

# Asynchronous Stochastic Variational Inference

Saad Mohamad  
Department of Computing,  
Bournemouth University  
Poole, UK  
smohamad@bournemouth.ac.uk

Abdelhamid Bouchachia  
Department of Computing,  
Bournemouth University  
Poole, UK  
abouchachia@bournemouth.ac.uk

Moamar Sayed-Mouchaweh  
Department of Informatics and  
Automatics, IMT Lille Douai  
Douai, France  
moamar.sayed-mouchaweh@  
imt-lille-douai.fr

## ABSTRACT

Stochastic variational inference (SVI) employs stochastic optimization to scale up Bayesian computation to massive data. Since SVI is at its core a stochastic gradient-based algorithm, horizontal parallelism can be harnessed to allow larger scale inference. We propose a lock-free parallel implementation for SVI which allows distributed computations over multiple slaves in an asynchronous style. We show that our implementation leads to linear speed-up while guaranteeing an asymptotic ergodic convergence rate  $O(1/\sqrt{T})$  given that the number of slaves is bounded by  $\sqrt{T}$  ( $T$  is the total number of iterations). The implementation is done in a high-performance computing (HPC) environment using message passing interface (MPI) for python (MPI4py). The extensive empirical evaluation shows that our parallel SVI is lossless, performing comparably well to its counterpart serial SVI with linear speed-up.

## KEYWORDS

Distributed Variational Inference, Probabilistic Modelling, Topic Mining, HPC

## 1 INTRODUCTION

Probabilistic models with latent variables have grown into a backbone in many modern machine learning applications such as text analysis, computer vision, time series analysis, network modelling, and others. The main challenge in such models is to compute the posterior distribution over some hidden variables encoding hidden structure in the observed data. Generally, computing the posterior is intractable and approximation is required. Markov chain Monte Carlo (MCMC) sampling has been the dominant paradigm for posterior computation. It constructs a Markov chain on the hidden variables whose stationary distribution is the desired posterior. Hence, the approximation is based on sampling for a long time to (hopefully) collect samples from the posterior [2].

Recently, variational inference (VI) has become widely used as a deterministic alternative approach to MCMC sampling. In general, VI tends to be faster than MCMC which makes it more suitable for problems with large data sets. VI turns the inference problem to an optimization problem by positing a simpler family of distributions and finding the member of the family that is closest to the true posterior distribution [19]. Hence, the inference task boils down to an optimization problem of a non-convex objective function. This allows us to bring sophisticated tools from optimization literature to tackle the performance problems. Recently, stochastic optimisation has been applied to VI in order to cope with massive data [8]. While VI requires repeatedly iterating over the whole data set before updating the variational parameters (parameters of the variational

objective), stochastic variational inference (SVI) updates the parameters every time a single data example is processed. Therefore, by the end of one pass through the dataset, the parameters will have been updated multiple times. Hence, the model parameters converge faster, while using less computational resources. The idea of SVI is to move the variational parameters at each iteration in the direction of a noisy estimate of the variational objective's natural gradient based on a couple of examples [8]. Following these stochastic gradients with certain conditions on the (decreasing) learning rate schedule, SVI provably converges to a local optimum [17].

Although stochastic optimization improves the performance of VI, its serial employment prevents scaling up the inference and harnessing distributed resources. Since, SVI is at its core a stochastic gradient-based optimisation algorithm, horizontal parallelism is straightforward. That is, computing stochastic gradients of a batch of data samples can be done locally in parallel given that the parameters update is synchronised. Such synchronisation limits the scalability by requiring slaves to send their stochastic gradients to the master prior to each parameter update. Hence, synchronous methods suffer from the curse of the last reducer; that is, a single slow slave can dramatically slow down the whole performance. Thus, asynchronous parallel optimization is an interesting alternative provided it maintains comparable convergence rate to its synchronous counterpart. Indeed, asynchronous parallel stochastic gradient-based optimisation algorithms have recently received broad attention [1, 6, 12, 16, 20].

Authors in [1] show that for smooth stochastic convex problems the asynchronisation effects are asymptotically negligible and order-optimal convergence results can be achieved. Since, the resulting objective function of the SVI is non-convex, we are particularly interested in the asynchronous parallel stochastic gradient algorithm (ASYSG) for smooth non-convex optimization [3]. A recent study [10] breaks the usual convexity assumption taken by [1]. Nonetheless, theoretical guarantees (convergence and speed-up) for many recent successes of ASYSG are provided. In this paper, we use the ASYSG algorithm proposed in [1] to come up with an asynchronous stochastic variational inference (ASYSVI) algorithm for a wide family of Bayesian models. We also adapt the theoretical studies of ASYSG for smooth non convex optimization from [10] to explain ASYSVI's convergence and speed-up properties. This paper proposes a novel contribution that allows to linearly speeding up SVI by distributing its stochastic natural gradient computations in an asynchronous way while guaranteeing an ergodic convergence rate  $O(1/\sqrt{T})$  under some assumptions. We take latent Dirichlet allocation (LDA) as a case study to empirically evaluate ASYSVI.

