

Modified Step Size for Enhanced Stochastic Gradient Descent: Convergence and Experiments

Mahsa Soheil Shamaee*^{} and Sajad Fathi Hafshejani^{}

Abstract

This paper introduces a novel approach to enhance the performance of the stochastic gradient descent (SGD) algorithm by incorporating a modified decay step size based on $\frac{1}{\sqrt{t}}$. The proposed step size integrates a logarithmic term, leading to the selection of smaller values in the final iterations. Our analysis establishes a convergence rate of $O(\frac{\ln T}{\sqrt{T}})$ for smooth non-convex functions without the Polyak-Łojasiewicz condition. To evaluate the effectiveness of our approach, we conducted numerical experiments on image classification tasks using the Fashion-MNIST and CIFAR10 datasets, and the results demonstrate significant improvements in accuracy, with enhancements of 0.5% and 1.4% observed, respectively, compared to the traditional $\frac{1}{\sqrt{t}}$ step size. The source code can be found at <https://github.com/Shamaeem/LNSQRTStepSize>.

Keywords: Stochastic gradient descent, Decay step size, Convergence rate.

2020 Mathematics Subject Classification: 68T05, 68T20, 68T40.

How to cite this article

M. Soheil Shamaee and S. Fathi Hafshejani, Modified step size for enhanced stochastic gradient descent: convergence and experiments, *Math. Interdisc. Res.* 9 (3) (2024) 237–253.

1. Introduction

Stochastic gradient descent (SGD) has a rich historical background, originating from the influential work by Robbins and Monro [1]. In the realm of modern

*Corresponding author (E-mail: soheilshamaee@kashanu.ac.ir)

Academic Editor: Mehrdad Lakestani

Received 18 July 2023, Accepted 20 August 2023

DOI: 10.22052/MIR.2023.253279.1426

© 2024 University of Kashan

 This work is licensed under the Creative Commons Attribution 4.0 International License.

machine learning, SGD has emerged as a fundamental optimization algorithm for training deep neural networks (DNNs), which have achieved remarkable performance across diverse domains such as image classification [2, 3], object detection [4], and machine translation [5].

The selection of an appropriate step size, often referred to as the learning rate, plays a pivotal role in the convergence behavior of SGD. If the step size value is too large, it can prevent SGD iterations from reaching the optimal point, leading to instability and divergence. On the other hand, excessively small step size values can result in slow convergence and hinder the algorithm's ability to escape suboptimal local minima [6]. To tackle these challenges, researchers have proposed various schemes to determine the step size dynamically during the optimization process.

One notable approach is the Armijo line search method, initially introduced by Vaswani et al. [7], provide theoretical guarantees for strong-convex, convex, and non-convex objective functions. Another strategy, proposed by Gower et al. [8], combines a constant learning rate with a decreasing learning rate schedule. Their algorithm starts with a fixed learning rate and transitions to a decreasing schedule after a specified number of iterations, often determined by the problem's condition number. While this technique ensures convergence for strongly convex functions, it necessitates prior knowledge of the condition number and is not directly applicable to non-convex problems.

Decay step size is a commonly employed strategy in SGD to improve the convergence of optimization algorithms [9, 10]. By gradually reducing the step size over iterations, decay step size methods facilitate finer adjustments in parameter updates, leading to improved convergence behavior and enhanced optimization performance [11, 12]. Among the various decay step sizes used in SGD, the $\frac{1}{\sqrt{t}}$ step size has been widely used due to its ease of implementation and the significant advantage of not requiring derivative information. For instance, this step size has been recognized for its excellent performance in binary classification, making it an effective choice [12]. Additionally, it exhibits favorable efficiency in the context of deep neural networks.

During the training of deep neural networks, the use of this step size encounters a fundamental challenge. While the step size value decreases appropriately in the early iterations, it remains excessively large in the final iterations. This leads to the algorithm's inability to converge to the optimal point. As a result, the SGD algorithm with $\frac{1}{\sqrt{t}}$ step size fails to achieve the desired accuracy for deep neural networks. To address this limitation, we propose an enhanced version of the $\frac{1}{\sqrt{t}}$ step size that incorporates the $\ln t$ function into its definition. By introducing this modification, our goal is to improve the accuracy and loss function compared to the original $\frac{1}{\sqrt{t}}$ step size commonly used in SGD.

Smith [13] introduced the efficient method of setting the step size, known as the cyclical learning rate. Utilizing cyclical learning rates for training neural networks can yield substantial enhancements in accuracy, eliminating the need for manual

tuning and often requiring fewer iterations for convergence [13]. Building on this concept, Loshchilov and Hutter presented a warm restart technique for SGD in [10]. This approach eliminates the necessity of computing gradient information for adjusting the step size in each iteration. Warm restarts operate by initializing the learning rate to a specific value η_0 at each restart, scheduling its subsequent decrease [10]. Additionally, studies have revealed that warm-restarted SGD exhibits significantly improved efficiency, taking notably less time compared to traditional learning rate adjustment strategies [14]. Over recent years, a variety of step sizes accompanied by warm restarts have been proposed [6, 15]. Extending the notion of cosine step size, Vrbančić introduced three distinct step sizes accompanied by warm restarts [14].

