

# Projected Forward Gradient-Guided Frank-Wolfe Algorithm via Variance Reduction

Mohammadreza Rostami, Solmaz S. Kia, *Senior member, IEEE*

**Abstract**—This paper aims to enhance the use of the Frank-Wolfe (FW) algorithm for training deep neural networks. Similar to any gradient-based optimization algorithm, FW suffers from high computational and memory costs when computing gradients for DNNs. This paper introduces the application of the recently proposed projected forward gradient (Projected-FG) method to the FW framework, offering reduced computational cost similar to backpropagation and low memory utilization akin to forward propagation. Our results show that trivial application of the Projected-FG introduces non-vanishing convergence error due to the stochastic noise that the Projected-FG method introduces in the process. This noise results in a non-vanishing variance in the Projected-FG estimated gradient. To address this, we propose a variance reduction approach by aggregating historical Projected-FG directions. We demonstrate rigorously that this approach ensures convergence to the optimal solution for convex functions and to a stationary point for non-convex functions. Simulations demonstrate our results.

**Index Terms**—Frank-Wolfe Algorithm, Projected Forward Gradient, Deep Neural Network, Backpropagation

## I. INTRODUCTION

Deep neural networks (DNNs) have become essential tools for solving function approximation problems over the past decade [1]. Training DNNs involves optimizing the parameter vector  $\theta \in \mathbb{R}^d$  (weights and biases) to minimize the loss between predictions and training data. To improve generalization, enhance interpretability, and reduce computational costs, it is desirable to promote sparsity in the DNN parameters. Traditionally, sparsity is achieved through post-training pruning. However, recent methods demonstrate that incorporating convex constraints directly into the training optimization problem enables one-shot pruning-aware training [2], [3]. These constraints impose sparsity and boundedness on  $\theta$  [4, Table 1]. For constrained DNN training, the Frank-Wolfe (FW) algorithm [5] is often used, as it avoids costly projections onto constraint sets, making it suitable for high-dimensional spaces.

As with any first-order optimization method, the FW algorithm requires computing the gradient of the training cost function. However, computing the gradient in DNNs is not a trivial task; it demands special attention due to high-dimensional matrix multiplications and the efficient management of computational and memory resources needed for computing each layer's activations. This letter conducts a critical analysis of applying a newly proposed approximate gradient computation method with low computational and memory costs, known as

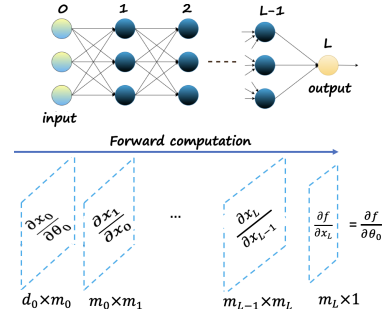


Fig. 1: Using forward propagation to compute  $\frac{\partial f}{\partial \theta_0}$  in a network with  $L - 1$  hidden layers, each with  $n$  nodes.

the *Projected Forward Gradient (Projected-FG)* [6]–[9], within the FW optimization framework.

For a DNN with cost  $f(\theta)$  with  $L + 1$  layers and one output shown in Fig. 1, the gradient  $\mathbf{g}(\theta) = \nabla f(\theta) = (\frac{\partial f}{\partial \theta_0}, \dots, \frac{\partial f}{\partial \theta_L})$ , where  $\theta_i \in \mathbb{R}^{d_i}$  is the parameter vector of the  $i$ th layer,  $i \in \mathcal{L} = \{0, 1, 2, \dots, L\}$ , is given by

$$\mathbf{g}_i(\theta) = \frac{\partial \mathbf{x}_i}{\partial \theta_i} \times \frac{\partial \mathbf{x}_{i+1}}{\partial \mathbf{x}_i} \times \dots \times \frac{\partial \mathbf{x}_L}{\partial \mathbf{x}_{L-1}} \times \frac{\partial f}{\partial \mathbf{x}_L}. \quad (1)$$

Here,  $\mathbf{g}_i(\theta) = \frac{\partial f}{\partial \theta_i}$  and  $\mathbf{x}_i = h_i(\mathbf{x}_{i-1}, \theta_i) \in \mathbb{R}^{m_i}$  is the activation (output) vector of  $i$ th layer and  $\times$  denotes matrix multiplication. The gradient (1) in DNNs is typically computed using either *forward propagation* (Fig. 1) or *backpropagation*. In forward propagation, starting from  $\frac{\partial f}{\partial \theta_i}$ , the gradient calculation involves multiplying  $L - 1$  Jacobian matrices of dimensions  $m_i \times m_{i-1}$  from left to right up to the final term  $\frac{\partial f}{\partial \mathbf{x}_L}$ , a vector of size  $m_{L-1}$ . Although computationally expensive due to high-dimensional matrix multiplications, this method effectively utilizes memory since  $\mathbf{x}_i$  of each layer  $i \in \mathcal{L}$  can be computed sequentially and overwritten at the same memory location. Backpropagation starts at the output layer's  $\frac{\partial f}{\partial \mathbf{x}_L}$  (a vector) and propagates backward, computing the gradient by backward multiplication between a vector and a high-dimensional matrix. While computationally efficient, backpropagation requires substantial memory to store all layer activations  $\mathbf{x}_i$ ,  $i \in \mathcal{L}$  [10].

Backpropagation is favored for its relatively low computational cost, and efforts have been made to reduce its memory usage. One approach involves checkpointing, where only a subset of activations is stored as checkpoints instead of saving every intermediate activation [11], [12]. During the backward pass, these checkpoints are used to recompute the necessary intermediate activations that were not stored, reducing the mem-

The authors are with the Department of Mechanical and Aerospace Engineering, University of California Irvine, Irvine, CA 92697, {mrostam2, solmaz}@uci.edu.

ory footprint by performing recomputation on-the-fly. Another strategy is to eliminate the need to store activations altogether in reversible networks. In reversible networks, the input to each layer can be reconstructed from the output using the inverse function. Studies such as [13], [14] exploit this property by recomputing activations backward during backpropagation. While these methods achieve memory savings, they come with an increased computational cost.

The *Projected-FG* [6]–[9] is a directional derivative computed in a forward fashion but with low computational cost:

**Definition 1** (Projected Forward Gradient). *Give a function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  the Projected-FG is defined as defined as  $\hat{\mathbf{g}}(\boldsymbol{\theta}) = \langle \mathbf{g}(\boldsymbol{\theta}), \mathbf{u} \rangle \mathbf{u}$ , where  $\mathbf{u} \in \mathbb{R}^d$  is a random vector with each entry  $u_i$  independently and identically distributed (i.i.d) with zero mean and unit variance.*  $\square$

Using left-to-right forward multiplication starting with vector  $\mathbf{u}_i \in \mathbb{R}^{d_i \times 1}$ ,  $i \in \mathcal{L}$ , results in subsequent multiplications between a row vector and a matrix in

$$\hat{\mathbf{g}}_i(\boldsymbol{\theta}) = \left( \underbrace{\mathbf{u}_i^\top}_{1 \times d_i} \times \frac{\partial \mathbf{x}_i}{\partial \boldsymbol{\theta}_i} \times \frac{\partial \mathbf{x}_{i+1}}{\partial \mathbf{x}_i} \times \dots \times \frac{\partial \mathbf{x}_L}{\partial \mathbf{x}_{L-1}} \times \frac{\partial f}{\partial \mathbf{x}_L} \right) \times \mathbf{u}_i,$$

providing the computational efficiency of backpropagation with the memory-efficient strategies of forward propagation.

