

# Sub 8-Bit Quantization of Streaming Keyword Spotting Models for Embedded Chipsets

Lu Zeng\*, Sree Hari Krishnan Parthasarathi\*, Yuzong Liu, Alex Escott,  
Santosh Cheekatmalla, Nikko Strom, Shiv Vitaladevuni

Alexa, Amazon, USA

{luzeng, sparta, liuyuzon, escottal, cheekatm,  
nikko, shivnaga}@amazon.com

**Abstract.** We propose a novel 2-stage sub 8-bit quantization aware training algorithm for all components of a 250K parameter feedforward, streaming, state-free keyword spotting model. For the 1<sup>st</sup>-stage, we adapt a recently proposed quantization technique using a non-linear transformation with  $\tanh(\cdot)$  on dense layer weights. In the 2<sup>nd</sup>-stage, we use linear quantization methods on the rest of the network, including other parameters (bias, gain, batchnorm), inputs, and activations. We conduct large scale experiments, training on 26,000 hours of de-identified production, far-field and near-field audio data (evaluating on 4,000 hours of data). We organize our results in two embedded chipset settings: a) with commodity ARM NEON instruction set and 8-bit containers, we present accuracy, CPU, and memory results using sub 8-bit weights (4, 5, 8-bit) and 8-bit quantization of rest of the network; b) with off-the-shelf neural network accelerators, for a range of weight bit widths (1 and 5-bit), while presenting accuracy results, we project reduction in memory utilization. In both configurations, our results show that the proposed algorithm can achieve: a) parity with a full floating point model’s operating point on a detection error tradeoff (DET) curve in terms of false detection rate (FDR) at false rejection rate (FRR); b) significant reduction in compute and memory, yielding up to 3 times improvement in CPU consumption and more than 4 times improvement in memory consumption.

**Index Terms:** 2-stage quantization, keyword spotting, embedded chipsets

## 1 Introduction

Wakeword detection, also known as keyword spotting (KWS), detects words or phrases of interest from streaming audio and plays a vital role in voice assistants [1,2]. KWS models are based on neural network architectures and are processed on device. A challenge for KWS systems is to attain high accuracy under tight resource constraints such as model size, runtime memory footprint, and power consumption. To address some of the constraints, previous work has

---

\*Equal contribution.

explored knowledge distillation [3], low rank approximation [4,5], and computationally efficient architectures [6,7,8]. As an orthogonal direction, quantization has been applied to KWS, where the components of the model are converted from 32-bit floating points to lower bit width representations.

Quantization is often applied post-training, resulting in performance degradation [9]. To mitigate this, quantization aware training (QAT) is applied to reduce errors [10,11]. QAT, including very low bit width (even binary quantization), is an established technique [12,13]. Since the quantization function is discrete (and therefore the gradient is zero almost everywhere), a fake quantizer with straight through estimator (STE) [14] or a Gumbel-Softmax trick [15] are used to estimate the parameters. Since accuracy of 1-bit models is still a challenge, in the context of speech processing, 4 to 6-bit QAT has been studied for event detection [10], speech recognition [16], and KWS [11].

Furthermore, commercially realizable KWS models have additional challenges often ignored in research focused on “model size vs accuracy” tradeoff: a) sub 8-bit quantization requires hardware support<sup>1</sup>; b) need to run in streaming mode<sup>2</sup>; c) cannot use corruptible memory<sup>3</sup>. Our paper happens in the context of: a) using sub 8-bit representations to avoid overflow errors with 8-bit containers in commodity or off-the-shelf platforms; b) addressing accuracy challenges with sub 8-bit, including 1-bit representations.

**Contributions:** In this paper, we tackle the problem of sub 8-bit quantization of on-device, small footprint, streaming, state-free KWS models, that can execute on commodity or off-the-shelf hardware platforms. We propose a novel 2-stage QAT algorithm: for the 1<sup>st</sup>-stage, we adapt a non-linear quantization method on weights [17], while for the 2<sup>nd</sup>-stage, we use linear quantization methods on other components of the network. We conduct large scale experiments, training on 26K hours of de-identified production audio, collected from a mix of far-field devices and mobile phones (evaluating on 4K hours of data). We show the efficacy of our methods by presenting accuracy and compute results (CPU<sup>4</sup>, memory, and model size) for sub 8-bit models (4, 5, 8-bit) on ARM NEON chipset, while projecting memory gains for a range of weight bit widths (1 and 5-bit) on off-the-shelf neural network accelerators.

