



AALBORG UNIVERSITY
DENMARK

Aalborg Universitet

Stochastic Discriminative EM

Masegosa, Andres

Published in:

Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence Conference (UAI)

Publication date:

2014

Document Version

Early version, also known as pre-print

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Masegosa, A. (2014). Stochastic Discriminative EM. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence Conference (UAI)* (pp. 573-582). AUAI Press. <http://arxiv.org/abs/1410.1784>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Stochastic Discriminative EM

Andrés R. Masegosa^{†‡}

[†] Dept. of Computer and Information Science
Norwegian University of Science and Technology
Trondheim, Norway

[‡] Dept. of Computer Science and A. I.
University of Granada
Granada, Spain

Abstract

Stochastic discriminative EM (sdEM) is an online-EM-type algorithm for discriminative training of probabilistic generative models belonging to the exponential family. In this work, we introduce and justify this algorithm as a stochastic *natural gradient* descent method, i.e. a method which accounts for the information geometry in the parameter space of the statistical model. We show how this learning algorithm can be used to train probabilistic generative models by minimizing different discriminative loss functions, such as the negative conditional log-likelihood and the Hinge loss. The resulting models trained by sdEM are always generative (i.e. they define a joint probability distribution) and, in consequence, allows to deal with missing data and latent variables in a principled way either when being learned or when making predictions. The performance of this method is illustrated by several text classification problems for which a multinomial naive Bayes and a latent Dirichlet allocation based classifier are learned using different discriminative loss functions.

1 INTRODUCTION

Online learning methods based on stochastic approximation theory [21] have been a promising research direction to tackle the learning problems of the so-called Big Data era [1, 10, 12]. Stochastic gradient descent (SGD) is probably the best known example of this kind of techniques, used to solve a wide range of learning problems [9]. This algorithm and other versions [29] are usually employed to train discriminative models such as logistic regression or SVM [10].

There also are some successful examples of the use of SGD for discriminative training of probabilistic generative models, as is the case of deep belief networks [19]. However,

this learning algorithm cannot be used directly for the discriminative training of general generative models. One of the main reasons is that statistical estimation or risk minimization problems of generative models involve the solution of an optimization problem with a large number of *normalization constraints* [26], i.e. those which guarantee that the optimized parameter set defines a valid probabilistic model. Although successful solutions to this problem have been proposed [16, 22, 26, 32], they are based on ad-hoc methods which cannot be easily extended to other statistical models, and hardly scale to large data sets.

Stochastic approximation theory [21] has also been used for maximum likelihood estimation (MLE) of probabilistic generative models with latent variables, as is the case of the online EM algorithm [13, 30]. This method provides efficient MLE estimation for a broad class of statistical models (i.e. exponential family models) by sequentially updating the so-called *expectation parameters*. The advantage of this approach is that the resulting iterative optimization algorithm is fairly simple and amenable, as it does not involve any normalization constraints.

In this paper we show that the derivation of Sato's online EM [30] can be extended for the discriminative learning of generative models by introducing a novel interpretation of this algorithm as a natural gradient algorithm [3]. The resulting algorithm, called stochastic discriminative EM (sdEM), is an online-EM-type algorithm that can train generative probabilistic models belonging to the exponential family using a wide range of discriminative loss functions, such as the negative conditional log-likelihood or the Hinge loss. In opposite to other discriminative learning approaches [26], models trained by sdEM can deal with missing data and latent variables in a principled way either when being learned or when making predictions, because at any moment they always define a joint probability distribution. sdEM could be used for learning using large scale data sets due to its stochastic approximation nature and, as we will show, because it allows to compute the *natural gradient* of the loss function with no extra cost [3]. Moreover, if allowed by the generative model and the discriminative loss

function, the presented algorithm could potentially be used interchangeably for classification or regression or any other prediction task. But in this initial work, sdEM is only experimentally evaluated in classification problems.

The rest of this paper is organized as follows. Section 2 provides the preliminaries for the description of the sdEM algorithm, which is detailed in Section 3. A brief experimental evaluation is given in Section 4, while Section 5 contains the main conclusions of this work.

2 PRELIMINARIES

2.1 MODEL AND ASSUMPTIONS

We consider generative statistical models for prediction tasks, where Y denotes the random variable (or the vector-value random variable) to be predicted, X denotes the predictive variables, and y^* denotes a prediction, which is made according to $y^* = \arg \max_y p(y, x|\theta)$.

Assumption 1. *The generative data model belongs to the exponential family with a natural (or canonical) parametrization*

$$p(y, x|\theta) \propto \exp(\langle s(y, x), \theta \rangle - A_l(\theta))$$

where θ is the so-called natural parameter which belongs to the so-called natural parameter space $\Theta \in \mathbb{R}^K$, $s(y, x)$ is the vector of sufficient statistics belonging to a convex set $\mathcal{S} \subseteq \mathbb{R}^K$, $\langle \cdot, \cdot \rangle$ denotes the dot product and A_l is the log partition function.

Assumption 2. *We are given a conjugate prior distribution $p(\theta|\alpha)$ of the generative data model*

$$p(\theta|\alpha) \propto \exp(\langle s(\theta), \alpha \rangle - A_g(\alpha))$$

where the sufficient statistics are $s(\theta) = (\theta, -A_l(\theta))$ and the hyperparameter α has two components $(\bar{\alpha}, \nu)$. ν is a positive scalar and $\bar{\alpha}$ is a vector also belonging to \mathcal{S} [6].

2.2 DUAL PARAMETERIZATION AND ASSUMPTIONS

The so-called *expectation parameter* $\mu \in \mathcal{S}$ can also be used to parameterize probability distributions of the exponential family. It is a dual set of the model parameter θ [2]. This *expectation parameter* μ is defined as the expected vector of sufficient statistics with respect to θ :

$$\begin{aligned} \mu &\triangleq E[s(y, x)|\theta] = \int s(y, x)p(y, x|\theta)dydx \\ &= \partial A_l(\theta)/\partial \theta \end{aligned} \quad (1)$$

The transformation between θ and μ is one-to-one: μ is a dual set of the model parameter θ [2]. Therefore, Equation (1) can be inverted as: $\theta = \theta(\mu)$. That is to say, for each $\theta \in \Theta$ we always have an associated $\mu \in \mathcal{S}$ and both parameterize the same probability distribution.

For obtaining the *natural parameter* θ associated to an *expectation parameter* μ , we need to make use of the negative of the entropy,

$$\begin{aligned} H(\mu) &\triangleq \int p(y, x|\theta(\mu)) \ln p(y, x|\theta(\mu)) dydx \\ &= \sup_{\theta \in \Theta} \langle \mu, \theta \rangle - A_l(\theta) \end{aligned} \quad (2)$$

Using the above function, the natural parameter θ can be explicitly expressed as

$$\theta = \theta(\mu) = \partial H(\mu)/\partial \mu \quad (3)$$

Equations (1), (2), (3) define the Legendre-Fenchel transform.

Another key requirement of our approach is that it should be possible to compute the transformation from μ to θ in closed form:

Assumption 3. *The transformation from the expectation parameter μ to the natural parameter θ , which can be expressed as*

$$\theta(\mu) = \arg \max_{\theta \in \Theta} \langle \mu, \theta \rangle - A_l(\theta) \quad (4)$$

is available in closed form.

The above equation is also known as the *maximum likelihood function*, because $\theta(\frac{1}{n} \sum_{i=1}^n s(y_i, x_i))$ gives the maximum likelihood estimation θ^* for a data set with n observations $\{(y_1, x_1), \dots, (y_n, x_n)\}$.

For later convenience, we show the following relations between the Fisher Information matrices $I(\theta)$ and $I(\mu)$ for the probability distributions $p(y, x|\theta)$ and $p(y, x|\theta(\mu))$, respectively [25]:

$$I(\theta) = \frac{\partial^2 A_l(\theta)}{\partial \theta \partial \theta} = \frac{\partial \mu}{\partial \theta} = I(\mu)^{-1} \quad (5)$$

$$I(\mu) = \frac{\partial^2 H(\mu)}{\partial \mu \partial \mu} = \frac{\partial \theta}{\partial \mu} = I(\theta)^{-1} \quad (6)$$

2.3 THE NATURAL GRADIENT

Let $\mathcal{W} = \{w \in \mathbb{R}^K\}$ be a parameter space on which the function $L(w)$ is defined. When \mathcal{W} is a Euclidean space with an orthonormal coordinate system, the negative gradient points in the direction of steepest descent. That is, the negative gradient $-\partial L(w)/\partial w$ points in the same direction as the solution to:

$$\arg \min_{dw} L(w + dw) \text{ subject to } \|dw\|^2 = \epsilon^2 \quad (7)$$

for sufficiently small ϵ , where $\|dw\|^2$ is the squared length of a small increment vector dw connecting w and $w + dw$. This justifies the use of the classical gradient descent method for finding the minimum of $L(w)$ by taking steps (of size ρ) in the direction of the negative gradient:

$$w_{t+1} = w_t - \rho \frac{\partial L(w_t)}{\partial w} \quad (8)$$

However, when \mathcal{W} is a Riemannian space [4], there are no orthonormal linear coordinates, and the squared length of vector dw is defined by the following equation,

$$\|dw\|^2 = \sum_{ij} g_{ij}(w) dw_i dw_j \quad (9)$$

where the $K \times K$ matrix $G = (g_{ij})$ is called the Riemannian metric tensor, and it generally depends on w . G reduces to the identity matrix in the case of the Euclidean space [4].

In a Riemannian space, the steepest descent direction is not anymore the traditional gradient. That is, $-\partial L(w)/\partial w$ is not the solution of Equation (7) when the squared length of the distance of dw is defined by Equation (9). Amari [3] shows that this solution can be computed by pre-multiplying the traditional gradient by the inverse of the Riemannian metric G^{-1} ,

Theorem 1. *The steepest descent direction or the natural gradient of $L(w)$ in a Riemannian space is given by*

$$-\frac{\tilde{\partial}L(w)}{\tilde{\partial}w} = -G^{-1}(w) \frac{\partial L(w)}{\partial w} \quad (10)$$

where $\tilde{\partial}L(w)/\tilde{\partial}w$ denotes the *natural gradient*.

As argued in [3], in statistical estimation problems we should use gradient descent methods which account for the natural gradient of the parameter space, as the parameter space of a statistical model (belonging to the exponential family or not) is a Riemannian space with the Fisher information matrix of the statistical model $I(w)$ as the tensor metric [2], and this is the only invariant metric that must be given to the statistical model [2].

2.4 SATO'S ONLINE EM ALGORITHM

Sato's online EM algorithm [30] is used for maximum likelihood estimation of missing data-type statistical models. The model defines a probability distribution over two random or vector-valued variables X and Z , and is assumed to belong to the exponential family:

$$p(z, x|\theta) \propto \exp(\langle s(z, x), \theta \rangle - A_t(\theta))$$

where (z, x) denotes a so-called *complete data* event. The key aspect is that we can only observe x , since z is an unobservable event. In consequence, the loss function $\ell(x, \theta)$ ¹ is defined by marginalizing z : $\ell(x, \theta) = -\ln \int p(z, x) dz$.

The online setting assumes the observation of a non-finite data sequence $\{(x_t)\}_{t \geq 0}$ independently drawn according to the unknown data distribution π . The objective function that EM seeks to minimize is given by the following expectation: $L(\theta) = E[\ell(x, \theta)|\pi]$.

¹We derive this algorithm in terms of minimization of a loss function to highlight its connection with sdEM.