However, the use of *Projected-FG* in optimization algorithms is not straightforward, as  $\hat{\mathbf{g}}(\boldsymbol{\theta})$  is a noisy estimate of  $\mathbf{g}(\boldsymbol{\theta})$ . In this letter, we demonstrate that the straightforward application of the *Projected-FG* method within the FW algorithm results in convergence to only a neighborhood of the optimal solution. We attribute this non-exact convergence to the non-vanishing variance in estimating the gradient with  $\hat{\mathbf{g}}(\boldsymbol{\theta})$ . To address the stochastic errors in gradient approximation, we propose an algorithm that reduces variance by averaging historical *Projected-FG* directions. Inspired by momentum-based methods that accelerate gradient descent [15], our approach uniquely incorporates the *Projected-FG* method to systematically reduce noise in gradient estimates. Unlike traditional momentum methods, which primarily use past gradients for faster convergence [16]–[18], our algorithm aggregates historical *Projected-FG* directions, enhancing the convergence of the FW algorithm. Through rigorous proofs, we show that our method provides a biased variance-reduced *Projected-FG* direction that converges to the true gradient with diminishing variance, ensuring convergence to the optimal solution for convex functions and to a stationary point for non-convex functions. A numerical example demonstrates the results.

## II. OBJECTIVE STATEMENT

Consider the constrained DNN training optimization problem

$$\min_{\boldsymbol{\theta} \in \mathcal{C}} f(\boldsymbol{\theta}), \quad (2)$$

where  $\mathcal{C} \subset \mathbb{R}^d$  denotes a compact convex constraint set.

<sup>1</sup>‘Projected’ here refers to the operation  $\langle \mathbf{g}(\boldsymbol{\theta}), \mathbf{u} \rangle \mathbf{u}$ , which projects the gradient  $\mathbf{g}(\boldsymbol{\theta})$  on vector  $\mathbf{u}$ .

---

## Algorithm 1 Projected-FG Frank-Wolfe (FGFW)

---

```

1: Require:  $\alpha_k$ 
2: Input:  $\boldsymbol{\theta}_0 \in \mathcal{C}$ 
3: Output:  $\boldsymbol{\theta}_{K+1}$ 
4: for  $k = 1, 2, \dots, K$  do
5:   Sample  $\mathbf{u}_k \sim \mathcal{N}(\mathbf{0}, I_d)$  and compute  $\hat{\mathbf{g}}(\boldsymbol{\theta}_k) = \langle \mathbf{g}(\boldsymbol{\theta}_k), \mathbf{u}_k \rangle \mathbf{u}_k$ 
6:   Compute  $\mathbf{s}_k \in \operatorname{argmin}_{\mathbf{s} \in \mathcal{C}} \langle \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbf{s} \rangle$ 
7:    $\boldsymbol{\theta}_{k+1} = (1 - \alpha_k) \boldsymbol{\theta}_k + \alpha_k \mathbf{s}_k$ 
8: end for

```

---

**Assumption 1** (Properties of the constraint set  $\mathcal{C}$ ). *The convex set  $\mathcal{C}$  is bounded with diameter  $D \in \mathbb{R}_{>0}$ , i.e.,*

$$\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\| \leq D, \quad \forall \boldsymbol{\theta}, \hat{\boldsymbol{\theta}} \in \mathcal{C}.$$

**Assumption 2** (Assumption on  $M$ -smoothness of the cost function). *The cost function is  $M$ -smooth (or have an  $M$ -Lipschitz continuous gradient) in  $\mathcal{C}$ , i.e, there exists a constant  $M \in \mathbb{R}_{>0}$  such that*

$$\|\mathbf{g}(\boldsymbol{\theta}) - \mathbf{g}(\hat{\boldsymbol{\theta}})\| \leq M \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|, \quad \forall \boldsymbol{\theta}, \hat{\boldsymbol{\theta}} \in \mathcal{C}.$$

Recall that for a  $M$ -smooth  $f(\boldsymbol{\theta})$ , we have  $f(\boldsymbol{\theta}) \leq f(\hat{\boldsymbol{\theta}}) + \langle \mathbf{g}(\hat{\boldsymbol{\theta}}), \boldsymbol{\theta} - \hat{\boldsymbol{\theta}} \rangle + \frac{M}{2} \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|^2$ , for any  $\boldsymbol{\theta}, \hat{\boldsymbol{\theta}} \in \mathcal{C}$  [19].

To solve optimization problem (2), Algorithm 1 presents the FW algorithm utilizing the *Projected-FG* gradient estimation, where the computation of gradient  $\mathbf{g}$  in line 5 is replaced by  $\hat{\mathbf{g}}$ . Our objective in this paper is to investigate the impact of using the *Projected-FG* estimate of the gradient within the FW algorithm and devise strategies to counteract any performance degradation caused by the use of the noisy  $\hat{\mathbf{g}}$ . To aid in our convergence analysis and algorithm design, we begin by reviewing some auxiliary results about the *Projected-FG*.

**Lemma II.1** (Unbiasedness of the *PF-Gradient* [8]). *The *PF-Gradient* in Definition 1 is an unbiased estimate of  $\mathbf{g}(\boldsymbol{\theta})$ , i.e.  $\mathbb{E}[\hat{\mathbf{g}}(\boldsymbol{\theta})] = \mathbf{g}(\boldsymbol{\theta})$ .*  $\square$

**Lemma II.2** (Upper bound on the *PF-Gradient* [8]). *The *PF-Gradient* in Definition 1 satisfies*

$$\mathbb{E}[\|\hat{\mathbf{g}}(\boldsymbol{\theta})\|^2] \leq (n+4)\|\mathbf{g}(\boldsymbol{\theta})\|^2. \quad (3)$$

**Lemma II.3** (Variance of the *PF-Gradient*). *Under Assumptions 1 and 2, the *PF-Gradient* in Definition 1 satisfies*

$$\mathbb{E}[\|\hat{\mathbf{g}}(\boldsymbol{\theta}) - \mathbf{g}(\boldsymbol{\theta})\|^2] \leq G = (n+3)M^2D^2, \quad \forall \boldsymbol{\theta} \in \mathcal{C}. \quad (4)$$

In what follows,  $f(\boldsymbol{\theta}^*)$  is the optimal value of (2). Moreover, the proofs of the formal results are presented in the appendix.

