BOLD: Boolean Logic Deep Learning

Van Minh Nguyen* Cristian Ocampo Aymen Askri Louis Leconte Ba-Hien Tran

Mathematical and Algorithmic Sciences Laboratory, Huawei Paris Research Center, France vanminh.nguyen@huawei.com

Abstract

Computational intensiveness of deep learning has motivated low-precision arithmetic designs. However, the current quantized/binarized training approaches are limited by: (1) significant performance loss due to arbitrary approximations of the latent weight gradient through its discretization/binarization function, and (2) training computational intensiveness due to the reliance on full-precision latent weights. This paper proposes a novel mathematical principle by introducing the notion of Boolean variation such that neurons made of Boolean weights and/or activations can be trained —for the first time—natively in Boolean domain instead of latent-weight gradient descent and real arithmetic. We explore its convergence, conduct extensively experimental benchmarking, and provide consistent complexity evaluation by considering chip architecture, memory hierarchy, dataflow, and arithmetic precision. Our approach achieves baseline full-precision accuracy in ImageNet classification and surpasses state-of-the-art results in semantic segmentation, with notable performance in image super-resolution, and natural language understanding with transformer-based models. Moreover, it significantly reduces energy consumption during both training and inference.

1 Introduction

Deep learning [59] has become the *de facto solution* to a wide range of tasks. However, running deep learning models for *inference* demands significant computational resources, yet it is just the tip of the iceberg. *Training* deep models is even much more intense. The extensive literature on this issue can be summarized into four approaches addressing different sources of complexity. These include: (1) model compression, pruning [38, 20, 109] and network design [94, 44, 96] for large model dimensions; (2) arithmetic approximation [58, 94, 15] for intensive multiplication; (3) quantization techniques like post-training [106, 34], quantization-aware training [37, 114, 51], and quantized training to reduce precision [16, 93]; and (4) hardware design [18, 87, 105, 101, 35, 112] to overcome computing bottleneck by moving computation closer to or in memory.

Aside from hardware and dataflow design, deep learning designs have primarily focused on the number of compute operations (OPs), such as FLOPS or BOPS, as a complexity measure [33, 83] rather than the consumed energy or memory, and particularly in inference tasks. However, it has been demonstrated that OPs alone are inadequate and even detrimental as a measure of system complexity. Instead, energy consumption provides a more consistent and efficient metric for computing hardware

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

^{*}Van Minh developed the mathematical principle, designed and implemented the Boolean deep learning framework first in C++ and later in PyTorch, validated the concept on the MNIST and CIFAR-10 benchmarks, and led all aspects of the project. Cristian evaluated the design using computer vision benchmarks. Aymen analyzed and evaluated the computational complexity. Louis proposed the incorporation of plasticity effects in the Boolean optimizer and developed the convergence analysis. Ba Hien contributed to the design and evaluation of Boolean attention mechanisms and enhanced the quality of the paper's presentation.

[94, 95, 108, 92]. Data movement, especially, dominates energy consumption and is closely linked to system architecture, memory hierarchy, and dataflow [57, 89, 110, 19]. Therefore, efforts aimed solely at reducing OPs are inefficient.

Quantization-aware training, notably binarized neural networks (BNNs) [24, 47], have garnered significant investigation [see, e.g. 83, 34, and references therein]. BNNs typically binarize weights and activations, forming principal computation blocks in binary. They learn binary weights, \mathbf{w}_{bin} , through *full-precision* (FP) *latent weights*, \mathbf{w}_{fp} , leading to no memory or computation savings during training. For example, a binarized linear layer is operated as $s = \alpha \cdot \mathbf{w}_{bin}^{\mathsf{T}} \mathbf{x}_{bin}$, where s is the output and α is a FP scaling factor, $\mathbf{w}_{bin} = \mathrm{sign}(\mathbf{w}_{fp})$, and $\mathbf{x}_{bin} = \mathrm{sign}(\mathbf{x}_{fp})$ is the binarized inputs. The weights are updated via common gradient descent backpropagation, i.e. $\mathbf{w}_{bin} = \mathrm{sign}(\mathbf{w}_{fp} - \eta \cdot \mathbf{g}_{\mathbf{w}_{fp}})$ with a learning rate η , and FP gradient signal $\mathbf{g}_{\mathbf{w}_{fp}}$. Gradient approximation of binarized variables often employs a differentiable proxy of the binarization function sign, commonly the identity proxy. Various approaches treat BNN training as a constrained optimization problem [7, 2, 3, 65], exploring methods to derive binary weights from real-valued latent ones. BNNs commonly suffers notable accuracy drops due to reduced network capacity and the use of proxy FP optimizers [61] instead of operating directly in binary domain [83, 77, 36]. Recent works mitigate this by incorporating multiple FP components in the network, retaining only a few binary dataflows [68]. Thus, while binarization aids in reducing inference complexity, it increases network training complexity and memory usage.

In contrast to binarizing FP models like BNNs, designing native binary models not relying on FP latent weight has been explored. For example, Expectation Backpropagation [91], although operating on full-precision training, was proposed for this purpose. Statistical physics-inspired [8, 9] and Belief Propagation [10] algorithms utilize integer latent weights, mainly applied to single perceptrons, with unclear applicability to deep models. Evolutionary algorithms [75, 50] are also an alternative but encounter performance and scalability challenges.

Summary: No scalable and efficient algorithm currently exists for *natively* training deep models in binary. The challenge of significantly reducing the training complexity while maintaining high performance of deep learning models remains open.

Contributions. For the aforementioned challenge, we propose a novel framework — Boolean Logic Deep Learning (B \oplus LD) — which relies on Boolean notions to define models and training:

- We introduce the notion of variation to the Boolean logic and develop a new mathematical framework of function variation (see § 3.2). One of the noticeable properties is that Boolean variation has the chain rule (see Theorem 3.12) similar to the continuous gradient.
- Based on the proposed framework, we develop a novel Boolean backpropagation and optimization method allowing for a deep model to support native Boolean components operated solely with Boolean logic and trained directly in Boolean domain, eliminating the need for gradient descent and FP latent weights (see § 3.3). This drastically cuts down memory footprint and energy consumption during *both training and inference* (see, e.g., Fig. 1).
- We provide a theoretical analysis of the convergence of our training algorithm (see Theorem 3.17).
- We conduct an extensive experimental campaign using modern network architectures such as convolutional neural networks (CNNs) and Transformers [100] on a wide range of challenging tasks including image classification, segmentation, super-resolution and natural language understanding (see § 4). We rigorously evaluate analytically the complexity of B⊕LD and BNNs. We demonstrate the superior performance of our method in terms of both accuracy and complexity compared to the state-of-the-art (see, e.g., Table 4, Table 5, Table 7).

2 Are Current Binarized Neural Networks Really Efficient?

Our work is closely related with the line of research on binarized neural networks (BNNs). The concept of BNNs traces back to early efforts to reduce the complexity of deep learning models. BINARYCONNECT [24] is one of the pioneering works that introduced the idea of binarizing FP weights during training, effectively reducing memory footprint and computational cost. Similarly, BINARYNET [47], XNOR-NET [86] extended this approach to binarize both weights and activations, further enhancing the efficiency of neural network inference. However, these early BNNs struggled with maintaining accuracy comparable to their full-precision counterparts. To address this issue,

Table 1: A summary of SOTA BNNs compared to our method. The notation X indicates the nonexistence of a specific requirement (column) within a given method (row). The colors denoting the methods shall be used consistently throughout the paper.

Method	Bitwidth (Weight-Act.)	Specialized Architecture	Mandatory FP Components	Multi-stage or KD Training	Weight Updates	Backprop	Training Arithmetic
(•) BINARYNET [47]	1-1	Х	Х	Х	FP latent-weights	Gradient	FP
(a) BINARYCONNECT [24]	1-32	X	X	X	FP latent-weights	Gradient	FP
(•) XNOR-NET [86]	1-1	X	X	X	FP latent-weights	Gradient	FP
BI-REALNET [69]	1-1	✓	✓	✓	FP latent-weights	Gradient	FP
(a) REAL2BINARY [73]	1-1	✓	✓	✓	FP latent-weights	Gradient	FP
REACTNET [68]	1-1	✓	✓	✓	FP latent-weights	Gradient	FP
(0) MELIUS-NET [11]	1-1	✓	✓	✓	FP latent-weights	Gradient	FP
(a) BNEXT [36]	1-1	✓	✓	✓	FP latent-weights	Gradient	FP
POKEBNN [117]	1-1	✓	✓	✓	FP latent-weights	Gradient	FP
(B⊕LD [Ours]	1-1	Х	Х	×	Native Boolean weights	Logic	Logic

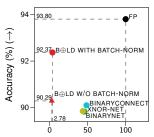
significant advances have been made on BNNs [see, e.g., 83, 85, and references therein], which can be categorized into three main aspects.

① Binarization strategy. The binarization strategy aims to efficiently convert real-valued data such as activations and weights into binary form $\{-1,1\}$. The sign function is commonly used for this purpose, sometimes with additional constraints such as clipping bounds in state-of-the-art (SOTA) methods like POKEBNN [117] or BNEXT [36]. REACTNET [68] introduces rsign as a more general alternative to the sign function, addressing potential shifts in the distribution of activations and weights. Another approach [99] explores the use of other two real values instead of strict binary to enhance the representational capability of BNNs.

② Optimization and training strategy. BNN optimization relies totally on latent-weight training, necessitating a differentiable proxy for the backpropagation. Moreover, latent-weight based training methods have to store binary and real parameters during training and often requires multiple sequential training stages, where one starts with training a FP model and only later enables binarization [36, 117, 107]. This further increases the training costs. Furthermore, knowledge distillation (KD) has emerged as a method to narrow the performance gap by transferring knowledge from a full-precision teacher model to a binary model [68, 107, 117, 60]. While a single teacher model can sufficiently improve the accuracy of the student model, recent advancements like multi-KD with multiple teachers, such as BNEXT [36], have achieved unprecedented performance. However, the KD approach often treats network binarization as an add-on to full-precision models. In addition, this training approach relies on specialized teachers for specific tasks, limiting adaptability on new data. Lastly, [103, 41] proposed some heuristics and improvements to the classic BNN latent-weight based optimizer.

(3) Architecture design. Architecture design in BNNs commonly utilizes RESNET [69, 68, 11, 36] and MOBILENET [68, 36] layouts. These methodologies often rely on heavily modified basic blocks, including additional shortcuts, automatic channel scaling via squeeze-and-excitation (SE) [117], block duplication with concatenation in the channel domain [68, 36]. Recent approaches introduce initial modules to enhance input adaptation for binary dataflow [107] or replace convolutions with lighter point-wise convolutions [66].

Another related line of research involves logic gate networks (LGNs) [79]. In LGNs, each neuron functions as a binary logic gate and, consequently, has only two inputs. Unlike traditional neural networks, LGNs do not utilize weights; instead, they are parameterized by selecting a specific logic gate for each neuron, which can be learned. Compared to the standard neural networks or our proposed method, LGNs are sparse because each neuron receives only two inputs rather than multiple inputs. Recently, [5] expanded on LGNs by incorporating flexible and differentiable lookup tables. While these advancements



Energy Consum. w.r.t. FP (%) (\leftarrow)

Figure 1: Our method against notable BNNs on CIFAR10 using VGG-SMALL. The energy is analytically evaluated considering a hypothetical V100 equivalence with native 1-bit support; cf § 4 for details.

show promises, adapting them to modern neural network architectures such as CNN or Transformers is challenging. Furthermore, these approaches have not been validated on large-scale datasets like IMAGENET or on tasks that require high precision, such as image segmentation or super-resolution, as demonstrated in our work.

Summary. Table 1 shows key characteristics of SOTA BNN methods. These methods will be considered in our experiments. Notice that all these techniques indeed have to involve operations on FP latent weights during training, whereas our proposed method works directly on native Boolean weights. In addition, most of BNN methods incorporate FP data and modules as mandatory components. As a result, existing BNNs consume much more training energy compared to our B⊕LD method. An example is shown in Fig. 1, where we consider the VGG-SMALL architecture [24, 90] on CIFAR10 dataset. In § 4 we will consider much larger datasets and networks on more challenging tasks. We can see that our method achieves $36\times$ and more than $15\times$ energy reduction compared to the FP baseline and BINARYNET, respectively, while yielding better accuracy than BNNs. Furthermore, BNNs are commonly tied to specialized network architecture and have to employ costly multi-stage or KD training. Meanwhile, our Boolean framework is completely orthogonal to these BNN methods. It is generic and applicable for a wide range of network architectures, and its training procedure purely relies on Boolean logic from scratch. Nevertheless, we stress that it is not obligatory to use all Boolean components in our proposed framework as it is flexible and can be extended to architectures comprised of a mix of Boolean and FP modules. This feature further improves the superior performance of our method as can be seen in Fig. 1, where we integrate batch normalization (BN) [49] into our Boolean model, and we will demonstrate extensively in our experiments.

3 Proposed Method

3.1 Neuron Design

Boolean Neuron. For the sake of simplicity, we consider a linear layer for presenting the design. Let $w_0, (w_1, \ldots, w_m)$, and (x_1, \ldots, x_m) be the bias, weights, and inputs of a neuron of input size $m \ge 1$. In the core use case of our interest, these variables are all Boolean numbers. Let L be a logic gate such as AND, OR, XOR, XNOR. The neuron's pre-activation output is given as follows:

$$s = w_0 + \sum_{i=1}^{m} L(w_i, x_i), \tag{1}$$

where the summation is understood as the counting of TRUES.

Mixed Boolean-Real Neuron. To allow for flexible use and co-existence of this Boolean design with real-valued parts of a deep model, two cases of mixed-type data are considered including Boolean weights with real-valued inputs, and real-valued weights with Boolean inputs. These two cases can be addressed by the following extension of Boolean logic to mixed-type data. To this end, we first introduce the essential notations and definitions. Specifically, we denote $\mathbb{B} := \{T, F\}$ equipped with the Boolean logic. Here, T and F indicate TRUE and FALSE, respectively.

Definition 3.1 (Three-valued logic). Define $\mathbb{M} \stackrel{\mathrm{def}}{=} \mathbb{B} \cup \{0\}$ with logic connectives defined according to those of Boolean logic as follows. First, the negation is: $\neg T = F$, $\neg F = T$, and $\neg 0 = 0$. Second, let L be a logic connective, denote by $L_{\mathbb{M}}$ and $L_{\mathbb{B}}$ when it is in \mathbb{M} and in \mathbb{B} , respectively, then $L_{\mathbb{M}}(a,b) = L_{\mathbb{B}}(a,b)$ for $a,b \in \mathbb{B}$ and $L_{\mathbb{M}}(a,b) = 0$ otherwise.

Notation 3.2. Denote by \mathbb{L} a logic set (e.g., \mathbb{B} or \mathbb{M}), \mathbb{R} the real set, \mathbb{Z} the set of integers, \mathbb{N} a numeric set (e.g., \mathbb{R} or \mathbb{Z}), and \mathbb{D} a certain set of \mathbb{L} or \mathbb{N} .

Definition 3.3. For $x \in \mathbb{N}$, its logic value denoted by x_{logic} is given as $x_{\text{logic}} = T \Leftrightarrow x > 0$, $x_{\text{logic}} = F \Leftrightarrow x < 0$, and $x_{\text{logic}} = 0 \Leftrightarrow x = 0$.

Definition 3.4. The magnitude of a variable x, denoted |x|, is defined as its usual absolute value if $x \in \mathbb{N}$. And for $x \in \mathbb{L}$: |x| = 0 if x = 0, and |x| = 1 otherwise.

Definition 3.5 (Mixed-type logic). For L a logic connective of L and variables a, b, operation c = L(a,b) is defined such that |c| = |a||b| and $c_{\text{logic}} = L(a_{\text{logic}},b_{\text{logic}})$.

Using Definition 3.5, neuron formulation Eq. 1 directly applies to the mixed Boolean-real neurons.

Forward Activation. It is clear that there can only be one unique family of binary activation functions, which is the threshold function. Let τ be a scalar, which can be fixed or learned, the forward Boolean activation is given as: y = T if $s \ge \tau$ and y = F otherwise where s is the preactivation. The backpropagation throughout this activation will be described in Appendix C.3.

3.2 Mathematical Foundation

In this section we describe the mathematical foundation for our method to train Boolean weights directly in the Boolean domain without relying on FP latent weights. Due to the space limitation, essential notions necessary for presenting the main results are presented here while a comprehensive treatment is provided in Appendix A.

Definition 3.6. Order relations '<' and '>' in \mathbb{B} are defined as follows: F < T, and T > F.

Definition 3.7. For $a, b \in \mathbb{B}$, the variation from a to b, denoted $\delta(a \to b)$, is defined as: $\delta(a \to b) \stackrel{\text{def}}{=} T$ if b > a, $\stackrel{\text{def}}{=} 0$ if b = a, and $\stackrel{\text{def}}{=} F$ if b < a.

Throughout the paper, $\mathcal{F}(\mathbb{S}, \mathbb{T})$ denotes the set of all functions from source \mathbb{S} to image \mathbb{T} .

Definition 3.8. For $f \in \mathcal{F}(\mathbb{B}, \mathbb{D})$, $\forall x \in \mathbb{B}$, write $\delta f(x \to \neg x) := \delta(f(x) \to f(\neg x))$. The variation of f w.r.t. x, denoted f'(x), is defined as: $f'(x) \stackrel{\text{def}}{=} \mathbf{xnor}(\delta(x \to \neg x), \delta f(x \to \neg x))$.

Remark 3.9. The usual notation of continuous derivative f' is intentionally adopted here for Boolean variation for convenience and notation unification. Its underlying meaning, i.e., continuous derivative or Boolean variation, can be understood directly from the context where function f is defined.

Intuitively, the variation of f w.r.t. x is T if f varies in the same direction with x.

Example 3.10. Let $a \in \mathbb{B}$, $f(x) = \mathbf{xor}(x, a)$ for $x \in \mathbb{B}$, the variation of f w.r.t. x can be derived by establishing a truth table (see Table 8 in Appendix A.1) from which we obtain $f'(x) = \neg a$.

For $f \in \mathcal{F}(\mathbb{Z}, \mathbb{N})$, its derivative, also known in terms of *finite differences*, has been defined in the literature as f'(x) = f(x+1) - f(x), see e.g. [52]. With the logic variation as introduced above, we can make this definition more generic as follows.

Definition 3.11. For $f \in \mathcal{F}(\mathbb{Z}, \mathbb{D})$, the variation of f w.r.t. $x \in \mathbb{Z}$ is defined as $f'(x) \stackrel{\text{def}}{=} \delta f(x \to x+1)$, where δf is in the sense of the variation defined in \mathbb{D} .

Theorem 3.12. The following properties hold:

- 1. For $f \in \mathcal{F}(\mathbb{B}, \mathbb{B})$: $(\neg f)'(x) = \neg f'(x), \forall x \in \mathbb{B}$.
- 2. For $f \in \mathcal{F}(\mathbb{B}, \mathbb{N})$, $\alpha \in \mathbb{N}$: $(\alpha f)'(x) = \alpha f'(x)$, $\forall x \in \mathbb{B}$.
- 3. For $f, g \in \mathcal{F}(\mathbb{B}, \mathbb{N})$: $(f+g)'(x) = f'(x) + g'(x), \forall x \in \mathbb{B}$.
- 4. For $\mathbb{B} \xrightarrow{f} \mathbb{B} \xrightarrow{g} \mathbb{D}$: $(g \circ f)'(x) = \mathbf{xnor}(g'(f(x)), f'(x)), \forall x \in \mathbb{B}$.
- 5. For $\mathbb{B} \xrightarrow{f} \mathbb{Z} \xrightarrow{g} \mathbb{D}$, $x \in \mathbb{B}$, if $|f'(x)| \le 1$ and g'(f(x)) = g'(f(x) 1), then: $(g \circ f)'(x) = \mathbf{xnor}(g'(f(x)), f'(x))$.