Building upon the insights from previous research, we present a novel approach in this work that utilizes a novel step size combined with the warm restarts technique for SGD. The key contributions of this paper can be summarized as follows:

- The new step size exhibits a distinct behavior compared to the $\frac{1}{\sqrt{t}}$ step size. By incorporating both $\frac{1}{\sqrt{t}}$ and $\ln t$, the step length gradually decreases in the final iterations, leading to convergence towards the optimal point. The impact of this modification will be demonstrated through the numerical results.
- We demonstrate the convergence rate of $O(\frac{\ln T}{\sqrt{T}})$ for smooth non-convex functions, without requiring the Polyak-Łojasiewicz (PL) condition.
- We evaluate the performance of the new step size through extensive experiments on two popular image classification datasets, that is, Fashion-MNIST and CIFAR10. The results indicate significant improvements in accuracy, with enhancements of 0.5% and 1.4% observed, respectively, when compared to the traditional $\frac{1}{\sqrt{t}}$ step size. In addition, we conduct SGD experiments on binary classification tasks using five diverse datasets: a1a, a2a, mushrooms, rcv1, and w1a. The results demonstrate that the new proposed step size consistently outperforms other step sizes in terms of accuracy and loss function.

The paper is organized as follows: Section 2 introduces the new step size, providing an overview of its formulation and properties. In Section 3, we analyze the convergence rates of the proposed step size on smooth non-convex functions, demonstrating its impressive $O(\frac{\ln T}{\sqrt{T}})$ convergence rate. Section 4 presents and discusses the numerical results obtained using the new decay step size, highlighting its effectiveness in improving optimization performance. Finally, Section 5 concludes the paper by summarizing the findings and drawing insightful conclusions from our study.

In this paper, we use the following notational conventions: The Euclidean norm of a vector is denoted by $\|\cdot\|$. The non-negative orthant and positive orthant of \mathbb{R}^d are denoted by \mathbb{R}_+^d and \mathbb{R}_{++}^d , respectively. We also use the notation $f(t) = O(g(t))$

to indicate that there exists a positive constant ω such that $f(t) \leq \omega g(t)$ for all $t \in \mathbb{R}_{++}$.

2. New step size

In this section, we briefly introduce the main optimization problem and state some assumptions. Afterward, we will present the new step size and algorithm.

We consider the following optimization problem:

$$\min_{x \in \mathbb{R}^d} f(x) = \min_{x \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n f_i(x), \quad (1)$$

where $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is the loss function for the i -th training sample over the variable $x \in \mathbb{R}^d$ and n denotes the number of samples. This minimization problem plays a key role in machine learning. Several iterative approaches for solving Equation (1) are known [16], and SGD is particularly popular when the dimensionality, n , is extremely large [1, 17]. SGD uses a random training sample $i_k \in \{1, 2, \dots, n\}$ to update x using the rule:

$$x_{k+1} = x_k - \eta_k \nabla f_{i_k}(x_k), \quad (2)$$

in which η_k is the step size used in iteration k and $\nabla f_{i_k}(x)$ is the (average) gradient of the loss function(s) [7].

2.1. Assumptions

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be the objective function, and consider the SGD algorithm. We make the following assumptions [11]:

- **A1:** The function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is L -smooth, which implies that for all x and y in the domain of f , we have:

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{L}{2} \|y - x\|^2, \quad (3)$$

where, $\nabla f(x)$ denotes the gradient of f at point x , and L is a positive constant representing the Lipschitz constant of f .

- **A2:** For any iteration $t \in \{1, 2, \dots, T\}$ of the SGD algorithm, we assume that the expected square norm of the difference between the stochastic gradient g_t and the true gradient $\nabla f(x_t)$ at the current iterate x_t is bounded as follows:

$$\mathbb{E}_t [\|g_t - \nabla f(x_t)\|^2] \leq \sigma^2, \quad (4)$$

where σ^2 is a positive constant.

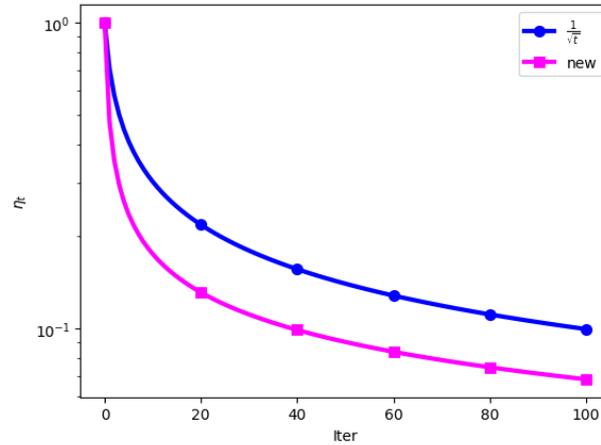


Figure 1: The comparison between the $\frac{1}{\sqrt{t}}$ step size and the modified step size incorporating $\ln t$ function.

2.2. The new step size

In this paper, we address a limitation associated with the $\frac{1}{\sqrt{t}}$ step size, where its value fails to decrease adequately during the final iterations. This characteristic poses a challenge in reaching the optimal point in some problems. To overcome this limitation, we propose a modified step size approach that combines the $\frac{1}{\sqrt{t}}$ function with the logarithmic function, $\ln t$, in an effort to effectively reduce the step size.

