Towards Explainable Bit Error Tolerance of Resistive RAM-Based Binarized Neural Networks

Sebastian Buschjäger*‡, Jian-Jia Chen†, Kuan-Hsun Chen†, Mario Günzel†‡, Christian Hakert†,
Katharina Morik*, Rodion Novkin†, Lukas Pfahler*‡, and Mikail Yayla†‡
*Artificial Intelligence Group, TU Dortmund University, Germany
†Design Automation for Embedded Systems Group, TU Dortmund University, Germany
‡These authors contributed equally to this research
{sebastian.buschjaeger, jian-jia.chen, kuan-hsun.chen, mario.guenzel, christian.hakert,
katharina.morik, rodion.novkin, lukas.pfahler, mikail.yayla}@tu-dortmund.de

arXiv:2002.00909v1 [cs.LG] 3 Feb 2020

***Abstract*—Non-volatile memory, such as resistive RAM (RRAM), is an emerging energy-efficient storage, especially for low-power machine learning models on the edge. It is reported, however, that the bit error rate of RRAMs can be up to 3.3% in the ultra low-power setting, which might be crucial for many use cases. Binary neural networks (BNNs), a resource efficient variant of neural networks (NNs), can tolerate a certain percentage of errors without a loss in accuracy and demand lower resources in computation and storage. The bit error tolerance (BET) in BNNs can be achieved by flipping the weight signs during training, as proposed by Hirtzlin et al., but their method has a significant drawback, especially for fully connected neural networks (FCNN): The FCNNs overfit to the error rate used in training, which leads to low accuracy under lower error rates. In addition, the underlying principles of BET are not investigated. In this work, we improve the training for BET of BNNs and aim to explain this property. We propose straight-through gradient approximation to improve the weight-sign-flip training, by which BNNs adapt less to the bit error rates. To explain the achieved robustness, we define a metric that aims to measure BET without fault injection. We evaluate the metric and find that it correlates with accuracy over error rate for all FCNNs tested. Finally, we explore the influence of a novel regularizer that optimizes with respect to this metric, with the aim of providing a configurable trade-off in accuracy and BET.**

I. Introduction

In the age of ubiquitous computing, sensors and computing facilities are embedded into various physical environments for data collection. Small devices apply machine learning models on data streams directly on the edge. Since the edge devices have resource constraints, such as in computation and storage, these models need to be efficient in execution and memory usage. Binary neural networks (BNNs) are one resource-efficient variant of neural networks (NNs), which are especially well suited for small embedded devices. Their weight parameters are stored as binary values, and the convolution operations are computed with XNOR followed by population count (POPCOUNT) instructions, which count the number of set bits. The trade-off for resource-efficiency is a decrease in accuracy by a few percentage points compared to full-precision neural networks. The efficient execution of BNNs has been researched in several recent works [1], [2], [3], but the memory type to use for BNN models in a low-power setting has received limited attention so far, despite the energy saving potential.

Non-volatile memories (NVMs), such as resistive random-access memory (RRAM), are emerging memory technologies for low-power storage. They are expected to be deployed in future computing systems with resource constraints [4]. Because of their non-volatility, they make normally-off computing efficient: The device is only powered on if there is computation to be done. This is especially convenient for inference on the low-power edge. RRAM, which stores information in the form of non-volatile resistive states, has comparable performance to DRAM, but uses less energy because no refreshs are needed and the read/write energy is lower [4]. Moreover, when using an ultra low-power setting for RRAM cell programming, the energy consumption can be lowered further, up to 30 times for the programming energy, as reported by Hirtzlin et al. [5]. This addresses one the major disadvantages of RRAMs: Cell lifetime. The lower programming energy stresses the cells less, which leads to increased lifetime.

The crucial drawback of the ultra low-power RRAM setting is, however, the high bit error rate of $\sim 3.3\%$. Hirtzlin et al. propose to use this setting for in-memory processing of BNNs, i.e. executing the BNN operations inside the memory, and show that BNNs can be trained to be more error tolerant, up to a rate of 4% without a significant accuracy drop. The increased bit error tolerance (BET) lowers the requirements on the memory and makes possible the use of the ultra low-power setting for BNNs. RRAM is therefore a highly promising low-power memory for BNNs.

