

Large-scale gradient-based training of Mixtures of Factor Analyzers

Alexander Gepperth
 Fulda University of Applied Sciences
 Fulda, Germany
 alexander.gepperth@cs.hs-fulda.de

Abstract—Gaussian Mixture Models (GMMs) are a standard tool in data analysis. However, they face problems when applied to high-dimensional data (e.g., images) due to the size of the required full covariance matrices (CMs), whereas the use of diagonal or spherical CMs often imposes restrictions that are too severe. The Mixture of Factor analyzers (MFA) model is an important extension of GMMs, which allows to smoothly interpolate between diagonal and full CMs based on the number of *factor loadings* l . MFA has successfully been applied for modeling high-dimensional image data [1]. This article contributes both a theoretical analysis as well as a new method for efficient high-dimensional MFA training by stochastic gradient descent, starting from random centroid initializations. This greatly simplifies the training and initialization process, and avoids problems of batch-type algorithms such as Expectation-Maximization (EM) when training with huge amounts of data. In addition, by exploiting the properties of the matrix determinant lemma, we prove that MFA training and inference/sampling can be performed based on precision matrices, which does not require matrix inversions after training is completed. At training time, the method requires the inversion of $l \times l$ matrices only. Besides the theoretical analysis and proofs, we apply MFA to typical image datasets such as SVHN and MNIST, and demonstrate the ability to perform sample generation and outlier detection.

Index Terms—Gaussian Mixture Models, Stochastic Gradient Descent, Mixture of Factor Analyzers, Mixture Models

I. INTRODUCTION

This contribution focuses on generative machine learning models applied to images. In the original sense of the term (see [2]), this implies that a learning algorithm aims at explicitly modeling the distribution of the data. This includes sampling from this distribution, but is by no means limited to this function. Most prominently, the direct modeling of the distribution itself allows outlier detection and inference.

Mixture models figure prominently among generative models, since their approach is to represent the overall data distribution as a linear combination of elementary distributions $\Psi_k(\mathbf{x})$, each of which has parameters ξ_k :

$$p(\text{vec } \mathbf{x}) = \sum_k \pi_k \Psi(\mathbf{x}; \xi_k) \equiv \sum_k \pi_k \Psi_k(\mathbf{x}) \quad (1)$$

Mixture modeling allows the derivation of important analytical results, e.g., if the Ψ_k are mathematically well-understood. A particular case is the use of multi-variate normal distributions:

$$\Psi_k(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \equiv \mathcal{N}_k(\mathbf{x}), \quad (2)$$

which gives rise to so-called *Gaussian Mixture Models* (GMMs).

MFA represent a special case of GMMs: for describing d -dimensional variables, MFA assumes the existence of an l -dimensional *latent space* with $l \ll d$ for each mixture component k . In the latent space, variables $\mathbb{R}^l \in \mathbf{z}_k$ follow a simple normal distribution: $\mathbf{z}_k \sim \mathcal{N}(0, \mathbf{I} \in \mathbb{R}^{l \times l})$. The relation between latent and full space is given by a simple generative model:

$$\mathbb{R}^d \ni \mathbf{x}_k = \boldsymbol{\mu}_k + \boldsymbol{\Lambda}_k \mathbf{z}_k + \boldsymbol{\epsilon}_k \quad (3)$$

with $\boldsymbol{\epsilon}_k \sim \mathcal{N}(0, \mathbf{D}_k)$, $\mathbf{D}_k \in \mathbb{R}^{d \times d}$ diagonal and $\boldsymbol{\Lambda}_k \in \mathbb{R}^{d \times l}$. From this, it directly follows that $\mathbf{x}_k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ with

$$\boldsymbol{\Sigma}_k \equiv \mathbf{D}_k + \boldsymbol{\Lambda}_k \boldsymbol{\Lambda}_k^T. \quad (4)$$

Inversely, the probability that a sample \mathbf{x} was generated by mixture component k is given as:

$$\ln p_k(\mathbf{x}) = -\frac{1}{2} \left\{ d \log(2\pi) + \log \det \boldsymbol{\Sigma}_k + \tilde{\mathbf{x}}_k^T \boldsymbol{\Sigma}_k^{-1} \tilde{\mathbf{x}}_k \right\}. \quad (5)$$

