

# On the Global Convergence of (Fast) Incremental EM Methods

Belhal Karimi<sup>1,2</sup>, Hoi-To Wai<sup>3</sup>, Eric Moulines<sup>2</sup> and Marc Lavielle<sup>1</sup>  
INRIA<sup>1</sup>, École Polytechnique<sup>2</sup>, Chinese University of Hong Kong<sup>3</sup>



ÉCOLE  
POLYTECHNIQUE  
UNIVERSITÉ PARIS-SACLAY

Inria  
INVENTEURS DU MONDE NUMÉRIQUE



## Maximum Likelihood Estimation (MLE)

- Given a set of  $n$  observations  $y = (y_i, i \in [n])$ .
- Goal: fitting the parametric model  $g(\cdot, \theta)$ .
- Maximum Likelihood Estimation of  $\theta$**   

$$\theta^* = \arg \max_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n \log g(y_i, \theta)$$
- $g(y, \theta) = \int_{\mathcal{Z}} f(z, y, \theta) \mu(dz)$  is a parametric model with a latent variable  $z$  – the function is generally intractable.
- We use the **EM** algorithm which takes advantage of the latent structure.

## Settings and Notation

- Regularized Empirical Risk Minimization:*

$$\min_{\theta \in \Theta} \bar{\mathcal{L}}(\theta) := R(\theta) + \mathcal{L}(\theta)$$

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\theta) := \frac{1}{n} \sum_{i=1}^n \{ -\log g(y_i; \theta) \}$$

- $\bar{\mathcal{L}}$  is possibly **nonconvex** and lower bounded
- For all  $i \in [n]$ ,  $f(z_i, y_i, \theta)$  and  $p(z_i|y_i, \theta)$  denote the complete likelihood and the posterior distribution
- Focus on the **Exponential Family Distribution**:  

$$f(z_i, y_i, \theta) = h(z_i, y_i) \exp(\langle S(z_i, y_i) | \phi(\theta) \rangle - \psi(\theta))$$

## EM for Exponential Family

- We define the following EM-related operations.

★ **E-operation:** for any  $\theta \in \Theta$ ,

$$\bar{s}(\theta) := \frac{1}{n} \sum_{i=1}^n \bar{s}_i(\theta)$$

also, define  $\bar{s}_i(\theta) := \int_{\mathcal{Z}} S(z_i, y_i) p(z_i|y_i; \theta) \mu(dz_i)$ .

★ **M-operation:** for any  $\hat{s}$ ,

$$\hat{\theta} := \bar{\theta}(\hat{s}) := \arg \min_{\theta \in \Theta} \{ R(\theta) + \psi(\theta) - \langle \hat{s} | \phi(\theta) \rangle \}$$

**Batch EM Algorithm:** given  $\hat{\theta}^{(0)}$ , set  $k = 0$ ,

- (E-step)  $\hat{s}^{(k+1)} = \bar{s}(\hat{\theta}^{(k)})$
- (M-step)  $\hat{\theta}^{(k+1)} = \bar{\theta}(\hat{s}^{(k+1)})$

**For large  $n$ , the E-step is computationally expensive!**

## General Formulation of Stochastic EM (sEM)

**Idea:** We replace the **E-step** with a **stochastic/incremental E-step (sE-step)** that looks at 1 sample only.

★ **sE-step** (general form): w/ const. stepsize  $\gamma$ ,

$$\hat{s}^{(k+1)} = \hat{s}^{(k)} - \gamma(\hat{s}^{(k)} - \mathcal{S}^{(k+1)})$$

- iEM:**  $\mathcal{S}^{(k+1)} = \mathcal{S}^{(k)} + \frac{1}{n}(\bar{s}_{i_k}^{(k)} - \bar{s}_{i_k}^{(\tau_{i_k}^k)})$
- sEM-VR:**  $\mathcal{S}^{(k+1)} = \bar{s}^{(\ell(k))} + (\bar{s}_{i_k}^{(k)} - \bar{s}_{i_k}^{(\ell(k))})$
- fiEM:**  $\mathcal{S}^{(k+1)} = \bar{s}^{(k)} + (\bar{s}_{i_k}^{(k)} - \bar{s}_{i_k}^{(t_k^k)})$  and  $\bar{s}^{(k+1)} = \bar{s}^{(k)} + n^{-1}(\bar{s}_{j_k}^{(k)} - \bar{s}_{j_k}^{(t_k^k)})$

- Set the termination number  $K \sim \mathcal{U}\{1, 2, \dots, K_{\max}\}$
- For  $k > 0$ :
  - Draw index  $i_k \in [n]$  uniformly (and  $j_k \in [n]$  for fiEM).
  - Compute the surrogate sufficient statistics  $\mathcal{S}^{(k+1)}$ .
  - Compute  $\hat{s}^{(k+1)}$  via the sE-step (**see the left**).
  - Compute  $\hat{\theta}^{(k+1)}$  via the M-step (**same as batch EM**):  

$$\hat{\theta}^{(k+1)} = \bar{\theta}(\hat{s}^{(k+1)})$$

• **Return:**  $\hat{\theta}^{(K)}$ .

