



HAL
open science

Online but Accurate Inference for Latent Variable Models with Local Gibbs Sampling

Christophe Dupuy, Francis Bach

► **To cite this version:**

Christophe Dupuy, Francis Bach. Online but Accurate Inference for Latent Variable Models with Local Gibbs Sampling. 2016. hal-01284900v2

HAL Id: hal-01284900

<https://inria.hal.science/hal-01284900v2>

Preprint submitted on 14 Mar 2016 (v2), last revised 30 Jan 2018 (v4)

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.

Online but Accurate Inference for Latent Variable Models with Local Gibbs Sampling

Christophe Dupuy^{1,2} and Francis Bach^{1,3}

¹INRIA - Sierra project - team

²Technicolor R&D France

³Département Informatique de l'École Normale Supérieure Paris

Abstract

We study parameter inference in large-scale latent variable models. We first propose an unified treatment of online inference for latent variable models from a non-canonical exponential family, and draw explicit links between several previously proposed frequentist or Bayesian methods. We then propose a novel inference method for the frequentist estimation of parameters, that adapts MCMC methods to online inference of latent variable models with the proper use of local Gibbs sampling. Then, for latent Dirichlet allocation, we provide an extensive set of experiments and comparisons with existing work, where our new approach outperforms all previously proposed methods. In particular, using Gibbs sampling for latent variable inference is superior to variational inference in terms of test log-likelihoods. Moreover, Bayesian inference through variational methods perform poorly, sometimes leading to worse fits with latent variables of higher dimensionality.

1 Introduction

Probabilistic graphical models provide general modelling tools for complex data, where it is natural to include assumptions on the data generating process by adding latent variables in the model. Such latent variable models are adapted to a wide variety of unsupervised learning tasks (Koller & Friedman, 2009; Murphy, 2012). In this paper, we focus on parameter inference in such latent variable models where the main operation needed for the standard expectation-maximization (EM) algorithm is intractable, namely dealing with conditional distributions over latent variables given the observed variables; latent Dirichlet allocation (LDA) (Blei et al., 2003) is our motivating example, but many hierarchical models exhibit this behavior. For such models, there exist two main classes of methods to deal efficiently with intractable exact inference in large-scale situations: sampling methods or variational methods.

Sampling methods can handle arbitrary distributions and lead to simple inference algorithms while converging to exact inference. However it may be slow to converge and non scalable to big datasets in practice. In particular, although efficient implementations have been developed, for example for LDA (Zhao et al., 2014; Yan et al., 2009), MCMC methods may not deal efficiently yet with continuous streams of data for our general class of models.

On the other hand, *variational inference* builds an approximate model for the posterior distribution over latent variables—called variational—and infer parameters of the true model through this approximation. The fitting of this variational distribution is formulated as an optimization problem where efficient (deterministic) iterative techniques such as gradient or coordinate ascent methods apply. This approach leads to scalable inference schemes (Hoffman et al., 2013), but due to approximations, there always remains a gap between the variational posterior and the true posterior distribution, inherent to algorithm design, and that will not vanish when the number of samples and the number of iterations increase.

Beyond the choice of approximate inference techniques for latent variables, parameter inference may be treated either from the *frequentist* point of view, e.g., using maximum likelihood inference, or a *Bayesian* point of view, where the posterior distribution of the parameter given the observed data is approximated. With massive numbers of observations, this posterior distribution is typically peaked around the maximum likelihood estimate, and the two inference frameworks should not differ much (Van der Vaart, 2000).

In this paper, we focus on methods that make a single pass over the data to estimate parameters. We make the following contributions:

1. We review and compare existing methods for online inference for latent variable models from a non-canonical exponential family in Section 2, and draw explicit links between several previously proposed frequentist or Bayesian methods.
2. We propose in Section 3 a novel inference method for the frequentist estimation of parameters, that adapts MCMC methods to online inference of latent variable models with the proper use of “local” Gibbs sampling. In our online scheme, we apply Gibbs sampling to the current observation, which is “local”, as opposed to “global” batch schemes where Gibbs sampling is applied to the entire dataset.
3. After formulating LDA as a non-canonical exponential family in Section 4, we provide an extensive set of experiments in Section 5, where our new approach outperforms all previously proposed methods. In particular, using Gibbs sampling for latent variable inference is superior to variational inference in terms of test log-likelihoods. Moreover, Bayesian inference through variational methods perform poorly, sometimes leading to worse fits with latent variables of higher dimensionality.

2 Online EM

We consider an *exponential family* model on random variables (X, h) with parameter $\eta \in \mathcal{E} \subseteq \mathbb{R}^d$ and with density (Lehmann & Casella, 1998):

$$p(X, h|\eta) = a(X, h) \exp [\langle \phi(\eta), S(X, h) \rangle - \psi(\eta)]. \quad (1)$$

We assume that h is hidden and X is observed. The vector $\phi(\eta) \in \mathbb{R}^d$ represents the natural parameter, $S(X, h) \in \mathbb{R}^d$ is the vector of sufficient statistics, $\psi(\eta)$ is the log-normalizer, and $a(X, h)$ is the underlying base measure. We consider a *non-canonical* family as in many models (such as LDA), the natural parameter $\phi(\eta)$ does not coincide with the model parameter η , that is, $\phi(\eta) \neq \eta$; we however assume that ϕ is injective.

We consider N *i.i.d.* observations $(X_i)_{i=1, \dots, N}$ from a distribution $t(X)$, which may be of the form $P(X|\eta^*) = \int_h p(X, h|\eta^*)$ for our model above and a certain $\eta^* \in \mathcal{E}$ (well-specified model) or not (misspecified model). Our goal is to obtain a predictive density $r(X)$ built from the data and using the model defined in (1), with the maximal expected log-likelihood $\mathbb{E}_{t(X)} \log r(X)$.

2.1 Maximum likelihood estimation

In the frequentist perspective, the predictive distribution $r(X)$ is of the form $p(X|\hat{\eta})$, for a well-defined estimator $\hat{\eta} \in \mathcal{E}$. The most common method is the EM algorithm (Dempster et al., 1977), which is an algorithm that aims at maximizing the likelihood of the observed data, that is,

$$\max_{\eta \in \mathcal{E}} \sum_{i=1}^N \log p(X_i|\eta). \quad (2)$$

More precisely, the EM algorithm is an iterative process to find the maximum likelihood (ML) estimate given observations $(X_i)_{i=1, \dots, N}$ associated to hidden variables $(h_i)_{i=1, \dots, N}$. The EM algorithm may be seen as the iterative construction of lower bounds of the log-likelihood function (Bishop, 2006). In the exponential family setting (1), we have, by Jensen’s inequality, given the model defined by $\eta' \in \mathcal{E}$ from the previous

iteration, and for any parameter $\eta \in \mathcal{E}$:

$$\begin{aligned}
\log p(X_i|\eta) &= \log \int p(X_i, h_i|\eta) dh_i \\
&\geq \int p(h_i|X_i, \eta') \log \frac{p(X_i, h_i|\eta)}{p(h_i|X_i, \eta')} dh_i \\
&= \int p(h_i|X_i, \eta') (\langle \phi(\eta), S(X_i, h_i) \rangle - \psi(\eta)) dh_i - C_i(\eta') \\
&= \langle \phi(\eta), \mathbb{E}_{p(h_i|X_i, \eta')} [S(X_i, h_i)] \rangle - \psi(\eta) - C_i(\eta'),
\end{aligned}$$

for a certain constant $C_i(\eta')$, with equality if $\eta' = \eta$. Thus, EM-type algorithms build locally tight lower bounds of the log-likelihood in (2), which are equal to

$$\langle \phi(\eta), \sum_{i=1}^N s_i \rangle - N\psi(\eta) + \text{cst},$$

for appropriate values of $s_i \in \mathbb{R}^d$ obtained by computing conditional expectations of h_i given X_i for the current model defined by η' (E-step), i.e., $s_i = \mathbb{E}_{p(h_i|X_i, \eta')} [S(X_i, h_i)]$. Then this function of η is maximized to obtain the next iterate (M-step). In standard EM applications, these two steps are assumed tractable. In Section 3, we will only assume that the M-step is tractable while the E-step is intractable.

Standard EM will consider $s_i = \mathbb{E}_{p(h_i|X_i, \eta')} [S(X_i, h_i)]$ for the previous value of the parameter η for all i , and hence, at every iteration, all observations $X_i, i = 1, \dots, N$ are considered for latent variable inference, leading to a slow “batch” algorithm for large N .

Incremental EM (Neal & Hinton, 1998) will only update a single element s_i coming from a single observation X_i and update the corresponding part of the sum $\sum_{j=1}^N s_j$ without changing other elements. In the extreme case where a single pass over the data is made, then the M-step at iteration i maximizes

$$\langle \phi(\eta), \sum_{j=1}^i \mathbb{E}_{p(h_j|X_j, \eta_{j-1})} [S(X_j, h_j)] \rangle - i\psi(\eta),$$

with respect to η . In the next section, we provide a (known) other interpretation of this algorithm.

2.2 Stochastic approximation

Given our frequentist objective $\mathbb{E}_{t(X)} \log p(X|\eta)$ to maximize defined as an expectation, we may consider two forms of stochastic approximation (Kushner & Yin, 2003), where observations X_i sampled from $t(X)$ are processed only once. The first one is stochastic gradient ascent, of the form

$$\eta_i = \eta_i + \gamma_i \frac{\partial \log p(X_i|\eta)}{\partial \eta},$$

or appropriately renormalized version thereof, i.e., $\eta_i = \eta_i + \gamma_i H^{-1} \frac{\partial \log p(X_i|\eta)}{\partial \eta}$, with several possibilities for the $d \times d$ matrix H , such as the negative Hessian of the partial or the full log-likelihood, or the negative covariance matrix of gradients, which can be seen as versions of natural gradient—see Titterton (1984); Cappé & Moulines (2009). This either leads to slow convergence (without H) or expensive iterations (with H), with the added difficulty of choosing a proper scale and decay for the step-size γ_i .

A key insight of Cappé & Moulines (2009) is to use a different formulation of stochastic approximation, *not explicitly based on stochastic gradient ascent*. Indeed, they consider the stationary equation $\mathbb{E}_{t(X)} \left[\frac{\partial \log p(X|\eta)}{\partial \eta} \right] = 0$ and expand it using the exponential family model (1) as follows:

$$\begin{aligned}
\frac{\partial \log p(X|\eta)}{\partial \eta} &= \frac{\partial \log \int p(X, h|\eta) dh}{\partial \eta} \\
&= \phi'(\eta) \mathbb{E}_{p(h|X, \eta)} [S(X, h)] - \psi'(\eta).
\end{aligned}$$

Given standard properties of the exponential family, namely

$$\psi'(\eta) = \phi'(\eta) \mathbb{E}_{p(h, X|\eta)} [S(X, h)],$$

and assuming invertibility of $\phi'(\eta)$, this leads to the following stationary equation:

$$\mathbb{E}_{t(X)} [\mathbb{E}_{p(h|X, \eta)} [S(X, h)]] = \mathbb{E}_{p(h, X|\eta)} [S(X, h)].$$

This stationary equation states that at optimality the sufficient statistics have the same expectation for the full model $p(h, X|\eta)$ and the joint “model/data” distribution $t(X)p(h|X, \eta)$.

Another important insight of Cappé & Moulines (2009) is to consider the change of variable on sufficient statistics $s(\eta) = \mathbb{E}_{p(h, X|\eta)} [S(X, h)]$, which is equivalent to

$$\eta = \eta^*(s) \in \arg \max \langle \phi(\eta), s \rangle - \psi(\eta),$$

(which is the usual M-step update). See Cappé & Moulines (2009) for detailed assumptions allowing this inversion. We may then rewrite the equation above as

$$\mathbb{E}_{t(X)} (\mathbb{E}_{p(h|X, \eta^*(s))} [S(X, h)]) = s.$$

This is a non-linear equation in $s \in \mathbb{R}^d$, with an expectation with respect to $t(X)$ which is only accessed through i.i.d. samples X_i , and thus a good candidate for the Robbins-Monro algorithm to solve stationary equations (and not to minimize functions) (Kushner & Yin, 2003), which takes the simple form:

$$s_i = s_{i-1} - \gamma_i (s_{i-1} - \mathbb{E}_{p(h_i|X_i, \eta^*(s_{i-1}))} [S(X_i, h_i)]),$$

with a step-size γ_i . It may be rewritten as

$$\begin{cases} s_i = (1 - \gamma_i)s_{i-1} + \gamma_i \mathbb{E}_{p(h_i|X_i, \eta_{i-1})} [S(X_i, h_i)] \\ \eta_i = \eta^*(s_i), \end{cases} \quad (3)$$

which has a particularly simple interpretation: instead of computing the expectation for all observations as in full EM, this stochastic version keeps tracks of old sufficient statistics through the variable s_{i-1} which is updated towards the current value $\mathbb{E}_{p(h_i|X_i, \eta_{i-1})} [S(X_i, h_i)]$. The parameter η is then updated to the value $\eta^*(s_i)$. Cappé & Moulines (2009) show that this update is asymptotically equivalent to the natural gradient update with three main improvements: (a) no matrix inversion is needed, (b) the algorithm may be accelerated through Polyak-Ruppert averaging (Polyak & Juditsky, 1992), i.e., using the average $\bar{\eta}_N$ of all η_i instead of the last iterate η_N , and (c) the step-size is particularly simple to set, as we are taking *convex* combinations of sufficient statistics, and hence only the decay rate of γ_i has to be chosen, i.e., of the form $\gamma_i = i^{-\kappa}$, for $\kappa \in (0, 1]$, without any multiplicative constant.

Incremental view. For the specific stepsize $\gamma_i = 1/i$, the online EM algorithm (3) corresponds exactly to the incremental EM presented above (Neal & Hinton, 1998), as then $s_i = \frac{1}{i} \sum_{j=1}^i \mathbb{E}_{p(h_j|X_j, \eta_{j-1})} [S(X_j, h_j)]$. See Mairal (2014) for a detailed convergence analysis of incremental algorithms, in particular showing that step-sizes larger than $1/i$ are preferable (we observe this in practice in Section 5).

