



HAL
open science

Federated Majorize-Minimization: Beyond Parameter Aggregation

Aymeric Dieuleveut, Gersende Fort, Mahmoud Hegazy, Hoi-To Wai

► To cite this version:

Aymeric Dieuleveut, Gersende Fort, Mahmoud Hegazy, Hoi-To Wai. Federated Majorize-Minimization: Beyond Parameter Aggregation. 2025. <hal-05189632>

HAL Id: hal-05189632

<https://hal.science/hal-05189632v1>

Preprint submitted on 28 Jul 2025

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

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



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Federated Majorize-Minimization: Beyond Parameter Aggregation

Aymeric Dieuleveut

*CMAP, École polytechnique,
Insitut Polytechnique de Paris*

AYMERIC.DIEULEVEUT@POLYTECHNIQUE.EDU

Gersende Fort

*LAAS-CNRS, Université de Toulouse, CNRS,
Toulouse, France*

GERSENDE.FORT@LAAS.FR

Mahmoud Hegazy

*CMAP, École polytechnique,
Insitut Polytechnique de Paris*

MAHMOUD.HEGAZY@POLYTECHNIQUE.EDU

Hoi-To Wai

*Department of SEEM, The Chinese University of Hong Kong
Shatin, Hong Kong*

HTWAI@SE.CUHK.EDU.HK

Abstract

This paper proposes a unified approach for designing stochastic optimization algorithms that robustly scale to the federated learning setting. Our work studies a class of Majorize-Minimization (MM) problems, which possesses a linearly parameterized family of majorizing surrogate functions. This framework encompasses (proximal) gradient-based algorithms for (regularized) smooth objectives, the Expectation Maximization algorithm, and many problems seen as variational surrogate MM. We show that our framework motivates a unifying algorithm called Stochastic Approximation Stochastic Surrogate MM (SA-SSMM), which includes previous stochastic MM procedures as special instances. We then extend SA-SSMM to the federated setting, while taking into consideration common bottlenecks such as data heterogeneity, partial participation, and communication constraints; this yields FedMM. The originality of FedMM is to learn locally and then aggregate information characterizing the *surrogate majorizing function*, contrary to classical algorithms which learn and aggregate the *original parameter*. Finally, to showcase the flexibility of this methodology beyond our theoretical setting, we use it to design an algorithm for computing optimal transport maps in the federated setting.

1 Introduction

Consider the following optimization problem

$$\operatorname{argmin}_{\theta \in \mathbb{R}^d} W(\theta) := (f(\theta) + g(\theta)), \quad \text{with } f(\theta) := \mathbb{E}_{Z \sim \pi} [\ell(Z, \theta)]; \quad (1)$$

where $g : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ is a proper lower semi-continuous convex function with domain $\Theta \subseteq \mathbb{R}^d$, $\ell : \mathbb{R}^p \times \Theta \rightarrow \mathbb{R}$ is a measurable function such that $\mathbb{E}_{\pi} [|\ell(Z, \theta)|] < +\infty$ for any $\theta \in \Theta$ and π denotes a probability measure on \mathbb{R}^p . This setting covers both the *batch learning*

when π is the empirical distribution of a (possibly large) data set $\{Z_1, \dots, Z_N\}$ and we solve $\operatorname{argmin}_{\theta \in \mathbb{R}^d} (N^{-1} \sum_{j=1}^N \ell(Z_j, \theta) + g(\theta))$ and the *online learning* when we have to process an i.i.d. stream of data $\{Z_1, Z_2, \dots\}$ with the distribution π .

We are interested in the setting where *majorizing functions* can be computed on f , which can then be minimized. This is a framework referred to as Majorize-Minimization (MM) (Lange, 2013), or surrogate optimization (Lange et al., 2000). More concretely, starting from the parameter θ_t , a standard MM algorithm constructs a surrogate $U_t : \Theta \rightarrow \mathbb{R}$ such that $f(\theta_t) = U_t(\theta_t)$ and $f(\theta) \leq U_t(\theta)$ for any $\theta \in \Theta$. Then, a parameter is constructed such that $\theta_{t+1} \in \operatorname{argmin} U_t(\theta) + g(\theta)$. Many methods can be interpreted from this point of view, e.g., gradient-based and proximal methods (Beck and Teboulle, 2009a; Wright et al., 2009), Expectation-Maximization (EM) (Dempster et al., 1977), boosting (Collins et al., 2002; Della Pietra et al., 2001), and matrix factorization (MF) algorithms such as non-negative MF (Paatero and Tapper, 1994; Lee and Seung, 1999) or clustering (Jain et al., 1999).

In a series of papers that inspired our work, Mairal (2013, 2015) investigated the use of MM methods in large-scale machine learning. The author proposed to incrementally build the sequence of surrogates by leveraging U_1, \dots, U_t in constructing U_{t+1} . We build on this approach and demonstrate that in many examples, *the space of the surrogate functions can be parameterized by a finite-dimensional space \mathcal{S}* ; i.e., each $s \in \mathcal{S}$ represents the mirror parameter characterizing the surrogate function $U(\cdot, s)$. We then perform (stochastic) incremental updates in the space of surrogate functions instead of the space of parameters θ .

This paper is particularly drawn to the Federated Learning (FL) setting (Kairouz et al., 2021), where data is shared between several agents and a central server aggregates the contributions of each agent. More concretely, we consider the extension of (1) to the federated setting, where n different *agents* collaborate to solve

$$\operatorname{argmin}_{\theta \in \mathbb{R}^d} \sum_{i=1}^n \mu_i f_i(\theta) + g(\theta), \quad (2)$$

where $f_i(\theta) := \mathbb{E}_{\pi_i} [\ell(Z, \theta)]$ and $\mu_i > 0$ is a weighting parameter, $\sum_{i=1}^n \mu_i = 1$. Here, for $i \in \{1, \dots, n\}$, we denote by π_i the probability distribution of the data of agent $\#i$. In addition, for the i -th agent, we consider the surrogate $U : \Theta \times \mathcal{S} \rightarrow \mathbb{R}$ satisfying: for all θ , there exists $s_\theta^i \in \mathcal{S}$ such that $f_i(\theta) = U(\theta, s_\theta^i)$ and $f_i(\cdot) \leq U(\cdot, s_\theta^i)$ for all $\theta \in \Theta$.

Under this setting, our main motivation is that for many problems, the surrogate function $U(\theta, s)$ is linear in the mirror parameter s . Thus, a direct aggregation in the surrogate space corresponds to aggregation in the functional space of upper bounds. This observation enables us to solve (2) via aggregating the mirror parameters in a distributed manner, resulting in a MM-like algorithm for the FL setting. On the other hand, a naive FL extension of a centralized optimization algorithm through aggregating the clients' parameters θ (e.g., FedAvg (McMahan et al., 2017b)) may not result in a stable algorithm. From a methodological standpoint, our key message, which we further detail in Section 3.1, is:

In numerous FL problems, the natural way to aggregate agents' contributions is by first aggregating on the surrogate space to construct a single surrogate function and then to minimize/majorize the common surrogate on the server

In FL, the communication between agents and the central server has been identified as a crucial bottleneck, that has been widely tackled in the gradient-based literature (e.g.

McMahan et al., 2017a; Alistarh et al., 2017). Moreover, the heterogeneity between the multiple agents' distributions often hinders convergence (Li et al., 2020; Sattler et al., 2019). We leverage the stochastic surrogate MM approach in FL to propose a communication-efficient method. Following the works of (Mishchenko et al., 2019), we incorporate control variates to handle heterogeneity. This enables us to propose a generic FL-MM algorithm. In particular cases, these algorithms can be seen as modifications or extensions encompassing multiple existing federated algorithms such as Federated EM (Dieuleveut et al., 2021; Marfoq et al., 2021) and Federated MF (Wang and Chang, 2022).

Outline For the rest of this paper, Section 2 will start by considering the centralized case, where we will introduce an MM framework that unifies many common settings. Using this framework, we show that multiple celebrated algorithms can be seen through the lens of surrogate optimization. We provide SA-SSMM, a Stochastic Surrogate MM method relying on Stochastic Approximation (Robbins and Monro, 1951), for the centralized case. Then, we give the natural extension, named FedMM, of the algorithm to the federated setting in Section 3. Further, Section 4 discusses the related works in the Federated Learning literature. Then, in Section 5, we provide our theoretical results. To evaluate our method, we consider a federated dictionary learning problem in Section 6. Finally, Section 7 shows how our approach can be regarded as a more general methodology for the design of federated algorithm. To this end, we consider a federated optimal transport problem and provide an algorithm to tackle it.

Notations For two vectors $a, b \in \mathbb{R}^q$, $\langle a, b \rangle$ denotes the Euclidean scalar product; we set $\|a\| := \sqrt{\langle a, a \rangle}$. All vectors are column-vectors; for a matrix A , A^\top is its transpose. For two $p \times q$ matrices A and B , $\langle A, B \rangle := \text{Trace}(A^\top B)$; the Frobenius norm of A is also denoted by $\|A\|$. \mathcal{M}_K^+ is the set of the $K \times K$ positive semi-definite matrices; $A \succ 0$ denotes that A is positive-definite.

The subdifferential of a proper function $f : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ is denoted by ∂f . The Jacobian matrix of a differentiable vector-valued function $f : \mathbb{R}^d \rightarrow \mathbb{R}^q$ is a $q \times d$ matrix, denoted by J_f ; when $q = 1$, we denote by $\nabla f(\cdot)$ the $d \times 1$ gradient vector. For $n \in \mathbb{N}$, set $[n] := \{1, \dots, n\}$. $\text{int}\Theta$ is the interior of a set Θ .

2 Linearly Parameterized Surrogate functions

We first consider an MM framework where the surrogate functions are linearly parameterized by focusing on problems satisfying MM-1 and MM-2.

MM-1: there exist a convex subset \mathcal{S} of \mathbb{R}^q and measurable functions $\bar{S} : \mathbb{R}^p \times \Theta \rightarrow \mathcal{S}$, $\psi : \Theta \rightarrow \mathbb{R}$ and $\phi : \Theta \rightarrow \mathbb{R}^q$ such that for any $\tau \in \Theta$, $\mathbb{E}_\pi [\|\bar{S}(Z, \tau)\|] < \infty$ and

$$f(\cdot) \leq f(\tau) + \psi(\cdot) - \psi(\tau) - \langle \mathbb{E}_\pi [\bar{S}(Z, \tau)], \phi(\cdot) - \phi(\tau) \rangle \quad \text{on } \Theta. \quad (3)$$

MM-2: the application $\mathbb{T} : \mathcal{S} \rightarrow \Theta$ defined by $\mathbb{T}(s) := \text{argmin}_{\theta \in \mathbb{R}^d} (g(\theta) + \psi(\theta) - \langle s, \phi(\theta) \rangle)$, is well-defined (i.e., there is a unique minimizer for all $s \in \mathcal{S}$) and measurable.

In words, the surrogate function given by the right-hand side in (3) admits a *parametric form*.

This MM framework encompasses many examples of optimization problems. Hereafter, we give three examples, which are detailed in Appendix A.

Example 1 (Quadratic surrogate) Let $f : \Theta \rightarrow \mathbb{R}$ be continuously differentiable on a neighborhood of Θ and its gradient be L_f -Lipschitz. Assume that there exists a measurable function G such that $\nabla f(\cdot) = \mathbb{E}_\pi [G(Z, \cdot)]$ on Θ . Then, for any $\rho \in (0, 1/L_f]$, MM-1 and MM-2 are satisfied with

$$\psi(\theta) := \frac{1}{2\rho} \|\theta\|^2, \quad \phi(\theta) := \frac{1}{\rho} \theta, \quad \bar{S}(Z, \tau) := \tau - \rho G(Z, \tau). \quad (4)$$

\mathbb{T} computes $\mathbb{T}(s) = \text{prox}_{\rho g}(s)$ where $\text{prox}_{\rho g}(s) := \text{argmin}_{\theta \in \mathbb{R}^d} (\rho g(\theta) + 1/2 \|\theta - s\|^2)$ is the Moreau-proximity operator of the function $\rho g(\cdot)$ at s (Moreau, 1962).

Example 2 (Jensen surrogate) Assume that $\ell(Z, \theta)$ is a function of the form

$$\ell(Z, \theta) := -\log \int_{\mathcal{H}} p(Z, h, \theta) \mu(\mathrm{d}h), \quad (5)$$

where $(Z, h, \theta) \mapsto p(Z, h, \theta)$ is a measurable positive function, and μ is a σ -finite positive measure on the measurable set \mathcal{H} . Assume also that $-\log p(Z, h, \theta)$ can be written in the form $\psi(\theta) - \langle S(Z, h), \phi(\theta) \rangle$ up to an additive term independent of θ . Then under integrability conditions, MM-1 holds with the functions ψ, ϕ given by $-\log p$ and with

$$\bar{S}(Z, \tau) := \int_{\mathcal{H}} S(Z, h) \mu(\mathrm{d}h | Z, \tau), \quad (6)$$

where for any (Z, τ) , $\mu(\cdot | Z, \tau)$ is a probability distribution on \mathcal{H} defined by

$$\mu(\mathrm{d}h | Z, \tau) := \frac{\exp(\langle S(Z, h), \phi(\tau) \rangle)}{\int_{\mathcal{H}} \exp(\langle S(Z, u), \phi(\tau) \rangle) \mu(\mathrm{d}u)} \mu(\mathrm{d}h).$$

An instance of such a loss function $\ell(Z, \theta)$ is the negative log-likelihood of the observation Z in a latent variable model when the complete likelihood $\theta \mapsto p(Z, h, \theta)$ is from the vector exponential family (see e.g. Brown (1986)). $S(Z, h)$ is the sufficient statistics, and $\exp(\psi)$ acts as a normalizing constant and is fully determined by $S(Z, h)$ and ϕ . The distribution $\mu(\mathrm{d}h | Z, \tau)$ is the a posteriori distribution of the latent variable h given the observation Z , when the model parameter is τ .

Example 3 (Variational surrogate) Assume that $\ell(Z, \theta) := \min_{h \in \mathcal{H}} \tilde{\ell}(Z, h, \theta)$ with

$$\tilde{\ell}(Z, h, \theta) := \psi(\theta) - \langle S(Z, h), \phi(\theta) \rangle + \xi(Z, h);$$

and the optimization problem $(Z, \theta) \mapsto \mathbf{M}(Z, \theta) := \text{argmin}_{h \in \mathcal{H}} \tilde{\ell}(Z, \theta, h)$ possesses a unique solution. Then under integrability conditions, MM-1 is satisfied with $\bar{S}(Z, \tau) := S(Z, \mathbf{M}(Z, \tau))$.

2.1 MM as a root-finding algorithm

The first consequence of MM1&2 is that for any $\tau \in \Theta$, there exists a majorizing function of the objective function which equals $(f + g)$ at τ :

$$(f + g)(\cdot) \leq f(\tau) + \psi(\cdot) - \psi(\tau) - \langle \mathbb{E}_\pi [\bar{S}(Z, \tau)], \phi(\cdot) - \phi(\tau) \rangle + g(\cdot), \quad \text{on } \Theta.$$

Based on this property, the MM algorithm defines a Θ -valued sequence by repeating two steps: given the current iterate $\theta_t \in \Theta$, (i) identify a surrogate function by computing the *mirror* parameter $\mathbb{E}_\pi [\bar{S}(Z, \theta_t)]$; (ii) define the next iterate as the minimizer of the surrogate function

$$\theta_{t+1} = \mathsf{T} \left(\mathbb{E}_\pi [\bar{S}(Z, \theta_t)] \right). \quad (7)$$

This sequence satisfies $(f + g)(\theta_{t+1}) \leq (f + g)(\theta_t)$ (Lange, 2013, Chapter 12).

Equivalently, the MM algorithm defines a \mathcal{S} -valued sequence which characterizes a sequence of surrogate functions in the functional space spanned by $g + \psi$ and ϕ_1, \dots, ϕ_q : given $s_t \in \mathcal{S}$ of the form $s_t = \mathbb{E}_\pi [\bar{S}(Z, \tau)]$ for some $\tau \in \Theta$, (i) identify the *mirror parameter* $\mathsf{T}(s_t)$ defined as the minimizer of the surrogate function identified by s_t ; (ii) define the next surrogate function by computing its parameter

$$s_{t+1} = \mathbb{E}_\pi [\bar{S}(Z, \mathsf{T}(s_t))]. \quad (8)$$

This sequence satisfies $(f + g)(\mathsf{T}(s_{t+1})) \leq (f + g)(\mathsf{T}(s_t))$.

The limiting points of the two MM methods are equivalent. Define the *mean field function*

$$\mathbf{h} : \mathcal{S} \rightarrow \mathbb{R}^q \quad \mathbf{h}(s) := \mathbb{E}_\pi [\bar{S}(Z, \mathsf{T}(s))] - s; \quad (9)$$

note that a root of \mathbf{h} is a fixed point of the iterative scheme (8). Proposition 1 states that a root of \mathbf{h} can be used along the mapping T to find a fixed point of the iterative scheme (7). Conversely, a fixed point of the iterative scheme (7) can be used along the mapping $\mathbb{E}_\pi [\bar{S}(Z, \cdot)]$ to find a root of \mathbf{h} . This is the main message of Proposition 1, which is established under the following assumptions.

A1 For any $s \in \mathcal{S}$, $\theta \mapsto \psi(\theta) - \langle s, \phi(\theta) \rangle$ is convex. In addition, $\text{int}\Theta \neq \emptyset$ and ψ and ϕ are \mathcal{C}^1 on $\text{int}\Theta$.

A2 The function f is \mathcal{C}^1 on $\text{int}\Theta$.

Proposition 1 Assume MM-1, MM-2, A1 and A2. For any $\tau \in \text{int}\Theta$,

$$0 \in \nabla f(\tau) + \partial g(\tau) \iff \mathsf{T} \left(\mathbb{E}_\pi [\bar{S}(Z, \tau)] \right) - \tau = 0. \quad (10)$$

It holds that for any $s \in \mathcal{S}$ such that $\mathsf{T}(s) \in \text{int}\Theta$,

$$-J_\phi(\mathsf{T}(s)) \mathbf{h}(s) \in \nabla f(\mathsf{T}(s)) + \partial g(\mathsf{T}(s)). \quad (11)$$

In addition, if $s_\star \in \mathcal{S}$ satisfies $\mathbf{h}(s_\star) = 0$ and $\mathsf{T}(s_\star) \in \text{int}\Theta$, then $\theta_\star := \mathsf{T}(s_\star)$ satisfies $\mathsf{T} \left(\mathbb{E}_\pi [\bar{S}(Z, \theta_\star)] \right) - \theta_\star = 0$. Conversely, if $\mathsf{T} \left(\mathbb{E}_\pi [\bar{S}(Z, \theta_\star)] \right) - \theta_\star = 0$ and $\theta_\star \in \text{int}\Theta$, then $s_\star := \mathbb{E}_\pi [\bar{S}(Z, \theta_\star)]$ satisfies $\mathbf{h}(s_\star) = 0$.

We defer the proof of Proposition 1 to Appendix B. A similar result was proved in Nguyen et al. (2023, Lemma 1).

2.2 A Stochastic Surrogate MM algorithm

In the large-scale learning setting considered here, the expectation $\mathbb{E}_\pi [\bar{S}(Z, \tau)]$ is not explicit but stochastic oracles are available, whatever $\tau \in \Theta$; on the other hand, the minimizer $\mathsf{T}(s)$ still has a closed-form expression whatever $s \in \mathcal{S}$. Thus, we propose a stochastic version of MM which replaces the deterministic fixed point algorithm by a Stochastic Approximation scheme (SA, see e.g. Robbins and Monro (1951), Benveniste et al. (1990); Borkar (2008)) designed to solve the root of a vector field $v : \mathbb{R}^u \rightarrow \mathbb{R}^u$; it repeats $x_{t+1} = x_t + \gamma_{t+1}U_{t+1}$ where $\{\gamma_t, t \geq 1\}$ is a sequence of positive deterministic step sizes and U_{t+1} is a random oracle of $v(x_t)$. Two properties of the oracles are crucial in the convergence analysis: its conditional bias given the past of the algorithm and its conditional variance (see e.g. Andrieu et al. (2005); Dieuleveut et al. (2023)). In light of MM-1 and MM-2, two strategies are possible for tackling (1): define a Θ -valued sequence targeting the roots of $\theta \mapsto \chi(\theta) := \mathsf{T}(\mathbb{E}_\pi [\bar{S}(Z, \theta)]) - \theta$; or define a \mathcal{S} -valued sequence targeting the roots of \mathbf{h} (see (9)). The main advantage of the second approach is to conserve the unbiasedness property of the oracle: when S is an unbiased oracle of $\mathbb{E}_\pi [\bar{S}(Z, \tau)]$, then $\mathsf{T}(\mathsf{S}) - \tau$ is a biased oracle of $\chi(\tau)$ - except in the specific case when T is linear; while when S is an unbiased oracle of $\mathbb{E}_\pi [\bar{S}(Z, \mathsf{T}(s))]$, then $\mathsf{S} - s$ is also an unbiased oracle of $\mathbf{h}(s)$. Biased oracles may ruin the convergence of the algorithm and have to be controlled with convenient design of the oracles, usually inducing a higher computational cost (see e.g. (Dieuleveut et al., 2023, Section IV.B.3.c)).

This yields the *Stochastic Approximation Stochastic Surrogate MM* (SA-SSMM) algorithm in Algorithm 1.

Algorithm 1: Stochastic Approximation Stochastic Surrogate MM (SA-SSMM)

Data: $T_{\max} \in \mathbb{N}^*$; a $(0, 1]$ -valued sequence $\{\gamma_t, t \in [T_{\max}]\}$; initial value $\hat{S}_0 \in \mathcal{S}$

Result: SA-SSMM sequence $\{\hat{S}_t, t \in \{0, \dots, T_{\max}\}\}$ and its mirror sequence $\{\mathsf{T}(\hat{S}_t), t \in \{0, \dots, T_{\max}\}\}$

- 1 **for** $t = 0, \dots, T_{\max} - 1$ **do**
 - 2 Sample S_{t+1} in \mathcal{S} , a random oracle for $\mathbb{E}_\pi[\bar{S}(Z, \mathsf{T}(\hat{S}_t))]$
 - 3 Set $\hat{S}_{t+1} = \hat{S}_t + \gamma_{t+1}(\mathsf{S}_{t+1} - \hat{S}_t)$
-

At the t -th iteration, Algorithm 1 utilizes a stochastic oracle of $\mathbf{h}(\hat{S}_t)$, the mean field function evaluated at the current iterate, defined by $\mathsf{S}_{t+1} - \hat{S}_t$, where S_{t+1} approximates $\mathbb{E}_\pi[\bar{S}(Z, \mathsf{T}(\hat{S}_t))]$. In *online learning*, given a new example Z_{t+1} , one can set $\mathsf{S}_{t+1} = \bar{S}(Z_{t+1}, \mathsf{T}(\hat{S}_t))$. In *large batch learning* with N examples, a possible oracle is $\mathsf{S}_{t+1} = \mathbf{b}^{-1} \sum_{j \in \mathcal{B}_{t+1}} \bar{S}(Z_j, \mathsf{T}(\hat{S}_t))$ where \mathcal{B}_{t+1} is a mini-batch of size \mathbf{b} sampled at random in $[N]$.

Since \mathcal{S} is convex (see MM-1) and $\gamma_{t+1} \in (0, 1]$, we have $\hat{S}_t + \gamma_{t+1}(\mathsf{S}_{t+1} - \hat{S}_t) \in \mathcal{S}$ and $\mathsf{T}(\hat{S}_{t+1})$ is well-defined. Note that in the special case of constant step sizes $\gamma_{t+1} = \gamma$ for all t , we have $\hat{S}_{t+1} = (1 - \gamma)^{t+1} \hat{S}_0 + \gamma \sum_{j=0}^t (1 - \gamma)^j \hat{S}_{t+1-j}$: the iterate \hat{S}_{t+1} forgets the past oracles at a geometric rate. When $\gamma_{t+1} = 1/(t+1)$, \hat{S}_{t+1} is the empirical expectation of the oracles: $\hat{S}_{t+1} = (t+1)^{-1} \sum_{j=1}^{t+1} \mathsf{S}_j$.

2.3 SA-SSMM on few examples and Related work

We now show that SA-SSMM is related to multiple stochastic algorithms in the literature.

Stochastic Gradient Descent and Stochastic Proximal Gradient Descent. The iterative algorithm (7) applied to Example 1 defines the deterministic sequence $\theta_{t+1} = \text{prox}_{\rho g}(\theta_t - \rho \nabla f(\theta_t))$ where $\text{prox}_{\rho g}$ is the proximity operator of the proper lower semicontinuous convex function ρg . When $g \equiv 0$, it is the Gradient Descent algorithm with constant step size; and otherwise, it is the Proximal Gradient algorithm (Beck and Teboulle, 2009b; Combettes and Pesquet, 2011; Parikh and Boyd, 2014). Stochastic versions of the deterministic algorithm (7) addressing the case when $\nabla f(\theta)$ is an expectation with no closed-form expression, are proposed in the literature (see e.g. Atchadé et al. (2017)).

The iterative algorithm (8) applied to Example 1 defines the deterministic sequence

$$s_{t+1} = s_{t+1/2} - \rho \nabla f(s_{t+1/2}), \quad \text{with} \quad s_{t+1/2} := \text{prox}_{\rho g}(s_t).$$

The update mechanism in line 3 of Algorithm 1 gets into $\hat{S}_{t+1} = \hat{S}_t - \gamma_{t+1}(\hat{S}_{t+1/2} - \rho G_{t+1} - \hat{S}_t)$ where $\hat{S}_{t+1/2} := \text{prox}_{\rho g}(\hat{S}_t)$ and G_{t+1} is a random oracle for the intractable quantity $\nabla f(\hat{S}_{t+1/2})$; G_{t+1} can cover many stochastic approximations of the exact gradient possibly defined to include variance reduction schemes.

