



Low-bit Shift Network for End-to-End Spoken Language Understanding

Anderson R. Avila, Khalil Bibi, Ruiheng Yang, Xinlin Li, Chao Xing, Xiao Chen

Huawei Noah's Ark Lab
Montreal, QC H3N 1X9, Canada

{anderson.avila, khalil.bibi, rui.heng.yang1, xinlin.li1, xingchao.ml,
chen.xiao2}@huawei.com

Abstract

Deep neural networks (DNN) have achieved impressive success in multiple domains. Over the years, the accuracy of these models has increased with the proliferation of deeper and more complex architectures. Thus, state-of-the-art solutions are often computationally expensive, which makes them unfit to be deployed on edge computing platforms. In order to mitigate the high computation, memory, and power requirements of inferring convolutional neural networks (CNNs), we propose the use of power-of-two quantization, which quantizes continuous parameters into low-bit power-of-two values. This reduces computational complexity by removing expensive multiplication operations and with the use of low-bit weights. ResNet is adopted as the building block of our solution and the proposed model is evaluated on a spoken language understanding (SLU) task. Experimental results show improved performance for shift neural network architectures, with our low-bit quantization achieving 98.76 % on the test set which is comparable performance to its full-precision counterpart and state-of-the-art solutions.

Index Terms: Spoken language understanding, quantization, edge computing

1. Introduction

With the advent of edge computing [1] and Tiny ML [2], interest in shrinking large deep neural network models has increased. The aim is to reduce computational complexity, allowing their deployment in small devices such as laptops, smartphones, and even less powerful gadgets such as wearables and hearing aids [3]. Deep neural networks (DNNs), in fact, have achieved breakthrough results in several areas. The outstanding performance of state-of-the-art models has been possible primarily with the development of larger and more complex architectures [3][4]. Training these models, however, is very demanding in terms of required computational resources which is associated with high energy consumption and carbon emission as discussed in [5]. This leads to undesirable financial and environmental costs that researchers have recently been trying to alleviate.

While such models have been generally deployed in the cloud, motivated by its computing power and data processing efficiency [1], there has been an increasing interest in alleviating the burden of cloud computing. This has several benefits for real-world applications. For instance, as more and more data is being generated at the edge of the network, processing such data on small devices can be more efficient than processing it in the cloud [1]. Moreover, aside from mitigating data safety and privacy issues, enabling computing on resource-limited devices

allows for latency reduction and minimization of data transmission through the network. This can ultimately benefit Internet of Things (IoT) applications [6].

$$2^{\tilde{p}x} = \begin{cases} x \ll \tilde{p}, & \text{for } \tilde{p} > 0 \\ x \gg \tilde{p}, & \text{for } \tilde{p} < 0, \text{ where } \tilde{p} \in \mathbb{Z} \\ x, \tilde{p} = 0 \end{cases} \quad (1)$$

We are particularly concerned with the deployment of end-to-end spoken language understanding (e2e SLU) systems on edge devices. While e2e SLU models can provide high accuracy in high-level hardware, deploying such models on edge and internet of things (IoT) devices is still a challenge [7, 8]. Inspired by recent publications [9], we investigate a low-bit multiplication-free neural network. Specifically, we explore the ternary shift network introduced in [10]. The method replaces computationally expensive operations (i.e. multiplication) by bitwise shift operations, which mathematically is the same as multiplying by a power of 2 [11], as shown in eq. (1). The authors in [10] also introduced a new training strategy that showed to be more effective than other shift network solutions, such as the one presented in [9], as it mitigates gradient vanishing and weight sign freezing problems in low-bit shift networks.

Although quantization is widely used to replace costly high-precision floating-point with low-precision fixed-point representation [12, 13], it still requires costly multiplication operations. In order to further improve the performance of fixed-point quantization, the proposed ternary shift network decomposes an 3 bits parameter into a sign-sparse-shift 3-fold manner. This way, we attain a low-bit network with weight dynamics that correspond to full-precision networks. As discussed in [11], bit-shifts can save roughly $196 \times$ memory and $24 \times$ energy costs over their multiplication counterpart.