The method proposed by Hirtzlin et al. [5] is simple: During training a certain percentage of bit errors are introduced into the network by randomly flipping weights. Even though this method is very effective, it has several drawbacks. First, the accuracy drop is considerable, especially for fully connected neural networks (FCNN). In their experiments, the FCNNs adapt to the error rates they were trained for, which means

This paper has been supported by Deutsche Forschungsgemeinschaft (DFG), as part of the Collaborative Research Center SFB 876 "Providing Information by Resource-Constrained Analysis", project A1 (http://sfb876.tu-dortmund.de) and Project OneMemory (CH 985/13-1).

that the accuracy of the BNNs drop in cases in which a lower percentage of errors is present. Secondly, faults have to be injected during training, which adds complexity to the training process. As discussed in [6], training NNs for general BET (i.e. without accuracy drops for lower error rates) is not a trivial task. The explanation of the underlying principles of the BET of BNNs is not explicitly studied in the literature as well.

In this work, we report on our progress and future directions for *achieving general bit error tolerance* and for *explaining* this property:

- We improve the BET training of the previous work by using straight-through gradient approximation, so that the NNs do not overfit to the error rates used during training.
- We present a metric that aims to measure the achieved BET of BNNs, without injecting faults.
- Based on this metric, we explore the impact of a regularizer with the goal of achieving general BET, which is a property independent from the error model.

The paper is organized as the following. Section II introduces BNNs formally. Section III formalizes the BET of NNs, whereas Section IV presents a novel regularizer to enhance the BET of NNs. Section VI surveys the related work, whereas Section V presents our experiments. Section VII concludes the paper.

II. Binarized Neural Networks

To train a binarized neural network with weights in $\mathbb{F}_2 = \{-1, +1\}$ we follow the approach presented by Hubara et al. [7] in which the weights are stored as floating point numbers, but both weights and activations are deterministically rounded to $\mathbb{F}_2$ during forward computation. The gradient updates are performed with full precision on the floating point weights.

A. Notation

Before we describe the training procedure in more detail, we introduce the notation used to describe neural networks. We assume a feed-forward network in which each layer l is associated with a weight tensor W^l. Each layer performs a generic operation $\circ$ to compute its output $h^l(X) := W^l \circ X$ given its input tensor X^{l-1} and weights in W^l. For example, a fully-connected layer computes the matrix product $h^l(X^{l-1}) = W^l X^{l-1}$ or the convolution layer computes a convolution with a number of filters defined by W denoted by $h^l(X^{l-1}) = W^l * X^{l-1}$. Between layers we apply an activation function $\sigma(h^l(X^l))$ in order to obtain a non-linear decision function. In BNNs it is common to use the sign function as activation function.

B. Training

Floating point networks are typically trained with gradient-based approaches such as mini-batched stochastic gradient descent (SGD) to minimize a loss function. Let $\mathcal{D} = \{(x_1, y_1), \ldots, (x_I, y_I)\}$ with $x_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$ denote the training data and let $\ell\colon \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ be the loss function. Let $W = (W^1, \ldots, W^L)$ denote the weight tensors of each layer in the neural network and let $f_W(x)$ be the output of the network given its weights W, then we aim to solve the following optimization problem

$$\arg\min_W \frac{1}{I} \sum_{(x,y)\in\mathcal{D}} \ell(f_W(x), y)$$

by a gradient descent strategy that computes the gradient $\nabla_W \ell$ using backpropagation. Unfortunately, in the case of binary neural networks we cannot perform gradient-based optimization directly. This is due to two reasons: First, the space of weights is discrete and thus the parameter-vector obtained by taking a small step in the opposite direction of the gradient is almost certainly not binary. Second, the sign-function is not differentiable and also its sub-differential is useless for optimization, as it is zero everywhere other than zero.

