

Fast and Slow Gradient Approximation for Binary Neural Network Optimization

Xinquan Chen¹, Junqi Gao^{1,2}, Biqing Qi^{2*}, Dong Li^{1,2}, Yiang Luo¹, Fangyuan Li¹, Pengfei Li¹

¹Harbin Institute of Technology, Harbin, P.R.China,

²Shanghai Artificial Intelligence Laboratory, Shanghai, P.R.China

{xinquanchen0117,gjunqi97,qibiqing7}@gmail.com,

{arvinlee826,normanluo668,jacklee19900212,lipengfei0208}@gmail.com

Abstract

Binary Neural Networks (BNNs) have garnered significant attention due to their immense potential for deployment on edge devices. However, the non-differentiability of the quantization function poses a challenge for the optimization of BNNs, as its derivative cannot be backpropagated. To address this issue, hypernetwork based methods, which utilize neural networks to learn the gradients of non-differentiable quantization functions, have emerged as a promising approach due to their adaptive learning capabilities to reduce estimation errors. However, existing hypernetwork based methods typically rely solely on current gradient information, neglecting the influence of historical gradients. This oversight can lead to accumulated gradient errors when calculating gradient momentum during optimization. To incorporate historical gradient information, we design a Historical Gradient Storage (HGS) module, which models the historical gradient sequence to generate the first-order momentum required for optimization. To further enhance gradient generation in hypernetworks, we propose a Fast and Slow Gradient Generation (FSG) method. Additionally, to produce more precise gradients, we introduce Layer Recognition Embeddings (LRE) into the hypernetwork, facilitating the generation of layer-specific fine gradients. Extensive comparative experiments on the CIFAR-10 and CIFAR-100 datasets demonstrate that our method achieves faster convergence and lower loss values, outperforming existing baselines.

Code — <https://github.com/two-tiger/FSG>

Introduction

In recent years, Deep Neural Networks (DNNs) have demonstrated remarkable achievements across various computer vision tasks (Girshick et al. 2014; He et al. 2016; Lecun, Bengio, and Hinton 2015; Everingham et al. 2015). However, the increasing number of layers and parameters in deep learning models poses significant challenges. These challenges often result in substantial model sizes and elevated computational demands, hindering deployment on resource-constrained devices such as smartphones, cameras, and drones. To address these limitations, a plethora of compression techniques have emerged, including network pruning (Ding et al. 2021; Luo, Wu, and Lin 2017), low-rank

approximation (Denton et al. 2014; Hayashi et al. 2019), architectural redesign (Chen et al. 2020; Iandola et al. 2016), and network quantization (Banner et al. 2018; Ajanthan et al. 2021). Among these, network quantization—which employs reduced-precision data types for weight storage and computations—holds significant promise. For instance, Binary Neural Networks (BNNs), which restrict weights to binary values (-1 and $+1$) (Courbariaux, Bengio, and David 2015), can achieve remarkable efficiency by saving approximately 32 times in memory and boosting inference speed by a factor of 58 compared to full-precision networks (Rastegari et al. 2016).

Due to the non-differentiable nature of the sign function, conventional backpropagation is hindered in updating gradients for BNNs. BinaryConnect (Courbariaux, Bengio, and David 2015) addresses this challenge by introducing the Straight-Through Estimator (STE), a technique used to approximate the gradient of the sign function during backpropagation. This method has become the fundamental paradigm for BNN optimization due to its effectiveness. However, a significant issue remains: the discrepancy between the binary quantization applied during forward propagation and the gradient estimation during backpropagation leads to a gradient mismatch. This mismatch introduces inherent estimation errors, which impede the convergence and optimization of BNNs. As gradient propagation continues, these errors accumulate, further limiting the effectiveness of BNNs optimization.

To address the aforementioned challenges, a variety of methods have been developed. These include redesigned straight-through estimators (Qin et al. 2020; Xu et al. 2021; Wu et al. 2023; Lin et al. 2020) and hypernetwork based methods, which utilize differentiable neural networks (Chen, Wang, and Pan 2019; Liu et al. 2020) to directly learn the derivatives of non-differentiable functions, converting non-differentiable parts into differentiable operations for optimization. However, the addition of neural networks introduces more complex operations and higher computational costs, which has led to the potential of hypernetwork based methods being overlooked. In fact, the ultimate goal of BNNs is to deploy on resource limited devices, and the training process for obtaining a quantized network version is not constrained by the computational limitations of the target device. Moreover, despite careful design, straight-

*Corresponding author

through estimators cannot perfectly approximate the derivative of the binary thresholding function, resulting in inherent estimation errors and gradient mismatch. Hypernetwork based methods offer a solution by dynamically adjusting the gradient approximation during training, thereby continually reducing estimation errors. This capability allows BNNs to optimize more flexibly and achieve higher performance. Nonetheless, the current hypernetwork based methods (Chen, Wang, and Pan 2019) use the current weight gradient and the full precision weight to generate the final gradient. This method of generating gradients uses limited information and ignores the guiding role of historical gradients. Yet the optimization objective of the hypernetwork aligns with that of the final task, rather than the objective of generating a more suitable weight gradient, and does not match the optimization goal of generating a more appropriate weight gradient.

Motivated by the concept of momentum in Stochastic Gradient Descent with Momentum (SGD-M), which uses historical gradients to inform the current gradient, we have designed a Historical Gradient Storage (HGS) module. Momentum stabilizes the direction of gradient descent and accelerates optimization. Our HGS module stores a portion of historical gradients, treating them as time-series data that encapsulate gradient change information. These historical gradient sequences are then input into the hypernetwork to generate weight gradients. Compared to using only the current gradient for learning, incorporating historical gradients provides additional context, enabling the network to generate gradients that more appropriate for BNN optimization. This approach enhances the optimization process for BNNs, facilitating more effective and efficient training.