Sato [30] derived the stochastic updating equation of online EM by relying on the free energy formulation, or lower bound maximization, of the EM algorithm [24] and on a discounting averaging method. Using our own notation, this updating equation is expressed as follows,

$$\begin{aligned} \mu_{t+1} &= (1 - \rho_t)\mu_t + \rho_t E_z[s(z, x_t|\theta(\mu_t))] \\ &= \mu_t + \rho_t (E_z[s(z, x_t|\theta(\mu_t))] - \mu_t) \\ &= \mu_t + \rho_t \frac{\partial \ell(x_t, \theta(\mu_t))}{\partial \theta} \end{aligned} \quad (11)$$

where $E_z[s(z, x_t|\theta(\mu_t))]$ denotes the expected sufficient statistics, $E_z[s(z, x_t|\theta(\mu_t))] = \int s(z, x_t) p(z|x_t, \theta(\mu_t)) dz$.

He proved the convergence of the above iteration method by casting it as a second order stochastic gradient descent using the following equality,

$$\frac{\partial \ell(x, \theta)}{\partial \theta} = \frac{\partial \mu}{\partial \theta} \frac{\partial \ell(x, \theta(\mu))}{\partial \mu} = I(\mu)^{-1} \frac{\partial \ell(x, \theta(\mu))}{\partial \mu} \quad (12)$$

This equality is obtained by firstly applying the chain rule, followed by the equality shown in Equation (5). It shows that online EM is equivalent to a stochastic gradient descent with $I(\mu_t)^{-1}$ as coefficient matrices [9].

Sato noted that that the third term of the equality in Equation (12) resembles a natural gradient (see Theorem 1), but he did not explore the connection. But the key insights of the above derivation, which were not noted by Sato, is that Equation (12) is also valid for other loss functions different from the marginal log-likelihood; and that the convergence of Equation (11) does not depend on the formulation of the EM as a "lower bound maximization" method [24].

3 STOCHASTIC DISCRIMINATIVE EM

3.1 THE sdEM ALGORITHM

We consider the following supervised learning setup. Let us assume that we are given a data set D with n observations $\{(y_1, x_1), \dots, (y_n, x_n)\}$. We are also given a *discriminative loss function*² $\ell(y_i, x_i, \theta)$. For example, it could be the negative conditional log-likelihood (NCLL) $\ell(y_i, x_i, \theta) = -\ln p(y_i, x_i|\theta) + \ln \int p(y, x_i|\theta) dy = -\ln p(y_i|x_i, \theta)$. Our learning problem consists in minimizing the following objective function:

$$\begin{aligned} L(\theta) &= \sum_{i=1}^n \ell(y_i, x_i, \theta) - \ln p(\theta|\alpha) \\ &= E[\ell(y, x, \theta)|\pi] - \frac{1}{n} \ln p(\theta|\alpha) \end{aligned} \quad (13)$$

where π is now the empirical distribution of D and $E[\ell(y, x, \theta)|\pi]$ the empirical risk. Although the above

²The loss function is assumed to satisfy the mild conditions given in [9]. E.g., it can be a non-smooth function, such as the Hinge Loss.

loss function is not standard in the machine learning literature, we note that when ℓ is the negative log-likelihood (NLL), we get the classic *maximum a posterior estimation*. This objective function can be seen as an extension of this framework.

sdEM is presented as a generalization of Sato’s online EM algorithm for finding the minimum of an objective function in the form of Equation (13) (i.e. the solution to our learning problem). The stochastic updating equation of sdEM can be expressed as follows,

$$\mu_{t+1} = \mu_t - \rho_t I(\mu_t)^{-1} \frac{\partial \bar{\ell}(y_t, x_t, \theta(\mu_t))}{\partial \mu} \quad (14)$$

where (y_t, x_t) denotes the t -th sample, randomly generated from π , and the function $\bar{\ell}$ has the following expression: $\bar{\ell}(y_t, x_t, \theta(\mu_t)) = \ell((y_t, x_t, \theta(\mu_t))) + 1/n \ln p(\theta(\mu_t))$. We note that this loss function satisfies the following equality, which is the base for a stochastic approximation method [21], $E[\bar{\ell}(y_t, x_t, \theta(\mu)) | \pi] = L(\theta(\mu))$.

Similarly to Amari’s natural gradient algorithm [3], the main problem of sdEM formulated as in Equation (14) is the computation of the inverse of the Fisher information matrix at each step, which becomes even prohibitive for large models. The following result shows that this can be circumvented when we deal with distributions of the exponential family:

Theorem 2. *In the exponential family, the natural gradient of a loss function with respect to the expectation parameters equals the gradient of the loss function with respect to the natural parameters,*

$$I(\mu)^{-1} \frac{\partial \bar{\ell}(y, x, \theta(\mu))}{\partial \mu} = \frac{\partial \bar{\ell}(y, x, \theta)}{\partial \theta}$$

Sketch of the proof. We firstly need to prove that $I(\mu)$ is a valid Riemannian tensor metric and, hence, the expectation parameter space has a Riemannian structure defined by the metric $I(\mu)$ and the definition of the natural gradient makes sense. This can be proved by the invariant property of the Fisher information metric to one-to-one reparameterizations or, equivalently, transformations in the system of coordinates [2, 4]. $I(\mu)$ is a Riemannian metric because it is the Fisher information matrix of the reparameterized model $p(y, x | \theta(\mu))$, and the reparameterization is one-to-one, as commented in Section 2.2.

The equality stated in the theorem follows directly from Sato’s derivation of the online EM algorithm (Equation (12)). This derivation shows that we can avoid the computation of $I(\mu)^{-1}$ by using the natural parameters instead of the expectation parameters and the function $\theta(\mu)$. \square

Theorem 1 simplifies the sdEM’s updating equation to,

$$\mu_{t+1} = \mu_t - \rho_t \frac{\partial \bar{\ell}(y_t, x_t, \theta(\mu_t))}{\partial \theta} \quad (15)$$

sdEM can be interpreted as a stochastic gradient descent algorithm iterating over the *expectation parameters* and guided by the natural gradient in this Riemannian space.

Algorithm 1 Stochastic Discriminative EM (sdEM)

Require: D is randomly shuffled.

- 1: $\mu_0 = \bar{\alpha}$; (initialize according to the prior)
 - 2: $\theta_0 = \theta(\mu_0)$;
 - 3: $t = 0$;
 - 4: **repeat**
 - 5: **for** $i = 1, \dots, n$ **do**
 - 6: **E-Step:** $\mu_{t+1} = \mu_t - \frac{1}{(1+\lambda t)} \frac{\partial \bar{\ell}(y_i, x_i, \theta_t)}{\partial \theta}$;
 - 7: **Check-Step:** $\mu_{t+1} = \text{Check}(\mu_{t+1}, \mathcal{S})$;
 - 8: **M-Step:** $\theta_{t+1} = \theta(\mu_{t+1})$;
 - 9: $t = t + 1$;
 - 10: **end for**
 - 11: **until** convergence
 - 12: **return** $\theta(\mu_t)$;
-

An alternative proof to Theorem 2 based on more recent results on information geometry has been recently given in [27]. The results of that work indicate that sdEM could also be interpreted as a mirror descent algorithm with a Bregman divergence as a proximity measure. It is beyond the scope of the paper to explore this relevant connection.

3.2 CONVERGENCE OF sdEM

In this section we do not attempt to give a formal proof of the convergence of sdEM, since very careful technical arguments would be needed for this purpose [9]. We simply go through the main elements that define the convergence of sdEM as an stochastic approximation method [21].

According to Equation (14), sdEM can be seen as a stochastic gradient descent method with the inverse of the Fisher information matrix $I(\mu)^{-1}$ as a coefficient matrix [9]. As we are dealing with exponential families, these matrices are always positive-definite. Moreover, if the gradient $\partial \bar{\ell}(y, x, \theta) / \partial \theta$ can be computed exactly (in Section 3.4 we discuss what happens when this is not possible), from Theorem 2, we have that it is an unbiased estimator of the natural gradient of the $L(\theta(\mu))$ defined in Equation 13,

$$E \left[\frac{\partial \bar{\ell}(y, x, \theta)}{\partial \theta} | \pi \right] = I(\mu)^{-1} \frac{\partial L(\theta(\mu))}{\partial \mu} \quad (16)$$

However, one key difference in terms of convergence between online EM and sdEM can be seen in Equation (11): μ_{t+1} is a convex combination between μ_t and the expected sufficient statistics. Then, $\mu_{t+1} \in \mathcal{S}$ during all the iterations. As will be clear in the next section, we do not have this same guarantee in sdEM, but we can take advantage of the log prior term of Equation (13) to avoid this problem. This term plays a dual role as both “regularization”

Table 1: sdEM updating equations for fully observed data (Section 3.3) .

Loss	sdEM equation
NLL	$\mu_{t+1} = (1 - \rho_t(1 + \frac{\nu}{n}))\mu_t + \rho_t (s(y_t, x_t) + \frac{1}{n}\bar{\alpha})$
NCLL	$\mu_{t+1} = (1 - \rho_t\frac{\nu}{n})\mu_t + \rho_t (s(y_t, x_t) - E_y[s(y, x_t) \theta(\mu_t)] + \frac{1}{n}\bar{\alpha})$
Hinge	$\mu_{t+1} = (1 - \rho_t\frac{\nu}{n})\mu_t + \rho_t \begin{cases} \frac{1}{n}\bar{\alpha} & \text{if } \ln \frac{p(y_t, x_t \theta)}{p(\bar{y}_t, x_t \theta)} > 1 \\ s(y_t, x_t) - s(\bar{y}_t, x_t) + \frac{1}{n}\bar{\alpha} & \text{otherwise} \end{cases}$ <p style="text-align: center;">where $\bar{y}_t = \arg \max_{y \neq y_t} p(y, x_t \theta)$</p>

term and log-barrier function [31] i.e. a continuous function whose value increases to infinity as the parameter approaches the *boundary of the feasible region* or the support of $p(\theta(\mu)|\alpha)$ ³. Then, if the step sizes ρ_t are small enough (as happens near convergence), sdEM will always stay in the feasible region \mathcal{S} , due to the effect of the log prior term. The only problem is that, in the initial iterations, the step sizes ρ_t are large, so one iteration can jump out of the boundary of \mathcal{S} . The method to avoid that depends on the particular model, but for the models examined in this work it seems to be a simple check in every iteration. For example, as we will see in the experimental section when implementing a multinomial Naive Bayes, we will check at every iteration that each sufficient statistic or “word count” is always positive. If a “word count” is negative at some point, we will set it to a very small value. As mentioned above, this does not hurt the convergence of sdEM because in the limit this problem disappears due the effect of the log-prior term.

The last ingredient required to assess the convergence of a stochastic gradient descent method is to verify that the sequence of step sizes satisfies: $\sum \rho_t = \infty$, $\sum \rho_t^2 < \infty$.

So, if the sequence $(\mu_t)_{t \geq 0}$ converges, it will probably converge to the global minimum $(\mu^*, \theta^* = \theta(\mu^*))$ if $L(\theta)$ is convex, or to a local minimum if $L(\theta)$ is not convex [9].

Finally, we give an algorithmic description of sdEM in Algorithm 1. Following [11], we consider steps sizes of the form $\rho_t = (1 + \lambda t)^{-1}$, where λ is a positive scalar⁴. As mentioned above, the “Check-Step” is introduced to guarantee that μ_t is always in \mathcal{S} . Like the online EM algorithm [30, 13], Algorithm 1 resembles the classic *expectation maximization* algorithm [15] since, as we will see in the next section, the gradient is computed using *expected*

sufficient statistics. Assumption 3 guarantees that the maximization step can be performed efficiently. This step differentiates sdEM from classic stochastic gradient descent methods, where such a computation does not exist.

3.3 DISCRIMINATIVE LOSS FUNCTIONS

As we have seen so far, the derivation of sdEM is complete except for the definition of the loss function. We will discuss now how two well known *discriminative loss functions* can be used with this algorithm.

