



One-pass Multiple Conformer and Foundation Speech Systems Compression and Quantization Using An All-in-one Neural Model

Zhaoqing Li^{1,*}, Haoning Xu^{1,*}, Tianzi Wang¹, Shoukang Hu², Zengrui Jin¹, Shujie Hu¹,
Jiajun Deng¹, Mingyu Cui¹, Mengzhe Geng³, Xunying Liu¹

¹The Chinese University of Hong Kong; ²Nanyang Technological University;
³National Research Council Canada

zqli@se.cuhk.edu.hk, hnxu@se.cuhk.edu.hk, xyliu@se.cuhk.edu.hk

Abstract

We propose a novel one-pass multiple ASR systems joint compression and quantization approach using an all-in-one neural model. A single compression cycle allows multiple nested systems with varying Encoder depths, widths, and quantization precision settings to be simultaneously constructed without the need to train and store individual target systems separately. Experiments consistently demonstrate the multiple ASR systems compressed in a single all-in-one model produced a word error rate (WER) comparable to, or lower by up to 1.01% absolute (6.98% relative) than individually trained systems of equal complexity. A 3.4x overall system compression and training time speed-up was achieved. Maximum model size compression ratios of 12.8x and 3.93x were obtained over the baseline Switchboard-300hr Conformer and LibriSpeech-100hr fine-tuned wav2vec2.0 models, respectively, incurring no statistically significant WER increase.

Index Terms: ASR, model compression, model quantization, self-supervised learning, knowledge distillation, transformer

1. Introduction

Current automatic speech recognition (ASR) systems based on Transformer [1] and their variants [2, 3] are becoming increasingly complex for practical applications. Such trend is further exasperated by the emergence and rapid proliferation of self-supervised learning (SSL) based speech foundation models represented by wav2vec2.0 [4], HuBERT [5], data2vec [6], and WavLM [7]. Their practical deployment on ad-hoc edge devices such as mobile phones, self-driving cars, and smart-watches require their respective computational resource constraints to be met. To this end, there is a pressing need to develop high-performance speech AI model compression techniques that flexibly account for the fine-grained performance-efficiency trade-off granularity across diverse user devices.

Prior research for Transformer based speech processing models has largely evolved into two categories: **1) architecture compression** methods that aim to minimize the Transformer model structural redundancy measured by their depth, width, sparsity, or their combinations using techniques such as pruning [8–10], low-rank matrix factorization [11, 12] and distillation [13, 14]; and **2) low-bit quantization** approaches that use either uniform [15–18], or mixed precision [12, 19] settings. A combination of both architecture compression and low-bit quantization approaches has also been studied to produce larger model compression ratios [12]. Architecture compression techniques for the recently emerging SSL speech foundation models have also been developed using weight pruning [20–22],

distillation [23–30], or both [31]. Only very limited prior research [32–34] studied SSL speech model quantization.

The above existing researches suffer from the following limitations: **1) weak scalability** when being used to produce compressed systems of varying target complexity that are tailored for diverse user devices. The commonly adopted approach requires each target compressed system with the desired size to be individually constructed, for example, in [14, 15, 17] for Conformer models, and similarly for SSL foundation models such as DistilHuBERT [23], FitHuBERT [24], DPHuBERT [31], PARP [20], and LightHuBERT [30] (no more than 3 systems of varying complexity were built). **2) limited scope of system complexity attributes** covering only a small subset of architecture hyper-parameters based on either network depth or width alone [8, 9, 11, 35, 36], or both [10, 13, 14, 37], while leaving out the task of low-bit quantization, or vice versa [15–19, 32–34]. This is particularly the case with the recent **HuBERT** model distillation research [23–25, 28–31] that are focused on architectural compression alone. **3) restricted application** to either standard supervised-learning based Transformers trained on a limited quantity of data [8–19] or their SSL foundation models alone [20–34]. Hence, a more comprehensive study on both model types is desired to offer further insights into the efficacy of deploying current compression and quantization techniques on speech foundation models such as Whisper [38]. In addition, statistical significance tests are missing in most existing studies that are important for measuring the efficacy of model compression in terms of ASR performance degradation.

To this end, this paper proposes a novel one-pass multiple ASR systems compression and quantization approach using an all-in-one neural model for both conventional supervised-learning based Conformer and SSL based wav2vec2.0 models. A single system compression cycle allows a large number of weight-sharing target systems (e.g., up to 8) with different compression ratios achieved by using varying Encoder depths, widths, and quantization precision settings to be simultaneously constructed without the need for separate training and storing individual target systems one by one. Experiments conducted using the 300-hr Switchboard dataset trained Conformer and LibriSpeech-100hr fine-tuned wav2vec2.0 SSL models consistently demonstrate the multiple ASR systems being compressed in a single pass produced word error rate (WER) comparable to, or lower by up to 1.01% absolute (6.98% relative) than those obtained using individually compressed baseline systems of equal complexity. An overall system compression and training time speed-up ratio of 3.4x was achieved. Maximum model size compression ratios of 12.8 and 3.93 times were obtained on the two tasks over the 32-bit full-precision baseline Conformer and wav2vec2.0 models, respectively, while incurring non-statistically significant or negligible WER increase.

