

Algorithmes Majoration-Minoration stochastiques pour l'Apprentissage Statistique grande échelle

Gersende Fort

CNRS

Institut de Mathématiques de Toulouse, France



Journées de la SFDS, 2022, Lyon.

In collaboration with

- Aymeric Dieuleveut, Ecole Polytechnique, CMAP, France
- Eric Moulines, Ecole Polytechnique, CMAP, France
- Geneviève Robin, CNRS, LaMME, France
- Hoi-To Wai, Chinese Univ. of Hong-Kong, Hong-Kong

Partly funded by

Fondation Simone et Cino Del Duca, Project OpSiMorE

I. Theme and Motivation

General theme: Stochastic Optimization

- Composite objective function

$$\text{Argmin}_{\mathbb{R}^d} (f(\theta) + g(\theta)) \quad f(\theta) := \mathbb{E}_{Z \sim \pi} [\ell(Z, \theta)]$$

- the function $f : \mathbb{R}^d \rightarrow \mathbb{R}$
 - $\mathbb{E}_{\pi}[\cdot]$ **can not** be explicitly evaluated
 - *mean value* of a loss over the examples (observations, data)
 - possibly **non convex**, continuously differentiable
- the function $g : \mathbb{R}^d \rightarrow (0, +\infty]$
 - explicit evaluation of $g(\theta)$
 - a regularization / penalization term, constraints on θ
 - $\Theta := \{\theta \in \mathbb{R}^d : g(\theta) < \infty\}$
 - convex, proper, lower semi-continuous

Computational Statistical Learning

$$\text{Argmin}_{\mathbb{R}^d} (f(\theta) + g(\theta)) \quad f(\theta) := \mathbb{E}_{Z \sim \pi} [\ell(Z, \theta)] \text{ noté aussi } \mathbb{E}_{\pi} [\ell(Z, \theta)]$$

- Large batch learning: empirical loss, finite sum setting

$$\pi := \frac{1}{N} \sum_{j=1}^N \delta_{Z_j} \quad f(\theta) = \frac{1}{N} \sum_{j=1}^N \ell(Z_j, \theta)$$

- Online learning: expected loss

$$\text{data stream } \{Z_j, j \geq 0\} \sim \pi \quad f(\theta) = \int \ell(z, \theta) \pi(\mathrm{d}z)$$

- Examples of loss functions

- quadratic $\|Z - \Xi(\theta)\|^2$
- linear regression, quadratic loss $\|Y - X\theta\|^2$, $Z = (Y, X)$
- negative log-likelihood $-\text{loglike}(Z; \theta)$

What we intend to do

Table of contents


- Among the family of **Majorize-Minimization** (MM) optimization algorithm
 - what is MM ?
 - why MM ?
- Stochastic Variance reduction techniques for large scale learning
 - why "stochastic" is required ?
 - A **mirror** definition of MM and its stochastic version
 - Why "variance reduction"
 - Case "large batch": A novel Variance-Reduced MM
- Federated Learning (FL)
 - what is FL ?
 - Stochastic MM in the FL setting

Optimization tool: Majorize-Minimization

See e.g. K. Lange "Optimization", Springer-Verlag, 2013.

The Majorize-Minimization algorithm

Iterative algorithm $\{\theta_k, k \geq 0\}$



- Given the current value θ_k
 - define a majorizing function (red)
 - tangent at θ_k
 - minimizer of the majorizing function $\rightarrow \theta_{k+1}$

- Key property:

$$f(\theta_{k+1}) \leq f(\theta_k)$$

Proof:

$$\begin{aligned} f(\theta_{k+1}) &\leq M(\theta_{k+1}; \theta_k) && M(\cdot; \theta_k) \text{ majorizes } f \\ &\leq M(\theta_k; \theta_k) && \theta_{k+1} \text{ minimizer of } M(\cdot; \theta_k) \\ &= f(\theta_k) && \text{tangent property: } f(\theta_k) = M(\theta_k; \theta_k) \end{aligned}$$

- MM will be used to majorize f ; we then deduce a majorizing function for $f + g$.

Example 1: gradient-based algorithms

Assume that the loss function f is smooth with L_f -Lipschitz gradient on Θ .

- Ex. In large scale learning: $\nabla f(\theta) = \mathbb{E}_\pi [G(Z, \theta)]$

- **Majorizing function of f , tangent at $\theta_k \in \Theta$:**

$$\begin{aligned} f(\theta) &\leq f(\theta_k) + \langle \nabla f(\theta_k), \theta - \theta_k \rangle + \frac{L_f}{2} \|\theta - \theta_k\|^2 \\ &\leq C_k + \frac{1}{\gamma} \left\{ \frac{1}{2} \|\theta\|^2 - \langle s_k, \theta \rangle \right\} \quad s_k := \theta_k - \gamma \nabla f(\theta_k) = \theta_k - \gamma \mathbb{E}_\pi [G(Z, \theta_k)] \quad \gamma \leq 1/L_f \end{aligned}$$

- **Minimization step.**

$$\begin{aligned} T(s_k) &= \text{prox}_{\gamma g}(s_k) \\ &= \text{prox}_{\gamma g}(\theta_k - \gamma \nabla f(\theta_k)) \end{aligned}$$

Gradient-based algorithms are among MM algorithms

With quadratic majorizing fcts.

Example 2: Expectation Maximization algorithms Dempster et al. (1977); Wu (1983)

Assume that the loss function:

$$\ell(Z, \theta) := -\log \int_{\mathcal{H}} p(Z, h, \theta) \nu(\mathrm{d}h), \quad \log p(Z, h, \theta) = \langle S(Z, h), \phi(\theta) \rangle - \psi(\theta)$$

- Ex.: Negative log-like, Latent variable model, complete data model in the curved exponential family.
- **Majorizing function of f , tangent at $\theta_k \in \Theta$. (E-step)**

Jensen's inequality:

$$\begin{aligned} \ell(Z, \theta) &\leq C_k - \int_{\mathcal{H}} \log p(Z, h, \theta) \nu(\mathrm{d}h|Z, \theta_k) \\ &\leq C_k + \psi(\theta) - \langle \int_{\mathcal{H}} S(Z, h) \nu(\mathrm{d}h|Z, \theta_k), \phi(\theta) \rangle. \end{aligned}$$

- **Optimization step. (M-step)**

$$T(s_\tau) = \operatorname{argmin}_{\theta} g(\theta) + \psi(\theta) - \langle s_k, \phi(\theta) \rangle$$

$$s_k := \mathbb{E}_{\pi} \left[\int_{\mathcal{H}} S(Z, h) \nu(\mathrm{d}h|Z, \theta_k) \right]$$

EM algorithms are among
MM algorithms
Majorizing fct **not**
quadratic in g^{al}

Example 3: Variational Surrogates

Assume that the loss function

$$\ell(Z, \theta) = \min_{h \in \mathcal{H}} \tilde{\ell}(Z, h, \theta), \quad \tilde{\ell}(Z, h, \theta) := \psi(\theta) - \langle S(Z, h), \phi(\theta) \rangle$$

Set $\mathbf{M}(Z, \theta) := \operatorname{argmin}_{h \in \mathcal{H}} \tilde{\ell}(Z, h, \theta)$.