## III. FRANK WOLFE ALGORITHM WITH PROJECTED FORWARD GRADIENT

Starting with  $\boldsymbol{\theta}_0 \in \mathcal{C}$ , in each iteration, FW consults an oracle to identify a descent direction in  $\mathcal{C}$  from  $\mathbf{s}_k = \operatorname{argmin}_{\mathbf{s} \in \mathcal{C}} \langle \mathbf{g}(\boldsymbol{\theta}_k), \mathbf{s} \rangle$ . The process proceeds with the update  $\boldsymbol{\theta}_{k+1} = (1 - \alpha_k) \boldsymbol{\theta}_k + \alpha_k \mathbf{s}_k$ ,  $\alpha_k \in [0, 1]$ , which is a convex combination of two vectors in  $\mathcal{C}$  resulting in  $\boldsymbol{\theta}_{k+1}$  remaining feasible within  $\mathcal{C}$ . Consequently, this method avoids projecting the updated point back into  $\mathcal{C}$ , thereby attributing a projection-free nature to the FW algorithm. The convergence of FW for convex functions is achieved by a rate of  $O(1/k)$  and  $O(1/\sqrt{k})$

---

**Algorithm 2** Averaged Projected FG Frank Wolfe (AFGFW)

---

1: **Require:**  $\alpha_k, \gamma_k$   
2: **Input:**  $\theta_0 \in \mathcal{C}, \mathbf{v}_0 = \mathbf{0}$   
3: **Output:**  $\theta_{K+1}$   
4: **for**  $k = 1, 2, \dots, K$  **do**  
5:   Sample  $\mathbf{u}_k \sim \mathcal{N}(\mathbf{0}, I_d)$  and compute  $\hat{\mathbf{g}}(\theta_k) = \langle \mathbf{g}(\theta_k), \mathbf{u}_k \rangle \mathbf{u}_k$   
6:   Compute  $\mathbf{v}_k = (1 - \gamma_k)\mathbf{v}_{k-1} + \gamma_k \hat{\mathbf{g}}(\theta_k)$   
7:   Compute  $\mathbf{s}_k \in \operatorname{argmin}_{\mathbf{s} \in \mathcal{C}} \langle \mathbf{v}_k, \mathbf{s} \rangle$   
8:    $\theta_{k+1} = (1 - \alpha_k)\theta_k + \alpha_k \mathbf{s}_k$   
9: **end for**

---

for non-convex functions [20], [21]. Next, we analyze the convergence behavior of Algorithm 1, showing that incorporating *Projected-FG* naively in FW framework introduces a convergence error. To simplify the study, we focus on smooth and convex costs.

**Lemma III.1** (Trajectories of Algorithm 1 remain in  $\mathcal{C}$ ). *Let  $f$  in (2) be differentiable. For any convex set  $\mathcal{C}$ , Algorithm 1, initialized at  $\theta_0 \in \mathcal{C}$  results in  $\theta_k \in \mathcal{C}$  for all  $k \in \mathbb{Z}_{\geq 0}$ .  $\square$*

Lemma III.1’s proof follows the same reasoning as the original FW algorithm, using mathematical induction [21]. Since  $\theta_0 \in \mathcal{C}$  and, by construction,  $\mathbf{s}_k \in \mathcal{C}$  for all  $k \in \mathbb{Z}_{\geq 0}$  (see line 6 of Algorithm 1), therefore,  $\theta_{k+1}$ , being a convex combination of two vectors in  $\mathcal{C}$ , will also reside in the convex set  $\mathcal{C}$  for all  $k \in \mathbb{Z}_{\geq 0}$ . Lemma III.1 demonstrates that Algorithm 1 preserves the FW algorithm’s “interior point property.” However, as we will show next, exact convergence is not maintained.

**Theorem III.1** (Convergence bound of Algorithm 1 for convex functions). *Let  $f$  in (2) be convex and Assumptions 1 and 2 hold. Let  $\alpha_k$  be a decaying function of  $k$  with an order of  $O(1/k)$  and  $\psi_1 = MD^2\sqrt{n+3}$ . Then, Algorithm 1 results in*

$$\lim_{k \rightarrow \infty} \mathbb{E}[f(\theta_k) - f(\theta^*)] = \frac{\psi_1}{1 - \beta}, \quad (5)$$

for  $0 < \beta < 1$ , with a convergence rate of order  $O(1/k)$ .  $\square$

Theorem III.1 shows Algorithm 1 converges with a sublinear rate to an error neighborhood of size  $\psi_1$  of the optimal function value, indicating *Projected-FG* cannot be trivially implemented within FW framework. To remove this convergence error, we propose Algorithm 2, which we present next.

#### IV. VARIANCE-REDUCED FRANK WOLFE ALGORITHM WITH PROJECTED FORWARD GRADIENT

By examining the proof of Theorem III.1, the error component  $\psi_1$  is due to the second term in equation (10), i.e.,  $\sqrt{\mathbb{E}[\|\mathbf{g}(\theta_k) - \hat{\mathbf{g}}(\theta_k)\|^2]}$ . This connects the convergence error in Algorithm 1 to the non-vanishing variance of the *Projected-FG* gradient estimator. To address this error, we introduced Algorithm 2, which improved the convergence characteristics by estimating the gradient using aggregated historical *F-Gradients*. This was achieved through  $\mathbf{v}_k = (1 - \gamma_k)\mathbf{v}_{k-1} + \gamma_k \hat{\mathbf{g}}(\theta_k)$  in line 6 of Algorithm 2. Rather than interpreting this line as a method for extrapolating the iterates [22], [23], we view  $\mathbf{v}_k$  as a biased variance-reduced gradient estimator. This perspective is formally confirmed in Lemma IV.1 below.

**Lemma IV.1** (The variance of the averaged projected forward gradient estimator of Algorithm 2 converges to zero). *Consider optimization problem (2). Let Assumption 2 holds, and  $\alpha_k$  to be a decaying function of  $k$  with an order of  $O(1/k)$ , and  $\eta_k = \frac{\alpha_k}{\gamma_k}$ . Then, if  $\lim_{k \rightarrow \infty} \eta_k = 0$ , Algorithm 2 results in*

$$\lim_{k \rightarrow \infty} \mathbb{E}[\|\mathbf{v}_k - \mathbf{g}(\theta_k)\|^2] = 0, \quad (6)$$

with the rate of  $\min\{O(\eta_k^2), O(\gamma_k)\}$ .  $\square$

The next results show that the asymptotic variance reduction of the gradient estimator  $\mathbf{v}_k$  leads to Algorithm 2 achieving exact convergence. Here, we rely on the fact that since  $\mathbf{v}_k \in \mathcal{C}$  a similar argument, as in Lemma III.1, can be established for Algorithm 2 (The details are omitted due to space limitations.)

**Theorem IV.1** (Convergence analysis of Algorithm 2 for convex functions). *Let  $f$  in (2) be convex and Assumptions 1 and 2 hold. Let  $\alpha_k$  be a decaying function of  $k$  with an order of  $O(1/k)$ . Then, Algorithm 2 results in  $\lim_{k \rightarrow \infty} \mathbb{E}[f(\theta_k) - f(\theta^*)] \rightarrow 0$  with a sublinear convergence rate of order  $\min\{O(\frac{\alpha_k}{\gamma_k}), O(\gamma_k)\}$ .  $\square$*

According to Theorem IV.1, by selecting  $\gamma_k = O(1/\sqrt{k})$ , Algorithm 2 achieves its optimal convergence rate of  $O(1/\sqrt{k})$ . This improves over Algorithm 1 by reducing noise through the incorporation of past *Projected-FG* directions.

In DNNs, often the cost function is non-convex. The following theorem presents the convergence guarantee of Algorithm 2 for non-convex functions. For non-convex constrained optimization problems, the convergence criterion typically used is the so-called *FW gap*, defined as

$$\mathcal{K}(\theta) = \max_{\mathbf{s} \in \mathcal{C}} \langle \mathbf{s} - \theta, -\mathbf{g}(\theta) \rangle,$$

where  $\mathcal{K}(\theta) = 0$  if and only if  $\theta$  is a stationary point [20].