To generate gradients more precisely, we have designed the Fast and Slow gradient Generation (FSG) method, inspired by the optimization method of SGD-M. This method employs two hypernetworks, referred to as fast-net and slow-net. Slow-net utilizes models with strong long-sequence modeling capabilities, such as Mamba (Gu and Dao 2023) and LSTM (Hochreiter and Schmidhuber 1997), to capture the essence of historical gradient sequences in order to generate gradients consistent with the concept of momentum. This process requires summarizing accumulated historical data, similar to slow evolution. While fast-net utilizes a Multi-Layer Perceptron (MLP) to extract high-dimensional features from the current gradient, which can quickly generate the current weight gradient. The final gradient is obtained by combining these two components. Gradient generation is layer-specific, so in order to further guide slow-net to generate appropriate momentum, we add Layer Recognition Embeddings (LRE). Specifically, we add a learnable layer embedding vector as the recognition of layer information. Compared with previous methods using a single hypernetwork, FSG enriches the process of gradient generation, allowing historical gradients to participate in gradient generation and providing richer information for gradient generation.

We conduct comprehensive experiments on the CIFAR-10 and CIFAR-100 datasets. The results demonstrate that our method outperforms various competitive benchmarks. In

summary, our contributions are as follows:

- Inspired by the idea of momentum, we examine the impact of historical gradients on gradient generation. To this end, we propose a HGS module for storing historical gradient sequences. By modeling these sequences, we generate gradient momentum, which guides the network in producing more suitable gradients.
- We design a FSG method to guide the gradient generation process, fully leveraging the information from both historical and current gradients to produce more refined gradients. Additionally, we propose LRE to assist the hypernetwork in learning layer-specific gradients.
- Extensive experiments on the CIFAR-10 and CIFAR-100 datasets demonstrate that our proposed FSG method achieves superior performance and faster convergence.

Related Work

The origins of BNNs can be traced back to Courbariaux’s groundbreaking work (Courbariaux et al. 2016), where they employ a sign function to binarize network weights. To address the challenge of the sign function’s non-differentiability during backpropagation, BinaryConnect (Courbariaux, Bengio, and David 2015) introduces the STE as an effective approximation technique. The outstanding performance of STE has since become the cornerstone of optimization strategies in this field.

Later work primarily focuses on constraining the key aspect of BNN optimization: the derivative of the sign function. For instance, IR-Net (Qin et al. 2020) employs an Error Decay Estimator (EDE) in a two-stage process to gradually approximate the sign function and reduce gradient error. RBNN (Lin et al. 2020) introduces a training-aware approximation of the sign function for gradient computation during backpropagation. FDA (Xu et al. 2021) utilizes combinations of sine functions in the Fourier frequency domain to estimate the gradient of the sign function during BNN training and employ neural networks to learn the approximation error. ReSTE (Wu et al. 2023) leverages power functions to design a correction estimator that balances estimation error and gradient stability.

Despite efforts to design approximation functions, inherent estimation errors persist between these approximations and the sign function’s derivatives, which are fundamentally uneliminable. Consequently, some recent approaches employ hypernetworks to transform the non-differentiable components of binary neural networks into differentiable operations. For instance, MetaQuant (Chen, Wang, and Pan 2019) proposes incorporating a neural network into the backpropagation process to learn the non-differentiable gradients. Similarly, QuantNet (Liu et al. 2020) utilizes a differentiable subnetwork to directly binarize full-precision weights, avoiding reliance on the STE or any learnable gradient estimator. Hypernetwork based methods (Andrychowicz et al. 2016) offer the advantage of adaptively adjusting gradient approximations during training, thereby reducing approximation errors. Unlike manually designed STE, these methods can dynamically refine the gradient approximations. However, they are limited in their ability to gen-

erate precise gradients. While the network learns based on the final loss, the optimization objective of hypernetworks is to generate better gradients for the sign function, leading to differing optimization goals.

Preliminaries

Definition of BNNs

BNNs are often derived from fully accurate networks (Rastegari et al. 2016). For a neural network model f with n layers, parameterized by $\mathcal{W} = [W_1, W_2, \dots, W_n]$, where W_i represents the weight of the i -th layer, we define a pre-processing function $\mathcal{A}(\cdot)$ and a quantization function $Q(\cdot)$ to binarize the weights. This article investigates the pre-processing and quantization methods of DoReFa (Zhou et al. 2016). The pre-processing function $\mathcal{A}(\cdot)$ normalizes the weights, while the quantization function $Q(\cdot)$ converts each value in the weight matrix to $\{-1, +1\}$. We have,

$$\hat{\mathcal{W}} = \mathcal{A}(\mathcal{W}) = \frac{\tanh(\mathcal{W})}{2\max(|\tanh(\mathcal{W})|)} + \frac{1}{2}, \quad (1)$$

$$\mathcal{W}_b = Q(\hat{\mathcal{W}}) = 2 \frac{\text{round}[(2^k - 1)\hat{\mathcal{W}}]}{2^k - 1} - 1. \quad (2)$$

For a training set $\{\mathbf{x}, \mathbf{y}\}$ with \mathbf{N} instances, we incorporate quantization in the forward process. Consequently, the loss function for training a binary neural network is given by $\ell(f(Q(\mathcal{A}(\mathcal{W})); \mathbf{x}), \mathbf{y})$. During the backpropagation process, we compute the derivative of ℓ with respect to \mathcal{W} ,

$$\frac{\partial \ell}{\partial \mathcal{W}} = \frac{\partial \ell}{\partial Q(\mathcal{A}(\mathcal{W}))} \frac{\partial Q(\mathcal{A}(\mathcal{W}))}{\partial \mathcal{A}(\mathcal{W})} \frac{\partial \mathcal{A}(\mathcal{W})}{\partial \mathcal{W}}, \quad (3)$$

where $\frac{\partial Q(\mathcal{A}(\mathcal{W}))}{\partial \mathcal{A}(\mathcal{W})}$ is the Dirac delta function. This introduces a discontinuity in the backpropagation process, causing it to be interrupted at this point.