The proof is provided in Appendix A.1. These results are extended to the multivariate case in a straightforward manner. For instance, for multivariate Boolean functions it is as follows.

Definition 3.13. For $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{B}^n$, denote $\mathbf{x}_{\neg i} := (x_1, \dots, x_{i-1}, \neg x_i, x_{i+1}, \dots, x_n)$ for $n \ge 1$ and $1 \le i \le n$. For $f \in \mathcal{F}(\mathbb{B}^n, \mathbb{B})$, the (partial) variation of f w.r.t. x_i , denoted $f'_i(\mathbf{x})$ or $\delta f(\mathbf{x})/\delta x_i$, is defined as: $f'_i(\mathbf{x}) \equiv \delta f(\mathbf{x})/\delta x_i \stackrel{\text{def}}{=} \mathbf{xnor}(\delta(x_i \to \neg x_i), \delta f(\mathbf{x} \to \mathbf{x}_{\neg i}))$.

Proposition 3.14. Let $f \in \mathcal{F}(\mathbb{B}^n, \mathbb{B})$, $n \geq 1$, and $g \in \mathcal{F}(\mathbb{B}, \mathbb{B})$. For $1 \leq i \leq n$:

$$(g \circ f)_i'(\mathbf{x}) = \mathbf{xnor}(g'(f(\mathbf{x})), f_i'(\mathbf{x})), \quad \forall \mathbf{x} \in \mathbb{B}^n.$$
 (2)

Example 3.15. From Example 3.10, we have $\delta \mathbf{xor}(x, a)/\delta x = \neg a$ for $a, x \in \mathbb{B}$. Using Theorem 3.12-(1) we have: $\delta \mathbf{xnor}(x, a)/\delta x = a$ since $\mathbf{xnor}(x, a) = \neg \mathbf{xor}(x, a)$.

Example 3.16. Apply Theorem 3.12-(3) to s from Eq. 1: $\delta s/\delta w_i = \delta L(w_i, x_i)/\delta w_i$ and $\delta s/\delta x_i = \delta L(w_i, x_i)/\delta x_i$. Then, for $L = \mathbf{xnor}$ as an example, we have: $\delta s/\delta w_i = x_i$ and $\delta s/\delta x_i = w_i$.

3.3 BackPropagation

With the notions introduced in § 3.2, we can write signals involved in the backpropagation process as shown in Fig. 2. Therein, layer l is a Boolean layer of consideration. For the sake of presentation

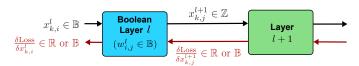


Figure 2: Illustration of backpropagation signals with a Boolean linear layer. Notice that the subsequent layer can be any FP/Boolean layers or activation functions.

simplicity, layer l is assumed a fully-connected layer, and:

$$x_{k,j}^{l+1} = w_{0,j}^l + \sum_{i=1}^m \mathcal{L}(x_{k,i}^l, w_{i,j}^l), \quad 1 \le j \le n,$$
(3)

where L is the utilized Boolean logic, k denotes sample index in the batch, m and n are the usual layer input and output sizes. Layer l is connected to layer l+1 that can be an activation layer, a batch normalization, an arithmetic layer, or any others. The nature of $\delta \text{Loss}/\delta x_{k,j}^{l+1}$ depends on the property of layer l+1. It can be the usual gradient if layer l+1 is a real-valued input layer, or a Boolean variation if layer l+1 is a Boolean-input layer. Given $\delta \text{Loss}/\delta x_{k,j}^{l+1}$, layer l needs to optimize its Boolean weights and compute signal $\delta \text{Loss}/\delta x_{k,i}^{l}$ for the upstream. Hereafter, we consider $L = \mathbf{xnor}$ when showing concrete illustrations of the method.

Atomic Variation. First, using Theorem 3.12 and its extension to the multivariate case by Proposition 3.14 in the same manner as shown in Example 3.16, we have:

$$\frac{\delta x_{k,j}^{l+1}}{\delta w_{i,j}^{l}} = \frac{\delta \mathcal{L}(x_{k,i}^{l}, w_{i,j}^{l})}{\delta w_{i,j}^{l}} \stackrel{\mathbf{L}=\mathbf{xnor}}{=} x_{k,i}^{l}, \quad \frac{\delta x_{k,j}^{l+1}}{\delta x_{k,i}^{l}} = \frac{\delta \mathcal{L}(x_{k,i}^{l}, w_{i,j}^{l})}{\delta x_{k,i}^{l}} \stackrel{\mathbf{L}=\mathbf{xnor}}{=} w_{i,j}^{l}. \tag{4}$$

Using the chain rules given by Theorem 3.12–(4 & 5), we have:

$$q_{i,j,k}^{l} := \frac{\delta \text{Loss}}{\delta w_{i,j}^{l}}|_{k} = \mathbf{xnor}(\frac{\delta \text{Loss}}{\delta x_{k,j}^{l+1}}, \frac{\delta x_{k,j}^{l+1}}{\delta w_{i,j}^{l}}) \stackrel{\mathbf{L}=\mathbf{xnor}}{=} \mathbf{xnor}(\frac{\delta \text{Loss}}{\delta x_{k,j}^{l+1}}, x_{k,i}^{l}), \tag{5}$$

$$g_{k,i,j}^{l} := \frac{\delta \text{Loss}}{\delta x_{k,i}^{l}}|_{j} = \mathbf{xnor}(\frac{\delta \text{Loss}}{\delta x_{k,j}^{l+1}}, \frac{\delta x_{k,j}^{l+1}}{\delta x_{k,i}^{l}}) \stackrel{\mathbf{L}=\mathbf{xnor}}{=} \mathbf{xnor}(\frac{\delta \text{Loss}}{\delta x_{k,j}^{l+1}}, w_{i,j}^{l}). \tag{6}$$

Aggregation. Atomic variation $q_{i,j,k}^l$ is aggregated over batch dimension k while $g_{k,i,j}^l$ is aggregated over output dimension j. Let $\mathbf{1}(\cdot)$ be the indicator function. For $b \in \mathbb{B}$ and variable x, define: $\mathbf{1}(x=b)=1$ if $x_{\mathrm{logic}}=b$ and $\mathbf{1}(x=b)=0$ otherwise. Atomic variations are aggregated as:

$$q_{i,j}^{l} := \frac{\delta \text{Loss}}{\delta w_{i,j}^{l}} = \sum_{k} \mathbf{1} (q_{i,j,k}^{l} = T) |q_{i,j,k}^{l}| - \sum_{k} \mathbf{1} (q_{i,j,k}^{l} = F) |q_{i,j,k}^{l}|,$$
(7)

$$g_{k,i}^{l} := \frac{\delta \text{Loss}}{\delta x_{k,i}^{l}} = \sum_{j} \mathbf{1} (g_{k,i,j}^{l} = \mathbf{T}) |g_{k,i,j}^{l}| - \sum_{j} \mathbf{1} (g_{k,i,j}^{l} = \mathbf{F}) |g_{k,i,j}^{l}|.$$
(8)

Boolean Optimizer. With $q_{i,j}^l$ obtained in Eq. 7, the rule for optimizing $w_{i,j}^l$ subjected to making the loss decreased is simply given according to its definition as:

$$w_{i,j}^l = \neg w_{i,j}^l \text{ if } \mathbf{xnor}(q_{i,j}^l, w_{i,j}^l) = \mathbf{T}.$$

$$(9)$$

Eq. 9 is the core optimization logic based on which more sophisticated forms of optimizer can be developed in the same manner as different methods such as Adam, Adaptive Adam, etc. have been developed from the basic gradient descent principle. For instance, the following is an optimizer that accumulates $q_{i,j}^l$ over training iterations. Denote by $q_{i,j}^{l,t}$ the optimization signal at iteration t, and by $m_{i,j}^{l,t}$ its accumulator with $m_{i,j}^{l,0} := 0$ and:

$$m_{i,j}^{l,t+1} = \beta^t m_{i,j}^{l,t} + \eta^t q_{i,j}^{l,t+1}, \tag{10}$$

where η^t is an accumulation factor that can be tuned as a hyper-parameter, and β^t is an autoregularizing factor that expresses the system's state at time t. Its usage is linked to brain plasticity [31] and Hebbian theory [40] forcing weights to adapt to their neighborhood. For the chosen weight's neighborhood, for instance, neuron, layer, or network level, β^t is given as:

$$\beta^t = \frac{\text{Number of unchanged weights at } t}{\text{Total number of weights}}.$$
 (11)

In the experiments presented later, β^t is set to per-layer basis. Finally, the learning process is as described in Algorithm 1. We encourage the readers to check the detailed implementations, practical considerations, and example codes of our proposed method, available in Appendix B and Appendix C.

Convergence Analysis. The following result describes how the iterative logic optimization based on Eq. 9 minimizes a predefined loss f, under the standard non-convex assumption. The technical assumptions and the proof are given in Appendix A.2.

Theorem 3.17. *Under the specified assumptions, Boolean optimization logic converges at:*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[\|\nabla f(w_t)\|^2 \right] \le \frac{A^*}{T\eta} + B^* \eta + C^* \eta^2 + Lr_d,$$
(12)

where $A^*=2(f(w_0)-f_*)$ with f_* being uniform lower bound assumed exists, $B^*=2L\sigma^2$, $C^*=4L^2\sigma^2\frac{\gamma}{(1-\gamma)^2}$ in which L is the assumed Lipschitz constant, and $r_d=d\kappa/2$.

The bound in Theorem 3.17 contains four terms. The first is typical for a general non-convex target and expresses how initialization affects the convergence. The second and third terms depend on the fluctuation of the minibatch gradients. There is an "error bound" of $2Ld\kappa$ independent of T. This error bound is the cost of using discrete weights as part of the optimization algorithm. Previous work with quantized models also includes such error bounds [61, 62].

```
Algorithm 1: Illustration with a FC layer.
      Input : Learning rate \eta, nb iterations T;
Initialize: m_{i,j}^{1,0} = 0; \ \beta^0 = 1;
1 for t = 0, \dots, T-1 do

| /* 1. Forward */
               Compute x^{l+1,t} following Eq. 3;
               /* 2. Backward */
              Receive \frac{\delta \text{Loss}}{\delta x_{k,j}^{l+1,t}} from downstream layer;
               /* 2.1 Backpropagation */
               Compute and backpropagate g^{l,t} of Eq. 8;
               /* 2.2 Weight update process */
               N_{\text{tot}} := 0, N_{\text{unchanged}} := 0;
              Update m_{i,j}^{l,t+1} = \beta^t m_{i,j}^{l,t} + \eta^t q_{i,j}^{l,t+1}; N_{\text{tot}} \leftarrow N_{\text{tot}} + 1; if \mathbf{xnor}(m_{i,j}^{l,t+1}, w_{i,j}^{l,t}) = \mathbf{T} then \begin{vmatrix} w_{i,j}^{l,t+1} = \neg w_{i,j}^{l,t} / * \text{ invert } */\\ m_{i,j}^{l,t+1} = 0; \end{vmatrix} else
 9
               \begin{vmatrix} w_{i,j}^{l,t+1} = w_{i,j}^{l,t} \;; & /* \text{ keep } */\\ N_{\text{unchanged}} \leftarrow N_{\text{unchanged}} + 1; \\ \text{Update } \eta^{t+1}, \beta^{t+1} = N_{\text{unchanged}}/N_{\text{tot}} \;; \\ \end{vmatrix} 
14
15
```

Regularization. Exploding and vanishing gradients are two well-known issues when it comes to train deep neural networks. During Boolean Logic training, our preliminary experiments indicated that the backpropagation signal also experiences similar phenomenon, resulting in unstable training. Our idea is to scale the backpropagation signals so as to match their variance. Thanks to the Boolean structure, assumptions on the involved signals can be made so that the scaling factor can be analytically derived in closed-form without need of learning or batch statistics computation. Precisely, through linear layers of output size m, the backpropagation signal is scaled with $\sqrt{2/m}$, and for convolutional layers of output size c_{out} , stride v and kernel sizes k_x , k_y , the backpropagation signal is scaled with $2\sqrt{2v/c_{\text{out}}k_xk_y}$ if maxpooling is applied, or $\sqrt{2v/c_{\text{out}}k_xk_y}$ otherwise. The full detail is presented in Appendix C.

4 Experiments

Our B⊕LD method achieves extreme compression by using both Boolean activations and weights. We rigorously evaluate its performance on challenging precision-demanding tasks, including *image* classification on CIFAR10 [55] and IMAGENET [56], as well as super-resolution on five popular datasets. Furthermore, recognizing the necessity of deploying efficient lightweight models for edge computing, we delve into three fine-tuning scenarios, showcasing the adaptability of our approach. Specifically, we investigate fine-tuning Boolean models for image classification on CIFAR10 and CIFAR100 using VGG-SMALL. For segmentation tasks, we study DEEPLABV3 [17] fine-tuned on

CITYSCAPES [23] and PASCAL VOC 2012 [29] datasets. The backbone for such a model is our Boolean RESNET18 [39] network trained from scratch on IMAGENET. Finally, we consider an evaluation in the domain of *natural language understanding*, fine-tuning BERT [26], a transformer-based [100] language model, on the GLUE benchmark [102].

Experimental Setup. To construct our B⊕LD models, we introduce Boolean weights and activations and substitute full-precision (FP) arithmetic layers with Boolean equivalents. Throughout all benchmarks, we maintain the general network design of the chosen FP baseline, while excluding FP-specific components like ReLU, PReLU activations, or batch normalization (BN) [49], unless specified otherwise. Consistent with the common setup in existing literature [see, e.g., 86, 21, 11], only the first and last layers remain in FP and are optimized using an Adam optimizer [54]. Comprehensive experiment details for reproducibility are provided in Appendix D.

Complexity Evaluation. It has been demonstrated that relying solely on FLOPS and BOPS for complexity assessment is inadequate [94, 95, 108, 92]. In fact, these metrics fail to capture the actual load caused by propagating FP data through the BNNs. Instead, energy consumption serves as a crucial indicator of efficiency. Given the absence of native Boolean accelerators, we estimate analytically energy consumption by analyzing the arithmetic operations, data movements within storage/processing units, and the energy cost of each operation. This approach is implemented for the Nvidia GPU (Tesla V100) and Ascend [63] architectures. Further details are available in Appendix E.

4.1 Image Classification

Our B⊕LD method is tested on two network configurations: *small & compact* and *large & deep*. In the former scenario, we utilize the VGG-SMALL [90] baseline trained on CIFAR10. Evaluation of our Boolean architecture is conducted both without BN [24], and with BN including activation from [69]. These designs achieve 90.29±0.09% (estimated over six repetitions) and 92.37±0.01% (estimated over five repetitions) accuracy, respectively (see Table 2). Notably, without BN, our results align closely with BINARYCONNECT [24], which employs 32-bit activations during both inference and training. Furthermore, BN brings the accuracy within 1 point of the FP baseline. Additional results are provided in the supplementary material for VGG-SMALL models ending with 1 FC layer.

Our method requires much less energy than the FP baseline. In particular, it consumes less than 5% of energy for our designs with and without BN respectively. These results highlight the remarkable energy efficiency of our B⊕LD method in both inference and training, surpassing latent-weight based training methods [24, 86, 48] reliant on FP weights. Notably, despite a slight increase in energy consumption, utilizing BN yields superior accuracy. Even with BN, our approach maintains superior efficiency compared to alternative methods, further emphasizing the flexibility of our approach in training networks with a blend of Boolean and FP components.

In the *large & deep* case, we consider the RESNET18 baseline trained from scratch on IMAGENET. We compare our approach to methods employing the same baseline, larger architectures, and additional training strategies such as KD with a RESNET34 teacher or FP-based shortcuts [69]. Our method consistently achieves the highest accuracy across all categories, ranging from the standard model (51.8% accuracy) to larger configurations (70.0% accuracy), as shown in Table 5. Additionally, our B \oplus LD method exhibits the smallest energy consumption in most categories, with a remarkable 24.45% for our large architecture with and without KD. Notably, our method outperforms the FP baseline when using $4\times$ filter enlargement (base 256), providing significant energy reduction (24.45%). Furthermore, it surpasses the SOTA POKEBNN [117], utilizing RESNET50 as a teacher.

For completeness, we also implemented neural gradient quantization, utilizing INT4 quantization with a logarithmic round-to-nearest approach [21] and statistics-aware weight binning [22]. Our experiments on IMAGENET confirm that 4-bit quantization is sufficient to achieve standard FP performances, reaching 67.53% accuracy in 100 epochs (further details provided in Appendix D.1.4).

4.2 Image Super-resolution

Next, we evaluate the efficacy of our B⊕LD method to synthesize data. We use a compact EDSR network [64] as our baseline, referred to as SMALL EDSR, comprising eight residual blocks. Our

²VGG-SMALL with the last FP layer mapping to 100 classes.

Table 2: Results with VGG-SMALL on CIFAR 10. 'Cons.' is the energy consumption w.r.t. the FP baseline, evaluated on 1 training iteration.

Method	W/A	Acc.(%)	Cons.(%) Ascend	Cons. (%) Tesla V100
Full-precision [114]	32/32	93.80	100.00	100.00
(BINARYCONNECT [24]	1/32	90.10	38.59	48.49
(*) XNOR-NET [86]	1/1	89.83	34.21	45.68
(•) BINARYNET [47]	1/1	89.85	32.60	43.61
(B⊕LD w/o BN [Ours]	1/1	90.29	3.64	2.78
(●) B⊕LD with BN [Ours]	1/1	92.37	4.87	3.71

Table 3: Super-resolution results measured in PSNR (dB) (\(^\)), using the EDSR baseline [64].

Task	Method	SET5 [12]	SET14 [113]	BSD100 [46]	URBAN100 [72]	DIV2K [1, 97]
×2	FULL EDSR (FP) SMALL EDSR (FP) (B B LD [Ours]	38.11 38.01 37.42	33.92 33.63 33.00	32.32 32.19 31.75	32.93 31.60 30.26	35.03 34.67 33.82
×3	FULL EDSR (FP) SMALL EDSR (FP) (a) B⊕LD [Ours]	34.65 34.37 33.56	30.52 30.24 29.70	29.25 29.10 28.72	==	31.26 30.93 30.22
×4	FULL EDSR (FP) SMALL EDSR (FP) (•) B⊕LD [Ours]	32.46 32.17 31.23	28.80 28.53 27.97	27.71 27.62 27.24	26.64 26.14 25.12	29.25 29.04 28.36

Table 4: Image segmentation results.

Dataset	Model	mIoU (%) (†)		
	FP BASELINE	70.7		
CITYSCAPES	BINARY DAD-NET [30]	58.1		
	(●) B⊕LD [Ours]	67.4		
	FP BASELINE	72.1		
PASCAL VOC 2012	(●) B⊕LD [Ours]	67.3		

Table 5: Results with RESNET18 baseline on IMAGENET. 'Base' is the mapping dimension of the first layer. Energy consumption is evaluated on 1 training iteration. 'Cons.' is the energy consumption w.r.t. the FP baseline.