3 Online EM with intractable models

The online EM updates in (3) lead to a scalable algorithm for optimization when the local E-step is tractable. However, in many latent variable models—e.g., LDA, or hierarchical Dirichlet processes (Teh et al., 2006)—it is intractable to compute the conditional expectation $\mathbb{E}_{p(h|X, \eta)} [S(X, h)]$.

Following Rohde & Cappé (2011), we propose to leverage the scalability of online EM updates (3) and locally approximate the conditional distribution $p(h|X, \eta)$ in the case this distribution is intractable to compute. We will however consider different approximate methods, namely Gibbs sampling or variational inference. Our method is thus restricted to models where the hidden variable h may naturally be splitted in two or more groups of simple random variables. Our algorithm is described in Algorithm 1 and may be instantiated with two approximate inference schemes which we now describe.

Algorithm 1 Gibbs / Variational online EM

Input: $\eta_0, s_0, \kappa \in (0, 1]$.

for $i = 1, \dots, N$ **do**

- Collect observation X_i ,
- Estimate $p(h_i|X_i, \eta_{i-1})$ with sampling (G-OEM) or variational inference (V-OEM),
- Apply (3) to sufficient statistics s_i and parameter η_i with $\gamma_i = 1/i^\kappa$,

end for

Output: $\bar{\eta}_N = \frac{1}{N} \sum_{i=1}^N \eta_i$ or η_N .

3.1 Variational inference: V-OEM

While variational inference had been considered before for online estimation of latent variable models, in particular for LDA for incremental EM (Sato et al., 2010), using it for online EM (which is empirically faster) had not been proposed and allows to use bigger step-sizes (e.g., $\kappa = 1/2$). These methods are based on maximizing the negative variational “free-energy”

$$\mathbb{E}_{q(h|\eta)} \log \frac{p(X, h|\eta)}{q(h|\eta)}, \quad (4)$$

with respect to $q(h|\eta)$ having a certain factorized form adapted to the model at hand, so that efficient coordinate ascent may be used. See, e.g., Hoffman et al. (2013). We now denote online EM with variational approximation of the conditional distribution $p(h|X, \eta)$ as V-OEM.

3.2 Sampling methods: G-OEM

MCMC methods to approximate the conditional distribution of latent variables with online EM have been considered by Rohde & Cappé (2011), who apply locally Metropolis-Hasting (M-H) algorithm (Metropolis et al., 1953; Hastings, 1970), and show results on simple synthetic datasets. While Gibbs sampling is widely used for many models such as LDA due to its simplicity and lack of external parameters, M-H requires a proper proposal distribution with frequent acceptance and fast mixing, which may be hard to find in high dimensions. We provide a different simpler local scheme based on Gibbs sampling (thus adapted to a wide variety of models), and propose a thorough favorable comparison on synthetic and real datasets with existing methods.

The Gibbs sampler is used to estimate posterior distributions by alternatively sampling parts of the variables given the other ones (see Casella & George, 1992, for details), and is standard and easy to use in many common latent variable models. In the following, the online EM method with Gibbs estimation of the conditional distribution $p(h|X, \eta)$ is denoted G-OEM.

As mentioned above, the online EM updates correspond to a stochastic approximation algorithm and thus are robust to random noise in the local E-step. As a result, our sampling method is particularly adapted as it is a random estimate of the E-step—see a theoretical analysis by Rohde & Cappé (2011), and thus we only need to compute a few Gibbs samples for the estimation of $p(h|X_i, \eta_{i-1})$. A key contribution of our paper is to reuse sampling techniques that have proved competitive in the batch set-up and to compare them to existing variational approaches.

3.3 “Boosted” inference

As the variational and MCMC estimations of $p(h|X_i, \eta_{i-1})$ are done with iterative methods, we can boost the inference of Algorithm 1 by applying the update in the parameter η in (3) after each iteration of the estimation of $p(h|X_i, \eta_{i-1})$. In the context of LDA, this was proposed by Sato et al. (2010) for incremental EM and we extend it to all versions of online EM. With this boost, we expect that the global parameters η converge faster, as they are updated more often. In the following, we denote by G-OEM++ (resp. V-OEM++) the method G-OEM (resp. V-OEM) augmented with this boost.

3.4 Variational Bayesian estimation

In the Bayesian perspective where η is seen as a random variable, we either consider a distribution based on model averaging, e.g., $r(X) = \int p(X|\eta)q(\eta)d\eta$ where $q(\eta) \propto \prod_{i=1}^N p(X_i|\eta)p(\eta)$ is the posterior distribution, or

$$r(X) = p(X|\bar{\eta}),$$

where $\bar{\eta}$ is the summary (e.g., the mean) of the posterior distribution $q(\eta)$, or of an approximation, which is usually done in practice (see, e.g., Hoffman & Blei, 2015) and is asymptotically equivalent when N tends to infinity.

The main problem is that, *even when the conditional distribution of latent variables is tractable*, it is intractable to manipulate the *joint* posterior distribution over the latent variables h_1, \dots, h_N and the parameter η . Variational inference techniques consider an approximation where hidden variables are independent of the parameter η , i.e., such that $p(\eta, h_1, \dots, h_N|X_1, \dots, X_N) \approx q(\eta) \prod_{i=1}^N q(h_i)$, which corresponds to the maximization of the following lower bound—called Evidence Lower BOund (ELBO)—on the log-likelihood $\log p(X_1, \dots, X_n)$ (Bishop, 2006):

$$\int q(\eta) \prod_{i=1}^N q(h_i) \log \frac{p(\eta) \prod_{i=1}^n p(X_i, h_i|\eta)}{q(\eta) \prod_{i=1}^N q(h_i)} d\eta dh_1 \dots dh_N.$$

The key insight from Hoffman et al. (2010); Broderick et al. (2013) is to consider the variational distribution $q(\eta)$ as the global parameter, and the cost function above as a sum of *local* functions that depend on the data X_i and the variational distribution $q(h_i)$. Once the local variational distribution $q(h_i)$ is maximized out, the sum structure may be leveraged in similar ways than for frequentist estimation, either by direct (natural) stochastic gradient (Hoffman et al., 2010) or incremental techniques that accumulate sufficient statistics (Broderick et al., 2013). A nice feature of these techniques is that they extend directly to models with intractable latent variable inference, by making additional assumptions on $q(h_i)$ (see for example the LDA situation in Section 4).

In terms of actual updates, they are similar to online EM in Section 3.1, with a few changes, but which turn out to lead to significant differences in practice. The similarity comes from the expansion of the ELBO as

$$\mathbb{E}_{q(\eta)} \left[\sum_{i=1}^N \mathbb{E}_{q(h_i)} \log \frac{p(X_i, h_i|\eta)}{q(h_i)} \right] + \mathbb{E}_{q(\eta)} \left[\log \frac{p(\eta)}{q(\eta)} \right].$$

The left hand side has the same structure than the variational EM update in (4), thus leading to similar updates, while the right hand side corresponds to the “Bayesian layer”, and the maximization with respect to $q(\eta)$ is similar to the M-step of EM (where η is seen as a parameter).

Like online EM techniques presented in Section 3, approximate inference for latent variable is used, but, when using Bayesian stochastic variational inference techniques, there are two additional sources of inefficiencies: (a) extra assumptions regarding the independence of η and h_1, \dots, h_N , and (b) the lack of explicit formulation as the minimization of an expectation, which prevents the simple use of the most efficient stochastic approximation techniques (together with their guarantees). While (b) can simply slow down the algorithm, (a) may lead to results which are far away from exact inference, even for large numbers of samples (see examples in Section 5).

Beyond variational inference, Gibbs sampling has been recently considered by Gao et al. (2016): their method consists in sampling hidden variables for the current document given current parameters, but (a) only some of the new parameters are updated by incrementally aggregating the samples of the current document with current parameters, and (b) the method is slower than G-OEM (see Section 5).

4 Application to LDA

LDA (Blei et al., 2003) is a probabilistic model that infers hidden topics given a text corpus where each document of the corpus can be represented as topic probabilities. In particular, the assumption behind LDA is that each document is generated from a mixture of topics and the model infers the hidden topics and the topic proportions of each document. In practice, inference is done using Bayesian variational EM (Blei et al., 2003), Gibbs sampling (Griffiths & Steyvers, 2004; Wallach, 2006) or stochastic variational inference (Hoffman et al., 2010; Broderick et al., 2013; Sato et al., 2010).

Hierarchical probabilistic model. Let D be the number of documents of a corpus $\mathcal{C} = \{X_1, \dots, X_D\}$, V the number of words in our vocabulary and K the number of latent topics in the corpus. Each topic β^k corresponds to a discrete distribution on the V words (that is an element of the simplex in V dimensions). A hidden discrete distribution θ_i over the K topics (that is an element of the simplex in K dimensions) is attached to each document X_i . LDA is a generative model applied to a corpus of text documents which assumes that each word of the i^{th} document X_i is generated as follows:

- Choose $\theta_i \sim \text{Dirichlet}(\alpha)$,
- For each word $x_n \in X_i = (x_1, \dots, x_{N_{X_i}})$:
 - Choose a topic $z_n \sim \text{Multinomial}(\theta_i)$,
 - Choose a word $x_n \sim \text{Multinomial}(\beta^{z_n})$.

In our settings, an observation is a document $X_i = (x_1, \dots, x_{N_{X_i}})$ with $x_n \in \{0, 1\}^V$ and $\sum_{v=1}^V x_{nv} = 1$. Each observation X_i is associated with the hidden variables h_i , with $h_i \equiv (Z_i = (z_1, \dots, z_{N_{X_i}}), \theta_i)$. The vector θ_i represents the topic proportions of document X_i and Z_i is the vector of topic assignments of each word of X_i . The variable h_i is local, i.e., attached to one observation X_i . The parameters of the model are global, represented by $\eta \equiv (\beta, \alpha)$, where β represents the topic matrix and α represents the Dirichlet prior on topic proportions.

We derive the LDA model in Appendix A.1 to find ϕ , S , ψ and a such that the joint probability $p(Z, \theta | X, \alpha, \beta)$ is in a non-canonical exponential family (1).

We may then readily apply all algorithms from Section 3 by estimating the conditional expectation $\mathbb{E}_{Z, \theta | X, \alpha, \beta}[S(X, Z, \theta)]$ with either variational inference (V-OEM) or Gibbs sampling (G-OEM). See Appendices A.2 and A.3 for online EM derivations. Note that the key difficulty of LDA is the presence of two interacting hidden variables Z and θ .

Variational inference. See Hoffman et al. (2013) for detailed derivations of variational inference for LDA, in the Bayesian setting (from which the updates in the frequentist setting may be easily obtained). See also Appendix A.2.

Sampling. As shown in Appendix A.3, for LDA, we use the collapsed Gibbs sampler to marginalize out θ , using the fact that $Z | \alpha$ follows a Dirichlet-multinomial distribution (or multivariate Pólya distribution) and the following relation (Griffiths & Steyvers, 2004):

$$p(z_{nk} = 1 | z_{-n}, X, \beta, \alpha) \propto \beta_{x_n}^k \times \frac{N_{-n,k} + \alpha_k}{(N_X - 1) + \sum_j \alpha_j}, \quad (5)$$

Algorithm 2 Gibbs sampling scheme to approximate $p(z_{nk} = 1|X, \beta, \alpha)$

Input: β, α, X .

Initialization: $Z_n^0 \sim \text{Mult}([\bar{\beta}_{x_n}^k]_{k=1,\dots,K})$, with $\bar{\beta}_{x_n}^k = \frac{\beta_{x_n}^k}{\sum_j \beta_{x_n}^j} \forall n \in \llbracket 1, N_X \rrbracket$.

for $t = 1, 2, \dots, P$ **do**

 Compute random permutation σ^t on $\llbracket 1, N_X \rrbracket$,

for $n = 0, 1, \dots, N_X$ **do**

 • Set $Z_{-n}^t = \left\{ (z_{\sigma^t(i)}^t)_{1 \leq i < n}, (z_{\sigma^t(i)}^{t-1})_{n < i \leq N_X} \right\}$,

 • Compute $\forall k, p(z_{\sigma^t(n)k} = 1|Z_{-n}^t, X, \beta, \alpha)$ with Equation (5),

 • Sample $z_{\sigma^t(n)}^t \sim \text{Mult} [p(z_{\sigma^t(n)}^t|Z_{-n}^t, X, \beta, \alpha)]$,

end for

end for

for $n = 0, 1, \dots, N_X$ **do**

 Set $Z_{-n}^t = \{(z_i^t)_{1 \leq i < n}, (z_i^t)_{n < i \leq N_X}\}$ **for** $t \geq \frac{3}{4}P$,

$p(z_{nk} = 1|X, \beta, \alpha) \leftarrow \frac{4}{P} \sum_{t=\frac{3}{4}P}^P p(z_{nk} = 1|Z_{-n}^t, X, \beta, \alpha)$

end for

Output: $\forall k, \forall n: (z_n^t)_{t=1,\dots,P}, p(z_{nk} = 1|X, \beta, \alpha)$.

with $N_{-n,k}$ the number of words assigned to topic k in the current document, except index n . In order to illustrate the simplicity of the local Gibbs sampling algorithms, we present it in Algorithm 2 (note that we average over the last quarter of samples, which corresponds to a short burn-in phase of three-quarters of the updates).

