



Quantized Guided Pruning for Efficient Hardware Implementations of Convolutional Neural Networks

Ghouthi Boukli Hacene, Vincent Gripon, Matthieu Arzel, Nicolas Farrugia, Yoshua Bengio

► To cite this version:

Ghouthi Boukli Hacene, Vincent Gripon, Matthieu Arzel, Nicolas Farrugia, Yoshua Bengio. Quantized Guided Pruning for Efficient Hardware Implementations of Convolutional Neural Networks. 2018. hal-01965304

HAL Id: hal-01965304

<https://hal.science/hal-01965304>

Preprint submitted on 25 Dec 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Quantized Guided Pruning for Efficient Hardware Implementations of Convolutional Neural Networks

Ghouthi Boukli Hacene^{1,2}, Vincent Gripon^{1,2}, Matthieu Arzel², Nicolas Farrugia² and Yoshua Bengio¹

¹ Université de Montréal, MILA, ²IMT Atlantique, Lab-STICC

(Paper submitted to ISCAS 2019 on October 31, 2018)

Abstract—Convolutional Neural Networks (CNNs) are state-of-the-art in numerous computer vision tasks such as object classification and detection. However, the large amount of parameters they contain leads to a high computational complexity and strongly limits their usability in budget-constrained devices such as embedded devices. In this paper, we propose a combination of a new pruning technique and a quantization scheme that effectively reduce the complexity and memory usage of convolutional layers of CNNs, and replace the complex convolutional operation by a low-cost multiplexer. We perform experiments on the CIFAR10, CIFAR100 and SVHN and show that the proposed method achieves almost state-of-the-art accuracy, while drastically reducing the computational and memory footprints. We also propose an efficient hardware architecture to accelerate CNN operations. The proposed hardware architecture is a pipeline and accommodates multiple layers working at the same time to speed up the inference process.

Index Terms—convolutional neural networks, pruning, weight binarization, hardware implementation.

I. INTRODUCTION AND RELATED WORK

For the past few years, Deep Neural Networks (DNNs), and especially Convolutional Neural Networks (CNNs) [1], have received considerable attention thanks to their remarkable accuracy in computer vision tasks [2]–[5] such as classification and detection [6]. However, their need for intensive computations and memory has meant that most of the implementations are based on GPUs, while providing efficient hardware implementations is still a very active and prospective field of research. Therefore, the deployment of CNNs in embedded systems is complex and potentially blocking for many potential applications.

To address this limitation, multiple approaches have been proposed in the literature. For example the authors of [7] and [8] propose to reduce DNNs’ memory footprint by compressing their weights. In these two approaches, the obtained DNN is not retrained after compression, leading to potentially sub-optimal solutions. Following this lead, the authors of [9] have showed that training and compressing weights simultaneously can lead to better accuracy. In the same vein, in [10] and [11], the authors propose to binarize the weights during the learning phase. As a result, the obtained DNN contains only weights whose values are 1 or -1 , while suffering from a very limited drop in accuracy compared to state-of-the-art solutions. These works have been improved later in [12], where the authors proposed to add a scaling factor per layer and per kernel, as a mean to offer better diversity

to binary networks, with almost no impact on memory usage or complexity. Other approaches have proposed to limit weights to three or more values (-1 , 0 , and 1) [13]–[15]. These approaches demonstrate that using slightly more bits to encode weights enable to improve accuracy by a significant amount, but they also require much more memory and other hardware components to compute non-binary operations. In recent works [16], [17], authors have proposed to binarize both weights and activations in CNNs, resulting in potentially very efficient hardware implementations. However, these methods end up with a significant lower accuracy than state-of-the-art ones.

Once a binary neural network has been trained, efficient implementations can advantageously benefit from simplified operations. For example, multiplications in binary neural networks can be replaced by simple low-cost multiplexers. Efficient solutions have been proposed in [18], [19]. However, even binary neural networks still require significant computational power and memory. These solutions also typically lead to a considerable latency, which may be an issue for some applications. In another line of work, authors have been aiming at reducing the number of trainable parameters in DNNs. In [7], [20], the authors successfully apply pruning techniques to fully connected layers of DNNs. However, state-of-the-art CNNs are using more and more convolutional layers nowadays: in a typical modern architectures like ResNet18, about 99% of the connections are in convolutional layers, and thus pruning connections only in fully connected layers has almost no impact on the overall complexity and memory usage of the architecture.

