

# FIXED-POINT QUANTIZATION AWARE TRAINING FOR ON-DEVICE KEYWORD-SPOTTING

Sashank Macha\*    Om Oza\*    Alex Escott\*    Francesco Calivá\*  
Robbie Armitano\*    Santosh Kumar Cheekatmalla\*  
Sree Hari Krishnan Parthasarathi†    Yuzong Liu\*

\*Alexa Perceptual Technologies, †Alexa AI, Amazon

## ABSTRACT

Fixed-point (FXP) inference has proven suitable for embedded devices with limited computational resources, and yet model training is continually performed in floating-point (FLP). FXP training has not been fully explored and the non-trivial conversion from FLP to FXP presents unavoidable performance drop. We propose a novel method to train and obtain FXP convolutional keyword-spotting (KWS) models. We combine our methodology with two quantization-aware-training (QAT) techniques – squashed weight distribution and absolute cosine regularization for model parameters, and propose techniques for extending QAT over transient variables, otherwise neglected by previous paradigms. Experimental results on the Google Speech Commands v2 dataset show that we can reduce model precision up to 4-bit with no loss in accuracy. Furthermore, on an in-house KWS dataset, we show that our 8-bit FXP-QAT models have a 4-6% improvement in relative false discovery rate at fixed false reject rate compared to full precision FLP models. During inference we argue that FXP-QAT eliminates q-format normalization and enables the use of low-bit accumulators while maximizing SIMD throughput to reduce user-perceived latency. We demonstrate that we can reduce execution time by 68% without compromising KWS model’s predictive performance or requiring model architectural changes. Our work provides novel findings that aid future research in this area and enable accurate and efficient models.

**Index Terms**— keyword-spotting, quantization-aware-training, fixed-point-arithmetic, on-device

## 1. INTRODUCTION

We are interested in a small memory footprint and low latency keyword spotting (KWS) system for embedded on-the-go (OTG) devices. Modern KWS algorithms use deep neural networks (DNN) and have shown to achieve human-level performance [1]. Despite their success, DNN deployment on-device is challenging due to computational and bandwidth demand. These challenges can be greatly mitigated by converting expensive floating-point (FLP) operations to fixed-point (FXP) while simultaneously improving inference speed [2]. FXP multiplications consume  $18.5\times$  less energy and half the memory compared to FLP [3]. However, DNN training is usually conducted in FLP due to operation parallelism on commercial GPUs, and wide numerical range that allows for better gradient updates. While similar parallelization for FXP can be realized on FPGAs [4], the lack of a processor comparable to GPUs makes training significantly slow. Most approaches thereby resort to converting models trained in FLP to FXP introducing bias that degrades performance [5, 6]. We propose a novel method that addresses the limitations of FXP training and uses FLP workflows

to obtain FXP models.

Computational requirements can be reduced further using low-precision inference via quantization, which allows increased operations per accessed memory byte [5, 7]. Such quantization is typically achieved by means of post-training-quantization (PTQ) [8], which however causes severe information loss affecting model accuracy. To maintain overall accuracy for quantized DNNs, quantization can be incorporated in the training phase leading to quantization-aware-training (QAT). QAT introduces quantization noise during training by means of deterministic rounding [9, 10, 11], reparametrization [12, 13] or regularization [14, 15] among few techniques, allowing DNNs to adapt to inference quantization. Notable work has shown that with QAT model parameters can be learned at binary and ternary precision [16, 17].


In this paper, we discuss methods through which we enable FXP training and combine it with QAT. Our approach allows training up to ultra-low-precision (4-bit) FXP models that are accurate, and perform on-par with full-precision (32-bit) FLP models demonstrated on publicly available Google Speech Commands v2 dataset [18]. On an inhouse KWS dataset, our 8-bit FXP-QAT models outperform the 32-bit FLP baseline by a relative 4-6%. We additionally propose key inference optimizations that make KWS on constrained OTG devices possible. Our proposed approach achieves 68% faster inference while maintaining similar accuracy to full-precision FLP models. Furthermore, we argue that with QAT, we can eliminate  $\sim 10\%$  CPU cycles otherwise spent in normalizing activation precision. Our contributions can be summarized as follows -

