



**HAL**  
open science

# Scalable magnetic resonance fingerprinting: Incremental inference of high dimensional elliptical mixtures from large data volumes

Geoffroy Oudoumanessah, Thomas Coudert, Carole Lartizien, Michel Dojat,  
Thomas Christen, Florence Forbes

## ► To cite this version:

Geoffroy Oudoumanessah, Thomas Coudert, Carole Lartizien, Michel Dojat, Thomas Christen, et al.. Scalable magnetic resonance fingerprinting: Incremental inference of high dimensional elliptical mixtures from large data volumes. 2024. hal-04836864

**HAL Id: hal-04836864**

**<https://hal.science/hal-04836864v1>**

Preprint submitted on 13 Dec 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.



Distributed under a Creative Commons Attribution 4.0 International License

# SCALABLE MAGNETIC RESONANCE FINGERPRINTING: INCREMENTAL INFERENCE OF HIGH DIMENSIONAL ELLIPTICAL MIXTURES FROM LARGE DATA VOLUMES

BY GEOFFROY OUDOUMANESSAH<sup>1,3,2,a</sup> , THOMAS COUDERT<sup>3,d</sup>   
 CAROLE LARTIZIEN<sup>2,c</sup> , MICHEL DOJAT<sup>1,3,e</sup>   
 THOMAS CHRISTEN<sup>3,f</sup>  AND FLORENCE FORBES<sup>1,b</sup> 

<sup>1</sup>Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France, <sup>a</sup>[geoffroy.oudoumanessah@inria.fr](mailto:geoffroy.oudoumanessah@inria.fr);  
<sup>b</sup>[florence.forbes@inria.fr](mailto:florence.forbes@inria.fr)

<sup>2</sup>Univ. Lyon, CNRS, Inserm, INSA Lyon, UCBL, CREATIS, UMR5220, U1294, F-69621, Villeurbanne, France,  
<sup>c</sup>[carole.lartizien@creatis.insa-lyon.fr](mailto:carole.lartizien@creatis.insa-lyon.fr)

<sup>3</sup>Univ. Grenoble Alpes, Inserm U1216, CHU Grenoble Alpes, Grenoble Institut des Neurosciences, 38000 Grenoble, France,  
<sup>d</sup>[thomas.coudert@inserm.fr](mailto:thomas.coudert@inserm.fr); <sup>e</sup>[michel.dojat@inserm.fr](mailto:michel.dojat@inserm.fr); <sup>f</sup>[thomas.christen@univ-grenoble-alpes.fr](mailto:thomas.christen@univ-grenoble-alpes.fr)

Magnetic Resonance Fingerprinting (MRF) is an emerging technology with the potential to revolutionize radiology and medical diagnostics. In comparison to traditional magnetic resonance imaging (MRI), MRF enables the rapid, simultaneous, non-invasive acquisition and reconstruction of multiple tissue parameters, paving the way for novel diagnostic techniques. In the original *matching* approach, reconstruction is based on the search for the best matches between *in vivo* acquired signals and a dictionary of high-dimensional simulated signals (fingerprints) with known tissue properties. A critical and limiting challenge is that the size of the simulated dictionary increases exponentially with the number of parameters, leading to an extremely costly subsequent matching. In this work, we propose to address this scalability issue by considering probabilistic mixtures of high-dimensional elliptical distributions, to learn more efficient dictionary representations. Mixture components are modelled as flexible elliptic shapes in low dimensional subspaces. They are exploited to cluster similar signals and reduce their dimension locally cluster-wise to limit information loss. To estimate such a mixture model, we provide a new incremental algorithm capable of handling large numbers of signals, allowing us to go far beyond the hardware limitations encountered by standard implementations. We demonstrate, on simulated and real data, that our method effectively manages large volumes of MRF data with maintained accuracy. It offers a more efficient solution for accurate tissue characterization and significantly reduces the computational burden, making the clinical application of MRF more practical and accessible.

**1. Introduction.** Traditional Magnetic Resonance (MR) imaging relies on an analytical resolution of dynamical equations using conventional tuning of the MR hardware through sequences of pulses, each characterized by different values of parameters such as the flip angle and repetition time. Standard quantitative MRI (qMRI) methods are based on a single sequence for a single parameter measurement at a time. This leads to high scan times for multi-parametric protocols as each parameter estimate involves one MR sequence. A recent approach named Magnetic Resonance Fingerprinting (MRF, [Ma et al. \(2013\)](#)) has been developed to overcome these limitations. The MRF protocol involves fast undersampled acquisitions with time-varying parameters defining the MRF sequence that produces temporal signal evolutions (named *fingerprints*) in each voxel. In the original proposal, a dictionary

---

*Keywords and phrases:* Dimension reduction, Clustering, Incremental learning, High dimensional mixture models, Elliptical distributions, Expectation Maximization algorithm.

search approach is used to compare the *in vivo* fingerprints with millions of numerical simulations of MR signals for which the associated parameters are known. These millions of simulated signals compose the so-called dictionary. The values of the parameters corresponding to the closest simulated signals or *matches* are then assigned to the associated *in vivo* voxels, allowing the simultaneous reconstruction of multiple quantitative maps (images) from extremely undersampled raw images, using only one single sequence, thus saving considerable acquisition time (Poorman et al., 2020; McGivney et al., 2020). From this, standard relaxometry MRF allows reconstructing parameter maps for relaxation times  $T_1$  and  $T_2$ , over the whole human brain ( $1 \text{ mm}^3$  spatial resolution) in 3 min (Ye et al., 2017; Gu et al., 2018) compared to 30 min for a standard T1/T2 exam. Moreover, the flexibility of the numerical simulations enables correction of system imperfections as well as some patient motions by including them in the model and post-processing pipelines (Bipin Mehta et al., 2019). Thus, MRF could be a game changer for emergency patients who need to complete exams in a few minutes. The power of the MRF approach is not limited to the estimation of relaxation times, in theory, it allows the measurement of any parameter that influences nuclear magnetization (e.g., microvascular networks), and could be added to the simulation model (Wang et al., 2019a; Coudert et al., 2024). However, increasing the number of estimated parameters, even moderately, induces the design of more complex sequences and increased reconstruction times, from hours to days. This limits the clinical application of high-dimensional MRF and necessitates the development of innovative processing methods. Consequently, a significant focus is on improving MRF reconstruction methods, as reviewed by Tippareddy et al. (2021) and Monga et al. (2024). In this work, we propose to focus on reducing the reconstruction times of MRF when more than the main two parameters,  $T_1$ ,  $T_2$ , are involved, including the addition of  $\delta f$  the frequency offset, the sensitivity of the magnetic field  $B_1$ , cerebral blood volume (CBV), and microvascular geometry (e.g., vessel radius denoted as  $R$ ).

MRF reconstruction is first recast as an inverse problem that can be solved using different approaches as recalled in the next section. All approaches make use, at some stage, of a dictionary of simulated pairs (parameters, signal), which represent our knowledge of the link between tissue parameters and MR time series, through a so-called *direct* or *forward* model. The dictionary is then either used to learn an *inversion* operator, from signal to tissue parameters, or to search for the best fits between observed signals and simulated ones. The approaches scalability relies thus greatly on their ability to extract efficiently the information encoded in simulations. Efficiency has different aspects: for search-type methods, the dictionary should not be too big, while for learning-based methods, the dictionary should be informative enough. These potentially opposite requirements call for efficient representations of simulated data. We propose to explore a divide-and-conquer strategy by introducing the framework of High Dimensional Mixtures of Elliptical Distributions (HD-MED). Probabilistic mixtures of high-dimensional elliptical distributions allow us to learn more efficient dictionary representations. Mixture components are modeled as flexible elliptical shapes in low-dimensional subspaces. They are exploited to cluster similar signals and reduce their dimension locally, at the cluster level, to limit information loss. To estimate such a mixture model, we provide a new incremental algorithm capable of handling large numbers of signals, allowing us to go far beyond the hardware limitations encountered by standard implementations. We demonstrate, on simulated and real data, that our method effectively manages large volumes of MRF data with maintained accuracy. It offers a more efficient solution for accurate tissue characterization and significantly reduces the computational burden, making the clinical application of MRF more practical and accessible.

In the rest of the paper, we first recall in Section 2 the two main types of approaches that have been investigated in the literature on MRF reconstruction. We then review related work on efficient dictionary representations and specify our contributions in Section 3. The

mixtures of elliptical distributions are presented in Section 4 with their use for dimension reduction in Section 5. The incremental algorithm proposed for the mixture estimation is presented in Section 6 and illustrated on MRF in Section 7.

**2. MRF reconstruction as an inverse problem.** An inverse problem refers to a situation where one aims to determine the causes of a phenomenon from experimental observations of its effects. In MRF, the goal is to infer from an observed signal a set of tissue characteristics or parameter values that best account for the signal. Such a resolution generally starts with modeling the phenomenon under consideration, which is called the "direct" or "forward problem." It is generally assumed that at least the numerical evaluation of the forward model is available because experts have designed equations that can be solved either analytically or numerically. The most common use of the forward model is via a simulator that allows the creation of a database  $\mathcal{D}_f$ , usually referred to as a dictionary, of  $N$  signals  $\mathbf{y}_1, \dots, \mathbf{y}_N$  with  $\mathbf{y}_i \in \mathcal{Y} \subset \mathbb{R}^M$ , generated (stored or computed on the fly) by running the theoretical (physical) model  $f$  for many different tissue parameter values  $\mathbf{t}_1, \dots, \mathbf{t}_N$  with  $\mathbf{t}_i \in \mathcal{T} \subset \mathbb{R}^L$  and  $\mathbf{y}_i = f(\mathbf{t}_i)$ . The generated tissue parameter values then only partially represent the full space  $\mathcal{T}$  of possible values and correspond to a discrete grid in the full space. In this context, we can distinguish two types of methods, referred to below as optimization and learning approaches.

*2.1. Optimization vs learning (or regression) approaches.* Optimization approaches include the most used method in MRF, namely dictionary matching, often called, in other domains, grid search, look-up-table or k-nearest neighbors. They consist of minimizing over parameters  $\mathbf{t}$  a merit function  $d$  expressing the similarity between the observed signal  $\mathbf{y}_{obs}$  and simulated signal  $f(\mathbf{t})$ ,

$$(1) \quad \hat{\mathbf{t}} \in \arg \min_{\mathbf{t} \in \mathcal{T}} d(\mathbf{y}_{obs}, f, \mathbf{t}).$$

Typically,  $d(\mathbf{y}_{obs}, f, \mathbf{t}) = d(\mathbf{y}_{obs}, f(\mathbf{t})) = \|\mathbf{y}_{obs} - f(\mathbf{t})\|^2$ . Solutions are searched in the full  $\mathcal{T}$  space but solutions could also be penalized as done in grid search methods. Indeed, in grid search, the previous full search is replaced by a simpler look-up or matching operation making use of the database  $\mathcal{D}_f$ , often created beforehand offline. The search space is significantly reduced from a continuous space  $\mathcal{T}$  to a discrete and finite  $\mathcal{D}_f$ . The speed gain is significant in comparison to traditional optimization methods as retrieving a value from memory is often faster than undergoing an expensive usually iterative computation. Their disadvantage is the instability of solutions. Many questions remain on how to choose the merit function, how many  $\mathbf{y}_n$  in the look up table have to be kept to estimate parameters, how to choose the look-up table, *etc.* When the number of parameters is small, grid search is suitable and can provide very good predictions. However, for even moderate numbers of parameters, the required number of elements in the dictionary renders grid search either intractable or inaccurate. The technique is not amortized, for each new  $\mathbf{y}_{obs}$ , we have to compute the matching score  $d(\mathbf{y}_{obs}, \mathbf{y}_n)$  for all  $\mathbf{y}_n$  in the dictionary. When the dimension of  $\mathbf{t}$  increases, the dictionary (and  $N$ ) has to be larger too for better accuracy. The computation of  $N$  matching scores can become time consuming.

Regression or learning methods are more efficient in that sense and usually have better amortization properties. In contrast to the previous ones, this category of methods has the advantage of adapting easily to provide tractable solutions in the case of massive inversions of high-dimensional data. The main principle is to transfer the computational cost and time from individual pointwise predictions to the learning of a global inverse operator from  $\mathcal{D}_f$ . The advantage is that once the operator is learned, it can be used, at negligible cost, for very large numbers of new signals. Then, the nature of the inverse operator needs to be specified.

Traditional learning or regression methods are not specifically designed for high dimensional data but there has been a large literature covering this case, see *e.g.* Giraud (2014) for a review. In MRF, a popular approach from Cohen, Zhu and Rosen (2018), uses a four-layer fully connected network (DRONE) to learn the dictionary signals for reconstructing  $T_1$  and  $T_2$  parameters. Even though DRONE provides good results, this perceptron-based model loses the temporal coherence of the signal. Recently, Cabini et al. (2024) proposed a Recurrent Neural Network (RNN) with long-short term memory (LSTM) blocks, which yields better reconstruction results for  $T_1$  and  $T_2$  with more robustness towards noisy acquisitions. However, when more parameters need to be estimated, such as vessel oxygenation or radius (Christen, Bolar and Zaharchuk, 2013; Christen et al., 2014), the dimensionality of the signals dramatically increases. Barrier et al. (2024) demonstrated that a simple RNN is prone to catastrophic forgetting, where the RNN well estimates the beginning of the signals but learns the end less effectively. To mitigate this issue, they proposed using bidirectional LSTM (bi-LSTM) blocks within the RNN architecture, ensuring that both the beginning and the end of the signal are efficiently learned by the new bi-LSTM blocks. Despite the good results in reconstructing  $T_1$ ,  $T_2$ , and vascular parameters, bi-LSTM still faces challenges in learning the middle part of the signals.

Generally, the dictionary is seen as a collection of simulated signals with no particular spatial correlation. Another way to simulate signals is to acquire real parametric maps from different subjects and then simulate the MRF images. This approach was first proposed by Soyak et al. (2021) and later improved by Gu et al. (2024). Both studies utilize a UNet (Ronneberger, Fischer and Brox, 2015) to infer  $T_1$  and  $T_2$  directly using the entire MRF image as an input to the network, preserving valuable spatial information. Given the high dimensionality of the signals, the authors proposed adding attention layers (Vaswani et al., 2017) to focus on the most important dimensions of the signals. More recently, Li and Hu (2024) highlighted the limitations of using CNN, which have a restricted receptive field and capture spatial information only locally. To overcome these limitations, the authors proposed using a Local-Global vision Transformer to capture spatial information globally as well. However, capturing spatial information has a cost. Indeed, one needs to acquire from a large group of subjects multiple  $T_1$  and  $T_2$  maps, which takes about 30 min for a complete exam, making the data acquisition process much more costly than using more conventional dictionaries. Additionally, since  $T_1$  and  $T_2$  need to be acquired at two different times, this method introduces more errors from the registration of the acquired maps.

Finally, Boux et al. (2021) also proposed to use GLLiM (Deleforge, Forbes and Horaud, 2015), a model that casts MRF reconstruction into a Bayesian inverse problem and then solves it using a learning approach. This method takes into account the high-dimensional property of MRF signals by defining the low-dimensional variables as the regressors, which in our case are the tissue parameters. By doing so, they start learning the *low-to-high* regression model from which they can derive the forward model parameters and then the *high-to-low* regression model, from MRF signals to tissue parameters, as desired.

**3. Related work and positioning.** In practice, acquired MRF acquisitions come as 4D matrix, made of a time series of 3D MRI images where each voxel contains the acquired, potentially long, fingerprint signal. In this work, we focus on designing efficient representations of large highly precise grids of simulated signals counterparts.

3.1. *Parcimonious representations of dictionaries.* The curse of dimensionality goes with what is often called the bless of dimensionality, which refers to the fact that in high dimensional data sets, useful information actually leaves in much smaller dimensional parts of the data space. One approach to the MRF reconstruction problem is then to reduce the

dimension of the dictionary beforehand to reduce the matching or learning cost. The first to propose such an approach for efficient matching were McGivney et al. (2014), who applied Singular Value Decomposition (SVD) to the dictionary of signals. Once the decomposition is learned, one can project any new acquisition into the SVD low-dimensional subspace. However, when the size of the dictionary increases, computing the SVD becomes costly as it requires loading the complete dictionary into fast-access memory (*e.g.* the RAM). To address this, Yang et al. (2018) proposed using randomized SVD (Halko, Martinsson and Tropp, 2011). Golbabaee et al. (2019) also suggested applying SVD to the dictionary before training a neural network. Other methods proposed a non-Euclidian analysis of the dictionary space projecting the signals into a lower-dimensional manifold (Li and Hu, 2023). Reducing the dimension of high-dimensional dictionaries assumes that most of the information in the signals can be captured and represented in a much lower-dimensional subspace. Classical techniques include principal component analysis (PCA, Jolliffe and Cadima (2016)), probabilistic principal component analysis (PPCA, Tipping and Bishop (1999a)), factor analyzers (FA), sparse models (Zou, Hastie and Tibshirani, 2006; d’Aspremont et al., 2007; Archambeau and Bach, 2008), and newer methods such as diffusion maps (Coifman and Lafon, 2006). More flexible approaches are based on mixtures of the previous ones, such as mixtures of factor analyzers (MFA, McLachlan, Peel and Bean (2003)) introduced by Ghahramani and Hinton (1997) and extended by McLachlan, Peel and Bean (2003); Baek and McLachlan (2011), and mixtures of PPCA (MPPCA, Tipping and Bishop (1999b); Xu, Balzano and Fessler (2023)) with recent generalizations (Hong et al., 2023; Xu, Balzano and Fessler, 2023). Another mixture approach is called HDDC in Bouveyron, Girard and Schmid (2007) and HD-GMM in Bouveyron and Brunet-Saumard (2014) for High Dimensional Gaussian Mixture Models, which encompass many forms of MFA and MPPCA and generalize them. In particular HD-GMM can be used to obtain multiple low-dimensional subspaces of different dimensions. For a review on high-dimensional clustering via mixtures, see Bouveyron and Brunet-Saumard (2014). Such a divide-and-conquer strategy has been used in MRF by first creating a global clustering of the signals and then reducing the dimension. This method was applied by Cauley et al. (2015), who performed K-way partitioning of the dictionary before using multiple cluster-associated PCA to reduce the dimension in each part. When matching a new signal, one first determines which cluster it belongs to and then applies the projection learned by the associated PCA model. More recently, Ullah et al. (2023) proposed a simpler approach: applying a clustering algorithm to the dictionary before utilizing GPUs for the *dictionary matching* enabling fast matching without the need for dimensionality reduction. However this method still faces issues when the number of parameters to estimate is larger than 3.