Hypernetwork Based Method

Hypernetwork based methods (Chen, Wang, and Pan 2019) introduce a shared neural network \mathcal{M} , which is trained simultaneously with the binary neural network. During the backpropagation process, the hypernetwork \mathcal{M} receives two inputs, $g_{\mathcal{W}} = \frac{\partial \ell}{\partial \mathcal{W}}$ and $\hat{\mathcal{W}}$, and generates the derivative of ℓ to the quantization function $Q(\mathcal{A}(\cdot))$, i.e.

$$\frac{\partial \ell}{\partial \mathcal{A}(\cdot)} = \mathcal{M}_{\phi}(g_{\mathcal{W}_b}, \hat{\mathcal{W}}), \quad (4)$$

the pre-processing function is differentiable, so the final derivative is,

$$\frac{\partial \ell}{\partial \mathcal{W}} = \mathcal{M}_{\phi}(g_{\mathcal{W}}, \hat{\mathcal{W}}) \frac{\partial \mathcal{A}(\mathcal{W})}{\partial \mathcal{W}}, \quad (5)$$

please note that in the first iteration, the hypernetwork is not working as the gradient has not yet been generated. At this point, the STE function will be temporarily used as a substitute.

Methodology

Motivation: Optimization Requires Historical Gradients

Gradient descent, a cornerstone optimization algorithm as described in (Saad 1998), relies on the gradient of the objective function with respect to the parameters. However, its reliance on instantaneous gradient can lead to premature convergence in local minima, hindering escape. To address this limitation, the momentum method incorporates exponential weighting from moving averages. By averaging past and current gradients with exponentially decaying weights, the method imparts inertia to the update direction. This inertia helps the algorithm to escape local minima more effectively, leading to improved convergence. SGD-M has the following update formula,

$$\begin{cases} \mathbf{v}_{t+1} = \beta \mathbf{v}_t - \alpha \nabla f(\mathbf{x}_t), \\ \mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{v}_{t+1}, \end{cases} \quad (6)$$

where \mathbf{v}_{t+1} is the momentum, β is the momentum constant, α is the step size, also known as the learning rate, $\nabla f(\mathbf{x}_t)$ is the gradient of f at \mathbf{x}_t , and the momentum \mathbf{v}_{t+1} can be expanded to:

$$\begin{aligned} \mathbf{v}_{t+1} &= \beta \mathbf{v}_t - \alpha \nabla f(\mathbf{x}_t) \\ &= \beta^2 \mathbf{v}_{t-1} - \beta \alpha \nabla f(\mathbf{x}_{t-1}) - \alpha \nabla f(\mathbf{x}_t) \\ &= \beta^3 \mathbf{v}_{t-2} - \beta^2 \alpha \nabla f(\mathbf{x}_{t-2}) - \beta \alpha \nabla f(\mathbf{x}_{t-1}) \\ &\quad - \alpha \nabla f(\mathbf{x}_t) \dots \end{aligned} \quad (7)$$

From the above expansion of momentum, we can see that momentum is a linear combination of historical gradients, which means that the direction and magnitude of historical gradients can guide the direction of the current update and the step size in optimization. Therefore, it can be considered that when optimizing non-differentiable functions, it is also possible to benefit from historical gradients. We focus on the role of historical gradients and introduce a module for storing historical gradients, called the Historical Gradient Storage (HGS) module, which saves the most recent l step gradients.

We treat the saved historical gradient sequence as a time-series data containing gradient change information. Mamba is a state-space model with a selection mechanism that can effectively filter out noise in historical gradient sequences, efficiently integrate historical gradients and provide more accurate momentum predictions. Therefore, we use a shared Mamba block (Gu and Dao 2023) to complete the modeling of the gradient sequence of each layer. We explain in detail the reasons for choosing Mamba as the momentum generator in the Appendix based on the principles of Mamba, and compare the performance of different sequence modeling models as momentum generators in Ablation Experiment.

Fast and Slow Gradient Generation Method

Historical Gradient Storage Module. In order to save the historical gradient, we introduce a module for storing historical gradients, termed the HGS module, and denote as $\mathcal{H} = \{h_1, h_2, \dots, h_n\}$, where h_i represents the historical gradient snapshot of the i -th layer. For simplicity, we assume

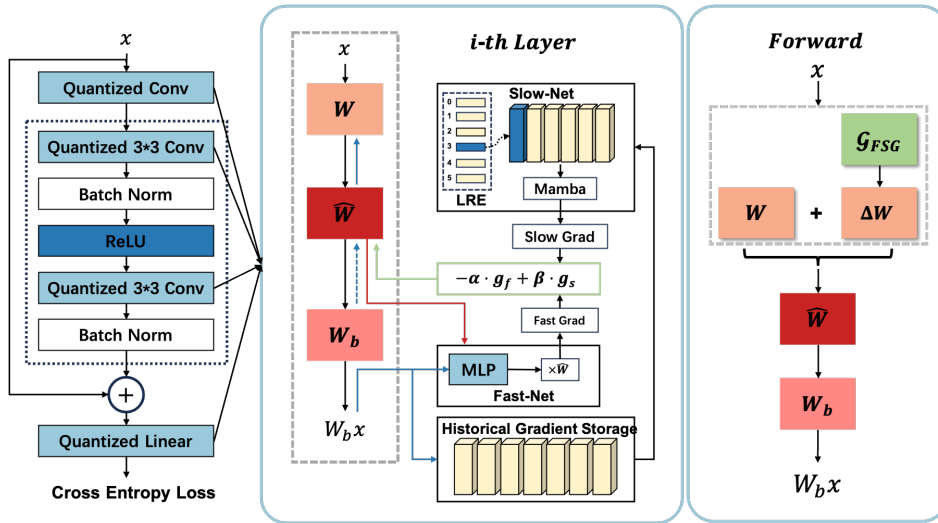


Figure 1: Fast and Slow Gradient Generation Illustration. Take ResNet as an example. During the backpropagation, the weight gradients from the previous iteration step are fed into the HGS and fast-net. Fast-net uses MLP to learn the scale of the weights, thereby obtaining the fast grad. The slow-net receives the historical gradient sequence from HGS and adds a LRE vector at the front of the sequence, then uses the mamba block to generate the slow grad. Ultimately, the slow grad and fast grad are combined through a weighted sum to generate the final gradients, replacing the non-differentiable parts (indicated by the blue dashed arrows). The forward process will be explained in Section Training of FSG.