Algorithm 1 Binarized forward pass $f_W(x)$.

```
function FORWARD(model, x)
    for l ∈ {1, ..., L} do
        x ← B(B(W^l) ∘ x)
    return x
```

To mitigate the first problem, Hubara et al. [7] propose a scheme that during training stores weights as floating point numbers constrained to values between -1 and 1 and 'binarizes' the network during the forward pass. More formally, let $b\colon \mathbb{R} \to \mathbb{F}_2$ be a binarization function with

$$b(x) = \begin{cases} 1 & x > 0 \\ -1 & \text{else} \end{cases}$$

and let $B(W^l)$ denote the element-wise application of b to W^l. We summarize the forward-pass in Algorithm 1. The outer application of B in Algorithm 1 acts as the activation function and that b can also be interpreted as an un-smooth version of the $tanh$ activation function. Therefore it is sometimes called hard-tanh or $Htanh$.

Then, during the backward pass they use full floating point precision. To mitigate the second problem – b is not differentiable – they replace the gradient of b with the so-called straight-through estimator. Consider the forward computation $Y = B(X)$. Let $\nabla_Y \ell$ denote the gradient with respect to Y. The straight-through estimator approximates

$$\nabla_X \ell := \nabla_Y \ell, \tag{1}$$

essentially pretending that B is the identity function. Using this gradient approximation we can apply standard stochastic gradient descent techniques with the small addition that all floating point weights are clipped to be between -1 and 1 after each gradient update.

For faster and more reliable training, we use the standard deep learning technique Batch Normalization. We insert batch normalization layers between layers and the following activation functions. The batch normalization layers shift and scale the outputs computed by the respective layers, then the sign function is applied. Hence the forward pass can

still be computed using only binary arithmetics, the batch normalization just shifts the threshold of the binary activation from zero to a data-dependent number. One peculiarity of our models is that we apply normalization also after the last linear layer, before the outputs are fed into a softmax-layer with subsequent cross-entropy loss. While this seems counter-intuitive at first, we find that it improves loss and eases training substantially. We suspect that this is due to the rescaling of outputs: Plain binary networks output large integer activations on the last layer which, fed into softmax activations, often result in vanishing gradients.

III. Bit Error Tolerance of BNNs

To understand the error tolerance of BNNs we propose to formalize it using a metric that is calculated on the neuron level with only one pass over the evaluation set. To do so we focus our efforts on CNNs. Please note that all of our definition are also applicable for FCNN if we view their inputs as 1×1 images. We first define the local error tolerance of a $2d$ feature map in a CNN. Then we leverage this definition into the error tolerance of a single neuron, which finally enables us to define the error tolerance of the whole network.

Consider a CNN and let n be the index of one neuron. Recall that the output of a neuron is a $2d$ feature map with height U and width V. We define the neuron's *local error tolerance* $T_{i,n,u,v}$ at position u, v by modeling the number of weight sign flips it can tolerate without a change of its output given the input x_i. To do so, let $h_{i,n,u,v}$ be the output of n-th neuron *before* applying the activation function. For neurons which are not in the first layer, we note two things: First, each neuron's output is computed by a weighted sum of inputs that are ± 1 with weights that are also ± 1. Second, the sign function is applied to this output. Thus as long as weight flips do not change the sign of the weighted sum, a neuron is error tolerant. Formally, we can quantify the error tolerance of a neuron given the input x_i by the distance of its output from 0:

$$T_{i,n,u,v} = \left| h_{i,n,u,v} - s_n - \tfrac{1}{2} \right|. \tag{2}$$

We include s_n to account for activation shifts due to the Batch Normalization Layer (without BatchNorm $s_n = 0$), and to avoid ambiguity at 0 we subtract $\frac{1}{2}$. Note that $T_{i,n,u,v}$ is a measure for the worst case error tolerance, in the sense that at least $\lfloor \frac{T_{i,n,u,v}}{2} \rfloor + 1$ weight sign flips, are necessary. With each weight sign flip $h_{i,n,u,v}$ can get closer to s_n and finally flip the output.

The definition of $T_{i,n,u,v}$ yields the following theorem:

Theorem 1. *Let $b \in \mathbb{R}_{\geq 0}$. If $T_{i,n,u,v} \geq b$ for all u, v then the neuron can tolerate at least $\lfloor \frac{b}{2} \rfloor$ bitflips, i.e. any bitflip of $\lfloor \frac{b}{2} \rfloor$ weights of the neuron does not affect its output.*

The proof can be found in the appendix.

