

Adaptive Learning Rate and Momentum for Training Deep Neural Networks

Zhiyong Hao¹(✉), Yixuan Jiang¹, Huihua Yu¹, and Hsiao-Dong Chiang¹

Cornell University, Ithaca NY 14850, USA
{zh272,yj373,hy437,hc63}@cornell.edu

Abstract. Recent progress on deep learning relies heavily on the quality and efficiency of training algorithms. In this paper, we develop a fast training method motivated by the nonlinear Conjugate Gradient (CG) framework. We propose the Conjugate Gradient with Quadratic line-search (CGQ) method. On the one hand, a quadratic line-search determines the step size according to current loss landscape. On the other hand, the momentum factor is dynamically updated in computing the conjugate gradient parameter (like Polak-Ribiere). Theoretical results to ensure the convergence of our method in strong convex settings is developed. And experiments in image classification datasets show that our method yields faster convergence than other local solvers and has better generalization capability (test set accuracy). One major advantage of the paper method is that tedious hand tuning of hyperparameters like the learning rate and momentum is avoided.

Keywords: Optimization Algorithm · Line Search · Deep Learning.

1 Introduction

We consider the minimization problem commonly used in the machine learning setting with the following finite-sum structure:

$$\min_{\theta} f(\theta) = \frac{1}{n} \sum_{i=1}^n f_i(\theta) \quad (1)$$

where $f : D \subset \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable and D is the search space. A point $\theta^* \in D$ is called a local minimum if $f(\theta^*) \leq f(\theta)$ for all $\theta \in D$ with $\|\theta - \theta^*\| < \sigma$ for $\sigma > 0$. A general training procedure is in Algorithm 1, where key steps are determining the update direction p_t and step size α_t in each iteration.

Several methods propose various choices of search directions. The Gradient Descent (GD) method directly takes the negative gradient $-\nabla_{\theta} f(\theta)$. Newton’s method selects p_t as $-(\nabla_{\theta}^2 f(\theta))^{-1} \nabla_{\theta} f(\theta)$ (where $\nabla_{\theta}^2 f(\theta)$ is the Hessian matrix). On the other hand, momentum [23, 19] is used to stabilize p_t in stochastic optimization. Adam [12] formalizes p_t as a moving average of first moments divided by that of second moments. Other methods focus on the choice of α_t . Various classic line-search methods were proposed, like the Newton-Raphson,

Algorithm 1 Iterative Approach for Optimizing (1)

```

1: procedure OPTIMIZE( $f(\theta)$ ,  $maxiter$ )
2:   Initialize  $\theta_0$ 
3:   for  $t \leftarrow 1$  to  $maxiter$  do
4:     Choose a search direction  $p_t$ ;
5:     Determine a update step size  $\alpha_t$ ;
6:      $\theta_t \leftarrow \theta_{t-1} + \alpha_t p_t$ ;
7:   return  $\theta_t$ 

```

Armijo, and Wolfe line search. Recently, [31] applied the Armijo rule in SGD. In this paper, we propose the Conjugate Gradient with Quadratic line-search method that jointly optimizes both p_t and α_t for fast convergence.

The performance of SGD is well-known to be sensitive to step size α_t (or the learning rate called in machine learning). Algorithms, such as AdaGrad [7] and Adam [12], ease this limitation by tuning the learning rate dynamically using past gradient information. Line-search strategy is another option to adaptively tune step size α_t . Exact line-search requires solving a simpler sub-optimization problem, such as applying quadratic interpolation [17]. An inexact line-search method can also be applied, such as the Armijo rule [1], to guarantee a sufficient decrease at each step. A stochastic variant of Armijo rule was proposed in [31] to set the step size for SGD.

Conjugate Gradient method (CG) is also popular for solving nonlinear optimization problems. In CG, p_t is chosen as $(-\nabla f(\theta) + \beta_t p_{t-1})$, where β_t is the conjugate parameter computed according to various formulas. α_t is determined by a line search ($\arg \min_{\alpha} f(\theta_t + \alpha p_t)$). In fact, CG can be understood as a Gradient Descent with an adaptive step size and dynamically updated momentum. For the classic CG method, step size is determined by the Newton-Raphson method or the Secant method, both of which need to exactly compute or approximate the Hessian matrix. This makes it less appealing to apply classic CG in deep learning.

In this paper, we propose the Conjugate Gradient with the Quadratic line-search method (CGQ), which uses quadratic interpolation line-search to choose the step size and adaptively tunes momentum term according to the Conjugate Gradient formula. CGQ requires no Hessian information and can be viewed as a variant of SGD with adaptive learning rate and momentum. For quadratic line-search, we propose two variants, *i.e.*, the 2-point method and the least squares method, which are introduced in detail in Section 3.2. And we illustrate how we dynamically determine the momentum term during training in Section 3.3. Our analyses indicate that SGD with 2-point quadratic line-search converges under convex and smooth settings. sCGQ is further proposed to improve efficiency and generalization on large models. Through various experiments, CGQ exhibits fast convergence, and is free from manually tuning the learning rate and momentum term.

Our major contributions are summarized as follows:

- A quadratic line-search method is proposed to dynamically adjust the learning rate of SGD.
- An adaptive scheme for computing the effective momentum factor of SGD is designed.
- Theoretical results are developed to prove the convergence of SGD with quadratic line-search.
- Line search on subset of batches is proposed to improve efficiency and generalization capability.

2 Related Work

In modern machine learning, SGD and its variants [7, 12, 16, 24], are the most prevalent optimization methods used to train large-scale neural networks despite their simplicity. Theoretically, the performance of SGD with a constant learning rate has been analyzed that under some assumptions, such as, the Strong Growth Condition, has linear convergence if the learning rate is small enough [26]. However, there are still two non-negligible limitations: (1) the theoretical convergence rate is not as satisfying as the full gradient descent method; and (2) the performance is very sensitive to hyperparameters, especially the learning rate.

In SGD, random sampling introduces randomness and degrades the theoretical convergence of SGD. One direction of research focuses on reducing the variance of random sampling so that larger learning rates are allowed. Based on this idea, a family of SGD variants called Variance Reduction methods were proposed [25, 11, 27, 6]. These methods have been proved to possess impressive theoretical convergence on both convex and non-convex optimization problems. However, the applications of these methods in machine learning are limited to optimization of simple logistic regression models because the additional memory and computational cost of these methods are unacceptable when training large-scale deep neural networks.