The motivation behind incorporating the $\ln t$ function lies in its gradual growth pattern, which enables a more controlled reduction in the step size when compared to the original $\frac{1}{\sqrt{t}}$ step size. By introducing the $\ln t$ function into the formulation, we aim to achieve a more refined and optimized step size throughout the optimization process. In this regard, we define the new step size as:

$$\eta_t = \frac{\eta_0}{\sqrt{t} + \ln t}, \quad \eta_0 \in (0, 1], \quad t = 1, 2, \dots, T. \tag{5}$$

Figure 1 illustrates the behavior of two step sizes: the original $\frac{1}{\sqrt{t}}$ step size and the newly proposed step size. The graph visually demonstrates how the new step size consistently selects smaller values, particularly during the final iterations of the SGD algorithm.

2.3. Algorithm

In this paper, we employ the warm restart Algorithm 1 with the same number of epochs in the inner loop, i.e., T . Algorithm 1 is initiated with the provided

initial step size η_0 , the number of inner iterations T , the number of outer epochs l , and the initial point x_0 . The algorithm consists of the outer and the inner loop. In each inner loop, the SGD with the new step size is executed and the point is updated. It is important to note that [Algorithm 1](#) was introduced in [\[11\]](#). It becomes evident that when $l = 1$, [Algorithm 1](#) transforms into the SGD algorithm.

Algorithm 1 SGD with warm restarts based on the new step size.

Input: Initial step size η_0 , initial point x_0 , the number of outer and inner iterations, i.e., l and T .

```

for  $i = 1, \dots, l$  do
  | for  $t = 1, \dots, T$  do
  | | Run SGD with the new step size  $\frac{\eta_0}{\sqrt{t} + \ln t}$ 
  | end
end

```

3. Convergence

In this section, we demonstrate that [Algorithm 1](#) using the newly proposed step size achieves a convergence rate of $O(\frac{\ln T}{\sqrt{T}})$ for smooth non-convex functions without the PL condition. Note that, the PL condition initially proposed by Polyak [\[18\]](#) and Łojasiewicz [\[19\]](#), stands as a fundamental cornerstone in demonstrating linear convergence rates for non-convex functions [\[11\]](#). To establish the convergence results, we initially demonstrate an $O(\frac{\ln T}{\sqrt{T}})$ convergence rate for a single outer iteration of [Algorithm 1](#), which corresponds to the SGD algorithm. Subsequently, we extend the proof to encompass l outer iterations. In this regard, we first introduce two preliminary lemmas [\[9, 11\]](#).

Lemma 3.1. *For the new step size given by (5), we have:*

$$\sum_{t=1}^T \eta_t \geq \eta_0(\sqrt{T} - 1).$$

Proof. To prove this lemma, we utilize the fact that $\ln t < \sqrt{t}$ for all $t \in [1, T]$. Thus, we have:

$$\sum_{t=1}^T \eta_t = \eta_0 \sum_{t=1}^T \frac{1}{\sqrt{t} + \ln t} \geq \eta_0 \sum_{t=1}^T \frac{1}{2\sqrt{t}} = \eta_0 \int_1^T \frac{1}{2\sqrt{t}} dt = \eta_0(\sqrt{T} - 1), \quad (6)$$

where the first inequality is derived from $\ln t \leq \sqrt{t}$ for all $t \geq 1$. □

Lemma 3.2. For the new step size given by (5), we have:

$$\sum_{t=1}^T \eta_t^2 \leq \eta_0^2 \ln T.$$

Proof. To prove this lemma, we use the fact that $\ln t \geq 0$ and $\sqrt{t} + \ln t \geq \sqrt{t}$ for all $t \geq 1$. Hence, we have:

$$\sum_{t=1}^T \eta_t^2 = \eta_0^2 \sum_{t=1}^T \frac{1}{(\sqrt{t} + \ln t)^2} \leq \eta_0^2 \sum_{t=1}^T \frac{1}{t} = \eta_0^2 \ln T, \tag{7}$$

where the first inequality is obtained from the fact that $\ln t \geq 0$ for all $t \geq 1$. \square

These preliminary lemmas provide important insights and bounds that will be used to establish the convergence results for the modified step size in subsequent sections.

Lemma 3.3 (Lemma 7.1 in [9]). Assuming that f is an L -smooth function and Assumption (A2) is satisfied, if $\eta_t \leq \frac{1}{cL}$, then SGD guarantees:

$$\frac{\eta_t}{2} \mathbb{E}[\|\nabla f(x_t)\|^2] \leq \mathbb{E}[f(x_t)] - \mathbb{E}[f(x_{t+1})] + \frac{L\eta_t^2 \sigma^2}{2}. \tag{8}$$

The following theorem provides $O(\frac{\ln T}{\sqrt{T}})$ rate of convergence for smooth non-convex functions without PL condition.

Theorem 3.4. Under Assumptions A1 and A2, and for $c > 1$, a single outer iteration of Algorithm 1, which corresponds to the SGD algorithm with the new proposed step sizes using $\eta_0 = \frac{1}{cL}$ guarantees the following inequality:

$$\mathbb{E}[\|\nabla f(\bar{x}_T)\|^2] \leq \frac{\ln T}{\sqrt{T} - 1} \left[\frac{2Lc(f(x_1) - f^*)}{\ln T} + \frac{\sigma^2}{Lc} \right],$$

where \bar{x}_T is a random iterate drawn from the sequence $\{x_t\}_{t=1}^T$ with probability $\mathbb{P}[\bar{x}_T = x_t] = \frac{\eta_t}{\sum_{t=1}^T \eta_t}$.