* Equal Contribution

The main contributions of the paper are summarized below:

1) To the best of our knowledge, this is the first work that develops an all-in-one neural model for joint architecture compression and parameter quantization of multiple nested Conformer or wav2vec2.0 ASR systems. 2) The one-pass joint compression-quantization approach can be applied to optimize a rich set of system complexity attributes covering network depth, width, precision, or their combinations. 3) The efficacy of our one-pass joint compression-quantization approach is demonstrated consistently for both Conformer and wav2vec2.0 models, offering a more comprehensive study than prior research restricted to only one of them [9–18, 20–22, 26, 33]. 4) Extensive experimental results demonstrate that multiple ASR sub-networks contained in all-in-one models produced WERs comparable to, or lower by up to 1.01% absolute (6.98% relative) than, those obtained using individually compressed baseline systems of equal complexity. An overall system compression and training time speed-up ratio of 3.4x was achieved.

2. Conformer AED System and Wav2vec2.0 Foundation Model

A Conformer [2] Encoder comprises multiple stacked layers. For example, in Fig 1(a), the Encoder could contain up to 12 layers, which we refer to as **model depth**. Each Encoder layer consists of the following modules sequentially: a position-wise feed-forward module (FFN), a multi-head self-attention module (MHSA), a convolution module (Conv), and a second FFN module (macaron-like) in the end. Among these modules, FFN accounts for most of the Encoder parameters. Particularly, it consists of two consecutive linear networks, where the first one expands the model dimensionality (e.g., from 256 to 2048, as shown in Fig 1(b)), and the second one restores it to before. Thus, we refer to **model width** as the intermediate dimensionality of FFN modules.

To train an attention-based encoder-decoder (AED) Conformer ASR system, an effective way is using multitask criterion interpolation [39] between the loss of CTC and the attention Decoder (i.e., a Transformer decoder). This is given by

$$\mathcal{L}_{conformer} = (1 - \lambda)\mathcal{L}_{att} + \lambda\mathcal{L}_{ctc}, \quad (1)$$

where λ is a constant and empirically set as 0.2 in this paper.

Speech SSL models such as wav2vec2.0 [4], HuBERT [5], and WavLM [7] share similar Transformer backbones with supervised models. For example, wav2vec2.0 consists of a CNN based feature extractor and a Transformer Encoder, where each Encoder layer contains an MHSA module and an FFN module. In this paper, we fine-tune a pre-trained wav2vec2.0 model with a pure CTC Decoder, namely $\mathcal{L}_{w2v2} = \mathcal{L}_{ctc}$.

In this paper, we perform compression and quantization on the entire Conformer Encoder and the Transformer Encoder excluding the CNNs for wav2vec2.0 models.

3. One-Pass Multiple Systems Compression

3.1. All-In-One Weight-Sharing Neural Model

Similar to once-for-all techniques [40–42], an all-in-one model refers to a neural network containing various sub-networks, where all small networks share parameters with larger ones. For example, in Fig. 1, an 8-layer 1024-dim 4-bit small model is nested in the 12-layer 2048-dim 8-bit all-in-one model. Once an all-in-one model is trained, the target sub-networks are automatically contained, resulting in no retraining and storage overhead to compress the model to a smaller one. In this paper, we propose an all-in-one model that offers flexibility in model

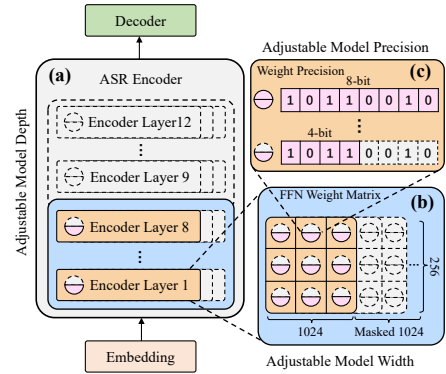


Figure 1: Diagram of an all-in-one ASR Encoder. A partly colored area denotes a sub-network, and all sub-networks share weights with larger ones. (a) Sub-networks can skip top layers. (b) Sub-networks can mask partial FFN weight matrices. (c) Sub-networks can be directly quantized to a lower bit-width.

depths, widths, quantization precision, and their combinations for both Conformer and wav2vec2.0.

3.2. Multiple Weight-Sharing Systems Training

To obtain an all-in-one neural network, a one-pass multitask training approach is needed to ensure all the sub-networks are also well trained. The training criterion is given by

$$\mathcal{L} = \mathcal{L}_{max} + \sum_{i \in \mathcal{M}} \lambda_i^1 \mathcal{L}_i \quad (2)$$

where \mathcal{L}_{max} is the loss of the largest network, \mathcal{M} is the set of other target sub-network structures, and \mathcal{L}_i is the loss of sub-network i , with a constant coefficient λ_i^1 . In Eq. 2, the first term is exactly the same as conventional single system training, and the second term is to enable sub-networks. Note that in our default setup, the largest model is the same size as our chosen baselines since a major goal of this paper is model compression.