We use ResNet based architecture in our experiments and the models are evaluated on the Fluent Speech Command (FSC) dataset [14]. Compared to the same architecture with 32 bits, as well as to its quantized version with 16, 8 and 3 bits, experimental results show that our solution can provide accuracy as high as 98 % for intent classification, thus showing competitive performance to its full-precision counterpart and quantized versions. To the best of our knowledge, this is the first work proposing a multiplication-free low-bit representation as an end-to-end SLU solution.

This paper is organized as follows. Section 2 presents the related work. In Section 3, we give details about the Ternary




Figure 1: *ResNet-18 based on shift operations.*

shift network. Section 4 introduces the SLU task. In Section 5 we present the experimental setup, followed by the result discussion in Section 6.

2. Related Work

In order to enable the deployment of DNN-based models in resource-constrained scenarios, significant efforts have been made to minimize the computational complexity of neural networks without compromising inference accuracy. Most of these efforts focus on low-precision fixed-point multiplications and to further improve these techniques shift network has been also investigated. In [4], for instance, the authors introduced a binarization scheme called BinaryConnect. The method replaces multiply-accumulate operations with simple accumulations, allowing to train DNN models with binary weights without compromising the precision of stored weights. In another similar work [15], the authors present an efficient and accurate way of approximating CNNs by binarizing the weights. Two quantized CNNs, namely Binary-Weight-Network and XNOR-Network, are investigated. In the former, CNN filters represented with binary values leads to memory reduction in the order of $32\times$. Different from the binary approach, where weights are constrained between two values (e.g., 0 and 1), ternary weight networks (TWN) are introduced in [16]. The weights in this work are constrained to three values: -1, 0, and 1. According to the authors, their model offers higher accuracy and higher model compression rate (i.e., $16\times$, $32\times$), requiring fewer multiplications compared to its full precision version [16]. In a more recent work [9], the authors introduced DeepShift. The solution presents two new operations, namely convolutional shift and fully-connected shift, that minimize the number of multiplications by introducing bitwise shift and sign flipping during both training and inference. During inference, for instance, only 5 bits (or less) are required to represent the weights.

While memory efficient and hardware friendly, training shift networks is challenging. The authors in [10], pointed out two major problems. First is the performance sensitivity to weight initialization which can lead to accuracy degradation. The second problem is the fact that quantizers are non-differentiable functions and gradient approximators may lead to gradient vanishing and weight sign freezing. To mitigate these problems the ternary shift network, proposed in [10], is discussed next.

3. Ternary Shift Neural Networks and S^3 Re-parameterization

Shift neural network [17, 9, 18] is a special fixed-point quantized CNN that restricts weight values to the range of zero or positive/negative power of two numbers, thus $w_{\text{shift}} \in \{0\} \cup \{\pm 2^p\}$. The multiplication between a fixed-point integer and a power of two can be implemented with the bit-shift operator. The bit-shift operator’s power consumption and hardware area cost are significantly lower than that of the multiplication operator under the same precision. Therefore, Shift networks outperform fixed-point quantized CNN in terms of energy consumption and computational efficiency. Gudovskiy et al. [17] show that shift networks can achieve 4x energy saving and 2.5x hardware resources saving compared to fixed-point quantized networks on FPGA.

Despite the competitive inference performance of Shift networks, training has been a challenge. All previous training methods [17, 9, 18] cannot achieve the same accuracy as the full-precision baselines on the ImageNet classification task and require initialization from a pre-trained full-precision checkpoint. Recently, Li et al. [10] proposed S^3 re-parameterization, a new method for training low-bit shift networks to tackle this challenge. This method re-parameterizes the discrete weights of shift networks to the product of multiple binary parameters. The low-bit shift networks trained with S^3 re-parameterization achieve the same accuracy as regular full-precision CNN in the ImageNet classification task with *ResNet* based architecture. Moreover, unlike previous methods, the S^3 method achieves high accuracy without using complex initialization or training strategies.