5 Evaluation

We evaluate our method by computing the likelihood on held-out documents, that is $p(X|\beta, \alpha)$ for any test document X . For LDA, the likelihood is intractable to compute. We approximate $p(X|\beta, \alpha)$ with the “left-to-right” evaluation algorithm (Wallach et al., 2009) applied to each test document. This algorithm is a mix of particle filtering and Gibbs sampling. On any experiments, this leads essentially to the same log-likelihood than Gibbs sampling with sufficiently enough samples—e.g., 200. In the following, we present results in terms of log-perplexity, defined as the opposite of the log-likelihood $-\log p(X|\eta)$. The lower the log-perplexity, the better the corresponding model. In our experiments, we compute the average test log-perplexity on N_t documents. We compare six different methods:

- G-OEM (our main algorithm): Gibbs online EM. Online EM algorithm with Gibbs estimation of the conditional distribution $p(h|X, \eta)$ (Algorithm 2). Frequentist approach and step-size $\gamma_i = 1/\sqrt{i}$;
- V-OEM++: variational online EM (also a new algorithm). Online EM algorithm with variational estimation of the conditional distribution $p(h|X, \eta)$, augmented with inference boosting from Section 3.3. Frequentist approach and step-size $\gamma_i = 1/\sqrt{i}$;
- OLDA: online LDA (Hoffman et al., 2010). Bayesian approach which maximizes the ELBO from Section 3.4, with natural stochastic gradient ascent and a step-size $\gamma_i = 1/\sqrt{i}$;
- SVB: streaming variational Bayes (Broderick et al., 2013). A variational Bayesian equivalent of V-OEM with step-size $\gamma_i = 1/i$;
- SPLDA: single pass LDA (Sato et al., 2010). The difference with V-OEM++ is that $\gamma_i = 1/i$ and the updates in α done with a Gamma prior (see Appendix F);

- **SGS:** streaming Gibbs sampling (Gao et al., 2016). This method is related to G-OEM with $\gamma_i = 1/i$. In this method, α is not optimized and set to a constant C_α . For comparison purposes, for each dataset, we set C_α to be the averaged final parameter $\hat{\alpha}$ obtained with G-OEM on the same dataset: $C_\alpha = \frac{1}{K} \sum_k \hat{\alpha}_k$. For each observation, only the last Gibbs sample is considered, leading to extra noise in the output.

For existing variational methods—OLDA, SVB, SPLDA— β is a random variable with prior $q(\beta)$. We estimate the likelihood $p(X|\hat{\beta}, \alpha)$ with the “left-to-right” algorithm by setting $\hat{\beta} = \mathbb{E}_q[\beta]$ for Bayesian methods. Due to space constraints we only present our results obtained with G-OEM and V-OEM++. Indeed, the inference boost presented in Section 4 is only beneficial for V-OEM. A detailed analysis is presented in Appendix B.1.

5.1 General settings

Initialization. We initialize randomly $\eta \equiv (\beta, \alpha)$. For a given experiment, we initialize all the methods with the same values of (β, α) for fair comparison, except SPLDA that has its own initialization scheme—see Sato et al. (2010) for more details.

Minibatch. We consider minibatches of size 100 documents for each update in order to reduce noise (Liang & Klein, 2009). In the case of online EM in Equation (3), we estimate an expectation for each observation of the minibatch. We update the new sufficient statistics s towards the average of the expectations over the minibatch. We do the same averaging for all the presented methods.

Number of local updates. For all the presented methods, we set the number of passes through each minibatch to $P = 20$. For G-OEM, this means that we perform 20 Gibbs sampling for each word of the minibatch. All other methods access each document 20 times (e.g., 20 iterations of variational inference on each document).

Datasets. We apply the methods on six different datasets, summarized below ($\overline{N_X}$ is the average length of documents). Following Blei et al. (2003), the synthetic dataset has been generated from 10 topics and the length of each document drawn from a Poisson(60).

DATASET	#DOCUMENTS	$\overline{N_X}$	#WORDS
SYNTHETIC	1,000,000	60	1,000
WIKIPEDIA ¹	1,010,000	162.3	7702
IMDB ²	614,589	82.2	10,000
AMAZON MOVIES ³	338,565	75.4	10,000
NEW YORK TIMES ⁴	299,877	287.4	44,228
PUBMED ⁴	2,100,000	82.0	113,568

The words in the datasets IMDB, Wikipedia, New York Times, Pubmed and Amazon movies are filtered by removing the stop-words and we select the most frequent words of the datasets. For the synthetic dataset, IMDB, Pubmed and Amazon movies, the size of the test sets is $N_t = 5,000$ documents. For Wikipedia and New York Times, the test sets contain $N_t = 2,000$ documents. We run the methods on 11 different train/test splits of each dataset. For all the presented results, we plot the median from the 11 experiments as a line—solid or dashed. For the sake of readability, we only present the same plots with error bars between the third and the seventh decile in Appendix D and Appendix E.

¹Code available from Hoffman et al. (2010)

²Dataset described in Diao et al. (2014)

³Data from Leskovec & Krevl (2014)

⁴UCI dataset (Lichman, 2013)

Table 1: Average computational time (in hours) for each method.

	IMDB	WIKIPEDIA	NYT	PUBMED
G-OEM	13H	55H	30H	58H
V-OEM++	9H	37H	20H	54H
OLDA	7H	33H	8H	30H
SVB	7H	34H	9H	30H
SPLDA	9H	37H	20H	54H
SGS	11H	48H	27H	50H

Computation time. For each presented method and dataset, the computational time is reported in Table 1. Although all methods have the same running-time complexities, coded in Python, sampling methods (G-OEM and SGS) need an actual loop over all documents while variational methods (OLDA, SVB, SPLDA and V-OEM++) may use vector operations, and may thus be up to twice faster. This could be mitigated by using efficient implementations of Gibbs sampling on minibatches (Yan et al., 2009; Zhao et al., 2014; Gao et al., 2016). Note also that to attain a given log-likelihood, our method G-OEM is significantly faster and often attains log-likelihoods not attainable by other methods (e.g., New York Times).

Step-size. In the following, we compare the results of our methods G-OEM and V-OEM++ with $\kappa = 1/2$, i.e., the step-size $\gamma_i = 1/\sqrt{i}$, without averaging. Detailed analysis of different settings of our method can be found in Appendix B. In particular, we compare different step-sizes and the effect of averaging over all iterates.

5.2 Results

Results obtained with the presented methods on different datasets for different values of the number K of topics are presented in Figure 1. Performance through iterations (i.e., as the number of documents increases) is presented in Figure 3. We first observe that for all experiments, our new methods G-OEM and V-OEM++ perform significantly better than all existing methods.

Influence of the number of topics K . As shown in Figure 1, for synthetic data in plot (a), although the true number of topics is $K^* = 10$, SPLDA and OLDA perform slightly better with $K = 20$, while G-OEM has the better fit for the correct value of K ; moreover, SVB has very similar performances for any value of K , which highlights the fact that this method does not capture more information with a higher value of K .

On non-synthetic datasets in plots (b)-(f), while the log-perplexity of G-OEM and V-OEM++ and SPLDA decreases with K , the log-perplexity of Bayesian methods—OLDA, and SVB—does not decrease significantly with K . As explained below, our interpretation is that the actual maximization of the ELBO does not lead to an improvement in log-likelihood.

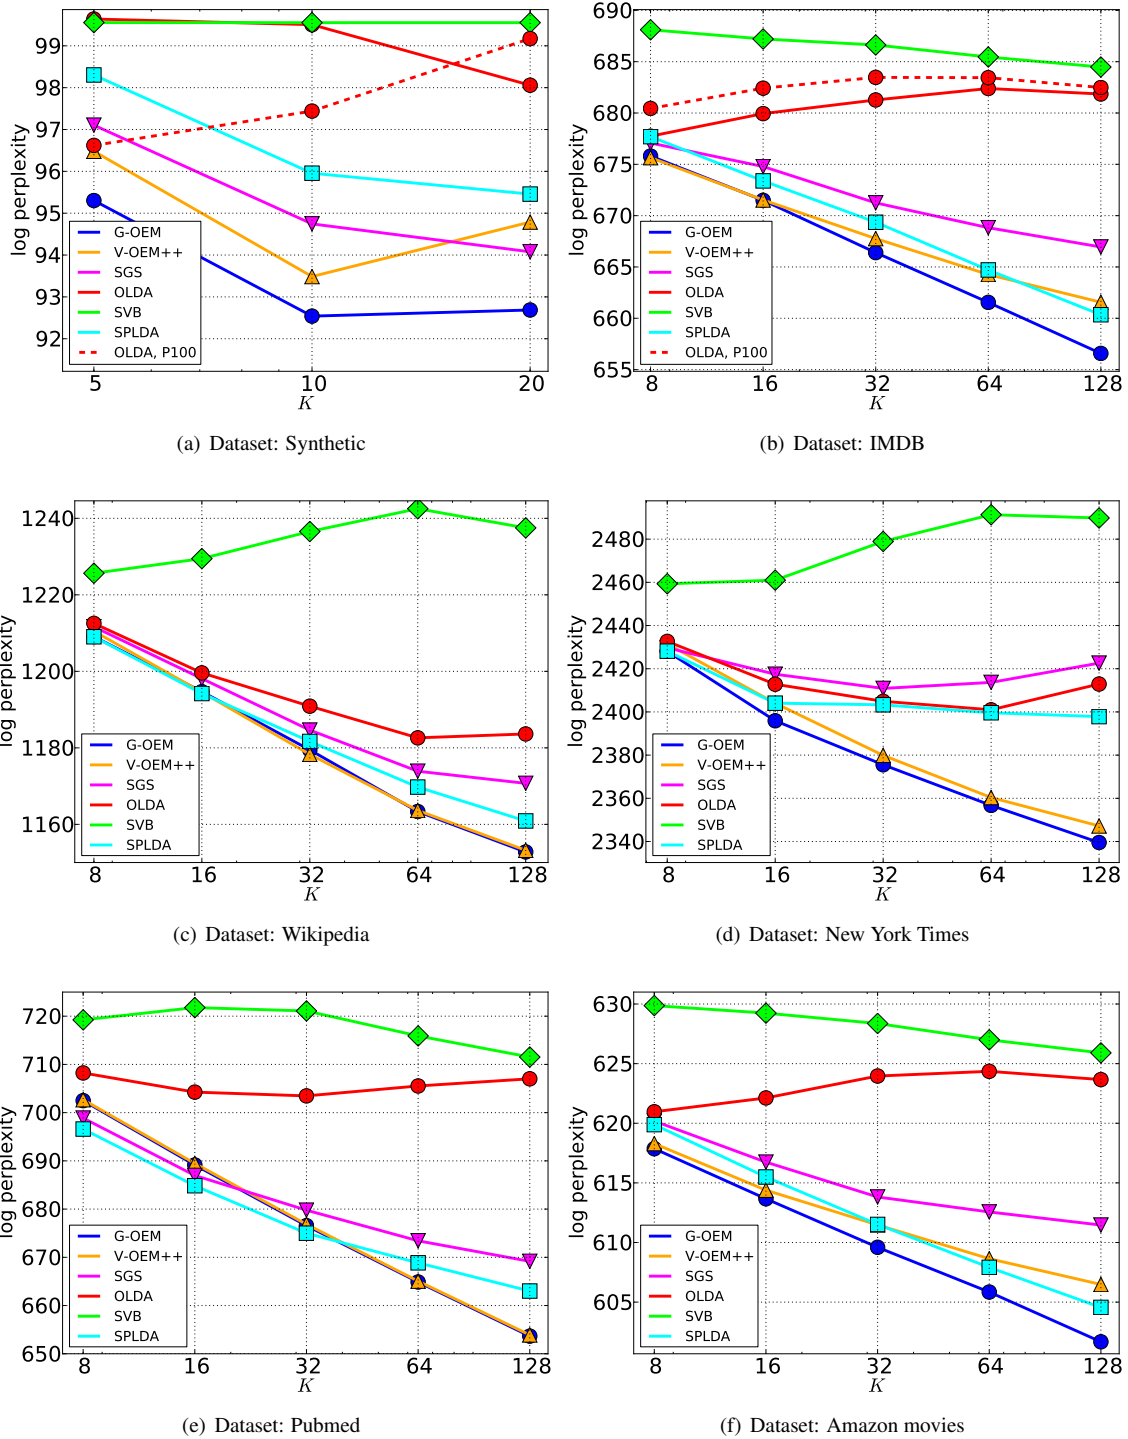


Figure 1: Perplexity on different test sets as a function of K , the number of topics inferred. OLDA, P100 corresponds to OLDA with $P = 100$ internal iterations. Best seen in color.

Performance through iterations. As shown in Figure 3, for synthetic data in plot (a), after only few dozens of iterations—few thousands of documents seen— G -OEM and V -OEM++ outperform the other presented methods. Bayesian methods again do converge but to a worse parameter value. On real datasets in plots (b)-(f), G -OEM is still significantly faster; we can indeed still observe that after around 100 iterations—10,000 documents seen— G -OEM performs better than other methods on all the datasets except Pubmed, where the performances of G -OEM, V -OEM++ and SPLDA are similar.

Variational vs. sampling. Our method G -OEM directly optimizes the likelihood with a consistent approximation, and performs better than its variational counterparts SPLDA and V -OEM++ in all experiments.

Frequentist vs. Bayesian. In all our experiments we observe that frequentist methods — G -OEM, V -OEM++ and SPLDA—outperform Bayesian methods—OLDA and SVB. As described in Section 3.4, Bayesian methods maximize the ELBO, which makes additional strong independence assumptions and here leads to poor results. For example, as the number K of topics increases, the log-likelihood goes down for some datasets. In order to investigate if this is an issue of slow convergence, we show on Figure 1 (dotted red line) that running $P = 100$ internal updates in OLDA to get a finer estimate of the ELBO for each document may deteriorate the performance. Moreover, Figure 2 presents the evolution of the ELBO, which does always increase when K increases, showing that the online methods do optimize correctly the ELBO (while not improving the true log-likelihood). See Appendix C for additional results on the convergence of the ELBO.