Negative Conditional Log-likelihood (NCLL)

As mentioned above, this loss function is defined as follows:

$$\ell_{CL}(y_t, x_t, \theta) = -\ln p(y_t, x_t|\theta) + \ln \int p(y, x_t|\theta) dy$$

And its gradient is computed as

$$\frac{\partial \ell_{CL}(y_t, x_t, \theta)}{\partial \theta} = -s(y_t, x_t) + E_y[s(y, x_t)|\theta]$$

where the sufficient statistic $s(y_t, x_t)$ comes from the gradient of the $\ln p(y_t, x_t|\theta)$ term in the NCLL loss, and the expected sufficient statistic $E_y[s(y, x_t)|\theta] = \int s(y, x_t)p(y|x_t, \theta)dy$, comes from the gradient of the $\ln \int p(y, x_t|\theta)dy$ term in the NCLL loss. As mentioned above, the computation of the gradient is similar to the *expectation step* of the classic EM algorithm.

The iteration equation of sdEM for the NCLL loss is detailed in Table 1. We note that in the case of multi-class prediction problems the integrals of the updating equation are replaced by sums over the different classes of the class variable Y . We also show the updating equation for the negative log-likelihood (NLL) loss for comparison purposes.

³The prior p would need to be suitably chosen.

⁴Our experiments suggest that trying $\lambda \in \{1, 0.1, 0.01, 0.001, \dots\}$ suffices for obtaining a quick convergence.

Table 2: sdEM updating equations for partially observed data (Section 3.4)

Loss	sdEM equation
NLL	$\mu_{t+1} = (1 - \rho_t(1 + \frac{\nu}{n}))\mu_t + \rho_t (E_z[s(y_t, z, x_t) \theta(\mu_t)] + \frac{1}{n}\bar{\alpha})$
NCLL	$\mu_{t+1} = (1 - \rho_t\frac{\nu}{n})\mu_t + \rho_t (E_z[s(y_t, z, x_t) \theta(\mu_t)] - E_{y_z}[s(y, z, x_t) \theta(\mu_t)] + \frac{1}{n}\bar{\alpha})$
Hinge	$\mu_{t+1} = (1 - \rho_t\frac{\nu}{n})\mu_t + \rho_t \begin{cases} \frac{1}{n}\bar{\alpha} & \text{if } \ln \frac{\int p(y_t, z, x_t \theta) dz}{\int p(\bar{y}_t, z, x_t \theta) dz} > 1 \\ E_z[s(y_t, z, x_t) \theta(\mu_t)] - E_z[s(\bar{y}_t, z, x_t) \theta(\mu_t)] + \frac{1}{n}\bar{\alpha} & \text{otherwise} \end{cases}$ <p style="text-align: center;">where $\bar{y}_t = \arg \max_{y \neq y_t} \int p(y, z, x_t \theta) dz$</p>

The Hinge loss

Unlike the previous loss which is valid for continuous and discrete (and vector-valued) predictions, this loss is only valid for binary or multi-class classification problems.

Margin-based loss functions have been extensively used and studied by the machine learning community for binary and multi-class classification problems [5]. However, in our view, the application of margin-based losses (different from the negative conditional log-likelihood) for discriminative training of probabilistic generative models is scarce and based on ad-hoc learning methods which, in general, are quite sophisticated [26]. In this section, we discuss how sdEM can be used to minimize the empirical risk of one of the most used margin-based losses, the Hinge loss, in binary and multi-class classification problems. But, firstly, we discuss how Hinge loss can be defined for probabilistic generative models.

We build on LeCun et al.'s ideas [23] about energy-based learning for prediction problems. LeCun et al. [23] define the Hinge loss for energy-based models as follows,

$$\max(0, 1 - (E(\bar{y}_t, x_t, w) - E(y_t, x_t, w)))$$

where $E(\cdot)$ is the energy function parameterized by a parameter vector w , $E(y_t, x_t, w)$ is the energy associated to the correct answer y_t and $E(\bar{y}_t, x_t, w)$ is the energy associated to the most offending incorrect answer, $\bar{y}_t = \arg \min_{y \neq y_t} E(y, x_t, w)$. Predictions y^* are made using $y^* = \arg \min_y E(y, x_t, w^*)$ when the parameter w^* that minimizes the empirical risk is found.

In our learning settings we consider the minus logarithm of the joint probability, $-\ln p(y_t, x_t|\theta)$, as an energy function. In consequence, we define the hinge loss as follows

$$\ell_{hinge}(y_t, x_t, \theta) = \max(0, 1 - \ln \frac{p(y_t, x_t|\theta)}{p(\bar{y}_t, x_t|\theta)}) \quad (17)$$

where \bar{y}_t denotes here too the most offending incorrect answer, $\bar{y}_t = \arg \max_{y \neq y_t} p(y, x_t|\theta)$.

The gradient of this loss function can be simply computed as follows

$$\frac{\partial \ell_{hinge}(y_t, x_t, \theta)}{\partial \theta} = \begin{cases} 0 & \text{if } \ln \frac{p(y_t, x_t|\theta)}{p(\bar{y}_t, x_t|\theta)} > 1 \\ -s(y_t, x_t) + s(\bar{y}_t, x_t) & \text{otherwise} \end{cases}$$

and the iteration equation for minimizing the empirical risk of the Hinge loss is also given in Table 1.

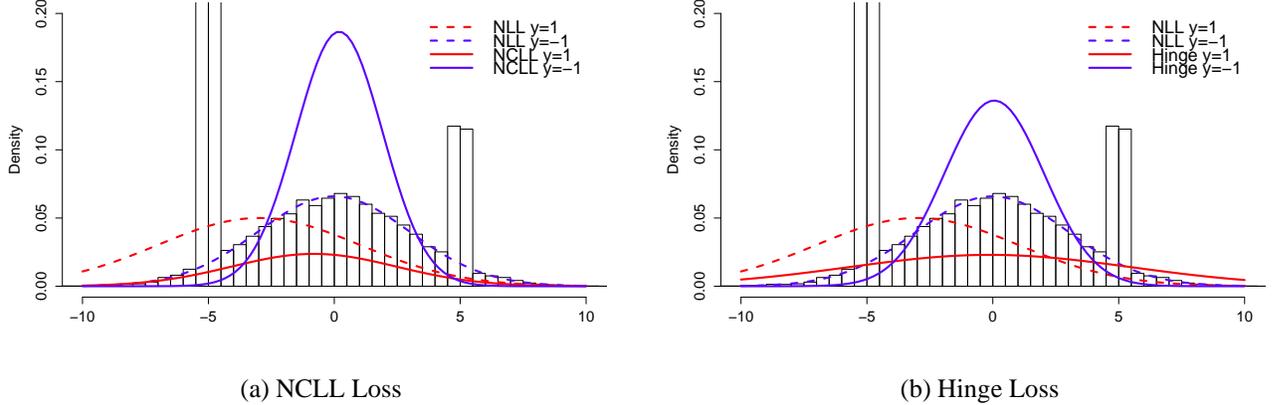
3.4 PARTIALLY OBSERVABLE DATA

The generalization of sdEM to partially observable data is straightforward. We denote by Z the vector of non-observable variables. sdEM will handle statistical models which define a probability distribution over (y, z, x) which belongs to the exponential family (Assumption 1). Assumption 2 and 3 remain unaltered.

The tuple (y, z, x) will denote the complete event or complete data, while the tuple (y, x) is the observed event or the observed data. So we assume that our given data set D with n observations is expressed as $\{(y_1, x_1), \dots, (y_n, x_n)\}$. So sdEM's Equation (14) and (15) are the same, with the only difference that the natural gradient is now defined using the inverse of the Fisher information matrix for the statistical model $p(y, z, x|\theta(\mu))$. The same happens for Theorem 2.

The NCLL loss and the Hinge loss are equally defined as in Section 3.3, with the only difference that the computation of $p(y_t, x_t|\theta)$ and $p(x_t|\theta)$ requires marginalization over z , $p(y_t, x_t|\theta) = \int p(y_t, z, x_t|\theta) dz$, $p(x_t|\theta) = \int p(y, z, x_t|\theta) dy dz$. The updating equations for sdEM under partially observed data for the NCLL and Hinge loss are detailed in Table 2. New expected sufficient statistics need to be computed,

Figure 1: Toy example (Section 4.1). The result using the NLL loss (i.e. MLE estimation) is plotted with dashed lines which represent the densities $p(y = k)N(x, \mu^{(k)}, \sigma^{(k)})$ for both classes (i.e. when the red line is higher than the blue line we predict the red class and vice versa). The estimated prediction accuracy of the MLE model is 78.6%. Solid lines represent the same estimation but using the NCLL and the Hinge loss. Their estimated prediction accuracies are 90.4% and 90.6%, respectively.



$E_z[s(y_t, z, x_t)|\theta] = \int s(y_t, z, x_t)p(z|y_t, x_t, \theta)dz$ and $E_{yz}[s(y, z, x_t)|\theta] = \int s(y, z, x_t)p(y, z|x_t, \theta)dydz$. As previously, we also show the updating equation for the negative log-likelihood (NLL) loss for comparison purposes.

3.5 sdEM AND APPROXIMATE INFERENCE

For many interesting models [8], the computation of the expected sufficient statistics in the iteration equations shown in Table 1 and 2 cannot be computed in closed form. This is not a problem as far as we can define *unbiased estimators* for these expected sufficient statistics, since the equality of Equation (16) still holds. As it will be shown in the next section, we use sdEM to discriminatively train *latent Dirichlet allocation* (LDA) models [8]. Similarly to [28], for this purpose we employ collapsed Gibbs sampling to compute the expected sufficient statistics, $E_z[s(y_t, z, x_t)|\theta]$, as it guarantees that at convergence samples are i.i.d. according to $p(z|y_t, x_t, \theta)$.

4 EXPERIMENTAL ANALYSIS

4.1 TOY EXAMPLE

We begin the experimental analysis of sdEM by learning a very simple Gaussian naive Bayes model composed by a binary class variable Y and a single continuous predictor X . Hence, the conditional density of the predictor given the class variable is assumed to be normally distributed. The interesting part of this toy example is that the training data is generated by a different model: $\pi(y = -1) = 0.5$, $\pi(x|y = -1) \sim N(0, 3)$ and $\pi(x|y = 1) \sim$

$0.8 \cdot N(-5, 0.1) + 0.2 \cdot N(5, 0.1)$. Figure 1 shows the histogram of the 30,000 samples generated from the π distribution. The result is a mixture of 3 Gaussians, one in the center with a high variance associated to $y = -1$ and two narrow Gaussians on both sides associated to $y = 1$.

sdEM can be used by considering 6 (non-minimal) sufficient statistics: $N^{(-1)}$ and $N^{(1)}$ as “counts” associated to both classes, respectively; $S^{(-1)}$ and $S^{(1)}$ as the “sum” of the x values associated to classes $y = -1$ and $y = 1$, respectively; and $V^{(-1)}$ and $V^{(1)}$ as the “sum of squares” of the x values for each class. We also have five parameters which are computed from the sufficient statistics as follows: Two for the prior of class $p(y = -1) = p^{(-1)} = N^{(-1)}/(N^{(-1)} + N^{(1)})$ and $p^{(1)} = N^{(1)}/(N^{(-1)} + N^{(1)})$; and four for the two Gaussians which define the conditional of X given Y , $\mu^{(-1)} = S^{(-1)}/N^{(-1)}$, $\sigma^{(-1)} = \sqrt{V^{(-1)}/N^{(-1)} - (S^{(-1)}/N^{(-1)})^2}$, and equally for $\mu^{(1)}$ and $\sigma^{(1)}$.