## 2 Small Footprint, Streaming, State-Free KWS Models

While QAT algorithm can be applied to convolutional models, we use a model that is a feedforward network (FFN) with 250k parameters using a bottleneck architecture [18]. The full precision version of this model has been optimized for low power, small footprint, streaming and state-free setting and serves as our baseline. It is trained on a set of positive and negative examples (i.e. positive examples contain the wakeword, while negative examples do not).

---

<sup>1</sup>An instruction set that can efficiently carry out matrix-vector multiplications.

<sup>2</sup>Models have to run with low latency – i.e., cannot use large buffers.

<sup>3</sup>Since the models are running continuously, they cannot get into a “bad” state.

<sup>4</sup>We use CPU cycles as a proxy for power consumption.

The model architecture has 250K learnable parameters, and it operates on 20-dimensional log mel filter bank energy (LFBE) features, computed with an analysis window size and shift of 25 ms and 10 ms respectively; the input to the model is 81 frames, downsampled by a factor of 3. The architecture consists of five fully connected layers with batch normalization [19] and ReLU [20] being used with all hidden layers. The output is a binary classification layer trained with cross entropy loss, representing the posterior probability of “wakeword” and “non-wakeword” [18]. Later, to investigate quantized models with lower bit width, we explore two other model sizes with different layer sizes, keeping the other architectural aspects the same. During training, the wakeword is consistently center aligned in the input window [21]. Adam optimizer is used to update the model parameters during training. During inference, the posterior estimates corresponding to the wakeword are smoothed by a windowed smoothing average (WMA) filter and then thresholded to infer the wakeword hypothesis.

### 3 Relevant Embedded Chipsets

In this section, we describe hardware considerations in two settings: a) commodity ARM NEON instruction set; b) off-the-shelf neural network accelerators.

**Commodity ARM NEON Instruction Set Overview:** In 2021 over 90% of mobile phones use ARM-based chipsets <sup>5</sup>. The majority of these include a Single-Instruction-Multiple-Data (SIMD) extension known as NEON; this is available by default in ARM’s chipsets for mobile phones (*armv8a*, *aarch64* <sup>6</sup>). The NEON instruction set is based on parallelizing arithmetic operations on vectors stored in 128-bit registers, allowing packing of multiple scalar values in the register – i.e. 4x32-bit floating point, 4x32-bit integer, 8x16 bit-integer, or 16x8-bit integer – then perform an arithmetic operation on every value in the register in a clock cycle. For NEON, 8-bit is the smallest supported container for Multiply-And-Accumulate (MAC) operations.

The compute requirement for a KWS model can be reduced by using lower bit widths for weights, input and and activations. Our QAT algorithm produces a model that can utilize the SIMD instruction set available with NEON. While NEON uses 8-bit containers for MAC operations, using sub 8-bit weights can yield benefits: a) prevent overflow when performing computations; b) lower model size for over the network deployments.

**Off-the-shelf Neural Network Accelerators:** The smallest container in ARM NEON architecture is 8-bit, so we do not observe a memory reduction with sub 8-bit weights. Some off-the-shelf accelerators can utilize sub 8-bit weights <sup>7,8</sup>. In this paper, we use data sheets from them to project memory savings for our QAT algorithm.

<sup>5</sup><https://www.counterpointresearch.com/global-smartphone-ap-market-share/>

<sup>6</sup><https://www.arm.com/blogs/blueprint/android-64bit-future-mobile>

<sup>7</sup><https://datasheets.maximintegrated.com/en/ds/MAX78000.pdf>

<sup>8</sup><https://www.syntiant.com/post/syntiant-introduces-second-generation-ndp120-deep-learning-processor-for-audio-and-sensor-apps>

## 4 Proposed 2-Stage QAT Algorithm

In this section, we describe the design of our approach. Section 4.1 gives an overview of our 2-stage training algorithm. We provide detailed descriptions of the first and second stages in Sections 4.2 and 4.3.