Let us discuss what the mirror sequence $\theta_t := \mathbb{T}(\hat{S}_t) = \text{prox}_{\rho g}(\hat{S}_t)$ is. By iterating line 3 in Algorithm 1, we have

$$\hat{S}_{t+1} = \Gamma_{1:t+1} \hat{S}_0 + \sum_{j=1}^t \gamma_j \Gamma_{j+1:t+1} \mathbf{S}_j + \gamma_{t+1} \mathbf{S}_{t+1}, \quad \Gamma_{r:r'} := \prod_{\ell=r}^{r'} (1 - \gamma_\ell). \quad (12)$$

From (12), we have $\theta_{t+1} = \text{prox}_{\rho g}(\tilde{\theta}_t - \rho \tilde{G}_{t+1})$ where $\tilde{\theta}_t := \sum_{j=1}^t \gamma_j \Gamma_{j+1:t+1} \theta_{j-1} + \gamma_{t+1} \theta_t$ is a weighted average of the previous mirror parameters, and $\tilde{G}_{t+1} := \Gamma_{1:t+1} \hat{S}_0 + \sum_{j=1}^t \gamma_j \Gamma_{j+1:t+1} G_j + \gamma_{t+1} G_{t+1}$ is a weighted average of the past oracles G_1, \dots, G_t and the current one G_{t+1} . As a conclusion, the mirror sequence $\{\mathbb{T}(\hat{S}_t), t \geq 0\}$ produced by Algorithm 1 is among the stochastic Proximal-Gradient algorithms, here the gradient step takes benefit of the full history of the algorithm.

Stochastic Approximation Expectation Maximization. The iterative procedure (7) applied to Example 2 defines the sequence

$$\theta_{t+1} = \text{argmax}_{\theta} \left\langle \mathbb{E}_{\pi} \left[\int S(Z, h) \mu(\mathrm{d}h | Z, \theta_t) \right], \phi(\theta) \right\rangle - g(\theta) - \psi(\theta). \quad (13)$$

When $g = 0$, this algorithm is known as the Expectation Maximization algorithm (EM) (Dempster et al., 1977) for the case of complete log-likelihood from a vector exponential family; the case $g \neq 0$ is the Maximum A Posteriori EM (also known as Bayesian EM, (McLachlan and Krishnan, 2008, Section 6.5)). The computation of the double expectation is the E-step, while the optimization step is the M-step.

EM is a popular algorithm for maximizing the likelihood of observations in a latent variable model (see e.g. McLachlan and Krishnan (2008)). We defer to Appendix C two concrete examples of EM, stressing they are instances of Example 2. In many situations, the E-step has no closed form expressions. It happens that the inner expectation w.r.t. the distribution $\mu(\mathrm{d}h | Z, \theta)$ is not explicit typically when this a posteriori distribution is known up to a normalizing constant. The outer expectation w.r.t. π also happens to be intractable: for

example in online learning, and in large batch learning when π is the empirical distribution $N^{-1} \sum_{i=1}^N \delta_{Z_i}$ with Dirac mass at the N independent examples Z_i . Stochastic versions of (13) were proposed: they consist in replacing the intractable expectations with a stochastic oracle (see e.g. the Stochastic EM [Celeux and Diebolt \(1992\)](#); and the Monte Carlo EM, [Wei and Tanner \(1990\)](#); [Fort and Moulines \(2003\)](#)).

If the deterministic sequence (7) corresponds to the usual description of the Bayesian-EM algorithm which outputs the successive M-steps, the deterministic sequence given by (8) is the output of the successive E-steps.

Algorithm 1 applied to Example 2 covers many popular stochastic EM algorithms: let us cite the Stochastic Approximation EM proposed by [Delyon et al. \(1999\)](#) when the inner expectation is not explicit; in the case $\gamma_t = 1$ for any t , the Incremental EM by [Neal and Hinton \(1998\)](#) when the outer expectation is not explicit; in the case $\gamma_t \in (0, 1)$, the Online EM proposed by [Sato and Ishii \(2000\)](#) and generalized by [Cappé and Moulines \(2009\)](#) when the outer expectation is not explicit; and many possible oracles proposed e.g. by [Chen et al. \(2018\)](#); [Fort et al. \(2021a,b\)](#); [Fort and Moulines \(2023\)](#), when possibly both inner and outer expectations are intractable and a variance reduction scheme is included.

Dictionary Learning and Matrix Factorization by Variational Surrogate. Applications of the variational surrogate setting are, among others, the Huber loss regression (e.g., [Mairal, 2015](#), Section 2.3.4), and the Dictionary Learning and Matrix Factorization (MF) problems. In the dictionary learning problem, we want to learn a dictionary $\theta \in \mathbb{R}^{p \times K}$ by solving

$$\text{Argmin}_{\theta \in \mathbb{R}^{p \times K}} \left(\mathbb{E}_{\pi} \left[\min_{h \in \mathbb{R}^K} (\mathbf{L}(Z, \theta h) + \chi(h)) \right] + g(\theta) \right), \quad (14)$$

where $\mathbf{L} : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ is a loss function, and $\chi : \mathbb{R}^K \rightarrow (-\infty, +\infty]$ and $g : \mathbb{R}^{p \times K} \rightarrow (-\infty, +\infty]$ are penalty and/or constraint functions. In MF, the vectors of coefficients h are also learnt. Again, π is the distribution of the i.i.d. stream of data $\{Z_1, Z_2, \dots\}$ in online learning, and π is the empirical distribution $N^{-1} \sum_{i=1}^N \delta_{Z_i}$ in batch learning with N i.i.d. observations $\{Z_1, \dots, Z_N\}$.

When $\mathbf{L}(z, z')$ is the quadratic loss $\|z - z'\|^2$, the problem (14) is of the form (1) and is an example of variational surrogate (see Example 3) with

$$\tilde{\ell}(Z, h, \theta) := Z^\top Z + \chi(h) - 2Z^\top \theta h + h^\top \theta^\top \theta h,$$

which is of the form $\psi(\theta) - \langle S(Z, h), \phi(\theta) \rangle + \xi(Z, h)$ by setting

$$\psi(\theta) = 0, \quad S(Z, h) := \begin{bmatrix} hh^\top \\ Zh^\top \end{bmatrix}, \quad \phi(\theta) := \begin{bmatrix} -\theta^\top \theta \\ 2\theta \end{bmatrix}, \quad \xi(Z, h) := Z^\top Z + \chi(h). \quad (15)$$

$S(Z, h)$ and $\phi(\theta)$ collect two matrices: a $K \times K$ one and a $p \times K$ one respectively. Here,

$$\mathbf{M}(Z, \tau) := \text{argmin}_{h \in \mathbb{R}^K} \left(\chi(h) + \text{Trace}(\tau^\top \tau h h^\top) - 2 \langle \tau^\top Z, h \rangle \right), \quad (16)$$

assuming that the minimizer exists and is unique. In MM-1, set $\mathcal{S} := \mathcal{M}_K^+ \times \mathbb{R}^{p \times K}$. Finally, for $s = (s^{(1)}, s^{(2)}) \in \mathcal{S}$,

$$\mathbf{T}(s) := \text{argmin}_{\theta \in \mathbb{R}^{p \times K}} \left(g(\theta) + \text{Trace}(\theta^\top \theta s^{(1)}) - 2 \text{Trace}(\theta^\top s^{(2)}) \right). \quad (17)$$

Let us describe a definition of oracles in SA-SSMM. At iteration $\#(t+1)$, choose a minibatch of observations $\{Z_i, i \in \mathcal{B}_{t+1}\}$ of size \mathbf{b} , where \mathcal{B}_{t+1} is sampled at random in $\{1, \dots, N\}$ in batch learning, or collects new observations in online learning; then, define the oracle $\mathbf{S}_{t+1} = (\mathbf{S}_{t+1}^{(1)}, \mathbf{S}_{t+1}^{(2)})$ in \mathcal{S} by

$$\mathbf{S}_{t+1}^{(1)} := \mathbf{b}^{-1} \sum_{i \in \mathcal{B}_{t+1}} H_{i,t} H_{i,t}^\top, \quad \mathbf{S}_{t+1}^{(2)} := \mathbf{b}^{-1} \sum_{i \in \mathcal{B}_{t+1}} Z_i H_{i,t}^\top, \quad \text{with } H_{i,t} := \mathbf{M}(Z_i, \theta_t). \quad (18)$$

In a specific setting of SA-SSMM for online learning, the mirror sequence $\{\theta_t, t \geq 0\}$ defined by $\theta_t := \mathbb{T}(\hat{S}_t)$ is the *online dictionary learning* proposed by Mairal et al. (2010, Algorithm 1). This setting is the case when $\chi(h) := \lambda \|h\|_1$ for some $\lambda > 0$, $g(\theta)$ is the $\{0, +\infty\}$ -valued indicator function of the set $\Theta := \{\theta = (\theta_{\cdot 1}, \dots, \theta_{\cdot K}) \in \mathbb{R}^{p \times K} : \|\theta_{\cdot k}\|^2 < 1 \text{ for all } k \in [K]\}$; $\mathbf{b} = 1$ and $\gamma_{t+1} := 1/(t+2)$. Note that in this setting, $\mathbf{M}(Z, \theta) := \operatorname{argmin}_h (\|Z - \theta h\|^2 + \lambda \|h\|_1)$ corresponds to a L^1 -regularized linear least-squares problem, an oracle for $\mathbf{M}(Z, \theta)$ is obtained by applying LARS-Lasso (Efron et al., 2004), or Proximal-Gradient Descent (see e.g. Section 6) among examples.

Structured tangent majorizing functions of $\ell(Z, \cdot)$ Assume that there exists a convex subset \mathcal{S} of \mathbb{R}^q and measurable functions $\bar{S} : \mathbb{R}^p \times \Theta \rightarrow \mathcal{S}$, $\psi : \Theta \rightarrow \mathbb{R}$ and $\phi : \Theta \rightarrow \mathbb{R}^q$ such that for any $\tau \in \Theta$, $\mathbb{E}_\pi [\|\bar{S}(Z, \tau)\|] < \infty$ and

$$\ell(z, \cdot) \leq \ell(z, \tau) + \psi(\cdot) - \psi(\tau) - \langle \bar{S}(z, \tau), \phi(\cdot) - \phi(\tau) \rangle \quad \text{on } \Theta; \quad (19)$$

see e.g. Example 2 (Eq. (34) in Appendix A.2) and Example 3 (Eq. (37) in Appendix A.3). Then the assumption MM-1 is satisfied. Assume also that MM-2 holds and consider the mirror sequence $\theta_t := \mathbb{T}(\hat{S}_t)$ of the sequence $\{\hat{S}_t, t \geq 0\}$ given by Algorithm 1. We have $\theta_{t+1} = \operatorname{argmin}_{\theta \in \mathbb{R}^d} \bar{U}_{t+1}(\theta)$ where

$$\begin{aligned} \bar{U}_{t+1}(\cdot) &:= g(\cdot) + \psi(\cdot) - \langle \hat{S}_{t+1}, \phi(\cdot) \rangle \\ &= (1 - \gamma_{t+1}) \left(g(\cdot) + \psi(\cdot) - \langle \hat{S}_t, \phi(\cdot) \rangle \right) + \gamma_{t+1} \left(g(\cdot) + \psi(\cdot) - \langle \mathbf{S}_{t+1}, \phi(\cdot) \rangle \right) \\ &= (1 - \gamma_{t+1}) \bar{U}_t(\cdot) + \gamma_{t+1} U_{t+1}(\cdot) \end{aligned}$$

where $U_{t+1}(\cdot) := g(\cdot) + \psi(\cdot) - \langle \mathbf{S}_{t+1}, \phi(\cdot) \rangle$. This discussion evidences that, in the setting (19), SA-SSMM defines iteratively a sequence of surrogate functions $\{\bar{U}_t, t \geq 0\}$.

Let us consider the case of online learning, when π in (1) is an unknown distribution but a stream of i.i.d. examples Z_1, Z_2, \dots is available, $Z_1 \sim \pi$. Define the oracle \mathbf{S}_{t+1} by $\bar{S}(Z_{t+1}, \theta_t)$. Upon noting that $\operatorname{argmin}_{\theta \in \mathbb{R}^d} \bar{U}_{t+1}(\theta)$ is invariant if a constant is added to \bar{U}_{t+1} , we can assume without loss of generality that the function U_{t+1} is tangent to the function $\theta \mapsto \ell(Z_{t+1}, \theta) + g(\theta)$ at the point θ_t : the surrogate function \bar{U}_{t+1} is updated at each iteration as a convex combination of the previous surrogate function \bar{U}_t and a novel majorizing and tangent function of the penalized loss function $\ell(Z_{t+1}, \cdot) + g(\cdot)$. This establishes that the mirror sequence $\{\theta_t, t \geq 0\}$ of the output of Algorithm 1 defines the sequence $\{\theta_t, t \geq 0\}$ given in Mairal (2013, Algorithm 1) for the minimization of $\theta \mapsto \mathbb{E}_\pi [\ell(Z, \theta)] + g(\theta)$.

3 Federated MM

In the following, we extend our scope to the Federated Learning (FL) framework in which n different *agents* (or *workers*) hold parts of the data. For $i \in [n]$, we denote by π_i the probability distribution of the data of agent $\#i$; we focus on the heterogeneous regime, in which π_i (strongly) depends on i . In the finite sample regime, let N_i be the number of examples on agent $\#i$ and $N := \sum_{i \in [n]} N_i$ be the total number of examples in the system; $\mu_i := N_i/N$ is the fraction of examples on agent $\#i$. In the online case (where each worker can access an infinite number of points), we set by default $\mu_i := n^{-1}$, the uniform weight over the n agents. The objective function can be written

$$\operatorname{argmin}_{\theta \in \mathbb{R}^d} \left(\sum_{i \in [n]} \mu_i f_i(\theta) + g(\theta) \right), \quad \text{with } f_i(\theta) := \mathbb{E}_{Z \sim \pi_i} [\ell(Z, \theta)]. \quad (20)$$

This federated setting, applied to Examples 1 to 3, corresponds to solving minimization problems by a MM approach based on respectively a quadratic surrogate, a Jensen surrogate or a variational surrogate, by using the union of examples over all workers while keeping locally each batch of observations.

3.1 Aggregate in the \mathcal{S} -space or in the Θ -space ?

Let's consider an ideal setting where all clients are always online, there are no communication constraints, and exact expectation w.r.t. π_i can be computed and the assumptions MM1&2 hold as follows: MM1 is satisfied for each function $\theta \mapsto \mathbb{E}_{\pi_i} [\ell(Z, \theta)]$, and $(\phi, \psi, S, \mathcal{S})$ do not depend on i ; MM2 is satisfied and \mathbb{T} does not depend on i .

Following the Majorize-Minimization approach, there are two ways to solve eq. (20) in a federated manner depending on how the aggregation step is performed. The first one relies on an aggregation in the parameter (Θ) space: (1) broadcast the current Θ -valued parameter θ_t to all clients, (2) compute $\theta_{t+1,i} := \mathbb{T}(\mathbb{E}_{\pi_i} [\bar{S}(Z, \theta_t)])$ at each client, (3) aggregate by setting

$$\theta_{t+1} := \sum_{i \in [n]} \mu_i \theta_{t+1,i} = \sum_{i \in [n]} \mu_i \mathbb{T}(\mathbb{E}_{\pi_i} [\bar{S}(Z, \theta_t)]), \quad \text{for } t \geq 0. \quad (21)$$

Here, the optimization steps are solved by each local agent but could also be run by the central server prior to aggregation. Note that special cases of (21) have appeared in (Marfoq et al., 2021; Wang and Chang, 2022). However, the scheme (21) no longer reduces to the centralized case and the sequence $\{\theta_t, t \geq 0\}$ may not generally converge. See Remark 1.

The second one relies on aggregation in the surrogate (\mathcal{S}) space: (1) from the current \mathcal{S} -valued parameter s_t , compute the mirror $\theta_{t+1} = \mathbb{T}(s_t)$ and broadcast it to all clients, (2) compute $s_{t+1,i} := \mathbb{E}_{\pi_i} [\bar{S}(Z, \mathbb{T}(s_t))]$ at each client, (3) aggregate by setting

$$s_{t+1} := \sum_{i \in [n]} \mu_i s_{t+1,i} = \sum_{i \in [n]} \mu_i \mathbb{E}_{\pi_i} [\bar{S}(Z, \mathbb{T}(s_t))], \quad \text{for } t \geq 0. \quad (22)$$

Note that the optimization step is solved by the central server. As a result, due to the linearity of the surrogate parameterization, this approach is equivalent to centrally solving the problem with the mixture distribution $\pi = \sum_{i \in [n]} \mu_i \pi_i$.

The following remark shows that the natural way to aggregate each worker's participation is in the space \mathcal{S} , and that aggregation should not be performed in the space Θ .

Remark 1 As a toy example, consider the real-valued loss function $\ell(Z, \theta) := Z\theta + 1/\theta$ and the constraint $g(\theta) := \iota_{\mathbb{R}_{>0}}(\theta)$ (the convex indicator of $\mathbb{R}_{>0}$) i.e. $\theta \in (0, +\infty)$; assume also that $Z > 0$ π_i -almost surely for all i . Here, the federated objective corresponds to

$$\operatorname{argmin}_{\theta \in \mathbb{R}_{>0}} \left(\sum_{i \in [n]} \mu_i \mathbb{E}_{\pi_i} [Z] \right) \theta + \frac{1}{\theta}$$

whose solution is $\theta_\star := \sqrt{1/\sum_{i \in [n]} \mu_i \mathbb{E}_{\pi_i} [Z]}$. The assumption MM-1 is satisfied for all the functions $\theta \mapsto \mathbb{E}_{\pi_i} [\ell(Z, \theta)]$ by setting $\phi(\theta) := -\theta$, $\psi(\theta) := 1/\theta$, $\bar{S}(Z, \theta) := Z$ and $\mathcal{S} := \mathbb{R}_{>0}$. The assumption MM-2 is satisfied with $\mathsf{T}(s) := 1/\sqrt{s}$. Aggregation in the \mathcal{S} -space defines a constant sequence $s_t := \sum_{i \in [n]} \mu_i \mathbb{E}_{\pi_i} [Z]$ whose mirror sequence is $\theta_t^{\text{mir}} := \mathsf{T}(s_t) = \theta_\star$. Aggregation in the Θ -space defines a constant sequence given by $\theta_t := \sum_{i \in [n]} \mu_i / \sqrt{\mathbb{E}_{\pi_i} [Z]}$ (see (21)); we have $\theta_t \neq \theta_\star$.

3.2 FedMM: a Federated Majorize-Minimization algorithm

Our idea is to develop a Federated Majorize-Minimization (FedMM) algorithm based on (22). To develop a practical algorithm, several challenges have to be overcome.

First, it is generally impossible to compute $\mathbb{E}_{\pi_i} [\bar{S}(Z, \cdot)]$, because the number of observations (on each worker) is large or even infinite; we thus have to rely on stochastic updates, for example as in Algorithm 1.

Second, the communication cost between workers is typically an important bottleneck. In the classical federated-optimization framework, two main solutions have been proposed to reduce this communication cost. One can either perform local iterations as in *LocalSGD* (Stich, 2019; Woodworth et al., 2020) and *FedAvg* (McMahan et al., 2017a), or reduce the communication cost either by performing a compression step (Alistarh et al., 2017; Konečný et al., 2016) or by sub-sampling workers at each step (referred to as Partial Participation - PP). In both situations, the data heterogeneity hinders the convergence and new algorithms have been proposed to ensure the convergence despite heterogeneity (for local iterations, see e.g. (Li et al., 2020; Karimireddy et al., 2020), for compressed FL, see e.g. (Wangni et al., 2018; Karimireddy et al., 2019; Mishchenko et al., 2019; Philippenko and Dieuleveut, 2021; Khaled and Richtárik, 2020; Richtárik et al., 2021)).

In the following, we propose a unified algorithm, FedMM that has the following features: (i) the aggregation step is performed in the space \mathcal{S} ; (ii) the algorithm uses stochastic oracles on $\mathbb{E}_{\pi_i} [\bar{S}(Z, \mathsf{T}(\hat{S}_t))]$, e.g., computed on a batch of examples at each iteration; (iii) Partial Participation (PP) is incorporated, and there is compression of the information computed by the active workers before communication to the central server. In the PP setting, data heterogeneity across clients poses a significant challenge. The global fixed point s (where $\mathbb{E}_\pi [\bar{S}(Z, \mathsf{T}(s))] = s$ i.e. $\mathbf{h}(s) = 0$) is an average over client-specific expectations, meaning that for any individual client i , it does not generally hold that $\mathbb{E}_{\pi_i} [\bar{S}(Z, \mathsf{T}(s))] = s$. If only a subset of clients participates in an update, a naive aggregation would be biased towards the participating clients' local fixed points, potentially destabilizing the algorithm. Similarly, lossy compression of client updates, especially if the compression error correlates with client drift from the global objective, can further exacerbate the impact of this underlying

Algorithm 2: Federated Majorize-Minimization FedMM

Data: $T_{\max} \in \mathbb{N}_{>0}$; $\alpha > 0$; $\mathfrak{p} \in (0, 1]$; a $(0, 1]$ -valued sequence $\{\gamma_t, t \in [T_{\max}]\}$; a distribution \mathbf{P} on the subsets of $\{1, \dots, n\}$; initial values $\hat{S}_0 \in \mathcal{S}$ and $V_{0,i} \in \mathbb{R}^q$ for all $i \in [n]$; a probability distribution $\{\mu_i, i \in [n]\}$.

Result: Sequence: $\{\hat{S}_t, t \in [T_{\max}]\}$

- 1 Set $V_0 = \sum_{i=1}^n \mu_i V_{0,i}$
- 2 **for** $t = 0, \dots, T_{\max} - 1$ **do** // for all iterations
- 3 Sample $\mathcal{A}_{t+1} \sim \mathbf{P}$
- 4 Send \hat{S}_t and $\mathbb{T}(\hat{S}_t)$ to the active workers
- 5 **for** $i \in \mathcal{A}_{t+1}$ **do** // on active worker #i
- 6 Sample $S_{t+1,i}$, oracle for $\mathbb{E}_{\pi_i}[\bar{S}(Z, \mathbb{T}(\hat{S}_t))]$
- 7 Set $\Delta_{t+1,i} = S_{t+1,i} - \hat{S}_t - V_{t,i}$
- 8 Set $V_{t+1,i} = V_{t,i} + \alpha/\mathfrak{p} \text{Quant}_{t+1,i}(\Delta_{t+1,i})$.
- 9 Send $\text{Quant}_{t+1,i}(\Delta_{t+1,i})$ to the central server
- 10 **for** $i \notin \mathcal{A}_{t+1}$ **do** // on inactive worker #i
- 11 Set $V_{t+1,i} = V_{t,i}$ // no update
- 12 // on the central server
- 13 Set $H_{t+1} = V_t + 1/\mathfrak{p} \sum_{i \in \mathcal{A}_{t+1}} \mu_i \text{Quant}_{t+1,i}(\Delta_{t+1,i})$
- 14
- 15 Compute $\hat{S}_{t+1/2} = \hat{S}_t + \gamma_{t+1} H_{t+1}$
- 16 Project it on \mathcal{S} , given $\mathbf{B}_t \succ 0$: $\hat{S}_{t+1} = \operatorname{argmin}_{s \in \mathcal{S}} \left(s - \hat{S}_{t+1/2} \right)^\top \mathbf{B}_t \left(s - \hat{S}_{t+1/2} \right)$
- 17 Set $V_{t+1} = V_t + \alpha/\mathfrak{p} \sum_{i \in \mathcal{A}_{t+1}} \mu_i \text{Quant}_{t+1,i}(\Delta_{t+1,i})$

heterogeneity by distorting the aggregated information. To counteract this, (iv) control variates are introduced. These client-specific variables track and correct for the deviation of local updates from a global perspective, effectively acting as a form of client-wise variance reduction. This mechanism helps to mitigate the adverse effects of heterogeneity when employing PP or compression, as explored in works like (Mishchenko et al., 2019; Dieuleveut et al., 2021). Our approach takes its roots in DIANA (Mishchenko et al., 2019), an algorithm for heterogeneity mitigation for federated optimization, and in FedEM (Dieuleveut et al., 2021), a specific instance of the general MM framework addressed in this paper.

Description of FedMM, Algorithm 2. At iteration $(t + 1)$, a set of active agents \mathcal{A}_{t+1} is selected at random from a probability distribution \mathbf{P} on the subsets of $\{1, \dots, n\}$, such that the mean number of active workers at each iteration is $n\mathfrak{p}$ ($\mathfrak{p} \in (0, 1]$). The server broadcasts back \hat{S}_t and its mirror $\mathbb{T}(\hat{S}_t)$ to the active workers. It runs the optimization step \mathbb{T} ; in practice, the computation of the map \mathbb{T} at line 4 can be replaced by an approximated optimization step, e.g., a few gradient steps. Nonetheless, our convergence analysis requires the exact computation of \mathbb{T} (see Theorem 1).

Then, each active agent $\#i$ samples a local surrogate oracle $S_{t+1,i}$ of $\mathbb{E}_{\pi_i}[\bar{S}(Z, \mathbb{T}(\hat{S}_t))]$. The agent's goal is to construct and communicate a low-variance, compressed estimate of its

contribution to the global update direction, which is related to the mean field $\mathfrak{h}(\hat{S}_t)$ from (9). To achieve this, the agent first computes its local update estimate, $S_{t+1,i} - \hat{S}_t$. However, due to data heterogeneity, this local estimate is a biased estimator of the global mean field $\mathfrak{h}(\hat{S}_t)$. To correct for this client drift, a control variate $V_{t,i}$ is subtracted. This results in the error-corrected vector $\Delta_{t+1,i} = S_{t+1,i} - \hat{S}_t - V_{t,i}$ (line 7). It is this locally-debiased quantity that is then compressed using a (potentially random) operator $\text{Quant}_{t+1,i}$ and sent to the server. These local subtractions are critical as they construct a more accurate update vector *before* the lossy compression step. The client’s control variate is then updated using the compressed information (line 8), with a step size $\alpha > 0$. Inactive agents (i.e. agents $\#i$ with $i \notin \mathcal{A}_{t+1}$) do not perform any updates.