- we propose novel FXP model training that leverages workflows optimized for FLP;
- we extend two QAT techniques to FXP and train low-bit KWS models with minimal accuracy degradation;
- we reduce inference execution time using low-bit accumulators enabled through FXP-QAT;
- we propose a two-tier accumulator-buffer approach to reduce overflow/saturation typical of large DNN layers;

The remainder of the paper is organized as follows: in sections 2 and 3, we present related work, and introduce QAT focusing on two methods based on squashed weight distribution and absolute-cosine regularization. In section 4, we propose and describe our contributions to these methods, which allow us to achieve fixed-point training. We describe our experimental setup, and results in section 5. After discussion, we conclude the paper with our vision on the future of FXP-QAT for small footprint devices.

## 2. RELATED WORK

Among inference setups aimed at enabling DNN deployment for low-power platforms, FXP systems play a key role [19]. FXP of-



**Fig. 1.** Proposed two-stage quantization approach to learn FXP weights in FLP training framework.


ferred the most promising trade-off between accuracy and computational complexity among different number systems [20]. Horowitz *et al.* [3] found corroborative evidence with FXP multiply-accumulate (MAC) units compared to FLP on 45nm technology. Notable work has been done in FXP inference — Umuroglu *et al.* [7] built a framework that mapped binary weights to hardware, while other foundational works proposed to design schemes for FXP inference [21, 22]. On another front, there has been work on efficient model topologies [4], better approaches to quantize parameters to FXP with low errors [6] and methods to address limitations of FXP training [23, 24, 25]. While the majority of the work has been centered around developing suitable models [26] or efficient hardware, QAT improves model adaptation to FXP. Lin *et al.* [27] used SQNR based quantization to achieve better accuracy with 8-bit FXP but suffered from data overflow during inference. Anwar *et al.* [28] presented an exhaustive L2-error minimization based quantization but required pre-training with full-precision. In the context of speech processing, 4 to 6-bit QAT have been studied for event detection [29], speech recognition [14] and KWS [9] but limited to floating-point. All the above methods either require pretraining of model parameters or have runtime limitations and have limited discussion on inference computations.

### 3. QUANTIZATION AWARE TRAINING

Quantized models are typically obtained by full-precision training followed by quantization. Nevertheless, for aggressive compression, low-bit quantization or cross format conversion from FLP to FXP, such quantization degrades accuracy. With QAT, we introduce inference specific quantization-noise during training to avoid inference time errors. Our FXP training is agnostic to QAT methodology and can extend to different techniques. We choose two recent works in this domain and describe how we enable FXP-QAT for different model components.

#### 3.1. Weight quantization

Using **squashed weight distribution (SQWD)**, weights can be reparametrized to a finite range (*e.g.*  $[-1, 1]$ ). This can be achieved by means of a non-linear function (*i.e.*  $\tanh$ ) that maps weight ( $w$ ) to  $(\hat{w})$  in uniform distribution [12]. A regularization loss limits the distribution-spread to asymptotes of  $\tanh$  for efficient mapping. Subsequently,  $\hat{w}$  is quantized to desired  $b$ -bit precision using eq. 1, where



**Fig. 2.** Activation normalization between two consecutive layers with PTQ. QAT limits the use of different q-formats eliminating normalization.

$w_{(flp,b)}$  refers to FLP weight  $w$  with  $b$ -bit precision.

$$w_{(flp,b)} = (\text{round}[2^{b-1}(\hat{w} + 1) - 0.5] + 0.5) * 2^{b-1} - 1 \quad (1)$$

Using **absolute cosine regularization (ACR)**, QAT can be defined as a regularization problem. Most QAT techniques do not include explicit quantization loss in back-propagation [14]. Conversely with ACR, the model parameters are weighted using an absolute-cosine function that represents quantization loss. The cosine function has  $2^b$  zero-points for  $b$ -bit quantization. We constrain our weights to the same finite range as SQWD to establish a fair comparison.