In this paper we propose to combine an efficient pruning technique, which can be effectively leveraged at implementation stage, with binary neural networks. We apply the proposed pruning technique on convolutional layers, resulting in very lightweight convolutions that can be implemented with simple multiplexers. The proposed method approaches state-of-art accuracy on the CIFAR10, CIFAR100 and SVHN dataset. We also propose a hardware implementation which uses very few resources and computational power. This implementation can compute more than one layer at a time and uses a simple multiplexer to perform convolutional operations. As such, it provides significantly smaller latency than existing counterparts.

The outline of the paper is as follows: in Section II we describe the proposed method and describe experiments on

the CIFAR10, CIFAR100 and SVHN dataset. In Section III we present the proposed hardware implementation and show hardware implementation results. Section IV concludes.

II. PROPOSED METHOD

In this section, we introduce a method to efficiently prune connections in convolutional layers. Note that pruning may have two different aims: a) to decrease the number of parameters to be trained in a given architecture, thus resulting in lesser chance of overfitting and b) to decrease the memory usage and complexity of a given architecture, so that it becomes lighter to implement in a budget-restricted configuration. If some author (e.g. [20]) argue they do both, we believe this is questionable as the reduction of the number of trainable weights they obtain on the one hand is balanced by the increasing complexity of identifying which connections are kept and which are lost in the process.

The proposed method has the double of interest of decreasing the number of parameters to be trained while keeping a simple deterministic way of identifying which connections are kept and which are disregarded.

A. Details of the Proposed Method

Let us denote by \mathbf{x} (resp. \mathbf{y} or \mathbf{w}) the input (resp. output or kernel) tensor of a given convolutional layer. We index \mathbf{x} (resp. \mathbf{y}) using three indices i, j, k (resp. ℓ), where $0 \leq i < i_{\max}$ and $0 \leq j < j_{\max}$ correspond to 2D coordinates and $0 \leq k < k_{\max}$ (resp. $0 \leq \ell < \ell_{\max}$) indexes a feature map. Similarly, we index \mathbf{w} using four indices: $0 \leq \iota \leq \iota_{\max}$ and $0 \leq \lambda \leq \lambda_{\max}$ correspond to 2D coordinates, and k and ℓ are as introduced above. So, an element of the input tensor is written $x_{i,j,k}$, an element of the kernel tensor is written $w_{\iota,\lambda,k,\ell}$ and an element of the output tensor is written $y_{i,j,\ell}$.

The idea we propose consists of removing most of the connections in each slice $\mathbf{w}_{\cdot,\cdot,k,\ell}$ of the kernel tensor. The connections to be kept are chosen according to a deterministic rule agnostic of the initialization and of the training dataset. Namely, we choose to only keep the connections $w_{\iota,\lambda,k,\ell}$ for which

$$\iota + \lambda_{\max} = k \pmod{\iota_{\max}\lambda_{\max}}. \quad (1)$$

When considering 3×3 kernels for example, we remove 89% of the connections in the convolutional layer. The reason for choosing this scheme is quite straightforward: we want diversity in the connections we keep to be sure our kernels do not simplify to a simple 1×1 convolution and still cover the initial kernel to its full extent (providing at least 9 feature maps are used).

We then perform the training on the remaining connections, disregarding the other ones. Using this method, the convolution of each slice of the kernel tensor is replaced by a simple multiplication.

To further benefit from the reduced complexity of this pruning technique, we combine it with a weight binarization method. Here, we use Binary Connect (BC) [10]. Once remaining connections have been binarized, it is possible to replace the multiplication operation by a multiplexer.

B. Results

To evaluate the performance of our proposed method, we use the CIFAR10 vision benchmark made of tiny 32×32 images. We compare various modern CNN architectures such as Resnet [21], Wide-Resnet [22], Densenet [23], and Mobilenet [24]. Note that these architectures contain 1×1 and 3×3 convolutional kernels only. Thus we apply the proposed method on the 3×3 kernels.

