

# Coupling-based Convergence Diagnostic and Step-size Scheme for Stochastic Gradient Descent

Xiang Li and Qiaomin Xie

Department of Industrial and Systems Engineering, University of Wisconsin–Madison  
xiang.li2@wisc.edu, qiaomin.xie@wisc.edu

## Abstract

The convergence behavior of Stochastic Gradient Descent (SGD) crucially depends on the step-size configuration. When using a constant step-size, the SGD iterates form a Markov chain, enjoying fast convergence during the initial transient phase. However, when reaching stationarity, the iterates oscillate around the optimum without making further progress. In this paper, we study the convergence diagnostics for SGD with constant step-size, aiming to develop an effective dynamic step-size scheme. We propose a novel *coupling*-based convergence diagnostic procedure, which monitors the distance of two coupled SGD iterates for stationarity detection. Our diagnostic statistic is simple and is shown to track the transition from transience to stationarity theoretically. We conduct extensive numerical experiments and compare our method against various existing approaches. Our proposed coupling-based step-size scheme is observed to achieve superior performance across a diverse set of convex and non-convex problems. Moreover, our results demonstrate the robustness of our approach to a wide range of hyperparameters.

## 1 Introduction

Stochastic Gradient Descent (SGD) aims to minimize an objective function by iteratively updating the model parameters based on noisy gradients of the function (Robbins and Monro 1951). SGD is well suited for large-scale machine learning scenarios, where the data are abundant and the computation is limited, as it processes data points one at a time or in small batches (Bottou 2010, 2012).

Mathematically, the SGD recursion is given by

$$\theta_{k+1} = \theta_k - \gamma_k \nabla f_k(\theta_k), \quad (1)$$

Here,  $\theta_k$  is the model parameters at iteration  $k$ ,  $\nabla f_k(\theta_k)$  is the noisy gradient of the objective function  $f$  at  $\theta_k$ , and  $\gamma_k > 0$  is the step-size (also known as the learning rate) for the  $k$ th iteration.

The step-size configuration plays a critical role in the convergence of SGD. A traditional approach is to use diminishing step-size, such as  $\gamma_k = \gamma_0/k$  or  $\gamma_0/\sqrt{k}$  (Robbins and Monro 1951; Bach 2014). These decay step-size schedules have been well studied, with non-asymptotic convergence guarantees (Moulines and Bach 2011; Lacoste-Julien,

Schmidt, and Bach 2012; Rakhlin, Shamir, and Sridharan 2011; Shamir and Zhang 2013; Hazan and Kale 2014). However, these schedules are often less robust to ill-conditioning and tend to suffer from a slow convergence rate.

In practice, constant step-size SGD is widely used, due to faster convergence and easy tuning. Prior work on constant step-size SGD (Needell, Ward, and Srebro 2014) provides a classical bias-variance decomposition of the mean-square error. Recent work provides a more precise characterization of its properties (Dieuleveut, Durmus, and Bach 2020; Yu et al. 2021; Lauand and Meyn 2023; Merad and Gaïffas 2023), by leveraging the fact that the SGD iterates  $(\theta_k)_{k \geq 0}$  defined by (1) form a Markov chain with a constant step-size  $\gamma_k \equiv \gamma > 0$ . The Markov chain  $(\theta_k)_{k \geq 0}$  is shown to converge at an exponential rate that is proportional to the step-size  $\gamma$ . However, the chain does not converge to the optimal solution  $\theta^*$  but exhibits as a random walk within a vicinity of  $\theta^*$  of radius  $\mathcal{O}(\sqrt{\gamma})$  as  $k \rightarrow \infty$ .

The insights imply that fast convergence can be achieved by using a large constant step-size; when the iterates reach stationarity, reduce the step-size and continue the iteration to achieve a better accuracy at the next saturation, upon which the procedure is repeated. Indeed a common practice for SGD and deep learning algorithms is to use epoch-wise constant step-size and periodically reduce it (He et al. 2016; Krizhevsky, Sutskever, and Hinton 2012; Wang, Magnússon, and Johansson 2021), but the epoch lengths are mostly predetermined and hand-picked. To make the best use of the procedure to achieve the best of both worlds, it is critical to develop a saturation diagnostic that effectively and promptly detects stationarity, while being computationally efficient.

The problem of detecting saturation in constant step-size SGD has been studied since the seminal work of Pflug (1983), who proposed using the running average of the inner products of successive gradients to monitor convergence. Pflug’s statistic has been used to develop several diagnostic algorithms for SGD (Chee and Toulis 2018; Yaida 2018). However, recent work has shown that Pflug’s method may fail even for quadratic functions, due to the large variance of the diagnostic statistic (Pesme, Dieuleveut, and Flammarion 2020). In particular, the noisy signal fails to accurately indicate whether saturation has been reached, leading to frequent early step-size reduction. Improving upon the diagnostic proposed by Yaida (2018), Lang, Xiao, and Zhang (2019)

developed a more robust detection rule based on gradients of a mini-batch. Their approach, however, tends to suffer from the opposite issue, i.e., being overly conservative and leading delays in stationary detection.

The more recent work in Pesme, Dieuleveut, and Flammarion (2020) introduced an alternative method based on the distance between the initial iterate  $\theta_0$  and the  $k$ th iterate  $\theta_k$ . The idea is that  $\|\theta_k - \theta^*\|^2$  is expected to reach saturation at approximately the same time as  $\|\theta_k - \theta_0\|^2$ . Their approach outperforms Pflug’s methods empirically, but lacks theoretical justifications for general convex problems. A closely related work in Sordello, He, and Su (2020) considers splitting the SGD recursion into multiple threads, which are initialized from the same  $\theta_0$  but use independent data points to calculate the gradients. Near-orthogonality of these gradients is used as an indicator of stationarity. Note that each thread only uses a fraction of the available data points.