### 4.1 Quantization Overview

A layer in a feedforward neural network typically takes the following form:

$$\mathbf{y} = \psi_{bn}(\phi_{relu}(\mathbf{W}\mathbf{x} + \mathbf{b})) \quad (1)$$

where  $\mathbf{x}$  is the input to the dense layer,  $\mathbf{W}$  is the dense layer weight matrix,  $\mathbf{b}$  is the bias,  $\phi_{relu}$  and  $\psi_{bn}$  represent ReLU and batch norm transformations. Using a gain  $\alpha$  and a linear quantizer ( $q_z(\cdot)$ ) that discretizes the values to  $z$  levels ( $\log_2(z)$  bits), let  $\mathbf{W}_q$ ,  $\mathbf{x}_q$ ,  $\alpha_q$ ,  $\mathbf{b}_q$  be the corresponding quantized representations. Quantization of such a layer can be done in the following steps:

$$\begin{aligned} \mathbf{y}_q &\leftarrow \alpha_q(\mathbf{W}_q \cdot \mathbf{x}_q + \mathbf{b}_q) \\ \mathbf{y}_q^{(\phi_q)} &\leftarrow q_z(\phi_{relu}(\mathbf{y}_q)) \\ \mathbf{y}_q^{(\phi_q, \psi_q)} &\leftarrow q_z(\psi_{bn}(\mathbf{y}_q^{(\phi_q)})) \end{aligned}$$

Our QAT algorithm for the layer follows a 2-stage procedure. During the 1<sup>st</sup>-stage we adapt the  $\tanh(\cdot)$  based quantization technique introduced in [17], for obtaining  $\mathbf{W}_q$ . In the 2<sup>nd</sup>-stage of training, we propose specific linear quantization techniques for the remaining components: gain ( $\alpha_q$ ), bias ( $\mathbf{b}_q$ ), batch normalization ( $q_z(\psi_{bn}(\cdot))$ ), activations ( $q_z(\phi_{relu}(\cdot))$ ), and inputs ( $\mathbf{x}_q$ ).

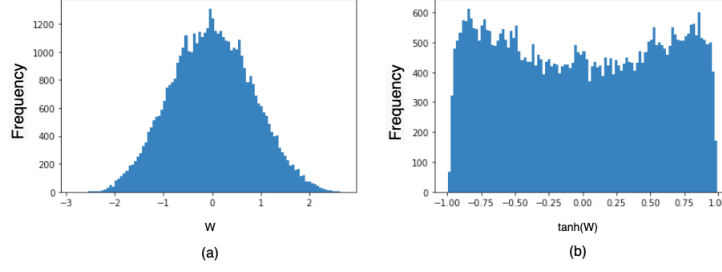
### 4.2 1<sup>st</sup> Stage: $\tanh(\cdot)$ Quantization of Weights

Linear quantization [22] is widely used for model compression. However, the weights of a neural network are typically Gaussian distributed. Linear quantization is not efficient for Gaussian distributions. To mitigate this problem, we adopt [17]: it encourages the weights of a neural network to become more uniform distributed, by first applying  $\tanh(\cdot)$  on  $\mathbf{W}$ , and then applying a linear quantizer to obtain  $\mathbf{W}_q \leftarrow q_z(\tanh(\mathbf{W}))$ ; here  $\mathbf{W}_q$  is the quantized weight matrix.

Fig. 1(a) shows an example on the original weight distribution  $\mathbf{W} \sim \mathcal{N}(0, \sigma_w^2)$ . In order to push the weight distribution towards  $\mathcal{U}(-1, 1)$ , the weights are initialized with a target distribution  $\mathcal{N}(0, \sigma_t^2)$  and a regularizer  $\lambda_\sigma(\sigma_w - \sigma_t)^2 + \lambda_\mu\mu_w^2$  is added during training to penalize deviations from the target weight distribution, where  $\sigma_w$  and  $\mu_w$  are the standard deviation and mean of the weights. An example of the resultant distribution is shown in Fig. 1(b).

### 4.3 2<sup>nd</sup> Stage: Linear Quantization of Full Network

In this section we describe the second stage, providing details on quantization of gain and bias, batch norm, activations, and lastly the inputs. We then perform one epoch of training to reestimate parameters.



**Fig. 1.** Distributions for (a) original dense layer weights with a nearly Gaussian distribution; (b) dense layer weights with a flatter distribution after  $\tanh(\cdot)$ .

**Quantization of BatchNorm.** Recall that BatchNorm (BN) has four sets of parameters,  $\gamma$ ,  $\beta$ ,  $\mu_B$  and  $\sigma_B^2$ . While  $\gamma$  and  $\beta$  are trainable parameters,  $\mu_B$  and  $\sigma_B^2$  are the empirical estimates of means and variances. These are running estimates of mean and standard deviation during training, while they are replaced by global estimates during inference. In this work, we quantize all four sets of parameters.