Training Modality	Method	Acc. (%)	Cons. (%) Ascend	Cons. (%) Tesla V100
FP BASELINE	RESNET18 [39]	69.7	100.00	100.00
	(a) BINARYNET [47] (b) XNOR-NET [86] (c) B⊕LD + BN (Base 64)	42.2 51.2 51.8	 8.77	 3.87
FP SHORTCUT	(a) BI-REALNET: 18 [69]	56.4	46.60	32.99
LARGE MODELS	(®) BI-REALNET:34 [69] (®) BI-REALNET:152 [69] (®) MELIUS-NET:29 [11] (®) B⊕LD (Base 256)	62.2 64.5 65.8 66.9	80.00 38.82	58.24 24.45
KD: RESNET34	(®) REAL2BINARY [73] (®) REACTNET-RESNET18 [68] (®) BNEXT:18 [36] (®) B@LD + BN (Base 192)) (®) B@LD (Base 256)	65.4 65.5 68.4 65.9 70.0	45.43 45.48 26.91 38.82	77.89 37.51 16.91 24.45
KD: RESNET50	() POKEBNN-RESNET18 [117]	65.2		

Table 6: Results with VGG-SMALL baseline finetuned on CIFAR10 and CIFAR100. 'FT' means 'Fine-Tuning'.

Ref.	Method	Model Init.	Train./FT Dataset	Bitwidth W/A/G	Acc. (%)
A	FP BASELINE	Random	CIFAR10	32/32/32	95.27
В	FP BASELINE 2	Random	CIFAR 100	32/32/32	77.27
C	(■) B⊕LD	Random	CIFAR 10	1/1/16	90.29
D	(■) B⊕LD ²	Random	CIFAR 100	1/1/16	68.43
Е	FP BASELINE ²	A	CIFAR 100	32/32/32	76.74
F	(●) B⊕LD ²	C	CIFAR 100	1/1/16	68.37
G	FP BASELINE	В	CIFAR 10	32/32/32	95.77
H	(■) B⊕LD	D	CIFAR 10	1/1/16	92.09

B⊕LD model employs Boolean residual blocks without BN. Results, presented in Table 3, based on the official implementation and benchmark³, reveal remarkable similarity to the FP reference at each scale. Particularly noteworthy are the prominent results achieved on SET14 and BSD100 datasets. Our method consistently delivers high PSNR for high-resolution images, such as DIV2K, and even higher for low-resolution ones, like SET5. However, akin to EDSR, our approach exhibits a moderate performance reduction at scale $4\times$. These findings highlight the capability of our method to perform adequately on detail-demanding tasks while exhibiting considerable robustness across image resolutions.

4.3 Adaptability on New Data

Image classification fine-tuning. We aim to assess the adaptability of our method to similar problems but different datasets, a common scenario for edge inference tasks. We employ the VGG-SMALL architecture without BN under two training configurations. Firstly, the B⊕LD model is trained from scratch with random initialization on CIFAR10 (REF. C) and CIFAR100 (REF. D). Secondly, we fine-tune the trained networks on CIFAR100 (REF. F) and CIFAR10 (REF. H), respectively. Notably, in Table 6, fine-tuning our trained model on CIFAR100 (REF. F) results in a model almost identical to the model trained entirely from scratch (REF. D). Additionally, a noteworthy result is observed with our model (REF. H), which achieves higher accuracy than the model trained from scratch (REF. C).

Image segmentation fine-tuning. Next, we expand the scope of the aforementioned fine-tuning experiment to encompass a larger network and a different task. The baseline is the DEEPLABV3 network for semantic segmentation. It consists of our Boolean RESNET18 (without BN) as the backbone, followed by the Boolean atrous pyramid pooling (ASPP) module [17]. We refrain from utilizing auxiliary loss or knowledge distillation techniques, as these methods introduce additional computational burdens, which are contrary to our objective of efficient on-device training. As demonstrated in Table 4, our method achieves a notable 67.4% mIoU on CITYSCAPES (see Fig. 3 for prediction examples). This result surpasses the SOTA, BINARY DAD-NET [30], and approaches

³https://github.com/sanghyun-son/EDSR-PyTorch

the performance of the FP baseline. Likewise, on PASCAL VOC 2012, our methodology nears the performance of the FP baseline. Importantly, these improvements are attained without the intermediate use of FP parameters during training, highlighting the efficiency and effectiveness of our approach. This shows that our method not only preserves the inherent lightweight advantages of highly quantized neural networks but also significantly enhances performance in complex segmentation tasks.

BERT fine-tuning for NLU tasks. Finally, we consider fine-tuning BERT [26], a transformer-based model [100], on the GLUE benchmark [102]. We follow the standard experimental protocol as in [26, 6, 82]. Our model and the chosen baselines are employed with 1-bit bitwidth for both weights and activations. Our Boolean BERT model is inspired by BIT [67] for binarizing activations and incorporating KD during training, where the FP teacher guides the student in a layer-wise manner. We follow the experimental setup of BIT, including using the same method for binarizing activations and backpropagation for softmax and attention in the BERT model. As shown in Table 7, all methods suffer from performance drop compared to the FP model as extreme binarization of transformer-based model is not trivial. Nevertheless, our method yields results comparable to BIT [67], the SOTA method on this task, outperforming BINARYBERT [6] and BIBERT [82] on average. This is remarkable as our method natively uses Boolean weights during the training, whereas the baselines heavily rely on FP latent weights. These findings indicate potential for energy-efficient large language models (LLMs) using our method for both training and inference.

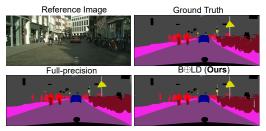


Figure 3: An example of CITYSCAPES.

				4		
Table 7.	RFRT	models	results	Source	code [67]	

Method	GLUE Benchmark (Accuracy, †)								
	MNLI	QQP	QNLI	sst-2	COLA	SST-B	MRPC	RTE	Avg.
FP BERT	84.9	91.4	92.1	93.2	59.7	90.1	86.3	72.2	83.9
BINARYBERT	35.6	66.2	51.5	53.2	0.0	6.1	68.3	52.7	41.0
BIBERT	66.1	84.8	72.6	88.7	25.4	33.6	72.5	57.4	63.2
BIT	77.1	82.9	85.7	87.7	25.1	71.1	79.7	58.8	71.0
BIT (Reprod.†)	76.8	87.2	85.6	87.5	24.1	70.5	78.9	58.8	69.7
(●) B⊕LD	75.6	85.9	84.1	88.7	27.1	68.7	78.4	58.8	70.9

5 Conclusions

We introduced the notion of Boolean variation and developed a first framework of its calculus. This novel mathematical principle enabled the development of Boolean logic backpropagation and Boolean optimization replacing gradient backpropagation and gradient descent for binary deep learning. Deep models can be built with native Boolean weights and/or Boolean activations, and trained in Boolean natively by this principled exact Boolean optimization. That brings a key advantage to the existing popular quantized/binarized training approach that suffers from critical bottlenecks – (i) performance loss due to an arbitrary approximation of the latent weight gradient through its discretization/binarization function, (ii) training computational intensiveness due to full-precision latent weights. We have extensively explored its capabilities, highlighting: (i) both training and inference are now possible in binary; (iv) deep training complexity can be drastically reduced to unprecedented levels. (iii) Boolean models can handle finer tasks beyond classification, contrary to common belief; (ii) in some applications, suitablely enlarging Boolean model can recover FP performance while still gaining significant complexity reduction.

Limitations. Due to current computing accelerators, such as GPUs, primarily designed for real arithmetic, our method could not be assessed on native Boolean accelerator. Nevertheless, its considerable potential may inspire the development of new logic circuits and architectures utilizing Boolean logic processing. It also remains as an open question the approximation capacity of Boolean neural networks. A mathematical result equivent to the existing universal approximation theory of real-valued neural networks would provide a solid guarantee.

Acknowledgments

This work has been made possible through the invaluable support of various departments and colleagues within Huawei. Van Minh Nguyen would like to extend his sincere thanks to everyone involved, with special appreciation to Jean-Claude Belfiore and the former students who have participated in this project over the years: Valentin Abadie, Tom Huix, Tianyu Li, and Youssef Chaabouni.

References

- E. Agustsson and R. Timofte. NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, July 2017.
- [2] T. Ajanthan, P. K. Dokania, R. Hartley, and P. H. Torr. Proximal Mean-field for Neural Network Quantization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [3] T. Ajanthan, K. Gupta, P. Torr, R. Hartley, and P. Dokania. Mirror Descent View for Neural Network Quantization. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130, pages 2809–2817. PMLR, 13–15 Apr 2021.
- [4] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic. QSGD: Communication-Efficient SGD via Gradient Quantization and Encoding. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017.
- [5] A. T. L. Bacellar, Z. Susskind, M. Breternitz Jr, E. John, L. K. John, P. M. V. Lima, and F. M. França. Differentiable weightless neural networks. In R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 2277–2295. PMLR, 21–27 Jul 2024.
- [6] H. Bai, W. Zhang, L. Hou, L. Shang, J. Jin, X. Jiang, Q. Liu, M. Lyu, and I. King. BinaryBERT: Pushing the Limit of BERT Quantization. In C. Zong, F. Xia, W. Li, and R. Navigli, editors, *Proceedings of the* 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 4334–4348. Association for Computational Linguistics, 2021.
- [7] Y. Bai, Y.-X. Wang, and E. Liberty. ProxQuant: Quantized Neural Networks via Proximal Operators. In *International Conference on Learning Representations*, 2018.
- [8] C. Baldassi. Generalization Learning in a Perceptron with Binary Synapses. *Journal of Statistical Physics*, 136(5):902–916, Sep 2009.
- [9] C. Baldassi and A. Braunstein. A Max-Sum Algorithm for Training Discrete Neural Networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(8):P08008, 2015.
- [10] C. Baldassi, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina. Subdominant Dense Clusters Allow for Simple Learning and High Computational Performance in Neural Networks with Discrete Synapses. *Physical Review Letters*, 115(12):128101, 2015.
- [11] J. Bethge, C. Bartz, H. Yang, Y. Chen, and C. Meinel. MeliusNet: An Improved Network Architecture for Binary Neural Networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1439–1448, January 2021.
- [12] M. Bevilacqua, A. Roumy, C. Guillemot, and M. line Alberi Morel. Low-Complexity Single-Image Super-Resolution based on Nonnegative Neighbor Embedding. In *Proceedings of the British Machine Vision Conference*, pages 135.1–135.10. BMVA Press, 2012.
- [13] S. Bianco, R. Cadene, L. Celona, and P. Napoletano. Benchmark Analysis of Representative Deep Neural Network Architectures. *IEEE Access*, 6:64270–64277, 2018.
- [14] A. Canziani, A. Paszke, and E. Culurciello. An Analysis of Deep Neural Network Models for Practical Applications. arXiv preprint arXiv:1605.07678, 2016.
- [15] H. Chen, Y. Wang, C. Xu, B. Shi, C. Xu, Q. Tian, and C. Xu. AdderNet: Do We Really Need Multiplications in Deep Learning? In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), June 2020.

- [16] J. Chen, Y. Gai, Z. Yao, M. W. Mahoney, and J. E. Gonzalez. A Statistical Framework for Low-bitwidth Training of Deep Neural Networks. In *Advances in Neural Information Processing Systems*, volume 33, pages 883–894. Curran Associates, Inc., 2020.
- [17] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam. Rethinking Atrous Convolution for Semantic Image Segmentation. arXiv preprint arXiv:1706.05587, 2017.
- [18] Y.-H. Chen, J. Emer, and V. Sze. Eyeriss: A Spatial Architecture for Energy-Efficient Dataflow for Convolutional Neural Networks. SIGARCH Computer Architecture News, 44(3):367–379, jun 2016.
- [19] Y.-H. Chen, T. Krishna, J. S. Emer, and V. Sze. Eyeriss: An Energy-efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks. *IEEE Journal of Solid-state Circuits*, 52(1):127–138, 2016.
- [20] Y. Cheng, D. Wang, P. Zhou, and T. Zhang. Model Compression and Acceleration for Deep Neural Networks: The Principles, Progress, and Challenges. *IEEE Signal Processing Magazine*, 35(1):126–136, 2018
- [21] B. Chmiel, R. Banner, E. Hoffer, H. B. Yaacov, and D. Soudry. Logarithmic Unbiased Quantization: Simple 4-bit Training in Deep Learning. arXiv:2112.10769, 2021.
- [22] J. Choi, P. I.-J. Chuang, Z. Wang, S. Venkataramani, V. Srinivasan, and K. Gopalakrishnan. Bridging the Accuracy Gap for 2-Bit Quantized Neural Networks (QNN). arXiv preprint arXiv:1807.06964, 2018.
- [23] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [24] M. Courbariaux, Y. Bengio, and J.-P. David. BinaryConnect: Training Deep Neural Networks with Binary Weights during Propagations. In Advances in Neural Information Processing Systems, volume 28. Curran Associates, Inc., 2015.
- [25] C. De Sa, M. Leszczynski, J. Zhang, A. Marzoev, C. R. Aberger, K. Olukotun, and C. Ré. High-Accuracy Low-Precision Training. arXiv preprint arXiv:1803.03383, 2018.
- [26] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171–4186. Association for Computational Linguistics, 2019.
- [27] R. Ding, T.-W. Chin, Z. Liu, and D. Marculescu. Regularizing Activation Distribution for Training Binarized Deep Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [28] Z. Du, R. Fasthuber, T. Chen, P. Ienne, L. Li, T. Luo, X. Feng, Y. Chen, and O. Temam. ShiDianNao: Shifting Vision Processing Closer to the Sensor. In *Proceedings of the 42nd Annual International Symposium on Computer Architecture*, pages 92–104, 2015.
- [29] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascalnetwork.org/challenges/VOC/voc2012/workshop/index.html.
- [30] A. Frickenstein, M.-R. Vemparala, J. Mayr, N.-S. Nagaraja, C. Unger, F. Tombari, and W. Stechele. Binary DAD-Net: Binarized Driveable Area Detection Network for Autonomous Driving. In 2020 IEEE International Conference on Robotics and Automation (ICRA), pages 2295–2301. IEEE, 2020.
- [31] E. Fuchs, G. Flügge, et al. Adult Neuroplasticity: More than 40 Years of Research. *Neural plasticity*, 2014, 2014.
- [32] Y. Gao, Y. Liu, H. Zhang, Z. Li, Y. Zhu, H. Lin, and M. Yang. Estimating GPU Memory Consumption of Deep Learning Models. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1342–1352, 2020.
- [33] E. García-Martín, C. F. Rodrigues, G. Riley, and H. Grahn. Estimation of Energy Consumption in Machine Learning. *Journal of Parallel and Distributed Computing*, 134:75–88, 2019.
- [34] A. Gholami, S. Kim, Z. Dong, Z. Yao, M. W. Mahoney, and K. Keutzer. A Survey of Quantization Methods for Efficient Neural Network Inference. In *Low-Power Computer Vision*, pages 291–326. Chapman and Hall/CRC, 2022.

- [35] C. Grimm and N. Verma. Neural Network Training on In-Memory-Computing Hardware With Radix-4 Gradients. IEEE Transactions on Circuits and Systems I: Regular Papers I, 69(10):4056–4068, 2022.
- [36] N. Guo, J. Bethge, C. Meinel, and H. Yang. Join the High Accuracy Club on ImageNet with A Binary Neural Network Ticket. arXiv preprint arXiv:2211.12933, 2022.
- [37] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan. Deep Learning with Limited Numerical Precision. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 1737–1746, Lille, France, 07–09 Jul 2015. PMLR.
- [38] S. Han, H. Mao, and W. J. Dally. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. In *International Conference on Learning Representations*, 2015.
- [39] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2016.
- [40] D. O. Hebb. The Organization of Behavior: A Neuropsychological Theory. Psychology press, 2005.
- [41] K. Helwegen, J. Widdicombe, L. Geiger, Z. Liu, K.-T. Cheng, and R. Nusselder. Latent Weights Do Not Exist: Rethinking Binarized Neural Network Optimization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [42] M. Horowitz. 1.1 Computing's Energy Problem (and What We Can Do about It). In 2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC), pages 10–14, 2014.
- [43] L. Hou, Q. Yao, and J. T. Kwok. Loss-aware Binarization of Deep Networks. In *International Conference on Learning Representations*, 2016.
- [44] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. arXiv preprint arXiv:1704.04861, 2017.
- [45] L. Hoyer, D. Dai, and L. Van Gool. DAFormer: Improving Network Architectures and Training Strategies for Domain-Adaptive Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9924–9935, 2022.
- [46] J.-B. Huang, A. Singh, and N. Ahuja. Single Image Super-Resolution from Transformed Self-Exemplars. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2015.
- [47] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio. Binarized Neural Networks. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [48] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio. Quantized Neural Networks: Training Neural Networks with Low Precision Weights and Activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- [49] S. Ioffe and C. Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR.
- [50] R. Ito and T. Saito. Dynamic Binary Neural Networks and Evolutionary Learning. In The 2010 International Joint Conference on Neural Networks, pages 1–5. IEEE, 2010.
- [51] Q. Jin, J. Ren, R. Zhuang, S. Hanumante, Z. Li, Z. Chen, Y. Wang, K. Yang, and S. Tulyakov. F8Net: Fixed-Point 8-bit Only Multiplication for Network Quantization. In *International Conference on Learning Representations*, 2021.
- [52] C. Jordan. Calculus of Finite Differences. Chelsea Publishing Company, New York, 2nd edition, 1950.
- [53] S. P. Karimireddy, Q. Rebjock, S. Stich, and M. Jaggi. Error Feedback Fixes SignSGD and other Gradient Compression Schemes. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 3252–3261. PMLR, 09–15 Jun 2019.
- [54] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.
- [55] A. Krizhevsky and G. Hinton. Learning Multiple Layers of Features from Tiny Images. *Master's thesis, Department of Computer Science, University of Toronto*, 2009.

- [56] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In Advances in Neural Information Processing Systems, volume 25. Curran Associates, Inc., 2012.
- [57] H. Kwon, P. Chatarasi, M. Pellauer, A. Parashar, V. Sarkar, and T. Krishna. Understanding Reuse, Performance, and Hardware Cost of DNN Dataflow: A Data-centric Approach. In *Proceedings of the* 52nd Annual IEEE/ACM International Symposium on Microarchitecture, pages 754–768, 2019.
- [58] I. Lavi, S. Avidan, Y. Singer, and Y. Hel-Or. Proximity Preserving Binary Code Using Signed Graph-Cut. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 34, pages 4535–4544, April 2020.
- [59] Y. LeCun, Y. Bengio, and G. Hinton. Deep Learning. *Nature*, 521(7553):436–444, 2015.
- [60] C. Lee, H. Kim, E. Park, and J.-J. Kim. INSTA-BNN: Binary Neural Network with INSTAnce-aware Threshold. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17325–17334, October 2023.
- [61] H. Li, S. De, Z. Xu, C. Studer, H. Samet, and T. Goldstein. Training Quantized Nets: A Deeper Understanding. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017.
- [62] Z. Li and C. M. De Sa. Dimension-Free Bounds for Low-Precision Training. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [63] H. Liao, J. Tu, J. Xia, H. Liu, X. Zhou, H. Yuan, and Y. Hu. Ascend: a Scalable and Unified Architecture for Ubiquitous Deep Neural Network Computing: Industry Track Paper. In 2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA), pages 789–801, 2021.
- [64] B. Lim, S. Son, H. Kim, S. Nah, and K. Mu Lee. Enhanced Deep Residual Networks for Single Image Super-Resolution. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2017.
- [65] M. Lin, R. Ji, Z. Xu, B. Zhang, Y. Wang, Y. Wu, F. Huang, and C.-W. Lin. Rotated Binary Neural Network. In Advances in Neural Information Processing Systems, volume 33, pages 7474–7485. Curran Associates, Inc., 2020.
- [66] C. Liu, W. Ding, P. Chen, B. Zhuang, Y. Wang, Y. Zhao, B. Zhang, and Y. Han. RB-Net: Training Highly Accurate and Efficient Binary Neural Networks with Reshaped Point-wise Convolution and Balanced activation. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(9):6414–6424, 2022.
- [67] Z. Liu, B. Oguz, A. Pappu, L. Xiao, S. Yih, M. Li, R. Krishnamoorthi, and Y. Mehdad. BiT: Robustly Binarized Multi-distilled Transformer. In *Advances in Neural Information Processing Systems*, volume 35, pages 14303–14316. Curran Associates, Inc., 2022.
- [68] Z. Liu, Z. Shen, M. Savvides, and K.-T. Cheng. ReActNet: Towards Precise Binary Neural Network with Generalized Activation Functions. In *Proceedings of the European Conference on Computer Vision* (ECCV), August 2020.
- [69] Z. Liu, B. Wu, W. Luo, X. Yang, W. Liu, and K.-T. Cheng. Bi-Real Net: Enhancing the Performance of 1-bit CNNs with Improved Representational Capability and Advanced Training Algorithm. In *Proceedings* of the European Conference on Computer Vision (ECCV), September 2018.
- [70] J. Long, E. Shelhamer, and T. Darrell. Fully Convolutional Networks for Semantic Segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2015.
- [71] I. Loshchilov and F. Hutter. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*, 2017.
- [72] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision (ICCV), July 2001.
- [73] B. Martinez, J. Yang, A. Bulat, and G. Tzimiropoulos. Training Binary Neural Networks with Real-tobinary Convolutions. In *International Conference on Learning Representations*, 2020.
- [74] X. Mei, K. Zhao, C. Liu, and X. Chu. Benchmarking the Memory Hierarchy of Modern GPUs. In *IFIP International Conference on Network and Parallel Computing*, pages 144–156. Springer, 2014.