that n layers of convolutional layers need to be generated, and the weights of the neural network are represented as $\mathcal{W} = [W_1, W_2, \dots, W_n]$, where $W_i \in \mathbb{R}^{C_{out}^i \times C_{in}^i \times K^i \times K^i}$ is the weight of the i -th convolutional layer has a corresponding gradient $g_{W_i} \in \mathbb{R}^{C_{out}^i \times C_{in}^i \times K^i \times K^i}$. We flatten the gradient to one dimension and save it as $g_{w_i} \rightarrow \bar{g}_{w_i} \in \mathbb{R}^{\xi^i \times 1}$, where $\xi^i = C_{out}^i \cdot C_{in}^i \cdot K^i \cdot K^i$, to control memory usage, we only retain the gradients from the most recent l steps,

$$h_i^t = \text{Concat}[\bar{g}_{w_i}^{t-l+1}, \dots, \bar{g}_{w_i}^{t-1}, \bar{g}_{w_i}^t] \in \mathbb{R}^{\xi^i \cdot l \times 1}, \quad (8)$$

where Concat denotes the concatenation along the appropriate dimension to form a vector of length $\xi^i \cdot l$, $\bar{g}_{w_i}^t$ is the flattened gradient of the weight W_i at step t .

Fast and Slow Gradient Generation Mechanism. As shown in the Fig. 1, we introduce two hypernetworks shared by each layer \mathcal{M}_f and \mathcal{M}_s , named fast-net and slow-net respectively. Fast-net receives the full-precision weights \hat{W}_i^t of the previous step and the gradients $g_{W_i}^t$ of the previous step to generate the gradients $g_{W_i}^{t+1}$ for updating. Slow-net receives the historical gradient sequence h_i^t provided by the HGS module to generate the gradient momentum v_{t+1} for updating. Therefore, there is the following update form, for the i -th layer,

$$\mathbf{W}_i^{t+1} = \mathbf{W}_i^t - \alpha \mathcal{M}_f(g_{W_i}^t, \hat{W}_i^t) + \beta \mathcal{M}_s(h_i^t) \quad (9)$$

Slow-net is responsible for modeling historical gradient sequences to generate gradients, which is consistent with the gradient momentum composed of historical gradients. The generation of momentum requires the accumulation of historical information, similar to a slow evolution. Fast-net, on

the other hand, learns the high-dimensional features of the current gradient and is a fast gradient generation method.

Layer Recognition Embedding. To enable the Mamba block to discern the layer-specific information and prevent the gradients from confusing each other during the training process, inspired by ViM (Zhu et al. 2024) and BERT (Devlin et al. 2018), we propose LRE, which initializing a learnable embedding vector for each layer of the model $\mathbb{E} \in \mathbb{R}^{n \times d}$, d is the dimension of the embedding space. For a given layer index i , LRE will return the corresponding embedding vector, which we define as $\mathbf{t}_i \in \mathbb{R}^d$. Define $\bar{h} = [\bar{h}_1, \bar{h}_2, \dots, \bar{h}_n]$, where \bar{h}_i is the input vector formed by concatenating the vector \mathbf{t}_i with the projected historical gradient.

$$\begin{aligned} \bar{h}_i^t &= [\mathbf{t}_i; \bar{g}_{w_i}^{t-l} \mathbf{W}_a, \dots, \bar{g}_{w_i}^{t-1} \mathbf{W}_a, \bar{g}_{w_i}^t \mathbf{W}_a] \\ &= [\mathbf{t}_i; h_i \mathbf{W}_a] \in \mathbb{R}^{(\xi^i \cdot l + 1) \times d} \end{aligned} \quad (10)$$

where $\mathbf{W}_a \in \mathbb{R}^{1 \times d}$ is a projection matrix that maps each flattened gradient $\bar{g}_{w_i}^{t-k}$ onto a d dimensional space. Next, we input \bar{h}_i^t into the Mamba block, we slice the output vector, take the last ξ^i sequences, and project them to a 1-dimensional vector,

$$\bar{g}_Q = \text{Mamba}(\bar{h}_i^t)[:, -\xi^i :, :] \mathbf{W}_b \in \mathbb{R}^{\xi^i \times 1}, \quad (11)$$

where $\mathbf{W}_b \in \mathbb{R}^{d \times 1}$ is a projection matrix. We reshape the output vector $\bar{g}_Q \rightarrow g_Q \in \mathbb{R}^{C_{out}^i \times C_{in}^i \times K^i \times K^i}$ and obtain the derivative of quantized function. Compared to using only the current gradient information to generate gradients, this approach uses historical gradients to guide the generation of the current gradient, fully utilizing the information.