Intuitively, a neuron is error tolerant if it is robust across all positions. Thus, we may demand that each position has a local error tolerance of at least b. More formally, the *error tolerance* $T^b_{i,n}$ of a neuron n given the input x_i is defined as:

$$T^b_{i,n} = \frac{1}{UV} \sum_{u=1}^{U} \sum_{v=1}^{V} \mathbb{1}\{T_{i,n,u,v} \geq b\}. \tag{3}$$

The error tolerance of the whole network can then be defined as the average error tolerance across all neurons:

$$T^b_i = \frac{1}{N} \sum_{n=1}^{N} T^b_{i,n}. \tag{4}$$

We determine T^b for a network by evaluating the BNN on the full data set:

$$T^b = \frac{1}{I} \sum_{i=1}^{I} T^b_i. \tag{5}$$

For neurons in the first layer, we assume that the inputs are not in $\{\pm 1\}$ but $\{0, \ldots, Z\}$. Thus we have to scale the local error tolerance:

$$T_{i,n,u,v} = \frac{\left| h_{i,n,u,v} - s_n - \frac{1}{2} \right|}{Z}. \tag{6}$$

With the definition of T^b, we aim to explain the robustness of BNNs against bit errors without fault injection.

IV. Training Bit Error Tolerant Neural Networks

In this section we propose two different ways to regularize the BNN training objective to account for bit errors during training and achieve bit error tolerance. The first approach is based directly on the insights in Section III, and the second one is based on flip-training as proposed by Hirtzlin et al. [5].

A. Direct Regularization

As discussed in Section III, a high T^b-value for a neuron indicates that many weight signs can flip without changing the activation of the neuron. The quantity T^b is not a differentiable function but essentially a count. However, we can still construct a regularizer that punishes those neurons that do not have a flip-tolerance of at least b. We rely on the well-known hinge function to build a convex and sub-differentiable regularizer. For a given bit-flip tolerance level b, we propose to regularize each neuron n for each input example i using the hinge-function

$$R^b_{n,u,v}(x_i) = \max(0, b - T_{i,n,u,v}). \tag{7}$$

Whenever a neuron has a $T_{i,n,u,v}$-value of at least b, the minimum of R is achieved. To regularize the whole network, we compute the mean of all neuron regularizers. We weight the regularizer with $\lambda > 0$ and add it to the loss.

B. Flip Regularization

Flip regularization is a technique first proposed by Hirtzlin et al. [5]. The idea is simple: To make the network robust against bit errors, we simulate those errors already during training time. During each forward-pass computation, we generate a random bitflip-mask and apply it to the binary weights. However, there are two ways to implement this. Let M denote a random bitflip mask with entries ± 1 of the same

Name	# Train	# Test	# Dim	# classes
FashionMNIST[1]	60000	10000	(1,28,28)	10
CIFAR10	50000	10000	(3,32,32)	10

TABLE I: Datasets used for experiments.

Parameter	Range
Regularization	$\lambda \in \{10^{-4}, 10^{-3}, 10^{-2}\}$
Flip probability	$p \in \{0.01, 0.05, 0.1, 0.2\}$
Robustness	$b \in \{32, 64, 128\}$
Fashion FCNN	In $\to$ FC 2048 $\to$ FC 2048 $\to$ 10
Fashion CNN	In $\to$ C64 $\to$ MP 2 $\to$ C64 $\to$ MP 2 $\to$ FC2048 $\to$ FC2048 $\to$ 10
CIFAR10 CNN	In $\to$ C128 $\to$ C128 $\to$ MP 2 $\to$ C256 $\to$ C256 $\to$ MP 2 $\to$ C256 $\to$ C256 $\to$ MP 2 $\to$ FC 2048 $\to$ FC 2048 $\to$ 10

TABLE II: Parameters used for experiments.

