

---

# A Stochastic Approximation Approach for Efficient Decentralized Optimization on Random Networks

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 A challenging problem in decentralized optimization is to develop algorithms with  
2 fast convergence on random and time varying topologies under unreliable and  
3 bandwidth-constrained communication network. This paper studies a stochastic  
4 approximation approach with a Fully Stochastic Primal Dual Algorithm (FSPDA)  
5 framework. Our framework relies on a novel observation that randomness in time  
6 varying topology can be incorporated in a stochastic augmented Lagrangian for-  
7 mulation, whose expected value admits saddle points that coincide with stationary  
8 solutions of the decentralized optimization problem. With the FSPDA framework,  
9 we develop two new algorithms supporting efficient sparsified communication on  
10 random time varying topologies — FSPDA-SA allows agents to execute multiple  
11 local gradient steps depending on the time varying topology to accelerate conver-  
12 gence, and FSPDA-STORM further incorporates a variance reduction step to improve  
13 sample complexity. For problems with smooth (possibly non-convex) objective  
14 function, within  $T$  iterations, we show that FSPDA-SA (resp. FSPDA-STORM) finds  
15 an  $\mathcal{O}(1/\sqrt{T})$ -stationary (resp.  $\mathcal{O}(1/T^{2/3})$ ) solution. Numerical experiments show  
16 the benefits of the FSPDA algorithms.

## 17 1 Introduction

18 Consider  $n$  agents that communicate on an undirected and connected graph/network  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with  
19  $\mathcal{V} = [n] := \{1, \dots, n\}$ ,  $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ . Each agent  $i \in [n]$  has access to a continuously differentiable  
20 (possibly non-convex) local objective function  $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$  and maintains a local decision variable  
21  $\mathbf{x}_i \in \mathbb{R}^d$ . Denote  $\mathbf{x} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top]^\top \in \mathbb{R}^{nd}$ . Our aim is to tackle:

$$\min_{\mathbf{x} \in \mathbb{R}^{nd}} \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}_i) \quad \text{s.t.} \quad \mathbf{x}_i = \mathbf{x}_j, \forall (i, j) \in \mathcal{E}. \quad (1)$$

22 In other words, (1) seeks a  $\mathbf{x}^* \in \mathbb{R}^d$  that minimizes  $F(\mathbf{x}) := (1/n) \sum_{i=1}^n f_i(\mathbf{x})$ . We are interested  
23 in the stochastic optimization setting where each  $f_i(\mathbf{x}_i)$  is given by (with slight abuse of notation)

$$f_i(\mathbf{x}_i) := \mathbb{E}_{\xi_i \sim \mathbb{P}_i} [f_i(\mathbf{x}_i; \xi_i)] \quad (2)$$

24 where  $\mathbb{P}_i$  represents the  $i$ -th data distribution. Problem (1) is relevant to the distributed learning  
25 problem especially in the decentralized case where a central server is absent. Prior works [Nedic and  
26 Ozdaglar, 2009, Lian et al., 2017, Nedic et al., 2017, Qu and Li, 2017] demonstrated that *decentralized*  
27 algorithms can tackle (1) efficiently through repeated message exchanges among the neighbors and  
28 local stochastic gradient updates.

29 Towards an efficient decentralized algorithm for (1), an important direction is to consider a *time*  
30 *varying graph topology* setting where the *active edge set* in  $\mathcal{G}$  changes over time. This is a generic  
31 setting covering cases when the communication links are unreliable, or the agents choose not to  
32 communicate in a certain round (a.k.a. local updates) [Koloskova et al., 2019a, Nadiradze et al., 2021].

Prior Works	SG	TV	w/o BH	Rate
Prox-GPDA [Hong et al., 2017]	✗	✗	✓	Asympt.
NEXT [Lorenzo and Scutari, 2016]	✗	✓	✓	Asympt.
DSGD [Koloskova et al., 2020]	✓	✓	✗	$\mathcal{O}(\sigma/\sqrt{nT})$
Swarm-SGD [Nadiradze et al., 2021]	✓	✓	✗	$\mathcal{O}(\sigma^2/\sqrt{T})$
CHOCO-SGD [Koloskova et al., 2019a]	✓	✗ <sup>‡</sup>	✗	$\mathcal{O}(\sigma/\sqrt{nT})$
Decen-Scaffnew [Mishchenko et al., 2022]	✓	✗ <sup>†</sup>	✓	$\mathcal{O}(\sigma/\sqrt{nT})$
Local-GT [Liu et al., 2024]	✓	✗ <sup>†</sup>	✓	$\mathcal{O}(\sigma/\sqrt{nT})$
LED [Alghunaim, 2024]	✓	✗ <sup>†</sup>	✓	$\mathcal{O}(\sigma/\sqrt{nT})$
FSPDA-SA (This Work)	✓	✓	✓	$\mathcal{O}(\sigma/\sqrt{nT})$
FSPDA-STORM (This Work)	✓	✓	✓	$\mathcal{O}(\sigma^{2/3}/T^{2/3})$

Table 1: Comparison of decentralized algorithms for **non-convex** optimization. In the table, ‘SG’ is ‘Stochastic Gradient’, ‘TV’ is ‘Time Varying Graph’, ‘w/o BH’ is ‘Without Bounded Heterogeneity’, and ‘Rate’ is the expected squared gradient norm  $\mathbb{E}[\|\nabla F(\bar{\mathbf{x}})\|^2]$  after  $T$  iterations. Note that  $\sigma^2$  is the variance of stochastic gradient. <sup>‡</sup>CHOCO-SGD incorporates broadcast gossip as a special case of compression. <sup>†</sup>ProxSkip, Local-GT, LED consider local updates with periodic communication.

By assuming that a random topology is drawn at each iteration, the convergence of decentralized stochastic gradient (DSGD) has been studied in [Lobel and Ozdaglar, 2010, Nadiradze et al., 2021] and is later on unified by [Koloskova et al., 2020] with tighter bounds for local updates, periodic sampling, etc. An alternative [Ram et al., 2010] is to analyze DSGD for the  $B$ -connectivity setting which requires the union of every  $B$  consecutive time varying topologies to yield a connected graph. Nevertheless, these works focused on vanilla DSGD that may have slow convergence (in transient stage) and is limited to bounded data heterogeneity. The prior restrictions can be relaxed using advanced algorithms such as gradient tracking [Qu and Li, 2017], EXTRA [Shi et al., 2015] and primal-dual framework [Hong et al., 2017, Hajinezhad and Hong, 2019, Yi et al., 2021].

As noted by [Koloskova et al., 2021], analyzing the convergence of sophisticated algorithms with time varying topology, such as gradient tracking [Qu and Li, 2017] is challenging due to the non-symmetric product of two (or more) mixing matrices. Existing works considered various restrictions on the time varying topology  $\mathcal{G}^{(t)} = (\mathcal{V}, \mathcal{E}^{(t)})$  and/or the problem (1): [Koloskova et al., 2021, Liu et al., 2024] studied gradient tracking with local updates that essentially takes  $\mathcal{E}^{(t)} = \mathcal{E}$  periodically and  $\mathcal{E}^{(t)} = \emptyset$  otherwise, also see [Mishchenko et al., 2022, Guo et al., 2023, Alghunaim, 2024] for a similar result and note that such algorithms require extra synchronization overhead; [Kovalev et al., 2021, 2024] considered a setting where  $\mathcal{G}^{(t)}$  is connected for any  $t$ ; [Nedic et al., 2017, Li and Lin, 2024] focused on (accelerated) gradient tracking with deterministic gradient when  $F(\mathbf{x})$  is (strongly) convex; [Lorenzo and Scutari, 2016] also considered deterministic gradient with possibly non-convex  $F(\mathbf{x})$  but only provides asymptotic convergence guarantees; [Lei et al., 2018, Yau and Wai, 2023] considered asymptotic convergence guarantees in the case of strictly (or strongly) convex  $F(\mathbf{x})$ . We provide a non-exhaustive list summarizing the convergence of existing works in Table 1.