In this paper, we propose a new, coupling-based convergence diagnostic procedure that is simple, flexible, and data-efficient. We apply this method to develop an effective dynamic stepsize scheme for SGD. Our main contributions are summarized as follows. First, building on the fact that the constant stepsize SGD iterates  $(\theta_k)_{k \geq 0}$  form a Markov chain, we design a stationarity diagnostic statistic via *Markov chain coupling*. Specifically, our coupling-based method maintains two SGD iterates  $(\theta_k^{(1)})_{k \geq 0}$  and  $(\theta_k^{(2)})_{k \geq 0}$  using the same stepsize and data points at each iteration (“coupling”), but with different initialization. To assess whether the iterates have reached stationary phase, we perform a diagnosis based on the ratio of coupled iterates difference  $\|\theta_k^{(1)} - \theta_k^{(2)}\|^2 / \|\theta_0^{(1)} - \theta_0^{(2)}\|^2$ . Once stationarity is detected, the stepsize is reduced, and the procedure is repeated. Our stationarity detection is simple and easy to implement. Moreover, we prove that our diagnostic statistic is valid for convergence detection in general convex problems. Furthermore, we conduct extensive experiments to evaluate our coupling-based dynamic stepsize scheme against various existing methods. Our approach consistently achieves superior performance across a range of convex and non-convex problems, such as logistic regression and ResNet-18, as well as in stochastic approximation with Markovian data. Lastly, our results demonstrate the robustness of our approach to a wide range of hyper-parameters.

## 2 Preliminaries

In this paper, we focus on the classical SGD algorithm for minimizing an objective function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ . With an initialization  $\theta_0 \in \mathbb{R}^d$ , the SGD recursion (1) can be written equivalently as

$$\theta_{k+1} = \theta_k - \gamma_k [\nabla f(\theta_k) + \varepsilon_k(\theta_k)], \quad k \geq 0, \quad (2)$$

where  $\varepsilon_k : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is the random field corresponding to the stochasticity in  $\nabla f_k(\cdot)$  as an estimate of the true gradient  $\nabla f(\cdot)$ , such that  $\nabla f_k(\cdot) = \nabla f(\cdot) + \varepsilon_k(\cdot)$ .

For our analysis, we consider the following assumptions on the loss function  $f$  and the noise functions  $(\varepsilon_k)_{k \geq 0}$ .

**Assumption A.1** (Smoothness). *The function  $f$  is  $L$ -smooth with  $L \geq 0$ , i.e., for all  $\theta, \theta' \in \mathbb{R}^d$ ,*

$$\|\nabla f(\theta) - \nabla f(\theta')\| \leq L\|\theta - \theta'\|.$$

**Assumption A.2** (Strong convexity). *The function  $f$  is strongly convex with parameter  $\mu > 0$ , i.e.,  $\forall \theta, \theta' \in \mathbb{R}^d$ ,*

$$\langle \nabla f(\theta) - \nabla f(\theta'), \theta - \theta' \rangle \geq \mu\|\theta - \theta'\|^2.$$

Assumption A.2 implies that the objective function  $f$  admits a unique global optimum  $\theta^*$ .

**Assumption A.3** (Zero-mean noise). *There exists a filtration  $(\mathcal{F}_k)_{k \geq 0}$  on some probability space such that for all  $k \in \mathbb{N}$  and  $\theta \in \mathbb{R}^d$ ,  $\varepsilon_k(\theta)$  is  $\mathcal{F}_k$ -measurable and  $\mathbb{E}[\varepsilon_k(\theta) | \mathcal{F}_{k-1}] = 0$ . In addition,  $(\varepsilon_k)_{k \geq 0}$  are independent and identically distributed random fields.*

Assumption A.3 implies that  $\nabla f_k(\theta) := \nabla f(\theta) + \varepsilon_k(\theta)$  is an unbiased estimator of the true gradient  $\nabla f(\theta)$ .

**Assumption A.4** (Co-coercivity). *The gradient is co-coercive in expectation, i.e.,  $\langle \nabla f(\theta) - \nabla f(\theta'), \theta - \theta' \rangle \geq (1/L)\mathbb{E}[\|\nabla f_k(\theta) - \nabla f_k(\theta')\|^2], \forall k > 0, \theta, \theta' \in \mathbb{R}$ , where  $L$  is given in Assumption A.1. Moreover, there exists a constant  $\sigma \geq 0$  such that for all  $k > 0$ ,  $\mathbb{E}[\|\varepsilon_k(\theta^*)\|^2] \leq \sigma^2$ .*

Assumption A.4 is standard and appeared in Dieuleveut, Durmus, and Bach (2020); Merad and Gaïffas (2023).

**Transience-stationarity under constant stepsize.** Constant stepsize SGD has recently gained increasing attention and its properties have been well understood for well-conditioned problems. In particular, the classical result (Needell, Ward, and Srebro 2014) decomposes the bound on  $\mathbb{E}[\|\theta_k - \theta^*\|^2]$  into bias and variance terms: the bias term is proportional to the initial condition  $\|\theta_0 - \theta^*\|$  and decays at an exponential rate; the variance term is determined by the gradient noise. In general, the iterates are not converging to the global optimum  $\theta^*$ .

Recent work has provided a more precise characterization of the bias-variance trade-off through the lens of Markov chain theory (Dieuleveut, Durmus, and Bach 2020; Yu et al. 2021). In particular, with a constant stepsize  $\gamma_k \equiv \gamma > 0$ , the SGD iterates  $(\theta_k)_{k \geq 0}$  given by (2) form a Markov chain. The Markov chain is shown to converge to a unique stationary distribution  $\pi_\gamma$  under appropriate conditions. Let  $\mathcal{P}_2(\mathbb{R}^d)$  denote the set of probability measure on  $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$  with a finite second moment, where  $\mathcal{B}(\mathbb{R}^d)$  is the Borel  $\sigma$ -field of  $\mathbb{R}^d$ .

**Proposition 1** (Proposition 2 in Dieuleveut, Durmus, and Bach 2020). *Suppose that Assumptions A.1–A.4 hold. With constant stepsize  $\gamma \in (0, 2/L)$ , the Markov chain  $(\theta_k)_{k \geq 0}$  given by the recursion (2) satisfies:*

$$W_2^2(P_\gamma^k(\theta_0, \cdot), \pi_\gamma) \leq \rho^k \mathbb{E}_{\theta \sim \pi_\gamma} [\|\theta_0 - \theta\|^2], \quad (3)$$

where  $P_\gamma^k$  is the  $k$ -step Markov kernel for the chain  $(\theta_k)_{k \geq 0}$ ,  $W_2(\nu, \nu')$  is the Wasserstein distance of order two between measures  $\nu, \nu' \in \mathcal{P}_2(\mathbb{R}^d)$ , and  $\rho := 1 - 2\gamma\mu(1 - \gamma L/2)$ .