$$w_{(fxp,b)} = \text{round}[2^{b-1} * (w_{(flp,32)} + 1.0)] \quad (2)$$

$$w_{(flp,b)} = w_{(fxp,b)} * 2^{b-1} - 1.0 \quad (3)$$

**Proposed approach with FXP** Either in SQWD or ACR, a FLP quantizer is used to quantize a weight ( $w$ ) from full-precision FLP ( $w_{(flp,32)}$ ) to  $b$ -bit FLP ( $w_{(flp,b)}$ ). For a FXP inference with no quantization error, a trivial solution would be to replace the FLP quantizer in eq. 1 with a FXP quantizer. However, FXP arithmetic has limited support on GPUs which makes training large ML models challenging. To combat this, we propose a two-stage quantizer (shown in figure 1) that maps  $w$  from full-precision FLP ( $w_{(flp,32)}$ ) to  $b$ -bit FXP ( $w_{(fxp,b)}$ ) using eq. 2 and to  $b$ -bit FLP ( $w_{(flp,b)}$ ) using eq. 3. This facilitates training in FLP arithmetic, albeit using FXP quantization. During back-propagation we use straight-through-estimators [30] for gradient calculation, to overcome the non-differentiability of the round  $[\cdot]$  operation. For inference, we use eq. 2 to export parameters to FXP. With a one-to-one mapping between quantized FLP and FXP representations, we can eliminate errors from PTQ.

#### 3.2. Input and activation quantization

Inputs to various model components form the remainder of inference arithmetic. In the context of hidden layer activations, deep models have a cascaded bit-growth problem due to the large number of multiply and accumulate (MAC) operations. To mitigate this, we placed fake quantization nodes between consecutive layers to scale down precision, and used a clipped ReLU with linear scaling to squash activations to smaller  $[-1, 1]$  range, prior to quantization with eq. 2 and eq. 3. On the other hand, to accommodate wider numerical range of feature inputs, we traded-off bit precision on fractional part. The number of bits used to represent fractional part is termed as  $q$ -format.

A slightly modified quantization approach with  $q$ -format for a feature  $f$  is proposed in eq. 4 and 5, where  $c = 0.5$  if  $f < 0$  else  $-0.5$ .

$$\lfloor b \rfloor f_{(fp,b,q)} = \min[\max(f * 2^q + c, -2^{b-1}), 2^{b-1} - 1] \quad (4)$$

$$f_{(flp,b,q)} = f_{(fp,b,q)} * 2^{-q} \quad (5)$$

#### 4. FIXED-POINT INFERENCE


This section describes how QAT improves FXP inference by eliminating a key trade-off typical of post-training-quantized models, and leverages increased parallelization with low-bit inference on SIMD architectures. Additionally, we describe and propose a solution for archetypal saturations in low-bit accumulators. The discussion is tailored in the scope of ARM NEON chipsets, a popular choice for FXP inference.

##### 4.1. q-format normalization

FXP inference setups - specifically designed for PTQ assign a different  $q$ -format ( $q_c$ ) per axis for model parameters to reduce quantization-error, resulting in mixed representations within the model. The number of representations scale with number of layers, and incur additional memory to store  $q$ -format mappings. To prevent this, activations are normalized during runtime to a uniform  $q_*$  prior to subsequent layer. These normalization operations roughly consume 10% cycles in inference. With QAT, we impose training time constraints to learn model parameters with a common  $q$ -format ( $\hat{q}$ ) thus eliminating additional cycles spent in normalization, speeding up inference. This difference between PTQ and QAT is illustrated in figure 2.

##### 4.2. Increased parallelization with QAT on SIMD

ARMv7a/v8a is the most adopted architecture for embedded devices using FXP arithmetic. The advanced-SIMD extension (NEON) in these architectures accelerates processing by parallelizing load, store, multiply, and accumulate operations. The `vmlal_s8` NEON intrinsic [31] operating on 128-bit wide NEON registers supports up to eight operations in a single cycle. During inference, MACs make up to 90% of computations, and we leverage parallelization capability of NEON to parallelize MACs and achieve faster inference. To fully utilize the  $8 \times$  parallelization, we limit the accumulator size in MAC operations to low 16-bits, since  $16 \times 8 = 128$ , shown in figure 3. The use of low-bit accumulator is enabled through FXP-QAT by training low-bit precision models. The 16-bit accumulator is scaled-down prior to next layer to avoid bit-growth explosion.