Most of these methods are designed for batch data and are thus sensitive to hardware limits such as memory, restricting the amount of data they can process. For instance, some dictionaries exceed terabyte sizes. A simple solution is to down-sample data sets before processing, potentially losing useful information. Another approach is to design incremental, also referred to as online, variants that handle data sequentially in smaller groups. A number of incremental approaches exist for dimension reduction techniques, see the recent SHASTA-PCA (Gilman et al., 2023) and references therein, or Balzano, Chi and Lu (2018) for a review. To our knowledge, much fewer solutions exist for mixtures. Estimation of such models is generally based on maximum likelihood estimation via the Expectation-Maximization (EM) algorithm (McLachlan and Krishnan (2008)). A preliminary attempt for an incremental MPPCA can be found in Bellas et al. (2013) but it is based on heuristic approximations of the EM steps.

**3.2. Contribution.** Herein, we propose to explore this divide-and-conquer strategy by introducing the framework of High Dimensional Mixtures of Elliptical Distributions (HD-MED) to get over the hump of estimating more than 2 parameters, *e.g.*  $T_1$ ,  $T_2$ ,  $\delta f$ ,  $B_1$ , CBV,

*R.* Considering the reconstruction of 6 parameters makes the associated dictionaries both large in size (an order of 13 terabyte of signals to considerably represent the tissues heterogeneity) and dimension, making our proposal twofold. Building on HD-MED, a generalization of HD-GMM, we show how they can be used to simultaneously compress and cluster large-scale high-dimensional MRF dictionaries, and more generally any dataset. We derive a new incremental algorithm, based on a principled EM framework, to learn such a model from potentially very large data volumes. We demonstrate the effectiveness of our approach on MRF reconstruction, showing results comparable to the high-dimensional dictionary matching referred to as *full-matching* in the next sections. This approach allows us to exceed the resolution, size used in current implementations and to reconstruct a larger number of MR parameter maps with improved accuracy, thereby advancing the clinical feasibility of MRF.

**4. High Dimensional Mixtures of Elliptical Distributions.** Elliptical distributions represent a family of distributions that contains Gaussian distributions but also heavy-tailed distributions, such as the Student distributions, that are often used as a more robust alternative to the Gaussian family. In this paper, we consider a sub-class of elliptical distributions, which can be expressed as infinite mixtures of Gaussian distributions.

4.1. *Gaussian scale mixture (GSM) distributions.* Scale mixtures of multivariate Gaussian distributions are an important subclass of elliptical distributions whose definition is recalled below. Scale mixtures of Gaussians share good properties with Gaussian distributions. They are tractable, lead to tractable inference procedures and provide more robust results, in contrast to Gaussian distributions that usually suffer from sensitivity to outliers.

DEFINITION 4.1 (Elliptical Distributions (ED)). A continuous random vector  $\mathbf{Y} \in \mathbb{R}^M$  follows a multivariate elliptical symmetric distribution if its probability density function (pdf)  $p(\mathbf{y})$  is of the following form (see [Cambanis, Huang and Simons \(1981\)](#) or [Kelker \(1970\)](#)),

$$(2) \quad p(\mathbf{y}) = C_{p,g} |\Sigma|^{-1/2} g\left((\mathbf{y} - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{y} - \boldsymbol{\mu})\right),$$

where  $\Sigma \in \mathbb{R}^{M \times M}$  is the scale matrix with determinant  $|\Sigma|$ ,  $\boldsymbol{\mu} \in \mathbb{R}^M$  is the location or mean vector,  $C_{p,g}$  is a normalizing constant such that the pdf  $p(\mathbf{y})$  integrates to one. The non-negative function  $g$  is called the density generator and determines the shape of the pdf. When  $\mathbf{Y}$  has density (2), we write  $\mathbf{Y} \sim \mathcal{E}_M(\boldsymbol{\mu}, \Sigma, g)$ .

Note that the scale matrix  $\Sigma$  is not necessarily equal to the covariance matrix,  $\Sigma$  is proportional to the covariance matrix if the latter exists. The pdf of the multivariate normal distribution is a special case of ED with  $g(u) = \exp(-u^2/2)$ . Another member of the elliptical family is the multivariate Student distribution. This distribution is well-studied in the literature ([Kotz and Nadarajah, 2004](#)) and admits a useful representation as a Gaussian scale mixture. Denoting  $\mathcal{N}_M(\mathbf{y}; \boldsymbol{\mu}, \Sigma)$  a  $M$ -variate Gaussian distribution with mean  $\boldsymbol{\mu}$  and covariance matrix  $\Sigma$ , a Gaussian scale mixture distribution is a distribution of the following form.

DEFINITION 4.2 (Gaussian scale mixture distributions (GSM)). If  $\boldsymbol{\mu}$  is a  $M$ -dimensional vector,  $\Sigma$  is a  $M \times M$  positive definite symmetric matrix and  $f$  is a pdf of a univariate positive variable  $W \in \mathbb{R}^+$ , then the  $M$ -dimensional density given by

$$(3) \quad p(\mathbf{y}) = \int_{\mathbb{R}^+} \mathcal{N}_M\left(\mathbf{y}; \boldsymbol{\mu}, \frac{\Sigma}{w}\right) f(w) dw$$

is said to be an infinite mixture of scaled Gaussians or Gaussian scale mixture (GSM) with mixing distribution function  $f$ . If vector  $\mathbf{Y}$  has density (3), we still write  $\mathbf{Y} \sim \mathcal{E}_M(\boldsymbol{\mu}, \Sigma, f)$  and refer to  $W$  as the mixing variable.

In practice, we will consider mixing distribution  $f_\theta$  that depends on some parameter  $\theta$  and also write  $\mathbf{Y} \sim \mathcal{E}_M(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \theta)$ . As already mentioned, famous GSM distributions include, the multivariate Student distribution (when  $f_\theta$  is the pdf of a Chi2 variable), the Pearson type VII distribution (when  $f_\theta$  is a gamma distribution) and the generalized Gaussian (when a power of  $W$  follows the gamma distribution). It is straightforward to see that GSM are elliptical distributions. However, not all elliptical distributions can be reduced to scale mixtures. For the previous reason, characterization and a way to represent elliptical distributions as GSM are very valuable. In [Gómez-Sánchez-Manzano, Gómez-Villegas and Marín \(2006\)](#), conditions are given under which elliptical distributions are GSMs. The issue of finding the corresponding mixing distribution is also addressed. An illustration of these results for generalized Gaussian distributions is given by [Gomez, Gomez-Villegas and Marin \(2008\)](#).

**4.2. Mixtures of High Dimensional GSM.** Consider a data set of  $N$  independent observations  $\{\mathbf{y}_1, \dots, \mathbf{y}_N\} \in \mathbb{R}^M$  assumed to be *i.i.d.* realizations of a random vector  $\mathbf{Y} \in \mathbb{R}^M$ . In addition, the data set is assumed to be made of  $K$  groups to be discovered. In this work, we consider finite mixtures of GSM, assuming that the underlying subsets are distributed according to different GSM in different proportions. In addition, to handle potentially high-dimensional observations with GSM, we propose a specific parameterization of the scale matrices.

**DEFINITION 4.3 (Finite mixture model).**  $\mathbf{Y}$  follows a finite mixture model if its pdf writes as

$$(4) \quad p(\mathbf{y}) = \sum_{k=1}^K \pi_k f_k(\mathbf{y}),$$

where  $\pi_k \in [0, 1]$  are the mixing weights that sum to one, and  $f_k$  is the conditional pdf of the  $k^{\text{th}}$  mixture component.

To simplify the notation, a finite mixture model where each pdf  $f_k$  is the pdf of a GSM distribution (3),  $\mathcal{E}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \boldsymbol{\theta}_k)$ , is referred to as a mixture of ED (MED). The number of parameters in MED grows quadratically with the dimension  $M$  due to the scale matrices  $\boldsymbol{\Sigma}_k$ . For large  $M$ , this can be problematic for the mixture estimation from a data set. In the Gaussian mixture case, to reduce the number of parameters, [Bouveyron, Girard and Schmid \(2007\)](#) proposed a family of parsimonious Gaussian mixture models, using the eigenvalues decomposition of the covariance matrices. We extend this idea to MED by reparametrizing the scale matrices  $\boldsymbol{\Sigma}_k$  as follows,

$$(5) \quad \boldsymbol{\Sigma}_k = \mathbf{D}_k \mathbf{A}_k \mathbf{D}_k^T,$$

where  $\mathbf{D}_k$  is a  $M \times M$  orthogonal matrix which contains the eigenvectors of  $\boldsymbol{\Sigma}_k$  and  $\mathbf{A}_k$  is a  $M \times M$  diagonal matrix that contains the associated eigenvalues in decreasing order. The key idea introduced by [Bouveyron, Girard and Schmid \(2007\)](#) is to consider that each cluster lies in a low-dimensional subspace of dimension  $d_k < M$ , which can be expressed by assuming that

$$(6) \quad \mathbf{A}_k = \text{diag}(a_{k1}, \dots, a_{kd_k}, b_k, \dots, b_k),$$

where  $a_{k1}, \dots, a_{kd_k}$  are the  $d_k$  largest eigenvalues of  $\boldsymbol{\Sigma}_k$  and  $b_k$  is a small negligible value. The  $d_k$  eigenvectors associated to the first  $d_k$  eigenvalues  $\{a_{k1}, \dots, a_{kd_k}\}$  define a cluster-specific subspace  $\mathbb{E}_k$ , which captures the main cluster shape. The orthogonal subspace is denoted by  $\mathbb{E}_k^\perp$ . Let  $\tilde{\mathbf{D}}_k$  consists of the  $d_k$  first columns of  $\mathbf{D}_k$  supplemented by  $(M - d_k)$

zero columns and  $\bar{\mathbf{D}}_k = (\mathbf{D}_k - \tilde{\mathbf{D}}_k)$ . It follows that  $P_k(\mathbf{y}) = \tilde{\mathbf{D}}_k \tilde{\mathbf{D}}_k^T (\mathbf{y} - \boldsymbol{\mu}_k) + \boldsymbol{\mu}_k$  and  $P_k^\perp(\mathbf{y}) = \bar{\mathbf{D}}_k \bar{\mathbf{D}}_k^T (\mathbf{y} - \boldsymbol{\mu}_k) + \boldsymbol{\mu}_k$  are the projections of  $\mathbf{y}$  on  $\mathbb{E}_k$  and  $\mathbb{E}_k^\perp$  respectively. This parameterization allows to handle high dimensional data in a computationally efficient way. For instance the quadratic form of the Mahalanobis distance, appearing in the generator  $g$  in (2), writes

$$\begin{aligned}
(\mathbf{y} - \boldsymbol{\mu}_k) \mathbf{D}_k \mathbf{A}_k^{-1} \mathbf{D}_k^T (\mathbf{y} - \boldsymbol{\mu}_k) &= (\mathbf{y} - \boldsymbol{\mu}_k)^T \tilde{\mathbf{D}}_k \mathbf{A}_k^{-1} \tilde{\mathbf{D}}_k^T (\mathbf{y} - \boldsymbol{\mu}_k) \\
&\quad + (\mathbf{y} - \boldsymbol{\mu}_k)^T \bar{\mathbf{D}}_k \mathbf{A}_k^{-1} \bar{\mathbf{D}}_k^T (\mathbf{y} - \boldsymbol{\mu}_k) \\
(7) \qquad \qquad \qquad &= \|\boldsymbol{\mu}_k - P_k(\mathbf{y})\|_{\tilde{\boldsymbol{\Sigma}}_k^{-1}}^2 + \frac{1}{b_k} \|\mathbf{y} - P_k(\mathbf{y})\|^2,
\end{aligned}$$

where  $\|\cdot\|_{\tilde{\boldsymbol{\Sigma}}_k^{-1}}^2$  is the norm defined by  $\|\mathbf{y}\|_{\tilde{\boldsymbol{\Sigma}}_k^{-1}}^2 = \mathbf{y}^T \tilde{\boldsymbol{\Sigma}}_k^{-1} \mathbf{y}$  with  $\tilde{\boldsymbol{\Sigma}}_k^{-1} = \tilde{\mathbf{D}}_k \mathbf{A}_k^{-1} \tilde{\mathbf{D}}_k^T$ . Equation (7) uses the definitions of  $P_k$  and  $P_k^\perp$  and  $\|\boldsymbol{\mu}_k - P_k^\perp(\mathbf{y})\|^2 = \|\mathbf{y} - P_k(\mathbf{y})\|^2$ . The gain comes from the fact that (7) does not depend on  $P_k^\perp$  and thus does not require the computation of the  $(M - d_k)$  latest columns of  $\mathbf{D}_k$ , the eigenvectors associated to the smallest eigenvalues. Similarly, determinants can be efficiently computed as  $\log(|\boldsymbol{\Sigma}_k|) = (\sum_{m=1}^{d_k} \log(a_{km})) + (M - d) \log(b_k)$ . This efficient parameterization, that now only depends on matrix  $\tilde{\mathbf{D}}_k$  and not on the complete matrix  $\mathbf{D}_k$ , is indicated by denoting the corresponding ED as  $\mathcal{HE}_{Md_k}(\boldsymbol{\mu}_k, \tilde{\mathbf{D}}_k^*, \mathbf{a}_k, b_k, \boldsymbol{\theta}_k)$ , where  $\tilde{\mathbf{D}}_k^*$  the matrix  $\tilde{\mathbf{D}}_k$  with the last zeros  $M - d_k$  columns omitted. We then refer to a MED coupled with this parameterization as a high-dimensional MED (HD-MED).

**DEFINITION 4.4 (HD-MED).** A random vector  $\mathbf{Y} \in \mathbb{R}^M$  follows a HD-MED distribution if for all  $k \in [1 : K]$ , the conditional pdf of the  $k^{\text{th}}$  mixture component  $f_k$  is an ED  $\mathcal{HE}_{Md_k}(\boldsymbol{\mu}_k, \tilde{\mathbf{D}}_k^*, \mathbf{a}_k, b_k, \boldsymbol{\theta}_k)$  with reparameterization given by (5) and (6). We denote

$$\mathbf{Y} \sim \mathcal{MHED}_{Md} \left( (\pi_k, \boldsymbol{\mu}_k, \tilde{\mathbf{D}}_k^*, \mathbf{a}_k, b_k, \boldsymbol{\theta}_k)_{k=1}^K \right).$$

With  $\mathbf{d} = (d_1, \dots, d_k)$ ,  $\mathbf{a}_k = (a_{k1}, \dots, a_{kd_k})$ .

## 5. Dimension reduction with HD-MED.

**5.1. Latent variable dimension reduction.** Standard PCA is defined without referring to a probabilistic model. Given a set of observations in  $\mathbb{R}^M$ , their  $M \times M$  empirical covariance matrix is decomposed into eigenvalues and eigenvectors and a number  $d \ll M$  of them are retained. For any observation  $\mathbf{y} \in \mathbb{R}^M$ , a lower dimensional representation can then be obtained by considering its projection to a lower dimensional subspace  $\hat{\mathbf{y}} = \hat{\boldsymbol{\Sigma}}_d^T \mathbf{y}$  where  $\hat{\boldsymbol{\Sigma}}_d$  is the matrix containing the  $d$  first eigenvectors of the empirical covariance matrix. If needed, its reconstruction in  $\mathbb{R}^M$ , optimal in the sense of the squared reconstruction error, can be obtained with  $\tilde{\mathbf{y}} = \hat{\boldsymbol{\Sigma}}_d \hat{\mathbf{y}}$ . Alternatively, if  $\mathbf{y}$  is assumed to be a realization of a random vector  $\mathbf{Y} \sim \mathcal{HE}_{Md}(\boldsymbol{\mu}, \tilde{\mathbf{D}}^*, \mathbf{a}, b, \boldsymbol{\theta})$ , a low dimensional representation of  $\mathbf{y}$  can be justified using the following latent variable model representation of  $\mathbf{Y}$ .