Proposition 1 shows that during the transient phase, i.e., before reaching stationarity, the Markov chain converges exponentially fast at a rate proportional to the stepsize  $\gamma$ . Consequently, using a larger stepsize achieves faster transient

convergence. However, when reaching stationary phase, the iterates exhibit a random walk around  $\theta^*$ , incurring a *saturated* expected error  $\mathbb{E}_{\pi_\gamma}[\theta] - \theta^* = A\gamma + O(\gamma^2)$  (Dieuleveut, Durmus, and Bach 2020, Theorem 4). Therefore, a larger stepsize suffers from a larger approximation error. Similar transience-stationarity properties have also been characterized for general stochastic approximation algorithms with Markovian data, which covers SGD as a special case (Lauand and Meyn 2023; Huo et al. 2024).

This above insight on the transient phase convergence rate and the stationary phase error naturally indicates an *adaptive constant stepsize scheduling* procedure to achieve the best of both worlds: Start with a large constant stepsize for fast transient convergence; when the iterates have approximately reached stationarity, decrease the stepsize and continue the SGD process to achieve smaller approximation error.

The key challenge here is how to detect the stationarity of the iterates effectively, which has remained an open problem in the literature (Pasupathy, Honnappa, and Hunter 2019). Several statistical tests have been proposed for stationarity diagnostic (Chee and Toulis 2018; Yaida 2018; Pesme, Dieuleveut, and Flammarion 2020). However, these methods either lead to poor empirical results or lack of theoretical analysis. There is a critical need for a more efficient and principled tool to assess the convergence of SGD iterates with constant stepsize. In the next section, we present our approach that uses Markov chain coupling algorithmically.

### 3 A Coupling-based Statistic for Stationarity Diagnostic

In this section, we introduce a novel coupling-based statistic for convergence diagnostic and present the theoretical foundation for justify its effectiveness.

Our approach builds on the fact that the SGD iterates  $(\theta_k)_{k \geq 0}$  forms a Markov chain under a constant stepsize. We propose a diagnostic that detects the convergence of the Markov chain  $(\theta_k)_{k \geq 0}$  by *coupling*. Stochastic coupling techniques have served as a powerful *analytical* tool for characterizing stochastic system performance. In particular, it has been widely employed to establish Markov chain convergence (Dieuleveut, Durmus, and Bach 2020; Huo, Chen, and Xie 2023; Foss and Tweedie 1998; Lauand and Meyn 2023), which serves as an inspiration for our approach.

In particular, consider two SGD iterates  $(\theta_k^{(1)})_{k \geq 0}$  and  $(\theta_k^{(2)})_{k \geq 0}$  that use the same stepsize  $\gamma$  and the same mini-batch data at every iteration, i.e., *sharing the same noise*, but with different initialization  $\theta_0^{(1)}$  and  $\theta_0^{(2)}$ . This coupling of the two Markov chains  $(\theta_k^{(1)})_{k \geq 0}$  and  $(\theta_k^{(2)})_{k \geq 0}$  has played an important role in showing the convergence of the iterates to their unique stationary distribution  $\pi_\gamma$ ; e.g., see the proof of Proposition 2 by Dieuleveut, Durmus, and Bach (2020).

Here we propose to employ the coupling technique algorithmically to construct a statistical test for stationarity diagnostic. Let  $D_k := \theta_k^{(1)} - \theta_k^{(2)}$  denote their difference. Intuitively, the distributional convergence of the iterates  $(\theta_k^{(1)})_{k \geq 0}$  and  $(\theta_k^{(2)})_{k \geq 0}$  implies that  $\|D_k\| \approx 0$  if

$\theta_k^{(1)}$  and  $\theta_k^{(2)}$  reach the stationary phase, thanks to the coupling procedure. This observation suggests a natural criterion for the stationarity of the primary sequence  $\theta_k^{(1)}$  via tracking the ratio of the differences,  $\|D_k\|/\|D_0\|$ , with the auxiliary sequence  $\theta_k^{(2)}$ . This naturally leads to the following convergence diagnostic: if  $\|D_k\|/\|D_0\|$  is smaller than some threshold, then the stepsize is decreased by a factor of  $r \in (0, 1)$ , upon which we reinitialize one sequence  $\theta_k^{(2)}$  and repeat this procedure.

Our intuition is clearly illustrated in Figure 1. Importantly, we observe that when ratio of the distance  $\|D_k\|/\|D_0\|$  decreases to a small value, the SGD iterates  $(\theta_k^{(1)})_{k \geq 0}$  approximately enter the stationary phase. A natural question arises: Can we theoretically justify that the behavior of the difference  $D_k$  accurately reflects the transition of iterates  $(\theta_k^{(1)})_{k \geq 0}$  between transience and stationarity?




Figure 1: Evolution of  $\|\theta_k^{(1)} - \theta^*\|^2$  and  $\|\theta_k^{(2)} - \theta^*\|^2$  and  $\|\theta_k^{(1)} - \theta_k^{(2)}\|^2$  under least squares regression with two different constant stepsizes.

**Quadratic Setting** To answer the above question for our coupling approach, we first consider the case of a quadratic objective function, as stated in the following assumption.

**Assumption A.5** (Quadratic semi-stochastic setting). *There exists a symmetric positive definite matrix  $H$ , a vector  $a$  and a constant  $c$  such that  $f(\theta) = \frac{1}{2}\theta^\top H\theta + a^\top\theta + c$ . The gradient noise  $\varepsilon_i(\theta) = \xi_i$  is independent of  $\theta$ , where  $(\xi_i)_{i \geq 0}$  are i.i.d and satisfy  $\mathbb{E}[\xi_i] = 0$  and  $\mathbb{E}[\xi_i^\top \xi_i] = C$ .*

Under this assumption, we can compute  $\mathbb{E}[\|\theta_k^{(1)} - \theta_k^{(2)}\|^2]$  in a closed-form, as stated in the following proposition. See the proof in Appendix A, provided separately.<sup>1</sup>

**Proposition 2.** *Suppose Assumptions A.1, A.2, and A.5 hold. Let  $\gamma \in (0, 1/L)$ . We have that for all  $k \geq 0$ :*

$$\mathbb{E}[\|D_k\|^2] = \mathbb{E}[D_0^\top (I - \gamma H)^{2k} D_0] \leq (1 - \gamma\mu)^{2k} \mathbb{E}[\|D_0\|^2].$$

Proposition 2 states that the expected squared distance decays exponentially and eventually converges to zero,

**Proposition 3** (Quadratic). *Under Assumptions A.1, A.2, and A.5, we have the following result with constant stepsize*

<sup>1</sup>An extended version can be found at Li and Xie (2024).