The above discussion highlights a gap in the existing literature —

*Is there any algorithm that achieves fast convergence on time varying (random) topology?*

This paper gives an affirmative answer through developing the Fully Stochastic Primal Dual Algorithm (FSPDA) framework that leads to efficient decentralized algorithms tackling (1) in its general form. The framework features the design of a new stochastic augmented Lagrangian function.

As pointed out by [Chang et al., 2020], many decentralized algorithms (including gradient tracking) can be interpreted as primal-dual algorithms finding a saddle point of the augmented Lagrangian function. However, its extension to time varying topology is not straightforward due to the inconsistency in dual variables updates. To overcome this challenge, we propose a stochastic equality constrained reformulation of (1) to model randomness in topology. Then, the latter yields a stochastic augmented Lagrangian function. Applying stochastic approximation (SA) to solve the latter leads to the FSPDA framework. Our contributions are

- We propose two new algorithms: (i) FSPDA-SA is derived by vanilla SA that applies primal-dual stochastic gradient descent-ascent on the stochastic augmented Lagrangian, (ii) FSPDA-STORM uses an additional control variate / momentum term to reduce the drift term's variance in a recursive manner. Both algorithms are fully stochastic as the random time varying topology is treated as a part of randomness. Additionally, our framework supports sparsified communication, i.e., the agents can choose to communicate a subset of primal coordinates at each iteration.
- We show that after  $T$  iterations, FSPDA-SA (resp. FSPDA-STORM) finds in expectation a solution whose squared gradient norm is  $\mathcal{O}(1/\sqrt{T})$  (resp.  $\mathcal{O}(1/T^{2/3})$ ). The convergence analysis is derived from a new Lyapunov function design that involves an unsigned inner product term and incorporates a variance condition on the random time varying topologies. Interestingly, we show empirically that using momentum in dual updates benefits the consensus error convergence.
- We also demonstrate that both FSPDA-SA and FSPDA-STORM can be implemented in a fully asynchronous manner, i.e., the agents can communicate and compute at different time slots, and supports local update as the algorithms allow for arbitrary time varying topology. That said, we remark that the convergence rates with local updates of FSPDA-SA and FSPDA-STORM are only suboptimal.

We provide numerical experiments to show that FSPDA-SA and FSPDA-STORM outperform existing algorithms in terms of iteration and communication complexity.

**Notations.** Let  $\mathbf{W} \in \mathbb{R}^{d \times d}$  be a symmetric (not necessarily positive semidefinite) matrix, the  $\mathbf{W}$ -weighted (semi) inner product of vectors  $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$  is denoted as  $\langle \mathbf{a} | \mathbf{b} \rangle_{\mathbf{W}} := \mathbf{a}^\top \mathbf{W} \mathbf{b}$ . Similarly, the  $\mathbf{W}$ -weighted (semi) norm is denoted by  $\|\mathbf{a}\|_{\mathbf{W}}^2 := \langle \mathbf{a} | \mathbf{a} \rangle_{\mathbf{W}}$ . The subscript notation is omitted for  $\mathbf{I}$ -weighted inner products. For any square matrix  $\mathbf{X}$ ,  $(\mathbf{X})^\dagger$  denotes its pseudo inverse.

## 2 The Fully Stochastic Primal Dual Algorithm (FSPDA) Framework

This section develops the FSPDA framework for tackling (1) and describes two variants of the framework leading to decentralized stochastic optimization of (1). Let  $\tilde{\mathbf{A}} \in \{-1, 0, 1\}^{|\mathcal{E}| \times n}$  be an incidence matrix of  $\mathcal{G}$ . By defining  $\mathbf{A} = \tilde{\mathbf{A}} \otimes \mathbf{I}_d \in \{-1, 0, 1\}^{|\mathcal{E}|d \times nd}$ , we observe that the consensus constraint in (1) is equivalent to  $\mathbf{A}\mathbf{x} = \mathbf{0}$ .

Our first step is to model the randomness in the time varying topology using the random variable (r.v.)  $\xi_a \sim \mathbb{P}_a$ . For each realization  $\xi_a$ , we define the random incidence matrix  $\mathbf{A}(\xi_a) := \mathbf{I}(\xi_a)\mathbf{A} \in \{-1, 0, 1\}^{|\mathcal{E}|d \times nd}$  where  $\mathbf{I}(\xi_a) \in \{0, 1\}^{|\mathcal{E}|d \times |\mathcal{E}|d}$  is a binary diagonal matrix. In addition to selecting each edge of  $\mathcal{G}$  randomly,  $\mathbf{I}(\xi_a)$  selects a random subset of  $d$  coordinates. As we will see later, this allows our approach to simultaneously achieve random sparsification for communication compression.

Assume that  $\mathbb{E}_{\xi_a \sim \mathbb{P}_a}[\mathbf{I}(\xi_a)]$  is a positive diagonal matrix, (1) is equivalent to:

$$\min_{\mathbf{x} \in \mathbb{R}^{nd}} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\xi_i \sim \mathbb{P}_i} [f_i(\mathbf{x}_i; \xi_i)] \quad \text{s.t.} \quad \mathbb{E}_{\xi_a \sim \mathbb{P}_a} [\mathbf{A}(\xi_a)] \mathbf{x} = \mathbf{0}. \quad (3)$$

Denote  $\xi = (\xi_1, \dots, \xi_n, \xi_a)$ , FSPDA hinges on the following *augmented Lagrangian* function of (3):

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) &:= \mathbb{E}_{\xi} [\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}; \xi)] \\ \text{with } \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}; \xi) &:= \sum_{i=1}^n f_i(\mathbf{x}_i; \xi_i) + \tilde{\eta} \langle \boldsymbol{\lambda} | \mathbf{A}(\xi_a) \mathbf{x} \rangle + \frac{\tilde{\gamma}}{2} \|\mathbf{A}(\xi_a) \mathbf{x}\|^2, \end{aligned} \quad (4)$$

where  $\tilde{\eta} > 0, \tilde{\gamma} > 0$  are penalty parameters. It can be verified that the saddle points of  $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$  correspond to the KKT points of (1) [Bertsekas, 2016]. For brevity, in the rest of this paper, we may drop the subscript in  $\xi$  whenever the notation is clear from the context.

FSPDA is developed from applying stochastic approximation (SA) to seek a saddle point of (4). By recognizing  $\mathbf{A}(\xi)^\top \mathbf{A}(\xi) = \mathbf{A}^\top \mathbf{A}(\xi)$ , we consider the stochastic gradients:

$$\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}; \xi) := \nabla \mathbf{f}(\mathbf{x}; \xi) + \tilde{\eta} \mathbf{A}^\top \boldsymbol{\lambda} + \tilde{\gamma} \mathbf{A}^\top \mathbf{A}(\xi) \mathbf{x}, \quad \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}; \xi) := \tilde{\eta} \mathbf{A}(\xi) \mathbf{x}, \quad (5)$$

where  $\nabla \mathbf{f}(\mathbf{x}; \xi) = [\nabla f_1(\mathbf{x}_1; \xi_1); \dots; \nabla f_n(\mathbf{x}_n; \xi_n)] \in \mathbb{R}^{nd}$ . Notice that to facilitate algorithm development, we have taken a deterministic  $\mathbf{A}$  for the term in  $\nabla_{\mathbf{x}} \mathcal{L}$  related to  $\boldsymbol{\lambda}$ . Now observe the  $i$ th  $d$ -dimensional block of  $\mathbf{A}^\top \mathbf{A}(\xi) \mathbf{x}$  which can be aggregated within  $\mathcal{N}_i(\xi)$  the neighborhood of the  $i$ th agent as:

$$[\mathbf{A}^\top \mathbf{A}(\xi) \mathbf{x}]_i = \sum_{j \in \mathcal{N}_i(\xi)} \mathbf{C}_{ij}(\xi) (\mathbf{x}_j - \mathbf{x}_i), \quad (6)$$

where  $\mathbf{C}_{ij}(\xi) \in \{0, 1\}^{d \times d}$  is diagonal and depends on the selected coordinates for the edge  $(i, j)$  under randomness  $\xi$ . Eq. (6) only relies on  $\mathbf{x}_j$  from neighbor  $j$  that is connected on the time varying

topology  $\mathcal{G}(\xi)$ . For illustration, an example of the above random graph model is given by Figure 3 in Appendix A. Importantly, (5) shows that with the stochastic augmented Lagrangian function, the time varying topology can be treated implicitly as a part of the randomness in the stochastic primal-dual gradients. The framework is thus described as being *fully stochastic* as in [Bianchi et al., 2021], and departs from [Liu et al., 2024, Alghunaim, 2024] that treat the topology as fixed during the derivation of primal-dual algorithm(s). From (5), (6), we derive *two* variants of FSPDA.