**PROPOSITION 5.1 (HD-ED latent variable model).** Let  $d \leq M - 1$ ,  $\mathbf{Y} \in \mathbb{R}^M$ ,  $\mathbf{X} \in \mathbb{R}^d$ ,  $\mathbf{E} \in \mathbb{R}^M$ ,  $W \in \mathbb{R}^+$  be random variables,  $\mathbf{V} \in \mathbb{R}^{M \times d}$  a matrix of linearly independent

columns,  $\boldsymbol{\mu} \in \mathbb{R}^M$  a vector and  $f_\theta$  the pdf of a positive univariate random variable defined by some parameter  $\theta$ . Assume that

$$\begin{aligned}\mathbf{Y} &= \mathbf{V}\mathbf{X} + \boldsymbol{\mu} + \mathbf{E} \\ (\mathbf{X}|W = w) &\sim \mathcal{N}(\mathbf{0}_d, w^{-1}\mathbf{I}_d) \\ (\mathbf{E}|W = w) &\sim \mathcal{N}(\mathbf{0}_M, b w^{-1}\mathbf{I}_M) \\ W &\sim f_\theta\end{aligned}$$

then,

$$\mathbf{Y} \sim \mathcal{HE}_{Md}(\boldsymbol{\mu}, \tilde{\mathbf{D}}^*, \mathbf{a}, b, \theta),$$

with  $b\mathbf{I}_M + \mathbf{V}\mathbf{V}^T = \mathbf{D}\mathbf{A}\mathbf{D}^T$  and  $\mathbf{A} = \text{diag}(a_1, \dots, a_d, b, \dots, b)$ .

Additionally, denoting by  $\mathbf{U} = b\mathbf{I}_d + \mathbf{V}^T\mathbf{V}$ , we have,

$$(8) \quad (\mathbf{X}|\mathbf{Y} = \mathbf{y}, W = w) \sim \mathcal{N}(\mathbf{U}^{-1}\mathbf{V}^T(\mathbf{y} - \boldsymbol{\mu}), w^{-1}b\mathbf{U}^{-1}).$$

It follows that  $\mathbb{E}[\mathbf{Y}|\mathbf{X} = \mathbf{x}] = \mathbf{V}\mathbf{x} + \boldsymbol{\mu}$  and  $\mathbb{E}[\mathbf{X}|\mathbf{Y} = \mathbf{y}] = \mathbf{U}^{-1}\mathbf{V}^T(\mathbf{y} - \boldsymbol{\mu})$ .

PROOF. It comes from the first three assumptions that

$$(\mathbf{Y}|W = w) \sim \mathcal{N}(\boldsymbol{\mu}, w^{-1}(\mathbf{V}\mathbf{V}^T + b\mathbf{I}_M))$$

and from the GSM definition 4.2 that

$$\mathbf{Y} \sim \mathcal{E}_M(\boldsymbol{\mu}, \mathbf{V}\mathbf{V}^T + b\mathbf{I}_M, f_\theta).$$

Since  $\mathbf{V}\mathbf{V}^T$  is of rank  $d$ ,  $\mathbf{V}\mathbf{V}^T + b\mathbf{I}_M$  admits an eigenvalue decomposition  $\mathbf{D}\mathbf{A}\mathbf{D}^T$  as stated.

Distribution (8) follows from standard Gaussian vectors properties. Using the tower property, it comes then  $\mathbb{E}[\mathbf{Y}|\mathbf{X} = \mathbf{x}] = \mathbb{E}[\mathbb{E}[\mathbf{Y}|\mathbf{X} = \mathbf{x}, W]] = \mathbf{V}\mathbf{x} + \boldsymbol{\mu}$  and  $\mathbb{E}[\mathbf{X}|\mathbf{Y} = \mathbf{y}] = \mathbb{E}[\mathbb{E}[\mathbf{X}|\mathbf{Y} = \mathbf{y}, W]] = \mathbf{U}^{-1}\mathbf{V}^T(\mathbf{y} - \boldsymbol{\mu})$ .  $\square$

The previous proposition states that a realization  $\mathbf{y} \in \mathbb{R}^M$  from an HD-ED, can also be seen as originating from a generative model with a lower-dimensional latent variable  $\mathbf{X} \in \mathbb{R}^d$ . Hence, a natural alternative to the standard PCA projection, is the conditional mean  $\hat{\mathbf{y}} = \mathbb{E}[\mathbf{X}|\mathbf{Y} = \mathbf{y}]$ , that is

$$(9) \quad \hat{\mathbf{y}} = Q(\mathbf{y}) = \mathbf{U}^{-1}\mathbf{V}^T(\mathbf{y} - \boldsymbol{\mu}),$$

and as a reconstruction or an approximation of the original information, the conditional mean  $\tilde{\mathbf{y}} = \mathbb{E}[\mathbf{Y}|\mathbf{X} = \hat{\mathbf{y}}]$ , that is  $\tilde{\mathbf{y}} = \mathbf{V}\hat{\mathbf{y}} + \boldsymbol{\mu}$ . Using the previous formulas, lower dimensional representations can thus be obtained using  $\mathbf{U}$  and  $\mathbf{V}$  but when estimating the parameters of the HD-ED, we get estimates for  $\mathbf{A}$  and  $\mathbf{D}$  instead. However, using that  $\mathbf{V}$  is of rank  $d$  and  $\mathbf{D}\mathbf{A}\mathbf{D}^T = b\mathbf{I}_M + \mathbf{V}\mathbf{V}^T$ , we can set

$$(10) \quad \mathbf{V} = \tilde{\mathbf{D}}^* \sqrt{\text{diag}(a_1, \dots, a_d) - b\mathbf{I}_d},$$

and deduce  $\mathbf{U}$  straightforwardly.

5.2. *Cluster globally, reduce locally.* The proposed dimension reduction and reconstruction, using the HD-ED model, generalizes to GSM, PPCA (Tipping and Bishop, 1999c) and robust PPCA (Archambeau, Delannay and Verleysen, 2006), the former using Gaussian and the latter Student distributions. Both have been extended to account for potential heterogeneity in data, considering mixtures, with the MPPCA model (Tipping and Bishop, 1999d) and its Student-based robust version (Archambeau, Delannay and Verleysen, 2008). In these models, all clusters are assumed to live in subspaces of the same dimension  $d$ . In the Gaussian case, an extension, allowing varying reduced dimensions  $d_k$  across clusters, have been proposed by Bouveyron, Girard and Schmid (2007) with their high dimensional Gaussian mixture model (HD-GMM). The proposed HD-GMM parameterization allows to handle high dimensional data in a computationally efficient way. However, it does not provide an actual lower dimensional representation of the data. While such a reduced-dimensional representation may often not be needed, it may be crucial to deal with hardware or software limitations. Originally, HD-GMM have not been designed for dimension reduction or compression but rather for clustering and density estimation in high-dimensional heterogeneous settings. Our HD-MED model uses the same efficient decomposition of the covariance matrix but generalizes it to the scale matrix of a GSM distribution. In addition, we describe how it can be further exploited as a dimension reduction technique. As finite mixture models, HD-MED can be used for clustering data into  $K$  clusters. For any possible observation  $\mathbf{y}$ , a HD-MED model provides a probability  $r_k(\mathbf{y})$  that  $\mathbf{y}$  is assigned to cluster  $k$  for each  $k \in [1:K]$ . Using (9), a reduced-dimension representation  $\hat{\mathbf{y}}_k$  of  $\mathbf{y}$ , for each of the  $K$  different subspaces, is denoted by  $\hat{\mathbf{y}}_k = Q_k(\mathbf{y})$  and given by,

$$(11) \quad \hat{\mathbf{y}}_k = Q_k(\mathbf{y}) = \mathbf{U}_k^{-1} \mathbf{V}_k^T (\mathbf{y} - \boldsymbol{\mu}_k),$$

while its reconstruction  $\check{\mathbf{y}}_k$  in the original space is given by

$$\check{\mathbf{y}}_k = \mathbf{V}_k \hat{\mathbf{y}}_k + \boldsymbol{\mu}_k.$$

In practice, it is reasonable to use as a reduced-dimension representation of  $\mathbf{y}$  only the one corresponding to the most probable cluster  $k$ , *i.e.* with the highest  $r_k(\mathbf{y})$ . In this setting, HD-MED acts as a divide-and-conquer paradigm by initially clustering the data into  $K$  clusters and then performing cluster-specific data reduction. The divide step allows a much more effective reduction than if a single subspace was considered, while in the conquer step, little information is lost, as for any new observation  $\mathbf{y}$ , cluster assignment probabilities  $r_k(\mathbf{y})$  can be straightforwardly computed to decide on the best reduced representation to be used. However, for subsequent processing, it is important to keep track of clustering information for each observation. The reduced representations cannot be pooled back altogether, as they are likely to become impossible to distinguish across clusters. Also, as summarized in Algorithm 1 and illustrated in Section 7, each reduced cluster may have to be processed separately but this additional cost is negligible compared to the hardware and software gain of a more efficient representation.

In practice, the Expectation-Maximization (EM) algorithm (Dempster, Laird and Rubin (1977)), that iteratively computes and maximizes the conditional expectation of the complete-data log-likelihood, is commonly used to infer the parameters of finite mixture models. The number of clusters  $K$ , and their respective inner dimension  $d_k$  are hyper-parameters that need to be tuned prior to the EM steps. As a simple solution, we use the Bayesian Information Criterion (BIC) to tune  $K$  and the  $d_k$ 's at the same time. In contrast to MPPCA solutions which assume the same subspace dimension  $d$  for all clusters, the possibility to handle different  $d_k$ 's and to allow non Gaussian cluster shapes is important for the target applications involving datasets that are very large. Using different dimensions across clusters is likely to yield a more efficient reduced representation of the data as illustrated in Section 7.

However, the standard, or batch EM algorithm needs all the dataset to be loaded in a fast access memory (*e.g.* the RAM) which is often limited as this kind of memory is expensive. In the case of large-scale dataset, the RAM is often overloaded and supported by a slow-memory which makes the iterations of the EM very slow. Batch sizes are then limited by resource constraints, so that very large data sets need either to be downsampled or to be handled in an incremental manner. Incremental versions of EM exist and can be adapted to our setting. In section 6, we provide a way to deal with this large-scale case by using an online version of the EM algorithm.

## 6. Online Learning of High Dimensional Mixtures of Elliptical Distributions.

6.1. *Online EM, main assumptions.* When the data volume is too large the EM algorithm becomes slow because of multiple data transfer between the RAM and the store of the computer. A way to handle large volumes is to use online learning. Online learning refers to procedures able to deal with data acquired sequentially. Online variants of EM, among others, are described in Cappé and Moulines (2009); Maire, Moulines and Lefebvre (2017); Karimi et al. (2019a,b); Fort, Moulines and Wai (2020); Kuhn, Matias and Rebafka (2020); Nguyen, Forbes and McLachlan (2020). As an archetype of such algorithms, we consider the online EM of Cappé and Moulines (2009) which belongs to the family of stochastic approximation algorithms (Borkar (2009)). This algorithm has been well theoretically studied and extended. However, it is designed only for distributions that admit a data augmentation scheme, or a latent variable formulation, yielding a complete likelihood of the exponential family form, see (12) below. This case is already very broad, including *e.g.* Gaussian, gamma, Student distributions and mixtures of those. We recall the main assumptions required and the online EM iteration, based on a latent variable formulation.

Assume  $(\mathbf{Y}_i)_{i=1}^N$  is a sequence of  $N$  *i.i.d.* replicates of a random variable  $\mathbf{Y} \in \mathcal{Y} \subset \mathbb{R}^M$ , observed one at a time. Extension to successive mini-batches of observations is straightforward (Nguyen, Forbes and McLachlan (2020)). In addition,  $\mathbf{Y}$  is assumed to be the visible part of  $(\mathbf{Y}, \mathbf{Z})$ , where  $\mathbf{Z} \in \mathbb{R}^l$  is a latent variable, *e.g.* the unknown component label in a mixture model or a mixing weight in a GSM formulation, and  $l \in \mathbb{Z}_+$ . For  $i \in [1 : N]$  then, each  $\mathbf{Y}_i$  is the visible part of  $(\mathbf{Y}_i, \mathbf{Z}_i)$ . Suppose  $\mathbf{Y}$  arises from some data generating process (DGP) characterised by a probability density function  $f(\mathbf{y}; \Theta_0)$ , with unknown parameters  $\Theta_0 \in \mathbb{T} \subseteq \mathbb{R}^p$ , for  $p \in \mathbb{Z}_+$ .

Using the sequence  $(\mathbf{Y}_i)_{i=1}^N$ , the method of Cappé and Moulines (2009) allows to sequentially estimate  $\Theta_0$  provided the following assumptions are met:

(A1) The complete-data likelihood for  $(\mathbf{Y}, \mathbf{Z})$  is of the exponential family form:

$$(12) \quad f_c(\mathbf{y}, \mathbf{z}; \Theta) = h(\mathbf{y}, \mathbf{z}) \exp \left\{ [\mathbf{s}(\mathbf{y}, \mathbf{z})]^\top \phi(\Theta) - \psi(\Theta) \right\},$$

with  $h : \mathbb{R}^{M+l} \rightarrow [0, \infty)$ ,  $\psi : \mathbb{R}^p \rightarrow \mathbb{R}$ ,  $\mathbf{s} : \mathbb{R}^{M+l} \rightarrow \mathbb{R}^q$ ,  $\phi : \mathbb{R}^p \rightarrow \mathbb{R}^q$ , for  $q \in \mathbb{Z}_+$ .

(A2) The function

$$(13) \quad \bar{\mathbf{s}}(\mathbf{y}; \Theta) = \mathbb{E}[\mathbf{s}(\mathbf{Y}, \mathbf{Z}) | \mathbf{Y} = \mathbf{y}; \Theta]$$

is well-defined for all  $\mathbf{y}$  and  $\Theta \in \mathbb{T}$ , where  $\mathbb{E}[\cdot | \mathbf{Y} = \mathbf{y}; \Theta]$  is the conditional expectation when  $\mathbf{X}$  arises from the DGP characterised by  $\Theta$ .

(A3) There is a convex  $\mathbb{S} \subseteq \mathbb{R}^q$ , satisfying: (i) for all  $\gamma \in (0, 1)$ ,  $\mathbf{s} \in \mathbb{S}$ ,  $\mathbf{y} \in \mathcal{Y}$ , and  $\Theta \in \mathbb{T}$ ,  $(1 - \gamma)\mathbf{s} + \gamma\bar{\mathbf{s}}(\mathbf{y}; \Theta) \in \mathbb{S}$ ; and (ii) for any  $\mathbf{s} \in \mathbb{S}$ , the function  $Q(\mathbf{s}; \Theta) = \mathbf{s}^\top \phi(\Theta) - \psi(\Theta)$  has a unique global maximizer on  $\mathbb{T}$  denoted by

$$(14) \quad \bar{\Theta}(\mathbf{s}) = \arg \max_{\Theta \in \mathbb{T}} Q(\mathbf{s}; \Theta).$$

Let  $(\gamma_i)_{i=1}^N$  be a sequence of learning rates in  $(0, 1)$  and let  $\Theta^{(0)} \in \mathbb{T}$  be an initial estimate of  $\Theta_0$ . For each  $i \in [1 : N]$ , the online EM of [Cappé and Moulines \(2009\)](#) proceeds by computing

$$(15) \quad \mathbf{s}^{(i)} = \gamma_i \bar{\mathbf{s}}(\mathbf{y}_i; \Theta^{(i-1)}) + (1 - \gamma_i) \mathbf{s}^{(i-1)},$$

and

$$(16) \quad \Theta^{(i)} = \bar{\Theta}(\mathbf{s}^{(i)}),$$

where  $\mathbf{s}^{(0)} = \bar{\mathbf{s}}(\mathbf{y}_1; \Theta^{(0)})$ . It is shown in Theorem 1 of [Cappé and Moulines \(2009\)](#) that when  $N$  tends to infinity, the sequence  $(\Theta^{(i)})_{i=1:N}$  of estimators of  $\Theta_0$  satisfies a convergence result to stationary points of the likelihood (cf. [Cappé and Moulines \(2009\)](#) for a more precise statement).

**6.2. Online EM for HD-ED.** In this subsection, we derive the online EM algorithm for HD-ED. The extension to mixtures of those (HD-MED) is straightforward and is detailed in [Nguyen and Forbes \(2022\)](#) or in the supplementary Section 1. The weight  $W$  distribution  $f$  in the GSM formulation (3) is assumed to belong to the exponential family. This case may seem restrictive but it encompasses a number of ED as the Gaussian, Student, Normal Inverse Gamma with no skewness *etc.* distributions.

