



HAL
open science

Fast Incremental Expectation Maximization for finite-sum optimization: nonasymptotic convergence

Gersende Fort, Pierre Gach, Eric Moulines

► **To cite this version:**

Gersende Fort, Pierre Gach, Eric Moulines. Fast Incremental Expectation Maximization for finite-sum optimization: nonasymptotic convergence. *Statistics and Computing*, In press, 31 (48), 10.1007/s11222-021-10023-9 . hal-02617725v2

HAL Id: hal-02617725

<https://hal.science/hal-02617725v2>

Submitted on 21 Dec 2020

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.

Fast Incremental Expectation Maximization for finite-sum optimization: nonasymptotic convergence. *

G. Fort¹, P. Gach², and E. Moulines³

¹Institut de Mathématiques de Toulouse & CNRS, France;

gersende.fort@math.univ-toulouse.fr

²Institut de Mathématiques de Toulouse & Université Toulouse 3, France;

pierre.gach@math.univ-toulouse.fr

³CMAP & Ecole Polytechnique, France; eric.moulines@polytechnique.edu

December 21, 2020

Abstract

Fast Incremental Expectation Maximization (FIEM) is a version of the EM framework for large datasets. In this paper, we first recast FIEM and other incremental EM type algorithms in the *Stochastic Approximation within EM* framework. Then, we provide nonasymptotic bounds for the convergence in expectation as a function of the number of examples n and of the maximal number of iterations K_{\max} . We propose two strategies for achieving an ϵ -approximate stationary point, respectively with $K_{\max} = O(n^{2/3}/\epsilon)$ and $K_{\max} = O(\sqrt{n}/\epsilon^{3/2})$, both strategies relying on a random termination rule before K_{\max} and on a constant step size in the Stochastic Approximation step. Our bounds provide some improvements on the literature. First, they allow K_{\max} to scale as \sqrt{n} which is better than $n^{2/3}$ which was the best rate obtained so far; it is at the cost of a larger dependence upon the tolerance ϵ , thus making this control relevant for small to medium accuracy with respect to the number of examples n . Second, for the $n^{2/3}$ -rate, the numerical illustrations show that thanks to an optimized choice of the step size and of the bounds in terms of quantities characterizing the optimization problem at hand, our results design a less conservative choice of the step size and provide a better control of the convergence in expectation.

*This work is partially supported by the *Fondation Simone et Cino Del Duca* through the project OpSiMorE; by the French *Agence Nationale de la Recherche* (ANR), project under reference ANR-PRC-CE23 MASDOL and Chair ANR of research and teaching in artificial intelligence - SCAI Statistics and Computation for AI; and by the Russian Academic Excellence Project '5-100'.

keywords: Computational Statistical Learning Large Scale Learning Incremental Expectation Maximization algorithm Momentum Stochastic Approximation Finite-sum optimization.

Mathematics Subject Classification (2010) MSC: 65C60 68Q32 65K10

1 Introduction

The Expectation Maximization (EM) algorithm was introduced by Dempster et al. (1977) to solve a non-convex optimization problem on $\Theta \subseteq \mathbb{R}^d$ when the objective function F is defined through an integral:

$$F(\theta) \stackrel{\text{def}}{=} -\frac{1}{n} \log \int_{\mathbf{Z}_n} G(z; \theta) \mu_n(dz), \quad (1)$$

for $n \in \mathbb{N} \setminus \{0\}$, a positive function G and a σ -finite positive measure μ_n on a measurable space $(\mathbf{Z}_n, \mathcal{Z}_n)$. EM is a Majorize-Minimization (MM) algorithm which, based on the current value of the iterate θ_{curr} , defines a majorizing function $\theta \mapsto \mathbf{Q}(\theta, \theta_{\text{curr}})$ given, up to an additive constant, by

$$\mathbf{Q}(\theta, \theta_{\text{curr}}) \stackrel{\text{def}}{=} -\frac{1}{n} \int_{\mathbf{Z}_n} \log G(z; \theta) \quad G(z; \theta_{\text{curr}}) \exp(n F(\theta_{\text{curr}})) \mu_n(dz) \quad .$$

The next iterate is chosen to be the/a minimum of $\mathbf{Q}(\cdot, \theta_{\text{curr}})$. Each iteration of EM is divided into two steps. In the E step (expectation step), a surrogate function is computed. In the M step (minimization step), the surrogate function is minimized. The computation of the Q function is straightforward when there exist functions $\Phi : \Theta \rightarrow \mathbb{R}^q$ and $S : \mathbf{Z}_n \rightarrow \mathbb{R}^q$ such that $n^{-1} \log G(z; \theta) = \langle S(z), \Phi(\theta) \rangle$; this yields $\mathbf{Q}(\theta, \theta_{\text{curr}}) = -\langle \bar{s}(\theta_{\text{curr}}), \Phi(\theta) \rangle$ where $\bar{s}(\theta_{\text{curr}})$ denotes the expectation of the function S with respect to (w.r.t.) the probability measure

$$\pi_{\theta_{\text{curr}}}(dz) \stackrel{\text{def}}{=} G(z; \theta_{\text{curr}}) \exp(n F(\theta_{\text{curr}})) \mu_n(dz) .$$

This paper is concerned with the case \mathbf{Z}_n is the n -fold Cartesian product of the set \mathbf{Z} (denoted by \mathbf{Z}^n), $z = (z_1, \dots, z_n) \in \mathbf{Z}^n$, $S(z) = n^{-1} \sum_{i=1}^n s_i(z_i)$, μ_n is the tensor product of the σ -finite positive measure μ on the measurable space $(\mathbf{Z}, \mathcal{Z})$. This

implies that

$$\begin{aligned}
 F(\theta) &= \frac{1}{n} \sum_{i=1}^n f_i(\theta) , \\
 f_i(\theta) &\stackrel{\text{def}}{=} -\log \int_{\mathbf{Z}} \exp(\langle s_i(z), \Phi(\theta) \rangle) \mu(\mathrm{d}z) ; \\
 \bar{s}(\theta_{\text{curr}}) &\propto \frac{1}{n} \sum_{i=1}^n \int_{\mathbf{Z}} s_i(z) \frac{\exp(\langle s_i(z), \Phi(\theta_{\text{curr}}) \rangle) \mu(\mathrm{d}z)}{\int_{\mathbf{Z}} \exp(\langle s_i(z'), \Phi(\theta_{\text{curr}}) \rangle) \mu(\mathrm{d}z')} .
 \end{aligned}$$

This finite-sum framework is motivated by large scale learning problems. In such case, n is the number of observations, assumed to be independent; the function f_i stands for a possibly non-convex loss associated to the observation $\#i$ and can also include a penalty (or a regularization) term. In the statistical context, F is the negative normalized log-likelihood of the n observations in a latent variable model, and G is the complete likelihood; when $\log G(z; \theta) \propto \langle S(z), \Phi(\theta) \rangle$, it belongs to the curved exponential family (see e.g. Brown (1986) and Sundberg (2019)).

When n is large, the computation of $\bar{s}(\theta_{\text{curr}})$ is computationally costly and should be avoided. We consider incremental algorithms which use, at each iteration, a mini-batch of examples. The computational complexity of these procedures typically displays a trade-off between the loss of information incurred by the use of a subset of the observations, and a faster progress toward the solutions since the parameters can be updated more often.

A pioneering work in this direction is the *incremental EM* by Neal and Hinton (1998): the data set is divided into B blocks and a single block is visited between each parameter update. The \mathbf{Q} function of *incremental EM* is again a sum over n terms, but each E step consists in updating only a block of terms in this sum (see Ng and McLachlan (2003)).

The *Online EM* algorithm by Cappé and Moulines (2009) was originally designed to process data streams. It replaces the computation of $\bar{s}(\theta_{\text{curr}})$ by an iteration of a Stochastic Approximation (SA) algorithm (see Robbins and Monro (1951)). *Online EM* in the finite sum setting is closely related to *Stochastic Gradient Descent*. Improved versions were considered by Chen et al. (2018) and by Karimi et al. (2019b) which introduced respectively *Stochastic EM with Variance Reduction (sEM-vr)* and *Fast Incremental EM (FIEM)* as variance reduction techniques within *Online EM* as an echo to *Stochastic Variance Reduced Gradient (SVRG, Johnson and Zhang (2013))* and *Stochastic Averaged Gradient (SAGA, Defazio et al. (2014))* introduced as variance reduction techniques within *Stochastic Gradient Descent*.

In this paper, we aim to study such incremental EM methods combined with a SA approach. The first goal of this paper is to cast *Online EM, incremental EM*

and *FIEM* into a framework called hereafter *Stochastic Approximation within EM* approaches; see subsection 2.2. We show that the E step of *FIEM* can be seen as the combination of an SA update and of a control variate; we propose to optimize the trade-off between update and variance reduction, which yields to the *opt-FIEM* algorithm (see also section 4 for a numerical exploration).

The second and main objective of this paper, is to derive nonasymptotic upper bounds for the convergence in expectation of *FIEM* (see section 3).

Following Ghadimi and Lan (2013) (see also Allen-Zhu and Hazan (2016), Reddi et al. (2016), Fang et al. (2018), Zhou et al. (2018) and Karimi et al. (2019b)), we propose to fix a maximal length K_{\max} and terminate a path $\{\theta^k, k \geq 0\}$ of the algorithm at some random time K uniformly sampled from $\{0, \dots, K_{\max} - 1\}$ prior the run and independently of it; our bounds control the expectation $\mathbb{E} [\|\nabla F(\theta^K)\|^2]$ and as a corollary, we discuss how to fix K_{\max} as a function of the sample size n in order to reach an ϵ -approximate stationary point i.e. to find $\hat{\theta}^{K,\epsilon}$ such that $\mathbb{E} [\|\nabla F(\hat{\theta}^{K,\epsilon})\|^2] \leq \epsilon$. Such a property is sometimes called ϵ -accuracy in expectation (see e.g. (Reddi et al., 2016, Definition 1)).

Karimi et al. (2019b) established that *incremental EM*, which picks at random one example per iteration, reaches ϵ -accuracy by choosing $K_{\max} = O(n\epsilon^{-1})$: even if the algorithm is terminated at a random time K , this random time is chosen as a function of K_{\max} which has to increase linearly with the size n of the data set. They also prove that for *FIEM*, ϵ -approximate stationarity is reached with $K_{\max} = O(n^{2/3}\epsilon^{-1})$ - here again, with one example picked at random per iteration. For these reasons, *FIEM* is preferable especially when n is large (see section 5 for a numerical illustration). Our major contribution in this paper is to show that for *FIEM*, the rate depends on the choice of some design parameters. By choosing a constant step size sequence in the SA step, depending upon n as $O(n^{-2/3})$, then ϵ -accuracy requires $K_{\max} = O(n^{2/3}\epsilon^{-1})$; we provide a choice of the step size (with an explicit dependence on the constants of the problem) and an explicit expression of the upper bound, which improve the results reported in Karimi et al. (2019b) (see subsection 3.2; see also section 4 for illustration). We then prove in subsection 3.3 that ϵ -accuracy can be achieved with $K_{\max} = O(\sqrt{n}\epsilon^{-3/2})$ iterations using another strategy for the definition of the step size. Finally, we go beyond the uniform distribution for the random termination time K by considering a large class of distributions on the set $\{0, \dots, K_{\max} - 1\}$ (see subsection 3.4).

Notations. $\langle a, b \rangle$ denotes the standard Euclidean scalar product on \mathbb{R}^ℓ , for $\ell \geq 1$; and $\|a\|$ the associated norm. For a matrix A , A^T is its transpose. By convention, vectors are column vectors.

For a smooth function ϕ , $\dot{\phi}$ denotes its gradient; for a smooth real-valued function

of several variables \mathbf{L} , $\partial_\tau^k \mathbf{L}$ stands for the partial derivative of order k with respect to the variable τ .

For a non negative integer n , $[n] \stackrel{\text{def}}{=} \{0, \dots, n\}$ and $[n]^\star \stackrel{\text{def}}{=} \{1, \dots, n\}$. $a \wedge b$ is the minimum of two real numbers a, b . The big O notation is used to leave out constants. For a random variable U , $\sigma(U)$ denotes the sigma algebra generated by U .

2 Incremental EM algorithms for finite-sum optimization

2.1 EM in the expectation space

This paper deals with EM-based algorithms to solve

$$\text{Argmin}_{\theta \in \Theta} F(\theta), \quad F(\theta) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\theta) + \mathbf{R}(\theta), \quad (2)$$

where

$$\mathcal{L}_i(\theta) \stackrel{\text{def}}{=} -\log \int_{\mathbf{Z}} \tilde{h}_i(z) \exp(\langle s_i(z), \phi(\theta) \rangle) \mu(\mathrm{d}z), \quad (3)$$

under the following assumption:

A1. $\Theta \subseteq \mathbb{R}^d$ is a measurable convex subset. $(\mathbf{Z}, \mathcal{Z})$ is a measurable space and μ is a σ -finite positive measure on $(\mathbf{Z}, \mathcal{Z})$. The functions $\mathbf{R} : \Theta \rightarrow \mathbb{R}$, $\phi : \Theta \rightarrow \mathbb{R}^q$ and $\tilde{h}_i : \mathbf{Z} \rightarrow \mathbb{R}_+$, $s_i : \mathbf{Z} \rightarrow \mathbb{R}^q$ for $i \in [n]^\star$ are measurable functions. Finally, for any $\theta \in \Theta$ and $i \in [n]^\star$, $-\infty < \mathcal{L}_i(\theta) < \infty$.

Under A1, for any $\theta \in \Theta$ and $i \in [n]^\star$, the quantity $p_i(z; \theta) \mu(\mathrm{d}z)$ where

$$p_i(z; \theta) \stackrel{\text{def}}{=} \tilde{h}_i(z) \exp(\langle s_i(z), \phi(\theta) \rangle + \mathcal{L}_i(\theta)),$$

defines a probability measure on $(\mathbf{Z}, \mathcal{Z})$. We assume that

A2. For all $\theta \in \Theta$ and $i \in [n]^\star$, the expectation

$$\bar{s}_i(\theta) \stackrel{\text{def}}{=} \int_{\mathbf{Z}} s_i(z) p_i(z; \theta) \mu(\mathrm{d}z)$$

exists and is computationally tractable.

For any $\theta \in \Theta$, define

$$\bar{s}(\theta) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \bar{s}_i(\theta). \quad (4)$$

The framework defined by (2) and (3) covers many computational learning problems such as empirical risk minimization with non-convex losses: R may include a regularization condition on the parameter θ , \mathcal{L}_i is the loss function associated to example $\#i$ and $n^{-1} \sum_{i=1}^n \mathcal{L}_i$ is the empirical loss. This framework includes negative log-likelihood inference in latent variable model (see e.g. Little and Rubin (2002)), when the complete data likelihood is from a curved exponential family; in this framework, $z \mapsto p_i(z; \theta) \mu(dz)$ is the a posteriori distribution of the latent variable $\#i$.

Given $\theta' \in \Theta$, define the function $\bar{F}(\cdot, \theta') : \Theta \rightarrow \mathbb{R}$ by

$$\begin{aligned} \bar{F}(\theta, \theta') &\stackrel{\text{def}}{=} -\langle \bar{s}(\theta'), \phi(\theta) \rangle + R(\theta) + \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\theta'), \\ \mathcal{C}_i(\theta') &\stackrel{\text{def}}{=} \mathcal{L}_i(\theta') + \langle \bar{s}_i(\theta'), \phi(\theta') \rangle. \end{aligned}$$

It is well known (see McLachlan and Krishnan (2008); Lange (2016); see also section 7 in the supplementary material) that $\{\bar{F}(\cdot, \theta'), \theta' \in \Theta\}$ is a family of majorizing function of the objective function F from which a Majorize-Minimization approach for solving (2) can be derived. Define

$$L(s, \cdot) : \theta \mapsto -\langle s, \phi(\theta) \rangle + R(\theta) \tag{5}$$

and consider the following assumption:

A3. For any $s \in \mathbb{R}^q$, $\theta \mapsto L(s, \theta)$ has a unique global minimum on Θ denoted $T(s)$.

In most successful applications of the EM algorithm, the function $\theta \mapsto L(s, \theta)$ is strongly convex. Strong convexity is however not required here. Starting from the current point θ^k , the EM iterative scheme $\theta^{k+1} = T \circ \bar{s}(\theta^k)$ first computes a point in $\bar{s}(\Theta)$ through the expectation \bar{s} , and then apply the map T to obtain the new iterate θ^{k+1} . It can therefore be described in the $\bar{s}(\Theta)$ -space, a space sometimes called the *expectation space*: define the sequence $\{\bar{s}^k, k \in \mathbb{N}\}$ by $\bar{s}^0 \in \mathbb{R}^q$ and for any $k \geq 0$

$$\bar{s}^{k+1} \stackrel{\text{def}}{=} \bar{s} \circ T(\bar{s}^k). \tag{6}$$

Sufficient conditions for the characterization of the limit points of any instance $\{\bar{s}^k, k \geq 0\}$ as the critical points of $F \circ T$, for the convergence of the functional along the sequence $\{F \circ T(\bar{s}^k), k \geq 0\}$, or for the convergence of the iterates $\{\bar{s}^k, k \geq 0\}$ exist in the literature (see e.g. Wu (1983); Lange (1995); Delyon et al. (1999) in the EM context and Zangwill (1967); Csiszár and Tusnády (1984); Gunawardana and Byrne (2005); Parizi et al. (2019) for general iterative MM algorithms). Proposition 1 characterizes the fixed points of $T \circ \bar{s}$ and of $\bar{s} \circ T$ under a set of conditions which will be adopted for the convergence analysis in Section 3.

A4. (i) The functions ϕ and \mathbb{R} are continuously differentiable on Θ^v where $\Theta^v \stackrel{\text{def}}{=} \Theta$ if Θ is open, or Θ^v is a neighborhood of Θ otherwise. \mathbb{T} is continuously differentiable on \mathbb{R}^q .

(ii) The function F is continuously differentiable on Θ^v and for any $\theta \in \Theta$, we have

$$\dot{F}(\theta) = - \left(\dot{\phi}(\theta) \right)^T \bar{s}(\theta) + \dot{\mathbb{R}}(\theta) .$$

(iii) For any $s \in \mathbb{R}^q$, $B(s) \stackrel{\text{def}}{=}} (\dot{\phi} \circ \mathbb{T})(s)$ is a symmetric $q \times q$ matrix with positive minimal eigenvalue.

Under A1 to A4-(i) and the assumption that Θ and $\phi(\Theta)$ are open subsets of resp. \mathbb{R}^d and \mathbb{R}^q , then 7 shows that A4-(ii) holds and the functions \mathcal{L}_i are continuously differentiable on Θ for all $i \in [n]^*$.

Under A1, A3 and the assumptions that (i) \mathbb{T} is continuously differentiable on \mathbb{R}^q and (ii) for any $s \in \mathbb{R}^q$, $\tau \mapsto \mathbb{L}(s, \tau)$ (see (5)) is twice continuously differentiable on Θ^v (defined in A4-(i)), then for any $s \in \mathbb{R}^q$, $\partial_\tau^2 \mathbb{L}(s, \mathbb{T}(s))$ is positive-definite and

$$B(s) = \left(\dot{\mathbb{T}}(s) \right)^T \partial_\tau^2 \mathbb{L}(s, \mathbb{T}(s)) \left(\dot{\mathbb{T}}(s) \right) ;$$

see (Delyon et al., 1999, Lemma 2). Therefore, $B(s)$ is a symmetric matrix and if $\text{rank}(\dot{\mathbb{T}}(s)) = q = q \wedge d$, its minimal eigenvalue is positive.

Proposition 1. Assume A1, A2 and A3. Define the measurable functions $V : \mathbb{R}^q \rightarrow \mathbb{R}$ and $h : \mathbb{R}^q \rightarrow \mathbb{R}^q$ by

$$V(s) \stackrel{\text{def}}{=} F \circ \mathbb{T}(s) , \quad h(s) \stackrel{\text{def}}{=} \bar{s} \circ \mathbb{T}(s) - s .$$

1. If s^* is a fixed point of $\bar{s} \circ \mathbb{T}$, then $\mathbb{T}(s^*)$ is a fixed point of $\mathbb{T} \circ \bar{s}$. Conversely, if θ^* is a fixed point of $\mathbb{T} \circ \bar{s}$ then $\bar{s}(\theta^*)$ is a fixed point of $\bar{s} \circ \mathbb{T}$.
2. Assume also A4. For all $s \in \mathbb{R}^q$, we have $\dot{V}(s) = -B(s) h(s)$, and the zeros of h are the critical points of V .

The proof is in subsection 6.1.1. As a conclusion, the EM algorithm summarized in Algorithm 10, is designed to converge to the zeros of

$$s \mapsto h(s) \stackrel{\text{def}}{=} \bar{s} \circ \mathbb{T}(s) - s , \tag{7}$$

which, for some models, are the critical points of $F \circ \mathbb{T}$.

Data: $K_{\max} \in \mathbb{N}, \bar{s}^0 \in \mathbb{R}^q$
Result: The EM sequence: $\bar{s}^k, k \in [K_{\max}]$
1 for $k = 0, \dots, K_{\max} - 1$ **do**
2 $\bar{s}^{k+1} = \bar{s} \circ \mathsf{T}(\bar{s}^k)$

Algorithm 1: EM in the expectation space

2.2 Stochastic Approximation within EM

In the finite-sum framework, the number of expectation evaluations \bar{s}_i per iteration of EM is the number n of examples (see Line 2 of Algorithm 10 and (4)). It is therefore very costly in the large scale learning framework. We review in this section few alternatives of EM which all substitute the EM update $\bar{s}^{k+1} = \bar{s} \circ \mathsf{T}(\bar{s}^k)$ (see Line 2 in Algorithm 10) with an update $\widehat{S}^k \rightarrow \widehat{S}^{k+1}$ of the form

$$\widehat{S}^{k+1} = \widehat{S}^k + \gamma_{k+1} s^{k+1}, \tag{8}$$

where $\{\gamma_k, k \geq 1\}$ is a deterministic positive sequence of *step sizes* (also called *learning rates*) chosen by the user and s^{k+1} is an approximation of $h(\widehat{S}^k) = \bar{s} \circ \mathsf{T}(\widehat{S}^k) - \widehat{S}^k$. When it is a random approximation, the iterative algorithm described by (8) is a SA algorithm designed to target the zeros of the mean field $s \mapsto h(s)$ (see (7)); see e.g. Benveniste et al. (1990); Borkar (2008) for a general review on SA. Many stochastic approximations of EM can be described by (8): let us cite for example the **Stochastic EM** by Celeux and Diebolt (1985), the **Monte Carlo EM** (MCEM, introduced by Wei and Tanner (1990) and studied by Fort and Moulines (2003)) which corresponds to $\gamma_{k+1} = 1$ and the **Stochastic Approximation EM** (SAEM) introduced by Delyon et al. (1999).

In the finite-sum framework, observe from (7) that for any $s \in \mathbb{R}^q$,

$$h(s) = \mathbb{E}[\bar{s}_I \circ \mathsf{T}(s) + W] - s, \tag{9}$$

where I is a uniform random variable on $[n]^*$ and W is a zero-mean random vector. Such an expression gives insights for the definition of SA schemes, including the combination with a variance reduction techniques through an adequate choice of W (see e.g. (Glasserman, 2004, Section 4.1.) for an introduction to control variates). We review below recent EM-based algorithms, designed for the finite-sum setting.

2.2.1 The Fast Incremental EM algorithm

Fast Incremental EM (FIEM) was introduced by Karimi et al. (2019b); it is given in Algorithm 2. Lines 4 to 7 are a recursive computation of $n^{-1} \sum_{i=1}^n \mathbf{S}_{k+1,i}$, stored

<p>Data: $K_{\max} \in \mathbb{N}$, $\widehat{S}^0 \in \mathbb{R}^q$, $\gamma_k \in (0, \infty)$ for $k \in [K_{\max}]^*$</p> <p>Result: The FIEM sequence: $\widehat{S}^k, k \in [K_{\max}]$</p> <ol style="list-style-type: none"> 1 $S_{0,i} = \bar{s}_i \circ \mathsf{T}(\widehat{S}^0)$ for all $i \in [n]^*$; 2 $\widetilde{S}^0 = n^{-1} \sum_{i=1}^n S_{0,i}$; 3 for $k = 0, \dots, K_{\max} - 1$ do 4 Sample I_{k+1} uniformly from $[n]^*$; 5 $S_{k+1,i} = S_{k,i}$ for $i \neq I_{k+1}$; 6 $S_{k+1,I_{k+1}} = \bar{s}_{I_{k+1}} \circ \mathsf{T}(\widehat{S}^k)$; 7 $\widetilde{S}^{k+1} = \widetilde{S}^k + n^{-1} (S_{k+1,I_{k+1}} - S_{k,I_{k+1}})$; 8 Sample J_{k+1} uniformly from $[n]^*$; 9 $\widehat{S}^{k+1} = \widehat{S}^k + \gamma_{k+1}(\bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \widehat{S}^k + \widetilde{S}^{k+1} - S_{k+1,J_{k+1}})$
--

Algorithm 2: Fast Incremental EM

in \widetilde{S}^{k+1} , where for $k \geq 0$,

$$S_{k+1,i} \stackrel{\text{def}}{=} \begin{cases} \bar{s}_{I_{k+1}} \circ \mathsf{T}(\widehat{S}^k) & \text{if } i = I_{k+1} \text{ ,} \\ S_{k,i} & \text{otherwise .} \end{cases} \quad (10)$$

This procedure avoids the computation of a sum with n terms at each iteration of FIEM, but at the price of a memory footprint since the \mathbb{R}^q -valued vectors $S_{k,i}$ for $i \in [n \wedge K_{\max}]^*$ have to be stored. Line 9 is of the form (8) with s^{k+1} equal to the sum of two terms: $\bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \widehat{S}^k$ is an oracle for $\mathbb{E}[\bar{s}_I \circ \mathsf{T}(s) - s]$ evaluated at $s = \widehat{S}^k$; and $W \stackrel{\text{def}}{=} \widetilde{S}^{k+1} - S_{k+1,J_{k+1}}$ acts as a control variate, which conditionally to the past $\mathcal{F}_{k+1/2} \stackrel{\text{def}}{=} \sigma(\widehat{S}^0, I_1, J_1, \dots, I_k, J_k, I_{k+1})$, is centered. A natural extension, which is not addressed in this paper, is to replace the draws I_{k+1}, J_{k+1} by mini-batches of examples sampled in $[n]^*$ - uniformly, with or without replacement.

The introduction of such a variable W is inherited from the **Stochastic Averaged Gradient** (SAGA, by Defazio et al. (2014)). The convergence analysis of FIEM was given in Karimi et al. (2019b): they derive nonasymptotic convergence results in expectation. The theoretical contribution of our paper, detailed in section 3, is to complement and improve these results.

On the computational side, each iteration of FIEM requires two draws from $[n]^*$, two expectation evaluations of the form $\bar{s}_i(\theta)$ and a maximization step; there is a space complexity through the storage of the auxiliary quantity $S_{k,\cdot}$ - its size being proportional to $q(2K_{\max} \wedge n)$ (in some specific situations, the size can be reduced - see the comment in (Schmidt et al., 2017, Section 4.1)). The initialization step also requires a maximization step and n expectation evaluations.