**FSPDA-SA Algorithm.** The first variant of FSPDA is derived from a direct application of stochastic gradient descent-ascent (SGDA) updates. Take  $\alpha > 0, \beta > 0$  as the step sizes, we have

$$\mathbf{x}^{t+1} = \mathbf{x}^t - \alpha \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^t, \boldsymbol{\lambda}^t; \xi^t), \quad \boldsymbol{\lambda}^{t+1} = \boldsymbol{\lambda}^t + \beta \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}^t, \boldsymbol{\lambda}^t; \xi^t). \quad (7)$$

Taking the variable substitution  $\hat{\boldsymbol{\lambda}} := \mathbf{A}^\top \boldsymbol{\lambda}$  yields the following recursion:

**FSPDA-SA:** for any  $t \geq 0$  and any  $i \in [n]$ ,

$$\mathbf{x}_i^{t+1} = \mathbf{x}_i^t - \alpha \nabla f_i(\mathbf{x}_i^t; \xi_i^t) - \eta \hat{\boldsymbol{\lambda}}_i^t + \gamma \sum_{j \in \mathcal{N}_i(\xi_a^t)} \mathbf{C}_{ij}(\xi_a^t)(\mathbf{x}_j^t - \mathbf{x}_i^t), \quad (8a)$$

$$\hat{\boldsymbol{\lambda}}_i^{t+1} = \hat{\boldsymbol{\lambda}}_i^t + \beta \sum_{j \in \mathcal{N}_i(\xi_a^t)} \mathbf{C}_{ij}(\xi_a^t)(\mathbf{x}_j^t - \mathbf{x}_i^t). \quad (8b)$$

Note that  $\mathbf{x}^0, \hat{\boldsymbol{\lambda}}^0$  can be initialized arbitrarily.

**FSPDA-STORM Algorithm.** The second variant of FSPDA reduces the variance of the stochastic gradient term in (5) using the recursive momentum variance reduction technique [Cutkosky and Orabona, 2019]. Herein, the key idea is to utilize a control variate in estimating the (primal-dual) gradients of  $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ . Take  $\alpha, \beta > 0$  and  $a_x, a_\lambda \in [0, 1]$  as the momentum parameters, we have  $\mathbf{x}^{t+1} = \mathbf{x}^t - \alpha \mathbf{m}_x^t, \boldsymbol{\lambda}^{t+1} = \boldsymbol{\lambda}^t + \beta \mathbf{m}_\lambda^t$  as the primal-dual updates, and

$$\begin{aligned} \mathbf{m}_x^{t+1} &= \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^{t+1}, \boldsymbol{\lambda}^{t+1}; \xi^{t+1}) + (1 - a_x)(\mathbf{m}_x^t - \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^t, \boldsymbol{\lambda}^t; \xi^{t+1})), \\ \mathbf{m}_\lambda^{t+1} &= \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}^{t+1}, \boldsymbol{\lambda}^{t+1}; \xi^{t+1}) + (1 - a_\lambda)(\mathbf{m}_\lambda^t - \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}^t, \boldsymbol{\lambda}^t; \xi^{t+1})). \end{aligned} \quad (9)$$

The aim of  $\mathbf{m}_x^{t+1}$  is to estimate  $\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^{t+1}, \boldsymbol{\lambda}^{t+1})$ . Now, instead of the straightforward estimator  $\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^{t+1}, \boldsymbol{\lambda}^{t+1}; \xi^{t+1})$ , we include an extra zero-mean term  $\mathbf{m}_x^t - \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^t, \boldsymbol{\lambda}^t; \xi^{t+1})$  to reduce the variance of the stochastic gradient estimation. The latter is a control variate that is computed recursively. Particularly, it has been shown in [Cutkosky and Orabona, 2019] that it can effectively reduce variance with a carefully designed parameter  $a_x$ , provided that the stochastic gradient map satisfies a mean-square Lipschitz condition. We summarize the algorithm as follows.

**FSPDA-STORM:** for any  $t \geq 0$  and any  $i \in [n]$ ,

$$\mathbf{x}_i^{t+1} = \mathbf{x}_i^t - \alpha \mathbf{m}_{x,i}^t, \quad (10a)$$

$$\hat{\boldsymbol{\lambda}}_i^{t+1} = \hat{\boldsymbol{\lambda}}_i^t + \beta \mathbf{m}_{\lambda,i}^t, \quad (10b)$$

$$\begin{aligned} \mathbf{m}_{x,i}^{t+1} &= (1 - a_x) [\mathbf{m}_{x,i}^t + \nabla f_i(\mathbf{x}_i^t; \xi_i^{t+1}) - \eta \hat{\boldsymbol{\lambda}}_i^t + \gamma \sum_{j \in \mathcal{N}_i(\xi_a^{t+1})} \mathbf{C}_{ij}(\xi_a^{t+1})(\mathbf{x}_j^t - \mathbf{x}_i^t)] \\ &\quad + \nabla f_i(\mathbf{x}_i^{t+1}; \xi_i^{t+1}) - \eta \hat{\boldsymbol{\lambda}}_i^{t+1} + \gamma \sum_{j \in \mathcal{N}_i(\xi_a^{t+1})} \mathbf{C}_{ij}(\xi_a^{t+1})(\mathbf{x}_j^{t+1} - \mathbf{x}_i^{t+1}) \end{aligned} \quad (10c)$$

$$\begin{aligned} \mathbf{m}_{\lambda,i}^{t+1} &= (1 - a_\lambda) [\mathbf{m}_{\lambda,i}^t + \sum_{j \in \mathcal{N}_i(\xi_a^{t+1})} \mathbf{C}_{ij}(\xi_a^{t+1})(\mathbf{x}_j^t - \mathbf{x}_i^t)] \\ &\quad + \sum_{j \in \mathcal{N}_i(\xi_a^{t+1})} \mathbf{C}_{ij}(\xi_a^{t+1})(\mathbf{x}_j^{t+1} - \mathbf{x}_i^{t+1}) \end{aligned} \quad (10d)$$

Note that to achieve the theoretical performance (see later in Sec. 3),  $\mathbf{x}^0, \hat{\boldsymbol{\lambda}}^0, \mathbf{m}_x^0, \mathbf{m}_\lambda^0$  shall be initialized as  $\mathbf{x}_i^0 = \bar{\mathbf{x}}^0, \hat{\boldsymbol{\lambda}}_i^0 = (\alpha/\eta)n^{-1}(\nabla F(\bar{\mathbf{x}}^0) - \nabla f_i(\bar{\mathbf{x}}^0)), \mathbf{m}_{x,i}^0 = \nabla F(\bar{\mathbf{x}}^0), \mathbf{m}_{\lambda,i}^0 = \mathbf{0}$  according to (23). We remark that a simple initialization choice  $\hat{\boldsymbol{\lambda}}^0 = \mathbf{m}_x^0 = \mathbf{m}_{\lambda,i}^0 = \mathbf{0}$  works well in practice.

Both FSPDA-SA and FSPDA-STORM are decentralized algorithms that can be implemented on random time varying topology, and support randomized sparsification for further communication compression. The key is to observe that in (8), (10), the only information required for agent  $i$  is to obtain  $\sum_{j \in \mathcal{N}_i(\xi_a^t)} \mathbf{C}_{ij}(\xi_a^t)(\mathbf{x}_j^t - \mathbf{x}_i^t)$ , and in addition  $\sum_{j \in \mathcal{N}_i(\xi_a^t)} \mathbf{C}_{ij}(\xi_a^t)(\mathbf{x}_j^{t-1} - \mathbf{x}_i^{t-1})$  for FSPDA-STORM, at iteration  $t$ .