size as W that we multiply component-wise to the binarized weights. We first consider computing the bit-flip operation as $H = (B(W) \cdot M) \circ X$. Standard backpropagation on a loss ℓ that is a function of H yields the following gradient of ℓ with respect to $B(W)$

$$\nabla_{B(W)}\ell = M \cdot \nabla_{B(W)\cdot M}\ell$$

which e.g. for fully connected layers amounts to a gradient update

$$\nabla_{B(W)}\ell = M \cdot (\nabla_H \ell \; X^T).$$

We see that an update computed this way is aware of the bit-flips that were performed and accounts for them. We propose instead to use a special flip-operator with straight-through gradient approximation. We denote by e_p the bit error function that flips its input with probability p and let E_p denote its component-wise counterpart. During training we change the forward pass such that it computes

$$X^{l+1} := B(E_p(B(W^l)) \circ X^l).$$

We replace the gradient of E_p with a straight-through approximation as in (1). This way, in the example above we now have $H = E_p(B(W)) \circ X$ with gradient updates $\nabla_{B(W)}\ell = \nabla_{E_p(B(W))}\ell$ which for fully connected layers yields the update

$$\nabla_{B(W)}\ell = \nabla_H \ell \; X^T$$

which is unaware of bit-flips and just uses the corrupted outputs H.

We believe that the approach using straight-through gradient approximation is superior and that the problems reported by Hirtzlin et al. [5] can be sourced to them using the native implementation. Particularly as we will see in Section V, our implementation does not overfit to a particular error probability.

V. Experiments

In this section we present our experiment results. We evaluate fully connected neural networks (FCNNs) and convolutional neural networks (CNNs) in the configurations shown in Table II FashionMNIST and CIFAR10 (see Table I). In all experiments we run the Adam optimizer for 100 epochs for FashionMNIST and 250 epochs for CIFAR10 to minimize the cross entropy loss. We use a batch size of 128 and an initial learning rate of 10^{-3}. To stabilize training we exponentially decrease the learning rate every 25 epochs by 50 percent. All experiments are repeated 5 times. First, we plot the accuracy over bit error rate for NNs trained with straight-through gradient approximation in the top row of Figure 1. Then, we show the correlation between T^b and accuracy over bit error rate in the bottom row of Figure 1. Finally, we evaluate the impact of our proposed direct regularizer on the error tolerance in Figure 2.

We notice that flip regularization improves the accuracy when bit errors are introduced. This effect is stronger for FCNNs. Moreover we see that we can trade a high accuracy at small error rates with a high accuracy at larger error rates. However we do not observe an overfitting to a particular bit error probability. Second, in the case of FCNNs trained on FashionMNIST and CNNs trained on CIFAR10, we observe that the accuracy over different error rates indeed correlates with T^b. For the case of CNNs on FashionMNIST, a correlation cannot be observed. Overall we see that CNNs are more brittle than FCNNs. This is likely due to the weight-sharing in CNNs, where a flip in a convolution filter has effects at every position in the feature map. This difference is also reflected in the T^b values. In conclusion we see that the T^b measure is better suited for fully connected networks.

Figure 2 depicts the results for the direct regularization training introduced in Section IV-A. We observe that this training method does not increase the accuracy over error rate, although the T^b values are high. Instead regularizing the training objective this way decreases accuracy at any error rate. Similar curve progressions can be observed for other hyperparameter settings and the CIFAR10 dataset. For smaller regularization scalings λ, the observed curves approach the unregularized curves, however we never obtain higher accuracies at any error rate. We conclude that our regularizer is currently unusable: While it effectively increases T^b, it does so by sacrificing accuracy thereby rendering the resulting models useless.