**Fig. 3.** Proposed two-tier accumulator-buffer mechanism that prevent saturations in low-bit accumulator.

##### 4.3. Saturations in low-bit accumulators

While smaller accumulators speed up inference, they saturate faster. The 16-bit NEON accumulator can saturate in as little as 2 and

**Table 1.** Number of activations corrupted by saturations at different convolutional layers with and without 2T-AB.

kernel size	# MACs per activation	# activations	# activations corrupted by saturations at flushing cadence			
			None	256 MACs	128 MACs	64 MACs
(3, 4, 32)	384	10,880	2,598	237	47	0
(4, 4, 32)	512	2,240	2,090	276	27	0
(7, 4, 40)	1,120	128	128	12	0	0
(1, 1, 128)	128	160	39	6	0	0
(1, 1, 160)	160	3	3	3	2	0
TOTAL			4,858	534	76	0

4 MACs using 8-bit and 7-bit models. Although we can employ QAT to train models up-to ultra-low-bit precision (*i.e.* 4-bit), model topologies influence saturations. Specifically, wide models involve greater MACs per activation and are more likely to saturate accumulators. Saturations are a significant problem for the model performance since they corrupt downstream model propagation. For the keyword-spotting problem we are interested in, we observed that 0.1% saturations caused a  $\sim 30\%$  degradation in keyword detection score. As a solution, we propose a two-tier accumulator-buffer (2T-AB) mechanism that mitigates this problem while maintaining benefits from  $8 \times$  SIMD parallelism. As opposed to performing all the required MACs for an activation using a single accumulator, we propose to periodically flush the 16-bit accumulator to a wider 32-bit fixed-width buffer. In table 1, we profile 6-bit KWS model and show 2T-AB reduces the number of saturations. We observe that sparsely flushing (*i.e.* every 256 MACs) reduces saturations by 89% while flushing every 64 MACs completely eliminates saturations.

#### 5. EXPERIMENTAL STUDY

We study the proposed fixed-point QAT technique using a KWS model for resource constrained devices. We evaluate our approach using quantized models trained with our approach against an unquantized full-precision model in terms of *i)* accuracy *ii)* latency.

##### 5.1. Datasets

We conducted our experimentation using an in-house KWS dataset ( $D_1$ ) and the Google speech commands v2 dataset ( $D_2$ ) [18].  $D_1$  consisted of 12.5K hours of de-identified labeled audio from far-field and near-field mobile phones and was split into training (10K hours) and testing (2K hours) while  $D_2$  comprised of 105K one-second audio snippets sampled at 16kHz, containing 35 different keyword utterances recorded in natural environments [18]. We used the official train/validation/test splits for  $D_2$ .

The audio in  $D_1$  and  $D_2$  was down-sampled to 16kHz, converted to single channel, and segmented using a 25ms analysis window with a 10ms shift. While audio files have varying length, we used a segment length of 76 to generate fixed length inputs. Subsequently, we extracted and concatenated 64-dimensional log-mel filter bank energies (LFBE-64) to form 2-D feature inputs to our convolutional KWS model.

##### 5.2. Model configurations

We use a 2-D fully convolutional model for all experiments. The model contains five trainable consecutive convolutional blocks preceded by a non-trainable normalization layer. Convolutional blocks vary across the network in terms of kernel size, yet each convolutional layer is followed by batch normalization and ReLU nonlinear-ity as activation. The output is a softmax classification layer representing the posterior probability of keyword detection. The model has a total of 199K learnable parameters and uses 125K flops.

**Table 2.** Relative false discovery rate (FDR) of different QAT models on dataset  $D_1$  and their relative execution times (s) as measured on ARM CortexA53.

model	bit-width	inference-format	relative FDR	relative exec-time
baseline	32-bit	FLP	1.0	1.0
optimized-PTQ	16-bit	FXP	1.109	0.45
SQWD (QAT)	8-bit	FXP	0.944	0.312
ACR (QAT)	8-bit	FXP	0.960	0.312

### 5.3. Training configurations

We standardize data prior to training using statistics computed from the training set and retain them for evaluation. We train all models in TensorFlow (2.4) for 1M steps and use Radam optimizer [32] to update model parameters. We increase learning rate to 1e-3 for the first 10% of total steps, and decay it linearly to 1e-5 thereafter. We use cross-entropy loss as our primary training objective, along with additional loss terms specific to different QAT paradigms. We use a batch size of 1500, and distribute training across 8 V100 GPUs.