## 2.1 Implementation Details and Connection to Existing Works

We discuss several features of the FSPDA algorithms and their connections to existing works.

**Local & Asynchronous Updates.** The *local update* scheme where each agent  $i$  is allowed to update its own local variables  $\mathbf{x}_i, \lambda_i$  for multiple iterations without a communication step is a common practice in decentralized optimization [Liu et al., 2024, Li and Lin, 2024, Alghunaim, 2024, Mishchenko et al., 2022]. As discussed before, such scheme can be seen as a special case of the FSPDA framework where the time varying topology  $\mathcal{E}^{(t)}$  is chosen such that the latter alternates between  $\mathcal{E}^{(t)} = \mathcal{E}$  and  $\mathcal{E}^{(t)} = \emptyset$ .

Furthermore, FSPDA-SA allows for the general case of *asynchronous* updates. This is done so by taking the stochastic gradient as  $\nabla f_i(\mathbf{x}_i^t; \xi^t) = b_i(\xi^t) \bar{b}_i \nabla f_i(\mathbf{x}_i^t; \xi^t)$  such that  $b_i(\xi^t) \in \{0, 1\}$  with  $\mathbb{E}[b_i(\xi^t)] = 1/\bar{b}_i$  for some constant  $\bar{b}_i > 0$ . Detailed discussions for a fully asynchronous implementation of FSPDA-SA can be found in Appendix A.

**Connection to Existing Works.** Evaluating  $\mathbf{x}^{t+2} - \mathbf{x}^{t+1}$  from the FSPDA-SA sequence and observe that the combination of (8a) and (8b) is equivalent to the second order recursion:

$$\begin{aligned} \mathbf{x}^{t+2} = & 2 \left( \mathbf{I} - \frac{\gamma}{2} \mathbf{A}^\top \mathbf{A}(\xi^{t+1}) \right) \mathbf{x}^{t+1} - (\mathbf{I} - (\gamma - \eta\beta) \mathbf{A}^\top \mathbf{A}(\xi^t)) \mathbf{x}^t \\ & - \alpha (\nabla \mathbf{f}(\mathbf{x}^{t+1}; \xi^{t+1}) - \nabla \mathbf{f}(\mathbf{x}^t; \xi^t)). \end{aligned} \quad (11)$$

This reduces the FSPDA-SA recursion into a primal-only sequence by eliminating the dual sequence  $\lambda^t$ . In the deterministic optimization setting when  $\mathbf{A}(\xi) \equiv \mathbf{A}$  and  $\nabla \mathbf{f}(\mathbf{x}; \xi) \equiv \nabla \mathbf{f}(\mathbf{x})$ , (11) is equivalent to the EXTRA algorithm [Shi et al., 2015] using the mixing matrix  $\tilde{\mathbf{W}} = \mathbf{I} - \gamma \text{Diag}(\tilde{\mathbf{W}}\mathbf{1}) + \gamma \tilde{\mathbf{W}}$  where  $\tilde{\mathbf{W}}$  is the 0-1 adjacency matrix of  $\mathcal{G}$ . Here, with an appropriate choice of  $\gamma$ ,  $\tilde{\mathbf{W}}$  will be doubly stochastic and satisfies the convergence requirement in [Shi et al., 2015]. Similar observations have been made in [Nedic et al., 2017] for the gradient tracking and DIGing algorithms.

On the other hand, for stochastic optimization on random networks, (11) suggests each agent to keep the current and previous iterates received from neighbors in the corresponding time varying topology. In this case, (11) yields an extension of the EXTRA/GT algorithms to time varying topology.

## 3 Convergence Analysis of FSPDA

This section presents the convergence rate analysis of FSPDA for (1). Unless otherwise specified, we focus on the case with smooth but possibly non-convex objective function. Specifically, we consider:

**Assumption 3.1.** Each  $f_i$  is  $L$ -smooth, i.e., for  $i = 1, \dots, n$ ,

$$\|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y})\| \leq L \|\mathbf{x} - \mathbf{y}\| \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d. \quad (12)$$

There exists  $f_\star > -\infty$  such that  $f_i(\mathbf{x}) \geq f_\star$  for any  $\mathbf{x} \in \mathbb{R}^d$ .

Note this implies that the global objective function  $F(\cdot)$  is  $L$ -smooth but possibly non-convex.

We further assume that the random network  $\mathcal{G}(\xi_a)$  is connected in expectation, yet each realization  $\mathcal{G}(\xi_a)$  may not be connected. Let  $\mathbf{R} = \mathbb{E}[\mathbf{I}(\xi_a)]$ , this leads to the following property concerning the expected graph Laplacian matrix  $\mathbf{A}^\top \mathbf{R} \mathbf{A} = \mathbb{E}[\mathbf{A}(\xi_a)^\top \mathbf{A}]$ . Defining the matrix  $\mathbf{K} := (\mathbf{I}_n - \mathbf{1}\mathbf{1}^\top/n) \otimes \mathbf{I}_d$ , we have

**Assumption 3.2.** There exists  $\rho_{\max} \geq \rho_{\min} > 0$  and  $\bar{\rho}_{\max} \geq \bar{\rho}_{\min} > 0$  such that

$$\rho_{\min} \mathbf{K} \preceq \mathbf{A}^\top \mathbf{R} \mathbf{A} \preceq \rho_{\max} \mathbf{K} \quad \text{and} \quad \bar{\rho}_{\min} \mathbf{K} \preceq \mathbf{A}^\top \mathbf{A} \preceq \bar{\rho}_{\max} \mathbf{K}. \quad (13)$$

It holds that  $\mathbf{A}^\top \mathbf{R} \mathbf{A} \mathbf{K} = \mathbf{A}^\top \mathbf{R} \mathbf{A} = \mathbf{K} \mathbf{A}^\top \mathbf{R} \mathbf{A}$ . The above assumption can be satisfied if  $\mathcal{G}$  is connected [Yi et al., 2021], [Yi et al., 2018, Lemma 2] and  $\text{diag}(\mathbf{R}) > \mathbf{0}$  such that each edge is selected with a positive probability. As an important consequence, if  $\gamma \leq \rho_{\min}/\rho_{\max}^2$ , we have

$$\|(\mathbf{I} - \gamma \mathbf{A}^\top \mathbf{R} \mathbf{A}) \mathbf{x}\|_{\mathbf{K}}^2 \leq (1 - \gamma \rho_{\min}) \|\mathbf{x}\|_{\mathbf{K}}^2, \quad \forall \mathbf{x} \in \mathbb{R}^{nd}.$$

We thus observe that the operator  $(\mathbf{I} - \gamma \mathbf{A}^\top \mathbf{R} \mathbf{A})$  serves a similar purpose as the mixing matrix in a average consensus algorithms and  $\rho_{\min}$  can be interpreted as the spectral radius of  $\mathcal{G}$  similar

182 to [Koloskova et al., 2020, Eq. (12)]. Moreover, if we define  $\mathbf{Q} := (\mathbf{A}^\top \mathbf{R} \mathbf{A})^\dagger$  such that it holds  
 183  $\mathbf{Q} \mathbf{A}^\top \mathbf{R} \mathbf{A} = \mathbf{A}^\top \mathbf{R} \mathbf{A} \mathbf{Q} = \mathbf{K}$ , Assumption 3.2 implies that  $\rho_{\max}^{-1} \mathbf{K} \preceq \mathbf{Q} \preceq \rho_{\min}^{-1} \mathbf{K}$ .