2.2.2 An optimized FIEM algorithm, opt-FIEM

From (9), line 9 of Algorithm 2 and the control variate technique, we explore here the idea to modify the original FIEM as follows (compare to line 9 in Algorithm 2)

$$\begin{aligned} \widehat{S}^{k+1} = \widehat{S}^k + \gamma_{k+1} \left(\bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \widehat{S}^k \right. \\ \left. + \lambda_{k+1} \left(\widetilde{S}^{k+1} - \mathbf{S}_{k+1, J_{k+1}} \right) \right) \end{aligned} \quad (11)$$

where $\lambda_{k+1} \in \mathbb{R}$ is chosen in order to minimize the conditional fluctuation

$$\gamma_{k+1}^{-2} \mathbb{E} \left[\|\widehat{S}^{k+1} - \widehat{S}^k\|^2 | \mathcal{F}_{k+1/2} \right] .$$

Upon noting that $\mathbb{E} \left[\widehat{S}^{k+1} - \widehat{S}^k | \mathcal{F}_{k+1/2} \right] = \gamma_{k+1} h(\widehat{S}^k)$, it is easily seen that equivalently, λ_{k+1} is chosen as the minimum of the conditional variance

$$\mathbb{E} \left[\|\gamma_{k+1}^{-1} \left(\widehat{S}^{k+1} - \widehat{S}^k \right) - h(\widehat{S}^k)\|^2 | \mathcal{F}_{k+1/2} \right] .$$

We will refer to this technique as the optimized FIEM (opt-FIEM) below; FIEM corresponds to the choice $\lambda_{k+1} = 1$ for any $k \geq 0$ and Online EM corresponds to the choice $\lambda_{k+1} = 0$ for any $k \geq 0$ (see Algorithm 3).

Upon noting that, given two random variables U, V such that $\mathbb{E}[\|V\|^2] > 0$, the function $\lambda \mapsto \mathbb{E}[\|U + \lambda V\|^2]$ reaches its minimum at a unique point given by $\lambda_\star \stackrel{\text{def}}{=} -\mathbb{E}[U^T V] / \mathbb{E}[\|V\|^2]$, the optimal choice for λ_{k+1} is given by (remember that conditionally to $\mathcal{F}_{k+1/2}$, $\widetilde{S}^{k+1} - \mathbf{S}_{k+1, J_{k+1}}$ is centered),

$$\lambda_{k+1}^\star \stackrel{\text{def}}{=} - \frac{\text{Tr Cov} \left(\bar{s}_J \circ \mathsf{T}(\widehat{S}^k), \widetilde{S}^{k+1} - \mathbf{S}_{k+1, J} | \mathcal{F}_{k+1/2} \right)}{\text{Tr Var} \left(\widetilde{S}^{k+1} - \mathbf{S}_{k+1, J} | \mathcal{F}_{k+1/2} \right)} \quad (12)$$

where J is a uniform random variable on $[n]^\star$, independent of $\mathcal{F}_{k+1/2}$, Tr denotes the trace of a matrix, and Cov, Var are resp. the covariance and variance matrices. With this optimal value, we have from (11)

$$\begin{aligned} & \gamma_{k+1}^{-2} \mathbb{E} \left[\|\widehat{S}^{k+1} - \widehat{S}^k\|^2 | \mathcal{F}_{k+1/2} \right] \\ &= \text{Tr Var} \left(\bar{s}_J \circ \mathsf{T}(\widehat{S}^k) - \widehat{S}^k | \mathcal{F}_{k+1/2} \right) \cdots \\ & \times \left(1 - \text{Corr}^2 \left(\bar{s}_J \circ \mathsf{T}(\widehat{S}^k), \widetilde{S}^{k+1} - \mathbf{S}_{k+1, J} | \mathcal{F}_{k+1/2} \right) \right) , \end{aligned} \quad (13)$$

where

$$\text{Corr}(U, V) \stackrel{\text{def}}{=} \text{TrCov}(U, V) / \{\text{TrVar}(U) \text{TrVar}(V)\}^{1/2}.$$

If the **opt-FIEM** algorithm $\{(\widehat{S}^k, \mathbf{S}_{k,\cdot}), k \geq 0\}$ were converging to $(s^*, \mathbf{S}_{\star,\cdot})$, we would have $n^{-1} \sum_{i=1}^n \mathbf{S}_{\star,i} = s^* = \bar{s} \circ \mathsf{T}(s^*)$ and $\mathbf{S}_{\star,i} = \bar{s}_i \circ \mathsf{T}(s^*)$ thus giving intuition that asymptotically when $k \rightarrow \infty$, $\lambda_k^* \approx 1$ (which implies that the correlation is 1 in (13)). The value $\lambda = 1$ is the value proposed in the original **FIEM**: therefore, asymptotically **opt-FIEM** and **FIEM** should be equivalent and **opt-FIEM** should have a better behavior in the first iterations of the algorithm. We will compare numerically **FIEM**, **opt-FIEM** and **Online EM** in section 4.

Upon noting that

$$\begin{aligned} \lambda_{k+1}^* &= - \frac{n^{-1} \sum_{j=1}^n \langle \bar{s}_j \circ \mathsf{T}(\widehat{S}^k), \widetilde{S}^{k+1} - \mathbf{S}_{k+1,j} \rangle}{n^{-1} \sum_{j=1}^n \|\widetilde{S}^{k+1} - \mathbf{S}_{k+1,j}\|^2}, \\ &= - \frac{n^{-1} \sum_{j=1}^n \langle \bar{s}_j \circ \mathsf{T}(\widehat{S}^k), \widetilde{S}^{k+1} - \mathbf{S}_{k+1,j} \rangle}{n^{-1} \sum_{j=1}^n \|\mathbf{S}_{k+1,j}\|^2 - \|\widetilde{S}^{k+1}\|^2}, \end{aligned}$$

the computational cost of λ_{k+1}^* is proportional to n : it is therefore an intractable quantity in the large scale learning setting considered in this paper. A numerical approximation has to be designed: for example, a Monte Carlo approximation of the numerator; and a recursive approximation (along the iterations k) of the denominator, mimicking the same idea as the recursive computation of the sum $\widetilde{S}^k = n^{-1} \sum_{i=1}^n \mathbf{S}_{k,i}$ in **FIEM**.

2.2.3 Online EM

Online EM is given by Algorithm 3; this description is a natural extension of the algorithm by Cappé and Moulines (2009) which was designed to process a stream of data. **Online EM** is of the form (8) with $s^{k+1} \stackrel{\text{def}}{=} \bar{s}_{I_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \widehat{S}^k$ which corresponds

Data: $K_{\max} \in \mathbb{N}$, $\widehat{S}^0 \in \mathbb{R}^q$, $\gamma_k \in (0, \infty)$ for $k \in [K_{\max}]^*$
Result: The Online EM sequence: $\widehat{S}^k, k \in [K_{\max}]$

- 1 **for** $k = 0, \dots, K_{\max} - 1$ **do**
- 2 Sample I_{k+1} uniformly from $[n]^*$;
- 3 $\widehat{S}^{k+1} = \widehat{S}^k + \gamma_{k+1} \left(\bar{s}_{I_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \widehat{S}^k \right).$

Algorithm 3: Online EM

to a natural oracle for (9) when $W = 0$. Conditionally to the past \widehat{S}^k , s^{k+1} is an unbiased approximation of $h(\widehat{S}^k)$.

Each iteration requires one draw in $[n]^*$, one expectation evaluation and one maximization step. Instead of sampling one observation per iteration, a mini-batch of examples can be used: line 3 would get into

$$\widehat{S}^{k+1} = \widehat{S}^k + \gamma_{k+1} \left(b^{-1} \sum_{i \in \mathcal{B}_{k+1}} \bar{s}_i \circ \mathbb{T}(\widehat{S}^k) - \widehat{S}^k \right)$$

where \mathcal{B}_{k+1} is a set of integers of cardinality b , sampled uniformly from $[n]^*$, with or without replacement.

Almost-sure convergence of the iterates in the long-time behavior ($K_{\max} \rightarrow \infty$) for **Online EM** was addressed in Cappé and Moulines (2009); similar convergence results in the mini-batch case for the ML estimation of exponential family mixture models were recently established by Nguyen et al. (2020). Nonasymptotic rates for the convergence in expectation are derived in Karimi et al. (2019a).

2.2.4 The incremental EM algorithm

The **Incremental EM (iEM)** algorithm is described by Algorithm 4. This description generalizes the original incremental EM proposed by Neal and Hinton (1998), which corresponds to the case $\gamma_{k+1} = 1$ and to a deterministic visit to the successive examples. As for FIEM, Lines 4 to 7 are a recursive computation of $\widetilde{S}^{k+1} =$

<p>Data: $K_{\max} \in \mathbb{N}$, $\widehat{S}^0 \in \mathbb{R}^q$, $\gamma_k \in (0, \infty)$ for $k \in [K_{\max}]^*$</p> <p>Result: The iEM sequence: $\widehat{S}^k, k \in [K_{\max}]$</p> <ol style="list-style-type: none"> 1 $S_{0,i} = \bar{s}_i \circ \mathbb{T}(\widehat{S}^0)$ for all $i \in [n]^*$; 2 $\widetilde{S}^0 = n^{-1} \sum_{i=1}^n S_{0,i}$; 3 for $k = 0, \dots, K_{\max} - 1$ do 4 Sample I_{k+1} uniformly from $[n]^*$; 5 $S_{k+1,i} = S_{k,i}$ for $i \neq I_{k+1}$; 6 $S_{k+1,I_{k+1}} = \bar{s}_{I_{k+1}} \circ \mathbb{T}(\widehat{S}^k)$; 7 $\widetilde{S}^{k+1} = \widetilde{S}^k + n^{-1} (S_{k+1,I_{k+1}} - S_{k,I_{k+1}})$; 8 $\widehat{S}^{k+1} = \widehat{S}^k + \gamma_{k+1} (\widetilde{S}^{k+1} - \widehat{S}^k)$
--

Algorithm 4: incremental EM

$n^{-1} \sum_{i=1}^n S_{k+1,i}$; and the update mechanism in Line 8 is of the form (8) with $s^{k+1} \stackrel{\text{def}}{=} \widetilde{S}^{k+1} - \widehat{S}^k$. Conditionally to the past $\sigma(\widehat{S}^0, I_1, \dots, I_k)$, s^{k+1} is a *biased* approximation of $h(\widehat{S}^k)$.

Algorithm 4 can be adapted in order to use a mini-batch of examples per iteration: the data set is divided into B blocks prior running **iEM**. Ng and McLachlan (2003) provided a numerical analysis of the role of B when **iEM** is applied to fitting a normal mixture model with fixed number of components; Gunawardana and Byrne (2005) provided sufficient conditions for the convergence in likelihood in the case the B blocks are visited according to a deterministic cycling.

Per iteration, the computational cost of **iEM** is one draw, one expectation evaluation and one maximization step. As for **FIEM**, there is a memory footprint for the storage of the \mathbb{R}^q -valued vectors $S_{k,i}$ for $i \in [n \wedge K_{\max}]^*$. The initialization requires n expectation evaluations and one maximization step.

3 Nonasymptotic bounds for convergence in expectation

The bounds are obtained by strengthening A4 with the following assumptions

A5. (i) *There exist $0 < v_{\min} \leq v_{\max} < \infty$ such that for all $s \in \mathbb{R}^q$, the spectrum of $B(s)$ is in $[v_{\min}, v_{\max}]$; $B(s)$ is defined in A4.*

(ii) *For any $i \in [n]^*$, $\bar{s}_i \circ \top$ is globally Lipschitz on \mathbb{R}^q with constant L_i .*

(iii) *The function $s \mapsto \dot{V}(s) = -B(s)h(s)$ is globally Lipschitz on \mathbb{R}^q with constant $L_{\dot{V}}$.*

3.1 A general result

Finding a point $\hat{\theta}^\epsilon$ such that $F(\hat{\theta}^\epsilon) - \min F \leq \epsilon$ is NP-hard in the non-convex setting (see Murty and Kabadi (1987)). Hence, in non-convex deterministic optimization of a smooth function F , convergence is often characterized by the quantity $\inf_{1 \leq k \leq K_{\max}} \|\nabla F(\theta^k)\|$ along a path of length K_{\max} ; in non-convex stochastic optimization, the quantity $\inf_{1 \leq k \leq K_{\max}} \mathbb{E} [\|\nabla F(\theta^k)\|^2]$ is sometimes considered when the expectation is w.r.t. the randomness introduced to replace intractable quantities with oracles. Nevertheless, in many frameworks such as the finite-sum optimization one we are interested in, such a criterion can not be used to define a termination rule for the algorithm since ∇F is intractable.

For EM-based methods in the expectation space, 1 and (7) imply that the convergence can be characterized by a "distance" of the path $\{\widehat{S}^k, k \geq 0\}$ to the set of the roots of h . We therefore introduce the following criteria: given a maximal number of

iterations K_{\max} , and a random variable K taking values in $[K_{\max} - 1]$, define

$$\begin{aligned} \mathbf{E}_0 &\stackrel{\text{def}}{=} \frac{1}{v_{\max}^2} \mathbb{E} \left[\|\dot{V}(\widehat{S}^K)\|^2 \right], \\ \mathbf{E}_1 &\stackrel{\text{def}}{=} \mathbb{E} \left[\|h(\widehat{S}^K)\|^2 \right], \\ \mathbf{E}_2 &\stackrel{\text{def}}{=} \mathbb{E} \left[\|\widetilde{S}^{K+1} - \bar{s} \circ \mathsf{T}(\widehat{S}^K)\|^2 \right], \end{aligned}$$

where K is chosen independently of the path. Upper bounds of these quantities provide a control of convergence in expectation for FIEM stopped at the random time K . Below K is the uniform r.v. on $[K_{\max} - 1]$, except in subsection 3.4.

The quantities \mathbf{E}_0 and \mathbf{E}_1 are classical in the literature: they stand for a measure of resp. a distance to a stationary point of the objective function $V = F \circ \mathsf{T}$, and a distance to the fixed points of EM. \mathbf{E}_2 is specific to FIEM: it quantifies how far the control variate \widetilde{S}^{k+1} is from the intractable mean $\bar{s} \circ \mathsf{T}(\widehat{S}^k)$ (see subsubsection 2.2.1 for the definition of \widetilde{S}^{k+1}). Under our assumptions, \mathbf{E}_0 and \mathbf{E}_1 are related as stated in 2, which is a straightforward consequence of 1.

Proposition 2. *Assume A1, A2, A3, A4 and A5-(i). For any $s \in \mathbb{R}^q$, we have $\langle h(s), \dot{V}(s) \rangle \leq -v_{\min} \|h(s)\|^2$ and $\mathbf{E}_0 \leq \mathbf{E}_1$.*

Theorem 3 is a general result for the control of quantities of the form

$$\sum_{k=0}^{K_{\max}-1} \left\{ \alpha_k \mathbb{E} \left[\|h(\widehat{S}^k)\|^2 \right] + \delta_k \mathbb{E} \left[\|\widetilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\widehat{S}^k)\|^2 \right] \right\}$$

where $\alpha_k \in \mathbb{R}$ and $\delta_k > 0$. In subsection 3.2 and subsection 3.3, we discuss how to choose the step sizes $\{\gamma_k, k \geq 1\}$ such that for any $k \in [K_{\max} - 1]$, α_k is non-negative and such that $A_{K_{\max}} \stackrel{\text{def}}{=} \sum_{k=0}^{K_{\max}-1} \alpha_k$ is positive. We then deduce from Theorem 3 an upper bound for

$$\sum_{k=0}^{K_{\max}-1} \frac{\alpha_k}{A_{K_{\max}}} \mathbb{E} \left[\|h(\widehat{S}^k)\|^2 \right] + \sum_{k=0}^{K_{\max}-1} \frac{\delta_k}{A_{K_{\max}}} \mathbb{E} \left[\|\widetilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\widehat{S}^k)\|^2 \right] \quad (14)$$

such that the larger $A_{K_{\max}}$ is, the better the bound is. (14) is then used to obtain upper bounds on \mathbf{E}_1 and \mathbf{E}_2 ; which provide in turn an upper bound on \mathbf{E}_0 by 2.

Theorem 3. *Assume A1, A2, A3, A4 and A5. Define $L^2 \stackrel{\text{def}}{=} n^{-1} \sum_{i=1}^n L_i^2$.*

Let K_{\max} be a positive integer, $\{\gamma_k, k \in \mathbb{N}\}$ be a sequence of positive step sizes and $\widehat{S}^0 \in \mathbb{R}^q$. Consider the FIEM sequence $\{\widehat{S}^k, k \in [K_{\max}]\}$ given by Algorithm 2. Set $\Delta V \stackrel{\text{def}}{=} \mathbb{E} \left[V(\widehat{S}^0) \right] - \mathbb{E} \left[V(\widehat{S}^{K_{\max}}) \right]$.

We have

$$\sum_{k=0}^{K_{\max}-1} \alpha_k \mathbb{E} \left[\|h(\widehat{S}^k)\|^2 \right] + \sum_{k=0}^{K_{\max}-1} \delta_k \mathbb{E} \left[\|\widetilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\widehat{S}^k)\|^2 \right] \leq \Delta V ,$$

with, for any $k \in [K_{\max} - 1]$,

$$\begin{aligned} \alpha_k &\stackrel{\text{def}}{=} \gamma_{k+1} v_{\min} - \gamma_{k+1}^2 (1 + \Lambda_k L^2) \frac{L_{\dot{V}}}{2} , \\ \delta_k &\stackrel{\text{def}}{=} \gamma_{k+1}^2 \left(1 + \frac{\Lambda_k \beta_{k+1} L^2}{(1 + \beta_{k+1})} \right) \frac{L_{\dot{V}}}{2} , \end{aligned}$$

where β_{k+1} is any positive number, and for $k \in [K_{\max} - 2]$,

$$\Lambda_k \stackrel{\text{def}}{=} \left(1 + \frac{1}{\beta_{k+1}} \right) \cdots \times \sum_{j=k+1}^{K_{\max}-1} \gamma_{j+1}^2 \prod_{\ell=k+2}^j \left(1 - \frac{1}{n} + \beta_{\ell} + \gamma_{\ell}^2 L^2 \right) .$$

By convention, $\Lambda_{K_{\max}-1} = 0$.

Proof. The detailed proof is in Section 6.2; let us give here a sketch of proof. Define H_{k+1} such that $\widehat{S}^{k+1} = \widehat{S}^k + \gamma_{k+1} H_{k+1}$. V is regular enough so that

$$V(\widehat{S}^{k+1}) - V(\widehat{S}^k) - \gamma_{k+1} \langle H_{k+1}, \dot{V}(\widehat{S}^k) \rangle \leq \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \|H_{k+1}\|^2 .$$

Then, the next step is to prove that

$$\begin{aligned} &\mathbb{E} \left[V(\widehat{S}^{k+1}) \right] - \mathbb{E} \left[V(\widehat{S}^k) \right] + \gamma_{k+1} \left(v_{\min} - \gamma_{k+1} \frac{L_{\dot{V}}}{2} \right) \mathbb{E} \left[\|h(\widehat{S}^k)\|^2 \right] \\ &\leq \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \mathbb{E} \left[\|H_{k+1} - \mathbb{E} [H_{k+1} | \mathcal{F}_{k+1/2}] \|^2 \right] , \end{aligned}$$

which, by summing from $k = 0$ to $k = K_{\max} - 1$, yields

$$\begin{aligned} &\sum_{k=0}^{K_{\max}-1} \gamma_{k+1} \left(v_{\min} - \gamma_{k+1} \frac{L_{\dot{V}}}{2} \right) \mathbb{E} \left[\|h(\widehat{S}^k)\|^2 \right] \\ &\leq \mathbb{E} \left[V(\widehat{S}^0) \right] - \mathbb{E} \left[V(\widehat{S}^{K_{\max}}) \right] \\ &\quad + \frac{L_{\dot{V}}}{2} \sum_{k=0}^{K_{\max}-1} \gamma_{k+1}^2 \mathbb{E} \left[\|H_{k+1} - \mathbb{E} [H_{k+1} | \mathcal{F}_{k+1/2}] \|^2 \right] . \end{aligned}$$

The most technical part is to prove that the last term on the RHS is upper bounded by

$$\frac{L\dot{V}}{2} \sum_{k=0}^{K_{\max}-1} \gamma_{k+1}^2 L^2 \left\{ \Lambda_k \mathbb{E} \left[\|h(\widehat{S}^k)\|^2 \right] - (1 + (1 + \beta_{k+1}^{-1})^{-1} \Lambda_k) \mathbb{E} \left[\|\widetilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\widehat{S}^k)\|^2 \right] \right\} .$$

This concludes the proof. \square

In the Stochastic Gradient Descent literature, complexity is evaluated in terms of *Incremental First-order Oracle* introduced by Agarwal and Bottou (2015), that is, roughly speaking, the number of calls to an oracle which returns a pair $(f_i(x), \nabla f_i(x))$. In our case, the equivalent cost is the number of expectation evaluations $\bar{s}_i(\theta)$ and the number of optimization steps $s \mapsto \mathsf{T}(s)$. K_{\max} iterations of FIEM calls $2K_{\max}$ evaluations of such expectations and K_{\max} optimization steps. As a consequence, the complexity analyses consist in discussing how K_{\max} has to be chosen as a function of n and ϵ in order to reach an ϵ -approximate stationary point defined by $E_1 \leq \epsilon$.

3.2 A uniform random stopping rule for a $n^{2/3}$ -complexity

The main result of this section establishes that by choosing a constant step size and a termination rule K sampled uniformly from $[K_{\max} - 1]$, an ϵ -approximate stationary point can be reached before

$$K_{\max} = O(n^{2/3} \epsilon^{-1} L_{\dot{V}}^{1/3} L^{2/3})$$

iterations.

For $\lambda \in (0, 1)$, $C > 0$ and n such that $n^{-1/3} < \lambda/C$, define

$$f_n(C, \lambda) \stackrel{\text{def}}{=} \left(\frac{1}{n^{2/3}} + \frac{C}{\lambda - C/n^{1/3}} \left(\frac{1}{n} + \frac{1}{1 - \lambda} \right) \right) . \quad (15)$$

Proposition 4 (application of Theorem 3). *Let $\mu \in (0, 1)$. Choose $\lambda \in (0, 1)$ and $C \in (0, +\infty)$ such that*

$$\sqrt{C} f_n(C, \lambda) = 2\mu v_{\min} \frac{L}{L_{\dot{V}}} . \quad (16)$$

Let $\{\widehat{S}^k, k \in \mathbb{N}\}$ be the FIEM sequence given by Algorithm 2 run with the constant step size

$$\gamma_\ell = \gamma_{\text{FGM}} \stackrel{\text{def}}{=} \frac{\sqrt{C}}{n^{2/3} L} = \frac{2\mu v_{\min}}{f_n(C, \lambda) n^{2/3} L_{\dot{V}}} . \quad (17)$$

For any $n > (C/\lambda)^3$ and $K_{\max} \geq 1$, we have

$$\mathbf{E}_1 + \frac{\mu}{(1-\mu)f_n(C, \lambda)n^{2/3}}\mathbf{E}_2 \leq \frac{n^{2/3}}{K_{\max}} \frac{L_{\dot{V}} f_n(C, \lambda)}{2\mu(1-\mu)v_{\min}^2} \Delta V, \quad (18)$$

where the errors \mathbf{E}_i are defined with a random variable K sampled uniformly from $[K_{\max} - 1]$.

The proof of 4 is in subsection 6.2.2. The first suggestion to solve the equation (16) is to choose $\lambda = C$ and $C \in (0, 1)$ such that

$$\sqrt{C} f_n(C, C) = 2\mu v_{\min} L / L_{\dot{V}}.$$

This equation possesses an unique solution C^* in $(0, 1)$ which is upper bounded by C^+ given by

$$C^+ \stackrel{\text{def}}{=} \frac{\sqrt{1 + 16\mu^2 v_{\min}^2 L^2 L_{\dot{V}}^{-2}} - 1}{4\mu v_{\min} L L_{\dot{V}}^{-1}}.$$

The consequence is that, given $\varepsilon \in (0, 1)$, by setting

$$M \stackrel{\text{def}}{=} \frac{L_{\dot{V}}}{2\mu(1-\mu)v_{\min}^2} f_n(C^*, C^*) \leq \frac{L_{\dot{V}}}{2\mu(1-\mu)v_{\min}^2} f_2(C^+, C^+),$$

we have

$$K_{\max} = M n^{2/3} \varepsilon^{-1} \implies \mathbf{E}_1 + \frac{L_{\dot{V}}}{2(1-\mu)L} \frac{\sqrt{C^*}}{v_{\min} n^{2/3}} \mathbf{E}_2 \leq \varepsilon \Delta V;$$

see subsection 8.1 in the supplementary material for a detailed proof of this comment.

Another suggestion is to exploit how (15) behaves when $n \rightarrow +\infty$; we prove in the supplementary material (subsection 8.1) that there exists N_* depending only upon $L, L_{\dot{V}}, v_{\min}$ such that for any $n \geq N_*$,

$$\mathbf{E}_1 + \frac{1}{3n^{2/3}} \left(\frac{L_{\dot{V}}}{L v_{\min}} \right)^{2/3} \mathbf{E}_2 \leq \frac{n^{2/3}}{K_{\max}} \frac{8}{3} \frac{L}{v_{\min}} \left(\frac{L_{\dot{V}}}{L v_{\min}} \right)^{1/3} \Delta V,$$

by choosing $C \leftarrow 0.25 (v_{\min} L / L_{\dot{V}})^{2/3}$ in the definition of the step size γ_{FGM} .

The conclusions of 4 confirm and improve previous results in the literature: (Karimi et al., 2019b, Theorem 2) proved that for FIEM applied with the constant step size

$$\gamma_K \stackrel{\text{def}}{=} \frac{v_{\min} n^{-2/3}}{\max(6, 1 + 4v_{\min}) \max(L_{\dot{V}}, L_1, \dots, L_n)}, \quad (19)$$

there holds

$$\mathbf{E}_1 \leq \frac{n^{2/3}}{K_{\max}} \Delta V \frac{(\max(6, 1 + 4v_{\min}))^2 \max(L_{\dot{V}}, L_1, \dots, L_n)}{v_{\min}^2}. \quad (20)$$

We improve this result. Firstly, we show that the RHS in (18) controls a larger quantity than \mathbf{E}_1 . Secondly, numerical explorations (see e.g. section 4) show that γ_{FGM} is larger than γ_{K} thus providing a more aggressive step size which may have a beneficial effect on the efficiency of the algorithm. Thirdly, these numerical illustrations also show that 4 provides a tighter control of the convergence in expectation. In both contributions however, the step size depends upon n as $O(n^{-2/3})$ and the bounds depend on n and K_{\max} resp. as the increasing function $n \mapsto n^{2/3}$ and the decreasing function $K_{\max} \rightarrow 1/K_{\max}$. The dependence upon n of the step size is the same as what was observed for Stochastic Gradient Descent (see e.g. Allen-Zhu and Hazan (2016)).