### 5.4. Evaluation metrics

We measure model accuracy using end-to-end Detection Error Tradeoff (DET) curves, which describe models’ False Rejection Rate (FRR) vs. False Discovery Rate (FDR). Similar to [33] we normalize the axes of DET curves and report relative FDR to full-precision floating-point models’ FDR when relative FRR is set to 1.0. We measure inference speed using a single core CortexA53 based MT8163 processor specifically compiled for ARMv7a. We use an audio stream of length 90s containing multiple instances of keyword, and run it continually on the processor against our KWS model to measure execution time. Similar to FDR, we normalize and report relative execution times to floating point model running on non-SIMD architectures.

**Table 3.** Relative execution time (s) of different inference models on ARM CortexA53 versus flushing cadences to study 2T-AB.

model	inference-setup (format, architecture)	flushing cadence	degree of parallelization	relative exec-time (%change)
baseline	FLP; non-SIMD	None	1x	1.0 (NA)
optimized-PTQ	FXP; SIMD	None	4x	0.450 (55%)
		None		0.312 (68.8%)
		512 MACs		0.309 (69.1%)
6-bit QAT (SQWD)	FXP; SIMD	256 MACs		0.309 (69.1%)
		128 MACs	8x	0.312 (68.8%)
		64 MACs		0.356 (64.4%)
		32 MACs		0.363 (63.7%)
		16 MACs		0.407 (59.3%)

### 5.5. Results and discussion

We train our KWS model at different precision using proposed QAT approach and refrain from changing model topology to compensate for performance at low-precision. We first trained a model in full-precision FLP (*baseline*), and performed post-training-quantization that includes common PTQ optimizations not limited to dynamic-q-factor selection, run-time normalization to reduce quantization losses (*optimized-PTQ*). We also trained two models using the two proposed FXP-QAT techniques (*i.e.* SQWD, ACR) at 8-bit precision. To run FXP inference, we export the model parameters to FXP with eq 2 and use a MT8163 processor compiled for ARMv7a. In table 2, we report relative FDR numbers for the above 4 models on dataset  $D_1$ . We observe that despite the optimizations, the PTQ models’ performance deteriorates compared to the baseline, while the models trained with fixed-point QAT perform better than the baseline

**Table 4.** Model accuracy on dataset  $D_2$  at different weight (w) and activation (a) precision (in bits). Table a) employs QAT using ACR, and table b) uses SQWD. The full-precision floating point accuracy of the model on this dataset is 90.7%.

w	a	8	7	6	5	4	w	a	8	7	6	5	4
		8	92.8	92.8	92.3	91.4			91.4	8	91.6	91.4	91.4
7	92.4	92.9	92.7	91.9	91.6	7	92.0	91.2	91.5	91.1	90.9		
6	92.5	92.2	92.7	92.1	91.2	6	91.3	91.6	91.7	91.3	91.3		
5	92.6	92.5	92.8	92.3	90.9	5	91.3	91.5	92.3	91.3	91.9		
4	93.1	92.7	92.5	92.3	91.8	4	91.5	91.7	91.6	91.6	91.3		

(a)

(b)

at half the precision. We attribute the performance improvement to additional regularization introduced by QAT. We also report the execution times of these four models (FLP using *x86* and FXP using *ARM CortexA53*) and observe that 16-bit FXP-PTQ models, and FXP-QAT models are 55% and 68% faster compared to FLP baseline. Previously in table 1 we showed how proposed 2T-AB mechanism reduces saturations, but we note that the occasional flushing adds additional instructions during inference; although, the overhead is minimal compared to 8 $\times$  speed up that we are able to realize. From table 3 we see that a sparse frequency of 512, 256 or 128 MACs there is no observable slowdown, but a dense flushing at 16 MACs increases by a relative 10%. Using dataset  $D_2$  we present some ablation studies with in table 4 and show that we are able to achieve comparable performance to floating-point accuracy at low-bit (6-bit) and ultra-low-bit (4-bit) precision with either SQWD or ACR. We also report that we could trade off precision on model weights than activations increased sensitivity towards activations. We confirm through our that we are able to successfully train FXP-QAT models that are faster and accurate than FLP counterparts.