In this work, the weight value range limits to $\{0, \pm 1, \pm 2, \pm 4\}$. With S^3 training method, the low-bit discrete weights are re-parameterized into one sign parameter w_{sign} , one sparse parameter w_{sparse} , and two shift parameters w_{s1} , w_{s2} . During forward propagation, all parameters are binarized by the Heaviside function $\mathbb{1}(\cdot)$, and calculate the discrete 3-bit shift network weight values based on eq. (2). During backward propagation, the straight-through estimator [19] is utilized to approximate the derivative of the Heaviside function.

$$w_{\text{shift}} = 2^S \mathbb{1}(w_{\text{sparse}}) \{2\mathbb{1}(w_{\text{sign}}) - 1\} \quad (2)$$

$$S = \mathbb{1}(w_{s2}) \{ \mathbb{1}(w_{s1}) + 1 \}$$

Following the S^3 training method, the dense weight regu-




Figure 2: 3-bit shift network with S3 re-parameterization.

larizer applies to the sparse parameter $\mathbf{w}_{\text{sparse}}$ during training.

$$\mathcal{R}_{\text{sparse}}(\mathbf{w}_{\text{sparse}}) = \|\max(-\mathbf{w}_{\text{sparse}}, 0)\|_1 \quad (3)$$

4. Spoken Language Understanding

Voice command recognition enables users to express their intentions via their voices. This is important for modern AI human computer interaction as it allows voice control of smart homes, smart speakers, and other devices such as phones and tablets. Recently, researchers proposed the end-to-end spoken language understanding [14] to extract the structure representation of user’s intention directly from the speech signals. An end-to-end SLU aims at classifying the observed utterance into one of the predefined semantic classes $L = \{l_1, \dots, l_k\}$ [20]. Thus, a semantic classifier is trained to maximize the class-posterior probability for a given observation, $X = \{x_1, x_2, \dots, x_j\}$, representing a sequence of acoustic frames. This is achieved by the following probability:

$$L^* = \arg \max_L P(L|X, \theta) \quad (4)$$

where θ is the parameters of the end-to-end neural network model. Compared to the tandem ASR + NLU architecture, the end-to-end approach solves the so-called ASR error propagation problem, but usually requires large amounts of data to deal with the high variability present in the speech signal.

5. Experimental Setup

5.1. Dataset

The FSC dataset comprises single-channel audio clips sampled at 16 kHz. The data was collected using crowdsourcing, with participants requested to cite random phrases for each intent twice. It contains about 19 hours of speech, providing a total of 30,043 utterances cited by 97 different speakers. The data is split in such a way that the training set contains 14.7 hours of data, totaling 23,132 utterances from 77 speakers. Validation and test sets comprise 1.9 and 2.4 hours of speech, leading to 3,118 utterances from 10 speakers and 3,793 utterances from other 10 speakers, respectively. The dataset comprises a total of 31 unique intent labels resulted in a combination of three slots per audio: action, object, and location. The latter can be either

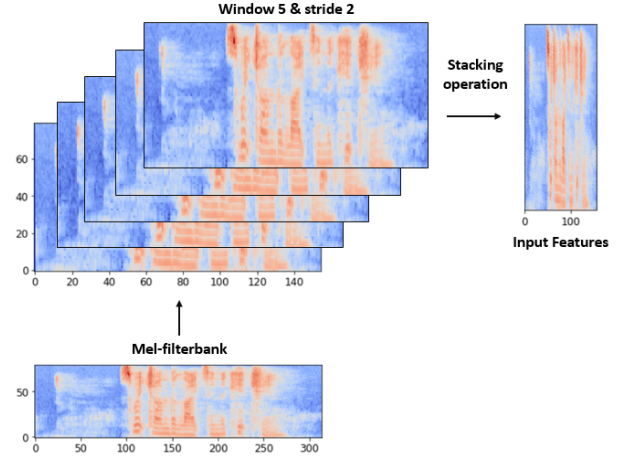


Figure 3: Stacking operation applied after feature extraction.

“none”, “kitchen”, “bedroom”, “washroom”, “English”, “Chinese”, “Korean”, or “German”. In our experiments, we defined intent as the combination of action and object, which led to a total of 15 different intent labels. Location was defined as slot, which led to a total of 8 different slot labels. More details about the dataset can be found in [14].

5.2. Features

In this work, each speech signal is (re)sampled at 16 kHz. We then extract 80-dimensional log Mel-Filterbank features. To extract the Mel features, the audio signal is processed in frames of 320 samples (i.e., 20-ms window length), with a step size of 160 samples (that is, 10-ms hop-size). Global Cepstral Mean and Variance Normalization (CMVN) are applied in order to mitigate the mismatch between training and testing data. After extracting mel-filterbank features, a stacking operation is performed. Specifically, we stacked 5 input frames with stride 2.