On the other hand, significant efforts seek to ameliorate SGD by designing an effective method to adaptively tune the learning rate as training. [21] divides the optimization process of SGD into transient phase and stationary phase, while the learning rate will be reduced after entering the stationary phase. In [29], scheduling learning rate cyclically for each iteration empirically boosts the convergence of SGD. Meta-learning [2] techniques that optimize the learning rate according to the hyper-gradient also have an outstanding performance in experiments. However, none of these techniques gives solid theoretical proof of fast convergence. In pursuit of an optimal learning rate schedule with a theoretical convergence guarantee, Vaswani *et al.* hybrid SGD with classic optimization techniques like Armijo line-search [31] and Polyak step size [15, 3], and the developed methods not only have sound theory foundation but also empirically show much faster convergence than many widely-used optimization methods. Our work is inspired by Quadratic Interpolation (QI) technique used in line-search [30]. We developed two versions of quadratic line-search method to assist SGD to automatically determine an ideal learning rate schedule.

To accelerate convergence of SGD, heavy ball algorithm (SGD with fixed momentum) [23] is proposed. In practical neural network training, Nesterov accelerated momentum [19] is often used. As discussed in [4], heavy ball method is highly related to conjugate gradient (CG) method, and many efforts have been put into improving SGD by injecting the spirit of CG method. CG method presents promising performance in training an auto-encoder comparing with SGD and LBFGS [14]. A CG-based ADAM [13] was proposed and showed comparative performance to Adam, AdaGrad in solving text and image classification problems, but no generalization performance was provided. Jin *et al.* [10] proposed a stochastic CG with Wolfe line-search that was tested on low-dimensional datasets. All of these related works motivate us to build a bridge between SGD and CG method and create a dynamic acceleration scheme for SGD. Instead of tuning momentum term manually, we dynamically adjust it during training according to Polak-Ribiere formula [22] which shows the best based on our experiments.

Most recently, [18] independently proposed a line search method based on parabolic approximation (PAL). Our major differences are: 1). CGQ introduces feedback to check the estimation quality that improves reliability. 2). CGQ eliminates the momentum hyperparameter, as discussed in Section 3.3; 3). For stochastic settings, CGQ does line search on a subset of batches to improve the runtime and generalization; 4). Convergence analysis of CGQ is conducted with no quadratic objective assumption.

3 The CGQ Method

3.1 Overview

The Conjugate Gradient with Quadratic Line-Search (CGQ) method follows the iterative training framework in Algorithm 1, and eliminates hand-tuning two major hyperparameters in SGD: the learning rate and momentum factor.

In each training iteration, CGQ optimizes the learning rate α_t by estimating the one-dimensional loss landscape as a parabola. Two variants, the 2-point quadratic interpolation and the Least Squares estimation, are illustrated in Section 3.2. For direction p_t , methods from Conjugate Gradient are adopted to dynamically adjust the momentum factor in SGD instead of using a hand-tuned fixed momentum, which is described in section 3.3.

3.2 Dynamically Adjusting the Learning Rate

Two-point Quadratic Interpolation Given a twice differentiable loss function $f(\theta)$ ($\theta \in \mathbb{R}^m$ is the parameter vector), its second order Taylor expansion around θ_0 is expressed as:

$$\begin{aligned} f(\theta) &= f(\theta_0) + \nabla f(\theta_0)^T(\theta - \theta_0) \\ &\quad + \frac{1}{2}(\theta - \theta_0)^T \nabla^2 f(\theta_0)(\theta - \theta_0) \\ &\quad + O((\theta - \theta_0)^2). \end{aligned} \tag{2}$$

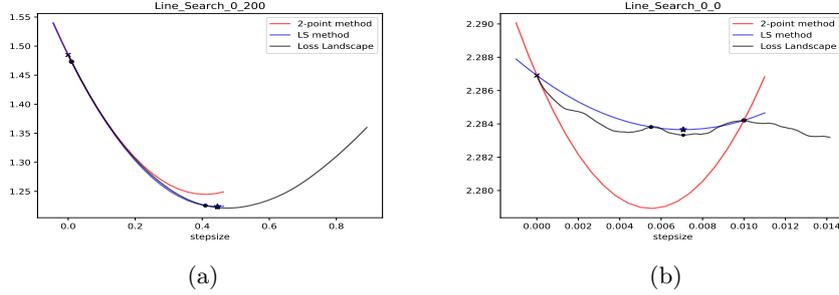


Fig. 1: Comparison of the two quadratic estimation algorithms (2-point method and least squares method). (a) When the landscape is smooth, both estimate well. (b) When the landscape is rough, least squares tend to capture more details. It is also observed in experiments that roughness only occurs at the beginning of a training, after which the landscapes are almost always like (a).

Along direction $p_t \in \mathbb{R}^m$, the 1D loss landscape $\phi(\alpha)$ w.r.t. the step size $\alpha \in \mathbb{R}$ can be described as:

$$\begin{aligned} \phi(\alpha) &= f(\theta - \alpha p_t) \\ &= f(\theta_0) - \alpha \cdot \nabla f(\theta_0)^T p_t \\ &\quad + \frac{1}{2} \alpha^2 \cdot p_t^T \nabla^2 f(\theta_0) p_t + O(\alpha^2). \end{aligned} \quad (3)$$

When the residual term $O(\alpha^2)$ is small enough, the line-search function can be expressed as:

$$\phi(\alpha) \approx q(\alpha) \triangleq A\alpha^2 + B\alpha + C \quad (4)$$

where $A = \frac{1}{2} p_t^T \nabla^2 f(\theta_0) p_t$, $B = -\nabla f(\theta_0)^T p_t$, and $C = f(\theta_0)$. When $A > 0$, $q(\alpha)$ reaches minimum at $\alpha^* = -B / (2A)$.

To get A, B, C at each training iteration without evaluating the Hessian matrix $\nabla^2 f(\theta_0)$, we propose to apply the 2-point quadratic interpolation method. Given two sample points $(0, q(0))$, $(\alpha_1, q(\alpha_1))$ and the slope $q'(0)$, parabola parameters can be calculated as:

$$\begin{aligned} A &= \frac{q(\alpha_1) - q(0) - q'(0) \cdot \alpha_1}{\alpha_1^2} \\ B &= q'(0) \\ C &= q(0) \end{aligned} \quad (5)$$

The stopping criteria and details are shown in Algorithm 2.