$\gamma \in (0, 1/L)$  :

$$\begin{aligned} & W_2^2(P_\gamma^k(\theta_0^{(1)}, \cdot), \pi_\gamma) \\ & \leq \frac{\mathbb{E}[\|D_k\|^2]}{\mathbb{E}[(D_0^\top q_{\max})^2]} \cdot \mathbb{E}_{\theta \sim \pi_\gamma} [\|\theta_0^{(1)} - \theta\|^2], \end{aligned} \quad (4)$$

where  $q_{\max}$  is the eigenvector associated with the largest eigenvalue of the matrix  $I - \gamma H$ .

*Proof.* By centering, we may assume without loss of generality that  $a = 0$  and  $c = 0$ , so  $f(\theta) = \frac{1}{2}\theta^\top H\theta$ . For this setting, we have the following claim on the distributional convergence of the Markov chain  $(\theta_k)_{k \geq 0}$ . The proof of the claim follows similar lines as that of Proposition 2 in Dieuleveut, Durmus, and Bach (2020). We provide the proof in Appendix A for completeness.

**Claim 1.** *Under the setting of Proposition 3, the Markov chain  $(\theta_k)_{k \geq 0}$  given by the recursion (2) satisfies:*

$$W_2^2(P_\gamma^k(\theta_0, \cdot), \pi_\gamma) \leq (1 - \gamma\lambda_{\min})^{2k} \mathbb{E}_{\theta \sim \pi_\gamma} [\|\theta_0 - \theta\|^2], \quad (5)$$

where  $\lambda_{\min}$  being the smallest eigenvalue of  $H$ .

Note that the largest eigenvalues of the matrix  $(I - \gamma H)$  is  $1 - \gamma\lambda_{\min}$ . Let  $q_{\max}$  be the corresponding eigenvector. By Proposition 2, we have:

$$\begin{aligned} \mathbb{E}[\|D_k\|^2] &= \mathbb{E}[D_0^\top (I - \gamma H)^{2k} D_0] \\ &\geq (1 - \gamma\lambda_{\min})^{2k} \mathbb{E}[(D_0^\top q_{\max})^2]. \end{aligned}$$

If the initial difference vector  $D_0$  satisfies  $\mathbb{E}[(D_0^\top q_{\max})^2] \neq 0$ , combining the above inequality with the inequality (5) from Claim 1 yields the desired result.  $\square$

From Proposition 3, we note that a small distance between the two iterates  $\theta_k^{(1)}, \theta_k^{(2)}$  implies that the iterates  $\theta_k^{(1)}$  approximately converges to its stationary distribution  $\pi_\gamma$ . Since the eigenvector  $q_{\max}$  is unknown, we propose to track the ratio of the iterates difference  $\|D_k\| / \|D_0\|$  as an approximation, which is easy to compute.

**General Convex Setting** We next generalize our analysis to non-quadratic setting.

**Theorem 1 (General Convex).** *Suppose Assumptions A.1–A.4 hold and  $\gamma \in (0, \gamma_0)$ , where  $\gamma_0 = \min\{\frac{1}{4L}, \frac{2L}{\mu}\}$ . Then  $\forall \tau \in [\frac{4L}{\mu}, \infty)$ , we have the following result:*

$$\begin{aligned} & W_2^2(P_\gamma^k(\theta_0^{(1)}, \cdot), \pi_\gamma) \\ & \leq \left( \frac{\mathbb{E}[\|D_k\|^2]}{\mathbb{E}[\|D_0\|^2]} \right)^{1/\tau} \cdot \mathbb{E}_{\theta \sim \pi_\gamma} [\|\theta_0^{(1)} - \theta\|^2]. \end{aligned}$$

*Proof.* Consider the SGD update rule in eq. (1) and fix an arbitrary integer  $k \geq 0$ . The expected squared norm of the parameter difference after one iteration is given by:

$$\begin{aligned} & \mathbb{E}\|\theta_{k+1}^{(1)} - \theta_{k+1}^{(2)}\|^2 \\ &= \mathbb{E}\|(\theta_k^{(1)} - \gamma \nabla f_k(\theta_k^{(1)})) - (\theta_k^{(2)} - \gamma \nabla f_k(\theta_k^{(2)}))\|^2 \\ &= \mathbb{E}\|\theta_k^{(1)} - \theta_k^{(2)}\|^2 + \gamma^2 \mathbb{E}\|\nabla f_k(\theta_k^{(1)}) - \nabla f_k(\theta_k^{(2)})\|^2 \\ &\quad - 2\gamma \mathbb{E}\langle \nabla f_k(\theta_k^{(1)}) - \nabla f_k(\theta_k^{(2)}), \theta_k^{(1)} - \theta_k^{(2)} \rangle. \end{aligned} \quad (6)$$

By the smoothness assumption of  $f$ , we have:

$$\langle \nabla f(\theta_k^{(1)}) - \nabla f(\theta_k^{(2)}), \theta_k^{(1)} - \theta_k^{(2)} \rangle \leq L \|\theta_k^{(1)} - \theta_k^{(2)}\|^2.$$

Therefore, we can lower bound the LHS of Equation (6) as

$$\begin{aligned} \mathbb{E}\|\theta_{k+1}^{(1)} - \theta_{k+1}^{(2)}\|^2 &\geq \mathbb{E}\|\theta_k^{(1)} - \theta_k^{(2)}\|^2 \\ &\quad - 2\gamma L \mathbb{E}\|\theta_k^{(1)} - \theta_k^{(2)}\|^2 \\ &\quad + \gamma^2 \mathbb{E}\|\nabla f_1(\theta_k^{(1)}) - \nabla f(\theta_k^{(2)})\|^2. \end{aligned}$$

Moreover, by strong convexity of  $f$  together with the Cauchy-Schwarz inequality, we obtain

$$\|\nabla f(\theta_k^{(1)}) - \nabla f(\theta_k^{(2)})\| \geq \mu \|\theta_k^{(1)} - \theta_k^{(2)}\|.$$

It follows that

$$\mathbb{E}\|\theta_{k+1}^{(1)} - \theta_{k+1}^{(2)}\|^2 \geq (1 - 2\gamma L + \gamma^2 \mu^2) \mathbb{E}\|\theta_k^{(1)} - \theta_k^{(2)}\|^2.$$