**PROPOSITION 6.1 (HD-ED exponential form).** *Let  $\mathbf{Y}$  be a HD-ED distributed variable,  $\mathbf{Y} \sim \mathcal{HE}_{Md}(\boldsymbol{\mu}, \tilde{\mathbf{D}}^*, \mathbf{a}, b, \boldsymbol{\theta})$ , and  $W$  a weight variable with pdf  $f_\theta$ . The set of parameters is denoted by  $\Theta = (\boldsymbol{\mu}, \tilde{\mathbf{D}}^*, \mathbf{a}, b, \boldsymbol{\theta})$ , with  $\tilde{\mathbf{D}}^*$  defined by its column vectors  $\tilde{\mathbf{D}}^* = [\mathbf{d}_1, \dots, \mathbf{d}_d]$ . If  $W$  belongs to the exponential family, i.e.  $f_\theta(w) = h_w(w) \exp[\mathbf{s}_w(w)^T \boldsymbol{\phi}_w(\boldsymbol{\theta}) - \psi_w(\boldsymbol{\theta})]$ , the complete data likelihood*

$$(17) \quad f_c(\mathbf{y}, w; \Theta) = f_\theta(w) \mathcal{N}_M(\mathbf{y}; \boldsymbol{\mu}, w^{-1} \mathbf{D} \mathbf{A} \mathbf{D}^T),$$

can be expressed in an exponential family form (12) with

$$(18) \quad \mathbf{s}(\mathbf{y}, w) = \begin{bmatrix} w\mathbf{y} \\ w \text{vec}(\mathbf{y}\mathbf{y}^T) \\ w\mathbf{y}^T \mathbf{y} \\ w \\ \mathbf{s}_w(w) \end{bmatrix}, \quad \boldsymbol{\phi}(\Theta) = \begin{bmatrix} \sum_{m=1}^d \left( \frac{1}{a_m} - \frac{1}{b} \right) \mathbf{d}_m \mathbf{d}_m^T \boldsymbol{\mu} + \frac{1}{b} \boldsymbol{\mu} \\ \frac{1}{2} \sum_{m=1}^d \left( \frac{1}{b} - \frac{1}{a_m} \right) \text{vec}(\mathbf{d}_m \mathbf{d}_m^T) \\ -\frac{1}{2b} \\ \frac{1}{2} \sum_{m=1}^d \left( \frac{1}{b} - \frac{1}{a_m} \right) \boldsymbol{\mu}^T \mathbf{d}_m \mathbf{d}_m^T \boldsymbol{\mu} - \frac{1}{2b} \boldsymbol{\mu}^T \boldsymbol{\mu} \\ \boldsymbol{\phi}_w(\boldsymbol{\theta}) \end{bmatrix},$$

$$\psi(\Theta) = \frac{1}{2} \sum_{m=1}^d \log a_m + \frac{M-d}{2} \log b + \psi_w(\boldsymbol{\theta}),$$

where  $\text{vec}$  denotes the vectorization operator ([Schott \(2016\)](#)).

**PROOF.** The proof is detailed in supplementary Section 2. □

The online EM algorithm (OEM) consists, as the batch EM, of two steps, the first one is the computation of the sufficient statistics (13), and the second one is the maximization of the likelihood (14). For the first step, we need to compute  $\bar{\mathbf{s}}(\mathbf{y}; \Theta) = \mathbb{E}[\mathbf{s}(\mathbf{Y}, \mathbf{Z}) | \mathbf{Y} = \mathbf{y}; \Theta]$ . This quantity requires to compute the following expectations  $\mathbb{E}[W | \mathbf{Y} = \mathbf{y}; \Theta]$ , and  $\mathbb{E}[s_w(W) | \mathbf{Y} = \mathbf{y}; \Theta]$ .

PROPOSITION 6.2 (Expectations of sufficient statistics). *Let  $\mathbf{Y} \sim \mathcal{HE}_{Md}(\boldsymbol{\mu}, \tilde{\mathbf{D}}^*, \mathbf{a}, b, \boldsymbol{\theta})$ ,  $\boldsymbol{\Theta} = (\boldsymbol{\mu}, \tilde{\mathbf{D}}^*, \mathbf{a}, b, \boldsymbol{\theta})$ , and  $W$  the mixing variable  $W \sim f_{\boldsymbol{\theta}}$ . Then  $\mathbf{Y}$  has density (2) with generator  $g$  and Mahalanobis distance  $u = (\mathbf{y} - \boldsymbol{\mu})^T \mathbf{D} \mathbf{A}^{-1} \mathbf{D}^T (\mathbf{y} - \boldsymbol{\mu})$  defined as in (7). It follows that*

$$(19) \quad \mathbb{E}[W | \mathbf{Y} = \mathbf{y}; \boldsymbol{\Theta}] = -\frac{2}{(2\pi)^{M/2}} \frac{g'(u)}{g(u)}.$$

where  $g'$  denotes the derivative of  $g$ .

PROOF. According to equations (2) and (3), we have

$$\begin{aligned} \mathbb{E}[W | \mathbf{Y} = \mathbf{y}; \boldsymbol{\Theta}] &= \frac{1}{(2\pi)^{M/2} g(u)} \int_{\mathbb{R}^+} w f_{\boldsymbol{\theta}}(w) \exp\left(-\frac{w}{2}u\right) dw \\ &= -\frac{2}{(2\pi)^{M/2}} \frac{g'(u)}{g(u)}, \end{aligned}$$

which proves equation (19).  $\square$

In contrast, there is no general formula for the expectation of  $s_w(W)$ , which depends on the mixing distribution  $f_{\boldsymbol{\theta}}$ . Once the expectation of the sufficient statistics in (13) is computed, we can update it following (15).

The next OEM step is the maximization step described in (14), which gives an estimation of the parameters at each iteration. The solution for  $\boldsymbol{\theta}$  varies with  $f_{\boldsymbol{\theta}}$ , but solutions for  $\boldsymbol{\mu}$ ,  $\mathbf{A}$ , and  $\tilde{\mathbf{D}}^*$  can be derived as follows. Let  $\bar{\boldsymbol{\Theta}}(\mathbf{s})$  be defined as the unique maximizer of function  $Q(\mathbf{s}, \boldsymbol{\Theta}) = \mathbf{s}^T \boldsymbol{\phi}(\boldsymbol{\Theta}) - \psi(\boldsymbol{\Theta})$  with  $\mathbf{s}$  a vector that matches the definition and dimension of  $\boldsymbol{\phi}(\boldsymbol{\Theta})$  in (18), and can be conveniently written as

$$(20) \quad \mathbf{s} = \begin{bmatrix} \mathbf{s}_1 \\ \text{vec}(\mathbf{S}_2) \\ s_3 \\ s_4 \\ s_5 \end{bmatrix},$$

with  $\mathbf{s}_1$  a  $M$ -dimensional vector,  $\mathbf{S}_2$  a  $M \times M$  matrix, and  $s_3, s_4, s_5$  three scalar values. Parameters are updated by maximizing  $Q$  with respect to  $\boldsymbol{\Theta}$ .  $\bar{\boldsymbol{\Theta}}(\mathbf{s})$  is defined as the root of the first-order condition

$$(21) \quad \mathbf{J}_{\boldsymbol{\phi}}(\boldsymbol{\Theta})\mathbf{s} - \frac{\partial \psi}{\partial \boldsymbol{\Theta}}(\boldsymbol{\Theta}) = \mathbf{0},$$

where  $\mathbf{J}_{\boldsymbol{\phi}}(\boldsymbol{\Theta}) = \frac{\partial \boldsymbol{\phi}}{\partial \boldsymbol{\Theta}}$  is the Jacobian of  $\boldsymbol{\phi}$ , with respect to  $\boldsymbol{\Theta}$ . Computing gradients leads to  $\bar{\boldsymbol{\Theta}}(\mathbf{s}) = (\bar{\boldsymbol{\mu}}(\mathbf{s}), \tilde{\mathbf{D}}^*(\mathbf{s}), \bar{\mathbf{A}}(\mathbf{s}), \bar{\boldsymbol{\theta}}(\mathbf{s}))$ , where  $\bar{\boldsymbol{\mu}}(\mathbf{s})$ , and  $\bar{\mathbf{A}}(\mathbf{s})$  are closed form, and  $\tilde{\mathbf{D}}^*(\mathbf{s})$  can be found using Riemannian optimization.

PROPOSITION 6.3 (Maxima). *We proceed in a ECM-like procedure (Meng and Rubin (1993)) by optimizing each parameters separately and incorporating them during the optimization of each other parameters.  $(\boldsymbol{\mu}, \mathbf{A})$  can be optimized easily using (21) and computing the gradients gives*

$$(22) \quad \bar{\boldsymbol{\mu}} = \frac{\mathbf{s}_1}{s_4},$$

$$(23) \quad \bar{a}_m = \mathbf{d}_m^T (\mathbf{S}_2 + s_4 \bar{\boldsymbol{\mu}} \bar{\boldsymbol{\mu}}^T - 2 \bar{\boldsymbol{\mu}} \mathbf{s}_1^T) \mathbf{d}_m \quad \text{for } m \in [1 : d],$$

$$(24) \quad \bar{b} = \frac{1}{M - d} \left( s_4 \bar{\boldsymbol{\mu}}^T \bar{\boldsymbol{\mu}} + s_3 - 2 \bar{\boldsymbol{\mu}}^T \mathbf{s}_1 - \sum_{m=1}^d \bar{a}_m \right).$$

The maximisation in  $\tilde{\mathbf{D}}^*$  has to take into account that  $\tilde{\mathbf{D}}^* \in St(M, d)$ , the Stiefel manifold of the  $M \times d$  matrices. Plugin-in the expressions of  $\boldsymbol{\mu}$  and  $\mathbf{A}$  above and omitting parts that depend on  $\boldsymbol{\theta}$  we have,

$$(25) \quad \tilde{\mathbf{D}}^* = \arg \max_{\tilde{\mathbf{D}}^* \in St(M, d)} \sum_{m=1}^d \left( \frac{1}{\bar{a}_m} - \frac{1}{\bar{b}} \right) \mathbf{d}_m^T (2 \bar{\boldsymbol{\mu}} \mathbf{s}_1^T - \mathbf{S}_2 - s_4 \bar{\boldsymbol{\mu}} \bar{\boldsymbol{\mu}}^T) \mathbf{d}_m.$$

For  $\bar{\boldsymbol{\theta}}$ , a general closed-form expression is not available, but if there are no particular constraints it results from solving the following equation

$$(26) \quad s_5 \frac{\partial \phi_w}{\partial \boldsymbol{\theta}}(\bar{\boldsymbol{\theta}}) - \frac{\partial \psi_w}{\partial \boldsymbol{\theta}}(\bar{\boldsymbol{\theta}}) = \mathbf{0}.$$

PROOF. We compute the gradients and use (21) for  $\boldsymbol{\mu}$  and  $\mathbf{A}$ , however for the vector  $\tilde{\mathbf{D}}^*$  we only plug-in the optimized values  $\bar{\boldsymbol{\mu}}$ , and  $\bar{\mathbf{A}}$  in (17) and solve

$$(27) \quad \arg \max_{\tilde{\mathbf{D}}^* \in St(M, d)} \mathbf{s}^T \phi(\bar{\boldsymbol{\mu}}, \bar{\mathbf{A}}, \tilde{\mathbf{D}}^*) - \psi(\bar{\boldsymbol{\mu}}, \bar{\mathbf{A}}, \tilde{\mathbf{D}}^*).$$

□

---

### Algorithm 1 Divide & Conquer high dimensional matching for MRF reconstruction

---

**Input** Dictionary of (signal, parameters) pairs  $\mathcal{D}_f = \{\mathbf{y}_i, \mathbf{t}_i\}_{i=1:N}$ ,  $N \gg 1$ ,  $\mathbf{t}_i \in \mathbb{R}^L$ ,  $\mathbf{y}_i \in \mathbb{R}^M$ ,  $M \gg 1$ .  
 In vivo acquired signals  $\{\tilde{\mathbf{y}}_j\}_{j=1:\tilde{N}}$ ,  $\tilde{\mathbf{y}}_j \in \mathbb{R}^M$ .

**1: Reduced dimension representation of the dictionary:**  $\{\hat{\mathbf{y}}_i, \mathbf{t}_i, r(\mathbf{y}_i)\}_{i=1:N}$

1.1 **Online HD-MED inference from  $\{\mathbf{y}_i\}_{i=1:N}$ :**  $K$  clusters,  $d_k < M$  for  $k = 1 : K \implies$  cluster assignment probabilities and cluster-wise projections  $(\mathbf{r}, \mathbf{Q}) = \{r_k(\cdot), Q_k(\cdot)\}_{k=1:K}$

1.2 **Cluster-wise fingerprint reductions:**  $\{\mathbf{y}_i\}_{i=1:N}, \mathbf{r}, \mathbf{Q} \implies \{\hat{\mathbf{y}}_i = Q_k(\mathbf{y}_i), i \in I_k\}$  with  $I_k = \{i, s.t. k = \arg \max_{\ell} r_{\ell}(\mathbf{y}_i)\}$ , for  $k = 1 : K$

**2: Cluster-wise matching of acquired signals:**

2.1 **Cluster-wise in vivo signal reductions:** Use learned  $(\mathbf{r}, \mathbf{Q})$  from step 1.1 to obtain  $\{Q_k(\tilde{\mathbf{y}}_j), j \in \tilde{I}_k\}$  with  $\tilde{I}_k = \{j, s.t. k = \arg \max_{\ell} r_{\ell}(\tilde{\mathbf{y}}_j)\}$ , for  $k = 1 : K$

2.2 **Matching:** For  $k = 1 : K$ , for  $j \in \tilde{I}_k$ , determine  $i(\tilde{\mathbf{y}}_j) = \arg \min_{i \in I_k} d(Q_k(\tilde{\mathbf{y}}_j), \hat{\mathbf{y}}_i)$  and set  $\tilde{\mathbf{t}}_j = \mathbf{t}_{i(\tilde{\mathbf{y}}_j)}$

**Return** Matched tissue properties  $\{\tilde{\mathbf{t}}_j\}_{j=1:\tilde{N}}$ ,  $\tilde{\mathbf{t}}_j \in \mathbb{R}^L$

---

**7. Application to magnetic resonance fingerprinting (MRF) reconstruction.** MRF is able to provide multiple quantitative tissue parameters images from shorter acquisition times, thanks to the simultaneous application of transient states excitation and highly under-sampled  $k$ -space read-outs. These two aspects have a combined impact on acquisition times

and image reconstruction accuracy. More undersampling allows more parameter estimations in reasonable acquisition times but is also responsible for larger undersampling errors, noise and artifacts, reducing map reconstruction accuracy.

In some earlier work (Oudoumanessah et al., 2024), the Gaussian version of our procedure, referred to as HD-GMM, was evaluated, for the reconstruction, in an ideal setting, of fully sampled acquisitions, targetting  $L = 3$  parameters. In this scenario, it was reported that HD-GMM, coupled with the online EM algorithm, achieved results comparable to full dictionary matching while significantly reducing reconstruction times. However, as we report in this section, HD-GMM performance degrades when dealing with largely undersampled acquisitions and becomes inadequate for estimating six parameters. In contrast, performance can be maintained by considering an elliptical version of our procedure (HD-MED) less sensitive to outliers.

In this section, we thus show how we can accurately reconstruct six parameter maps from *in vivo* undersampled acquisitions (Section 7.1) by leveraging an extensive high-resolution dictionary (Section 7.2) and the HD-MED model. Algorithm 1 provides a schematic summary of our procedure. Figure 1 provides an illustration of our matching strategy or step 2 in Algorithm 1, once the HD-MED model has been estimated. Each acquired signal is first assigned to one of the learned clusters, and reduced accordingly to be then matched to the best dictionary signal in the corresponding cluster. It leads to improvements in both memory management complexity and reconstruction speed for parameter maps when compared to traditional dictionary matching, see Section 7.5.

We compare two instances of HD-MED, namely HD-GMM and HD-STM, where the mixture components are respectively set to Gaussian and Student distributions. The computations for the corresponding OEM are detailed in Section 3 of the Appendix. All the experiments are performed with a Python code using the JAX library (Bradbury et al., 2018) with Nvidia V100-32gb GPU except for the dictionary generation part that is done with a mix of Matlab and Python. The Python code is available at <https://github.com/geoffroyO/HD-MED>.

**7.1. Undersampled MRF acquisitions.** *In vivo* acquisitions were conducted on 6 healthy volunteers ( $28 \pm 5.5$  years old, 3 males and 3 females) using a 32-channel head receiver array on a Philips 3T Achieva dStream MRI at the IRMaGe facility (MAP-IRMaGe protocol, NCT05036629). This study was approved by the local medical ethics committee and informed consent was obtained from each volunteer prior to image acquisition. The imaging pulse sequence was based on an IR-bSSFP acquisition. 260 repetitions were acquired following the parameters proposed in Coudert et al. (2024). The acquisitions were performed using quadratic variable density spiral sampling (12 interleaves out of 13), matrix size=192x192x(4-5), voxel size=1.04x1.04x3.00 mm<sup>3</sup> for a total scan duration of 2 minutes per slice. While the acquisition time may appear high for an MRF context compared to Gu et al. (2018), where full T1 and T2 maps are generated in 2 minutes, the longer sequences in Coudert et al. (2024) account for vascular parameters without requiring contrast agents.

Spiral sampling in MRI involves acquiring data in a spiral trajectory through k-space, covering the center first and gradually moving outward, which allows for faster data acquisition and more efficient use of the scanner’s gradients. However, this method leads to undersampling noise, manifesting as artifacts, because it captures less information about the image, reducing the ability to accurately reconstruct fine details (Körzdörfer et al., 2019). It results in a number  $\tilde{N}_s \approx 140,000 - 180,000$  of *in vivo* MR signals per subject  $s$ , each of dimension  $M = 260$ , and a total number  $\tilde{N}$  of approximately  $\tilde{N} = 960,000$  signals to be matched for the reconstruction of  $L = 6$  parameter maps for 6 subjects.

