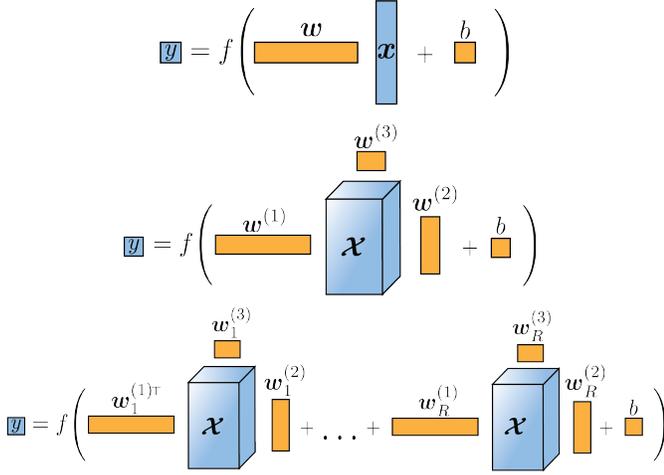


Fig. 1: (Top) GLM for vector-valued data. The data and model parameters are depicted in blue and orange, respectively. (Middle-bottom) Extensions of GLM to tensor-valued data. The bottom representation allows the encoding of more complex behaviors as the weight tensor is encoded as a sum of rank-1 tensors.

Following (11), as shown in [11], [12], the model can be generalized to M -dimensional tensor-valued data

$$y = \langle \mathcal{X}, \sum_{r=1}^R \mathbf{w}_r^{(1)} \circ \dots \circ \mathbf{w}_r^{(M)} \rangle + b + \epsilon, \quad (13)$$

$$= \langle \mathcal{X}, \mathbf{W} \rangle + b + \epsilon,$$

therefore taking the underlying structure of the data into account.

Given a dataset of M -dimensional tensor inputs and corresponding outputs $\{\mathcal{X}_n, y_n\}_{n=1}^N$, the parameters of the tensor ridge regression (TRR) model (13) are learned by maximizing its likelihood or equivalently its corresponding log-likelihood

$$\ell(\mathbf{W}, b, \sigma) = \sum_{n=1}^N \log \mathcal{N}(y_n | \langle \mathcal{X}_n, \mathbf{W} \rangle + b, \sigma^2), \quad (14)$$

where σ^2 is the variance of the zero-mean Gaussian random variable ϵ . By using the inner product equality (7), one can observe that the model (13) is linear in $\mathbf{W}^{(m)} = [\mathbf{w}_1^{(m)} \dots \mathbf{w}_R^{(m)}]$ individually, so that the parameters $\{\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(M)}, b\}$ can be learned by optimizing a sequence of generalized linear models (see [11], [12] for details). Therefore, by adding a zero-mean Gaussian prior on the weight tensor, equivalent to the regularization term $-\lambda_{\mathbf{W}} \sum_{m=1}^M \|\mathbf{W}^{(m)}\|_F^2$, to the log-likelihood function, the bias b and factor matrices $\mathbf{W}^{(m)}$ are updated at each iteration until convergence with

$$\text{vec}(\mathbf{W}^{(m)}) \leftarrow (\Phi^\top \Phi + \lambda_{\mathbf{W}} \mathbf{I})^{-1} \Phi^\top (\mathbf{y} - b \mathbf{1}), \quad (15)$$

$$b \leftarrow \frac{1}{N} \sum_{n=1}^N y_n - \langle \mathcal{X}_n, \mathbf{W} \rangle, \quad (16)$$

where the n -th row of the matrix Φ is equal to $\text{vec}(\mathcal{X}_{n,(m)} \mathbf{W}^{(-m)})$, the n -th element of the vector \mathbf{y} is y_n and $\mathbf{1} \in \mathbb{R}^N$ is a vector of N ones. Note that other types of regularization are also proposed in [11].

D. Tensor logistic regression

In the classical multi-class logistic regression model, the posterior probability of the class \mathcal{C}_i is given by the softmax function

$$p(\mathcal{C}_i | \mathbf{x}, \boldsymbol{\theta}) = \frac{\exp(\mathbf{x}^\top \mathbf{v}_i + a_i)}{\sum_{j=1}^C \exp(\mathbf{x}^\top \mathbf{v}_j + a_j)}, \quad (17)$$

where $\boldsymbol{\theta}$ denotes the parameters of the model and C the number of classes. Similarly as ridge regression, the logistic regression model can be extended to tensor-valued data by encoding the tensor of weights as a sum of R rank-one tensors, thus leading to the following tensor-valued softmax function [13], [14]

$$\pi_i = p(\mathcal{C}_i | \mathcal{X}, \boldsymbol{\theta}) = \frac{\exp(\langle \mathcal{X}, \mathbf{V}_i \rangle + a_i)}{\sum_{j=1}^C \exp(\langle \mathcal{X}, \mathbf{V}_j \rangle + a_j)}, \quad (18)$$

where $\mathbf{V}_i = \sum_{r=1}^R \mathbf{v}_{i,r}^{(1)} \circ \dots \circ \mathbf{v}_{i,r}^{(M)}$. Similarly to TRR, the tensor logistic regression (TLR) model takes into account the underlying structure of the data and reduces the number of parameters in the model compared to a vector-based representation of the tensor-valued data.

Given a dataset of inputs and corresponding unit vector outputs $\{\mathcal{X}_n, y_n\}_{n=1}^N$ where $y_{n,i} = 1$ indicates that the n -th data belong to the i -th class, the log-likelihood of the multivariate tensor logistic regression model is

$$\begin{aligned} \ell(\{\mathbf{V}_i, a_i\}_{i=1}^C) &= \log \prod_{n=1}^N \prod_{i=1}^C \pi_i^{y_{n,i}} \\ &= \sum_{n=1}^N \left(\sum_{i=1}^C y_{n,i} (\langle \mathcal{X}_n, \mathbf{V}_i \rangle + a_i) - \log \sum_{i=1}^C \exp(\langle \mathcal{X}_n, \mathbf{V}_i \rangle + a_i) \right). \end{aligned} \quad (19)$$

Note that a regularization term $-\lambda_{\mathbf{V}} \sum_{i=1}^C \sum_{m=1}^M \|\mathbf{V}_i^{(m)}\|_F^2$ can be added to the log-likelihood function to avoid the problem of overfitting. The parameters $\{\mathbf{V}_i^{(1)}, \dots, \mathbf{V}_i^{(M)}, a_i\}_{i=1}^C$ can be learned by minimizing the negative log-likelihood of the model via any gradient-based optimizer, e.g., Newton's method or limited memory BFGS. By exploiting (7), the gradients of the regularized negative log-likelihood used in the optimization process can be computed as

$$\frac{\delta(-\ell(\{\mathbf{V}_i, a_i\}_{i=1}^C))}{\delta \mathbf{V}_i^{(m)}} = \sum_{n=1}^N (\pi_{n,i} - y_{n,i}) \text{vec}(\mathcal{X}_{n,(m)} \mathbf{V}_i^{(-m)}) + 2 \lambda_{\mathbf{V}} \text{vec}(\mathbf{V}_i^{(m)}), \quad (20)$$

$$\frac{\delta(-\ell(\{\mathbf{V}_i, a_i\}_{i=1}^C))}{\delta a_i} = \sum_{n=1}^N (\pi_{n,i} - y_{n,i}), \quad (21)$$

where $\pi_{n,i} = p(\mathcal{C}_i | \mathcal{X}_n, \boldsymbol{\theta})$.