- Prior works:** iEM–Neal & Hinton, 1998; sEM-VR–Chen et al., 2018 (local convergence); fiEM–inspired by SAGA.

**How to analyze their global convergence? And which algorithm is faster?**

## iEM as an incremental MM Scheme

- iEM can be interpreted as **incremental MM** (Mairal, 2015) with the **upper bound surrogate function**:

$$Q_i(\theta; \theta') := - \int_{\mathcal{Z}} \{ \log f(z_i, y_i; \theta) - \log p(z_i|y_i; \theta') \} p(z_i|y_i; \theta') \mu(dz_i), \text{ where } Q_i(\theta; \theta') \geq \mathcal{L}_i(\theta), \forall \theta.$$

- Incremental MM:** at every iteration  $k$ , we obtain  $\hat{\theta}^{(k+1)} = \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n Q_i(\theta; \hat{\theta}^{(k)})$ .
- Convergence Analysis:** with exponential family model,  $Q_i(\theta; \theta') - \mathcal{L}_i(\theta)$  is  **$L_e$ -smooth for all  $i$** :

**Theorem (iEM)** For any  $K_{\max} \geq 1$ ,  $K \sim \mathcal{U}([0, K_{\max} - 1])$  independent of the  $\{i_k\}_{k=0}^{K_{\max}}$ , we have the **global rate**:

$$\mathbb{E}[\|\nabla \bar{\mathcal{L}}(\hat{\theta}^{(K)})\|^2] \leq n \frac{2L_e}{K_{\max}} \mathbb{E}[\bar{\mathcal{L}}(\hat{\theta}^{(0)}) - \bar{\mathcal{L}}(\hat{\theta}^{(K_{\max})})], \quad (1)$$

## sEM-VR/fiEM are Scaled Gradient Methods and Faster than iEM

- Unlike iEM, the sEM-VR and fiEM methods can be analyzed as **scaled gradients** methods. Consider:

$$\min_{s \in \mathcal{S}} V(s) := \bar{\mathcal{L}}(\bar{\theta}(s)) = R(\bar{\theta}(s)) + \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\bar{\theta}(s)).$$

- Variance-reduced scaled gradient:** the sE-step update  $\hat{s}^{(k)}$  by  $\hat{s}^{(k)} - \mathcal{S}^{(k+1)}$ , we can show

$$\langle \mathbb{E}[\hat{s}^{(k)} - \mathcal{S}^{(k+1)}] | \nabla V(\hat{s}^{(k)}) \rangle \geq v_1 \|\mathbb{E}[\hat{s}^{(k)} - \mathcal{S}^{(k+1)}]\|^2 \geq v_2 \|\nabla V(\hat{s}^{(k)})\|^2 \text{ for some } v_1, v_2 > 0.$$

$\therefore$  the sEM-VR/fiEM methods are **variance-reduced, scaled gradient** updates of the sufficient statistics.

- Convergence Analysis:** with exponential family model, **the function  $V(s)$  is  $\bar{L}_v$ -smooth**,

**Theorem (sEM-VR)**  $\gamma = \frac{\mu v_{\min}}{\alpha L_v n^{2/3}}$  & epoch  $m = \frac{n}{2\mu^2 v_{\min}^2 + \mu}$ :


$$\mathbb{E}[\|\nabla V(\hat{s}^{(K)})\|^2] \leq n^{\frac{2}{3}} \frac{2\bar{L}_v}{\mu K_{\max}} \frac{v_{\max}^2}{v_{\min}^2} \mathbb{E}[V(\hat{s}^{(0)}) - V(\hat{s}^{(K_{\max})})].$$

**Theorem (fiEM)**  $\gamma = \frac{v_{\min}}{\alpha L_v n^{2/3}}$  &  $\alpha = \max\{6, 1 + 4v_{\min}\}$ :


$$\mathbb{E}[\|\nabla V(\hat{s}^{(K)})\|^2] \leq n^{\frac{2}{3}} \frac{\alpha^2 \bar{L}_v v_{\max}^2}{K_{\max} v_{\min}^2} \mathbb{E}[V(\hat{s}^{(0)}) - V(\hat{s}^{(K_{\max})})].$$

## Fitting a Gaussian Mixture Model

- Goal:** fitting a GMM with a penalization.
- Faster rate for sEM-VR and fiEM ( $n = 10^5$ )