- **Majorizing function of f , at $\theta_k \in \Theta$.**

$$\ell(Z, \theta) \leq C_k + \psi(\theta) - \langle S(Z, \mathbf{M}(Z, \theta_k)), \phi(\theta) \rangle$$

- **Optimization step.**

$$\mathbf{T}(s_k) = \operatorname{argmin}_{\theta} g(\theta) + \psi(\theta) - \langle s_k, \phi(\theta) \rangle$$

$$s_k := \mathbb{E}_{\pi} [S(Z, \mathbf{M}(Z, \theta_k))]$$

- **Example:**

Dictionary Learning

$$f(\theta) = \mathbb{E}_{\pi} \left[\min_{h \in \mathcal{H}} \|Z - \theta h\|^2 \right]$$

Z : observations $d \times 1$

θ : dictionary $d \times K$

h : code $K \times 1$

All these problems share

- A parametric majorizing function of $f + g$ at $\tau \in \Theta$, of the form

$$\theta \mapsto C_\tau + g(\theta) + \psi(\theta) - \langle \mathbb{E}_\pi [\bar{S}(Z, \tau)], \phi(\theta) \rangle$$

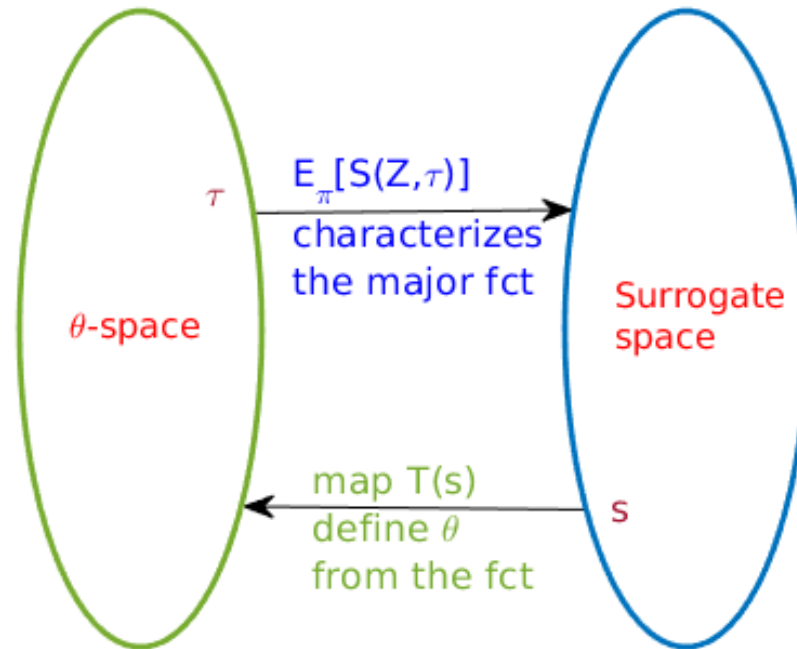
- The iterative process

$$\cdot \rightarrow \theta_k \rightarrow \mathbb{E}_\pi [\bar{S}(Z, \theta_k)] \xrightarrow{\mathbb{T}} \theta_{k+1} \rightarrow \mathbb{E}_\pi [\bar{S}(Z, \theta_{k+1})] \xrightarrow{\mathbb{T}} \theta_{k+2} \rightarrow \dots$$

or equivalently

$$\cdot \rightarrow \theta_k \rightarrow \underbrace{\mathbb{E}_\pi [\bar{S}(Z, \theta_k)]}_{s_k} \xrightarrow{\mathbb{T}} \underbrace{\theta_{k+1}}_{\mathbb{T}(s_k)} \rightarrow \underbrace{\mathbb{E}_\pi [\bar{S}(Z, \mathbb{T}(s_k))]}_{s_{k+1}} \xrightarrow{\mathbb{T}} \underbrace{\theta_{k+2}}_{\mathbb{T}(s_{k+1})} \rightarrow \underbrace{\ddots}_{s_{k+2}}$$

A unifying point of view Dieuleveut, F., Wai (2022); Nguyen, Forbes, F., Cappé (2022)



- ... "the surrogate-space, in the foreground ! the θ -space in the background"
- the θ -space is the *mirror*
- **Novel and unified approaches** for (i) large scale learning, (ii) federated learning.

A specific structure for the majorizing functions

$$\text{Argmin}_{\mathbb{R}^d} (f(\theta) + g(\theta)) \quad f(\theta) := \mathbb{E}_{Z \sim \pi} [\ell(Z, \theta)]$$

We assume hereafter

- Hyp. **MM-1** There exist $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$, $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^q$, $\bar{S} : \mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathcal{S} \subseteq \mathbb{R}^q$ s.t.

$$\forall \tau, \quad \forall \theta \in \Theta : \quad f(\theta) \leq f(\tau) + \psi(\theta) - \psi(\tau) - \langle \mathbb{E}_{\pi} [\bar{S}(Z, \tau)], \phi(\theta) - \phi(\tau) \rangle$$

i.e. for any $\tau \in \Theta$, there exists a majorizing function for f , tangent at τ

The majorizing function of $(f + g)$ is in a *parametric* family of functions

$$\theta \mapsto C_{\tau} + g(\theta) + \psi(\theta) - \langle s_{\tau}, \phi(\theta) \rangle \quad \text{where } s_{\tau} := \mathbb{E}_{\pi} [\bar{S}(Z, \tau)]$$

- Hyp. **MM-2** For any $s \in \mathcal{S}$,

$$\mathbf{T}(s) := \text{argmin}_{\theta} (g(\theta) + \psi(\theta) - \langle s, \phi(\theta) \rangle)$$

exists and is unique.

Stationary points in the θ -space / stationary points in the surrogate space

- Under regularity assumptions, for any $s \in \mathcal{S}$, there exists $p_s \in \partial g(\mathbb{T}(s))$,

$$p_s + \nabla f(\mathbb{T}(s)) = -J_\phi(\mathbb{T}(s)) \mathbf{h}(s) \quad \mathbf{h}(s) := \mathbb{E}_\pi [\bar{S}(Z, \mathbb{T}(s))] - s.$$

Dieuleveut, F., Wai (2022)

If $s_\star \in \mathcal{S}$ satisfies $\mathbf{h}(s_\star) = 0$ then $\theta_\star := \mathbb{T}(s_\star)$ is a stationary point of $f + g$.

And conversely.