184 Next we consider several assumptions on the noise variance of the random quantities in FSPDA:

185 **Assumption 3.3.** For any fixed  $\mathbf{x}_i \in \mathbb{R}^d$ ,  $i \in [n]$ , there exists  $\sigma_i \geq 0$  such that

$$\mathbb{E}_{\xi_i \sim \mathbb{P}_i} [\|\nabla f_i(\mathbf{x}_i; \xi_i) - \nabla f_i(\mathbf{x}_i)\|^2] \leq \sigma_i^2. \quad (14)$$

186 To simplify notations, we define  $\bar{\sigma}^2 := (1/n) \sum_{i=1}^n \sigma_i^2$ .

187 **Assumption 3.4.** For any fixed  $\mathbf{x} \in \mathbb{R}^{nd}$ , there exists  $\sigma_A \geq 0$  such that

$$\mathbb{E}_{\xi_a \sim \mathbb{P}_a} [\|\mathbf{A}(\xi_a)^\top \mathbf{A} \mathbf{x} - \mathbf{A}^\top \mathbf{R} \mathbf{A} \mathbf{x}\|^2] \leq \sigma_A^2 \|\mathbf{x}\|_{\mathbf{K}}^2. \quad (15)$$

188 Assumption 3.3 is standard. Meanwhile for Assumption 3.4, the variance term  $\sigma_A^2$  measures the  
 189 quality of the random topology  $\mathcal{G}(\xi_a)$  in approximating the expected graph Laplacian  $\mathbf{A}^\top \mathbf{R} \mathbf{A}$ . The  
 190 latter is important as it contributes to the variance in the drift term of FSPDA. Observe that  $\sigma_A^2$   
 191 decreases with the proportion of edges selected in each random subgraph  $\mathcal{G}(\xi_a)$ .

192 To facilitate our discussions, we define the following quantities:

$$\bar{\mathbf{x}}^t := \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^t, \quad \sum_{i=1}^n \|\mathbf{x}_i^t - \bar{\mathbf{x}}^t\|^2 = \|\mathbf{x}^t\|_{\mathbf{K}}^2. \quad (16)$$

193 **Convergence of FSPDA-SA.** We summarize the convergence rate for FSPDA-SA as follows. The proof  
 194 can be found in Appendix C:

**Theorem 3.5.** Under Assumptions 3.1, 3.2, 3.3, 3.4. Suppose that the step sizes satisfy the conditions defined in (46). Then, for any  $T \geq 1$  with the random stopping iteration  $\mathsf{T} \sim \text{Unif}\{0, \dots, T-1\}$ , the iterates generated by FSPDA-SA satisfy

$$\mathbb{E} [\|\nabla F(\bar{\mathbf{x}}^{\mathsf{T}})\|^2] \leq \frac{F_0 - f_\star}{\alpha T/8} + 8\alpha \mathbb{C}_\sigma \frac{\bar{\sigma}^2}{n}, \quad (17)$$

$$\mathbb{E} [\sum_{i=1}^n \|\mathbf{x}_i^{\mathsf{T}} - \bar{\mathbf{x}}^{\mathsf{T}}\|^2] \leq \frac{F_0 - f_\star}{\alpha \gamma \rho_{\min} T/8} + \frac{8\alpha^2 \mathbb{C}_\sigma \bar{\sigma}^2}{\alpha \gamma \rho_{\min} n}, \quad (18)$$

for any  $\alpha > 0$ , where  $F_0, \mathbb{C}_\sigma$  are defined in (44), (50).

195  
 196 Setting  $\alpha = \mathcal{O}(n/\sqrt{T\bar{\sigma}^2})$ ,  $\alpha = \sqrt{n/(T\bar{\sigma}^2)}$  (and assuming  $\bar{\sigma} > 0$ ), we have

$$\mathbb{E} [\|\nabla F(\bar{\mathbf{x}}^{\mathsf{T}})\|^2] = \mathcal{O}(\bar{\sigma}/\sqrt{nT}), \quad (19)$$

197 which is the same asymptotic convergence rate as a centralized SGD algorithm that takes  $n$  stochastic  
 198 gradient samples uniformly from each agent, i.e., linear speedup [Lian et al., 2017]. Also, using  
 199  $\alpha = 1$ , the consensus error converges as a rate of  $\mathbb{E} [\sum_{i=1}^n \|\mathbf{x}_i^{\mathsf{T}} - \bar{\mathbf{x}}^{\mathsf{T}}\|^2] = \mathcal{O}(n^2 \sigma_A^2 \rho_{\max}/(T \rho_{\min}^2))$   
 200 under the same step size choice used in (19). Notice that for  $T \gg 1$ , the effect of random topology  
 201 only degrades the convergence of consensus error, keeping the transient rate in (19) unaffected. If  
 202 the gradients are deterministic ( $\bar{\sigma} = 0$ ), setting  $\alpha = (L^2 \eta_\infty \rho_{\min})^{1/3}$ ,  $\alpha = \alpha_\infty$  will yield a better  
 203 convergence rate as  $\mathbb{E} [\|\nabla F(\bar{\mathbf{x}}^{\mathsf{T}})\|^2] = \mathcal{O}(\sigma_A^4 \sqrt{n}/T)$ . Without a transient phase, the error due to  
 204 random graph and coordinate sparsification is persistent through  $\sigma_A^4$  in the above convergence rate.

205 We further show that the convergence of FSPDA-SA can be accelerated if the objective function of (1)  
 206 satisfies the Polyak-Lojasiewicz (PL) condition:

207 **Assumption 3.6.** There exists a constant  $\mu > 0$  such that  $2\mu(F(\mathbf{x}) - f_\star) \leq \|\nabla F(\mathbf{x})\|^2$ ,  $\forall \mathbf{x} \in \mathbb{R}^d$ .

208 Assumption 3.6 includes strongly convex functions as a special case, but also includes other non-  
 209 convex functions; see [Karimi et al., 2016]. We observe:

**Corollary 3.7.** Suppose the assumptions and step size conditions in Theorem 3.5 hold. Furthermore, with Assumption 3.6, there exists  $\delta \in (0, 1)$  such that for any  $t \geq 0$ ,

$$\mathbb{E}_t [F_{t+1} - f_\star] \leq (1 - \delta)(F_t - f_\star) + \mathbb{C}_\sigma \alpha^2 \bar{\sigma}^2 / n \quad (20)$$

for  $F_t, \mathbb{C}_\sigma$  defined in (44), (70), and  $\delta = \min\{\alpha\mu/4, \gamma\rho_{\min}/16, \eta\beta/(3\rho_{\min}), \eta/12\}$ .



The proof can be found in Appendix C.6. By setting  $\alpha = c \ln(T)/(n^2 T)$  in (20), with a carefully chosen  $c$  and a sufficiently large  $T$  such that  $\alpha \leq \alpha_\infty$ , we can ensure that

$$\mathbb{E} [F(\bar{\mathbf{x}}^T) - f_\star + \|\mathbf{x}^T\|_{\mathbf{K}}^2] = \mathcal{O}(\bar{\sigma}^2 \ln(T)/(\mu n T)) \quad (21)$$

In the case of deterministic gradient, i.e.,  $\bar{\sigma}^2 = 0$ , by setting  $\alpha = \alpha_\infty$ , (20) ensures a linear convergence rate of  $\mathbb{E} [F(\bar{\mathbf{x}}^T) - f_\star + \|\mathbf{x}^T\|_{\mathbf{K}}^2] = \mathcal{O}((1 - \delta)^T)$ , which shows that the performance of FSPDA-SA is on par with [Nedic et al., 2017, Xu et al., 2017], despite it only requires one round of (sparsified) transmission per iteration.

**Convergence of FSPDA-STORM.** To exploit the benefits of control variates, we need an additional assumption on the stochastic gradient map:

**Assumption 3.8.** Each stochastic function  $f_i(\cdot; \xi)$  is  $L_\xi$ -smooth in expectation, i.e., for  $i = 1, \dots, n$ ,

$$\mathbb{E}_\xi [\|\nabla f_i(\mathbf{x}; \xi) - \nabla f_i(\mathbf{y}; \xi)\|^2] \leq L_\xi^2 \|\mathbf{x} - \mathbf{y}\|^2 \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d. \quad (22)$$

The above assumption is also known as the mean-square smoothness condition, see [Cutkosky and Orabona, 2019], which is strictly stronger than Assumption 3.1. We observe the following convergence guarantee for FSPDA-STORM, whose proof can be found in Appendix D.