We will use the convenient shorthand $\tilde{\mathbf{x}}_k \equiv \mathbf{x} - \boldsymbol{\mu}_k$ throughout this article. The goal of MFA training is to find the *loading matrices* $\boldsymbol{\Lambda}_k$, the component noise matrices \mathbf{D}_k , the component means $\boldsymbol{\mu}_k$ and the component weights π_k . This is achieved by maximizing the MFA log-likelihood:

$$\mathcal{L} = \sum_n \log \sum_k p_k(\mathbf{x}_n) \quad (6)$$

Optimization of \mathcal{L} is typically done using the Expectation-Maximization (EM) algorithm which is applicable to many latent-variable models. It relies on the iterative optimization of a lower-bound to the model log-likelihood, which can be shown to be tight. Since EM is not conceptually based on gradient descent, it does not require learning rates or step sizes to be set and is thus very easy to handle.

A. Motivation

While EM training is feasible for small amounts of low-dimensional data, SGD is a preferable alternative in other cases:

Efficient training for large-scale problems When the number of training samples is high, EM as a batch-type algorithm becomes infeasible since it requires a pass through the whole dataset for each iteration. Stochastic generalizations

of EM to mini-batch type learning exist [3], but involve several new hard-to-tune hyper-parameters. Here, SGD offers a principled alternative. Being based on mini-batches and having a solid theoretical foundation in the Robins-Monro procedure [4], it can be applied to arbitrarily large datasets.

Training from random initial conditions EM convergence is strongly dependent on initialization, and thus initializing EM by a clustering algorithm such as k-means is common. This, too, becomes problematic for large-scale problems since k-means is a batch-type algorithm as well. We therefore propose training MFA by SGD from random initial conditions, which greatly simplifies the training procedure and removes the necessity to process the whole dataset for initialization.

Efficiency of MFA for high-dimensional data MFA training involves the inversion of large matrices for high latent-space dimensionalities. since the log-likelihood 5 contains a determinant. However, in 5, covariance matrices must be inverted even when no training is performed. Formulating MFA in terms of precision matrices (i.e., inverse to covariance matrices) will remove this requirement.

B. Related Work

The MFA model was introduced in [5], together with the basic idea of using the Woodbury matrix inversion theorem and the matrix determinant lemma to avoid inverting large matrices during MFA training. Application of MFA to high-dimensional (image) data was proposed by [6], [7], using the same mathematical ideas for ensuring efficiency even for high data dimensions. In addition to treating high-dimensional image data, [1] proposes training MFA on large-scale datasets like the CelebA face database and demonstrates excellent image generation performance. In this article, stochastic gradient descent is used instead of EM due to the large amount of samples, although the model is still needs to be initialized by k-means.

Extending MFA to a deep model was first proposed by [8], where one MFA instance was trained on the latent variables extracted from data by another one. Advantages of this approach when training on high-dimensional data are described, although it is mentioned that more than two layers are rarely useful. This idea is expanded upon by [9], which introduces deeper MFA models based on a similar principle, although application to image data is not described. Both [8] and [9] exclusively use batch-type EM for training.

All described works on MFA employ the variance-based description of MFA, meaning that the quantities to be optimized are covariance matrices as opposed to precisions.

C. Goals and contributions

The goals of this article are to establish MFA training by SGD as a simple and scalable alternative to EM and sEM in streaming scenarios with potentially high-dimensional data. For doing so, we build upon previous work on the optimization of GMMs by SGD [10]. The main novel contributions are:

- Mathematical analysis to prove that optimization of MFA log-likelihood by SGD is feasible
- Proof that MFA can be formulated in terms of precision matrices
- Procedure to train MFA by SGD from random initial conditions
- Demonstration of MFA sampling and outlier detection on typical image datasets (MNIST, SVHN)

Additionally, we provide a TensorFlow implementation.¹

II. DATA

We use the following datasets:

MNIST [11] is the common benchmark for computer vision systems and classification problems. It consists of 60 000 28×28 gray scale images of handwritten digits (0-9). **FashionMNIST** [12] consists of images of clothes in 10 categories and is structured like the MNIST. It should be more challenging than the MNIST dataset.