5.3. Baselines

The proposed method is compared to several baselines, including 32-bits full-precision, integer quantization at 16, 8, 4, 3 and 2 bits, all based on kernel multiplication. These baselines are referred to as $FP32$, $Q16$, $Q8$, $Q4$, $Q3$ and $Q2$. We also included Deepshift as another shift network solution as baseline [9], and we refer to this solution as $D4$, $D3$, $D2$ for quantization at 4, 3 and 2 bits. The proposed shift network is named $S4$, $S3$ and $S2$, also for quantization at 4, 3 and 2 bits. We also included as baseline the work proposed in [14], in which the SLU model is first pre-trained to predict phonemes on the first layers and words, thus optimizing the feature representation for SLU.

ResNet-18 is used as our backbone and they are trained from scratch. The architecture adopted in this work is based on the model proposed in [21], with the adaptation of having only one input channel as our input representation is based on mel-filterbank.

5.4. Experimental Settings

Our network is trained on mini-batches of 8 samples over a total of 200 epochs. In [10], for instance, the authors pointed out the importance of training the model for 200 epochs. In our




Figure 4: Performance in terms of (a) loss and (b) accuracy of the ResNet-18 explored in different versions, including its 32-bits full-precision (FP32), integer quantization (Q3), Deepshift (D3) and proposed ternary shift network (S3), all quantized at 3 bits.

experiments, we found Adam optimizer converging faster, but providing poorer generalization performance compared to the SGD optimizer. Thus, except for DeepShift, trained with the Adam optimizer, all models were trained using the SGD optimizer. The initial learning rate was set to 0.0001 and cosine annealing was also applied. The FSC dataset is split into training, validation, and test sets, and the hyperparameters and final model were selected based on the performance of the validation set. All reported results are based on the accuracy of validation and test sets.

6. Results

Figure 4-a and Figure 4-b show the performance of 3 baselines (i.e. FP32, Q3, D3) and the proposed low-bit shift network, referred to as S3, in terms of loss and accuracy, respectively. The losses are attained from the train set whereas the accuracy is based on the test set. We can observe that the quantized models, Q3, S3, follow similar trend as the full-precision model regarding the loss and their accuracy with S3 offering slightly better results. While these models have a monotonic decay of their losses and an increasing accuracy towards the end of training, D3 enters an overfitting regime around epoch 100.

In Table 1, we present the performance in terms of accuracy and other quantization settings are explored. Results are based on the best model and not on the final model. For D3, for instance, we chose a model around epoch 100 which gave the best accuracy on the validation set. The same was done for the other experiments in Table 1. Results are compatible to the baseline results presented in [14], with the full-precision (FP32) even outperforming it with accuracy as high as 92.34 % and 98.97 % for the validation and test set, respectively, followed by our proposed solution quantized with 4 bits, S4, which provides 92.24 % and 98.76 %, respectively, for the validation and test sets as well.

Quantization had a mild impact on the overall performance. For integer quantization, for example, quantizing with 16 bits showed to be slightly beneficial compared to quantizing with 2 bits with latter dropping accuracy only 0.06 % on the test set. Similar trend was found for the ternary shift network. While quantizing with 4 bits gives 98.76 % accuracy, quantizing with 2 bits provides 98.41 % accuracy, representing a decay of only

Table 1: Experimental results for ResNet-18. Performance is reported in terms of accuracy (%).

Kernel Operation	Method	# Bits	Val	Test
-	[14]	-	89.50	98.80
Multiplication	FP32	32	92.36	98.97
	Q16	16	90.73	98.47
	Q8	8	90.69	98.57
	Q4	4	91.24	98.52
	Q3	3	91.43	98.28
	Q2	2	91.08	98.39
Shift	D4	4	79.76	93.16
	D3	3	82.39	94.83
	D2	2	80.46	93.80
	S4 (ours)	4	92.24	98.76
	S3 (ours)	3	90.31	98.41
	S2 (ours)	2	90.66	98.41

0.35 %. These results show that compressing a model with low-bit representation is feasible for the SLU task and with the ternary shift network fast inference will be possible with the absence of expensive multiplication operations.