From the 1<sup>st</sup>-stage, we inherit floating point representations of the four sets of parameters. During this stage of training (i.e., 2<sup>nd</sup>), we quantize these parameters. To quantize the BN parameters, given a preferred bit-width, we apply linear quantization on  $\gamma$ ,  $\beta$ ,  $\mu_B$  and  $\sigma_B^2$  separately. Following [23], BN parameters, especially  $\mu_B$  and  $\sigma_B^2$ , tend to have large dynamic range. We also observe that the parameters continuously shift during training. We start with  $\mu_B$  and  $\sigma_B^2$ : to reduce the dynamic range, we introduce a scale factor  $C_{BN}$  and quantize  $\mathbf{x}$ ,  $\mu_B$  and  $\sigma_B^2$  with a linear quantizer ( $q_z(\cdot)$ ):

$$\begin{aligned} \mathbf{x}^{\text{norm}} &= \frac{\mathbf{x} - \mu_B}{\sqrt{\sigma_B^2}} = \frac{\frac{\mathbf{x}}{C_{BN}} - \frac{\mu_B}{C_{BN}}}{\sqrt{\frac{\sigma_B^2}{C_{BN}^2}}} \\ \mathbf{x}_q^{\text{norm}} &\leftarrow \frac{q_z\left(\frac{\mathbf{x}}{C_{BN}}\right) - q_z\left(\frac{\mu_B}{C_{BN}}\right)}{\sqrt{q_z\left(\frac{\sigma_B^2}{C_{BN}^2}\right)}} \end{aligned} \quad (2)$$

We then perform the remaining steps in quantizing a BN transformation, with the application of the  $q_z(\cdot)$  on  $\gamma$  and  $\beta$  to obtain  $q_z(\psi_{bn}(\mathbf{x}))$ :

$$\begin{aligned} \gamma_q &\leftarrow q_z(\gamma) \\ \beta_q &\leftarrow q_z(\beta) \\ q_z(\psi_{bn}(\mathbf{x}_q^{\text{norm}})) &\leftarrow \gamma_q \cdot \mathbf{x}_q^{\text{norm}} + \beta_q \end{aligned}$$

**Quantization of Gain and Bias.** We apply a linear quantizer on gain ( $\alpha$ ) and bias ( $\mathbf{b}$ ) to obtain the respective quantized representations  $\mathbf{b}_q$  and  $\alpha_q$ .

**Quantization of Activation Functions.** A standard dense layer is typically followed by a non-linear activation function and a BN layer in order. Our initial experiments showed that a direct linear quantization of ReLU can lead to a large drop in accuracy. To mitigate this, since ReLU is unbounded on the positive domain, we experimented with bounded (clipped ReLU, sigmoid, and tanh) as well as with smoother activation functions (GeLU, SiLU)[24,25]. Clipped ReLU was the most promising activation function in our initial study; experiments with GeLU and SiLU did not yield conclusive results; while sigmoid and tanh yielded worse performance than with a direct quantization of ReLU (presumably require more training updates due to gradient saturation). The output of clipped ReLU is processed with the linear quantizer to obtain  $q_z(\phi_{relu}(\cdot))$ . We also found that with quantization, the order of BN and the activation matters. Specifically, we switch the order of BN and clipped ReLU activations.

**Quantization of Inputs.** The input LFBE features are processed using global mean and variance normalization. Subsequent to this step, the normalized input is processed similar to Section 4.3, such that  $\mathbf{x}_q^{\text{input}} \leftarrow q_z(\frac{\mathbf{x}^{\text{input}}}{C_{input}})$ .

## 5 Experimental Setup

In this section we describe our training and test datasets; we also discuss the models and evaluation metrics. All experiments in this paper were conducted on de-identified production datasets.

**Datasets:** For our experiments we used a fully labeled training dataset consisted of 26K hours of audio. The training dataset contains both far-field audio and near-field mobile phone audio. We used two test sets in this work: (a) a validation test set (referred to as VAL), which consisted of about 4K hours of audio data, (b) an independent test set (referred to as TEST), which also consisted of about 4K hours of audio data. Both VAL and TEST data contain far-field audio and near-field mobile phone audio. TEST data was collected from a wider range of commodity devices.