**Theorem IV.2** (Convergence analysis of Algorithm 2 for non-convex functions). *Consider Algorithm 2 under Assumptions 1 and 2. Let  $\alpha_k$  and  $\gamma_k$  be decaying functions of  $k$  with orders of  $O(1/k)$  and  $O(1/\sqrt{k})$ , respectively. Then,  $\lim_{K \rightarrow \infty} \mathbb{E}[\mathcal{K}(\theta_K)] = 0$  with an order of  $O\left(\frac{1}{\log K}\right)$ , satisfying the upper bound*

$$\mathbb{E}[\mathcal{K}(\theta_K)] \leq \frac{f(\theta_0) - f(\theta^*)}{\log K} + E, \quad k \in \{1, 2, \dots, K\},$$

where  $f(\theta^*)$  is the optimal value of problem (2) and  $E$  is a decaying term of the order  $O\left(\frac{1}{\log K}\right)$ .  $\square$

Theorem IV.2 shows Algorithm 2 achieves a convergence rate of  $O\left(\frac{1}{\log K}\right)$  to a stationary point for non-convex functions.

#### V. NUMERICAL EVALUATION

We evaluate the training loss and memory reduction of Algorithms 1 and 2 when training a DNN with constrained parameters for a multinomial classification task. The data set we use, the MNIST dataset [24], consists of  $28 \times 28$  pixel grayscale images of handwritten digits ranging from 0 to 9. Here, the purpose of training is to classify each digit correctly, aiming to create a model that can precisely predict the digit represented in any input image. We utilize a DNN with a depth

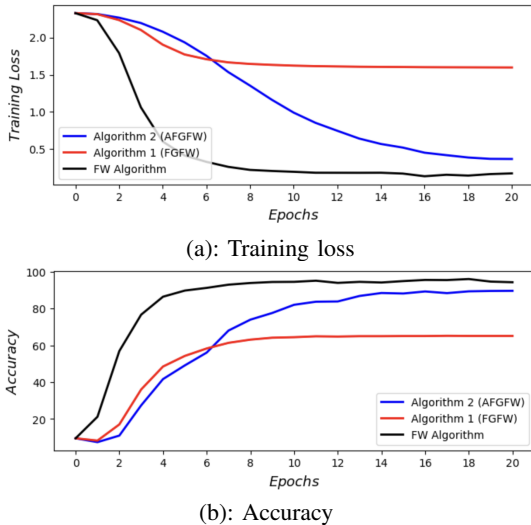


Fig. 2: Training loss and accuracy with respect to the number of epochs for Algorithm 1 and 2 over 20 epochs.

of 5 layers, each containing 10 neurons. A ReLU activation function is applied across all layers. The model is trained for 50 epochs. To encourage sparsity in the model, we use an  $\ell_1$ -norm constraint on the parameters in all layers, whose bound is set at  $10^{-4}$ . To implement Algorithms 1 and 2, we choose  $\gamma_k = \frac{1}{\sqrt{k}}$  and  $\alpha_k = \frac{1}{k}$ . For Algorithm 1, we set  $\alpha_k = \frac{2}{k+2}$  which is a standard step size in the FW Algorithm.

Figure 2(a) shows the training loss for the FW algorithm and Algorithms 1 and 2. As predicted by Theorem III.1, the naive adaptation of the *Projected-FG* method within the FW framework in Algorithm 1 results in a steady-state error due to the non-vanishing noise in the gradient estimate  $\hat{g}(\theta_k)$ .

On the other hand, Algorithm 2 shows convergence behavior as proved in Theorem IV.2, albeit slower than Algorithm FW algorithm. But, FW algorithm achieves faster convergence by using the full gradient computed via backpropagation, which requires higher memory usage. The maximum memory usage during training with *Projected-FG* in both Algorithms 1 and 2 is approximately  $990 \text{ MiB} \pm 10\%$  (where 1 mebibyte (MiB) = 1024 kibibytes (KiB)) per epoch. In contrast, using backpropagation to calculate the gradient results in a memory usage of about  $2150 \text{ MiB} \pm 10\%$  per epoch. This demonstrates the advantage of *Projected-FG* methods in training DNNs by significantly reducing memory consumption compared to the backpropagation method used in FW algorithm.

Figure 2(b) shows the test dataset accuracy for the FW algorithm, Algorithm 1, and Algorithm 2, with accuracies of 94.29%, 65.13%, and 89.63%, respectively, after 20 epochs. Notably, enforcing the  $\ell_1$ -norm constraint sparsified the DNN parameters, setting 2,387 of 8,290 elements of the optimized  $\theta$  to zero, which improved generalization to unseen datasets. Without this constraint, accuracy drops to 86.31%, highlighting the benefits of constraints in DNN training.

## VI. CONCLUSION

This paper presented an enhanced approach to the FW algorithm for training DNNs by integrating the *Projected-FG*

method. Our study demonstrated that while the straightforward application of *Projected-FG* within the FW framework reduced computational and memory costs, it introduced non-vanishing convergence errors due to stochastic noise. To address this challenge, we proposed a novel algorithm that incorporated variance reduction by averaging historical *Projected-FG* directions. This approach not only mitigated the noise but also ensured convergence to the optimal solution for convex functions and to a stationary point for non-convex functions. Through theoretical analysis and numerical simulations, we validated the effectiveness and efficiency of our proposed method. The results highlighted the potential of using *Projected-FG* to achieve memory-efficient training of DNNs without compromising convergence properties. Future work will explore the proposed method in distributed settings aligned with network topologies. We will also investigate using other distributions for sampling  $u_k$ , including the symmetric and binary Rademacher distribution, known for its variance reduction attribute. This can lead to more stable performance in algorithms as demonstrated for example in [25], [26].

## REFERENCES

- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.
- [2] M. Lu, X. Luo, T. Chen, W. Chen, D. Liu, and Z. Wang, "Learning pruning-friendly networks via frank-wolfe: One-shot, any-sparsity, and no retraining," in *International Conference on Learning Representations*, 2022.
- [3] M. Zimmer, C. Spiegel, and S. Pokutta, "Compression-aware training of neural networks using Frank-Wolfe," *arXiv*, 2022.
- [4] M. Jaggi, "Revisiting Frank-Wolfe: Projection-free sparse convex optimization," in *International conference on machine learning*, pp. 427–435, 2013.
- [5] S. Pokutta, C. Spiegel, and M. Zimmer, "Deep neural network training with frank-wolfe. arxiv e-prints, art," *arXiv preprint arXiv:2010.07243*, 2020.
- [6] D. Silver, A. Goyal, I. Danihelka, M. Hessel, and H. van Hasselt, "Learning by directional gradient descent," in *International Conference on Learning Representations*, 2021.
- [7] D. Silver, A. Goyal, I. Danihelka, M. Hessel, and H. van Hasselt, "Learning by directional gradient descent," in *International Conference on Learning Representations*, 2022.
- [8] A. G. Baydin, B. A. Pearlmutter, D. Syme, F. Wood, and P. Torr, "Gradients without backpropagation," *arXiv*, 2022.
- [9] D. Krylov, A. Karamzade, and R. Fox, "Moonwalk: Inverse-forward differentiation," *arXiv*, 2024.
- [10] G. Novikov, D. Bershatsky, J. Gusak, A. Shonenkov, D. Dimitrov, and I. Oseledets, "Few-bit backward: Quantized gradients of activation functions for memory footprint reduction," 2022.
- [11] J. Martens and I. Sutskever, "Training deep and recurrent networks with hessian-free optimization," in *Neural Networks: Tricks of the Trade: Second Edition*, pp. 479–535, Springer, 2012.
- [12] A. Gruslys, R. Munos, I. Danihelka, M. Lanctot, and A. Graves, "Memory-efficient backpropagation through time," *Advances in neural information processing systems*, vol. 29, 2016.
- [13] A. N. Gomez, M. Ren, R. Urtasun, and R. B. Grosse, "The reversible residual network: Backpropagation without storing activations," *Advances in neural information processing systems*, vol. 30, 2017.
- [14] K. Mangalam, H. Fan, Y. Li, C.-Y. Wu, B. Xiong, C. Feichtenhofer, and J. Malik, "Reversible vision transformers," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10830–10840, 2022.
- [15] N. Qian, "On the momentum term in gradient descent learning algorithms," *Neural networks*, vol. 12, no. 1, pp. 145–151, 1999.
- [16] C. Liu and M. Belkin, "Accelerating SGD with momentum for overparameterized learning," *arXiv*, 2018.
- [17] W. Liu, L. Chen, Y. Chen, and W. Zhang, "Accelerating federated learning via momentum gradient descent," *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 8, pp. 1754–1766, 2020.