Training of FSG

Similar to (Chen, Wang, and Pan 2019), we will incorporate shared fast-net \mathcal{M}_f and slow-net \mathcal{M}_s into the forward process of the model, as shown in Fig. 1. Eq. 9 is used to merge the gradients generated by FSG and the weight W_i of i -th layer. Then we have the forward process in the i -th convolutional layer,

$$\begin{aligned}\hat{W}_i^{t+1} &= \mathcal{A}(W_i^{t+1}) \\ &= \mathcal{A}\left[W_i^t - \alpha \mathcal{M}_f(\mathcal{G}_{W_i}^t, \hat{W}_i^t) \frac{\partial \mathcal{A}(W_i^t)}{\partial W_i^t} + \beta \mathcal{M}_s(\hat{h}_i^t)\right]\end{aligned}\quad (12)$$

At the same time, we register the generated gradient $\mathcal{G}_{FSG} = \beta \mathcal{M}_s(\hat{h}_i^t) - \alpha \mathcal{M}_f(\mathcal{G}_{W_i}^t, \hat{W}_i^t) \frac{\partial \mathcal{A}(W_i^t)}{\partial W_i^t}$ in the optimizer to replace $g_W = \frac{\partial \ell}{\partial W}$, and we calculate the cross entropy loss \mathcal{L} based on the results generated by quantifying weights:

$$\mathcal{L} = \ell\left(f(Q(\hat{W}_i^{t+1}); \mathbf{x}), \mathbf{y}\right), \quad (13)$$

therefore, \mathcal{M}_f and \mathcal{M}_s are associated with the final loss and can be used for backpropagation to update parameters. We assume that \mathcal{M}_f and \mathcal{M}_s are parametered by ϕ_f and ϕ_s , respectively, during the backpropagation process,

$$\frac{\partial \ell}{\partial \phi_f^{t+1}} = \frac{\partial \ell}{\partial \hat{W}_i^{t+1}} \frac{\partial \hat{W}_i^{t+1}}{\partial \phi_f^{t+1}} = -\alpha \cdot \mathcal{G}_{FSG} \frac{\partial \mathcal{M}_f(\mathcal{G}_{W_i}^t, \hat{W}_i^t)}{\partial \phi_f^{t+1}} \quad (14)$$

$$\frac{\partial \ell}{\partial \phi_s^{t+1}} = \frac{\partial \ell}{\partial \hat{W}_i^{t+1}} \frac{\partial \hat{W}_i^{t+1}}{\partial \phi_s^{t+1}} = \beta \cdot \mathcal{G}_{FSG} \frac{\partial \mathcal{M}_s(\hat{h}_i^t)}{\partial \phi_s^{t+1}} \quad (15)$$

After the hypernetwork updates, we use the registered gradient \mathcal{G}_{FSG} to update all parameters with the optimizer.

Convergence Analysis of FSG

In this section, we analyze the convergence of the proposed fast and slow gradient generation mechanism under the general convex function condition. Randomly sample a dataset $\mathcal{D} = \{(a_1, b_1), (a_2, b_2), \dots, (a_N, b_N)\}$. The task is to predict the label b given the input a , i.e. determine an optimal function ϕ that minimizes the expected risk $\mathbb{E}[L(\phi(a), b)]$, where $L(\cdot, \cdot)$ represents the loss function and the function ϕ is the prediction function in a certain function space.

In practice, in order to reduce the range of the objective function, it is necessary to parameterize $\phi(\cdot)$ to $\phi(\cdot; x)$, where x is a parameter. Using empirical risk to approximate expected risk requires solving the following minimization problem:

$$\min_x \frac{1}{N} \sum_{i=1}^N L(\phi(a_i; x), b_i) = \mathbb{E}_{(a,b) \sim \hat{P}} [L(\phi(a; x), b)], \quad (16)$$

define $f_i(x) = L(\phi(a_i; x), b_i)$, then we only need to consider the following stochastic optimization problem:

$$\min_{x \in \mathbb{R}^n} f(x) \stackrel{\text{def}}{=} \frac{1}{N} \sum_{i=1}^N f_i(x) \quad (17)$$

Theorem 1. (Convergence of FSG) Let formula 9 run t iterations. By setting $\alpha = \frac{C}{\sqrt{t+1}}$,

$$\begin{aligned}& \mathbb{E}[f(\hat{\mathbf{x}}_t) - f(\mathbf{x}^*)] \\ & \leq \frac{\beta}{(1-\beta)(t+1)} (f(\mathbf{x}_0) - f(\mathbf{x}^*)) + \frac{(1-\beta)\|\mathbf{x}_0 - \mathbf{x}^*\|^2}{2C\Omega\kappa\sqrt{t+1}} \\ & \quad + \frac{C\Theta\rho(G^2 + \delta^2)}{2\Omega\kappa(1-\beta)\sqrt{t+1}}\end{aligned}\quad (18)$$

Where $C, \Omega, \kappa, \Theta, \rho$ is a normal number, \mathbf{x}^* is the optimal solution and $\hat{\mathbf{x}}_t = \sum_{k=0}^t \mathbf{x}_k / (t+1)$.

Remark. Theorem 1 describes the convergence of FSG, i.e. the expected error with the optimal solution decreases as the number of iterations increases for the iterative formula t times, and the convergence order is $1/\sqrt{t+1}$. The convergence speed is slow at the beginning of the iteration. As the number of iterations increases, the convergence speed will increase. It can almost guarantee convergence.

The proof is given in the Appendix.

Experiment

In this section, we follow the setting of (Chen, Wang, and Pan 2019) and explore whether FSG can effectively learn the derivatives of quantization functions on the CIFAR-10/100 dataset. Then we follow the setting of (Chen, Wang, and Pan 2019; Qin et al. 2020) to compare it with several State-Of-The-Art (SOTA) methods.

Experiment Setup

In this work, we use CIFAR-10 and CIFAR-100 as benchmark datasets. CIFAR-10 (Krizhevsky, Hinton et al. 2009) and CIFAR-100 (Krizhevsky, Hinton et al. 2009) are commonly used datasets for image classification, consisting of 50000 training images and 10000 test images, with 10 and 100 categories respectively. Each image has a size of 32×32 and includes an RGB color channel. For CIFAR-10, we conduct experiments using ResNet-20/32/44 architecture, and for CIFAR-100, we conduct experiments using ResNet-56/110 architecture. For fair comparison, except for the first convolutional layer and the last fully connected layer, other layers are binarized in this experiment. See the Appendix for more experimental settings and details.