**Theorem 3.9.** Under Assumptions 3.1, 3.2, 3.3, 3.4, 3.8. Suppose that the step sizes satisfy the conditions in (184) - (214). Then, for any  $T \geq 1$  with the random stopping iteration  $T \sim \text{Unif}\{0, \dots, T-1\}$ , the iterates generated by FSPDA-STORM satisfy

$$\mathbb{E} [\|\nabla F(\bar{\mathbf{x}}^T)\|^2] \leq \frac{F_0 - f_\star}{T\alpha/4} + \frac{(\mathbf{e} \cdot 2a_x^2 + \mathbf{f} \cdot 4a_x^2 n)\bar{\sigma}^2}{\alpha/4}, \quad (23)$$

$$\mathbb{E} [\sum_{i=1}^n \|\mathbf{x}_i^T - \bar{\mathbf{x}}^T\|^2] \leq \frac{F_0 - f_\star}{T\mathbf{a}\gamma\rho_{\min}/8} + \frac{(\mathbf{e} \cdot 2a_x^2 + \mathbf{f} \cdot 4a_x^2 n)\bar{\sigma}^2}{\mathbf{a}\gamma\rho_{\min}/8}, \quad (24)$$

where the constants  $F_0, \mathbf{a}, \mathbf{e}, \mathbf{f}$  are defined in (110).

Setting  $\alpha = \mathcal{O}(\bar{\sigma}^{-2/3} T^{-1/3})$ ,  $\eta = \mathcal{O}(n)$ ,  $\gamma = \mathcal{O}(T^{-1/3})$ ,  $\beta = \mathcal{O}(n^{-1} T^{-2/3})$ ,  $a_x = \mathcal{O}(\bar{\sigma}^{-4/3} T^{-2/3})$ ,  $a_\lambda = \mathcal{O}(T^{-1/3})$ ,  $\mathbf{f} = \mathcal{O}(n^{-1} T^{1/3})$  (see (111) - (117)), and initializing the algorithm such that  $\|\mathbf{v}^0\|_{\mathbf{K}}^2 = \mathcal{O}(T^{-2/3})$ ,  $\|\bar{\mathbf{m}}_x^0 - (1/n)\mathbf{1}_\otimes^\top \nabla \mathbf{f}(\mathbf{x}^0)\|^2 = \mathcal{O}(T^{-1/3})$  and  $\|\bar{\mathbf{m}}_x^0 - \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^0, \lambda^0)\|^2 = \mathcal{O}(T^{-1/3})$ , we have

$$\mathbb{E} [\|\nabla F(\bar{\mathbf{x}}^T)\|^2] = \mathcal{O}(\bar{\sigma}^{2/3}/T^{2/3}). \quad (25)$$

In regard to the order of  $\bar{\sigma}$  and  $T$ , provided that  $n$  is small, the convergence rate of FSPDA-STORM matches the lower bound [Arjevani et al., 2023] for non-convex functions under the same smoothness assumption. Moreover, by the same choice of step sizes, the consensus error converges at the rate of  $\mathbb{E} [\sum_{i=1}^n \|\mathbf{x}_i^T - \bar{\mathbf{x}}^T\|^2] = \mathcal{O}(\bar{\sigma}^{2/3} n \rho_{\min}^{-1} T^{-2/3})$ . We remark that in (25), the rate remains constant as  $n$  increases such that FSPDA-STORM does not offer the same *linear speedup* observed in Theorem 3.5 for FSPDA-SA. Nevertheless, as  $T \gg 1$ , the rate of FSPDA-STORM will surpass that of FSPDA-SA and other decentralized algorithms on time varying topologies.

Lastly, we provide detailed discussions on the convergence rates above, e.g., transient time, effects of random topology, etc., in Appendix B.

### 3.1 Insight from Analysis: Fixed Point Iteration of FSPDA-SA

From (8a), the following recursive relationship holds for  $\bar{\mathbf{x}}^t$ : using the relation  $\mathbf{1}^\top \mathbf{A}^\top = \mathbf{0}$ , we have

$$\bar{\mathbf{x}}^{t+1} = \bar{\mathbf{x}}^t - \frac{\alpha}{n} \sum_{i=1}^n \nabla f_i(\mathbf{x}_i^t; \xi_i^t). \quad (26)$$

This shows that the evolution of  $\{\bar{\mathbf{x}}^t\}_{t \geq 0}$  is similar to that of ‘centralized’ SGD applied on (1) except that the local gradients are evaluated on the local iterates. However, it is still not straightforward to analyze the convergence of FSPDA-SA as the update of  $\mathbf{x}^t$  involves the dual variable  $\lambda^t$  which lacks an intuitive interpretation for constructing the right Lyapunov function.

To this end, we study the fixed point(s) of (8) to gain insights. Suppose that for some  $t_\star$ , the fixed point conditions  $\mathbb{E}[\lambda^{t_\star+1} | \xi^{t_\star}] = \lambda^{t_\star}$ ,  $\mathbb{E}[\mathbf{x}^{t_\star+1} | \xi^{t_\star}] = \mathbf{x}^{t_\star}$  hold. Since  $\mathbf{R}$  is a diagonal matrix with positive diagonal elements, we observe

$$\mathbb{E}[\lambda^{t_\star+1} | \xi^{t_\star}] = \lambda^{t_\star} \iff \mathbf{R} \mathbf{A} \mathbf{x}^{t_\star} = \mathbf{0} \iff \mathbf{A} \mathbf{x}^{t_\star} = \mathbf{0}, \quad (27)$$

246 On the other hand, the primal update yields

$$\mathbb{E}[\mathbf{x}^{t_*+1} \mid \xi^{t_*}] = \mathbf{x}^{t_*} - \alpha \nabla \mathbf{f}(\mathbf{x}^{t_*}) - \eta \mathbf{A}^\top \boldsymbol{\lambda}^{t_*}. \quad (28)$$

247 Since  $\mathbf{x}_1^{t_*} = \mathbf{x}_2^{t_*} = \dots = \mathbf{x}_n^{t_*}$  at the fixed point (due to (27)), by the consensus condition across two  
248 time steps, it implies

$$\begin{aligned} \mathbb{E}[\mathbf{x}^{t_*+1} \mid \xi^{t_*}] - \mathbf{x}^{t_*} &= (\mathbf{1} \otimes \mathbf{I}_d)(\bar{\mathbf{x}}^{t_*+1} - \bar{\mathbf{x}}^{t_*}) \\ &\iff \alpha \nabla \mathbf{f}(\mathbf{x}^{t_*}) + \eta \mathbf{A}^\top \boldsymbol{\lambda}^{t_*} = \frac{\alpha}{n} (\mathbf{1}\mathbf{1}^\top \otimes \mathbf{I}_d) \nabla \mathbf{f}(\mathbf{x}^{t_*}) \\ &\iff \eta \mathbf{A}^\top \boldsymbol{\lambda}^{t_*} = \alpha \left( \frac{1}{n} \mathbf{1}\mathbf{1}^\top - \mathbf{I}_n \right) \otimes \mathbf{I}_d \nabla \mathbf{f}((\mathbf{1} \otimes \mathbf{I}) \bar{\mathbf{x}}^{t_*}). \end{aligned} \quad (29)$$

249 From (29), we see that  $\hat{\boldsymbol{\lambda}}^t$  shall converge to the difference between global and local gradient. Inspired  
250 by the above, to facilitate the analysis later, we define

$$\mathbf{v}^t := \mathbf{A}^\top \boldsymbol{\lambda}^t + \frac{\alpha}{\eta} \nabla \mathbf{f}((\mathbf{1} \otimes \mathbf{I}) \bar{\mathbf{x}}^t), \quad (30)$$

251 for any  $t \geq 0$ . In particular, we see that  $\|\mathbf{v}^t\|_{\mathbf{K}}^2$  measures the violation of (29) in tracking the average  
252 deterministic gradient using the dual variables. The latter will be instrumental in analyzing the  
253 consensus error bound, as revealed in Lemma C.2.