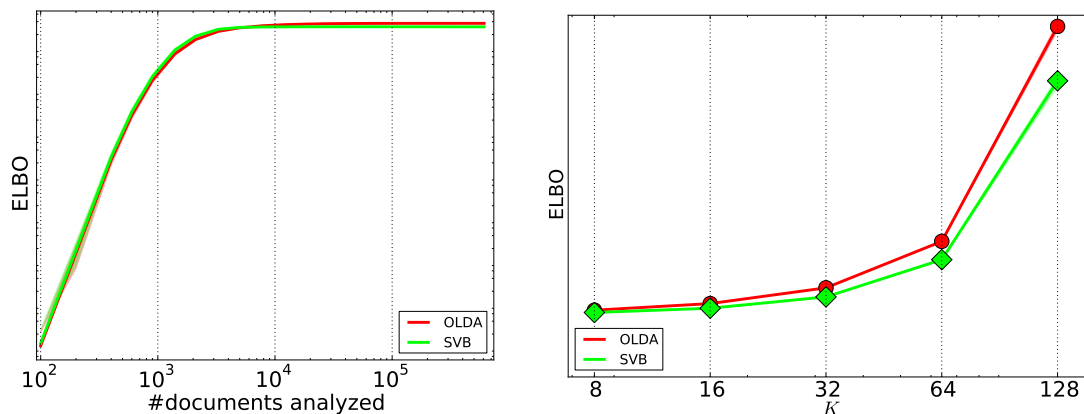
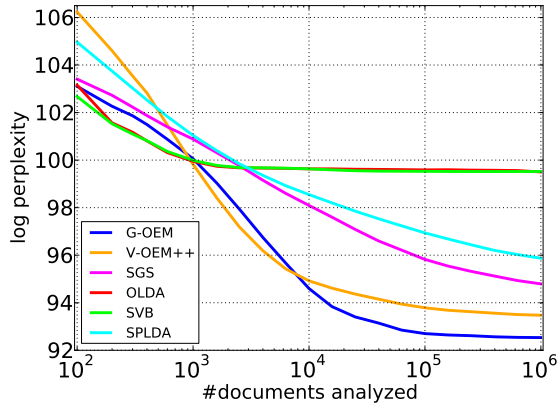
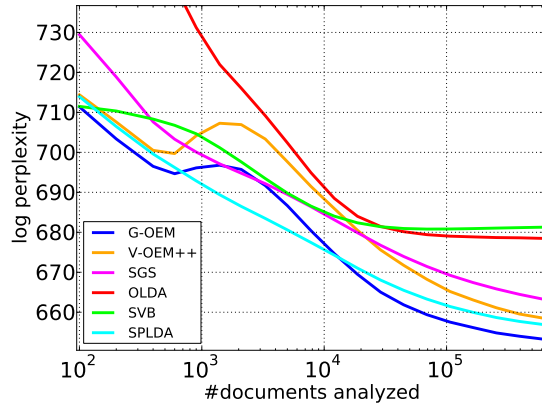


Figure 2: Dataset: IMDB. Evidence Lower Bound (ELBO) computed on test sets (20 internal iterations and 1 pass over the dataset). *Left*: ELBO through iterations with $K = 128$. *Right*: ELBO as a function of the number of topics K .

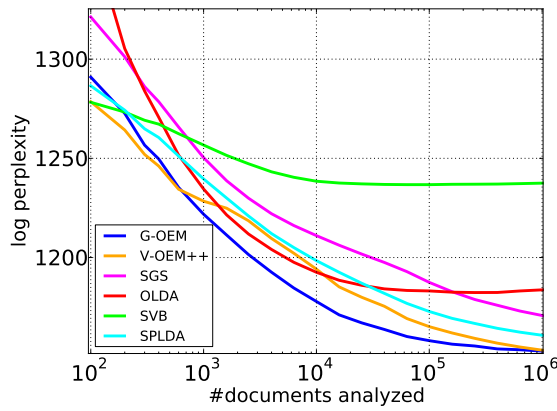
Small step-sizes vs. large step-sizes. SPLDA is also a variational method which is equivalent to V -OEM++, but with a step-size $1/i$, which is often slower than bigger step-sizes (Mairal, 2014), which we do observe—see Appendix B.2 for a further analysis on the effect of the choice of step-sizes as $1/i^k$ on G -OEM.



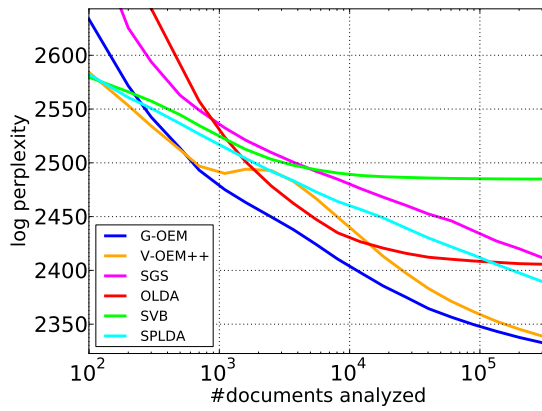
(a) Synthetic, $K = 10$



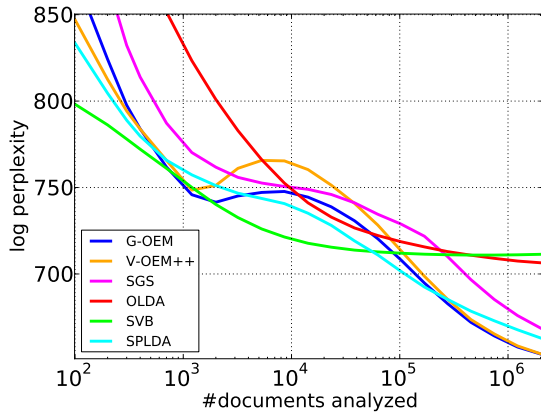
(b) IMDB, $K = 128$



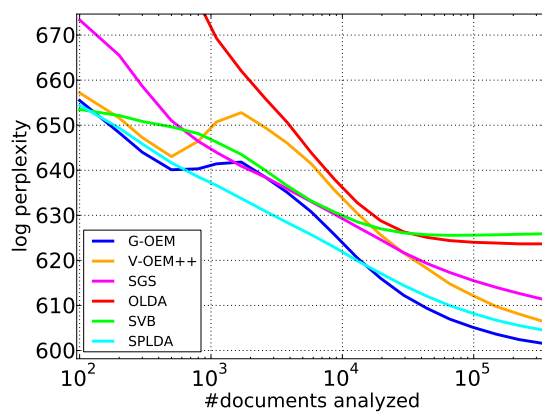
(c) Wikipedia, $K = 128$



(d) New York Times, $K = 128$



(e) Pubmed, $K = 128$



(f) Amazon movies, $K = 128$

Figure 3: Perplexity through iterations on different test sets with the presented methods. Best seen in color.

Qualitative results. Due to lack of space, we propose in Appendix G an empirical qualitative study of the extracted topics on IMDB. Our method G -OEM, on top of obtaining better log-likelihoods, also extracts more consistent topics and “qualitative” topics (i.e., with only positive/negative words) that other methods are not able to extract.

6 Conclusion

We have developed an online inference scheme to handle intractable conditional distributions of latent variables, with a proper use of local Gibbs sampling within online EM, that leads to significant improvements over variational methods and Bayesian estimation procedures. It would be interesting to explore distributed large-scale settings (Broderick et al., 2013; Yan et al., 2009; Gao et al., 2016) and potentially larger (e.g., constant) step-sizes that have proved efficient in supervised learning (Bach & Moulines, 2013).

Acknowledgements

We would like to thank David Blei, Olivier Cappé and Nicolas Flammarion for helpful discussions related to this work.

References

- Bach, F. and Moulines, E. Non-strongly-convex smooth stochastic approximation with convergence rate $O(1/n)$. In *Adv. NIPS*, 2013.
- Bird, S., Klein, E., and Loper, E. *Natural language processing with Python*. ” O’Reilly Media, Inc.”, 2009.
- Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer, 2006.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. Latent Dirichlet allocation. *JMLR*, 3:993–1022, 2003.
- Bottou, L. Online learning and stochastic approximations. *On-line learning in neural networks*, 17:9, 1998.
- Broderick, T., Boyd, N., Wibisono, A., Wilson, A. C., and Jordan, M. I. Streaming variational Bayes. *Adv. NIPS*, 2013.
- Cappé, O. and Moulines, E. Online EM algorithm for latent data models. *Journal of the Royal Statistical Society*, 71(3):593–613, 2009.
- Casella, G. and George, E. I. Explaining the Gibbs sampler. *The American Statistician*, 46(3):167–174, 1992.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society. Series B (methodological)*, 39(1):1–38, 1977.
- Diao, Q., Qiu, M., Wu, C.-Y., Smola, A. J., Jiang, J., and Wang, C. Jointly modeling aspects, ratings and sentiments for movie recommendation (JMARS). In *Proc. ACM SIGKDD*, 2014.
- Gao, Yang, Chen, Jianfei, and Zhu, Jun. Streaming gibbs sampling for LDA model. *arXiv preprint arXiv:1601.01142*, 2016.
- Griffiths, T. L. and Steyvers, M. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(suppl 1):5228–5235, 2004.
- Hastings, K. W. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

- Hoffman, M. D. and Blei, D. M. Structured stochastic variational inference. In *Proc. AISTATS*, 2015.
- Hoffman, M. D., Blei, D. M., and Bach, F. R. Online learning for latent Dirichlet allocation. *Adv. NIPS*, 2010.
- Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. Stochastic variational inference. *JMLR*, 14(1): 1303–1347, 2013.
- Koller, D. and Friedman, N. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- Kushner, H. J. and Yin, G. G. *Stochastic approximation and recursive algorithms and applications*. Springer-Verlag, second edition, 2003.
- Lehmann, Erich Leo and Casella, George. *Theory of point estimation*, volume 31. Springer Science & Business Media, 1998.
- Leskovec, J. and Krevl, A. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>, June 2014.
- Liang, P. and Klein, D. Online EM for unsupervised models. In *Proceedings of human language technologies: The 2009 annual conference of the North American chapter of the association for computational linguistics*, pp. 611–619. Association for Computational Linguistics, 2009.
- Lichman, M. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Mairal, Julien. Incremental majorization-minimization optimization with application to large-scale machine learning. *arXiv preprint arXiv:1402.4419*, 2014.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- Minka, T. P. Estimating a dirichlet distribution. Technical report, 2000.
- Murphy, K. P. *Machine learning: a probabilistic perspective*. MIT Press, 2012.
- Neal, R. M. and Hinton, G. E. A view of the EM algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*, pp. 355–368. Springer, 1998.
- Polyak, B. T. and Juditsky, A. B. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- Rohde, D. and Cappé, O. Online maximum-likelihood estimation for latent factor models. In *Statistical Signal Processing Workshop (SSP), 2011 IEEE*. IEEE, 2011.
- Sato, I., Kurihara, K., and Nakagawa, H. Deterministic single-pass algorithm for LDA. *Adv. NIPS*, 2010.
- Teh, Y. W., Jordan, M. I., Beal, M. J., and Blei, D. M. Hierarchical Dirichlet processes. *Journal of the american statistical association*, 101(476):1566–1581, 2006.
- Titterton, M. D. Recursive parameter estimation using incomplete data. *Journal of the Royal Statistical Society. Series B (Methodological)*, 46(2):257–267, 1984.
- Van der Vaart, A. W. *Asymptotic Statistics*, volume 3. Cambridge University Press, 2000.
- Wallach, H. M. Topic modeling: beyond bag-of-words. In *Proc. ICML*, 2006.
- Wallach, H. M., Murray, I., Salakhutdinov, R., and Mimno, D. Evaluation methods for topic models. *Proc. ICML*, 2009.

Yan, F., Xu, N., and Qi, Y. Parallel inference for latent Dirichlet allocation on graphics processing units. In *Adv. NIPS*, 2009.

Zhao, H., Jiang, B., and Canny, J. Same but different: Fast and high-quality Gibbs parameter estimation. *arXiv preprint arXiv:1409.5402*, 2014.

A Application of online EM to LDA

Let D be the number of documents of a corpus $\mathcal{C} = \{X_1, \dots, X_D\}$, V the number of words in our vocabulary and K the number of latent topics in the corpus. Each topic β^k corresponds to a discrete distribution on the V words (that is an element of the simplex in V dimensions). A hidden discrete distribution θ_i over the K topics (that is an element of the simplex in K dimensions) is attached to each document i . LDA is a generative model applied to a corpus of text documents which assumes that each word of the i^{th} document X_i is generated as follows:

- Choose $\theta^i \sim \text{Dirichlet}(\alpha)$
- For each word $x_n \in X^i = (x_1, \dots, x_{N_{X^i}})$:
 - Choose a topic $z_n \sim \text{Multinomial}(\theta^i)$,
 - Choose a word $x_n \sim \text{Multinomial}(\beta^{z_n})$.

In our settings, an observation is a document X_i , associated with the hidden variables $h_i \equiv (Z_i, \theta_i)$, with $Z_i = (z_1, \dots, z_{N_{X^i}})$. The vector θ_i represents the topic proportions of document X_i and Z_i is the vector of topic assignments of each word of X_i . The variable h_i is local, as attached to one observation. The parameters of the model are global, represented by $\eta \equiv (\beta, \alpha)$, where β represents the topic matrix and α represents the Dirichlet prior on topic proportions.