Proof. Using the definition of \bar{x}_T and Lemma 3.3, we have:

$$\begin{aligned} \mathbb{E}[\|\nabla f(\bar{x}_T)\|^2] &= \frac{\eta_t \mathbb{E}[\|\nabla f(x_t)\|^2]}{\sum_{t=1}^T \eta_t} \leq \frac{2 \sum_{t=1}^T [\mathbb{E}[f(x_t)] - \mathbb{E}[f(x_{t+1})]]}{\sum_{t=1}^T \eta_t} + \frac{L\sigma^2 \sum_{t=1}^T \eta_t^2}{\sum_{t=1}^T \eta_t} \\ &\leq \frac{2(f(x_1) - f^*)}{\sum_{t=1}^T \eta_t} + \frac{L\sigma^2 \sum_{t=1}^T \eta_t^2}{\sum_{t=1}^T \eta_t} \\ &\leq \frac{2(f(x_1) - f^*)}{\eta_0(\sqrt{T} - 1)} + \frac{L\sigma^2 \eta_0 \ln T}{(\sqrt{T} - 1)} \\ &= \frac{2cL}{(\sqrt{T} - 1)} (f(x_1) - f^*) + \frac{\sigma^2 \ln T}{Lc(\sqrt{T} - 1)}, \end{aligned} \tag{9}$$

where the third inequality is obtained by using [Lemmas 3.1](#) and [3.2](#). The expression [\(9\)](#) can be rewritten as follows:

$$\mathbb{E}[\|\nabla f(\bar{x}_T)\|^2] \leq \frac{\ln T}{\sqrt{T}-1} \left[\frac{2Lc(f(x_1) - f^*)}{\ln T} + \frac{\sigma^2}{Lc} \right]. \quad (10)$$

□

Remark 1. [Theorem 3.4](#) implies that:

$$\mathbb{E}[\|\nabla f(\bar{x}_T)\|^2] \leq O\left(\frac{\ln T}{\sqrt{T}}\right).$$

This result demonstrates that a single outer iteration of [Algorithm 1](#), which corresponds to the SGD based on the new modified step size enjoys an $O(\frac{\ln T}{\sqrt{T}})$ rate of convergence for smooth non-convex functions without the PL condition. Remarkably, this rate of convergence matches that of the traditional $\frac{1}{\sqrt{t}}$ step size [\[9\]](#).

Utilizing the outcomes derived from [Theorem 3.4](#), we can now calculate the convergence rate for the warm restart SGD algorithm.

Corollary 3.5. (*SGD with warm restarts*): Under Assumptions A1 and A2, for a given value of T and $\eta_0 = \frac{1}{cL}$, [Algorithm 1](#) guarantees the following convergence:

$$\mathbb{E}|\nabla f(\bar{x}_T)|^2 \leq \frac{\ln T}{\sqrt{T}-1} \left(\frac{2lcL}{\ln T} (f(\tilde{x}_1) - f^*) + \frac{\sigma^2 l}{Lc} \right),$$

where $f(\tilde{x}_1) = \max_i f(x_{1_i})$ for $i = 1, 2, \dots, l$.

Proof. [Theorem 3.4](#) is true for all $i = 1, 2, \dots, l$. Therefore, we have:

$$\begin{aligned} \min_i (\mathbb{E}\|\nabla f(\bar{x}_T)\|^2) &\leq \sum_{i=1}^l \mathbb{E}\|\nabla f(\bar{x}_{T_i})\|^2 \\ &\leq \sum_{i=1}^l \left(\frac{2cL}{\sqrt{T}-1} (f(x_{1_i}) - f^*) + \frac{\sigma^2 \ln T}{Lc(\sqrt{T}-1)} \right) \\ &\leq l \max_i \left(\frac{2cL}{\sqrt{T}-1} (f(x_{1_i}) - f^*) + \frac{\sigma^2 \ln T}{Lc(\sqrt{T}-1)} \right) \\ &= \frac{2lcL}{\sqrt{T}-1} (f(\tilde{x}_1) - f^*) + \frac{\sigma^2 l \ln T}{Lc(\sqrt{T}-1)} \\ &= \frac{\ln T}{\sqrt{T}-1} \left(\frac{2lcL}{\ln T} (f(\tilde{x}_1) - f^*) + \frac{\sigma^2 l}{Lc} \right), \end{aligned}$$

in which $f(\tilde{x}_1) = \max_i \{f(x_{1_i})\}$.

□

4. Numerical results

In this section, we performed two sets of experiments to assess the effectiveness of our proposed scheme. The first series of experiments involved classifying images on two different datasets: Fashion-MNIST and CIFAR10. These datasets are commonly used in computer vision research for image classification tasks. The second series of experiments focused on the binary classification of patterns, using five different datasets: a1a, a2a, mushrooms, rcv1, and w1a. These datasets cover a diverse range of patterns and are commonly used in machine learning research for binary classification tasks. To assess the performance of our proposed approach, we compared it with state-of-the-art methods through experimental studies. By conducting these comparisons, we gain insights into the effectiveness of our approach and how it stacks up against existing techniques. Now, let's dive deeper into the mentioned methods, datasets, and the learning model applied for the classification task.