Therefore, by induction on  $k$ , we obtain the lower bound:

$$\mathbb{E}\|\theta_k^{(1)} - \theta_k^{(2)}\|^2 \geq \varrho^k \mathbb{E}\|\theta_0^{(1)} - \theta_0^{(2)}\|^2, \quad (7)$$

where  $\varrho := 1 - 2\gamma L + \gamma^2 \mu^2$ . Applying Lemma 1 from Appendix A.1 yields:

$$\begin{aligned} \mathbb{E}\|D_k\|^2 / \mathbb{E}\|D_0\|^2 &\geq \varrho^k \geq (1 - \gamma\mu)^{k \cdot k_0} \geq (1 - \gamma\mu)^{k \cdot \tau}, \end{aligned}$$

where the last inequality follows from the fact that  $\tau \geq k_0 := 4L/\mu$ . Combining with Proposition 1, we obtain

$$\begin{aligned} W_2^2(P_\gamma^k(\theta_0, \cdot), \pi_\gamma) &\leq \rho^k \mathbb{E}_{\theta \sim \pi_\gamma} [\|\theta_0 - \theta\|^2] \\ &\leq (1 - \gamma\mu)^k \mathbb{E}_{\theta \sim \pi_\gamma} [\|\theta_0 - \theta\|^2] \\ &\leq (\mathbb{E}\|D_k\|^2 / \mathbb{E}\|D_0\|^2)^{1/\tau} \cdot \mathbb{E}_{\theta \sim \pi_\gamma} [\|\theta_0 - \theta\|^2]. \end{aligned} \quad \square$$

Built on Proposition 3 and Theorem 1, we conclude that the ratio of distances,  $\|D_k\|^2 / \|D_0\|^2$  naturally serves as an efficient statistic that detects the transition of the Markov chain from transience to stationarity. In particular, when the statistic falls below a predefined threshold, it indicates that  $\theta_k^{(1)}$  approximately saturates. Then the algorithm triggers a stepsize decay and repeats.

We formally describe our method in Algorithm 1, named *Coupling-based SGD*. Our diagnostic algorithm aims to detect the saturation of constant stepsize SGD  $(\theta_k^{(1)})_{k \geq 0}$  by tracking its distance to an auxiliary SGD sequence  $(\theta_k^{(2)})_{k \geq 0}$ . Given two initial points, two SGD iterates progress simultaneously using the same stepsize  $\gamma$  and mini-batch of data. At each iteration  $k$ , we calculate the ratio of distances,  $\|\theta_k^{(1)} - \theta_k^{(2)}\|^2 / \|\theta_0^{(1)} - \theta_0^{(2)}\|^2$ . If the ratio falls below a certain threshold  $\beta$ , the stepsize is decreased by a factor  $r$ , and the auxiliary iterate  $\theta_k^{(2)}$  is *re-initialized*. Setting an arbitrary re-initialization might take the auxiliary iterate far away from the vicinity of the optimum solution  $\theta^*$ . Consequently, it would take the auxiliary iterates  $\theta_k^{(2)}$  long time to re-entering the vicinity and sync with the primary iterates  $\theta_k^{(1)}$ . To address this issue, we propose to set the new initial point using the  $b$ -step backward iterate  $\theta_{k-b}^{(2)}$ , which stays close to the stationary neighborhood with an appropriate  $b$ . Our experiment in Section 4.3 demonstrates the robustness of our approach to the parameter  $b$ .

---

**Algorithm 1: Coupling-based SGD**

---

**Input:**  $\theta_0^{(1)}, \theta_0^{(2)}$   
**Parameters:** initial stepsize  $\gamma$ , stepsize decay factor  $r$ , backward steps  $b$ , threshold  $\beta$

- 1: **while**  $k \in \{1, 2, 3, \dots, n\}$  **do**
- 2:  $\theta_k^{(1)} \leftarrow \theta_{k-1}^{(1)} - \gamma \nabla f_{k-1}(\theta_{k-1}^{(1)})$
- 3:  $\theta_k^{(2)} \leftarrow \theta_{k-1}^{(2)} - \gamma \nabla f_{k-1}(\theta_{k-1}^{(2)})$
- 4:  $S \leftarrow \|\theta_k^{(1)} - \theta_k^{(2)}\|^2 / \|\theta_0^{(1)} - \theta_0^{(2)}\|^2$  {distance between two SGD sequences  $\approx \rho^k$ }
- 5: **if**  $S < \beta$  **then**
- 6:  $\gamma \leftarrow r \cdot \gamma$  {reduce the stepsize}
- 7:  $\theta_k^{(2)} \leftarrow \theta_{k-b}^{(2)}$  {update the new initial point}
- 8: **end if**
- 9: **end while**

---



---

**Algorithm 2: Adaptive Coupling-based SGD**

---

**Input:**  $\theta_0^{(1)}, \theta_0^{(2)}$   
**Parameter:** initial stepsize  $\gamma$ , stepsize decay factor  $r$ , backward steps  $b$ , initial threshold  $\beta$ , threshold decay factor  $\eta$ .

- 1: **while**  $k \in \{1, 2, 3, \dots, n\}$  **do**
- 2:  $\theta_k^{(1)} \leftarrow \theta_{k-1}^{(1)} - \gamma \nabla f_{k-1}(\theta_{k-1}^{(1)})$
- 3:  $\theta_k^{(2)} \leftarrow \theta_{k-1}^{(2)} - \gamma \nabla f_{k-1}(\theta_{k-1}^{(2)})$
- 4:  $S \leftarrow \|\theta_k^{(1)} - \theta_k^{(2)}\|^2 / \|\theta_0^{(1)} - \theta_0^{(2)}\|^2$  {distance between two SGD sequences  $\approx \rho^k$ }
- 5: **if**  $S < \beta$  **then**
- 6:  $\gamma \leftarrow r \cdot \gamma$  {reduce the stepsize}
- 7:  $\beta \leftarrow \eta \cdot \beta$  {reduce the threshold}
- 8:  $\theta_k^{(2)} \leftarrow \theta_{k-b}^{(2)}$  {update the new initial point}
- 9: **end if**
- 10: **end while**

---

**Adaptive Coupling-based SGD** As illustrated in Figure 1, as the stepsize  $\gamma$  gets smaller, the iterates in stationarity incur smaller approximation error and variance. Indeed, the stationary distribution  $\pi_\gamma$  is shown to satisfy  $\mathbb{E}_{\pi_\gamma}[\theta] - \theta^* = A\gamma + O(\gamma^2)$  and  $\mathbb{E}_{\pi_\gamma}[\|\theta - \theta^*\|^2] = A'\gamma + O(\gamma^2)$  (Dieuleveut, Durmus, and Bach 2020, Theorem 4). With a smaller  $\gamma$ , the stationary distribution  $\pi_\gamma$  lives in a smaller vicinity of  $\theta^*$ . Consequently, to accurately detect distributional convergence of the iterates, it makes sense to employ a more stringent criterion with smaller stepsize  $\gamma$ . To this end, we propose using an adaptive threshold  $\beta$  for the diagnostic statistic, as presented in Algorithm 2. In particular, as  $\gamma$  decreases, we also reduce the threshold  $\beta$  by a factor  $\eta \in (0, 1)$ .