III. TENSOR-VARIATE MIXTURE OF EXPERTS

A mixture of experts (ME) regression model [16] aims at solving a nonlinear supervised learning problem by combining the predictions of a set of experts. The model is composed of a gate determining a soft division of the input space and several experts making predictions in the different regions of

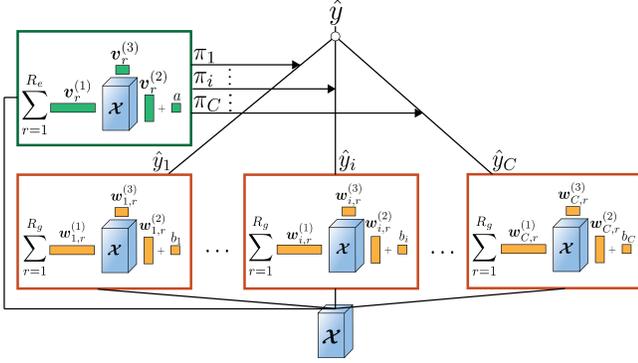


Fig. 2: TME model. The gate is represented in the green box and the experts are represented in orange boxes.

the input space. In this section, we propose to generalize the ME regression model to tensor-variate data by using tensor-variate models for the experts and for the gate.

A. TME model

Given a tensor-variate input \mathcal{X} and an output \mathbf{y} , the tensor-variate mixture distribution is

$$p(\mathbf{y}|\mathcal{X}, \boldsymbol{\theta}) = \sum_{i=1}^C p(i|\mathcal{X}, \boldsymbol{\theta}_g) p(\mathbf{y}|i, \mathcal{X}, \boldsymbol{\theta}_e), \quad (22)$$

where C is the number of experts, $p(i|\mathcal{X}, \boldsymbol{\theta}_g)$ is the probability of the i -th expert to be activated (gate's rating) and $p(\mathbf{y}|i, \mathcal{X}, \boldsymbol{\theta}_e)$ is the model of the i -th expert. $\boldsymbol{\theta} = \{\boldsymbol{\theta}_g, \boldsymbol{\theta}_e\}$, where $\boldsymbol{\theta}_g$ and $\boldsymbol{\theta}_e$ denote the parameters of the gate and the set of experts, respectively.

Similarly to the original ME model, we define the gate of the TME model by the tensor-variate softmax function, so that

$$p(i|\mathcal{X}, \boldsymbol{\theta}_g) = \pi_i, \quad (23)$$

with π_i defined by (18), $\mathbf{v}_i = \sum_{r=1}^{R_g} \mathbf{v}_{i,r}^{(1)} \circ \dots \circ \mathbf{v}_{i,r}^{(M)}$ and R_g the rank of the weight tensors \mathbf{v}_i . The experts follow the Gaussian model

$$p(\mathbf{y}|i, \mathcal{X}, \boldsymbol{\theta}_e) = \mathcal{N}(\mathbf{y}|\boldsymbol{\psi}_i(\mathcal{X}) + \mathbf{b}_i, \boldsymbol{\Sigma}_i), \quad (24)$$

where $\boldsymbol{\psi}_i(\mathcal{X}) = \begin{pmatrix} \langle \mathcal{X}, \mathcal{W}_{i,1} \rangle \\ \vdots \\ \langle \mathcal{X}, \mathcal{W}_{i,D} \rangle \end{pmatrix}$, $\mathcal{W}_{i,d} = \sum_{r=1}^{R_{e_i}} \mathbf{w}_{i,d,r}^{(1)} \circ \dots \circ$

$\mathbf{w}_{i,d,r}^{(M)}$ and R_{e_i} is the rank of the weight tensors $\mathcal{W}_{i,d}$. Note that one weight tensor $\mathcal{W}_{i,d}$ is defined for each element of $\mathbf{y} \in \mathbb{R}^D$. This is similar to the vector case, where different vectors \mathbf{w}_i weight the input for each element of the output,

so that $\mathbf{y} = \begin{pmatrix} \mathbf{w}_1^\top \\ \vdots \\ \mathbf{w}_D^\top \end{pmatrix} \mathbf{x} + \mathbf{b}$.

Figure 2 illustrates the proposed model. Single predictions are computed by using the expectation of the TME model (22), so that

$$\hat{\mathbf{y}} = \sum_{i=1}^C \pi_i (\boldsymbol{\psi}_i(\mathcal{X}) + \mathbf{b}_i). \quad (25)$$

B. Training of TME

Similarly to ME, the TME model can be trained using the expectation-maximization (EM) algorithm. By introducing a set of binary latent variables $\{z_{n,i}\}$ where $z_{n,i} = 1$ indicate that the data n was generated by the i -th mixture component, the expected complete data log-likelihood is given by

$$Q(\boldsymbol{\theta}) = \sum_{n=1}^N \sum_{i=1}^C r_{n,i} \log \left(\pi_{n,i} \mathcal{N}(\mathbf{y}_n | \boldsymbol{\psi}_i(\mathcal{X}_n) + \mathbf{b}_i, \boldsymbol{\Sigma}_i) \right), \quad (26)$$

where $r_{n,i}$ denotes the responsibility of the i -th component for the n -th data point so that $r_{n,i} = p(z_{n,i} = 1 | \mathcal{X}_n, \boldsymbol{\theta})$. In the E-step, the responsibilities $r_{n,i}$ are computed using

$$r_{n,i} = \frac{\pi_{n,i} \mathcal{N}(\mathbf{y}_n | \boldsymbol{\psi}_i(\mathcal{X}_n) + \mathbf{b}_i, \boldsymbol{\Sigma}_i)}{\sum_{j=1}^C \pi_{n,j} \mathcal{N}(\mathbf{y}_n | \boldsymbol{\psi}_j(\mathcal{X}_n) + \mathbf{b}_j, \boldsymbol{\Sigma}_j)}. \quad (27)$$