Finally, the central server aggregates all contributions and compensates the local subtraction of control variates by adding V_t (line 13): it computes H_{t+1} , an oracle of $\mathbb{E}_\pi \left[\bar{S}(Z, \mathbb{T}(\hat{S}_t)) \right] - \hat{S}_t$ (line 13). Then it performs an SA-type update step on the iterate and obtains $\hat{S}_{t+1/2}$ (see line 15) and projects back onto the space \mathcal{S} (line 16). Crucially, this projection may not be performed according to the Euclidean geometry but possibly according to the geometry defined through a positive definite matrix B_t , allowed to depend on the iteration index t and on \hat{S}_t . In our experiments, we will simply set $B_t = I$, the identity matrix, recovering the Euclidean geometry; nonetheless, our convergence result (see Theorem 1) requires a finer specification of this matrix.

4 MM within FL: Related works in the framework MM1&2

The crucial difference between FedMM and competing algorithms in the literature is that the aggregation steps are performed in the \mathcal{S} -space (see Section 3.1). In most federated algorithms to perform personalization, EM, or matrix factorization, aggregation steps are performed in the initial parameter space Θ .

Federated Composite Optimization. The general optimization problem (1), structured as $\text{argmin}_{\theta \in \mathbb{R}^d} f(\theta) + g(\theta)$, is a form of composite optimization. In instances where f is smooth and g is a proper, lower semi-continuous convex function, our MM framework, as exemplified by the quadratic surrogate in Example 1, naturally leads to updates where the $\mathbb{T}(s)$ operator becomes the proximal operator of $g(\theta)$. In the literature on federated optimization, methods like ProxSkip/Scaffnew by Mishchenko et al. (2022) and RandProx/RandProx-FL by Condat and Richtárik (2023) investigate scenarios where regularization functions and their proximal operators are used primarily to enforce the consensus mechanism among clients. These works propose performing the proximal step (which ensures consensus) periodically, rather than at every communication round, to reduce communication overhead. While their primary focus is on consensus, this framework can be conceptually extended: the periodic server-side proximal step, which enforces consensus, shares a methodological similarity with our FedMM, where the \mathbb{T} operator (which itself could be a proximal operator for $g(\cdot)$ in specific MM instances - see e.g. Example 1) is also evaluated centrally at the server, and not as part of the local client activity. However, the fundamental premise of these methods, even when extended, remains distinct from our MM-based surrogate aggregation, as their core design centers on periodic randomized consensus enforcement to reduce communication overhead rather

than the iterative minimization of global majorizing functions derived from client-specific surrogate parameters.

A different perspective on aggregation for composite problems is offered by Yuan et al. (2021). They observe that for problems with sparsity-inducing regularizers, such as $g(\cdot)$ being the ℓ_1 norm, if clients compute their local model updates (involving a local client-side proximal step for g) and these updated model parameters are then averaged in the Θ -space, the resulting global model may lose its intended sparsity. As a remedy, they propose FedDualAvg, where aggregation occurs in what they term the "dual space". In the context of ℓ_1 -regularization, the dual variables often correspond to the subgradient of the ℓ_1 -norm, which directly relates to the parameters defining a particular type of surrogate or upper bound on the regularizer. Aggregating these dual variables is thus conceptually akin to aggregating parameters in a surrogate space \mathcal{S} , aligning with our paper’s methodological argument against naive Θ -space aggregation when a more natural, structure-preserving aggregation space exists. Our FedMM framework provides a more general and direct approach by explicitly identifying and aggregating in the functional space of linearly parameterized functions.

Federated EM. The Expectation-Maximization (EM) algorithm (see Section 2.3 and Appendix C) has been adapted by several works for the federated learning setting. Marfoq et al. (2021) considered a federated setting where the data at each agent comes from a mixture of L underlying distributions. These L distributions are shared among clients, and data heterogeneity arises from each client having a different mixture weight. The primary motivation is to facilitate federated personalization by assigning each client a distinct mixture weight and, consequently, varying their models. The authors demonstrate that several federated personalization approaches can be viewed as specific instances of their formulation, such as (Canh et al., 2020; Dinh et al., 2021; Hanzely et al., 2020). Most interestingly, the FedEM algorithm of Marfoq et al. (2021) is presented as a particular instance of a more general Federated Surrogate optimization framework (Algorithm 3 and Appendix F, Marfoq et al., 2021). However, in their approach, the parametric structure of the surrogate space (i.e., the E-step sufficient statistics) is leveraged only locally by each client for their surrogate optimization step with a locally (depending only on local client data) executed approximation of the T map. The subsequent aggregation of updates occurs in the original parameter space Θ , rather than in the surrogate space \mathcal{S} . Marfoq et al. (2021) provide convergence guarantees that assume a bound on the heterogeneity between client data distributions (Assumption 7, Marfoq et al., 2021).

In contrast, the Federated EM algorithm proposed by Dieuleveut et al. (2021) directly employs \mathcal{S} -space aggregation. Their method introduces and analyzes a federated Majorization-Minimization algorithm specifically for EM, where Assumption MM-1 is satisfied with an unconstrained surrogate space $\mathcal{S} = \mathbb{R}^q$. Consequently, their algorithm can be seen as a specific instance of our FedMM without the projection step (line 16 in Algorithm 2).

More recently, Tian et al. (2024) (with FedGrEM) tackled a similar formulation to Marfoq et al. (2021) and extended the federated EM framework to manage adversarial data and outlier tasks. Their "federated gradient EM" involves clients executing local E-steps to compute sufficient statistics, which are then used in local gradient-based M-steps (approximating a local version of the T operator at client $\#i$, depending on client’s data distribution.). The

resulting client-specific model parameters (in Θ) are subsequently aggregated at the server, differing from FedMM’s central \mathbb{T} operation post- \mathcal{S} -aggregation. Zhixu et al. (2024) studied federated EM for Mixtures of Linear Regressions (FMLR), characterizing convergence rates based on client-to-data ratios. While focused on statistical rates for this specific model, their analysis (Zhixu et al., 2024, Propositions 2 and 3) implicitly involves \mathcal{S} -space aggregation of sufficient statistics.

Federated Matrix Factorization (FMF). In batch learning, FMF consists in solving (see (14))

$$\operatorname{argmin}_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^n \sum_{j=1}^{N_i} \operatorname{argmin}_{h_{ij} \in \mathcal{H}_i} \{ \|Z_{ij} - \theta h_{ij}\|^2 + \chi(h_{ij}) \} + g(\theta), \quad (23)$$

where $Z_{ij} \in \mathbb{R}^p$ is the j -th observation of the client $\#i$, N_i is the number of observations available at agent $\#i$, $\Theta \subseteq \mathbb{R}^{p \times K}$, $\mathcal{H}_i \subseteq \mathbb{R}^K$, and χ, g are penalty terms on the coefficients h and on the dictionary θ respectively; $N := \sum_{i=1}^n N_i$. The problem of FMF has attracted a lot of attention (Flanagan et al., 2020; Gao et al., 2020; Hegedús et al., 2019; Singhal et al., 2021), an application being federated clustering (clustering agents *or* clustering observations); such clusters may constitute another way to perform personalization (e.g., Mansour et al., 2020; Sattler et al., 2020; Ghosh et al., 2020).

As discussed in Section 2.3 (see (15) to (18)), the assumptions MM1-2 are satisfied and FedMM applies: at iteration $\#t$ given the current model $\theta_t := \mathbb{T}(\hat{S}_t)$, the active local agents compute an oracle of the matrices $N_i^{-1} \sum_{j=1}^{N_i} h_{ij}^{*,t} (h_{ij}^{*,t})^\top$ and $N_i^{-1} \sum_{j=1}^{N_i} Z_{ij} (h_{ij}^{*,t})^\top$ where $h_{ij}^{*,t}$ solves

$$\operatorname{argmin}_{h \in \mathcal{H}_i} (\|Z_{ij} - \theta_t h\|^2 + \chi(h)). \quad (24)$$

FedMGS, introduced by Wang and Chang (2022) with an application to unsupervised data clustering via K -means through specific choices of \mathcal{H}_i , χ and g , is an alternative to FedMM. In FedMGS, the oracle of the matrices is obtained by substituting the exact solution $h_{ij}^{*,t}$ for the output of few iterations of a gradient algorithm designed to solve (24). FedMGS is based on *gradient sharing* principles: the central server aggregates the oracles of the matrices so that FedMGS aggregates in the \mathcal{S} -space similarly to FedMM. Contrary to FedMM, FedMGS does not include a control variate technique and a quantization step run by local agents, but the central server carries out a control variate scheme for heterogeneity reduction; the selected number of agents at each iteration is constant; an agent which is not selected still updates locally the vectors h_{ij} ; the local agent ensures that the quantities sent to the central server are in \mathcal{S} . For both FedMM and FedMGS, the optimization step MM-2 is run by the central server; in FedMGS, \mathbb{T} is approximated by many iterations of a projected gradient algorithm.

Similar Methodologies: Aggregation of Functional Approximations. The principle of aggregating functional parameters extends beyond our specific Majorize-Minimization framework to the more general setting of surrogate optimization, which relaxes our MM setting by allowing surrogate functions to not necessarily be global majorizing functions. From our methodological standpoint, the key to designing federated algorithms lies in preserving the structure of the underlying surrogate-based method. An ideal federated surrogate algorithm, under perfect conditions (e.g., full client participation, no communication constraints, exact computation of \mathbb{T} operator), should reduce exactly to its centralized counterpart. Our linearly

parameterized MM framework achieves this naturally: aggregating the surrogate parameters in the space \mathcal{S} is equivalent to averaging the surrogate functions themselves, ensuring the central optimization step operates on a true global surrogate.

This perspective helps contextualize other methods that follow a similar philosophy. For example, second-order (Newton-type) methods in FL, as explored by [Islamov et al. \(2021\)](#), can be viewed through this lens. Clients compute local gradients and Hessians, which are the parameters defining local quadratic approximations of the objective function. These functional parameters are sent to the server and aggregated to form a global quadratic model. Only then is the optimization step, the Newton update, is performed on the server using this aggregated information. This process is directly analogous to our approach of aggregating parameters of local majorizing functions (our surrogates) in \mathcal{S} -space before the central minimization step is performed using the \mathbb{T} operator.

5 Theoretical Analysis and Open Problems

We now proceed to provide a convergence analysis along with the required assumptions.

Stochasticity and Compression: We assume that the oracles $\{\mathbf{S}_{k+1,i}, k \geq 0\}$ available at each local agent $\#i$ satisfy the following assumption. First, the local oracles are independent between local agents; and second, given an agent $\#i$, the oracles are unbiased with a variance allowed to be specific to each agent.

A3 For all $t \geq 0$, the oracles $\{\mathbf{S}_{t+1,i}, i \in [n]\}$ are conditionally independent of the past history of the algorithm \mathcal{F}_t ¹. Moreover, for any $i \in [n]$, $\mathbb{E}[\mathbf{S}_{t+1,i} | \mathcal{F}_t] = \mathbb{E}_{\pi_i}[\bar{S}(Z, \mathbb{T}(\hat{S}_t))]$ and there exists $\sigma_i^2 > 0$ such that for any $k \geq 0$, $\mathbb{E}[\|\mathbf{S}_{t+1,i} - \mathbb{E}_{\pi_i}[\bar{S}(Z, \mathbb{T}(\hat{S}_t))]\|^2 | \mathcal{F}_t] \leq \sigma_i^2$.

Next, we make the following assumption on the compression operator and on partial participation. In the analysis, we consider a *specific* model on the partial participation, that assumes that the agents are selected independently at random at each step.

A4 (ω) For all $s \in \mathbb{R}^q$, the random variables $\{\text{Quant}_{t+1,i}(s), t \geq 0, i \in [n]\}$ are i.i.d. with the same distribution as $\text{Quant}(s)$. There exists $\omega \geq 0$ such that for all $s \in \mathbb{R}^q$,

$$\mathbb{E}[\text{Quant}(s)] = s, \quad \mathbb{E}[\|\text{Quant}(s) - s\|^2] \leq \omega \|s\|^2. \quad (25)$$

Each compression is applied independently of any other randomness in the past of the algorithm.

The case $\omega = 0$ corresponds to the case of no compression.

A4(ω) encompasses operators that have been widely used in the literature, for example Sparsification ([Wangni et al., 2018](#)) and random (scalar or vector) quantization ([Alistarh et al., 2017](#); [Dai et al., 2019](#); [Leconte et al., 2021](#)). In particular (25) describes the individual behavior of each compression operator, specifically the fact that it is unbiased and has a relatively bounded variance (i.e., the variance of $\text{Quant}(s)$ scales as $\omega \|s\|^2$). See e.g ([Dieuleveut et al., 2021](#), Supplemental Section B) for the block- p -quantization example.

1. see Appendix D.3 for a rigorous definition of the filtration.

A5 (p) *At each step, each worker is selected with probability \mathbf{p} , $\mathbf{p} \in (0, 1]$, independently of the past and of the other workers.*

Independent compression steps in A4(ω) and independent selection in A5(\mathbf{p}) allow us to tackle the workers' selection as a particular form of unbiased compression. It is shown in Appendix D.2 that compression and partial participation are equivalent: partial participation can be modeled as another compression operation on top of Quant, which sends the output of Quant with probability \mathbf{p} and nothing with probability $1 - \mathbf{p}$.

Model. We state our regularity assumptions on the functions $\mathbf{h}_i(\cdot)$ for all $i \in [n]$ on the set \mathcal{S} , and on $\phi(\mathbf{T}(\cdot))$.

A6 (i) *For all $i \in [n]$, the function $\mathbf{h}_i(s) := \mathbb{E}_{\pi_i} [\bar{S}(Z, \mathbf{T}(s))] - s$ is globally L_i -Lipschitz on \mathcal{S} . (ii) $\sup_{\mathcal{S}} \|\mathbf{h}(\cdot)\| < \infty$.*

A7 (i) *There exist positive constants C_\star and $C_{\star\star}$ and for any $s \in \mathcal{S}$, there exists a $q \times q$ positive-definite matrix $\mathbf{B}(s)$ such that for any $s' \in \mathcal{S}$,*

$$\begin{aligned} \|\phi(\mathbf{T}(s')) - \phi(\mathbf{T}(s)) - \mathbf{B}(s)(s' - s)\| &\leq C_\star \|s' - s\|^2, \\ |\langle s' - s, \phi(\mathbf{T}(s')) - \phi(\mathbf{T}(s)) \rangle| &\leq C_{\star\star} \|s' - s\|^2. \end{aligned}$$

(ii) *There exist $0 < v_{\min} \leq v_{\max}$ such that for any $s \in \mathcal{S}$ and $u \in \mathbb{R}^q$, $v_{\min} \|u\|^2 \leq u^\top \mathbf{B}(s) u \leq v_{\max} \|u\|^2$.*

We verify A7 for a few examples, instances of the general settings described in Example 1 and Example 2. See the discussions in Appendix E.

The projection. Finally, we proceed to specify the required assumptions on the projection in Algorithm 2 line 16.

A8 (i) \mathcal{S} *is a closed convex subset of \mathbb{R}^q . (ii) Choose $\mathbf{B}_t := \mathbf{B}(\hat{S}_t)$ in FedMM, where \mathbf{B} is given by A 7.*

To facilitate our discussions later, we set

$$\Pi(\hat{s}, \mathbf{B}) := \operatorname{argmin}_{s \in \mathcal{S}} (s - \hat{s})^\top \mathbf{B} (s - \hat{s}), \quad \hat{s} \in \mathbb{R}^q; \quad \|u\|_{\mathbf{B}} := \sqrt{u^\top \mathbf{B} u}. \quad (26)$$

Finite time analysis. We now provide an explicit finite time analysis convergence of FedMM.

To quantify the convergence of the latter, we concentrate on the control of

$$\bar{\mathcal{E}}_{t+1} := \frac{\|\Pi(\hat{S}_t + \gamma_{t+1} \mathbf{h}(\hat{S}_t), \mathbf{B}_t) - \hat{S}_t\|_{\mathbf{B}_t}^2}{\gamma_{t+1}^2}$$

which is equal to $\|\mathbf{h}(\hat{S}_t)\|_{\mathbf{B}_t}^2$ when $\Pi(s, \mathbf{B}) = s$ and is lower bounded by $v_{\min} \|\mathbf{h}(\hat{S}_t)\|^2$ under A 7; $\bar{\mathcal{E}}_{t+1}$ measures how far the current iterate \hat{S}_t is from the set $\{s \in \mathcal{S} : \mathbf{h}(s) = 0\}$. Two cases are considered in turn: first, the case when there is no compression ($\omega = 0$ in A4) and the full participation regime is considered ($\mathbf{p} = 1$ in A5); second, the general case when $\omega > 0$ or $\mathbf{p} \in (0, 1)$. The proof of Theorem 1 is in Appendix D.

Theorem 1 Assume A1 to 8. Let $\{\hat{S}_t, t \geq 0\}$ be the sequence given by FedMM. Set

$$\omega_p := \omega + (1 + \omega) \frac{1 - p}{p}, \quad \bar{C} := C_{**} + \sup_S \|h(\cdot)\| C_*, \quad \sigma^2 := n \sum_{i=1}^n \mu_i^2 \sigma_i^2, \quad L^2 := n \sum_{i=1}^n \mu_i^2 L_i^2.$$

(A) Case $\omega_p = 0$. Run FedMM with $\alpha \in [0, 1]$ and step sizes $\gamma_{t+1} \in (0, v_{\min}/(2\bar{C}))$. For any $T \in \mathbb{N}_{>0}$

$$\sum_{t=1}^T \frac{\gamma_t}{4} \left(1 - \frac{2\bar{C}}{v_{\min}} \gamma_t\right) \mathbb{E} [\bar{\mathcal{E}}_t] \leq \mathbb{E} \left[(f + g)(\mathbb{T}(\hat{S}_0)) - \min(f + g) \right] + v_{\max} \frac{\sigma^2}{n} \sum_{t=1}^T \gamma_t.$$

(B) Case $\omega_p > 0$. Run FedMM with $\alpha \in (0, 1/(1 + \omega_p)]$ and step sizes satisfying

$$\gamma_{t+1} \leq \gamma_t \leq \gamma_1 \quad 1 - \frac{2\bar{C}}{v_{\min}} \gamma_t - 64 \frac{\omega_p v_{\max} L^2}{\alpha^2 v_{\min} n} \gamma_t^2 > 0, \quad \gamma_1 \leq \frac{\alpha}{4\sqrt{\omega_p}} \frac{\sqrt{v_{\min}}}{\sqrt{v_{\max}}} \frac{\sqrt{n}}{L}.$$

For any $T \in \mathbb{N}_{>0}$,

$$\begin{aligned} & \sum_{t=1}^T \frac{\gamma_t}{4} \left(1 - \frac{2\bar{C}}{v_{\min}} \gamma_t - 64 \frac{\omega_p v_{\max} L^2}{\alpha^2 v_{\min} n} \gamma_t^2\right) \mathbb{E} [\bar{\mathcal{E}}_t] \\ & \leq \mathbb{E} \left[(f + g)(\mathbb{T}(\hat{S}_0)) - \min(f + g) \right] + \frac{4\omega_p v_{\max}}{\alpha} \gamma_1 \sum_{i=1}^n \mu_i^2 \mathbb{E} \left[\|V_{0,i} - \mathbf{h}_i(\hat{S}_0)\|^2 \right] \\ & \quad + 2v_{\max}(1 + 5\omega_p) \frac{\sigma^2}{n} \sum_{t=1}^T \gamma_t. \end{aligned} \tag{27}$$

In the case the step sizes are constant ($\gamma_t = \gamma$ for all t), we obtain the following corollary. Here, γ is chosen small enough to ensure that the LSH is positive.

Corollary 1 (of Theorem 1, when $\omega_p > 0$.) Let $\gamma > 0$ be such that

$$1 - \frac{2\bar{C}}{v_{\min}} \gamma - 64 \frac{\omega_p v_{\max} L^2}{\alpha^2 v_{\min} n} \gamma^2 \geq \frac{1}{2}.$$

Let τ be a uniform random variable on $\{1, \dots, T\}$, independent of the path of the algorithm.

$$\begin{aligned} \mathbb{E} [\bar{\mathcal{E}}_\tau] & \leq \frac{8}{T\gamma} \mathbb{E} \left[(f + g)(\mathbb{T}(\hat{S}_0)) - \min(f + g) \right] + \frac{32\omega_p v_{\max}}{T\alpha} \sum_{i=1}^n \mu_i^2 \mathbb{E} \left[\|V_{0,i} - \mathbf{h}_i(\hat{S}_0)\|^2 \right] \\ & \quad + 16v_{\max}(1 + 5\omega_p) \frac{\sigma^2}{n}. \end{aligned}$$

(Parameters and constants). Constants σ^2 (resp. L^2) are defined so that when $\mu_i = 1/n$ (i.e., in batch learning with equal number of data at each worker, or in online learning), and σ_i (resp. L_i) are constant, then $\sigma = \sigma_i$ (resp. $L = L_i$). The step sizes γ_t have to be chosen smaller than an absolute constant depending on the regularity of the problem and the compression factor. Theorem 1 does not require any convexity assumption on f in MM-1 though MM-2 is satisfied when the surrogate is strongly convex even if not explicitly assumed.

(*Convergence guarantee*). Controlling the quantity at a random stopping time τ , uniformly chosen over all iterations is classical to obtain convergence rates in non-convex optimization (Ghadimi and Lan, 2013). The bound can be decomposed as the sum of an initial quantity, divided by the number of iterations, and a noise term, divided by the number of workers and the batch size \mathbf{b} .

(*Heterogeneity*). The result is valid without any assumption on the data heterogeneity ($\pi_i \neq \pi_j$). In the batch learning case, a natural choice for the initial control variates is $V_{0,i} = \mathbf{h}_i(\hat{S}_0)$ (at the cost of evaluating the expectation once). Complexity guarantee is then not affected (thus robust to) by the heterogeneity.

(*Case $\alpha = 0$*). When $\omega_p = 0$, we can choose $\alpha = 0$: control variates play no role even in the heterogeneity setting. When $\omega_p > 0$, the bound is obtained under the assumption that $\alpha > 0$. Finite time analysis can be run in the case $\omega_p > 0, \alpha = 0$ but they provide bad conditions on the step sizes γ_t and poor upper bounds (see Appendix D.7). We also illustrate numerically that, in the heterogeneity setting, $\alpha = 0$ may lead to a worsened convergence rate, see Figure 2 in Section 6.

Challenges in the analysis. While Theorem 1 provides a first result, the Federated MM framework that we introduce raises numerous theoretical challenges.

- First and foremost, the surrogate space \mathcal{S} can be constrained (see, e.g., Examples 2 and 3 in Appendix E). The updates resulting from the aggregation of the workers' contributions may not stay within \mathcal{S} . We are thus required to perform a projection step. This issue is reinforced by variance reduction schemes, that often result in performing non-convex combinations (e.g., the difference of a candidate surrogate with a control variate). It is also reinforced by the compression step, the choice of the compression operator Quant may be impactful; for example, a scalar quantization of a positive semidefinite matrix is not ensured to be a positive semidefinite matrix. Yet, these variance reduction schemes and compression steps are necessary to reduce the impact of heterogeneity and the computational cost, and helpful to accelerate convergence.

- In order to guarantee that $\mathbb{T}(\hat{S}_{t+1})$ is well defined, FedMM requires a projection of the point $\hat{S}_t + \gamma_{t+1}H_{t+1}$ on \mathcal{S} , yielding \hat{S}_{t+1} . Quite often, this projection step is applied to ensure numerical stability but is not taken into account when deriving a theoretical analysis of the algorithm. On the contrary, the projection is part of our convergence analysis provided in Theorem 1. Nevertheless, our current analysis requires the projection to be performed w.r.t. the geometry of the matrix $\mathbf{B}(\hat{S}_t)$ which can be interpreted as a local approximation of the geometry of the mapping $s \mapsto \phi(\mathbb{T}(s))$ - see A7.

- From an analytical standpoint, the projection introduces bias to the oracle (roughly speaking, $\mathbb{E}[\Pi(\hat{S}, \mathbf{B})] \neq \Pi(\mathbb{E}[\hat{S}], \mathbf{B})$). This bias makes the analysis more complicated and degrades the convergence rate (the limit variance obtained in Theorem 1 actually resembles the one in (Gorbunov et al., 2021, Eq. (41)) for biased online non-convex federated optimization). Moreover, the (convex) space \mathcal{S} can also be considered as a Riemannian manifold: the projections could also be performed w.r.t. this geometry.

6 Dictionary Learning Experiments

We apply the FedMM algorithm to the following stochastic dictionary learning problem

$$\operatorname{argmin}_{\theta \in \mathbb{R}^{p \times K}} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\pi_i} \left[\min_{h \in \mathbb{R}^K} \left\{ \frac{1}{2} \|Z - \theta h\|^2 + \lambda \|h\|_1 \right\} + \eta \|\theta\|^2 \right], \quad (28)$$

where $Z \in \mathbb{R}^p$ is the random observation, $\|\cdot\|_1$ denotes the L_1 -norm, and $\lambda, \eta > 0$ are regularization parameters. As detailed in Section 2.3, this problem fits our variational surrogate framework (Example 3) and FedMM can be specialized by: (1) The surrogate space for aggregation is set to $\mathcal{S} := \mathcal{M}_K^+ \times \mathbb{R}^{p \times K}$; (2) At each iteration, a participating client receives $\mathbb{T}(\hat{S}_t)$ and uses it to compute its local surrogate oracle $\mathbf{S}_{t+1,i} = (\mathbf{S}_{t+1,i}^{(1)}, \mathbf{S}_{t+1,i}^{(2)})$ - see (18).

We consider three data settings, two synthetic and one on real data. For the synthetic setting, we consider two scenarios: homogeneous distributions $\pi_i = \pi$ and heterogeneous distributions $\pi_i \neq \pi_j$. Both start by generating a matrix $\theta_* \in \mathbb{R}^{p \times K}$ such that $(\theta_*)_{ij} \sim \mathcal{N}(0, 1)$. Then, we generate the set $\{Z_t := \theta_* h_t\}_{t=1}^{\text{tot}}$, where, for each t , h_t is a sparse vector with 20% of randomly selected non-zero entries, distributed as $\mathcal{N}(0, 1)$. For the scenario where clients have homogeneous distributions, we set $\text{tot} = 250$ and give each client a copy of the full data. For the heterogeneous case, with $\text{tot} = 5000$, constrained k-means (Bradley et al., 2000) is used to split $\{Z_t\}_{t=1}^{\text{tot}}$ into n equal clusters while maximizing the average distance between different clusters. Then, each client is assigned a unique cluster. For our experiments, we used $n = 20$, $K = 15$, $\lambda = 0.1$, and $\eta = 0.2$.