The rest of the paper is structured as follows. We briefly review variational and stochastic variational inference in Sec. 2. We describe our asynchronous stochastic variational inference algorithm along with its convergence analysis in Sec. 3. Latent Dirichlet allocation case study is developed in Sec. 4. Related work is discussed in Sec. 5. Empirical evaluation is presented in Sec. 6 and the paper concludes with a discussion in Sec. 7.

## 2 BACKGROUND

In the following, we derive the model family studied in this paper and review SVI. We follow the same pattern in [8].

**Model family.** Our family of models consists of three random variables: observations  $\mathbf{x} = \mathbf{x}_{1:n}$ , local hidden variables  $\mathbf{z} = \mathbf{z}_{1:n}$ , global hidden variables  $\boldsymbol{\beta}$  and fixed parameters  $\boldsymbol{\alpha}$ . The model assumes that the distribution of the  $n$  pairs of  $(\mathbf{x}_i, \mathbf{z}_i)$  is conditionally independent given  $\boldsymbol{\beta}$ . Further, their distribution and the prior distribution of  $\boldsymbol{\beta}$  are in an exponential family:

$$p(\boldsymbol{\beta}, \mathbf{x}, \mathbf{z} | \boldsymbol{\alpha}) = p(\boldsymbol{\beta} | \boldsymbol{\alpha}) \prod_{i=1}^n p(\mathbf{z}_i, \mathbf{x}_i | \boldsymbol{\beta}), \quad (1)$$

$$p(\mathbf{z}_i, \mathbf{x}_i | \boldsymbol{\beta}) = h(\mathbf{x}_i, \mathbf{z}_i) \exp(\boldsymbol{\beta}^T t(\mathbf{x}_i, \mathbf{z}_i) - a(\boldsymbol{\beta})), \quad (2)$$

$$p(\boldsymbol{\beta} | \boldsymbol{\alpha}) = h(\boldsymbol{\beta}) \exp(\boldsymbol{\alpha}^T t(\boldsymbol{\beta}) - a(\boldsymbol{\alpha})) \quad (3)$$

Here, we overload the notation for the base measures  $h(\cdot)$ , sufficient statistics  $t(\cdot)$  and log normalizer  $a(\cdot)$ . While the soul of the proposed approach is generic, for simplicity we assume a conjugacy relationship between  $(\mathbf{x}_i, \mathbf{z}_i)$  and  $\boldsymbol{\beta}$ . That is, the distribution  $p(\boldsymbol{\beta} | \mathbf{x}, \mathbf{z})$  is in the same family as the prior  $p(\boldsymbol{\beta} | \boldsymbol{\alpha})$ .

Note that this innocent looking family of models includes (but is not limited to) latent Dirichlet allocation [4], Bayesian Gaussian mixture, probabilistic matrix factorization, hidden Markov models, hierarchical linear and probit regression, and many Bayesian non-parametric models.

**Mean-field variational inference.** Variational inference (VI) approximates intractable posterior  $p(\boldsymbol{\beta}, \mathbf{z} | \mathbf{x})$  by positing a family of simple distributions  $q(\boldsymbol{\beta}, \mathbf{z})$  and find the member of the family that is closest to the posterior (closeness is measured with KL divergence). The resulting optimization problem is equivalent maximizing the evidence lower bound (ELBO):

$$\mathcal{L}(q) = E_q[\log p(\mathbf{x}, \mathbf{z}, \boldsymbol{\beta})] - E_q[\log p(\mathbf{z} | \boldsymbol{\beta})] \leq \log p(\mathbf{x}) \quad (4)$$

Mean-field is the simplest family of distribution, where the distribution over the hidden variables factorizes as follows:

$$q(\boldsymbol{\beta}, \mathbf{z}) = q(\boldsymbol{\beta} | \boldsymbol{\lambda}) \prod_{i=1}^n p(\mathbf{z}_i | \boldsymbol{\phi}_i) \quad (5)$$

Further, each variational distribution is assumed to come from the same family of the true one. Mean-field variational inference optimizes the new ELBO with respect to the local and global variational parameters  $\boldsymbol{\phi}$  and  $\boldsymbol{\lambda}$ :

$$\mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\phi}) = E_q \left[ \log \frac{p(\boldsymbol{\beta})}{q(\boldsymbol{\beta})} \right] + \sum_{i=1}^n E_q \left[ \log \frac{p(\mathbf{x}_i, \mathbf{z}_i | \boldsymbol{\beta})}{q(\mathbf{z}_i)} \right] \quad (6)$$