- [75] G. Morse and K. O. Stanley. Simple Evolutionary Optimization Can Rival Stochastic Gradient Descent in Neural Networks. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, pages 477–484, 2016.
- [76] V. M. Nguyen. Boolean Variation and Boolean Logic BackPropagation. arXiv:2311.07427, May 2024.
- [77] G. Nie, L. Xiao, M. Zhu, D. Chu, Y. Shen, P. Li, K. Yang, L. Du, and B. Chen. Binary Neural Networks as a General-propose Compute Paradigm for On-device Computer Vision. arXiv preprint arXiv:2202.03716, 2022.
- [78] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019.
- [79] F. Petersen, C. Borgelt, H. Kuehne, and O. Deussen. Deep differentiable logic gate networks. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 2006–2018. Curran Associates, Inc., 2022.
- [80] C. Philippenko and A. Dieuleveut. Bidirectional Compression in Heterogeneous Settings for Distributed or Federated Learning with Partial Participation: Tight Convergence Guarantees. arXiv preprint arXiv:2006.14591, 2020.
- [81] B. T. Polyak. Introduction to Optimization. Translations Series in Mathematics and Engineering. Optimization Software Inc. Publication Division, New York, 1987.
- [82] H. Qin, Y. Ding, M. Zhang, Q. Yan, A. Liu, Q. Dang, Z. Liu, and X. Liu. BiBERT: Accurate Fully Binarized BERT. In *International Conference on Learning Representations*, 2022.
- [83] H. Qin, R. Gong, X. Liu, X. Bai, J. Song, and N. Sebe. Binary Neural Networks: A Survey. Pattern Recognition, 105:107281, 2020.
- [84] H. Qin, R. Gong, X. Liu, M. Shen, Z. Wei, F. Yu, and J. Song. Forward and Backward Information Retention for Accurate Binary Neural Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [85] H. Qin, M. Zhang, Y. Ding, A. Li, Z. Cai, Z. Liu, F. Yu, and X. Liu. BiBench: Benchmarking and Analyzing Network Binarization. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 28351–28388. PMLR, 23–29 Jul 2023.
- [86] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks. In *Proceedings of the European Conference on Computer Vision* (ECCV), October 2016.
- [87] A. Sebastian, M. Le Gallo, R. Khaddam-Aljameh, and E. Eleftheriou. Memory Devices and Applications for in-memory Computing. *Nature Nanotechnology*, 15(7):529–544, 2020.
- [88] Y. S. Shao and D. Brooks. Energy Characterization and Instruction-Level Energy Model of Intels Xeon Phi Processor. In *International Symposium on Low Power Electronics and Design (ISLPED)*, pages 389–394, 2013.
- [89] J. Sim, S. Lee, and L.-S. Kim. An Energy-Efficient Deep Convolutional Neural Network Inference Processor With Enhanced Output Stationary Dataflow in 65-nm CMOS. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 28(1):87–100, 2019.
- [90] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In International Conference on Learning Representations, 2015.
- [91] D. Soudry, I. Hubara, and R. Meir. Expectation Backpropagation: Parameter-Free Training of Multilayer Neural Networks with Continuous or Discrete Weights. In Advances in Neural Information Processing Systems, volume 27. Curran Associates, Inc., 2014.
- [92] E. Strubell, A. Ganesh, and A. McCallum. Energy and Policy Considerations for Deep Learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy, Jul 2019. Association for Computational Linguistics.

- [93] X. Sun, N. Wang, C.-Y. Chen, J. Ni, A. Agrawal, X. Cui, S. Venkataramani, K. El Maghraoui, V. V. Srinivasan, and K. Gopalakrishnan. Ultra-Low Precision 4-bit Training of Deep Neural Networks. In Advances in Neural Information Processing Systems, volume 33, pages 1796–1807. Curran Associates, Inc., 2020.
- [94] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer. Efficient Processing of Deep Neural Networks: A Tutorial and Survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.
- [95] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer. How to Evaluate Deep Neural Network Processors: Tops/w (Alone) Considered Harmful. *IEEE Solid-State Circuits Magazine*, 12(3):28–41, 2020.
- [96] M. Tan and Q. Le. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR, 09–15 Jun 2019.
- [97] R. Timofte, E. Agustsson, L. Van Gool, M.-H. Yang, L. Zhang, B. Lim, et al. NTIRE 2017 Challenge on Single Image Super-Resolution: Methods and Results. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.
- [98] H. Touvron, A. Vedaldi, M. Douze, and H. Jegou. Fixing the Train-Test Resolution Discrepancy. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [99] Z. Tu, X. Chen, P. Ren, and Y. Wang. AdaBin: Improving Binary Neural Networks with Adaptive Binary Sets. In *Proceedings of the European Conference on Computer Vision (ECCV)*, October 2022.
- [100] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [101] N. Verma, H. Jia, H. Valavi, Y. Tang, M. Ozatay, L.-Y. Chen, B. Zhang, and P. Deaville. In-Memory Computing: Advances and Prospects. *IEEE Solid-State Circuits Magazine*, 11(3):43–55, 2019.
- [102] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *International Conference on Learning Representations*, 2019.
- [103] E. Wang, J. J. Davis, D. Moro, P. Zielinski, J. J. Lim, C. Coelho, S. Chatterjee, P. Y. Cheung, and G. A. Constantinides. Enabling Binary Neural Network Training on the Edge. In *Proceedings of the 5th International Workshop on Embedded and Mobile Deep Learning*, pages 37–38, 2021.
- [104] J. Wangni, J. Wang, J. Liu, and T. Zhang. Gradient sparsification for communication-efficient distributed optimization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 31. Curran Associates, Inc., 2018.
- [105] S. Williams, A. Waterman, and D. Patterson. Roofline: An Insightful Visual Performance Model for Multicore Architectures. *Communications of the ACM*, 52(4):65–76, apr 2009.
- [106] G. Xiao, J. Lin, M. Seznec, H. Wu, J. Demouth, and S. Han. SmoothQuant: Accurate and Efficient Post-Training Quantization for Large Language Models. In *Proceedings of the 40th International Conference* on Machine Learning, volume 202, pages 38087–38099. PMLR, 23–29 Jul 2023.
- [107] X. Xing, Y. Li, W. Li, W. Ding, Y. Jiang, Y. Wang, J. Shao, C. Liu, and X. Liu. Towards Accurate Binary Neural Networks via Modeling Contextual Dependencies. In *Proceedings of the European Conference on Computer Vision (ECCV)*, October 2022.
- [108] T.-J. Yang, Y.-H. Chen, J. Emer, and V. Sze. A Method to Estimate the Energy Consumption of Deep Neural Networks. In 2017 51st Asilomar Conference on Signals, Systems, and Computers, pages 1916–1920, October 2017.
- [109] T.-J. Yang, Y.-H. Chen, and V. Sze. Designing Energy-Efficient Convolutional Neural Networks Using Energy-Aware Pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [110] X. Yang, M. Gao, Q. Liu, J. Setter, J. Pu, A. Nayak, S. Bell, K. Cao, H. Ha, P. Raina, et al. Interstellar: Using Halide's Scheduling Language to Analyze DNN Accelerators. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 369–383, 2020.

- [111] Z. Yang, Y. Wang, K. Han, C. XU, C. Xu, D. Tao, and C. Xu. Searching for Low-Bit Weights in Quantized Neural Networks. In Advances in Neural Information Processing Systems, volume 33, pages 4091–4102. Curran Associates, Inc., 2020.
- [112] S. Yu and P.-Y. Chen. Emerging Memory Technologies: Recent Trends and Prospects. *IEEE Solid-State Circuits Magazine*, 8(2):43–56, 2016.
- [113] R. Zeyde, M. Elad, and M. Protter. On Single Image Scale-up using Sparse-Representations. In Curves and Surfaces: 7th International Conference, Avignon, France, June 24-30, 2010, Revised Selected Papers 7, pages 711–730. Springer, 2012.
- [114] D. Zhang, J. Yang, D. Ye, and G. Hua. LQ-Nets: Learned Quantization for Highly Accurate and Compact Deep Neural Networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [115] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. Mixup: Beyond Empirical Risk Minimization. In International Conference on Learning Representations, 2018.
- [116] R. Zhang, A. G. Wilson, and C. De Sa. Low-Precision Stochastic Gradient Langevin Dynamics. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pages 26624–26644. PMLR, 17–23 Jul 2022.
- [117] Y. Zhang, Z. Zhang, and L. Lew. PokeBNN: A Binary Pursuit of Lightweight Accuracy. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 12475–12485, June 2022.
- [118] Z. Zhang. Derivation of Backpropagation in Convolutional Neural Network (CNN). University of Tennessee, Knoxville, TN, 22:23, 2016.
- [119] B. Zhuang, C. Shen, M. Tan, L. Liu, and I. Reid. Structured Binary Neural Networks for Accurate Image Classification and Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

Appendix

Table of Contents

A	A. 1	Boolean Variation Method Boolean Variation Calculus	18 18 22
В	Cod	e Sample of Core Implementation	26
C	Trai	ning Regularization	28
	C .1	Assumptions and Notations	28
	C.2	Backpropagation Scaling	29
	C.3	BackPropagation Through Boolean Activation Function	31
D	Exp	erimental Details	31
	D.1	Image Classification	31
	D.2	Image Super-resolution	34
	D.3	Semantic Segmentation	34
	D.4	Boolean BERT Fine-tuning	40
E	Ene	rgy Estimation	40
	E.1	Hardware Specification	40
	E.2	Compute Energy	41
	E.3	Memory Energy	41
F	Broa	nder Impacts	44

A Details of the Boolean Variation Method

A.1 Boolean Variation Calculus

This section provides supplementary details of the Boolean Variation method, a sole development can be found in [76]. With reference to Example 3.10, Table 8 gives the variation truth table of $f(x) = \mathbf{xor}(a, x)$, for $a, x \in \mathbb{B}$.

Table 8: Variation truth table of $f(x) = \mathbf{xor}(a, x), a, x \in \mathbb{B}$.

a	x	$\neg x$	$\delta(x \to \neg x)$	f(a,x)	$f(a, \neg x)$	$\delta f(x \to \neg x)$	f'(x)
Т	Т	F	F	F	Т	T	F
${ m T}$	\mathbf{F}	\mathbf{T}	${ m T}$	${ m T}$	\mathbf{F}	\mathbf{F}	\mathbf{F}
\mathbf{F}	\mathbf{T}	F	\mathbf{F}	${ m T}$	\mathbf{F}	\mathbf{F}	${ m T}$
F	\mathbf{F}	Τ	${ m T}$	\mathbf{F}	${ m T}$	${ m T}$	${ m T}$

We now proceed to the proof of Theorem 3.12.

Definition A.1 (Type conversion). *Define:*

$$p: \mathbb{N} \to \mathbb{L}$$

$$x \mapsto p(x) = \begin{cases} T, & \text{if } x > 0, \\ 0, & \text{if } x = 0, \\ F, & \text{if } x < 0. \end{cases}$$
 (13)

$$e: \mathbb{L} \to \mathbb{N}$$

$$a \mapsto e(a) = \begin{cases} +1, & \text{if } a = T, \\ 0, & \text{if } a = 0, \\ -1, & \text{if } a = F. \end{cases}$$
 (14)

p projects a numeric type in logic, and e embeds a logic type in numeric. The following properties are straightforward:

Proposition A.2. *The following properties hold:*

- 1. $\forall x, y \in \mathbb{N}$: $p(xy) = \mathbf{xnor}(p(x), p(y))$.
- 2. $\forall a, b \in \mathbb{L}$: $e(\mathbf{xnor}(a, b)) = e(a) e(b)$.
- 3. $\forall x, y \in \mathbb{N}$: $x = y \Leftrightarrow |x| = |y|$ and p(x) = p(y).

In particular, property Proposition A.2(2) implies that by the embedding map $e(\cdot)$, we have:

$$(\{T, F\}, \mathbf{xor}) \cong (\{\pm 1\}, -\times), \tag{15}$$

$$(\{T, F\}, \mathbf{xnor}) \cong (\{\pm 1\}, \times), \tag{16}$$

where \cong and \times stand for isomorphic relation, and the real multiplication, resp. A consequence is that by $e(\cdot)$, a computing sequence of pointwise XOR/XNOR, counting, and majority vote is equivalent to a sequence of pointwise multiplications and accumulation performed on the embedded data. This property will be used in Appendices A.2 and C for studying Boolean method using some results from BNNs literature and real analysis.

Proposition A.3. The following properties hold:

- 1. $a \in \mathbb{L}$, $x \in \mathbb{N}$: $\mathbf{xnor}(a, x) = e(a)x$.
- 2. $x, y \in \mathbb{N}$: **xnor**(x, y) = xy.
- 3. $x \in \{\mathbb{L}, \mathbb{N}\}, y, z \in \mathbb{N}$: $\mathbf{xnor}(x, y + z) = \mathbf{xnor}(x, y) + \mathbf{xnor}(x, z)$.
- 4. $x \in \{\mathbb{L}, \mathbb{N}\}, y, \lambda \in \mathbb{N}$: $\mathbf{xnor}(x, \lambda y) = \lambda \mathbf{xnor}(x, y)$.
- 5. $x \in \{\mathbb{L}, \mathbb{N}\}, y \in \mathbb{N}$: $\mathbf{xor}(x, y) = -\mathbf{xnor}(x, y)$.

Proof. The proof follows definitions 3.5 and A.1.

- Following Definition 3.1 we have $\forall t \in \mathbb{M}$, $\mathbf{xnor}(T,t) = t$, $\mathbf{xnor}(F,t) = \neg t$, and $\mathbf{xnor}(0,t) = 0$. Put $v = \mathbf{xnor}(a,x)$. We have |v| = |x| and $p(v) = \mathbf{xnor}(a,p(x))$. Hence, $a = 0 \Rightarrow p(v) = 0 \Rightarrow v = 0$; $a = T \Rightarrow p(v) = p(x) \Rightarrow v = x$; $a = F \Rightarrow p(v) = \neg p(x) \Rightarrow v = -x$. Hence (1).
- The result is trivial if x = 0 or y = 0. For $x, y \neq 0$, put $v = \mathbf{xnor}(x, y)$, we have |v| = |x||y| and $p(v) = \mathbf{xnor}(p(x), p(y))$. According to Definition A.1, if $\mathrm{sign}(x) = \mathrm{sign}(y)$, we have $p(v) = T \Rightarrow v = |x||y| = xy$. Otherwise, i.e., $\mathrm{sign}(x) = -\mathrm{sign}(y)$, $p(v) = F \Rightarrow v = -|x||y| = xy$. Hence (2).
- (3) and (4) follow (1) for $x \in \mathbb{L}$ and follow (2) for $x \in \mathbb{N}$.
- For (5), write $u = \mathbf{xor}(x, y)$ and $v = \mathbf{xnor}(x, y)$, we have |u| = |v| and $p(u) = \mathbf{xor}(p(x), p(y)) = \neg \mathbf{xnor}(p(x), p(y)) = \neg p(v)$. Thus, $\operatorname{sign}(u) = -\operatorname{sign}(v) \Rightarrow u = -v$.

61930

Proposition A.4. For $f, g \in \mathcal{F}(\mathbb{B}, \mathbb{B})$, $\forall x, y \in \mathbb{B}$ the following properties hold:

1.
$$\delta f(x \to y) = \mathbf{xnor}(\delta(x \to y), f'(x)).$$

2.
$$(\neg f)'(x) = \neg f'(x)$$
.

3.
$$(g \circ f)'(x) = \mathbf{xnor}(g'(f(x)), f'(x)).$$

Proof. The proof is by definition:

1. $\forall x, y \in \mathbb{B}$, there are two cases. If y = x, then the result is trivial. Otherwise, i.e., $y = \neg x$, by definition we have:

$$f'(x) = \mathbf{xnor}(\delta(x \to \neg x), \delta f(x \to \neg x))$$

$$\Leftrightarrow \quad \delta f(x \to \neg x) = \mathbf{xnor}(\delta(x \to \neg x), f'(x)).$$

Hence the result.

2. $\forall x, y \in \mathbb{B}$, it is easy to verify by truth table that $\delta(\neg f(x \to y)) = \neg \delta f(x \to y)$. Hence, by definition.

$$(\neg f)'(x) = \mathbf{xnor}(\delta(x \to \neg x), \delta(\neg f(x \to \neg x)))$$

$$= \mathbf{xnor}(\delta(x \to \neg x), \neg \delta f(x \to \neg x))$$

$$= \neg \mathbf{xnor}(\delta(x \to \neg x), \delta f(x \to \neg x))$$

$$= \neg f'(x).$$

3. Using definition, property (i), and associativity of **xnor**, $\forall x \in \mathbb{B}$ we have:

$$\begin{split} (g \circ f)'(x) &= \mathbf{xnor}(\delta(x \to \neg x), \delta g(f(x) \to f(\neg x))) \\ &= \mathbf{xnor}(\delta(x \to \neg x), \mathbf{xnor}(\delta f(x \to \neg x), g'(f(x)))) \\ &= \mathbf{xnor}(g'(f(x)), \mathbf{xnor}(\delta(x \to \neg x), \delta f(x \to \neg x))) \\ &= \mathbf{xnor}(g'(f(x)), f'(x)). \end{split}$$

Proposition A.5. For $f \in \mathcal{F}(\mathbb{B}, \mathbb{N})$, the following properties hold:

1.
$$x, y \in \mathbb{B}$$
: $\delta f(x \to y) = \mathbf{xnor}(\delta(x \to y), f'(x))$.

2.
$$x \in \mathbb{B}$$
, $\alpha \in \mathbb{N}$: $(\alpha f)'(x) = \alpha f'(x)$.

3.
$$x \in \mathbb{B}, g \in \mathcal{F}(\mathbb{B}, \mathbb{N})$$
: $(f+g)'(x) = f'(x) + g'(x)$.

Proof. The proof is as follows:

1. For $x, y \in \mathbb{B}$. Firstly, the result is trivial if y = x. For $y \neq x$, i.e., $y = \neg x$, by definition:

$$f'(x) = \mathbf{xnor}(\delta(x \to \neg x), \delta f(x \to \neg x)).$$
 Hence, $|\delta f(x \to \neg x)| = |f'(x)|$ since $|\delta(x \to \neg x)| = 1$, and
$$\mathbf{p}(f'(x)) = \mathbf{xnor}(\delta(x \to \neg x), \mathbf{p}(\delta f(x \to \neg x)))$$
 $\Leftrightarrow \quad \mathbf{p}(\delta f(x \to \neg x)) = \mathbf{xnor}(\delta(x \to \neg x), \mathbf{p}(f'(x))),$

where $p(\cdot)$ is the logic projector Eq. 13. Thus, $\delta f(x \to \neg x) = \mathbf{xnor}(\delta(x \to \neg x), f'(x))$. Hence the result.