In the M-step, the parameters are updated to maximize the expected complete data log-likelihood (26). First, the parameters of the experts $\boldsymbol{\theta}_e$ are updated iteratively until convergence, based on (15) and (16), with

$$\text{vec}(\mathbf{W}_{i,d}^{(m)}) \leftarrow (\tilde{\boldsymbol{\Phi}}_{i,d}^\top \tilde{\boldsymbol{\Phi}}_{i,d} + \lambda \mathbf{W} \mathbf{I})^{-1} \tilde{\boldsymbol{\Phi}}_{i,d}^\top (\tilde{\mathbf{y}}_d - \mathbf{b}_{i,d} \mathbf{1}), \quad (28)$$

$$\mathbf{b} \leftarrow \frac{1}{N} \sum_{n=1}^N \tilde{\mathbf{y}}_{n,d} - \langle \tilde{\mathcal{X}}_n, \mathcal{W}_{i,d} \rangle, \quad (29)$$

where the n -th row of the matrix $\tilde{\boldsymbol{\Phi}}_{i,d}$ is equal to $\text{vec}(\tilde{\mathcal{X}}_{n,(m)} \mathbf{W}_{i,d}^{(-m)})$, the n -th element of the vector $\tilde{\mathbf{y}}$ is $\tilde{y}_{n,d}$ and $\tilde{\mathcal{X}}_n = \sqrt{r_{nk}} \mathcal{X}_n$, $\tilde{\mathbf{y}}_n = \sqrt{r_{nk}} \mathbf{y}_n$ are the scaled input tensors and output vectors, respectively. The covariance of the experts Gaussian model is then updated as

$$\boldsymbol{\Sigma}_i \leftarrow \frac{\sum_{n=1}^N r_{n,i} (\mathbf{y}_n - \boldsymbol{\psi}_i(\mathcal{X}_n) - \mathbf{b}_i)^\top (\mathbf{y}_n - \boldsymbol{\psi}_i(\mathcal{X}_n) - \mathbf{b}_i)}{\sum_{n=1}^N r_{n,i}}. \quad (30)$$

Finally, the gate parameters $\boldsymbol{\theta}_g$ are updated by maximizing the log-likelihood of the multivariate tensor logistic regression model

$$\ell(\{\mathbf{v}_i, a_i\}_{i=1}^C) = \log \prod_{n=1}^N \prod_{i=1}^C \pi_{n,i}^{r_{n,i}} - \lambda \mathbf{v} \sum_{i=1}^C \sum_{m=1}^M \|\mathbf{v}_i^{(m)}\|_F^2, \quad (31)$$

based on (19). Similarly to tensor logistic regression, a gradient-based optimizer is used to minimize the negative log-likelihood with gradients given by (20) and (21), where $y_{n,i}$ is replaced by $r_{n,i}$.

Note that regularization terms have been added in the M-step to avoid the overfitting problem. The E-step and M-step are iterated until convergence of the algorithm.

C. Model selection and initialization

Selecting the number of experts of ME is known as a difficult problem [17]. In the case where the structure of the application allows it, as in the experiments of Section IV, the number of experts can be determined by the experimenter. Otherwise, strategies used for ME, such as exhaustive search, growing or pruning models can be adapted for TME.

The TME model assumes fixed ranks R_g and R_{e_i} for the gate and experts weight tensors, respectively. The appropriate rank can be estimated using cross-validation or through usual model selection criterion, e.g., the Bayesian information criterion (BIC).

Previous works on TRR [11], [12] and TLR [14] showed that both TRR and TLR models converge to a similar solution independently of the initial weight values. Therefore, the weight tensors of TME can be initialized with random values. In order to facilitate the convergence, we initialized the weights of the expert model \mathcal{W}_i and of the gate \mathcal{V}_i as equal to the weights \mathcal{W} obtained from TRR.

IV. EXPERIMENTS

In this section, we first evaluate the functionality and the performance of the proposed TME on artificially generated data. The approach is then applied to the detection and recognition of hand movements from tactile myography (TMG) data. An offline experiment and a real-time teleoperation experiment, where participants controlled a robotic arm and hand based on their hand movements, illustrate the effectiveness of the proposed TME model. A video of the teleoperation experiment accompanies the paper. Source codes related to the experiments will be available after publication.

A. 2D shape example

In this illustrative example, we propose to evaluate the performance of the proposed TME model for different ranks under various sample sizes and signal strengths. To do so, we generate artificial data following the model (22) from known parameters θ and we evaluate the recovery of those parameters by the model. In this illustrative example, we consider matrix-variate inputs $\mathbf{X} \in \mathbb{R}^{64 \times 64}$ whose elements are independent and normally distributed. The output y is normally distributed with a mean given by a 2-classes TME model with zero biases

$$\hat{y} = \frac{\exp(\langle \mathbf{X}, \mathbf{V} \rangle)}{1 + \exp(\langle \mathbf{X}, \mathbf{V} \rangle)} \langle \mathbf{X}, \mathbf{W}_1 \rangle + \frac{1}{1 + \exp(\langle \mathbf{X}, \mathbf{V} \rangle)} \langle \mathbf{X}, \mathbf{W}_2 \rangle, \quad (32)$$

and a standard deviation σ . The weight matrices \mathbf{V} , \mathbf{W}_1 , $\mathbf{W}_2 \in \mathbb{R}^{64 \times 64}$ are equal to the binary 2D shapes represented in Figure 3a, where the black and white regions correspond to 1 and 0, respectively. The use of these 2D shapes was inspired by the illustrative example presented by [12].

We first examine the performances of the proposed TME model for ranks R_g and $R_e = R_{e_1} = R_{e_2}$ varying from 1 to 3 with a sample size $N = 1000$ equally divided between the two classes and a noise level σ equal to 10% of the standard deviation of the mean \hat{y} . The regularization terms $\lambda_{\mathbf{V}}$ and $\lambda_{\mathbf{W}}$ were fixed as equal to 0.1. Moreover, we compare the TME model with the standard ME regression model whose gate is defined by the softmax function (17) and experts follow a Gaussian model with a mean given by (12).

Figure 3 shows the original and recovered weight matrices by the ME and TME models along with the root-mean-square error (RMSE) for the estimation of the weight matrices and

| | (a) | (b) | (c) | (d) | (e) | (f) |
|------|------|-------|-------|-------|-------|-------|
| | | ME | TME | TME | TME | TME |
| | | | (2,1) | (1,2) | (2,2) | (3,3) |
| RMSE | 0.30 | 0.16 | 0.18 | 0.12 | 0.18 | |
| BIC | - | 10209 | 8940 | 9835 | 11216 | |

Fig. 3: (a) True weight matrices of the gate and experts functions (from top to bottom: \mathbf{V} , \mathbf{W}_1 and \mathbf{W}_2). (b) Recovered weight matrices by ME. (c-f) Recovered weight matrices by TME for different ranks (R_g, R_e).

the BIC value for TME. We observe that TME outperforms ME for all the tested rank values as the maximum RMSE value achieved by TME is 0.21 ($R_g = 1, R_e = 3$) versus 0.3 for the ME model. Moreover, we observe that the weight matrices retrieved by ME are noisier and the shapes of the experts weights are not clearly delimited and tend to be fused together compared to those retrieved by TME. Similar results were obtained for different sample sizes and noise levels.