A.1 LDA and exponential families

An observation X is a document of length N_X , where $X = (x_1, \dots, x_{N_X})$, $x_n \in \{0, 1\}^V$ and $\sum_{v=1}^V x_{nv} = 1$. Our corpus \mathcal{C} is a set of D observations $\mathcal{C} = (X^1, \dots, X^D)$. For each document X^i a hidden variable θ^i is associated, corresponding to the topic distribution of document X^i . For each word x_n of document X^i a hidden variable $z_n \in \{0, 1\}^K$ is attached, corresponding to the topic assignment of word x_n . We want to find ϕ , S , ψ and a such that, the joint probability is in the exponential family (1):

$$p(X, Z, \theta | \beta, \alpha) = a(X, Z, \theta) \exp [\langle \phi(\beta, \alpha), S(X, Z, \theta) \rangle - \psi(\beta, \alpha)],$$

given an observation X and hidden variables Z and θ . For the LDA model, we have:

$$\begin{aligned} p(X, Z, \theta | \beta, \alpha) &= \prod_{n=1}^{N_X} p(x_n | z_n, \beta) p(z_n | \theta) p(\theta | \alpha) \\ &= \prod_{n=1}^{N_X} \prod_{k=1}^K \prod_{v=1}^V [(\beta_v^k)^{x_{nv}} \theta_k]^{z_{nk}} p(\theta | \alpha), \end{aligned}$$

$$\begin{aligned} p(X, Z, \theta | \beta, \alpha) &= \exp \left[\sum_{n=1}^{N_X} \sum_{k=1}^K z_{nk} \log \theta_k \right] \\ &\times \exp \left[\sum_{n=1}^{N_X} \sum_{k=1}^K \sum_{v=1}^V x_{nv} z_{nk} \log \beta_v^k \right] \\ &\times \exp \left[\sum_{k=1}^K (\alpha_k - 1) \log \theta_k + B(\alpha) \right], \end{aligned}$$

with $B(\alpha) = \log \left[\Gamma \left(\sum_{i=1}^K \alpha_i \right) \right] - \sum_{i=1}^K \log[\Gamma(\alpha_i)]$. We deduce the non-canonical exponential family setting $\phi, S, \psi a$:

$$S(X, Z, \theta) = \begin{pmatrix} S_{kv}^1 \equiv \left[\sum_{n=1}^{N_X} z_{nk} x_{nv} \right]_{kv} \\ S_k^2 \equiv [\log \theta_k]_k \end{pmatrix}, \quad (6)$$

$$\phi(\beta, \alpha) = \begin{pmatrix} \phi_{kv}^1 \equiv [\log \beta_v^k]_{kv} \\ \phi_k^2 \equiv [\alpha_k]_k \end{pmatrix}, \quad (7)$$

with $S^1, \phi^1 \in \mathbb{R}^{K \times V}$ and $S^2, \phi^2 \in \mathbb{R}^K$,

$$\psi(\beta, \alpha) = \sum_{i=1}^K \log[\Gamma(\alpha_i)] - \log \left[\Gamma \left(\sum_{i=1}^K \alpha_i \right) \right], \quad (8)$$

and

$$a(X, Z, \theta) = \exp \left[\sum_{k=1}^K \left(\sum_{n=1}^{N_X} z_{nk} - 1 \right) \log \theta_k \right].$$

The one-to one mapping between the sufficient statistics $s = \begin{pmatrix} s_1 \\ s_2 \end{pmatrix}$ and (β, α) is defined by:

$$(\beta, \alpha)^*[s] = \begin{cases} \arg \max_{\beta \geq 0, \alpha \geq 0} \langle \phi(\beta, \alpha), s \rangle - \psi(\beta, \alpha) \\ \text{s.t.} \quad \beta^\top \mathbf{1} = \mathbf{1}. \end{cases}$$

The objective $\langle \phi(\beta, \alpha), s \rangle - \psi(\beta, \alpha)$ is concave in β from concavity of log and concave on any α_k for $\alpha \geq 0$ as the function $B(\alpha) = \log \left[\Gamma \left(\sum_{i=1}^K \alpha_i \right) \right] - \sum_{i=1}^K \log[\Gamma(\alpha_i)]$ is concave as the negative log-partition of the Dirichlet distribution. We use the Lagrangian method for β :

$$L(\beta, \lambda) = \sum_{k=1}^K \sum_{v=1}^V s_{kv}^1 \log \beta_v^k + \lambda^\top (\beta^\top \mathbf{1} - \mathbf{1}),$$

with $\lambda \in \mathbb{R}^K$. The derivative of L is cancelled when:

$$\forall(k, v), \frac{s_{kv}^1}{\beta_v^k} + \lambda_k = 0 \Rightarrow \lambda_k = - \sum_{v=1}^V s_{kv}^1,$$

as $\sum_{v=1}^V \beta_v^k = 1$. We then have $(\beta^*(s))_{kv} = \frac{s_{kv}^1}{\sum_j s_{kj}^1}$. This mapping satisfies the constraint $\beta \geq 0$ because for any observation X and hidden variable Z , we have $S^1(X, Z)_{kv} \geq 0$. This comes from (6) and the fact that $\forall(n, k, v), (x_{nv}, z_{nk}) \in \{0, 1\}^2$. We find the condition on α by setting the derivatives to 0, which gives $\forall k \in \llbracket 1, K \rrbracket$:

$$s_k^2 - \Psi([\alpha^*(s)]_k) + \Psi \left(\sum_{i=1}^K [\alpha^*(s)]_i \right) = 0,$$

where $\Psi : x \mapsto \frac{\partial}{\partial x} [\log \Gamma](x)$ is the digamma function. Finally, $(\alpha^*(s), \beta^*(s))$ satisfies $\forall(k, v)$:

$$\begin{cases} (\beta^*(s))_{kv} & \equiv \left[\frac{s_{kv}^1}{\sum_j s_{kj}^1} \right]_{kv} \\ \Psi([\alpha^*(s)]_k) - \Psi \left(\sum_{i=1}^K [\alpha^*(s)]_i \right) & = s_k^2. \end{cases}$$

α^* is usually estimated with gradient ascent method (Blei et al., 2003; Hoffman et al., 2010). We can also estimate α with the fixed point iteration (Minka, 2000) which consists in repeating the following update until convergence:

$$\alpha_k^{new} = \Psi^{-1} \left(\Psi \left(\sum_{i=1}^K \alpha_i^{old} \right) + s_k^2 \right).$$

We use the fixed point iteration to estimate α^* as it is more stable in practice. We study different updates for α in Appendix F.

We can now apply Algorithm 1 to LDA. The only missing step is the estimation of the conditional expectation $\mathbb{E}_{Z, \theta | X, \alpha_t, \beta_t} [S(X, Z, \theta)]$, with $X = (x_1, \dots, x_{N_X})$ and $Z = (z_1, \dots, z_{N_X})$. We explain how to approximate this expectation with variational inference and Gibbs sampling.

A.2 Variational online EM applied to LDA (V-OEM)

In this section we explain how to approximate $\mathbb{E}_{Z, \theta | X, \alpha_t, \beta_t} [S(X, Z, \theta)]$ with variational inference. The idea behind variational inference is to maximize the Evidence Lower Bound (ELBO), a lower bound on the probability of the observations:

$$p(X) \geq \text{ELBO}(X, p, q),$$

where q represents the variational model. In the case of LDA, the variational model is often set with a Dirichlet(γ) prior on θ and a multinomial prior on Z (Hoffman et al., 2013):

$$q(Z, \theta) = q(\theta | \gamma) \prod_{n=1}^{N_X} q(z_n | \zeta_n). \quad (9)$$

We then maximize the ELBO with respect to γ and ζ , which is equivalent to minimize the Kullback-Leibler (KL) divergence between the variational posterior and the true posterior:

$$\max_{\gamma, \zeta} \text{ELBO}(X, p, q) \Leftrightarrow \min_{\gamma, \zeta} \text{KL}[p(Z, \theta | X) || q(\theta, Z)]. \quad (10)$$

We solve this problem with block coordinate descent, which leads to iteratively updating γ and ζ as follows:

$$\zeta_{nk} \propto \prod_{v=1}^V (\beta_v^k)^{x_{nv}} \exp[\Psi(\gamma_k)], \quad (11)$$

$$\gamma_k = \alpha_k + \sum_{n=1}^{N_X} \zeta_{nk}. \quad (12)$$

We then approximate $\mathbb{E}_{Z, \theta | X, \alpha_t, \beta_t} [S(X, Z, \theta)]$ with the variational posterior. Given (6) and (9), we have:

$$\mathbb{E}_{p(Z, \theta | X)} [S(X, Z, \theta)] \approx \mathbb{E}_{q(Z, \theta)} [S(X, Z, \theta)] \quad (13)$$

$$= \left[\begin{array}{c} \left(\sum_{n=1}^{N_{X_{t+1}}} \zeta_{nk} x_{nv} \right)_{kv} \\ \left(\Psi(\gamma_k) - \Psi \left(\sum_{j=1}^K \gamma_j \right) \right)_k \end{array} \right]. \quad (14)$$

The variational approximation of $\mathbb{E}_{p(Z, \theta | X)} [S(X, Z, \theta)]$ is then done in two steps:

1. Iteratively update ζ with (11) and γ with (12),
2. $\mathbb{E}_{p(Z, \theta | X)} [S(X, Z, \theta)] \leftarrow \mathbb{E}_{q(Z, \theta | \gamma, \zeta)} [S(X, Z, \theta)]$, equation (14).

As γ and ζ are set to minimize the distance between the variational posterior and the true posterior (10) we expect that this approximation is close to the true expectation. However, as the variational model is a simplified version of the true model, there always remains a gap between the true posterior and the variational posterior.

A.3 Gibbs online EM applied to LDA (G-OEM)

In this section we explain how to approximate $\mathbb{E}_{Z,\theta|X,\alpha,\beta} [S(X, Z, \theta)]$ with Gibbs sampling.

Expectation of S^1 . Given (6), we have $\forall k \in \llbracket 1, K \rrbracket, \forall v \in \llbracket 1, V \rrbracket$:

$$\begin{aligned} \mathbb{E}_{Z,\theta|X,\alpha,\beta} [(S^1(X, Z))_{kv}] &= \mathbb{E}_{Z,\theta|X,\alpha,\beta} \left[\sum_{n=1}^{N_X} z_{nk} x_{nv} \right] \\ &= \sum_{n=1}^{N_X} \int_{Z,\theta} z_{nk} x_{nv} p(z_n, \theta | X, \beta, \alpha) d\theta dz \\ &= \sum_{n=1}^{N_X} x_{nv} p(z_{nk} = 1 | X, \beta, \alpha). \end{aligned}$$

We see that we only need the probability of z , and can thus use collapsed Gibbs sampling (Griffiths & Steyvers, 2004).

We have, following Bayes rule:

$$\begin{aligned} p(z_{nk} = 1 | z_{-n}, X, \beta, \alpha) \\ \propto p(x_n | z_{nk} = 1, \beta) p(z_{nk} = 1 | z_{-n}, \alpha), \end{aligned}$$

where z_{-n} is the topic assignments except index n . From the LDA model: $p(x_n | z_{nk} = 1, \beta) = \sum_{v=1}^V x_{nv} \beta_v^k$. In the following, we use the notation $\beta_{x_n}^k \equiv \sum_{v=1}^V x_{nv} \beta_v^k$ for the sake of simplicity. We then use the fact that the topic proportions θ has a Dirichlet(α) prior, which implies that $Z | \alpha$ follows a Dirichlet-multinomial distribution (or multivariate Pólya distribution). As a result, the conditional distribution is:

$$p(z_{nk} = 1 | z_{-n}, \alpha) = \frac{N_{-n,k} + \alpha_k}{(N_X - 1) + \sum_j \alpha_j},$$

with $N_{-n,k}$ the number of words assigned to topic k in the current document, except index n . Finally, we have the following relation (Griffiths & Steyvers, 2004):

$$p(z_{nk} = 1 | z_{-n}, X, \beta, \alpha) \propto \beta_{x_n}^k \times \frac{N_{-n,k} + \alpha_k}{(N_X - 1) + \sum_j \alpha_j}. \quad (15)$$

We estimate $p(z_{nk} = 1 | X, \beta, \alpha)$ with Gibbs sampling by iteratively sampling topic assignments z_n for each word, as detailed in Algorithm 2. We average over the last quarter of samples to reduce noise in the final output. We then incorporate the output in Algorithm 1 in the main text.

Expectation of S^2 . Given (6), we also have $\forall k \in \llbracket 1, K \rrbracket, \forall v \in \llbracket 1, V \rrbracket$:

$$\mathbb{E}_{Z,\theta|X,\alpha,\beta} [(S^2(X, Z))_k] = \mathbb{E}_{Z,\theta|X,\alpha,\beta} [\log \theta_k].$$

On the one hand, we have that:

$$\begin{aligned} p(Z, \theta | X, \beta, \alpha) &= p(Z | \theta, X, \beta, \alpha) p(\theta | X, \beta, \alpha) \\ &= C(\alpha) \prod_{k=1}^K \theta_k^{(\sum_{n=1}^{N_X} z_{nk}) + \alpha_k - 1}, \end{aligned}$$

with $C(\alpha) = \frac{\Gamma(\sum_{i=1}^K \alpha_i)}{\prod_{i=1}^K \Gamma(\alpha_i)}$. On the other hand:

$$p(Z, \theta | X, \beta, \alpha) \propto p(\theta | Z, \alpha) p(Z | X, \beta).$$

We deduce from the two identity that:

$$\begin{aligned}
p(\theta|Z, \alpha) &\propto \prod_{k=1}^K \theta_k^{\left(\sum_{n=1}^{N_X} z_{nk}\right) + \alpha_k - 1} \\
\Rightarrow \theta|Z, \alpha &\sim \text{Dirichlet}\left(\alpha + \sum_{n=1}^{N_X} z_n\right).
\end{aligned}$$

Finally, the expectation is:

$$\begin{aligned}
&\mathbb{E}_{Z, \theta|X, \alpha, \beta}[(S^2(X, Z))_k] = \mathbb{E}_{Z, \theta|X, \alpha, \beta}[\log \theta_k] \\
&= \mathbb{E}_{Z|X, \beta, \alpha} \left[\mathbb{E}_{\theta|Z, \alpha} [\log \theta_k] \right] \\
&= \mathbb{E}_{Z|X, \beta, \alpha} \left[\Psi \left([\alpha(s)]_k + \sum_{n=1}^{N_X} z_{nk} \right) \right] \\
&\quad - \Psi \left(\sum_{i=1}^K [\alpha(s)]_i + N_X \right),
\end{aligned}$$

as the prior on $\theta|Z, \alpha$ is Dirichlet $\left(\alpha + \sum_{n=1}^{N_X} z_n\right)$. We use the values of z sampled with Algorithm 2 to estimate this expectation. More precisely, keeping notations of Algorithm 2:

$$\begin{aligned}
&\mathbb{E}_{Z|X, \beta, \alpha} \left[\Psi \left([\alpha(s)]_k + \sum_{n=1}^{N_X} z_{nk} \right) \right] \\
&\approx \frac{1}{P} \sum_{t=1}^P \Psi \left([\alpha(s)]_k + \sum_{n=1}^{N_X} z_{nk}^t \right).
\end{aligned}$$

B Gibbs/variational online EM analysis

B.1 Effect of inference boosting on G-OEM and V-OEM

The effect of the inference boost as described in Section 3.3 on G-OEM and V-OEM with synthetic and IMDB datasets is presented in Figure 4 and in Figure 5. It leads to a minor improvement for G-OEM++ and a significant one for V-OEM++.