**Evaluation metrics:** During inference, we tuned the WMA values for the models on held-out datasets. We measure the model performance with DET curves having False Rejection Rate (FRR) on the x-axis and False Discovery Rate (FDR) on the y-axis. Similar to [18,26,21], we normalize the axes of DET curves and report relative FDR. In the interest of space, we do not present the full DET curves for all experiments; in such cases we only report relative degradation in FDR at the FRR of the baseline model’s operating point (OP), where the chosen OP corresponds to the 1.0 in relative FRR.

**Models:** All our models use the architecture described in Section 2. The proposed QAT technique, presented in Section 4, is applied to different components of the model, and the results are described in Section 6. We also study the effect of number of bits for parameters (weights, biases, gain, BN parameters), input and activations. Details of the models are summarized in Table 1. M0 is a full precision model. M1 uses 8-bit parameters and 16-bit activation and input. M2,

**Table 1.** Model architecture including hidden layer size, # param., # bits for weights, # bits for gain ( $\alpha$ ), bias ( $b$ ), and BN parameters ( $\mu, \sigma, \gamma, \beta$ ), and input and activation.

ID	Layer Size	# Param.	# bits weight ( $\alpha, b, BN$ )	# bits act.	# bits input
M0	{87, 400, 87, 400, 87, 400}	250K	32	32	32
M1	{87, 400, 87, 400, 87, 400}	250K	8	8	16
M2	{87, 400, 87, 400, 87, 400}	250K	8	8	8
M3	{87, 400, 87, 400, 87, 400}	250K	5	8	8
M4	{87, 400, 87, 400, 87, 400}	250K	4	8	8

M3 and M4 use weights quantized to 8, 5, and 4 bits respectively, while the activations and input are quantized to 8-bits. Models M0 to M4 have 250K learnable parameters.

## 6 Results

In this section, we provide a detailed study of the proposed 2-stage QAT approach, against an unquantized full precision model in terms of (a) accuracy, (b) memory and computation. Our results are presented in 2 groups: a) Non-binary, sub 8-bit models in Sections 6.1 and 6.2, for ARM NEON instruction set; b) Binary (1-bit) weight models in Section 6.3, for off-the-shelf accelerators.

### 6.1 Non-binary Sub-8 bit Models: 1<sup>st</sup>-Stage Training

To study the effectiveness of  $\tanh(\cdot)$ , we train models with 5-bit dense layer weights for 100K updates, with and without  $\tanh$ ; note that other components are not quantized in this experiment, and that all trainable parameters are updated. Table 2 presents the results, comparing against an unquantized baseline model. The table also presents results with two quantizers for weights: a) the non-linear quantization with  $\tanh(\cdot)$ ; b) a linear quantizer without  $\tanh(\cdot)$ . The model with  $\tanh(\cdot)$  achieves a 7.4% relative degradation in FDR, while the model without  $\tanh(\cdot)$  yields a 12.3 % degradation in FDR.

**Table 2.** Relative degradation in FDR at baseline model’s FRR on VAL. Models have 5-bit weights with or without  $\tanh(\cdot)$ .

Quantization	Rel. FDR (%)
Unquantized model	0.0 (Baseline)
With $\tanh(\cdot)$	7.4
Without $\tanh(\cdot)$	12.3

## 6.2 Non-binary Sub-8 bit Models: 2<sup>nd</sup>-Stage Training

In this section, we compare the performance of the quantized models against the baseline unquantized model. The baseline model is trained for 500K model updates. For the proposed QAT approach, the models were trained for 500K and 35K updates in the 1<sup>st</sup> and 2<sup>nd</sup>-stage training respectively. Firstly, we discuss the accuracy implications, and then present results in terms of memory and CPU gains (for ARM NEON instruction set).

**Evaluation of Accuracy.** From Table 3, M1 with the proposed 2-stage training algorithm yields a 3.1% relative degradation in FDR. To study the effect of number of bits for weights, we quantized M2, M3 and M4 with 2-stage training. Notice that going from M0 to M1, there is a very small increase in relative degradation in FDR at FRR (3.1% increase). Here the changes include: 32 to 8-bit weights; 32 to 8-bit quantization of other parameters (bias, gain, BN parameters); and 32 to 16-bit input and activations. Further reducing the inputs and activations from 16-bit to 8-bit (i.e., M1 to M2) does not lead to an increase in FDR (3.1% to 3.5%). A further change in weights from 8 to 5 or 4-bits (i.e., M2 to M3 or M4) leads to a small increase in FDR (3.5% to 7.4% or 8.7%).