3.3 A uniform random stopping rule for a \sqrt{n} -complexity

Here again, we consider an FIEM path run with a constant step size and stopped at a random time K sampled uniformly from $[K_{\max} - 1]$: we prove that an ϵ -stationary point can be reached before

$$K_{\max} = O(\sqrt{n}\epsilon^{-3/2})$$

iterations. Define

$$\tilde{f}_n(C, \lambda) \stackrel{\text{def}}{=} \frac{1}{(nK_{\max})^{1/3}} + C \left(\frac{1}{n} + \frac{1}{1 - \lambda} \right). \quad (21)$$

Proposition 5 (application of Theorem 3). *Let $\mu \in (0, 1)$. Choose $\lambda \in (0, 1)$ and $C > 0$ such that*

$$\sqrt{C} \tilde{f}_n(C, \lambda) = 2\mu v_{\min} \frac{L}{L_{\dot{V}}}. \quad (22)$$

Let $\{\hat{S}^k, k \in [K_{\max}]\}$ be the FIEM sequence given by Algorithm 2 run with the constant step size

$$\gamma_\ell = \tilde{\gamma}_{\text{FGM}} \stackrel{\text{def}}{=} \frac{\sqrt{C}}{n^{1/3} K_{\max}^{1/3} L_{\dot{V}}} \frac{2\mu v_{\min}}{\tilde{f}_n(C, \lambda) n^{1/3} K_{\max}^{1/3}}. \quad (23)$$

For any positive integers n, K_{\max} such that $n^{1/3} K_{\max}^{-2/3} \leq \lambda/C$, we have

$$\mathbf{E}_1 + \frac{\mu}{(1 - \mu)\tilde{f}_n(C, \lambda)} \frac{1}{(nK_{\max})^{1/3}} \mathbf{E}_2 \leq \frac{n^{1/3}}{K_{\max}^{2/3}} \frac{L_{\dot{V}} \tilde{f}_n(C, \lambda)}{2\mu(1 - \mu)v_{\min}^2} \Delta V,$$

where the errors E_i are defined with a random variable K sampled uniformly from $[K_{\max} - 1]$.

The proof of 5 is in subsection 6.2.3. From this upper bound, it can be shown (see subsection 8.2 in the supplementary material) that for any $\tau > 0$, there exists $M > 0$ depending upon $L, L_{\dot{V}}, v_{\min}, \mu$ and τ such that for any $\varepsilon > 0$,

$$K_{\max} \geq \left(\sqrt{n}\tau^{3/2} \right) \vee \left(M\sqrt{n}\varepsilon^{-3/2} \right) \implies \frac{n^{1/3}}{K_{\max}^{2/3}} \frac{L_{\dot{V}} \tilde{f}_n(\lambda\tau, \lambda)}{2\mu(1-\mu)v_{\min}^2} \leq \varepsilon.$$

To our best knowledge, this is the first result in the literature which establishes a nonasymptotic control for FIEM at such a rate: the upper bound depends on n as the increasing function of $n \mapsto n^{1/3}$ and depends on K_{\max} as the decreasing function of $K_{\max} \mapsto K_{\max}^{-2/3}$.

As a corollary of 4 and 5, we have two upper bounds of the errors E_1, E_2 : the first one is $O(n^{2/3}K_{\max}^{-1})$ and the second one is $O(n^{1/3}K_{\max}^{-2/3})$. The first or second strategy will be chosen depending on the accuracy level ε : if $\varepsilon = n^{-e}$ for some $e > 0$, then we have to choose $K_{\max} = O(n^{2/3}\varepsilon^{-1}) = O(n^{2/3+e})$ in the first strategy and $K_{\max} = O(\sqrt{n}\varepsilon^{-3/2}) = O(n^{1/2+3e/2})$ in the second one; if $e \in (0, 1/3)$, the second approach is preferable.

When $K_{\max} = A\sqrt{n}\varepsilon^{-3/2}$, then the constant step size is $\tilde{\gamma}_{\text{FGM}} = \sqrt{C}\varepsilon(LA^{1/3}\sqrt{n})^{-1}$. In the case $\sqrt{n}\varepsilon^{-3/2} < \tilde{A}n^{2/3}\varepsilon^{-1}$, we have $\tilde{\gamma}_{\text{FGM}} > \sqrt{C}/(LA^{1/3}\tilde{A}n^{2/3})$ thus showing that the step size is lower bounded by $O(n^{-2/3})$ (see γ_{FGM} in 4). We have $\tilde{\gamma}_{\text{FGM}} \propto 1/\sqrt{n}$ when $K_{\max} \propto \sqrt{n}$: the result of 5 is obtained with a slower step size (seen as a function of n) than what was required in 4.

We now discuss a choice for the pair (λ, C) which exploits how (21) behaves when $n \rightarrow +\infty$; we prove in subsection 8.2 in the supplementary material that for any $\tau > 0$, there exists N_{\star} depending only upon $L, L_{\dot{V}}, v_{\min}, \tau$ such that for any $N_{\star} \leq n \leq \tau^3 K_{\max}^2$,

$$\begin{aligned} E_1 + \frac{2^{10/3}(1-\lambda_{\star})^{-1/3}\mu^2}{\tilde{f}_n^2(\lambda_{\star}\tau, \lambda_{\star})} \left(\frac{Lv_{\min}}{L_{\dot{V}}} \right)^{2/3} \frac{1}{(nK_{\max})^{1/3}} E_2 \\ \leq \frac{n^{1/3}}{K_{\max}^{2/3}} \frac{4}{3} \left(\frac{2L^2L_{\dot{V}}}{v_{\min}^4} \right)^{1/3} (1-\lambda_{\star})^{-1/3} \Delta V, \end{aligned}$$

where λ_{\star} is the unique solution of $(v_{\min}L)^2 \tau^3(1-\lambda_{\star})^2 = (2L_{\dot{V}})^2 \lambda_{\star}^3$.

3.4 A non-uniform random termination rule

Given a distribution $p_0, \dots, p_{K_{\max}-1}$ for the r.v. K , we show how to fix the step sizes $\gamma_1, \dots, \gamma_{K_{\max}}$ in order to deduce from Theorem 3 a control of the errors E_1 and E_2 .

For $\lambda \in (0, 1)$, $C > 0$ and $n > (C/\lambda)^3$, define the function $F_{n,C,\lambda}$

$$F_{n,C,\lambda} : x \mapsto \frac{L_{\dot{V}}}{2L^2 n^{2/3}} x \left(v_{\min} \frac{2L}{L_{\dot{V}}} - x f_n(C, \lambda) \right),$$

where f_n is defined by (15). $F_{n,C,\lambda}$ is positive, increasing and continuous on $(0, v_{\min} L / (L_{\dot{V}} f_n(C, \lambda))]$.

Proposition 6 (application of Theorem 3). *Let K be a $[K_{\max} - 1]$ -valued random variable with positive weights $p_0, \dots, p_{K_{\max}-1}$. Choose $\lambda \in (0, 1)$ and $C > 0$ such that*

$$\sqrt{C} f_n(C, \lambda) = v_{\min} \frac{L}{L_{\dot{V}}}. \quad (24)$$

For any $n > (C/\lambda)^3$ and $K_{\max} \geq 1$, we have

$$\begin{aligned} \mathbf{E}_1 + \frac{L_{\dot{V}}^2}{v_{\min}^2} n^{2/3} \max_k p_k f_n(C, \lambda) \sum_{k=0}^{K_{\max}-1} \gamma_{k+1}^2 \mathbb{E} \left[\|\tilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\hat{S}^k)\|^2 \right] \\ \leq n^{2/3} \max_k p_k \frac{2L_{\dot{V}} f_n(C, \lambda)}{v_{\min}^2} \Delta V, \end{aligned}$$

where the FIEM sequence $\{\hat{S}^k, k \in [K_{\max}]\}$ is obtained with

$$\gamma_{k+1} = \frac{1}{n^{2/3} L} F_{n,C,\lambda}^{-1} \left(\frac{p_k}{\max_{\ell} p_{\ell}} \frac{v_{\min}^2}{2L_{\dot{V}} f_n(C, \lambda)} \frac{1}{n^{2/3}} \right).$$

The proof of 6 is in subsection 6.2.4. As already commented in subsection 3.2, if we choose $\lambda = C$, then (24) gets into

$$\sqrt{C} \left(\frac{1}{n^{2/3}} + \frac{1}{1 - n^{-1/3}} \left(\frac{1}{n} + \frac{1}{1 - C} \right) \right) = \frac{v_{\min} L}{L_{\dot{V}}}.$$

There exists a unique solution C^* , which is upper bounded by a quantity which only depends upon the quantities $L, L_{\dot{V}}, v_{\min}$; hence, so $f_n(C^*, C^*)$ is and the control of \mathbf{E}_i given in 6 depends on n at most as $n \mapsto n^{2/3}$ and on K_{\max} as $K_{\max} \mapsto \max_k p_k$.

If we choose $\lambda = 1/2$, the constant C satisfies $C \leq (v_{\min} L / (4L_{\dot{V}}))^{2/3}$ (see subsection 8.3 in the supplementary material), and the nonasymptotic control given by 6 is available for $8n > (v_{\min} L / L_{\dot{V}})^2$.

Since $\sum_k p_k = 1$, we have $\max_k p_k \geq 1/K_{\max}$ thus showing that among the distributions $\{p_j, j \in [K_{\max} - 1]\}$, the quantity $\max_k p_k$ is minimal with the uniform distribution. In that case, the results of 6 can be compared to the results of 4: both RHS are increasing functions of n at the rate $n^{2/3}$; both are decreasing functions of

K_{\max} at the rate $1/K_{\max}$; the constants C, λ solving the equality in (16) in the case $\mu = 1/2$ are the same as the constants C, λ solving (24): as a consequence,

$$\frac{2L_{\dot{V}}f_n(C, \lambda)}{v_{\min}^2} = \frac{L_{\dot{V}}f_n(C, \lambda)}{2\mu(1-\mu)v_{\min}^2}, \quad \mu = 1/2.$$

Finally, when $k \mapsto p_k$ is constant, the step sizes given by 6 are constant as in 4; and they are equal since

$$F_{n,C,\lambda}^{-1} \left(\frac{v_{\min}^2 n^{-2/3}}{2L_{\dot{V}}f_n(C, \lambda)} \right) = \sqrt{C} = \frac{v_{\min}L}{L_{\dot{V}}f_n(C, \lambda)}.$$

Hence 6 and 4 are the same when $p_k = 1/K_{\max}$ for any k .

4 A toy example

In this section, we consider a very simple optimization problem which could be solved without requiring the incremental EM machinery ¹

$\mathcal{N}_p(\mu, \Gamma)$ denotes a \mathbb{R}^p -valued Gaussian distribution, with expectation μ and covariance matrix Γ .

4.1 Description

n \mathbb{R}^y -valued observations are modeled as the realization of n vectors $Y_i \in \mathbb{R}^y$ whose distribution is described as follows: conditionally to (Z_1, \dots, Z_n) , the r.v. are independent with distribution $Y_i \sim \mathcal{N}_y(AZ_i, I_y)$ where $A \in \mathbb{R}^{y \times p}$ is a deterministic matrix and I_y denotes the $y \times y$ identity matrix; (Z_1, \dots, Z_n) are i.i.d. under the distribution $\mathcal{N}_p(X\theta, I_p)$, where $\theta \in \Theta \stackrel{\text{def}}{=} \mathbb{R}^q$ and $X \in \mathbb{R}^{p \times q}$ is a deterministic matrix. Here, X and A are known, and θ is unknown; we want to estimate θ , as a solution of a (possibly) penalized maximum likelihood estimator, with penalty term $\rho(\theta) \stackrel{\text{def}}{=} v\|\theta\|^2/2$ for some $v \geq 0$. If $v = 0$, it is assumed that the rank of X and AX are resp. $q = q \wedge y$ and $p = p \wedge y$. In this model, the r.v. (Y_1, \dots, Y_n) are i.i.d. with distribution $\mathcal{N}_y(AX\theta; I_y + AA^T)$. The minimum of the function $\theta \mapsto F(\theta) \stackrel{\text{def}}{=} -n^{-1} \log g(Y_{1:n}; \theta) + \rho(\theta)$, where $g(Y_{1:n}; \cdot)$ denotes the likelihood of

¹The numerical applications are developed in MATLAB by the first author of the paper. The code files are publicly available from <https://github.com/gfort-lab/OpSiMorE/tree/master/FIEM>

the vector (Y_1, \dots, Y_n) , is unique and is given by

$$\begin{aligned}\theta_\star &\stackrel{\text{def}}{=} \left(v\mathbf{I}_q + X^T A^T (\mathbf{I}_y + AA^T)^{-1} AX \right)^{-1} X^T A^T (\mathbf{I}_y + AA^T)^{-1} \bar{Y}_n, \\ \bar{Y}_n &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n Y_i.\end{aligned}$$

Nevertheless, using the above description of the distribution of Y_i , this optimization problem can be cast into the general framework described in Section 2.1. The loss function (see (3)) is the normalized negative log-likelihood of the distribution of Y_i and is of the form (3) with

$$\phi(\theta) \stackrel{\text{def}}{=} \theta, \quad \mathbf{R}(\theta) \stackrel{\text{def}}{=} \frac{1}{2} \theta^T (X^T X + v\mathbf{I}_q) \theta, \quad s_i(z) \stackrel{\text{def}}{=} X^T z.$$

Under the stated assumptions on X , the function $\theta \mapsto -\langle s, \phi(\theta) \rangle + R(\theta)$ is defined on \mathbb{R}^q and for any $s \in \mathbb{R}^q$, it possesses an unique minimum given by

$$\mathbf{T}(s) \stackrel{\text{def}}{=} (v\mathbf{I}_q + X^T X)^{-1} s.$$

Define

$$\begin{aligned}\Pi_1 &\stackrel{\text{def}}{=} X^T (\mathbf{I}_p + A^T A)^{-1} A^T \in \mathbb{R}^{q \times y}, \\ \Pi_2 &\stackrel{\text{def}}{=} X^T (\mathbf{I}_p + A^T A)^{-1} X (v\mathbf{I}_q + X^T X)^{-1} \in \mathbb{R}^{q \times q}.\end{aligned}$$

The a posteriori distribution $p_i(\cdot, \theta) d\mu$ of the latent variable Z_i given the observation Y_i is a Gaussian distribution

$$\mathcal{N}_p \left((\mathbf{I}_p + A^T A)^{-1} (A^T Y_i + X\theta), (\mathbf{I}_p + A^T A)^{-1} \right),$$

so that for all $i \in \{1, \dots, n\}$,

$$\begin{aligned}\bar{s}_i(\theta) &\stackrel{\text{def}}{=} X^T (\mathbf{I}_p + A^T A)^{-1} (A^T Y_i + X\theta) \\ &= \Pi_1 Y_i + X^T (\mathbf{I}_p + A^T A)^{-1} X \theta \in \mathbb{R}^q, \\ \bar{s}_i \circ \mathbf{T}(s) &= \Pi_1 Y_i + \Pi_2 s.\end{aligned}$$

Therefore, A1, A2, A3 and A4-(i), (ii) are satisfied. Since $\phi \circ \mathbf{T}(s) = \mathbf{T}(s)$ then $B(s) = (v\mathbf{I}_q + X^T X)^{-1}$ for any $s \in \mathbb{R}^q$, and A4-(iii) and A5-(i) hold with

$$\begin{aligned}v_{\min} &\stackrel{\text{def}}{=} \frac{1}{v + \max_{\text{eig}}(X^T X)}, \\ v_{\max} &\stackrel{\text{def}}{=} \frac{1}{v + \min_{\text{eig}}(X^T X)};\end{aligned}$$

here, \max_eig and \min_eig denote resp. the maximum and the minimum of the eigenvalues. $\bar{s}_i \circ T(s) = \Pi_1 Y_i + \Pi_2 s$ thus showing that A5-(ii) holds with the same constant $L_i = L$ for all i . Finally, $s \mapsto B^T(s) (\bar{s} \circ T(s) - s)$ is globally Lipschitz with constant

$$L_{\hat{V}} \stackrel{\text{def}}{=} \max |\text{eig}((\nu I_q + X^T X)^{-1}(\Pi_2 - I_q))|;$$

here eig denotes the eigenvalues. This concludes the proof of A5-(iii).

4.2 The algorithms

Given the current value \widehat{S}^k , one iteration of **EM**, **Online EM**, **FIEM** and **opt-FIEM** are given by Algorithm 5 and Algorithm 6.

Online EM requires K_{\max} random draws from $[n]^\star$ per run of length K_{\max} iterations; **FIEM** and **opt-FIEM** require $2 \times K_{\max}$ draws. For a fair comparison of the algorithms along one run, the same seed is used for all the algorithms when sampling the examples from $[n]^\star$. Such a protocol allows to compare the strategies by "freezing" the randomness due to the random choice of the examples, and to really explain the different behaviors only by the values of the design parameters (the step size, for example) or by the updating scheme which is specific to each algorithm.

All the paths, whatever the algorithms, are started at the same value \widehat{S}^0 .

<p>Data: $\widehat{S}^k \in \mathbb{R}^q$, Π_1, Π_2 and \overline{Y}_n Result: $\widehat{S}_{\text{EM}}^{k+1}$ $\widehat{S}_{\text{EM}}^{k+1} = \Pi_1 \overline{Y}_n + \Pi_2 \widehat{S}^k$</p>

Algorithm 5: Toy example: one iteration of EM.

<p>Data: $\widehat{S}^k \in \mathbb{R}^q$, $S \in \mathbb{R}^{qn}$, $\widetilde{S} \in \mathbb{R}^q$; a step size $\gamma_{k+1} \in (0, 1]$ and a coefficient λ_{k+1}; the matrices Π_1, Π_2; the examples Y_1, \dots, Y_n Result: $\widehat{S}_{\text{FIEM}}^{k+1}$ $\widehat{S}_{\text{FIEM}}^{k+1} = \widehat{S}^k + \gamma_{k+1} \left(\Pi_1 Y_{J_{k+1}} + \Pi_2 \widehat{S}^k - \widehat{S}^k + \lambda_{k+1} \left\{ \widetilde{S} - S_{J_{k+1}} \right\} \right)$</p>

Algorithm 6: Toy example: one iteration of **Online EM** ($\lambda_{k+1} = 0$), **FIEM** ($\lambda_{k+1} = 1$) and **opt-FIEM**.

4.3 Numerical analysis

We choose $Y_i \in \mathbb{R}^{15}$, $Z_i \in \mathbb{R}^{10}$ and $\theta_{\text{true}} \in \mathbb{R}^{20}$. The entries of the matrix A (resp. X) are obtained as a stationary Gaussian auto-regressive process: the first column is sampled from $\sqrt{1 - \rho^2} \mathcal{N}_{15}(0; \mathbf{I})$ (resp. from $\sqrt{1 - \tilde{\rho}^2} \mathcal{N}_{10}(0; \mathbf{I})$) with $\rho = 0.8$ (resp. $\tilde{\rho} = 0.9$). θ_{true} is sparse with 40% of the components set to zero; and the other ones are sampled uniformly from $[-5, 5]$.

The regularization parameter ν is set to 0.1.

FIEM: the step sizes and the nonasymptotic controls. The first analysis is to compare the nonasymptotic bounds and the constant step sizes provided by 4, 5 and (Karimi et al., 2019b, Theorem 2) (see also (19) and (20)): the bounds are of the form

$$\frac{n^a}{K_{\max}^b} \mathcal{B} \Delta V ;$$

the numerical results below correspond to $\Delta V = 1$ and are obtained with a data set of size $n = 1e6$. Figure 1 shows the value of the constant C solving (16) when λ is successively set to $\{0.25, 0.5, 0.75\}$ and as a function of $\mu \in (0.01, 0.9)$. Figure 2 shows the same analysis for the constant C solving (22). Figure 3 and Figure 4 display the quantity \mathcal{B} as a function of μ and when the pair (λ, C) is fixed to $\lambda \in \{0.25, 0.5, 0.75\}$ and C solves resp. (16) and (22). The role of λ looks quite negligible; the bound \mathcal{B} seems to be optimal with $\mu \approx 0.25$. Note that the constants C and \mathcal{B} given by 5 depend on K_{\max} : the results displayed here correspond to $K_{\max} = n$ but we observed that the plots are the same with $K_{\max} = 1e2n$ and $K_{\max} = 1e3n$ (remember that $n = 1e6$).

Figure 5 displays the step sizes as a function of $\mu \in (0.01, 0.9)$, when $\lambda = 1/2$ and for different strategies of K_{\max} : $K_{\max} \in \{n, 1e2n, 1e3n\}$. Figure 6 displays the quantity $n^a K_{\max}^{-b} \mathcal{B}$. **Case 1** (resp. **Case 2**) corresponds to the definition given in 4 (resp. 5). For **Case 1** and Karimi et al., $(a, b) = (2/3, 1)$ and for **Case 2**, $(a, b) = (1/3, 2/3)$. The first conclusion is that our results improve Karimi et al. (2019b): we provide a larger step size (improved by a factor up to 55, with the strategy **Case 1**, $\mu = 0.25$, $\lambda = 0.5$) and a tighter bound (reduced by a factor up to 235, with the strategy **Case 1**, $\mu = 0.25$, $\lambda = 0.5$). The second conclusion is about the comparison of 4 and 5: as already commented (see subsection 3.3), the first strategy is preferable when the tolerance level ϵ is small (w.r.t. $n^{-1/3}$).

Comparison of Online EM, FIEM and opt-FIEM. The algorithms are run with the same constant step size given by (17) when C solves (16) with $\mu = 0.25$ and $\lambda = 0.5$. The size of the data set is $n = 1e3$ and the maximal number of iterations is $K_{\max} = 20n$. Since the non asymptotic bounds are essentially based on the control




Figure 1: For $\lambda \in \{0.25, 0.5, 0.75\}$ and $\mu \in (0.01, 0.9)$, evolution of the constant C solving (16)

of $\gamma_{k+1}^{-2} \mathbb{E} [\|\widehat{S}^{k+1} - \widehat{S}^k\|^2]$ (see the sketch of proof of Theorem 3 in section 3), we first compare the algorithms through this criterion: the expectation is approximated by a Monte Carlo sum over $1e3$ independent runs. The second criterion for comparison is a distance of the iterates to the unique solution θ_* via the expectation $\mathbb{E} [\|\theta^k - \theta_*\|]$ and the standard deviation $\text{std}(\|\theta^k - \theta_*\|)$ again approximated by a Monte Carlo sum over the same $1e3$ independent runs.

Figure 7 displays the evolution of $k \mapsto \lambda_{k+1}^*$, the optimal coefficient given by (12); in this toy example, it is computed explicitly. As intuited in subsection 2.2.2, we obtain $\lambda_{k+1}^* \approx 1$ for large iteration indexes k ; **FIEM** and **opt-FIEM** have the same (or almost the same) update scheme $\widehat{S}^k \rightarrow \widehat{S}^{k+1}$. The ratio of the expectations $\mathbb{E} [\|\theta_{\text{opt-FIEM}}^k - \theta_*\|] / \mathbb{E} [\|\theta_{\text{Alg}}^k - \theta_*\|]$ and of the standard deviations $\text{std}(\|\theta_{\text{opt-FIEM}}^k - \theta_*\|) / \text{std}(\|\theta_{\text{Alg}}^k - \theta_*\|)$ are displayed on Figure 8 when Alg is **FIEM** and **Online EM**.




Figure 2: For $\lambda \in \{0.25, 0.5, 0.75\}$ and $\mu \in (0.01, 0.9)$, evolution of the constant C solving (22)

They are shown as a function of k for $k \in \{1e2, 5e2, 1e3, 1.5e3, \dots, 6e3, 7e3, \dots, 20e3\}$. When Alg is FIEM and the number of iterations k is large, we observe that both the ratio of the mean values and the ratio of the standard deviations tend to one: this is an echo to the previous comment $\lambda_k^* \approx 1$. Note also that when k is large, **Online EM** has a really poor behavior when compared to **opt-FIEM** (and therefore also to FIEM). For the first iterations of the algorithm, we observe first that **opt-FIEM** and **Online-EM** escape more rapidly from the (possibly bad) initial value than FIEM; **opt-FIEM** surpasses FIEM by reducing the variance up to 22%. Second, the plot also shows that **Online EM** may reduce the variability of **opt-FIEM** up to 18%, but **opt-FIEM** provides a drastic variability reduction in the first iterations. Since we advocate to stop FIEM at a random time K sampled in the range $\{0, \dots, K_{\max} - 1\}$, **opt-FIEM** gives insights on how to improve the behavior of incremental EM algorithms in the first




Figure 3: For $\lambda \in \{0.25, 0.5, 0.75\}$ and $\mu \in (0.01, 0.9)$, evolution of the quantity \mathcal{B} given by 4

iterations. Figure 9 shows $k \mapsto \gamma_{k+1}^{-2} \mathbb{E} \left[\|\widehat{S}^{k+1} - \widehat{S}^k\|^2 \right]$ for the three algorithms when $k \in [1.5e3, 5e3]$. The plot illustrates again that **opt-FIEM** improves **FIEM** during these first iterations; and improves drastically **Online EM**.

5 Mixture of Gaussian distributions

Notations. For two $p \times p$ matrices A, B , $\langle A, B \rangle$ is the trace of $B^T A$: $\langle A, B \rangle \stackrel{\text{def}}{=} \text{Tr}(B^T A)$. I_p stands for the $p \times p$ identity matrix. \otimes stands for the Kronecker product. \mathcal{M}_p^+ denotes the set of the invertible $p \times p$ covariance matrices. $\det(A)$ is the determinant of the matrix A .