For the real data, we used the MovieLens 1M dataset (Harper and Konstan, 2015), which consists of the movie ratings of 6000 users for 4000 movies. We subsample the dataset to have 5000 user ratings for 500 movies. Thus, we recover a dataset of 5000 vectors embedded in \mathbb{R}^p , for $p = 500$. To generate the federated clients' data, we split the 5000 vectors, using constrained k-means, among $n = 20$ clients and set the data distributions π_i for client i to be the uniform distribution over the $N_i = 250$ vectors received. Finally, we set the embedding dimension K to 50. Similar to the synthetic setting, we used $\lambda = 0.1$, and $\eta = 0.2$.

For all experiments, we used a decaying step size $\gamma_t = \beta / \sqrt{\beta + t}$ with β tuned for each algorithm and setting from the range $[0.001, 0.05]$. For the projection step, line (16), we set the projection matrix \mathbf{B}_t to the identity. Nonetheless, note that our theoretical analysis requires another selection of \mathbf{B}_t (see A7 and Appendix E.3).

To evaluate our methodology, we compare it with a naive algorithm, where at step $t + 1$, each active agent $\#i$ computes a new surrogate and optimizes it to obtain a new local parameter $\theta_{t+1}^i \in \Theta$. Then, it communicates θ_{t+1}^i to the server, which aggregates the parameters of active agents to construct θ_{t+1} . This naive algorithm exactly mirrors FedMM, except that the communications and the server aggregation step occur in the parameter space and not in the surrogate space.

Besides tracking the loss value in (28), we report two other metrics. First, we track the normalized surrogate update squared norm: at iteration $t + 1$, for FedMM, we report $\mathcal{E}_{t+1}^s := \|\hat{S}_{t+1} - \hat{S}_t\|^2 / \gamma_{t+1}^2$. Meanwhile, for the naive algorithm aggregating in the parameter space, we report $\mathcal{E}_{t+1}^{s,p} := \|\mathcal{T}(\theta_{t+1}) - \mathcal{T}(\theta_t)\|^2 / \gamma_{t+1}^2$, where $\mathcal{T}(\theta) := 1/n \sum_{i \in [n]} \mathbb{E}_{\pi_i} [\bar{S}(Z, \theta)]$. Similarly, we also track the parameter space update squared norm. In particular, for the algorithm aggregating in the parameter space, we track $\mathcal{E}_{t+1}^p := \|\theta_{t+1} - \theta_t\|^2 / \gamma_{t+1}^2$. Meanwhile, for FedMM, we track

$\mathcal{E}_{t+1}^{p,s} := \|\tau(\hat{S}_{t+1}) - \tau(\hat{S}_t)\|^2 / \gamma_{t+1}^2$. For all the figures, we report the average over 10 independent runs, obtained with different seeds.

Impact of Aggregation Space: Our first set of experiments is reported in Figure 1. We set the number of participating clients at each iteration to be 10 ($p = 0.5$). Each active client computes $S_{t+1,i}$ using 50 examples, sampled at random among the local examples. We considered a setting where the communication between the clients and the server is compressed using quantization with 8 bits per coordinate. We set the control variates stepsize to $\alpha = 0.01$.

First, for the objective value (28), we observe that, when using FedMM, it monotonically decays for all settings. Meanwhile, for the naive algorithm aggregating in the parameter space, it diverges in the synthetic heterogeneous setting. In addition, FedMM achieves better objective value across all settings of data.

For \mathcal{E}_{t+1}^p and $\mathcal{E}_{t+1}^{p,s}$, both the naive algorithm and FedMM decay. For the synthetic data, we observe faster decay for the algorithm designed to aggregate in the parameter space. Meanwhile, for the MovieLens dataset, FedMM initially achieves faster decay but is eventually beaten by the naive algorithm.

Finally, for \mathcal{E}_{t+1}^s and $\mathcal{E}_{t+1}^{s,p}$, FedMM decays across all datasets. Meanwhile, the naive algorithm diverges on synthetic data. In addition, for MovieLens dataset, the naive algorithm initially decays steadily but becomes more unstable as the number of communication rounds increases. We note that \mathcal{E}_{t+1}^s mirrors the quantity controlled by Theorem 1 (up to a different projection geometry). Thus, the experimental results mirror the theoretical conclusions on convergence up to an error term. In addition, we highlight that even though the naive algorithm converges w.r.t. \mathcal{E}_{t+1}^p , it diverges in the surrogate space (see $\mathcal{E}_{t+1}^{s,p}$).

Impact of Control Variates: Our second set of experiments aimed to investigate the effect of control variates in mitigating the consequences of data heterogeneity and partial participation; they are presented in Figure 2. In this set, we only run FedMM. To investigate and isolate the effect of control variates, we initialize the control variates to zero, i.e. $V_{0,i} = 0$ and set their update step size to $\alpha \in \{0, 0.01\}$. As a result, when $\alpha = 0$, control variates are not used. At each step t , 10 agents are active, i.e. $p = 0.5$. To remove the effect of stochasticity, apart from partial participation, we make each active agent $\#i$ use all their local examples, i.e. $S_{t+1,i} = \mathbb{E}_{\pi_i} [\bar{S}(Z, \theta_t)]$.

Across all data settings, we observe no effect of using control variates on the objective value. For the synthetic homogeneous dataset and $\alpha = 0.01$, given that all the clients share the same data and that there is no communication compression, the control variates exactly cancel out and result in no effect. Meanwhile, for the heterogeneous synthetic data and the MovieLens dataset, we observe that the use of control variates improves the rate of decay of both \mathcal{E}_{t+1}^s and $\mathcal{E}_{t+1}^{p,s}$ allowing them to reach significantly lower values.

7 Federated Optimal Transport Map with an MM Approach

We now move beyond the scope of the framework characterized by MM1&2 to show how our methodology of aggregating in the surrogate space may be generalized beyond linearly parameterized surrogates. In particular, to illustrate how our methodology may be used to design novel federated algorithms, we consider a federated version of the problem of

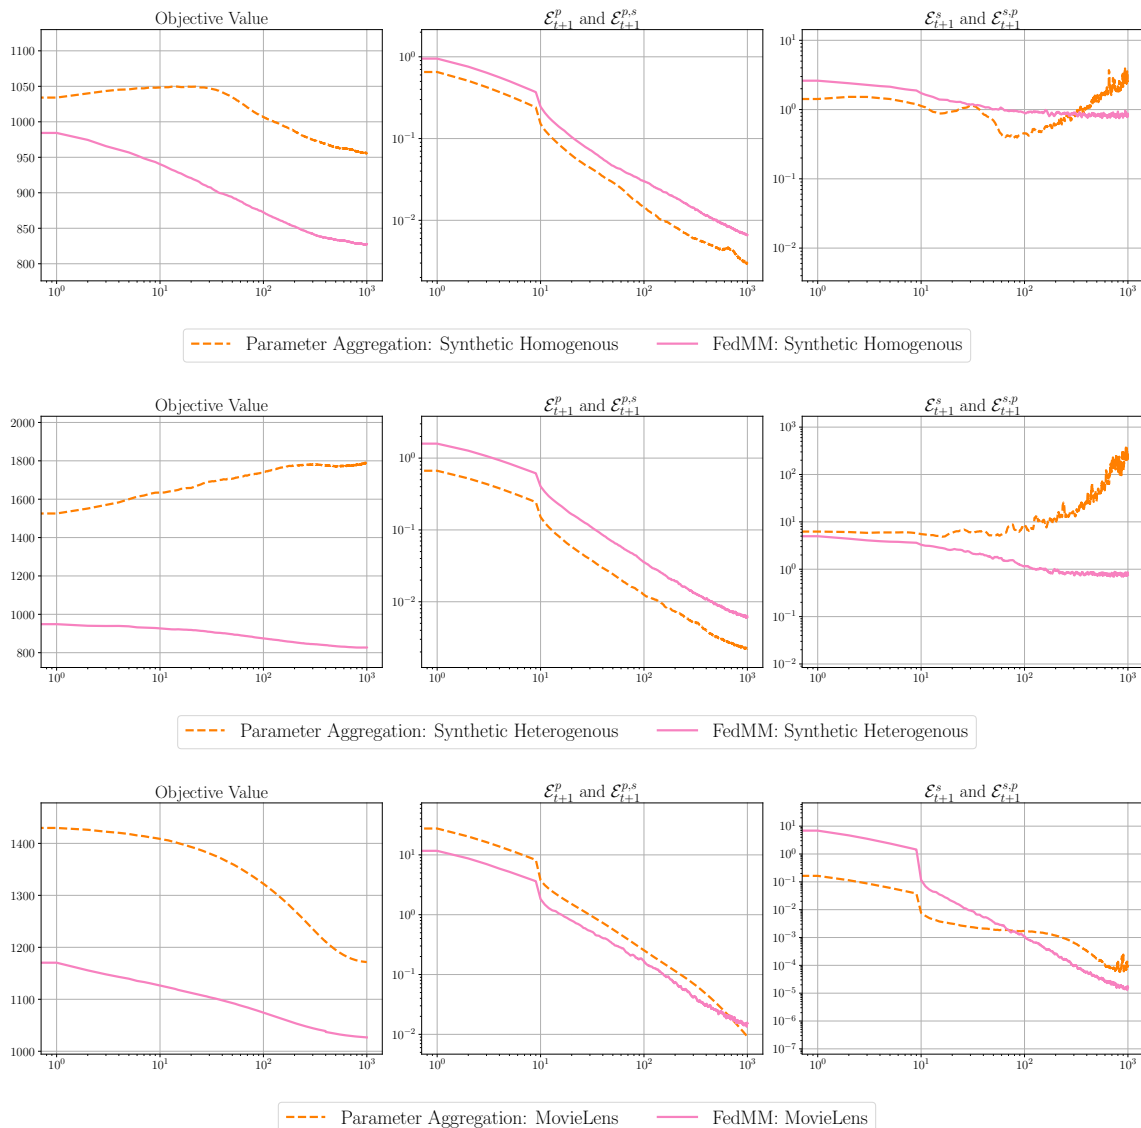


Figure 1: **Impact of Aggregation Space.** The columns (left to right) display the objective value, parameter update size (\mathcal{E}_t^p and $\mathcal{E}_t^{p,s}$), and surrogate update size (\mathcal{E}_t^s and $\mathcal{E}_t^{s,p}$) as a function of communication rounds, $t \in 1, \dots, T_{\max}$. The rows correspond to different data settings (from top to bottom): synthetic homogeneous, synthetic heterogeneous, and MovieLens 1M. Results are averaged over 10 runs, with the shaded regions indicating one standard deviation of the reported value across the 10 runs.

computing parametric optimal transport (OT) maps between two measures according to the framework of Korotin et al. (2021a). Before presenting the federated extension and our MM-based approach, we first describe the problem in the centralized setting. Consider two measures \mathcal{P} and \mathcal{Q} on $\mathcal{X} \subset \mathbb{R}^d$ and denote by $\mathbb{M}(\mathcal{X})$ the set of measurable functions on \mathcal{X}

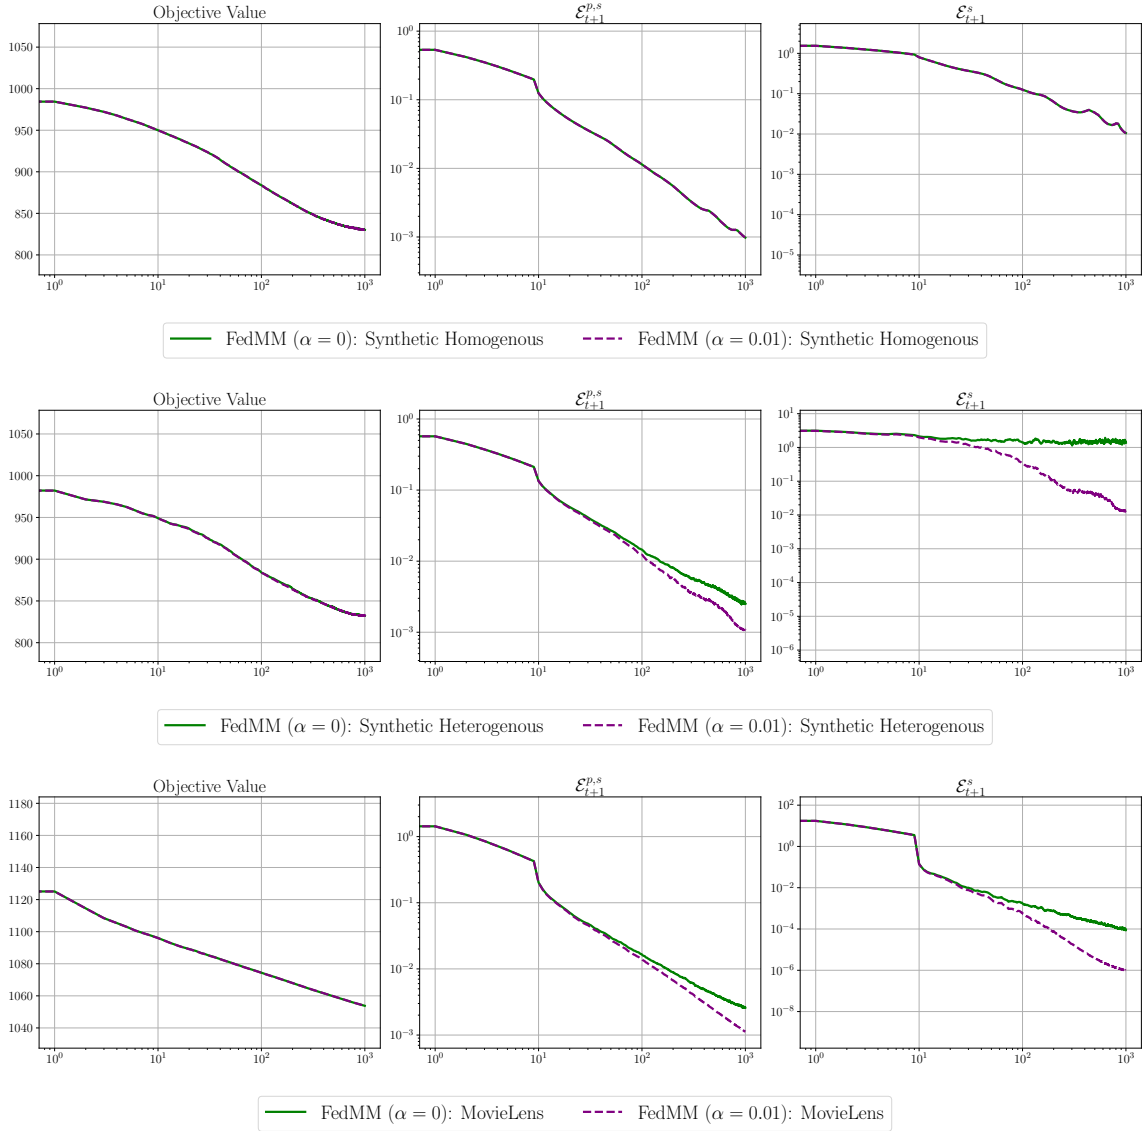


Figure 2: **Impact of Control Variates.** The columns (left to right) display the objective value, parameter update size ($\mathcal{E}_t^{p,s}$), and surrogate update size (\mathcal{E}_t^s) as a function of communication rounds, $t \in 1, \dots, T_{\max}$. The rows correspond to different data settings (from top to bottom): synthetic homogeneous, synthetic heterogeneous, and MovieLens 1M. Results are averaged over 10 runs, with the shaded regions indicating one standard deviation of the reported value across the 10 runs.

taking values in \mathcal{X} , the Wasserstein-2 transport map is defined as

$$m^* \in \underset{\substack{\{m \circ \mathcal{P} = \mathcal{Q}\} \\ m \in \mathbb{M}(\mathcal{X})}}{\operatorname{argmin}} \mathbb{E}_{\mathcal{P}} \left\{ \|X - m(X)\|^2 \right\}, \quad (29)$$

where the condition $\{m \circ \mathcal{P} = \mathcal{Q}\}$ indicates that the push-forward measure obtained by applying m to \mathcal{P} is equal to \mathcal{Q} , i.e. for any Borel set $B \subseteq \mathcal{X}$, it holds that $\mathcal{Q}(B) = \mathcal{P}(m^{-1}(B))$. In practice, m^* is computed by first solving the Kantorovich dual problem

$$f^* = \operatorname{argmax}_{f \in \text{CVX}(\mathcal{X})} -\mathbb{E}_{\mathcal{P}}[f(X)] - \mathbb{E}_{\mathcal{Q}}[f^c(Y)], \quad (30)$$

where $\text{CVX}(\mathcal{X})$ is the set of convex functions on \mathcal{X} and f^c is the convex conjugate of f . Under the condition that \mathcal{P} and \mathcal{Q} admit continuous probability density functions, it is known that the solution of the primal formulation may be constructed from the dual solution as $m^* = \nabla f^*$ (Villani, 2021). In addition, the inverse map $(m^*)^{-1}$ satisfies $(m^*)^{-1} = \nabla (f^*)^c$. Given the structure of the problem, Makkuva et al. (2020) aimed to find the optimal transport map as the gradient of an Input Convex Neural Network (ICNN) (Amos et al., 2017), leading to the problem

$$\omega^* \in \operatorname{argmin}_{\omega \in \Theta} \mathbb{E}_{\mathcal{P}, \mathcal{Q}} \left[f_{\omega}(X) + \max_{x \in \mathcal{X}} \{\langle x, Y \rangle - f_{\omega}(x)\} \right],$$

where $\mathbb{E}_{\mathcal{P}, \mathcal{Q}}$ indicates that $X \sim \mathcal{P}$ and $Y \sim \mathcal{Q}$. Here, Θ represents a parametric family of ICNN. Then, the transport map m^* can be approximated using the map $x \mapsto \nabla_x f_{\omega^*}(x)$. Korotin et al. (2021a) proposes to relax the min-max formulation by using a second ICNN to parameterize the convex conjugate leading to the minimization problem

$$\omega^*, \theta^* \in \operatorname{argmin}_{\omega, \theta \in \Theta^2} \{ \mathbb{E}_{\mathcal{P}}[f_{\omega}(X)] + \mathbb{E}_{\mathcal{Q}}[\langle \nabla_y f_{\theta}(Y), Y \rangle - f_{\omega}(\nabla_y f_{\theta}(Y))] \}, \quad (31)$$

where we used the fact that the convex conjugate f^c of a strictly convex function f satisfies

$$\nabla f^c(u) = \operatorname{argmax}_x (\langle u, x \rangle - f(x)).$$

However, a solution to (31) may not necessarily guarantee that f_{θ} is the convex conjugate of f_{ω} . Thus, Korotin et al. (2021a) further leverage the structure of the problem and add the regularizer

$$R_{\mathcal{Q}}(\omega, \theta) := \mathbb{E}_{\mathcal{Q}} \left[\|\nabla f_{\omega} \circ \nabla f_{\theta}(Y) - Y\|^2 \right],$$

and thus recovering the formulation

$$\omega^*, \theta^* \in \operatorname{argmin}_{\omega, \theta \in \Theta^2} \{ \mathbb{E}_{\mathcal{P}}[f_{\omega}(X)] + \mathbb{E}_{\mathcal{Q}}[\langle \nabla f_{\theta}(Y), Y \rangle - f_{\omega}(\nabla f_{\theta}(Y))] + R_{\mathcal{Q}}(\omega, \theta) \}. \quad (32)$$

7.1 Extension to the Federated Setting

We are interested in extending this problem to the federated setting. Consider a scenario where n clients, each possessing data from a distinct local distribution \mathcal{P}_i , aim to learn a single, shared optimal transport map to a common target distribution \mathcal{Q} . This formulation is highly relevant for federated domain adaptation. For instance, imagine a consortium of hospitals (clients) wishing to use a diagnostic AI model trained on a large, public dataset (with distribution \mathcal{Q}). Each hospital's local data has a local distribution \mathcal{P}_i , which may or may not be similar to other hospitals in the consortium. The goal is to leverage the consortium data to learn a transport map that transports the hospital's local distribution

onto \mathcal{Q} on which the model was trained, thus improving its performance without retraining it.

In this federated context, the objective is to find a single pair (ω^*, θ^*) that performs well on average across all clients:

$$\omega^*, \theta^* \in \underset{\omega, \theta \in \Theta^2}{\operatorname{argmin}} W(\omega, \theta) := \sum_{i \in [n]} \mu_i l_i(\omega, \theta) + \lambda \mathcal{R}_{\mathcal{Q}}(\omega, \theta), \quad (33)$$

where

$$l_i(\omega, \theta) := \{\mathbb{E}_{\mathcal{P}_i} [f_{\omega}(X)] + \mathbb{E}_{\mathcal{Q}} [\langle \nabla f_{\theta}(Y), Y \rangle - f_{\omega}(\nabla f_{\theta}(Y))]\}.$$

When ω and θ parameterize an ICNN, the objective in (33) is differentiable in (ω, θ) , making it directly tractable by available first-order federated algorithms (Reddi et al., 2021; Wang et al., 2021). Nonetheless, we will demonstrate a more effective alternative approach, which is based on our methodology of surrogate aggregation. In particular, the central thesis of our work is that for many optimization problems, the most effective federated approach is to have clients compute local surrogate functions, aggregate these surrogates on the server to form a global surrogate, and then perform the minimization step centrally. We now demonstrate how this design philosophy can be generalized to problems, like the OT formulation above, that do not strictly fit the linearly parameterized framework of MM1&2.

First, we note that the objective in (33) decouples the data dependencies between the two sets of parameters, θ and ω . In particular, once f_{ω} is fixed, the optimization problem for the conjugate network parameters θ , only depends on the shared public distribution \mathcal{Q} . This structural insight allows us to interpret the proposed alternating optimization scheme as a pseudo-MM algorithm that directly applies the principles developed earlier in this paper.

To formalize this, consider the objective function for a single client i , which we denote as $\min_{\omega, \theta} W_i(\omega, \theta)$, where $W_i(\omega, \theta) := l_i(\omega, \theta) + \lambda \mathcal{R}_{\mathcal{Q}}(\omega, \theta)$. This local objective can be alternatively viewed as a function of θ alone as it is equivalent to $\min_{\theta} \mathcal{W}_i(\theta)$, where $\mathcal{W}_i(\theta) := \min_{\omega} W_i(\omega, \theta)$. In particular, for an iterate θ_t , consider the client’s local best-response $\omega_i(\theta_t) := \operatorname{argmin}_{\omega} W_i(\omega, \theta_t)$. This step implicitly defines a majorizing function for the client’s true objective. Specifically, the function $U_{i,t}(\theta) \triangleq W_i(\omega_i(\theta_t), \theta)$ is a majorizing surrogate for $\mathcal{W}_i(\theta)$, as it satisfies the two key properties: (1) It is tangent at the point of construction: $U_{i,t}(\theta_t) = W_i(\omega_i(\theta_t), \theta_t) = \mathcal{W}_i(\theta_t)$; (2) It upper bounds \mathcal{W}_i as $U_{i,t}(\theta) = W_i(\omega_i(\theta_t), \theta) \geq \min_{\omega} W_i(\omega, \theta) = \mathcal{W}_i(\theta)$ for all θ .

Within our proposed framework, the parameter $\omega_i(\theta_t)$ plays a role directly analogous to the surrogate parameter $\mathbb{E}_{\pi_i} [\bar{S}(Z, \theta)]$ in the linearly parameterized setting in MM1&2: it is the finite-dimensional object that fully defines the client’s surrogate function $U_{i,t}(\cdot)$. Thus, analogously to FedMM, it is possible to develop an algorithm, which aims to aggregate $\omega_1(\theta_t), \dots, \omega_n(\theta_t)$ to construct a single surrogate at the server, which is then minimized at the server. We operationalize this approach in Algorithm 3, which exploits this idea to solve the federated objective (33).

Algorithm 3: Pseudo-MM for Federated OT Map (FedMM-OT)

Data: $T_{\max} \in \mathbb{N}_{>0}$; $\alpha > 0$; $\mathfrak{p} \in (0, 1]$; a $(0, 1]$ -valued sequence $\{\gamma_t, t \in [T_{\max}]\}$; a distribution \mathbf{P} on the subsets of $\{1, \dots, n\}$; initial values $\omega_0, \theta_0 \in \Theta$ and $V_{0,i} \in \Theta$ for all $i \in [n]$

Result: Sequence $\{(\omega_t, \theta_t), t \in [T_{\max}]\}$

- 1 Set $V_0 = \sum_{i=1}^n \mu_i V_{0,i}$
- 2 **for** $t = 0, \dots, T_{\max} - 1$ **do**
- 3 Sample $\mathcal{A}_{t+1} \sim \mathbf{P}$
- 4 Broadcast (ω_t, θ_t) to the agents $\#i, i \in \mathcal{A}_{t+1}$
- 5 **for** $i \in \mathcal{A}_{t+1}$ **do** *// on active worker $\#i$*
- 6 Compute $\omega_i(\theta_t) = \operatorname{argmin}_{\omega} W_i(\omega, \theta_t)$
- 7 Set $\Delta_{t+1,i} = \omega_i(\theta_t) - \omega_t - V_{t,i}$
- 8 Set $V_{t+1,i} = V_{t,i} + (\alpha/\mathfrak{p})\Delta_{t+1,i}$
- 9 Send $\Delta_{t+1,i}$ to the central server
- 10 **for** $i \notin \mathcal{A}_{t+1}$ **do** *// on inactive worker $\#i$*
- 11 $V_{t+1,i} = V_{t,i}$
- 12 *// on the central server*
- 13 Set $H_{t+1} = V_t + (1/\mathfrak{p}) \sum_{i \in \mathcal{A}_{t+1}} \mu_i \Delta_{t+1,i}$
- 14 Compute $\omega_{t+1} = \omega_t + \gamma_{t+1} H_{t+1}$
- 15 Project ω_{t+1} onto Θ
- 16 Update $\theta_{t+1} = \operatorname{argmin}_{\theta \in \Theta} W(\omega_{t+1}, \theta)$
- 17 Set $V_{t+1} = V_t + \alpha/\mathfrak{p} \sum_{i \in \mathcal{A}_{t+1}} \mu_i \Delta_{t+1,i}$

A single round of the algorithm is initiated by the server broadcasting the current global parameters, (ω_t, θ_t) . Each active client $i \in \mathcal{A}_{t+1}$ is tasked with constructing its local surrogate by computing the best-response parameter $\omega_i(\theta_t)$ for its local potential function. Then, clients transmit a control variate corrected version of their functional parameter update $\omega_i(\theta_t)$. The server, in turn, performs the aggregation and global minimization steps. It first aggregates the updates received from all active clients to form a new global potential parameter ω_{t+1} . Using this newly formed ω_{t+1} , the server then performs the global minimization for the conjugate parameter

$$\theta_{t+1} \in \operatorname{argmin}_{\theta \in \Theta} W(\omega_{t+1}, \theta).$$

In practice, we note that the exact minimization steps in both lines (6) and (16) can be relaxed and approximated by performing multiple gradient steps.