## 4 Empirical Study

In this section, we present an empirical study of our proposed diagnostic method on various tasks.

## 4.1 Setup

**Algorithms.** Our main baselines include the Pflug’s diagnostic-based approach ISGD<sup>1/2</sup> (Chee and Toulis 2018), and the distance-based algorithm (Pesme, Dieuleveut, and Flammarion 2020). For these baseline methods, we carefully tune their parameters to achieve their best performance.

To demonstrate the effectiveness and robustness of our proposed diagnostic method, we conduct extensive experiments on a diverse set of problems, including (1) Logistic regression, (2) Least squares regression, (3) The 18-layer ResNet model (He et al. 2016), (4) Linear stochastic approximation with Markovian data, (5) SVM, (6) Uniformly convex functions, (7) Lasso. Due to space constraint, here we focus on the experiment setup and results of the first three tasks and defer details of other settings to Appendix B.<sup>2</sup>

**Logistic Regression.** The objective function  $f$  is defined as  $f(\theta) = \mathbb{E}[\log(1 + e^{-y_i \langle x_i, \theta \rangle})]$ . The inputs  $x_i$  are i.i.d. drawn from a multivariate normal distribution  $\mathcal{N}(0, H)$ , where  $H$  is a diagonal matrix with dimension  $d$ . The outputs  $y_i \in \{-1, 1\}$  are generated according to the logistic probabilistic model. We include averaged-SGD with step-sizes  $\gamma_k = 1/\sqrt{k}$  as a baseline, which achieves the optimal rate of  $\mathcal{O}(1/n)$  (Bach 2014). In addition, we evaluate averaged-SGD with step-sizes  $\gamma_k = C/\sqrt{k}$ , tuning the parameter  $C$  to achieve the best performance.

**Least Squares Regression.** The objective function  $f$  is given by  $f(\theta) = \frac{1}{2} \mathbb{E}[(y_i - \langle x_i, \theta \rangle)^2]$ . The inputs  $x_i$  are the same as in the logistic regression model. The outputs  $y_i$  are generated according to:  $y_i = \langle x_i, \theta^* \rangle + \varepsilon_i$ , where  $\varepsilon_i$  are i.i.d. noise following a normal distribution  $\mathcal{N}(0, \sigma^2)$ . We include averaged-SGD with a constant stepsize  $\gamma = 1/2R^2$ , where  $R^2 = \text{Tr}H$ , due to its optimal convergence rate  $\mathcal{O}(\sigma^2 d/n)$  (Bach and Moulines 2013). We also evaluate SGD with step-size  $\gamma_k = 1/\mu k$ , which achieves a rate of  $1/\mu k$ .

**ResNet on CIFAR-10.** We consider an 18-layer ResNet model (He et al. 2016) and train it on the CIFAR-10 dataset (Krizhevsky 2009). The training is conducted using an initial stepsize of 0.01, and a batch size of 128. To employ our coupling-based algorithm, we use PyTorch’s CustomScheduler() scheduler, with a stepsize decay factor of 0.1, a patience parameter of 10, and a threshold of 0.95. Specifically, we initialize two different ResNet models, and monitor the distance between their parameters. When the distance falls below the threshold for a specific number of epochs, we reduce the stepsize by the decay factor.

## 4.2 Main Results

Our first set of experiments aim to demonstrate the efficiency of our coupling-based diagnostic for detecting stationarity of the iterates. We compare our algorithms against SGD with different constant step-sizes  $\gamma$  for the least squares regression setting. The results are presented in Figure 2. Note that a larger constant stepsize leads to a faster convergence, but

<sup>2</sup>The code is available at <https://github.com/XiangLi/coupling-based-sgd>

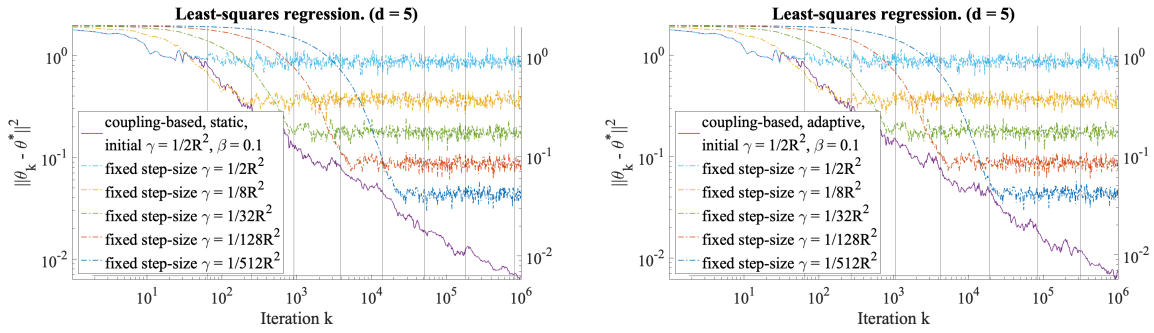


Figure 2: Effectiveness of our coupling-based statistic for stationarity diagnostic. Left: Algorithm 1 with static threshold; Right: Algorithm 2 with adaptive threshold. The vertical lines correspond to restarts of our coupling-based algorithms.

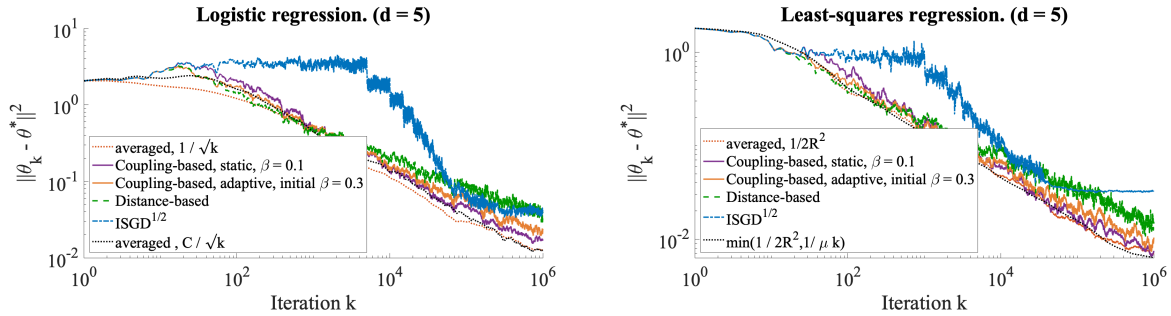


Figure 3: Logistic regression (left) and Least squares regression (right). The initial stepsize of coupling/distance-based and ISGD<sup>1/2</sup> is  $\gamma_0 = 4/R^2$  for logistic regression, and  $\gamma_0 = 1/2R^2$  for least squares. The errors are averaged over 10 replications.