It iteratively updates each variational parameter holding the other parameters fixed. With the assumptions taken so far, each update

has a closed form solution. The local parameters are a function of the global parameters:

$$\boldsymbol{\phi}(\boldsymbol{\lambda}_t) = \arg \max_{\boldsymbol{\phi}} \mathcal{L}(\boldsymbol{\lambda}_t, \boldsymbol{\phi}) \quad (7)$$

We are interested in the global parameters which summarises the whole dataset (clusters in Bayesian Gaussian mixture, topics in LDA):

$$\mathcal{L}(\boldsymbol{\lambda}) = \max_{\boldsymbol{\phi}} \mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\phi}) \quad (8)$$

To find the optimal value of  $\boldsymbol{\lambda}$  given that  $\boldsymbol{\phi}$  is fixed, we compute the natural gradient of  $\mathcal{L}(\boldsymbol{\lambda})$  and set it to zero by setting:

$$\boldsymbol{\lambda}^* = \boldsymbol{\alpha} + \sum_{i=1}^n E_{\boldsymbol{\phi}_i(\boldsymbol{\lambda}_t)} [t(\mathbf{x}_i, \mathbf{z}_i)] \quad (9)$$

Thus, the new optimal global parameters are  $\boldsymbol{\lambda}_{t+1} = \boldsymbol{\lambda}^*$ . The algorithm works by iterating between computing the optimal local parameters given the global ones (Eq. (7)) and computing the optimal global parameters given the local ones (Eq. (9)).

**Stochastic variational inference.** Rather than analysing all the data to compute  $\boldsymbol{\lambda}^*$  at each iteration, stochastic optimization can be used. Assuming that the data is uniformly at random selected from the dataset, an unbiased noisy estimator of  $\mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\phi})$  can be developed based on a single data point:

$$\mathcal{L}_i(\boldsymbol{\lambda}, \boldsymbol{\phi}_i) = E_q \left[ \log \frac{p(\boldsymbol{\beta})}{q(\boldsymbol{\beta})} \right] + n E_q \left[ \log \frac{p(\mathbf{x}_i, \mathbf{z}_i | \boldsymbol{\beta})}{q(\mathbf{z}_i)} \right] \quad (10)$$

The unbiased stochastic approximation of the ELBO as a function of  $\boldsymbol{\lambda}$  can be written as follows:

$$\mathcal{L}_i(\boldsymbol{\lambda}) = \max_{\boldsymbol{\phi}_i} \mathcal{L}_i(\boldsymbol{\lambda}, \boldsymbol{\phi}_i) \quad (11)$$

Following the same step in the previous section, we end up with a noisy unbiased estimate of Eq. (9):

$$\hat{\boldsymbol{\lambda}} = \boldsymbol{\alpha} + n E_{\boldsymbol{\phi}_i(\boldsymbol{\lambda}_t)} [t(\mathbf{x}_i, \mathbf{z}_i)] \quad (12)$$

At each iteration, we move the global parameters a step-size  $\rho_t$  (learning rate) in the direction of the noisy natural gradient:

$$\boldsymbol{\lambda}_{t+1} = (1 - \rho_t) \boldsymbol{\lambda}_t + \rho_t \hat{\boldsymbol{\lambda}} \quad (13)$$

With certain conditions on  $\rho_t$ , the algorithm converges ( $\sum_{t=1}^{\infty} \rho_t = \infty, \sum_{t=1}^{\infty} \rho_t^2 < \infty$ ) [17].

## 3 ASYNCHRONOUS STOCHASTIC VARIATIONAL INFERENCE

In this section, we describe our proposed parallel implementation of ASYSVI on computer cluster and study its convergence and speed-up properties. The steps of the algorithm follows from the original ASYSG in [1].

### 3.1 Algorithm Description

ASYSVI is presented analogously to ASYSG in [10] but in the context of VI. The architecture of the computer network on which ASYSVI is run is known as the *star-shaped* network. In this network, a master machine maintains the global variational parameter  $\boldsymbol{\lambda}$ , whereas other machines serve as slaves which independently and simultaneously compute the local variational parameters  $\boldsymbol{\phi}$  and stochastic gradients of ELBO  $\mathcal{L}(\boldsymbol{\lambda})$ . The slaves only communicate with the master to exchange information in which they access the