Crucially, the surrogate constructed at the server, parameterized by ω_{t+1} , is not guaranteed to be a majorizing function of the the global objective, and thus FedMM-OT does not strictly constitute an MM algorithm. As such, we term this approach as “pseudo-MM”. Nonetheless, it follows the identical principle advocated in this paper: local computations produce surrogate parameters, which are then aggregate centrally.

7.2 Evaluation of Optimal Transport Map

Evaluating OT map solutions is challenging due to the high sample complexity required for accurate Wasserstein distance estimation, specifically $O(\varepsilon^{-d})$ samples for $O(\varepsilon)$ accuracy

(Dudley, 1969), which becomes computationally prohibitive in high dimensions. To address this, we adopt the standardized benchmark methodology for evaluating continuous quadratic-cost OT maps of Korotin et al. (2021b).

This benchmark leverages ICNNs to construct pairs of continuous measures $(\mathcal{P}_d, \mathcal{Q}_d)$ where the ground truth OT map m_{true}^* is analytically known. In particular, for each dimension d , the benchmark starts by defining the following measures: \mathcal{P}_d is a Gaussian mixture with three components, while $\mathcal{Q}_d^1, \mathcal{Q}_d^2$, each a Gaussian mixture with ten components. Using two ICNNs, a push forward function $m_{\text{true}}^* : \mathcal{X} \rightarrow \mathcal{X}$ is learned to approximate $T^* \circ \mathcal{P} \approx 1/2 (\mathcal{Q}_d^1 + \mathcal{Q}_d^2)$. Finally, the benchmark sets $\mathcal{Q}_d = T^* \circ \mathcal{P}$ and, as such, the transport map between \mathcal{P}_d and \mathcal{Q}_d is explicitly known.

To quantify the performance of a fitted parametric transport map $m : \mathcal{X} \rightarrow \mathcal{X}$ against this ground truth, the benchmark employs the L_2 Unexplained Variance Percentage metric ($L_2^{\text{UVP}}(m)$), which is formally defined as

$$L_2^{\text{UVP}}(m) := 100 \cdot \frac{\|m - m_{\text{true}}^*\|_{L^2(\mathcal{P}_d)}^2}{\text{Var}(\mathcal{Q}_d)},$$

where $\text{Var}(\mathcal{Q}_d)$ is a scalar-valued notion of the spread of \mathcal{Q}_d . We use the original implementation² to have $\text{Var}(\mathcal{Q}_d)$ being the L_1 norm of the covariance matrix of \mathcal{Q}_d .

The L_2^{UVP} value serves as an indicator of how accurately the computed transport map m approximates the true optimal transport map m_{true}^* .

7.3 Experimental Results

To benchmark the effectiveness of Algorithm 3, we again compare against an algorithm that aims to solve the federated objective without aggregation of surrogate parameters. In particular, the objective in (33) is differentiable in (ω, θ) , making it amenable to federated first-order methods (McMahan et al., 2017b; Reddi et al., 2021). As such, we compare Algorithm 3 with FedAdam (Reddi et al., 2021).

We adapted the experimental setting and hyperparameters used by Korotin et al. (2021b, MMv2). In particular, we use a three-layer dense ICNN (Korotin et al., 2021a) for f_θ and f_ω . In addition, except for the learning rate, which we tune, we use the Adam (Kingma and Ba, 2017) hyperparameters used in the centralized setting in the code of Korotin et al. (2021b) for FedAdam.

For Algorithm 3, we relax the optimization procedure on the clients' side in line (6) to be a single gradient step. As such, both Algorithm 3 and FedAdam use similar computational resources on the client side. Finally, for the server side optimization process in line (16) to be ten gradient steps using the Adam optimizer.

We report our results on the benchmark from Korotin et al. (2021b) (described in Section 7.2), using the L_2^{UVP} metric. To adapt this benchmark to the federated setting, we generate \mathcal{P}_d as previously detailed. Then, for $n = 10$ clients, we create individual client distributions by splitting samples from \mathcal{P}_d using constrained k-means. Finally, in (33), we set $\mu_i = 1/n$. The results for different dimensions are presented in Figure 3. Across all dimensions, we observe that our approach significantly outperforms the direct application of FedAdam.

2. <https://github.com/iamalexkorotin/Wasserstein2Benchmark>

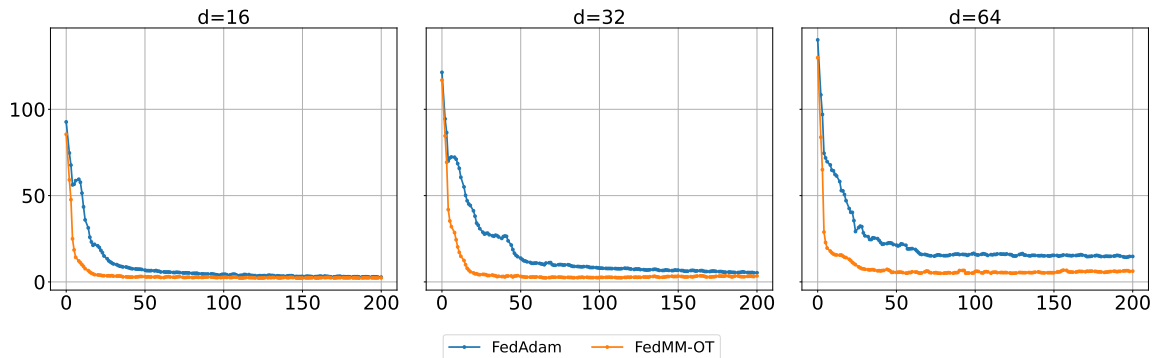


Figure 3: Performance is measured by (L_2^{UVP}) , where lower is better, versus communication rounds for dimensions $d = 16$, $d = 32$, and $d = 64$. In all settings, our proposed FedMM-OT algorithm converges faster than the FedAdam baseline.

8 Conclusion

This paper investigates Stochastic Surrogate MM methods that iteratively construct surrogate functions in a *parametric space*. This principle allows us to introduce FedMM, a unified framework for a novel class of FL algorithms designed to address the key bottlenecks of communication and heterogeneity. We believe that this simple yet elegant idea opens up many algorithmic directions, and simultaneously poses exciting theoretical difficulties.

Finally, a point of interest is the link with Mirror Descent (MD). Kunstner et al. (2021) showed that a (deterministic) MM - namely, the EM algorithm - can be seen as a MD algorithm. This remark strongly relies on a canonical parameterization $\phi(\theta) \equiv \theta$ of the surrogate functions, see MM-2 making the analysis simpler (eq. (3), Kunstner et al., 2021). Understanding the impact of the parameterization on the behavior of Federated MM is an interesting aspect left unexplored.

Acknowledgments

The work of Aymeric Dieuleveut and Mahmoud Hegazy is supported by French State aid managed by the French Agence Nationale de la Recherche (ANR) under France 2030 program with the reference ANR-23-PEIA-005 (REDEEM project), and by ANR-23-IACL-0005 (Hi! PARIS AI Cluster), in particular Hi! PARIS *FLAG* chair. The work of Gersende Fort is partly supported by ANR under the projects ANR-23-CE48-0009 and ANR-24-CE23-1529. The work of Hoi-To Wai is supported by project #MMT-p5-23 of the Shun Hing Institute of Advanced Engineering, The Chinese University of Hong Kong

Appendix A. The conditions MM-1 and MM-2

A.1 Example 1

Let $\rho \in (0, 1/L_f]$. For any $\tau \in \Theta$, it holds on Θ

$$f(\cdot) \leq f(\tau) + \langle \nabla f(\tau), \cdot - \tau \rangle + \frac{L_f}{2} \|\cdot - \tau\|^2 \leq f(\tau) + \langle \nabla f(\tau), \cdot - \tau \rangle + \frac{1}{2\rho} \|\cdot - \tau\|^2.$$

Since

$$\langle \nabla f(\tau), \theta - \tau \rangle + \frac{1}{2\rho} \|\theta - \tau\|^2 = \frac{1}{2\rho} (\|\theta\|^2 - \|\tau\|^2) - \frac{1}{\rho} \langle \tau - \rho \nabla f(\tau), \theta - \tau \rangle,$$

MM-1 is satisfied with

$$\psi(\theta) := \frac{1}{2\rho} \|\theta\|^2, \quad \phi(\theta) := \frac{1}{\rho} \theta, \quad \bar{S}(Z, \tau) := \tau - \rho G(Z, \tau).$$

Upon noting that $1/2\|\theta\|^2 - \langle s, \theta \rangle = 1/2\|\theta - s\|^2 - 1/2\|s\|^2$, it holds

$$\Upsilon(s) := \operatorname{argmin}_{\theta} \left(g(\theta) + \frac{1}{2\rho} \|\theta\|^2 - \frac{1}{\rho} \langle s, \theta \rangle \right) = \operatorname{argmin}_{\theta} \left(\rho g(\theta) + \frac{1}{2} \|\theta - s\|^2 \right).$$

A.2 Example 2

Let $\tau \in \Theta$. For any $\theta \in \Theta$, we write

$$\ell(Z, \theta) - \ell(Z, \tau) = -\log \int_{\mathcal{H}} \frac{p(Z, h, \theta)}{p(Z, h, \tau)} \mu(\mathrm{d}h|Z, \tau).$$

By the Jensen's inequality, we have for any $\theta \in \Theta$

$$\ell(Z, \theta) - \ell(Z, \tau) \leq -\int_{\mathcal{H}} \log \left(\frac{p(Z, h, \theta)}{p(Z, h, \tau)} \right) \mu(\mathrm{d}h|Z, \tau).$$

Since $-\log p(Z, h, \theta) = \psi(\theta) - \langle S(Z, h), \phi(\theta) \rangle$, this yields

$$\ell(Z, \theta) \leq \ell(Z, \tau) + \psi(\theta) - \psi(\tau) - \langle \bar{S}(Z, \tau), \phi(\theta) - \phi(\tau) \rangle \quad (34)$$

where $\bar{S}(Z, \tau)$ is given by (6). We now apply the expectation \mathbb{E}_{π} , and this concludes the proof of MM-1.

A.3 Example 3

Let $\tau \in \Theta$. First, observe that by definition of $\mathbf{M}(Z, \tau)$, it holds for any Z

$$\ell(Z, \tau) = \psi(\tau) - \langle S(Z, \mathbf{M}(Z, \tau)), \phi(\tau) \rangle + \xi(Z, \mathbf{M}(Z, \tau)). \quad (35)$$

We have $\ell(Z, \cdot) \leq \tilde{\ell}(Z, h, \cdot)$ for any $h \in \mathcal{H}$. Apply this inequality with $h \leftarrow \mathbf{M}(Z, \tau)$ and use (35): this yields

$$\ell(Z, \cdot) \leq \tilde{\ell}(Z, \mathbf{M}(Z, \tau), \cdot) = \psi(\cdot) - \langle S(Z, \mathbf{M}(Z, \tau)), \phi(\cdot) \rangle + \xi(Z, \mathbf{M}(Z, \tau)), \quad \forall Z \in \mathbb{R}^p, \quad (36)$$

$$= \ell(Z, \tau) + \psi(\cdot) - \psi(\tau) - \langle S(Z, \mathbf{M}(Z, \tau)), \phi(\cdot) - \phi(\tau) \rangle \quad (37)$$

From this inequality and (35), it follows under integrability conditions that

$$f(\cdot) \leq f(\tau) + \psi(\cdot) - \psi(\tau) - \langle \mathbb{E}_\pi [S(Z, \mathbf{M}(Z, \tau))], \phi(\cdot) - \phi(\tau) \rangle.$$

Therefore, MM-1 is satisfied with $\bar{S}(Z, \tau) := S(Z, \mathbf{M}(Z, \tau))$ and the functions ϕ, ψ defined by $\tilde{\ell}$.

Appendix B. Proof of Proposition 1

By definition of \mathbb{T} (see MM-2) and the Fermat's rule (see e.g. (Bauschke and Combettes, 2011, Theorem 16.3)), for any $s \in \mathcal{S}$, we have

$$0 \in \partial(g(\cdot) + \psi(\cdot) - \langle s, \phi(\cdot) \rangle)(\mathbb{T}(s)). \quad (38)$$

Let $s \in \mathcal{S}$. Under A1, both $g(\cdot)$ and $\psi(\cdot) - \langle s, \phi(\cdot) \rangle$ are proper lower semicontinuous convex functions, and $\Theta \cap \text{intdom}(\psi(\cdot) - \langle s, \phi(\cdot) \rangle) \supset \text{int}\Theta$ is non empty. Therefore, by (Bauschke and Combettes, 2011, Corollary 16.50), for any $\tau \in \Theta$

$$\partial(g(\cdot) + \psi(\cdot) - \langle s, \phi(\cdot) \rangle)(\tau) = \partial g(\tau) + \partial(\psi(\cdot) - \langle s, \phi(\cdot) \rangle)(\tau). \quad (39)$$

A1 and (Bauschke and Combettes, 2011, Proposition 17.31) imply that for any $\tau \in \Theta$,

$$\partial(\psi(\cdot) - \langle s, \phi(\cdot) \rangle)(\tau) = \{\nabla\psi(\tau) - J_\phi(\tau) s\}. \quad (40)$$

Therefore, from (38), (39) and (40), for any $s \in \mathcal{S}$, we have

$$J_\phi(\mathbb{T}(s)) s - \nabla\psi(\mathbb{T}(s)) \in \partial g(\mathbb{T}(s)). \quad (41)$$

On the other hand, MM-1 implies that for any $\tau \in \Theta$, $\theta \mapsto f(\theta) - f(\tau) - \psi(\theta) + \psi(\tau) + \langle \mathbb{E}_\pi [\bar{S}(Z, \tau)], \phi(\theta) - \phi(\tau) \rangle \leq 0$ and is equal to zero at $\theta = \tau$. Therefore, under A1 and A2, for any $\tau \in \text{int}\Theta$, it holds

$$\nabla f(\tau) - \nabla\psi(\tau) + J_\phi(\tau) \mathbb{E}_\pi [\bar{S}(Z, \tau)] = 0. \quad (42)$$

For any $\tau \in \text{int}\Theta$, we have

$$\begin{aligned} \nabla f(\tau) + \partial g(\tau) &\stackrel{(42)}{=} \nabla\psi(\tau) - J_\phi(\tau) \mathbb{E}_\pi [\bar{S}(Z, \tau)] + \partial g(\tau) \\ &\stackrel{(40)}{=} \partial(\psi(\cdot) - \langle s, \phi(\cdot) \rangle)(\tau) + \partial g(\tau) \\ &\stackrel{(39)}{=} \partial(g(\cdot) + \psi(\cdot) - \langle \mathbb{E}_\pi [\bar{S}(Z, \tau)], \phi(\cdot) \rangle)(\tau). \end{aligned}$$

Combined with the Fermat's rule, this implies that $0 \in \nabla f(\tau) + \partial g(\tau)$ iff τ is a minimizer of $g(\cdot) + \psi(\cdot) - \langle \mathbb{E}_\pi [\bar{S}(Z, \tau)], \phi(\cdot) \rangle$, i.e., under MM-2, $\tau = \mathbb{T}(\mathbb{E}_\pi [\bar{S}(Z, \tau)])$. This concludes the proof of (10).

On the other hand, the result (11) follows by observing that as $\mathbb{T}(s) \in \text{int}\Theta$,

$$\begin{aligned} -J_\phi(\mathbb{T}(s)) \mathbf{h}(s) &= J_\phi(\mathbb{T}(s)) (s - \mathbb{E}_\pi [\bar{S}(Z, \mathbb{T}(s))]) \\ &\stackrel{(42)}{=} \nabla f(\mathbb{T}(s)) - \nabla\psi(\mathbb{T}(s)) + J_\phi(\mathbb{T}(s)) s \\ &\stackrel{(40)}{\in} \nabla f(\mathbb{T}(s)) + \partial g(\mathbb{T}(s)). \end{aligned}$$

Let $s_\star \in \mathcal{S}$ be such that $\mathbf{h}(s_\star) = 0$ and $\mathbb{T}(s_\star) \in \text{int}\Theta$, then (10) and (11) imply that $\theta_\star = \mathbb{T}(s_\star)$. On the other hand, let $\theta_\star \in \text{int}\Theta$ be such that $\theta_\star = \mathbb{T}(\mathbb{E}_\pi [\bar{S}(Z, \theta_\star)])$, then (10) and (11) imply $\theta_\star = \mathbb{T}(s_\star)$, and furthermore $s_\star = \mathbb{E}_\pi [\bar{S}(Z, \mathbb{T}(s_\star))]$. This yields $\mathbf{h}(s_\star) = 0$.

Appendix C. Examples of Jensen’s surrogate: the Expectation Maximization algorithm

We derive the surrogate parameterization for two concrete EM examples in Section 2.3.

C.1 First Example: Poisson observations

Statistical model. Consider a Poisson random variable; conditionally to some hidden random variable h , the intensity of the Poisson variable is $\exp(\theta + h)$. We assume that the distribution of h is μ . Based on examples Z_1, \dots, Z_j, \dots , we want to estimate $\theta \in \mathbb{R}$ as the Maximum A Posteriori (MAP) when the prior is proportional to $\theta \mapsto \exp(-\lambda \exp(\theta))$ for some regularization parameter $\lambda > 0$. The complete likelihood i.e. the likelihood of the pair (Z, h) is proportional to $\theta \mapsto \exp(Z(\theta + h) - \exp(\theta + h))$. A MAP estimator is a minimizer of

$$\theta \mapsto \mathbb{E}_\pi \left[-\log \int_{\mathbb{R}} \exp(\theta Z - \exp(\theta) \exp(h)) \mu(dh) \right] + \lambda \exp(\theta); \quad (43)$$

it is of the form $\theta \mapsto f(\theta) + g(\theta)$ with $g(\theta) := \lambda \exp(\theta)$. In batch learning with i.i.d. observations Z_j , π is the empirical distribution of the N examples in the batch: $\pi(dz) := 1/N \sum_{j=1}^N \delta_{Z_j}(dz)$. In online learning with a stream of independent examples having the same distribution as Z , π denotes the distribution of Z .

Setting MM-1 and MM-2. From (43) and Example 2, MM-1 is satisfied with

$$\psi(\theta) := 0, \quad \phi(\theta) := \begin{bmatrix} \theta \\ \exp(\theta) \end{bmatrix}, \quad S(Z, h) := \begin{bmatrix} Z \\ -\exp(h) \end{bmatrix}, \quad \Theta := \mathbb{R}, \quad \mathcal{S} := \mathbb{R}_{>0} \times \mathbb{R}_{<0}.$$

MM-2 is satisfied: the optimization operator \mathbb{T} is defined on \mathcal{S} by

$$\mathbb{T}(s_1, s_2) := \operatorname{argmin}_{\theta \in \mathbb{R}} \{ \lambda \exp(\theta) - s_1 \theta - s_2 \exp(\theta) \} = \ln \left(\frac{s_1}{\lambda - s_2} \right).$$

Upon noting that the first component of $\mathbb{E}_\pi [\bar{S}(Z, h, \tau)]$ does not depend on τ and is equal to $\mathbb{E}_\pi[Z]$, let us propose a second strategy when this expectation is explicit (typically in batch learning) possibly at a large computational cost (when the size N of the batch is large). The function $\psi(\cdot)$ contains the terms depending upon θ which have a closed form expression.

Setting MM-1 and MM-2 - when $\mathbb{E}_\pi[Z]$ is explicit. MM-1 is satisfied with

$$\psi(\theta) := -\theta \mathbb{E}_\pi[Z], \quad \phi(\theta) := \exp(\theta), \quad S(Z, h) := -\exp(h), \quad \Theta := \mathbb{R}, \quad \mathcal{S} := \mathbb{R}_{<0}.$$

In MM-2, the optimization operator \mathbb{T} is defined on \mathcal{S} by

$$\mathbb{T}(s) := \operatorname{argmin}_{\theta \in \mathbb{R}} \{ \lambda \exp(\theta) - \mathbb{E}_\pi[Z] \theta - s \exp(\theta) \} = \ln \left(\frac{\mathbb{E}_\pi[Z]}{\lambda - s} \right).$$

C.2 Second example: mixture of Gaussian distributions

The statistical model. Consider a mixture model of L \mathbb{R}^p -valued Gaussian distributions. The means m_1, \dots, m_L are unknown; the weights ν_1, \dots, ν_L and the covariance matrices

$\Gamma_1, \dots, \Gamma_L$ are known. Based on observations Z_1, \dots, Z_j, \dots , we want to estimate $\theta := (m_1, \dots, m_L) \in \mathbb{R}^{p \times L}$ as a penalized Maximum Likelihood (ML) estimator that minimizes

$$\theta \mapsto \mathbb{E}_\pi \left[-\log \sum_{\ell=1}^L \frac{\nu_\ell}{\sqrt{\text{Det}(\Gamma_\ell)}} \exp \left(-\frac{1}{2} (Z - m_\ell)^\top \Gamma_\ell^{-1} (Z - m_\ell) \right) \right] + \frac{\lambda}{2} \sum_{\ell=1}^L m_\ell^\top m_\ell,$$

where $\lambda > 0$ is a regularization parameter. π is either $\pi(\mathbf{d}z) := N^{-1} \sum_{j=1}^N \delta_{Z_j}(\mathbf{d}z)$, the empirical distribution on the N i.i.d. examples Z_j in batch learning; or π denotes the distribution of Z in online learning with a stream of i.i.d. samples.

The criterion is of the form $\theta \mapsto f(\theta) + g(\theta)$ with $g(\theta) := \lambda/2 \sum_{\ell=1}^L m_\ell^\top m_\ell$ and

$$\begin{aligned} f(\theta) &:= \mathbb{E}_\pi \left[-\log \sum_{\ell=1}^L \frac{\nu_\ell}{\sqrt{\text{Det}(\Gamma_\ell)}} \exp \left(-\frac{1}{2} (Z - m_\ell)^\top \Gamma_\ell^{-1} (Z - m_\ell) \right) \right] \\ &= \mathbb{E}_\pi \left[-\log \int \frac{1}{\sqrt{\text{Det}(\Gamma_h)}} \exp \left(-\frac{1}{2} (Z - m_h)^\top \Gamma_h^{-1} (Z - m_h) \right) \mu(\mathbf{d}h) \right] \end{aligned}$$

where μ is the distribution on $\{1, \dots, L\}$ defined by $\mu(\{h\}) := \nu_h$.