As a first experiment, we aim at estimating the drop in performance caused by pruning connections. We thus randomly remove m connections per kernel slice. Figure 1 shows that the accuracy of the architecture is quite robust to this process, even when 8 out of the 9 connections in slices of 3×3 kernels are removed.

We then report in Table I the obtained results using Equation (1) to remove kernels connections. Note that contrary to the previous experiment, removed connections are not chosen randomly anymore but according to a deterministic scheme. As a consequence, the positions of removed connections does not have to be stored. We compare the accuracy obtained using baseline architectures, pruned ones, binarized ones, and our proposed method mixing pruning and BC. Note that BC offers a 32 compression factor in terms of memory used, and our method roughly multiply this factor by 9, achieving an almost 300 factor compression in total. We also perform experiments on SVHN (resp. CIFAR100) on Resnet18 (resp. WideResnet-40-10) and obtain 97%/96% (resp. 80%/77%) accuracy for Full-precision/pruning+BC.

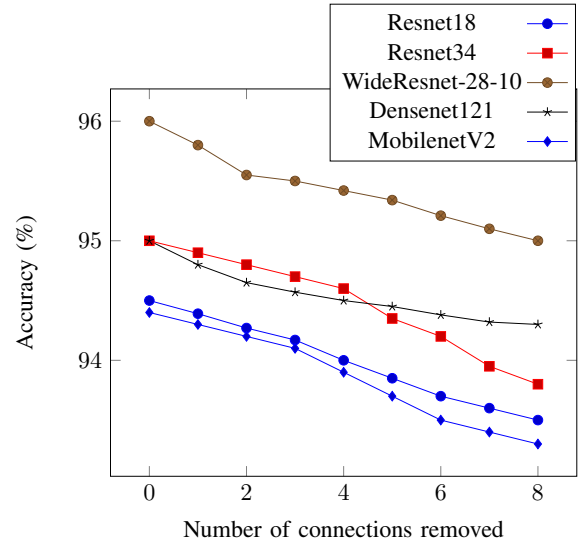


Fig. 1. Evolution of accuracy as function of number of connections removed per kernel slice.

III. HARDWARE IMPLEMENTATION

In this section, we first introduce the hardware architecture of the proposed method, its different components, and the way they are connected. Then, we present the hardware implementation of the proposed method, applied on ResNet18, on a Field Programmable Gate Array (FPGA).

TABLE I
COMPARISON OF ACCURACY BETWEEN BASELINE ARCHITECTURES, PRUNED ONES, BINARIZED ONES, AND THE PROPOSED METHOD ON CIFAR10.

	Resnet18	Resnet34	WideResnet-28-10	Densenet121	MobilenetV2
Full-precision	94.5%	95%	96%	95%	94.4%
Pruning	93.5%	93.8%	95%	94.3%	93.3%
BC	93.31%	93.64%	95.2%	94.5%	93%
Pruning + BC	91%	91.3%	94%	93%	91%

TABLE II
FPGA RESULTS FOR THE PROPOSED ARCHITECTURE ON VU13P (XCVU13P-FIGD2104-1-E).

	P	LUT	FF	BRAMs	Frequency	Processing Latency	Processing outflow	Power
Conv64 – 64	16	22424	22424	114	240MHz	52 μ s	19230 images/s	3.7W
4 \times Conv64 – 64	16	89746	75235	456	240MHz	208 μ s	19230 images/s	6.5W
3 \times Conv128 – 128	32	59780	45024	171	240MHz	154, 8 μ s	19379 images/s	4.8W
3 \times Conv128 – 128	64	134090	102552	171	240MHz	103, 2 μ s	29069 images/s	7.8W
3 \times Conv256 – 256	64	74067	52051	87	250MHz	147, 3 μ s	20366 images/s	5.5W
3 \times Conv256 – 256	128	154599	102723	87	218MHz	112, 8 μ s	26595 images/s	7.8W
3 \times Conv512 – 512	128	132155	52151	45	208MHz	177 μ	16949 images/s	7.9W

A. Hardware Architecture

In Figure 2, we depict the proposed hardware architecture for performing convolutions, which we name a “layer block”. This architecture uses a simple low-cost multiplexer. In more details, a layer block is made of two sub-blocks: a memory one and a processing unit one.