The sdEM’s updating equations for the NCLL loss can be written as follows

$$\begin{aligned} N_{t+1}^{(k)} &= N_t^{(k)} + \rho_t(I[y_t = k] - p_t(k|x_t)) + \frac{\rho_t}{n} \\ S_{t+1}^{(k)} &= (1 - \frac{\rho_t}{n})S_t^{(k)} + \rho_t x_t (I[y_t = k] - p_t(k|x_t)) \\ V_{t+1}^{(k)} &= (1 - \frac{\rho_t}{n})V_t^{(k)} + \rho_t x_t^2 (I[y_t = k] - p_t(k|x_t)) + \frac{\rho_t}{n} \end{aligned}$$

where k indexes both classes, $k \in \{-1, 1\}$, $I[\cdot]$ denotes the indicator function, $p_t(k|x_t)$ is an abbreviation of $p(y = k|x_t, \theta_t)$, and θ_t is the parameter vector computed from the sufficient statistics at the t -th iteration.

Figure 2: Convergence trade-off of the Hinge loss versus the NCLL loss and the perplexity for a multinomial naive Bayes model trained minimizing the Hinge loss using sdEM. Circle-lines, triangle-lines and cross-lines correspond to the results with 20NewsGroup, Cade and Reuters-R52 datasets, respectively.

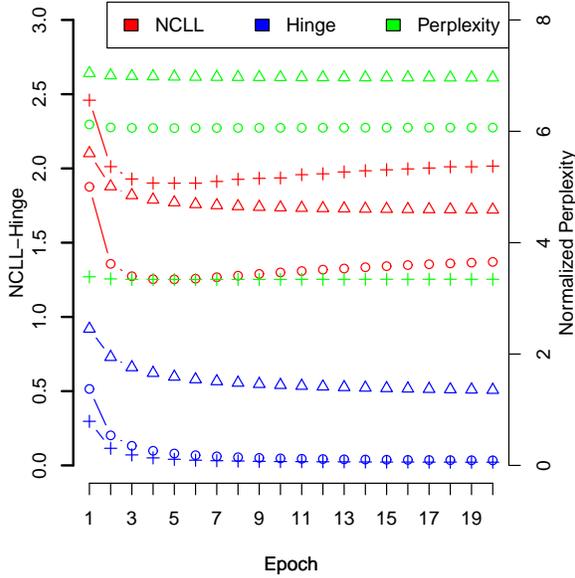
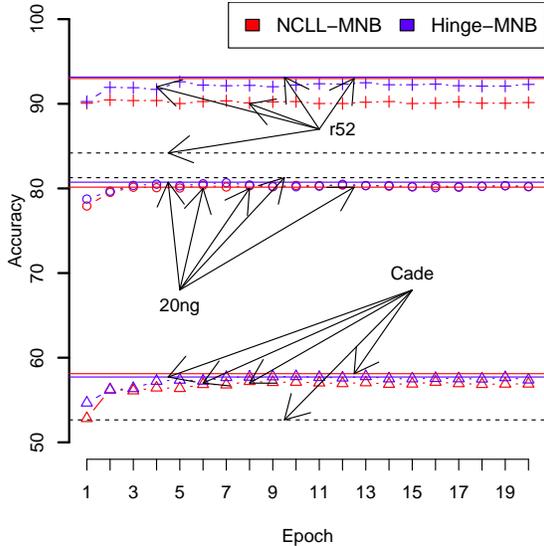


Figure 3: Convergence of the classification accuracy for a multinomial naive Bayes model trained minimizing the NCLL loss (NCLL-MNB) and the Hinge loss (Hinge-MNB) using sdEM. Red circle-lines, red triangle-lines and red cross-lines correspond to the results of NCLL-MNB with 20NewsGroup, Cade and Reuters-R52 datasets, respectively. Same for Hinge-MNB. The three blue and the three red solid lines detail the accuracy of logistic regression and SVM, respectively. The three dashed black lines detail the accuracy of plain MNB with a Laplace prior.



Similarly, the sdEM’s updating equations for the Hinge loss can be written as follows,

$$N_{t+1}^{(k)} = N_t^{(k)} + ky_t \rho_t I[\ln \frac{p_t(y_t|x_t)}{p_t(\bar{y}_t|x_t)} < 1] + \frac{\rho_t}{n}$$

$$S_{t+1}^{(k)} = (1 - \frac{\rho_t}{n})S_t^{(k)} + ky_t \rho_t x_t I[\ln \frac{p_t(y_t|x_t)}{p_t(\bar{y}_t|x_t)} < 1]$$

$$V_{t+1}^{(k)} = (1 - \frac{\rho_t}{n})V_t^{(k)} + ky_t \rho_t x_t^2 I[\ln \frac{p_t(y_t|x_t)}{p_t(\bar{y}_t|x_t)} < 1] + \frac{\rho_t}{n}$$

where the product ky_t is introduced in the updating equations to define the sign of the sum, and the indicator function $I[\cdot]$ defines when the hinge loss is null.

In the above set of equations we have considered as a conjugate prior for the Gaussians a three parameter Normal-Gamma prior, $\nu = 1$ and $\bar{\alpha}_1 = 0$ for $S^{(k)}$ and $\bar{\alpha}_2 = 1$ for $V^{(k)}$ [6, page 268], and a Beta prior with $\nu = 0$ and $\bar{\alpha} = 1$ for $N^{(k)}$. We note that these priors assign zero probability to “extreme” parameters $p^{(k)} = 0$ (i.e. $N^{(k)} = 0$) and $\sigma^{(k)} = 0$ (i.e. $V^{(k)}/N^{(k)} - (S^{(k)}/N^{(k)})^2 = 0$).

Finally, the “Check-step” (see Algorithm 1) performed before computing θ_{t+1} , and which guarantees that all sufficient statistics are correct, is implemented as follows:

$$N_{t+1}^{(k)} = \max(N_{t+1}^{(k)}, \frac{\rho_t}{n})$$

$$V_{t+1}^{(k)} = \max(V_{t+1}^{(k)}, \frac{(S_{t+1}^{(k)})^2}{N_{t+1}^{(k)}} + \frac{\rho_t}{n})$$

I.e., when the $N^{(k)}$ “counts” are negative or too small or when the $V^{(k)}$ values lead to negative or null deviations $\sigma^{(k)} \leq 0$, they are fixed with the help of the prior term.

The result of this experiment is given in Figure 1 and clearly shows the different trade-offs of both loss functions compared to *maximum likelihood estimation*. It is interesting to see how a generative model which does not match the underlying distribution is able to achieve a pretty high prediction accuracy when trained with a discriminative loss function (using the sdEM algorithm).

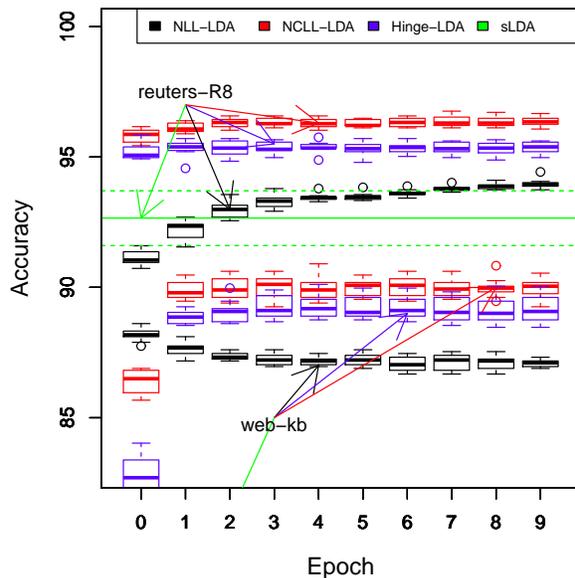
4.2 sdEM FOR TEXT CLASSIFICATION

Next, we briefly show how sdEM can be used to discriminatively train some generative models used for text classification, such as multinomial naive Bayes and a similar classifier based on latent Dirichlet allocation models [8]. Supplementary material with full details of these experiments and the Java code used in this evaluation can be download at: <http://sourceforge.net/projects/sdem/>

Multinomial Naive Bayes (MNB)

MNB assumes that words in documents with the same class or label are distributed according to an independent multinomial distribution. sdEM can be easily applied to train this

Figure 4: Convergence of the classification accuracy of LDA classification models trained by sdEM using different loss functions (NLL, NCLL and Hinge) over 10 different random initializations. The two dashed lines and the single solid line detail the maximum, minimum and mean accuracy of sLDA, respectively, over 10 random initializations.



model. The sufficient statistics are the “prior class counts” and the “word counts” for each class. The updating equations and the check step are the same as those of $N_t^{(k)}$ in the previous toy example. Parameters of the MNB are computed simply through normalization operations. Two different conjugate Dirichlet distributions were considered: A “Laplace prior” where $\bar{\alpha}_i = 1$; and a “Log prior” where $\bar{\alpha}_i =$ “logarithm of the number of words in the corpus”. We only report analysis for “Laplace prior” in the case of NCLL loss and for “Log prior” in the case of Hinge loss. Other combinations show similar results, although NCLL was more sensitive to the chosen prior.

We evaluate the application of sdEM to MNB with three well-known multi-class text classification problems: 20Newsgroup (20 classes), Cade (12 classes) and Reuters21578-R52 (52 classes). Data sets are stemmed. Full details about the data sets and the train/test data sets split used in this evaluation can be found in [14].

Figure 2 shows the convergence behavior of sdEM with $\lambda = 1e-05$ when training a MNB by minimizing the Hinge loss (Hinge-MNB). In this figure, we plot the evolution of the Hinge loss but also the evolution of the NCLL loss and the normalized perplexity (i.e. the perplexity measure [8] divided by the number of training documents) at each epoch. We can see that there is a trade-off between the different losses. E.g., Hinge-MNB decreases the Hinge loss (as expected) but tends to increase the NCLL loss, while it

only decreases perplexity at the very beginning.

Figure 3 displays the evolution of the classification accuracy of two MNBs trained minimizing the NCLL loss and the Hinge loss using sdEM. We compare them to: the standard MNB with a “Laplace prior”; the L2-regularized Logistic Regression; and the primal L2-regularized SVM. The two later methods were taken from the Liblinear toolkit v.18 [17]. As can be seen, sdEM is able to train simple MNB models with a performance very close to that provided by highly optimized algorithms.

Latent Dirichlet Allocation (LDA)

We briefly show the results of sdEM when discriminatively training LDA models. We define a classification model equal to MNB, but where the documents of the same class are now modeled using an independent LDA model. We implement this model by using, apart from the “prior class counts”, the standard sufficient statistics of the LDA model, i.e. “words per hidden topic counts”, associated to each class label. Similarly to [28], we used an online Collapsed Gibbs sampling method to obtain, at convergence, unbiased estimates of the expected sufficient statistics (see Table 2).

This evaluation was carried out using the standard train/test split of the Reuters21578-R8 (8 classes) and web-kb (4 classes) data sets [14], under the same preprocessing than in the MNB’s experiments. Figure 4 shows the results of this comparison using 2-topics LDA models trained with the NCLL loss (NCLL-LDA), the Hinge loss (Hinge-LDA), and also the NLL loss (NLL-LDA) following the updating equations of Table 2. We compared these results with those returned by supervised-LDA (sLDA) [7] using the same prior, but this time with 50 topics because less topics produced worse results. We see again how a simple generative model trained with sdEM outperforms much more sophisticated models.

5 CONCLUSIONS

We introduce a new learning algorithm for discriminative training of generative models. This method is based on a novel view of the online EM algorithm as a stochastic *natural gradient* descent algorithm for minimizing general discriminative loss functions. It allows the training of a wide set of generative models with or without latent variables, because the resulting models are always generative. Moreover, sdEM is comparatively simpler and easier to implement (and debug) than other ad-hoc approaches.

Acknowledgments

This work has been partially funded from the European Union’s Seventh Framework Programme for research, technological development and demonstration under grant agreement no 619209 (AMIDST project).