4.1. Methods

Here, we conduct a comprehensive comparison study to evaluate various step sizes. We consider the following step sizes:

- $\eta_t = \text{constant}$,
- $\eta_t = \frac{\eta_0}{1+\alpha\sqrt{t}}$,
- $\eta_t = \frac{\eta_0}{1+\alpha t}$,
- $\eta_t = \frac{\eta_0}{2} \left(1 + \cos \frac{t\pi}{T}\right)$,
- $\eta_t = \eta_0 \left(\frac{1}{1+\alpha(\sqrt{t}+\ln t)}\right)$.

We have various step size update strategies with the following names: SGD with constant step size, step size with $O(\frac{1}{\sqrt{t}})$ decay, step size with $O(\frac{1}{t})$ decay, cosine step size update, and the new step decay method. The parameter t represents the iteration number of the inner loop, and each outer iteration involves multiple iterations for training on mini-batches.

Additionally, we compare the results of the newly proposed step decay method with Adam [20], SGD+Armijo method [7], PyTorch's ReduceLROnPlateau scheduler5 (abbreviated as ReduceLROnPlateau), and stagewise step size. In this comparison, we refer to the points where the step size decreases in the stagewise step decay method as milestones. It's worth noting that since Nesterov momentum is used in all SGD variants, the stagewise step decay method essentially covers the performance of multistage accelerated algorithms (e.g., [21]).

4.2. Multi-class classification using deep networks

Fashion-MNIST is a dataset that includes a training set of 50,000 and a test set of 10,000 grayscale images. Each image in this dataset has a size of 28×28 pixels. For the classification task on this dataset, we employed a Convolutional Neural Network (CNN) model. Let me dissect the architecture of the CNN model we used for this task. It consists of two convolutional layers. The size of the filter used in each convolutional layer is 5×5 . We applied padding of 2 to ensure that the spatial dimensions of the output feature maps match the input size. The model also incorporates two max-pooling layers with a kernel size of 2×2 . Max-pooling reduces the spatial dimensions and helps in capturing important features while discarding unnecessary details. To further process the extracted features, the model includes two fully connected layers. Each of these layers has 1024 hidden nodes. The activation function used for the hidden nodes is the Rectified Linear Unit (ReLU), which helps introduce non-linearity into the model and allows it to learn complex patterns effectively.

In order to prevent overfitting, a dropout technique is applied with a probability of 0.5 in the hidden layer of the deep model. Dropout randomly sets a fraction of the input units to zero during training, forcing the network to learn robust representations. To evaluate and compare the performance of different algorithms, we utilized the cross-entropy function as the loss function.

The CIFAR10 dataset is composed of 60,000 color images, each with a size of 32×32 pixels. These images are divided into 10 different classes, and each class contains 6,000 images. The dataset is further split into a training set of 50,000 images and a test set of 10,000 images. During the training process on this dataset, a batch size of 128 is utilized. This means that each epoch of training comprises 390 iterations. To evaluate the performance of the algorithms on the CIFAR10 dataset, we employed a deep learning architecture known as the 20-layer Residual Neural Network (ResNet). ResNet was introduced by [22] and has proven to be highly effective in various computer vision tasks. The loss function used in this model is the cross-entropy loss.

A grid search was conducted to determine the initial values for parameters η_0 and α for Fashion-MNIST and CIFAR10, resulting in $\{0.05, 0.15\}$ and $\{0.0253, 0.025\}$, respectively. For the remaining step sizes, the initial values from [11] were employed.

4.3. Binary classification with kernels

This series of experiments aims to classify the data into two classes using the Radial Basis Function (RBF) kernel without introducing any regularization techniques. We experiment with five standard datasets: a1a, a2a, mushrooms, rcv1 and w1a from LIBSVM [23]. These datasets have been widely adopted in the machine learning community [14], which allows researchers to compare the performance of various classification algorithms on the same standardized data. To create these

Table 1: Details for binary classification datasets.

Data set	Dimension (d)	Training Set Size	Test Set Size	Kernel Bandwidth
a1a	123	24765	6191	1
a2a	123	24237	6060	1
mushrooms	112	6499	1625	0.5
rcv1	47236	16194	4048	0.25
w1a	300	37818	9454	1

datasets, we exclusively utilized the training sets provided by the LIBSVM library [23], and performed an 80 : 20 split, where 80% of the data was designated as the training set and the remaining 20% was set aside as the test set.

The a1a dataset is a widely used benchmark dataset in the field of machine learning and data mining. It consists of binary classification tasks where the goal is to predict whether a person’s income exceeds 50,000 based on various attributes such as age, education, marital status, occupation, etc. The a2a dataset is another well-known benchmark dataset that is used for binary classification tasks. It is similar to the a1a dataset in that it focuses on predicting income, but it contains additional attributes and a larger number of instances.