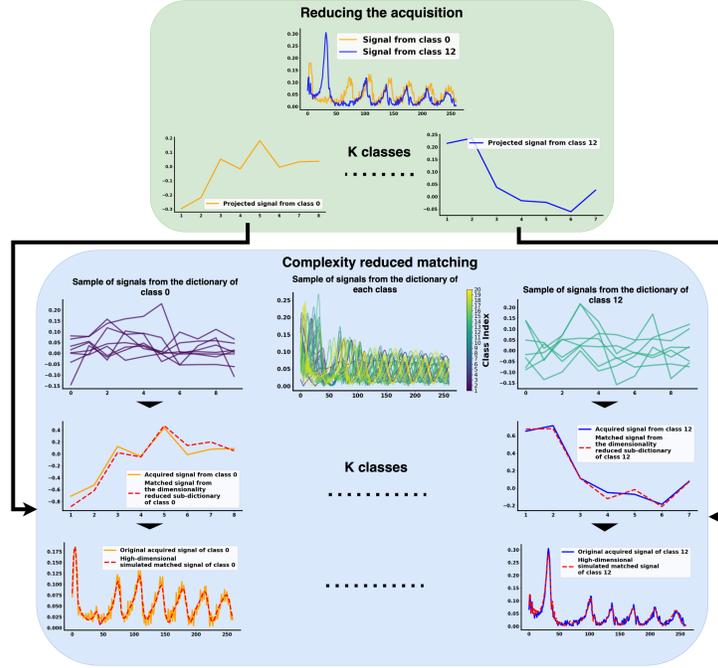


Fig 1: HD-MED-based low-dimensional signal matching step (step 2 of Algorithm 1) Green block: each acquired signal is assigned to one of the  $K$  HD-MED learned clusters and reduced accordingly using the assigned cluster projection. Blue block: reduced acquired signals are then matched with their closest reduced simulated counterpart in the cluster.

**7.2. MRF dictionary.** As introduced, the MRF sequence used here is a bSSFP-derived sequence. This type of sequence has been preferred because of its sensitivity to local frequency distributions related to microvascular structures in the imaging voxel. To account for this, the MRF dictionary is generated using the approach described in Coudert et al. (2024). First a base dictionary  $\mathcal{D}_{f_0}$  is simulated using Bloch-equations for different combinations of  $(T_1, T_2, B_1, \delta f)$  browsed through a 4-dimensional regular grid. Simulations were made at a magnetic field strength of 3.0T, on a regular parameter grid made of 20  $T_1$  values (from 200 to 3500ms), 20  $T_2$  values (from 10 to 600ms), 10  $B_1$  values (from 0.7 to 1.2) and 100 frequency offset  $\delta f$  values (from  $-50$  to 49 Hz with an increment of 1 Hz), keeping only signals for which  $T_1 > T_2$ , resulting in an initial 390,000 entries dictionary. These Bloch simulations are obtained using an in-house mix of Python and Matlab code based on initial implementation by B. Hargreaves (Hargreaves). The relaxometric parameters vary within the dictionary to allow their estimation.  $B_1$  is varied to ensure the realism of the dictionary, given the sequence's sensitivity to this parameter, which could otherwise bias the estimation of the other parameters. Finally, varying  $\delta f$  values are necessary to compute microvascular contributions, as detailed below.

To build signals that capture additional microvascular network's blood volume (CBV) and mean vessel radius (R) information, we follow a construction proposed by Coudert et al. (2024). Microvascular network segmentations are used to pre-compute 2500 different intra-voxel frequency distributions centred at  $\delta f$  values. A CBV and R value characterizes each distribution. Using these pre-computed distributions, new signals are obtained by summing signals from the base dictionary weighted according to each frequency distribution. It results an expanded 6-dimensional dictionary  $\mathcal{D}_f$  of almost  $N = 400,000,000$  signals of dimension  $M = 260$ , encoding for  $(T_1, T_2, B_1, \delta f, CBV, R)$ . Figure 2 gives an illustration on how

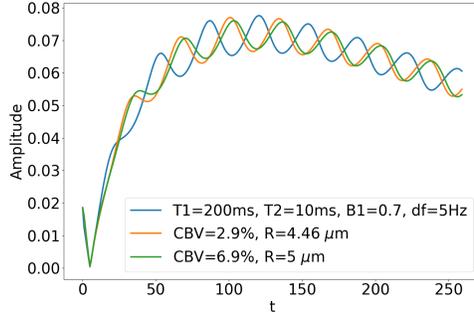


Fig 2: Illustration of the base dictionary expansion. The blue signal corresponding to a non-blood voxel with associated  $T_1$ ,  $T_2$ ,  $B_1$  and  $\delta f$  values. By convoluting this signal along the  $\delta f$  dimension with two different frequency distributions corresponding to different values of  $CBV$  and  $R$ , the resulting orange and green signals are then obtained for voxels with the same  $T_1$ ,  $T_2$ ,  $B_1$  and  $\delta f$  values as the blue signal but with different  $CBV$  and  $R$  values.

this operation changes a given signal in the base dictionary into 2 new signals depending on values of  $CBV$  and  $R$ .

*7.3. Model selection and initialization.* The only two hyperparameters that need to be set are the number of mixture components  $K$  and the desired reduced dimension vector  $\mathbf{d}$ . For mixture models,  $K$  can be selected using the Bayesian Information Criterion (BIC) (Schwarz, 1978), which requires running multiple models with varying  $K$  values and finding the elbow or the minimum of the BIC curve. The desired dimension vector  $\mathbf{d}$  is automatically set during an initialization step, which makes  $K$  the only parameter that the user needs to set prior to running the OEM algorithm.

The initialization step includes determining the reduced dimension vector  $\mathbf{d}$ . We randomly choose a subset of the dictionary that fits into memory and run a batch full covariance EM algorithm to determine a first estimation of the  $\Sigma_k$ 's. These estimations are decomposed into eigenvalues and eigenvectors to determine  $\mathbf{D}_k$ 's and  $\mathbf{A}_k$ 's. For each  $k$ ,  $d_k$  is then determined by applying a scree plot to the eigenvalues using the kneedle algorithm (Satopaa et al., 2011),  $\mathbf{a}_k$  is then initialized to the first  $d_k$  eigenvalues, and the last eigenvalues are averaged to initialize  $b_k$ . This kind of initialization of  $\mathbf{D}_k$  and  $\mathbf{A}_k$  is called spectral initialization and is similar to the one proposed by Hong et al. (2021), proving to be relatively stable compared to other types of initializations.

In Figure 3 (left), we show the different BIC curves for HD-GMM and HD-STM models trained on the dictionary of signals with  $K$  varying from 5 to 80. The elbows of the curves are found at  $K = 30$  for HD-GMM, and at  $K = 25$  for HD-STM. Figures 3 (middle and right) show the different vectors  $\mathbf{d}$  obtained for varying  $K$ . The larger the points, the higher the number of components having the reduced dimension indicated on the  $y$ -axis. Most of these dimensions are between 5 and 40, meaning that the original  $M = 260$  signal dimension can be reduced by a factor of approximately 7 to 10 leading to a reduction from 13 To to approximately 1.5 To of signals stored in memory. To check, how much is lost in this reduction, we compute the root mean square error (RMSE) between the dictionary  $\mathbf{y}_i$  signals and their reconstructions  $\check{\mathbf{y}}_i$  from their reduced representations (11),

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (\mathbf{y}_i - \check{\mathbf{y}}_i)^2}.$$

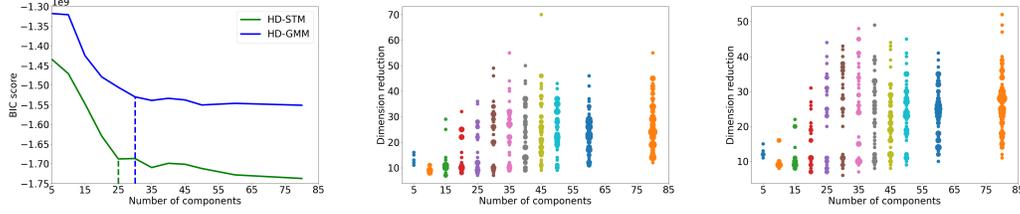


Fig 3: Model selection. BIC scores (left) and clusters dimensions, with respect to the number of components  $K$ , for HD-GMM (middle) and HD-STM (right) applied to signals of dimension  $M = 260$ . Points sizes reflect the proportion of clusters with a given dimension ( $y$ -axis).

Figure 4 shows as expected that the reconstruction RMSE decreases when  $K$  increases and is smaller for HD-STM in particular for small values of  $K$ . For the selected  $K$  values it represents about 2.5% of the signal.

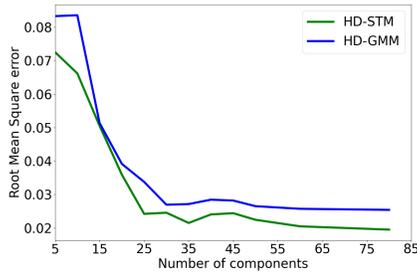


Fig 4: Dictionary reduction loss as measured by RMSE values over the dictionary signals, for HD-GMM and HD-STM, with respect to the number of components  $K$ .

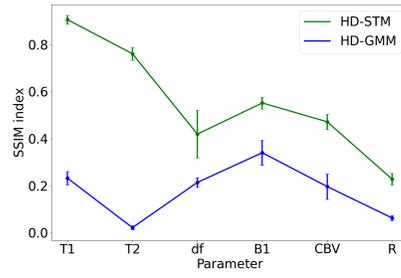


Fig 5: Parameter maps comparison. Average SSIM in  $[-1,1]$ , the higher the better, and standard deviation over subjects, between HD-GMM (resp. HD-STM) and full matching maps.

**7.4. Tissue and sensitivity parameters reconstruction.** For the six subjects, we compare the parameter maps obtained with our Algorithm 1 using HD-GMM and HD-STM to that obtained with traditional matching, referred to as *full matching* (*i.e.*, matching with uncompressed signals). While *full matching* cannot be considered the ground truth, it remains the reference method due to its robustness to the slightly undersampled MRF acquisitions used in this study, as demonstrated in (Coudert et al., 2024), despite its high computational and memory demands.

Table 1 reports for each of the 6 parameters, the distances to *full matching* parameter estimations, showing the average Mean Absolute Errors (*MAE*) across voxels for all  $n_{subject} = 6$  subjects and all slices. For another comparison that takes into account the image structure, Figure 5 shows the reconstruction quality as measured by the structural similarity index measure (SSIM) (Wang et al., 2004), for both HD-STM and HD-GMM when compare to full matching, for each parameter and averaged over subjects. The SSIM is a decimal value between -1 and 1, where 1 indicates perfect similarity, 0 indicates no similarity, and -1

indicates perfect anti-correlation.

$$MAE = \frac{1}{n_{\text{subject}}} \sum_{s=1}^{n_{\text{subject}}} \frac{1}{\tilde{N}_s} \sum_{j=1}^{\tilde{N}_s} |\tilde{t}_{j,s} - t_{j,s}^{\text{match}}|,$$

$$SSIM = \frac{1}{n_{\text{subject}}} \sum_{s=1}^{n_{\text{subject}}} SSIM \left( (\tilde{t}_{j,s})_{1 \leq j \leq \tilde{N}_s}, (t_{j,s}^{\text{match}})_{1 \leq j \leq \tilde{N}_s} \right).$$

TABLE 1

Parameter maps reconstruction. Average MAE and standard deviation, over voxels and subjects, for HD-GMM ( $K = 30$ ) and HD-STM ( $K = 25$ ) with respect to full matching. Best values in bold.

Parameter	$T_1$ (ms)	$T_2$ (ms)	$\delta f$ (Hz)	$B_1$ sensitivity ( $10^{-3}$ )	CBV (%)	$R$ ( $\mu\text{m}$ )
HD-GMM	$462 \pm 29$	$325 \pm 17$	$10 \pm 1$	$92 \pm 18$	$5 \pm 0.8$	$2 \pm 0.03$
HD-STM	<b><math>78 \pm 14</math></b>	<b><math>5 \pm 2</math></b>	<b><math>6 \pm 1</math></b>	<b><math>40 \pm 5</math></b>	<b><math>1.5 \pm 0.5</math></b>	<b><math>1.6 \pm 0.06</math></b>

HD-STM consistently outperforms HD-GMM across all parameters, as evidenced by Figures 5 and Table 1. HD-STM achieves superior SSIM and MAE values, indicating better structural correspondence. This advantage is particularly pronounced for parameters such as  $T_1$  and  $T_2$ , where HD-STM demonstrates robust reconstruction capabilities, whereas HD-GMM fails to accommodate too much undersampling noise. However, for parameters like CBV and  $R$ , HD-STM’s performance declines, suggesting that these parameters are inherently more challenging to model accurately. In these cases, HD-GMM yields the lowest SSIM values and the highest MAE values, further underscoring its limitations. Overall, while HD-STM proves to be the more reliable method, the significant variability across subjects—particularly for  $\delta f$  and CBV—and the observed SSIM drop for  $R$  highlight the need for further refinements.

For another assessment of the maps quality, Table 2 presents the mean values and standard deviations of  $T_1$ ,  $T_2$ ,  $CBV$ , and  $R$ , computed over voxels in white and gray matter ROIs delineated on  $T_1$  maps obtained from *full matching*. Compared to the ranges for healthy subjects reported by Wansapura et al. (1999); Bjørnerud and Emblem (2010); Delphin et al. (2023), HD-STM produces values that are both more consistent with expected ranges and closer to those obtained using full matching than to those derived from HD-GMM. In particular, with HD-GMM,  $T_2$  values significantly depart from the literature reference, as also visible in Figure 6.

Indeed, the conclusions drawn from Figure 5 and Tables 1–2 are further illustrated in Figure 6. This figure highlights HD-GMM’s failure to reconstruct parameters accurately, even when some estimates fall within healthy ranges reported in the literature from Table 2. In contrast, HD-STM consistently delivers parameter estimates for  $T_1$ ,  $T_2$ ,  $\delta f$ , and  $B_1$  that are nearly artifact-free and closely aligned with expected values. Additionally, HD-STM effectively captures primary structures and generates homogeneous maps for vascular parameters such as  $CBV$  and  $R$  even though minor residual *shim* artifacts persist, these are likely due to the high degree of undersampling in the acquired images.

**7.5. Computation times.** Experiments conducted on an Nvidia V100-32gb GPU are fast. It takes only 7 minutes to reconstruct 1 slice of a subject using both HD-GMM and HD-STM, compared to 45 minutes for *full matching*. This represents a 6-fold time reduction for an HD-MED model. However, in more realistic scenarios, calculations are performed locally

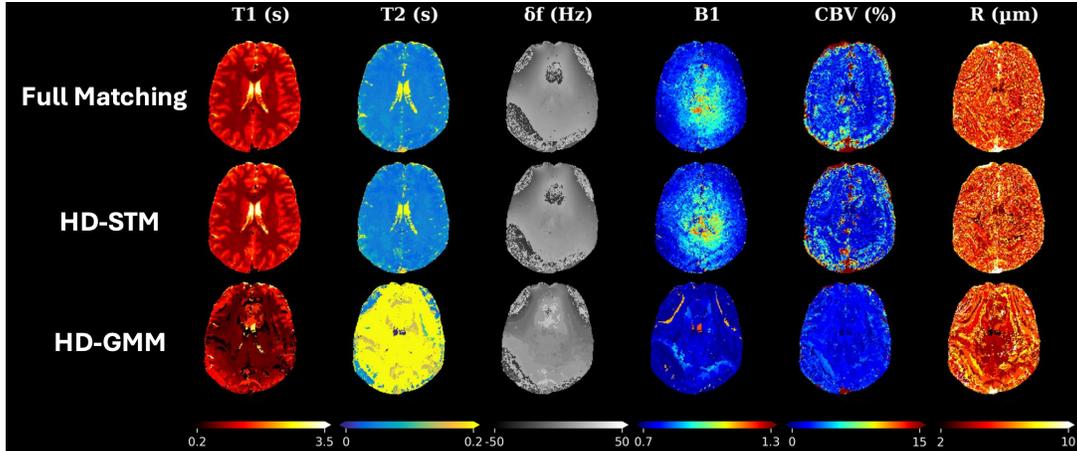


Fig 6:  $T_1$ ,  $T_2$ ,  $\delta f$ ,  $B_1$  sensitivity,  $CBV$ ,  $R$  maps (columns) with different methods (lines) - *full matching*, HD-STM ( $K = 25$ ) and HD-GMM ( $K = 30$ ). Results for all other slices of this subject are shown in Appendix Section 5.

TABLE 2

Mean  $T_1$ ,  $T_2$ ,  $CBV$ ,  $R$  values and standard deviations, in white (WM) and grey (GM) matters ROIs. Values significantly departing from literature reference (last column) are in red.

Parameter	ROI	Full matching	HD-GMM	HD-STM	Literature (Wang et al., 2019b)
$T_1$ (ms)	WM	$891 \pm 6$	$650 \pm 23$	$869 \pm 19$	690-1100
	GM	$1566 \pm 12$	$1650 \pm 57$	$1631 \pm 56$	1286-1393
$T_2$ (ms)	WM	$52 \pm 0.0$	$400 \pm 15$	$48 \pm 2$	56-80
	GM	$95 \pm 4$	$370 \pm 43$	$96 \pm 12$	78-117
$CBV$ (%)	WM	$5.4 \pm 0.6$	$2.0 \pm 0.1$	$4.8 \pm 0.8$	1.7 - 3.6
	GM	$9.2 \pm 1.1$	$4.2 \pm 0.8$	$7.6 \pm 1$	3.0 - 8.0
$R$ ( $\mu s$ )	WM	$5.7 \pm 0.1$	$5.6 \pm 0.1$	$5.7 \pm 0.1$	$6.8 \pm 0.3$
	GM	$6.2 \pm 0.1$	$5.4 \pm 0.4$	$6.1 \pm 0.1$	$7.3 \pm 0.3$

by a medical practitioner on a CPU. On an Apple M2 Pro CPU, *full matching* takes days and is prone to many memory issues that needs to be solve, while the HD-GMM and HD-STM variants require only 4h30.