The setting MM-1-MM-2. It holds

$$(Z - m_h)^\top \Gamma_h^{-1} (Z - m_h) = Z^\top \Gamma_h^{-1} Z + m_h^\top \Gamma_h^{-1} m_h - 2 \langle Z, \Gamma_h^{-1} m_h \rangle.$$

Since $1 = \mathbf{1}_{\{1\}}(h) + \dots + \mathbf{1}_{\{L\}}(h)$, we have

$$\langle \xi, \chi(h) \rangle = \left\langle \begin{bmatrix} \xi \mathbf{1}_{\{1\}}(h) \\ \dots \\ \xi \mathbf{1}_{\{L-1\}}(h) \end{bmatrix}, \begin{bmatrix} \chi(1) - \chi(L) \\ \dots \\ \chi(L-1) - \chi(L) \end{bmatrix} \right\rangle + \langle \xi, \chi(L) \rangle$$

from which we deduce that

$$\begin{aligned} \langle Z, \Gamma_h^{-1} m_h \rangle &= \left\langle \begin{bmatrix} Z \mathbf{1}_{\{1\}}(h) \\ \dots \\ Z \mathbf{1}_{\{L-1\}}(h) \end{bmatrix}, \begin{bmatrix} \Gamma_1^{-1} m_1 - \Gamma_L^{-1} m_L \\ \dots \\ \Gamma_{L-1}^{-1} m_{L-1} - \Gamma_L^{-1} m_L \end{bmatrix} \right\rangle + \langle Z, \Gamma_L^{-1} m_L \rangle. \\ m_h^\top \Gamma_h^{-1} m_h &= \left\langle \begin{bmatrix} \mathbf{1}_{\{1\}}(h) \\ \dots \\ \mathbf{1}_{\{L-1\}}(h) \end{bmatrix}, \begin{bmatrix} m_1^\top \Gamma_1^{-1} m_1 - m_L^\top \Gamma_L^{-1} m_L \\ \dots \\ m_{L-1}^\top \Gamma_{L-1}^{-1} m_{L-1} - m_L^\top \Gamma_L^{-1} m_L \end{bmatrix} \right\rangle + m_L^\top \Gamma_L^{-1} m_L. \end{aligned}$$

Therefore, from Example 2, **MM-1** is satisfied with $\psi(\theta) := -m_L^\top \Gamma_L^{-1} \mathbb{E}_\pi [Z] + 1/2 m_L^\top \Gamma_L^{-1} m_L$,

$$\phi(\theta) := \begin{bmatrix} \Gamma_1^{-1} m_1 - \Gamma_L^{-1} m_L \\ \dots \\ \Gamma_L^{-1} m_L - \Gamma_L^{-1} m_L \\ -1/2 (m_1^\top \Gamma_1^{-1} m_1 - m_L^\top \Gamma_L^{-1} m_L) \\ \dots \\ -1/2 (m_{L-1}^\top \Gamma_{L-1}^{-1} m_{L-1} - m_L^\top \Gamma_L^{-1} m_L) \end{bmatrix}, \quad S(Z, h) := \begin{bmatrix} Z \mathbf{1}_{\{1\}}(h) \\ \dots \\ Z \mathbf{1}_{\{L-1\}}(h) \\ \mathbf{1}_{\{1\}}(h) \\ \dots \\ \mathbf{1}_{\{L-1\}}(h) \end{bmatrix}.$$

The set \mathcal{S} collects the points of the form $s = (s_{1,1}, \dots, s_{1,L-1}, s_{2,1}, \dots, s_{2,L-1})$ where $s_{1,\ell} \in \mathbb{R}^p$, $s_{2,\ell} \in [0, 1]$ and $\sum_{\ell=1}^{L-1} s_{2,\ell} \leq 1$. In MM-2, for any $s \in \mathcal{S}$, $\mathbb{T}(s)$ is the minimizer of $g(\theta) + \psi(\theta) - \langle s, \phi(\theta) \rangle$. $\mathbb{T}(s) \in \mathbb{R}^{p \times L}$ and the column $\#\ell$, for $\ell = 1, \dots, L-1$, is given by

$$[\mathbb{T}(s)]_{:, \ell} := (s_{2,\ell} \mathbf{I}_p + \lambda \Gamma_\ell)^{-1} s_{1,\ell};$$

the column $\#L$ is given by

$$[\mathbb{T}(s)]_{:, L} := \left(\left(1 - \sum_{\ell=1}^{L-1} s_{2,\ell} \right) \mathbf{I}_p + \lambda \Gamma_L \right)^{-1} \left(\mathbb{E}_\pi[Z] - \sum_{\ell=1}^{L-1} s_{1,\ell} \right).$$

Appendix D. Proof of Theorem 1

D.1 Outline of the section

The proof of Theorem 1 can be carried for the *full participation regime* up to adjusting the properties of the compression operator i.e. when A5(1) holds, with an adjusted compression operator satisfying A4(ω_p) where $\omega_p := \omega + (1-p)/p(\omega + 1)$. Such a property is addressed in Appendix D.2 together with the statement of Algorithm 4, equivalent under A5 to Algorithm 2. This line of analysis has been frequently used in the literature, for example in Koloskova et al. (2019).

Stochastic quantities are successively introduced in the learning process: when selecting active workers; when producing oracles $S_{t+1,i}$ and possibly in the compression step. Finally, the local randomnesses are aggregated by the central server when updating the system \hat{S}_{t+1} . In Appendix D.3, a filtration is introduced in echo to these steps.

Finally, the proof of Theorem 1 is in Appendix D.5; technical lemmas conclude this appendix (see Appendix D.6).

D.2 From Partial Participation to Compression

Let

$$\widetilde{\text{Quant}}(s) := \frac{U}{p} \text{Quant}(s), \quad s \in \mathbb{R}^q,$$

where U is a Bernoulli random variable with probability of success p and independent of $\text{Quant}(s)$. Lemma 1 shows that if Quant satisfies A4(ω), then $\widetilde{\text{Quant}}$ satisfies A4(ω_p).

Lemma 1 *Let Quant be a compression operator satisfying (25) with the constant ω . Then $\widetilde{\text{Quant}}$ satisfies (25), with a constant ω_p given by*

$$\omega_p := \omega + \frac{1-p}{p}(\omega + 1). \quad (44)$$

Proof Let $s \in \mathbb{R}^q$. It holds:

$$\mathbb{E} \left[\widetilde{\text{Quant}}(s) \right] = \frac{1}{p} \mathbb{E} [U] \mathbb{E} [\text{Quant}(s)] = \mathbb{E} [\text{Quant}(s)]$$

where we used that U and Quant are independent, and the expectation of U is p . Since U and $\text{Quant}(s)$ are independent, it holds

$$\mathbb{E} \left[\left(\frac{U}{p} - 1 \right) \langle \text{Quant}(s), \text{Quant}(s) - s \rangle \right] = 0.$$

Therefore

$$\mathbb{E} \left[\|\widetilde{\text{Quant}}(s) - s\|^2 \right] = \mathbb{E} \left[\left(\frac{U}{p} - 1 \right)^2 \|\text{Quant}(s)\|^2 \right] + \mathbb{E} [\|\text{Quant}(s) - s\|^2].$$

Using again that U and $\text{Quant}(s)$ are independent and $\mathbb{E} [\text{Quant}(s)] = s$, the first term is equal to

$$\mathbb{E} \left[\left(\frac{U}{p} - 1 \right)^2 \right] \mathbb{E} [\|\text{Quant}(s)\|^2] = \frac{1-p}{p} \left(\mathbb{E} [\|\text{Quant}(s) - s\|^2] + \|s\|^2 \right).$$

By using (25), we have

$$\mathbb{E} \left[\|\widetilde{\text{Quant}}(s) - s\|^2 \right] \leq \frac{1-p}{p}(\omega+1)\|s\|^2 + \omega\|s\|^2.$$

Therefore, $\widetilde{\text{Quant}}$ satisfies (25) with the constant ω_p . ■

As a corollary, when clients are selected independently at random with probability p (see A5(p)), Algorithm 2 is equivalent to the algorithm where all the workers are active at all steps and the compression operator Quant is swapped with $\widetilde{\text{Quant}}$; in words, $\widetilde{\text{Quant}}$ encodes both the compression step and the partial participation step. Such an equivalence follows directly upon noting that (i) identifying a set of active agents \mathcal{A}_{t+1} whose size is np in expectation, is realized by sampling n independent Bernoulli random variables $(U_{t+1,1}, \dots, U_{t+1,n})$, with success probability p , (ii) it holds

$$\frac{1}{p} \sum_{i \in \mathcal{A}_{t+1}} a_i = \sum_{i=1}^n \frac{U_{t+1,i}}{p} a_i,$$

We set $\widetilde{\text{Quant}}_{t+1,i}(s) := p^{-1}U_{t+1,i} \text{Quant}_{t+1,i}(s)$ and move from Algorithm 2 to Algorithm 4. From Lemma 1, it is easily seen that the r.v. $\{\widetilde{\text{Quant}}_{t+1,i}(s), t \geq 0, i \in [n]\}$ satisfy A4(ω_p).

D.3 Filtrations

For a random variable U , we denote by $\sigma(U)$ the sigma-field generated by U i.e. the smallest sigma-field that makes U measurable. Set

$$\mathcal{F}_0 := \sigma \left(\hat{S}_0; V_{0,i} \text{ for } i \in [n] \right).$$

We define by induction the following filtration, which takes into account the randomness introduced by the algorithm at different stages and possibly at each local worker $\#i$:

$$\begin{aligned} \mathcal{F}_{t+1/3,i} &:= \mathcal{F}_t \vee \sigma(S_{t+1,i}) \\ \mathcal{F}_{t+2/3,i} &:= \mathcal{F}_{t+1/3,i} \vee \sigma \left(\widetilde{\text{Quant}}(\Delta_{t+1,i}) \right), \\ \mathcal{F}_{t+1} &:= \bigvee_{i=1}^n \mathcal{F}_{t+2/3,i}. \end{aligned}$$

$\mathcal{F}_{t+1/3,i}$ contains the past of the algorithm up to iteration $\#t$ included, and the randomness introduced by the oracles $S_{t+1,i}$. $\mathcal{F}_{t+2/3,i}$ adds to $\mathcal{F}_{t+1/3,i}$ the randomness introduced by the compression mechanism. \mathcal{F}_{t+1} contains all the randomness introduced by the n local workers until their iteration $\#(t+1)$ included. We have for all $i \in [n]$ and $t \geq 0$, $\mathcal{F}_t \subseteq \mathcal{F}_{t+1/3,i} \subseteq \mathcal{F}_{t+2/3,i} \subseteq \mathcal{F}_{t+1}$ and

$$\begin{aligned} S_{t+1,i} \in \mathcal{F}_{t+1/3,i}, \quad \Delta_{t+1,i} \in \mathcal{F}_{t+1/3,i}, \quad \widetilde{\text{Quant}}(\Delta_{t+1,i}) \in \mathcal{F}_{t+2/3,i}, \quad V_{t+1,i} \in \mathcal{F}_{t+2/3,i} \\ H_{t+1} \in \mathcal{F}_{t+1}, \quad \hat{S}_{t+1} \in \mathcal{F}_{t+1}, \quad V_{t+1} \in \mathcal{F}_{t+1}. \end{aligned}$$

Algorithm 4: FedMM (when Partial Participation is part of Quant)

Data: $T_{\max} \in \mathbb{N}_{>0}$; $\alpha > 0$; a $(0, 1]$ -valued sequence $\{\gamma_t, t \in [T_{\max}]\}$; initial values $\hat{S}_0 \in \mathcal{S}$ and $V_{0,i} \in \mathbb{R}^q$ for all $i \in [n]$; a probability distribution $\{\mu_i, i \in [n]\}$ on $[n]$.

Result: Sequence: $\{\hat{S}_t, t \in [T_{\max}]\}$

- 1 Set $V_0 = \sum_{i=1}^n \mu_i V_{0,i}$
- 2 **for** $t = 0, \dots, T_{\max} - 1$ **do** // for all iterations
- 3 Send \hat{S}_t and $\mathbb{T}(\hat{S}_t)$ to the workers
- 4 **for** $i = 1, \dots, n$ **do** // for all workers
- 5 Sample $S_{t+1,i}$, oracle for $\mathbb{E}_{\pi_i}[\bar{S}(Z, \mathbb{T}(\hat{S}_t))]$
- 6 Set $\Delta_{t+1,i} = S_{t+1,i} - \hat{S}_t - V_{t,i}$
- 7 Set $V_{t+1,i} = V_{t,i} + \alpha \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i})$
- 8 Send $\widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i})$ to the central server
- 9 // on the central server
- 10 Set $H_{t+1} = V_t + \sum_{i=1}^n \mu_i \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i})$
- 11 Compute $\hat{S}_{t+1/2} = \hat{S}_t + \gamma_{t+1} H_{t+1}$
- 12 Project it on \mathcal{S} , given $\mathbf{B}_t \succ 0$: $\hat{S}_{t+1} = \operatorname{argmin}_{s \in \mathcal{S}} (s - \hat{S}_{t+1/2})^\top \mathbf{B}_t (s - \hat{S}_{t+1/2})$
- 13 Set $V_{t+1} = V_t + \alpha \sum_{i=1}^n \mu_i \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i})$

Furthermore, note that (see Lemma 1)

$$\mathbb{E} \left[\widetilde{\text{Quant}}(\Delta_{t+1,i}) \mid \mathcal{F}_{t+1/3,i} \right] = \Delta_{t+1,i}, \quad (45)$$

and under A3,

$$\mathbb{E} [S_{t+1,i} \mid \mathcal{F}_t] = \mathbb{E}_{\pi_i} \left[\bar{S}(Z, \mathbb{T}(\hat{S}_t)) \right].$$

D.4 Notations

Define the functions

$$\mathbf{h}(s) := \mathbb{E}_{\pi} [\bar{S}(Z, \mathbb{T}(s))] - s = \sum_{i=1}^n \mu_i \mathbb{E}_{\pi_i} [\bar{S}(Z, \mathbb{T}(s))] - s, \quad \mathbf{h}_i(s) := \mathbb{E}_{\pi_i} [\bar{S}(Z, \mathbb{T}(s))] - s.$$

Set

$$\bar{C} := C_{\star\star} + \sup_S \|\mathbf{h}(\cdot)\| C_{\star}, \quad \omega_p := \omega + \frac{1-p}{p}(\omega+1), \quad \sigma^2 := n \sum_{i=1}^n \mu_i^2 \sigma_i^2, \quad L^2 := n \sum_{i=1}^n \mu_i^2 L_i^2.$$

For any $t \geq 1$, define

$$\bar{\mathcal{E}}_{t+1} := \frac{\|\Pi(\hat{S}_t + \gamma_{t+1} \mathbf{h}(\hat{S}_t), \mathbf{B}_t) - \hat{S}_t\|_{\mathbf{B}_t}^2}{\gamma_{t+1}^2};$$

note that when $\Pi(\hat{S}_t + \gamma_{t+1} \mathbf{h}(\hat{S}_t), \mathbf{B}_t) = \hat{S}_t + \gamma_{t+1} \mathbf{h}(\hat{S}_t)$, then $\bar{\mathcal{E}}_{t+1} = \|\mathbf{h}(\hat{S}_t)\|_{\mathbf{B}_t}^2$.

D.5 Proof of Theorem 1

We start by establishing Proposition 2, which provides a Lyapunov inequality for the function $(f + g) \circ \mathbb{T}$, for any timestep $t \geq 0$.

Proposition 2 *Assume MM-1, MM-2 and A7. For any $t \geq 0$, choose $\gamma_{t+1} \in (0, v_{\min}/(2\bar{C}))$. With probability 1, it holds for any $t \geq 0$*

$$\begin{aligned} \mathbb{E} \left[f(\mathbb{T}(\hat{S}_{t+1})) + g(\mathbb{T}(\hat{S}_{t+1})) \mid \mathcal{F}_t \right] &\leq f(\mathbb{T}(\hat{S}_t)) + g(\mathbb{T}(\hat{S}_t)) \\ &\quad - \frac{\gamma_{t+1}}{4} (1 - 2\bar{C}v_{\min}^{-1}\gamma_{t+1}) \mathbb{E} [\bar{\mathcal{E}}_{t+1} \mid \mathcal{F}_t] \\ &\quad + \gamma_{t+1} (1 - \bar{C}v_{\min}^{-1}\gamma_{t+1}) \mathbb{E} \left[\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbb{B}_t}^2 \mid \mathcal{F}_t \right]. \end{aligned}$$

\bar{C} , $\bar{\mathcal{E}}_{t+1}$ and \mathbf{h} are defined in Appendix D.4.

Proof We start with proving that for any $s, s' \in \mathcal{S}$, it holds

$$(f+g) \circ \mathbb{T}(s') \leq (f+g) \circ \mathbb{T}(s) - \langle \mathbf{h}(s), \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle + \langle s' - s, \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle. \quad (46)$$

Let $s, s' \in \mathcal{S}$. From MM-1, we have

$$f(\mathbb{T}(s')) \leq f(\mathbb{T}(s)) + \psi(\mathbb{T}(s')) - \psi(\mathbb{T}(s)) - \langle \mathbb{E}_\pi [\bar{S}(Z, \mathbb{T}(s))], \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle.$$

This yields

$$\begin{aligned} f(\mathbb{T}(s')) &\leq f(\mathbb{T}(s)) - \langle \mathbf{h}(s), \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle \\ &\quad + \psi(\mathbb{T}(s')) - \psi(\mathbb{T}(s)) - \langle s, \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle. \end{aligned}$$

Therefore

$$(f + g) \circ \mathbb{T}(s') \leq (f + g) \circ \mathbb{T}(s) - \langle \mathbf{h}(s), \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle + \Xi(s, s') \quad (47)$$

where $\Xi(s, s') := g(\mathbb{T}(s')) - g(\mathbb{T}(s)) + \psi(\mathbb{T}(s')) - \psi(\mathbb{T}(s)) - \langle s, \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle$. Let us prove that $\Xi(s, s') \leq \langle s' - s, \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle$. We write

$$\Xi(s, s') = \mathcal{L}(s', s') + \langle s' - s, \phi(\mathbb{T}(s')) \rangle - \mathcal{L}(s, s), \quad \mathcal{L}(s, s') := g(\mathbb{T}(s)) + \psi(\mathbb{T}(s)) - \langle s', \phi(\mathbb{T}(s)) \rangle.$$

The upper bound on $\Xi(s, s')$ follows from Lemma 2. This concludes the proof of (46); applied with $s' = \hat{S}_{t+1}$ and $s = \hat{S}_t$, this yields

$$\begin{aligned} f(\mathbb{T}(\hat{S}_{t+1})) + g(\mathbb{T}(\hat{S}_{t+1})) &\leq f(\mathbb{T}(\hat{S}_t)) + g(\mathbb{T}(\hat{S}_t)) - \langle \mathbf{h}(\hat{S}_t), \phi(\mathbb{T}(\hat{S}_{t+1})) - \phi(\mathbb{T}(\hat{S}_t)) \rangle \\ &\quad + \langle \hat{S}_{t+1} - \hat{S}_t, \phi(\mathbb{T}(\hat{S}_{t+1})) - \phi(\mathbb{T}(\hat{S}_t)) \rangle. \end{aligned}$$

By A7, we have

$$\langle \hat{S}_{t+1} - \hat{S}_t, \phi(\mathbb{T}(\hat{S}_{t+1})) - \phi(\mathbb{T}(\hat{S}_t)) \rangle \leq C_{**} \|\hat{S}_{t+1} - \hat{S}_t\|^2 \leq v_{\min}^{-1} C_{**} \gamma_{t+1}^2 \mathcal{E}_{t+1},$$

where we set $\mathcal{E}_{t+1} := \gamma_{t+1}^{-2} \|\hat{S}_{t+1} - \hat{S}_t\|_{\mathbb{B}_t}^2$. We write

$$-\left\langle \mathbf{h}(\hat{S}_t), \phi(\mathbb{T}(\hat{S}_{t+1})) - \phi(\mathbb{T}(\hat{S}_t)) \right\rangle \leq -\left\langle H_{t+1}, \hat{S}_{t+1} - \hat{S}_t \right\rangle_{\mathbb{B}_t} \quad (48)$$

$$+ \left| \left\langle H_{t+1} - \mathbf{h}(\hat{S}_t), \hat{S}_{t+1} - \hat{S}_t \right\rangle_{\mathbb{B}_t} \right| \quad (49)$$

$$+ \left| \left\langle \mathbf{h}(\hat{S}_t), \phi(\mathbb{T}(\hat{S}_{t+1})) - \phi(\mathbb{T}(\hat{S}_t)) - \mathbb{B}_t(\hat{S}_{t+1} - \hat{S}_t) \right\rangle \right|. \quad (50)$$

By Corollary 2, an upper bound for (48) is $-\gamma_{t+1}\mathcal{E}_{t+1}$. By using $|\langle a, b \rangle_{\mathbb{B}_t}| \leq \gamma_{t+1}\|a\|_{\mathbb{B}_t}^2/2 + \|b\|_{\mathbb{B}_t}^2/(2\gamma_{t+1})$ and $\|\hat{S}_{t+1} - \hat{S}_t\|_{\mathbb{B}_t} = \gamma_{t+1}^2\mathcal{E}_{t+1}$, an upper bound for (49) is

$$\frac{\gamma_{t+1}}{2}\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbb{B}_t}^2 + \frac{\gamma_{t+1}}{2}\mathcal{E}_{t+1}.$$

Under A7, an upper bound for (50) is

$$C_\star \sup_{\mathcal{S}} \|\mathbf{h}(\cdot)\| \|\hat{S}_{t+1} - \hat{S}_t\|^2 = v_{\min}^{-1} C_\star \sup_{\mathcal{S}} \|\mathbf{h}(\cdot)\| \gamma_{t+1}^2 \mathcal{E}_{t+1}$$

Apply Corollary 2 again to obtain a bound of $-\mathcal{E}_{t+1}$; then the proof is concluded by application of the conditional expectation. \blacksquare

The field H_{t+1} is a random approximation of $\mathbf{h}(\hat{S}_t)$. Proposition 3 shows that the oracle H_{t+1} is unbiased, conditionally to the past \mathcal{F}_t . The proof of Proposition 3 establishes that the error $H_{t+1} - \mathbf{h}(\hat{S}_t)$ is the sum of two conditionally centered and uncorrelated terms: the first one is the mean value of the local oracles, and the second one is the mean value of the local compression operators. Therefore, two terms contribute to the conditional variance of the oracle H_{t+1} : the first one is the weighted average of the conditional variance of the local oracles, and the second one is the variance of the weighted compression errors. The control variate $V_{t,i}$ plays a role in this second term.

Proposition 3 *Assume A3, A4(ω) and A5(\mathbf{p}). With probability one, it holds for any $t \geq 0$*

$$\mathbb{E}[H_{t+1} | \mathcal{F}_t] = \mathbf{h}(\hat{S}_t), \quad (51)$$

$$\mathbb{E}\left[\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbb{B}_t}^2 \mid \mathcal{F}_t\right] \leq (1 + \omega_{\mathbf{p}})v_{\max} \frac{\sigma^2}{n} + \omega_{\mathbf{p}}v_{\max} \sum_{i=1}^n \mu_i^2 \|V_{t,i} - \mathbf{h}_i(\hat{S}_t)\|^2. \quad (52)$$

\mathbf{h} , \mathbf{h}_i , $\omega_{\mathbf{p}}$ and σ^2 are defined in Appendix D.4.

Proof Let $t \geq 0$. We have $H_{t+1} = V_t + \sum_{i=1}^n \mu_i \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i})$, where, by definition of the filtration, $V_{t,i} \in \mathcal{F}_{t-1+2/3,i} \subset \mathcal{F}_t$, $\hat{S}_t \in \mathcal{F}_t$, $\Delta_{t+1,i} \in \mathcal{F}_{t+1/3,i}$, $\widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i}) \in \mathcal{F}_{t+2/3,i}$ and $V_t \in \mathcal{F}_t$. Therefore,

$$\begin{aligned} \mathbb{E}[H_{t+1} | \mathcal{F}_t] &= V_t + \sum_{i=1}^n \mu_i \mathbb{E}\left[\mathbb{E}\left[\widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i}) \mid \mathcal{F}_{t+1/3,i}\right] \mid \mathcal{F}_t\right] = V_t + \sum_{i=1}^n \mu_i \mathbb{E}[\Delta_{t+1,i} | \mathcal{F}_t] \\ &= V_t + \sum_{i=1}^n \mu_i \left(\mathbb{E}[S_{t+1,i} | \mathcal{F}_t] - \hat{S}_t - V_{t,i}\right) = \sum_{i=1}^n \mu_i \mathbb{E}_{\pi_i} \left[\bar{S}(Z, \mathbb{T}(\hat{S}_t))\right] - \hat{S}_t = \mathbf{h}(\hat{S}_t), \end{aligned}$$

where we used A3, A4(ω), Lemma 1 and Proposition 5. This concludes the proof of (51). We write, by using Proposition 5,

$$H_{t+1} - \mathbf{h}(\hat{S}_t) = V_t + \sum_{i=1}^n \mu_i \{ \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i}) - \mathbf{h}_i(\hat{S}_t) \} = \sum_{i=1}^n \mu_i \{ \Xi_{t+1,i}^{(1)} + \Xi_{t+1,i}^{(2)} \},$$

with

$$\Xi_{t+1,i}^{(1)} := \mathbf{S}_{t+1,i} - \mathbb{E}_{\pi_i} \left[\bar{S}(Z, \mathbb{T}(\hat{S}_t)) \right], \quad \Xi_{t+1,i}^{(2)} := \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i}) - \Delta_{t+1,i};$$

By A3, A4(ω) and Lemma 1,

$$\Xi_{t+1,i}^{(1)} \in \mathcal{F}_{t+1/3,i}, \quad \mathbb{E} \left[\Xi_{t+1,i}^{(1)} \mid \mathcal{F}_t \right] = 0, \quad \mathbb{E} \left[\Xi_{t+1,i}^{(2)} \mid \mathcal{F}_{t+1/3,i} \right] = 0. \quad (53)$$

Since the workers are independent conditionally to \mathcal{F}_t (see A5(p)) and $\mathbf{B}_t \in \mathcal{F}_t$, we have

$$\mathbb{E} \left[\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbf{B}_t}^2 \mid \mathcal{F}_t \right] = \sum_{i=1}^n \mu_i^2 \mathbb{E} \left[\|\Xi_{t+1,i}^{(1)} + \Xi_{t+1,i}^{(2)}\|_{\mathbf{B}_t}^2 \mid \mathcal{F}_t \right] = \sum_{i=1}^n \mu_i^2 \sum_{\ell=1}^2 \mathbb{E} \left[\|\Xi_{t+1,i}^{(\ell)}\|_{\mathbf{B}_t}^2 \mid \mathcal{F}_t \right]$$

where we used that for all $i \in [n]$, $\mathbb{E} \left[\left\langle \Xi_{t+1,i}^{(1)}, \Xi_{t+1,i}^{(2)} \right\rangle_{\mathbf{B}_t} \mid \mathcal{F}_t \right] = 0$ by (53). We have by A3

$$\mathbb{E} \left[\|\Xi_{t+1,i}^{(1)}\|_{\mathbf{B}_t}^2 \mid \mathcal{F}_t \right] = \mathbb{E} \left[\|\mathbf{S}_{t+1,i} - \mathbb{E}_{\pi_i} \left[\bar{S}(Z, \mathbb{T}(\hat{S}_t)) \right]\|_{\mathbf{B}_t}^2 \mid \mathcal{F}_t \right] \leq v_{\max} \sigma_i^2;$$

and by Lemma 1,

$$\mathbb{E} \left[\|\Xi_{t+1,i}^{(2)}\|_{\mathbf{B}_t}^2 \mid \mathcal{F}_t \right] \leq v_{\max} \mathbb{E} \left[\|\Xi_{t+1,i}^{(2)}\|^2 \mid \mathcal{F}_t \right] \leq \omega_{\mathbf{p}} v_{\max} \mathbb{E} \left[\|\Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right].$$

As a conclusion, we proved

$$\mathbb{E} \left[\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbf{B}_t}^2 \mid \mathcal{F}_t \right] \leq v_{\max} \sum_{i=1}^n \mu_i^2 \sigma_i^2 + v_{\max} \omega_{\mathbf{p}} \sum_{i=1}^n \mu_i^2 \mathbb{E} \left[\|\Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right]. \quad (54)$$

Let $i \in [n]$. By A3, $\mathbb{E}[\mathbf{S}_{t+1,i} \mid \mathcal{F}_t] = \mathbf{h}_i(\hat{S}_t) + \hat{S}_t$. In addition, $\hat{S}_t \in \mathcal{F}_t$ and $V_{t,i} \in \mathcal{F}_t$, thus implying that

$$\mathbb{E} \left[\left\langle \mathbf{h}_i(\hat{S}_t) - V_{t,i}, \mathbf{S}_{t+1,i} - \hat{S}_t - \mathbf{h}_i(\hat{S}_t) \right\rangle \mid \mathcal{F}_t \right] = 0.$$

Hence, we get by A3 again,

$$\begin{aligned} \mathbb{E} \left[\|\Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right] &= \mathbb{E} \left[\|\mathbf{S}_{t+1,i} - \hat{S}_t - V_{t,i}\|^2 \mid \mathcal{F}_t \right] \\ &= \|\mathbf{h}_i(\hat{S}_t) - V_{t,i}\|^2 + \mathbb{E} \left[\|\mathbf{S}_{t+1,i} - \hat{S}_t - \mathbf{h}_i(\hat{S}_t)\|^2 \mid \mathcal{F}_t \right] \\ &= \|\mathbf{h}_i(\hat{S}_t) - V_{t,i}\|^2 + \mathbb{E} \left[\|\mathbf{S}_{t+1,i} - \mathbb{E}_{\pi_i} \left[\bar{S}(Z, \mathbb{T}(\hat{S}_t)) \right]\|^2 \mid \mathcal{F}_t \right] \\ &\leq \|\mathbf{h}_i(\hat{S}_t) - V_{t,i}\|^2 + \sigma_i^2. \end{aligned}$$

This concludes the proof of (52). ■

Proposition 4 shows that the sequence of random variables $\{\sum_{i=1}^n \mu_i^2 \|V_{t,i} - \mathbf{h}_i(\hat{S}_t)\|^2, t \geq 0\}$ satisfies a Lyapunov inequality. The important property is that for small enough stepsizes and for convenient values of $\alpha \in (0, 1]$ depending on the coefficient ω_p , it is of the form

$$\mathbb{E} \left[\sum_{i=1}^n \mu_i^2 \|V_{t+1,i} - \mathbf{h}_i(\hat{S}_{t+1})\|^2 \middle| \mathcal{F}_t \right] \leq \lambda \sum_{i=1}^n \mu_i^2 \|V_{t,i} - \mathbf{h}_i(\hat{S}_t)\|^2 + \Upsilon_t$$

where $\lambda \in (0, 1)$ and $\Upsilon_t \geq 0$. Such an inequality, combined with Proposition 2 will be crucial for the proof of Theorem 1.