**CPU, Memory, and Model Size Gains.** Reduction in compute resources (memory, CPU) is dependent on hardware architecture. In Table 3, we present CPU, memory consumption on ARM NEON.

**CPU:** We observe a 45% reduction in CPU when using 8-bit activation and input versus 16-bit activation and input. This reduction in bandwidth is due to 8-bit SIMD processors achieving a throughput of 2x computation with 8x8 bit multiplications compared to 8x16 bit multiplications. Note that going from a full precision M0 model to M1 (a 8-bit parameters, 16-bit activations and input), lead to a 50% saving in CPU consumption. Also, since 8-bit containers are used, further reduction in bit widths (i.e. from M2 to M4), do not yield gains in MCPS.

**Memory and Model size:** We observe a 75% reduction in memory by using 8-bit weights, input and activations when compared to non-quantized, 32-bit models. The smallest data type in ARM NEON is 8-bit, so we don't observe any

**Table 3.** Performance of models in terms of rel. degradation in FDR at FRR against baseline full precision model on TEST data. CPU, memory consumption, and model size for ARM NEON in MCPS and KB (MCPS refers to Million Cycles Per Second).

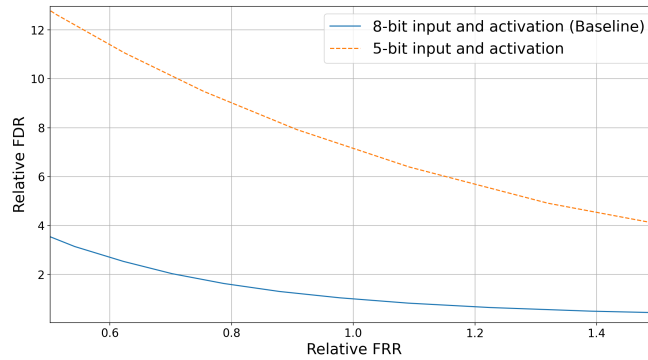
ID	# bits weight	# bits act.	# bits input	Quant. Method	Rel. FDR Degrad. (%)	CPU (MCPS)	Memory (kb)	Model size (kb)
M0	32	32	32	No Quant.	0.0 (Baseline)	37.6	912	912
M1	8	16	16	2-stage	3.1	18.8	228	228
M2	8	8	8	2-stage	3.5	10.3	228	228
M3	5	8	8	2-stage	7.4	10.3	228	142
M4	4	8	8	2-stage	8.7	10.3	228	114

**Table 4.** *Relative degradation in FDR at FRR against baseline on TEST, compressing weights, other parameters, input and activations to 8-bit.*

Quantization on	Rel. FDR Degrad. (%)
weight	3.8
weight + bias + gain	4.7
weight + input + activation	5.2
weight + $BN_1$	5.3
weight + $BN_2$	9.0

memory reduction by using sub-8-bit weights. However, we study and use sub 8-bit weights because this effectively provides more headroom for accumulation of weights, input and activations – reducing the likelihood of overflow, and using a lower-precision accumulator. We also note the model size, which decreases linearly with bit width.

**Some Ablation Studies.** In this segment, we present 3 sets of ablation studies: 1) quantizing different components of the network to 8-bits; 2) quantizing inputs and activations from 8-bit to 5-bit; 3) quantizing other components of the network from 8-bit to 5-bit (bias, gain, batch norm parameters).



**Fig. 2.** *On TEST: DET curves showing performance of quantized model with 8-bit input and activation (Baseline, Blue) and 5-bit quantization (Orange)*

**1) 8-bit models:** Fixing 8-bit quantization for all components of the network, we perform an ablation study for accuracy implications in terms of quantizing bias, gain, BN parameters, input and activations. The unquantized baseline model and the quantized models in this section were trained for 100K updates. Table 4 presents the results in terms of relative degradation in FDR at equal FRR compared to the baseline model on TEST data. With  $\tanh(\cdot)$  8-bit weight

quantization we obtain a 3.8 % relative degradation in FDR. With further quantization of input and activations (5.2%) or bias and gain (4.7%) or batch norm parameters (5.3%), we did not see a significant decrease in model performance. However, it is interesting to note from the Table 4 that the dynamic scaling of BN yields an improvement (i.e., from 9.0% to 5.3% relative FDR degradation).