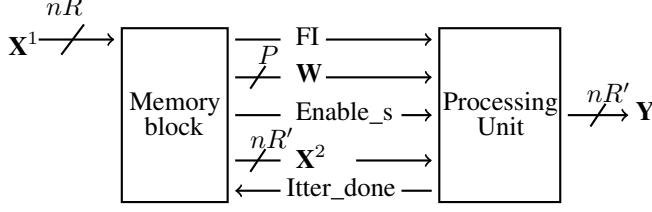


Fig. 2. Hardware architecture of a layer block.

The memory block contains two block RAMs (BRAMs) where content is encoded using n bits fixed point. The first is used to store the computed feature maps. Once they are all computed, the content of the first BRAM is copied to the second one, so that it becomes the input of the next layers. At the same time, the computed feature maps of another image can be stored in the first BRAM. We thus obtain a pipeline architecture, in which all implemented layers work at the same time to speed up the classification process.

To avoid data overflow, we process each row of a slice of the input tensor \mathbf{x} independently, and each slice of the kernel tensor independently. In more details, we copy from BRAM one to BRAM two a feature subvector $\mathbf{X}_{i,k}^2 = \{x_{i,1,k}^2, x_{i,2,k}^2, \dots, x_{i,R',k}^2\}$ made of R' values, instead of the whole feature vector $\mathbf{X}_{i,k}^1 = \{x_{i,1,k}^1, x_{i,2,k}^1, \dots, x_{i,R,k}^1\}$ made of $R > R'$ values (cf. Figure 2). This is to account for the border effects (padding). To simplify notations, we replace $\mathbf{X}_{i,k}^1$ (resp. $\mathbf{X}_{i,k}^2$) by \mathbf{X}^1 (resp. \mathbf{X}^2) in the following.

To compute a convolutional operation, kernels move along feature maps with a step which is called stride in CNNs. In a typical case in which stride value is 1, \mathbf{X}^2 represents either the first R' values, the middle R' ones or the last R' ones, where $R' = R - 2$, depending on the position of the nonzero kernel value. When stride value is 2 (c.f. Figure 4), only half of the values are copied from \mathbf{X}^1 to \mathbf{X}^2 by selecting either the odd or even values of j in $x_{i,j,k}^1$ using multiplexers. This process can be generalised to any stride value other than 1 or 2.

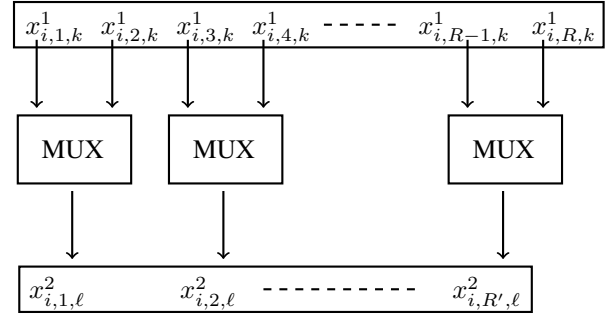


Fig. 4. Hardware architecture to copy the first BRAM contents to the second BRAM, when stride value is 2.

The processing unit uses \mathbf{X}^2 and a vector \mathbf{W} made of P values coded on 1 bit each corresponding to weights in the convolution kernel. It thus computes in parallel P feature vectors (cf. Figure 3). The First-Input signal (FI) is set to 1 when the first feature vector is read from the second BRAM to initialise registers by 0. To compute each feature vector p where $1 \leq p \leq P$, we use the corresponding W_p to add either \mathbf{X}^2 or $-\mathbf{X}^2$ to the content of register p . Once all input feature vectors have been read from the second BRAM of memory block, the signal Enable_s is set to 1, and the content of registers is written one by one into the first BRAM of the