2. Firstly $\forall x, y \in \mathbb{B}$, we have

$$\delta(\alpha f(x \to y)) = \alpha f(y) - \alpha f(x) = \alpha \delta f(x \to y).$$

Hence, by definition,

$$(\alpha f)'(x) = \mathbf{xnor}(\delta(x \to \neg x), \delta(\alpha f(x \to \neg x)))$$

$$= \mathbf{xnor}(\delta(x \to \neg x), \alpha \delta f(x \to \neg x))$$

$$= \alpha \mathbf{xnor}(\delta(x \to \neg x), \delta f(x \to \neg x)), \text{ due to Proposition A.3(4)}$$

$$= \alpha f'(x).$$

3. For $f, g \in \mathcal{F}(\mathbb{B}, \mathbb{N})$,

$$\begin{split} (f+g)'(x) &= \mathbf{xnor}(\delta(x \to \neg x), \delta(f+g)(x \to \neg x)) \\ &= \mathbf{xnor}(\delta(x \to \neg x), \delta f(x \to \neg x) + \delta g(x \to \neg x)) \\ &\stackrel{(*)}{=} \mathbf{xnor}(\delta(x \to \neg x), \delta f(x \to \neg x)) + \mathbf{xnor}(\delta(x \to \neg x), \delta g(x \to \neg x)), \\ &= f'(x) + g'(x), \end{split}$$

where (*) is due to Proposition A.3(3).

Proposition A.6 (Composition rules). The following properties hold:

1. For
$$\mathbb{B} \xrightarrow{f} \mathbb{B} \xrightarrow{g} \mathbb{D}$$
: $(g \circ f)'(x) = \mathbf{xnor}(g'(f(x)), f'(x)), \forall x \in \mathbb{B}$.

2. For
$$\mathbb{B} \xrightarrow{f} \mathbb{Z} \xrightarrow{g} \mathbb{D}$$
, $x \in \mathbb{B}$, if $|f'(x)| \le 1$ and $g'(f(x)) = g'(f(x) - 1)$, then: $(g \circ f)'(x) = \mathbf{xnor}(g'(f(x)), f'(x))$.

Proof. The proof is as follows.

- 1. The case of $\mathbb{B} \xrightarrow{f} \mathbb{B} \xrightarrow{g} \mathbb{B}$ is obtained from Proposition A.4(3). For $\mathbb{B} \xrightarrow{f} \mathbb{B} \xrightarrow{g} \mathbb{N}$, by using Proposition A.5(1), the proof is similar to that of Proposition A.4(3).
- 2. By definition, we have

$$(g \circ f)'(x) = \mathbf{xnor}(\delta(x \to \neg x), \delta g(f(x) \to f(\neg x))). \tag{17}$$

Using property (1) of Proposition A.5, we have:

$$f(\neg x) = f(x) + \delta f(x \to \neg x)$$

= $f(x) + \mathbf{xnor}(\delta(x \to \neg x), f'(x)).$ (18)

Applying Eq. 18 back to Eq. 17, the result is trivial if f'(x) = 0. The remaining case is |f'(x)| = 1 for which we have $\mathbf{xnor}(\delta(x \to \neg x), f'(x)) = \pm 1$. First, for $\mathbf{xnor}(\delta(x \to \neg x), f'(x)) = 1$, we have:

$$\delta g(f(x) \to f(\neg x)) = \delta g(f(x) \to f(x) + 1)$$

$$= g'(f(x))$$

$$= \mathbf{xnor}(g'(f(x)), 1)$$

$$= \mathbf{xnor}(g'(f(x)), \mathbf{xnor}(\delta(x \to \neg x), f'(x))). \tag{19}$$

Substitute Eq. 19 back to Eq. 17, we obtain:

$$\begin{split} (g \circ f)'(x) &= \mathbf{xnor}(\delta(x \to \neg x), \delta g(f(x) \to f(\neg x))) \\ &= \mathbf{xnor}(\delta(x \to \neg x), \mathbf{xnor}(g'(f(x)), \mathbf{xnor}(\delta(x \to \neg x), f'(x)))) \\ &= \mathbf{xnor}(g'(f(x)), f'(x)), \end{split}$$

where that last equality is by the associativity of **xnor** and that **xnor**(x, x) = T for $x \in \mathbb{B}$. Similarly, for **xnor** $(\delta(x \to \neg x), f'(x)) = -1$, we have:

$$\delta g(f(x) \to f(\neg x)) = \delta g(f(x) \to f(x) - 1)$$

$$= -g'(f(x) - 1)$$

$$= \mathbf{xnor}(g'(f(x) - 1), -1)$$

$$= \mathbf{xnor}(g'(f(x) - 1), \mathbf{xnor}(\delta(x \to \neg x), f'(x))). \tag{20}$$

Substitute Eq. 20 back to Eq. 17 and use the assumption that g'(f(x)) = g'(f(x) - 1), we have:

$$(g \circ f)'(x) = \mathbf{xnor}(\delta(x \to \neg x), \delta g(f(x) \to f(\neg x)))$$

= $\mathbf{xnor}(\delta(x \to \neg x), \mathbf{xnor}(g'(f(x) - 1), \mathbf{xnor}(\delta(x \to \neg x), f'(x))))$
= $\mathbf{xnor}(g'(f(x)), f'(x)).$

Hence the result.

The results of Theorem 3.12 are from Propositions A.4 to A.6.

П

A.2 Convergence Proof

To the best of our knowledge, in the quantized neural network literature and in particular BNN, one can only prove the convergence up to an irreducible error floor [61]. This idea has been extended to SVRG [25], and recently to SGLD in [116], which is also up to an error limit.

In this section we provide complexity bounds for Boolean Logic in a smooth non-convex environment. We introduce an abstraction to model its optimization process and prove its convergence.

A.2.1 Continuous Abstraction

Boolean optimizer is discrete, proving its convergence directly is a hard problem. The idea is to find a continuous equivalence so that some proof techniques existing from the BNN and quantized neural networks literature can be employed. We also abstract the logic optimization rules as compressor $Q_0(), Q_1()$, and define gradient accumulator a_t as $a_{t+1} = a_t + \varphi(q_t)$. When η is constant, we recover the definition (10) and obtain $m_t = \eta a_t$. Our analysis is based on the following standard non-convex assumptions on f:

- **A. 1.** Uniform Lower Bound: There exists $f_* \in \mathbb{R}$ s.t. $f(w) \geq f_*$, $\forall w \in \mathbb{R}^d$.
- **A. 2.** Smooth Derivatives: The gradient $\nabla f(w)$ is L-Lipschitz continuous for some L > 0, i.e., $\forall w, \forall v \in \mathbb{R}^d \colon \|\nabla f(w) \nabla f(v)\| \le L\|w v\|$.
- **A. 3.** Bounded Variance: The variance of the stochastic gradients is bounded by some $\sigma^2 > 0$, i.e., $\forall w \in \mathbb{R}^d \colon \mathbb{E}\left[\tilde{\nabla}f(w)\right] = \nabla f(w) \text{ and } \mathbb{E}\left[\|\tilde{\nabla}f(w)\|^2\right] \leq \sigma^2$.
- **A. 4.** Compressor: There exists $\gamma < 1$ s.t. $\forall w, \forall v \in \mathbb{R}^d$, $||Q_1(v, w) v||^2 \le \gamma ||v||^2$.
- **A. 5.** Bounded Accumulator: There exists $\kappa \in \mathbb{R}_+^*$ s.t. $\forall t$ and $\forall i \in [d]$, we have $|a_t|_i \leq \kappa$.
- **A. 6.** Stochastic Flipping Rule: For all $w \in \mathbb{R}^d$, we have $\mathbb{E}[Q_0(w)|w] = w$.

In existing frameworks, quantity $\tilde{\nabla} f(\cdot)$ denotes the stochastic gradient computed on a random minibatch of data. Boolean framework does not have the notion of gradient, it however has an optimization signal (i.e., $\varphi(\mathbf{q})$ or its accumulator m_t (10)) that plays the same role as $\tilde{\nabla} f(\cdot)$. Therefore, these two notions, i.e., continuous gradient and Boolean optimization signal, can be encompassed into a generalized notion. That is the root to the following continuous relaxation in which $\tilde{\nabla} f(\cdot)$ standards for the optimization signal computed on a random mini-batch of data.

Within this continuous relaxation framework, A. 6 expresses our assumption that the flipping rule is stochastic and unbiased. Note that this assumption is standard in the literature related to (stochastic) quantization, see e.g., [4, 81, 104].

For reference, the original Boolean optimizer as formulated in § 3 is summarized in Algorithm 2 in which $flip(w_t, m_{t+1})$ flips weight and $reset(w_t, m_{t+1})$ resets its accumulator when the flip condition is triggered.

Algorithm 2: Boolean optimizer

```
1 m_{t+1} \leftarrow \beta_t m_t + \eta \varphi(q_t);
2 w_{t+1} \leftarrow \text{flip}(w_t, m_{t+1});
3 m_{t+1} \leftarrow \text{reset}(w_t, m_{t+1});
```

Algorithm 3: Equivalent formulation of Boolean optimizer

```
\begin{array}{l} \textbf{Data:}\ Q_0,Q_1\ \text{quantizer} \\ \textbf{1}\ m_t \leftarrow \eta \tilde{\nabla} f(w_t) + e_t; \\ \textbf{2}\ \Delta_t \leftarrow Q_1(m_t,w_t); \\ \textbf{3}\ w_{t+1} \leftarrow Q_0(w_t - \Delta_t); \\ \textbf{4}\ e_{t+1} \leftarrow m_t - \Delta_t; \end{array}
```

Algorithm 3 describes an equivalent formulation of Boolean optimizer. Therein, Q_0 , Q_1 are quantizers which are specified in the following. Note that EF-SIGNSGD (SIGNSGD with Error-Feedback)

algorithm from [53] is a particular case of this formulation with $Q_0() = Identity()$ and $Q_1() = sign()$. For Boolean Logic abstraction, they are given by:

$$\begin{cases}
Q_1(m_t, w_t) = w_t(\text{ReLu}(w_t m_t - 1) + \frac{1}{2}\operatorname{sign}(w_t m_t - 1) + \frac{1}{2}), \\
Q_0(w_t) = \operatorname{sign}(w_t).
\end{cases}$$
(21)

The combination of Q_1 and Q_0 is crucial to take into account the reset property of the accumulator m_t . Indeed in practice, $\Delta_t := Q_1(m_t, w_t)$ is always equal to 0 except when $|m_t| > 1$ and $\mathrm{sign}(m_t) = \mathrm{sign}(w_t)$ (i.e., when the flipping rule is applied). As w_t has only values in $\{\pm 1\}$, Q_0 acts as identity function, except when Δ_t is non-zero (i.e., when the flipping rule is applied). With the choices (21), we can identify $\mathrm{flip}(w_t, m_t) = Q_0(w_t - Q_1(m_t, w_t))$. We do not have closed-form formula for $\mathrm{reset}(w_t, m_{t+1})$ from Algorithm 2, but the residual errors e_t play this role. Indeed, $e_{t+1} = m_t$ except when Δ_t is non-zero (i.e., when the flipping rule is applied and e_{t+1} is equal to 0).

The main difficulty in the analysis comes from the parameters quantization $Q_0()$. Indeed, we can follow the derivations in Appendix B.3 from [53] to bound the error term $\mathbb{E}[\|e_t\|^2]$, but we also have additional terms coming from the quantity:

$$h_t = Q_0(w_t - Q_1(m_t, w_t)) - (w_t - Q_1(m_t, w_t)).$$
(22)

As a consequence, assumptions 1 to 6 enable us to obtain $\mathbb{E}[h_t] = 0$ and to bound the variance of h_t . Remark A.7. Assumptions 1 to 3 are standard. Assumptions 4 to 6 are non-classic but dedicated to Boolean Logic strategy. A. 4 is equivalent to assuming Boolean Logic optimization presents at least one flip at every iteration t. A. 4 is classic in the literature of compressed SGD [53, 4, 80]. Moreover, A. 5 and A. 6 are not restrictive, but algorithmic choices. For example, rounding (Q_0 function) can be stochastic based on the value of the accumulator m_t . Similar to STE clipping strategy, the accumulator can be clipped to some pre-defined value κ before applying the flipping rule to verify A. 5.

Remark A.8. Our proof assumes that the step size η is constant over iterations. But in practice, we gently decrease the value of η at some time steps. Our proof can be adapted to this setting by defining a gradient accumulator a_t such that $a_{t+1} = a_t + \varphi(q_t)$. When η is constant we recover the definition (10) and we obtain $m_t = \eta a_t$. In the proposed algorithm, gradients are computed on binary weight w_t and accumulated in a_t . Then, one applies the flipping rule on the quantity $\tilde{w}_t = \eta a_t$ ($\tilde{w}_t = m_t$ when η is constant), and one (may) reset the accumulator a_t .

We start by stating a key lemma which shows that the residual errors e_t maintained in Algorithm 3 do not accumulate too much.

Lemma A.9. Under A. 3 and A. 4, the error can be bounded as $\mathbb{E}[\|e_t\|^2] \leq \frac{2\gamma}{(1-\gamma)^2}\eta^2\sigma^2$.

Proof. We start by using the definition of the error sequence:

$$||e_{t+1}||^2 = ||Q_1(m_t, w_t) - m_t||^2.$$

Next we make use of A. 4:

$$||e_{t+1}||^2 \le \gamma ||m_t||^2.$$

We develop the accumulator update:

$$||e_{t+1}||^2 \le \gamma ||e_t + \eta \tilde{\nabla} f(w_t)||^2.$$

We thus have a recurrence relation on the bound of e_t . Using Young's inequality, we have that for any $\beta > 0$,

$$||e_{t+1}||^2 \le \gamma(1+\beta)||e_t||^2 + \gamma(1+\frac{1}{\beta})\eta^2||\tilde{\nabla}f(w_t)||^2.$$

Rolling the recursion over and using A. 3 we obtain:

$$\mathbb{E}\left[\|e_{t+1}\|^{2}\right] \leq \gamma(1+\beta)\mathbb{E}\left[\|e_{t}\|^{2}\right] + \gamma(1+\frac{1}{\beta})\eta^{2}\mathbb{E}\left[\|\tilde{\nabla}f(w_{t})\|^{2}\right]$$

$$\leq \gamma(1+\beta)\mathbb{E}\left[\|e_{t}\|^{2}\right] + \gamma(1+\frac{1}{\beta})\eta^{2}\sigma^{2}$$

$$\leq \sum_{t}^{t}(\gamma(1+\beta))^{t}\gamma(1+\frac{1}{\beta})\eta^{2}\sigma^{2}$$

$$\leq \frac{\gamma(1+\frac{1}{\beta})}{1-\gamma(1+\beta)}\eta^{2}\sigma^{2}.$$

Take $\beta = \frac{1-\gamma}{2\gamma}$ and plug it in the above bounds gives:

$$\mathbb{E}[\|e_{t+1}\|^2] \le \frac{2\gamma}{(1-\gamma)^2}\eta^2\sigma^2.$$

Then, the next Lemma allows us to bound the averaged norm-squared of the distance between the Boolean weight and $w_t - Q_1(m_t, w_t)$. We make use of the previously defined quantity h_t (22) and have:

Lemma A.10. Under assumptions A. 5 and A. 6: $\mathbb{E}[\|h_t\|^2] \leq \eta d\kappa$.

Proof. Let consider a coordinate $i \in [d]$. $Q_0|_i$ as -1 or +1 for value with some probability $p_{i,t}$. For the ease of presentation, we will drop the subscript i. Denote $u_t := w_t - Q_1(m_t, w_t)$. Hence, h_t can take value $(1-u_t)$ with some probability p_t and $(-1-u_t)$ with probability $1-p_t$. Assumption A. 6 yields $2p_t - 1 = u_t$. Therefore, we can compute the variance of h_t as follows:

$$\mathbb{E}[\|h_t\|^2] = \mathbb{E}\left[\sum_{i=1}^{d} 1 + (w_t - Q_1(m_t, w_t))^2 - 2Q_0(w_t - Q_1(m_t, w_t)(w_t - Q_1(m_t, w_t))\right]$$

$$= \sum_{i=1}^{d} ((1 - u_t)^2 p_t + (-1 - u_t)^2 (1 - p_t))$$

$$= \sum_{i=1}^{d} (1 - u_t^2).$$

The definition of u_t leads to

$$1 - u_t^2 = 1 - (1 + Q_1(m_t, w_t)^2 - 2w_t Q_1(m_t, w_t))$$

= $Q_1(m_t, w_t)(2w_t - Q_1(m_t, w_t)).$

When $m_t < 1$, we directly have $Q_1(m_t, w_t)(2w_t - Q_1(m_t, w_t)) = 0 \le \eta \kappa$. When $m_t \ge 1$, we apply the definition of Q_1 to obtain:

$$Q_1(m_t, w_t)(2w_t - Q_1(m_t, w_t)) \le m_t(2 - m_t) \le \eta \kappa.$$

Therefore, we can apply this result to every coordinate, and conclude that:

$$\mathbb{E}[\|h_t\|^2] \le \eta d\kappa. \tag{23}$$

A.2.2 Proof of Theorem 3.17

We now can proceed to the proof of Theorem 3.17.

https://doi.org/10.52202/079017-1978

Proof. Consider the virtual sequence $x_t = w_t - e_t$. We have:

$$x_{t+1} = Q_0(w_t - \Delta_t) - (m_t - \Delta_t)$$

= $(Q_0(w_t - \Delta_t) + \Delta_t - e_t) - \eta \tilde{\nabla} f(w_t).$

Considering the expectation with respect to the random variable Q_0 and the gradient noise, we have:

$$\mathbb{E}[x_{t+1}|w_t] = x_t - \eta \nabla f(w_t).$$

We consider $\mathbb{E}_t[\cdot]$ the expectation with respect to every random process know up to time t. We apply the L-smoothness assumption A. 2, and assumptions A. 3, A. 6 to obtain:

$$\mathbb{E}_t[f(x_{t+1}) - f(x_t)] \le -\eta \langle \nabla f(x_t), \nabla f(w_t) \rangle + \frac{L}{2} \mathbb{E}_t \Big[\|(Q_0(w_t - \Delta_t) + \Delta_t) - \eta \tilde{\nabla} f(w_t) - w_t\|^2 \Big].$$

We now reuse h_t from (22) and simplify the above:

$$\mathbb{E}_{t}[f(x_{t+1}) - f(x_{t})] \leq -\eta \langle \nabla f(x_{t}), \nabla f(w_{t}) \rangle + \frac{L}{2} \mathbb{E}_{t} \Big[\|h_{t} - \eta \tilde{\nabla} f(w_{t})\|^{2} \Big] \\
\leq -\eta \langle \nabla f(x_{t}) - \nabla f(w_{t}) + \nabla f(w_{t}), \nabla f(w_{t}) \rangle + \frac{L}{2} \mathbb{E}_{t} \Big[\|h_{t} - \eta \tilde{\nabla} f(w_{t})\|^{2} \Big].$$

Using Young's inequality, we have that for any $\beta > 0$,

$$\mathbb{E}_{t}[f(x_{t+1}) - f(x_{t})] \leq -\eta \langle \nabla f(x_{t}) - \nabla f(w_{t}) + \nabla f(w_{t}), \nabla f(w_{t}) \rangle + \frac{L}{2} (1+\beta) \mathbb{E}_{t} [\|h_{t}\|^{2}] + \frac{L}{2} \eta^{2} (1+\frac{1}{\beta}) \sigma^{2}.$$