The mushrooms dataset is a popular dataset used in the field of classification. It contains attributes of various mushrooms, such as cap shape, cap color, odor, gill size, etc., and the target variable is whether the mushroom is edible or poisonous. The rcv1 dataset, also known as Reuters Corpus Volume 1, is a large collection of news articles from Reuters, a major news agency. It consists of over 800,000 documents categorized into topics such as business, politics, sports, health, etc. The rcv1 dataset is often used for tasks such as text classification, information retrieval, and natural language processing research. The w1a is derived from the web page dataset. It has two categories and 300 sparse binary keyword attributes. 2,477 examples are used, among which 72 examples are positive. Table 1 illustrates the details for each mentioned dataset.

For all the datasets mentioned, the initial values of parameters η_0 and α in the new proposed step size, used for binary classification with kernels, are set to 0.05 and 0.00001 respectively. On the other hand, the initial values from the study by [11] were utilized for the remaining step sizes.

4.4. Results and discussion

Based on Figure 2 and Table 2, the newly suggested step size demonstrates impressive results. In the Fashion-MNIST dataset, it achieves a training loss that is nearly zero, comparable to the performance of well-established methods like SGD+Armijo. Furthermore, it outperforms all other methods in terms of test accuracy. In the CIFAR10 dataset, SGD with the new step size outperforms the

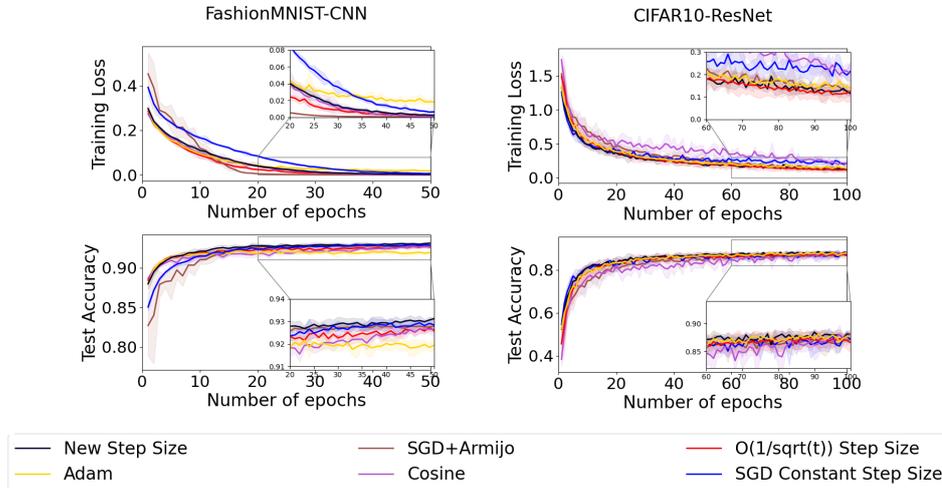


Figure 2: Comparison of new proposed step size and five other step sizes on Fashion-MNIST and CIFAR10 datasets.

previously studied method with a step size of $O(\frac{1}{\sqrt{t}})$, which is considered the best based on both training loss and test accuracy as illustrated in Figure 2. Table 3 also emphasises the superiority of the SGD with new step size over the $\frac{1}{\sqrt{t}}$ step decay in both Fashion-MNIST and CIFAR10 datasets.

Based on the observations made in Figures 3 and 4, it can be seen that the implementation of SGD with a new step decay consistently achieves the highest performance in a1a, a2a, mushrooms, rcv1, and w1a datasets. This is evident in terms of both the training loss and test accuracy in all mentioned datasets. Additionally, this method demonstrates a faster convergence to a satisfactory solution compared to the alternative approaches. Table 4 illustrates that the introduced step size resulted in a reduction of the loss functions by 0.01, 0.02, 0.03, 0.003, and 0.02 for the a1a, a2a, mushrooms, rcv1, and w1a datasets, respectively, in comparison to the $\frac{1}{\sqrt{t}}$ step size. Additionally, the new step size enhances the accuracy of the a1a dataset by 0.7% when contrasted with the $\frac{1}{\sqrt{t}}$ step size.

5. Conclusion

This paper introduced a novel approach to enhance the stochastic gradient descent (SGD) algorithm by modifying the decay step size based on $\frac{1}{\sqrt{t}}$. We established a convergence rate of $O(\frac{\ln T}{\sqrt{T}})$ for smooth non-convex functions without the Polyak-Łojasiewicz condition. The numerical experiments conducted on image classification tasks using the Fashion-MNIST, and CIFAR10 datasets demonstrated the

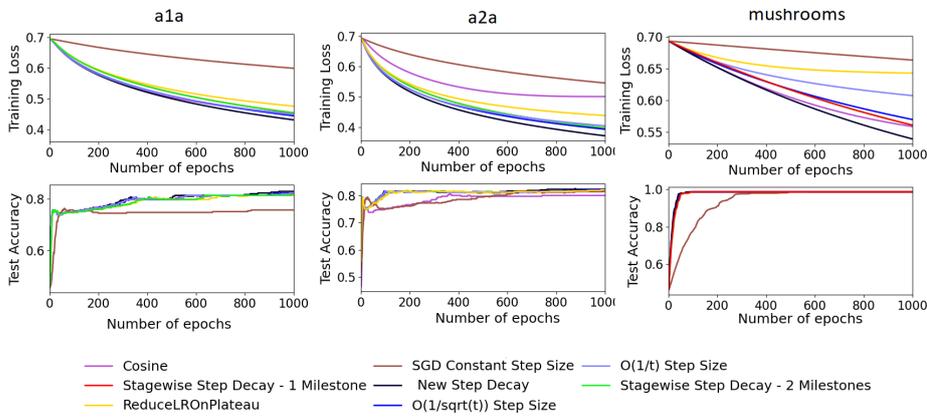


Figure 3: Comparison of new proposed step size and seven other step sizes on a1a, a2a, and mushrooms datasets.