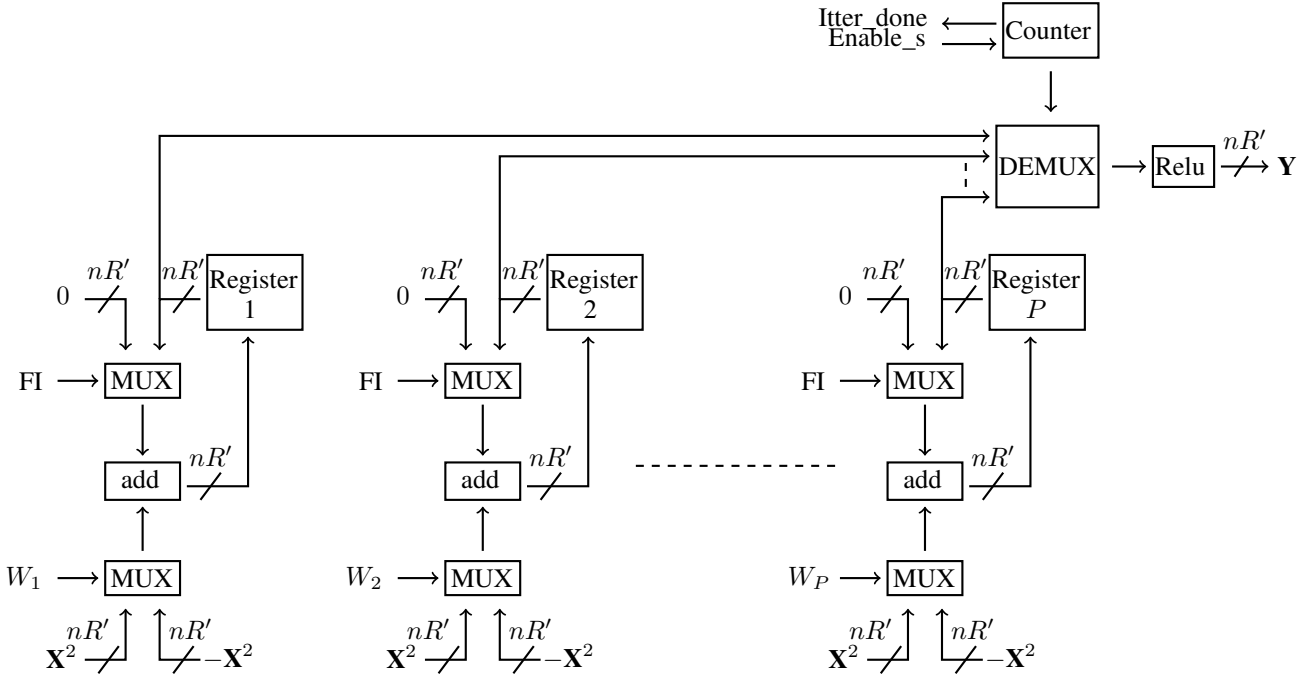


Fig. 3. Hardware architecture of a processing unit block.

memory block of the next layer. At the end of this process, the `Itter_done` signal is set to 1 in the processing unit block, so new data can be read from the memory block to process other feature vectors.

To achieve the computation associated with the layer block described in Figure 2, $k_{max}j_{max}$ clock cycles (CCs) are required to copy all contents from the first BRAM to the second one, $j_{max}k_{max}\ell_{max}/P$ CCs to compute all output feature vectors of one layer, and $j_{max}\ell_{max}$ CCs to write all computed feature vectors into the memory block of the next layer. Thus the total number of CCs required is:

$$CCs = j_{max}k_{max} + \frac{j_{max}k_{max}\ell_{max}}{P} + j_{max}\ell_{max} \quad (2)$$

This should be compared to [19], where the number of clock cycles becomes:

$$CCs = \frac{3j_{max}^2k_{max}\ell_{max}}{P} \quad (3)$$

We observe that the proposed architecture is $3j_{max}$ faster than [19], which can be significant when j_{max} is big. For instance with the CIFAR10 dataset, at the input layer of a CNN $j_{max} = 32$, and thus the proposed method is 96 times faster. In addition it is a pipeline architecture, so it can be $3Lj_{max}$ faster where L is the total number of layer blocks that fit in an FPGA.

Note that in the proposed architecture, P should be lower or equal to ℓ_{max} , otherwise reaching full parallelism would require to read more than one vector \mathbf{X}^2 , and as such would also require more BRAMs, resulting in a more complex architecture.