inducing a larger error when reaching saturation. We observe that, for both the static (Algorithm 1) and adaptive coupling-based methods (Algorithm 2), the restarts (solid vertical line) mostly happen when the constant stepsize iterates start reaching saturation. That is, our coupling-based diagnostic statistic almost perfectly detects the stationarity of the iterates, without early on restart nor delay. Compared with the static variant (Figure 2 left), the adaptive version (Figure 2 right) becomes more conservative with the decreasing of threshold  $\beta$ , leading to more accurate detection of stationarity; see, for example, the 5th restart.

We next compare our method to other algorithms under logistic regression and least squares regression. For both settings, we consider synthetic dataset with size  $n = 1e6$  and dimension  $d = (5, 20, 50, 100)$ . The results of  $d = 5$  are presented in Figure 3, and results of other dimensions are provided in Appendix B. We observe that both static and adaptive variant of our algorithms consistently achieve superior performance across all settings, with more prominent advantages in higher-dimensional cases. In particular, for logistic regression, our methods achieve comparable performance as the averaged-SGD with the best tuned diminishing stepsizes  $\gamma_k = C/\sqrt{k}$ . Under the least squares regression, the performance of our methods is close to that with stepsizes  $1/\mu k$ , but without knowing the problem parameter  $\mu$ . We note that the Pflug’s diagnostic-based method ISGD<sup>1/2</sup> (blue lines) prematurely detects convergence and restarts too often. Con-

sequently, it results in small stepsizes quickly and thus stopping further progress towards the optimal point early on.

**ResNet.** To demonstrate the general applicability of our coupling-based algorithm, we applied it to deep learning tasks. As shown in Figure 5, our method (blue line) achieved a 94% test accuracy, comparable to state-of-the-art results for ResNet18 on this benchmark. Although the test accuracy with constant stepsize (red line) appears to plateau around epoch 75, manually adjusting the learning rate at this point (green line) resulted in trivial improvements. This suggests that the adjustment was likely made too early. In contrast, our algorithm effectively identifies the optimal adjustment point at epoch 145, ensuring meaningful improvements.

### 4.3 Robustness Study

Our last set of experiments aim to study the sensitivity of our methods to algorithm parameters, including the stepsize decay factor  $r$ , the threshold  $\beta$  and backward steps  $b$  for re-initialization of auxiliary sequence. Here we focus on our static coupling-based algorithm for logistic regression and least squares regression with dimension  $d = 100$ . Additional robustness results for adaptive variant and other settings are provided in Appendix B.

The results are presented in Figure 4, where each row corresponds to different data models, and each column corresponds to variations of different hyper-parameters. Specifically, we first perform the experiments with different step-

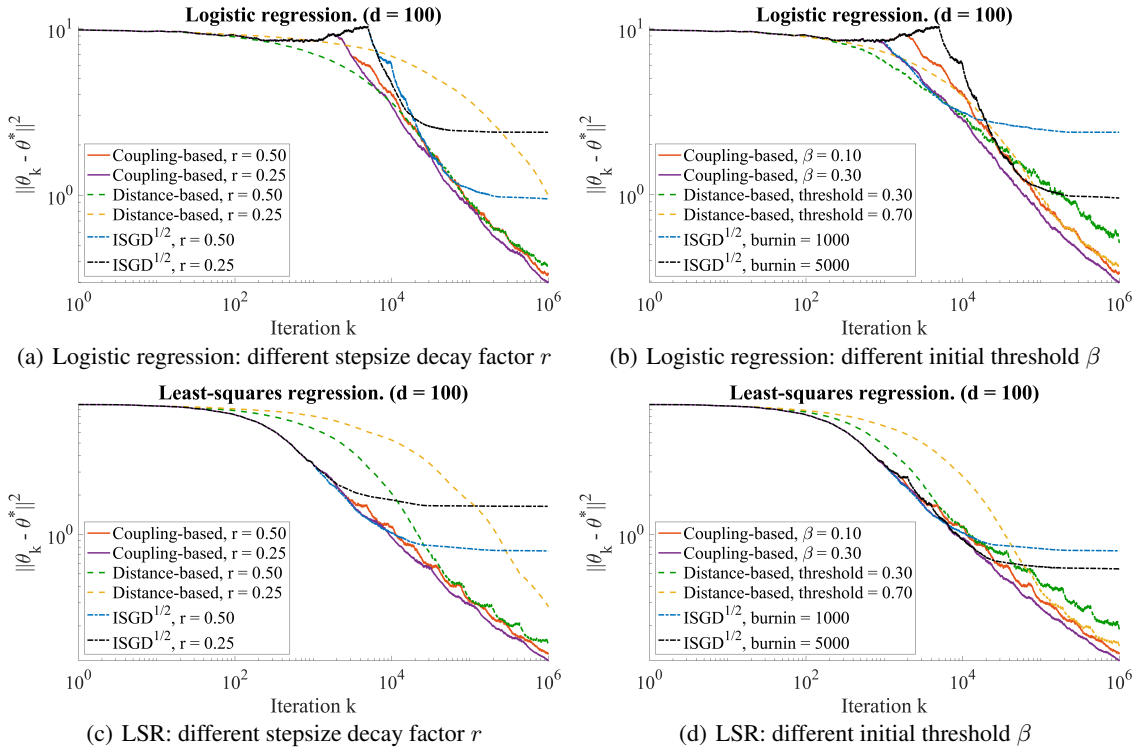


Figure 4: Robustness results under logistic regression and least squares regression (LSR) with  $d = 100$ .