In this section ², **FIEM** is applied to solve Maximum Likelihood inference in a

²see footnote 1




Figure 4: For $\lambda \in \{0.25, 0.5, 0.75\}$ and $\mu \in (0.01, 0.9)$, evolution of the quantity \mathcal{B} given by 5

mixture of g Gaussian distributions centered at μ_ℓ and sharing the same covariance matrix Σ (see Frühwirth-Schnatter et al. (2019) for a recent review on mixture models): given n \mathbb{R}^p -valued observations y_1, \dots, y_n , find a point $\hat{\theta}_n^{\text{ML}} \in \Theta$ satisfying $F(\hat{\theta}_n^{\text{ML}}) \stackrel{\text{def}}{=} R(\hat{\theta}_n^{\text{ML}}) + n^{-1} \sum_{i=1}^n \mathcal{L}_i(\hat{\theta}_n^{\text{ML}}) \leq F(\theta)$ for any $\theta \in \Theta$ where $\theta \stackrel{\text{def}}{=} (\alpha_1, \dots, \alpha_g, \mu_1, \dots, \mu_g, \Sigma)$,

$$\Theta \stackrel{\text{def}}{=} \left\{ \alpha_\ell \geq 0, \sum_{\ell=1}^g \alpha_\ell = 1 \right\} \times \mathbb{R}^{pg} \times (\mathcal{M}_p^+) \subseteq \mathbb{R}^{g+pg+(p \times p)}.$$

In addition,

$$R(\theta) + \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\theta) = -\frac{1}{n} \sum_{i=1}^n \log \sum_{\ell=1}^g \alpha_\ell \mathcal{N}_p(\mu_\ell, \Sigma)[y_i],$$




Figure 5: Value of the constant step size given by Karimi et al., 4 (Case 1) and 5 (Case 2). The step size is shown as a function of $\mu \in (0.01, 0.9)$. In Case 2, different strategies for K_{\max} are considered.

where we set (the term $p \log(2\pi)/2$ is omitted)

$$R(\theta) \stackrel{\text{def}}{=} \frac{1}{2} \log \det(\Sigma) + \frac{1}{2n} \sum_{i=1}^n y_i^T \Sigma^{-1} y_i = \frac{1}{2} \left(\left\langle \Sigma^{-1}, \frac{1}{n} \sum_{i=1}^n y_i y_i^T \right\rangle - \log \det(\Sigma^{-1}) \right).$$

In this example, $\mathcal{L}_i(\theta) = -\log \sum_{z=1}^g \exp(\langle s_i(z), \phi(\theta) \rangle)$ with

$$s_i(z) = \mathbf{A}_{y_i} \begin{bmatrix} 1_{z=1} \\ \dots \\ 1_{z=g} \end{bmatrix} \in \mathbb{R}^{g+pg}, \quad \mathbf{A}_y \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{I}_g \\ \mathbf{I}_g \otimes y \end{bmatrix}.$$

We use the MNIST dataset ³. The data are pre-processed as in Nguyen et al. (2020):

³available at <http://yann.lecun.com/exdb/mnist/>

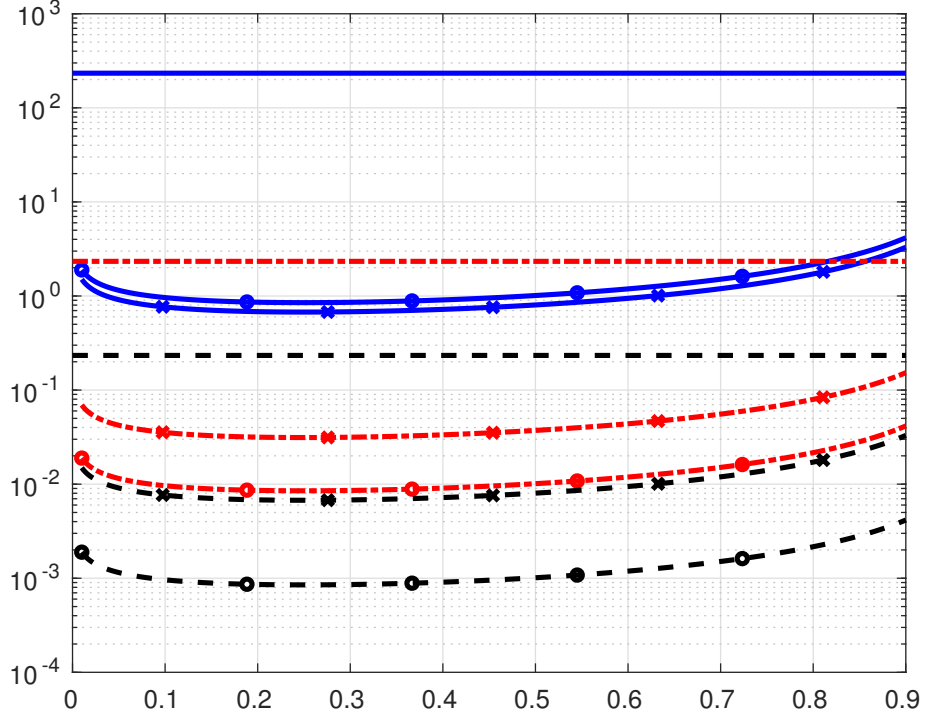


Figure 6: Value of the control $n^a K_{\max}^{-b} \mathcal{B}$ given by 4 (Case 1, with a circle), 5 (Case 2, with a cross) and Karimi et al. (no markers). The control is displayed as a function of $\mu \in (0.01, 0.9)$ and for different values of K_{\max} : $K_{\max} = n$ (solid line), $K_{\max} = 1e2 n$ (dash-dot line) and $K_{\max} = 1e3 n$ (dashed line).

the training set contains $n = 6e4$ images of size 28×28 ; among these 784 pixels, 67 are non informative since they are constant over all the pictures so they are removed yielding to n observations of length 717; each feature is centered and standardized (among the n observations) and a PCA of the associated 717×717 covariance matrix is applied in order to summarize the features by the first $p = 20$ principal components. In the numerical applications, we fix $g = 12$ components in the mixture.




Figure 7: The coefficient λ_k^* (see (12)) as a function of the number of iterations k ; it is a random variable, and the solid line is the mean value (the dashed lines are resp. the quantiles 0.25 and 0.75) over $1e3$ independent paths.

The maximization step $s \mapsto \mathbb{T}(s)$ is given by

$$\begin{aligned} \widehat{\alpha}_\ell &\stackrel{\text{def}}{=} \frac{s_\ell}{\sum_{u=1}^g s_u}, \ell \in [g]^*, \\ \widehat{\mu}_\ell &\stackrel{\text{def}}{=} \frac{s_{g+(\ell-1)p+1:g+\ell p}}{s_\ell}, \ell \in [g]^*, \\ \widehat{\Sigma} &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T - \sum_{\ell=1}^g s_\ell \widehat{\mu}_\ell \widehat{\mu}_\ell^T, \end{aligned}$$

where $s \stackrel{\text{def}}{=} (s_1, \dots, s_{g+pg})$; see subsection 10.1.3 in the supplementary material.

Since we want $\mathbb{T}(s) \in \Theta$, \mathbb{T} is defined at least on $\mathcal{S} \subset \mathbb{R}^{g+pg}$:

$$\mathcal{S} \stackrel{\text{def}}{=} \left\{ n^{-1} \sum_{i=1}^n A_{y_i} \rho_i : \right. \\ \left. \rho_i = (\rho_{i,1}, \dots, \rho_{i,g}), \rho_{i,\ell} \geq 0, \sum_{\ell=1}^g \rho_{i,\ell} = 1 \right\} ;$$

see subsection 10.1.4 in the supplementary material.

This model is used to go beyond the theoretical framework adopted in this paper. The first extension concerns the domain of \mathbb{T} : A3 assumes that \mathbb{T} is defined on \mathbb{R}^q (here, $q = g + pg$) while the above description shows that it is not always true. This gap between theory and application is classical for mixture of Gaussian distributions (while $\rho_{\ell,i}$ may be a signed quantity or while we may have $\sum_{\ell=1}^g \rho_{i,\ell} \neq 1$ for the considered algorithms (see subsection 10.2.2 to subsection 10.2.5 in the supplementary material for a detailed derivation), numerically we always obtained quantities \widehat{S}^k which were in \mathcal{S} .

The second extension concerns the use of mini-batches at each iteration of incremental EM algorithms: instead of sampling one example per iteration (see e.g. line 4, line 8 in Algorithm 2, line 2 in Algorithm 3 and line 4 in Algorithm 4), a mini-batch of size b is used - sampled at random from the n available examples, possibly with replacement. In the supplementary material, we provide in subsection 10.2 a description of **iEM**, **Online EM** and **FIEM** in the case $b > 1$.

EM, **iEM**, **Online EM** and **FIEM** are compared when used to solve the above Maximum Likelihood inference problem. All the paths of these algorithms are started from the same point $\theta^0 \in \Theta$ defined by the randomization scheme described in (Kwedlo, 2015, section 4); we then set $\widehat{S}^0 \stackrel{\text{def}}{=} n^{-1} \sum_{i=1}^n \bar{s}_i(\theta^0)$; the normalized log-likelihood $-F(\theta^0)$ is equal to -58.3097 (equivalently, the unnormalized log-likelihood is $-3.4986 \text{e}+6$). Note that, as mentioned below, the evaluation of the log-likelihood does not include the constant $+p \log(2\pi)/2$.

Each iteration of **iEM**, **Online EM** (resp. **FIEM**) calls a mini-batch of $b = 100$ examples (resp. 2 mini-batches of size $b = 100$ examples each) sampled uniformly from $[n]^*$ with replacement; for a fair comparison of the paths produced by these algorithms, the same seed is used.

The paths are seen as cycles of *epochs*, an epoch being defined as the processing of n examples: for **EM**, an epoch is one iteration; for **iEM** and **Online EM**, an epoch is n/b iterations; for **FIEM**, an epoch is $n/(2b)$ iterations. Below, the paths are run until $100n$ examples are processed, which means 100 iterations or epochs for **EM**, and $100n/b$ iterations (or 100 epochs) for both **iEM** and **Online EM**. Instead of a pure **FIEM** algorithm, we implement **h-FIEM**, an hybrid algorithm obtained by

first running `kswitch` epochs of `Online EM` and then switching to epochs of `FIEM`: we choose `kswitch = 6` so that `h-FIEM` processes $100n$ examples after $6n/b$ iterations (or 6 epochs) of `Online EM` and $94n/(2b)$ iterations (or 94 epochs) of `FIEM`. The use of `h-FIEM` is to explicitly illustrate the variance reduction of the `FIEM` iterations when compared to the `Online EM` ones.

`iEM` is run with the constant step size $\gamma_{k+1} = 1$; `Online EM` and `FIEM` are run with $\gamma_{k+1} = 5e-3$.

Figure 10 and Figure 11 display the normalized log-likelihood along a path of `EM`, `iEM`, `Online EM` and `h-FIEM`, resp. for the first epochs (from 1 to 25) and by discarding the first ones (from 15 to 100). The first conclusion is that the incremental methods forget the initial value far more rapidly than `EM`, which is the consequence of the incremental processing of the observations which allow many updates of the parameter θ^k (or equivalently, of the statistic \hat{S}^k) before the use of n examples (which is equivalent to the learning cost of one iteration of `EM`). The second conclusion is that the incremental EM-based methods perform a better maximization of the normalized log-likelihood $-F$. Finally, `Online EM` and `h-FIEM` are better than `iEM`: the log-likelihood converges resp. to $-1.9094e+6$, $-1.9080e+6$ and $-1.9100e+6$ (the plot displays the normalized log-likelihood); and it is clear that `h-FIEM` reduces the variability of the `Online EM` path. The same conclusions are drawn from different runs; the supplementary material provides a similar plot when the curves are the average over 10 independent paths; Table 1 reports the mean value and the standard deviation of the log-likelihood over these 10 runs.

Table 1: Normalized log-likelihood along a `EM`, `iEM`, `Online EM` and `h-FIEM` path, at epoch #1, 15, 25, 50, 100. The value is the average over 10 independent runs (the standard deviation is in parenthesis). The log-likelihood is obtained by multiplying by $n = 6e+4$.

	#1	#15	#25	#50	#100
<code>EM</code>	$-3.4102e+1$ -	$-3.2033e+1$ -	$-3.1896e+1$ -	$-3.1890e+1$ -	$-3.1889e+1$ -
<code>iEM</code>	$-3.3672e+1$ ($4.90e-3$)	$-3.1982e+1$ ($5.33e-2$)	$-3.1869e+1$ ($1.87e-2$)	$-3.1843e+1$ ($1.87e-2$)	$-3.1827e+1$ ($1.38e-2$)
<code>Online EM</code>	$-3.2999e+1$ ($2.67e-2$)	$-3.1872e+1$ ($5.14e-2$)	$-3.1828e+1$ ($4.67e-2$)	$-3.1823e+1$ ($4.68e-2$)	$-3.1823e+1$ ($4.50e-2$)
<code>h-FIEM</code>	$-3.2999e+1$ ($2.67e-2$)	$-3.1900e+1$ ($5.31e-2$)	$-3.1853e+1$ ($6.94e-2$)	$-3.1806e+1$ ($5.18e-2$)	$-3.1804e+1$ ($5.25e-2$)

A fluctuation of 1% (resp. 1 ‰) around the optimal normalized log-likelihood corresponds to a lower bound of -32.1184 (resp. -31.8322): for `EM` such an accuracy

is reached after 12 iterations (resp. is never reached); for **iEM**, it is reached after 11 epochs (resp. is never reached); for **Online EM**, after 4 epochs (resp. 23 epochs); for **h-FIEM**, after 4 epochs (resp. 34 epochs). An accuracy of 1 ‰ is never reached by **Online EM** and is reached after 36 epochs for **h-FIEM**.

Figure 12 shows the estimation of the $g = 12$ weights α_ℓ along a path of length 100 epochs. The comparison of **Online EM** (bottom left) and **h-FIEM** (bottom right) shows that **h-FIEM** acts as a variability reduction technique along the path, without slowing down the convergence rate. Figure 13 displays the limiting value of these paths i.e. the estimate of the weights $\alpha_1, \dots, \alpha_g$ defined as the value of the parameter at the end of 100 epochs; the weights are sorted in descending order. **Online EM** and **h-FIEM** provide similar estimates.

6 Proof

6.1 Proof of section 2

6.1.1 Proof of 1

(Proof of 1). The statements are trivial and we only prove the first claim: if $s^\star = \bar{s} \circ \mathbb{T}(s^\star)$ then by applying \mathbb{T} (under the uniqueness assumption A3), we have $\mathbb{T}(s^\star) = (\mathbb{T} \circ \bar{s}) \circ \mathbb{T}(s^\star)$ and the proof follows.

(Proof of 2). For $\theta \in \Theta^v$, set $D\phi(\theta) \stackrel{\text{def}}{=} \left(\dot{\phi}(\theta)\right)^T$. By A4-(ii) and a chain rule,

$$\dot{V}(s) = \left(\dot{\mathbb{T}}(s)\right)^T \left\{ \dot{\mathbb{R}}(\mathbb{T}(s)) - D\phi(\mathbb{T}(s)) \bar{s} \circ \mathbb{T}(s) \right\} .$$

Moreover, using A3 and A4-(i), the minimum $\mathbb{T}(s)$ is a critical point of $\theta \mapsto L(s, \theta)$: we have for any $s \in \mathbb{R}^q$, $\dot{\mathbb{R}}(\mathbb{T}(s)) - D\phi(\mathbb{T}(s)) s = 0$. Hence,

$$\dot{V}(s) = - \left(\dot{\mathbb{T}}(s)\right)^T D\phi(\mathbb{T}(s)) h(s) = - (B(s))^T h(s) .$$

A4-(iii) implies that $B^T = B$ and the zeros of h are the zeros of \dot{V} .

6.1.2 Auxiliary result

Lemma 7. *Assume that Θ and $\phi(\Theta)$ are open; and ϕ is continuously differentiable on Θ . Then for all $i \in \{1, \dots, n\}$, \mathcal{L}_i is continuously differentiable on Θ .*

If in addition A1, A2, A3 and A4-(i) hold, then F (resp. $V \stackrel{\text{def}}{=} F \circ \mathbb{T}$) is continuously differentiable on Θ (resp. on \mathbb{R}^q) and for any $\theta \in \Theta$,

$$\dot{F}(\theta) = - \left(\dot{\phi}(\theta)\right)^T \bar{s}(\theta) + \dot{\mathbb{R}}(\theta) .$$

Proof. A1 and (Sundberg, 2019, Proposition 3.8) (see also (Brown, 1986, Theorem 2.2.)) imply that $L_i : \tau \mapsto \int_Z h_i(z) \exp(\langle s_i(z), \tau \rangle) \mu(dz)$ is continuously differentiable on the interior of the set

$$\{\tau \in \mathbb{R}^q, \int_Z h_i(z) \exp(\langle s_i(z), \tau \rangle) \mu(dz) < \infty\}$$

and its derivative is

$$\int_Z s_i(z) h_i(z) \exp(\langle s_i(z), \tau \rangle) \mu(dz) .$$

This set contains $\phi(\Theta)$ under A1. The equality $\mathcal{L}_i = -\log(L_i \circ \phi)$ and the differentiability of composition of functions conclude the proof of the first item. The second one easily follows. \square

6.2 Proofs of section 3

For any $k \geq 0$ and $i \in [n]^*$, we define $\widehat{S}^{<k,i}$ such that

$$\widetilde{S}^k = \frac{1}{n} \sum_{i=1}^n \bar{s}_i \circ \mathsf{T}(\widehat{S}^{<k,i}) ;$$

it means $\widehat{S}^{<0,i} \stackrel{\text{def}}{=} \widehat{S}^0$ for all $i \in [n]^*$ and for $k \geq 0$,

$$\widehat{S}^{<k+1,i} = \widehat{S}^\ell, \tag{25}$$

with

$$\begin{cases} \ell = k & \text{if } I_{k+1} = i, \\ \ell \in [k-1]^* & \text{if } I_{k+1} \neq i, I_k \neq i, \dots, I_{\ell+1} = i, \\ \ell = 0 & \text{otherwise.} \end{cases} \tag{26}$$

Define the filtrations, for $k \geq 0$,

$$\begin{aligned} \mathcal{F}_k &\stackrel{\text{def}}{=} \sigma(\widehat{S}^0, I_1, J_1, \dots, I_k, J_k), \\ \mathcal{F}_{k+1/2} &\stackrel{\text{def}}{=} \sigma(\widehat{S}^0, I_1, J_1, \dots, I_k, J_k, I_{k+1}) ; \end{aligned}$$

note that $\widehat{S}^k \in \mathcal{F}_k$ and $S_{k+1,\cdot} \in \mathcal{F}_{k+1/2}$. Set

$$H_{k+1} \stackrel{\text{def}}{=} \bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \widehat{S}^k + \frac{1}{n} \sum_{i=1}^n S_{k+1,i} - S_{k+1,J_{k+1}} .$$

6.2.1 Proof of Theorem 3

By 1 and A5-(iii), \dot{V} is $L_{\dot{V}}$ -Lipschitz on \mathbb{R}^q , and we have

$$\begin{aligned} V(\widehat{S}^{k+1}) - V(\widehat{S}^k) &\leq \left\langle \widehat{S}^{k+1} - \widehat{S}^k, \dot{V}(\widehat{S}^k) \right\rangle + \frac{L_{\dot{V}}}{2} \|\widehat{S}^{k+1} - \widehat{S}^k\|^2 \\ &\leq \gamma_{k+1} \left\langle H_{k+1}, \dot{V}(\widehat{S}^k) \right\rangle + \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \|H_{k+1}\|^2. \end{aligned}$$

Taking the expectation yields, upon noting that $\widehat{S}^k \in \mathcal{F}_k$,

$$\begin{aligned} &\mathbb{E} \left[V(\widehat{S}^{k+1}) \right] - \mathbb{E} \left[V(\widehat{S}^k) \right] \\ &\leq \gamma_{k+1} \mathbb{E} \left[\left\langle \mathbb{E} [H_{k+1} | \mathcal{F}_k], \dot{V}(\widehat{S}^k) \right\rangle \right] + \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \mathbb{E} [\|H_{k+1}\|^2] \\ &\leq \gamma_{k+1} \mathbb{E} \left[\left\langle h(\widehat{S}^k), \dot{V}(\widehat{S}^k) \right\rangle \right] + \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \mathbb{E} [\|H_{k+1}\|^2] \\ &\leq -\gamma_{k+1} v_{\min} \mathbb{E} [\|h(\widehat{S}^k)\|^2] + \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \mathbb{E} [\|H_{k+1}\|^2] \\ &\leq -\gamma_{k+1} \left(v_{\min} - \gamma_{k+1} \frac{L_{\dot{V}}}{2} \right) \mathbb{E} [\|h(\widehat{S}^k)\|^2] \\ &\quad + \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \mathbb{E} [\|H_{k+1} - h(\widehat{S}^k)\|^2] \end{aligned}$$

where we used that $\mathbb{E} [H_{k+1} | \mathcal{F}_k] = h(\widehat{S}^k)$ and 2. Set

$$\begin{aligned} A_k &\stackrel{\text{def}}{=} \mathbb{E} [\|h(\widehat{S}^k)\|^2] , \\ B_{k+1} &\stackrel{\text{def}}{=} \mathbb{E} [\|\widetilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\widehat{S}^k)\|^2] . \end{aligned}$$

By 8 and 9, we have for any $k \geq 0$:

$$\begin{aligned} \mathbb{E} \left[V(\widehat{S}^{k+1}) \right] - \mathbb{E} \left[V(\widehat{S}^k) \right] &\leq -\gamma_{k+1} \left(v_{\min} - \gamma_{k+1} \frac{L_{\dot{V}}}{2} \right) A_k - \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} B_{k+1} \\ &\quad + \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \mathbb{E} [\|\bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \mathsf{S}_{k+1, J_{k+1}}\|^2] \\ &\leq T_{1,k} + T_{2,k+1} \end{aligned}$$

by setting

$$\begin{aligned} T_{1,k} &\stackrel{\text{def}}{=} -\gamma_{k+1} \left(v_{\min} - \gamma_{k+1} \frac{L_{\dot{V}}}{2} \right) A_k + \gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \sum_{j=0}^{k-1} \tilde{\Lambda}_{j+1,k} A_j \\ T_{2,k+1} &\stackrel{\text{def}}{=} -\gamma_{k+1}^2 \frac{L_{\dot{V}}}{2} \left\{ B_{k+1} + \sum_{j=0}^{k-1} \tilde{\Lambda}_{j+1,k} (1 + \beta_{j+1}^{-1})^{-1} B_{j+1} \right\} ; \end{aligned}$$

by convention, $\sum_{j=0}^1 a_j = 0$. By summing from $k = 0$ to $k = K_{\max} - 1$, we have

$$\begin{aligned} \sum_{k=0}^{K_{\max}-1} \gamma_{k+1} \left(v_{\min} - \gamma_{k+1} \frac{L\dot{V}}{2} \right) A_k - \frac{L\dot{V}L^2}{2} \sum_{k=0}^{K_{\max}-2} \gamma_{k+1}^2 \Lambda_k A_k \\ \leq \Delta V - \frac{L\dot{V}}{2} \sum_{k=0}^{K_{\max}-1} \gamma_{k+1}^2 (L^2 \Xi_k + 1) B_{k+1}, \end{aligned}$$

where for $0 \leq k \leq K_{\max} - 2$ and with the convention $\Lambda_{K_{\max}-1} = \Xi_{K_{\max}-1} = 0$,

$$\begin{aligned} \Lambda_k &\stackrel{\text{def}}{=} \left(1 + \frac{1}{\beta_{k+1}} \right) \sum_{j=k+1}^{K_{\max}-1} \gamma_{j+1}^2 \left(\frac{n-1}{n} \right)^{j-k} \prod_{\ell=k+2}^j (1 + \beta_\ell + \gamma_\ell^2 L^2) \\ &\leq \left(1 + \frac{1}{\beta_{k+1}} \right) \sum_{j=k+1}^{K_{\max}-1} \gamma_{j+1}^2 \prod_{\ell=k+2}^j \left(1 - \frac{1}{n} + \beta_\ell + \gamma_\ell^2 L^2 \right), \\ \Xi_k &\stackrel{\text{def}}{=} \left(1 + \frac{1}{\beta_{k+1}} \right)^{-1} \Lambda_k = \frac{\Lambda_k \beta_{k+1}}{1 + \beta_{k+1}}. \end{aligned}$$

Hence,

$$\begin{aligned} \sum_{k=0}^{K_{\max}-1} \left\{ \gamma_{k+1} \left(v_{\min} - \gamma_{k+1} \frac{L\dot{V}}{2} \right) - \gamma_{k+1}^2 \Lambda_k \frac{L\dot{V}L^2}{2} \right\} A_k \\ + \sum_{k=0}^{K_{\max}-1} \gamma_{k+1}^2 \{1 + \Xi_k L^2\} \frac{L\dot{V}}{2} B_{k+1} \leq \Delta V. \end{aligned}$$

6.2.2 Proof of 4

It is a follow-up of Theorem 3; the quantities $\alpha_k, \Lambda_k, \delta_k$ introduced in the statement of Theorem 3 are used below without being defined again. We consider the case when for $\ell \in [K_{\max}]^*$,

$$\beta_\ell \stackrel{\text{def}}{=} \frac{1-\lambda}{n^{\mathbf{b}}}, \quad \gamma_\ell^2 \stackrel{\text{def}}{=} \frac{C}{L^2 n^{2\mathbf{c}} K_{\max}^{2\mathbf{d}}},$$

for some $\lambda \in (0, 1)$, $C > 0$ and $\mathbf{b}, \mathbf{c}, \mathbf{d}$ to be defined in the proof in such a way that (i) $\alpha_k \geq 0$, (ii) $\sum_{k=0}^{K_{\max}-1} \alpha_k$ is positive and as large as possible. Since there will be a discussion on (n, C, λ) , we make more explicit the dependence of some constants upon these quantities: α_k will be denoted by $\alpha_k(n, C, \lambda)$.

With these definitions, we have

$$1 - \frac{\rho_n}{n} \stackrel{\text{def}}{=} 1 - \frac{1}{n} + \beta_\ell + \gamma_\ell^2 L^2 = 1 - \frac{1}{n} \left(1 - \frac{1-\lambda}{n^{\mathbf{b}-1}} - \frac{C}{n^{2\mathbf{c}-1} K_{\max}^{2\mathbf{d}}} \right),$$

and choose $(\mathbf{b}, \mathbf{c}, \mathbf{d}, \lambda, C)$ such that

$$\frac{1-\lambda}{n^{\mathbf{b}-1}} + \frac{C}{n^{2\mathbf{c}-1}K_{\max}^{2\mathbf{d}}} < 1, \quad (27)$$

which ensures that $\rho_n \in (0, 1)$. Hence, for any $k \in [K_{\max} - 2]$,

$$\begin{aligned} \Lambda_k &\leq n^{\mathbf{b}} \left(\frac{1}{n^{\mathbf{b}}} + \frac{1}{1-\lambda} \right) \frac{C}{L^2 n^{2\mathbf{c}} K_{\max}^{2\mathbf{d}}} \sum_{j=k+1}^{K_{\max}-1} \left(1 - \frac{\rho_n}{n} \right)^{j-k-1} \\ &\leq \left(\frac{1}{n^{\mathbf{b}}} + \frac{1}{1-\lambda} \right) \frac{C}{L^2 \rho_n} \frac{1}{n^{2\mathbf{c}-\mathbf{b}-1} K_{\max}^{2\mathbf{d}}}. \end{aligned}$$

From this upper bound, we deduce for any $k \in [K_{\max} - 1]$: $\alpha_k(n, C, \lambda) \geq \underline{\alpha}_n(C, \lambda)$ where

$$\underline{\alpha}_n(C, \lambda) \stackrel{\text{def}}{=} \frac{\sqrt{C}}{Ln^{\mathbf{c}}K_{\max}^{\mathbf{d}}} \left(v_{\min} - \frac{L\dot{V}}{2L} \frac{\sqrt{C}}{n^{\mathbf{c}}K_{\max}^{\mathbf{d}}} - \frac{L\dot{V}}{2L} \frac{C^{3/2}}{\rho_n n^{3\mathbf{c}-\mathbf{b}-1} K_{\max}^{3\mathbf{d}}} \left(\frac{1}{n^{\mathbf{b}}} + \frac{1}{1-\lambda} \right) \right). \quad (28)$$

From (27) and (28), we choose $\mathbf{b} = 1$, $\mathbf{c} = 2/3$, $\mathbf{d} = 0$; which yields for $n \geq 1$, since $\rho_n = \lambda - Cn^{-1/3}$

$$n^{2/3} \underline{\alpha}_n(C, \lambda) \geq \mathcal{L}_n(C, \lambda),$$

with

$$\begin{aligned} \mathcal{L}_n(C, \lambda) &\stackrel{\text{def}}{=} \frac{L\dot{V}\sqrt{C}}{2L^2} \left(v_{\min} \frac{2L}{L\dot{V}} - \sqrt{C} f_n(C, \lambda) \right), \\ f_n(C, \lambda) &\stackrel{\text{def}}{=} \frac{1}{n^{2/3}} + \frac{C}{\lambda - Cn^{-1/3}} \left(\frac{1}{n} + \frac{1}{1-\lambda} \right). \end{aligned}$$

Let $\mu \in (0, 1)$. Fix $\lambda \in (0, 1)$ and $C > 0$ such that (see (27) for the second condition)

$$\sqrt{C} f_n(C, \lambda) = 2\mu v_{\min} \frac{L}{L\dot{V}}, \quad \frac{1}{n^{1/3}} < \frac{\lambda}{C}. \quad (29)$$

This implies that $n^{2/3} \alpha_k(n, C, \lambda) \geq n^{2/3} \underline{\alpha}_n(C, \lambda) \geq n^{2/3} \alpha_{\star}(C) \stackrel{\text{def}}{=} \sqrt{C}(1-\mu)v_{\min}/L$. We obtain an upper bound on \mathbf{E}_1 by

$$\mathbf{E}_1 \leq \frac{1}{K_{\max} \alpha_{\star}(C)} \sum_{k=0}^{K_{\max}-1} \alpha_k(n, C, \lambda) \mathbb{E} \left[\|h(\widehat{S}^k)\|^2 \right].$$

For \mathbf{E}_2 , since $\delta_k \geq L_{\dot{V}} \gamma_{k+1}^2/2$,

$$\begin{aligned} & \frac{L_{\dot{V}} \sqrt{C}}{2L(1-\mu)n^{2/3} v_{\min}} \frac{1}{v_{\min}} \mathbf{E}_2 \\ & \leq \frac{L_{\dot{V}} C}{2L^2 n^{4/3} K_{\max} \alpha_{\star}(C)} \sum_{k=0}^{K_{\max}-1} \mathbb{E} \left[\|\tilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\hat{S}^k)\|^2 \right] \\ & \leq \frac{1}{K_{\max} \alpha_{\star}(C)} \sum_{k=0}^{K_{\max}-1} \delta_k \mathbb{E} \left[\|\tilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\hat{S}^k)\|^2 \right]. \end{aligned}$$

We then conclude by

$$\frac{1}{K_{\max} \alpha_{\star}(C)} = \frac{n^{2/3}}{K_{\max}} \frac{L}{\sqrt{C}(1-\mu)v_{\min}}, \quad (30)$$

and use $\sqrt{C}f_n(C, \lambda) = 2\mu v_{\min} L/L_{\dot{V}}$.