One of the major advantages of our method is that when moving closer to a local optimum as training proceeds, α^* will decay automatically, due to the nature of a convex parabola. This is also verified and visualized in our experiment section. Another issue is when the initial γ is too large compared to the estimated α^* . In such cases, $q(\alpha)$ is not a good surrogate to represent the local landscape

Algorithm 2 Quadratic Line-Search (2-point interpolation)

```

1: procedure QLS2( $f(\cdot), \theta_t, p_t, slope$ )
2:   Initialize:  $\gamma, K_{max}, \alpha_{max}, \sigma$ 
3:    $(x_0, y_0) \leftarrow (0, f(\theta_t))$ 
4:   for  $k \leftarrow 1$  to  $K_{max}$  do
5:      $(x_1, y_1) \leftarrow (\gamma, f(\theta_t + \gamma \cdot p_t))$ 
6:      $q_{A,B,C}(\cdot) \leftarrow \text{Interpolate}(x_0, y_0, x_1, y_1, slope)$ 
7:      $\alpha^* = \min(-B / 2A, \alpha_{max})$ 
8:     if  $A \leq 0$  then
9:       return  $\gamma$ 
10:    else if  $q_{ABC}(\alpha^*) > 0$  &  $f(\theta_t + \alpha^* p_t) < f(\theta_t)$  then
11:      return  $\alpha^*$ 
12:    else
13:       $\gamma \leftarrow \alpha^*$ 
14:    return  $\alpha^*$ 

```

of $\phi(\alpha)$ near the origin. This is handled by decaying the initial value of γ when it is consistently larger than the recent average of α^* . Through experiments, we observe that A is usually estimated positive in training neural networks, while the $A < 0$ case is also properly dealt with in Algorithm 2.

Convergence Analysis Our method is majorly motivated by empirical results. To give some insights on the theoretical side, we provide a convergence analysis of quadratic line-search under SGD update rule. To start with, we use the following assumptions on the loss function $f(\theta) = \sum_i f_i(\theta)$, which are commonly used for convergence analysis in deep learning [31, 15]. (a) f is *lower-bounded* by a finite value f^* ; (b) *interpolation*: This requires that $\nabla f(\theta^*) = 0$ implies $\nabla f_i(\theta^*) = 0$ for all components f_i ; (c) f is *L-smooth*, *i.e.*, the gradient ∇f is L -Lipschitz continuous; and (d) f is μ *strong-convex*. Proofs are given in the appendix.

The following lemma provides the bound of the step size returned by quadratic line-search at each iteration k .

Lemma 1. *Assuming (1) f_i 's are L_i -smooth, (2) μ_i strong-convexity of f_i , the step size α_k returned by quadratic line-search in algorithm 2 that is constrained within $(0, \alpha_{max}]$ satisfies:*

$$\alpha_k \in [\min\{\frac{1}{L_{ik}}, \alpha_{max}\}, \min\{\frac{1}{\mu_{ik}}, \alpha_{max}\}]. \quad (6)$$

Next, we provide the convergence of quadratic line-search under strong-convex and smooth assumptions.

Theorem 2. *Assuming (1) interpolation, (2) f_i 's are L_i -smooth, and (3) μ_i strong-convexity of f_i 's. When setting $\alpha_{max} \in (0, \min_i\{\frac{1}{L_i} + \frac{\mu_i}{L_i^2}\})$, the parameters of a neural network trained with quadratic line-search satisfies:*

$$\mathbb{E}[\|\theta_k - \theta^*\|^2] \leq M^k \|\theta_0 - \theta^*\|^2 \quad (7)$$

where $M = \mathbb{E}_i[\max\{g_i(\alpha_{max}), 1 - \frac{\mu_i}{L_i}\}] \in (0, 1)$ is the convergence rate, and $g_i(\alpha_k) = 1 - (\mu_i + L_i)\alpha_k + L_i^2\alpha_k^2$.

Least Squares Estimation Notice that theorem 2 assumes the loss landscape to be strongly convex and L -smooth. In reality, we find that such conditions are satisfied easier in the later than earlier phase of training. In other words, the loss landscape is more rough at the beginning, and the 2-point method sometimes counters difficulties getting a good estimation $q(\alpha)$. For instance, in rough surfaces like fig. 1b, the 2-point method may focus too much on the local information at the origin (both zero's order and first order moments.), which may lead to poor estimation. This motivates us to further develop a more robust version of the quadratic line-search using least squares estimation.

Instead of only using two samples (plus gradient), the Least Squares estimator takes more samples. And instead of finding a perfect match on the samples, Least Squares minimizes the mean squared error among all samples. In our case, Least Squares minimizes the following objective:

$$\min_{\omega} \|X \cdot \omega - Y\|_2^2 = \frac{1}{n} \sum_i (A\alpha_i^2 + B\alpha_i + C - \phi(\alpha_i))^2 \quad (8)$$

where

$$X = \begin{bmatrix} \alpha_1^2 & \alpha_1 & 1 \\ \alpha_2^2 & \alpha_2 & 1 \\ \dots & \dots & \dots \end{bmatrix} \in \mathbb{R}^{n \times 3} \quad (9)$$

is the sample matrix formulated by each step α_i , $\omega = [A, B, C]^T$ are the coefficients to be optimized, and $Y = [\phi(\alpha_1), \phi(\alpha_2), \dots]^T \in \mathbb{R}^n$ is the observation vector. The Least Squares estimator is: $\omega = (X^T X)^{-1} X^T Y$, where $(X^T X)^{-1} X^T$ is known as the Moore-Penrose inverse.

The algorithm is described in Algorithm 3. When the loss landscape $\phi(\alpha)$ is rough, as shown in Figure 1b, the slope of one point is not useful in order to capture the landscape trend. By using least squares estimation on multiple sample points without using slope, the quadratic interpolation becomes insensitive to gradient information near the origin, and seeks to accommodate more global information. Through experiments in Section 4, we show that using least squares is more robust.