References

- [1] Sungjin Ahn, Anoop Korattikara Balan, and Max Welling. Bayesian posterior sampling via stochastic gradient Fisher scoring. In *ICML*, 2012.
- [2] Shun-ichi Amari. *Differential-geometrical methods in statistics*. Springer, 1985.
- [3] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural Comput.*, 10(2):251–276, 1998.
- [4] Shun-ichi Amari and Hiroshi Nagaoka. *Methods of information geometry*, volume 191. American Mathematical Soc., 2007.
- [5] Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [6] José M Bernardo and Adrian FM Smith. *Bayesian theory*, volume 405. John Wiley & Sons, 2009.
- [7] David M Blei and Jon D McAuliffe. Supervised topic models. In *NIPS*, volume 7, pages 121–128, 2007.
- [8] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [9] L. Bottou. Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9), 1998.
- [10] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.
- [11] L. Bottou. Stochastic gradient descent tricks. In *Neural Networks: Tricks of the Trade*, pages 421–436. Springer, 2012.
- [12] L. Bottou and O. Bousquet. The tradeoffs of large scale learning. In *NIPS*, volume 4, page 2, 2007.
- [13] Olivier C. and E. Moulines. On-line expectation-maximization algorithm for latent data models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3):593–613, 2009.
- [14] Ana Cardoso-Cachopo. Improving Methods for Single-label Text Categorization. PdD Thesis, 2007.
- [15] Arthur P Dempster, Nan M Laird, Donald B Rubin, et al. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal statistical Society*, 39(1):1–38, 1977.
- [16] Greiner et al. Structural extension to logistic regression: Discriminative parameter learning of belief net classifiers. *Mach. Learning*, 59(3):297–322, 2005.
- [17] R.E. Fan, K.W. Chang, C.J. Hsieh, X.R. Wang, and C.J. Lin. Liblinear: A library for large linear classification. *JMLR*, 9:1871–1874, 2008.
- [18] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. The weka data mining software: an update. *ACM SIGKDD explorations newsletter*, 11(1):10–18, 2009.
- [19] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.
- [20] Thorsten Joachims. *Text categorization with support vector machines: Learning with many relevant features*. Springer, 1998.
- [21] Harold Joseph Kushner and G George Yin. *Stochastic approximation algorithms and applications*. Springer New York, 1997.
- [22] S. Lacoste-Julien, F. Sha, and M.I. Jordan. DiscLDA: Discriminative learning for dimensionality reduction and classification. In *NIPS*, volume 83, page 85, 2008.
- [23] Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, and F Huang. A tutorial on energy-based learning. *Predicting structured data*, 2006.
- [24] 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.
- [25] Frank Nielsen and Vincent Garcia. Statistical exponential families: A digest with flash cards. *arXiv preprint arXiv:0911.4863*, 2009.
- [26] F. Pernkopf, M. Wohlmayr, and S. Tschitschek. Maximum margin Bayesian network classifiers. *IEEE Trans. PAMI*, 34(3):521–532, 2012.
- [27] Garvesh Raskutti and Sayan Mukherjee. The information geometry of mirror descent. *arXiv preprint arXiv:1310.7780*, 2013.
- [28] D. Rohde and O. Cappe. Online maximum-likelihood estimation for latent factor models. In *Statistical Signal Processing Workshop (SSP), 2011 IEEE*, pages 565–568, June 2011.
- [29] David Ruppert. Efficient estimations from a slowly convergent Robbins-Monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988.
- [30] Masa-aki Sato. Convergence of on-line EM algorithm. In *Proc. of the Int. Conf. on Neural Information Processing*, volume 1, pages 476–481, 2000.

- [31] SJ Wright and J Nocedal. *Numerical optimization*, volume 2. Springer New York, 1999.
- [32] Jun Zhu, Amr Ahmed, and Eric P Xing. MedLDA: maximum margin supervised topic models for regression and classification. In *ICML*, pages 1257–1264. ACM, 2009.

Appendices

This supplementary material aims to extend, detail and complement the experimental evaluation of sdEM given in the main paper. The structure of this document is as follows. Section A details the experimental evaluation of the multinomial naive Bayes classifier and introduces new experiments comparing with the stochastic gradient descent algorithm [10]. The experimental evaluation of sdEM applied to latent Dirichlet allocation models is detailed and extended in Section B. Section C points to the software repository where all the software code used in this experimental evaluation can be downloaded to reproduce all these results.

A Multinomial Naive Bayes for text classification

Description of the algorithm

As commented in the main paper, a multinomial Naive Bayes (MNB) classifier assumes that the words of the documents with the same class labels are distributed according to an independent multinomial probability distribution. In this section we evaluate the use of sdEM to discriminatively train MNB models using the NCLL and the Hinge loss functions. In the first case, such a model would be related to a logistic regression model; while in the second case we will obtain a model directly related to a linear support vector machine classifier [20].

The general updating equations for this problem can be found in the main paper in Table 2. But a detailed pseudo-code description is now given in Algorithm 2 for the NCLL loss and in Algorithm 4 for the Hinge loss. In both cases, the sufficient statistics are the "prior class counts" stored in the matrix C and the "word counts per class" stored in the matrix N . Matrix M is introduced to allow efficient computations of the posterior probability of the class variable given a document d , $p(Y|d, N, M, C, \gamma)$. How this posterior is computed is detailed in Algorithm 3. In that way, the computational complexity of processing a label-document pair is linear in the number of words of the document. Finally, the function $Normalize(\cdot, \cdot)$ produces a multinomial probability by normalizing the vector of counts. The second argument contains the prior correction considered

in this normalization, i.e. the value which is added to each single component of the count vector to avoid null probabilities, similar to what is done in Algorithm 3.

In both algorithms, we consider a Dirichlet distribution prior for the multinomial distributions. As detailed in the header of these algorithms, two different priors are considered: prior $P1$, with Dirichlet's metaparameters $\alpha_k = 1$; and prior $P2$ with $\alpha_k = \ln |W|$, where $|W|$ denotes the total number of different words in the corpus. In both cases, the prior assigns null probability to parameters lying in the "border" of the parameter space (i.e. when a null probability is assigned to some word).

As commented in the "toy example" of the main paper in Section 4.1, the parametrization that we chose for this Dirichlet prior makes that the ν parameter, arising in the exponential family form of this prior (see Assumption 2 of the main paper), be equal to null, $\nu = 0$. This can be seen when expressing the Dirichlet distribution in the following exponential form:

$$\begin{aligned} Dir(\theta_1, \dots, \theta_K; \alpha_1, \dots, \alpha_K) &= \frac{\prod_k \Gamma(\alpha_k)}{\Gamma(\sum_k \alpha_k)} \theta_1^{\alpha_1-1} \dots \theta_K^{\alpha_K-1} \\ &= \exp\left(\sum_k (\alpha_k - 1) \ln \theta_k + \sum_k \ln \Gamma(\alpha_k) - \ln \Gamma(\sum_k \alpha_k)\right) \end{aligned}$$

The second and third terms inside the exponent in the above equation correspond to the log partition function $A_g(\alpha)$ of the prior. The first term correspond to the dot product between the sufficient statistics $(\ln \theta_1, \dots, \ln \theta_K)$ and the natural parameters $(\alpha_1 - 1, \dots, \alpha_K - 1)$. As can be seen, the ν parameter can be obviated in this definition, i.e. $\nu = 0$.

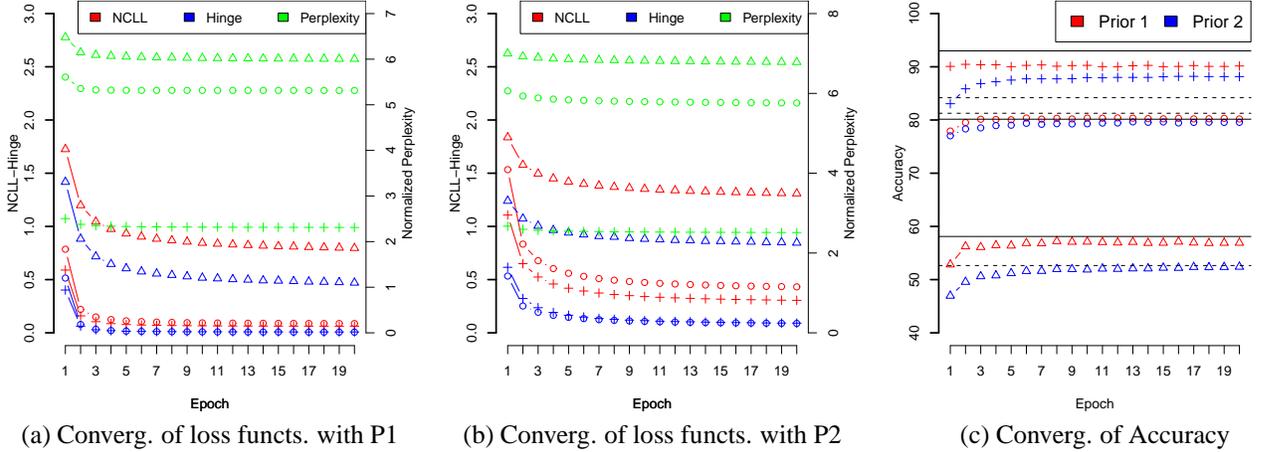
Table 3: Data sets statistics. $|W|$ denotes the number of different words in the data set, $|Y|$ denotes the number of labels of the class variable, D_{train} and D_{test} the number of documents in the training and the test set, respectively. These data sets can be downloaded, for example, from <http://sourceforge.net/projects/sgmweka/>.

Name	$ W $	$ Y $	D_{train}	D_{eval}
ACL-IMDB	89527	2	47000	3000
Amazon12	86914	2	267875	100556
Cade	157483	12	27322	13661
Reuters-R8	14575	8	2785	1396
Reuters-R52	16145	52	5485	2189
WebKb	7287	4	6532	2568
20 News-group	54580	20	11293	7528

Experimental Evaluation

As detailed in the main paper, we evaluate the application of sdEM to MNB with three well-known multi-

Figure 5: Convergence behavior of sdEM applied to a multinomial naive Bayes model (NCLL-MNB) with different priors. Circle-lines, triangle-lines and cross-lines correspond to the results with 20NewsGroup, Cade and Reuters-R52 datasets, respectively. In the third figure, the three solid lines detail the accuracy of logistic regression for these three data sets. The tree dashed lines detail the accuracy of plain MNB with P1 (MNB results with P2 are omitted because they are much worse).



class text classification problems: 20Newsgroup, Cade and Reuters21578-R52. These data sets are stemmed. Full details about the data sets and the train/test data sets split used in this evaluation can be found in [14]. Although Table 3 shows some of the main statistics of these data sets.

Figure 5 (a), Figure 5 (b), Figure 6 (a) and Figure 6 (b) show the convergence behavior of sdEM with $\lambda = 1e-05$ ⁵ when training the MNB by minimizing the NCLL loss (NCLL-MNB) and by minimizing the Hinge loss (Hinge-MNB), respectively. In both cases, we plot the evolution of the NCLL loss, the Hinge loss and the normalized perplexity (i.e. the perplexity measure [8] divided by the number of training documents) of the trained model at each epoch. We can see that there is a trade-off between the different losses. For example, Hinge-MNB decreases the Hinge loss (as expected) but tends to increase the NCLL loss, while it only decreases perplexity at the very beginning. This last trend is much stronger when considering the P1 prior. A similar behavior can be observed for NCLL-MNB, with the main difference that the NCLL loss is an upper bound of the Hinge loss, and then when NCLL-MNB minimizes the NCLL loss it also minimizes the Hinge loss. Here it can be also observed that the perplexity remains quite stable specially for P1.

Figure 5 (c) and Figure 6 (c) displays the evolution of the classification accuracy for the above models. We compare it to the standard MNB with a “Laplace prior”⁶ and

⁵Other values yield similar results and offer stable convergence, although at lower pace.