6.2.3 Proof of 5

It is a follow-up of Theorem 3; the quantities $\alpha_k, \Lambda_k, \delta_k$ introduced in the statement of Theorem 3 are used below without being defined again.

We consider the case when, for $\ell \in [K_{\max}]^*$,

$$\beta_{\ell} \stackrel{\text{def}}{=} \frac{1-\lambda}{n^{\mathbf{b}}}, \quad \gamma_{\ell}^2 \stackrel{\text{def}}{=} \frac{C}{L^2 n^{2\mathbf{c}} K_{\max}^{2\mathbf{d}}}$$

for some $\lambda \in (0, 1)$, $C > 0$ and $\mathbf{b}, \mathbf{c}, \mathbf{d}$ to be defined in the proof in such a way that (i) $\alpha_k \geq 0$, (ii) $\sum_{k=0}^{K_{\max}-1} \alpha_k$ is positive and as large as possible. Since there will be a discussion on (n, C, λ) , we make more explicit the dependence of some constants upon these quantities: α_k will be denoted by $\alpha_k(n, C, \lambda)$.

With these definitions, we have

$$\rho \stackrel{\text{def}}{=} 1 - \frac{1}{n} + \beta_{\ell} + L^2 \gamma_{\ell}^2 = 1 - \frac{1}{n} \left(1 - \frac{1-\lambda}{n^{\mathbf{b}-1}} - \frac{C}{n^{2\mathbf{c}-1} K_{\max}^{2\mathbf{d}}} \right),$$

and choose $(\mathbf{b}, \mathbf{c}, \mathbf{d}, \lambda, C)$ such that

$$\frac{1-\lambda}{n^{\mathbf{b}-1}} + \frac{C}{n^{2\mathbf{c}-1} K_{\max}^{2\mathbf{d}}} \leq 1, \quad (31)$$

which ensures that $\rho \in (0, 1]$. Hence, for any $k \in [K_{\max} - 2]$,

$$\begin{aligned} \Lambda_k &\leq n^b \left(\frac{1}{n^b} + \frac{1}{1-\lambda} \right) \frac{C}{L^2 n^{2c} K_{\max}^{2d}} \sum_{j=k+1}^{K_{\max}-1} \rho^{j-k-1} \\ &\leq \left(\frac{1}{n^b} + \frac{1}{1-\lambda} \right) \frac{C}{L^2 n^{2c-b} K_{\max}^{2d-1}}. \end{aligned}$$

From this upper bound, we obtain the following lower bound for any $k \in [K_{\max} - 1]$: $\alpha_k(n, C, \lambda) \geq \underline{\alpha}_n(C, \lambda)$ where

$$(n^c K_{\max}^d) \underline{\alpha}_n(C, \lambda) \stackrel{\text{def}}{=} \frac{\sqrt{C}}{L} \left(v_{\min} - \sqrt{C} \frac{L\dot{V}}{2L} \left\{ \frac{1}{n^c K_{\max}^d} + \frac{C}{n^{3c-b} K_{\max}^{3d-1}} \left(\frac{1}{n^b} + \frac{1}{1-\lambda} \right) \right\} \right).$$

Based on this inequality and on (31), we choose $b = 1$ and $c = d = 1/3$; which yields for $n \geq 1$,

$$\begin{aligned} (nK_{\max})^{1/3} \underline{\alpha}_n(C, \lambda) &= \mathcal{L}_n(C, \lambda) \stackrel{\text{def}}{=} \frac{\sqrt{C}L\dot{V}}{2L^2} \left(v_{\min} \frac{2L}{L\dot{V}} - \sqrt{C} \tilde{f}_n(C, \lambda) \right), \\ \tilde{f}_n(C, \lambda) &\stackrel{\text{def}}{=} \frac{1}{(nK_{\max})^{1/3}} + C \left(\frac{1}{n} + \frac{1}{1-\lambda} \right). \end{aligned}$$

Let $\mu \in (0, 1)$. Fix $\lambda \in (0, 1)$ and $C > 0$ such that (see (31) for the second condition)

$$\sqrt{C} \tilde{f}_n(C, \lambda) = 2\mu v_{\min} \frac{L}{L\dot{V}}, \quad \frac{n^{1/3}}{K_{\max}^{2/3}} \leq \frac{\lambda}{C}. \quad (32)$$

This implies that

$$\begin{aligned} (nK_{\max})^{1/3} \alpha_k(n, C, \lambda) &\geq (nK_{\max})^{1/3} \underline{\alpha}_n(C, \lambda) \\ &\geq (nK_{\max})^{1/3} \alpha_{\star}(C) \stackrel{\text{def}}{=} \sqrt{C} (1-\mu) v_{\min} / L. \end{aligned}$$

We obtain the upper bound on \mathbf{E}_1 by

$$\mathbf{E}_1 \leq \frac{1}{K_{\max} \alpha_{\star}(C)} \sum_{k=0}^{K_{\max}-1} \alpha_k(n, C, \lambda) \mathbb{E} \left[\|h(\hat{S}^k)\|^2 \right].$$

For \mathbf{E}_2 and since $\delta_k \geq L_{\dot{V}} \gamma_{k+1}^2 / 2$

$$\begin{aligned} \frac{L_{\dot{V}} \sqrt{C}}{2(1-\mu)Ln^{1/3}} \frac{1}{K_{\max}^{1/3} v_{\min}} \mathbf{E}_2 &\leq \frac{L_{\dot{V}} C}{2L^2 n^{2/3} K_{\max}^{2/3}} \frac{1}{K_{\max} \alpha_{\star}(C)} \sum_{k=0}^{K_{\max}-1} \mathbb{E} \left[\|\tilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\hat{S}^k)\|^2 \right] \\ &\leq \frac{1}{K_{\max} \alpha_{\star}(C)} \sum_{k=0}^{K_{\max}-1} \delta_k \mathbb{E} \left[\|\tilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\hat{S}^k)\|^2 \right]. \end{aligned}$$

We then conclude by

$$\frac{1}{K_{\max} \alpha_{\star}(C)} = \frac{n^{1/3}}{K_{\max}^{2/3}} \frac{L}{\sqrt{C}(1-\mu)v_{\min}}, \quad (33)$$

and use $\sqrt{C} \tilde{f}_n(C, \lambda) = 2\mu v_{\min} L / L_{\dot{V}}$.

6.2.4 Proof of 6

It is a follow-up of Theorem 3; the quantities $\alpha_k, \Lambda_k, \delta_k$ introduced in the statement of Theorem 3 are used below without being defined again.

Let $p_0, \dots, p_{K_{\max}-1}$ be positive real numbers such that $\sum_{k=0}^{K_{\max}-1} p_k = 1$. We consider the case when

$$\beta_\ell \stackrel{\text{def}}{=} \frac{1-\lambda}{n^{\mathbf{b}}}, \quad \gamma_\ell^2 \stackrel{\text{def}}{=} \frac{C_\ell}{L^2 n^{2\mathbf{c}} K_{\max}^{2\mathbf{d}}},$$

for $\lambda \in (0, 1)$, $C_\ell > 0$, and $\mathbf{b}, \mathbf{c}, \mathbf{d}$ to be defined in the proof.

The first step consists in the definition of a function \mathcal{A} and of a family \mathcal{C} of vectors $\underline{C} = (C_1, \dots, C_{K_{\max}}) \in (\mathbb{R}^+)^{K_{\max}}$ such that

$$\alpha_k \geq \mathcal{A}(C_{k+1}) \geq 0, \quad \sum_{\ell=0}^{K_{\max}-1} \mathcal{A}(C_{\ell+1}) > 0.$$

The second step proves that we can find $\underline{C} \in \mathcal{C}$ such that $p_k = \mathcal{A}(C_{k+1}) / \sum_{\ell=0}^{K_{\max}-1} \mathcal{A}(C_{\ell+1})$ for any $k \in [K_{\max} - 1]$.

Such a pair $(\mathcal{A}, \underline{C})$ is not unique, and among the possible ones, we indicate two strategies, all motivated by making the sum $\sum_{\ell=0}^{K_{\max}-1} \mathcal{A}(C_{\ell+1})$ as large as possible.

Step 1- Definition of the function \mathcal{A} . With the definition of the sequences γ_ℓ and β_ℓ , we have

$$1 - \frac{\rho_{n,\ell}}{n} \stackrel{\text{def}}{=} 1 - \frac{1}{n} + \beta_\ell + \gamma_\ell^2 L^2 = -\frac{1}{n} \left(1 - \frac{1-\lambda}{n^{\mathbf{b}-1}} - \frac{C_\ell}{n^{2\mathbf{c}-1} K_{\max}^{2\mathbf{d}}} \right)$$

and choose $(\mathbf{b}, \mathbf{c}, \mathbf{d}, \lambda, C_\ell)$ such that

$$\frac{1 - \lambda}{n^{\mathbf{b}-1}} + \frac{C_{\max}}{n^{2\mathbf{c}-1} K_{\max}^{2\mathbf{d}}} < 1, \text{ where } C_{\max} \stackrel{\text{def}}{=} \max_{\ell} C_{\ell}, \quad (34)$$

which ensures that $\rho_{n,\ell} \in (0, 1)$. Define

$$\rho_n \stackrel{\text{def}}{=} \min_{\ell} \rho_{n,\ell} = 1 - \frac{1 - \lambda}{n^{\mathbf{b}-1}} - \frac{C_{\max}}{n^{2\mathbf{c}-1} K_{\max}^{2\mathbf{d}}}.$$

Hence, for any $k \in [K_{\max} - 2]$,

$$\begin{aligned} \Lambda_k &\leq n^{\mathbf{b}} \left(\frac{1}{n^{\mathbf{b}}} + \frac{1}{1 - \lambda} \right) \frac{1}{L^2 n^{2\mathbf{c}} K_{\max}^{2\mathbf{d}}} \sum_{j=k+1}^{K_{\max}-1} C_{j+1} \left(1 - \frac{\rho_n}{n} \right)^{j-k-1} \\ &\leq \left(\frac{1}{n^{\mathbf{b}}} + \frac{1}{1 - \lambda} \right) \frac{C_{\max}}{L^2 \rho_n} \frac{1}{n^{2\mathbf{c}-\mathbf{b}-1} K_{\max}^{2\mathbf{d}}}. \end{aligned}$$

From this upper bound, we obtain the following lower bound on α_k , for any $k \in [K_{\max} - 1]$,

$$\alpha_k \geq \frac{\sqrt{C_{k+1}}}{L n^{\mathbf{c}} K_{\max}^{\mathbf{d}}} \left(v_{\min} - \frac{L_{\dot{V}} \sqrt{C_{k+1}}}{2L n^{\mathbf{c}} K_{\max}^{\mathbf{d}}} - \frac{L_{\dot{V}} C_{\max} \sqrt{C_{k+1}}}{2L \rho_n n^{3\mathbf{c}-\mathbf{b}-1} K_{\max}^{3\mathbf{d}}} \left(\frac{1}{n^{\mathbf{b}}} + \frac{1}{1 - \lambda} \right) \right).$$

Based on this inequality and on (34), we choose $\mathbf{b} = 1$, $\mathbf{c} = 2/3$, $\mathbf{d} = 0$: this yields $\rho_n = \lambda - C_{\max} n^{-1/3}$ and $\alpha_k \geq \underline{\alpha}_k$ with (see (15) for the definition of f_n)

$$\underline{\alpha}_k \stackrel{\text{def}}{=} \frac{\sqrt{C_{k+1}} L_{\dot{V}}}{2L^2 n^{2/3}} \left(v_{\min} \frac{2L}{L_{\dot{V}}} - \sqrt{C_{k+1}} f_n(C_{\max}, \lambda) \right); \quad (35)$$

the condition (34) gets into $n^{-1/3} < \lambda/C_{\max}$.

Define the quadratic function $x \mapsto \mathcal{A}(x) \stackrel{\text{def}}{=} Ax(v_{\min} - Bx)$ where

$$A \stackrel{\text{def}}{=} \frac{1}{L n^{2/3}}, \quad B \stackrel{\text{def}}{=} f_n(C, \lambda) \frac{L_{\dot{V}}}{2L}; \quad (36)$$

we have $\underline{\alpha}_k = \mathcal{A}(\sqrt{C_{k+1}})$. By 11 in the supplementary material, \mathcal{A} is increasing on $(0, v_{\min}/(2B)]$, reaches its maximum at $x_{\star} \stackrel{\text{def}}{=} v_{\min}/(2B)$ and its maximal value is $\mathcal{A}_{\star} \stackrel{\text{def}}{=} Av_{\min}^2/(4B)$. In addition, its inverse \mathcal{A}^{-1} exists on $(0, \mathcal{A}_{\star}]$.

Step 2- Choice of $C_1, \dots, C_{K_{\max}}$. We are now looking for $C_1, \dots, C_{K_{\max}}$ such that

$$p_k = \mathcal{A}(\sqrt{C_{k+1}}) / \sum_{\ell=0}^{K_{\max}-1} \mathcal{A}(\sqrt{C_{\ell+1}})$$

or equivalently

$$\frac{p_k}{p_I} = \frac{\mathcal{A}(\sqrt{C_{k+1}})}{\mathcal{A}(\sqrt{C_I})}, \quad I \in \operatorname{argmax}_k p_k. \quad (37)$$

It remains to fix $\mathcal{A}(\sqrt{C_I})$ in such a way that \mathcal{A} is invertible on $(0, \sqrt{C_I}]$. Since we also want $\sum_\ell \mathcal{A}(\sqrt{C_{\ell+1}}) = \mathcal{A}(\sqrt{C_I})/p_I$ as large as possible, and \mathcal{A} is increasing on $(0, x_\star]$, we choose

$$\sqrt{C_I} = \sqrt{C_{\max}} = x_\star = \frac{v_{\min}}{2B}. \quad (38)$$

Therefore, C_{\max} solves the equation $\sqrt{C_{\max}} = v_{\min}/(2B)$ or equivalently

$$\frac{v_{\min}L}{L_{\dot{Y}}} = \sqrt{C_{\max}} f_n(C_{\max}, \lambda), \quad (39)$$

under the constraint that $\lambda \in (0, 1)$ and $n^{-1/3} < \lambda/C_{\max}$. When C_{\max} is fixed, we set

$$\sqrt{C_{k+1}} \stackrel{\text{def}}{=} \mathcal{A}^{-1} \left(\frac{p_k}{\max_\ell p_\ell} \mathcal{A}(\sqrt{C_{\max}}) \right).$$

With these definitions, we have (see (37))

$$\frac{1}{\sum_{k=0}^{K_{\max}-1} \mathcal{A}(\sqrt{C_{k+1}})} = \frac{\max_\ell p_\ell}{\mathcal{A}(\sqrt{C_{\max}})}.$$

Remember that

$$\mathcal{A}(\sqrt{C_{\max}}) = \mathcal{A}(x_\star) = v_{\min} \sqrt{C_{\max}} / (2Ln^{2/3}).$$

Step 3. Lower bound on δ_k We write

$$\delta_k \geq \frac{L_{\dot{Y}}}{2} \gamma_{k+1}^2,$$

so that

$$\frac{\delta_k}{\sum_{k=0}^{K_{\max}-1} \mathcal{A}(\sqrt{C_{k+1}})} \geq \frac{L_{\dot{Y}}L}{v_{\min}} n^{2/3} \frac{\max_\ell p_\ell}{\sqrt{C_{\max}}} \gamma_{k+1}^2.$$

6.2.5 Auxiliary results

Lemma 8. *Assume A1, A2 and A3. For any $k \geq 0$,*

$$\mathbb{E} [\|H_{k+1}\|^2] = \mathbb{E} [\|H_{k+1} - h(\widehat{S}^k)\|^2] + \mathbb{E} [\|h(\widehat{S}^k)\|^2],$$

and

$$\begin{aligned} \mathbb{E} [\|H_{k+1} - h(\widehat{S}^k)\|^2] + \mathbb{E} [\|\widetilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\widehat{S}^k)\|^2] \\ = \mathbb{E} [\|\bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - S_{k+1, J_{k+1}}\|^2]. \end{aligned}$$

Proof. Since $\mathbb{E} [H_{k+1} | \mathcal{F}_{k+1/2}] = h(\widehat{S}^k)$, we have

$$\mathbb{E} [\|H_{k+1}\|^2] = \mathbb{E} [\|H_{k+1} - h(\widehat{S}^k)\|^2] + \mathbb{E} [\|h(\widehat{S}^k)\|^2] .$$

In addition, upon noting that $\mathbf{S}_{k+1,i} \in \mathcal{F}_{k+1/2}$ for any $i \in [n]^*$,

$$\begin{aligned} H_{k+1} - h(\widehat{S}^k) &= \bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \mathbf{S}_{k+1,J_{k+1}} - \bar{s} \circ \mathsf{T}(\widehat{S}^k) + \widetilde{S}^{k+1} \\ &= \bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \mathbf{S}_{k+1,J_{k+1}} - \mathbb{E} \left[\bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \mathbf{S}_{k+1,J_{k+1}} \middle| \mathcal{F}_{k+1/2} \right] , \end{aligned}$$

we have

$$\begin{aligned} \mathbb{E} [\|H_{k+1} - h(\widehat{S}^k)\|^2] + \mathbb{E} [\|\widetilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\widehat{S}^k)\|^2] \\ = \mathbb{E} [\|\bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \mathbf{S}_{k+1,J_{k+1}}\|^2] . \end{aligned}$$

□

Proposition 9. *Assume A1, A2, A3 and A5-(ii). Set $L^2 \stackrel{\text{def}}{=} n^{-1} \sum_{i=1}^n L_i^2$. Then*

$$\mathbb{E} [\|\bar{s}_{J_1} \circ \mathsf{T}(\widehat{S}^0) - \mathbf{S}_{1,J_1}\|^2] = 0 ,$$

and for any $k \geq 1$ and $\beta_1, \dots, \beta_k > 0$,

$$\begin{aligned} \mathbb{E} [\|\bar{s}_{J_{k+1}} \circ \mathsf{T}(\widehat{S}^k) - \mathbf{S}_{k+1,J_{k+1}}\|^2] \\ \leq \sum_{j=1}^k \widetilde{\Lambda}_{j,k} \left\{ \mathbb{E} [\|h(\widehat{S}^{j-1})\|^2] - \left(1 + \frac{1}{\beta_j}\right)^{-1} \mathbb{E} [\|\widetilde{S}^j - \bar{s} \circ \mathsf{T}(\widehat{S}^{j-1})\|^2] \right\} , \end{aligned}$$

where

$$\widetilde{\Lambda}_{j,k} \stackrel{\text{def}}{=} L^2 \left(\frac{n-1}{n}\right)^{k-j+1} \gamma_j^2 \left(1 + \frac{1}{\beta_j}\right) \prod_{\ell=j+1}^k (1 + \beta_\ell + \gamma_\ell^2 L^2) .$$

By convention, $\prod_{\ell=k+1}^k a_\ell = 1$.

Proof.

$$\mathbb{E} [\|\bar{s}_{J_1} \circ \mathsf{T}(\widehat{S}^0) - \mathbf{S}_{1,J_1}\|^2] = \frac{1}{n} \sum_{i=1}^n \mathbb{E} [\|\bar{s}_i \circ \mathsf{T}(\widehat{S}^0) - \mathbf{S}_{1,i}\|^2] = 0 .$$

Let $k \geq 1$. We write (see (10))

$$\begin{aligned} S_{k+1,i} &= S_{k,i}1_{I_{k+1} \neq i} + \bar{s}_i \circ \mathsf{T}(\widehat{S}^k)1_{I_{k+1}=i} \\ &= \bar{s}_i \circ \mathsf{T}(\widehat{S}^{\langle k,i \rangle})1_{I_{k+1} \neq i} + \bar{s}_i \circ \mathsf{T}(\widehat{S}^k)1_{I_{k+1}=i}, \end{aligned}$$

where $\widehat{S}^{\langle \ell,i \rangle}$ is defined by (25). This yields, by A5-(ii)

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\|\bar{s}_i \circ \mathsf{T}(\widehat{S}^k) - S_{k+1,i}\|^2 \right] &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\|\bar{s}_i \circ \mathsf{T}(\widehat{S}^k) - \bar{s}_i \circ \mathsf{T}(\widehat{S}^{\langle k,i \rangle})\|^2 1_{I_{k+1} \neq i} \right] \\ &\leq \Delta_k \stackrel{\text{def}}{=} \frac{n-1}{n^2} \sum_{i=1}^n L_i^2 \mathbb{E} \left[\|\widehat{S}^k - \widehat{S}^{\langle k,i \rangle}\|^2 \right]. \end{aligned} \quad (40)$$

We have

$$\Delta_k = \frac{n-1}{n^2} \sum_{i=1}^n L_i^2 \mathbb{E} \left[\|\widehat{S}^k - \widehat{S}^{k-1} + (\widehat{S}^{k-1} - \widehat{S}^{\langle k-1,i \rangle}) 1_{I_k \neq i}\|^2 \right]$$

where we used in the last inequality that

$$\widehat{S}^{\langle k,i \rangle} = \widehat{S}^{k-1} 1_{I_k=i} + \widehat{S}^{\langle k-1,i \rangle} 1_{I_k \neq i}.$$

Upon noting that $2\langle \tilde{U}, V \rangle \leq \beta^{-1} \|\tilde{U}\|^2 + \beta \|V\|^2$ for any $\beta > 0$, we have for any \mathcal{G} -measurable r.v. V

$$\mathbb{E} [\|U + V\|^2] \leq \mathbb{E} [\|U\|^2] + \beta^{-1} \mathbb{E} [\|\mathbb{E}[U|\mathcal{G}]\|^2] + (1 + \beta) \mathbb{E} [\|V\|^2].$$

Applying this inequality with $\beta \leftarrow \beta_k$, $U \leftarrow \widehat{S}^k - \widehat{S}^{k-1} = \gamma_k H_k$ and $\mathcal{G} \leftarrow \mathcal{F}_{k-1/2}$ yields

$$\begin{aligned} \Delta_k &\leq \gamma_k^2 \frac{n-1}{n} L^2 \mathbb{E} [\|H_k\|^2] + \frac{\gamma_k^2}{\beta_k} \frac{n-1}{n} L^2 \mathbb{E} [\|\mathbb{E}[H_k|\mathcal{F}_{k-1/2}]\|^2] \\ &\quad + (1 + \beta_k) \frac{n-1}{n^2} \sum_{i=1}^n L_i^2 \mathbb{E} \left[\|\widehat{S}^{k-1} - \widehat{S}^{\langle k-1,i \rangle}\|^2 1_{I_k \neq i} \right]. \end{aligned}$$

By Lemma 8 and (40), we have

$$\mathbb{E} [\|H_k\|^2] \leq \mathbb{E} [\|h(\widehat{S}^{k-1})\|^2] + \Delta_{k-1} - \mathbb{E} \left[\|\widehat{S}^k - \bar{s} \circ \mathsf{T}(\widehat{S}^{k-1})\|^2 \right];$$

for the second term, we use again $\mathbb{E} [H_k|\mathcal{F}_{k-1/2}] = h(\widehat{S}^{k-1})$; for the third term, since $I_k \in \mathcal{F}_{k-1/2}$, $\widehat{S}^{k-1} \in \mathcal{F}_{k-1}$, $\widehat{S}^{\langle k-1,i \rangle} \in \mathcal{F}_{k-1}$, then

$$\sum_{i=1}^n L_i^2 \mathbb{E} \left[\|\widehat{S}^{k-1} - \widehat{S}^{\langle k-1,i \rangle}\|^2 1_{I_k \neq i} \right] = n \Delta_{k-1}.$$

Therefore, we established

$$\begin{aligned} \Delta_k \leq & (1 + \beta_k + \gamma_k^2 L^2) \frac{n-1}{n} \Delta_{k-1} \\ & + \gamma_k^2 \left(1 + \frac{1}{\beta_k}\right) L^2 \frac{n-1}{n} \mathbb{E} \left[\|h(\widehat{S}^{k-1})\|^2 \right] \\ & - \gamma_k^2 L^2 \frac{n-1}{n} \mathbb{E} \left[\|\widetilde{S}^k - \bar{s} \circ \mathsf{T}(\widehat{S}^{k-1})\|^2 \right]. \end{aligned}$$

The proof is then concluded by standard algebra upon noting that $\Delta_0 = 0$. \square