3.3 Dynamically Adjusting Momentum

In this section, we propose an automatic computation of the momentum factor. In SGD or its variants, this is manually set and fixed in value. We build a bridge between the momentum in SGD and the conjugate gradient parameter β_t . (Recall that in the Nonlinear Conjugate Gradient (CG) method, the conjugate direction s_t at iteration t is updated as $s_t = -\nabla f + \beta_t \cdot s_{t-1}$, and parameters are updated by $x_{t+1} = x_t + \alpha_t \cdot s_t$.) β_t has a deterministic optimal value (s.t. s_t is conjugate to s_{t-1}) using the Hessian matrix in a quadratic objective. But for a general

Algorithm 3 Quadratic line-search (Least Squares)

```

1: procedure QLS2( $f(\cdot), \vec{\theta}_t, \vec{p}_t, K_{max}, \alpha_{max}, slope$ )
2:   Initialize:  $\gamma$ (sample distance);  $\alpha_0$ ;
3:    $X = \{(0, 1, 0), (0, 0, 1)\}$ ;  $Y = \{slope, f(\theta_t)\}$ 
4:   for  $k \leftarrow 1$  to  $K_{max}$  do
5:      $x_k \leftarrow (\gamma^2, \gamma, 1)$ 
6:      $X \leftarrow X \cup \{x_k\}$ 
7:      $Y \leftarrow Y \cup \{f(\theta_t + \gamma \cdot p_t)\}$ 
8:      $q_{A,B,C}(\cdot) \leftarrow \text{LeastSquare}(X, Y)$ 
9:     if  $A \leq 0$  then
10:      return  $\gamma$ 
11:     else if  $q_{ABC}(\alpha^*) > 0$  &  $f(\theta_t + \alpha^* p_t) < f(\theta_t)$  then
12:      return  $\alpha^*$ 
13:     else
14:        $\gamma \leftarrow \alpha^*$ 
15:   return  $\alpha^*$ 

```

nonlinear case, various heuristic formulas are proposed [9, 8, 22, 5, 28]. α_t can be determined using either the Newton–Raphson method or a line-search such as the Armijo rule [1] and Wolfe conditions [32, 33]. The Nonlinear CG method generalizes various gradient-based optimization methods. For instance, when $\beta_t = 0$ and α_t is constant, CG degenerates as the Gradient Descent, and when β_t is a non-zero constant, it instantiates the Gradient Descent with momentum, which is also discussed in [4]. Moreover, the CG framework can also be used in the stochastic settings, where gradients $\nabla_{\theta} c(\theta)$ are replaced by partial gradients $\nabla_{\theta} c_{batch}(\theta)$ in each iteration.

With our proposed Quadratic line-search, the local solver is able to adaptively change the learning rate for performance. We next show that using the CG framework, the momentum term in SGD (β_t in CG), which is usually constant in a SGD solver, can also be adaptively adjusted.

For a quadratic objective of the following form:

$$\min f(\theta) = \frac{1}{2} \theta^T Q \theta + b^T \theta \quad (10)$$

where $\theta \in \mathbb{R}^n$ is the model parameters, and $Q \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$ are problem-dependent matrices. The optimal β_t can be chosen as

$$\beta_t = \frac{\nabla f(\theta_t)^T \cdot Q \cdot p_{t-1}}{p_{t-1}^T Q p_{t-1}} \quad (11)$$

where p_{t-1} is the update direction at time $t - 1$. Then the next update direction $p_t = -\nabla f(\theta_t) + \beta_t \cdot p_{t-1}$ is guaranteed to be Q -conjugate to p_{t-1} , which accelerates training by eliminating influences on previous progress. For general nonlinear objective functions with no explicit Q and the Hessian is difficult to compute, different heuristics were proposed in [9, 8, 22, 5, 28]. Since the classic Gradient

Descent with momentum method is a stationary version of the Conjugate Gradient [4], a dynamic choice of β_t , instead of fixed, can potentially speed up the training process. It is also notable that to compute β_t using the aforementioned formula, no Hessian is required. For example, the Polak-Ribiere [22] formula:

$$\beta_t^{PR} = \frac{\nabla f(\theta_t)^T (\nabla f(\theta_t) - \nabla f(\theta_{t-1}))}{\nabla f(\theta_{t-1})^T \nabla f(\theta_{t-1})} \quad (12)$$

only requires the recent two gradients.

A bounded version of the Polak-Ribiere momentum is used, *i.e.*, to compute the momentum as $\beta'_t = \min\{\max\{0, \beta_t^{PR}\}, \beta_{max}\}$. Through experiments, it is shown that by applying such dynamic momentum β'_t , the performance is comparable to or even better than the best manual momentum. The major advantage is that another crucial hyperparameter is eliminated from hand tuning.

3.4 Optimizations for Large Datasets

Conceptually, CGQ performs a line search at every iteration. This unsurprisingly introduces more computation. To alleviate this issue on large systems, it is reasonable to use partially observed curvature (from batches) to estimate that of the whole dataset. This motivates us to implement the sCGQ that performs line search stochastically: For each training batch, the optimizer will perform a line search (that optimizes learning rate) with a certain probability p . For batches without line searches, the moving average of past line search results will be applied. fig. 2 show that this makes the run time of sCGQ with fixed momentum comparative to naive SGD, and sCGQ with dynamic momentum comparative to other line search methods. Experiments in section 4.1 and the runtime comparison in fig. 2 show that sCGQ benefits from both fast execution and good estimation of 1-D loss landscape (in terms of learning rate).

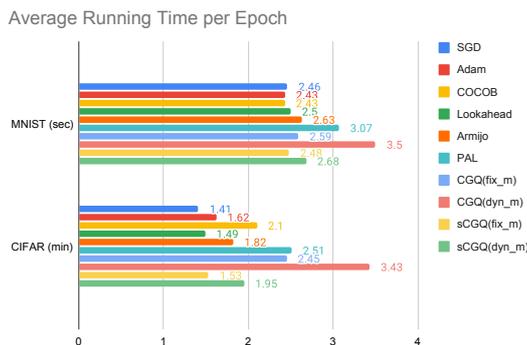


Fig. 2: Running time comparison. sCGQ: stochastic line search with probability 0.1. `fix_m`: fixed momentum; `dyn_m`: Polak-Ribiere momentum. 2-point interpolation and Least Squares takes similar time.

Method	Train Loss / Test Accuracy			
	MNIST	SVHN		
	MLP	VGG-16	ResNet-110	DenseNet-100
SGD	.076 / 97.24	.080 / 93.16	.217 / 92.68	.158 / 94.34
Adam [12]	.014 / 98.02	.026 / 93.14	.134 / 95.34	.093 / 95.35
Lookahead [34]	.131 / 95.99	.039 / 93.95	.157 / 93.14	.112 / 94.12
COCOBS [20]	.267 / 92.72	.036 / 91.81	.069 / 90.34	.035 / 91.30
Armijo [31]	.023 / 97.85	.008 / 92.42	–	–
PAL [18]	.021 / 97.60	–	.016 / 94.73	.017 / 95.10
CGQ(2pt)	.007/98.24	.011 / 94.70	.009 / 95.31	.004/96.11
CGQ(LS)	.007/98.24	.006 / 94.79	.012 / 95.81	.008 / 95.99
sCGQ(2pt)	.008 / 98.24	.020 / 93.89	.043 / 95.84	.029 / 94.23
sCGQ(LS)	.008 / 98.12	.016 / 94.93	.039 / 95.63	.032 / 95.43

Table 1: Performance comparison of different optimization algorithms on the MNIST and SVHN datasets. Training budget: 20 epochs (MNIST) / 40 epochs (SVHN). “–” means not converging.