**8. Conclusion.** In this work, we combine robust latent variable representations, clustering and incremental learning to propose a new tractable and accurate way to represent and handle large volumes of potentially heterogeneous high-dimensional data. To our knowledge, this combination and its use as a data compression strategy is novel. The clustering structure of HD-MED allows to address data heterogeneity, for a greater dimensionality reduction without increasing information loss. Incremental learning allows to handle large volumes, resulting in significant reduction in both computational costs and information losses. In terms of implementation, the procedure is flexible and easy to interpret. It depends on two main hyperparameters, the number of clusters  $K$  and the vector of reduced dimensions  $d$ , that can be set using conventional model selection criteria or chosen by the user to meet resource constraints. Typically increasing  $K$  to reduce the size of clusters or decreasing  $d$  to increase compression.

As an illustration, we focus on a Magnetic Resonance Fingerprinting (MRF) application. The proposed method drastically reduces the computational time required on standard hardware, such as CPUs, which are commonly used in clinical environments. Beyond this crucial processing time gain, dimensionality reduction has also an interesting impact on patient data acquisition time, as it helps mitigating the effect of noise in fast-acquired *in vivo* signals. This reductions are thus not only a technical achievement but also a significant step forward in making advanced MRF techniques more accessible and usable in everyday clinical practice.

In addition, although illustrated with a simple matching procedure, which is the current reference in the target MRF application, our procedure can be coupled with other simulation-based inference (SBI) approaches (Cranmer, Brehmer and Louppe, 2020), such as approximate Bayesian computation (ABC) or other neural techniques (Barrier et al., 2024). Like MRF, simulation-based inference has to face two opposite requirements, which are the need for large and high-dimensional simulated data sets to accurately capture information on the physics under study and the issue of handling such large volumes due to computational resources constraints. ABC uses distances between observations and simulations, and inference is based on these distances. A reduced data representation can be typically coupled with ABC techniques proposing automatic summary statistics selection such as in Forbes et al. (2022); Fearnhead and Prangle (2012). It can also be used with SBI neural approaches, e.g. Häggström et al. (2024), when coupled with a regression model for high-dimensional data, see Algorithm 1 in supplementary material. Compare to matching and standard ABC, regression-based procedures are more amortized, which further reduces the reconstruction time. In terms of MRF and other medical imaging applications, future work will thus involve combining HD-MED with a regression model or a neural network, such as in Deleforge, Forbes and Horaud (2015) or Golbabaee et al. (2019).

At last, none of the previously mentioned methods, currently make use of spatial information at the voxel level, to perform map reconstruction. The idea would be to exploit voxel proximity to either improve or accelerate parameter prediction. A relatively easy extension of mixture-based methods is to account for spatial information by adding a Markov dependence on the clusters, as done for instance in Deleforge et al. (2015).

**9. Financial disclosure.** G. Oudoumanessah was financially supported by the AURA region, and was granted access to the HPC resources of IDRIS under the allocation 2022-AD011013867 made by GENCI.

The MRI facility IRMaGe is partly funded by the French program “Investissement d’avenir” run by the French National Research Agency, grant “Infrastructure d’avenir en Biologie et Santé” [ANR-11-INBS-0006]. The project is supported by the French National Research Agency [ANR-20-CE19-0030 MRFUSE].

## REFERENCES

- ARCHAMBEAU, C. and BACH, F. R. (2008). Sparse probabilistic projections. In *21st International Conference on Neural Information Processing Systems* 73–80.
- ARCHAMBEAU, C., DELANNAY, N. and VERLEYSEN, M. (2006). Robust probabilistic projections. In *Proceedings of the 23rd International conference on machine learning* 33–40.
- ARCHAMBEAU, C., DELANNAY, N. and VERLEYSEN, M. (2008). Mixtures of robust probabilistic principal component analyzers. *Neurocomputing* **71** 1274–1282.
- BAEK, J. and MCLACHLAN, G. J. (2011). Mixtures of common t-factor analyzers for clustering high-dimensional microarray data. *Bioinformatics* **27** 1269–1276.
- BALZANO, L., CHI, Y. and LU, Y. M. (2018). Streaming PCA and Subspace Tracking: The Missing Data Case. In *Proceedings of the IEEE* **106** 1293–1310.

- BARRIER, A., COUDERT, T., DELPHIN, A., LEMASSON, B. and CHRISTEN, T. (2024). MARVEL: MR Fingerprinting with Additional micRoVascular Estimates using bidirectional LSTMs. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* 259–269. Springer.
- BELLAS, A., BOUVEYRON, C., COTTRELL, M. and LACAÏLLE, J. (2013). Model-based Clustering of High-dimensional Data Streams with Online Mixture of Probabilistic PCA. *Advances in Data Analysis and Classification* **7** 281–300.
- BIPIN MEHTA, B., COPPO, S., FRANCES MCGIVNEY, D., IAN HAMILTON, J., CHEN, Y., JIANG, Y., MA, D., SEIBERLICH, N., GULANI, V. and ALAN GRISWOLD, M. (2019). Magnetic resonance fingerprinting: a technical review. *Magnetic resonance in medicine* **81** 25–46.
- BJØRNERUD, A. and EMBLEM, K. E. (2010). A fully automated method for quantitative cerebral hemodynamic analysis using DSC–MRI. *Journal of Cerebral Blood Flow & Metabolism* **30** 1066–1078.
- BORKAR, V. S. (2009). *Stochastic approximation: a dynamical systems viewpoint* **48**. Springer, Gurgaon.
- BOUVEYRON, C. and BRUNET-SAUMARD, C. (2014). Model-based clustering of high-dimensional data: A review. *Computational Statistics & Data Analysis* **71** 52–78.
- BOUVEYRON, C., GIRARD, S. and SCHMID, C. (2007). High-dimensional data clustering. *Computational statistics & data analysis* **52** 502–519.
- BOUX, F., FORBES, F., ARBEL, J., LEMASSON, B. and BARBIER, E. L. (2021). Bayesian inverse regression for vascular magnetic resonance fingerprinting. *IEEE Transactions on Medical Imaging* **40** 1827–1837.
- BRADBURY, J., FROSTIG, R., HAWKINS, P., JOHNSON, M., LEARY, C., MACLAURIN, D., NECULA, G., PASZKE, A., VANDERPLAS, J., WANDERMAN-MILNE, S. and ZHANG, Q. (2018). JAX: composable transformations of Python+NumPy programs. <http://github.com/google/jax>.
- CABINI, R. F., BARZAGHI, L., CICOLARI, D., AROSIO, P., CARRAZZA, S., FIGINI, S., FILIBIAN, M., GAZZANO, A., KRAUSE, R., MARIANI, M. et al. (2024). Fast deep learning reconstruction techniques for preclinical magnetic resonance fingerprinting. *NMR in Biomedicine* **37** e5028.
- CAMBANIS, S., HUANG, S. and SIMONS, G. (1981). On the theory of elliptically contoured distributions. *Journal of Multivariate Analysis* **11** 368–385.
- CAPPÉ, O. and MOULINES, E. (2009). On-line expectation–maximization algorithm for latent data models. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **71** 593–613.
- CAULEY, S. F., SETSOMPOP, K., MA, D., JIANG, Y., YE, H., ADALSTEINSSON, E., GRISWOLD, M. A. and WALD, L. L. (2015). Fast group matching for MR fingerprinting reconstruction. *Magnetic resonance in medicine* **74** 523–528.
- CHRISTEN, T., BOLAR, D. S. and ZAHARCHUK, G. (2013). Imaging brain oxygenation with MRI using blood oxygenation approaches: methods, validation, and clinical applications. *American journal of neuroradiology* **34** 1113–1123.
- CHRISTEN, T., PANNETIER, N., NI, W. W., QIU, D., MOSELEY, M. E., SCHUFF, N. and ZAHARCHUK, G. (2014). MR vascular fingerprinting: A new approach to compute cerebral blood volume, mean vessel radius, and oxygenation maps in the human brain. *Neuroimage* **89** 262–270.
- COHEN, O., ZHU, B. and ROSEN, M. S. (2018). MR fingerprinting deep reconstruction network (DRONE). *Magnetic resonance in medicine* **80** 885–894.
- COIFMAN, R. R. and LAFON, S. (2006). Diffusion maps. *Applied and Computational Harmonic Analysis* **21** 5–30.
- COUDERT, T., DELPHIN, A., BARRIER, A., LEGRIS, L., WARNKING, J. M., LAMALLE, L., DONEVA, M., LEMASSON, B., BARBIER, E. L. and CHRISTEN, T. (2024). Relaxometry and contrast-free cerebral microvascular quantification using balanced Steady-State Free Precession MR Fingerprinting. <https://arxiv.org/abs/2411.03414>.
- CRANMER, K., BREHMER, J. and LOUPPE, G. (2020). The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences* **117** 201912789.
- D’ASPREMONT, A., ELGHAOUI, L., JORDAN, M. I. and LANCKRIET, G. (2007). A direct formulation for sparse PCA using semidefinite programming. *SIAM Review* **49** 434–48.
- DELEFORGE, A., FORBES, F. and HORAUD, R. (2015). High-Dimensional Regression with Gaussian Mixtures and Partially-Latent Response Variables. *Statistics & Computing* **25** 893–911.
- DELEFORGE, A., FORBES, F., BA, S. O. and HORAUD, R. (2015). Hyper-Spectral Image Analysis With Partially Latent Regression and Spatial Markov Dependencies. *IEEE J. Sel. Top. Signal Process.* **9** 1037–1048.
- DELPHIN, A., COUDERT, T., FAN, A., MOSELEY, M. E., ZAHARCHUK, G. and CHRISTEN, T. (2023). MR Vascular Fingerprinting with 3D realistic blood vessel structures and machine learning to assess oxygenation changes in human volunteers. In *2023 ISMRM & ISMRT Annual Meeting & Exhibition*.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society: series B (methodological)* **39** 1–22.

- FEARNHEAD, P. and PRANGLE, D. (2012). Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **74** 419–474.
- FORBES, F., NGUYEN, H. D., NGUYEN, T. and ARBEL, J. (2022). Summary statistics and discrepancy measures for approximate Bayesian computation via surrogate posteriors. *Statistics and Computing* **32** 85.
- FORT, G., MOULINES, E. and WAI, H.-T. (2020). A stochastic path-integrated differential estimator expectation maximization algorithm. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*.
- GHAHRAMANI, Z. and HINTON, G. E. (1997). The EM algorithm for mixtures of factor analyzers Technical Report, University of Toronto.
- GILMAN, K., HONG, D., FESSLER, J. A. and BALZANO, L. (2023). Streaming Probabilistic PCA for Missing Data with Heteroscedastic Noise. *arXiv preprint arXiv:2310.06277*.
- GIRAUD, C. (2014). *Introduction to high-dimensional statistics*. Chapman and Hall/CRC.
- GOLBABAEE, M., CHEN, D., GÓMEZ, P. A., MENZEL, M. I. and DAVIES, M. E. (2019). Geometry of deep learning for magnetic resonance fingerprinting. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 7825–7829. IEEE.
- GOMEZ, E., GOMEZ-VILLEGAS, M. and MARIN, J. (2008). Multivariate Exponential Power Distributions as Mixtures of Normal Distributions with Bayesian Applications. *Communications in Statistics-theory and Methods* **37**.
- GÓMEZ-SÁNCHEZ-MANZANO, E., GÓMEZ-VILLEGAS, M. and MARÍN, J. (2006). Sequences of elliptical distributions and mixtures of normal distributions. *Journal of multivariate analysis* **97** 295–310.
- GU, Y., WANG, C. Y., ANDERSON, C. E., LIU, Y., HU, H., JOHANSEN, M. L., MA, D., JIANG, Y., RAMOS-ESTEBANEZ, C., BRADY-KALNAY, S. et al. (2018). Fast magnetic resonance fingerprinting for dynamic contrast-enhanced studies in mice. *Magnetic resonance in medicine* **80** 2681–2690.
- GU, Y., PAN, Y., FANG, Z., MA, L., ZHU, Y., ANDROJNA, C., ZHONG, K., YU, X. and SHEN, D. (2024). Deep learning-assisted preclinical MR fingerprinting for sub-millimeter T1 and T2 mapping of entire macaque brain. *Magnetic Resonance in Medicine* **91** 1149–1164.
- HÄGGSTRÖM, H., RODRIGUES, P. L., OUDOUMANESSAH, G., FORBES, F. and PICCHINI, U. (2024). Fast, accurate and lightweight sequential simulation-based inference using Gaussian locally linear mappings. *arXiv preprint arXiv:2403.07454*.
- HALKO, N., MARTINSSON, P.-G. and TROPP, J. A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review* **53** 217–288.
- HARGREAVES, B. Bloch Equation Simulator. <http://www-mrsl.stanford.edu/~brian/blochsim/>.
- HONG, D., GILMAN, K., BALZANO, L. and FESSLER, J. A. (2021). HePPCAT: Probabilistic PCA for data with heteroscedastic noise. *IEEE Transactions on Signal Processing* **69** 4819–4834.
- HONG, D., YANG, F., FESSLER, J. A. and BALZANO, L. (2023). Optimally Weighted PCA for High-Dimensional Heteroscedastic Data. *SIAM Journal on Mathematics of Data Science* **5** 222–250.
- JOLLIFFE, I. T. and CADIMA, J. (2016). Principal component analysis: a review and recent developments. *Philosophical transactions of the royal society A: Mathematical, Physical and Engineering Sciences* **374** 20150202.
- KARIMI, B., MIASOJEDOW, B., MOULINES, E. and WAI, H.-T. (2019a). Non-asymptotic analysis of biased stochastic approximation scheme. *Proc. Mach. Learn. Res.* **99** 1-31.
- KARIMI, B., WAI, H.-T., MOULINES, R. and LAVIELLE, M. (2019b). On the global convergence of (fast) incremental expectation maximization methods. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS)*.
- KELKER, D. (1970). Distribution theory of spherical distributions and a location-scale parameter generalization. *Sankhya: The Indian Journal of Statistics, Series A* **32** 419–430.
- KÖRZDÖRFER, G., PFEUFFER, J., KLUGE, T., GEBHARDT, M., HENSEL, B., MEYER, C. H. and NITTKA, M. (2019). Effect of spiral undersampling patterns on FISP MRF parameter maps. *Magnetic resonance imaging* **62** 174–180.
- KOTZ, S. and NADARAJAH, S. (2004). *Multivariate t-distributions and their applications*. Cambridge University Press.
- KUHN, E., MATIAS, C. and REBAFKA, T. (2020). Properties of the stochastic approximation EM algorithm with mini-batch sampling. *Stat. Comp.* **30** 1725-1739.
- LI, P. and HU, Y. (2023). Learned Tensor Low-CP-Rank and Bloch Response Manifold Priors for Non-Cartesian MRF Reconstruction. *IEEE Transactions on Medical Imaging*.
- LI, P. and HU, Y. (2024). Deep magnetic resonance fingerprinting based on Local and Global Vision Transformer. *Medical Image Analysis* **95** 103198.
- MA, D., GULANI, V., SEIBERLICH, N., LIU, K., SUNSHINE, J. L., DUERK, J. L. and GRISWOLD, M. A. (2013). Magnetic resonance fingerprinting. *Nature* **495** 187–192.