References

- Agarwal A, Bottou L (2015) A lower bound for the optimization of finite sums. In: Bach F, Blei D (eds) Proceedings of the 32nd International Conference on Machine Learning, PMLR, Proceedings of Machine Learning Research, vol 37, pp 78–86
- Allen-Zhu Z, Hazan E (2016) Variance reduction for faster non-convex optimization. In: Balcan M, Weinberger K (eds) Proceedings of The 33rd International Conference on Machine Learning, PMLR, Proceedings of Machine Learning Research, vol 48, pp 699–707
- Benveniste A, Priouret P, Métivier M (1990) Adaptive Algorithms and Stochastic Approximations. Springer-Verlag, Berlin, Heidelberg
- Borkar V (2008) Stochastic Approximation: A Dynamical Systems Viewpoint. Cambridge University Press
- Brown LD (1986) Fundamentals of statistical exponential families with applications in statistical decision theory, Institute of Mathematical Statistics Lecture Notes—Monograph Series, vol 9. Institute of Mathematical Statistics, Hayward, CA
- Cappé O, Moulines E (2009) On-line Expectation Maximization algorithm for latent data models. *J Roy Stat Soc B Met* 71(3):593–613
- Celeux G, Diebolt J (1985) The SEM algorithm: a probabilistic teacher algorithm derived from the EM algorithm for the mixture problem. *Computational Statistics Quarterly* 2:73–82
- Chen J, Zhu J, Teh Y, Zhang T (2018) Stochastic Expectation Maximization with Variance Reduction. In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R (eds) Advances in Neural Information Processing Systems 31, Curran Associates, Inc., pp 7967–7977

- Csiszár I, Tusnády G (1984) Information geometry and alternating minimization procedures. In: Recent results in estimation theory and related topics, suppl. 1, Statist. Decisions, pp 205–237
- Defazio A, Bach F, Lacoste-Julien S (2014) SAGA: A Fast Incremental Gradient Method With Support for Non-Strongly Convex Composite Objectives. In: Ghahramani Z, Welling M, Cortes C, Lawrence ND, Weinberger KQ (eds) Advances in Neural Information Processing Systems 27, Curran Associates, Inc., pp 1646–1654
- Delyon B, Lavielle M, Moulines E (1999) Convergence of a Stochastic Approximation version of the EM algorithm. *Ann Statist* 27(1):94–128
- Dempster A, Laird N, Rubin D (1977) Maximum Likelihood from Incomplete Data via the EM Algorithm. *J Roy Stat Soc B Met* 39(1):1–38
- Fang C, Li C, Lin Z, Zhang T (2018) SPIDER: Near-Optimal Non-Convex Optimization via Stochastic Path-Integrated Differential Estimator. In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R (eds) Advances in Neural Information Processing Systems 31, Curran Associates, Inc., pp 689–699
- Fort G, Moulines E (2003) Convergence of the Monte Carlo Expectation Maximization for curved exponential families. *Ann Statist* 31(4):1220–1259
- Frühwirth-Schnatter S, Celeux G, Robert CP (eds) (2019) Handbook of mixture analysis. Chapman & Hall/CRC Handbooks of Modern Statistical Methods, CRC Press, Boca Raton, FL
- Ghadimi S, Lan G (2013) Stochastic First- and Zeroth-Order Methods for Nonconvex Stochastic Programming. *SIAM J Optim* 23(4):2341–2368
- Glasserman P (2004) Monte Carlo methods in financial engineering. Springer, New York
- Gunawardana A, Byrne W (2005) Convergence theorems for generalized alternating minimization procedures. *J Mach Learn Res* 6:2049–2073
- Johnson R, Zhang T (2013) Accelerating Stochastic Gradient Descent using Predictive Variance Reduction. In: Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ (eds) Advances in Neural Information Processing Systems 26, Curran Associates, Inc., pp 315–323
- Karimi B, Miasojedow B, Moulines E, Wai HT (2019a) Non-asymptotic Analysis of Biased Stochastic Approximation Scheme. In: Beygelzimer A, Hsu D (eds) Proceedings of the Thirty-Second Conference on Learning Theory, PMLR, Phoenix, USA, Proceedings of Machine Learning Research, vol 99, pp 1944–1974

- Karimi B, Wai HT, Moulines E, Lavielle M (2019b) On the Global Convergence of (Fast) Incremental Expectation Maximization Methods. In: Wallach H, Larochelle H, Beygelzimer A, d'Alché Buc F, Fox E, Garnett R (eds) *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., pp 2837–2847
- Kwedlo W (2015) A new random approach for initialization of the multiple restart EM algorithm for Gaussian model-based clustering. *Pattern Anal Applic* 18:757–770
- Lange K (1995) A Gradient Algorithm Locally Equivalent to the EM Algorithm. *J R Statist Soc B* 57(2):425–437
- Lange K (2016) *MM Optimization Algorithms*. Other Titles in Applied Mathematics, Society for Industrial and Applied Mathematics
- Little RJA, Rubin D (2002) *Statistical analysis with missing data*, 2nd edn. Wiley Series in Probability and Statistics, Wiley-Interscience [John Wiley & Sons], Hoboken, NJ
- McLachlan G, Krishnan T (2008) *The EM algorithm and extensions*. Wiley series in probability and statistics, Wiley
- Murty K, Kabadi S (1987) Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming* 39:117–129
- Neal RM, Hinton GE (1998) A View of the EM Algorithm that Justifies Incremental, Sparse, and other Variants. In: Jordan MI (ed) *Learning in Graphical Models*, Springer Netherlands, Dordrecht, pp 355–368
- Ng SK, McLachlan GJ (2003) On the choice of the number of blocks with the incremental EM algorithm for the fitting of normal mixtures. *Stat Comput* 13(1):45–55
- Nguyen H, Forbes F, McLachlan G (2020) Mini-batch learning of exponential family finite mixture models. *Stat Comput* 30:731–748
- Parizi SN, He K, Aghajani R, Sclaroff S, Felzenszwalb P (2019) Generalized Majorization-Minimization. In: C K, S R (eds) *Proceedings of the 36th International Conference on Machine Learning*, PMLR, Long Beach, California, USA, *Proceedings of Machine Learning Research*, vol 97, pp 5022–5031
- Reddi S, Sra S, Póczos B, Smola A (2016) Fast Incremental Method for Smooth Nonconvex Optimization. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp 1971–1977
- Robbins H, Monro S (1951) A stochastic approximation method. *Ann Math Statist* 22(3):400–407

- Schmidt M, Le Roux N, Bach F (2017) Minimizing finite sums with the stochastic average gradient. *Mathematical Programming* 162(1-2):83–112
- Sundberg R (2019) *Statistical Modelling by Exponential Families*. Cambridge University Press
- Wei G, Tanner M (1990) A Monte Carlo Implementation of the EM Algorithm and the Poor Man’s Data Augmentation Algorithms. *J Am Stat Assoc* 85(411):699–704
- Wu C (1983) On the Convergence Properties of the EM Algorithm. *Ann Statist* 11(1):95–103
- Zangwill WI (1967) Non-linear programming via penalty functions. *Management Sci* 13:344–358
- Zhou D, Xu P, Gu Q (2018) Stochastic nested variance reduced gradient descent for nonconvex optimization. In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R (eds) *Advances in Neural Information Processing Systems* 31, Curran Associates, Inc., pp 3921–3932




Figure 8: For $k \in \{1e2, 5e2, 1e3, 1.5e3, \dots, 6e3, 7e3, \dots, 20e3\}$, ratio of the expectations (Exp) $\mathbb{E} [\|\theta_{\text{opt-FIEM}}^k - \theta_\star\|] / \mathbb{E} [\|\theta_{\text{Alg}}^k - \theta_\star\|]$ when Alg is FIEM (solid line with circle) and then **Online EM** (solid line with cross); and the standard deviations (std) $\text{std}(\|\theta_{\text{opt-FIEM}}^k - \theta_\star\|) / \text{std}(\|\theta_{\text{Alg}}^k - \theta_\star\|)$ when Alg is FIEM (dashed line with circle) and then **Online EM** (dashed line with cross). The expectations and standard deviations are approximated by a Monte Carlo sum over $1e3$ independent runs.




Figure 9: Monte Carlo approximation (over $1e3$ independent runs) of $k \mapsto \gamma_{k+1}^{-2} \mathbb{E} \left[\|\widehat{S}^{k+1} - \widehat{S}^k\|^2 \right]$ for Online EM, FIEM and opt-FIEM.

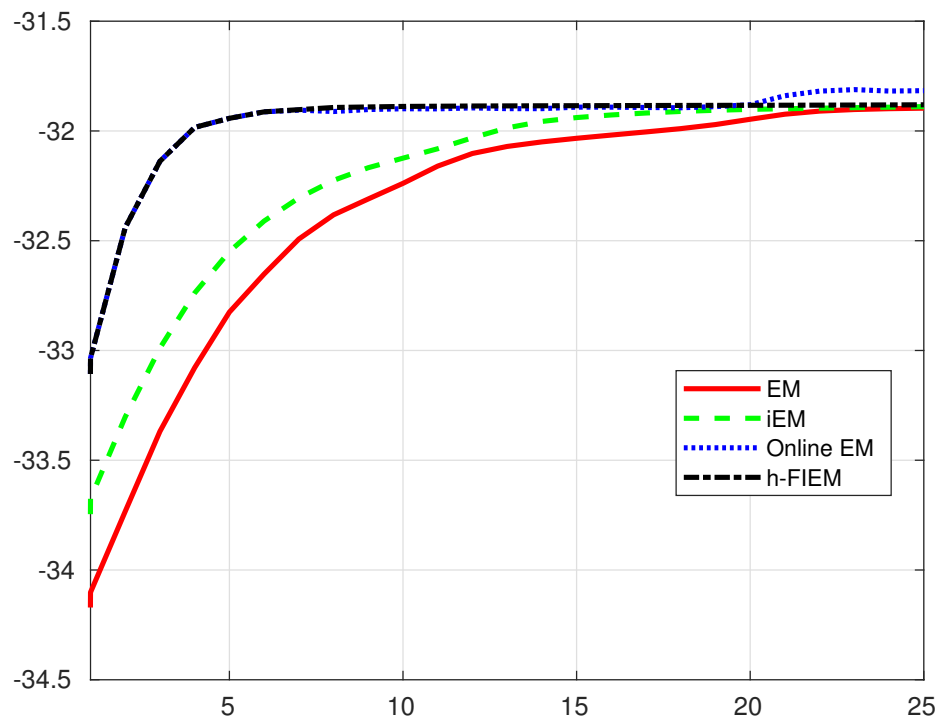


Figure 10: Evolution of the normalized log-likelihood along one path of length 100 epochs: only the epochs 1 to 25 are displayed. All the paths start from the same value at time $t = 0$, with a normalized log-likelihood equal to -58.31 .




Figure 11: Evolution of the normalized log-likelihood along one path of length 100 epochs: the first 14 epochs are discarded. All the paths start from the same value at time $t = 0$, with a normalized log-likelihood equal to -58.31 .

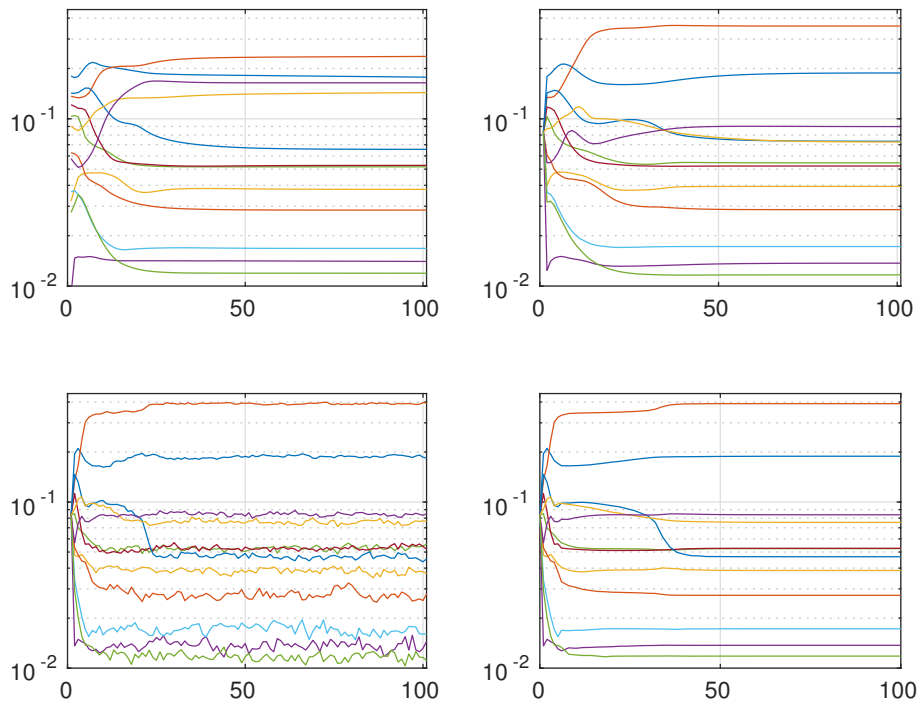


Figure 12: Evolution of the $g = 12$ weights along one path of length 100 epochs. All the paths start from the same value at time $t = 0$. EM (top left), iEM (top right), Online EM (bottom left) and h-FIEM (bottom right).

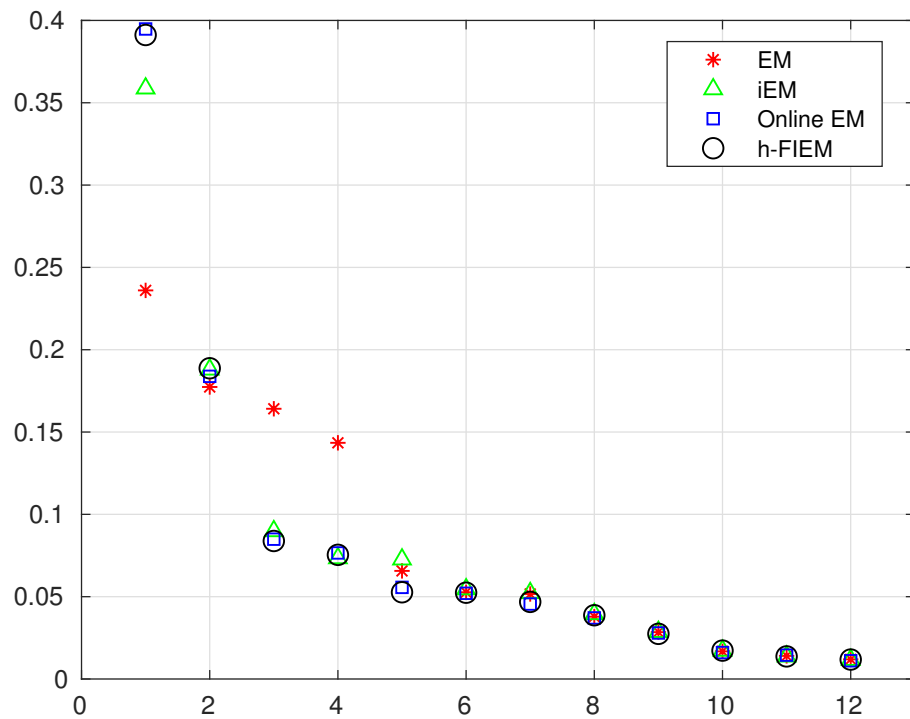


Figure 13: Estimation of the $g = 12$ weights of the mixture model. The estimator is the value of the parameter obtained at the end of a single path of length 100 epochs.

Supplementary material to "Fast Incremental Expectation Maximization for finite-sum optimization: nonasymptotic convergence"

This *supplementary material* provides

1. proofs of some comments.
2. details and additional analyses for the numerical illustration on Gaussian Mixture Models (section 5).

Notations. Vectors are column vectors. For $a, b \in \mathbb{R}^d$, $\langle a, b \rangle = a^T b$ is the Euclidean scalar product; $\langle a, b \rangle$ denotes the standard Euclidean scalar product on \mathbb{R}^ℓ , for $\ell \geq 1$; and $\|a\|$ the associated norm. For a matrix A , A^T is its transpose.

For a non negative integer n , $[n] \stackrel{\text{def}}{=} \{0, \dots, n\}$ and $[n]^\star \stackrel{\text{def}}{=} \{1, \dots, n\}$. $a \wedge b$ is the minimum of two real numbers a, b .

For two $p \times p$ matrices A, B , $\langle A, B \rangle$ is the trace of $B^T A$: $\langle A, B \rangle \stackrel{\text{def}}{=} \text{Tr}(B^T A)$. I_p stands for the $p \times p$ identity matrix. \otimes stands for the Kronecker product. \mathcal{M}_p^+ denotes the set of the invertible $p \times p$ covariance matrices. $\det(A)$ is the determinant of the matrix A .

$\mathcal{N}_p(\mu, \Gamma)$ denotes a \mathbb{R}^p -valued Gaussian distribution, with expectation μ and covariance matrix Γ .

7 EM as a Majorize-Minimization algorithm

The following result shows that $\{\bar{F}(\cdot, \theta'), \theta' \in \Theta\}$ is a family of majorizing function of the objective function F from which a Majorize-Minimization approach for solving (2) can be derived under A3. This MM algorithm is EM (see item 3).

Proposition 10. *Assume A1 and A2.*

1. For any $i \in [n]^\star$ and $\theta' \in \Theta$ we have

$$\mathcal{L}_i(\cdot) \leq -\langle \bar{s}_i(\theta'), \phi(\cdot) \rangle + C_i(\theta') .$$

2. For any $\theta' \in \Theta$, we have $F \leq \bar{F}(\cdot, \theta')$ and $\bar{F}(\theta', \theta') = F(\theta')$.

3. Assume also A3. Given $\theta^0 \in \Theta$, the sequence defined by $\theta^{k+1} \stackrel{\text{def}}{=} \top \circ \bar{s}(\theta^k)$ for any $k \geq 0$, satisfies $F(\theta^{k+1}) \leq F(\theta^k)$.

Proof. (proof of item 1). From the Jensen's inequality, it holds

$$\begin{aligned} \mathcal{L}_i(\theta) - \mathcal{L}_i(\theta') &\leq - \int_{\mathcal{Z}} \langle s_i(z), \phi(\theta) - \phi(\theta') \rangle p_i(z; \theta') \mu(dz) \\ &= - \langle \bar{s}_i(\theta'), \phi(\theta) - \phi(\theta') \rangle ; \end{aligned}$$

which concludes the proof.

(proof of item 2) From (2) and item 1, it holds

$$F(\theta) \leq - \langle \bar{s}(\theta'), \phi(\theta) \rangle + \frac{1}{n} \sum_{i=1}^n \mathcal{C}_i(\theta') + R(\theta) .$$

(proof of item 3) From item 2 and the definition of \mathbb{T} , it holds

$$F(\mathbb{T} \circ \bar{s}(\theta^k)) \leq \bar{F}(\mathbb{T} \circ \bar{s}(\theta^k), \theta^k) \leq \bar{F}(\theta^k, \theta^k) = F(\theta^k) .$$

□

8 Proof of the comments in section 3

8.1 Comments in Section subsection 3.2

• **The choice** $\lambda = C$. Since $n \geq 2$, the second condition in (29) is satisfied with $\lambda = C$. (30) is a decreasing function of C so that by the first condition in (29), C solves

$$\sqrt{C} \left(\frac{1}{n^{2/3}} + \frac{1}{1 - n^{-1/3}} \left(\frac{1}{n} + \frac{1}{1 - C} \right) \right) = 2\mu v_{\min} \frac{L}{L_{\dot{V}}} .$$

A solution exists in $(0, 1)$ and is unique (see 12); it is denoted by C^* . Since the LHS is lower bounded by $C \mapsto \sqrt{C}(1 - C)^{-1}$ on $(0, 1)$, C^* is upper bounded by $C^+ \in (0, 1)$ solving

$$\sqrt{C} = 2\mu v_{\min} \frac{L}{L_{\dot{V}}} (1 - C) .$$

This yields $C^+ = (\sqrt{1 + 4A^2} - 1)/(2A)$ with $A \stackrel{\text{def}}{=} 2\mu v_{\min} L/L_{\dot{V}}$. Note that $f_n(C^*, C^*) \leq f_2(C^*, C^*) \leq f_2(C^+, C^+)$; for the second inequality, 12 is used again.

• **Another choice, for any n large enough.** When $n \rightarrow \infty$, we have

$$\mathcal{L}_n(C, \lambda) \uparrow \mathcal{L}_{\infty}(C, \lambda) \stackrel{\text{def}}{=} \frac{L_{\dot{V}} \sqrt{C}}{2L^2} \left(v_{\min} \frac{2L}{L_{\dot{V}}} - \frac{C^{3/2}}{\lambda} \frac{1}{1 - \lambda} \right) .$$

By 13 applied with $A \leftarrow v_{\min}/L$ and $B \leftarrow 2L_{\dot{V}}/L^2$, we have $\mathcal{L}_\infty(C, \lambda) \leq \mathcal{L}_\infty(C_\star, \lambda_\star)$ where

$$\lambda_\star \stackrel{\text{def}}{=} \frac{1}{2}, \quad C_\star \stackrel{\text{def}}{=} \frac{1}{4} \left(\frac{v_{\min} L}{L_{\dot{V}}} \right)^{2/3},$$

$$\mathcal{L}_\infty(C_\star, \lambda_\star) = \frac{3 v_{\min}}{8 L} \left(\frac{v_{\min} L}{L_{\dot{V}}} \right)^{1/3}.$$

In the proof of 4, we established that for any $\lambda \in (0, 1)$ and $C > 0$ such that $\lambda - Cn^{-1/3} \in (0, 1)$, we have

$$n^{2/3} \alpha_k(n, C, \lambda) \geq n^{2/3} \underline{\alpha}_n(C, \lambda) \geq \mathcal{L}_n(C, \lambda).$$

Set $\tilde{N}_\star \stackrel{\text{def}}{=} (v_{\min} L / L_{\dot{V}})^2 / 8$; for any $n \geq \tilde{N}_\star$, we have $\lambda_\star - C_\star n^{-1/3} \in (0, 1)$ so that

$$n^{2/3} \alpha_k(n, C_\star, \lambda_\star) \geq n^{2/3} \underline{\alpha}_n(C_\star, \lambda_\star) \geq \mathcal{L}_n(C_\star, \lambda_\star).$$

This implies that for any $k \in [K_{\max} - 1]$,

$$\lim_n n^{2/3} \alpha_k(n, C_\star, \lambda_\star) \geq \lim_n n^{2/3} \underline{\alpha}_n(C_\star, \lambda_\star) \geq \mathcal{L}_\infty(C_\star, \lambda_\star) > 0,$$

thus showing that for any n large enough - let us say $n \geq N_\star$ (with N_\star which only depends upon $L, L_{\dot{V}}, v_{\min}$), we have for any $k \in [K_{\max} - 1]$,

$$n^{2/3} \alpha_k(n, C_\star, \lambda_\star) \geq \mathcal{L}_\infty(C_\star, \lambda_\star) > 0.$$

Therefore, we first write

$$\mathbf{E}_1 \leq \frac{1}{K_{\max}} \frac{n^{2/3}}{\mathcal{L}_\infty(C_\star, \lambda_\star)} \sum_{k=0}^{K_{\max}-1} \alpha_k(n, C_\star, \lambda_\star) \mathbb{E} \left[\|h(\hat{S}_k)\|^2 \right];$$

we then write, by using $\delta_k \geq \gamma_{k+1}^2 L_{\dot{V}} / 2$ and $\gamma_{k+1}^2 = C_\star / (L^2 n^{4/3})$,

$$\begin{aligned} \frac{1}{3n^{2/3}} \left(\frac{L_{\dot{V}}}{L v_{\min}} \right)^{2/3} \mathbf{E}_2 &= \frac{n^{2/3} L_{\dot{V}}}{2 \mathcal{L}_\infty(C_\star, \lambda_\star)} \gamma_{k+1}^2 \mathbf{E}_2 \\ &\leq \frac{1}{K_{\max}} \frac{n^{2/3}}{\mathcal{L}_\infty(C_\star, \lambda_\star)} \sum_{k=0}^{K_{\max}-1} \delta_k(n, C_\star, \lambda_\star) \mathbb{E} \left[\|\tilde{S}^{k+1} - \bar{s} \circ \mathsf{T}(\hat{S}^k)\|^2 \right] \end{aligned}$$

from which we obtain

$$\mathbf{E}_1 + \frac{1}{3n^{2/3}} \left(\frac{L_{\dot{V}}}{L v_{\min}} \right)^{2/3} \mathbf{E}_2 \leq \frac{1}{K_{\max}} \frac{n^{2/3}}{\mathcal{L}_\infty(C_\star, \lambda_\star)} \Delta V.$$

This concludes the proof.