## 6. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a method for training and developing FXP models that uses floating-point arithmetic during training, but has a one-one mapping and eliminates conversion errors post-training. We applied this method in conjunction with two different QAT techniques to learn low-precision model parameters and also propose methods that apply QAT over transient variables. We test our approach in the scope of keyword-spotting that often require accurate and low-latency inference. We initially demonstrate and show that our 8-bit FXP QAT models outperform 32-bit FLP counterparts. Encouraged by the promising results, we further reduced precision, and observed that our method has a strong generalization capability allowing us to train models up to ultra-low-bit precision (4-bit) with minimal accuracy degradation. Besides accurate, our model is more efficient than traditional models, achieving 68% and 30% speed up in execution time compared to floating-point inference, or the traditional combination of post training weight quantization, followed by fixed-point inference. In conclusion, we argue that it is possible to reduce on-device computation processing time, without sacrificing accuracy, which is especially important for always-on keyword spotting systems. In the next iteration of our work we plan to explore how our approach scales to large models and different problem spaces like computer vision and complex tasks like speech recognition. Additionally, we plan to combined our work with other compression techniques like pruning, distillation, and neural architecture search to generate efficient models. We believe that our findings will broadly impact the community focused on deep learning on resource constrained devices.

## 7. REFERENCES

- [1] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep learning*, MIT press, 2016.
- [2] Alexandros Kouris, Stylianos I. Venieris, and Christos-Savvas Bouganis, “Cascade cnn: Pushing the performance limits of quantisation in convolutional neural networks,” *2018 28th International Conference on FPGAs*.
- [3] Mark 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 Paper*, 2014.
- [4] Chun Yan Lo and Chiu-Wing Sham, “Energy efficient fixed-point inference system of convolutional neural network,” in *2020 IEEE 63rd International MWCAS*.
- [5] Ritchie Zhao, Weinan Song, Wentao Zhang, Tianwei Xing, Jeng-Hau Lin, Mani Srivastava, Rajesh Gupta, and Zhiru Zhang, “Accelerating binarized convolutional neural networks with software-programmable fpgas,” 2017, ACM.
- [6] Norbert Mitschke, Michael Heizmann, Klaus-Henning Noffz, and Ralf Wittmann, “A fixed-point quantization technique for convolutional neural networks based on weight scaling,” in *2019 IEEE ICIP*.
- [7] Yaman Umuroglu, Nicholas J. Fraser, Giulio Gambardella, Michaela Blott, Philip Heng Wai Leong, Magnus Jahre, and Kees A. Vissers, “Finn: A framework for fast, scalable binarized neural network inference,” *Proceedings of the 2017 ACM/SIGDA International Symposium on FPGAs*.
- [8] Ron Banner, Yury Nahshan, and Daniel Soudry, *Post Training 4-Bit Quantization of Convolutional Networks for Rapid Deployment*, Curran Associates Inc., 2019.
- [9] Yuriy Mishchenko, Yusuf Goren, Ming Sun, Chris Beauchene, Spyros Matsoukas, Oleg Rybakov, and Shiv Naga Prasad Vitaladevuni, “Low-bit quantization and quantization-aware training for small-footprint keyword spotting,” *2019 18th IEEE ICMLA*.
- [10] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew G. Howard, Hartwig Adam, and Dmitry Kalenichenko, “Quantization and training of neural networks for efficient integer-arithmetic-only inference,” *2018 IEEE/CVF Conference on CVPR*.
- [11] Hao Li, Soham De, Zheng Xu, Christoph Studer, Hanan Samet, and Tom Goldstein, “Training quantized nets: A deeper understanding,” in *NIPS*, 2017.