- [18] G. Nakerst, J. Brennan, and M. Haque, "Gradient descent with momentum—to accelerate or to super-accelerate?," *arXiv*, 2020.
- [19] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*, vol. 87. Springer Science & Business Media, 2013.
- [20] S. Lacoste-Julien, "Convergence rate of frank-wolfe for non-convex objectives," *ArXiv*, vol. abs/1607.00345, 2016.
- [21] M. Frank, P. Wolfe, *et al.*, "An algorithm for quadratic programming," *Naval research logistics quarterly*, vol. 3, no. 1–2, pp. 95–110, 1956.
- [22] Y. Liu, Y. Gao, and W. Yin, "An improved analysis of stochastic gradient descent with momentum," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [23] R. Kidambi, P. Netrapalli, P. Jain, and S. Kakade, "On the insufficiency of existing momentum schemes for stochastic optimization," in *2018 Information Theory and Applications Workshop*, pp. 1–9, 2018.
- [24] L. Deng, "The mnist database of handwritten digit images for machine learning research," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.
- [25] J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE transactions on automatic control*, vol. 37, no. 3, pp. 332–341, 1992.
- [26] M. Rostami and S. Kia, "Fedscalar: A communication efficient federated learning," *arXiv preprint arXiv:2410.02260*, 2024.
- [27] Y. Nesterov, *Lectures on convex optimization*, vol. 137. Springer, 2018.
- [28] H. K. Khalil, *Nonlinear Systems*. Englewood Cliffs, NJ: Prentice Hall, 3 ed., 2002.

## APPENDIX

*Proof:* [Proof of Lemma II.3] We note that

$$\begin{aligned} \mathbb{E}[\|\hat{\mathbf{g}}(\boldsymbol{\theta}) - \mathbf{g}(\boldsymbol{\theta})\|^2] &\stackrel{(a)}{=} \mathbb{E}[\|\hat{\mathbf{g}}(\boldsymbol{\theta})\|^2] - \|\mathbf{g}(\boldsymbol{\theta})\|^2 \stackrel{(b)}{\leq} \\ (n+3)\|\mathbf{g}(\boldsymbol{\theta})\|^2 &\stackrel{(c)}{\leq} (n+3)M^2\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|^2 \stackrel{(d)}{\leq} (n+3)M^2D^2. \end{aligned}$$

Right hand side (RHS) of equality (a) follows from Lemma II.1, RHS of inequality (b) follows from Lemma II.2, RHS of inequality (c) follows from Assumption 2 and RHS of inequality (d) follows from Assumption 1.

*Proof:* [Proof of Theorem III.1] Given Lemma III.1 that ensures  $\boldsymbol{\theta}_k \in \mathcal{C}$  and using Assumption 2, we can write  $\mathbb{E}[f(\boldsymbol{\theta}_{k+1})|\boldsymbol{\theta}_k] \leq f(\boldsymbol{\theta}_k) + \langle \mathbf{g}(\boldsymbol{\theta}_k), \mathbb{E}[\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k|\boldsymbol{\theta}_k] \rangle + \frac{M}{2} \mathbb{E}[\|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k\|^2|\boldsymbol{\theta}_k]$ . Then,  $\boldsymbol{\theta}_{k+1} = (1 - \alpha_k)\boldsymbol{\theta}_k + \alpha_k \mathbf{s}_k$  results in

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta}_{k+1})|\boldsymbol{\theta}_k] &\leq f(\boldsymbol{\theta}_k) + \alpha_k \langle \mathbf{g}(\boldsymbol{\theta}_k), \mathbb{E}[\mathbf{s}_k - \boldsymbol{\theta}_k|\boldsymbol{\theta}_k] \rangle \\ &+ \frac{M\alpha_k^2}{2} \mathbb{E}[\|\mathbf{s}_k - \boldsymbol{\theta}_k\|^2|\boldsymbol{\theta}_k] \\ &\stackrel{(a)}{=} f(\boldsymbol{\theta}_k) + \alpha_k \langle \mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbb{E}[\mathbf{s}_k - \boldsymbol{\theta}_k|\boldsymbol{\theta}_k] \rangle + \\ &\alpha_k \langle \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbb{E}[\mathbf{s}_k - \boldsymbol{\theta}_k|\boldsymbol{\theta}_k] \rangle + \frac{M\alpha_k^2}{2} \mathbb{E}[\|\mathbf{s}_k - \boldsymbol{\theta}_k\|^2|\boldsymbol{\theta}_k], \quad (7) \end{aligned}$$

where RHS of equality (a) in (7) follows from adding and subtracting  $\alpha_k \langle \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbb{E}[\mathbf{s}_k - \boldsymbol{\theta}_k|\boldsymbol{\theta}_k] \rangle$ . Applying total expectation on both sides of (7) we have

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta}_{k+1})] &\leq \mathbb{E}[f(\boldsymbol{\theta}_k)] + \alpha_k \mathbb{E}[\langle \mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbf{s}_k - \boldsymbol{\theta}_k \rangle] \\ &+ \alpha_k \mathbb{E}[\langle \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbf{s}_k - \boldsymbol{\theta}_k \rangle] + \frac{M\alpha_k^2}{2} \mathbb{E}[\|\mathbf{s}_k - \boldsymbol{\theta}_k\|^2] \\ &\stackrel{(a)}{\leq} \mathbb{E}[f(\boldsymbol{\theta}_k)] + \alpha_k \mathbb{E}[\langle \mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbf{s}_k - \boldsymbol{\theta}_k \rangle] \\ &+ \alpha_k \mathbb{E}[\langle \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \boldsymbol{\theta}^* - \boldsymbol{\theta}_k \rangle] + \frac{M\alpha_k^2}{2} \mathbb{E}[\|\mathbf{s}_k - \boldsymbol{\theta}_k\|^2], \quad (8) \end{aligned}$$

where in RHS of inequality (a) in (8),  $\langle \mathbf{s}_k, \hat{\mathbf{g}}(\boldsymbol{\theta}_k) \rangle$  is replaced by its upper bound  $\langle \boldsymbol{\theta}^*, \hat{\mathbf{g}}(\boldsymbol{\theta}_k) \rangle$  because  $\langle \boldsymbol{\theta}^*, \hat{\mathbf{g}}(\boldsymbol{\theta}_k) \rangle \geq$