4 Experiments

To empirically evaluate the performance of the proposed method, we design a thorough experiment consisting of two parts described in the following sections. We focus on comparing our method with different configurations in Section 4.1. In Section 4.2, we benchmark our method with six popular optimization methods: SGD+momentum, Adam [12], COCOBS [20], Lookahead optimizer [34], SGD+Armijo Line-Search [31], and PAL [18]. We manually tune their hyperparameters and the results summarized in the following sections are their best performances in the experiments to our best knowledge. Specifically for SGD and ADAM, the learning rates are fixed (SGD: 0.02 for VGG, 0.01 on MNIST, and 0.1 elsewhere; Adam: 0.001). There is no learning rate for COCOBS. For our own method, we tested four variants: CGQ(2pt)/CGQ(LS): complete line search with quadratic interpolation or least squares estimation; sCGQ(2pt)/sCGQ(LS): stochastic line search with quadratic interpolation/least squares estimation. We showcase the fast convergence of our method on training multi-class classifiers for widely used benchmark image datasets including MNIST, SVHN, CIFAR-10, and CIFAR-100. For fairness, model parameters are initialized the same among all solvers for each test case, and we run multiple times and show the average performance for each case. All experiments are conducted on a NVIDIA GTX-1080Ti GPU, and the program is developed under PyTorch 1.7.0 framework. Code is available in <https://github.com/zh272/CGQ-solver>.

4.1 Ablation Test

Dynamical Momentum Applying SGD with momentum to reach faster convergence has become common in deep neural network training. Heavy ball [23] and Nesterov momentum [19] are two classic static acceleration scheme that

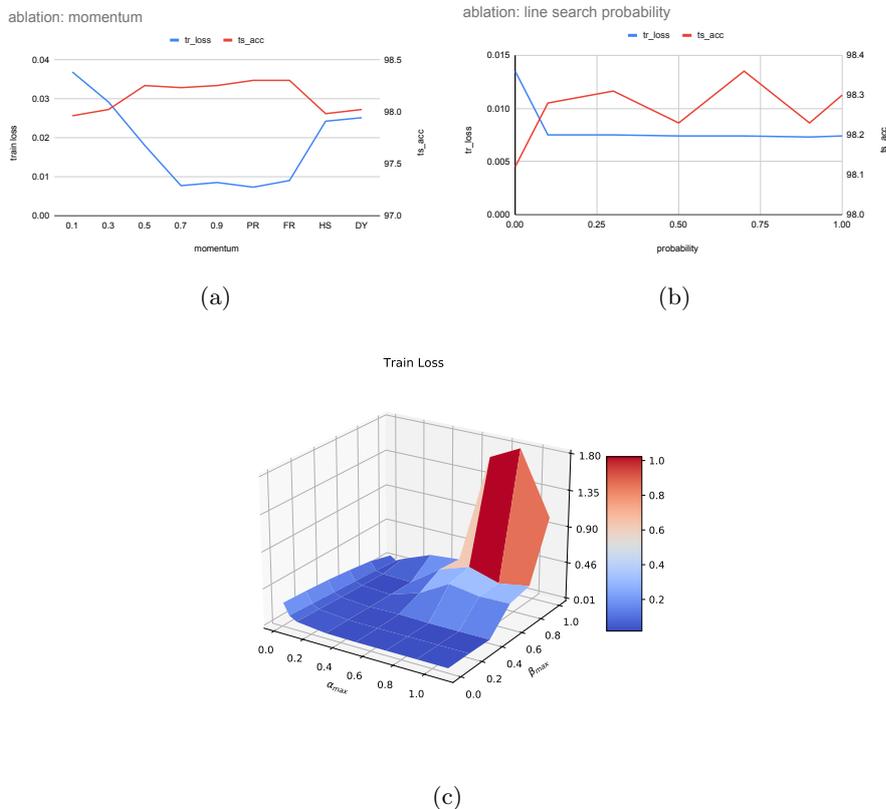


Fig. 3: Training loss (One-layer MLP trained on MNIST dataset) w.r.t.: (a) Different momentum settings. (b) Line search probability. “0”: naive SGD with no line search; “1”: line search on every mini-batch. (c) α_{max} and β_{max} .

consider the momentum term as a fixed hyperparameter. In this section, we compare the performances of SGD with quadratic line-search plus (1) Heavy ball momentum[23] with coefficients from 0.1 to 0.9, (2) Fletcher-Reeves momentum [8], (3) Polak-Ribiere momentum [22], (4) Hestenes-Stiefel momentum [9], and (5) Dai-Yuan momentum [5].

The benchmark experiment is run on the MNIST dataset. The comparison results are summarized in Figure 3a. On the basis of the results of the ablation test, it is promising to use Polak-Ribiere momentum as a dynamic acceleration scheme for our method so that the momentum term is no longer a hyperparameter that needs tuning manually.

Line Search Probability To evaluate the impact of the randomness on line search, we test our method with probability thresholds p varying from 0(no line search) to 1(complete line search). Results in fig. 3b show that line search

Method	Train Loss / Test Accuracy					
	VGG-16		ResNet-164		DenseNet-100	
	CIFAR-10	CIFAR-100	CIFAR-10	CIFAR-100	CIFAR-10	CIFAR-100
SGD	.116 / 89.64	.685 / 64.37	.271 / 86.06	.860 / 63.55	.255 / 86.71	.901 / 62.06
Adam	.044 / 88.06	.408 / 55.21	.144 / 90.06	.419 / 68.07	.127 / 90.42	.364 / 68.23
Lookahead	.069 / 89.21	.369 / 60.97	.155 / 88.33	.526 / 59.21	.153 / 88.50	.594 / 61.62
COCOB	.087 / 84.78	.772 / 44.83	.330 / 78.96	1.450 / 48.00	.205 / 80.63	1.103 / 51.56
Armijo	.298 / 81.18	.025 / 62.50	.003 / 93.64	.004 / 73.26	.009 / 92.72	.086 / 65.55
PAL	–	–	.017 / 92.04	.452 / 64.01	.023 / 90.99	.072 / 65.90
CGQ(2pt)	.008 / 88.21	.025 / 65.46	.007 / 93.04	.014 / 75.26	.009 / 91.49	.020 / 75.01
CGQ(LS)	.010 / 88.64	.019/66.49	.010 / 93.52	.019 / 73.43	.010 / 92.23	.029 / 74.35
sCGQ(2pt)	.035 / 89.16	.038 / 63.12	.017 / 93.75	.016 / 75.94	.020 / 93.71	.035 / 76.15
sCGQ(LS)	.041 / 88.88	.044 / 62.93	.016 / 93.56	.016 / 75.60	.028 / 93.13	.032 / 76.04

Table 2: Performance comparison for the CIFAR-10 and CIFAR-100 datasets. Training budget: 200 epochs. “–” means the method does not converge.

improves the final model performance even with small probability p , and it does not yield further improvements as p increases.