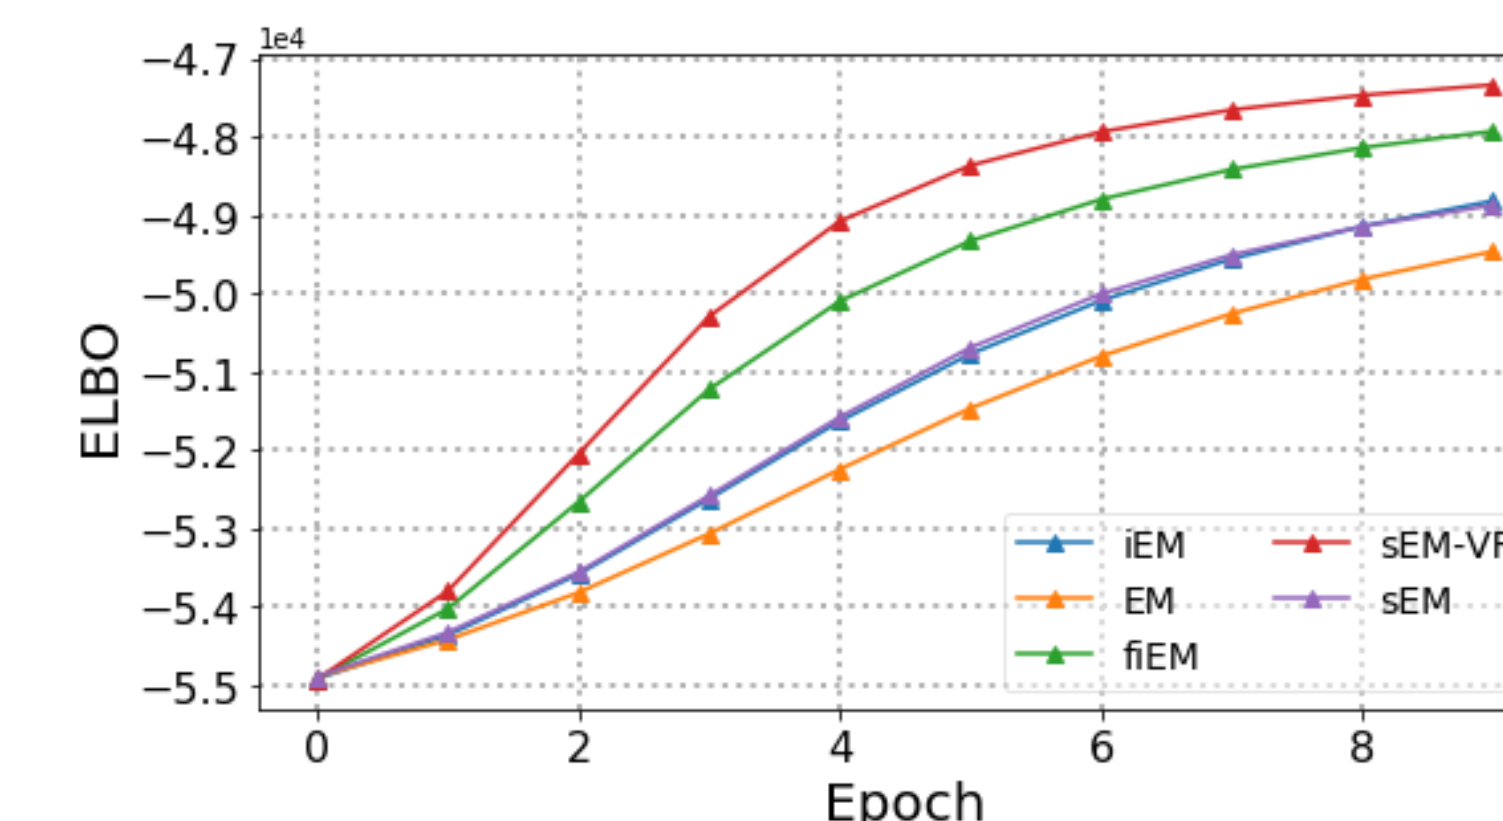
- Iteration number to reach  $\epsilon$ -accuracy vs.  $n$ .



- Reveal the linear  $[O(n/\epsilon)]$  and sublinear  $[O(n^{2/3}/\epsilon)]$  rates

## Probabilistic Latent Semantic Analysis (PLSA)

- Goal:** Classifying  $D$  docs into  $K$  topics
- FAO (UN Food and Agriculture Organization) datasets.



## References

- Jianfei Chen, Jun Zhu, Yee Whye Teh, and Tong Zhang. Stochastic expectation maximization with variance reduction. In *NeurIPS*, 2018.
- Julien Mairal. Incremental majorization-minimization optimization with application to large-scale machine learning. *SIAM Journal on Optimization*, 2015.
- Radford M Neal and Geoffrey E Hinton. A view of the EM algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*, pages 355–368. Springer, 1998.