- **We will design algorithms**
 - in the s -space,
 - targeting a zero of $\mathbf{h}(s)$
 - $\mathbf{h}(s)$ is an untractable expectation (w.r.t. π) !

**Stochastic surrogate MM
and
Variance reduction**

The Stochastic surrogate MM algorithm Dieuleveut, F., Wai (2022)

- Based on **Stochastic Approximation** Robbins and Monro (1951); Benveniste et al. (1990):

given learning rates $\{\gamma_k, k \geq 0\}$

$$\hat{S}_{k+1} = \hat{S}_k + \gamma_{k+1} H_{k+1} \quad H_{k+1} \approx \mathbf{h}(\hat{S}_k) = \mathbb{E}_\pi [\bar{S}(Z, \mathbf{T}(\hat{S}_k))] - \hat{S}_k.$$

Algorithm 1: Stochastic Surrogate MM (StoSur-MM)

Input: $k_{\max} > 0, \hat{S}_0 \in \mathcal{S}$

Result: StoSur-MM sequence $\{\hat{S}_k, k \leq k_{\max}\}$ and its mirror $\{\mathbf{T}(\hat{S}_k), k < k_{\max}\}$

for $k = 0, \dots, k_{\max} - 1$ **do**

 Compute $\mathbf{T}(\hat{S}_k)$

 Sample \mathbf{S}_{k+1} , a random oracle for $\mathbb{E}_\pi[\bar{S}(Z, \mathbf{T}(\hat{S}_k))]$

 Set $\hat{S}_{k+1} = \hat{S}_k + \gamma_{k+1}(\mathbf{S}_{k+1} - \hat{S}_k)$

- **Examples of oracles:**

- Large batch learning:

$$\mathbf{S}_{k+1} = \frac{1}{b} \sum_{j \in \mathcal{B}_{k+1}} \bar{S}(Z_j, \mathbf{T}(\hat{S}_k))$$

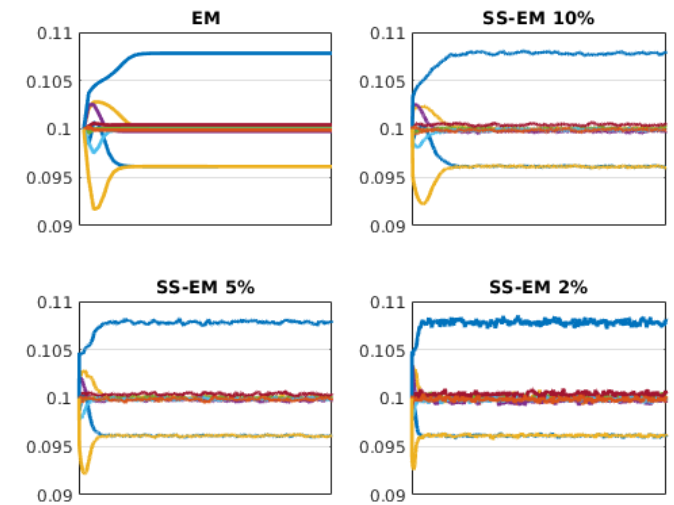
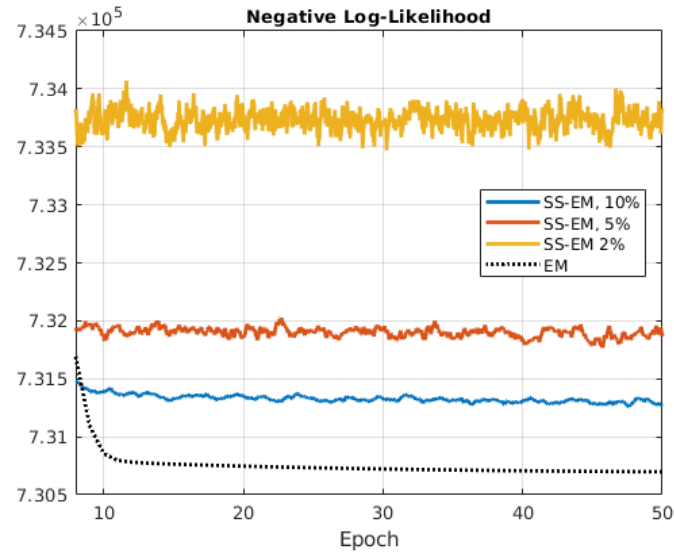
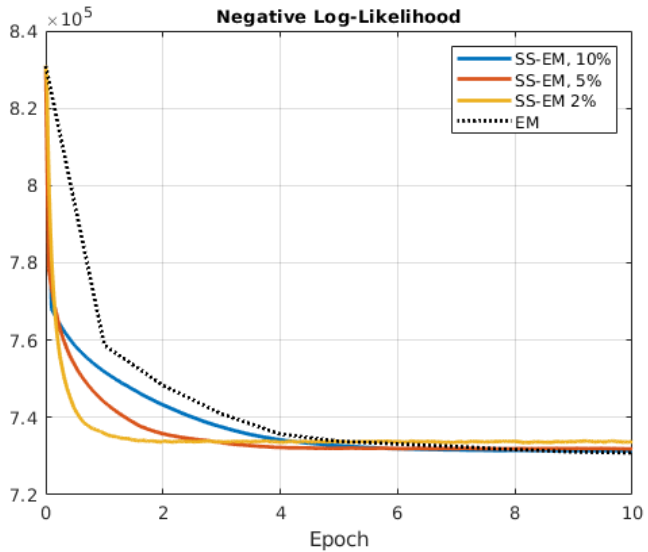
- Online learning:

$$\mathbf{S}_{k+1} = \bar{S}(Z_{k+1}, \mathbf{T}(\hat{S}_k))$$

- In the Gdt Case, known as "Stochastic Gdt". In the EM case: "Online EM" Cappé and Moulines (2009)

On a toy example

- Mixture of 10 Gaussian distributions in \mathbb{R}^{21} . $N = 6e4$ examples.
- θ collects the weights, the expectations and the covariance matrices.
- Toy example: EM could be applied; let us do as if it were not and run SS-EM