$\min_{\mathbf{s} \in \mathcal{C}} \langle \mathbf{s}, \hat{\mathbf{g}}(\boldsymbol{\theta}_k) \rangle = \langle \mathbf{s}_k, \hat{\mathbf{g}}(\boldsymbol{\theta}_k) \rangle$ . Adding and subtracting  $\alpha_k \mathbb{E}[\langle \mathbf{g}(\boldsymbol{\theta}_k), \boldsymbol{\theta}^* - \boldsymbol{\theta}_k \rangle]$  in (8) results in

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta}_{k+1})] &\leq \mathbb{E}[f(\boldsymbol{\theta}_k)] + \alpha_k \mathbb{E}[\langle \mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbf{s}_k - \boldsymbol{\theta}^* \rangle] \\ &+ \alpha_k \mathbb{E}[\langle \mathbf{g}(\boldsymbol{\theta}_k), \boldsymbol{\theta}^* - \boldsymbol{\theta}_k \rangle] + \frac{M\alpha_k^2}{2} \mathbb{E}[\|\mathbf{s}_k - \boldsymbol{\theta}_k\|^2], \\ &\stackrel{(a)}{\leq} \mathbb{E}[f(\boldsymbol{\theta}_k)] + \alpha_k \mathbb{E}[\langle \mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k), \mathbf{s}_k - \boldsymbol{\theta}^* \rangle] \\ &+ \alpha_k \mathbb{E}[f(\boldsymbol{\theta}^*) - f(\boldsymbol{\theta}_k)] + \frac{M\alpha_k^2}{2} \mathbb{E}[\|\mathbf{s}_k - \boldsymbol{\theta}_k\|^2], \\ &\stackrel{(b)}{\leq} \mathbb{E}[f(\boldsymbol{\theta}_k)] + \alpha_k D \mathbb{E}[\|\mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k)\|] \\ &+ \alpha_k \mathbb{E}[f(\boldsymbol{\theta}^*) - f(\boldsymbol{\theta}_k)] + \frac{M\alpha_k^2 D^2}{2}. \quad (9) \end{aligned}$$

RHS of inequality (a) in (9) follows from the convexity property of the cost function  $f$  (i.e.,  $f(\boldsymbol{\theta}^*) - f(\boldsymbol{\theta}_k) \geq \langle \mathbf{g}(\boldsymbol{\theta}_k), \boldsymbol{\theta}^* - \boldsymbol{\theta}_k \rangle$ ), RHS of inequality (b) in (9) follows from Cauchy-Schwarz inequality and Assumption 1. Adding and subtracting  $f(\boldsymbol{\theta}^*)$  on the both sides of (9), we get

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta}_{k+1}) - f(\boldsymbol{\theta}^*)] &\leq (1 - \alpha_k) \mathbb{E}[f(\boldsymbol{\theta}_k) - f(\boldsymbol{\theta}^*)] + \\ &\alpha_k D \mathbb{E}[\|\mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k)\|] + \frac{M\alpha_k^2 D^2}{2}. \end{aligned}$$

It follows from Lyapunov's inequality [27] that  $\mathbb{E}[\|\mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k)\|] \leq \sqrt{\mathbb{E}[\|\mathbf{g}(\boldsymbol{\theta}_k) - \hat{\mathbf{g}}(\boldsymbol{\theta}_k)\|^2]}$ ; therefore, using the bound in Lemma (II.3), we can write

$$\begin{aligned} \mathbb{E}[f(\boldsymbol{\theta}_{k+1}) - f(\boldsymbol{\theta}^*)] &\leq (1 - \alpha_k) \mathbb{E}[f(\boldsymbol{\theta}_k) - f(\boldsymbol{\theta}^*)] \\ &+ \alpha_k M D^2 \sqrt{n+3} + \frac{M\alpha_k^2 D^2}{2}. \quad (10) \end{aligned}$$

By defining  $\delta_{k+1} = \mathbb{E}[f(\boldsymbol{\theta}_{k+1}) - f(\boldsymbol{\theta}^*)]$ , then, we have

$$\delta_{k+1} \leq (1 - \alpha_k)\delta_k + \psi_1 \alpha_k + \psi_2 \alpha_k^2, \quad \forall k \geq 0,$$

where  $\psi_1 = M D^2 \sqrt{n+3}$  and  $\psi_2 = \frac{M D^2}{2}$ . Invoking Comparison Lemma [28] and Lyapunov convergence analysis, we consider  $V_k = |\delta_k|$  as our candidate Lyapunov function to write

$$\begin{aligned} V_{k+1} - V_k &= |(1 - \alpha_k)\delta_k + \alpha_k(\psi_1 + \psi_2 \alpha_k)| - |\delta_k| \\ &\leq -\alpha_k |\delta_k| + \alpha_k(\psi_1 + \psi_2 \alpha_k) \\ &= -\beta \alpha_k |\delta_k| - (1 - \beta)\alpha_k |\delta_k| + \alpha_k(\psi_1 + \psi_2 \alpha_k), \end{aligned}$$

where  $0 < \beta < 1$ . Then if  $|\delta_k| \geq \frac{\psi_1 + \psi_2 \alpha_k}{1 - \beta}$ , we have exponential convergence to the ball with radius  $\frac{\psi_1 + \psi_2 \alpha_k}{1 - \beta}$ . Therefore, we can write  $\lim_{k \rightarrow \infty} \mathbb{E}[f(\boldsymbol{\theta}_k) - f(\boldsymbol{\theta}^*)] = \frac{\psi_1}{1 - \beta}$ , with the rate determined by  $O(\alpha_k)$ .

*Proof:* [Proof of Lemma IV.1] In Algorithm 2, we have

$$\begin{aligned} \|\mathbf{v}_k - \mathbf{g}(\boldsymbol{\theta}_k)\|^2 &= \|(1 - \gamma_k)\mathbf{v}_{k-1} + \gamma_k \hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k)\|^2 \\ &= \|(1 - \gamma_k)(\mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_k)) + \gamma_k(\hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k))\|^2 \\ &= (1 - \gamma_k)^2 \|\mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_k)\|^2 + \gamma_k^2 \|\hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k)\|^2 \\ &\quad + 2\gamma_k(1 - \gamma_k)(\mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_k))^\top (\hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k)) \\ &= (1 - \gamma_k)^2 \|\mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_k)\|^2 \\ &\quad + \mathbf{g}(\boldsymbol{\theta}_{k-1}) - \mathbf{g}(\boldsymbol{\theta}_{k-1})\|^2 + \gamma_k^2 \|\hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k)\|^2 \\ &\quad + 2\gamma_k(1 - \gamma_k)(\mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_k))^\top (\hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k)). \quad (11) \end{aligned}$$

Conditional expectation on both sides of (11) leads to

$$\mathbb{E}[\|\mathbf{v}_k - \mathbf{g}(\boldsymbol{\theta}_k)\|^2 | \sigma_{0:k}] = (1 - \gamma_k)^2 \mathbb{E}[\|\mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_k)\|^2 | \sigma_{0:k}]$$

$$\begin{aligned}
& + \mathbf{g}(\boldsymbol{\theta}_{k-1}) - \mathbf{g}(\boldsymbol{\theta}_{k-1}) \|^2 | \boldsymbol{\sigma}_{0:k} ] \\
& + \gamma_k^2 \mathbb{E} [ \| (\hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k)) \|^2 | \boldsymbol{\sigma}_{0:k} ], \quad (12)
\end{aligned}$$

where  $\boldsymbol{\sigma}_{0:k}$  is a vector containing all the randomness up to iteration  $k$ .