## 254 4 Numerical Experiments

255 This section reports the numerical experiments on practical performance of FSPDA. For the time  
256 varying topology, we take an extreme setting where for each realization  $\mathcal{G}(\xi_a)$ , only one edge will  
257 be selected uniformly at random from  $\mathcal{G}$ . We evaluate the performance with the worst-agent metric,  
258 i.e., we present the training loss as  $\max_{i \in [n]} F(\mathbf{x}_i^t)$ , and the stationarity/gradient-norm measure as  
259  $\max_{i \in [n]} \|\nabla F(\mathbf{x}_i^t)\|^2$ . This captures the worst-case of the solutions produced by the algorithms.  
260 Unless otherwise specified, all algorithms are initialized with  $\mathbf{x}_i^0 = \bar{\mathbf{x}}^0$ , and for FSPDA we initialize  
261  $\hat{\boldsymbol{\lambda}}^0 = \mathbf{m}_{x,i}^0 = \mathbf{m}_{\lambda,i}^0 = \mathbf{0}$ , and the stochastic gradients are estimated with a batch size of 256. In the  
262 interest of space, omitted details and hyperparameters of the experiments can be found in Appendix F.

263 **MNIST Experiments.** The first set of experiments considers a moderate-scale setting of training a  
264 one hidden layer feed-forward neural network with 100 hidden neurons (total number of parameters  
265  $d = 79,510$ ) on the MNIST dataset with  $m = 60,000$  samples of 784-dimensional features.

266 In the first experiment, we consider the static topology  $\mathcal{G}$  as an Erdos-Renyi graph with connectivity of  
267  $p = 0.5$  and  $n = 10$  agents. We compare the proposed FSPDA-SA, FSPDA-STORM with six benchmark  
268 algorithms utilizing different types of time-varying topology. Among them, DSGD [Koloskova et al.,  
269 2020] and Swarm-SGD [Nadiradze et al., 2021] use the general time varying topology setting as FSPDA  
270 where each edge of  $\mathcal{G}(\xi_a)$  is active uniformly at random, in addition to random sparsification used  
271 FSPDA-SA and adaptive quantized used in Swarm-SGD; CHOCO-SGD [Koloskova et al., 2019b] takes  
272  $\mathcal{G}(\xi_a)$  as an broadcasting subgraph where one agent selects all his/her neighbors; Decen-Scaffnew  
273 [Mishchenko et al., 2022], LED [Alghunaim, 2024], and K-GT [Liu et al., 2024] utilize local updates  
274 where  $\mathcal{G}(\xi_a)$  is either taken as an empty topology, or as the static topology  $\mathcal{G}$ . We configure these  
275 algorithms such that they have the same communication cost (in terms of bits transmitted over  
276 network) *on average*. For instance, the local update algorithms (Decen-Scaffnew, LED, K-GT)  
277 only communicate once using  $\mathcal{G}$  every  $\mathcal{O}\left(\frac{|\mathcal{E}|d}{k}\right)$  iterations to match the communication cost of  
278  $k$ -coordinate sparse one-edge random graph used in FSPDA.

279 The local objective function held by each agent is the cross-entropy classification loss on a local  
280 dataset with  $m_i = 6000$  samples, plus a regularization loss  $\frac{\lambda}{2} \|\mathbf{x}_i\|^2$  with  $\lambda = 10^{-4}$ , where  $\mathbf{x}_i$  are the  
281 weight parameters of the feed-forward neural network classifier. We split the training set into  $n = 10$   
282 disjoint sets such that each set contains only one class label and assign each set to one agent as its  
283 local dataset. Note that as we do not shuffle the data samples across local datasets, the local objective  
284 function held by different agents will become highly heterogeneous.

285 Fig. 1 compares the squared gradient norm, training loss, consensus error of the benchmarked algo-  
286 rithms. We first note that both FSPDA algorithms have significantly outperformed DSGD, Swarm-SGD  
287 on the general time varying topology as well as CHOCO-SGD. Meanwhile, the performance of FSPDA  
288 is comparable to the local update algorithms Decen-Scaffnew, LED, K-GT. Notice that the latter



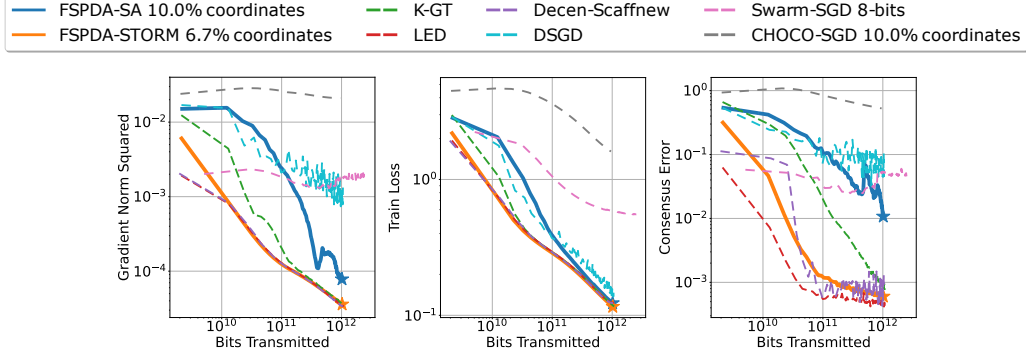


Figure 1: Feed-forward neural network classification training on MNIST using  $10^6$  iterations.




Figure 2: Resnet-50 classification training on Imagenet.

require additional synchronization steps which may not be suitable for random networks. Lastly, we notice that as  $T \gg 1$ , FSPDA-STORM can slightly outperform FSPDA-SA due to its  $\mathcal{O}(1/T^{2/3})$  rate as shown in our analysis. We further expand the experiments by a series of ablation studies over data heterogeneity, sparsity levels, graph topologies, gradient noise and dual momentum in Appendix E.

**Imagenet Experiments.** The second set of experiments consider a large-scale setting for training a Resnet-50 network (total number of parameters  $d = 25,557,032$ ) on the Imagenet dataset (training dataset of 1,281,168 images from 100 classes, re-scaled and cropped to  $256 \times 256$  image dimensions). We consider cross-entropy classification loss plus the same L2 norm regularization loss as in the previous setup. We split the dataset across a network of  $n = 8$  nodes where the static graph  $\mathcal{G}$  is taken as the fully connected topology. The performance metrics are measured at the network average iterate  $\bar{x}^t$ . Inspired by [Loshchilov and Hutter, 2016, Eq. (5)] we adopt a cosine learning rate scheduling with 5 epochs of linear warm up for every algorithm. In particular, the step sizes  $\alpha, \eta$  of FSPDA-SA are scheduled simultaneously such that  $\alpha_t/\eta_t$  remains constant, as illustrated in Appendix F. We draw a batch of 128 samples to estimate the stochastic gradient.

We focus on the communication efficiency and only compare FSPDA-SA, CHOCO-SGD, Swarm-SGD in this experiment due to limited resources. The results are reported in Figure 2 that compare the test accuracy and training loss against iteration number and bits transmitted. When compared with CHOCO-SGD, FSPDA-SA achieves almost the same accuracy using one-edge random graphs with at least 100x reduction in communication cost on 100 epoch training. Also notice that further compressing the communication to 0.1% sparse coordinates in FSPDA-SA requires more training epochs to recover the same level of accuracy.

**Conclusions.** This paper proposed a fully stochastic primal dual gradient algorithm (FSPDA) framework for decentralized optimization over arbitrarily time varying random networks. We utilize a new stochastic augmented Lagrangian function and apply SA to search for its saddle point. We develop two algorithms, one is by plain SA (FSPDA-SA), and one uses control variates for variance reduction (FSPDA-STORM). We prove that both algorithms achieve state-of-the-art convergence rates, while relaxing assumptions on both bounded heterogeneity and the type of time varying topologies.