VI. Related Work

Deep Nets offer remarkable performance in state of the art image classification tasks, but require immense computation power during training and during inference. Thus, a natural research question in this context is to ask, whether we can reduce the computation and memory requirements of Deep Nets without hurting its performance. A common approach to reduce both, memory and computation demands, is to quantize the weights of an already trained network after training is completed. This way, weights can be stored using fewer bits and fixed point arithmetics can be exploited during inference. However, this post-processing step usually degrades the classification performance, which leads to sub-optimal performance [8], [9], [10]. More evolved approaches

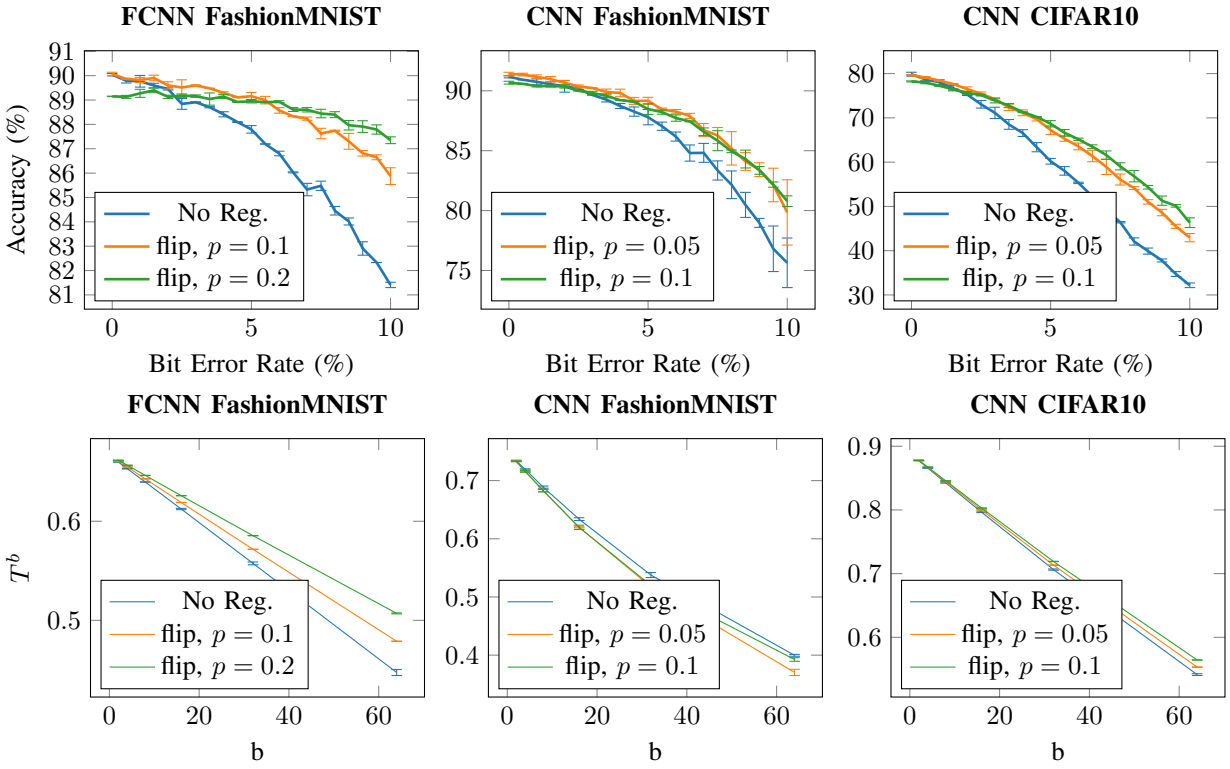

Fig. 1: The experiment results for our proposed flip-training. In the top row are the accuracies plotted over bit error rate. In the bottom row are the T^b values plotted over b.

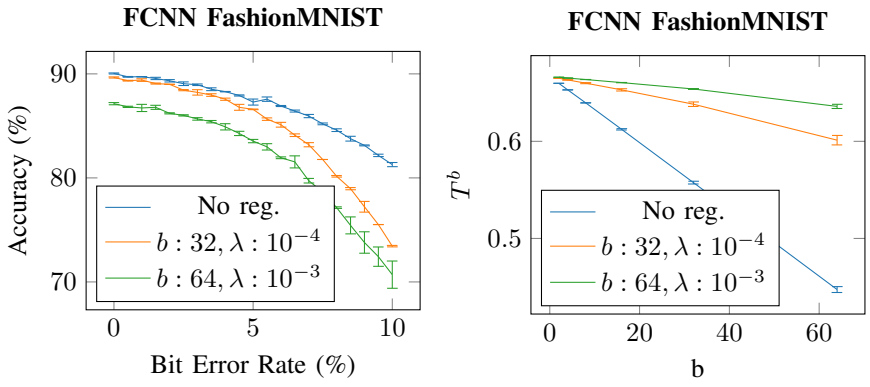

Fig. 2: The experiment results for the direct regularization.