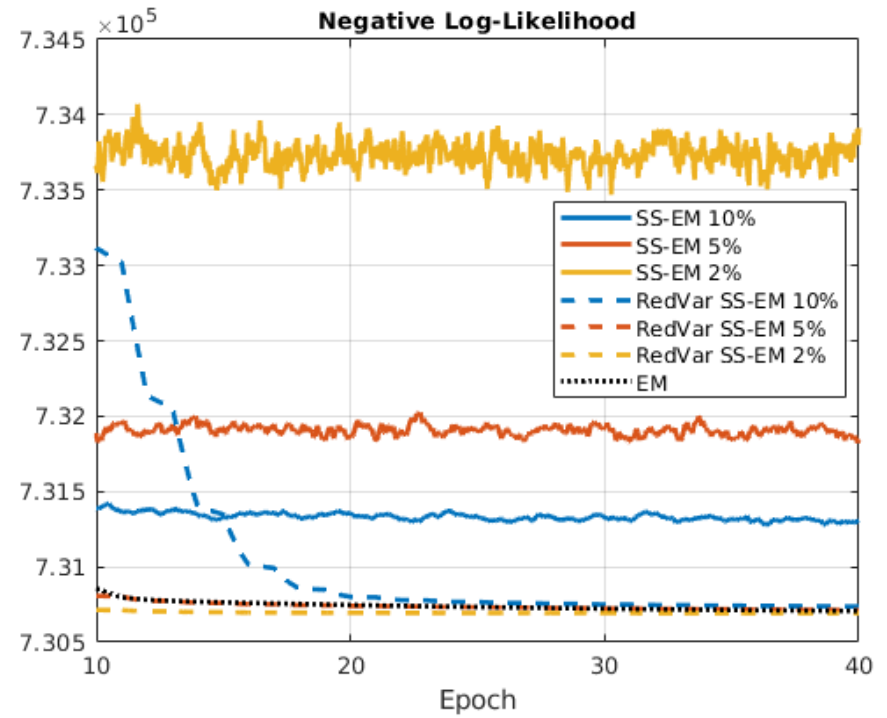
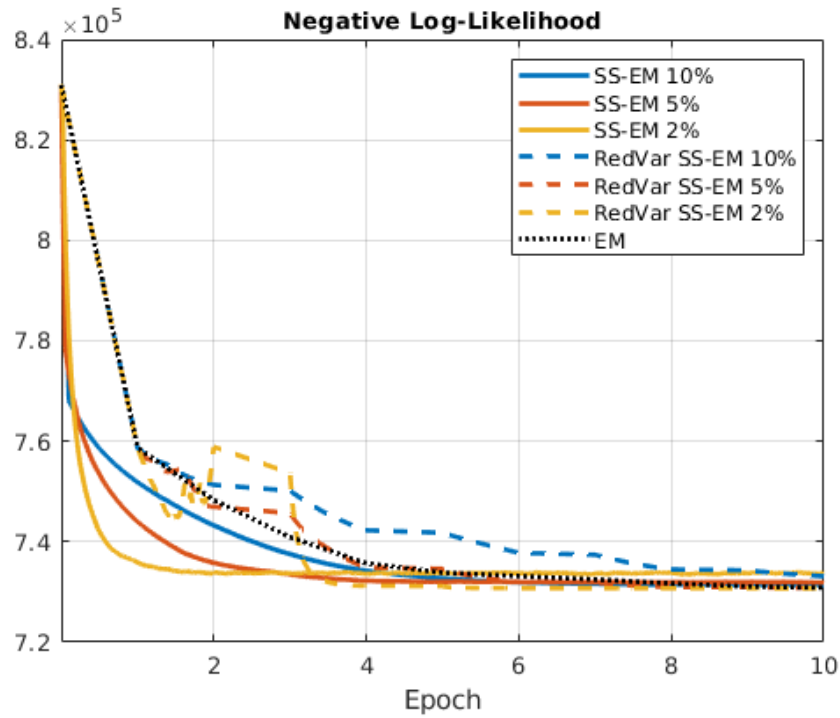


Stochastic Surrogate MM for different mini batch size.

Evolution of the negative log-likelihood, first epochs (left) and after 8 epochs (center).

Evolution of the 10 estimated weights vs the iteration index (right).

Toy example (to follow): let us introduce variance reduction



Stochastic Surrogate MM and **Variance reduced** Stochastic Surrogate MM for different mini batch size.

Evolution of the negative log-likelihood, first epochs (left) and after 10 epochs (right).

Large batch learning: Variance reduction within Stochastic Approximation (SA)

- Control variates

- If $\mathbf{h}(s) = \mathbb{E}[H_s]$ and $\mathbb{E}[V] = 0$, then $\mathbf{h}(s) = \mathbb{E}[H_s + V] \rightarrow$ many possibilities for the definition of SA
- "If U and V are negatively correlated, there exists c such that $\text{Var}(U + cV) < \text{Var}(U)$ ".

- Examples:

...

incremental EM

Neal and Hinton (1998)

Ng and McLachlan (2003)

SVRG

Johnson and Zhang (2013)

SEM-VR

Chen et al. (2018)

SAGA

Defazio et al. (2014)

FIEM

Karimi et al. (2019)


F., Gach, Moulines (2021)

SPIDER

Fang et al. (2018)

SPIDER-EM

F., Moulines, Wai (2020)



Nbr of processed examples required to reach convergence, as a fct of the problem size n (F., Moulines, Wai, 2020)

Algorithm 2: Stochastic Path-Integrated Differential Estimator - EM

Data: $k_{\text{in}} \in \mathbb{N}_*$, $k_{\text{out}} \in \mathbb{N}_*$, $\hat{S}_{\text{init}} \in \mathcal{S}$, $\{\gamma_{t,k+1}, t \geq 1, k \geq 0\}$ positive sequence.

Result: The SPIDER-EM sequence: $\hat{S}_{t,k}, t = 1, \dots, k_{\text{out}}$ and $k = 0, \dots, k_{\text{in}} - 1$

$$\hat{S}_{1,0} = \hat{S}_{1,-1} = \hat{S}_{\text{init}}, \quad \mathbf{S}_{1,0} = N^{-1} \sum_j \bar{S}(Z_j, \mathbb{T}(\hat{S}_{1,-1}))$$

for $t = 1, \dots, k_{\text{out}}$ **do**

for $k = 0, \dots, k_{\text{in}} - 2$ **do**

 Sample a mini-batch $\mathcal{B}_{t,k+1}$ in $\{1, \dots, N\}$ of size \mathbf{b} , with or without replacement

$$\mathbf{S}_{t,k+1} = \mathbf{S}_{t,k} + \mathbf{b}^{-1} \sum_{j \in \mathcal{B}_{t,k+1}} \left\{ \bar{S}(Z_j, \mathbb{T}(\hat{S}_{t,k})) - \bar{S}(Z_j, \mathbb{T}(\hat{S}_{t,k-1})) \right\}$$

$$\hat{S}_{t,k+1} = \hat{S}_{t,k} + \gamma_{t,k+1} (\mathbf{S}_{t,k+1} - \hat{S}_{t,k})$$

$$\hat{S}_{t+1,-1} = \hat{S}_{t,k_{\text{in}}-1}$$


$$\mathbf{S}_{t+1,0} = N^{-1} \sum_j \bar{S}(Z_j, \mathbb{T}(\hat{S}_{t+1,-1}))$$

$$\hat{S}_{t+1,0} = \hat{S}_{t,k_{\text{in}}-1} + \gamma_{t,k_{\text{in}}} (\mathbf{S}_{t+1,0} - \hat{S}_{t,k_{\text{in}}-1})$$


$$\mathbf{S}_{t,k+1} = \underbrace{\frac{1}{\mathbf{b}} \sum_{j \in \mathcal{B}_{t,k+1}} \bar{S}(Z_j, \mathbb{T}(\hat{S}_{t,k}))}_{\approx \mathbb{E}_\pi[\bar{S}(Z, \mathbb{T}(\hat{S}_{t,k-1}))]} + \underbrace{\mathbf{S}_{t,k}}_{\approx \mathbb{E}_\pi[\bar{S}(Z, \mathbb{T}(\hat{S}_{t,k-1}))]} - \frac{1}{\mathbf{b}} \sum_{j \in \mathcal{B}_{t,k+1}} \bar{S}(Z_j, \mathbb{T}(\hat{S}_{t,k-1}))$$

- natural random field (red) and the control variate (blue): correlated through $\mathcal{B}_{t,k+1}$
- unfortunately, biased approximation $\mathbf{S}_{t,k+1} \rightarrow$ restart every k_{in} iterations in order to cancel the bias.