## References

- Sulaiman A Alghunaim. Local exact-diffusion for decentralized optimization and learning. *IEEE Transactions on Automatic Control*, 2024.
- Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *Mathematical Programming*, 199(1): 165–214, 2023.
- Dimitri Bertsekas. *Nonlinear Programming*, volume 4. Athena Scientific, 2016.
- Pascal Bianchi, Walid Hachem, and Adil Salim. A fully stochastic primal-dual algorithm. *Optimization Letters*, 15(2):701–710, 2021.
- Tsung-Hui Chang, Mingyi Hong, Hoi-To Wai, Xinwei Zhang, and Songtao Lu. Distributed learning in the nonconvex world: From batch data to streaming and beyond. *IEEE Signal Processing Magazine*, 37(3):26–38, 2020.
- Ashok Cutkosky and Francesco Orabona. Momentum-based variance reduction in non-convex sgd. *Advances in neural information processing systems*, 32, 2019.
- Luyao Guo, Sulaiman A Alghunaim, Kun Yuan, Laurent Condat, and Jinde Cao. Revisiting decentralized proxskip: Achieving linear speedup. *arXiv preprint arXiv:2310.07983*, 2023.
- Davood Hajinezhad and Mingyi Hong. Perturbed proximal primal–dual algorithm for nonconvex nonsmooth optimization. *Mathematical Programming*, 176(1):207–245, 2019.
- Mingyi Hong, Davood Hajinezhad, and Ming-Min Zhao. Prox-pda: The proximal primal-dual algorithm for fast distributed nonconvex optimization and learning over networks. In *International Conference on Machine Learning*, pages 1529–1538. PMLR, 2017.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2):1–210, 2021.
- Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part I 16*, pages 795–811. Springer, 2016.
- Anastasia Koloskova, Tao Lin, Sebastian U Stich, and Martin Jaggi. Decentralized deep learning with arbitrary communication compression. In *International Conference on Learning Representations*, 2019a.
- Anastasia Koloskova, Sebastian Stich, and Martin Jaggi. Decentralized stochastic optimization and gossip algorithms with compressed communication. In *International Conference on Machine Learning*, pages 3478–3487. PMLR, 2019b.
- Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian Stich. A unified theory of decentralized sgd with changing topology and local updates. In *International Conference on Machine Learning*, pages 5381–5393. PMLR, 2020.
- Anastasiia Koloskova, Tao Lin, and Sebastian U Stich. An improved analysis of gradient tracking for decentralized machine learning. *Advances in Neural Information Processing Systems*, 34: 11422–11435, 2021.
- Dmitry Kovalev, Elnur Gasanov, Alexander Gasnikov, and Peter Richtarik. Lower bounds and optimal algorithms for smooth and strongly convex decentralized optimization over time-varying networks. *Advances in Neural Information Processing Systems*, 34:22325–22335, 2021.
- Dmitry Kovalev, Ekaterina Borodich, Alexander Gasnikov, and Dmitrii Feoktistov. Lower bounds and optimal algorithms for non-smooth convex decentralized optimization over time-varying networks. *arXiv preprint arXiv:2405.18031*, 2024.

363 Jinlong Lei, Han-Fu Chen, and Hai-Tao Fang. Asymptotic properties of primal-dual algorithm for  
364 distributed stochastic optimization over random networks with imperfect communications. *SIAM*  
365 *Journal on Control and Optimization*, 56(3):2159–2188, 2018.

366 Huan Li and Zhouchen Lin. Accelerated gradient tracking over time-varying graphs for decentralized  
367 optimization. *Journal of Machine Learning Research*, 25(274):1–52, 2024.

368 Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized  
369 algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic  
370 gradient descent. *Advances in neural information processing systems*, 30, 2017.

371 Yue Liu, Tao Lin, Anastasia Koloskova, and Sebastian U Stich. Decentralized gradient tracking with  
372 local steps. *Optimization Methods and Software*, pages 1–28, 2024.

373 Ilan Lobel and Asuman Ozdaglar. Distributed subgradient methods for convex optimization over  
374 random networks. *IEEE Transactions on Automatic Control*, 56(6):1291–1306, 2010.

375 Paolo Di Lorenzo and Gesualdo Scutari. Next: In-network nonconvex optimization. *IEEE Transac-*  
376 *tions on Signal and Information Processing over Networks*, 2(2):120–136, 2016.

377 Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv*  
378 *preprint arXiv:1608.03983*, 2016.

379 Songtao Lu, Xinwei Zhang, Haoran Sun, and Mingyi Hong. Gnsd: A gradient-tracking based  
380 nonconvex stochastic algorithm for decentralized optimization. In *2019 IEEE Data Science*  
381 *Workshop (DSW)*, pages 315–321. IEEE, 2019.

382 Konstantin Mishchenko, Grigory Malinovsky, Sebastian Stich, and Peter Richtárik. Proxskip: Yes!  
383 local gradient steps provably lead to communication acceleration! finally! In *International*  
384 *Conference on Machine Learning*, pages 15750–15769. PMLR, 2022.

385 Giorgi Nadiradze, Amirmojtaba Sabour, Peter Davies, Shigang Li, and Dan Alistarh. Asynchronous  
386 decentralized sgd with quantized and local updates. *Advances in Neural Information Processing*  
387 *Systems*, 34:6829–6842, 2021.

388 Angelia Nedic and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization.  
389 *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.

390 Angelia Nedic, Alex Olshevsky, and Wei Shi. Achieving geometric convergence for distributed  
391 optimization over time-varying graphs. *SIAM Journal on Optimization*, 27(4):2597–2633, 2017.

392 Shi Pu, Alex Olshevsky, and Ioannis Ch Paschalidis. A sharp estimate on the transient time of  
393 distributed stochastic gradient descent. *IEEE Transactions on Automatic Control*, 67(11):5900–  
394 5915, 2021.

395 Tiancheng Qin, S Rasoul Etesami, and César A Uribe. Communication-efficient decentralized local  
396 sgd over undirected networks. In *2021 60th IEEE Conference on Decision and Control (CDC)*,  
397 pages 3361–3366. IEEE, 2021.

398 Guannan Qu and Na Li. Harnessing smoothness to accelerate distributed optimization. *IEEE*  
399 *Transactions on Control of Network Systems*, 5(3):1245–1260, 2017.

400 S Sundhar Ram, Angelia Nedić, and Venugopal V Veeravalli. Distributed stochastic subgradient  
401 projection algorithms for convex optimization. *Journal of optimization theory and applications*,  
402 147:516–545, 2010.

403 Wei Shi, Qing Ling, Gang Wu, and Wotao Yin. Extra: An exact first-order algorithm for decentralized  
404 consensus optimization. *SIAM Journal on Optimization*, 25(2):944–966, 2015.

405 Jinming Xu, Shanying Zhu, Yeng Chai Soh, and Lihua Xie. Convergence of asynchronous distributed  
406 gradient methods over stochastic networks. *IEEE Transactions on Automatic Control*, 63(2):  
407 434–448, 2017.

- 408 Chung-Yiu Yau and Hoi-To Wai. Fully stochastic distributed convex optimization on time-varying  
409 graph with compression. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages  
410 145–150. IEEE, 2023.
- 411 Xinlei Yi, Lisha Yao, Tao Yang, Jemin George, and Karl H Johansson. Distributed optimization for  
412 second-order multi-agent systems with dynamic event-triggered communication. In *2018 IEEE*  
413 *Conference on Decision and Control (CDC)*, pages 3397–3402. IEEE, 2018.
- 414 Xinlei Yi, Shengjun Zhang, Tao Yang, Tianyou Chai, and Karl H Johansson. Linear convergence  
415 of first-and zeroth-order primal–dual algorithms for distributed nonconvex optimization. *IEEE*  
416 *Transactions on Automatic Control*, 67(8):4194–4201, 2021.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]



Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Due to limited computing resources and time constraints, we are unable to perform multiple runs of our algorithms and report the error bars. We will produce the error bar statistics if time permits.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: [NA]

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

### 13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

### 16. Declaration of LLM usage



726 Question: Does the paper describe the usage of LLMs if it is an important, original, or  
727 non-standard component of the core methods in this research? Note that if the LLM is used  
728 only for writing, editing, or formatting purposes and does not impact the core methodology,  
729 scientific rigorousness, or originality of the research, declaration is not required.

730 Answer: [NA]

731 Justification: [NA]

732 Guidelines:

- 733 • The answer NA means that the core method development in this research does not  
734 involve LLMs as any important, original, or non-standard components.
- 735 • Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>)  
736 for what should or should not be described.