**2) 8-bit to 5-bit change in input and activations:** While compressing input and activations from 16-bit to 8 bits (from M1 to M2 in Table 3) only results in a minor performance degradation, reducing further from 8-bit to 5-bit input and activations leads to a significant drop in performance (see Fig. 2).

**3) 8-bit to 5-bit change in parameters:** Similar to the small degradation in FDR in going from 8 bits to 5 bits for weights (Table 3), reducing the bit width from 8 to 5 for other parameters does not lead to a significant degradation in FDR (3.0% increase in relative FDR).

### 6.3 Binary Weight Models

In this section, we present results for models with 1-bit weights using 2-stage training; note that the other parameters, input and activations are non-binary (i.e., either quantized or full precision). The models were trained for 500K and 35K updates in the 1<sup>st</sup> and 2<sup>nd</sup>-stage training respectively.

We experimented with two quantized models, with 2-times and 6-times the number of parameters as the baseline model (250K parameters): a) M5 with 500K parameters – where the first, second, and the last layer weights are 5-bit, the other layer weights are 1-bit, and all other parameters, input and activations being 8-bit; b) M6 with 1.5M parameters – where all layer weights are 1-bit, and other parameters, input and activations being 8-bit. Results are reported in Table 5. In addition to FDR, we report the model size and project reduction in memory utilization with off-the-shelf neural network accelerators.

In our experiments, models having binary weights are sensitive to learning rates. Although Adam optimizer is employed in model training, we find that manually changing learning rate to a smaller number for the last few model updates of the 1<sup>st</sup>-stage improves the model performance for M5 and M6. In Table 5, we find that the learning rate change results in 13.6% and 6.6% reduction in relative degradation in FDR against baseline model (M0) for M5 and M6

**Table 5.** Performance of the models in terms of rel. degradation in FDR at FRR matching baseline full precision model on TEST. Hidden layer sizes of M5 are {123, 566, 123, 566, 123, 566}, while those of M6 are {1080, 566, 123, 566, 123, 566}.

ID #	Param.	Rel. FDR Degrad. (%)	Memory (kb)	Model size (kb)	Note
M0	250K	0.0	912	912	Baseline
M5	500k	14.0	128	128	with large LR
M6	1.5M	6.2	198	198	with large LR
M5	500k	0.4	128	128	with LR change
M6	1.5M	-0.4	198	198	with LR change

respectively. With the LR change, M5 and M6 are able to achieve performance parity with M0. With off-the-shelf neural network accelerators, projected memory utilization is equivalent to the model size, with 86.0% and 78.3% reduction in memory consumption with M5 and M6.

## 7 Conclusions

We proposed a 2-stage algorithm for sub 8-bit quantization of 250K parameter KWS models. For the 1<sup>st</sup>-stage, we adapted a recently proposed QAT technique using a non-linear transformation on weights. In the 2<sup>nd</sup>-stage, we used linear quantization methods on other parameters, input and activations. This paper happens in the setting of on-device, low footprint, streaming KWS models, that being explored on two embedded chipset settings, where we achieved parity in accuracy against a full precision model in terms of FDR at a chosen FRR. With sub 8-bit non-binary weight models, on an ARM NEON architecture, we match accuracy, and obtain up to 3 times improvement in CPU consumption and more than 4 times improvement in memory consumption. With a binary weight model (and other components being 8-bit), using off-the-shelf neural network accelerators, at accuracy parity with a full precision model, we project 4 to 7 times reduction in memory consumption and model size.