SPIDER EM for inference of Gaussian mixture (MNIST data set) F, Moulines, Wai (2020)




(top) Online EM, (bottom) SPIDER EM. Estimation of the weights, vs the nbr of epoch



Objective fct value vs the nbr of epoch



Quantile 0.25 of the distribution of $\|\mathbf{h}(\hat{S}_k)\|^2$



Quantile 0.75 of the distribution of $\|\mathbf{h}(\hat{S}_k)\|^2$

What kind of control of convergence ?

- From a Lyapunov inequality,

$$\mathbb{E} \left[(f + g)(\mathbf{T}(\hat{S}_{k+1})) \mid \mathcal{F}_k \right] \leq (f + g)(\mathbf{T}(\hat{S}_k)) - (D_{k+1})^2 + \mathbb{E} [R_{k+1} \mid \mathcal{F}_k]$$

- Almost-sure convergence when $k \rightarrow \infty$ (Robbins-Siegmund lemma)
- Explicit (upper) bounds

$$\sum_{k=0}^{k_{\max}-1} \mathbb{E} [D_{k+1}^2] \leq (f + g)(\mathbf{T}(\hat{S}_0)) - \min(f + g) + \sum_{k=0}^{k_{\max}-1} \mathbb{E} [R_{k+1}]$$

- Difficulties: it is **not** a gradient-SA algorithm

- Bounds for

- ϵ -stationarity (also designed to find the roots of \mathbf{h})
- non convex optimization

$$\mathbb{E} \left[\|\mathbf{h}(\hat{S}_\tau)\|^2 \right] \leq \bullet \quad \tau \text{ uniform on } \{1, \dots, k_{\max}\}$$

- scaling of the design parameters (k_{in} , k_{out} , \mathbf{b}) and the learning rate as a fct of N (data set size) and the accuracy level ϵ .

ϵ -stationarity for SPIDER-EM

- Theorem F., Moulines, Wai (2020): Explicit control of the mean error $\|\mathbf{h}(\hat{S}_t)\|^2$

Set $W(s) := (f + g)(\mathbf{T}(s))$. /* Lyapunov fct in the s -space; smooth ($L_{\dot{W}}$); $\nabla W(s) = -B(s)\mathbf{h}(s)$ */

Set $L^2 := N^{-1} \sum_{i=1}^N L_i^2$. /* Lipschitz constants related to $s \mapsto \bar{S}(Z_i, \mathbf{T}(s))$ */

Fix $k_{\text{out}}, k_{\text{in}}, \mathbf{b} \in \mathbb{N}_*$. Choose $\alpha \in (0, v_{\min}/\mu_*(k_{\text{in}}, \mathbf{b}))$ with

$$\mu_*(k_{\text{in}}, \mathbf{b}) := v_{\max} \frac{\sqrt{k_{\text{in}}}}{\sqrt{\mathbf{b}}} + \frac{L_{\dot{W}}}{2L}. \quad /*\text{spectrum of } B(s) \text{ in } [v_{\min}, v_{\max}] \text{ uniformly in } s */$$

Run the algorithm with $\xi_t = k_{\text{in}}$ and $\gamma_{t,k} := \alpha/L$. Then

$$\mathbb{E} \left[\|\mathbf{h}(\hat{S}_{\tau, \xi-1})\|^2 \right] \leq \left(\frac{1}{k_{\text{in}}} + \frac{\alpha^2}{\mathbf{b}} \right) \frac{1}{k_{\text{out}} \alpha \{v_{\min} - \alpha \mu_*(k_{\text{in}}, \mathbf{b})\}} \frac{2L}{\left(\mathbb{E} [W(\hat{S}_0)] - \min W \right)}$$

where (τ, ξ) is a uniform r.v. on $\{1, \dots, k_{\text{out}}\} \times \{0, \dots, k_{\text{in}} - 1\}$ indep of $\{\hat{S}_{t,k}\}$.

Complexity analysis

- How to choose k_{in} , k_{out} , \mathbf{b} in order to reach ϵ -stationarity, with an optimal computational cost ?
 - ϵ -stationarity for non-convex optimization

$$\mathbb{E} \left[\|\mathbf{h}(\hat{S}_\tau)\|^2 \right] \leq \epsilon \quad \text{random stopping time } \tau$$

- Computational cost: (i) Nbr of processed example, (ii) Nbr of optimization steps.

- Result:

$$k_{\text{in}} = \mathbf{b} = O(\sqrt{n}), \quad k_{\text{out}} = O\left(\frac{1}{(\epsilon k_{\text{in}})}\right)$$

- Nbr of optimization steps: $O(1/\epsilon)$





- Nbr of processed examples: $O(\sqrt{n} \epsilon^{-1})$

Algorithm	Complexity (proc. ex.)
Online-EM	ϵ^{-2}
i-EM	$n \epsilon^{-1}$
sEM-vr	$n^{2/3} \epsilon^{-1}$
FIEM	$n^{2/3} \epsilon^{-1} \wedge \sqrt{n} \epsilon^{-3/2}$
SPIDER-EM	$\sqrt{n} \epsilon^{-1}$

SPIDER-EM – state of the art !