Experimental Results and Analysis

In this experiment, we aim to investigate whether the FSG can produce effective gradients. To this end, we compare FSG with the STE (Zhou et al. 2016), FCGrad (Chen, Wang, and Pan 2019) and LSTMFC (Chen, Wang, and Pan 2019). Our experiments are conducted on both SGD and Adam optimizers.

Tab. 1 shows the experimental results of STE, FCGrad, LSTMFC and FSG using different optimization methods on the CIFAR-10 dataset. The results show that the accuracy of FSG is significantly better than other types of hypernetworks, and it has a smaller loss, with ResNet-32 having the smallest difference of only 0.81% compared to the results of full precision inference.

Backbone	Forward	Optimization	Backward	Train Acc (%)	Test Acc (%)	Loss	FP Acc (%)
ResNet-20	DoReFa	SGD	STE	55.469 ± 3.342	46.81 ± 4.893	1.451 ± 0.395	91.70
			FCGrad	92.084 ± 1.103	88.54 ± 0.291	0.220 ± 0.021	
			LSTMFC	92.070 ± 1.012	88.28 ± 0.810	0.228 ± 0.027	
		FSG(ours)	92.302 ± 1.247	88.67 ± 0.633	0.218 ± 0.012		
		Adam	STE	93.138 ± 0.976	87.96 ± 0.173	0.151 ± 0.042	
			FCGrad	94.331 ± 0.986	89.96 ± 0.068	0.148 ± 0.027	
LSTMFC	94.768 ± 1.023		89.98 ± 0.103	0.148 ± 0.029			
FSG(ours)	94.926 ± 0.747	91.00 ± 0.804	0.143 ± 0.039				
ResNet-32	DoReFa	SGD	STE	55.469 ± 2.948	47.12 ± 4.683	1.443 ± 0.397	92.13
			FCGrad	95.377 ± 1.588	89.93 ± 0.246	0.154 ± 0.024	
			LSTMFC	95.807 ± 1.643	90.40 ± 0.149	0.150 ± 0.021	
		FSG(ours)	96.898 ± 1.476	90.53 ± 0.192	0.145 ± 0.019		
		Adam	STE	97.056 ± 1.010	89.47 ± 0.097	0.124 ± 0.023	
			FCGrad	95.288 ± 1.648	90.28 ± 0.068	0.900 ± 0.018	
LSTMFC	95.140 ± 1.779		90.15 ± 0.074	0.907 ± 0.023			
FSG(ours)	97.632 ± 1.310	91.42 ± 0.140	0.083 ± 0.016				
ResNet-44	DoReFa	SGD	STE	59.375 ± 3.322	49.73 ± 4.628	1.240 ± 0.081	93.55
			FCGrad	92.228 ± 2.029	89.93 ± 0.246	0.711 ± 0.024	
			LSTMFC	94.548 ± 1.947	90.40 ± 0.149	0.723 ± 0.022	
		FSG(ours)	98.916 ± 1.398	91.91 ± 0.122	0.031 ± 0.015		
		Adam	STE	97.656 ± 0.996	90.21 ± 0.075	0.083 ± 0.012	
			FCGrad	90.584 ± 1.573	90.98 ± 0.068	0.083 ± 0.022	
LSTMFC	90.147 ± 1.392		90.95 ± 0.074	0.082 ± 0.023			
FSG(ours)	99.262 ± 1.322	92.78 ± 0.053	0.022 ± 0.011				

Table 1: Results on CIFAR-10 dataset. Loss is the cross-entropy loss value. FP Acc is the result of the full precision version of the backbone on this dataset. The bold font is utilized to denote the optimal results, the underline is applied to indicate the suboptimal outcomes.

Tab. 2 shows the experimental results of STE, FCGrad, LSTMFC, and FSG on the CIFAR-100 dataset using different optimization methods. Similar to CIFAR-10, FSG demonstrates superior performance over LSTMFC across a range of scenarios, achieving a higher accuracy rate of 68.04% compared to LSTMFC’s 66.30%. This indicates that FSG’s incorporation of historical gradient information significantly improves the optimization process by providing more accurate estimates of the gradients needed for learning quantization schemes.

Convergence Analysis

In this comparative study, we aim to evaluate the convergence speed of different optimizers, namely STE, FCGrad, LSTMFC, and FSG, by analyzing their loss curves throughout the training process, as depicted in Fig 2. The loss curves are plotted to visualize the performance of each method during training, using the DoReFa quantization function and the ResNet44 network architecture on the CIFAR-10 dataset. The loss curves reveal that FSG exhibit superior convergence characteristics, with a notably faster descent rate and consistently lower loss values compared to FCGrad, LSTMFC and STE. Furthermore, compared to traditional SGD optimizer, the Adam optimizer exhibits a more stable and consistent optimization process, which reflects that our momentum generation gradient strategy is more suitable for the Adam optimizer.

Ablation Experiment

In this section, we investigate the effects of various sequence models as slow-nets on performance and examine the influence of the balancing parameter α , which governs the blending of fast and slow gradients within the FSG method, alongside the impact of the hyperparameter l , which represents the

length of the historical gradient memory fed into the slow network.

Using different slow-net. In this experiment, we use two sequence modeling models, including LSTM and Mamba, as slow-nets. For the CIFAR10 and CIFAR100 datasets, we employed ResNet-44 and ResNet-56 as the respective backbone architectures. As evident from Table 3, utilizing Mamba as a momentum generator exhibited superior performance, indicating that Mamba possesses enhanced capabilities in modeling long sequences. Its distinctive selection mechanism was found to aid in the filtering of noise within the historical gradient sequence.