B. Hardware Results

We implemented one/few layers of Resnet18 on Xilinx Ultra Scale Vu13p (xcvu13p-figd2104-1-e) FPGA. The implemented layers are arranged in a pipeline, and their functionality has been verified comparing the output of each layer block with the ones obtained by software simulation over a batch of examples. Table II shows the required resources to implement one/few layers of Resnet18 trained on CIFAR10 dataset for different values of P . It also shows that the obtained architecture obtain a low processing latency to compute a valid output of one layer. Moreover, this processing latency increases when processing more than one layer, but processing outflow is maintained thanks to the pipeline design.

IV. CONCLUSION

In this paper, we proposed to extend pruning techniques to convolutional layers of DNNs. We introduced a deterministic pruning scheme that can be taken advantage of in implementations. We combined pruning with weight binarization to reduce both complexity and memory usage and showed the resulting neural network is still able to reach very high accuracy.

We implemented the proposed scheme using a low cost hardware architecture in which complex convolution operations are replaced by simple multiplexers. As a result, we were able to implement a considerable part of some complex CNNs such as Resnet18. Moreover, the architecture only consumes a few watts, making it a good solution for embedded applications. Future work will extend this method to all kernel shapes, and propose a low cost hardware architecture to handle other challenging vision datasets such as ImageNet.

REFERENCES

- [1] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [2] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size," *arXiv preprint arXiv:1602.07360*, 2016.
- [3] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [4] B. Graham, "Fractional max-pooling," *CoRR*, vol. abs/1412.6071, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6071>
- [5] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Re-thinking the inception architecture for computer vision," *arXiv preprint arXiv:1512.00567*, 2015.
- [6] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, p. 436, 2015.
- [7] 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.
- [8] Y.-D. Kim, E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin, "Compression of deep convolutional neural networks for fast and low power mobile applications," *arXiv preprint arXiv:1511.06530*, 2015.
- [9] L. Hou, Q. Yao, and J. T. Kwok, "Loss-aware binarization of deep networks," *arXiv preprint arXiv:1611.01600*, 2016.
- [10] 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*, 2015, pp. 3123–3131.
- [11] G. Soulié, V. Gripon, and M. Robert, "Compression of deep neural networks on the fly," in *International Conference on Artificial Neural Networks*. Springer, 2016, pp. 153–160.
- [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*. Springer, 2016, pp. 525–542.
- [13] Z. Lin, M. Courbariaux, R. Memisevic, and Y. Bengio, "Neural networks with few multiplications," *arXiv preprint arXiv:1510.03009*, 2015.
- [14] F. Li, B. Zhang, and B. Liu, "Ternary weight networks," *arXiv preprint arXiv:1605.04711*, 2016.
- [15] S. Zhou, Y. Wu, Z. Ni, X. Zhou, H. Wen, and Y. Zou, "Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients," *arXiv preprint arXiv:1606.06160*, 2016.
- [16] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1," *arXiv preprint arXiv:1602.02830*, 2016.
- [17] W. Tang, G. Hua, and L. Wang, "How to train a compact binary neural network with high accuracy?" in *AAAI*, 2017, pp. 2625–2631.
- [18] R. Andri, L. Cavigelli, D. Rossi, and L. Benini, "Yodann: An ultra-low power convolutional neural network accelerator based on binary weights," in *ISVLSI*, 2016, pp. 236–241.
- [19] A. Ardakani, C. Condo, and W. J. Gross, "A convolutional accelerator for neural networks with binary weights," in *Circuits and Systems (ISCAS), 2018 IEEE International Symposium on*. IEEE, 2018, pp. 1–5.
- [20] —, "Sparsely-connected neural networks: towards efficient vlsi implementation of deep neural networks," *arXiv preprint arXiv:1611.01427*, 2016.
- [21] 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*, 2016, pp. 770–778.
- [22] S. Zagoruyko and N. Komodakis, "Wide residual networks," *arXiv preprint arXiv:1605.07146*, 2016.
- [23] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *CVPR*, vol. 1, no. 2, 2017, p. 3.
- [24] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4510–4520.