α_{max} and β_{max} α_{max} (upper bound of learning rate) and β_{max} (upper bound of momentum), are two important hyperparameters in CGQ. In this section we perform an ablation test to show the feasible regions. We perform a grid search on the 2D space by α_{max} (0.01 to 1.1) and β_{max} (0 to 1). It is observed from Figure 3c that the model performs well when α_{max} and β_{max} are not both large. And the model will diverge when β_{max} is larger than 1 (thus not shown on figure). Empirically good choices are: $\beta_{max} = 0.8$; $\alpha_{max} = 0.3$ (ResNet) / 0.1 (DenseNet) / 0.05 (VGG).

4.2 Multi-class classification on image datasets

Recent evidence has shown that deep convolutional neural networks lead to impressive breakthroughs in multi-class image classification. We mainly select the three popular families of architectures of deep convolutional neural networks: ResNet, DenseNet and VGG to benchmark our method with the other powerful optimizers. The comparison results will be presented in the following two sections, according to different datasets used in experiments.

MNIST and SVHN The MNIST dataset contains 60,000 hand-written digit images for training, and 10,000 images for testing. Images in MNIST are grayscale images with size 28×28 . SVHN dataset consists of 73,257 images in the training set, 26032 images in the testing set and another 531,131 additional images provided for training. All the images in SVHN are colored digit images with a size of 32×32 . For MNIST, we construct a fully connected network with one hidden layer of 1000 neurons as a toy benchmark example. For SVHN, we

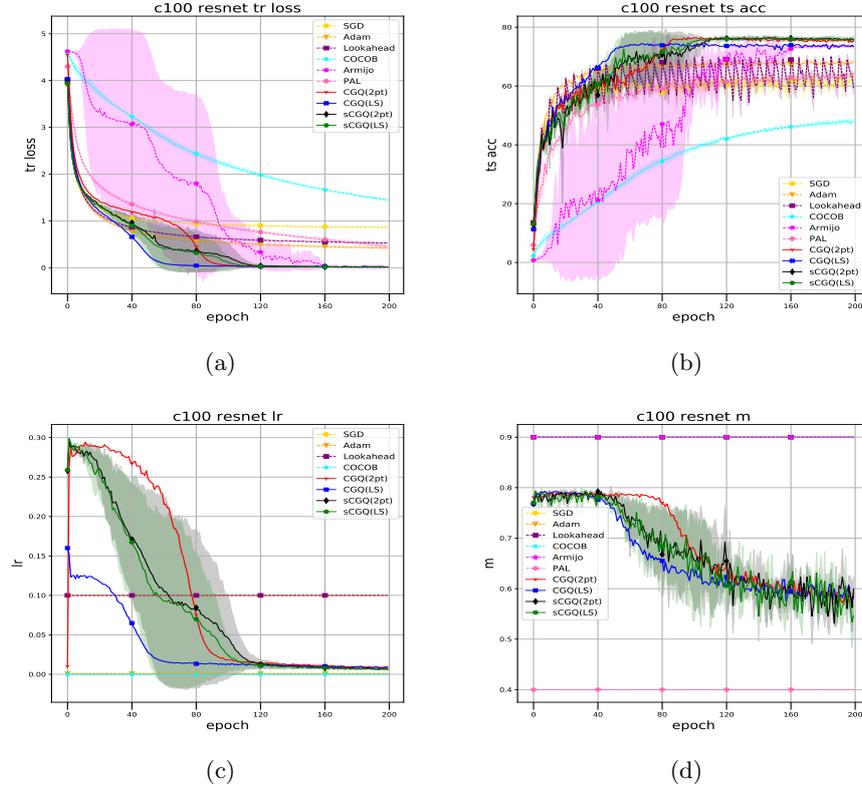


Fig. 4: Dynamics of training a ResNet-164 with various optimization methods on the CIFAR100. (a) Training Loss. (b) Test Accuracy. (c) Learning Rate. (d) Momentum

experiment with three more expressive architectures: VGG-16, ResNet-110, and DenseNet-100.

The main results are summarized in Table 1. We observed that with the same training budget, CGQ presents better performances in terms of both training loss and test accuracy. The stochastic version of We omit SGD with Armijo-line-search in the SVHN experiment since it appeared to be unstable and easy to diverge.

CIFAR-10 and CIFAR-100 The CIFAR-10 and CIFAR-100 datasets both consist of 50,000 images for training and 10,000 images for testing, while images are in 10 and 100 different classes, respectively. We run the experiments using VGG-16, ResNet-164, and DenseNet-100. Each of model-algorithm pair is trained for 200 epochs with same initialization points.

In Table 2, we record the training loss and test accuracy of each model and optimizer pair in the CIFAR-10 and CIFAR-100 experiments. In most of the cases, CGQ and sCGQ not only reach the best performance among all line

search methods, but also converges to a local minimum faster. In the the more challenging CIFAR-100 case, the margin is more significant. fig. 4 shows the dynamics for training ResNet on the CIFAR100 dataset. Each line plots the mean over multiple runs, while corresponding standard deviations are shown as shaded area around the mean. (a) and (b) tells that CGQ variants performs stably better in training loss and testing accuracy compared to other compared methods. Moreover, (c) and (d) demonstrates the automatic decaying effect on learning rate and momentum of the CGQ method.

It is observed from the tables that CGQ performs better than sCGQ on the training set, but sCGQ has better test time performance. In this sense, the sCGQ is preferable to CGQ for larger datasets, not only because of faster execution time, but also because of better generalization. It is also observed that 2-point interpolation method performs better than its Lease Squares variant on larger datasets (CIFAR). This could due to an underestimation of the smoothness for deep neural networks. This is also validated through a visualization of the loss landscape along the line search direction in fig. 1.