Express  $\mathbf{g}(\boldsymbol{\theta}_{k-1}) - \mathbf{g}(\boldsymbol{\theta}_k) = \zeta_k \mathbf{p}_k$ , where  $\mathbf{p}_k$  is a normalized vector ( $\|\mathbf{p}_k\| = 1$ ) and  $\zeta_k \in \mathbb{R}_{>0}$ . Next note that, from line 8 of Algorithm 2 we have  $\|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\| = \alpha_k \|\mathbf{s}_k - \boldsymbol{\theta}_k\| \leq \alpha_k D$  (recall line 8). Then, under Assumption 2 we can write  $\|\mathbf{g}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_{k-1})\| \leq M \|\boldsymbol{\theta}_k - \boldsymbol{\theta}_{k-1}\| \leq M \alpha_k D$ . Thus,  $\lim_{k \rightarrow \infty} \zeta_k = 0$  at least as fast as  $\alpha_k$ .

Hence, if we choose  $\alpha_k$  and  $\gamma_k$  such that  $\lim_{k \rightarrow \infty} \frac{\alpha_k}{\gamma_k} = 0$ , then, we can write  $\mathbf{g}(\boldsymbol{\theta}_{k-1}) - \mathbf{g}(\boldsymbol{\theta}_k) = \eta_k \gamma_k \mathbf{p}_k$  where  $\eta_k = \frac{\alpha_k}{\gamma_k} \in \mathbb{R}_{>0}$  and  $\lim_{k \rightarrow \infty} \eta_k = 0$ . Lastly, taking total expectation on both sides of (12), we have

$$\begin{aligned}
\mathbb{E} [\| \mathbf{v}_k - \mathbf{g}(\boldsymbol{\theta}_k) \|^2] & = (1 - \gamma_k)^2 \mathbb{E} [\| \mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_{k-1}) \| \\
& + \eta_k \gamma_k \mathbf{p}_k \|^2] + \gamma_k^2 \mathbb{E} [\| (\hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k)) \|^2] \\
& = (1 - \gamma_k)^2 (1 + \gamma_k)^2 \mathbb{E} [\| \frac{1}{1 + \gamma_k} (\mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_{k-1})) \\
& + \frac{\gamma_k}{1 + \gamma_k} \eta_k \mathbf{p}_k \|^2] + \gamma_k^2 \mathbb{E} [\| (\hat{\mathbf{g}}(\boldsymbol{\theta}_k) - \mathbf{g}(\boldsymbol{\theta}_k)) \|^2]. \quad (13)
\end{aligned}$$

Using the result from Lemma II.3 in (13), we have

$$\begin{aligned}
\mathbb{E} [\| \mathbf{v}_k - \mathbf{g}(\boldsymbol{\theta}_k) \|^2] & \leq (1 - \gamma_k)^2 (1 + \gamma_k)^2 \mathbb{E} [\| \frac{1}{1 + \gamma_k} (\mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_{k-1})) \\
& + \frac{\gamma_k}{1 + \gamma_k} \eta_k \mathbf{p}_k \|^2] + \gamma_k^2 \mathbf{G}. \quad (14)
\end{aligned}$$

Subsequently, invoking Jensen's inequality in (14), we have

$$\begin{aligned}
\mathbb{E} [\| \mathbf{v}_k - \mathbf{g}(\boldsymbol{\theta}_k) \|^2] & \leq (1 - \gamma_k)^2 (1 + \gamma_k)^2 \mathbb{E} [\| \mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_{k-1}) \|^2] \\
& + (1 - \gamma_k)^2 (1 + \gamma_k)^2 \gamma_k \eta_k^2 + \gamma_k^2 \mathbf{G} \\
& \stackrel{(a)}{\leq} (1 - \gamma_k) \mathbb{E} [\| \mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_{k-1}) \|^2] + (1 - \gamma_k) \gamma_k \eta_k^2 + \gamma_k^2 \mathbf{G} \\
& \stackrel{(b)}{\leq} (1 - \gamma_k) \mathbb{E} [\| \mathbf{v}_{k-1} - \mathbf{g}(\boldsymbol{\theta}_{k-1}) \|^2] + \gamma_k \eta_k^2 + \gamma_k^2 \mathbf{G}, \quad (15)
\end{aligned}$$

RHS of inequality (a) in (15) follows from  $(1 - \gamma_k)^2 (1 + \gamma_k) = (1 - \gamma_k)(1 - \gamma_k^2) \leq (1 - \gamma_k)$ , RHS of inequality (b) in (15) follows from  $(1 - \gamma_k) \leq 1$ . Denoting  $\delta_k = \mathbb{E} [\| \mathbf{v}_k - \mathbf{g}(\boldsymbol{\theta}_k) \|^2]$ , we have  $\delta_k \leq (1 - \gamma_k) \delta_{k-1} + \gamma_k \eta_k^2 + \gamma_k^2 \mathbf{G}$ .

Invoking Comparison Lemma [28] and Lyapunov convergence analysis, we consider  $V_{k-1} = |d_{k-1}|$  as our candidate Lyapunov function to write

$$\begin{aligned}
V_k - V_{k-1} & = |(1 - \gamma_k) d_{k-1} + \gamma_k \eta_k^2 + \gamma_k^2 \mathbf{G}| - |d_{k-1}| \\
& \leq -\gamma_k |d_{k-1}| + \gamma_k \eta_k^2 + \gamma_k^2 \mathbf{G} \\
& = -\beta \gamma_k |d_{k-1}| - (1 - \beta) \gamma_k |d_{k-1}| + \gamma_k \eta_k^2 + \gamma_k^2 \mathbf{G}
\end{aligned}$$

where  $0 < \beta < 1$ . Then if  $|d_{k-1}| \geq \frac{\eta_k^2 + \gamma_k \mathbf{G}}{(1 - \beta)}$ , we have exponential convergence to the ball with radius  $\frac{\eta_k^2 + \gamma_k \mathbf{G}}{(1 - \beta)}$ . Then, the overall rate of the convergence of  $\lim_{k \rightarrow \infty} \mathbb{E} [\| \mathbf{v}_k - \mathbf{g}(\boldsymbol{\theta}_k) \|^2] = 0$  is determined by  $\min\{O(\eta_k^2), O(\gamma_k)\}$ .