- MAIRE, F., MOULINES, E. and LEFEBVRE, S. (2017). Online EM for functional data. *Computational Statistics and Data Analysis* **111** 27–47.
- MC GIVNEY, D. F., PIERRE, E., MA, D., JIANG, Y., SAYBASILI, H., GULANI, V. and GRISWOLD, M. A. (2014). SVD compression for magnetic resonance fingerprinting in the time domain. *IEEE transactions on medical imaging* **33** 2311–2322.
- MC GIVNEY, D. F., BOYACIOĞLU, R., JIANG, Y., POORMAN, M. E., SEIBERLICH, N., GULANI, V., KEENAN, K. E., GRISWOLD, M. A. and MA, D. (2020). Magnetic resonance fingerprinting review part 2: Technique and directions. *Journal of Magnetic Resonance Imaging* **51** 993–1007.
- MCLACHLAN, G. J. and KRISHNAN, T. (2008). *The EM Algorithm and Extensions*. Wiley 2nd edition.
- MCLACHLAN, G. J., PEEL, D. and BEAN, R. (2003). Modelling high-dimensional data by mixtures of factor analyzers. *Computational Statistics & Data Analysis* **41** 379–388.
- MENG, X.-L. and RUBIN, D. B. (1993). Maximum likelihood estimation via the ECM algorithm: A general framework. *Biometrika* **80** 267–278.
- MONGA, A., SINGH, D., DE MOURA, H. L., ZHANG, X., ZIBETTI, M. V. and REGATTE, R. R. (2024). Emerging Trends in Magnetic Resonance Fingerprinting for Quantitative Biomedical Imaging Applications: A Review. *Bioengineering* **11** 236.
- NGUYEN, H. D., FORBES, F. and MCLACHLAN, G. J. (2020). Mini-batch learning of exponential family finite mixture models. *Statistics and Computing* **30** 731–748.
- NGUYEN, H. and FORBES, F. (2022). Global implicit function theorems and the online expectation–maximisation algorithm. *Australian & New Zealand Journal of Statistics* **64**.
- OUDOUMANESSAH, G., COUDERT, T., MEYER, L., DELPHIN, A., DOJAT, M., LARTIZIEN, C. and FORBES, F. (2024). Cluster globally, Reduce locally: Scalable efficient dictionary compression for magnetic resonance fingerprinting. <https://arxiv.org/abs/2411.07415>.
- POORMAN, M. E., MARTIN, M. N., MA, D., MC GIVNEY, D. F., GULANI, V., GRISWOLD, M. A. and KEENAN, K. E. (2020). Magnetic resonance fingerprinting Part I: Potential uses, current challenges, and recommendations. *Journal of Magnetic Resonance Imaging* **51** 675–692.
- RONNEBERGER, O., FISCHER, P. and BROX, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18* 234–241. Springer.
- SATOPAA, V., ALBRECHT, J., IRWIN, D. and RAGHAVAN, B. (2011). Finding a "kneedle" in a haystack: Detecting knee points in system behavior. In *2011 31st international conference on distributed computing systems workshops* 166–171. IEEE.
- SCHOTT, J. R. (2016). *Matrix analysis for statistics*. Wiley probability & statistics, Hoboken.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *The Annals of Statistics* 461–464.
- SOYAK, R., NAVRUZ, E., ERSOY, E. O., CRUZ, G., PRIETO, C., KING, A. P., UNAY, D. and OKSUZ, I. (2021). Channel attention networks for robust MR fingerprint matching. *IEEE Transactions on Biomedical Engineering* **69** 1398–1405.
- TIPPAREDDY, C., ZHAO, W., SUNSHINE, J. L., GRISWOLD, M., MA, D. and BADVE, C. (2021). Magnetic resonance fingerprinting: an overview. *European Journal of Nuclear Medicine and Molecular Imaging* **48** 4189–4200.
- TIPPING, M. E. and BISHOP, C. M. (1999a). Probabilistic principal component analysis. *Journal of the Royal Statistical Society B: Statistical Methodology* **61** 611–622.
- TIPPING, M. E. and BISHOP, C. M. (1999b). Mixtures of probabilistic principal component analyzers. *Neural computation* **11** 443–482.
- TIPPING, M. E. and BISHOP, C. M. (1999c). Probabilistic principal component analysis. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **61** 611–622.
- TIPPING, M. E. and BISHOP, C. M. (1999d). Mixtures of probabilistic principal component analyzers. *Neural computation* **11** 443–482.
- ULLAH, I., HASSAN, A. M., SAAD, R. M. and OMER, H. (2023). GPU accelerated grouped magnetic resonance fingerprinting using clustering techniques. *Magnetic Resonance Imaging* **97** 13–23.
- VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L., GOMEZ, A. N., KAISER, Ł. and POLOSUKHIN, I. (2017). Attention is all you need. *Advances in neural information processing systems* **30**.
- WANG, Z., BOVIK, A. C., SHEIKH, H. R. and SIMONCELLI, E. P. (2004). Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing* **13** 600–612.
- WANG, Z., ZHANG, J., CUI, D., XIE, J., LYU, M., HUI, E. S. and WU, E. X. (2019a). Magnetic Resonance Fingerprinting Using a Fast Dictionary Searching Algorithm: MRF-ZOOM. *IEEE trans. bio-medical engineering* **66** 1526–1535.
- WANG, C. Y., COPPO, S., MEHTA, B. B., SEIBERLICH, N., YU, X. and GRISWOLD, M. A. (2019b). Magnetic resonance fingerprinting with quadratic RF phase for measurement of T2\* simultaneously with  $\delta f$ , T1, and T2. *Magnetic resonance in medicine* **81** 1849–1862.

- WANSAPURA, J. P., HOLLAND, S. K., DUNN, R. S. and BALL JR, W. S. (1999). NMR relaxation times in the human brain at 3.0 tesla. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* **9** 531–538.
- XU, A. S., BALZANO, L. and FESSLER, J. A. (2023). HeMPPCAT: mixtures of probabilistic principal component analysers for data with heteroscedastic noise. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 1–5.
- YANG, M., MA, D., JIANG, Y., HAMILTON, J., SEIBERLICH, N., GRISWOLD, M. A. and MCGIVNEY, D. (2018). Low rank approximation methods for MR fingerprinting with large scale dictionaries. *Magnetic resonance in medicine* **79** 2392–2400.
- YE, H., CAULEY, S. F., GAGOSKI, B., BILGIC, B., MA, D., JIANG, Y., DU, Y. P., GRISWOLD, M. A., WALD, L. L. and SETSOMPOP, K. (2017). Simultaneous multislice magnetic resonance fingerprinting (SMS-MRF) with direct-spiral slice-GRAPPA (ds-SG) reconstruction. *Magnetic resonance in medicine* **77** 1966–1974.
- ZOU, H., HASTIE, T. and TIBSHIRANI, R. (2006). Sparse Principal Component Analysis. *Journal of Computational and Graphical Statistics* **15** 265–286.

**SUPPLEMENTARY MATERIAL**  
**SCALABLE MAGNETIC RESONANCE FINGERPRINTING: INCREMENTAL**  
**INFERENCE OF HIGH DIMENSIONAL ELLIPTICAL MIXTURES FROM**  
**LARGE DATA VOLUMES**

**1. Online EM for finite mixture model.** This section completes Section 6.1 of the main paper. We demonstrate that the online EM (OEM) algorithm for a finite mixture model can be derived from the OEM algorithm for a single component. We denote the mixture parameters as  $\Theta_{\mathcal{M}}$ , which include the pairs  $(\pi_k, \Theta_k)$  for  $k = 1, \dots, K$ , where  $\Theta_k$  are the component-wise parameters of the finite mixture model. Additionally, we consider the pair  $(\mathbf{Y}^T, Z)$ , where  $Z = 1 : K$  is a latent variable such that  $\mathbb{P}(Z = k) = \pi_k$ .

We assume that the complete-data likelihood for each component can be expressed in the exponential family form, as given by Equation (12) in the main paper, for each  $f_c(\mathbf{y}, z; \Theta_k)$ . Following ?, the complete-data likelihood for the mixture model  $f_c(\mathbf{y}, z, \Theta_{\mathcal{M}})$  can be written as:

$$(1) \quad f_c(\mathbf{y}, z, \Theta_{\mathcal{M}}) = h(\mathbf{y}, z) \exp \left[ \sum_{\xi=1}^K [\log \pi_{\xi} + [\mathbf{s}(\mathbf{y}, z)]^T \phi(\Theta_{\xi}) - \psi(\Theta_{\xi})] \right]$$

$$= h_{\mathcal{M}}(\mathbf{y}, z) \exp \left[ [\mathbf{s}_{\mathcal{M}}(\mathbf{y}, z)]^T \phi_{\mathcal{M}}(\Theta_{\mathcal{M}}) - \psi_{\mathcal{M}}(\Theta_{\mathcal{M}}) \right],$$

where  $h_{\mathcal{M}} = h$ ,  $\psi_{\mathcal{M}} = \mathbf{0}$ ,

$$(2) \quad \mathbf{s}_{\mathcal{M}}(\mathbf{y}, z) = \begin{bmatrix} \mathbf{1}_{z=1} \\ \mathbf{1}_{z=1} \mathbf{s}(\mathbf{y}, z) \\ \vdots \\ \mathbf{1}_{z=K} \\ \mathbf{1}_{z=K} \mathbf{s}(\mathbf{y}, z) \end{bmatrix}, \text{ and } \phi_{\mathcal{M}}(\Theta_{\mathcal{M}}) = \begin{bmatrix} \log \pi_1 - \psi(\Theta_1) \\ \phi(\Theta_1) \\ \vdots \\ \log \pi_K - \psi(\Theta_K) \\ \phi(\Theta_K) \end{bmatrix}.$$

We can conveniently write the vector  $\mathbf{s}_{\mathcal{M}}$  that matches the dimension of  $\mathbf{s}_{\mathcal{M}}(\mathbf{y}, z)$  as

$$(3) \quad \mathbf{s}_{\mathcal{M}}^T = (s_{01}, \mathbf{s}_{\mathcal{M}1}^T, \dots, s_{0K}, \mathbf{s}_{\mathcal{M}K}^T),$$

where  $\mathbf{s}_{\mathcal{M}z}$  corresponds to one component of the exponential family form of the mixture model distribution. Via Equation (2), the objective function  $Q_{\mathcal{M}}$  for the mixture complete-data likelihood as defined in Equation (14) of the main paper can be written as

$$(4) \quad Q_{\mathcal{M}}(\mathbf{s}_{\mathcal{M}}, \Theta_{\mathcal{M}}) = \mathbf{s}_{\mathcal{M}}^T \phi_{\mathcal{M}}(\Theta_{\mathcal{M}}) = \sum_{z=1}^K s_{0z} (\log \pi_z - \psi(\Theta_z)) + \mathbf{s}_{\mathcal{M}z}^T \phi(\Theta_z).$$

Whatever the form of the component *p.d.f.*, the maximisation with respect to  $\pi_z$  yields

$$(5) \quad \bar{\pi}_z(\mathbf{s}_{\mathcal{M}}) = \frac{s_{0z}}{\sum_{\xi=1}^K s_{0\xi}}.$$

Then, for each  $z \in \{1 \dots K\}$ ,

$$\frac{\partial Q_{\mathcal{M}}}{\partial \Theta_z}(\mathbf{s}_{\mathcal{M}}, \Theta_{\mathcal{M}}) = -s_{0z} \frac{\partial \psi}{\partial \Theta_z}(\Theta_z) + \mathbf{J}_{\phi}(\Theta_z) \mathbf{s}_{\mathcal{M}z}$$

$$\begin{aligned}
&= s_{0z} \left( \mathbf{J}_\phi(\boldsymbol{\Theta}_z) \begin{bmatrix} \mathbf{s}_{\mathcal{M}z} \\ s_{0z} \end{bmatrix} - \frac{\partial \psi}{\partial \boldsymbol{\Theta}_z} \right) \\
(6) \quad &= s_{0z} \frac{\partial Q}{\partial \boldsymbol{\Theta}_z} \left( \begin{bmatrix} \mathbf{s}_{\mathcal{M}z} \\ s_{0z} \end{bmatrix}, \boldsymbol{\Theta}_z \right),
\end{aligned}$$

where  $Q$  is the objective function of form Equation (14) from main paper corresponding to a single component *p.d.f.* Since  $s_{0z} > 0$ , for all  $z \in \{1 \dots K\}$ , it follows that the maximization of  $Q_{\mathcal{M}}$  can be conducted by solving

$$(7) \quad \frac{\partial Q}{\partial \boldsymbol{\Theta}_z} \left( \begin{bmatrix} \mathbf{s}_{\mathcal{M}z} \\ s_{0z} \end{bmatrix}, \boldsymbol{\Theta}_z \right) = \mathbf{0},$$

with respect to  $\boldsymbol{\Theta}_z$  for each  $z$ . Therefore, it is enough to show that for each component  $z$  of the finite mixture model that there exists a root of the equation above denoted as  $\bar{\boldsymbol{\Theta}}_z(\mathbf{s}_{\mathcal{M}z}/s_{0z})$  in order to find a solution for the maximizer of the mixture objective  $Q_{\mathcal{M}}$ . Then we can set

$$(8) \quad \bar{\boldsymbol{\Theta}}_{\mathcal{M}}(\mathbf{s}_{\mathcal{M}}) = \begin{bmatrix} \bar{\pi}_1(\mathbf{s}_{\mathcal{M}}) \\ \bar{\boldsymbol{\Theta}}_1(\mathbf{s}_{\mathcal{M}1}/s_{01}) \\ \vdots \\ \bar{\pi}_K(\mathbf{s}_{\mathcal{M}}) \\ \bar{\boldsymbol{\Theta}}_K(\mathbf{s}_{\mathcal{M}K}/s_{0K}) \end{bmatrix},$$

where the form of  $\bar{\pi}_K(\mathbf{s}_{\mathcal{M}})$  is given in (5).

To complete the OEM algorithm, we need to specify the quantity  $\bar{\mathbf{s}}_{\mathcal{M}}(\mathbf{y}; \boldsymbol{\Theta}_{\mathcal{M}}) = \mathbb{E}[\mathbf{s}_{\mathcal{M}}(\mathbf{Y}, Z) | \mathbf{Y} = \mathbf{y}; \boldsymbol{\Theta}_{\mathcal{M}}]$  as defined in Equation (13) of the main paper. Using the definition of  $\mathbf{s}_{\mathcal{M}}(\mathbf{y}, z)$  in Equation (2), we compute,

$$(9) \quad r_k = \mathbb{E}[\mathbf{1}_{\{Z=k\}} | \mathbf{Y} = \mathbf{y}; \boldsymbol{\Theta}_{\mathcal{M}}] = \frac{\pi_k f_k(\mathbf{y})}{\sum_{\xi=1}^K \pi_\xi f_\xi(\mathbf{y})},$$

and

$$(10) \quad \mathbb{E}[\mathbf{1}_{\{Z=k\}} \mathbf{s}(\mathbf{Y}, Z) | \mathbf{Y} = \mathbf{y}; \boldsymbol{\Theta}_{\mathcal{M}}] = r_k \mathbb{E}[\mathbf{s}(\mathbf{Y}) | \mathbf{Y} = \mathbf{y}, Z = k; \boldsymbol{\Theta}_{\mathcal{M}}],$$

where  $f_k$  is defined in Equation (4) of the main paper as the conditional pdf of the  $k^{\text{th}}$  mixture component.

Thus, in the finite mixture framework, deriving the OEM algorithm for a single component suffices to deduce the OEM algorithm for the entire mixture. In the next section, we derive the OEM algorithm for a single component of HD-MED. For simplicity, we omit further mention of the form of  $h$  in the exponential family representation, as it does not affect the OEM algorithm derivation.

## 2. HD-ED exponential family form: Proof of Proposition 6.1.

PROOF. Recall that  $\tilde{\mathbf{D}}$  denotes the  $d$  first columns of  $\mathbf{D}$  completed with zero columns and  $\bar{\mathbf{D}} = \mathbf{D} - \tilde{\mathbf{D}}$ . The following multivariate normal distribution can be decomposed as follows,

$$\mathcal{N}_M\left(\mathbf{y}; \boldsymbol{\mu}, \frac{\mathbf{D}\mathbf{A}\mathbf{D}^T}{w}\right) = \prod_{m=1}^d \mathcal{N}_1\left([\tilde{\mathbf{D}}^T(\mathbf{y} - \boldsymbol{\mu})]_m; 0, \frac{a_m}{w}\right) \prod_{m=d+1}^M \mathcal{N}_1\left([\bar{\mathbf{D}}^T(\mathbf{y} - \boldsymbol{\mu})]_m; 0, \frac{b}{w}\right),$$

where  $[\cdot]_m$  denotes the  $m^{\text{th}}$  element of the input vector. For  $m \leq d$  it comes,

$$[\tilde{\mathbf{D}}^T(\mathbf{y} - \boldsymbol{\mu})]_m^2 = \text{vec}(\mathbf{d}_m \mathbf{d}_m^T)^T \text{vec}(\mathbf{y} \mathbf{y}^T) + \text{vec}(\mathbf{d}_m \mathbf{d}_m^T)^T \text{vec}(\boldsymbol{\mu} \boldsymbol{\mu}^T) - 2\boldsymbol{\mu}^T \mathbf{d}_m \mathbf{d}_m^T \mathbf{y}.$$

And for  $d < m \leq M$  we have:

$$\prod_{m=d+1}^M \mathcal{N}_1\left([\bar{\mathbf{D}}^T(\mathbf{y} - \boldsymbol{\mu})]_m; 0, \frac{b}{w}\right) = \frac{1}{(2\pi b)^{\frac{M-d}{2}}} \exp\left(-\frac{w}{2b} \sum_{m=d+1}^M \bar{\mathbf{D}}^T(\mathbf{y} - \boldsymbol{\mu})_m^2\right).$$

It is then possible to remove the dependence on  $\bar{\mathbf{D}}$  by noting that

$$\begin{aligned} \sum_{m=d+1}^M [\bar{\mathbf{D}}^T(\mathbf{y} - \boldsymbol{\mu})]_m^2 &= \|\mathbf{y} - P(\mathbf{y})\|^2 \\ &= \mathbf{y}^T \mathbf{y} + \boldsymbol{\mu}^T \boldsymbol{\mu} - 2\mathbf{y}^T \boldsymbol{\mu} - \sum_{m=1}^d [\tilde{\mathbf{D}}^T(\mathbf{y} - \boldsymbol{\mu})]_m^2, \end{aligned}$$

using that  $P(\mathbf{y}) = \tilde{\mathbf{D}} \tilde{\mathbf{D}}^T(\mathbf{y} - \boldsymbol{\mu}) + \boldsymbol{\mu}$ . Rewriting expressions with the vec operator and combining both parts where  $m \leq d$  and  $d \leq m \leq M$  gives the exponential family form for the complete data likelihood (18) in main paper.  $\square$

**3. Gamma-distributed mixing variable.** In Section 6 of the main paper we proposed an online E-step with uncomplete calculations that depend on the exponential family form of mixing variable  $W$ . In this section we propose to specify these calculations when  $W$  follows a Gamma distribution  $f_{\boldsymbol{\theta}}(w) = \mathcal{G}(w; \alpha, \beta)$  with the two parameters  $\boldsymbol{\theta} = (\alpha, \beta)$  in  $\mathbb{R}^{+2}$ . More specifically, we need to compute the expectation (19),  $\mathbb{E}[W|\mathbf{Y}; \boldsymbol{\Theta}]$ , from the main paper, and  $\mathbb{E}[s_w(W)|\mathbf{Y}; \boldsymbol{\Theta}]$ . Then we need to derive the root equations for  $\alpha$  and  $\beta$ . This Gamma-distributed mixing  $W$  encompasses several distributions as indicated in Table 1.