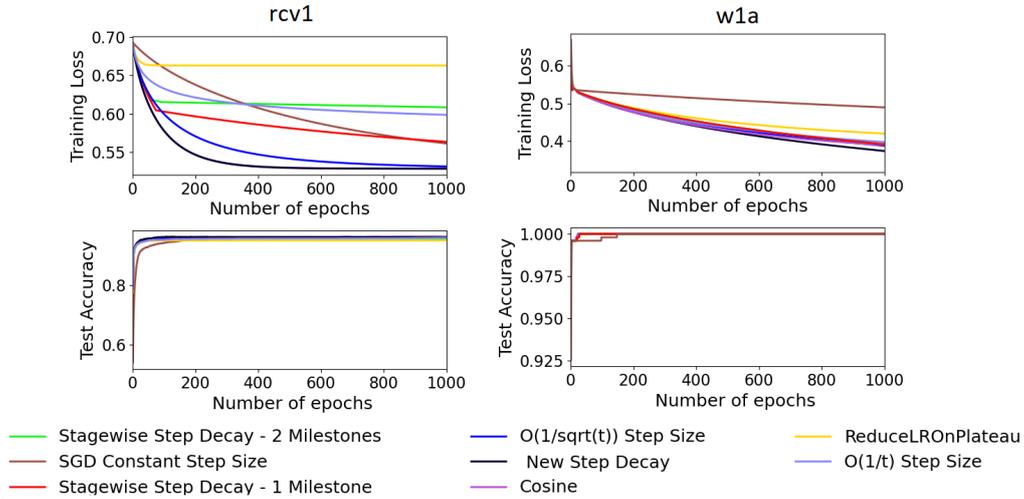


Figure 4: Comparison of new proposed step size and seven other step sizes on rcv1 and w1a datasets.

Table 2: The average final training loss and test accuracy on the Fashion-MNIST dataset, along with the 95% confidence intervals obtained from 5 runs starting from different random seeds.

Methods	Training loss	Test accuracy
Constant Step Size	0.0007 ± 0.0003	0.9299 ± 0.0016
$O(\frac{1}{\sqrt{t}})$ Step Size	0.0011 ± 0.0003	0.9261 ± 0.0007
Adam	0.0131 ± 0.0017	0.9166 ± 0.0019
SGD+Armijo	$6.73E-05 \pm 0.00$	0.9277 ± 0.0012
Cosine step size	$0.0004 \pm 1.1E - 05$	0.9284 ± 0.0005
New Step Size	0.002 ± 0.00	0.931 ± 0.00

Table 3: Performance comparison of $\frac{1}{\sqrt{t}}$ and new step size on Fashion-MNIST and CIFAR10 datasets.

Step sizes	$\frac{1}{\sqrt{t}}$	$\frac{1}{\sqrt{t}}$	$\frac{1}{\sqrt{t}+\ln t}$	$\frac{1}{\sqrt{t}+\ln t}$
Data set	Training loss	Test accuracy	Training loss	Test accuracy
FashionMNIST	0.00 ± 0.00	0.926 ± 0.00	0.00 ± 0.00	0.93 ± 0.00
CIFAR10	0.12 ± 0.02	0.87 ± 0.00	0.08 ± 0.01	0.89 ± 0.00

effectiveness of the proposed approach, with accuracy improvements of 0.5%, and 1.4%, respectively, over the traditional $\frac{1}{\sqrt{t}}$ step size. Furthermore, in the case of binary datasets, the introduced step size exhibited improvements in the loss function by 0.01, 0.02, 0.03, 0.003, and 0.02 for the a1a, a2a, mushrooms, rcv1, and w1a datasets, respectively, when compared to the $\frac{1}{\sqrt{t}}$ step size.

As a result of this paper, the combination of the $\frac{1}{\sqrt{t}}$ step size with the $\ln t$ function has led to an enhancement in the efficiency of the $\frac{1}{\sqrt{t}}$ step size. This finding suggests a potential avenue for future research where other step sizes could

Table 4: Performance comparison of $\frac{1}{\sqrt{t}}$ and new step size on datasets for binary classification task.

Step sizes	$\frac{1}{\sqrt{t}}$	$\frac{1}{\sqrt{t}}$	$\frac{1}{\sqrt{t}+\ln t}$	$\frac{1}{\sqrt{t}+\ln t}$
Data set	Training loss	Test accuracy	Training loss	Test accuracy
a1a	0.44	0.822	0.43	0.829
a2a	0.39	0.82	0.37	0.82
mushrooms	0.57	0.98	0.54	0.98
rcv1	0.531	0.96	0.528	0.96
w1a	0.39	1	0.37	1

be similarly combined with suitable functions to boost their efficiency.

Conflicts of Interest. The authors declare that they have no conflicts of interest regarding the publication of this article.