incorporate quantization directly into the training, so that nets can retain their accuracy. Here two approaches exists:

The first approach aims to perform all operations during training (including gradient computation) with fixed point arithmetics. This way, the network is always restricted to fixed point values and efficient accelerations of the training is enabled by the means of FPGAs and GPUs. However, such an approach must guarantee a certain numerical precision so that gradient updates are still meaningful and it has to make sure that gradient estimates during training are still unbiased [11]. The second approach only quantizes the network during the forward pass, but performs all gradient operations with full floating point precision. This way, gradient updates can be performed with full floating point precision, while the networks performance is based on its fixed-point weights. In its most extreme version this approach restricts all intermediate calculations and weights to only two values, e.g. '+1' and '-1'[12], [7]. Since gradient computations are still performed with floating-point precision, this approach still enables regular optimization with stochastic gradient descent and possible regularization, e.g. to enhance the robustness of neural networks against bit errors.

Nonetheless, the error tolerance training of BNNs for low-power memories has not received much attention yet. Recent

work related to BNNs on NVMs focuses more on the realization of the low-power in-memory processing of BNNs than on the error tolerance training aspect. E.g., in [13] Hirtzlin et al. propose to use Spin Torque Magnetoresistive RAM (ST-MRAM) for the in-memory implementation of NNs. In their work they highlight the inherent bit error tolerance of BNNs that were trained without robustness training. They show that half the energy can be saved without accuracy loss when using a low power setting to write to the memory cells.

VII. Conclusion

In this work, we improved the state-of-the-art bit error tolerance training for BNNs and evaluated a metric that aims to explain the achieved tolerance. We were able to eliminate the NNs' overfitting to the error rates by employing a special flip-operator with straight-through gradient approximation in the gradient computation. For BNNs trained with our flip-regularization, we evaluated the robustness metric and found that it correlates with accuracy over error rate for all FCNNs tested. CNNs trained on FashionMNIST with our improved flip-regularization do not show a high robustness value T^b; we hypothesize this is because of the weight sharing property of CNNs and their more complex layer structure.

We also tried to optimize the NNs with respect to the robustness metric T^b. Although we can achieve high T^b values, it does not lead to a better accuracy over error rate. We think that this is mainly due to regularizing each neuron equally in our model. This ignores the effect that a highly robust second layer can compensate low robustness of the first layer.

In the future, we aim to improve the robustness metric T^b, so that error tolerance is better described by it. In the flip-regularization, only the error rate can be configured, therefore we also aim to improve our direct regularization method, so that the error tolerance can be more finely tuned, e.g. with a configurable trade-off between accuracy and bit error tolerance.

References

[1] M. Courbariaux and Y. Bengio, "Binarynet: Training deep neural networks with weights and activations constrained to +1 or -1," *CoRR*, vol. abs/1602.02830, 2016.

[2] S. Mehta, M. Rastegari, L. G. Shapiro, and H. Hajishirzi, "Espnetv2: A light-weight, power efficient, and general purpose convolutional neural network," *CoRR*, vol. abs/1811.11431, 2018.

[3] H. Yang, M. Fritzsche, C. Bartz, and C. Meinel, "Bmxnet: An open-source binary neural network implementation based on mxnet," in *Proceedings of the 25th ACM International Conference on Multimedia*, MM '17, (New York, NY, USA), pp. 1209–1212, ACM, 2017.

[4] J. Boukhobza, S. Rubini, R. Chen, and Z. Shao, "Emerging nvm: A survey on architectural integration and research challenges," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 23, pp. 14:1–14:32, 2017.

[5] T. Hirtzlin, M. Bocquet, J. . Klein, E. Nowak, E. Vianello, J. . Portal, and D. Querlioz, "Outstanding bit error tolerance of resistive ram-based binarized neural networks," in *2019 IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, pp. 288–292, 2019.