The Gamma distribution has density  $\mathcal{G}(w; \alpha, \beta) = w^{\alpha-1} \frac{\beta^\alpha e^{-\beta w}}{\Gamma(\alpha)}$  with  $\Gamma$  the Gamma function. This distribution is easily put in the exponential family form with

$$(11) \quad \begin{aligned} \mathbf{s}_w(w) &= \begin{bmatrix} w \\ \log w \end{bmatrix}, \quad \boldsymbol{\phi}_w(\alpha, \beta) = \begin{bmatrix} -\beta \\ \alpha \end{bmatrix}, \\ \psi_w(\alpha, \beta) &= \log(\Gamma(\alpha)) - \alpha \log(\beta). \end{aligned}$$

The next proposition gives the necessary formulas to derive the online E-step for an ED with a Gamma-distributed mixing variable.

PROPOSITION 3.1 (Online E-step with Gamma mixing variable).

$$(12) \quad \mathbb{E}[W|\mathbf{Y} = \mathbf{y}] = \frac{\alpha + M}{u + 2\beta}$$

$$(13) \quad \mathbb{E}[\log(W)|\mathbf{Y} = \mathbf{y}] = \Psi^{(0)}\left(\frac{\alpha + 1}{2}\right) - \log\left(\frac{u + 2\beta}{2}\right),$$

with  $u = (\mathbf{y} - \boldsymbol{\mu})^T \mathbf{D} \mathbf{A}^{-1} \mathbf{D}^T (\mathbf{y} - \boldsymbol{\mu})$  and  $\Psi^{(0)}$  the diGamma function defined as the derivative of  $\log \Gamma(\cdot)$ .

PROOF. We have  $(\mathbf{Y}|W = w) \sim \mathcal{N}_M(\boldsymbol{\mu}, \boldsymbol{\Sigma}/w)$  and  $W \sim \mathcal{G}(\alpha, \beta)$ , then

$$\begin{aligned} p(w|\mathbf{Y} = \mathbf{y}) &\propto w^{\frac{\alpha}{2} - \frac{M}{2}} \exp\left(-\frac{w}{2} \left(\frac{u}{2} + \beta\right)\right) \\ &\propto \mathcal{G}\left(w; \frac{\alpha + M}{2}, \frac{u}{2} + \beta\right), \end{aligned}$$

which provides the values for both required expectations.  $\square$

For the online maximization step, we need to find the solution of the root equation (26) in the main paper. By conveniently writing  $\mathbf{s}_w^T = [s_{1w} \ s_{2w}]$  and differentiating with respect to  $\alpha$  and  $\beta$ , it comes the following system of equations,

$$(14) \quad s_{2w} - \Psi^{(0)}(\alpha) + \log \beta = 0$$

$$(15) \quad -s_{1w} + \frac{\alpha}{\beta} = 0.$$

Then substituting  $\alpha/s_{1w}$  to  $\beta$  in (14) and under the condition that  $s_{1w} \neq 0$ , (14) leads to an equation with  $\alpha$  only.

3.1. *Mixture of HD Student-t distributions.* As recalled in Table 1, the Student distribution is an ED with a Gamma-distributed mixing variable with  $\alpha = \beta = \frac{\nu}{2}$  where  $\nu$  is usually referred to as the degrees of freedom. Previous formulas apply and the root equation for  $\nu$  simplifies into,

$$s_{2w} - s_{1w} - \Psi^{(0)}\left(\frac{\nu}{2}\right) + \log\left(\frac{\nu}{2}\right) + 1 = 0.$$

TABLE 1  
Some EDs with Gamma mixing variable.

ED	Mixing variable
Student	$W \sim \mathcal{G}\left(\frac{\nu}{2}, \frac{\nu}{2}\right)$
Pearson Type VII	$W \sim \mathcal{G}\left(\alpha - \frac{p}{2}, \frac{1}{2}\right)$
Generalized Student	$W \sim \mathcal{G}\left(\frac{\nu}{2}, \frac{\alpha}{2}\right)$
Cauchy	$W \sim \mathcal{G}\left(\frac{1}{2}, \frac{1}{2}\right)$

**4. Divide & Conquer high dimensional inverse regression of large data volumes.** As mentioned in the main paper conclusion, the presented data compression procedure could also be combined with a regression model in place of the matching step. Algorithm 1 below is a variant of Algorithm 1 in the paper illustrating this possibility. Step 1 remains the same, Step 2 corresponds to the estimation of regression functions for each cluster and is new, Step 3 is the cluster-wise prediction of new signals using the  $K$  estimated regression functions. More specifically step 3.1 consists of projecting an observation  $\tilde{\mathbf{y}}_j$  on each of the  $K$  subspaces and of computing the probability of  $\tilde{\mathbf{y}}_j$  to be assigned to each cluster. Step 3.2 corresponds then to a cluster-wise prediction when only the regression function for the most probable cluster is used, *i.e.*  $\tilde{\mathbf{t}}_j = \tilde{\mathbf{t}}_j^k$  with  $k = \arg \max_l r_l(\tilde{\mathbf{y}}_j)$ . Step 3.2 is written in a slightly more general manner using  $\mathcal{L}$  to denote a function that would combine, in a way to be chosen, all predictions and all assignment probabilities.

**5. Application to MRF reconstruction.** We provide additional illustration by showing in Figure 1 the parameter maps obtained with full matching and HD-STM, for all the slices of the subject presented in main Figure 6. Both methods, HD-STM and full matching provide visually similar reconstructions.

**6. Generalized scale mixtures of Gaussian distributions.** In the main text, we focused on the case where the covariance matrix of each Gaussian component is scaled by a factor of  $\frac{1}{w}$ . However, as proposed in ?, it is possible to generalize Gaussian scale mixtures by using a more flexible link function  $k(\cdot)$ , leading to the formulation:

**Algorithm 1** Divide & Conquer high dimensional inverse regression of large data volumes

**Input** Data set of simulated (signal, parameters) pairs  $\mathcal{D}_f = \{\mathbf{y}_i, \mathbf{t}_i\}_{i=1:N}$ ,  $N \gg 1$ ,  $\mathbf{t}_i \in \mathbb{R}^L$ ,  $\mathbf{y}_i \in \mathbb{R}^M$ ,  $M \gg 1$ .  
Observed signals  $\{\tilde{\mathbf{y}}_j\}_{j=1:\tilde{N}}$ ,  $\tilde{\mathbf{y}}_j \in \mathbb{R}^M$ .

**1: Reduced dimension representation of simulated data:**  $\{\hat{\mathbf{y}}_i, \mathbf{t}_i, r(\mathbf{y}_i)\}_{i=1:N}$ .

**1.1 Online HD-MED inference from  $\{\mathbf{y}_i\}_{i=1:N}$ :**  $K$  clusters,  $d_k < M$  for  $k = 1 : K \implies$  cluster assignment probabilities and cluster-wise projections  $(\mathbf{r}, \mathbf{Q}) = \{r_k(\cdot), Q_k(\cdot)\}_{k=1:K}$ .

**1.2 Cluster-wise signal reduction:**  $\{\mathbf{y}_i\}_{i=1:N}, \mathbf{r}, \mathbf{Q} \implies \{\hat{\mathbf{y}}_i = Q_k(\mathbf{y}_i), i \in I_k\}$  with  $I_k = \{i, s.t. k = \arg \max_{\ell} r_{\ell}(\mathbf{y}_i)\}$ , for  $k = 1 : K$ .

**2: Cluster-wise regressions:** for  $k = 1 : K$ , learn a regression function  $l_k : \mathbb{R}^{d_k} \rightarrow \mathbb{R}^L$  from  $\{(\mathbf{t}_i, \hat{\mathbf{y}}_i), i \in I_k\}$ .

**3: Cluster-wise predictions for observed signals:**

**3.1 Cluster-wise observed signal reduction:** Use learned  $(\mathbf{r}, \mathbf{Q})$  from step 1.1 to obtain  $\{Q_k(\tilde{\mathbf{y}}_j), r_k(\tilde{\mathbf{y}}_j), j = 1 : \tilde{N}, k = 1 : K\}$

**3.2 Cluster-wise prediction:** For  $j = 1 : \tilde{N}, k = 1 : K$ , compute  $\tilde{\mathbf{t}}_j^k = l_k(Q_k(\tilde{\mathbf{y}}_j)) \implies$  predict  $\tilde{\mathbf{t}}_j = \mathcal{L}(\tilde{\mathbf{t}}_j^1, r_1(\tilde{\mathbf{y}}_j), \dots, \tilde{\mathbf{t}}_j^K, r_K(\tilde{\mathbf{y}}_j))$  e.g.  $\tilde{\mathbf{t}}_j = \tilde{\mathbf{t}}_j^k$  with  $k = \arg \max_{\ell} r_{\ell}(\tilde{\mathbf{y}}_j)$ .

**Return** Predicted parameters  $\{\tilde{\mathbf{t}}_j\}_{j=1:\tilde{N}}$ ,  $\tilde{\mathbf{t}}_j \in \mathbb{R}^L$

$$(16) \quad p(\mathbf{y}) = \int_{\mathbb{R}^+} \mathcal{N}_M(\mathbf{y}; \boldsymbol{\mu}, k(w)\boldsymbol{\Sigma}) f(w) dw.$$

This generalization encompasses a wider class of distributions, as summarized in Table 1 of ?. A possible continuation of our work is to derive properties from the main paper for this more general case, while still applying the proposed parcimonious parameterization of high dimensional scale matrices. It is noteworthy that we retain the ability to define a latent variable model for this extended class of distributions. Using the same notation, the following proposition holds for this new class, denoted by  $\mathcal{HE}$  as in the main paper:

**PROPOSITION 6.1.** *Let  $d \leq M - 1$ ,  $\mathbf{Y} \in \mathbb{R}^M$ ,  $\mathbf{X} \in \mathbb{R}^d$ ,  $\mathbf{E} \in \mathbb{R}^M$ ,  $W \in \mathbb{R}^+$  be random variables,  $\mathbf{V} \in \mathbb{R}^{M \times d}$  a matrix of linearly independent columns,  $\boldsymbol{\mu} \in \mathbb{R}^M$  a vector, a link function  $k : \mathbb{R} \rightarrow \mathbb{R}$  and  $f_{\boldsymbol{\theta}}$  the pdf of a positive univariate random variable defined by some parameter  $\boldsymbol{\theta}$ . Assume that*

$$\begin{aligned} \mathbf{Y} &= \mathbf{V}\mathbf{X} + \boldsymbol{\mu} + \mathbf{E} \\ (\mathbf{X}|W = w) &\sim \mathcal{N}(\mathbf{0}_d, k(w)\mathbf{I}_d) \\ (\mathbf{E}|W = w) &\sim \mathcal{N}(\mathbf{0}_M, b k(w)\mathbf{I}_M) \\ W &\sim f_{\boldsymbol{\theta}} \end{aligned}$$

then,

$$\mathbf{Y} \sim \mathcal{HE}_{Md}(\boldsymbol{\mu}, \tilde{\mathbf{D}}^*, \mathbf{a}, b, \boldsymbol{\theta}),$$

with  $b\mathbf{I}_M + \mathbf{V}\mathbf{V}^T = \mathbf{D}\mathbf{A}\mathbf{D}^T$  and  $\mathbf{A} = \text{diag}(a_1, \dots, a_d, b, \dots, b)$ .

Additionally, denoting by  $\mathbf{U} = b\mathbf{I}_d + \mathbf{V}^T\mathbf{V}$ , we have,

$$(17) \quad (\mathbf{X}|\mathbf{Y} = \mathbf{y}, W = w) \sim \mathcal{N}(\mathbf{U}^{-1}\mathbf{V}^T(\mathbf{y} - \boldsymbol{\mu}), k(w)b\mathbf{U}^{-1}).$$

It follows that  $\mathbb{E}[\mathbf{Y}|\mathbf{X} = \mathbf{x}] = \mathbf{V}\mathbf{x} + \boldsymbol{\mu}$  and  $\mathbb{E}[\mathbf{X}|\mathbf{Y} = \mathbf{y}] = \mathbf{U}^{-1}\mathbf{V}^T(\mathbf{y} - \boldsymbol{\mu})$ .

PROOF. The proof follows the same structure as Proposition 5.1 in the main paper, with the scaling factor  $\frac{1}{w}$  replaced by the link function  $k(w)$ .  $\square$

Regarding the OEM part, we can similarly derive propositions for this new class of distributions, assuming that  $\mathbf{W}$  follows a distribution in the exponential family. For all link functions  $k(w)$ , most calculations involve replacing the scaling factor  $\frac{1}{w}$  with  $k(w)$ . This opens the possibility of studying high-dimensional finite mixtures of distributions such as the slash, variance gamma, and Laplace distributions.

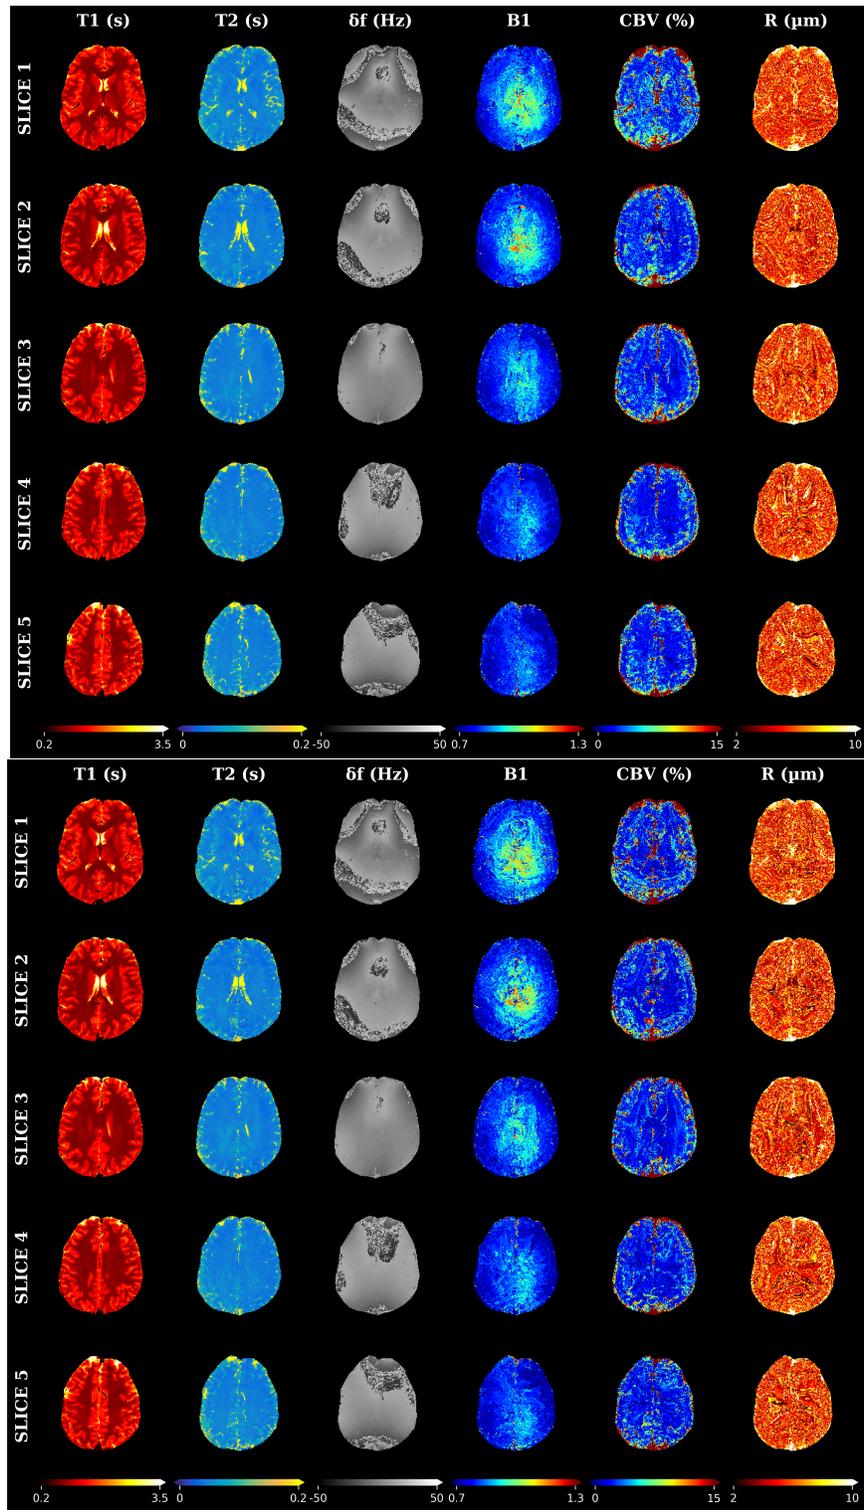


FIG 1.  $T_1$ ,  $T_2$ ,  $\delta f$ ,  $B_1$  sensitivity,  $CBV$ ,  $R$  maps (columns) with different methods- full matching (top), HD-STM with  $K = 25$  (bottom) for all the slices of the subject presented in Figure 6 of the main paper.