Due to their structure, a rank-2 setting is sufficient to capture a cross or a t-shape, while a low rank setting does not allow to exactly represent a disk shape. As expected, the cross and t-shape are fully recovered by TME in the cases where $R_g \geq 2$ and $R_e \geq 2$, respectively, while approximations of the shapes are obtained for lower ranks (see Figures 3c–3f). Moreover, while the disk shape is approximated by a square in a rank-1 setting (Figure 3c), it is already fairly recovered by a rank-2 or rank-3 setting (see Figures 3d–3f).

Consistently with the aforementioned observations, the minimum RMSE value is obtained by TME with a rank-(2, 2) setting. Moreover, TME with ranks $R_g = R_e = 3$ obtains a slightly higher RMSE than with ranks $R_g = R_e = 2$. This can be explained by the fact that approximating the cross and t-shape with a rank-3 setting, while a rank-2 setting is sufficient, leads to an overfitted estimation with a higher influence of the noise contained in the training data. According to the BIC values reported for the tested TME models, the model with $R_g = 1$ and $R_e = 2$ should be selected (lowest BIC cost). However, in practice, one may prefer the rank-(2,2) setting in this case, suggesting that other rank selection methods, such as cross-validation, may be used in function of the application. Note that similar observations were made for different sample sizes and noise levels.

As shown in Figure 4, the estimation accuracy increases with the sample size and decreases with the noise level σ , validating the consistency of the proposed method.

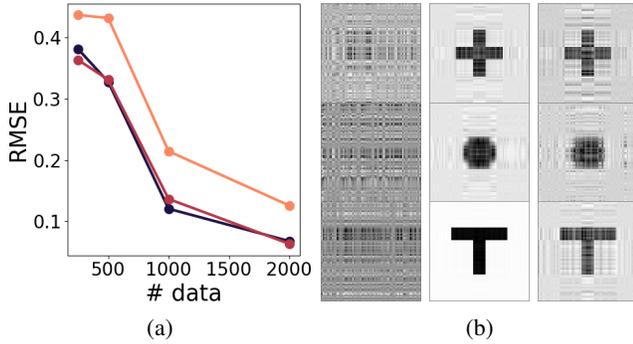


Fig. 4: (a) Evolution of the RMSE of the estimation of TME weight matrices in function of the sample size N for different noise levels. The curves corresponding to noise levels σ equal to 1, 10 and 50% of the standard deviation of the mean \hat{y} are represented in dark blue, red and yellow, respectively. The sample size is equally divided between the two classes. (b) Recovered weight matrices V , W_1 and W_2 by TME for ranks $R_g = R_e = 2$ with, from left to right, sample sizes $N = 500, 2000$ and 2000 and noise levels $\sigma = 1\%, 1\%$ and 50% of \hat{y} .

B. Detection of hand movements from tactile myography

In the context of prosthetic hands, tactile myography (TMG) has recently been proposed as a complementary or alternative approach to the traditional surface electromyography (sEMG) to achieve simultaneous and proportional control of multiples degrees of freedom (DOFs) of a hand prosthesis (see, e.g., [6], [21]). In this context, the aim of TMG is to measure the pressure related to the deformation induced by the muscles activity of the forearm. This signal is then used to determine the corresponding hand and wrist movements. Our TMG sensor, developed in [6] and showed in Figure 5a, is composed of 320 resistive taxels organized in a 8×40 array forming a bracelet. Therefore, the data provided by the sensor are intrinsically matrix-valued.

Previous studies showed that Ridge regression (RR) directly applied to the data of the bracelet allows the prediction of different finger and wrist movements [6], which could outperform detection using sEMG [22]. However, RR does not take into account the matrix structure of the TMG data as they are vectorized before the application of the regression method. Moreover, the same weight vectors are used independently of the activated movements, which may result in false positive detection of activations. Therefore, the motivations to use TME for this application are the following: (1) the structure of the data is taken into account in the regression process; (2) the problem is decomposed in two subparts, namely detecting which movements are activated and determining their individual level of activation; and (3) the low computational complexity to evaluate one test sample allows TME to be used for real-time detection of hand and wrist movements from TMG data.

In this experiment, we investigate the performance of TME on the dataset¹ presented in the second experiment of [23]. The dataset was gathered from 9 healthy participants requested to

replicate the movements of a stimulus in the form of a 9-DOF hand model while wearing the tactile bracelet. Ground truth was obtained from the values of the animated hand model displayed on a monitor. This method has the drawback of possibly reducing the precision of the prediction of the intended activations due to the delay required by the participant to replicate the displayed movement. However, this approach allows the association of intended activations with input signal patterns in the case of amputees (since ground truth data can obviously not be collected by other means in this case). Each participant executed three times a sequence of six movements, namely wrist flexion, wrist extension, wrist supination, thumb flexion, index flexion and little-finger flexion. The data were recorded during the whole cycle of the stimulus, namely transition, activation, transition and relaxation phases, in order to obtain the whole range of activation from rest to complete finger and wrist movements (see [23] for more details).

The training dataset is composed of data recorded at zero and full activation. The testing dataset is composed of data recorded during the transition parts, containing the whole range of intermediate activation levels. Therefore, the evaluation of the performance of the model is compatible with the evaluation in forecasted studies with amputees, as they cannot provide accurate intermediate training data.

We compared the performance of vector-based and tensor-based algorithms, namely RR, ME and Gaussian process regression (GPR), as well as TRR and TME, on this dataset. The TMG data were centered for all the methods. The regularization parameter of RR and ME were fixed to 0.1. GPR was used with a radial basis function (RBF) kernel whose parameters were optimized using GPy [24]. Note that GPR with RBF kernel with an Euclidean distance measure was proved to reach good performances on this dataset in [23]. The rank of TRR was determined using 5-fold cross-validation and the regularization parameter $\lambda_{\mathbf{W}}$ was fixed to 0.1. For TME, as for ME, one expert was considered for each of the six finger and wrist movements. In order to facilitate the training process, we considered a common value R_e for the ranks R_{e_i} . The ranks R_g and R_e were determined using 5-fold cross-validation for $2 \leq R_g, R_e \leq 6$. The regularization parameters of TME, $\lambda_{\mathbf{W}}$ and $\lambda_{\mathbf{V}}$, were fixed by the experimenter as equal to 0.1. Note that small variations of the regularization parameters did not change significantly the results for the different regression methods.