8.2 Comments in Section subsection 3.3

Complexity. For $\tau > 0$, set $C = \lambda\tau$. Then for any $\lambda \in (0, 1)$,

$$\sqrt{\lambda\tau}\tilde{f}_n(\lambda\tau, \lambda) = \frac{\sqrt{\lambda}\sqrt{\tau}}{(nK_{\max})^{1/3}} + \lambda^{3/2}\tau^{3/2} \left(\frac{1}{n} + \frac{1}{1-\lambda} \right),$$

which is a continuous increasing function of λ , which tends to zero when $\lambda \rightarrow 0$ and to $+\infty$ when $\lambda \rightarrow 1$. Hence, there exists a unique $\lambda_\star \in (0, 1)$, depending upon $L, L_{\dot{V}}, v_{\min}, \tau, \mu$ and n, K_{\max} such that

$$\sqrt{\lambda_\star\tau}\tilde{f}_n(\lambda_\star\tau, \lambda_\star) = 2\mu v_{\min}L/L_{\dot{V}}.$$

Note however that since $\sqrt{\lambda\tau}\tilde{f}_n(\lambda\tau, \lambda) \geq \lambda^{3/2}\tau^{3/2}/(1-\lambda)$ for any $\lambda \in (0, 1)$, then λ_\star is upper bounded by the unique solution $\lambda^+ \in (0, 1)$ satisfying $L_{\dot{V}}\lambda^{3/2}\tau^{3/2}/(2L(1-\lambda)) = \mu v_{\min}$ (see 14). Such a solution λ^+ only depends upon $L, L_{\dot{V}}, v_{\min}, \tau, \mu$. Hence, for any $\tau > 0$,

$$\tilde{f}_n(\lambda\tau, \lambda) \leq \sup_{n, K_{\max}} \tilde{f}_n(\lambda^+(\tau)\tau, \lambda^+(\tau))$$

and the RHS does not depend on n, K_{\max} . This inequality implies that

$$\frac{n^{1/3}}{K_{\max}^{2/3}} \frac{L\tilde{f}_n(\lambda\tau, \lambda)}{\mu(1-\mu)v_{\min}^2} \leq \frac{n^{1/3}}{K_{\max}^{2/3}} M \quad M \stackrel{\text{def}}{=} \frac{L}{\mu(1-\mu)v_{\min}^2} \sup_{n, K_{\max}} \tilde{f}_n(\lambda^+(\tau)\tau, \lambda^+(\tau)).$$

Hence, there exists $M > 0$ depending upon $L, L_{\dot{V}}, v_{\min}, \tau, \mu$ such that for any $\varepsilon > 0$,

$$K_{\max} \geq \left(\tau^{3/2}\sqrt{n} \right) \vee \left(M\sqrt{n}\varepsilon^{-3/2} \right) \implies \frac{n^{1/3}}{K_{\max}^{2/3}} \frac{L\tilde{f}_n(\lambda\tau, \lambda)}{\mu(1-\mu)v_{\min}^2} \leq \varepsilon.$$

Another choice of (λ, C) , for any n large enough. In this section, we consider that there exists $\tau > 0$ such that $\sup_{n, K_{\max}} n^{1/3}K_{\max}^{-2/3} \leq \tau$, that $n \rightarrow \infty$ and that $nK_{\max} \rightarrow \infty$. In this asymptotic, we have $\mathcal{L}_n(C, \lambda) \uparrow \mathcal{L}_\infty(C, \lambda)$ where

$$\mathcal{L}_\infty(C, \lambda) \stackrel{\text{def}}{=} \frac{\sqrt{C}}{L} \left(v_{\min} - \frac{L_{\dot{V}}}{2L} \frac{C^{3/2}}{1-\lambda} \right).$$

For any $(C, \lambda) \in \mathbb{R}^+ \times (0, 1)$ s.t. $\tau \leq \lambda/C$, we have $\mathcal{L}_\infty(C, \lambda) \leq \mathcal{L}_\infty(C_\star(\lambda), \lambda)$ where

$$C_\star(\lambda) \stackrel{\text{def}}{=} \left(\frac{v_{\min}L}{2L_{\dot{V}}} \right)^{2/3} (1-\lambda)^{2/3};$$

see 13. The condition $C\tau \leq \lambda$ implies that this inequality holds for any $\lambda \in [\lambda_*, 1)$ where λ_* is the unique solution of (see 14)

$$\left(\frac{v_{\min}L}{2L\dot{V}}\right)^2 (1 - \lambda_*)^2 = \lambda_*^3/\tau^3.$$

Since $\mathcal{L}_\infty(C_*(\lambda), \lambda) = \frac{3}{4} \left(\frac{v_{\min}^4}{2L^2L\dot{V}}\right)^{1/3} (1 - \lambda)^{1/3}$, this quantity is maximal by choosing $\lambda = \lambda_*$. Therefore, for any $(C, \lambda) \in \mathbb{R}^+ \times (0, 1)$, s.t. $\tau \leq \lambda/C$, we have

$$\lim_n n^{1/3} K_{\max}^{1/3} \underline{\alpha}_n(C_*(\lambda_*), \lambda_*) = \mathcal{L}_\infty(C_*(\lambda_*), \lambda_*) > 0.$$

For any n large enough (with a bound which only depends upon $L, L\dot{V}, v_{\min}, \tau$), we have

$$\frac{1}{K_{\max} \alpha_*(C_*, \lambda_*)} = \frac{n^{1/3}}{K_{\max}^{2/3}} \frac{4}{3} \left(\frac{2L^2L\dot{V}}{v_{\min}^4}\right)^{1/3} (1 - \lambda_*)^{-1/3}.$$

First, we write

$$\mathbb{E}_1 \leq \frac{n^{1/3}}{K_{\max}^{2/3} \mathcal{L}_\infty(C_*(\lambda_*), \lambda_*)} \sum_{k=0}^{K_{\max}-1} \alpha_k(n, C_*(\lambda_*), \lambda_*) \mathbb{E} \left[\|h(\widehat{S}^k)\|^2 \right].$$

Then we write that by using $\delta_k \geq \gamma_{k+1}^2 L\dot{V}/2$ and $\gamma_{k+1}^2 = C_*(\lambda_*)/(L^2 n^{2/3} K_{\max}^{2/3})$ that

$$\begin{aligned} & \frac{2^{10/3}(1 - \lambda_*)^{-1/3} \mu^2}{\tilde{f}_n^2(\lambda_* \tau, \lambda_*)} \left(\frac{Lv_{\min}}{L\dot{V}}\right)^{2/3} \frac{1}{(nK_{\max})^{1/3}} \mathbb{E}_2 \\ &= \frac{L\dot{V}}{2 \mathcal{L}_\infty(C_*(\lambda_*), \lambda_*)} \gamma_{k+1}^2 n^{1/3} K_{\max}^{1/3} \mathbb{E}_2 \\ &\leq \frac{n^{1/3} K_{\max}^{1/3}}{K_{\max} \mathcal{L}_\infty(C_*(\lambda_*), \lambda_*)} \sum_{k=0}^{K_{\max}-1} \delta_k(n, C_*(\lambda_*), \lambda_*) \mathbb{E} \left[\|\tilde{S}^{k+1} - \bar{s} \circ \Upsilon(\widehat{S}^k)\|^2 \right] \end{aligned}$$

We then conclude that

$$\begin{aligned} \mathbb{E}_1 + \frac{2^{10/3}(1 - \lambda_*)^{-1/3} \mu^2}{\tilde{f}_n^2(\lambda_* \tau, \lambda_*)} \left(\frac{Lv_{\min}}{L\dot{V}}\right)^{2/3} \frac{1}{(nK_{\max})^{1/3}} \mathbb{E}_2 \\ \leq \frac{n^{1/3}}{K_{\max}^{2/3} \mathcal{L}_\infty(C_*(\lambda_*), \lambda_*)} \Delta V. \end{aligned}$$

8.3 Comments in Section subsection 3.4

Case $\lambda = C$. A simple strategy is to choose $n \geq 2$ and $C_{\max} = \lambda$ solution of $v_{\min}/2 = \sqrt{C}f_n(C, C)$. This solution exists and is unique, and it is upper bounded by a quantity C^+ which depends only on $L, L_{\dot{V}}, v_{\min}$ - see subsection 8.1 for a similar discussion.

Case $\lambda = 1/2$. $f_n(C, \lambda)$ controls the errors E_i and we can choose $\lambda \in (0, 1)$ and then $C > 0$ such that this quantity is minimal; to make the computations easier, we minimize w.r.t. λ the function $\lim_n f_n(C, \lambda)$: it behaves like $\lambda^{-1}(1 - \lambda)^{-1}$ so that we set $\lambda = 1/2$. The equation $\sqrt{C}f_n(C, 1/2) = v_{\min}L/L_{\dot{V}}$ possesses an unique solution C_n in $(0, n^{1/3}/2)$.

Upon noting that $x \mapsto \sqrt{x}f_n(x, 1/2)$ is lower bounded by $x \mapsto 4x^{3/2}$, C_n satisfies

$$C_n \leq \left(\frac{v_{\min}L}{4L_{\dot{V}}} \right)^{2/3},$$

thus showing that the constraint $n^{-1/3} < \lambda/C_n = 1/(2C_n)$ is satisfied for any n such that $8n > (v_{\min}L/L_{\dot{V}})^2$.

9 Technical Lemmas

Lemma 11. *Let $A, B, v > 0$ and define $F(x) \stackrel{\text{def}}{=} Ax(v - Bx)$ on \mathbb{R} . Then the roots of F are $\{0, v/B\}$; F is positive on $(0, v/B)$; the maximal value of F is $Av^2/(4B)$ and it is reached at $x_{\star} \stackrel{\text{def}}{=} v/2B$.*

Lemma 12. *Let $a, b > 0$ and define F on $(0, 1)$ by $F(x) = \sqrt{x}(a + b/(1 - x))$. F is increasing on $(0, 1)$ and for any $v > 0$, there exists an unique $x \in (0, 1)$ such that $F(x) = v$.*

Proof. $x \mapsto F(x)$ is continuous and increasing on $(0, 1)$, tends to zero when $x \rightarrow 0$ and to $+\infty$ when $x \rightarrow 1$; therefore for any $v > 0$, there exists an unique $x \in (0, 1)$ such that $F(x) = v$. \square

Lemma 13. *Let $A, B > 0$. The function $F : x \mapsto Ax - Bx^4$ defined on $(0, \infty)$ reaches its unique maximum at $x_{\star} \stackrel{\text{def}}{=} A^{1/3}B^{-1/3}4^{-1/3}$ and $F(x_{\star}) = 3A^{4/3}/(B^{4/3})^{1/3}$.*

Proof. $F'(x) = A - 4Bx^3$ and $F''(x) = -12Bx^2 < 0$; hence, F' is decreasing. $F'(x) = 0$ iff $x^3 = A/(4B)$, showing $F' > 0$ on $(0, x_{\star})$ with $x_{\star} \stackrel{\text{def}}{=} A^{1/3}/(4B)^{1/3}$. Hence, F is increasing on $[0, x_{\star}]$ and then decreasing. \square

Lemma 14. For any $v > 0$, the function $x \mapsto (1-x)^2/x^3$ is decreasing on $(0, 1)$ and there exists a unique $x \in (0, 1)$ solving $(1-x)^2/x^3 = v$.

Proof. The derivative of $x \mapsto (1-x)^2/x^3$ is $-x^{-4}(x-3)(x-1)$ thus showing that the function is decreasing on $(0, 1)$; it tends to $+\infty$ when $x \rightarrow 0$ and to 0 when $x \rightarrow 1$. This concludes the proof. \square

10 Example: Mixture of multivariate Gaussian distributions.

Set

$$f(y) \stackrel{\text{def}}{=} \sum_{\ell=1}^g \alpha_{\ell} \mathcal{N}_p(\mu_{\ell}, \Sigma)[y]; \quad \Gamma \stackrel{\text{def}}{=} \Sigma^{-1}. \quad (41)$$

We write, up to the multiplicative constant $\sqrt{2\pi}^{-p}$,

$$\begin{aligned} f(y) &= \sum_{z=1}^g \alpha_z \sqrt{\det(\Gamma)} \exp\left(-\frac{1}{2}(y - \mu_z)^T \Gamma (y - \mu_z)\right) \\ &= \sqrt{\det(\Gamma)} \exp\left(-\frac{1}{2}y^T \Gamma y\right) \sum_{z=1}^g \exp\left(\sum_{\ell=1}^g 1_{z=\ell} \left\{ \ln \alpha_{\ell} - \frac{1}{2}\mu_{\ell}^T \Gamma \mu_{\ell} + \mu_{\ell}^T \Gamma y \right\}\right); \end{aligned}$$

Parametric statistical model. Set

$$\theta \stackrel{\text{def}}{=} (\alpha_1, \dots, \alpha_g, \mu_1, \dots, \mu_g, \Sigma),$$

and denote by \mathcal{M}_p^+ the set of the $p \times p$ positive definite matrices. Then we set

$$\Theta \stackrel{\text{def}}{=} \{\alpha_{\ell} \geq 0, \sum_{\ell=1}^g \alpha_{\ell} = 1\} \times (\mathbb{R}^p)^g \times \mathcal{M}_p^+.$$

Latent variable model in the exponential family. The density (41) is of the form

$$\exp(-R_y(\theta)) \sum_{z=1}^g \exp(\langle s_y(z), \phi(\theta) \rangle)$$

with

$$R_y(\theta) \stackrel{\text{def}}{=} \frac{1}{2}y^T \Gamma y - \frac{1}{2} \ln \det(\Gamma)$$

and $s_y(z) \stackrel{\text{def}}{=} \mathbf{A}_y \rho(z) \in \mathbb{R}^{g(1+p)}$ and $\phi = (\phi^{(1)}, \phi^{(2)})$

$$\mathbf{A}_y \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{I}_g \\ \mathbf{I}_g \otimes y \end{bmatrix} \in \mathbb{R}^{g(1+p) \times g} \quad (42)$$

$$\rho(z) \stackrel{\text{def}}{=} \begin{bmatrix} 1_{z=1} \\ \dots \\ 1_{z=g} \end{bmatrix} \in \mathbb{R}^g \quad (43)$$

$$\phi^{(1)}(\theta) \stackrel{\text{def}}{=} [\ln \alpha_\ell - 0.5 \mu_\ell^T \Gamma \mu_\ell]_{1 \leq \ell \leq g} \in \mathbb{R}^g \quad (44)$$

$$\phi^{(2)}(\theta) \stackrel{\text{def}}{=} [(\Gamma \mu_\ell)]_{1 \leq \ell \leq g} \in \mathbb{R}^{pg} . \quad (45)$$

Remember that $y^T A y = \text{Tr}(A y y^T)$ is the scalar product of A and $y y^T$.

10.1 The model

Let y_1, \dots, y_n be n \mathbb{R}^p -valued observations; they are modeled as the realization of a vector (Y_1, \dots, Y_n) with distribution

- conditionally to a $[g]^*$ -valued vector of random variables (Z_1, \dots, Z_n) , (Y_1, \dots, Y_n) are independent; and the conditional distribution of Y_i is $\mathcal{N}_p(\mu_{Z_i}, \Sigma)$.
- the r.v. (Z_1, \dots, Z_n) are i.i.d., Z_1 takes values on $[g]^*$ with weights $\alpha_1, \dots, \alpha_g$.

Equivalently, the random variables (Y_1, \dots, Y_n) are independent with distribution $\sum_{\ell=1}^g \alpha_\ell \mathcal{N}_p(\mu_\ell, \Sigma)$.

The goal is to estimate the parameter $\theta \in \Theta$ by a Maximum Likelihood approach.

10.1.1 The expression of \mathcal{L}_i for $i \in [n]^*$ and \mathbf{R}

We want to minimize on Θ

$$\theta \mapsto \frac{1}{n} \sum_{i=1}^n \left(\mathbf{R}_{y_i}(\theta) - \ln \sum_{z=1}^g \exp(\langle s_i(z), \phi(\theta) \rangle) \right)$$

which is of the form $n^{-1} \sum_{i=1}^n \mathcal{L}_i(\theta) + \mathbf{R}(\theta)$ with

$$\mathcal{L}_i(\theta) \stackrel{\text{def}}{=} - \ln \sum_{z=1}^g \exp(\langle s_i(z), \phi(\theta) \rangle) , \quad \mathbf{R}(\theta) \stackrel{\text{def}}{=} \frac{1}{2} \text{Tr} \left(\Gamma \frac{1}{n} \sum_{i=1}^n y_i y_i^T \right) - \frac{1}{2} \ln \det(\Gamma) ;$$

s_i is a shorthand notation

$$s_i(z) \stackrel{\text{def}}{=} s_{y_i}(z) = \mathbf{A}_{y_i} \rho(z) . \quad (46)$$

10.1.2 The expression of $p_i(z, \theta)$ and $\bar{s}_i(\theta)$ for $i \in [n]^*$

We have for any $u \in [g]^*$,

$$p_i(u, \theta) \stackrel{\text{def}}{=} \frac{\alpha_u \mathcal{N}_p(\mu_u, \Sigma)[y_i]}{\sum_{\ell=1}^g \alpha_\ell \mathcal{N}_p(\mu_\ell, \Sigma)[y_i]}$$

so that

$$\bar{s}_i(\theta) = \mathbf{A}_{y_i} \bar{\rho}_i(\theta), \quad \bar{\rho}_i(\theta) = \begin{bmatrix} \bar{\rho}_{i,1}(\theta) \\ \dots \\ \bar{\rho}_{i,g}(\theta) \end{bmatrix} \stackrel{\text{def}}{=} \begin{bmatrix} p_i(1, \theta) \\ \dots \\ p_i(g, \theta) \end{bmatrix} \in \mathbb{R}^g.$$

10.1.3 The expression of \mathbb{T} .

Let $s = (s^{(1)}, s^{(2)}) \in \mathbb{R}^g \times \mathbb{R}^{pg}$; we write $\langle s, \phi(\theta) \rangle = \sum_{j=1}^2 \langle s^{(j)}, \phi^{(j)}(\theta) \rangle$ where $\phi^{(j)}$ are defined by (44)-(45).

Remember that $\mathbb{T}(s) = \operatorname{argmin}_{\theta \in \Theta} -\langle s, \phi(\theta) \rangle + \mathbf{R}(\theta)$. We obtain $\mathbb{T}(s) = \{\alpha_1, \dots, \alpha_g, \mu_1, \dots, \mu_g, \Sigma\}$ with

$$\begin{aligned} \alpha_\ell &\stackrel{\text{def}}{=} \frac{s^{(1),\ell}}{\sum_{u=1}^g s^{(1),u}}, \\ \mu_\ell &\stackrel{\text{def}}{=} \frac{s^{(2),\ell}}{s^{(1),\ell}}, \\ \Sigma &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T - \sum_{\ell=1}^g s^{(1),\ell} \mu_\ell \mu_\ell^T. \end{aligned}$$

The expressions of α_ℓ, μ_ℓ are easily obtained; we provide details for the covariance matrix. We have for any symmetric matrix H

$$\begin{aligned} \ln \frac{\det(\Gamma + H)}{\det(\Gamma)} &= \ln \det(I + \Gamma^{-1}H) = \ln(1 + \operatorname{Tr}(\Gamma^{-1}H) + o(\|H\|)) \\ &= \operatorname{Tr}(\Gamma^{-1}H) + o(\|H\|) = \langle H, \Gamma^{-1} \rangle + o(\|H\|) \end{aligned}$$

$\mathbb{T}(s)$ depends on Γ through the function

$$G(\Gamma) \stackrel{\text{def}}{=} -\frac{1}{2} \ln \det(\Gamma) + \frac{1}{2} \left\langle \Gamma, \frac{1}{n} \sum_{i=1}^n y_i y_i^T + \sum_{\ell=1}^g s^{(1),\ell} \mu_\ell \mu_\ell^T \right\rangle - \left\langle \Gamma, \sum_{\ell=1}^g \mu_\ell (s^{(2),\ell})^T \right\rangle.$$

Therefore

$$G(\Gamma + H) - G(\Gamma) = \frac{1}{2} \langle H, \Gamma^{-1} \rangle + \frac{1}{2} \left\langle H, \frac{1}{n} \sum_{i=1}^n y_i y_i^T + \sum_{\ell=1}^g s^{(1),\ell} \mu_\ell \mu_\ell^T \right\rangle - \left\langle H, \sum_{\ell=1}^g \mu_\ell (s^{(2),\ell})^T \right\rangle.$$

This yields the update

$$\begin{aligned}\Sigma &= \Gamma^{-1} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T + \sum_{\ell=1}^g s^{(1),\ell} \mu_\ell \mu_\ell^T - 2 \sum_{\ell=1}^g \mu_\ell (s^{(2),\ell})^T \\ &= \frac{1}{n} \sum_{i=1}^n y_i y_i^T - \sum_{\ell=1}^g s^{(1),\ell} \mu_\ell \mu_\ell^T\end{aligned}$$

by using $\mu_\ell = s^{(2),\ell} / s^{(1),\ell}$.

10.1.4 The domain of \mathbb{T} .

We will prove in the following sections that our algorithms **all** require the computation of $\mathbb{T}(\hat{S})$ for \hat{S} of the form $n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \hat{\rho}_i$. Therefore, let us restrict our attention to the case

$$s = \frac{1}{n} \sum_{i=1}^n \mathbf{A}_{y_i} \rho_i, \quad \rho_i = \begin{bmatrix} \rho_{i,1} \\ \dots \\ \rho_{i,g} \end{bmatrix} \in \mathbb{R}^g,$$

and let us formulate sufficient conditions on $\rho_{i,\ell}$ so that $\mathbb{T}(s) \in \Theta$.

The weights. For all $\ell \in [g]^*$, we want $\alpha_\ell \in [0, 1]$. Therefore, it is required

$$\forall \ell \in [g]^*, \quad \frac{\sum_{i=1}^n \rho_{i,\ell}}{\sum_{i=1}^n \sum_{u=1}^g \rho_{i,u}} \in [0, 1].$$

The expectations. Upon noting that the expression of the log-likelihood of an observation y is unchanged if $y \leftarrow y - c$ and $\mu_\ell \leftarrow \mu_\ell - c$ for any $c \in \mathbb{R}^p$, we must have

$$\frac{1}{n} \sum_{i=1}^n y_i y_i^T - \sum_{\ell=1}^g s^{(1),\ell} \mu_\ell \mu_\ell^T = \frac{1}{n} \sum_{i=1}^n (y_i - c)(y_i - c)^T - \sum_{\ell=1}^g s^{(1),\ell} (\mu_\ell - c)(\mu_\ell - c)^T$$

for any $c \in \mathbb{R}^p$. This yields

$$\sum_{\ell=1}^g s^{(1),\ell} = 1, \quad \sum_{\ell=1}^g s^{(1),\ell} \mu_\ell = \frac{1}{n} \sum_{i=1}^n y_i.$$

Equivalently

$$\frac{1}{n} \sum_{i=1}^n \sum_{\ell=1}^g \rho_{i,\ell} = 1, \quad \frac{1}{n} \sum_{i=1}^n y_i = \frac{1}{n} \sum_{i=1}^n \sum_{\ell=1}^g \rho_{i,\ell} y_i. \quad (47)$$

The covariance matrix. Finally, Σ has to be definite positive: we have

$$\begin{aligned}\Sigma &= \frac{1}{n} \sum_{i=1}^n y_i y_i^T - \frac{1}{n} \sum_{\ell=1}^g \left(\sum_{i=1}^n \rho_{i,\ell} \right) \left(\sum_{i=1}^n \frac{\rho_{i,\ell}}{\sum_{j=1}^n \rho_{j,\ell}} y_i \right) \left(\sum_{i=1}^n \frac{\rho_{i,\ell}}{\sum_{j=1}^n \rho_{j,\ell}} y_i \right)^T \\ &= \frac{1}{n} \sum_{i=1}^n \left(1 - \sum_{\ell=1}^g \rho_{i,\ell} \right) y_i y_i^T \\ &\quad + \frac{1}{n} \sum_{i=1}^n \sum_{\ell=1}^g \rho_{i,\ell} \left(y_i - \sum_{j=1}^n \frac{\rho_{j,\ell}}{\sum_{r=1}^n \rho_{r,\ell}} y_j \right) \left(y_i - \sum_{j=1}^n \frac{\rho_{j,\ell}}{\sum_{r=1}^n \rho_{r,\ell}} y_j \right)^T.\end{aligned}$$

As a conclusion, these conditions are satisfied if

$$\rho_{i,\ell} \geq 0, \quad \sum_{\ell=1}^g \rho_{i,\ell} = 1.$$

Therefore, the domain of \mathbb{T} contains

$$\mathcal{S} \stackrel{\text{def}}{=} \left\{ \frac{1}{n} \sum_{i=1}^n \mathbf{A}_{y_i} \rho_i : \rho_i = (\rho_{i,1}, \dots, \rho_{i,g}) \in (\mathbb{R}_+)^g, \sum_{\ell=1}^g \rho_{i,\ell} = 1 \right\}.$$

10.2 Algorithms

10.2.1 Notations

Given $\theta \in \Theta$, define the a posteriori distribution for all $i \in [n]^*$ and $u \in [g]^*$,

$$p_i(u, \theta) \stackrel{\text{def}}{=} \frac{\alpha_u \mathcal{N}_p(\mu_u, \Sigma)[y_i]}{\sum_{\ell=1}^g \alpha_\ell \mathcal{N}_p(\mu_\ell, \Sigma)[y_i]}.$$

For all $i \in [n]^*$, set

$$\bar{s}_i(\theta) \stackrel{\text{def}}{=} \mathbf{A}_{y_i} \rho_i(\theta) \quad \rho_i(\theta) = \begin{bmatrix} \rho_{i,1}(\theta) \\ \dots \\ \rho_{i,g}(\theta) \end{bmatrix} \stackrel{\text{def}}{=} \begin{bmatrix} p_i(1, \theta) \\ \dots \\ p_i(g, \theta) \end{bmatrix}.$$

For a subset $\mathcal{B} \subseteq \{1, \dots, n\}$ of size b , define

$$\bar{s}_{\mathcal{B}}(\theta) \stackrel{\text{def}}{=} \frac{1}{b} \sum_{i \in \mathcal{B}} \bar{s}_i(\theta) = \frac{1}{b} \sum_{i=1}^n \mathbf{A}_{y_i} \rho_i(\theta) 1_{i \in \mathcal{B}}.$$

10.2.2 The EM algorithm

Input. the current value of the parameter θ^k .

One iteration. Compute the statistic

$$\bar{s}(\theta^k) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \mathbf{A}_{y_i} \rho_i(\theta^k).$$

Update the parameter $\theta^{k+1} = \mathbb{T}(\bar{s}(\theta^k))$.

Is the statistic in the domain of \mathbb{T} ? We have $\theta^{k+1} = \mathbb{T}(\bar{s}(\theta^k))$ with $\bar{s}(\theta^k) = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \rho_i(\theta^k)$. It is easily seen that

$$\rho_{i,\ell}(\theta^k) \geq 0, \quad \sum_{\ell=1}^g \rho_{i,\ell}(\theta^k) = 1,$$

which implies that $\widehat{\mathcal{S}}^{k+1} \in \mathcal{S}$ and therefore, $\theta^{k+1} \in \Theta$.

Update the parameters. for $\ell \in [g]^*$

$$\begin{aligned} \alpha_\ell^{k+1} &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \rho_{i,\ell}(\theta^k), \\ \mu_\ell^{k+1} &\stackrel{\text{def}}{=} \frac{\sum_{i=1}^n \rho_{i,\ell}(\theta^k) y_i}{\sum_{i=1}^n \rho_{i,\ell}(\theta^k)} = \frac{1}{n \alpha_\ell^{k+1}} \sum_{i=1}^n \rho_{i,\ell}(\theta^k) y_i, \\ \Sigma^{k+1} &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T - \sum_{\ell=1}^g \alpha_\ell^{k+1} \mu_\ell^{k+1} (\mu_\ell^{k+1})^T. \end{aligned}$$

10.2.3 The iEM algorithm

Input:

- the current value of the parameter θ^k
- a step size $\gamma_{k+1} \in [0, 1]$.
- the current value of the statistic $\widehat{\mathcal{S}}^k = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \hat{\rho}_i^k$ where $\hat{\rho}_{i,\ell}^k \geq 0$ and $\sum_{\ell=1}^g \hat{\rho}_{i,\ell}^k = 1$.
- the current memory vectors $\mathbf{S}_i^k = \mathbf{A}_{y_i} \tilde{\rho}_i^k$ for $i = 1, \dots, n$, where $\sum_{\ell=1}^g \tilde{\rho}_{i,\ell}^k = 1$.
- the current mean of this vector $\tilde{\mathcal{S}}^k = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \tilde{\rho}_i^k$.

<p>Data: $K_{\max} \in \mathbb{N}$, $\theta^{\text{init}} \in \Theta$, $\Sigma_{\star} \stackrel{\text{def}}{=} n^{-1} \sum_{i=1}^n y_i y_i^T$</p> <p>Result: The EM sequence: $(\widehat{S}^k, \theta^k)$, $k \in [K_{\max}]$</p> <p>1 /* Initialization */ ;</p> <p>2 Compute $\widehat{S}^0 = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \rho_i(\theta^{\text{init}})$ for $j = 1, 2$;</p> <p>3 Set $\theta^0 = \theta^{\text{init}}$;</p> <p>4 for $k = 0, \dots, K_{\max} - 1$ do</p> <p>5 /* Update the statistics */ ;</p> <p>6 $\widehat{S}^{k+1} = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \rho_i(\theta^k)$;</p> <p>7 /* Update the parameter θ^{k+1} */ ;</p> <p>8 $\alpha_{\ell}^{k+1} = n^{-1} \sum_{i=1}^n \rho_{i,\ell}(\theta^k)$ for $\ell \in [g]^{\star}$;</p> <p>9 $\mu_{\ell}^{k+1} = n^{-1} \sum_{i=1}^n \rho_{i,\ell}(\theta^k) y_i / \alpha_{\ell}^{k+1}$ for $\ell \in [g]^{\star}$;</p> <p>10 $\Sigma^{k+1} = \Sigma_{\star} - \sum_{\ell=1}^g \alpha_{\ell}^{k+1} \mu_{\ell}^{k+1} (\mu_{\ell}^{k+1})^T$</p>
--

Algorithm 7: The EM algorithm. Total number of calls to the examples: $K_{\max} \times n$.