III. MATHEMATICAL ANALYSIS

A. Efficient formulation of MFA

As shown in [1], [5], the component log-likelihoods can be expressed without having to store and invert d -dimensional matrices by exploiting the Woodbury matrix inversion lemma and the matrix determinant lemma:

$$\begin{aligned} \log \det \Sigma_k &= \log \det \mathbf{L}_k + \log \det \mathbf{D}_k \\ \Sigma_k^{-1} &= (\Lambda_k \Lambda_k^T + \mathbf{D}_k)^{-1} \\ &= \mathbf{D}_k^{-1} - \mathbf{D}_k^{-1} \Lambda_k \mathbf{L}_k^{-1} \Lambda_k^T \mathbf{D}_k^{-1} \end{aligned} \quad (7)$$

where $\mathbf{L}_k \equiv \mathbf{I} + \Lambda_k^T \mathbf{D}_k^{-1} \Lambda_k$ is an $l \times l$ matrix and thus very efficient to compute and process.

B. Proof that MFA can be performed based on precisions

Eq. (5) could be directly used for gradient descent, which is however difficult in practice due to the explicit matrix inversions. Although the computational load of the matrix inversions in Eq. (7) is reduced w.r.t. inverting full covariance matrices, is still significant and grows as $\mathcal{O}(l^3)$. Worse still, matrix inversion may be numerically problematic in SGD for initial solutions that are far from the optimum (unlike [1] which performs SGD starting from a k-means initialization). Lastly, Eq. (5) requires matrix inversions not only for training, but even when executing a trained model.

We therefore resort to a more efficient and numerically robust strategy which avoids explicit matrix inversions. To this effect, we parameterize multi-variate normal densities by precision matrices $\mathbf{P}_k \equiv \Sigma_k^{-1}$. While the parameters of the generative model Eq. (3) have a simple relation to the learned covariance matrices, the same is not generally true for precision matrices. The question is thus how the parameters Λ_k, \mathbf{D}_k of the generative model Eq. (3) can be recovered from a trained precision matrix.

¹<https://github.com/anon-scientist/mfa>

For achieving this, we parameterize precision matrices differently from Eq. (4), using diagonal $d \times d$ matrices \mathbf{E}_k and $d \times l$ matrices $\mathbf{\Gamma}_k$ as

$$\mathbf{P}_k = \mathbf{\Sigma}_k^{-1} = \mathbf{E}_k - \mathbf{\Gamma}_k \mathbf{\Gamma}_k^T. \quad (8)$$

With this definition, the log-likelihoods can be computed in a more efficient fashion than it would be possible when using covariance matrices, compare to Eq. (5):

$$\begin{aligned} \log p_k(\mathbf{x}) &= -\frac{1}{2} \left\{ d \log(2\pi) - \log \det \mathbf{P}_k + \tilde{\mathbf{x}}^T \mathbf{P}_k \tilde{\mathbf{x}} \right\} \\ &= -\frac{1}{2} \left\{ d \log(2\pi) - \log \det \mathbf{P}_k + \tilde{\mathbf{x}}^T \mathbf{E} \tilde{\mathbf{x}} - \right. \\ &\quad \left. - (\mathbf{\Gamma}_k^T \tilde{\mathbf{x}})^2 \right\} \end{aligned} \quad (9)$$

By the matrix determinant lemma, we again have

$$\log \det \mathbf{P}_k = \log \det \mathbf{M}_k + \log \det \mathbf{E}_k \quad (10)$$

where $\mathbf{M}_k \equiv \mathbf{I} - \mathbf{\Gamma}_k^T \mathbf{E}_k^{-1} \mathbf{\Gamma}_k$ is again an $l \times l$ matrix.

The minus sign in Eq. (8) ensures that, by the Woodbury matrix inversion lemma, we can express the covariance matrices $\mathbf{\Sigma}_k$ as a function of \mathbf{E}_k and $\mathbf{\Gamma}_k$, which in turns allows the determination of the generative model parameters:

$$\mathbf{\Sigma}_k = \mathbf{P}_k^{-1} = (\mathbf{E}_k - \mathbf{\Gamma}_k \mathbf{\Gamma}_k^T)^{-1} = \mathbf{E}_k^{-1} + \mathbf{E}_k^{-1} \mathbf{\Gamma}_k \mathbf{M}_k^{-1} \mathbf{\Gamma}_k^T \mathbf{E}_k^{-1} \quad (11)$$