B.2 Step-sizes and averaging

We apply G-OEM with different stepsizes $\gamma_i = \frac{1}{i^\kappa}$. Note that because we average sufficient statistics, there is no needed proportionality constants. We first compare the performance of the last iterate η_N (without averaging) and the average of the iterates $\bar{\eta}_N = \frac{1}{N} \sum_{i=0}^N \eta_i$ (with averaging) for different values of κ .

Results are presented in Figure 6 on the synthetic data and in Figure 7 on the IMDB dataset. For $\kappa \in [0, \frac{1}{2}[$, averaging improves the performance while for $\kappa \in]\frac{1}{2}, 1]$, averaging deteriorates the performance. For $\kappa = \frac{1}{2}$, averaging is only slightly beneficial on IMDB dataset. For constant stepsizes $\kappa = 0$ the averaging improves significantly the performance, as the iterates do not converge and tend to oscillate around a local optimum (Bach & Moulines, 2013). We can expect the same effect for $\kappa \in [0, \frac{1}{2}[$ as the function $n \mapsto \frac{1}{n^\kappa}$ decreases slowly for such values of κ . For $\kappa \in]\frac{1}{2}, 1]$, the stochastic gradient ascent scheme is guaranteed to converge to a local optimum (Bottou, 1998). The averaging then deteriorates the performance as it incorporates the first iterates, which gets the last iterate away from local optimum. However, the stepsize $1/i$ ($\kappa = 1$) is not competitive. The performance with $\kappa = 0.75$ is only slightly better on IMDB dataset. The setting $\kappa = \frac{1}{2}$ represents a good balance between first and last iterates. For this step-size, performances with or without averaging are similar but results without averaging seem to be more stable, hence our choice for all our other simulations.

C Evolution of the ELBO

Figure 8 presents the evolution of the ELBO for online LDA (OLDA) and SVB on different test sets. We compute the ELBO on test documents as described by Hoffman et al. (2010). This plot helps us to observe that even if the ELBO reaches a local maximum (i.e., it stabilizes), the quality of the model in terms of perplexity is not controllable. We can also see in Figure 8 that the ELBO is much better optimized with $K = 128$ than with other values of K for both SVB and OLDA, that is, as expected, latent variables of higher dimensionality lead to better fits for the cost function which is optimized. However, for several datasets the performance in terms of perplexity is better with low values of K ($K = 8$ or $K = 16$) than with high dimensional variables ($K = 64$ or $K = 128$).

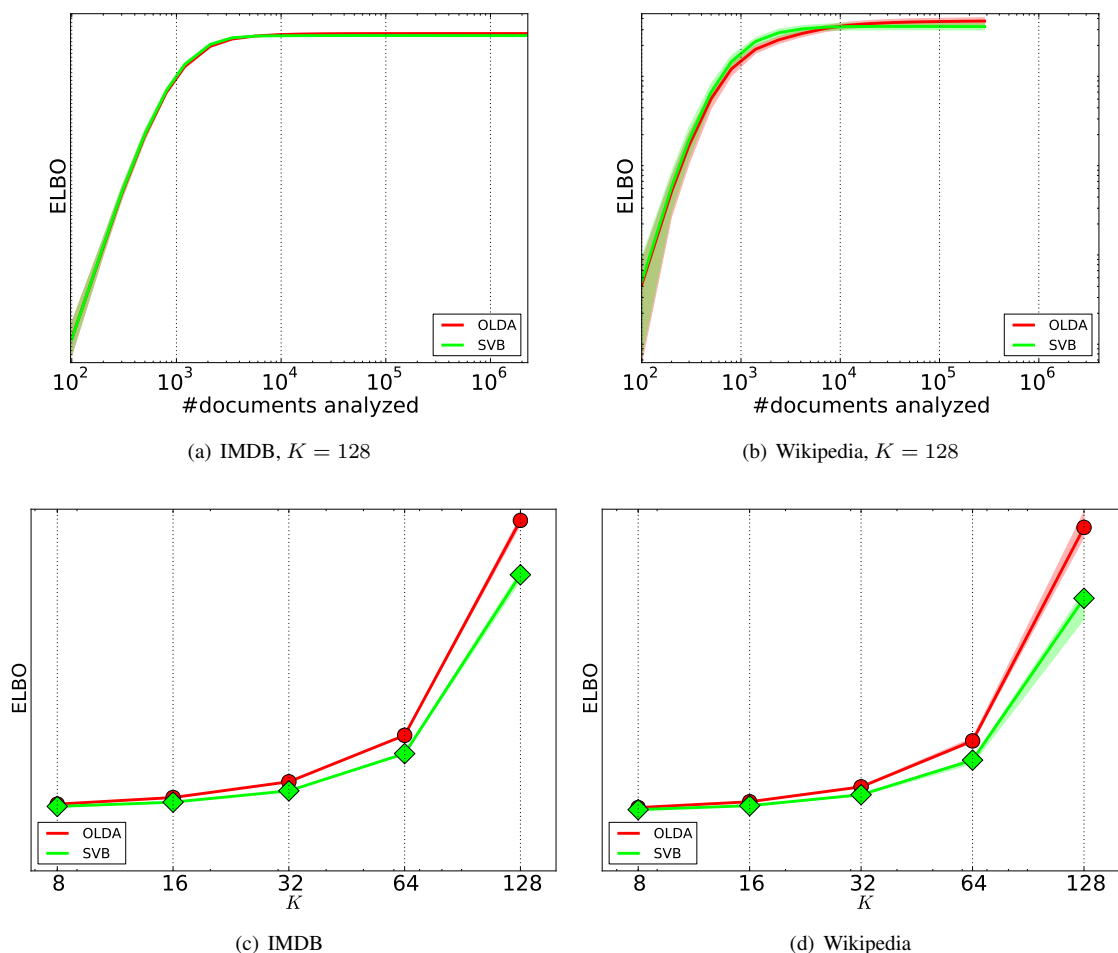


Figure 8: Evidence Lower BOUND (ELBO) computed on different test sets. Top: ELBO through iterations, with 4 passes over each dataset and 200 internal iterations. Bottom: ELBO as a function of the number of topics K , with 20 internal iterations and 1 pass over each dataset. Best seen in color.

In order to check if more internal iterations could help variational Bayesian methods, we present in Table 2 the values of perplexity reached by OLDA when running 4 passes over each dataset with $P = 200$ internal iterations 1 pass over each dataset with $P = 20$ internal iterations. We observe that the ELBO converges

Table 2: Comparison of log-perplexity levels reached with OLDA and SVB on IMDB dataset.

	$P = 200, 4 \text{ PASSES}$	$P = 20, 1 \text{ PASS}$
OLDA	682.6 ± 3.7	681.9 ± 3.9
SVB	683.8 ± 3.8	684.5 ± 3.8

quickly to a local optimum and doing ten times more internal iterations does not change significantly the final performance.

D Performance with different K , with error bars

The performance of the presented methods for different values of K on the different datasets is presented in Figure 9. We plot the median from the 11 experiments as a line—solid or dashed—and a shaded region between the third and the seventh decile.

E Performance through iterations, with error bars

The performance through iterations of the presented methods on the different datasets is presented in Figure 10. We plot the median from the 11 experiments as a line—solid or dashed—and a shaded region between the third and the seventh decile.

F Updates in α

In this section we compare the different types of updates for α . Figure 11 presents results obtained on synthetic dataset for fixed point iteration algorithm (Minka, 2000) and by putting a gamma prior on α (Sato et al., 2010). We observe that the fixed point method leads to better performance for G-OEM and G-OEM++. For V-OEM, the gamma updates better perform for $\kappa = \frac{1}{2}$. The performances of the gamma updates and the fixed point method are very similar for V-OEM++. Note that the algorithm V-OEM++ with $\kappa = 1$ and gamma updates on α is exactly equivalent to SPLDA (Sato et al., 2010). The performance of this method can be improved by setting $\kappa = \frac{1}{2}$ with any update on α .

We also observe that fixing α to α_{true} that generated the data does not necessarily lead to better performance.

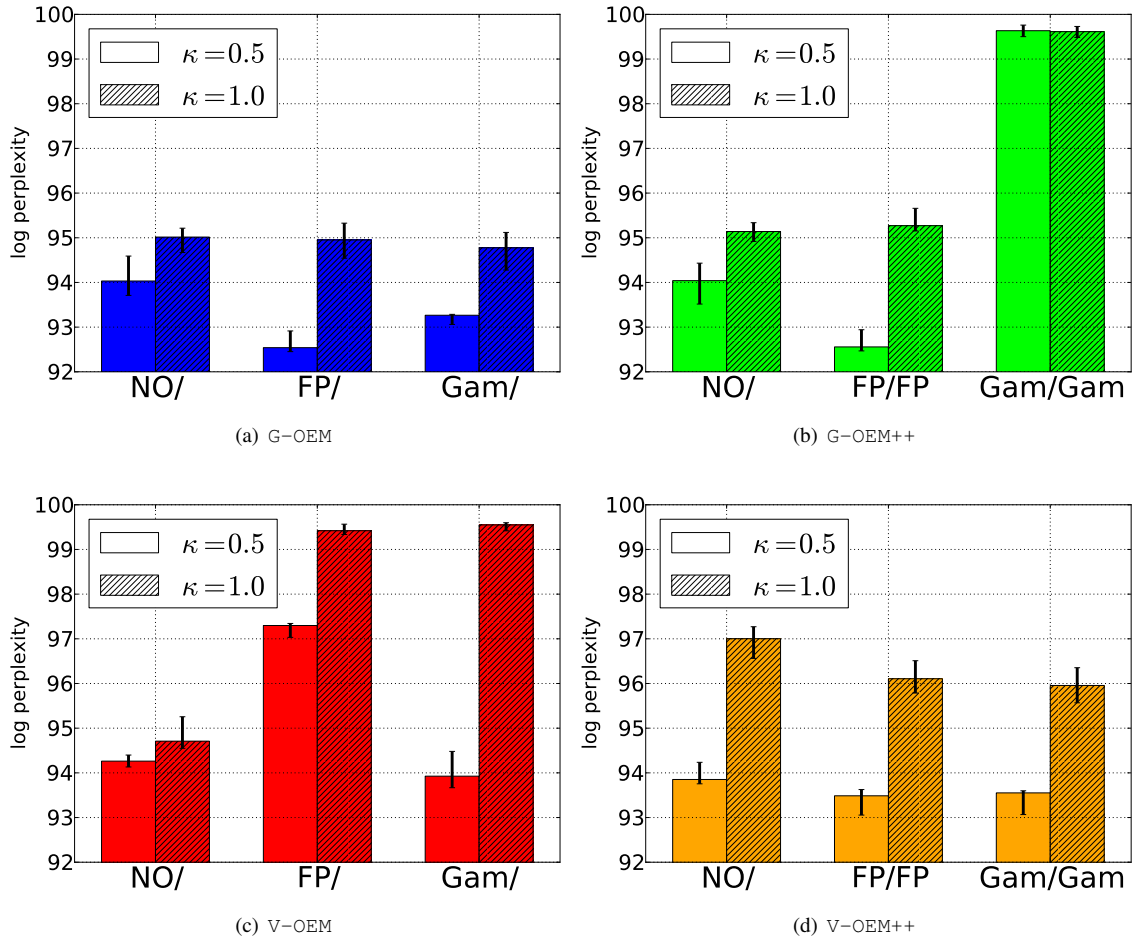


Figure 11: Dataset: Synthetic, $K = 10$. Perplexity on different test sets for different types of updates for α ; for boosted methods, we use the same inference for α for local and global updates. NO: α is fixed and set to α_{true} that generated the data; FP: fixed point iteration; Gam: gamma prior on α (Sato et al., 2010). Best seen in color.

G Empirical analysis

In this section we provide a qualitative empirical analysis on the topics extracted with the different methods. Examples of eight topics extracted with G-OEM and OLDA on the IMDB dataset of movie reviews are presented in Table 3.

We first compute the KL divergence between the $K = 128$ topics extracted with G-OEM and the $K = 128$ topics extracted with OLDA. We run the Hungarian algorithm on the resulting distance matrix to assign each topic extracted with G-OEM to a single topic of OLDA. We choose manually eight topics extracted with G-OEM that are representative of the usual behavior, and display the eight corresponding topics of OLDA assigned with the above method.

We observe that the topics extracted with G-OEM are more consistent than topics extracted with OLDA: topics of G-OEM precisely describe only one aspect of the reviews while the topics of OLDA tend to mix several aspects in each topic. For instance, the words of topic 1 extracted with G-OEM are related to *horror* movies. The words of the corresponding topic extracted with OLDA mix *horror* movies — e.g., *horror*, *scary*

— and *ghost* movies — e.g., *ghost*, *haunt*. In this OLDA topic 1, we can also observe less relevant words, like *effective*, *mysterious*, which are not directly linked with *horror* and *ghost* vocabularies. We can make the same remarks with topic 2 and topic 3, respectively related to *comedy* movies and *romantic comedy* movies. In topic 2 extracted with G-OEM, the least related words to *comedy* are names of characters/actors — i.e., *steve* and *seth* — while the words not related to *comedy* in topic 25 of OLDA are more general, belonging to a different lexical field — e.g., *sport*, *site*, *progress*, *brave*, *definition*. In topic 3 of G-OEM, all the presented words are related to *romantic comedy* while in topic 3 of OLDA, the words *old*, *hard* and *review* are not related to this genre.