Making use again of smoothness and Young's inequality we have:

$$\mathbb{E}_{t}[f(x_{t+1}) - f(x_{t})] \leq -\eta \|\nabla f(w_{t})\|^{2} - \eta \langle \nabla f(x_{t}) - \nabla f(w_{t}), \nabla f(w_{t}) \rangle
+ \frac{L}{2} (1 + \beta) \mathbb{E}_{t} [\|h_{t}\|^{2}] + \frac{L}{2} \eta^{2} (1 + \frac{1}{\beta}) \sigma^{2}
\leq -\eta \|\nabla f(w_{t})\|^{2} + \frac{\eta \rho}{2} \|\nabla f(w_{t})\|^{2} + \frac{\eta}{2\rho} \|\nabla f(x_{t}) - \nabla f(w_{t})\|^{2}
+ \frac{L}{2} (1 + \beta) \mathbb{E}_{t} [\|h_{t}\|^{2}] + \frac{L}{2} \eta^{2} (1 + \frac{1}{\beta}) \sigma^{2}
\leq -\eta \|\nabla f(w_{t})\|^{2} + \frac{\eta \rho}{2} \|\nabla f(w_{t})\|^{2} + \frac{\eta L^{2}}{2\rho} \underbrace{\|x_{t} - w_{t}\|^{2}}_{\|e_{t}\|^{2}}
+ \frac{L}{2} (1 + \beta) \mathbb{E}_{t} [\|h_{t}\|^{2}] + \frac{L}{2} \eta^{2} (1 + \frac{1}{\beta}) \sigma^{2}.$$

Under the law of total expectation, we make use of Lemma A.9 and Lemma A.10 to obtain:

$$\mathbb{E}[f(x_{t+1})] - \mathbb{E}[f(x_t)] \le -\eta(1 - \frac{\rho}{2})\mathbb{E}[\|\nabla f(w_t)\|^2] + \frac{\eta L^2}{2\rho} \frac{4\gamma}{(1 - \gamma)^2} \eta^2 \sigma^2 + \frac{L}{8}\eta(1 + \beta)d\kappa + \frac{L}{2}\eta^2(1 + \frac{1}{\beta})\sigma^2.$$

Rearranging the terms and averaging over t gives for $\rho < 2$ (we can choose for instance $\rho = \beta = 1$):

$$\frac{1}{T+1} \sum_{t=0}^{T} \|\nabla f(w_t)\|^2 \le \frac{2(f(w_0) - f_*)}{\eta(T+1)} + 2L\sigma^2 \eta + 4L^2 \sigma^2 \frac{\gamma}{(1-\gamma)^2} \eta^2 + \frac{L}{2} d\kappa.$$

The bound in Theorem 3.17 contains 4 terms. The first term is standard for a general non-convex target and expresses how initialization affects convergence. The second and third terms depend on the fluctuation of the minibatch gradients. Another important aspect of the rate determined by Theorem 3.17 is its dependence on the quantization error. Note that there is an "error bound" of $\frac{L}{2}d\kappa$ that remains independent of the number of update iterations. The error bound is the cost of using discrete weights as part of the optimization algorithm. Previous work with quantized models also includes error bounds [61, 62].

B Code Sample of Core Implementation

In this section we provide example codes in Python of a Boolean linear layer that employs xor logic kernel. This implementation is in particular based on PyTorch [78]. The class of Boolean linear layer is defined in Algorithm 4; and its backpropagation mechanism, overwritten from autograd, is shown in Algorithm 5. Here, we consider both cases of the received backpropagation signal Z as described in Fig. 2, which are either Boolean (see Algorithm 6) or real-valued (see Algorithm 7). An example code of the Boolean optimizer is provided in Algorithm 8.

Notice that our custom XORLinear layer can be flexibly combined with any FP PyTorch modules to define a model. The parameters of this layer are optimized by the BooleanOptimizer, whereas those of FP layers are optimized by a common FP optimizer like Adam [54].

Algorithm 4: Python code of XOR linear layer

```
import torch
  from torch import Tensor, nn, autograd
 from typing import Any, List, Optional, Callable
7 class XORLinear(nn.Linear):
      def __init__(self, in_features: int, out_features: int, bool_bprop: bool, **kwargs):
          super(XORLinear, self).__init__(in_features, out_features, **kwargs)
self.bool_bprop = bool_bprop
10
13
      def reset_parameters(self):
           self.weight = nn.Parameter(torch.randint(0, 2, self.weight.shape))
           if self.bias is not None:
16
               self.bias = nn.Parameter(torch.randint(0, 2, (self.out_features,)))
18
19
      def forward(self, X):
           return XORFunction.apply(X, self.weight, self.bias, self.bool_bprop)
```

Algorithm 5: Python code of the backpropagation logic of XOR linear layer

```
class XORFunction(autograd.Function):
       Ostaticmethod
       def forward(ctx, X, W, B, bool_bprop: bool):
           ctx.save_for_backward(X,W,B)
           ctx.bool_bprop = bool_bprop
           # Elementwise XOR logic
           S = torch.logical_xor(X[:,None,:], W[None,:,:])
10
11
           # Sum over the input dimension
           S = S.sum(dim=2) + B
           \# O-centered for use with BatchNorm when preferred
           S = S - W.shape[1]/2
16
           return S
       @staticmethod
20
21
22
       def backward(ctx, Z):
           if ctx.bool_bprop:
    G_X, G_W, G_B = backward_bool(ctx, Z)
                G_X, G_W, G_B = backward_real(ctx, Z)
           return G_X, G_W, G_B, None
```

Algorithm 6: Backpropagation logic with Boolean received backpropagation

```
def backward_bool(ctx, Z):
       Variation of input:
           - delta(xor(x,w))/delta(x) = neg w
- delta(Loss)/delta(x) = xnor(z,neg w) = xor(z,w)
       Variation of weights:
           - delta(xor(x,w))/delta(w) = neg x
           - delta(Loss)/delta(x) = xnor(z,neg x) = xor(z,x)
       Variation of bias:
           - bias = xnor(bias,True) ==> Variation of bias is driven in
10
             the same basis as that of weight with xnor logic and input True.
       Aggregation:
           - Count the number of TRUEs = sum over the Boolean data
           - Aggr = TRUEs - FALSEs = TRUEs - (TOT - TRUEs) = 2TRUES - TOT
             where TOT is the size of the aggregated dimension
      X, W, B = ctx.saved_tensors
18
19
       # Boolean variation of input
       G_X = torch.logical_xor(Z[:,:,None], W[None,:,:])
21
22
23
       # Aggregate over the out_features dimension G_X = 2 * G_X.sum(dim=1) - W.shape[0]
       # Boolean variation of weights
       G_W = torch.logical_xor(Z[:,:,None], X[:,None,:])
27
       # Aggregate over the batch dimension G_W = 2 * G_W.sum(dim=0) - X.shape[0]
29
30
31
       # Boolean variation of bias
       if B is not None:
           # Aggregate over the batch dimension
34
           G_B = 2 * Z.sum(dim=0) - Z.shape[0]
35
36
       # Return
       return G_X, G_W, G_B
```

Algorithm 7: Backpropagation logic with real received backpropagation

```
def backward_real(ctx, Z):
    X, W, B = ctx.saved_tensors

"""

Boolean variation of input processed using torch avoiding loop:
    -> xor(Z: Real, W: Boolean) = -Z * emb(W)
    -> emb(W): T->1, F->-1 => emb(W) = 2W-1
    => delta(Loss)/delta(X) = Z*(1-2W) """

G_X = Z.mm(1-2*W)

"""

Boolean variation of weights processed using torch avoiding loop:
    -> xor(Z: Real, X: Boolean) = -Z * emb(X)
    -> emb(X): T->1, F->-1 => emb(X)
    -> emb(X): T->1, F->-1 => emb(X)

"""
G_W = Z.t().mm(1-2*X)

""" Boolean variation of bias """
if B is not None:
    G_B = Z.sum(dim=0)

# Return
return G_X, G_W, G_B
```

Algorithm 8: Python code of Boolean optimizer

```
class BooleanOptimizer(torch.optim.Optimizer):
            super(BooleanOptimizer, self).__init__(params, dict(lr=lr))
for param_group in self.param_groups:
    param_group['accums'] = [torch.zeros_like(p.data) for p in param_group['params']]
            def __init__(self, params, lr: float):
                          param_group['ratios'] = [0 for p in param_group['params']]
           @property
def nb_flips(self):
    n = self._nb_flips
    self._nb_flips = 0
10
11
15
16
17
            def step(self):
                   for param_group in self.param_groups:
    for idx, p in enumerate(param_group['params']):
        self.update(p, param_group, idx)
18
           def update(self, param: Tensor, param_group: dict, idx: int):
    accum = param_group['ratios'][idx] * param_group['accums'][idx] + param_group['lr
'] * param.grad.data
21
22
                   param_group['accums'][idx] = accum
                   param_to_flip = accum * (2*param.data-1) >= 1
param.data[param_to_flip] = torch.logical_not(param.data[param_to_flip])
25
26
                   param_group['accums'][idx][param_to_flip] = 0.
param_group['ratios'][idx] = 1 - param_to_flip.float().mean()
self._nb_flips += float(param_to_flip.float().sum())
27
```

C Training Regularization

C.1 Assumptions and Notations

We use capital letters to denote tensors. So, W^l , X^l , S^l , and Z^l denote weight, input, pre-activation, and received backpropagation tensors of a layer l. We also use Proposition A.2 and the discussion therein for studying Boolean logic in its embedding binary domain. It follows that in this section \mathbb{B} denotes the binary set $\{\mp 1\}$.

Let $\mathcal{N}(\mu, \sigma^2)$ denote the Gaussian distribution of mean μ and variance σ^2 , and $\mathbf{B}(p)$ denote Bernoulli distribution on \mathbb{B} with $p \in [0, 1]$. Note that for $X \sim \mathbf{B}(p)$, if $\mathbb{E}[X] = 0$, then $\mathbb{E}[X^2] = 1$.

We assume that the backpropagation signal to each layer follows a Gaussian distribution. In addition, according to experimental evidence, cf. Fig. 4, we assume that their mean μ can be neglected w.r.t. their variance σ^2 .

It follows that:

$$Z^l \sim \mathcal{N}(\mu, \sigma^2), \quad \text{with} \quad \mu \ll \sigma.$$
 (24)

We also assume that signals are mutually independent, i.e., computations are going to be made on random vectors, matrices and tensors, and it is assumed that the coefficients of these vectors, matrices and tensors are mutually independent. We assume that forward signals, backpropagation signals, and weights (and biases) of the layers involved are independent.

Consider a Boolean linear layer of input size n and output size m, denote:

- Boolean weights: $W^l \in \mathbb{B}^{m \times n}$;
- Boolean inputs: $X^l \in \mathbb{B}^n$;
- Pre-activation and Outputs: $S^l \in [|-n,n|]^m, X^{l+1} \in \mathbb{B}^m$;
- Downstream backpropagated signal: $Z^l \in \mathbb{R}^m$;
- Upstream backpropagated signal: $Z^{l-1} \in \mathbb{R}^n$.

In the forward pass, we have: $S^l = \mathbf{xor}(W^l, X^l)$, and $X^{l+1} = \mathbf{maj}(S^l)$.

With the following notations:

- $W^l = (W^l_{ij})_{i < m, j < n} \sim (\mathbf{B}(p_{ij}))_{i < m, j < n} \text{ with } p_{ij} \in [0, 1];$
- $X^{l} = (X_{i}^{l})_{i < n} \sim (\mathbf{B}(q_{i}))_{i < n} \text{ with } q_{i} \in [0, 1];$
- $Z^l = (Z_i^l)_{i < m} \sim (\mathcal{N}(\mu_i, \sigma_i^2))_{i < m};$
- \tilde{Z} stands for the truncated (with derivative of tanh) version of Z for sake of simplicity.

And assuming that $\forall i, \sigma_i = \sigma, \mu_i = \mu$ and $\mu \ll \sigma$, we can derive scaling factors for linear layers in the next paragraph.

Remark C.1. The scaling factor inside a convolutional layer behaves in a similar fashion except that the scalar dot product is replace by a **full** convolution with the 180-rotated kernel matrix.

C.2 Backpropagation Scaling

We now compute the variance of the upstream backpropagated signal \mathbb{Z}^{l-1} with respect to the number of neurons and the variance of the downstream backpropagated signal:

$$\forall j < n, \operatorname{Var}(Z_j^{l-1}) = \sum_{i=1}^{m} \operatorname{Var}(W_{ij}^l \tilde{Z}_i^l), \quad (W, Z \text{ are self mutually independent})$$
 (25)

$$=\sum_{i}^{m} \mathbb{E}\left[W_{ij}^{l}^{2} \tilde{Z}_{i}^{l^{2}}\right] - \mathbb{E}\left[W_{ij}^{l} \tilde{Z}_{i}^{l}\right]^{2} \tag{26}$$

$$= \sum_{i}^{m} \mathbb{E}\left[W_{ij}^{l^{2}}\right] \mathbb{E}\left[\tilde{Z_{i}}^{l^{2}}\right] - \mathbb{E}\left[W_{ij}^{l}\right]^{2} \mathbb{E}\left[\tilde{Z}_{i}^{l}\right]^{2}, \quad (W, Z \text{ are independent}) \quad (27)$$

$$= \sum_{i}^{m} \mathbb{E}\left[\tilde{Z}_{i}^{l^{2}}\right] - (2p_{ij} - 1)^{2} \mathbb{E}\left[\tilde{Z}_{i}^{l}\right]^{2}$$
 (28)

$$\simeq m\mathbb{E}\left[\tilde{Z}^{l^2}\right],$$
 $(\mu \ll \sigma)$ (29)

$$= m\mathbb{E}\Big[Z^{l^2}\Big]\mathbb{E}\Big[\frac{\partial \tanh}{\partial u}(u = \alpha W^l \cdot x^l)^2\Big], \qquad \text{(independence assump.)} \quad (30)$$

Ratio mean over standard deviation of the backprop over the first batch of all epochs

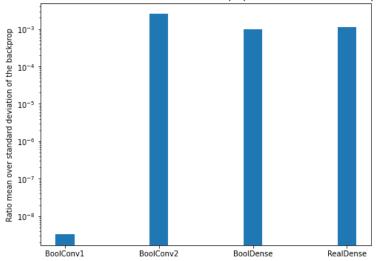


Figure 4: Empirical ratio of the mean to standard deviation of the backpropagation signal, experimented with CNN composed of BoolConv - BoolConv - BoolDense - RealDense layers and MNIST dataset.

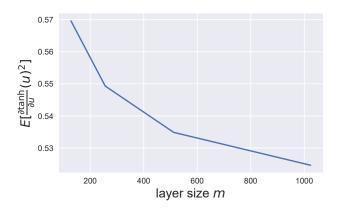


Figure 5: Expected value of \tanh derivative with integer values as input, for several output sizes m.

Let us focus on the term $\mathbb{E}\big[\frac{\partial \tanh}{\partial u}(u)^2\big]$, where $u\in[|-m,m|]$ for m even. The probability $\mathbb{P}(u=l)$ is computed thanks to enumeration. The event "u=l" indicates that we have k+l scalar values at level "+1" and k at level "-1" such that 2k+l=m. Hence,

$$\mathbb{P}(u=l) = \binom{m}{\frac{m-l}{2}} \frac{1}{2} \frac{m-\frac{m-l}{2}}{2}.$$
(31)

Thus,

$$\mathbb{E}\left[\frac{\partial \tanh}{\partial u}(u)^2\right] = \sum_{u=-m}^{m} \frac{\partial \tanh}{\partial u}(u)^2 p(u)$$
(32)

$$=2\sum_{u=0}^{m}\frac{\partial\tanh}{\partial u}(u)^{2}p(u),\qquad\text{(with symmetry)}\tag{33}$$

$$= \frac{1}{2^{m-1}} \sum_{u=0}^{m} {m \choose \frac{m-l}{2}} (1 - \tanh^{2}(\alpha u)).$$
 (34)

The latter can be easily pre-computed for a given value of output layer size m, see Figure 5.

The above figure suggests that for reasonable layer sizes m, $\mathbb{E}\left[\frac{\partial \tanh}{\partial u}(u=\alpha W^l\cdot x^l)^2\right]\simeq \frac{1}{2}$. As a consequence we can make use of (30), and approximate the variance of the backpropagated signal as:

$$\operatorname{Var}(Z^{l-1}) = \frac{m}{2} \operatorname{Var}(Z^l). \tag{35}$$

Remark C.2. The backpropagation inside a convolutional layer behaves in a similar fashion except that the tensor dot product is replace by a **full** convolution with the 180-rotated kernel matrix. For a given stride v and kernel sizes k_x , k_y , the variance of the backpropagated signal is affected as follow:

$$\operatorname{Var}(Z^{l-1}) = \frac{mk_x k_y}{2v} \operatorname{Var}(Z^l). \tag{36}$$

Let denote by MP the maxpooling operator (we assume its size to be (2,2)). In the backward pass, one should not forget the impact of the $\tanh'(\alpha\Delta)$, and the MP operator so that:

$$Z^{l-1} = \operatorname{Conv}_{\text{full}}(W^{l}_{\text{rot180}}, Z^{l}) \cdot \frac{\partial \tanh}{\partial u}(u = \operatorname{MP}[\operatorname{Conv}(\alpha W^{l}, X^{l})]) \cdot \frac{\partial \operatorname{MP}}{\partial u}(u = \operatorname{Conv}(\alpha W^{l}, X^{l})).$$
(37)

Let us focus on the last term: $\frac{\partial \mathrm{MP}}{\partial u}(u = \mathrm{Conv}(\alpha W^l, X^l)) = \mathbf{1}(u = \mathrm{max}(\mathrm{Conv}(\alpha W^l, X^l)))$. Hence,

$$\mathbb{E}\left[\frac{\partial \mathrm{MP}}{\partial u}(u = \mathrm{Conv}(\alpha W^l, X^l))^2\right] = \mathbb{E}\left[\frac{\partial \mathrm{MP}}{\partial u}(u = \mathrm{Conv}(\alpha W^l, X^l))\right]$$
(38)

$$= \frac{1}{4} \times 1 + 0. \tag{39}$$

As a consequence, for a given stride v and kernel sizes k_x, k_y , the variance of the backpropagated signal is affected as follow:

$$\operatorname{Var}(Z^{l-1}) = \frac{1}{4} \frac{m k_x k_y}{2v} \operatorname{Var}(Z^l). \tag{40}$$

C.3 BackPropagation Through Boolean Activation Function

Due to the binary activation, the effect on the loss function by an action on weight w diminishes with the distance $\Delta:=|s-\tau|$ from threshold τ to pre-activation s to which w contributes. Throughout the step activation function, the backpropagation signal can be optionally re-weighted by a function which is inversely proportional to Δ , for instance, $\tanh'(\Delta)$, $(1+\Delta)^{-2}$, $\exp(-\Delta)$, or any other having this property. In our study, $\tanh'(\alpha\Delta)$ turns out to be a good candidate in which α is used to match the spreading in terms of standard deviation of this function and that of the pre-activation distribution.