Examining Eq. (11) and remembering that, as a consequence of the formulation of MFA as a generative model, we have $\mathbf{\Sigma}_k = \mathbf{D}_k + \mathbf{\Lambda}_k \mathbf{\Lambda}_k^T$, we can compare terms and arrive at the following identifications:

$$\begin{aligned} \mathbf{D}_k &\equiv \mathbf{E}_k^{-1} \\ \mathbf{\Lambda}_k &\equiv \mathbf{E}_k^{-1} \mathbf{\Gamma}_k \mathcal{M}_k^{-0.5} \mathbf{O}_k \\ \mathbf{O}_k \mathcal{M}_k^{-1} \mathbf{O}_k^T &= \mathbf{M}_k^{-1} \end{aligned} \quad (12)$$

The "square root" of the inverse of \mathbf{M}_k is performed by first diagonalizing \mathbf{M}_k as $\mathbf{M}_k = \mathbf{O}^T \mathcal{M}_k \mathbf{O}$, with \mathbf{O} orthogonal and \mathcal{M}_k diagonal. The inverse of the diagonal matrix \mathcal{M}_k is trivial, and we obtain $\mathbf{M}_k^{-1} = \mathbf{O}^T \mathcal{M}_k^{-1/2} \mathcal{M}_k^{-1/2} \mathbf{O}_k$. The orthogonal matrix \mathbf{O}_k can actually be omitted from $\mathbf{\Lambda}_k$ here because an orthogonal transformation, when applied to the random latent vector, will produce another random vector of unit variance and zero mean. This defines all required parameters of the generative model in terms of the learned precision matrices, which allows us to sample efficiently in the low-dimensional space.

C. Properties of \mathbf{M}_k

The following properties of the $l \times l$ -matrices \mathbf{M}_k are relevant to SGD optimization. Here, we will just present some proofs and comments on why certain properties are desirable:

Symmetry: $\mathbf{M}_k^T = \mathbf{M}_k$. This automatically follows from the definition of \mathbf{M}_k :

$$\begin{aligned} \mathbf{M}_k^T &= (\mathbf{I} - \mathbf{\Gamma}_k \mathbf{E}_k^{-1} \mathbf{\Gamma}_k^T)^T \\ &= \mathbf{I} - (\mathbf{\Gamma}_k^T)^T (\mathbf{E}_k^{-1})^T \mathbf{\Gamma}_k^T \\ &= \mathbf{I} - \mathbf{\Gamma}_k \mathbf{E}_k^{-1} \mathbf{\Gamma}_k^T = \mathbf{M}_k \end{aligned} \quad (13)$$

since the inverse of the diagonal matrix \mathbf{E}_k is diagonal as well and thus symmetric. Symmetry is important since the eigenvalues of a symmetric real matrix are real, and, by positive-definiteness, must be strictly positive. As a consequence, the definiteness of the \mathbf{M}_k can be monitored via their eigenvalues. These are cheap to compute since l is usually small.

Diagonality This is not strictly required but facilitates the computation of eigenvalues and, above all, ensures linear independency of the columns of the loading matrices $\mathbf{\Gamma}_k$, see Sec. IV. This can be proven as follows: for any two columns $\mathbf{\Gamma}_{k:i}$, $\mathbf{\Gamma}_{k:j}$ of the loading matrices, we know that $\mathbf{M}_{ij} = \mathbf{I}_{ij} - \mathbf{\Gamma}_{k:i}^T \mathbf{E}_k^{-1} \mathbf{\Gamma}_{k:j}$. Diagonality of \mathbf{M} implies that $\mathbf{M}_{ij} = 0$ for $i \neq j$. If, on the other hand, $\mathbf{\Gamma}_{k:i}$ and $\mathbf{\Gamma}_{k:j}$, $i \neq j$ were linearly dependent ($\mathbf{\Gamma}_{k:i} = k \mathbf{\Gamma}_{k:j}$, $k \in \mathbb{R}$), we would have $\mathbf{M}_{ij} = \mathbf{I}_{ij} - k^2 \mathbf{\Gamma}_{k:i}^T \mathbf{E}_k^{-1} \mathbf{\Gamma}_{k:i}$ which cannot be zero due to the positive-definiteness of \mathbf{E}_k . Thus, it is shown that diagonalizing the \mathbf{M}_k leads to linear independency of the columns in the loading matrices.