Federated learning through MM

Federated Learning



- The central server coordinates the participation of the local devices/clients/workers
- Local training data sets, **never** uploaded to the server
- FL reduces privacy and security risks

- Global model maintained by the central server: sent to the devices
- Each worker updates an information on the global model
- Only this update is communicated to the central server; aggregation by the central server

● Methodological developments under the 'constraints'

- Local data sets, heterogeneous, unbalanced
- Partial participation of the clients (charged devices, plugged-in, free wi-fi connection, ...)
- Massively distributed: large nbr devices w.r.t. the size of the local data sets
- Communication cost \gg Computational cost: compression

Federated MM: the data

- n local workers with their own data set (\rightarrow distribution π_i)
- Local worker $\#i$: has a local objective function

$$\theta \mapsto f_i(\theta) := \mathbb{E}_{\pi_i} [\ell(Z, \theta)]$$

and the central server targets

$$\theta \mapsto f(\theta) := \sum_{i=1}^n \mu_i \mathbb{E}_{\pi_i} [\ell(Z, \theta)], \quad \mu_i \text{ weights, unbalanced local data sets } (= 1/n \text{ for online learning})$$

- The central server runs a stochastic surrogate MM, by aggregation of local oracles provided by the local agents

$$\mathbf{S}_{k+1,i} \approx \mathbb{E}_{\pi_i} [\bar{S}(Z, \mathbf{T}(\hat{S}_k))] \longrightarrow \mathbf{S}_{k+1} \approx \mathbb{E}_{\pi} [\bar{S}(Z, \mathbf{T}(\hat{S}_k))]$$

and then, perform the optimization step: $\mathbf{T}(\hat{S}_{k+1})$ and broadcast it to the local agents.

Federated MM: partial participation and quantization

- At each iteration, a local worker $\#i$ is active with probability $p \in (0, 1)$.
- If active at iteration $\#(k + 1)$, the local worker $\#i$
 - learns an oracle $S_{k+1,i}$ of the parameter $\mathbb{E}_{\pi_i} [\bar{S}(Z, \mathbb{T}(\hat{S}_k))]$ characterizing the local majorizing function
 - quantize the information before sending it to the central server: $\text{Quant}(\cdot)$ see e.g. Alistarh et al. (2018), Horvath et al. (2019)
 - the quantized information is **not the oracle**, but the difference with the previously sent information $V_{k,i}$

$$\Delta_{k+1,i} := S_{k+1,i} - \hat{S}_k - V_{k,i}$$

- it updates this control variate $V_{\cdot,i}$'s adapted from Mishchenko et al., 2019

$$V_{k+1,i} = V_{k,i} + \alpha \text{Quant}(\Delta_{k+1,i})$$

Algorithm 3: Federated MM with partial participation (PP)

Data: $k_{\max} \in \mathbb{N}^*$; for $i \in [n]^*$, $V_{0,i} \in \mathbb{R}^q$; $\hat{S}_0 \in \mathbb{R}^q$; a positive sequence $\{\gamma_{k+1}, k \in [k_{\max} - 1]\}$; $\alpha > 0$ and $p \in (0, 1)$.

Result: The sequence: $\{\hat{S}_k, k \in [k_{\max}]\}$ and its mirror $\mathbb{T}(\hat{S}_k)$

Set $V_0 = n^{-1} \sum_{i=1}^n V_{0,i}$

for $k = 0, \dots, k_{\max} - 1$ **do**

 Sample \mathcal{A}_{k+1} , the set of active workers /* each worker active with probability p , independently */

for $i \in \mathcal{A}_{k+1}$ **do**

 Sample $\mathbf{S}_{k+1,i}$, an approximation of $\mathbb{E}_{\pi_i} [\bar{S}(Z, \mathbb{T}(\hat{S}_k))]$

 Set $\Delta_{k+1,i} = \mathbf{S}_{k+1,i} - \hat{S}_k - V_{k,i}$

 Set $V_{k+1,i} = V_{k,i} + \alpha \text{Quant}(\Delta_{k+1,i})$. /* $0 < \alpha \leq 1/(1 + \omega)$. V_k : control variate */

 Send $\text{Quant}(\Delta_{k+1,i})$ to the central server /* random, unbiased, $\mathbb{E}[\|\text{Quant}(x) - x\|^2] \leq \omega \|x\|^2$ */

for $i \notin \mathcal{A}_{k+1}$ **do**

 Set $V_{k+1,i} = V_{k,i}$ (no update)

 Set $H_{k+1} = V_k + p^{-1} \sum_{i \in \mathcal{A}_{k+1}} \mu_i \text{Quant}(\Delta_{k+1,i})$ /* compensate for (i) the V_k 's and (ii) the PP */


 Set $\hat{S}_{k+1} = \hat{S}_k + \gamma_{k+1} H_{k+1}$

 Set $V_{k+1} = V_k + \alpha \sum_{i \in \mathcal{A}_{k+1}} \mu_i \text{Quant}(\Delta_{k+1,i})$ /* learn the aggregated control variates */

 Send \hat{S}_{k+1} and $\mathbb{T}(\hat{S}_{k+1})$ to the n workers

Toy example: inference of a \mathbb{R}^2 -valued Gaussian mixture model with 2 components (1/2)

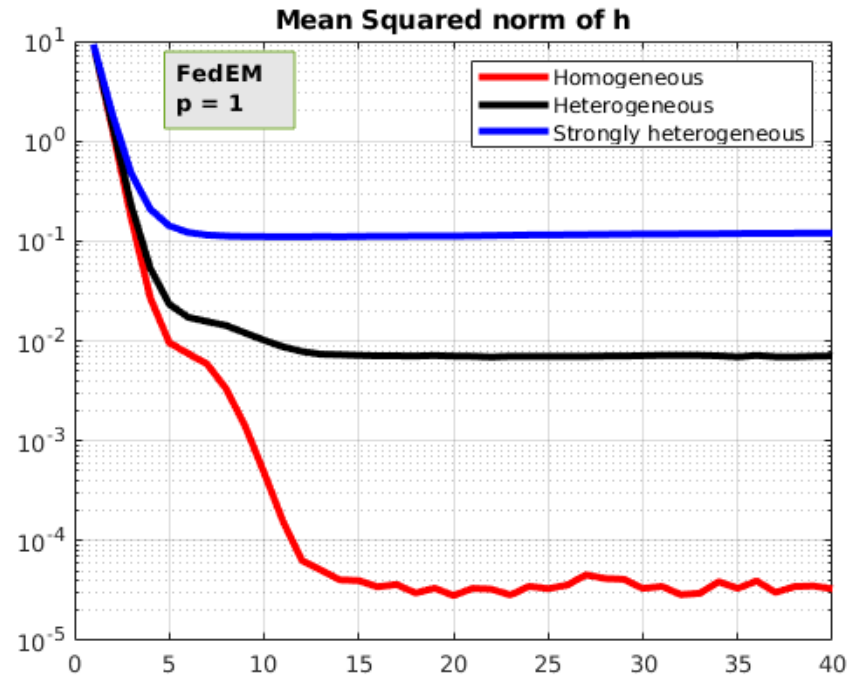
- Robustness to partial participation



$k \mapsto \mathbb{E} [\|h(\hat{S}_k)\|^2]$ vs the nbr of epochs.