Table I shows the mean and standard deviation over the 9 participants of the RMSE between the ground truth and the prediction for the aforementioned methods. We observe that taking into account the structure of the data in the regression process improves the quality of predictions as both TRR and TME outperform their vector counterparts. Moreover, taking into account the structure of the data allows a linear method (TRR) to achieve performances comparable to those obtained by non-linear methods. GP and TME achieve the best performance compared to the other methods, with TME obtaining a slightly lower RMSE (0.303 ± 0.074) than GP (0.305 ± 0.060). Moreover, TME obtained the minimum RMSE for 5 participants out of 9. Note that GP with RBF kernel slightly outperformed GP with linear ($0.455 \pm$), Matérn 32

¹The dataset is available online at <http://www.idiap.ch/paper/mdpi/data/exp2/>.

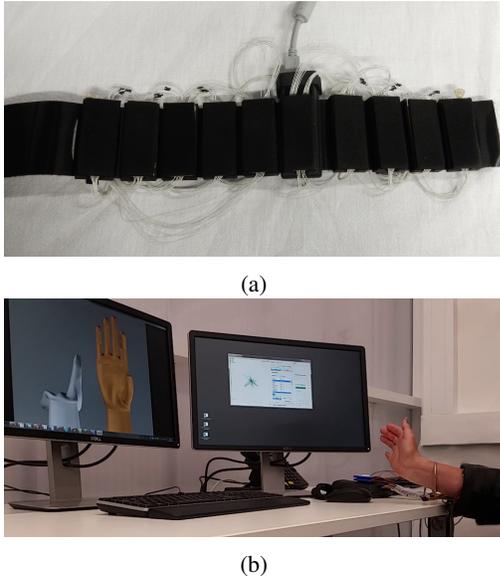


Fig. 5: (a) The TMG sensor used for the experiments. The bracelet is here rolled out, showing its 10 modules of 8×4 resistive cells. (b) Data collection [23]. The participant, wearing the tactile bracelet, imitates the grey animated hand model.

TABLE I: Performance comparison in terms of RMSE between different regression methods to predict fingers and wrist movements from TMG data.

| RR | ME | GP | TRR | TME |
|-----------------|-----------------|-----------------|-----------------|-----------------|
| 0.45 ± 0.07 | 0.33 ± 0.07 | 0.30 ± 0.06 | 0.35 ± 0.11 | 0.30 ± 0.07 |

(0.308 ± 0.061) and Matérn 52 (0.306 ± 0.059) kernels, therefore all the results are presented with RBF kernel.

Figure 6 shows an example of original and recovered activations for all the movements over the time for GP and TME. We observe that TME is generally recovering a more stable signal than GP when one movement is not activated. In the zones of zero activations, the signal recovered by GP tends to oscillate around zero. However, the signal recovered by TME can have a bigger delay than GP to detect an activation different than zero (see, e.g., Fig. 6, wrist extension).

Table II shows the average training and testing computation time for the tested regression methods. The computation times were measured using a non-optimized Python code on a laptop with 2.7GHz CPU and 32 GB of RAM. Despite the fact that TME converged with less than 10 iterations of the EM algorithm for all the participants, we observe that this method has the highest computation time for training. This is mainly due to the fact that a TLR model is optimized at each step of the EM algorithm. Therefore, the training time of TME could be improved by optimizing the training time of TLR in the same way as machine learning libraries do for standard logistic regression. However, this training time remains reasonable for the method to be applied in real-time with amputees. Moreover, a testing time of 1 ms is also reasonable for real-time applications with TME allowing predictions at a frequency > 50 Hz, as usually targeted by real-time detection of hand movements. As opposed to GP, the

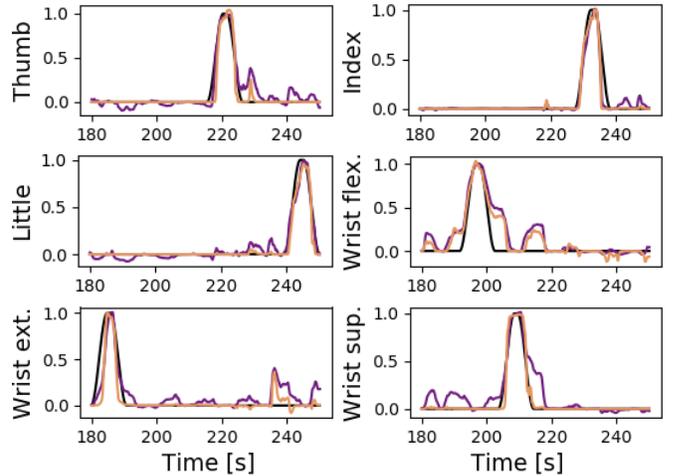


Fig. 6: Original and recovered activations of the different fingers and wrist movements over time. The whole range of activation is represented from 0 to 1 on the vertical axis for each movement. The ground truth is shown by black curves, while the signals recovered by TME and GP are displayed in yellow and purple, respectively.

TABLE II: Average training and testing computation time for the different regression methods. The methods are trained on ~ 1000 data samples. The training times of TRR and TME are computed for fixed ranks. The testing computation time is measured for one data sample. Training and testing computation time are expressed in seconds and milliseconds, respectively.

| | RR | ME | GP | TRR | TME |
|--------------|--------|------|-----|-------|-----|
| Training [s] | 0.02 | 22 | 3 | 15 | 560 |
| Testing [ms] | 0.0065 | 0.08 | 1.5 | 0.035 | 1.0 |

computation testing time of TME is independent of the number of training data and depends only on the number of experts. Therefore, TME would be adapted to real-time predictions independently on the number of provided training data.

C. Real-time teleoperation with tactile myography

To evaluate our method in a scenario closer to the end-user case, we conducted a real-time teleoperation experiment where 11 non-amputated participants (one female and ten males) controlled a robotic hand and arm based on the activation of the muscles on their forearm.

In the first part of the experiment, a protocol similar to the data collection of the experiment of Section IV-B was applied to collect TMG data associated with the hand postures of the participants. The tactile bracelet was placed on the forearm of the participant with the closing gap on the ulna bone. The participant, wearing the tactile bracelet and sitting in front of a monitor, was asked to replicate the movements of a model of the 24-DOFs dexterous motor hand of the Shadow Robot Company [25]. Similarly to the previous experiment, ground truth was obtained from the values of the animated hand model. Each participant executed four times the sequence of four movements, namely wrist flexion, wrist extension, power grasp and fingers extension. The participants were asked to perform the different movements in a relaxed way, particularly the fingers were relaxed during wrist movements.

Each stimulus follows a cycle of 14 s composed of a transition phase (2 s), an activation phase (6 s), a transition phase (2 s) and a relaxing phase (4 s). The data collected during the activation and relaxing phases, i.e., at zero and full activations, were used to train the regression models.