Positive-definiteness The matrices $\mathbf{M}_k = \mathbf{I} - \mathbf{\Gamma}_k^T \mathbf{E}_k^{-1} \mathbf{\Gamma}_k$ are not positive-definite by construction, although this property is a requirement of the model. Since we showed in the previous paragraph that linearly independent columns of the $\mathbf{\Gamma}_k$ achieve diagonality of the \mathbf{M}_k , all we need to show is that there exist choices of particular $\mathbf{\Gamma}_k$ and \mathbf{E}_k for which one diagonal element of $\mathbf{M}_k < 0$. Such a choice is, e.g.: $(\mathbf{\Gamma}_k)_{1,:} = [2, 0, \dots]^T$ and $\mathbf{E}_k = \mathbf{I}$. In this case, we have $M_{11} = -3$ which proves our proposition. The remaining columns of $\mathbf{\Gamma}_k$ are arbitrary as long as they are independent from the first one. In order to maintain positive-definiteness, the eigenvalues of the diagonal \mathbf{M}_k must therefore be monitored and the appropriate columns in $\mathbf{\Gamma}_k$ modified in case eigenvalues drop below 0.

D. Sampling from a trained MFA model

Sampling is now a straightforward procedure which is conducted in two steps: initially, a component k^* is selected by drawing from a multinomial distribution parameterized by the component weights π_k : $k^* \sim M(\pi_1, \dots, \pi_K)$. For the selected component k^* , a realization of the latent variable $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$ is drawn and then transformed to the high-dimensional space using the generative model of Eq. (3):

$$\begin{aligned} \mathbf{x}_{k^*} &= \boldsymbol{\mu}_{k^*} + \mathbf{\Lambda}_{k^*} \mathbf{z} + \boldsymbol{\epsilon}_{k^*} \\ \boldsymbol{\epsilon}_{k^*} &\sim \mathcal{N}(0, \mathbf{D}_{k^*}). \end{aligned} \quad (14)$$

Here, we compute the loading matrices $\mathbf{\Lambda}_{k^*}$ and the diagonal covariance matrices \mathbf{D}_{k^*} appearing in the generative model from the results of precision-based training as indicated in

Eq. (12). As stated there, we omit orthogonal matrices from the definition of the loading matrices, and finally use:

$$\mathbf{\Lambda}_{k^*} \equiv \mathbf{E}_{k^*}^{-1} \mathbf{\Gamma}_{k^*} \mathcal{M}_{k^*}^{-0.5} \quad (15)$$

This two-step sampling procedure is equivalent to that of a GMM, except for the fact that the full-rank covariance matrices in GMMs are here expressed by the loading matrices of lower rank: $\mathbf{\Sigma}_{k^*} = \mathbf{\Lambda}_{k^*} \mathbf{\Lambda}_{k^*}^T + \mathbf{D}_{k^*}$. This is a general property of MFA models, irrespectively of the way they are trained.

IV. PROPOSAL FOR CONSTRAINED SGD OPTIMIZATION

When optimizing the precision-based MFA model by gradient descent, we use the procedure for training GMMs by SGD as described in [10]. In addition, several MFA-specific constraints are enforced in addition to the usual GMM constraint $\sum_k \pi_k = 1$:

Linearly independent columns of the loading matrices

$\mathbf{\Gamma}_k$ Although it is not formally required by the model, the columns of the $\mathbf{\Gamma}_k$ should be at least linearly independent. If they were not, the dependent columns would not contain additional information about the data and could thus be discarded. Indeed, it can be shown that a solution where two columns of the loading matrix are dependent or identical constitutes a local extremal point of the loss (to be avoided for SGD). For addressing the independence constraints, we first observe that, if the diagonal precision matrices are spherical, $\mathbf{E}_k = c\mathbf{I}$, the diagonalization of \mathbf{M}_k achieves orthogonality of the columns in $\mathbf{\Gamma}_k$. We therefore initialize diagonal precisions to $\mathbf{E}_k = c\mathbf{I}$ and diagonalize \mathbf{M}_k after each SGD step. This will ensure orthogonality of the columns in the loading matrix in the initial phases of SGD. As the \mathbf{E}_k evolve under the influence of SGD, loading matrix columns will no longer be orthogonal but still linearly independent by virtue of diagonalizing the \mathbf{M}_k , see Sec. III-B.