7. Conclusion

In this work, we propose a low-bit deep neural network for end-to-end (e2e) spoken language understanding (SLU). For that, a power-of-two quantization is applied on continuous weights from a full-precision ResNet-18. The low-bit representation is further decomposed in a sign-sparse-shift 3-fold manner. The method reduces computation complexity by removing expensive multiplication operations and using low-bit weights. We evaluate the performance of the proposed model on the Fluent Speech Command dataset. Our solution can provide accuracy as high as 98.76 %, showing that it can achieve comparable performance to its full-precision counterpart and state of the art models. As future work, we plan to investigate the use of our solution on recurrent neural networks as they might further improved inference latency for the e2e SLU task.

8. References

- [1] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE internet of things journal*, vol. 3, no. 5, pp. 637–646, 2016.
- [2] C. R. Banbury, V. J. Reddi, M. Lam, W. Fu, A. Fazel, J. Holleman, X. Huang, R. Hurtado, D. Kanter, A. Lokhmotov *et al.*, "Benchmarking tinyml systems: Challenges and direction," *arXiv preprint arXiv:2003.04821*, 2020.
- [3] S. Braun, H. Gamper, C. K. Reddy, and I. Tashev, "Towards efficient models for real-time deep noise suppression," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 656–660.
- [4] 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.
- [5] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in nlp," *arXiv preprint arXiv:1906.02243*, 2019.
- [6] M. S. Aslanpour, A. N. Toosi, C. Cicconetti, B. Javadi, P. Sbarski, D. Taibi, M. Assuncao, S. S. Gill, R. Gaire, and S. Dustdar, "Serverless edge computing: vision and challenges," in *2021 Australasian Computer Science Week Multiconference*, 2021, pp. 1–10.
- [7] A. Saade, J. Dureau, D. Leroy, F. Caltagirone, A. Coucke, A. Ball, C. Doumouro, T. Lavril, A. Caulier, T. Bluche *et al.*, "Spoken language understanding on the edge," in *2019 Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing-NeurIPS Edition (EMC2-NIPS)*. IEEE, 2019, pp. 57–61.
- [8] S. Rongali, B. Liu, L. Cai, K. Arkoudas, C. Su, and W. Hamza, "Exploring transfer learning for end-to-end spoken language understanding," *arXiv preprint arXiv:2012.08549*, 2020.
- [9] M. Elhoushi, Z. Chen, F. Shafiq, Y. H. Tian, and J. Y. Li, "Deepshift: Towards multiplication-less neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2359–2368.
- [10] X. Li, B. Liu, Y. Yu, W. Liu, C. Xu, and V. P. Nia, "s3: Sign-sparse-shift reparametrization for effective training of low-bit shift networks," *arXiv preprint arXiv:2107.03453*, 2021.
- [11] H. You, X. Chen, Y. Zhang, C. Li, S. Li, Z. Liu, Z. Wang, and Y. Lin, "Shiftaddnet: A hardware-inspired deep network," *Advances in Neural Information Processing Systems*, vol. 33, pp. 2771–2783, 2020.
- [12] A. Bie, B. Venkitesh, J. Monteiro, M. Haidar, M. Rezagholizadeh *et al.*, "A simplified fully quantized transformer for end-to-end speech recognition," *arXiv preprint arXiv:1911.03604*, 2019.
- [13] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2704–2713.
- [14] L. Lugosch, M. Ravanelli, P. Ignato, V. S. Tomar, and Y. Bengio, "Speech model pre-training for end-to-end spoken language understanding," *arXiv preprint arXiv:1904.03670*, 2019.
- [15] 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.
- [16] F. Li, B. Zhang, and B. Liu, "Ternary weight networks," *arXiv preprint arXiv:1605.04711*, 2016.
- [17] D. A. Gudovskiy and L. Rigazio, "Shiftcnn: Generalized low-precision architecture for inference of convolutional neural networks," *arXiv preprint arXiv:1706.02393*, 2017.
- [18] A. Zhou, A. Yao, Y. Guo, L. Xu, and Y. Chen, "Incremental network quantization: Towards lossless cnns with low-precision weights," *arXiv preprint arXiv:1702.03044*, 2017.
- [19] Y. Bengio, N. Léonard, and A. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," *arXiv preprint arXiv:1308.3432*, 2013.
- [20] R. Masumura, Y. Ijima, T. Asami, H. Masataki, and R. Higashinaka, "Neural confnet classification: Fully neural network based spoken utterance classification using word confusion networks," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 6039–6043.
- [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.