One iteration. Sample at random a set \mathcal{B}_{k+1} of b integers in $[n]^{\star}$, possibly with replacement.

Update the memory quantities: for $i \notin \mathcal{B}_{k+1}$, $\mathbf{S}_i^{k+1} = \mathbf{S}_i^k$ and otherwise for all $i \in \mathcal{B}_{k+1}$, $\mathbf{S}_i^{k+1} \stackrel{\text{def}}{=} \mathbf{A}_{y_i} \rho_i(\theta^k)$.

Update its mean

$$\begin{aligned} \widetilde{S}^{k+1} &\stackrel{\text{def}}{=} \widetilde{S}^k + \frac{1}{n} \left(\sum_{i \in \mathcal{B}_{k+1}} \mathbf{S}_i^{k+1} - \sum_{i \in \mathcal{B}_{k+1}} \mathbf{S}_i^k \right) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{A}_{y_i} \left((1 - 1_{i \in \mathcal{B}_{k+1}}) \widetilde{\rho}_i^k + \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} \right) . \end{aligned}$$

Update the statistics \widehat{S} by setting

$$\begin{aligned} \widehat{S}^{k+1} &\stackrel{\text{def}}{=} (1 - \gamma_{k+1}) \widehat{S}^k + \gamma_{k+1} \widetilde{S}^{k+1} \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{A}_{y_i} \left((1 - \gamma_{k+1}) \widehat{\rho}_i^k + \gamma_{k+1} (1 - 1_{i \in \mathcal{B}_{k+1}}) \widetilde{\rho}_i^k + \gamma_{k+1} \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} \right) . \end{aligned}$$

Induction assumption on the expression of \widetilde{S}^{k+1} and \widehat{S}^{k+1} . \widetilde{S}^{k+1} is of the form $n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \widetilde{\rho}_i^{k+1}$ with, for any $i \in [n]^{\star}$,

$$\widetilde{\rho}_i^{k+1} \stackrel{\text{def}}{=} (1 - 1_{i \in \mathcal{B}_{k+1}}) \widetilde{\rho}_i^k + \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} ;$$

is easily seen that $\sum_{\ell=1}^g \hat{\rho}_{i,\ell}^{k+1} = 1$ and $\rho_{i,\ell}^{k+1} \geq 0$.

\hat{S}^{k+1} is of the form $n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \hat{\rho}_i^{k+1}$ with, for any $i \in [n]^*$,

$$\hat{\rho}_i^{k+1} \stackrel{\text{def}}{=} (1 - \gamma_{k+1}) \hat{\rho}_i^k + \gamma_{k+1} (1 - \mathbf{1}_{i \in \mathcal{B}_{k+1}}) \tilde{\rho}_i^k + \gamma_{k+1} \rho_i(\theta^k) \mathbf{1}_{i \in \mathcal{B}_{k+1}} .$$

Since $\sum_{\ell=1}^g \rho_{i,\ell}(\theta^k) = \sum_{\ell=1}^g \hat{\rho}_{i,\ell} = \sum_{\ell=1}^g \tilde{\rho}_{i,\ell}^k = 1$ then $\sum_{\ell=1}^g \hat{\rho}_{i,\ell}^{k+1} = 1$. In addition, since $\rho_i, \hat{\rho}_{i,\ell}^k$ and $\tilde{\rho}_{i,\ell}^k$ are non negative and $\gamma_{k+1} \in [0, 1]$, then $\hat{\rho}_{i,\ell}^{k+1}$ is non negative.

It the statistic \hat{S}^{k+1} in the domain of \mathbb{T} ? We have just established that $\hat{\rho}_{i,\ell}^{k+1} \geq 0$ and $\sum_{\ell=1}^g \hat{\rho}_{i,\ell}^{k+1} = 1$. Consequently $\hat{S}^{k+1} \in \mathcal{S}$ which implies that $\mathbb{T}(\hat{S}^{k+1}) \in \Theta$.

Update the parameters:

$$\begin{aligned} \alpha_\ell^{k+1} &\stackrel{\text{def}}{=} \hat{S}_\ell = \frac{1}{n} \sum_{i=1}^n \hat{\rho}_{i,\ell}^{k+1} , \\ \mu_\ell^{k+1} &\stackrel{\text{def}}{=} \frac{\hat{S}_\ell^{k+1}}{\hat{S}_\ell^{k+1}} = \frac{1}{n \alpha_\ell^{k+1}} \sum_{i=1}^n y_i \hat{\rho}_{i,\ell}^{k+1} , \\ \Sigma^{k+1} &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T - \sum_{\ell=1}^g \alpha_\ell^{k+1} \mu_\ell^{k+1} (\mu_\ell^{k+1})^T . \end{aligned}$$

10.2.4 The Online-EM algorithm

Input.

- the current value of the parameter θ^k
- a step size $\gamma_{k+1} \in [0, 1]$
- the current value of the statistics $\hat{S}^k = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \hat{\rho}_i^k$ such that $\hat{\rho}_{i,\ell}^k \geq 0$ and $n^{-1} \sum_{i=1}^n \sum_{\ell=1}^g \hat{\rho}_{i,\ell}^k = 1$;

One iteration. Sample at random a set \mathcal{B}_{k+1} of b integers in $\{1, \dots, n\}$, with NO replacement; and compute the statistics

$$\begin{aligned} \hat{S}^{k+1} &\stackrel{\text{def}}{=} (1 - \gamma_{k+1}) \hat{S}^k + \gamma_{k+1} \frac{1}{b} \sum_{i \in \mathcal{B}_{k+1}} \mathbf{A}_{y_i} \rho_i(\theta^k) , \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{A}_{y_i} \left((1 - \gamma_{k+1}) \hat{\rho}_i^k + \gamma_{k+1} \frac{n}{b} \rho_i(\theta^k) \mathbf{1}_{i \in \mathcal{B}_{k+1}} \right) . \end{aligned}$$

```

Data:  $K_{\max} \in \mathbb{N}$ ,  $\theta^{\text{init}} \in \Theta$ ,  $\gamma \in (0, 1]$ ,  $b \in \mathbb{N}$ ,  $\Sigma_{\star} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T$ .
Result: The incremental EM sequence:  $(\widehat{S}^k, \theta^k)$ ,  $k \in [K_{\max}]$ 
1 /* Initialization */ ;
2 Compute  $S_{0,i} = A_{y_i} \rho_i(\theta^{\text{init}})$  for  $i \in [n]^*$  ;
3 Compute  $\widetilde{S}^0 = \widehat{S}^0 = n^{-1} \sum_{i=1}^n A_{y_i} \rho_i(\theta^{\text{init}})$  ;
4 Set  $\theta^0 = \theta^{\text{init}}$  ;
5 for  $k = 0, \dots, K_{\max} - 1$  do
6   Sample a mini-batch  $\mathcal{B}_{k+1}$  of size  $b$  ;
7   /* Update the memory quantities */ ;
8    $S_{k+1,i} = S_{k,i}$  for  $i \notin \mathcal{B}_{k+1}$ ;
9    $S_{k+1,i} = A_{y_i} \rho_i(\theta^k)$  for  $i \in \mathcal{B}_{k+1}$  ;
10   $\widetilde{S}^{k+1} = \widetilde{S}^k + n^{-1} \sum_{i \in \mathcal{B}_{k+1}} (S_{k+1,i} - S_{k,i})$  ;
11  /* Update the statistics */ ;
12   $\widehat{S}^{k+1} = (1 - \gamma)\widehat{S}^k + \gamma\widetilde{S}^{k+1}$  ;
13  /* Update the parameter  $\theta^{k+1}$  */ ;
14   $\alpha_{\ell}^{k+1} = \widehat{S}_{\ell}^{k+1}$  for  $\ell \in [g]^*$ ;
15   $\mu_{\ell}^{k+1} = \widehat{S}_{g+(\ell-1)p+1:g+\ell p}^{k+1} / \widehat{S}_{\ell}^{k+1}$  for  $\ell \in [g]^*$  ;
16   $\Sigma^{k+1} = \Sigma_{\star} - \sum_{\ell=1}^g \alpha_{\ell}^{k+1} \mu_{\ell}^{k+1} (\mu_{\ell}^{k+1})^T$ 

```

Algorithm 8: The incremental EM algorithm. Total number of calls to the examples: $n + K_{\max} \times b$.

It is of the form $\widehat{S}^{k+1} = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \widehat{\rho}_i^{k+1}$ with

$$\widehat{\rho}_i^{k+1} \stackrel{\text{def}}{=} (1 - \gamma_{k+1}) \widehat{\rho}_i^k + \gamma_{k+1} \frac{n}{b} \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} .$$

Induction assumption on the expression of \widehat{S}^{k+1} . Since $\rho_{i,\ell}(\theta^k) \geq 0$, $\widehat{\rho}_{i,\ell}^k \geq 0$ and $\gamma_{k+1} \in [0, 1]$, then $\widehat{\rho}_{i,\ell}^{k+1} \geq 0$.

Since \mathcal{B}_{k+1} is sampled with NO replacement, we have $\sum_{i=1}^n 1_{i \in \mathcal{B}_{k+1}} = b$ thus implying, by using the induction assumption $n^{-1} \sum_{i=1}^n \sum_{\ell=1}^g \widehat{\rho}_{i,\ell}^k = 1$ that

$$n^{-1} \sum_{i=1}^n \sum_{\ell=1}^g \widehat{\rho}_{i,\ell}^{k+1} = 1 .$$

Is the statistic \widehat{S}^{k+1} in the domain of \mathbb{T} ? We have $\theta^{k+1} = \mathbb{T}(\widehat{S}^{k+1})$ with $\widehat{S}^{k+1} = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \widehat{\rho}_i^{k+1}$. Unfortunately, we can not prove (even by induction) that $\sum_{\ell=1}^g \widehat{\rho}_{i,\ell}^{k+1} = 1$ for all $i \in [n]^*$. Therefore, we do not have necessarily $\widehat{S}^{k+1} \in \mathcal{S}$.

Update the parameters. For $\ell \in [g]^*$,

$$\begin{aligned} \alpha_\ell^{k+1} &\stackrel{\text{def}}{=} \widehat{S}_\ell^{k+1} = \frac{1}{n} \sum_{i=1}^n \widehat{\rho}_{i,\ell}^{k+1} , \\ \mu_\ell^{k+1} &\stackrel{\text{def}}{=} \frac{\widehat{S}_\ell^{k+1}}{\widehat{S}_\ell^{k+1}} = \frac{1}{n \alpha_\ell^{k+1}} \sum_{i=1}^n \widehat{\rho}_{i,\ell}^{k+1} y_i , \\ \Sigma^{k+1} &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T - \sum_{\ell=1}^g \alpha_\ell^{k+1} \mu_\ell^{k+1} (\mu_\ell^{k+1})^T . \end{aligned}$$

10.2.5 The FIEM algorithm

Input.

- the current value of the parameter θ^k
- a step size $\gamma_{k+1} \in [0, 1]$.
- the current value of the statistics $\widehat{S}^k = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \widehat{\rho}_i^k$
- the current memory vectors $S_i^k = \mathbf{A}_{y_i} \widehat{\rho}_i^k$, with $\sum_{\ell=1}^g \widehat{\rho}_{i,\ell}^k = 1$.
- the current mean of this vector $\widetilde{S}^k = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \widehat{\rho}_i^k$, with $\sum_{\ell=1}^g \widehat{\rho}_{i,\ell}^k = 1$ and $\widehat{\rho}_{i,\ell}^k \geq 0$.

<p>Data: $K_{\max} \in \mathbb{N}$, $\theta^{\text{init}} \in \Theta$, $\gamma \in (0, 1]$, $b \in \mathbb{N}$, $\Sigma_{\star} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T$.</p> <p>Result: The Online EM sequence: $(\widehat{S}^k, \theta^k)$, $k \in [K_{\max}]$</p> <pre> 1 /* Initialization */ ; 2 Compute $\widehat{S}^0 = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \rho_i(\theta^{\text{init}})$ for $j = 1, 2$; 3 Set $\theta^0 = \theta^{\text{init}}$; 4 for $k = 0, \dots, K_{\max} - 1$ do 5 Sample a mini-batch \mathcal{B}_{k+1} of size b with no replacement ; 6 /* Update the statistics */ ; 7 $\bar{s}_{\mathcal{B}_{k+1}}(\theta^k) = b^{-1} \sum_{i \in \mathcal{B}_{k+1}} \mathbf{A}_{y_i} \rho_i(\theta^k)$; 8 $\widehat{S}^{k+1} = (1 - \gamma)\widehat{S}^k + \gamma \bar{s}_{\mathcal{B}_{k+1}}(\theta^k)$; 9 /* Update the parameter θ^{k+1} */ ; 10 $\alpha_{\ell}^{k+1} = \widehat{S}_{\ell}^{k+1}$ for $\ell \in [g]^{\star}$; 11 $\mu_{\ell}^{k+1} = \widehat{S}_{g+(\ell-1)p+1:g+\ell p}^{k+1} / \alpha_{\ell}^{k+1}$ for $\ell \in [g]^{\star}$; 12 $\Sigma^{k+1} = \Sigma_{\star} - \sum_{\ell=1}^g \alpha_{\ell}^{k+1} \mu_{\ell}^{k+1} (\mu_{\ell}^{k+1})^T$ </pre>

Algorithm 9: The Online EM algorithm. Total number of calls to the examples: $n + K_{\max} \times b$.

One iteration. Sample a mini-batch \mathcal{B}_{k+1} of size b and update the memory quantities: for $i \notin \mathcal{B}_{k+1}$, $S_i^{k+1} = S_i^k$ and otherwise for $i \in \mathcal{B}_{k+1}$, $S_i^{k+1} \stackrel{\text{def}}{=} \mathbf{A}_i \rho_i(\theta^k)$. Update the mean of this memory quantity

$$\begin{aligned} \widetilde{S}^{k+1} &\stackrel{\text{def}}{=} \widetilde{S}^k + \frac{1}{n} \sum_{i \in \mathcal{B}_{k+1}} (S_i^{k+1} - S_i^k) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{A}_{y_i} \left((1 - 1_{i \in \mathcal{B}_{k+1}}) \widetilde{\rho}_i^k + \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} \right) . \end{aligned}$$

Sample a second mini-batch \mathcal{B}'_{k+1} of size b . Compute

$$\bar{s}_{\mathcal{B}'_{k+1}}(\theta^k) \stackrel{\text{def}}{=} \frac{1}{b} \sum_{i \in \mathcal{B}'_{k+1}} \mathbf{A}_{y_i} \rho_i(\theta^k) .$$

Set $\mathbf{S}_{\mathcal{B}'_{k+1}} \stackrel{\text{def}}{=} \mathfrak{b}^{-1} \sum_{j \in \mathcal{B}'_{k+1}} \mathbf{S}_j^{k+1}$, and update the statistics $\widehat{\mathbf{S}}^{k+1}$

$$\begin{aligned} \widehat{\mathbf{S}}^{k+1} &\stackrel{\text{def}}{=} \widehat{\mathbf{S}}^k + \gamma_{k+1} \left(\bar{\mathbf{s}}_{\mathcal{B}'_{k+1}}(\theta^k) - \widehat{\mathbf{S}}^k + \widetilde{\mathbf{S}}^{k+1} - \mathbf{S}_{\mathcal{B}'_{k+1}} \right), \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{A}_{y_i} \left((1 - \gamma_{k+1}) \hat{\rho}_i^k + \gamma_{k+1} \frac{n}{\mathfrak{b}} \rho_i(\theta^k) 1_{i \in \mathcal{B}'_{k+1}} + \gamma_{k+1} (1 - 1_{i \in \mathcal{B}_{k+1}}) \tilde{\rho}_i^k \right. \\ &\quad \left. + \gamma_{k+1} \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} - \gamma_{k+1} \frac{n}{\mathfrak{b}} \left(\tilde{\rho}_i^k 1_{i \notin \mathcal{B}_{k+1}} + \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} \right) 1_{i \in \mathcal{B}'_{k+1}} \right). \end{aligned}$$

$\widehat{\mathbf{S}}^{k+1}$ is of the form $n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \hat{\rho}_i^{k+1}$ with

$$\begin{aligned} \hat{\rho}_i^{k+1} &\stackrel{\text{def}}{=} (1 - \gamma_{k+1}) \hat{\rho}_i^k + \gamma_{k+1} \frac{n}{\mathfrak{b}} \rho_i(\theta^k) 1_{i \in \mathcal{B}'_{k+1}} + \gamma_{k+1} (1 - 1_{i \in \mathcal{B}_{k+1}}) \tilde{\rho}_i^k \\ &\quad + \gamma_{k+1} \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} - \gamma_{k+1} \frac{n}{\mathfrak{b}} \left(\tilde{\rho}_i^k 1_{i \notin \mathcal{B}_{k+1}} + \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}} \right) 1_{i \in \mathcal{B}'_{k+1}}. \end{aligned}$$

Induction assumption on $\widetilde{\mathbf{S}}^k$ and $\widehat{\mathbf{S}}^k$. Since $\sum_{\ell=1}^g \rho_{i,\ell} = \sum_{\ell=1}^g \hat{\rho}_{i,\ell}^k = 1$, then it is easily seen that $\sum_{\ell=1}^g \hat{\rho}_{i,\ell}^{k+1} = 1$.

We have $\widetilde{\mathbf{S}}^{k+1} = n^{-1} \sum_{i=1}^n \mathbf{A}_{y_i} \tilde{\rho}_i^{k+1}$ with $\tilde{\rho}_i^{k+1} = (1 - 1_{i \in \mathcal{B}_{k+1}}) \tilde{\rho}_i^k + \rho_i(\theta^k) 1_{i \in \mathcal{B}_{k+1}}$. It is easily seen that under the induction assumption $\sum_{\ell=1}^g \tilde{\rho}_{i,\ell}^{k+1} = 1$ and $\tilde{\rho}_{i,\ell}^{k+1} \geq 0$.

Is the statistic $\widehat{\mathbf{S}}^{k+1}$ in the domain of Υ ? The property $\hat{\rho}_{i,\ell}^{k+1} \geq 0$ may fail even assuming that $\hat{\rho}_{i,\ell}^k \geq 0$.

Update the parameters

$$\begin{aligned} \alpha_\ell^{k+1} &\stackrel{\text{def}}{=} \widehat{\mathbf{S}}_\ell^{k+1}, \\ \mu_\ell^{k+1} &\stackrel{\text{def}}{=} \frac{\widehat{\mathbf{S}}_\ell^{k+1}}{\widehat{\mathbf{S}}_\ell^{k+1}}, \\ \Sigma^{k+1} &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n y_i y_i^T - \sum_{\ell=1}^g \alpha_\ell^{k+1} \mu_\ell^{k+1} (\mu_\ell^{k+1})^T. \end{aligned}$$

```

Data:  $K_{\max} \in \mathbb{N}$ ,  $\theta^{\text{init}} \in \Theta$ ,  $\gamma \in (0, 1]$ ,  $b \in \mathbb{N}$ ,  $\Sigma_{\star} \stackrel{\text{def}}{=} n^{-1} \sum_{i=1}^n y_i y_i^T$ .
Result: The FIEM sequence:  $(\widehat{S}^k, \theta^k)$ ,  $k \in [K_{\max}]$ 
1 /* Initialization */ ;
2 Compute  $S_{0,i} = A_{y_i} \rho_i(\theta^{\text{init}})$  for  $i = 1, \dots, n$  ;
3 Compute  $\widetilde{S}^0 = \widehat{S}^0 = n^{-1} \sum_{i=1}^n A_{y_i} \rho_i(\theta^{\text{init}})$  ;
4 Set  $\theta^0 = \theta^{\text{init}}$  ;
5 for  $k = 0, \dots, K_{\max} - 1$  do
6   Sample independently two mini-batches  $\mathcal{B}_{k+1}$  and  $\mathcal{B}'_{k+1}$ , both of size  $b$  ;
7   /* Update the memory quantities */ ;
8    $S_{k+1,i} = S_{k,i}$  for  $i \notin \mathcal{B}_{k+1}$  ;
9    $S_{k+1,i} = A_{y_i} \rho_i(\theta^k)$  for  $i \in \mathcal{B}_{k+1}$  ;
10   $\widetilde{S}^{k+1} = \widetilde{S}^k + n^{-1} \sum_{i \in \mathcal{B}_{k+1}} (S_{k+1,i} - S_{k,i})$  ;
11  /* Update the statistics */ ;
12   $\bar{s}_{\mathcal{B}'_{k+1}}(\theta^k) = b^{-1} \sum_{r \in \mathcal{B}'_{k+1}} A_{y_r} \rho_r(\theta^k)$  ;
13   $S_{k+1, \mathcal{B}'_{k+1}} = b^{-1} \sum_{r \in \mathcal{B}'_{k+1}} S_{k+1,r}$  ;
14   $\widehat{S}^{k+1} = (1 - \gamma) \widehat{S}^k + \gamma \left( \bar{s}_{\mathcal{B}'_{k+1}}(\theta^k) + \widetilde{S}_{k+1} - S_{k+1, \mathcal{B}'_{k+1}} \right)$  ;
15  /* Update the parameter  $\theta^{k+1}$  */ ;
16   $\alpha_{\ell}^{k+1} = \widehat{S}_{\ell}^{k+1}$  for  $\ell \in [g]^*$  ;
17   $\mu_{\ell}^{k+1} = \widehat{S}_{g+(\ell-1)p+1:g+\ell p}^{k+1} / \widehat{S}_{\ell}^{k+1}$  for  $\ell \in [g]^*$  ;
18   $\Sigma^{k+1} = \Sigma_{\star} - \sum_{\ell=1}^g \alpha_{\ell}^{k+1} \mu_{\ell}^{k+1} (\mu_{\ell}^{k+1})^T$ 

```

Algorithm 10: The FIEM algorithm. Total number of calls to the examples: $n + K_{\max} \times 2b$.

10.3 Additional plots




Figure 14: Evolution of the normalized log-likelihood: average over 10 independent paths of length 100 epochs. The first 25 epochs are displayed. All the paths start from the same value at time $t = 0$, with a normalized log-likelihood equal to -58.31 .




Figure 15: Evolution of the normalized log-likelihood: average over 10 independent paths of length 100 epochs. The first 14 epochs are discarded. All the paths start from the same value at time $t = 0$, with a normalized log-likelihood equal to -58.31 .




Figure 16: Evolution of the $p = 20$ components of one of the expectation vector μ_ℓ along one path of length 100 epochs. All the paths start from the same value at time $t = 0$. EM (top left), iEM (top right), Online EM (bottom left) and h-FIEM (bottom right).




Figure 17: Estimation of the $p = 20$ components of one of the expectation vector μ_ℓ . The estimator is the value of the parameter obtained at the end of a single path of length 100 epochs.

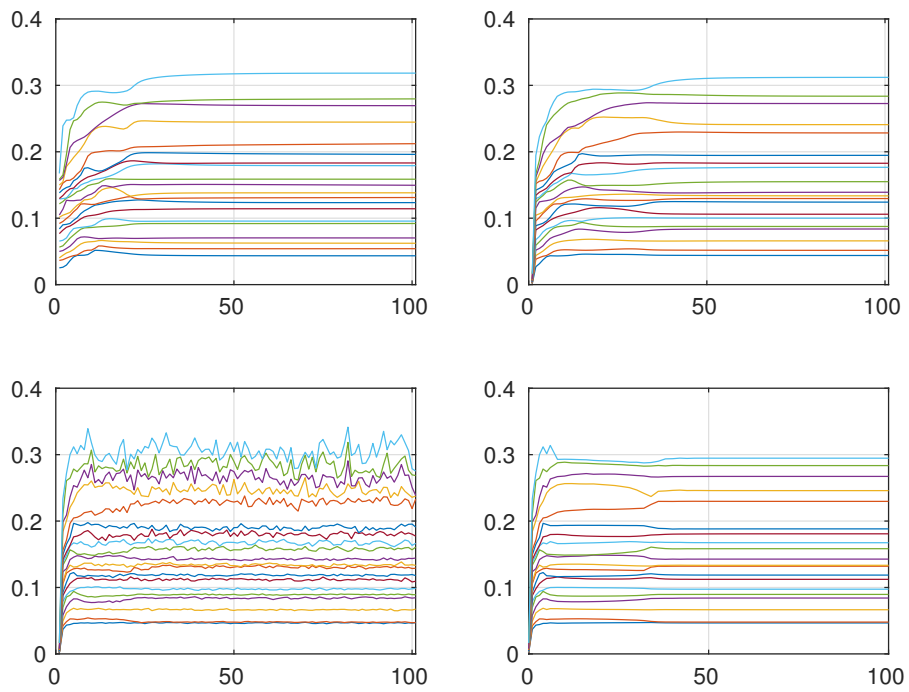


Figure 18: Evolution of the $p = 20$ eigenvalues of the covariance matrix Σ along one path of length 100 epochs. All the paths start from the same value at time $t = 0$. EM (top left), iEM (top right), Online EM (bottom left) and h-FIEM (bottom right).

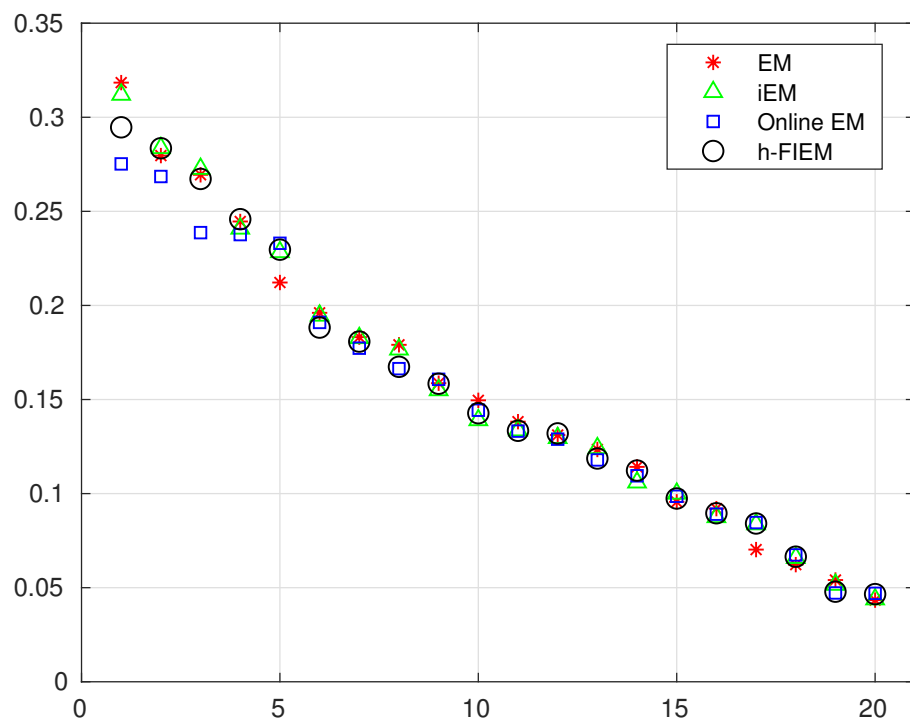


Figure 19: Estimation of the $p = 20$ eigenvalues of the covariance matrix Σ . The estimator is the value of the parameter obtained at the end of a single path of length 100 epochs.