Positive-definiteness of the \mathbf{M}_k If this were not the case, the logarithm in Sec. III-B would be undefined. The \mathbf{E}_k must be positive-definite in any case, which we achieve by re-parameterizing them as the square of a diagonal matrix: $\mathbf{E}_k = \mathcal{E}_k^2$. Thus, the matrices \mathbf{P}_k are positive-definite as well. For maintaining positive-definiteness of \mathbf{M}_k , we must ensure that $\det \mathbf{M}_k > 0$, which is not guaranteed, see Sec. III-C. Since the \mathbf{M}_k are diagonalized after each SGD step (see previous paragraph), their eigenvalues can be read off from their diagonal. Maintaining positive-definiteness then amounts to preventing negative eigenvalues during SGD. Where an eigenvalue $(\mathbf{M}_k)_{ii}$ is below a threshold M_{\min} , we multiply the corresponding vector $\mathbf{\Gamma}_{k:i}$ in the loading matrix $\mathbf{\Gamma}_k$ by a factor that ensures that $(\mathbf{M}_k)_{ii} = \mathbf{\Gamma}_{k:i}^T \mathbf{E}^{-1} \mathbf{\Gamma}_{k:i} = M_{\min}$. Multiplication by a constant factor is the only simple operation that preserves the diagonality of the \mathbf{M}_k and the linear independence of the columns of the loading matrices.

V. EXPERIMENTS

We created a publicly available implementation of SGD-based MFA training based on Tensorflow 2.7 [13], in particular

its keras package. This implementation is used for all experiments described here, with acceleration provided GPUs of the type "nVidia GTX 2080 Super". Each experiment for which metrics are presented is repeated 10 times, and we always report means and standard deviations of these metrics.

Unless otherwise stated, we always use $K = 25$ components and a latent dimension of $l = 4$. Otherwise, default parameters and initialization from [10] are used. In particular, centroids are always initialized to uniform random values between -0.1 and 0.1. The diagonal precision matrices \mathcal{E}_k are clipped from above at a value of $D_{\max} = 20$ in order to avoid unbounded precision values for pixels that have no variability. Throughout all experiments, a mini-batch size of 100 is used for SGD.

In order to avoid undesirable local optima during early phases of training, it is imperative that centroids μ_k converge before the precision and loading matrices $\mathbf{E}_k, \mathbf{\Gamma}_k$. In practice, this can be achieved by giving different weights to the gradients $\nabla_{\mu_k} \mathcal{L}$, $\nabla_{\mathbf{E}_k} \mathcal{L}$ and $\nabla_{\mathbf{\Gamma}_k} \mathcal{L}$. Sensible values for these weights are $\lambda_{\Gamma} = \lambda_{\mathbf{E}} = 0.1$ and $\lambda_{\mu} = 1$, although in this case the precision matrices will converge rather slowly. A workaround to artificially accelerate training is to conduct separate training phases: in phase I, only the μ_k are adapted, whereas all quantities are adapted together with equal weights in phase II. Unless otherwise stated, we use 15 epochs for phase I and 50 epochs for phase II.

A. Training procedure and basic feasibility

In order to demonstrate the feasibility of MFA training by constrained SGD, we train the MFA model as in Sec. V on MNIST and FashionMNIST. To avoid cluttered and over-complex results, we restrict training to classes 0,1 and 2 (although it works just fine with all classes). For both datasets, we report the final centroids, precisions and loading matrices in Fig. 1. As stated in [10], the centroids are initialized to random values between -0.1 and 0.1, variances are uniformly initialized to $D_{\max} = 20$ and loading matrices are initialized such that $\mathbf{M}_k = 0.0001\mathbf{I}$. The component weights are uniformly initialized to $\pi_k = \frac{1}{4}$.

When repeating this experiment 10 times, we observe that convergence is always achieved, although of course the precise converged values vary due to random initial conditions. When inspecting the centroids, the SOM-like self-organization by similarity is apparent, which is an artifact of the training process, see Fig. 1.

It is visually apparent from Fig. 1 that the columns of the loading matrices are all distinct (due to diagonalizing \mathbf{M}_k), and capture major directions of variations. We also note that the strength of variations decreases with higher l , reminiscent of principal directions in PCA.