We also observe that G-OEM extracts strongly “qualitative” topics — topic 4 and topic 5 — which is not done with OLDA. Indeed, it is difficult to group the top words of topic 4 or topic 5 of OLDA in the same lexical field. Except *dialogue* and *suppose*, all the top words of topic 4 of G-OEM are negative words. These two words may appear in a lot of negative sentences, leading to a high weight in this topic. In topic 5 of G-OEM, the words *absolutely* and *visual* are non strictly positive words while the thirteen other words in this topic convey a positive opinion. The word *absolutely* is an adverb much more employed in positive sentences than negative or neutral sentences, which can explain its high weight in topic 5.

The topic 6 of both G-OEM and OLDA can be considered as a “junk” topic, as for both method, most of its top words are contractions of modal verbs or frequent words — e.g., *didn't*, *isn't*, *wait*, *bad*. The contractions are not filtered when removing the stop words as they are not included in the list of words removed¹.

For both G-OEM and OLDA, the top words of topic 7 are general words about movies. These words are usually employed to describe a movie as a whole — e.g., *narrative*, *filmmaker*.

Finally, the top words of topic 8 of G-OEM are related to the situation of the scenes. We could not find such topic in the other presented methods and we can see that the top words of topic 8 of OLDA — supposedly close to topic 8 of G-OEM — are related to *family* movies. Each word of topic 8 of G-OEM — except *group* and *beautiful* — are related to a *spatial location*, and may help answer the question “*where does the scene take place?*”.

¹See NLTK toolbox (Bird et al., 2009) for the exhaustive list of stop words.

Table 3: Comparison of topics extracted on IMDB dataset, $K = 128$. The 15 top words of eight different topics extracted with G-OEM and OLDA are presented.

		G-OEM								OLDA							
#	TOPIC 1	TOPIC 2	TOPIC 3	TOPIC 4	TOPIC 5	TOPIC 6	TOPIC 7	TOPIC 8	#	TOPIC 1	TOPIC 2	TOPIC 3	TOPIC 4	TOPIC 5	TOPIC 6	TOPIC 7	TOPIC 8
1	VIOLENCE	COMEDY	ROMANTIC	BAD	BRILLIANT	DIDN'T	NARRATIVE	TOWN	1	HORROR	HILARIOUS	LOVE	BAD	GREAT	WRONG	VISUAL	YOUNG
2	VIOLENT	FUNNY	COMEDY	WORST	PERFECT	I'VE	ULTIMATELY	LOCAL	2	NIGHT	ROMANCE	ENJOY	ACTION	BEST	DIDN'T	FOCUS	FAMILY
3	DISTURBING	LAUGH	LOVE	WASTE	BEAUTIFUL	WASN'T	CINEMATIC	MOUNTAIN	3	DEAD	CLEVER	PRETTY	INTERESTING	STAR	I'VE	REALITY	CHILD
4	BRUTAL	JOKE	ROMANCE	BORING	MASTERPIECE	ISN'T	APPROACH	VILLAGE	4	TWIST	SMART	OLD	ORIGINAL	LONG	WAIT	DIFFICULT	FATHER
5	MURDER	HILARIOUS	FUNNY	AWFUL	AMAZING	WE'RE	PROTAGONIST	LOCATION	5	SCARY	INTRIGUING	FUNNY	FAR	JOHN	CATCH	FILMMAKER	SON
6	GRAPHIC	COMIC	CHARMING	POOR	SUPERB	I'LL	SEEMINGLY	ROAD	6	EFFECTIVE	COMEDIC	COMEDY	SPECIAL	EXCELLENT	EXACTLY	IMAGE	AGE
7	KILLER	COMEDIC	CHEMISTRY	DIALOGUE	STUNNING	COULDN'T	NATURE	JOURNEY	7	MYSTERIOUS	FUNNIEST	FUN	FIGHT	CLASSIC	WASN'T	NARRATIVE	WHOSE
8	TORTURE	STEVE	SWEET	WORSE	WONDERFUL	WOULDN'T	NE	GROUP	8	BLOODY	PROGRESS	HARD	HERO	BEAUTIFUL	HUGE	INTELLIGENT	TALE
9	VICTIM	AMUSING	ENJOY	DULL	ABSOLUTELY	PRETTY	FILMMAKER	TRAVEL	9	GHOST	SPORT	PERFECT	ENTIRE	DRAMA	I'LL	ACCEPT	DISCOVER
10	RAPE	FUN	HEART	FAIL	BEST	BAD	WHOSE	COUNTRY	10	HAUNT	SITE	LAUGH	HALF	WONDERFUL	CHOICE	IMPRESSION	EASILY
11	KILL	GAG	NICE	MESS	INCREDIBLE	HAVEN'T	CRAFT	LANDSCAPE	11	FEAR	BRAVE	ENTERTAINING	SAVE	MICHAEL	ETC	EXTREME	DREAM
12	HORROR	SETH	CHARM	RIDICULOUS	BRILLIANTLY	GUESS	CONTEMPORARY	LAND	12	EVIL	DREADFUL	WORTH	DIALOGUE	HEART	SERIOUSLY	CENTRAL	INTRODUCE
13	BLOODY	FUNNIEST	GREAT	SUPPOSE	BEAUTIFULLY	AREN'T	MANNER	BEAUTIFUL	13	NIGHTMARE	SHOULDER	NICE	FULL	FORGET	NOTICE	MAINTAIN	MARRY
14	REVENGE	CAMEO	FUN	TERRIBLE	VISUAL	ENJOY	SERVE	AREA	14	GORY	GIMMICK	FAVORITE	VIOLENCE	ROBERT	RIDICULOUS	DENY	RAISE
15	BLOOD	SITUATION	WONDERFUL	UNFORTUNATELY	PERFECTLY	REVIEW	MATERIAL	TRIP	15	MASK	DEFINITION	REVIEW	EXAMPLE	EARLY	ANSWER	NAIL	RULE

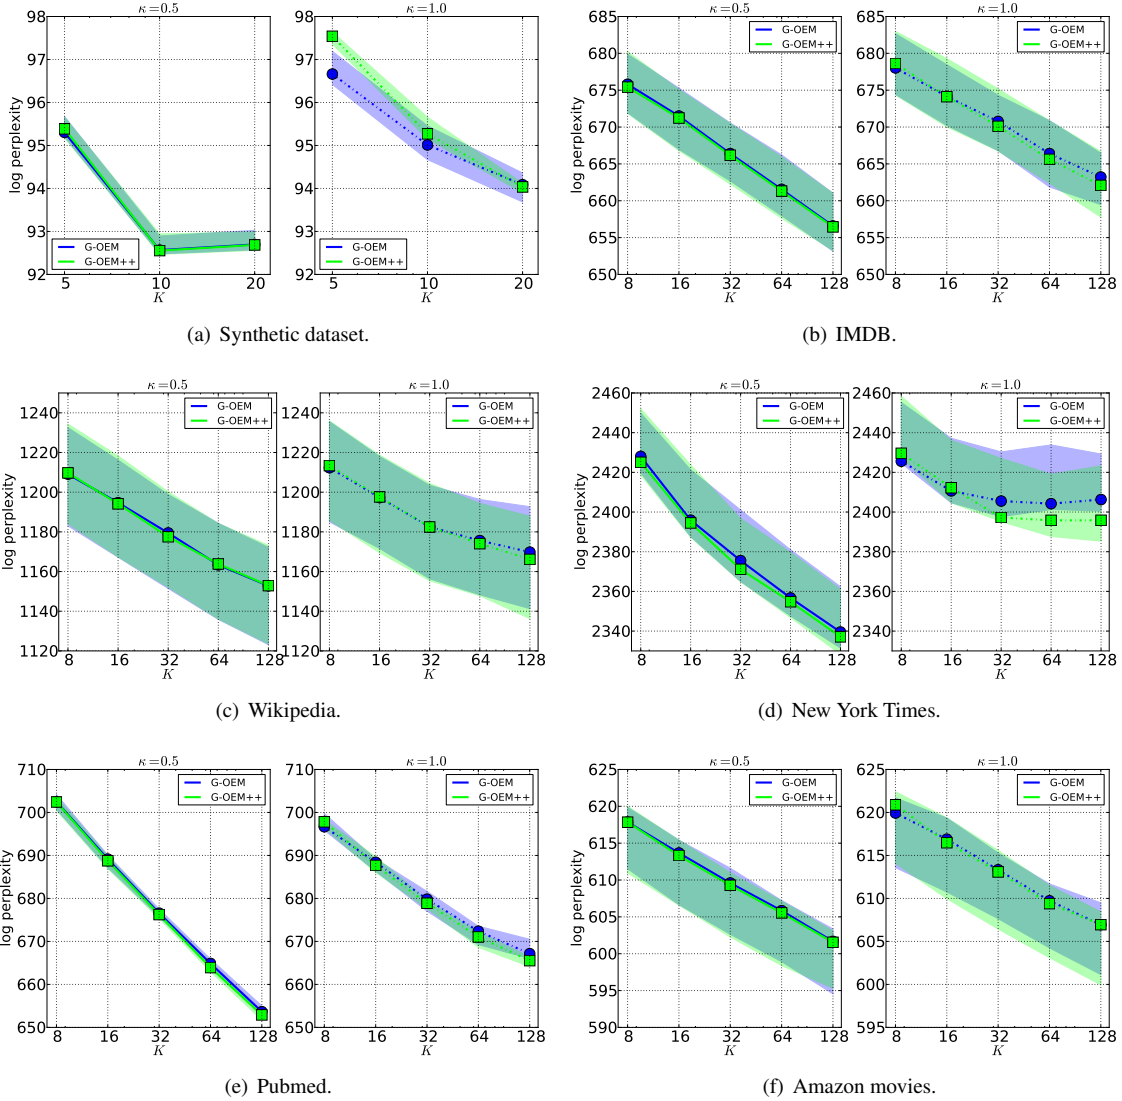


Figure 4: G-OEM. Perplexity on different test sets as a function of the number of topics K for regular EM and boosted EM (++). We observe that for almost all datasets, there is no significant improvement when boosting the inference. Our interpretation is that each Gibbs sample is noisy and does not provide a stable boost. Best seen in color.

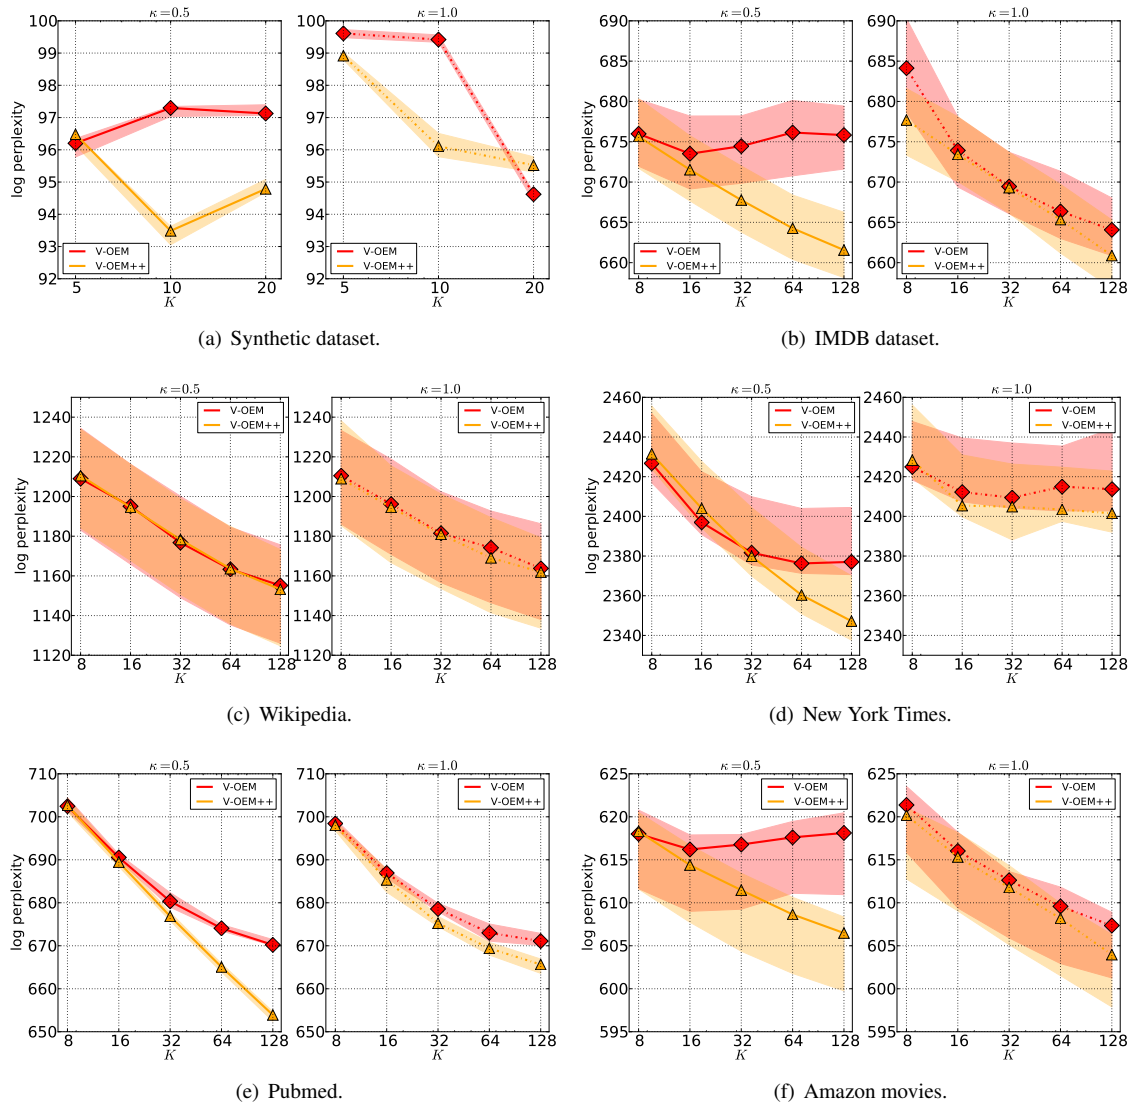


Figure 5: V-OEM. Perplexity on different test sets as a function of the number of topics K for regular EM and boosted EM (++). We observe that boosting inference improves significantly the results on all the datasets excepted on Wikipedia where V-OEM and V-OEM++ have similar performances. The variational estimation of the posterior is finer and finer through iterations. When updating the parameters at each iteration of the posterior estimation, the inference is indeed boosted. Best seen in color.