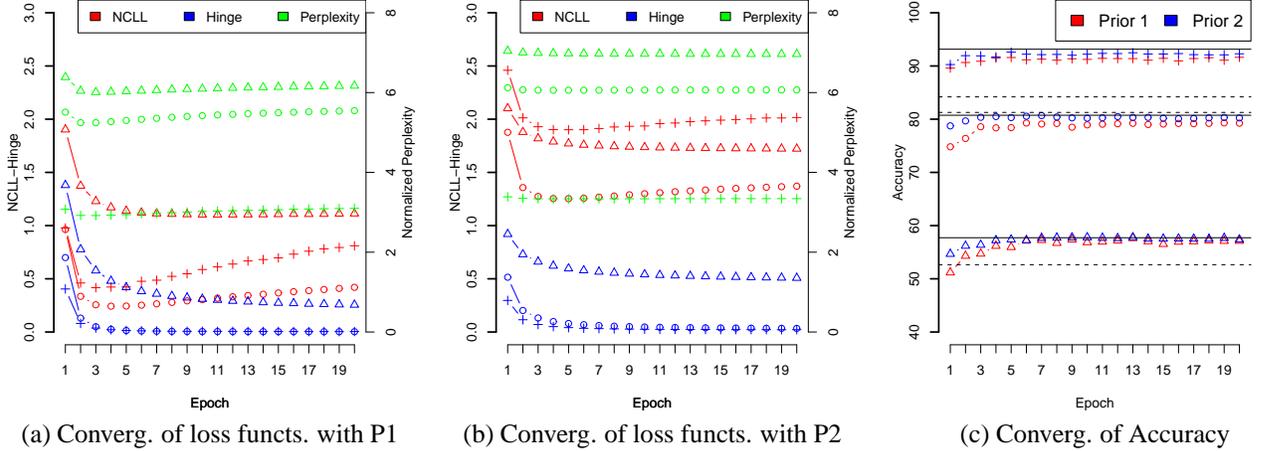
⁶A “Log prior” was also evaluated but reported much worse results.

with L2-regularized Logistic Regression and primal L2-regularized SVM implemented in the Liblinear toolkit v.18 [17]. For the case of the NCLL loss, the models seem to be more dependent of the chosen prior, specially for the Cade dataset. In any case, we can see that sdEM is able to train simple MNB models with a performance very close to that provided by highly optimized algorithms.

A new set of experiments is included in this analysis comparing the MNB models learnt with sdEM with the classic stochastic gradient descent (SGD) algorithm. This evaluation is made using the Amazon12 and ACL-IMDB data sets (whose main details can be found in Table 3). We choose these data sets because they are binary classification problems, which are very well defined problems for logistic regression and linear SVM models. How SGD is used to train this model can be seen in [11].

In this evaluation we simply plot the evolution of the classification accuracy of the SGD algorithm when training a linear classifier using the NCLL loss (NCLL-SGD) and the Hinge loss (Hinge-SGD) with a L2 regularization for different learning rates or decreasing steps ρ_t . SGD is implemented as detailed in [11], where the weight of the regularized term is fixed to $1e-4$. As recommended in [11], learning rates ρ_t for SGD are computed as follows: $\rho_t = \frac{\lambda}{1+\lambda \cdot 0.0001 \cdot t}$. We also look at the evolution of the classification accuracy of NCLL-MNB and Hinge-MNB with different priors in these two data sets and using different learning rates. In both cases, the plotted learning rates ρ_t are selected by using different λ values of the form $\lambda \in \{1, 0.1, 0.01, 0.001, 0.0001, 0.00001, \dots\}$. These results are shown in Figures 7 and 8. In each case, we con-

Figure 6: Convergence behavior of sEM applied to a multinomial naive Bayes model (Hinge-MNB) with different priors. Circle-lines, triangle-lines and cross-lines correspond to the results with 20NewsGroup, Cade and Reuters-R52 datasets, respectively. In the third figure, the three solid lines detail the accuracy of SVM for these three data sets. The tree dashed lines detail the accuracy of plain MNB with P1 (MNB results with P2 are omitted because they are much worse).



sider the 5 consecutive λ values with the quickest convergence speed.

B Latent Dirichlet Allocation (LDA) for text classification

Description of the algorithm

As commented in the main paper, we depart from a classification model similar to MNB, but where the documents of the same class are now modeled using an independent LDA model instead of a multinomial distribution. The generative process of each label-document pair in the corpus would be as follows [8]:

1. Choose class label $y \sim p(y|\theta_Y)$, a multinomial probability.
2. Choose $N \sim Poisson(\xi_y)$, the length of the document follows a Poisson distribution.
3. Choose $\phi_y \sim Dir(\alpha_y)$, a Dirichlet distribution with dimension $|Z|$ (the meta-parameters are set to $1/|Z|$ in the experimental evaluation).
4. For each of the N words w_n :
 - (a) Choose a topic $z_n \sim Multinomial(\phi_y)$ with dimension $|Z|$.
 - (b) Choose a word $w_n \sim p(w_n|z_n, \beta_y)$, a multinomial probability conditioned on the topic z_n .

In our case the unknown parameters are the β_y for each class label, which defines the multinomial distribution of the step 4 (b) and the parameter θ_Y which defines the prior.

We denote by d to a document as a bag of words $d = \{w_1, \dots, w_N\}$ and we denote by z_d to a particular hidden topic assignment vector for the words in d . Then the sufficient statistics for this model would be a three dimensional matrix indexed by $k \in \{1, \dots, |Y|\}$, $z \in \{1, \dots, |Z|\}$ and $w \in \{1, \dots, |W|\}$, where $|W|$ denotes again the total number of different words in the corpus. The (k, z, w) -th component of this sufficient statistics matrix is computed as follows:

$$s_{k,z,w}(y, z_d, d) = I[y = k] \sum_n I[z_n = z] I[w_n = w]$$

As previously commented in the main paper, these sufficient statistics would correspond to the "words per hidden topic counts". By adding the "prior class counts", we would complete all the sufficient statistics that define this classification model.

As also commented in the main paper, similarly to [28], we used an online Collapsed Gibbs sampling method to obtain, at convergence, unbiased estimates of the expected sufficient statistics (see Section 3.5 in the main paper). This collapsed Gibbs sampling method makes use of the analytical marginalization of the parameter ϕ_y and samples in turn each of the indicator variables z_1, \dots, z_N . The probability of an indicator variable z_n conditioned on all the words of the document and all the other indicators variables can be computed as follows:

$$p(z_n|y, \{z_{n'}\}_{n' \neq n}, d) \propto \beta_{y,z_n,w_n} \cdot (S_{z_n}^{(-w_n)} + \alpha) \quad (18)$$

Figure 7: Convergence of the classification accuracy for NCLL-SGD and NCLL-MNB for different priors and different learning rates in the Amazon12 and ACL-IMDB data set. Solid lines detail the accuracy of the aforementioned Liblinear’s logistic regression and dashed lines detail the accuracy of the plain MNB with the corresponding prior. The numbers of the bottom right legends correspond to different λ values (see Section 3.2 of the main paper), which defines how the sEM’s learning rates ρ_t decreases over time.

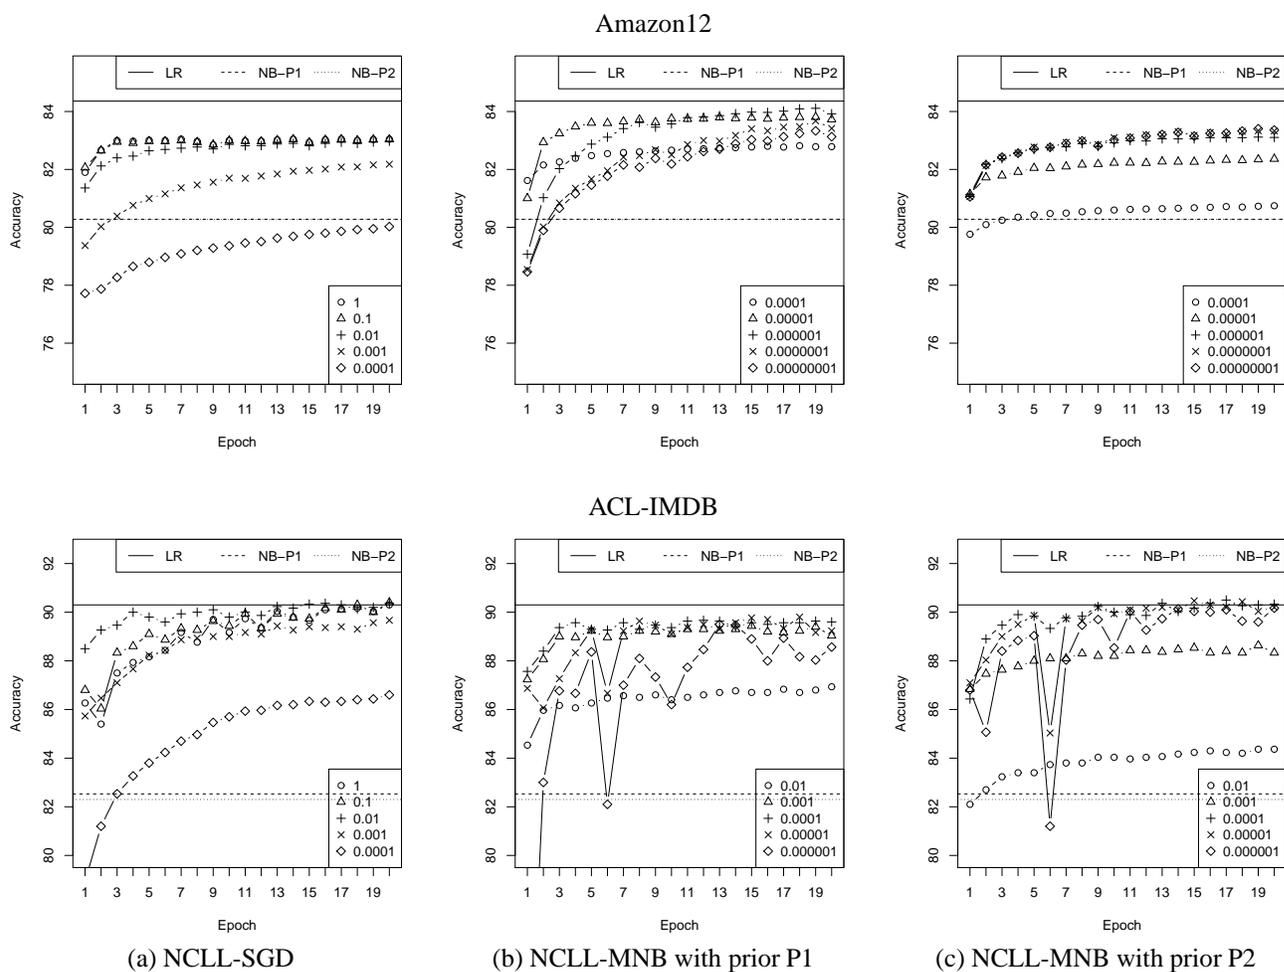
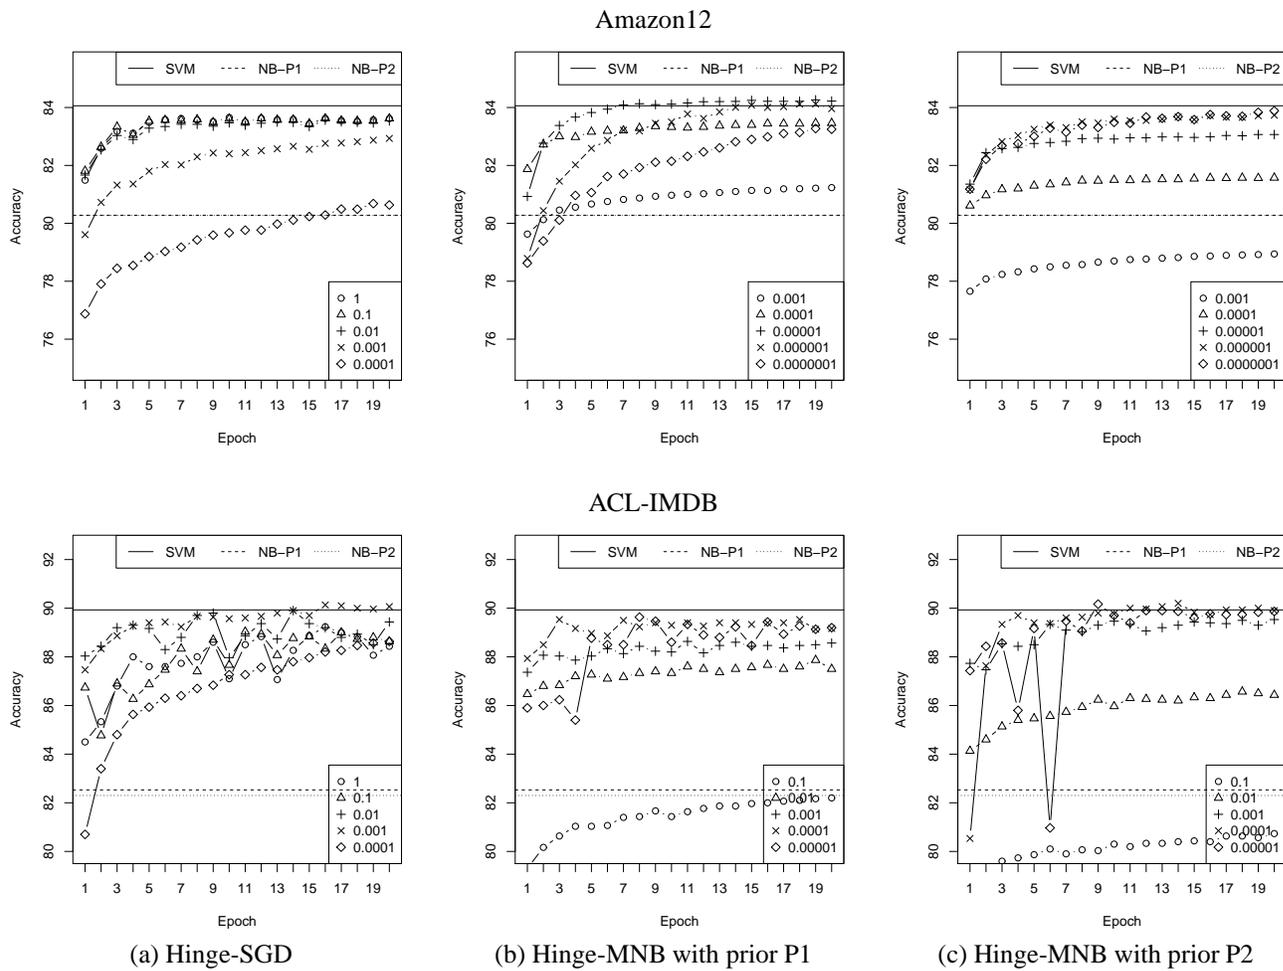