[6] C. Torres-Huitzil and B. Girau, "Fault and error tolerance in neural networks: A review," *IEEE Access*, vol. 5, pp. 17322–17341, 2017.

[7] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks," in *Advances in neural information processing systems*, pp. 4107–4115, 2016.

[8] D. Lin, S. Talathi, and S. Annapureddy, "Fixed point quantization of deep convolutional networks," in *International Conference on Machine Learning*, pp. 2849–2858, 2016.

[9] M. Courbariaux, J.-P. David, and Y. Bengio, "Low precision storage for deep learning," *arXiv preprint arXiv:1412.7024*, 2014.

[10] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding," *arXiv preprint arXiv:1510.00149*, 2015.

[11] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan, "Deep learning with limited numerical precision," in *International Conference on Machine Learning*, pp. 1737–1746, 2015.

[12] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "Xnor-net: Imagenet classification using binary convolutional neural networks," in *European Conference on Computer Vision*, pp. 525–542, Springer, 2016.

[13] T. Hirtzlin, B. Penkovsky, J.-O. Klein, N. Locatelli, A. F. Vincent, M. Bocquet, J.-M. Portal, and D. Querlioz, "Implementing binarized neural networks with magnetoresistive ram without error correction," 2019.

Appendix

Proof for Theorem 1. At first we consider the n-th neuron and assume that it is not in the first layer. Let u, v be a position for the convolution result. As described in (2), we know that $T_{i,n,u,v} = |h_{i,n,u,v} - s_n - \frac{1}{2}|$. For improved readability, we write h for $h_{i,n,u,v}$ and s for s_n. By construction of the activation function, the output of the neuron at position u, v is $+1$ for $h - s - \frac{1}{2} > 0$ and -1 for $h - s - \frac{1}{2} < 0$ since h and s are assumed to be integer values. Furthermore, the case $h - s - \frac{1}{2} = 0$ does not occur. If $T_{i,n,u,v} \geq b$ then either $h - s - \frac{1}{2} \geq b$ or $h - s - \frac{1}{2} \leq -b$.

In the first case we have $h - s - \frac{1}{2} \geq b \geq 0$ and the output at u, v is $+1$. We denote by $\tilde{y}$ the value of h after the bitflips of up to $\lfloor \frac{b}{2} \rfloor$ weights. By definition, h is a weighted sum where each summand is one input of the neuron multiplied with one weight. Since each summand is in $\{\pm 1\}$, changing one sign changes h by 2. Therefore $\tilde{h}$ can differ by up to $2 \cdot \lfloor \frac{b}{2} \rfloor$ from h and is in $[h - \lfloor b \rfloor, h + \lfloor b \rfloor]$. For $\tilde{h}$ we still have $\tilde{h} - s - \frac{1}{2} \geq b - \lfloor b \rfloor \geq 0$ which causes an output of $+1$ at u, v.

The second case is proven analogously: We have $h - s - \frac{1}{2} \leq -b \leq 0$ and the output of the neuron at u, v is -1. Changing the sign of $\lfloor \frac{b}{2} \rfloor$ summands of h can increase the value of h by up to $\lfloor b \rfloor$. After at most $\lfloor \frac{b}{2} \rfloor$ bitflips, we obtain a new value $\tilde{h} \in [h - \lfloor b \rfloor, h + \lfloor b \rfloor]$. We have $\tilde{h} - s - \frac{1}{2} \leq -b + \lfloor b \rfloor \leq 0$ and the output of this neuron at u, v is still -1.

If the n-th neuron is in the first layer we have $T_{i,n,u,v} = \frac{|h - s - \frac{1}{2}|}{Z}$ by (6). Since the summands of h are in $[-Z, Z]$, the value of h can change by up to $2 \cdot Z$ per bitflip. Thus, after $\lfloor \frac{b}{2} \rfloor$ bitflips the value of $\tilde{h}$ is in the interval $[h - Z \cdot \lfloor b \rfloor, h + Z \cdot \lfloor b \rfloor]$. But the proof still works as above since $|h - s - \frac{1}{2}| \geq Z \cdot b$ by assumption. □