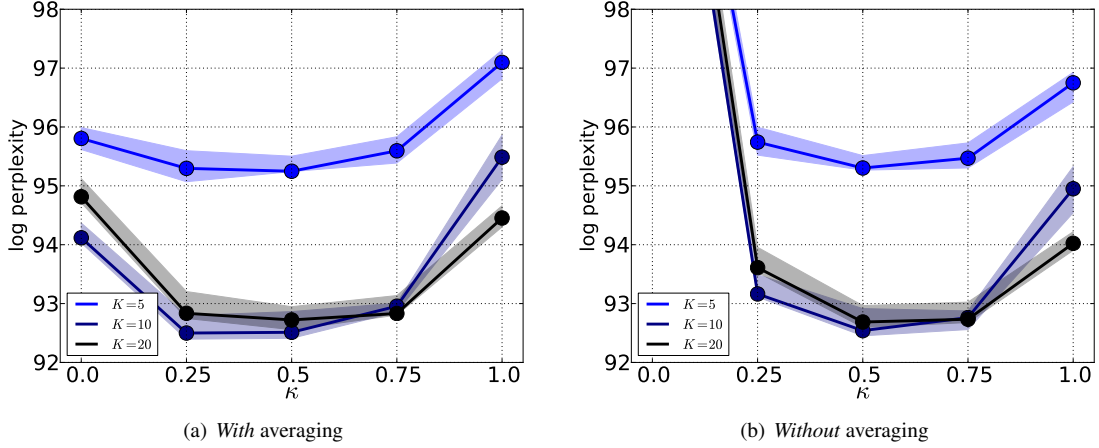


Figure 6: Dataset: synthetic. Perplexity on different test sets as a function of the exponent κ —the corresponding stepsize is $\gamma_i = \frac{1}{i^\kappa}$ —for G-OEM with averaging (left) and without averaging (right). The number of topics inferred K goes from 5 (the lightest) to 20 (the darkest). Best seen in color.

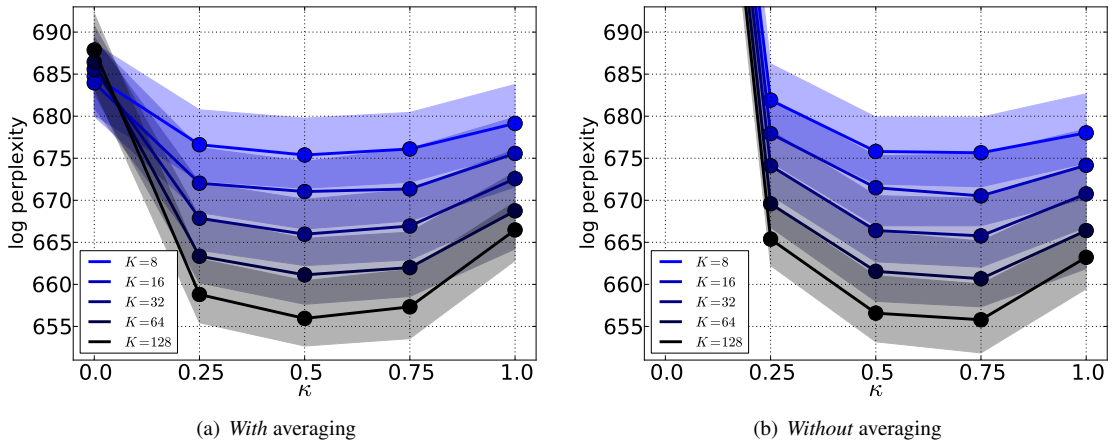
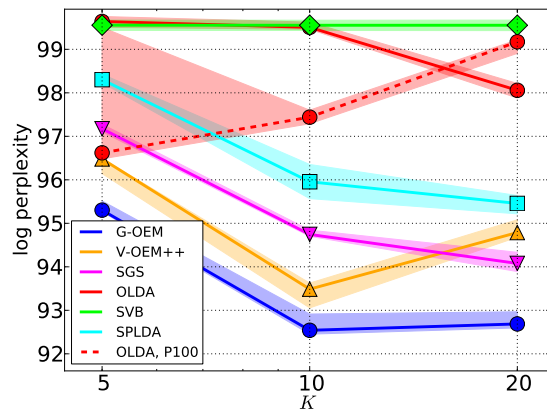
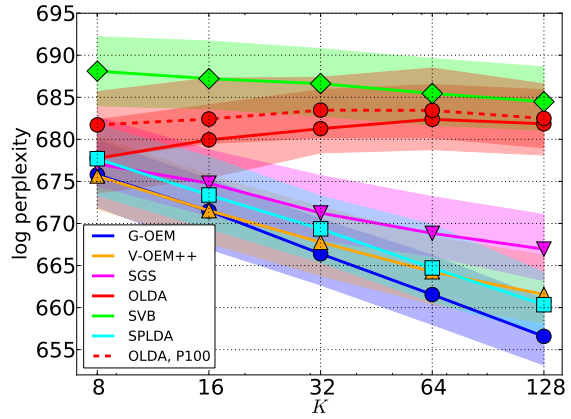


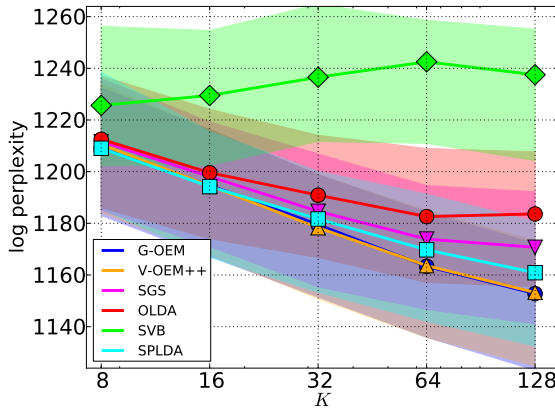
Figure 7: Dataset: IMDB. Perplexity on different test sets as a function of the exponent κ —the corresponding stepsize is $\gamma_i = \frac{1}{i^\kappa}$ —for G-OEM with averaging (left) and without averaging (right). The number of topics inferred K goes from 8 (the lightest) to 128 (the darkest). Best seen in color.



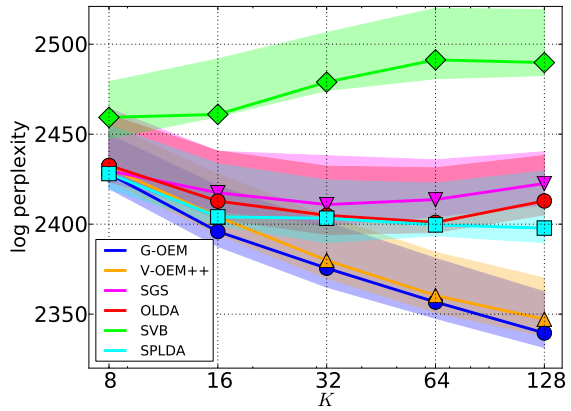
(a) Dataset: Synthetic



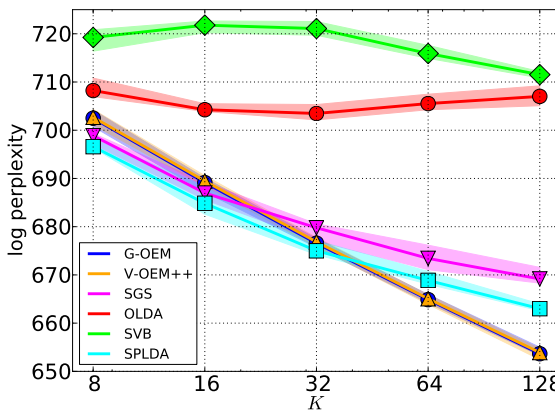
(b) Dataset: IMDB



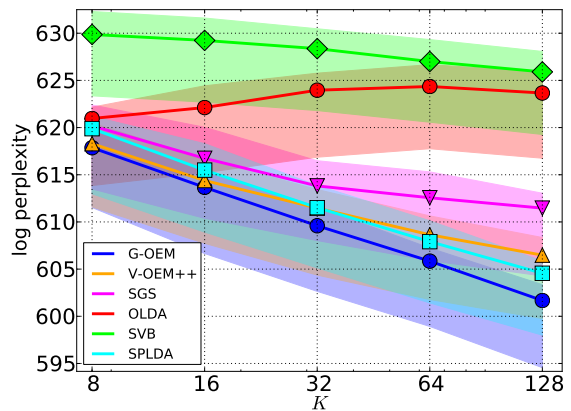
(c) Dataset: Wikipedia



(d) Dataset: New York Times

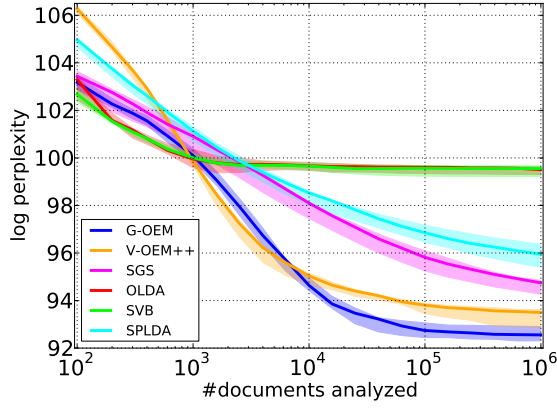


(e) Dataset: Pubmed

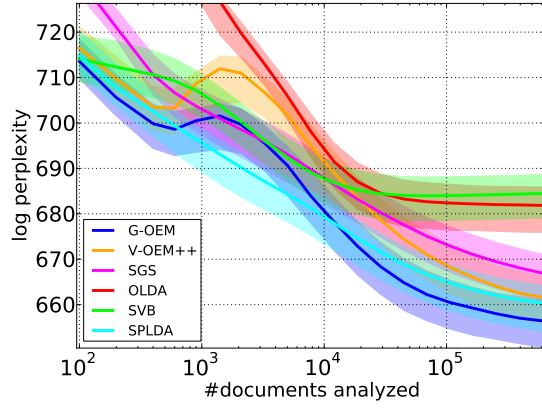


(f) Dataset: Amazon movies

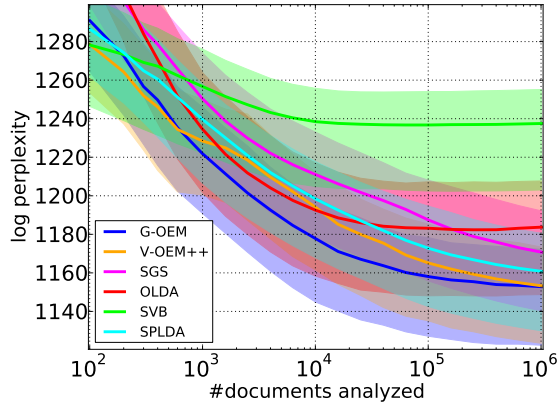
Figure 9: Perplexity on different test sets as a function of K , the number of topics inferred. Same as Figure 1, but with error bars. Best seen in colors.



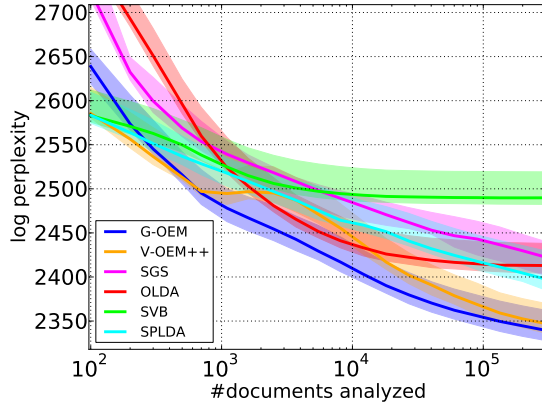
(a) Synthetic, $K = 10$



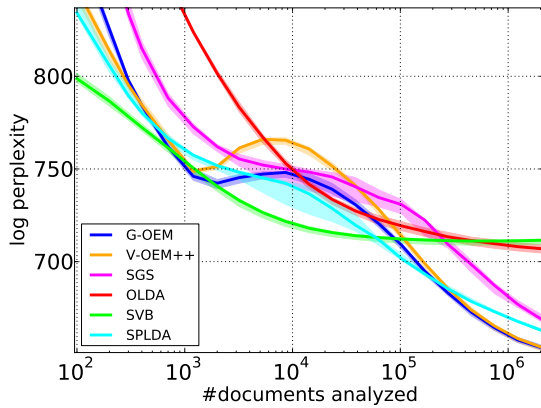
(b) IMDB, $K = 128$



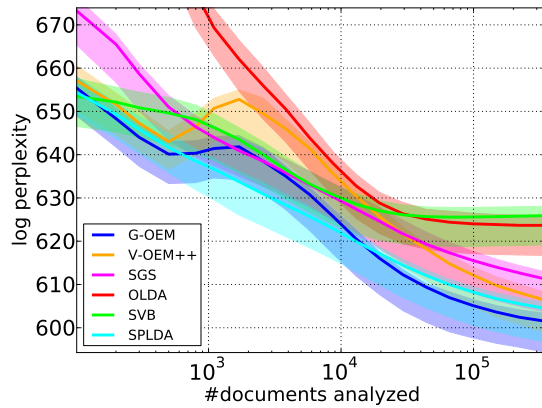
(c) Wikipedia, $K = 128$



(d) New York Times, $K = 128$



(e) Pubmed, $K = 128$



(f) Amazon movies, $K = 128$

Figure 10: Perplexity through iterations on different test sets with the presented methods. Same as Figure 3, but with error bars. Best seen in colors.