References

- [1] H. Robbins and S. Monro, A stochastic approximation method, *Ann. Math. Stat.* (1951) 400 – 407.
- [2] A. Krizhevsky and G. Hinton, *Learning Multiple Layers of Features from Tiny Images*, University of Toronto, Tech. rep. 2009.
- [3] A. Krizhevsky, I. Sutskever and G. E. Hinton, Imagenet classification with deep convolutional neural networks, *Commun. ACM* **60** (2017) 84 – 90, <https://doi.org/10.1145/3065386>.
- [4] J. Redmon and A. Farhadi, Yolo9000: better, faster, stronger, *In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, (2017) 7263 – 7271.
- [5] J. Zhang and C. Zong, Deep neural networks in machine translation: an overview, *IEEE Intell. Syst.* **30** (2015) 16 – 25, <https://doi.org/10.1109/MIS.2015.69>.
- [6] P. Mishra and K. Sarawadekar, Polynomial learning rate policy with warm restart for deep neural network, *In: TENCN 2019-2019 IEEE Region 10 Conf. (TENCN)*, (2019) 2087 – 2092.
- [7] S. Vaswani, A. Mishkin, I. Laradji, M. Schmidt, G. Gidel and S. Lacoste-Julien, Painless stochastic gradient: interpolation, line-search, and convergence rates, *In. Proc. of the 33rd Int. Conf. on Neural Information Processing Systems*, (2019) 3732 – 3745.
- [8] R. M. Gower, N. Loizou, X. Qian, A. Sailanbayev, E. Shulgin and P. Richtárik, Sgd: general analysis and improved rates, *In: Int. Conf. on Machine Learning*, California, (2019) 5200 – 5209.
- [9] X. Wang, S. Magnússon and M. Johansson, On the convergence of step decay step-size for stochastic optimization, *Adv. Neural Inf. Process. Syst.* **34** (2021) 14226 – 14238.
- [10] I. Loshchilov and F. Hutter, SGDR: Stochastic gradient descent with warm restarts, *arXiv preprint arXiv:1608.03983*, 2016.
- [11] X. Li, Z. Zhuang and F. Orabona, A second look at exponential and cosine step sizes: simplicity, adaptivity and performance, *In: Int. Conf. on Machine Learning*, (2021) 6553 – 6564 PMLR.

- [12] W. Tao, S. Long, G. Wu and Q. Tao, The role of momentum parameters in the optimal convergence of adaptive Polyak's heavy-ball methods, *arXiv preprint arXiv:2102.07314* 2021.
- [13] L. N. Smith, Cyclical learning rates for training neural networks, *In: 2017 IEEE Winter Conf. on Applications of Computer Vision (WACV)*, (2017) 464 – 472, <https://doi.org/10.1109/WACV.2017.58>.
- [14] G. Vrbančič and V. Podgorelec, Efficient ensemble for image-based identification of pneumonia utilizing deep CNN and SGD with warm restarts, *Expert Syst. Appl.* **187** (2022) p. 115834, <https://doi.org/10.1016/j.eswa.2021.115834>.
- [15] G. Xu, H. Cao, Y. Dong, C. Yue and Y. Zou, Stochastic gradient descent with step cosine warm restarts for pathological lymph node image classification via pet/ct images, *In 2020 IEEE 5th International Conference on Signal and Image Processing (ICSIP)* (2020) 490 – 493.
- [16] J. Nocedal and S. J. Wright, Numerical optimization, Springer New York, NY, 1999
- [17] A. Nemirovski, A. Juditsky, G. Lan and A. Shapiro, Robust stochastic approximation approach to stochastic programming, *SIAM J. Optim.* **19** (2009) 1574 – 1609, <https://doi.org/10.1137/070704277>.
- [18] B. T. Polyak, Gradient methods for minimizing functionals, *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki*, **3** (1963) 864 – 878, [https://doi.org/10.1016/0041-5553\(63\)90382-3](https://doi.org/10.1016/0041-5553(63)90382-3).
- [19] S. Łojasiewicz, A topological property of real analytic subsets, *Coll. du CNRS, Les équations aux dérivées partielles*, **117** (1963) p. 2.
- [20] D. P. Kingma and J. Ba, Adam: a method for stochastic optimization, *arXiv preprint arXiv:1412.6980* (2014).
- [21] N. Serhat Aybat, A. R. Fallah, M. Gurbuzbalaban and A. Ozdaglar, A universally optimal multistage accelerated stochastic gradient method, *In. Proc. of the 33rd Int. Conf. on Neural Information Processing Systems*, (2019) 8525 – 8536.
- [22] K. He, X. Zhang, S. Ren and J. Sun, Deep residual learning for image recognition, *In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, (2016) 770 – 778.
- [23] C. C. Chang and C. J. Lin, Libsvm: a library for support vector machines, *ACM Trans. Intell. Syst. Technol.* **2** (2011) 1 – 27, <https://doi.org/10.1145/1961189.1961199>.

Mahsa Soheil Shamaee
Department of Computer Science,
Faculty of Mathematical Science,
University of Kashan, Kashan, Iran
e-mail: soheilshamaee@kashanu.ac.ir

Sajad Fathi Hafshejani
Department of Applied Mathematics,
Shiraz University of Technology,
Shiraz, I. R. Iran e-mail: s.fathi@sutech.ac.ir