B. MFA sampling

We use the trained models from Sec. V-A to perform sampling (see Sec. III-D) from mixture components 0,6 and 15 (MNIST) and 1,3,13 (FashionMNIST). As described in Sec. III-D, the mixture component to sample is chosen randomly according to the component weights, but here we

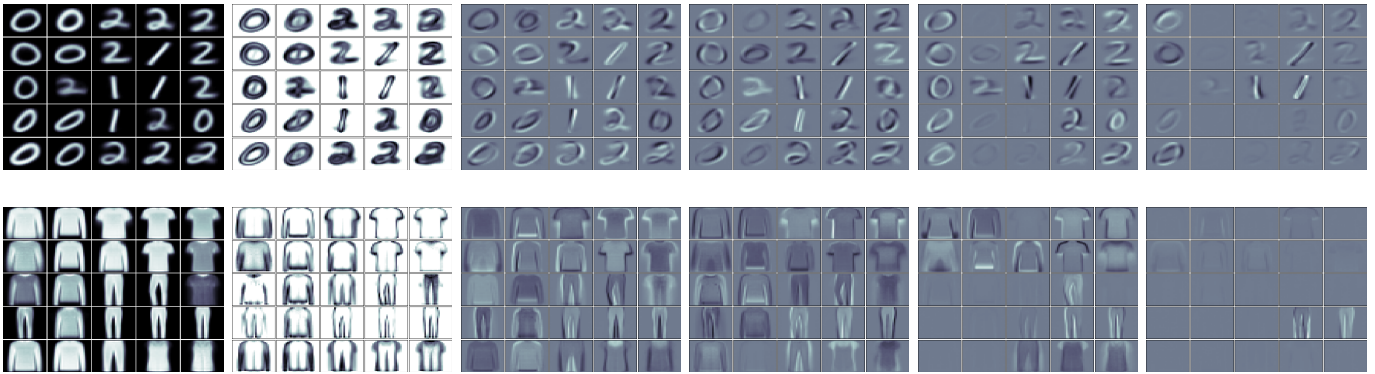


Fig. 1. Centroids μ_k (left), precisions E_k (second from left) and loading matrices Γ_k for $l = 0, 1, 2, 3$ (four right-most images) for precision-based MFA performed on MNIST (upper row) and FashionMNIST(lower row). Each tile in the loading matrix images belongs to the centroid at the corresponding tile position. Centroids are scaled in the $[0, 1]$ -range, loading matrices in the $[-1, 1]$ -range and precisions between 18 and 20.




Fig. 2. Sampling from MNIST(upper row) and FashionMNIST(lower row) from selected mixture components. Selected components (to be compared to centroids in Fig. 1) are 0,6,15 for MNIST and 1,3,13 for FashionMNIST. Please enlarge the figure in order to observe that each sample is distinct in shape from the others. This is most notable for the rightmost MNIST samples of digit class 2. But also for other classes, differences in slant and strokes are observable.

manually select components to sample from, in order to illustrate how MFA introduces variability into sampling from the same component. Sampling results are presented in Fig. 2. We observe that, despite the fact of originating from the same centroid, all samples are subtly different due to the multi-dimensional latent space that captures variations along the directions stored in the factor loadings.

C. Outlier detection and comparison to GMMs

For demonstrating the outlier detection capacity and thus the validity of precision-based MFA, we train it on MNIST classes 0-8 and record the test test log-likelihoods on the remaining class 9. For this experiment, we use $K = 49$ components but leave the experimental procedure and parameters of Sec. V

dataset	GMM	MFA
MNIST	90.1	92.2
FashionMNIST	83.7	85.3

TABLE I
AUC MEASURES COMPUTED FOR OUTLIER DETECTION BY PRECISION-BASED MFA AND VANILLA GMM. DATASETS ARE MNIST AND FASHION MNIST, WHERE CLASSES 0-8 WERE USED FOR TRAINING.

untouched otherwise. Since class 9 has not been used for training, it constitutes an outlier class and should be recognized as such. This is realized by a threshold θ applied to each log-likelihood $\mathcal{L}(x_n)$, where $\mathcal{L}(x_n) < \theta$ indicates an outlier since log-likelihoods are maximized.

For a given value of θ , we can compose two values: the

percentage α of true inliers (class 0-8 samples) that are recognized, and the percentage β of true outliers (class 9 samples) that are rejected. By varying θ , we obtain the ROC-like plots, for which we compute the area-under-the-curve measure (AUC). AUCs for MNIST and FashionMNIST, both for a GMM (same number of components) and precision-based MFA, are given in Tab. I. We observe that MFA slightly improves upon GMM performance.