Figure 5: ResNet-18 Test Accuracy on CIFAR-10 dataset. The initial stepsizes are 0.01.

size decay factor  $r$ , while keeping other parameters fixed. As illustrated in Figure 4(a) and 4(c), our algorithm is insensitive to the stepsize decay factor  $r$ , and consistently achieves the best performance for both logistic regression and least squares regression. In contrast, the performance of the distance-based algorithm and the  $\text{ISGD}^{1/2}$  algorithm varies significantly with  $r$ . We also test the sensitivity of our method and the distance-based algorithm w.r.t. the threshold parameter for convergence diagnostic, and the  $\text{ISGD}^{1/2}$  algorithm w.r.t. the burn-in parameter. Figure 4(b) and 4(d) show that our coupling-based method achieves similar superior performance with different  $\beta$  across various settings.

## 5 Conclusion and Future Work

In this paper, we study the convergence diagnostic of stochastic gradient descent with constant stepsize. We propose a novel diagnostic statistic based on the distance between two coupled SGD sequences. The distance is easy to compute and is shown to track the distributional convergence of SGD iterates under constant stepsize theoretically. A key advantage of our method is that the two coupled iterates can be run in parallel with the same data, which avoids an increase in the wall-clock time despite the additional computational cost.

We conduct an extensive range of experiments to numerically examine the properties of our proposed coupling-based methods. Our algorithms consistently outperform existing approaches. Moreover, our method can be applied to other stochastic optimization frameworks beyond SGD. For applications like reinforcement learning, where data is often limited, our adaptive stepsize approach can make more efficient use of data and accelerate convergence.

There are several research directions one can take to extend our work. While our empirical results demonstrate robustness of our approach to the threshold for convergence diagnostic, it would be valuable to develop a principled rule for choosing this parameter. Another future direction worth pursuing is to extend our analysis for other problem classes, particularly non-convex settings that satisfy structural conditions such as dissipativity or the generalized Polyak-Lojasiewicz condition, and general stochastic approximation frameworks.

## Acknowledgments

This project is supported in part by National Science Foundation (NSF) grants CNS-1955997, EPCN-2432546 and CAREER Award 2339794.

## References

- Bach, F. 2014. Adaptivity of averaged stochastic gradient descent to local strong convexity for logistic regression. *arXiv:1303.6149*.
- Bach, F.; and Moulines, E. 2013. Non-strongly-convex smooth stochastic approximation with convergence rate  $O(1/n)$ . *Advances in neural information processing systems*, 26.
- Bottou, L. 2010. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010: 19th International Conference on Computational Statistics Paris France, August 22-27, 2010 Keynote, Invited and Contributed Papers*, 177–186. Springer.
- Bottou, L. 2012. Stochastic gradient descent tricks. In *Neural Networks: Tricks of the Trade: Second Edition*, 421–436. Springer.
- Chee, J.; and Toulis, P. 2018. Convergence diagnostics for stochastic gradient descent with constant step size. *arXiv:1710.06382*.
- Dieuleveut, A.; Durmus, A.; and Bach, F. 2020. Bridging the gap between constant step size stochastic gradient descent and Markov chains. *The Annals of Statistics*, 48(3): 1348 – 1382.
- Foss, S. G.; and Tweedie, R. L. 1998. Perfect simulation and backward coupling. *Communications in Statistics. Stochastic Models*, 14(1-2): 187–203.
- Hazan, E.; and Kale, S. 2014. Beyond the Regret Minimization Barrier: Optimal Algorithms for Stochastic Strongly-Convex Optimization. *Journal of Machine Learning Research*, 15(71): 2489–2512.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huo, D.; Chen, Y.; and Xie, Q. 2023. Bias and extrapolation in Markovian linear stochastic approximation with constant stepsizes. In *Abstract Proceedings of the 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 81–82.
- Huo, D.; Zhang, Y.; Chen, Y.; and Xie, Q. 2024. The Collusion of Memory and Nonlinearity in Stochastic Approximation With Constant Stepsize. *arXiv preprint arXiv:2405.16732*.
- Krizhevsky, A. 2009. Learning Multiple Layers of Features from Tiny Images. *Master's thesis, University of Tront*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Lacoste-Julien, S.; Schmidt, M.; and Bach, F. 2012. A simpler approach to obtaining an  $O(1/t)$  convergence rate for the projected stochastic subgradient method. *arXiv preprint arXiv:1212.2002*.
- Lang, H.; Xiao, L.; and Zhang, P. 2019. Using statistics to automate stochastic optimization. *Advances in neural information processing systems*, 32.
- Lauand, C. K.; and Meyn, S. 2023. The curse of memory in stochastic approximation. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, 7803–7809. IEEE.
- Li, X.; and Xie, Q. 2024. Coupling-based Convergence Diagnostic and Stepsize Scheme for Stochastic Gradient Descent. *arXiv:2412.11341*.
- Merad, I.; and Gaïffas, S. 2023. Convergence and concentration properties of constant step-size SGD through Markov chains. *arXiv preprint arXiv:2306.11497*.
- Moulines, E.; and Bach, F. 2011. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. *Advances in neural information processing systems*, 24.
- Needell, D.; Ward, R.; and Srebro, N. 2014. Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm. *Advances in neural information processing systems*, 27.
- Pasupathy, R.; Honnappa, H.; and Hunter, S. R. 2019. Open Problem—Adaptive Constant-Step Stochastic Approximation. *Stochastic Systems*, 9(3): 307–310.
- Pesme, S.; Dieuleveut, A.; and Flammarion, N. 2020. On Convergence-Diagnostic based Step Sizes for Stochastic Gradient Descent. *CoRR*, abs/2007.00534.
- Pflug, G. C. 1983. On the determination of the step size in stochastic quasigradient methods.
- Rakhlin, A.; Shamir, O.; and Sridharan, K. 2011. Making gradient descent optimal for strongly convex stochastic optimization. *arXiv preprint arXiv:1109.5647*.
- Robbins, H.; and Monro, S. 1951. A stochastic approximation method. *The annals of mathematical statistics*, 400–407.
- Shamir, O.; and Zhang, T. 2013. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International conference on machine learning*, 71–79. PMLR.
- Sordello, M.; He, H.; and Su, W. 2020. Robust Learning Rate Selection for Stochastic Optimization via Splitting Diagnostic. *arXiv:1910.08597*.
- Wang, X.; Magnússon, S.; and Johansson, M. 2021. On the Convergence of Step Decay Step-Size for Stochastic Optimization. *arXiv:2102.09393*.
- Yaida, S. 2018. Fluctuation-dissipation relations for stochastic gradient descent. *arXiv:1810.00004*.
- Yu, L.; Balasubramanian, K.; Volgushev, S.; and Erdogdu, M. A. 2021. An analysis of constant step size SGD in the non-convex regime: Asymptotic normality and bias. *Advances in Neural Information Processing Systems*, 34: 4234–4248.