---

**Algorithm 1** ASYSVI-Master

---

- 1: **Input:** number of iteration  $T$  and step-size  $\{\rho_t\}_{t=0, \dots, T-1}$
  - 2: **initialize:**  $\lambda^0$  randomly and  $t$  to 0
  - 3: **while** ( $t < T$ ) **do**
  - 4:   Aggregate  $M$  stochastic natural gradients  $\hat{\nabla} \mathcal{L}_1(\lambda^{t-\tau_{t,1}}), \dots, \hat{\nabla} \mathcal{L}_M(\lambda^{t-\tau_{t,M}})$  from the slaves
  - 5:   Average the  $M$  stochastic natural gradients.  $G_M^t = \sum_m \hat{\nabla} \mathcal{L}_m(\lambda^{t-\tau_{t,m}})$
  - 6:   Update the current estimate of the global variational parameter.  $\lambda^{t+1} = \lambda^t + \rho_t G_M^t$
  - 7:    $t = t + 1$
  - 8: **end while**
- 

**Algorithm 2** ASYSVI-Slave

---

- 1: **Input:** data size  $D$
  - 2: **while** (True) **do**
  - 3:   Sample a data point  $\mathbf{x}_i$  uniformly from the data set
  - 4:   Pull a global variational parameter  $\lambda^*$  from the master
  - 5:   Compute the local variational parameters  $\phi_i^*(\lambda^*)$  corresponding to the data point  $\mathbf{x}_i$  and the global variational parameter  $\lambda^*$ ,  $\phi_i^*(\lambda^*) = \arg \max_{\phi_i} \mathcal{L}_i(\lambda^*, \phi_i)$
  - 6:   Compute the stochastic natural gradient with respect to the global parameter  $\lambda$ ,  $\mathbf{g}_i(\lambda) = \alpha + DE_{\phi_i(\lambda)}[t(\mathbf{x}_i, \mathbf{z}_i)] - \lambda$
  - 7:   Push  $\mathbf{g}_i(\lambda^*)$  to the master
  - 8: **end while**
- 

state of the global variational parameter and provide the master with the stochastic gradients. These gradients are computed with respect to  $\lambda$  based on few (mini-batched or single) data points acquired from distributed sources. The master aggregates predefined amounts of stochastic gradients from slaves nonchalantly about the sources of the collected stochastic gradients. Then, it updates its current global variational parameter. The update step is performed as an atomic operation where slaves cannot read the value of the global variational parameter during this step. However, vertical parallelism can be achieved by adopting the ASYSG algorithm proposed in [16]. Furthermore, a hybrid horizontal-vertical parallelism could be achieved by combining the mechanism used in [15] with ASYSVI (more details in Sec. 5 and Sec. 7)

The key difference between ASYSVI and the synchronous parallel SVI is that ASYSVI does not lock the slaves until the master's update step is done. That is, the slaves might compute some stochastic gradients based on early value of the global variational parameter. By allowing delayed and asynchronous updates, one might expect slower convergence if any. In the next section, we apply the study of [10] on SVI to show that the effect of stochastic gradients delay will vanish asymptotically. The algorithms of ASYSVI-master and ASYSVI-slave are shown in Alg. 1 and Alg. 2. We denote by  $\tau_{t,m}$  the delays between the current iteration  $t$  and the one when the slave pulled the global variational parameter at which it computed the stochastic gradient.

### 3.2 Convergence Analysis

Following [1, 10], we take the same assumptions, but replace the gradient with the natural gradient:

- Unbiased gradient: the expectation of the stochastic natural gradient of Eq. (11) is equivalent to the natural gradient of Eq. (8):

$$\hat{\nabla} \mathcal{L}(\lambda) = E[\hat{\nabla} \mathcal{L}_i(\lambda)] \quad (14)$$

where  $\hat{\nabla}$  denotes natural gradient. This assumption already holds in SVI problems for the family of models shown in Sec. 2.

- Bounded variance: the variance of the stochastic natural gradient is bounded for all  $\lambda \in \mathcal{G}$ ,  $E[|\hat{\nabla} \mathcal{L}_i(\lambda) - \hat{\nabla} \mathcal{L}(\lambda)|^2] \leq \sigma^2$ . By applying SVI natural gradient, we end up with the following formulation:

$$E[|nE_{\phi_i(\lambda)}[t(\mathbf{x}_i, \mathbf{z}_i)] - \sum_{i=1}^n E_{\phi_i(\lambda)}[t(\mathbf{x}_i, \mathbf{z}_i)]|^2] \leq \sigma^2 \quad (15)$$

- Lipschitz-continuous gradient: the natural gradient is  $L$ -Lipschitz-continuous for all  $\lambda \in \mathcal{G}$  and  $\lambda' \in \mathcal{G}$ ,  $\|\hat{\nabla} \mathcal{L}(\lambda) - \hat{\nabla} \mathcal{L}(\lambda')\| \leq L\|\lambda - \lambda'\|$ . By applying SVI natural gradient, we end up with the following formulation:

$$\|\sum_{i=1}^n E_{\phi_i(\lambda)}[t(\mathbf{x}_i, \mathbf{z}_i)] - \lambda - \sum_{i=1}^n E_{\phi_i(\lambda')}[t(\mathbf{x}_i, \mathbf{z}_i)] + \lambda'\| \leq L\|\lambda - \lambda'\| \quad (16)$$