We start by computing the variance of the pre-activation signal S^l with respect to the number of neurons, without considering the influence of backward \tanh' :

$$\forall j < n, \operatorname{Var}(S_j^l) = \sum_{i=1}^{m} \operatorname{Var}(W_{ij}^l X_i^l), \quad (W, X \text{ are self mutually independent})$$
 (41)

$$= \sum_{i}^{m} \mathbb{E}\left[W_{ij}^{l}^{2} X_{i}^{l^{2}}\right] - \mathbb{E}\left[W_{ij}^{l} X_{i}^{l}\right]^{2}\right] \tag{42}$$

$$= \sum_{i}^{m} \mathbb{E}\left[W_{ij}^{l^2}\right] \mathbb{E}\left[X_i^{l^2}\right] - \mathbb{E}\left[W_{ij}^{l}\right]^2 \mathbb{E}\left[X_i^{l}\right]^2, \quad (W, X \text{ are independent}) \quad (43)$$

$$= \sum_{i}^{m} \mathbb{E}\left[X_{i}^{l^{2}}\right] - (2p_{ij} - 1)^{2} \mathbb{E}\left[X_{i}^{l}\right]^{2}$$
(44)

$$\simeq m\mathbb{E}\left[X^{l^2}\right], \quad \text{since } \mu \ll \sigma$$
 (45)

$$=m.$$
 (46)

This leads to

$$\alpha = \frac{\pi}{2\sqrt{3m}}\tag{47}$$

in order to have $\mathrm{Var}(\alpha S) = \frac{\pi^2}{12}$, where m is the range of the pre-activation, e.g., $m = c_{\mathrm{in}} \times k_x \times k_y$ for a 2D convolution layer of filter dimensions $[c_{\mathrm{in}}, k_x, k_y, c_{\mathrm{out}}]$.

D Experimental Details

D.1 Image Classification

D.1.1 Training Setup

The presented methodology and the architecture of the described Boolean neural networks (NNs) were implemented in PyTorch [78] and trained on 8 Nvidia Tesla V100 GPUs. The networks thought predominantly Boolean, also contain a fraction of FP parameters that were optimized using the Adam optimizer [54] with learning rate 10^{-3} . For learning the Boolean parameters we used the Boolean optimizer (see Algorithm 8). Training the Boolean networks for image classification was conducted

with learning rates $\eta=150$ and $\eta=12$ (see Equation 10), for architectures with and without batch normalization, respectively. The hyper-parameters were chosen by grid search using the validation data. During the experiments, both optimizers used the cosine scheduler iterating over 300 epochs.

We employ data augmentation techniques when training low bitwidth models which otherwise would overfit with standard techniques. In addition to techniques like random resize crop or random horizontal flip, we used RandAugment, lighting [68] and Mixup [115]. Following [98], we used different resolutions for the training and validation sets. For IMAGENET, the training images were 192×192 px and 224×224 px for validation images. The batch size was 300 for both sets and the cross-entropy loss was used during training.

D.1.2 CIFAR10

VGG-SMALL is found in the literature with different fully-connected fully-connected (FC) layers. Several works take inspiration from the classic work of [24], which uses 3 FC layers. Since other BNN methodologies only use a single FC layer, Table 9 presents the results with the modified VGG-SMALL.

Table 9: Top-1 accuracy for different binary methodologies using the modified VGG-SMALL (ending with 1 FC layer) on the CIFAR10 dataset.

Method	Forward Bit-width (W/A)	Training Bit-width (W/G)	Acc. (%)
FP	32/32	32/32	93.8
XNOR-NET [86]	1/1	32/32	87.4
Lab [43]	1/1	32/32	87.7
RAD [27]	1/1	32/32	90.0
IR-NET [84]	1/1	32/32	90.4
RBNN [65]	1/1	32/32	91.3
SLB [111]	1/1	32/32	92.0
B⊕LD [Ours]	1/1	1/16	90.8

D.1.3 Ablation Study on Image Classification

The final block design for image classification was established after iterating over two models. The Boolean blocks examined were evaluated using the RESNET18 baseline architecture and adjusting the training settings to improve performance. Figure 6 presents the preliminary designs.

The Boolean Block I, Figure 6a, is similar to the original RESNET18 block in that BN operations are removed and ReLUs are replaced by the Boolean activation. This design always includes a convolution in the shortcut with spatial resolution being handled by the stride. Notice that for this block we add a Boolean activation after the Maxpool module in the baseline (also for the final baseline architecture). The Boolean Block II, Figure 6b, is composed by two stacked residual modules. For downsampling blocks we use the reshaping operation to reduce the spatial resolution and enlarge the channel dimensions both by a factor of 2. The shortcut is modified accordingly with different operations in order to guarantee similar spatial dimensions before the summation.

Table 10 summarizes the results obtained with the proposed designs on IMAGENET. During our experimentation, we validated the hypothesis that increasing network capacity on the convolutional layers yielded higher accuracy values. However, similar to FP CNNs, we confirmed there is a limit by which the hypothesis ceases to be true, leading to overfitting. Incorporating a more severe training strategy had a sustained positive impact. Even so, for larger configurations, the compromise between accuracy and size can be cumbersome.

Among the strategies to reduce overfitting during training we included: mixup data-augmentation [115], image illumination tweaking, rand-augment and smaller input resolution for training than for validation [98]. All combined, increased the accuracy by \sim 3 points (check results for Block II + base channel 230 with and w/o additional data augmentation).

Compared to Block II, notice that the data streams in Block I are predominantly Boolean throughout the design. This is because it makes use of lightweight data types such as integer (after convolutions) and binary (after activations). In addition, it avoids the need of using a spatial transformation that may affect the data type and data distribution. In that regard, Block II requires 4 times more parameters

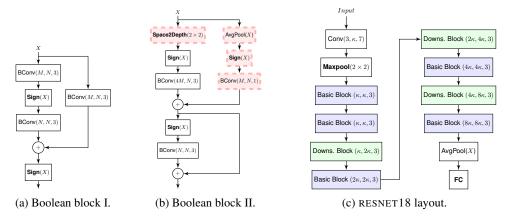


Figure 6: Preliminary designs for the baseline architecture and the Boolean basic blocks. The dashed and red-shaded operations in the Boolean block II are introduced for downsampling blocks.

for the convolution after reshaping, than the corresponding operation in Block I. This is exacerbated in upper layer convolutions, where the feature maps are deeper. Therefore, it makes sense to use Block I, as it is lighter and less prone to overfitting when the network capacity is expanded.

D.1.4 Neural Gradient Quantization

For completeness, we also implemented neural gradient quantization to quantize it by using INT4 quantization with logarithmic round-to-nearest approach [21] and statistics aware weight binning [22]. Statistics aware weight binning is a method that seeks for the optimal scaling factor, per layer, that minimizes the quantization error based on the statistical characteristics of neural gradients. It involves per layer additional computational computations, but stays negligible with respect to other (convolution) operations. On IMAGENET, we recover the findings from [21]: 4 bits quantization is enough to recover standard backpropagation performances.

D.1.5 Basic Blocks SOTA BNNs for Classification

Recent BNN methodologies have proposed different mechanisms to improve performance. Most of them exploit full-precision operations to adjust datastreams within the network, like shift and scaling factors before binary activations [68] or channel scaling through Squeeze-and-Excitation modules [73, 36]. Figure 7 shows the basic blocks of three methodologies that perform particularly well in IMAGENET. Together with BN and regular activations, those techniques not only add an additional level of complexity but also lead to heavier use of computational resources and latency delays.

For comparison we also show the proposed block (Figure 7a) used in our experiments for Image Classification, Image Segmentation and Image Super-Resolution. Our block is compact in the sense that it only includes Boolean convolutions and Boolean activations, strategically placed to keep the input and output datastreams Boolean.

Table 10: Evaluation of the proposed blocks in IMAGENET and their respective configurations during training.

Block Design	Base Channel	1 st Conv. Bit-width	Shortcut Fil. Size	Data Augmentation	Acc. (%)
	128	32	1×1	Random Crop, Random Flip	53.35
	192	32	1×1	Random Crop, Random Flip	56.79
Block I	192	32	1×1	Lighting, Mixup, RandAugment and [98]	61.90
	256	32	1×1	Lighting, Mixup, RandAugment and [98]	64.32
	256	32	3×3	Lighting, Mixup, RandAugment and [98]	66.89
	128	1	1×1	Random Crop, Random Flip	56.05
	128	32	1×1	Random Crop, Random Flip	58.38
Block II	192	32	1×1	Random Crop, Random Flip	61.10
DIOCK II	192	32	1×1	Lighting, Mixup, RandAugment and [98]	63.21
	230	32	1×1	Random Crop, Random Flip	61.22
	230	32	1×1	Lighting, Mixup, RandAugment and [98]	64.41

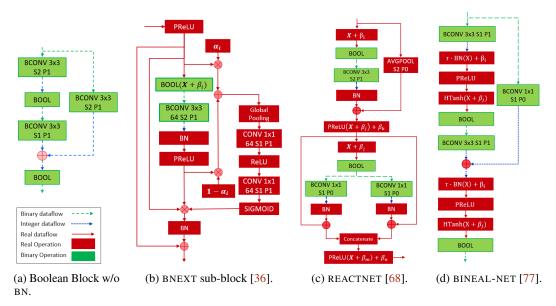


Figure 7: Comparative graph of popular BNN techniques and our Boolean module. Notice how multiple full-precision operations like BN, PReLU, or Squeeze-and-Excitation are overly used on each BNN block.

D.2 Image Super-resolution

The seminal EDSR [64] method for super-resolution was used together with our Boolean methodology. In particular, the residual blocks are directly replaced by our Boolean basic block, see Figure 8. For all three tasks in super-resolution, (i.e. $\times 2, \times 3, \times 4$), training was carried out with small patches of 96×96 px (40 of them extracted randomly from each single image in the DIV2K dataset) and validated with the original full-resolution images. The learning rate for real and boolean parameters were 10^{-4} and $\eta=36$, respectively. The networks were trained by minimizing the L_1 -norm between the ground-truth and the predicted upsampled image while using the Adam optimizer and Boolean optimizer (see Algorithm 8). In our experiments the batch size was 20. Some example images generated by our methodology are showed in Figures 9 and 10.

D.3 Semantic Segmentation

D.3.1 Network architecture

Our Boolean architecture is based on DEEPLABV3 [17], which has shown great success in semantic segmentation. It is proven that using dilated or atrous convolutions, which preserve the large feature maps, instead of strided convolutions is prominent for this task. In our Boolean model with RESNET18 layout, we replace the strided convolutions in the last two RESNET18 layers with the non-strided version, and the dilated convolutions are employed to compensate for the reduced receptive field. Thus, the images are $8\times$ downsampled instead of $32\times$, preserving small object features and allowing more information flow through the Boolean network. As shown in Figure 6a, in the Boolean basic block, a 3×3 convolution instead of 1×1 convolution is used to ensure the comparable dynamic range of pre-activations between the main path and the shortcut. Keeping these Boolean convolutional layers non-dilated naturally allows the backbone to extract multi-scale features without introducing additional computational cost.

The Atrous Spatial Pyramid Pooling (ASPP) consists of multiple dilated convolution layers with different dilation rates and global average pooling in parallel, which effectively captures multi-scale information. In the Boolean ASPP (BOOL-ASPP), we use one 1×1 Boolean convolution and three 3×3 Boolean dilated convolution with dilation rates of $\{12, 24, 36\}$ following by Boolean activation functions. The global average pooling (GAP) branch in ASPP captures image-level features, which is crucial for global image understanding as well as large object segmenting accuracy. However, in BOOL-ASPP, as shown in Figure 12c, the Boolean input X leads to significant information loss

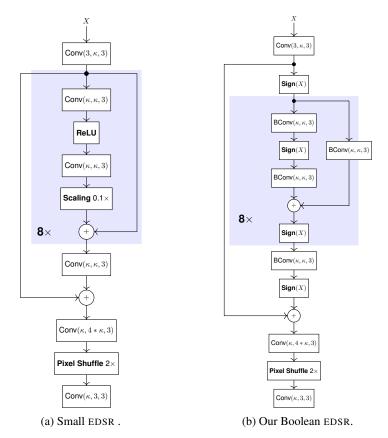


Figure 8: Small EDSR for single scale $\times 2$ super-resolution and our Boolean version with Boolean residual blocks. In both architectures the channels dimensions are $\kappa=256$ and the shaded blocks are repeated $8\times$.

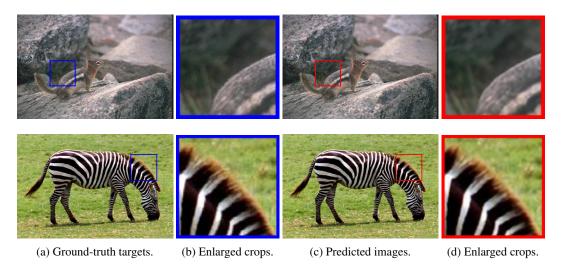


Figure 9: Ground-truth high resolution images and the output of our Boolean super-resolution methodology. First row: image "013" from BSD100, with PSNR: 35.54 dB. Second row: image "014" from Set14, with PSNR: 33.92 dB.

before the global average pooling may cause performance degradation on large objects. Therefore, we keep the inputs integer for the GAP branch as demonstrated in Figure 12d. To prevent numerical instability, batch normalization is used in the GAP branch before each activation function. Using



Figure 10: Ground-truth high resolution target image (top) and the output of our Boolean super-resolution methodology (bottom). Image "0810" from the validation set of DIV2K, with PSNR: 34.90 dB

BOOL-ASPP enhances the multi-scale feature extraction and avoids parameterized upsampling layers, e.g. transposed convolution.

D.3.2 Training setup

The model was trained on the CITYSCAPES dataset for 400 epochs with a batch size of 8. The AdamW optimizer [71] with an initial learning rate of 5×10^{-4} and the Boolean logic optimizer (see Algorithm 8) with a learning rate of $\eta=12$ were used respectively for real and Boolean parameters. At the early training stage, parameters could easily be flipped due to the large backward signal; thus, to better benefit from the IMAGENET-pretrained backbone, we reduce the learning rate for Boolean parameters in the backbone to $\eta=6$. We employed the polynomial learning rate policy with p=0.9 for all parameters. The cross-entropy loss was used for optimization. We did not employ auxiliary

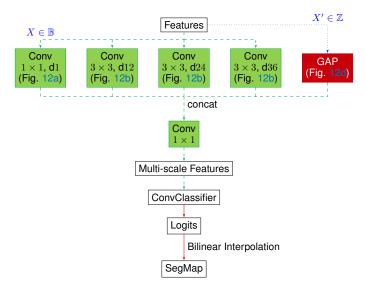


Figure 11: Boolean segmentation architecture.

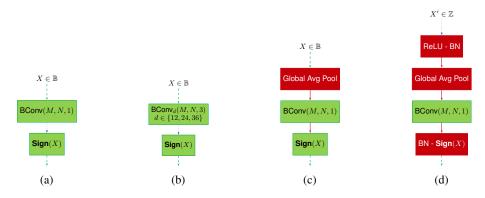


Figure 12: Boolean Atrous Spatial Pyramid Pooling (BOOL-ASPP) architecture. (a) 1×1 Conv branch. (b) 3×3 dilated Conv branch with dilation rate of d. (c) Naive global average pooling branch. (d) Global average pooling branch.

loss or knowledge distillation as these training techniques require additional computational cost, which is not in line with our efficient on-device training objective.

D.3.3 Data sampling and augmentation

We aim to reproduce closely full-precision model performance in the semantic segmentation task with Boolean architecture and Boolean logic training. Due to the nature of the Boolean network, the common regularization method, e.g., weight decay, is not applicable. Moreover, with more trainable parameters, the Boolean network can suffer from over-fitting. In particular, as shown in Table 11, the imbalanced dataset for semantic segmentation aggravates the situation. There is a significant performance gap for several classes which has low occurrence rate, including *rider* (9.5%), *motor* (11.2%), bus (9.5%), truck (6.9%), train (17.0%). We argue that the performance gap is due to the similarity between classes and the dataset's low occurrence rate, which is confirmed as shown in Figure 13.

Data augmentation and sampling are thus critical for Boolean model training. Regarding data augmentation, we employed multi-scale scaling with a random scaling factor ranging from 0.5 to 2. We adopted a random horizontal flip with probability p=0.5 and color jittering. In addition, we used rare class sampling (RCS) [45] to avoid the model over-fitting to frequent classes. For class c,

Table 11: Class per image and performance gap occurrence rates in CITYSCAPES training set with naive BOOL-ASPP design. Class with low performance gap[†] and class with high performance gap*.

	Image ratio (%)	Δ mIoU (%)
Road [†]	98.62	0.0
Sideway [†]	94.49	0.7
Building [†]	98.62	0.6
Wall	32.61	7.4
Fence	43.56	3.8
Pole	99.13	1.8
Light	55.73	6.5
Sign [†]	94.39	2.8
Vegetation [†]	97.18	0.1
Terrain [†]	55.60	0.8
Sky^{\dagger}	90.29	-0.2
Person	78.76	1.5
Rider*	34.39	7.4
Car [†]	95.19	0.3
Truck*	12.07	6.9
Bus*	9.21	12.8
Train*	4.77	17.0
Motor*	17.24	15.3
bike	55.33	2.0

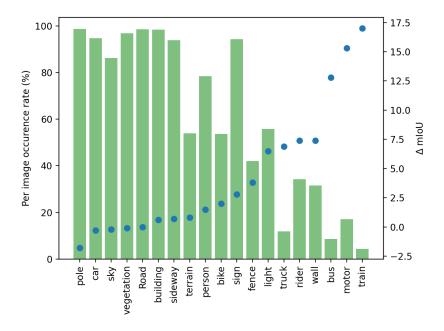


Figure 13: Class per image occurrence ratio and performance gap with naive BOOL-ASPP design.

the occurrence frequency in image f_c is given by:

$$f_c = \frac{\sum_{i=1}^{N} \mathbf{1}(c \in y_i)}{N},\tag{48}$$

where N is the number of samples and y_i is the set of classes existing in sample i. The sampling probability of class c is thus defined as:

$$p_c = \frac{\exp\left(\frac{1-f_c}{T}\right)}{\sum_{c'=1}^K \exp\left(\frac{1-f_{c'}}{T}\right)},\tag{49}$$

where K is the number of classes, and T is a hyper-parameter for sampling rate balancing. In particular, for the CITYSCAPES dataset, we selected T=0.5.

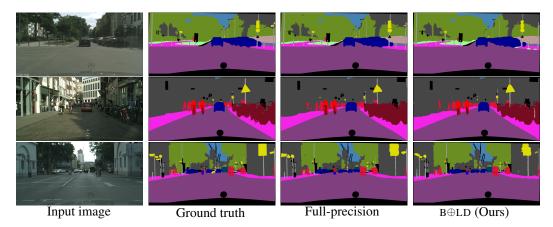


Figure 14: Qualitative comparison of Boolean model on CITYSCAPES validation set.

D.3.4 Qualitative analysis on cityscapes validation set

The qualitative results of our Boolean network and the full-precision based are demonstrated in Figure 14. Despite the loss of model capacity, the proposed Boolean network trained with Boolean logic optimizer has comparable performance with large objects in the frequent classes, even in the complicated scene.

D.3.5 More experiments on semantic segmentation

We evaluated the effectiveness of BOOL-ASPP by investigating the per-class performance gap to the full-precision model. As demonstrated in Table 12, a significant gap exists between the Boolean architecture with naive BOOL-ASPP design; i.e., using Boolean activations for ASPP module as illustrated in Figure 12c. However, the gap could be reduced by using BOOL-ASPP and RCS. In particular, the BOOL-ASPP improves the IoU of *truck* from 54.5% to 64.1% and *bus* from 65.5% to 68.8%, *bike* from 68.8% to 69.1% and *motor* from 42.8% to 46.8%. This indicates that combining proposed BOOL-ASPP and RCS improves the model performance on low occurrence classes as well as similar classes with which are easy to be confused.

D.3.6 Validation on PASCAL VOC 2012 dataset

We also evaluated our Boolean model on the 21-class PASCAL VOC 2012 dataset with augmented additional annotated data containing 10,582,1,449, and 1,456 images in training, validation, and test set, respectively. The same setting is used as in the experiments on the CITYSCAPES dataset, except the model was trained for 60 epochs.