## References

1. G. Chen, C. Parada, and G. Heigold, “Small-footprint keyword spotting using deep neural networks,” in *Proc. of ICASSP*, 2014.
2. S. Panchapagesan, M. Sun, A. Khare, S. Matsoukas, A. Mandal, B. Hoffmeister, and S. Vitaladevuni, “Multi-task learning and weighted cross-entropy for dnn-based keyword spotting.” in *Proc. of Interspeech*, 2016.
3. G. Tucker, M. Wu, M. Sun, S. Panchapagesan, G. Fu, and S. Vitaladevuni, “Model compression applied to small-footprint keyword spotting.” in *Proc. of Interspeech*, 2016.
4. R. Prabhavalkar, O. Alsharif, A. Bruguier, and L. McGraw, “On the compression of recurrent neural networks with an application to lvsr acoustic modeling for embedded speech recognition,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5970–5974.
5. M. Sun, D. Snyder, Y. Gao, V. K. Nagaraja, M. Rodehorst, S. Panchapagesan, N. Strom, S. Matsoukas, and S. Vitaladevuni, “Compressed time delay neural network for small-footprint keyword spotting.” in *Proc. of Interspeech*, 2017.
6. S. Mittermaier, L. Kürzinger, B. Waschneck, and G. Rigoll, “Small-footprint keyword spotting on raw audio data with sinc-convolutions,” in *Proc. of ICASSP*, 2020.
7. X. Li, X. Wei, and X. Qin, “Small-footprint keyword spotting with multi-scale temporal convolution,” *arXiv preprint arXiv:2010.09960*, 2020.
8. P. Blouw, G. Malik, B. Morcos, A. R. Voelker, and C. Eliasmith, “Hardware aware training for efficient keyword spotting on general purpose and specialized hardware,” *arXiv preprint arXiv:2009.04465*, 2020.

9. M. Vandersteegen, K. Van Beeck, and T. Goedemé, “Integer-only cnns with 4 bit weights and bit-shift quantization scales at full-precision accuracy,” *Electronics*, 2021.
10. B. Shi, M. Sun, C.-C. Kao, V. Rozgic, S. Matsoukas, and C. Wang, “Compression of acoustic event detection models with low-rank matrix factorization and quantization training,” *arXiv preprint arXiv:1905.00855*, 2019.
11. Y. Mishchenko, Y. Goren, M. Sun, C. Beauchene, S. Matsoukas, O. Rybakov, and S. N. P. Vitaladevuni, “Low-bit quantization and quantization-aware training for small-footprint keyword spotting,” in *Proc. of IEEE International Conference On Machine Learning And Applications(ICMLA)*, 2019.
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. 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.
14. Y. Bengio, N. Léonard, and A. Courville, “Estimating or propagating gradients through stochastic neurons for conditional computation,” *arXiv preprint arXiv:1308.3432*, 2013.
15. E. Jang, S. Gu, and B. Poole, “Categorical reparameterization with gumbel-softmax,” *arXiv preprint arXiv:1611.01144*, 2016.
16. H. D. Nguyen, A. Alexandridis, and A. Mouchtaris, “Quantization aware training with absolute-cosine regularization for automatic speech recognition.” in *Proc. of Interspeech*, 2020.
17. N. Strom, H. Khan, and W. Hamza, “Squashed weight distribution for low bit quantization of deep models,” in *Submitted to Proceedings of Interspeech*, 2022.
18. Y. Gao, N. D. Stein, C.-C. Kao, Y. Cai, M. Sun, T. Zhang, and S. Vitaladevuni, “On front-end gain invariant modeling for wake word spotting,” *arXiv preprint arXiv:2010.06676*, 2020.
19. S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International conference on machine learning*, 2015, pp. 448–456.
20. A. F. Agarap, “Deep learning using rectified linear units (relu),” *arXiv preprint arXiv:1803.08375*, 2018.
21. C. Jose, Y. Mishchenko, T. Senechal, A. Shah, A. Escott, and S. Vitaladevuni, “Accurate detection of wake word start and end using a cnn,” in *Interspeech*, 2020.
22. V. Vanhoucke, A. Senior, and M. Z. Mao, “Improving the speed of neural networks on cpus,” in *Deep Learning and Unsupervised Feature Learning Workshop, NIPS 2011*, 2011.
23. R. Banner, I. Hubara, E. Hoffer, and D. Soudry, “Scalable methods for 8-bit training of neural networks,” *Advances in neural information processing systems*, vol. 31, 2018.
24. D. Hendrycks and K. Gimpel, “Gaussian error linear units (gelus),” *arXiv preprint arXiv:1606.08415*, 2016.
25. S. Elfving, E. Uchibe, and K. Doya, “Sigmoid-weighted linear units for neural network function approximation in reinforcement learning,” *Neural Networks*, vol. 107, pp. 3–11, 2018.
26. M. Sun, D. Snyder, Y. Gao, V. Nagaraja, M. Rodehorst, S. Panchapagesan, N. Ström, S. Matsoukas, and S. Vitaladevuni, “Compressed time delay neural network for small-footprint keyword spotting,” in *Interspeech*, 2017.