- Bounded delay: All delay variables  $\tau_{t,m}$  are bounded:

$$\max_{t,m} \tau_{t,m} \leq B \quad (17)$$

In addition to these assumptions, authors [1, 10] assume that each slave receives a stream of independent data points. Although this assumption might not be satisfied strictly in practice, we follow the same assumption for analysis purpose. Thus, the same theoretical results obtained by [10] can be applied for ASYSVI, namely, an ergodic convergence rate  $O(1/\sqrt{MT})$  provided that  $T$  is greater than  $O(B^2)$ . The results also show that, since the number of slaves is proportional to  $B$ , the ergodic convergence rate is achieved as long as the number of slaves is bounded by  $O(\sqrt{T/M})$ . Note that  $O(1/\sqrt{MT})$  is consistent with the serial stochastic gradient (SG) and the stochastic variational inference (SVI). Thus, ASYSG and ASYSVI allow for a linear speed-up if  $B \leq O(\sqrt{T/M})$ .

## 4 CASE STUDY: LATENT DIRICHLET ALLOCATION

Latent Dirichlet allocation (LDA) is an instance of the family of models described in Sec 2 where the global, local, observed variables and their distributions are set as follows:

- the global variables  $\{\beta\}_{k=1}^K$  are the topics in LDA. A topic is a distribution over the vocabulary, where the probability of a word  $w$  in topic  $k$  is denoted by  $\beta_{k,w}$ . Hence, the prior distribution of  $\beta$  is a Dirichlet distribution  $p(\beta) = \prod_k \text{Dir}(\beta_k; \eta)$
- the local variables are the topic proportions  $\{\theta_d\}_{d=1}^D$  and the topic assignments  $\{\{z_{d,w}\}_{d=1}^D\}_{w=1}^W$  which index the topic that generates the observations. Each document is associated with a topic proportion which is a distribution over topics,  $p(\theta) = \prod_d \text{Dir}(\theta_d; \alpha)$ . The assignments  $\{\{z_{d,w}\}_{d=1}^D\}_{w=1}^W$  are indices, generated by  $\theta_d$ , that couple topics with words,  $p(z_d|\theta) = \prod_w \theta_{d,z_{d,w}}$
- the observations  $x_d$  are the words of the documents which are assumed to be drawn from topics  $\beta$  selected by indices  $z_d$ ,  $p(x_d|z_d, \beta) = \prod_w \beta_{z_{d,w}, x_{d,w}}$

The basic idea of LDA is that documents are represented as random mixtures over latent topics, where each topic is characterized by a distribution over words [4]. LDA assumes the following generative process:

- 1 Draw topics  $\beta_k \sim \text{Dir}(\eta, \dots, \eta)$  for  $k \in \{1, \dots, K\}$
- 2 Draw topic proportions  $\theta_d \sim \text{Dir}(\alpha, \dots, \alpha)$  for  $d \in \{1, \dots, D\}$
- 2.1 Draw topic assignments  $z_{d,w} \sim \text{Mult}(\theta_d)$  for  $w \in \{1, \dots, W\}$
- 2.1.1 Draw word  $x_{d,w} \sim \text{Mult}(\beta_{z_{d,w}})$

According to Sec. 2, each variational distribution is assumed to come from the same family of the true one. Hence,  $q(\beta_k|\lambda_k) = \text{Dir}(\lambda_k)$ ,  $q(\theta_d|\gamma_d) = \text{Dir}(\gamma_d)$  and  $q(z_{d,w}|\phi_{d,w}) = \text{Mult}(\phi_{d,w})$ . To compute the stochastic natural gradient  $g_i$  in Alg. 2 for LDA, we need to find the sufficient statistic  $t(\cdot)$  presented in Eq. (2). By writing the likelihood of LDA in the form of Eq. (2), we obtain  $t(x_d, z_d) = \sum_{w=1}^W \mathbf{I}_{z_{d,w}, x_{d,w}}$ , where  $\mathbf{I}_{i,j}$  is equal to 1 for entry  $(i, j)$  and 0 for all the rest. Hence, the stochastic natural gradient  $g_i(\lambda_k)$  can be written as follows:

$$g_i(\lambda_k) = \eta + D \sum_{w=1}^W \phi_{i,w}^k \mathbf{I}_{k, x_{i,w}} - \lambda_k \quad (18)$$

Details on how to compute the local variational parameters  $\phi_i^*(\lambda^*)$  in Alg. 2 can be found in [8].