During the second part of the experiment, the participant teleoperated a Shadow robot hand mounted on a 7-DOFs Mitsubishi PA10 robot arm. S/he was sitting in front of the robotic system with the palm of the Shadow robot hand facing right, as showed in Figure 7a. The different movements taught to the model in the first part of the experiments were mapped to the robotic system as follows: wrist flexion and extension were used to move the arm forward (in the direction of the palm) and backward (in the direction of the back of the hand), respectively. Power grasp and fingers extension were used to close and open the Shadow robot hand. When wrist flexion or extension was detected above a certain activation threshold, the velocity of the robot arm was incremented in the corresponding direction proportionally to the detected activation. Similarly, the posture of the robotic hand was incremented proportionally to the activation of power grasp or fingers extension if they were detected above a predefined threshold. The detected activations were also displayed on the Shadow robot hand model as in the first part of the experiment.

At the beginning of the second part of the experiment, the participant could get used to the learned mapping by controlling the simulated Shadow robot hand for a few minutes. Then, while teleoperating the real robotic system, the participant was asked to control the arm in order to approach it close to an object placed on a cube, to grasp this object and to bring it to a specific location on the left (A) or on the right (B) side and to release it. The complete setup is showed in Figure 7a. Three objects with different diameters were considered, namely a chips cylinder ($\varnothing 75$ mm), a thin woodstick ($\varnothing 21$ mm) and a PET bottle ($\varnothing 63$ mm), as shown in Figure 7b. A total of 8 tasks were executed by each participant. The first six tasks consisted of bringing each object to A and then to B. Once a contact with the object was detected by the tactile fingertip sensors of the Shadow robot hand, the grasp pose was automatically maintained by the hand so that the subject could relax his/her fingers and focus on the wrist motion to steer the arm. The grasp pose was released as soon as a fingers extension command was detected. The maintenance of the grasp and the release were announced verbally by the system. The two last tasks consisted of bringing the PET bottle to A and to B without any holding assistance by the robot system. The time to complete each task was limited to two minutes. In case the object felt from the cube or was released out of the desired area, the experimenter replaced it at the initial position and the participant continued to execute the task in the remaining time. In case the participant could not control both the arm and the hand, e.g., if the arm was drifting continuously in one direction, the arm commands were disabled and the participant was requested to maintain a grasp on the object for 10 s before releasing it. Each participant tried to complete the 8 tasks with two different regression methods, namely TME and RR, trained on the data collected in the first part of the experiment. RR was chosen for comparison since it

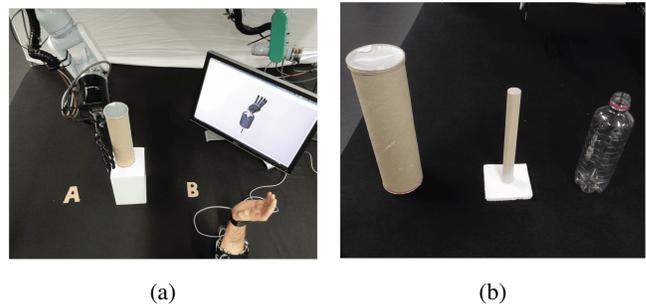


Fig. 7: (a) Setup of the teleoperation experiment. The participant is requested to grasp the chips cylinder and to place it on A or B. The distance between the initial position of the object and A or B is approximately 15 cm. The detected activations are displayed on the monitor by an animated hand model. (b) Objects used during the teleoperation experiment.

is considered as the baseline method for regression with TMG data. The order in which the two methods were tested was alternated between the participants.

Figure 8 shows snapshots of a participant executing different tasks. 6, respectively 7, out of the 11 participants were able to control both the arm and the hand during the whole experiment by using TME and RR, respectively. Note that the cases where the arm commands had to be disabled occurred mainly for the second tested method (4 participants out of 5 testing TME as second method and 3 out of 4 participants testing RR as second method), suggesting a decreasing of performance over time. One participant was not able to control both arm and hands for both methods and an other participant was able to control them for the 4 first tasks of the first tested method only.

The success rates, or ratios of successful tasks, for the case in which the participants controlled both robot arm and hand are presented in Table III. We observe that TME outperformed the performance RR by 15% when all the objects and both locations A and B are considered. The time needed to accomplish successful tasks were 55.6 ± 31.1 s and 53.9 ± 31.5 s for TME and RR, respectively, showing almost no difference between the two methods.

For both methods, the tasks involving the woodstick result in the lowest success rate. Due to the small diameter of this object, the arm had to be positioned very precisely and a complete grasp activation had to be detected in order to perform a successful grasp. Therefore, the tasks involving the woodstick were the hardest to complete for the participants. In the case of TME, the success rate for the chips cylinder is lower than for the PET bottle. This can be explained by the fact that a small activation of the grasp movement was sometimes detected when the participants were flexing their wrist to make the arm move in the direction of the object. However, the robotic hand had to be completely opened to be able to be placed around the chips cylinder before grasping, while it could still be placed around the bottle if a small grasping activation was detected. In the case of RR, the success rate diminishes for the bottle compared to the chips cylinder. This may suggest a stronger decreasing of performance over

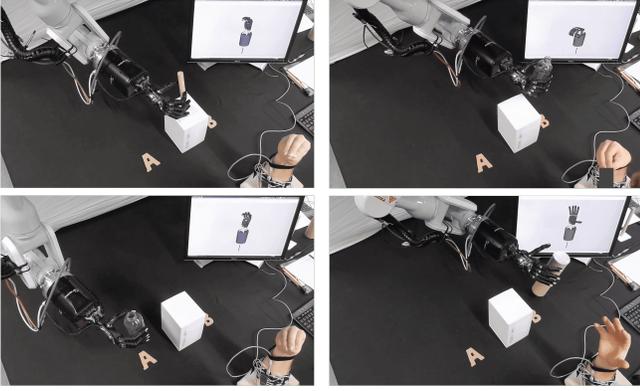


Fig. 8: Snapshots of the teleoperation experiment for different tasks. A participant is (a) grasping the woodstick, (b) bringing the bottle to B, (c) reaching A while holding the bottle, (d) releasing the chips cylinder at B. The detected activations are displayed by a Shadow robot hand model on the monitor.

TABLE III: Success rates over all the tasks and for each object in the case where the participants teleoperated both the robot arm and hand. The success rates are given in percent [%].

| | Total | Chips cylinder | Woodstick | Bottle | Bottle (no HA) |
|-----|-------|----------------|-----------|--------|----------------|
| TME | 45.8 | 41.7 | 16.7 | 66.7 | 58.3 |
| RR | 30.6 | 44.4 | 12.5 | 35.7 | 28.6 |

time with this method.