5 Conclusion

In this paper, we propose CGQ method to replace hand tuning two crucial hyperparameters in SGD, *i.e.*, learning rate and momentum. The learning rate is determined by our quadratic line-search method, and momentum is adaptively computed using the bounded Polak-Ribiere formula.

Experiments on modern convolutional neural networks show that CGQ converges faster due to its adaptiveness to the loss landscape. Theoretical results are also provided towards the convergence of quadratic line search under convex and smooth assumptions. To improve efficiency on larger models, we further propose the sCGQ method that performs line search on fewer iterations. This modification improves not only the run time, but also the generalization capability in terms of test accuracy. In most cases, the CGQ method outperforms other local methods.

A Proof of Lemma 1

Suppose $h_{ik}(\alpha) = f_{ik}(\theta_k - \alpha \nabla f_{ik}(\theta_k))$. Besides $\alpha = 0$, sample the other point $\alpha = \alpha_1$ and get the quadratic interpolant of f_{ik} as follows,

$$q(\alpha) = \frac{\|\nabla f_{ik}(\theta_k)\|^2 \alpha_1 + h_{ik}(\alpha_1) - h_{ik}(0)}{\alpha_1^2} \alpha^2 - \|\nabla f_{ik}(\theta_k)\|^2 \alpha + h_{ik}(0).$$

Thus, the step size α_k returned by a quadratic line-search is

$$\alpha_k = \frac{\|\nabla f_{ik}(\theta_k)\|^2 \alpha_1^2}{2(\|\nabla f_{ik}(\theta_k)\|^2 \alpha_1 + h_{ik}(\alpha_1) - h_{ik}(0))}.$$

According to the assumption on the smoothness of f_{ik} and the stochastic gradient descent update rule, the following inequality holds:

$$h_{ik}(\alpha_1) \leq h_{ik}(0) + \nabla f_{ik}(\theta_k)'(-\nabla f_{ik}(\theta_k)\alpha_1) + \frac{L_{ik}}{2} \|\nabla f_{ik}(\theta_k)\alpha_1\|^2.$$

Thus,

$$h_{ik}(\alpha_1) - h_{ik}(0) \leq (-\alpha_1 + \frac{L_{ik}}{2}\alpha_1^2) \|\nabla f_{ik}(\theta_k)\|^2.$$

As a result,

$$\alpha_k \geq \frac{\|\nabla f_{ik}(\theta_k)\|^2 \alpha_1^2}{2(\|\nabla f_{ik}(\theta_k)\|^2 \alpha_1 + (-\alpha_1 + \frac{L_{ik}}{2}\alpha_1^2) \|\nabla f_{ik}(\theta_k)\|^2)} = \frac{1}{L_{ik}}.$$

Similarly, based on the assumption on the strong convexity of f_{ik} and the stochastic gradient descent update rule,

$$h_{ik}(\alpha_1) - h_{ik}(0) \geq (-\alpha_1 + \frac{\mu_{ik}}{2}\alpha_1^2) \|\nabla f_{ik}(\theta_k)\|^2.$$

Therefore,

$$\alpha_k \leq \frac{1}{\mu_{ik}}.$$

By definition, $\alpha_k \in (0, \alpha_{max}]$, and thus, the step size returned by quadratic line-search satisfies the following relation:

$$\alpha_k \in [\min\{\frac{1}{L_{ik}}, \alpha_{max}\}, \min\{\frac{1}{\mu_{ik}}, \alpha_{max}\}].$$

□

B Proof of Theorem 1

Suppose the loss function $f(\theta)$ is minimized at θ^* . According to the stochastic gradient descent update rule,

$$\|\theta_{k+1} - \theta^*\|^2 = \|\theta_k - \alpha_k \nabla f_{ik}(\theta_k) - \theta^*\|^2 = \|\theta_k - \theta^*\|^2 - 2\alpha_k \langle \nabla f_{ik}(\theta_k), \theta_k - \theta^* \rangle + \alpha_k^2 \|\nabla f_{ik}(\theta_k)\|^2.$$

Based on the assumption on the strong convexity of $f_{ik}(\cdot)$,

$$-\langle \nabla f_{ik}(\theta_k), \theta_k - \theta^* \rangle \leq f_{ik}(\theta^*) - f_{ik}(\theta_k) - \frac{\mu_{ik}}{2} \|\theta_k - \theta^*\|^2.$$

Thus,

$$\|\theta_{k+1} - \theta^*\|^2 \leq (1 - \mu_{ik}\alpha_k) \|\theta_k - \theta^*\|^2 + 2\alpha_k [f_{ik}(\theta^*) - f_{ik}(\theta_k)] + \alpha_k^2 \|\nabla f_{ik}(\theta_k)\|^2.$$

Based on the assumption on the smoothness of $f_{ik}(\cdot)$,

$$f_{ik}(\theta^*) - f_{ik}(\theta_k) \leq -\frac{1}{2L_{ik}} \|\nabla f_{ik}(\theta_k)\|^2$$

$$\|\nabla f_{ik}(\theta_k) - \nabla f_{ik}(\theta^*)\|^2 \leq L_{ik}^2 \|\theta_k - \theta^*\|^2.$$

In the interpolation setting, $\nabla f_{ik}(\theta^*) = 0$; therefore,

$$\|\nabla f_{ik}(\theta_k)\|^2 \leq L_{ik}^2 \|\theta_k - \theta^*\|^2.$$

As a result,

$$\|\theta_{k+1} - \theta^*\|^2 \leq [1 - (\mu_{ik} + L_{ik})\alpha_k + L_{ik}^2\alpha_k^2] \|\theta_k - \theta^*\|^2$$

Let $g_{ik}(\alpha_k) = 1 - (\mu_{ik} + L_{ik})\alpha_k + L_{ik}^2\alpha_k^2$. According to the assumptions, it is trivial to get $L_{ik} > \mu_{ik} > 0$, and thus,

$$g_{ik}(\alpha_k) > 1 - 2L_{ik}\alpha_k + L_{ik}^2\alpha_k^2 \geq 0.$$

For each f_i , define a constant $C_i = \frac{1}{L_i} + \frac{\mu_i}{L_i^2}$. If $\alpha_{max} \in (0, C_{min})$, $g_{ik}(\alpha_{max}) < 1$. Based on Lemma 1,

$$0 \leq g_{ik}(\alpha_k) \leq \max\{g_{ik}(\alpha_{max}), g_{ik}(\frac{1}{L_{ik}})\} = \max\{g_{ik}(\alpha_{max}), 1 - \frac{\mu_{ik}}{L_{ik}}\} < 1.$$