Having computed the elements needed to run ASYSVI's algorithms 1 and 2, we move to the convergence analysis. Since the data is assumed to be subsampled uniformly, the gradient unbiased assumption holds for LDA. We can always find a constant variable such that the bounded variance is satisfied. At the worst case, the variance of the stochastic natural gradient of LDA can be bounded by  $DW(\max_{i,w}(\phi_{i,w}^k)^2 - \min_{i',w'}(\phi_{i',w'}^k)^2)$ ,  $\forall k$ . Therefore, it can be bounded by  $O((DW)^2)$ . It is clear that the Lipschitz-continuous gradient can be satisfied for any class of the family models proposed in Sec. 2 and hence, for LDA. Finally, the bounded delay can be guaranteed through the implementation. Therefore, ASYSVI of LDA can converge since the aforementioned assumptions can be satisfied.

## 5 RELATED WORK

Few work has been proposed to scale variational inference to large datasets. We can distinguish two major classes. The first class is based on the Bayesian filtering approach [5, 9]. That is, the sequential nature of Bayes theorem is exploited to recursively update an approximation of the posterior. Particularly, variational inference is used between the updates to approximate the posterior which becomes the prior of the next step. Author in [9] employs forgetting factors to decay the contributions from old data points in favour of a new better one. The algorithm proposed in [5] considers a sequence of data batches. It iterates over the data points in the batch until convergence. The computation of the batches posterior is done in a distributed and asynchronous manner. That is, the algorithm applies VI by performing asynchronous Bayesian updates to the posterior as batches of data arrive continuously. Similar to our approach, master-slave architecture is used.

The second class of work is based on optimization [8, 14, 15]. As we already discussed, SVI proposed by [8] employs stochastic optimization to scale up Bayesian computation to massive data. SVI is inherently serial and requires the model parameters to fit in the memory of a single processor. Authors in [14] presents an inference algorithm, where the data is divided across several slaves and each of them perform VI updates in parallel. However at each iteration, the slaves are synchronized to combine their obtained parameters. Such synchronisation limits the scalability and decreases the speed of the update to that of the slowest slave. To avoid bulk synchronization, authors in [15] propose an asynchronous and lock-free update. In this update, vertical parallelism is adopted, where each processor asynchronously updates a subset of the parameters based on a subset of the data attributes. In contrast, we adopt horizontal parallelism update based on few (mini-batched or single) data points acquired from distributed sources. The update steps are, then, aggregated to form the global update. Note that the proposed approach can make use of the mechanism proposed by [15] to achieve a hybrid horizontal-vertical parallelism. On the contrary to [15], our approach is not customised for LDA and can be simply applied to any model from the family of models presented in Sec. 2

## 6 EXPERIMENTAL RESULTS

In the following, we demonstrate the usefulness of distributing the computation of SVI, mainly the speed-up advantages of ASYSVI. For this purpose, we compare the speed-up of ASYSVI LDA against serial SVI LDA (online LDA [7]). The two versions of LDA are evaluated on three datasets consisting of very large collections of documents. We also evaluate ASYSVI LDA in the streaming setting where new documents arrive in the form of stream. The algorithm implementation is available in Python <sup>1</sup>.

The performance evaluation is done using held-out perplexity as a measure of model fit. Perplexity is defined as the geometric mean of the inverse marginal probability of each word in the held-out set of documents [4]. To validate the speed-up properties, following [10], we compute the running time speed-up (TSP):

$$TSP = \frac{\text{running time of online LDA}}{\text{running time of asynchronous LDA}} \quad (19)$$

<sup>1</sup>Code will be uploaded later

Table 1: Parameters settings

Data sets	Enron emails				NYTimes news articles				Wikipedia articles			
<i>batch</i>	16	64	256	1024	16	64	256	1024	16	64	256	1024
$\kappa$	0.7	0.7	0.5	0.5	0.7	0.7	0.5	0.5	0.7	0.7	0.5	0.5
$\tau_0$	1024	24	24	1	1024	24	24	1	1024	1024	1024	1024
perplexity	5919	5348	5264	4771	11989	10156	9015	5501	1446	1390	1355	1332




Figure 1: Comparing ASYSVI LDA to online LDA on *Enron* dataset




Figure 2: TSP and RSP with respect to streaming samples on *Enron* dataset

**Datasets:** we perform all comparisons and evaluations on three corpora of documents. The first two corpora are available on [11]. The third corpus was used in [7].

- *Enron* emails: The corpus contains 39,861 email messages from about 150 users, mostly senior management of Enron.

Data is proceeded before usage by removing all words not in a vocabulary dictionary of 28,102 words.

- *NYTimes* news articles: The corpus contains 300,000 news articles from the *New York Times*. Data is proceeded before




Figure 3: TSP and RSP with respect to streaming samples on *NYTimes* dataset




Figure 4: TSP and RSP with respect to streaming samples on *Wikipedia* dataset

usage by removing all words not in a vocabulary dictionary of 102,660 words.