We observe that the success rates for the bottle are similar with and without the holding assistance activated for both methods. This result is particularly interesting as it shows that combinations of hand and wrist movements, namely grasping with wrist flexion and grasping with wrist extension in this experiment, can be detected while training only on individual movements. Moreover, both aforementioned combinations were equally detected as the number of completed tasks for each location A or B was similar, i.e., 4 and 3 successful tasks for A and B with TME and 2 for each location with RR. Moreover, some of the participant did not wait that the contact with the object was detected before bringing it to its final location. Therefore, they managed to complete other tasks without using the holding assistance.

Table IV shows the proportion of the failed tasks for which the time ran out during each of the task steps, namely grasping, moving and releasing the object, for TME and RR. We observe that the proportions are similar for both methods with the grasping step being the main cause of failure, followed by the moving step. Failures during grasping occurred mainly because the detected grasp activation was not sufficient to grasp the object or when it was activated too soon, therefore resulting in the object being pushed out of the support box. Failures while moving were due to difficulties to detect wrist flexion and extension or to the object falling down while the participant was trying to reach A or B. Finally, if the fingers extension movement was not detected properly while the holding assistance was activated, the opening of the hand was not triggered, resulting in failure to release the object.

The success rates for the case in which the arm commands

TABLE IV: Proportion of the failed tasks due to running out of time during the grasp, the displacement and the release of the object [%].

| | Grasping | Moving | Releasing |
|-----|----------|--------|-----------|
| TME | 57.7 | 30.8 | 11.5 |
| RR | 58.1 | 32.6 | 9.3 |

TABLE V: Success rates over all the tasks and for each object for the case in which the commands of the arm had to be disabled. The success rates are given in percent [%].

| | Total | Chips cylinder | Woodstick | Bottle | Bottle (no HA) |
|-----|-------|----------------|-----------|--------|----------------|
| TME | 38.8 | 40.0 | 0 | 40.0 | 40.0 |
| RR | 53.8 | 50.0 | 66.7 | 50.0 | 50.0 |

where disabled and the participant was only requested to grasp the object are presented in Table V. In this case, RR outperforms TME, especially for the tasks involving the woodstick. This can be explained by the fact that, in most of the cases, the arm commands had to be disabled for RR because the arm was drifting on the left or on the right, while it had to be disabled for TME because no activation was detected to move the arm, so that the participant could not position it to grasp the object. Generally, the detected activation of the grasp movement was also limited, therefore some tasks were difficult to achieve, especially those involving the woodstick.

V. DISCUSSION

The proposed TME model allows the structure of tensor-valued data to be taken into account in the regression problem. Overfitting can then be reduced, which is particularly important when only few tensor-valued training data are available. We showed the effectiveness of the approach to detect hand movements from TMG data, outperforming the other tested methods in an offline experiment and allowing participants to teleoperate a robotic arm and hand in real-time. Particularly, the method was able to successfully detect intermediate and combination of activation, while trained only with zero and complete individual movements. Notably, participants managed to activate wrist flexion or extension along with power grasp. This indicates that a holding assistance may not be required. However, some participants reported that the holding assistance was helpful as a feedback indicating that the grasp was effective or to make them feel more comfortable while teleoperating the arm as they could focus on one movement only.

A decrease of performance over time, indicated by the necessity of deactivating the arm commands while testing the second method, seemed to have occurred during the real-time experiment. Moreover, some participants reported that they felt that the control of the robotic arm and hand was harder to perform over time. Moreover, we qualitatively observed that this problem seemed to occur particularly for participants who were trying to apply high forces to execute the different movements. We hypothesize that this is due to small displacements of the TMG bracelet over time, inducing a shift of the testing data compared to the training data. This

problem could be overcome by improving the placement of the bracelet and by adapting the model over time. Techniques such as covariate shift adaptation in the case of sEMG [26] could be investigated.

It is important to emphasize the fact that the participants were able to adapt in some extent to the predictions of the method. They slightly modified their hand movements in order to obtain the desired action of the robotic arm and hand. Therefore, we observe a form of active learning, where the method learned from the training data, while the participants learned from the method in order to achieve the desired performance.

In both experiments, the ranks of the experts were given by a common value. The performance of TME may be further improved by selecting a specific rank for each expert. However, to avoid increasing computation time, automatic rank selection procedures have to be investigated. In particular, the automatic rank selection presented in [11] could be exploited to determine the rank of the expert models. The suggested method uses a $\ell_{1,2}$ norm regularization and optimizes the model with iteratively reweighted least squares (IRLS) algorithm. This approach seems promising as the authors reported in their experiments that the automatic procedure provided the same rank as the one selected by cross-validation.

VI. CONCLUSION

This paper presented an extension of mixture of experts to tensor-valued data. Our method brings together the advantages and robustness of mixture models and tensor methods. Therefore, it allows an efficient combination of predictions from experts specialized in different regions of the input space, while taking into account the structure of tensor-valued data in the soft space division and in the predictions of the experts. The data are efficiently exploited, so that a model trained with a small amount of training data is able to achieve good performances, while overcoming the overfitting problem. This is particularly important in robotics as the amount of training data is often small compared to the dimensionality of the data. The effectiveness of our model was illustrated with artificially generated data and in two experiments aiming at detecting hand movements by measuring the pressure induced by the muscles activity of the forearm with tactile myography. We showed that the testing computational time of the proposed model is low, due to a computational cost independent of the number of training data, therefore making it compatible with real-time robotic applications.

Future work will investigate automatic rank selection procedures with the objective of automatically determining all the ranks of the model, therefore avoiding the use of cross-validation in the training process. Moreover, we will investigate extensions of the proposed tensor-variate mixture of experts (TME) model to other applications in robotics, as well as to more complex models, such as hierarchical TME [17]. It is worth noting that the proposed TME model permits to incorporate structural information of the data as a special case of neural network. Extensions of this model could then also lead to interesting perspectives in the development of neural

networks for tensor data that would have better interpretability, that would allow training with smaller amount of data and that would provide better generalization results by avoiding overfitting.

ACKNOWLEDGMENT

We thank Guillaume Walck for his significant help on the preparation and execution of the experiment on the real robotic setup and for his suggestions during the preparation of the manuscript.