Estimated by Monte Carlo

- Robustness to heterogeneity




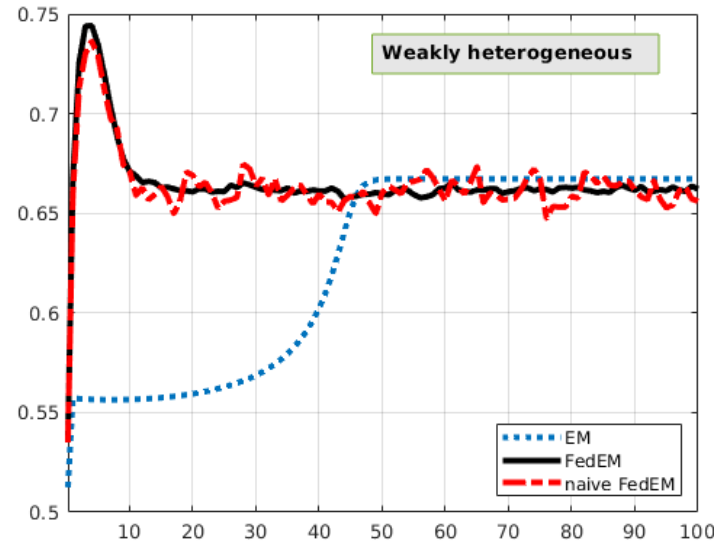
$k \mapsto \mathbb{E} [\|h(\hat{S}_k)\|^2]$ vs the nbr of epochs.

Estimated by Monte Carlo

Toy example: inference of a \mathbb{R}^2 -valued Gaussian mixture model with 2 components (2/2)

- Federated MM vs naive-Federated MM ?

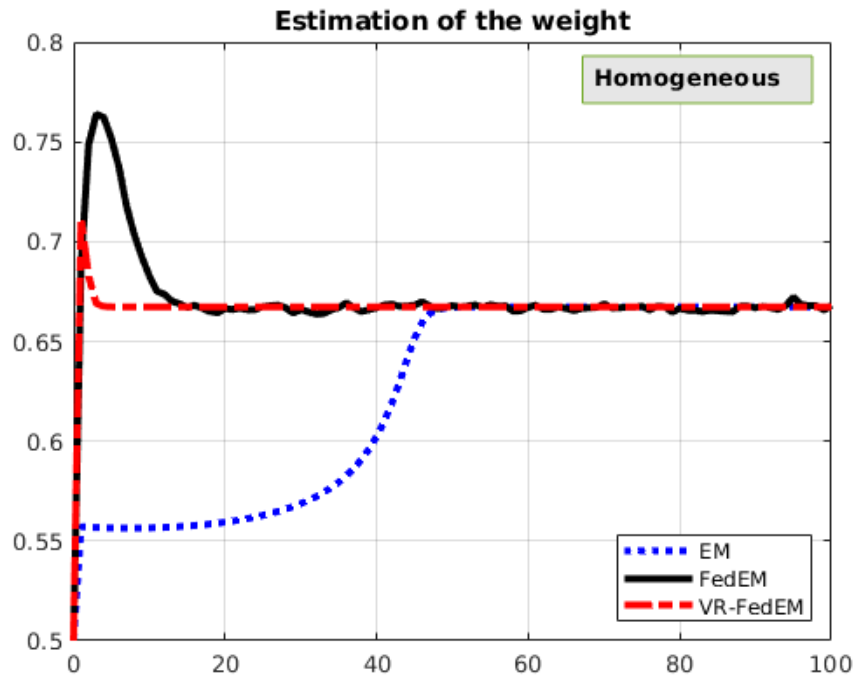
In naive-Federated MM:
remove the variables $V.c$'s
– i.e. the control variates
introduced to control the
variance of the quantiza-
tion step.



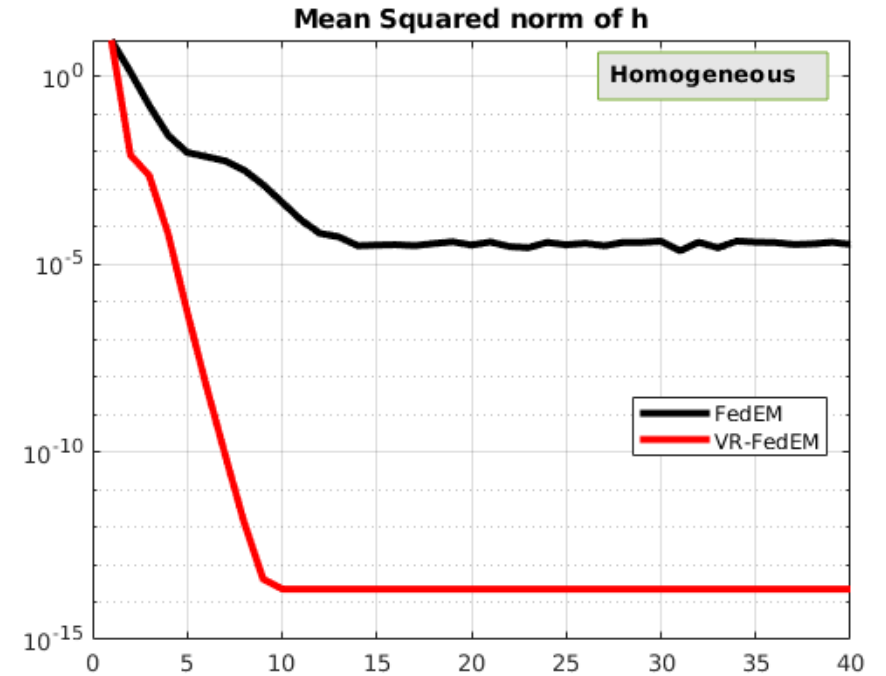
Estimation of the weight vs the nbr epoch; Case "homogeneous" and case "strongly heterogeneous"

Variance Reduction in Federated MM

- Mix the SPIDER variance reduction idea and the federated learning algorithm
- Case of EM (toy example, to follow)



Estimation of the weight vs the nbr epoch



$k \mapsto \mathbb{E} [\|h(\hat{S}_k)\|^2]$ vs the nbr of epochs.

Estimated by Monte Carlo

Bounds for convergence

- Theorem (adapted from) Dieuleveut, F., Moulines, Robin (2021): Explicit control of the mean error $\|\mathbf{h}(\hat{S}_\cdot)\|^2$

Run the algorithm with $\alpha := (1 + \omega)^{-1}$ /* $\mathbb{E}[\|\text{Quant}(x) - x\|^2] \leq \omega \|x\|^2$ */

and $\gamma_k = \gamma \in (0, \gamma_{\max}]$, where

$$\gamma_{\max} := \frac{v_{\min}}{2L_{\dot{W}}} \wedge \frac{p\sqrt{n}}{2\sqrt{2}L(1+\omega)\sqrt{\omega + (1-p)(1+\omega)/p}}.$$

Set $\sigma^2 := n \sum_{i=1}^n \mu_i^2 \sigma_i^2$. /* σ_i^2 : variance of the oracles at agent # i */

Denote by τ the uniform random variable on $\{0, \dots, k_{\max} - 1\}$. Then,

$$\begin{aligned} v_{\min} \left(1 - \gamma \frac{L_{\dot{W}}}{v_{\min}}\right) \mathbb{E} \left[\|\mathbf{h}(\hat{S}_\tau)\|^2 \right] &\leq \frac{(W(\hat{S}_0) - \min W)}{\gamma k_{\max}} \\ &+ \frac{\gamma}{k_{\max}} 2L_{\dot{W}} \frac{\omega}{\alpha} \sum_{i=1}^n \mu_i^2 \|V_{0,i} - \mathbf{h}_i(\hat{S}_0)\|^2 \\ &+ \gamma L_{\dot{W}} \frac{1 + 5(\omega + (1-p)(1+\omega)/p)}{n} \sigma^2. \end{aligned}$$

Complexity analysis in the case $p = 1$

Given an accuracy level ϵ , how to choose the design parameters in order to minimize the number of optimization ?