- *Wikipedia* articles: this corpus contains 1M documents downloaded from *Wikipedia*. Data is proceeded before usage by removing all words not in a vocabulary dictionary of 7,700 words.

**Settings the parameters:** In all experiments,  $\alpha$  and  $\eta$  are fixed at 0.01 and the number of topics  $K = 50$ . We evaluated a range of settings of the learning parameters,  $\kappa$ ,  $\tau_0$ , and  $batch$  on all four corpora. The parameters  $\kappa$  and  $\tau_0$ , defined in [7], control the learning step-size  $\rho_t$ . For each corpora, we use 29,861 emails from *Enron* dataset, 50,000 news articles from *NYTimes* dataset and 300,000 documents from *Wikipedia* dataset as training sets. We also reserve 5,000 documents as a validation set and another 5,000 documents as a testing set. The online LDA is run (one time per corpus) on the training sets

of each corpus for  $\kappa \in \{0.5, 0.7, 0.9\}$ ,  $\tau_0 \in \{1, 24, 256, 1024\}$ , and  $batch \in \{16, 64, 256, 1024\}$ . Table 1 summarises the best settings of each  $batch$  along with the resulting perplexity on the test set for each corpus.

**Comparing Serial online LDA and asynchronous LDA:** for each dataset, we set the parameters setting that give the best performance (least perplexity). ASYSVI LDA is then compared against serial SVI LDA using the same parameters setting.

The code is implemented on high-performance computing (HPC) environment using message passing interface (MPI) for python (MPI4py). The cluster consists of 10 nodes, excluding the head node, with each node is a four-core processor. We run AYSVI LDA on *Enron* dataset for number of workers  $nW \in \{2, 4, 6, 8, 9, 18, 27, 36\}$ ,  $B$  is set to 5. The number of employed nodes is equal to  $nW$  as long as  $nW$  is less than 9. As  $nW$  becomes higher than the available

nodes, the processors' cores of nodes are employed as slaves until all the cores of all the nodes are used i.e.,  $9 \times 4 = 36$ . Since the batch size is fixed to 1024, each slave processes a batch of data of size  $S = 1024/M$  per iteration, where  $M$  is fixed to 36. Thus, the gradient computed by each slave will be multiplied by  $D/S$ . Hence, line number 6 of Alg. 2 becomes,  $\mathbf{g}_i(\lambda) = \alpha + (D/S)E_{\phi_i(\lambda)}[t(\mathbf{x}_i, \mathbf{z}_i)] - \lambda$ .

Figure 1 summarises the total speed up (i.e., TSP measured at the end of the algorithm) as well as the ratio of serial LDA perplexity to parallel LDA (RSP) on the test set for *Enron* dataset. It shows TSP and RSP results with respect to the number of slaves. It is clear that as long as each node is assigned one slave, the speed-up is linear which demonstrates the convergence analysis done in Sec. 3. Linear speed-up slowly converts to sub-linear as solo machines host more than one slave. The main reason of such behaviour is the communication delay caused by the increase of the network traffic. Hence, TSP is affected by the hardware. The communication cost starts affecting the algorithm speed-up when it becomes comparable to the local computation. Hence, increasing the local computation by increasing the batch size can be adopted to soften the communication effect. However, this decreases the convergence rate and increase local memory load. Hence, a balanced trade-off should be considered. RSP in Figure 1 shows that although the speed of online LDA has been increased up to 15 times, performance is not seriously affected. We also evaluate TSP and RSP on *NYTimes* and *Wikipedia* for  $nW = 27$ . The processing speed of Online LDA on *NYTimes* has been increased 19.29 times,  $TSP = 19.29$ , with slight loss of performance,  $RSP = 0.97$ . For *Wikipedia*,  $TSP = 18.58$  and  $RSP = 0.94$ .

Figures 2, 3 and 4 present TSP and RSP with respect to streaming samples from *Enron*, *NYTimes* and *Wikipedia* datasets. These figures show the performance of ASYSVI in a true online setting where the algorithm continually collects samples from the hard driver for the case of *Enron* and *NYTimes* or by downloading online in the case of *Wikipedia*. The perplexity is obtained online on the coming batches before being used to update the model parameters. The plots in the figures are slightly softened using a low-pass filter in order to make them easy to read. These plots show that the speed-up becomes invariant as more samples are processed. The poor speed-up in the beginning is normally caused by initialization and loading process. It can be noticed that the performance of ASYSVI LDA suffers at the beginning then it becomes comparable to online LDA after certain number of iterations. This behaviour can be explained by the convergence condition shown in Sec. 3 ( $T$  is greater than  $O(B^2)$ ). Thus, as the number of iterations increases, the convergence of ASYSVI LDA is guaranteed and its performance becomes comparable to that of online LDA. Hence, RSP approaches 1.