Figure 8: Convergence of the classification accuracy for Hinge-SGD and Hinge-MNB and different learning rates for different priors in the Amazon12 and ACL-IMDB data set. Solid lines detail the accuracy of the aforementioned Liblinear’s SVM classifier and dashed lines detail the accuracy of the plain MNB with the corresponding prior. The numbers of the bottom right legends correspond to different λ values (see Section 3.2 of the main paper), which defines how the sEM’s learning rates ρ_t decreases over time.



where $S_z^{(-w_n)} = \sum_{n' \neq n} I[z_{n'} = z]$ and β_{y, z_n, w_n} is the component of the β parameter vector which defines the probability that the n -th word in document d is equal to w_n given that hidden topic is z_n and the class label of the document is y , $p(w_n | z_n, \beta_y)$.

The above equation defines a Markov chain that when it is run generates unbiased samples from its stationary distribution, $p(z_n | d, y)$ (after discarding the first burn-in samples). So, we could then compute the expected sufficient statistics required to apply the sdEM algorithm over these models. Let us note that under our online settings the β parameter of Equation (18) is fixed to the values β_{t-1} estimated in the previous step and the this online collapsed Gibbs sampler only requires that the simulation is conditioned to the latent variables of the current observed document (i.e. it does not involve the hidden topics of the other documents in the corpus as happens with its batch counterpart).

In Algorithm 5 and Algorithm 7, we give a pseudocode description of the sdEM algorithm when applied to the this LDA classification model when using the NCLL and the Hinge loss functions, respectively. As can be seen, this algorithms does not directly relate to the standard LDA implementation, because we employ the same simplification used in the implementation⁷ of the sLDA algorithm [7] for multi-class prediction. This simplification assumes that all the occurrences of the same word in a document share the same hidden topic. The first effect of this assumption is that the number of hidden variables is reduced and the algorithm is much quicker. Whether this simplifying assumption has a positive or negative effect in the classification performance of the models is not evaluated here.

Let us also see in the pseudo-code of these two algorithms, that Hinge-LDA will tend to be computationally more efficient than NCLL-LDA, because Hinge-LDA does not update any parameter when it classifies a document with a margin higher than 1. However, NCLL-LDA always updates all the parameters. When we deal with a high number of classes, this may imply a great difference in the computational performance. But this is something which is not evaluated in this first experimental study.

We also use a heuristic method⁸ to initialize the hidden topics variables z_n of the incoming document which consists in sampling the hidden topics according to Equation 18, where $S_{z_n}^{(-w_n)}$ is computed on-the-fly i.e. for the first word is a vector of zeros and, then, it is updated according to the sampled topics. It is similar to running collapsed Gibbs sampling for one iteration.

We emphasis again that these algorithms are based on the updating equations given in the Table 2 of the main paper.

Experimental Evaluation

As previously commented in the paper, this evaluation was carried out using the standard train/test split of the Reuters21578-R8 and Web-KB data sets [14], under the same preprocessing than in the MNB’s experiments. In Table 3 some statistics about these data sets are given.

We used sdEM to train 2-topics LDA classification by minimizing the NCLL loss (NCLL-LDA), by minimizing the Hinge loss (Hinge-LDA), and also by minimizing the negative log-likelihood loss (NLL-LDA), following the updating equations of Table 2 in the main paper. We remind that at Figure 3 in the main paper, we show the results of the comparison of the classification accuracy of these models with the results obtained by supervised-LDA (sLDA) [7] using the same prior, but using 50 topics because with less topics it produced worse results.

We plot here at Figure 9, the convergence behavior at the training phase of the above models. The aim is to highlight how there is again a similar trade-off between the different losses when we train this model by minimizing a discriminative loss function such as NCLL or Hinge loss w.r.t. when we train this same model by minimizing a ”generative loss” such as the negative log-likelihood (NLL).

Looking at these figures we can see like neither NCLL-LDA nor Hinge-LDA decrease the perplexity loss in opposite to NLL-LDA. We can also see that NLL-LDA does decrease either the NCLL or the Hinge loss but not so successfully as NCLL-LDA or Hinge-LDA.

C sdEM Java Code

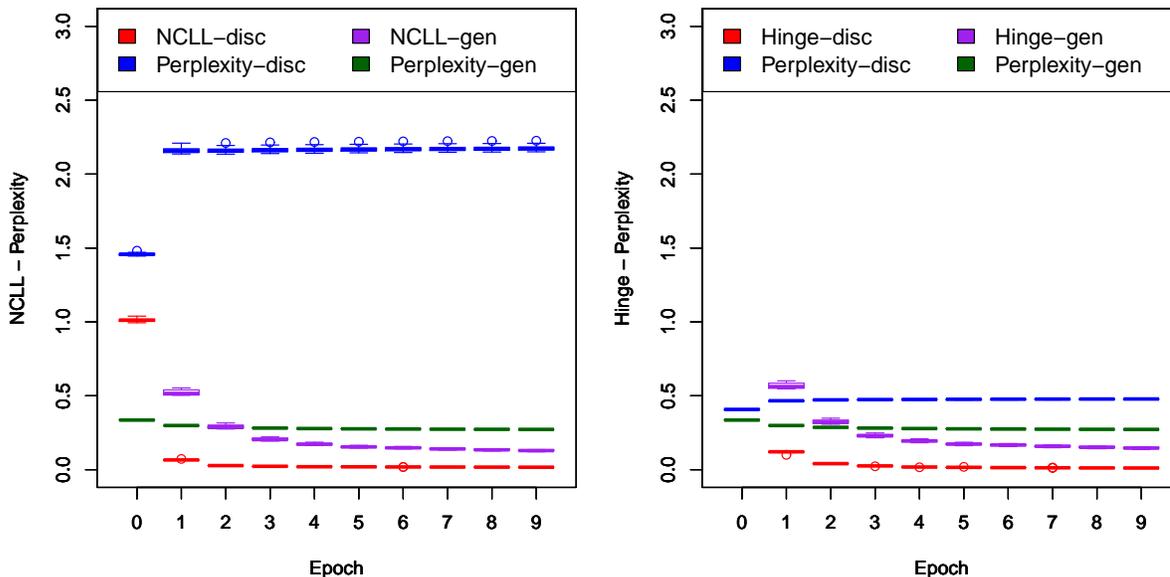
All the code used to build all the experiments presented in this supplemental material or in the main paper can be downloaded from the following code repository ”<https://sourceforge.net/projects/sdem/>” (in ”Files” tab). This code is written in Java and mostly builds on Weka [18] data structures.

⁷Code available at <http://www.cs.cmu.edu/~chongw/slida/>

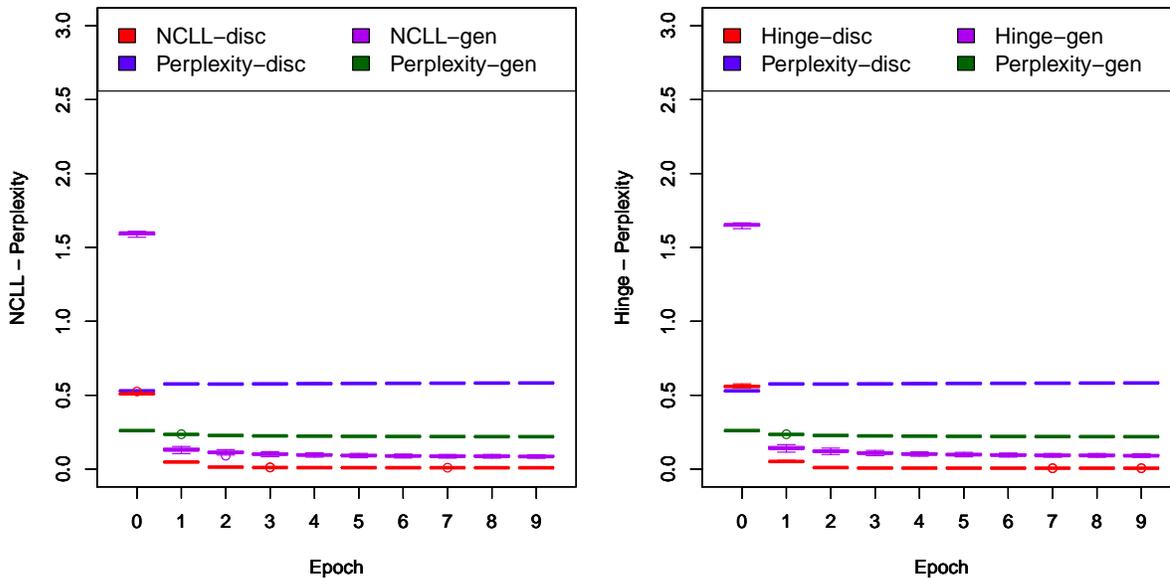
⁸It is proposed in <http://shuyo.wordpress.com/2011/06/27/collapsed-gibbs-sampling-estimation-for-latent-dirichlet-allocation-3/>.