VI. DISCUSSION

The mathematical and experimental results presented in Sec. III and Sec. V indicate that, first of all, precision-based MFA can be successfully trained by SGD, from random initial conditions, and on large-scale datasets. We showed that MFA models trained by SGD can be successfully used for sampling and outlier detection, two important functionalities of unsupervised learning. Here, we will discuss the wider implications of these results:

Efficiency for streaming data In conventional machine learning settings where a model is first trained, then deployed/applied without further adaptation, it is a feasible strategy to pre-compute the inverse covariance matrix after training has finished, and thus to avoid matrix inversions. In situations where the model is updated continuously while being applied (evaluate-then-train strategy, see [14]), this is no longer feasible since the covariance matrix will change continuously. Here, the precision-based approach will be superior since it requires matrix inversion for the training step only.

Simplicity The proposed SGD approach to MFA has the advantage of being extremely simple. In particular, no complex initialization of centroids by, e.g., k-means, is required, and neither is the even more complex initialization of covariance matrices as, e.g., used in [1]. Instead, centroids are initialized to random small values with guaranteed convergence, as shown here and in [10].

Processing of large and high-dimensional datasets Due to the "trick" of computing determinants and matrix inverses on the low-dimensional matrices M_k , L_k , MFA can be applied to high-dimensional data, at least as long as l is small. This has been shown in previous works [5], [6] using EM for optimization. However, MFA training on large datasets or streaming data has been problematic due to the batch-type nature of EM, which causes memory requirements to grow linearly with the number of samples. Stochastic variants of EM only partially fix this problem since they introduce several new and unintuitive hyper-parameters that must be tuned by grid search. Here, SGD offers a principled alternative, since its memory requirements only depend on the chosen mini-batch size, and since the choice of the single learning-rate parameter is well-understood.

Independent factor loadings In contrast to the original formulation of MFA [5] which imposes no constraints on the loading matrices, we demand that factor loadings be independent. We do not observe any problems related to this additional constraint, as convergence was universal in

all conducted experiments. Conversely, we did observe a few cases where SGD was stuck in local extremal points with partially identical factor loadings.

Small- l regime MFA is strongly related to principal components analysis (PCA), since the factor loadings aim to capture, for each mixture component, the directions that best explain that component's variance. In contrast to PCA, we do not require individual directions to be orthogonal, only independent. As well-known fact is that the number of directions required to explain a large part of the variance is rather small, especially for images. Thus, running MFA in a "small l regime" seems feasible.

VII. CONCLUSION AND OUTLOOK

This article has a mathematical as well as practical contribution, showing how precision-based MFA can be trained by SGD in theory, and then validating the mathematical proofs by experiments using an own keras-based implementation. Future work will include a convolutional generalization of MFA, and the stacking into deep MFA hierarchies for realistic sampling of complex images.

REFERENCES

- [1] Eitan Richardson and Yair Weiss. On GANs and GMMs. *Advances in Neural Information Processing Systems*, 2018-December(NeurIPS):5847–5858, 2018.
- [2] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- [3] Olivier Cappé and Eric Moulines. On-line expectation-maximization algorithm for latent data models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3):593–613, 2009.
- [4] Herbert Robbins and Sutton Monro. A stochastic approximation method. *Annals of Math.Stat.*, 22, 1951.
- [5] Zoubin Ghahramani and Geoffrey E. Hinton. The EM Algorithm for Mixtures of Factor Analyzers. *Compute*, pages 1–8, 1997.
- [6] G. J. McLachlan, D. Peel, and R. W. Bean. Modelling high - dimensional data by mixtures of factor analyzers. *Computational Statistics and Data Analysis*, 41(3-4):379–388, 2003.
- [7] Geoffrey McLachlan and David Peel. Mixtures of Factor Analyzers. pages 238–256, jan 2005.
- [8] Yichuan Tang, Ruslan Salakhutdinov, and Geoffrey Hinton. Deep mixtures of factor analysers. *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, 1:505–512, 2012.
- [9] Cinzia Viroli and Geoffrey J. McLachlan. Deep Gaussian mixture models. *Statistics and Computing*, 29(1):43–51, 2019.
- [10] A Gepperth and B Pfülb. Gradient-based training of gaussian mixture models for high-dimensional streaming data. *Neural Processing Letters*, 2021. accepted.
- [11] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323, 1998.
- [12] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. pages 1–6, 2017.
- [13] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- [14] A Gepperth and B Hammer. Incremental learning algorithms and applications. 2016.