## 7 CONCLUSION AND DISCUSSION

We have introduced ASYSVI, an asynchronous parallel implementations for SVI on computer cluster. ASYSVI leads to linear speed-up, while guaranteeing an asymptotic convergence rate given some assumptions involving the number of the slaves and iterations. Empirical results using latent Dirichlet allocation topic model as a case study have demonstrated the advantages of ASYSVI over SVI, particularly with respect to the key issue of speeding-up the computation while maintaining comparable performance to SVI.

In future work, vertical parallelism can be adopted along with the proposed horizontal one leading to a hybrid horizontal-vertical parallelism. In such case, multi-core processors will be used for the vertical parallelism, while horizontal parallelism is achieved on a multi-node machine. Another avenue of interest is to derive an algorithm for streaming, distributed, asynchronous inference where the number of instances is not known. Moreover, it is interesting to apply ASYSVI on very large scale problems and particularly on other models of the family discussed in Sec. 2 and studying the effect of the statistical properties of those models.

## ACKNOWLEDGMENT

S. Mohamad and A. Bouchachia were supported by the European Commission under the Horizon 2020 Grant 687691 related to the project: PROTEUS: Scalable Online Machine Learning for Predictive Analytics and Real-Time Interactive Visualization.

## REFERENCES

- [1] Alekh Agarwal and John C Duchi. 2011. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems*. 873–881.
- [2] Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I Jordan. 2003. An introduction to MCMC for machine learning. *Machine learning* 50, 1-2 (2003), 5–43.
- [3] Dimitri P Bertsekas and John N Tsitsiklis. 1989. *Parallel and distributed computation: numerical methods*. Vol. 23. Prentice hall Englewood Cliffs, NJ.
- [4] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *the Journal of machine Learning research* 3 (2003), 993–1022.
- [5] Tamara Broderick, Nicholas Boyd, Andre Wibisono, Ashia C Wilson, and Michael I Jordan. 2013. Streaming variational bayes. In *Advances in Neural Information Processing Systems*. 1727–1735.
- [6] Hamid Reza Feysmahdavian, Arda Aytekin, and Mikael Johansson. 2016. An asynchronous mini-batch algorithm for regularized stochastic optimization. *IEEE Trans. Automat. Control* 61, 12 (2016), 3740–3754.
- [7] Matthew Hoffman, Francis R Bach, and David M Blei. 2010. Online learning for latent dirichlet allocation. In *advances in neural information processing systems*. 856–864.
- [8] Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. 2013. Stochastic variational inference. *The Journal of Machine Learning Research* 14, 1 (2013), 1303–1347.
- [9] Antti Honkela and Harri Valpola. 2003. On-line variational Bayesian learning. In *4th International Symposium on Independent Component Analysis and Blind Signal Separation*. 803–808.
- [10] Xiangru Lian, Yijun Huang, Yuncheng Li, and Ji Liu. 2015. Asynchronous parallel stochastic gradient for nonconvex optimization. In *Advances in Neural Information Processing Systems*. 2737–2745.
- [11] M. Lichman. 2013. UCI Machine Learning Repository. (2013). <http://archive.ics.uci.edu/ml>
- [12] Horia Mania, Xinghao Pan, Dimitris Papaliopoulos, Benjamin Recht, Kannan Ramchandran, and Michael I Jordan. 2015. Perturbed iterate analysis for asynchronous stochastic optimization. *arXiv preprint arXiv:1507.06970* (2015).
- [13] James McInerney, Rajesh Ranganath, and David M Blei. 2015. The Population Posterior and Bayesian Inference on Streams. *arXiv preprint arXiv:1507.05253* (2015).
- [14] Willie Neiswanger, Chong Wang, and Eric Xing. 2015. Embarrassingly Parallel Variational Inference in Nonconjugate Models. *arXiv preprint arXiv:1510.04163* (2015).
- [15] Parameswaran Raman, Jiong Zhang, Hsiang-Fu Yu, Shihao Ji, and SVN Vishwanathan. 2016. Extreme Stochastic Variational Inference: Distributed and Asynchronous. *arXiv preprint arXiv:1605.09499* (2016).
- [16] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. 2011. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*. 693–701.
- [17] Herbert Robbins and Sutton Monro. 1951. A stochastic approximation method. *The annals of mathematical statistics* (1951), 400–407.
- [18] Lucas Theis and Matthew D Hoffman. 2015. A trust-region method for stochastic variational inference with applications to streaming data. *arXiv preprint arXiv:1505.07649* (2015).
- [19] Martin J Wainwright and Michael I Jordan. 2008. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning* 1, 1-2 (2008), 1–305.
- [20] Ruiliang Zhang and James T Kwok. 2014. Asynchronous Distributed ADMM for Consensus Optimization. In *ICML*. 1701–1709.