As shown in Table 13, our model with fully Boolean logic training paradigm, i.e., without any additional intermediate latent weight, achieved comparable performance as the state-of-the-art latent-weight-based method. Our Boolean model improved performance by incorporating multi-resolution feature extraction modules to 67.3% mIoU.

sike mIoU 79.8 90.1 49.4 90.0 57.5 92.4 74.3 54.6 91.2 78.3 58.0 70.8 70.7 97.3 48.5 55.0 59.2 69. 61.4 66.6 Naive Bool ASPP 97.3 79 1 89.5 51.2 52.7 90.1 56.7 92.6 47.2 91.5 66.3 66.2 -1.8 -0.1 0.8 7.4 6.9 17.0 4.4 0.7 3.8 6.5 2.8 -0.21.5 -0.312.8 15.3 2.0 0.0 0.6 B⊕LD [Ours] 53.3 **5.9** 45.1 9.5

Table 12: Class-wise IoU performance on CITYSCAPES validation set.

Table 13: Performance on PASCAL VOC 2012 val set.

Seg. head	Model	mIoU (%)	Δ
FP baseline FCN-32s [70] GROUP-NET [119] B⊕LD [Ours]		64.9 60.5 60.1	- 4.4 4.8
DEEPLABV3 [17]	FP baseline B⊕LD [Ours]	72.1 67.3	4.8

D.4 Boolean BERT Fine-tuning

We conducted experiments on the BERT model [26] drawing on the experimental framework proposed in BIT [67]. For this experiment, our goal was to fine-tune the original pre-trained BERT model using Boolean precision. To validate our method we used the GLUE benchmark [102] with 8 datasets. For our experiments, we modified the baseline FP architecture using the proposed methodology Algorithm 8. That is, core operations within the transformer are substituted by native Boolean components, e.g. Boolean activations and Boolean linear layers. The FP parameters were optimized using the Adam optimizer with the learning rates proposed in [67]. Correspondingly for Boolean weights, we used our Boolean optimizer with learning rate $\eta = 100$.

E Energy Estimation

Energy consumption is a fundamental metric for measuring hardware complexity. However, it requires specific knowledge of computing systems and makes it hard to estimate. Few results are available, though experimental-based and limited to specific tested models [see, e.g., 32, 88, 74, 13, 14, 33]. Although experimental evaluation is precise, it requires considerable implementation efforts while not generalizing. In addition, most relevant works are only limited to inference and not training [19, 57, 110].

E.1 Hardware Specification

Ascend architecture. We intend to estimate the training energy consumption on Ascend chip architecture introduced in [63] and dedicated to DNN computing. The core design of Ascend is described in [63]. Essentially, it introduces a 3D (cube) computing unit, providing the bulk of high-intensity computation and increasing data reuse. On the other hand, it provides multiple levels of on-chip memory. In particular, memory L0, which is nearest to the computing cube, is tripled to boost further near-memory computing capability, namely L0-A dedicated to the left-hand-side (LHS) input data, L0-B dedicated to RHS input data, and L0-C for the output. For instance, in a convolution, L0-A, L0-B, and L0-C correspond to the input feature maps (IFMAPS), FILTERS, and output feature maps (OFMAPS), respectively. In addition, the output results going through L0-C can be processed by a Vector Unit for in-place operations such as normalization and activation. Table 14 shows energy efficiency and capacity of the memory hierarchy of a commercial Ascend architecture [63].

Table 14: Memory hierarchy and energy efficiency (EE) of an Ascend core [63] used in our evaluation.

	L3 (DRAM)	L2	L1	L0-A	L0-B	L0-C
EE [GBPS/mW]	0.02	0.2	0.4	4.9	3.5	5.4
Capacity [KB]	_	8192	1024	64	64	256

Nvidia architecture. We also have estimated the energy consumption for the Nvidia GPU (Tesla V100). Similarly, this architecture utilizes different memory levels with varying read and write energy characteristics. For instance, the L2 cache size is 6 MB per GPU. The L1 cache is 64 KB per Streaming Multiprocessor (SM). Each Tesla V100 GPU has 80 SMs, so the total L1 cache is 5120 KB (5 MB). However, specific details on the read and write energy consumption for each memory level of the Tesla V100 GPU are proprietary and not publicly disclosed by Nvidia. Thus, we show

the normalized energy consumption relative to the computation of a MAC at the arithmetic logic unit (ALU). Table 15 shows the numbers for each storage level, which are extracted from a commercial 65 nm process [18].

Table 15: Normalized energy cost relative to the computation of one MAC operation at ALU.

DRAM	L2	L1	RF	1 MAC at ALU
200×	$6 \times$	$2\times$	$1 \times$	1×

E.2 Compute Energy

Energy consumption is the sum of compute and memory energies. Compute energy is simply given by the number of arithmetic operations multiplied by their unit cost. The number of arithmetic operations is directly determined from the layer's parameters. For the Ascend architecture, their unit cost is obtained by considering the compute efficiency at 1.7 TOPS/W [63]. For Boolean logic operations, we follow the usual estimation that ADD INT-n costs (2n-1) logic operations where n stands for bitwidth.

E.3 Memory Energy

On the other hand, *memory energy* is all consumed for moving data between their storage through memory levels and the computing unit during the entire lifetime of the process. Since energy consumed at each memory level is given by the number of data accesses to that level times peraccess energy cost, it consists in determining the number of accesses to each level of all data streams (i.e., LHS, RHS, Output). Besides taking into account the hardware architecture and memory hierarchy of chip, our approach to quantifying memory energy is based on existing methods [19, 94, 57, 110, 42, 108] for dataflow and energy evaluation. Given the layer parameters and memory hierarchy, it amounts to:

- 1. Tiling: determining the tiling strategy for allocating data streams on each memory level.
- Movement: specifying how data streams are reused or kept stationary to determine their access numbers.

In the following, we present our method for the forward and backward passes by taking the example of a convolution layer, as convolutions are the main components of CNNs and the primary source of complexity due to their high data reuse. The parameters of 2D convolution layer are summarized in Table 16. Here, we denote IFMAPS, FILTERS, and OFMAPS by I, F, and O, respectively.

Table 16: Shape parameters of a convolution layer.

Parameter	Description		
N	batch size		
M	number of OFMAPS channels		
C	number of IFMAPS channels		
H^I/W^I	IFMAPS plane height/width		
H^F/W^F	FILTERS plane height/width		
H^O/W^O	OFMAPS plane height/width		

E.3.1 Tiling

Since the IFMAPS and FILTERS are usually too large to be stored in buffers, the tiling strategy is aimed at efficiently transferring them to the computing unit. Determining tiling parameters, which are summarized in Table 17, is an NP-Hard problem [110].

Table 17: Tiling parameters of a convolution layer.

Parameter	Description		
M_2	number of tiling weights in L2 buffer		
M_1	number of tiling weights in L1 buffer		
M_0	number of tiling weights in L0-B buffer		
N_2	number of tiling IFMAPS in L2 buffer		
N_1	number of tiling IFMAPS in L1 buffer		
N_0	number of tiling IFMAPS in LO-A buffer		
H_2^I/W_2^I	height/width of tiling IFMAPS in L2 buffer		
H_1^I/W_1^I	height/width of tiling IFMAPS in L2 buffer		
H_0^I/W_0^I	height/width of tiling IFMAPS in LO-A buffer		

An iterative search over possibilities subject to memory capacity constraint provides tiling combinations of IFMAPS and FILTERS on each memory level. Different approaches can be used and Algorithm 9 shows an example that explores the best tiling parameters subjected to maximizing the buffer utilization and near compute stationary (i.e., as much reuse as possible to reduce the number of accesses to higher levels). Therein, the amount of data stored in level i is calculated as:

$$Q_i^I = N_i \times C_i \times H_i^I \times W_i^I \times b^I,$$

$$Q_i^F = M_i \times C_i \times H^F \times W^F \times b^F,$$
(50)

where Q_i^I/Q_i^F and b^I/b^F represent the memory and bitwidth of IFMAPS/FILTERS, respectively.

Algorithm 9: Loop tiling strategy in the *i*th level

```
Input: tiling parameters of IFMAPS and FILTERS at level i+1, and buffer capacity of level i. Output: tiling parameters of IFMAPS and FILTERS at level i.
```

```
1 Initialize
           \mathcal{E}^{min} := \infty;
3 for M_i \leftarrow M_{i+1} to 1 do
           for N_i \leftarrow N_{i+1} to 1 do
 4
                 \begin{array}{l} \textbf{for } H_i^I \leftarrow H_{i+1}^I \ \textbf{to } H^F \ \textbf{do} \\ \mid \ \textbf{for } W_i^I \leftarrow W_{i+1}^I \ \textbf{to } W^F \ \textbf{do} \end{array}
 5
 6
                              Calculate Q_i, the required amount of IFMAPS and FILTERS to be stored in the ith
 7
                               level of capacity Q_i^{\text{max}};
                              Calculate \mathcal{E}_i, the energy cost of moving IFMAPS and FILTERS from the ith level;
 8
                              if (Q_i \leq Q_i^{max}) and (\mathcal{E}_i < \mathcal{E}^{min}) then
                                    Retain tiling parameters as best;
10
                                    \mathcal{E}^{\min} \leftarrow \mathcal{E}_i;
12 return Best tiling parameters
```

E.3.2 Data movement

For data movement, at level L0, several data stationary strategies, called *dataflows*, have been proposed in the literature, notably weight, input, output, and row stationary [19]. Since Ascend chip provides tripled L0 buffers, partial sums can be directly stationary in the computing cube, hence equivalent to output stationary whose implementation is described in [28]. For the remaining levels, our question of interest is how to move IFMAPS block $[N_{i+1}, C_{i+1}, H^I_{i+1}, W^I_{i+1}]$ and FILTERS block $[M_{i+1}, C_{i+1}, H^F, W^F]$ from level i+1 to level i efficiently. Considering that:

- IFMAPS are reused by the FILTERS over output channels,
- FILTERS are reused over the IFMAPS spatial dimensions,
- FILTERS are reused over the batch dimension,

• IFMAPS are usually very large whereas FILTERS are small,

the strategy that we follow is to keep FILTERS stationary on level i and cycle through IFMAPS when fetching them from level i+1 as shown in Algorithm 10. Therein, FILTERS and IFMAPS are read block-by-block of their tiling sizes, i.e., FILTERS block $[M_i, C_i, H^F, W^F]$ and IFMAPS block $[N_i, C_i, H_i^I, W_i^I]$. Hence, the number of filter accesses to level i+1 is 1 whereas the number of IFMAPS accesses to level i+1 equals the number of level-i FILTERS blocks contained in level i+1. Following this method, the number of accesses to memory levels of each data stream can be determined. Hence, denote by:

- n_i^d : number of accesses to level i of data d,
- ε_i : energy cost of accessing level i, given as the inverse of energy efficiency from Table 14.

Following [19], the energy cost of moving data d from DRAM (L3) into the cube is given as:

$$\mathcal{E}^d = n_3^d \varepsilon_3 + n_3^d n_2^d \varepsilon_2 + n_3^d n_2^d n_1^d \varepsilon_1 + n_3^d n_2^d n_1^d n_0^d \varepsilon_0.$$
 (51)

Regarding the output partial sums, the number of accumulations at each level is defined as the number of times each data goes in and out of its lower-cost levels during its lifetime. Its data movement energy is then given as:

$$\mathcal{E}^{O} = (2n_3^{O} - 1)\varepsilon_3 + 2n_3^{O}(n_2^{O} - 1)\varepsilon_2 + 2n_3^{O}n_2^{O}(n_1^{O} - 1)\varepsilon_1 + 2n_3^{O}n_2^{O}n_1^{O}(n_0^{O} - 1)\varepsilon_0, \quad (52)$$

where factor of 2 accounts for both reads and writes and the subtraction of 1 is because we have only one write in the beginning [19].

```
Algorithm 10: Data movement from i + 1 to i levels
```

```
Input: tiling parameters of IFMAPS and FILTERS at levels i+1 and i.

1 repeat

2 read next FILTERS block of size [M_i, C_i, H^F, W^F] from levels i+1 to i;

3 repeat

4 read next IFMAPS block of size [N_i, C_i, H_i^I, W_i^I] from levels i+1 to i;

5 let the data loaded to i be processed;

6 until all IFMAPS are read into level i;

7 until all FILTERS are read into level i;
```

E.3.3 Forward

In the forward pass, there are three types of input data reuse:

- For an $H^I \times W^I$ IFMAP, there are $H^O \times W^O$ convolutions performed with a single $H^F \times W^F$ filter to generate a partial sum. The filter is reused $H^O \times W^O$ times, and this type of reuse is defined as *filter convolutional reuse*. Also, each feature in the IFMAPS is reused $H^F \times W^F$ times, and this is called *feature convolutional reuse*.
- Each IFMAP is further reused across M filters to generate M output channels. This is called IFMAPS reuse.
- Each FILTER is further reused across the batch of N IFMAPS. This type of reuse is called
 filter reuse.

From the obtained tiling parameters, the number of accesses that is used for (51) and (52) is determined by taking into account the data movement strategy as shown in Algorithm 10. As a result, Table 18 summarizes the number of accesses to memory levels for each data type in the forward pass. Therein, $\alpha^v = H^O/H^I$, $\alpha^h = W^O/W^I$, H^O_i/W^O_i define the height/width of tiling OFMAPS in Li buffers, $\alpha^v_i = H^O_i/H^I_i$, and $\alpha^h_i = W^O_i/W^O_i$ for i=2,1, and 0.

Table 18: Numbers of accesses at different memory levels of forward convolution.

Data	DRAM (L3)	L2	L1	L0
$I\left(n_{i}^{I}\right)$	$\left\lceil \frac{M}{M_2} \right\rceil \times \frac{\alpha^v}{\alpha_2^v} \times \frac{\alpha^h}{\alpha_2^h}$	$\left\lceil rac{M_2}{M_1} ight ceil imes rac{lpha_2^v}{lpha_1^v} imes rac{lpha_2^h}{lpha_1^h}$	$\left\lceil \frac{M_1}{M_0} \right ceil imes rac{lpha_1^v}{lpha_0^v} imes rac{lpha_1^h}{lpha_0^h}$	$H^F \times W^F \times \alpha_0^v \times \alpha_0^h$
$F(n_i^F)$	1	$\left\lceil \frac{N}{N_2} \right\rceil \times \left\lceil \frac{H^O}{H_2^O} \right\rceil \times \left\lceil \frac{W^O}{W_2^O} \right\rceil$	$\left\lceil \frac{N_2}{N_1} \right ceil imes \left\lceil \frac{H_2^O}{H_1^O} \right ceil imes \left\lceil \frac{W_2^O}{W_1^O} \right\rceil$	$\left\lceil \frac{N_1}{N_0} \right ceil imes \left\lceil \frac{H_1^O}{H_0^O} ight ceil imes \left\lceil \frac{W_1^O}{W_0^O} ight ceil$
$O(n_i^O)$	1	1	1	1

E.3.4 Backward

For the backward pass, given that $\partial \text{Loss}/\partial O$ is backpropagated from the downstream, it consists in computing $\partial \text{Loss}/\partial F$ and $\partial \text{Loss}/\partial I$. Following the derivation of backpropagation in CNNs by [118], it is given that:

$$\partial \text{Loss}/\partial F = \text{Conv}(I, \partial \text{Loss}/\partial O),$$
 (53)

$$\partial \operatorname{Loss}/\partial I = \operatorname{Conv}(\operatorname{Rot}_{\pi}(F), \partial \operatorname{Loss}/\partial O),$$
 (54)

where $\mathrm{Rot}_\pi(F)$ is the filter rotated by 180-degree. As a result, the backward computation structure is also convolution operations, hence follows the same process as detailed above for the forward pass. For instance, Table 19 summarizes the number of accesses at each memory level in the backward pass when calculating the gradient $G^I = \partial \mathrm{Loss}/\partial I$. Therein, C_i defines the number of tiling IFMAPS in $\mathrm{L}i$ buffer, $\beta^v = H^I/H^O$, $\beta^h = W^I/W^O$, $\beta^v_i = H^I_i/H^O_i$, and $\beta^h_i = W^I/W^O_i$ for i=2,1, and 0.

Table 19: Numbers of accesses at different memory levels for $\partial \text{Loss}/\partial I$.

Data	DRAM (L3)	L2	L1	L0
$O(n_i^O)$	$\left\lceil \frac{C}{C_2} \right\rceil \times \frac{\beta^v}{\beta_2^v} \times \frac{\beta^h}{\beta_2^h}$	$\left\lceil \frac{C_2}{C_1} \right\rceil imes \frac{\beta_2^v}{\beta_1^v} imes \frac{\beta_2^h}{\beta_1^h}$	$\left\lceil \frac{C_1}{C_0} \right\rceil \times \frac{\beta_1^v}{\beta_0^v} \times \frac{\beta_1^h}{\beta_0^h}$	$H^F \times W^F \times \beta_0^v \times \beta_0^h$
$F(n_i^F)$	1	$\left\lceil \frac{N}{N_2} \right\rceil \times \left\lceil \frac{H^I}{H_2^I} \right\rceil \times \left\lceil \frac{W^I}{W_2^I} \right\rceil$	$\left\lceil \frac{N_2}{N_1} \right\rceil \times \left\lceil \frac{H_2^I}{H_1^I} \right\rceil \times \left\lceil \frac{W_2^I}{W_1^I} \right\rceil$	$\left\lceil \frac{N_1}{N_0} \right\rceil \times \left\lceil \frac{H_1^I}{H_0^I} \right\rceil \times \left\lceil \frac{W_1^I}{W_0^I} \right\rceil$
$G^{I}\left(n_{i}^{G^{I}}\right)$	1	1	1	1

F Broader Impacts

Our multidomain comprehensive examination confirms that it is possible to create high performing binary deep neural networks thanks to the proposed method. Our findings suggest the positive impacts in many domains, making deep learning more environmentally friendly, in particular reduce the complexity of huge models like LLMs, and enabling new applications like online, incremental, on-device training, and user-centric AI models. Given the prevalence of LLMs, our approach can facilitate faster predictions on more affordable devices, contributing to the democratization of deep learning. On the other hand, computing architectures have been so far pushed far from its native logic by high-precision arithmetic applications, e.g., 16-bit floating-point is currently a most popular AI computing architecture. Boolean logic deep learning would motivate new software optimization and hardware accelerator architectures in the direction of bringing them back to the native Boolean logic computing. The proposed mathematical notion and its calculus could also benefit other fields such as circuit theory, binary optimization, etc.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The contributions claimed in the abstract and introduction are fully based on the method described in § 3 and performance gains benchmarked in § 4.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss limitations, in particular, of using our method in the current computing architectures in the paragraph "Limitations" in § 5, and provide details about the convergence assumptions in Appendix A.2.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The assumption of each theoretical result is explicty and formally stated. However, due to space constraint, all the proofs are postponed to Appendix A.2, and in particular the long list of convergence analysis assumptions are also reported in the appendix due to the space constraint.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All the training details are described in § 4, in addition the utilized training regularization techniques are also fully described in Appendix C. Python code implementation is also provided in Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in

some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We used publicly common datasets in our experiments. In Appendix we provide example codes and experimental details for reproduction purpose.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All details can be found in § 4, training regularization techniques in Appendix C and the code in Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Standard deviation are provided (when computationally doable) in § 4.

Guidelines

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide information on the computer resources in Appendix D and Appendix E.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: It has been and continues to be our first priority to be fully compliant with the scientific research spirit, and with NeurIPS Code of Ethics in particular.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: See § 5.

Guidelines:

• The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We used publicly common datasets and code package. We fully cited the original paper that produced those assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]
Justification: NA
Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]
Justification: NA
Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]
Justification: NA
Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.