REFERENCES

- [1] T. Kolda and B. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, 2009.
- [2] Q. Zhao, G. Zhou, L. Zhang, and A. Cichocki, "Tensor-Variate Gaussian Processes Regression and Its Application to Video Surveillance," in *IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2014, pp. 1265–1269.
- [3] S. Calinon, "A tutorial on task-parameterized movement learning and retrieval," *Intelligent Service Robotics*, vol. 9, no. 1, pp. 1–29, January 2016.
- [4] F. Miwakeichi, E. Martínez-Montes, P. A. Valdés-Sosa, N. Nishiyama, H. Mizuhara, and Y. Yamaguchi, "Decomposing EEG data into space-time-frequency components using Parallel Factor Analysis," *NeuroImage*, vol. 22, no. 3, pp. 1035–1045, 2004.
- [5] Y. Washizawa, H. Higashi, T. Rutkowski, T. Tanaka, and A. Cichocki, "Tensor Based Simultaneous Feature Extraction and Sample Weighting for EEG Classification," in *Intl. Conf. on Neural Information Processing (ICONIP)*, 2010, pp. 26–33.
- [6] F. Kõiva, E. Riedenklaus, C. Viegas, and C. Castellini, "Shape Conformable High Spatial Resolution Tactile Bracelet for Detecting Hand and Wrist Activity," in *IEEE Intl. Conf. on Rehabilitation Robotics (ICORR)*, 2015, pp. 157–162.
- [7] A. Ozdemir, A. Zare, M. A. Iwen, and S. Aviyente, "Extension of PCA to higher order data structures: An introduction to tensors, tensor decompositions, and tensor PCA," *Proceedings of the IEEE*, vol. 106, no. 8, pp. 1341–1358, 2018.
- [8] H. Lu, K. N. Plataniotis, and A. N. Venetsanopoulos, "MPCA: Multilinear principal component analysis of tensor objects," *IEEE Transactions on Neural Networks*, vol. 19, no. 1, pp. 18–39, 2008.
- [9] D. Tao, X. Li, X. Wu, and S. J. Maybank, "General Tensor Discriminant Analysis and Gabor Features for Gait Recognition," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 29, no. 10, pp. 1700–1715, 2007.
- [10] Y. Tang, R. Salakhutdinov, and G. Hinton, "Tensor analyzers," in *Proceedings of the 30th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 28, no. 3, 2013, pp. 163–171.
- [11] W. Guo, I. Kotsia, and I. Patras, "Tensor learning for regression," *IEEE Transactions on Image Processing*, vol. 21, no. 2, pp. 816–827, 2012.
- [12] H. Zhou, L. Li, and H. Zhu, "Tensor Regression with Applications in Neuroimaging Data Analysis," *Journal of the American Statistical Association*, vol. 108, no. 502, pp. 540–552, 2013.
- [13] H. Hung and C. Wang, "Matrix variate logistic regression model with application to eeg data," *Biostatistics*, vol. 14, no. 1, pp. 189–202, 2013.
- [14] X. Tan, Y. Zhang, S. Tang, J. Shao, F. Wu, and Y. Zhuang, "Logistic Tensor Regression for Classification," in *Intelligent Science and Intelligent Data Engineering*, 2013, pp. 573–581.
- [15] M. Signorello, L. De Lathauwer, and J. A. K. Suykens, "A Kernel-Based Framework to Tensorial Data Analysis," *Neural Networks*, vol. 24, no. 8, pp. 861–874, 2011.
- [16] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, "Adaptive Mixture of Local Experts," *Neural Computation*, vol. 3, no. 1, pp. 79–87, 1991.
- [17] S. E. Yuksel, J. N. Wilson, and P. D. Gader, "Twenty years of mixture of experts," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, no. 8, pp. 1177–1193, 2012.
- [18] J. D. Carroll and J. J. Chang, "Analysis of individual differences in multidimensional scaling via an n-way generalization of "Eckart-Young" decomposition," *Psychometrika*, vol. 35, no. 3, pp. 283–319, 1970.

- [19] R. A. Harshman, "Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multimodal factor analysis," *UCLA Working Papers in Phonetics*, vol. 16, pp. 1–84, 1970.
- [20] J. M. Hahne, F. Bießmann, N. Jiang, H. Rehbaum, D. Farina, F. C. Meinecke, K. R. Müller, and L. C. Parra, "Linear and nonlinear regression techniques for simultaneous and proportional myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 2, pp. 269–279, 2014.
- [21] S. L. Phillips and W. Craelius, "Residual kinetic imaging: a versatile interface for prosthetic control," *Robotica*, vol. 23, no. 3, pp. 277–282, 2005.
- [22] C. Nissler, M. Connan, M. Nowak, and C. Castellini, "Online tactile myography for simultaneous and proportional hand and wrist myocontrol," in *Proceedings of Myoelectric Control Symposium (MEC)*, 2017.
- [23] N. Jaquier, M. Connan, C. Castellini, and S. Calinon, "Tensor decompositions and applications," *Technologies*, vol. 4, no. 5, 2017.
- [24] GPY, "GPY: A Gaussian process framework in python," <http://github.com/SheffieldML/GPy>, since 2012.
- [25] "Shadow robot company," <https://www.shadowrobot.com/>.
- [26] M. Vidovic, H. Hwang, S. Amsüss, J. M. Hahne, D. Farina, and K.-R. Müller, "Improving the Robustness of Myoelectric Pattern Recognition for Upper Limb Prostheses by Covariate Shift Adaptation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 24, no. 9, 2016.



Noémie Jaquier received the B.Sc on Microengineering and the M.Sc in Robotics with a Minor in Computational Neurosciences from the Ecole Polytechnique Federale de Lausanne (EPFL) in 2014 and 2016, respectively. She is a Ph.D student at the Idiap Research Institute. Her research interest cover human-robot interactions, machine learning and differential geometry in robotics. Personal webpage: <http://njaquier.ch>



Robert Haschke received the diploma and PhD in Computer Science from the University of Bielefeld, Germany, in 1999 and 2004, working on the theoretical analysis of oscillating recurrent neural networks. Since then, his work focuses more on robotics, still employing neural networks and machine learning methods wherever possible. Robert is currently heading the Robotics Group within the Neuroinformatics Group, striving to enrich the dexterous manipulation skills of bimanual, multi-fingered robots through interactive learning. His fields of research include neural networks, machine learning, cognitive bimanual robotics, grasping and manipulation with multi-fingered dexterous hands and multimodal feedback (vision + tactile), planning, and software integration.



Sylvain Calinon received the Ph.D. degree in robotics from Learning Algorithms and Systems Laboratory, Ecole Polytechnique Fdrale de Lausanne (EPFL), Lausanne, Switzerland, in 2007. He is a Senior Researcher at the Idiap Research Institute, and a Lecturer at the EPFL. From 2009 to 2014, he was a Team Leader at the Department of Advanced Robotics, Italian Institute of Technology. From 2007 to 2009, he was a Postdoc at EPFL. Dr Calinon is the author of 100+ publications in robot learning and human-robot interaction, with recognition including Best Paper Awards in the journal of Intelligent Service Robotics (2017) and at the IEEE Ro-Man2007, and Best Paper Award Finalist at ICRA2016, ICIRA2015, IROS2013 and Humanoids2009. He currently serves as an Associate Editor in IEEE Transactions on Robotics and IEEE Robotics and Automation Letters. Website: <http://calinon.ch>