- [12] Nikko Strom, Haidar Khan, and Wael Hamza, “Squashed weight distribution for low bit quantization of deep models,” *Proceedings of InterSpeech*, 2022.
- [13] Robin Dupont, Hichem Sahbi, and Guillaume Michel, “Weight reparametrization for budget-aware network pruning,” *2021 IEEE ICIP*.
- [14] Hieu Duy Nguyen, Anastasios Alexandridis, and Athanasios Mouchtaris, “Quantization aware training with absolute-cosine regularization for automatic speech recognition,” in *InterSpeech*, 2020.
- [15] Ruizhou Ding, Ting-Wu Chin, Zeye Dexter Liu, and Diana Marculescu, “Regularizing activation distribution for training binarized deep networks,” *2019 IEEE Conference on CVPR*.
- [16] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David, “Binaryconnect: Training deep neural networks with binary weights during propagations,” *Advances in NeurIPS*, 2015.
- [17] Wei Zhang, Lu Hou, Yichun Yin, Lifeng Shang, Xiao Chen, Xin Jiang, and Qun Liu, “Ternarybert: Distillation-aware ultra-low bit bert,” in *EMNLP*, 2020.
- [18] Pete Warden, “Speech commands: A dataset for limited-vocabulary speech recognition,” *arXiv preprint arXiv:1804.03209*, 2018.
- [19] Vasanthakumar Rajagopal, Chandra Kumar Ramasamy, Ashok Vishnoi, Raj Narayana Gadde, Narasinga Rao Miniskar, and Sirish Kumar Pasupuleti, “Accurate and efficient fixed point inference for deep neural networks,” in *2018 25th IEEE ICIP*.
- [20] Patrick Judd, Jorge Albericio, Tayler Hetherington, Tor M. Aamodt, Natalie Enright Jerger, and Andreas Moshovos, “Proteus: Exploiting numerical precision variability in deep neural networks,” 2016.
- [21] Chen Zhang, Peng Li, Guangyu Sun, Yijin Guan, Bingjun Xiao, and Jason Cong, “Optimizing fpga-based accelerator design for deep convolutional neural networks,” 2015.
- [22] Ritchie Zhao, Weinan Song, Wentao Zhang, Tianwei Xing, Jeng-Hau Lin, Mani B. Srivastava, Rajesh K. Gupta, and Zhiru Zhang, “Accelerating binarized convolutional neural networks with software-programmable fpgas,” *Proceedings of the 2017 ACM/SIGDA International Symposium on FPGAs*.
- [23] Sujeong Jo, Hanmin Park, Gunhee Lee, and Kiyoun Choi, “Training neural networks with low precision dynamic fixed-point,” in *2018 IEEE 36th ICCD*.
- [24] Yasufumi Sakai, “Quantization for deep neural network training with 8-bit dynamic fixed point,” in *2020 7th ISCM*.
- [25] Donghyeon Han, Dongseok Im, Gwangtae Park, Youngwoo Kim, Seokchan Song, Juhyoung Lee, and Hoi-Jun Yoo, “Hnpu: An adaptive dnn training processor utilizing stochastic dynamic fixed-point and active bit-precision searching,” *IEEE Journal of Solid-State Circuits*, 2021.
- [26] Ming Sun, Varun Nagaraja, Björn Hoffmeister, and Shiv Vitaladevuni, “Model shrinking for embedded keyword spotting,” in *2015 IEEE 14th ICMLA*.
- [27] Darryl Dexu Lin, Sachin S. Talathi, and V. Sreekanth Annareddy, “Fixed point quantization of deep convolutional networks,” in *ICML*, 2016.
- [28] Sajid Anwar, Kyuyeon Hwang, and Wonyong Sung, “Fixed point optimization of deep convolutional neural networks for object recognition,” in *2015 IEEE ICASSP*, 2015.
- [29] Bowen Shi, Ming Sun, Chieh-Chi Kao, Viktor Rozgic, Spyridon Matsoukas, and Chao Ching Wang, “Compression of acoustic event detection models with low-rank matrix factorization and quantization training,” *ArXiv*, 2019.
- [30] Yoshua Bengio, Nicholas Léonard, and Aaron Courville, “Estimating or propagating gradients through stochastic neurons for conditional computation,” *arXiv preprint arXiv:1308.3432*, 2013.
- [31] “Neon programmers guide, vector multiply-accumulate,” .
- [32] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han, “On the variance of the adaptive learning rate and beyond,” *ArXiv*, 2020.
- [33] Lu Zeng, Sree Hari Krishnan Parthasarathi, Yuzong Liu, Alex Escott, Santosh Kumar Cheekatmalla, Nikko Strom, and Shiv Naga Prasad Vitaladevuni, “Sub 8-bit quantization of streaming keyword spotting models for embedded chipsets,” in *TSD*, 2022.