Figure 9: Convergence behavior of sdEM when applied to the LDA classification model. Left side figures consider the NCLL-LDA model, i.e. when sdEM minimizes the NCLL loss function. Then, the series NCLL-disc and Perplexity-disc display the evolution of these two losses for the NCLL-LDA model. Right side figures consider the Hinge-LDA model, i.e. when sdEM minimizes the Hinge loss function. Then, the series Hinge-disc and Perplexity-disc display the evolution of these two losses for the Hinge-LDA model. Series NCLL-gen, Hinge-gen and Perplexity-gen show the evolution of the NCLL, Hinge and perplexity losses, respectively, for the NLL-LDA model, i.e. when sdEM minimizes the negative log-likelihood (NLL) loss function.

Reuters-R8 data set



Web-KB data set



(a) NCLL-LDA

(b) Hinge-LDA

Algorithm 2 sdEM for Multinomial Naive Bayes with the NCLL loss. $|d|$ denotes the number of different words in the current document d and $|w|_d$ to the number of times word w appears in document d . $|W|$ denotes the total number of different words in the corpus.

Require: D is randomly shuffled.

Require: α value as prior count for each word. Two values are considered $\alpha = 1$ and $\alpha = \ln |W|$.

```

1:  $\forall k, w \ N[k][w] = \alpha; C[k] = 1.0; M[k] = \alpha * |W|;$ 
2:  $t = 0;$ 
3:  $\gamma = 0;$ 
4: repeat
5:   for each label-document pair  $(y, d)$  do
6:      $t = t + 1;$ 
7:      $\rho = \frac{1}{1 + \lambda \cdot t}$ 
8:      $\gamma = \gamma + \alpha \cdot \frac{\rho}{n}$ 
9:     for each distinct word  $w$  in the document  $d$  do
10:       $N[y][w] = N[y][w] + \rho \cdot |w|_d \cdot (1 - p(Y = y|d, N, M, C, \gamma));$ 
11:       $M[y] = M[y] + \rho \cdot |w|_d \cdot (1 - p(Y = y|d, N, M, C, \gamma));$ 
12:      for  $k = 1, \dots, |Y| : k \neq y$  do
13:         $oldVal = N[k][w];$ 
14:         $N[k][w] = N[k][w] - \rho \cdot |w|_d \cdot ps(Y = k|d, N, M, C, \gamma);$ 
15:         $N[k][w] = \max(N[k][w], 0);$ 
16:         $M[k] = M[k] + (N[k][w] - oldVal);$ 
17:      end for
18:    end for
19:     $C[y] = C[y] + \rho \cdot (1 - p(Y = y|d, N, M, C, \gamma));$ 
20:    for  $k = 1, \dots, |Y| : k \neq y$  do
21:       $C[k] = C[k] - \rho \cdot p(Y = k|d, N, M, C, \gamma);$ 
22:       $C[k] = \max(C[k], 0);$ 
23:    end for
24:  end for
25: until convergence
26:  $\bar{N} = \text{Normalize}(N, \gamma);$ 
27:  $\bar{C} = \text{Normalize}(C, \gamma);$ 
28: return  $\bar{N}$  and  $\bar{C};$ 

```

Algorithm 3 Compute predictions $P(Y = k|d, N, M, C, \gamma)$ with Multinomial Naive Bayes. $|d|$ denotes the number of different words in the current document d and $|w|_d$ denotes the number of times word w appears in document d . The function "Logs2Probs" simply exponentiate the log values and then normalize.

Require: N, M, C, γ with non-negative values.

```

1:  $\forall k \ LogDC[k] = 0.0;$ 
2: for  $k = 1, \dots, |Y|$  do
3:    $LogDC[k] = \ln(C[k] + \gamma);$ 
4:    $sumW = 0;$ 
5:   for each distinct word  $w$  in the document  $d$  do
6:      $LogDC[k] = LogDC[k] + |w|_d \cdot \ln(N[y][w] + \gamma);$ 
7:      $sumW = sumW + |w|_d;$ 
8:   end for
9:    $LogDC[k] = LogDC[k] - sumW \cdot \ln(M[k] + \gamma \cdot |d|);$ 
10: end for
11: return  $\text{Logs2Probs}(\text{LogDC});$ 

```

Algorithm 4 sdEM for Multinomial Naive Bayes with the Hinge loss. $|d|$ denotes the number of different words in the current document d and $|w|_d$ denotes the number of times word w appears in document d . $|W|$ denotes the total number of different words in the corpus.

Require: D is randomly shuffled.

Require: α value as prior count for each word. Two values are considered $\alpha = 1$ and $\alpha = \ln |W|$.

```

1:  $\forall k, w \ N[k][w] = \alpha; C[k] = 1.0; M[k] = \alpha * |d|;$ 
2:  $t = 0;$ 
3:  $\gamma = 0;$ 
4: repeat
5:   for each label-document pair  $(y, d)$  do
6:      $t = t + 1;$ 
7:      $\rho = \frac{1}{1 + \lambda \cdot t}$ 
8:      $\gamma = \gamma + \alpha \cdot \frac{\rho}{n}$ 
9:      $\bar{y} = \arg \max_{y' \neq y} p(Y = y' | x);$ 
10:    if  $(\ln p(Y = y | x) - \ln p(Y = \bar{y} | x)) > 1$  then
11:      Go for the next document;
12:    end if
13:    for each distinct word  $w$  in the document  $d$  do
14:       $N[y][w] = N[y][w] + \rho \cdot |w|_d \cdot (1 - p(Y = y | d, N, M, C, \gamma));$ 
15:       $M[y] = M[y] + \rho \cdot |w|_d \cdot (1 - p(Y = y | d, N, M, C, \gamma));$ 
16:       $oldVal = N[\bar{y}][w];$ 
17:       $N[\bar{y}][w] = N[\bar{y}][w] - \rho \cdot |w|_d \cdot p(Y = \bar{y} | d, N, M, C, \gamma);$ 
18:       $N[\bar{y}][w] = \max(N[\bar{y}][w], 0);$ 
19:       $M[k] = M[k] + (N[\bar{y}][w] - oldVal);$ 
20:    end for
21:     $C[y] = C[y] + \rho \cdot (1 - p(Y = y | d, N, M, C, \gamma));$ 
22:     $C[\bar{y}] = C[\bar{y}] - \rho \cdot p(Y = \bar{y} | d, N, M, C, \gamma);$ 
23:     $C[\bar{y}] = \max(C[\bar{y}], 0);$ 
24:  end for
25: until convergence
26:  $\bar{N} = \text{Normalize}(N, \gamma);$ 
27:  $\bar{C} = \text{Normalize}(C, \gamma);$ 
28: return  $\bar{N}$  and  $\bar{C};$ 

```

Algorithm 5 sdEM for the LDA based classifier using the NCLL loss. $|d|$ denotes the number of different words in the current document d , $|w|_d$ denotes the number of times word w appears in document d and $|Z|$ denotes the number of hidden topics in the LDA model.

Require: D is randomly shuffled.

Require: η defines the prior for the "word per topic counts". In the experiments, it is fixed to $\eta = 0.1$.

```

1:  $\forall k, z, w \ N[k][z][w] = \frac{\eta}{|Z|}; C[k] = 1.0; M[k][z] = |W| \frac{\eta}{|Z|};$ 
2:  $t = 0;$ 
3:  $\gamma = 0;$ 
4: repeat
5:   for each label-document pair  $(y, d)$  do
6:      $t = t + 1;$ 
7:      $\rho = \frac{1}{1+\lambda \cdot t}$ 
8:      $\gamma = \gamma + \frac{\eta}{|Z|} \cdot \frac{\rho}{n}$ 
9:      $\text{OnlineLDA}(d, N[y], M[y], \rho, \gamma, \varpi = (1 - p(Y = y|d, N, M, C, \gamma)));$ 
10:    for  $k = 1, \dots, |Y| : k \neq y$  do
11:       $\text{OnlineLDA}(d, N[k], M[k], \rho, \gamma, \varpi = -p(Y = k|d, N, M, C, \gamma));$ 
12:    end for
13:     $C[y] = C[y] + \rho \cdot (1 - p(Y = y|d, N, M, C, \gamma));$ 
14:    for  $k = 1, \dots, |Y| : k \neq y$  do
15:       $C[k] = C[k] - \rho \cdot p(Y = k|d, N, M, C, \gamma);$ 
16:       $C[k] = \max(C[k], 0);$ 
17:    end for
18:  end for
19: until convergence
20:  $\bar{N} = \text{Normalize}(N, \gamma);$ 
21:  $\bar{C} = \text{Normalize}(C, \gamma);$ 
22: return  $\bar{N}$  and  $\bar{C};$ 

```

Algorithm 6 $\text{OnlineLDA}(d, N, M, \rho, \gamma, \varpi)$. The vector s would correspond to the expected sufficient statistics for d computed by online collapsed Gibbs sampling.

Require: $d, N, M, \rho, \gamma, \varpi$ properly computed.

```

1:  $s = \text{OnlineCollapsedGibbsSampling}(d, N, M, \gamma);$ 
2: for each distinct word  $w$  in the document  $d$  do
3:   for  $z=1, \dots, |Z|$  do
4:      $\text{oldVal} = N[z][w];$ 
5:      $N[z][w] = N[z][w] + \rho \cdot \varpi \cdot s[z][w];$ 
6:      $N[z][w] = \max(N[z][w], 0);$ 
7:      $M[z] = M[z] + (N[z][w] - \text{oldVal});$ 
8:   end for
9: end for

```

Algorithm 7 sdEM for the LDA based classifier using the Hinge loss. $|d|$ denotes the number of different words in the current document d , $|w|_d$ denotes the number of times word w appears in document d and $|Z|$ denotes the number of hidden topics in the LDA model.

Require: D is randomly shuffled.

Require: η defines the prior for the "word per topic counts". In the experiments, it is fixed to $\eta = 0.1$.

- 1: $\forall k, z, w N[k][z][w] = \frac{\eta}{|Z|}$; $C[k] = 1.0$; $M[k][z] = |W| \frac{\eta}{|Z|}$;
- 2: $t = 0$;
- 3: $\gamma = 0$;
- 4: **repeat**
- 5: **for** each label-document pair (y, d) **do**
- 6: $t = t + 1$;
- 7: $\rho = \frac{1}{1 + \lambda \cdot t}$
- 8: $\gamma = \gamma + \frac{\eta}{|Z|} \cdot \frac{\rho}{n}$
- 9: $\bar{y} = \arg \max_{y' \neq y} p(Y = y' | x)$;
- 10: **if** $(\ln p(Y = y | x) - \ln p(Y = \bar{y} | x)) > 1$ **then**
- 11: Go for the next document;
- 12: **end if**
- 13: OnlineLDA($d, N[y], M[y], \rho, \gamma, \varpi = (1 - p(Y = y | d, N, M, C, \gamma))$);
- 14: OnlineLDA($d, N[\bar{y}], M[\bar{y}], \rho, \gamma, \varpi = -p(Y = \bar{y} | d, N, M, C, \gamma)$);
- 15: $C[y] = C[y] + \rho \cdot (1 - p(Y = y | d, N, M, C, \gamma))$;
- 16: $C[\bar{y}] = C[\bar{y}] - \rho \cdot p(Y = \bar{y} | d, N, M, C, \gamma)$;
- 17: $C[\bar{y}] = \max(C[\bar{y}], 0)$;
- 18: **end for**
- 19: **until** convergence
- 20: $\bar{N} = \text{Normalize}(N, \gamma)$;
- 21: $\bar{C} = \text{Normalize}(C, \gamma)$;
- 22: **return** \bar{N} and \bar{C} ;