- The number of optimization is k_{\max} chosen in order to reach the accuracy level ϵ :

$$\mathcal{K}_{\text{opt}}(\epsilon) = O\left(\frac{1}{\epsilon^2} \frac{(1+\omega)\sigma^2}{n}\right) \vee O\left(\frac{1}{\epsilon \gamma_{\max}}\right)$$

1st term is leader iff $\epsilon \ll \gamma_{\max}(1+\omega)\sigma^2/n$ (high noise regime)

- Compression effect:**

γ impacted by compression iff $n \ll \omega^3$.

On \mathcal{K}_{opt} : see table

	Complexity regime:	$\frac{(1+\omega)\sigma^2}{n\epsilon^2}$	$\frac{1}{\gamma_{\max}\epsilon}$
γ_{\max} regime:	E.g. case when	High noise σ^2 , small ϵ	Low σ^2 larger ϵ
$\frac{v_{\min}}{2L\dot{W}} \frac{1}{\sqrt{n}}$	large ratio n/ω^3	$\times \omega$	$\times 1$
$\frac{1}{2\sqrt{2}L(1+\omega)\sqrt{\omega}}$	low ratio n/ω^3	$\times \omega$	$\times \omega^{3/2}/\sqrt{n}$

Conclusion

Contributions

- A unifying point of view: surrogate MM covering popular algorithms.
- "surrogate in the foreground": efficient approach to (i) summarize the data through "sufficient statistics", (ii) in FL: interpretation of the quantity aggregated at the central server.
- Novel algorithms for large scale learning (large batch, online); and federated learning.
- Explicit bound of convergence, complexity analysis for the Stochastic Surrogate MM, possibly with variance reduction, and Federated Surrogate MM.
- Non convex optimization. Algorithms and some theoretical results, for the non smooth optimization case.

Open theoretical analyses

- The map T : when no closed form.
- The constraints $s \in \mathcal{S} \rightarrow$ Stochastic Approximation restricted to a structured subset of \mathbb{R}^d .
- Other strategies for variance reduction in Federated Learning

Bibliography

Talk based on the publications and preprint:

G. Fort, E. Moulines, and H.-T. Wai.

A Stochastic Path Integral Differential Estimator Expectation Maximization Algorithm.

In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16972–16982. Curran Associates, Inc., 2020.

A. Dieuleveut, G. Fort, E. Moulines and G. Robin.

Federated-EM with heterogeneity mitigation and variance reduction.

In A. Beygelzimer and Y. Dauphin and P. Liang and J. Wortman Vaughan editors, *Advances in Neural Information Processing Systems*, 2021.

A. Dieuleveut, G. Fort and H.-T. Wai.

Federated Majorize-Minimization for large scale learning
Submitted.

H.D. Nguyen, F. Forbes, G. Fort and O. Cappé

An online Minorization-Maximization algorithm.

Accepted for publication in IFCS 2022 proceedings.

Other references

D. Alistarh, T. Hoefler, M. Johansson, N. Konstantinov, S. Khirirat, and C. Renggli.

The convergence of sparsified gradient methods.

In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 5973–5983. Curran Associates, Inc., 2018.

A. Benveniste, M. Métivier, and P. Priouret.

Adaptive Algorithms and Stochastic Approximations.

Springer Verlag, 1990.

Cappé, O. and Moulines, E.

On-line Expectation Maximization algorithm for latent data models.

J Roy Stat Soc B Met 71(3):593–613, 2009.

Chen, J. and Zhu, J. and Teh, Y. and Zhang, T.

Stochastic Expectation Maximization with Variance Reduction.

In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R (eds) *Advances in Neural Information Processing Systems* 31, Curran Associates, Inc., pp 7967–7977; 2018.

Defazio, A. and Bach, F. and Lacoste-Julien, S.

SAGA: A Fast Incremental Gradient Method With Support for Non-Strongly Convex Composite Objectives.

In: Ghahra-mani Z, Welling M, Cortes C, Lawrence ND, Weinberger KQ (eds) *Advances in Neural Information Processing Systems* 27, Curran Associates, Inc., pp 1646–1654; 2014.

A. Dempster, N. Laird, and D. Rubin.
Maximum Likelihood from Incomplete Data via the EM Algorithm.
J. Roy. Stat. Soc. B Met., 39(1):1–38, 1977.

Fang, C. and Li, C. and Lin, Z. and Zhang, T.
SPIDER: Near-Optimal Non-Convex Optimization via Stochastic Path-Integrated Differential Estimator.
In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R (eds) Advances in Neural Information Processing Systems 31, Curran Associates, Inc., pp 689–69; 2018.

Fort, G. and Gach, P. and Moulines, E.
The Fast Incremental Expectation Maximization for finite-sum optimization: asymptotic convergence.
Statistics and Computing, 31, 2021.

Horváth, S. and Kovalev, D. and Mishchenko, D. and Stich, S. and Richtárik, P.
Stochastic distributed learning with gradient quantization and variance reduction.
arXiv preprint arXiv:1904.05115, 2019.

Johnson, R. and Zhang, T.
Accelerating Stochastic Gradient Descent using Predictive Variance Reduction.
In: Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ (eds) Advances in Neural Information Processing Systems 26, Curran Associates, Inc., pp 315–323; 2013.

Karimi, B. and Wai, H.T. and Moulines, E. and Lavielle, M.
On the Global Convergence of (Fast) Incremental Expectation Maximization Methods.
In: Wallach H, Larochelle H, Beygelzimer A, d'Alché Buc F, Fox E, Garnett R (eds) Advances in Neural Information Processing Systems 32, Curran Associates, Inc., pp 2837–2847; 2019.

K. Lange.
MM Optimization Algorithms.
SIAM-Society for Industrial and Applied Mathematics, 2016.

Mishchenko, K. and Gorbunov, E. and Takáč, M. and Richtárik, P.
Distributed Learning with Compressed Gradient Differences
arXiv, 2019.