Influence of β . In this experiments, we explore the sensitivity of the FSG method to the combination parameter β , selecting values of 0.9, 0.7, 0.5, 0.3, and 0.1 for evaluation on the CIFAR-100 dataset. We utilize the ResNet56 architecture and the DoReFa quantization function for our analysis. From the result in the Fig. 2, it can be found that the highest accuracy is achieved when β is set to 0.3. This indicates that the current gradient generated by fast-net dominates the update, which helps to respond quickly to changes in gradient and speed up convergence. At the same time, the gradient momentum generated by the historical gradient sequence can still reduce the gradient oscillation caused by noise to a certain extent. Based on these results, we recommend starting with a smaller value of β (around 0.3).

Influence of l . For l , we select 3, 4, 5, 6, and 7 to conduct experiments on CIFAR-100, using ResNet56 as the network architecture and DoReFa as the quantization function. From the result in the Fig. 2, it can be found that the highest accuracy is achieved when l is set to 6. This indicates that historical gradient memory requires an appropriate length, and shorter memory lengths cannot accumulate enough sequence information. A longer memory length in-

Backbone	Forward	Optimization	Backward	Train Acc (%)	Test Acc (%)	Loss	FP Acc (%)
ResNet-56	DoReFa	SGD	STE	51.126 ± 7.374	42.27 ± 8.143	1.890 ± 0.178	71.22
			FCGrad	73.934 ± 1.639	63.25 ± 0.935	0.861 ± 0.012	
			LSTMFC	73.934 ± 1.454	62.95 ± 2.183	0.901 ± 0.010	
		FSG(ours)	74.239 ± 1.691	64.02 ± 1.344	0.822 ± 0.012		
		Adam	STE	81.250 ± 0.642	64.78 ± 0.533	0.618 ± 0.032	
			FCGrad	79.944 ± 0.923	66.56 ± 0.351	0.684 ± 0.022	
LSTMFC	79.336 ± 0.931		66.30 ± 0.793	0.686 ± 0.024			
FSG(ours)	81.452 ± 0.873	69.48 ± 0.979	0.604 ± 0.019				
ResNet-110	DoReFa	SGD	STE	53.159 ± 9.793	43.42 ± 8.902	1.783 ± 0.211	72.54
			FCGrad	82.379 ± 1.553	65.15 ± 2.490	0.602 ± 0.023	
			LSTMFC	81.074 ± 1.457	64.75 ± 2.850	0.619 ± 0.027	
		FSG(ours)	84.733 ± 2.773	65.23 ± 0.822	0.601 ± 0.012		
		Adam	STE	85.156 ± 1.344	66.84 ± 1.205	0.548 ± 0.016	
			FCGrad	85.299 ± 1.284	68.74 ± 0.363	0.483 ± 0.014	
LSTMFC	85.699 ± 1.273		67.14 ± 1.286	0.503 ± 0.013			
FSG(ours)	85.738 ± 1.031	68.15 ± 0.973	0.462 ± 0.011				

Table 2: Results on CIFAR-100 dataset. Loss is the cross-entropy loss value. FP Acc is the result of the full precision version of the backbone on this dataset. The bold font is utilized to denote the optimal results, the underline is applied to indicate the suboptimal outcomes.

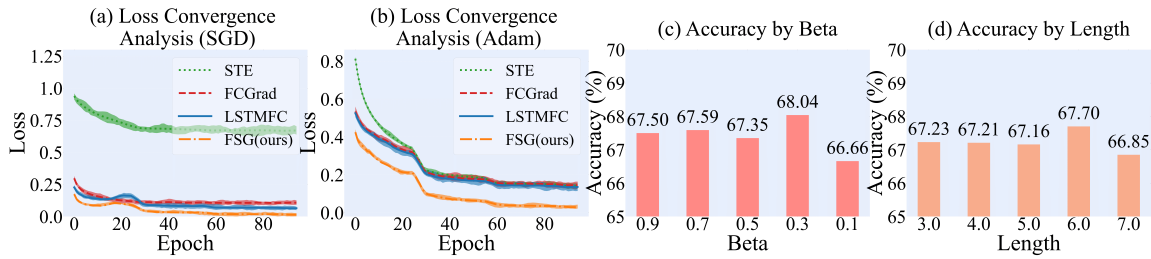


Figure 2: (a) Loss Curve of ResNet44 on CIFAR-10 Dataset with SGD Optimizer. (b) Loss Curve of ResNet44 on CIFAR-10 Dataset with Adam Optimizer. (c) Accuracy by Hyperparameter Beta. (d) Accuracy by Hyperparameter Length.

Backbone	Optimization	Slow-net	Test Acc (%)	FP Acc (%)
ResNet-44	CIFAR10	LSTM	92.07 ± 0.066	93.55
		Mamba	92.63 ± 0.052	
ResNet-56	CIFAR100	LSTM	67.60 ± 0.670	71.22
		Mamba	68.04 ± 0.948	

Table 3: Performance comparison of different sequence modeling models as slow-net.

Dataset	CIFAR-10	CIFAR-10	CIFAR-100
Backbone	ResNet-20	ResNet-44	ResNet-56
FP	91.70	93.55	71.22
DoReFa	87.96 ± 0.173	90.21 ± 0.075	64.78 ± 0.533
ReSTE	89.26 ± 0.884	90.25 ± 0.783	65.59 ± 0.733
LSTMFC	90.03 ± 0.299	91.91 ± 0.538	66.48 ± 0.808
IR-Net	90.80 ± 0.971	92.37 ± 0.532	68.94 ± 0.967
RBNN	90.89 ± 0.921	91.59 ± 0.874	67.10 ± 0.933
FSG(ours)	91.00 ± 0.804	92.78 ± 0.053	69.48 ± 0.979

Table 4: Performance comparison with SOTA methods. FP is the full-precision version of the backbone. The bold font is utilized to denote the optimal results, the underline is applied to indicate the suboptimal outcomes.

roduces more historical gradient noise, which is not conducive to the generation of gradient momentum.