Therefore,

$$\|\theta_{k+1} - \theta^*\|^2 \leq \max\{g_{ik}(\alpha_{max}), 1 - \frac{\mu_{ik}}{L_{ik}}\} \|\theta_k - \theta^*\|^2.$$

Take the expectation with respect to i_k :

$$\mathbb{E}[\|\theta_{k+1} - \theta^*\|^2] \leq \mathbb{E}_{i_k}[\max\{g_{ik}(\alpha_{max}), 1 - \frac{\mu_{ik}}{L_{ik}}\}] \|\theta_k - \theta^*\|^2$$

where $\mathbb{E}_{i_k}[\max\{g_{ik}(\alpha_{max}), 1 - \frac{\mu_{ik}}{L_{ik}}\}] \in (0, 1)$. □

References

1. Armijo, L.: Minimization of functions having lipschitz continuous first partial derivatives. *Pacific J. Math.* **16**(1), 1–3 (1966)
2. Baydin, A.G., Cornish, R., Rubio, D.M., Schmidt, M., Wood, F.: Online learning rate adaptation with hypergradient descent. *ICLR* (2018)
3. Berrada, L., Zisserman, A., Kumar, M.P.: Training Neural Networks for and by Interpolation. *International Conference on Machine Learning* (2020)
4. Bhaya, A., Kaszkurewicz, E.: Steepest descent with momentum for quadratic functions is a version of the conjugate gradient method. *Neural Networks* **17**, 65–71 (2004). [https://doi.org/10.1016/S0893-6080\(03\)00170-9](https://doi.org/10.1016/S0893-6080(03)00170-9)
5. Dai, Y.H., Yuan, Y.X.: A nonlinear conjugate gradient method with a strong global convergence property. *SIAM J. Optimization* **10**(1), 177–182 (1999)
6. Defazio, A., Bach, F., Lacoste-Julien, S.: Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*. vol. 27, pp. 1646–1654. Curran Associates, Inc. (2014)

7. Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* **12**(61), 2121–2159 (2011)
8. Fletcher, R., Reeves, C.M.: Function minimization by conjugate gradients. *The Computer Journal* **7**(2), 149–154 (01 1964). <https://doi.org/10.1093/comjnl/7.2.149>
9. Hestenes, M.R., Stiefel, E.: Methods of conjugate gradients for solving linear systems. *Journal of research of the National Bureau of Standards* **49**, 409–436 (1952)
10. Jin, X.B., Zhang, X.Y., Huang, K., Geng, G.G.: Stochastic Conjugate Gradient Algorithm With Variance Reduction. *IEEE Transactions on Neural Networks and Learning Systems* **30**(5), 1360–1369 (2019). <https://doi.org/10.1109/TNNLS.2018.2868835>
11. Johnson, R., Zhang, T.: Accelerating stochastic gradient descent using predictive variance reduction. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*. vol. 26, pp. 315–323. Curran Associates, Inc. (2013)
12. Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)* (2015)
13. Kobayashi, Y., Iiduka, H.: Conjugate-gradient-based Adam for stochastic optimization and its application to deep learning (2020), <http://arxiv.org/abs/2003.00231>
14. Le, Q.V., Ngiam, J., Coates, A., Lahiri, A., Prochnow, B., Ng, A.Y.: On optimization methods for deep learning. In: *ICML* (2011)
15. Loizou, N., Vaswani, S., Laradji, I., Lacoste-Julien, S.: Stochastic polyak step-size for sgd: An adaptive learning rate for fast convergence. In: *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*. pp. 1–33 (2021)
16. Loshchilov, I., Hutter, F.: Fixing weight decay regularization in adam. *arXiv:1711.05101* (2017)
17. M. J. D. Powell: An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Computer Journal* **7**(2), 155–162 (1964)
18. Mutschler, M., Zell, A.: Parabolic approximation line search for dnns. In: *NeurIPS* (2020)
19. Nesterov, Y.: A method of solving a convex programming problem with convergence rate $o(1/k^2)$. *Soviet Mathematics Doklady* **27**, 372–376 (1983)
20. Orabona, F., Pal, D.: Coin betting and parameter-free online learning. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29, pp. 577–585. Curran Associates, Inc. (2016)
21. Pesme, S., Dieuleveut, A., Flammarion, N.: On convergence-diagnostic based step sizes for stochastic gradient descent. In: *Proceedings of the International Conference on Machine Learning 1 pre-proceedings (ICML 2020)* (2020)
22. Polak, E., Ribiere, G.: Note sur la convergence de méthodes de directions conjuguées. *ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique* **3**(R1), 35–43 (1969)
23. Polyak, B.T.: Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics* **4**(5), 1–17 (1964)
24. Reddi, S.J., Kale, S., Kumar, S.: On the convergence of adam and beyond. In: *The 35th International Conference on Machine Learning (ICML)* (2018)
25. Schmidt, M., Le Roux, N., Bach, F.: Minimizing finite sums with the stochastic average gradient. *Mathematical Programming* **162**(1), 83–112 (2017)

26. Schmidt, M., Roux, N.L.: Fast convergence of stochastic gradient descent under a strong growth condition (2013)
27. Shalev-Shwartz, S., Zhang, T.: Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research* **14**(1), 567599 (2013)
28. Shewchuk, J.R.: An introduction to the conjugate gradient method without the agonizing pain (August 1994), <http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf>
29. Smith, L.N.: Cyclical learning rates for training neural networks. In: *IEEE Winter Conference on Applications of Computer Vision (WACV)* (2017)
30. Vandebogert, K.: Method of quadratic interpolation (September 2017), https://people.math.sc.edu/kellerlv/Quadratic_Interpolation.pdf
31. Vaswani, S., Mishkin, A., Laradji, I., Schmidt, M., Gidel, G., Lacoste-Julien, S.: Painless stochastic gradient: Interpolation, line-search, and convergence rates. In: *Advances in Neural Information Processing Systems*. pp. 3727–3740 (2019)
32. Wolfe, P.: Convergence conditions for ascent methods. *SIAM Review* **11**(2), 226–000 (1969)
33. Wolfe, P.: Convergence conditions for ascent methods. ii: Some corrections. *SIAM Review* **13**(2), 185–000 (1969)
34. Zhang, M., Lucas, J., Ba, J., Hinton, G.E.: Lookahead optimizer: k steps forward, 1 step back. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 32, pp. 9597–9608 (2019)