*Proof:* [Proof of Theorem IV.1] Following the same procedure as in the proof of Theorem III.1, we obtain  $\mathbb{E}[f(\mathbf{x}_{k+1}) - f(\boldsymbol{\theta}^*)] \leq (1 - \alpha_k) \mathbb{E}[f(\mathbf{x}_k) - f(\boldsymbol{\theta}^*)] + \alpha_k D \sqrt{\mathbb{E}[\| \mathbf{g}(\boldsymbol{\theta}_k) - \mathbf{v}_k \|^2]} + \frac{M \alpha_k^2 D^2}{2}$ . By defining  $\delta_{k+1} = \mathbb{E}[f(\boldsymbol{\theta}_{k+1}) - f(\boldsymbol{\theta}^*)]$  we can write  $\delta_{k+1} \leq (1 - \alpha_k) \delta_k + \psi_1 \alpha_k + \psi_2 \alpha_k^2$ ,  $\forall k \geq 0$  where  $\psi_1 =$

$D \sqrt{\mathbb{E}[\| \mathbf{g}(\boldsymbol{\theta}_k) - \mathbf{v}_k \|^2]}$  and  $\psi_2 = \frac{M D^2}{2}$ . Same as the proof in theorem III.1, choosing  $V_k = |\delta_k|$  as our candidate Lyapunov function, we obtain

$$V_{k+1} - V_k \leq -\beta \alpha_k |\delta_k| - (1 - \beta) \alpha_k |\delta_k| + \alpha_k (\psi_1 + \psi_2 \alpha_k),$$

where  $0 < \beta < 1$ . Then if  $|\delta_k| \geq \frac{\psi_1 + \psi_2 \alpha_k}{1 - \beta}$ , we have exponential convergence to the ball with radius  $\frac{\psi_1 + \psi_2 \alpha_k}{1 - \beta}$ . Hence, using Lemma IV.1, the overall rate of convergence of  $\lim_{k \rightarrow \infty} \mathbb{E}[f(\boldsymbol{\theta}_k) - f(\boldsymbol{\theta}^*)] = 0$  is determined by  $\min\{O(\frac{\alpha_k}{\gamma_k}), O(\gamma_k)\}$ .

*Proof:* [Proof of Theorem IV.2] Under Assumption 2, we can write  $\mathbb{E}[f(\boldsymbol{\theta}_{k+1}) | \boldsymbol{\theta}_k] \leq f(\boldsymbol{\theta}_k) + \langle \mathbf{g}(\boldsymbol{\theta}_k), \mathbb{E}[\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k | \boldsymbol{\theta}_k] \rangle + \frac{M}{2} \mathbb{E} [\| \boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k \|^2 | \boldsymbol{\theta}_k]$ . Substituting the update rule (line 8) of Algorithm 2 in this relationship leads to

$$\begin{aligned}
f(\boldsymbol{\theta}_{k+1}) & \leq f(\boldsymbol{\theta}_k) + \langle \mathbf{g}(\boldsymbol{\theta}_k), \boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k \rangle + \frac{M}{2} \|\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k\|^2 \\
& = f(\boldsymbol{\theta}_k) + \langle \mathbf{g}(\boldsymbol{\theta}_k), \alpha_k (\mathbf{s}_k - \boldsymbol{\theta}_k) \rangle + \frac{M}{2} \|\alpha_k (\mathbf{s}_k - \boldsymbol{\theta}_k)\|^2 \\
& \stackrel{(a)}{\leq} f(\boldsymbol{\theta}_k) + \alpha_k \langle \mathbf{g}(\boldsymbol{\theta}_k), \mathbf{s}_k - \boldsymbol{\theta}_k \rangle + \frac{M D^2 \alpha_k^2}{2}, \quad (16)
\end{aligned}$$

RHS of inequality (a) in (16) follows from Assumption 1. Adding and subtracting  $\alpha_k \langle \mathbf{v}_k, \mathbf{s}_k - \boldsymbol{\theta}_k \rangle$  into (16) results in

$$\begin{aligned}
f(\boldsymbol{\theta}_{k+1}) & \leq f(\boldsymbol{\theta}_k) + \alpha_k \langle \mathbf{v}_k, \mathbf{s}_k - \boldsymbol{\theta}_k \rangle \\
& + \alpha_k \langle \mathbf{g}(\boldsymbol{\theta}_k) - \mathbf{v}_k, \mathbf{s}_k - \boldsymbol{\theta}_k \rangle + \frac{M D^2 \alpha_k^2}{2} \\
& \stackrel{(a)}{\leq} f(\boldsymbol{\theta}_k) - \alpha_k \mathcal{K}(\boldsymbol{\theta}_k) + D \alpha_k \|\mathbf{g}(\boldsymbol{\theta}_k) - \mathbf{v}_k\| + \frac{M D^2 \alpha_k^2}{2}. \quad (17)
\end{aligned}$$

RHS of inequality (a) in (17) follows from Assumption 1, use of the definition of  $\mathcal{K}(\boldsymbol{\theta}_k)$  and Cauchy-Schwarz inequality. Taking total expectations on both sides, we obtain

$$\begin{aligned}
\mathbb{E}[f(\boldsymbol{\theta}_{k+1})] & \leq \mathbb{E}[f(\boldsymbol{\theta}_k)] - \alpha_k \mathbb{E}[\mathcal{K}(\boldsymbol{\theta}_k)] \\
& + D \alpha_k \mathbb{E}[\|\mathbf{g}(\boldsymbol{\theta}_k) - \mathbf{v}_k\|] + \frac{M D^2 \alpha_k^2}{2} \\
& \leq \mathbb{E}[f(\boldsymbol{\theta}_k)] - \alpha_k \mathbb{E}[\mathcal{K}(\boldsymbol{\theta}_k)] + D \alpha_k \sqrt{\mathbb{E}[\| \mathbf{g}(\boldsymbol{\theta}_k) - \mathbf{v}_k \|^2]} + \\
& \frac{M D^2 \alpha_k^2}{2} \leq \mathbb{E}[f(\boldsymbol{\theta}_k)] - \alpha_k \mathbb{E}[\mathcal{K}(\boldsymbol{\theta}_k)] + O(\frac{1}{k^{1.25}}) + O(\frac{1}{k^2}), \quad (18)
\end{aligned}$$

The last summands are used as a shorthand notation, for brevity, to represent decaying terms (Recall Lemma IV.1). Rearranging and summing  $k$  from 1 to  $K$ , we have

$$\begin{aligned}
\sum_{k=1}^K \alpha_k \mathbb{E}[\mathcal{K}(\boldsymbol{\theta}_k)] & \leq \frac{f(\boldsymbol{\theta}_0) - \mathbb{E}[f(\boldsymbol{\theta}_{K+1})]}{\sum_{k=1}^K \alpha_k} \\
& + \frac{\sum_{k=1}^K O(\frac{1}{k^{1.25}})}{\sum_{k=1}^K \alpha_k} + \frac{\sum_{k=1}^K O(\frac{1}{k^2})}{\sum_{k=1}^K \alpha_k} \\
& \leq \frac{f(\boldsymbol{\theta}_0) - f(\boldsymbol{\theta}^*)}{\sum_{k=1}^K \alpha_k} + \frac{\sum_{k=1}^K O(\frac{1}{k^{1.25}})}{\sum_{k=1}^K \alpha_k} + \frac{\sum_{k=1}^K O(\frac{1}{k^2})}{\sum_{k=1}^K \alpha_k}, \quad (19)
\end{aligned}$$

where both  $\sum_{k=1}^K O(\frac{1}{k^{1.25}})$  and  $\sum_{k=1}^K O(\frac{1}{k^2})$  are convergent series which can be verified using the integral test. From (19) and substituting  $\sum_{k=1}^K \alpha_k \approx \log K$  for large  $K$ , we have

$$\mathbb{E}[\mathcal{K}(\boldsymbol{\theta}_K)] \leq \frac{f(\boldsymbol{\theta}_0) - f(\boldsymbol{\theta}^*)}{\log K} + E \quad (20)$$

where  $E$  is of order  $O(\frac{1}{\log K})$ , which concludes the proof.