Comparison with SOTA Methods

In order to verify the performance of FSG, we conducted performance studies on it compared to other BNN optimization methods. We follow the experimental setup of (Chen, Wang, and Pan 2019; Qin et al. 2020). A series of SOTA methods are compared on the CIFAR-10 and CIFAR-100 datasets to validate performance, including DoReFa (Zhou et al. 2016), LSTMFC (Chen, Wang, and Pan 2019), ReSTE (Wu et al. 2023), RBNN (Lin et al. 2020), IR-Net (Qin et al. 2020). The experimental results are shown in Tab. 4. From the table, it can be seen that our method performs well, surpassing all other methods in terms of final accuracy. On CIFAR-100, it exceeds the second-best baseline by 0.54%.

Conclusion

This work investigated the generation of non-differentiable function gradients based on hypernetwork methods in BNN optimization. We re-examined the process of gradient generation from the inspiration of momentum, modeled the historical gradient sequence to generate the first-order momentum required for optimization, and introduced the FSG method, using a combination of gradient and momentum to generate the final gradient. In addition, we had added LRE to the slow-net to help generate layer specific fine gradients. The experiments on the CIFAR-10 and CIFAR-100 datasets showed that our method has faster convergence speed, lower loss values, and better performance than competitive baselines.

Acknowledgments

This work is Supported by Shanghai Artificial Intelligence Laboratory.

References

- Ajanthan, T.; Gupta, K.; Torr, P.; Hartley, R.; and Dokania, P. 2021. Mirror descent view for neural network quantization. In *International conference on artificial intelligence and statistics*, 2809–2817. PMLR.
- Andrychowicz, M.; Denil, M.; Gomez, S.; Hoffman, M. W.; Pfau, D.; Schaul, T.; Shillingford, B.; and De Freitas, N. 2016. Learning to learn by gradient descent by gradient descent. *Advances in neural information processing systems*, 29.
- Banner, R.; Hubara, I.; Hoffer, E.; and Soudry, D. 2018. Scalable methods for 8-bit training of neural networks. *Advances in neural information processing systems*, 31.
- Chen, H.; Wang, Y.; Xu, C.; Shi, B.; Xu, C.; Tian, Q.; and Xu, C. 2020. AdderNet: Do we really need multiplications in deep learning? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1468–1477.
- Chen, S.; Wang, W.; and Pan, S. J. 2019. Metaquant: Learning to quantize by learning to penetrate non-differentiable quantization. *Advances in Neural Information Processing Systems*, 32.
- Courbariaux, M.; Bengio, Y.; and David, J.-P. 2015. Binaryconnect: Training deep neural networks with binary weights during propagations. *Advances in neural information processing systems*, 28.
- Courbariaux, M.; Hubara, I.; Soudry, D.; El-Yaniv, R.; and Bengio, Y. 2016. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*.
- Denton, E. L.; Zaremba, W.; Bruna, J.; LeCun, Y.; and Fergus, R. 2014. Exploiting linear structure within convolutional networks for efficient evaluation. *Advances in neural information processing systems*, 27.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ding, X.; Hao, T.; Tan, J.; Liu, J.; Han, J.; Guo, Y.; and Ding, G. 2021. Resrep: Lossless cnn pruning via decoupling remembering and forgetting. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4510–4520.
- Everingham, M.; Eslami, S. A.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2015. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111: 98–136.
- Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 580–587.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. *arXiv preprint arXiv:2312.00752*.
- Hayashi, K.; Yamaguchi, T.; Sugawara, Y.; and Maeda, S.-i. 2019. Exploring unexplored tensor network decompositions for convolutional neural networks. *Advances in Neural Information Processing Systems*, 32.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Computation*, 9(8): 1735–1780.
- Iandola, F. N.; Han, S.; Moskewicz, M. W.; Ashraf, K.; Dally, W. J.; and Keutzer, K. 2016. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and 0.5 MB model size. *arXiv preprint arXiv:1602.07360*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *nature*, 521(7553): 436–444.
- Lin, M.; Ji, R.; Xu, Z.; Zhang, B.; Wang, Y.; Wu, Y.; Huang, F.; and Lin, C.-W. 2020. Rotated binary neural network. *Advances in neural information processing systems*, 33: 7474–7485.
- Liu, J.; Wen, D.; Wang, D.; Tao, W.; Chen, T.-W.; Osa, K.; and Kato, M. 2020. QuantNet: Learning to quantize by learning within fully differentiable framework. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, 38–53. Springer.
- Luo, J.-H.; Wu, J.; and Lin, W. 2017. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, 5058–5066.
- Qin, H.; Gong, R.; Liu, X.; Shen, M.; Wei, Z.; Yu, F.; and Song, J. 2020. Forward and backward information retention for accurate binary neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2250–2259.
- Rastegari, M.; Ordonez, V.; Redmon, J.; and Farhadi, A. 2016. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, 525–542. Springer.
- Saad, D. 1998. Online algorithms and stochastic approximations. *Online Learning*, 5(3): 6.
- Wu, X.-M.; Zheng, D.; Liu, Z.; and Zheng, W.-S. 2023. Estimator meets equilibrium perspective: A rectified straight through estimator for binary neural networks training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17055–17064.
- Xu, Y.; Han, K.; Xu, C.; Tang, Y.; Xu, C.; and Wang, Y. 2021. Learning frequency domain approximation for binary neural networks. *Advances in Neural Information Processing Systems*, 34: 25553–25565.
- Zhou, S.; Wu, Y.; Ni, Z.; Zhou, X.; Wen, H.; and Zou, Y. 2016. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*.

Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision Mamba: Efficient Visual Representation Learning with Bidirectional State Space Model. *arXiv preprint arXiv:2401.09417*.