Proposition 4 *Assume A3, A4(ω), A6 and A7. Define for all $t \geq 0$,*

$$G_t := \sum_{i=1}^n \mu_i^2 \|V_{t,i} - \mathbf{h}_i(\hat{S}_t)\|^2.$$

Choose $\alpha \in [0, 1/(1 + \omega_p)]$ and $\beta > 1$. With probability one, for any $t \geq 0$ it holds

$$\begin{aligned} \mathbb{E}[G_{t+1} | \mathcal{F}_t] \leq & \beta \left(1 - \alpha + \frac{2\omega_p}{\beta - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) G_t + \frac{2\beta}{\beta - 1} v_{\min}^{-1} \frac{L^2}{n} \gamma_{t+1}^2 \mathbb{E}[\bar{\mathcal{E}}_{t+1} | \mathcal{F}_t] \\ & + \beta \left(\alpha + \frac{2(1 + \omega_p)}{\beta - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) \frac{\sigma^2}{n}. \end{aligned} \quad (55)$$

ω_p , \mathbf{h}_i , L^2 , σ^2 and $\bar{\mathcal{E}}_{t+1}$ are defined in Appendix D.4.

Proof Let $t \geq 0$ and $i \in [n]$. For any $\beta > 1$, using that $\|a + b\|^2 \leq \beta/(\beta-1)\|a\|^2 + \beta\|b\|^2$, we have

$$\begin{aligned} \|V_{t+1,i} - \mathbf{h}_i(\hat{S}_{t+1})\|^2 & \leq \beta \|V_{t+1,i} - \mathbf{h}_i(\hat{S}_t)\|^2 + \beta/(\beta-1) \|\mathbf{h}_i(\hat{S}_{t+1}) - \mathbf{h}_i(\hat{S}_t)\|^2 \\ & \leq \beta \|V_{t+1,i} - \mathbf{h}_i(\hat{S}_t)\|^2 + \beta/(\beta-1) v_{\min}^{-1} L_i^2 \gamma_{t+1}^2 \mathcal{E}_{t+1}, \end{aligned}$$

where $\mathcal{E}_{t+1} := \|\hat{S}_{t+1} - \hat{S}_t\|_{\mathbb{B}_t}^2 / \gamma_{t+1}^2$. We used A6 in the last inequality. By Lemma 4 and the condition $\alpha(1 + \omega_p) - 1 \leq 0$, we have

$$\begin{aligned} \mathbb{E} \left[\|V_{t+1,i} - \mathbf{h}_i(\hat{S}_{t+1})\|^2 \middle| \mathcal{F}_t \right] & \leq \beta(1 - \alpha) \|V_{t,i} - \mathbf{h}_i(\hat{S}_t)\|^2 \\ & \quad + \beta\alpha\sigma_i^2 + \frac{\beta}{(\beta - 1)} v_{\min}^{-1} L_i^2 \gamma_{t+1}^2 \mathbb{E}[\mathcal{E}_{t+1} | \mathcal{F}_t]. \end{aligned}$$

We now multiply by μ_i^2 and sum in i ; this yields

$$\mathbb{E}[G_{t+1} | \mathcal{F}_t] \leq \beta(1 - \alpha) G_t + \beta\alpha \frac{\sigma^2}{n} + \frac{\beta}{\beta - 1} v_{\min}^{-1} \frac{L^2}{n} \gamma_{t+1}^2 \mathbb{E}[\mathcal{E}_{t+1} | \mathcal{F}_t].$$

By using Corollary 2

$$\begin{aligned} \mathbb{E}[G_{t+1} | \mathcal{F}_t] & \leq \beta(1 - \alpha) G_t + \beta\alpha \frac{\sigma^2}{n} + \frac{2\beta}{\beta - 1} v_{\min}^{-1} \frac{L^2}{n} \gamma_{t+1}^2 \mathbb{E}[\bar{\mathcal{E}}_{t+1} | \mathcal{F}_t] \\ & \quad + \frac{2\beta}{\beta - 1} v_{\min}^{-1} \frac{L^2}{n} \gamma_{t+1}^2 \mathbb{E} \left[\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbb{B}_t}^2 \middle| \mathcal{F}_t \right]. \end{aligned}$$

The proof of (55) is concluded by Proposition 3. ■

Conclusion: Proof of Theorem 1 in the case $\omega_p = 0$. Let $\alpha \in [0, 1]$. From Proposition 2 and Proposition 3, it holds

$$\begin{aligned} \mathbb{E} \left[f(\mathbb{T}(\hat{S}_{t+1})) + g(\mathbb{T}(\hat{S}_{t+1})) \mid \mathcal{F}_t \right] &\leq f(\mathbb{T}(\hat{S}_t)) + g(\mathbb{T}(\hat{S}_t)) \\ &\quad - \frac{\gamma_{t+1}}{4} (1 - 2\bar{C}v_{\min}^{-1}\gamma_{t+1}) \mathbb{E} [\bar{\mathcal{E}}_{t+1} \mid \mathcal{F}_t] + \gamma_{t+1} (1 - \bar{C}v_{\min}^{-1}\gamma_{t+1}) v_{\max} \frac{\sigma^2}{n}. \end{aligned}$$

By summing from $t = 0$ to $t = T - 1$ and computing the expectation, we obtain the result.

Proof of Theorem 1 in the case $\omega_p > 0$. From Proposition 2, Proposition 3 and Proposition 4, we have for any positive sequence $\{\rho_{t+1}, t \geq 0\}$

$$\begin{aligned} &\mathbb{E} \left[f(\mathbb{T}(\hat{S}_{t+1})) + g(\mathbb{T}(\hat{S}_{t+1})) \mid \mathcal{F}_t \right] + \rho_{t+1} \mathbb{E} [G_{t+1} \mid \mathcal{F}_t] \leq f(\mathbb{T}(\hat{S}_t)) + g(\mathbb{T}(\hat{S}_t)) + \rho_t G_t \\ &+ \gamma_{t+1} \left(\frac{\rho_{t+1}}{\gamma_{t+1}} \beta \left(1 - \alpha + \frac{2\omega_p}{\beta - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) + (1 - \bar{C}v_{\min}^{-1}\gamma_{t+1}) \omega_p v_{\max} - \frac{\rho_t}{\gamma_{t+1}} \right) G_t \\ &- \frac{\gamma_{t+1}}{4} \left(1 - 2\bar{C}v_{\min}^{-1}\gamma_{t+1} - \frac{8\beta}{\beta - 1} v_{\min}^{-1} \frac{L^2}{n} \rho_{t+1} \gamma_{t+1} \right) \mathbb{E} [\bar{\mathcal{E}}_{t+1} \mid \mathcal{F}_t] \\ &+ \gamma_{t+1} \left(\frac{\rho_{t+1}}{\gamma_{t+1}} \beta \left(\alpha + \frac{2(1 + \omega_p)}{\beta - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) + (1 - \bar{C}v_{\min}^{-1}\gamma_{t+1}) (1 + \omega_p) v_{\max} \right) \frac{\sigma^2}{n} \end{aligned}$$

Let $\alpha \in (0, 1/(1 + \omega_p)]$. Choose $\beta > 1$ and $\rho_{t+1} = \rho\gamma_{t+1}$. Then,

$$\begin{aligned} &\mathbb{E} \left[f(\mathbb{T}(\hat{S}_{t+1})) + g(\mathbb{T}(\hat{S}_{t+1})) \mid \mathcal{F}_t \right] + \rho\gamma_{t+1} \mathbb{E} [G_{t+1} \mid \mathcal{F}_t] \leq f(\mathbb{T}(\hat{S}_t)) + g(\mathbb{T}(\hat{S}_t)) + \rho\gamma_t G_t \\ &+ \gamma_{t+1} \left(\rho\beta \left(1 - \alpha + \frac{2\omega_p}{\beta - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) + (1 - \bar{C}v_{\min}^{-1}\gamma_{t+1}) \omega_p v_{\max} - \rho \frac{\gamma_t}{\gamma_{t+1}} \right) G_t \\ &- \frac{\gamma_{t+1}}{4} \left(1 - 2\bar{C}v_{\min}^{-1}\gamma_{t+1} - \frac{8\rho\beta}{\beta - 1} v_{\min}^{-1} \frac{L^2}{n} \gamma_{t+1}^2 \right) \mathbb{E} [\bar{\mathcal{E}}_{t+1} \mid \mathcal{F}_t] \\ &+ \gamma_{t+1} \left(\rho\beta \left(\alpha + \frac{2(1 + \omega_p)}{\beta - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) + (1 - \bar{C}v_{\min}^{-1}\gamma_{t+1}) (1 + \omega_p) v_{\max} \right) \frac{\sigma^2}{n} \end{aligned}$$

Now, choose β and ρ such that

$$1 < \beta \leq 2, \quad \beta(1 - \alpha) \leq 1 - \alpha/2, \quad \beta/(\beta - 1) \leq 2/\alpha, \quad \rho := 4\omega_p v_{\max}/\alpha;$$

for example, $\beta = 2$ if $\alpha \geq 2/3$ and $\beta(1 - \alpha) = 1 - \alpha/2$ otherwise. This yields, by using $\gamma_{t+1} \leq \gamma_t \leq \gamma_1$,

$$\begin{aligned} \rho\beta \left(1 - \alpha + \frac{2\omega_p}{\beta - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) &\leq \rho \left(1 - \frac{\alpha}{2} \right) + \rho^4 \frac{\omega_p}{\alpha} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_1^2 \leq \rho - \rho \frac{\alpha}{4}; \\ (1 - \bar{C}v_{\min}^{-1}\gamma_{t+1}) \omega_p v_{\max} - \rho \frac{\gamma_t}{\gamma_{t+1}} &\leq \omega_p v_{\max} - \rho = \rho \frac{\alpha}{4} - \rho; \\ 8\rho \frac{\beta}{\beta - 1} &\leq \frac{64}{\alpha^2} \omega_p v_{\max}; \\ \rho\beta \left(\alpha + \frac{2(1 + \omega_p)}{\beta - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) &\leq 4\omega_p v_{\max} \left(2 \wedge \frac{1 - \alpha/2}{1 - \alpha} \right) + (1 + \omega_p) v_{\max}. \end{aligned}$$

Hence,

$$\begin{aligned}
 & \mathbb{E} \left[f(\mathbb{T}(\hat{S}_{t+1})) + g(\mathbb{T}(\hat{S}_{t+1})) \mid \mathcal{F}_t \right] + \frac{4\omega_p v_{\max}}{\alpha} \gamma_{t+1} \mathbb{E} [G_{t+1} \mid \mathcal{F}_t] \\
 & \leq f(\mathbb{T}(\hat{S}_t)) + g(\mathbb{T}(\hat{S}_t)) + \frac{4\omega_p v_{\max}}{\alpha} \gamma_t G_t \\
 & \quad - \frac{\gamma_{t+1}}{4} \left(1 - 2\bar{C} v_{\min}^{-1} \gamma_{t+1} - 64 \frac{\omega_p}{\alpha^2} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) \mathbb{E} [\bar{\mathcal{E}}_{t+1} \mid \mathcal{F}_t] \\
 & \quad + 2\gamma_{t+1} v_{\max} \left(2\omega_p \left(2 \wedge \frac{1 - \alpha/2}{1 - \alpha} \right) + (1 + \omega_p) \right) \frac{\sigma^2}{n}.
 \end{aligned}$$

Observe that

$$2\gamma_{t+1} v_{\max} \left(2\omega_p \left(2 \wedge \frac{1 - \alpha/2}{1 - \alpha} \right) + (1 + \omega_p) \right) \frac{\sigma^2}{n} \leq 2\gamma_{t+1} v_{\max} (1 + 5\omega_p) \frac{\sigma^2}{n}.$$

By summing from $t = 0$ to $t = T - 1$ (with the convention that $\gamma_0 = \gamma_1$) and computing the expectations, we obtain the result.

D.6 Technical lemmas

Lemma 2 *Assume MM-1 and MM-2. For $s, s' \in \mathcal{S}$, set $\mathcal{L}(s, s') := g(\mathbb{T}(s)) + \psi(\mathbb{T}(s)) - \langle s', \phi(\mathbb{T}(s)) \rangle$. Then, for any $s, s' \in \mathcal{S}$, it holds*

$$\mathcal{L}(s', s') \leq \mathcal{L}(s, s) - \langle s' - s, \phi(\mathbb{T}(s)) \rangle, \quad (56)$$

$$\langle s' - s, \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) \rangle \geq 0. \quad (57)$$

Proof Let $s, s' \in \mathcal{S}$. By definition of $\mathbb{T}(s')$ (see MM-2), we have for any $\theta \in \Theta$

$$\begin{aligned}
 \mathcal{L}(s', s') &= g(\mathbb{T}(s')) + \psi(\mathbb{T}(s')) - \langle s', \phi(\mathbb{T}(s')) \rangle \leq g(\theta) + \psi(\theta) - \langle s', \phi(\theta) \rangle \\
 &= g(\theta) + \psi(\theta) - \langle s, \phi(\theta) \rangle - \langle s' - s, \phi(\theta) \rangle.
 \end{aligned}$$

Applied with $\theta = \mathbb{T}(s) \in \Theta$, this yields (56). By symmetry, we have $\mathcal{L}(s, s) \leq \mathcal{L}(s', s') + \langle s' - s, \phi(\mathbb{T}(s')) \rangle$ which, added to (56), yields (57). \blacksquare

Lemma 3 generalizes classical results about projection on a closed convex set (see e.g. (Polyak, 2021, Lemma 2.11)) to the case \mathbf{B} is not the identity matrix.

Lemma 3 *Let \mathcal{B} be a closed convex subset of \mathbb{R}^q and \mathbf{B} be a $q \times q$ positive definite matrix. Set $\langle a, b \rangle_{\mathbf{B}} := \langle \mathbf{B}a, b \rangle$, $\|a\|_{\mathbf{B}} := (a^\top \mathbf{B}a)^{1/2}$ and*

$$\Pi(\hat{s}, \mathbf{B}) := \operatorname{argmin}_{s \in \mathcal{B}} (s - \hat{s})^\top \mathbf{B} (s - \hat{s}).$$

Then for any $\hat{s}, s' \in \mathcal{B}$: $\langle \Pi(\hat{s}, \mathbf{B}) - \hat{s}, s' - \Pi(\hat{s}, \mathbf{B}) \rangle_{\mathbf{B}} \geq 0$. In addition, for any $\hat{s}_1, \hat{s}_2 \in \mathcal{B}$:

$$\|\hat{s}_1 - \hat{s}_2\|_{\mathbf{B}} \geq \|\Pi(\hat{s}_1, \mathbf{B}) - \Pi(\hat{s}_2, \mathbf{B})\|_{\mathbf{B}}.$$

Proof Let $\hat{s} \in \mathcal{B}$. A descent direction at $\Pi(\hat{s}, \mathbf{B})$ of the function $s' \mapsto (s' - \hat{s})^\top \mathbf{B}(s' - \hat{s})$ does not exist so that

$$\forall s' \in \mathcal{B}, \quad \langle \Pi(\hat{s}, \mathbf{B}) - \hat{s}, s' - \Pi(\hat{s}, \mathbf{B}) \rangle_{\mathbf{B}} \geq 0. \quad (58)$$

This concludes the proof of the first claim. This property implies

$$\langle \Pi(\hat{s}_2, \mathbf{B}) - \hat{s}_2, \Pi(\hat{s}_1, \mathbf{B}) - \Pi(\hat{s}_2, \mathbf{B}) \rangle_{\mathbf{B}} \geq 0, \quad \langle \Pi(\hat{s}_1, \mathbf{B}) - \hat{s}_1, \Pi(\hat{s}_2, \mathbf{B}) - \Pi(\hat{s}_1, \mathbf{B}) \rangle_{\mathbf{B}} \geq 0,$$

from which we deduce

$$\langle \hat{s}_1 - \hat{s}_2, \Pi(\hat{s}_1, \mathbf{B}) - \Pi(\hat{s}_2, \mathbf{B}) \rangle_{\mathbf{B}} \geq \|\Pi(\hat{s}_1, \mathbf{B}) - \Pi(\hat{s}_2, \mathbf{B})\|_{\mathbf{B}}^2.$$

The Cauchy-Schwarz inequality concludes the proof. \blacksquare

Corollary 2 (of Lemma 3) *For any $t \geq 0$, define*

$$\mathcal{E}_{t+1} := \frac{\|\hat{S}_{t+1} - \hat{S}_t\|_{\mathbf{B}_t}^2}{\gamma_{t+1}^2} = \frac{\|\Pi(\hat{S}_{t+1/2}, \mathbf{B}_t) - \hat{S}_t\|_{\mathbf{B}_t}^2}{\gamma_{t+1}^2}.$$

Then, for any $t \geq 0$, $\gamma_{t+1}\mathcal{E}_{t+1} \leq \langle \hat{S}_{t+1} - \hat{S}_t, H_{t+1} \rangle_{\mathbf{B}_t}$ and

$$\bar{\mathcal{E}}_{t+1} \leq 2\mathcal{E}_{t+1} + 2\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbf{B}_t}, \quad \mathcal{E}_{t+1} \leq 2\bar{\mathcal{E}}_{t+1} + 2\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbf{B}_t}.$$

Proof Let $t \geq 0$. We write

$$\begin{aligned} \langle \Pi(\hat{S}_{t+1/2}, \mathbf{B}_t) - \hat{S}_t, \hat{S}_{t+1/2} - \hat{S}_t \rangle_{\mathbf{B}_t} &= \langle \Pi(\hat{S}_{t+1/2}, \mathbf{B}_t) - \hat{S}_t, \hat{S}_{t+1/2} - \Pi(\hat{S}_{t+1/2}, \mathbf{B}_t) \rangle_{\mathbf{B}_t} \\ &\quad + \|\Pi(\hat{S}_{t+1/2}, \mathbf{B}_t) - \hat{S}_t\|_{\mathbf{B}_t}^2 \end{aligned}$$

and by Lemma 2, it holds

$$\langle \Pi(\hat{S}_{t+1/2}, \mathbf{B}_t) - \hat{S}_t, \hat{S}_{t+1/2} - \hat{S}_t \rangle_{\mathbf{B}_t} \geq \|\Pi(\hat{S}_{t+1/2}, \mathbf{B}_t) - \hat{S}_t\|_{\mathbf{B}_t}^2 = \gamma_{t+1}^2 \mathcal{E}_{t+1},$$

from which we deduce the first inequality upon noting that $\Pi(\hat{S}_{t+1/2}, \mathbf{B}_t) = \hat{S}_{t+1}$ and $\hat{S}_{t+1/2} - \hat{S}_t = \gamma_{t+1}H_{t+1}$. Lemma 2 applied with $\hat{s}_1 := \hat{S}_t + \gamma_{t+1}H_{t+1}$ and $\hat{s}_2 := \hat{S}_t + \gamma_{t+1}\mathbf{h}(\hat{S}_t)$ yields

$$\|\hat{S}_{t+1} - \Pi(\hat{S}_t + \gamma_{t+1}\mathbf{h}(\hat{S}_t), \mathbf{B}_t)\|_{\mathbf{B}_t} \leq \gamma_{t+1}\|H_{t+1} - \mathbf{h}(\hat{S}_t)\|_{\mathbf{B}_t}.$$

The second and third inequalities follow from $\|a + b\|_{\mathbf{B}}^2 \leq 2\|a\|_{\mathbf{B}}^2 + 2\|b\|_{\mathbf{B}}^2$. \blacksquare

Proposition 5 proves that the control variate V_t computed by the central server is equal to the mean value of the control variates computed locally by the workers.

Proposition 5 *With probability one, for any $t \geq 0$, we have $V_t = \sum_{i=1}^n \mu_i V_{t,i}$.*

Proof By definition of V_0 , the property holds true when $t = 0$. Assume this holds true for some $t \geq 0$. By definition of V_{t+1} at the central server in Algorithm 4, and by the induction assumption we write

$$V_{t+1} = \sum_{i=1}^n \mu_i V_{t,i} + \alpha \sum_{i=1}^n \mu_i \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i}).$$

The definition of $V_{t+1,i}$ at each local worker yields

$$V_{t+1} = \sum_{i=1}^n \mu_i V_{t,i} + \sum_{i=1}^n \mu_i (V_{t+1,i} - V_{t,i}) = \sum_{i=1}^n \mu_i V_{t+1,i}.$$

This concludes the induction. \blacksquare

Lemma 4 *Assume A3 and A4(ω). Choose $\alpha \in [0, 1]$. With probability one, for any $t \geq 0$ and $i \in [n]$, it holds*

$$\begin{aligned} \mathbb{E} \left[\|V_{t+1,i} - \mathbf{h}_i(\hat{S}_t)\|^2 \mid \mathcal{F}_t \right] &\leq (1 - \alpha) \|V_{t,i} - \mathbf{h}_i(\hat{S}_t)\|^2 + \alpha \sigma_i^2 \\ &\quad + \alpha (\alpha(1 + \omega_p) - 1) \mathbb{E} \left[\|\Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right]. \end{aligned}$$

ω_p and \mathbf{h}_i are defined in Appendix D.4.

Proof Let $t \geq 0$ and $i \in [n]$. We write,

$$V_{t+1,i} - \mathbf{h}_i(\hat{S}_t) = V_{t,i} - \mathbf{h}_i(\hat{S}_t) + \alpha \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i})$$

which yields

$$\mathbb{E} \left[\|V_{t+1,i} - \mathbf{h}_i(\hat{S}_t)\|^2 \mid \mathcal{F}_t \right] = \mathbb{E} \left[\|V_{t,i} - \mathbf{h}_i(\hat{S}_t) + \alpha \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i})\|^2 \mid \mathcal{F}_t \right]. \quad (59)$$

Upon noting that by Lemma 1, $\mathbb{E} \left[\widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i}) - \Delta_{t+1,i} \mid \mathcal{F}_{t+1/3,i} \right] = 0$, and $V_{t,i} - \mathbf{h}_i(\hat{S}_t) + \alpha \Delta_{t+1,i} \in \mathcal{F}_{t+1/3,i}$, we first write

$$\begin{aligned} &\mathbb{E} \left[\|V_{t,i} - \mathbf{h}_i(\hat{S}_t) + \alpha \widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i})\|^2 \mid \mathcal{F}_t \right] \\ &= \mathbb{E} \left[\|V_{t,i} - \mathbf{h}_i(\hat{S}_t) + \alpha \Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right] + \alpha^2 \mathbb{E} \left[\|\widetilde{\text{Quant}}_{t+1,i}(\Delta_{t+1,i}) - \Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right] \\ &\leq \mathbb{E} \left[\|V_{t,i} - \mathbf{h}_i(\hat{S}_t) + \alpha \Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right] + \alpha^2 \omega_p \mathbb{E} \left[\|\Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right] \end{aligned} \quad (60)$$

by using A4(ω) and Lemma 1. Second, we use the equality: for $x, y \in \mathbb{R}^q$ and $\alpha \in \mathbb{R}$, $\|(1 - \alpha)x + \alpha y\|^2 = (1 - \alpha)\|x\|^2 + \alpha\|y\|^2 - \alpha(1 - \alpha)\|x - y\|^2$; this yields by using A3

$$\begin{aligned} &\mathbb{E} \left[\|V_{t,i} - \mathbf{h}_i(\hat{S}_t) + \alpha \Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right] \\ &= \mathbb{E} \left[\|(1 - \alpha)(V_{t,i} - \mathbf{h}_i(\hat{S}_t)) + \alpha(S_{t+1,i} - \hat{S}_t - \mathbf{h}_i(\hat{S}_t))\|^2 \mid \mathcal{F}_t \right] \\ &\leq (1 - \alpha) \|V_{t,i} - \mathbf{h}_i(\hat{S}_t)\|^2 + \alpha \sigma_i^2 - \alpha(1 - \alpha) \mathbb{E} \left[\|\Delta_{t+1,i}\|^2 \mid \mathcal{F}_t \right]. \end{aligned} \quad (61)$$

The proof follows from (59), (60) and (61). \blacksquare

D.7 Theorem 1 in the case $\omega_p > 0$ and $\alpha = 0$

Apply Proposition 4 with a time varying coefficient β , denoted by β_{t+1} . Combined with Proposition 2 and Proposition 3, this yields

$$\mathbb{E} \left[f(\mathbb{T}(\hat{S}_{t+1})) + g(\mathbb{T}(\hat{S}_{t+1})) \mid \mathcal{F}_t \right] + \rho_{t+1} \mathbb{E} [G_{t+1} \mid \mathcal{F}_t] \leq f(\mathbb{T}(\hat{S}_t)) + g(\mathbb{T}(\hat{S}_t)) + \rho_t G_t \quad (62)$$

$$+ \gamma_{t+1} \left(\frac{\rho_{t+1}}{\gamma_{t+1}} \beta_{t+1} \left(1 + \frac{2\omega_p}{\beta_{t+1} - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^2 \right) + (1 - \bar{C} v_{\min}^{-1} \gamma_{t+1}) \omega_p v_{\max} - \frac{\rho_t}{\gamma_{t+1}} \right) G_t \quad (63)$$

$$- \frac{\gamma_{t+1}}{4} \left(1 - 2\bar{C} v_{\min}^{-1} \gamma_{t+1} - \frac{8\beta_{t+1}}{\beta_{t+1} - 1} v_{\min}^{-1} \frac{L^2}{n} \rho_{t+1} \gamma_{t+1} \right) \mathbb{E} [\bar{\mathcal{E}}_{t+1} \mid \mathcal{F}_t] \quad (64)$$

$$+ \gamma_{t+1} \left(\rho_{t+1} \beta_{t+1} \frac{2(1 + \omega_p)}{\beta_{t+1} - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1} + (1 - \bar{C} v_{\min}^{-1} \gamma_{t+1}) (1 + \omega_p) v_{\max} \right) \frac{\sigma^2}{n}. \quad (65)$$

Following the same lines as in the proof of the case $\omega_p > 0, \alpha > 0$, choose $\gamma_{t+1} > 0, \rho_{t+1} > 0$ and $\beta_{t+1} > 1$ such that the term in (63) is non positive, the term between parenthesis in (64) is positive and the term in (65) is proportional to γ_{t+1}^u for some $u \geq 1$.

★ When $\alpha = 0$, we can not choose $\gamma_{t+1} = \gamma$ and $\rho_{t+1} = \rho\gamma$ as in the case $\alpha > 0$. Observe indeed that with this strategy, (63) gets into

$$\rho\beta_{t+1} \left(1 + \frac{2\omega_p}{\beta_{t+1} - 1} \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma^2 \right) + (1 - \bar{C} v_{\min}^{-1} \gamma) \omega_p v_{\max} - \rho$$

and it can not be made non positive whatever γ, ρ positive (we have $\beta_{t+1} > 1$).

★ Let us consider the case when $\lim_t \gamma_{t+1} = 0$. Let $\epsilon > 0, u \in (0, 1)$ and choose

$$(1 + 2\epsilon)\gamma_{t+1}^u \leq \gamma_t^u, \quad \rho_{t+1} := \gamma_{t+1}^u, \quad \beta_{t+1} := \frac{1}{1 + \epsilon} \frac{\gamma_t^u}{\gamma_{t+1}^u}; \quad (66)$$

by convention, set $\gamma_0 := (1 + 2\epsilon)^{1/u} \gamma_1$. Then, (63) is upper bounded by

$$\left(\frac{\gamma_t^u}{1 + \epsilon} + 2\omega_p(2 + \epsilon^{-1}) \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^{2+u} + \omega_p v_{\max} \gamma_{t+1} - \gamma_t^u \right) G_t,$$

where we used that $\beta_{t+1}/(\beta_{t+1} - 1)$ is upper bounded by $2 + \epsilon^{-1}$. This term is nonpositive by choosing γ_1 small enough; details are left to the interested reader. (64) is upper bounded by

$$- \frac{\gamma_{t+1}}{4} \left(1 - 2\bar{C} v_{\min}^{-1} \gamma_{t+1} - 8(2 + \epsilon^{-1}) v_{\min}^{-1} \frac{L^2}{n} \gamma_{t+1}^{1+u} \right) \mathbb{E} [\bar{\mathcal{E}}_{t+1} \mid \mathcal{F}_t]$$

which is upper bounded by $-\gamma_{t+1}/8 \mathbb{E} [\bar{\mathcal{E}}_{t+1} \mid \mathcal{F}_t]$ as soon as γ_1 is chosen small enough. Finally, (65) is upper bounded by

$$\gamma_{t+1} \left(2(1 + \omega_p)(2 + \epsilon^{-1}) \frac{v_{\max}}{v_{\min}} \frac{L^2}{n} \gamma_{t+1}^{1+u} + (1 - \bar{C} v_{\min}^{-1} \gamma_{t+1}) (1 + \omega_p) v_{\max} \right) \frac{\sigma^2}{n},$$

which is upper bounded by $2\gamma_{t+1}(1 + \omega_{\mathbf{p}})v_{\max}\sigma^2/n$ by choosing γ_1 small enough. Taking the expectations and summing from $t = 0$ to $t = T - 1$ yields

$$\begin{aligned} \frac{1}{8} \sum_{t=1}^T \gamma_t \mathbb{E} [\bar{\mathcal{E}}_t] &\leq \mathbb{E} \left[(f + g) \circ \mathbb{T}(\hat{S}_0) \right] - \min(f + g) + (1 + 2\epsilon)^2 \gamma_1 G_0 \\ &\quad + 2(1 + \omega_{\mathbf{p}})v_{\max} \frac{\sigma^2}{n} \sum_{t=1}^T \gamma_t. \end{aligned}$$

In this strategy, the stepsize sequence $\{\gamma_t, t \geq 1\}$ is majorized by a sequence decreasing to zero at a geometric rate (see (66)) and $\lim_{T \rightarrow \infty} \sum_{t=1}^T \gamma_t < \infty$. Dividing the LHS and the RHS by $\sum_{t=1}^T \gamma_t$ shows that the RHS can not be made small, even when the number of iterations T is large.

Appendix E. The assumption A7 on few examples

E.1 Quadratic surrogate

Let us prove A7 for Example 1. Since g is a proper lower semicontinuous convex function, ∂g is a monotone operator (see e.g. (Bauschke and Combettes, 2011, Example 20.3)) which implies that the proximal operator $\text{prox}_{\rho g}$ is firmly nonexpansive. Therefore, for all s, s' , we have

$$\| (s' - \text{prox}_{\rho g}(s')) - (s - \text{prox}_{\rho g}(s)) \|^2 \leq \|s' - s\|^2, \quad (67)$$

and

$$|\langle \text{prox}_{\rho g}(s') - \text{prox}_{\rho g}(s), s' - s \rangle| \leq \|s' - s\| \|\text{prox}_{\rho g}(s') - \text{prox}_{\rho g}(s)\| \leq \|s' - s\|^2. \quad (68)$$

Since $\phi(\mathbb{T}(s)) = \rho^{-1} \text{prox}_{\rho g}(s)$, we deduce from (67) and (68) that A7 is verified with

$$\mathbf{B}(s) := \rho^{-1} \mathbf{I}, \quad v_{\min} = v_{\max} = \rho^{-1}, \quad C_{\star} := \rho^{-2}, \quad C_{\star\star} := \rho^{-1}.$$

E.2 Jensen surrogate

EM is a special case of Jensen surrogate (see Appendix C).

We check A7 for EM applied to Poisson observations when $\mathbb{E}_{\pi}[Z]$ is explicit (see Appendix C.1). We choose \mathcal{S} to be equal to the compact interval $[-M, 0]$ for some $M > 0$. This implies that $\lambda - s \in [\lambda, \lambda + M]$. For any $s, s' \in \mathcal{S}$, we write

$$\begin{aligned} \phi(\mathbb{T}(s')) - \phi(\mathbb{T}(s)) &= \mathbb{E}_{\pi}[Z] \left(\frac{1}{\lambda - s'} - \frac{1}{\lambda - s} \right) \\ &= \frac{\mathbb{E}_{\pi}[Z]}{\lambda - s} \frac{(s' - s)}{\lambda - s'} = \frac{\mathbb{E}_{\pi}[Z]}{(\lambda - s)^2} (s' - s) + \frac{\mathbb{E}_{\pi}[Z]}{(\lambda - s)^2} \frac{(s' - s)^2}{\lambda - s'}. \end{aligned}$$

Hence A7 is verified with

$$\mathbf{B}(s) := \frac{\mathbb{E}_{\pi}[Z]}{(\lambda - s)^2}, \quad v_{\min} := \frac{1}{(\lambda + M)^2}, \quad v_{\max} := \frac{1}{\lambda^2}, \quad C_{\star} := \frac{\mathbb{E}_{\pi}[Z]}{\lambda^3}, \quad C_{\star\star} := \frac{\mathbb{E}_{\pi}[Z]}{\lambda^2}.$$

E.3 Variational surrogate

Let us consider the variational surrogate applied to Dictionary Learning / Matrix Factorization (see Section 2.3). We have

$$\phi(\theta) = \begin{bmatrix} -\theta^\top \theta \\ 2\theta \end{bmatrix}$$

and $s = (s^{(1)}, s^{(2)}) \in \mathcal{M}_K^+ \times \mathbb{R}^{p \times K}$. Here, $\mathbf{B}(s)$ is not an invertible matrix so that the condition $v_{\min} > 0$ does not hold. To be convinced, consider the case when $p = K = 1$ so that $\phi : \mathbb{R} \rightarrow \mathbb{R}^2$ is equal to $(-u^2, 2u)$; and $s := (s^{(1)}, s^{(2)}) \in \mathbb{R}_{>0} \times \mathbb{R}$. Then $\mathbf{T}(s) = s^{(2)}/s^{(1)}$. It is easily checked that we have

$$\mathbf{B}(s) = \begin{bmatrix} \frac{1}{s^{(1)}} & -\frac{s^{(2)}}{(s^{(1)})^2} \\ -\frac{s^{(2)}}{(s^{(1)})^2} & \frac{(s^{(2)})^2}{(s^{(1)})^3} \end{bmatrix} = \frac{1}{s^{(1)}} \begin{bmatrix} 1 \\ -\frac{s^{(2)}}{s^{(1)}} \end{bmatrix} \begin{bmatrix} 1 \\ -\frac{s^{(2)}}{s^{(1)}} \end{bmatrix}^\top.$$

We developed the convergence analysis under the assumption that $\mathbf{B}(s)$ is invertible (with a positive minimal eigenvalue uniformly lower bounded for $s \in \mathcal{S}$). When this assumption does not hold, the proof has to be generalized in order to define projections and scalar product in a space defined by the image of $\mathbf{B}(\hat{\mathcal{S}}_t)$. This makes the proof more intricate, it is out of the scope of this paper.

References

- D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic. QSGD: Communication-Efficient SGD via Gradient Quantization and Encoding. *Advances in Neural Information Processing Systems*, 30:1709–1720, 2017.
- B. Amos, L. Xu, and J. Kolter. Input Convex Neural Networks. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 146–155. PMLR, 06–11 Aug 2017.
- C. Andrieu, E. Moulines, and P. Priouret. Stability of Stochastic Approximation under Verifiable Conditions. *SIAM Journal on Control and Optimization*, 44(1):283–312, 2005.
- Y. Atchadé, G. Fort, and E. Moulines. On Perturbed Proximal Gradient Algorithms. *Journal of Machine Learning Research*, 18(10):1–33, 2017.
- H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer Publishing Company, Incorporated, 2011.
- A. Beck and M. Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009a.
- A. Beck and M. Teboulle. *Gradient-based algorithms with applications to signal-recovery problems*, page 42–88. Cambridge University Press, 2009b.
- A. Benveniste, M. Métivier, and P. Priouret. *Adaptive Algorithms and Stochastic Approximations*. Springer Verlag, 1990.
- V. S. Borkar. *Stochastic Approximation. A Dynamical Systems Viewpoint*. Cambridge University Press, Cambridge; Hindustan Book Agency, New Delhi, 2008.
- P. Bradley, K. Bennett, and A. Demiriz. Constrained K-Means Clustering. *Microsoft Research, Redmond*, 20(0):0, 2000.
- L. D. Brown. *Fundamentals of Statistical Exponential Families with Applications in Statistical Decision Theory*. Lecture notes-monograph series Fundamentals of statistical exponential families. Institute of Mathematical Statistics, 1986.
- T. Canh, T. Nguyen, and N. Josh. Personalized Federated Learning with Moreau Envelopes. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21394–21405. Curran Associates, Inc., 2020.
- O. Cappé and É. Moulines. On-Line Expectation–Maximization Algorithm for Latent Data Models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3): 593–613, 2009.
- G. Celeux and J. Diebolt. A stochastic approximation type EM algorithm for the mixture problem. *Stochastics and Stochastic Reports*, 41(1-2):119–134, 1992.

- J. Chen, J. Zhu, Y. Teh, and T. Zhang. Stochastic Expectation Maximization with Variance Reduction. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 7967–7977, 2018.
- M. Collins, R. E. Schapire, and Y. Singer. Logistic Regression, AdaBoost and Bregman Distances. *Machine Learning*, 48(1):253–285, 2002.
- P. L. Combettes and J.-C. Pesquet. *Proximal Splitting Methods in Signal Processing*, pages 185–212. Springer New York, New York, NY, 2011.
- L. Condat and P. Richtárik. RandProx: Primal-Dual Optimization Algorithms with Randomized Proximal Updates. In *The Eleventh International Conference on Learning Representations*. OpenReview.net, 2023.
- X. Dai, X. Yan, K. Zhou, H. Yang, K. K. Ng, J. Cheng, and Y. Fan. Hyper-Sphere Quantization: Communication-Efficient SGD for Federated Learning. Technical report, arXiv:1911.04655, 2019.
- S. Della Pietra, V. Della Pietra, and J. Lafferty. Duality and Auxiliary Functions for Bregman Distances. Technical report, Carnegie-Mellon University, 2001.
- B. Delyon, M. Lavielle, and E. Moulines. Convergence of a Stochastic Approximation Version of the EM Algorithm. *Ann. Statist.*, 27(1):94–128, 1999.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society: Series B*, 39:1–38, 1977.
- A. Dieuleveut, G. Fort, E. Moulines, and G. Robin. Federated-EM with heterogeneity mitigation and variance reduction. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 29553–29566. Curran Associates, Inc., 2021.
- A. Dieuleveut, G. Fort, E. Moulines, and H.-T. Wai. Stochastic Approximation Beyond Gradient for Signal Processing and Machine Learning. *IEEE Transactions on Signal Processing*, 71:3117–3148, 2023.
- C. Dinh, T. T. Vu, N. H. Tran, M. N. Dao, and H. Zhang. FedU: A Unified Framework for Federated Multi-Task Learning with Laplacian Regularization. Technical report, arXiv:2102.07148, 2021.
- R. M. Dudley. The speed of mean glivenko-cantelli convergence. *The Annals of Mathematical Statistics*, 40(1):40–50, 1969.
- B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least Angle Regression. *The Annals of Statistics*, 32(2), April 2004.
- A. Flanagan, W. Oyomno, A. Grigorievskiy, K. E. Tan, S. Khan, and M. Ammad-Ud-Din. Federated Multi-view Matrix Factorization for Personalized Recommendations. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 324–347. Springer, 2020.

- G. Fort and E. Moulines. Convergence of the Monte Carlo Expectation Maximization for Curved Exponential Families. *Ann. Statist.*, 31(4):1220–1259, 2003.
- G. Fort and E. Moulines. Stochastic variable metric proximal gradient with variance reduction for non-convex composite optimization. *Stat Comput*, 33(65), 2023.
- G. Fort, P. Gach, and E. Moulines. Fast incremental expectation maximization for finite-sum optimization: nonasymptotic convergence. *Statistics and Computing*, 31, 2021a.
- G. Fort, E. Moulines, and H.-T. Wai. Geom-Spider-EM: Faster Variance Reduced Stochastic Expectation Maximization for Nonconvex Finite-Sum Optimization. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3135–3139, 2021b.
- D. Gao, B. Tan, C. Ju, V. W. Zheng, and Q. Yang. Privacy Threats Against Federated Matrix Factorization. Technical report, arXiv:2007.01587, 2020.
- S. Ghadimi and G. Lan. Stochastic First- and Zeroth-Order Methods for Nonconvex Stochastic Programming. *SIAM J. Optimiz.*, 23(4):2341–2368, 2013.
- A. Ghosh, J. Chung, D. Yin, and K. Ramchandran. An Efficient Framework for Clustered Federated Learning. *Advances in Neural Information Processing Systems*, 33:19586–19597, 2020.
- E. Gorbunov, Z. Burlachenko, K. P. and Li, and P. Richtárik. MARINA: Faster Non-Convex Distributed Learning with Compression. In *International Conference on Machine Learning*, pages 3788–3798. PMLR, 2021.
- F. Hanzely, S. Hanzely, S. Horváth, and P. Richtárik. Lower Bounds and Optimal Algorithms for Personalized Federated Learning. *Advances in Neural Information Processing Systems*, 33:2304–2315, 2020.
- F. Harper and J. A. Konstan. The Movielens Datasets: History and context. *ACM Transactions on interactive intelligent systems*, 5(4):1–19, 2015.
- I. Hegedűs, G. Danner, and M. Jelasity. Decentralized Recommendation Based on Matrix Factorization: A Comparison of Gossip and Federated Learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 317–332. Springer, 2019.
- R. Islamov, X. Qian, and P. Richtárik. Distributed Second Order Methods with Fast Rates and Compressed Communication. In *International Conference on Machine Learning*, pages 4617–4628. PMLR, 2021.
- A. K. Jain, M. N. Murty, and P. J. Flynn. Data clustering: a review. *ACM computing surveys (CSUR)*, 31(3):264–323, 1999.
- P. Kairouz, H. McMahan, et al. *Advances and Open Problems in Federated Learning*. New Foundations and Trends, 2021.

- S. P. Karimireddy, Q. Rebjock, S. Stich, and M. Jaggi. Error Feedback Fixes SignSGD and other Gradient Compression Schemes. In *International Conference on Machine Learning*, pages 3252–3261. PMLR, May 2019. ISSN: 2640-3498.
- S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In *International Conference on Machine Learning*, pages 5132–5143. PMLR, 2020.
- A. Khaled and P. Richtárik. Gradient Descent with Compressed Iterates. Technical report, arXiv:1909.04716, 2020.
- D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*, January 2017. arXiv: 1412.6980.
- A. Koloskova, S. Stich, and M. Jaggi. Decentralized Stochastic Optimization and Gossip Algorithms with Compressed Communication. In *International conference on machine learning*, pages 3478–3487. PMLR, 2019.
- J. Konečný, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon. Federated Learning: Strategies for Improving Communication Efficiency. In *NIPS Workshop on Private Multi-Party Machine Learning*, 2016.
- A. Korotin, V. Egiazarian, A. Asadulaev, A. Safin, and E. Burnaev. Wasserstein-2 Generative Networks. In *International Conference on Learning Representations*, 2021a.
- A. Korotin, L. Li, A. Genevay, J. Solomon, A. Filippov, and E. Burnaev. Do Neural Optimal Transport Solvers Work? A Continuous Wasserstein-2 Benchmark. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021b.
- F. Kunstner, R. Kumar, and M. Schmidt. Homeomorphic-Invariance of EM: Non-Asymptotic Convergence in KL Divergence for Exponential Families via Mirror Descent. In *International Conference on Artificial Intelligence and Statistics*, pages 3295–3303. PMLR, 2021.
- K. Lange. *Optimization*. Springer Texts in Statistics. Springer New York, 2013.
- K. Lange, D. Hunter, and I. Yang. Optimization Transfer Using Surrogate Objective Functions. *Journal of computational and graphical statistics*, 9(1):1–20, 2000.
- L. Leconte, A. Dieuleveut, E. Oyallon, E. Moulines, and G. Pages. DoStoVoQ: Doubly Stochastic Voronoi Vector Quantization SGD for Federated Learning. Technical report, openreview.net, 2021.
- D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.
- X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang. On the Convergence of FedAvg on Non-IID Data. In *International Conference on Learning Representations*, 2020.

- J. Mairal. Stochastic Majorization-Minimization Algorithms for Large-Scale Optimization. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- J. Mairal. Incremental Majorization-Minimization Optimization with Application to Large-Scale Machine Learning. *SIAM Journal on Optimization*, 25(2):829–855, 2015.
- J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *J. Mach. Learn. Res.*, 11:19–60, 2010.
- A. Makkua, A. Taghvaei, S. Oh, and J. Lee. Optimal transport mapping via input convex neural networks. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6672–6681. PMLR, 2020.
- Y. Mansour, M. Mohri, J. Ro, and A. Suresh. Three Approaches for Personalization with Applications to Federated Learning. Technical report, arXiv:2002.10619, 2020.
- O. Marfoq, G. Neglia, A. Bellet, L. Kamani, and R. Vidal. Federated Multi-Task Learning under a Mixture of Distributions. *Advances in Neural Information Processing Systems*, 34, 2021.
- G. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. Wiley series in probability and statistics. Wiley, 2008.
- B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Aguera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017a.
- B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, April 2017b. ISSN: 2640-3498.
- K. Mishchenko, E. Gorbunov, M. Takáč, and P. Richtárik. Distributed Learning with Compressed Gradient Differences. Technical report, arXiv:1901.09269, 2019.
- K. Mishchenko, G. Malinovsky, S. Stich, and P. Richtarik. ProxSkip: Yes! Local gradient steps provably lead to communication acceleration! Finally! In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 15750–15769, 2022.
- J.-J. Moreau. Fonctions convexes duales et points proximaux dans un espace hilbertien. *C. R. Acad. Sci., Paris*, 255:2897–2899, 1962.
- R. M. Neal and G. E. Hinton. A View of the EM Algorithm that Justifies Incremental, Sparse, and other Variants. In M. I. Jordan, editor, *Learning in Graphical Models*, pages 355–368. Springer Netherlands, Dordrecht, 1998.

- H. Nguyen, F. Forbes, G. Fort, and O. Cappé. An Online Minorization-Maximization Algorithm. In P. Brito, J. G. Dias, B. Lausen, A. Montanari, and R. Nugent, editors, *Classification and Data Science in the Digital Age*, pages 263–271. Springer International Publishing, 2023.
- P. Paatero and U. Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2):111–126, 1994.
- N. Parikh and S. Boyd. Proximal Algorithms. *Found. Trends Optim.*, 1(3):127–239, 2014.
- C. Philippenko and A. Dieuleveut. Preserved central model for faster bidirectional compression in distributed settings. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- R. Polyak. *Introduction to Continuous Optimization*. Springer Optimization and Its Applications. Springer, April 2021.
- S. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, K. S., and H. McMahan. Adaptive Federated Optimization. In *International Conference on Learning Representations*, 2021.
- P. Richtárik, I. Sokolov, and I. Fatkhullin. EF21: A New, Simpler, Theoretically Better, and Practically Faster Error Feedback. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- H. Robbins and S. Monro. A Stochastic Approximation Method. *Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- M.-A. Sato and S. Ishii. On-line EM Algorithm for the Normalized Gaussian Network. *Neural Computation*, 12(2):407–432, 2000.
- F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek. Robust and Communication-Efficient Federated Learning From Non-i.i.d. Data. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–14, 2019.
- F. Sattler, K.-R. Müller, and W. Samek. Clustered Federated Learning: Model-Agnostic Distributed Multitask Optimization under Privacy Constraints. *IEEE Transactions on Neural Networks and Learning Systems*, 32(8):3710–3722, 2020.
- K. Singhal, H. Sidahmed, Z. Garrett, S. Wu, J. Rush, and S. Prakash. Federated Reconstruction: Partially Local Federated Learning. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 11220–11232. Curran Associates, Inc., 2021.
- S. Stich. Local SGD Converges Fast and Communicates Little. In *International Conference on Learning Representations*, 2019.
- Y. Tian, H. Weng, and Y. Feng. Towards the Theory of Unsupervised Federated Learning: Non-asymptotic Analysis of Federated EM Algorithms. In *Forty-first International Conference on Machine Learning*, 2024.

- C. Villani. *Topics in Optimal Transportation*, volume 58. American Mathematical Soc., 2021.
- J. Wang, Z. Charles, et al. A Field Guide to Federated Optimization. Technical report, arXiv:2107.06917, 2021.
- S. Wang and T.-H. Chang. Federated Matrix Factorization: Algorithm Design and Application to Data Clustering. *IEEE Transactions on Signal Processing*, 70:1625–1640, 2022.
- J. Wangni, J. Wang, J. Liu, and T. Zhang. Gradient Sparsification for Communication-Efficient Distributed Optimization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- G. Wei and M. Tanner. A Monte Carlo Implementation of the EM Algorithm and the Poor Man’s Data Augmentation Algorithms. *J. Am. Stat. Assoc.*, 85(411):699–704, 1990.
- B. Woodworth, K. Patel, S. Stich, Z. Dai, B. Bullins, B. McMahan, O. Shamir, and N. Srebro. Is Local SGD Better than Minibatch SGD? In *International Conference on Machine Learning*, pages 10334–10343. PMLR, 2020.
- S. J. Wright, R. Nowak, and M. Figueiredo. Sparse Reconstruction by Separable Approximation. *IEEE Transactions on signal processing*, 57(7):2479–2493, 2009.
- H. Yuan, M. Zaheer, and S. Reddi. Federated Composite Optimization. In *International Conference on Machine Learning*, pages 12253–12266. PMLR, 2021.
- T. Zhixu, C. Rajita, and K. Sanjeev. On the Convergence of a Federated Expectation-Maximization Algorithm. Technical report, arXiv:2408.05819, 2024.