Since all sub-networks are nested in larger ones, their losses are obtained by forwarding through only a partial of the largest network. Specifically, when compressing through depth, the several top layers can be skipped (Fig. 1(a)). For width, a certain portion of rows (or columns) of FFN weight matrices can be masked¹ with zero before forwarding (Fig. 1(b)). For quantization, the weights can be directly quantized to the target bit-width before forwarding (Fig. 1(c)). Therefore, with such methods, we only need to store the largest network in the end while multiple sub-networks are already automatically contained.

3.3. Kullback–Leibler Divergence Regularization

Like the effect of knowledge distillation, when simultaneously training multiple systems, larger models are implicitly providing guidance to smaller ones, while we can make this effect explicit by adding an extra Kullback–Leibler (KL) divergence regularization. Denote p_{max} , and p_i , ($i \in \mathcal{M}$) as the output distribution of the largest network and sub-network i , respectively. The KL regularization term is given by

$$\Omega_i = D_{KL}(SG(p_{max}) || p_i), \quad (3)$$

where $SG(\cdot)$ denotes the stop-gradient operation preventing the gradient from flowing back to the largest structure. Finally, we can write the criterion with KL regularization:

$$\mathcal{L}_{kl} = \mathcal{L}_{max} + \sum_{i \in \mathcal{M}} (\lambda_i^1 \mathcal{L}_i + \lambda_i^2 \Omega_i), \quad (4)$$

where λ_i^1 and λ_i^2 are constant coefficients.

¹For wav2vec2.0, rows (columns) having small L_2 -norm values are considered of less importance and masked with high-priority.

4. Experiments

4.1. Experimental Setup

Baseline Models. The baseline Conformer system is configured using the ESPnet [43] recipe². For wav2vec2.0 models, the wav2vec2-base-100h is downloaded from Huggingface³. We fine-tuned wav2vec2-base-100h for 5 epochs as our baseline.

Data. The Conformer systems are trained with 300-hr benchmark Switchboard corpus [44] and evaluated on NIST Hub5'00, RT02, and RT03 evaluation sets. We fine-tune the wav2vec2-base-100h model on LibriSpeech's 100-hour clean subset [45].

Training. All the Conformer systems, including one-pass and individually compressed ones, are trained from scratch with exactly the same learning schedule as the baseline Conformer recipe. For wav2vec2.0 systems, we fine-tune wav2vec2-base-100h with a batch size of around 202 seconds of audios for 5 epochs, which is consistent with our baseline. We use the AdamW optimizer with a learning rate of 3e-5. A linear warmup is used for the first 10% training steps, followed by a linear decay to zero. All the experiments in this paper are conducted on a single NVIDIA A40 (48G) GPU.

Quantization. In this paper, we apply symmetric quantization on model weights using quantization-aware training (QAT) with a straight-through estimator and a learnable scale factor [46].

4.2. Results of Conformer Systems

Single Attribute Compression: We first conduct experiments on Conformer systems with all-in-one models sharing weights through each network attribute separately. We study two configurations for each attribute, i.e., {8, 12} layers for model depth, {1024, 2048} for model width, and {4-bit, 8-bit} for weight precision. Results evaluated on Hub5'00, RT02, and RT03 sets are shown in Tab. 1. For a clearer comparison, we also reported the average WER over all three test sets, where several trends can be found: **1)** On whichever attribute, all the sub-networks contained in all-in-one models produce significantly lower word error rate (WER) than individually compressed systems of equal network architecture. (e.g., in Tab. 1, Sys. 2s and 3s vs. Sys. 1; Sys. 9s vs. Sys. 7, etc.). **2)** On whichever attribute, the largest sub-network (with the same size as the baseline) of all-in-one model produces significantly lower WER than the baseline system (i.e., Sys. 2l, 3l, 5l, 6l, 9l, 10l vs. Sys. 0 in Tab. 1). Remarkably, some small sub-networks even have performance better than (i.e., Sys. 6s and 10s vs. Sys. 0, in Tab. 1) or comparable to (i.e., Sys. 5s and 9s vs. Sys. 0, in Tab. 1) the baseline. **3)** On whichever attribute, the KL regularization⁴ further improves the performance of all-in-one models (e.g., Sys. 3l vs. Sys. 2l; Sys. 6s vs. Sys. 5s, etc.).

Multiple Attribute Compression: On top of compressing separately, multiple attribute compression of Conformer systems with all-in-one models is studied in Tab. 2, where the architecture configuration list is the combination of attributes in Tab. 1, and different attributes can be compressed simultaneously. For Conformer all-in-one models, when training, we only forward through 3 sub-networks (i.e., the largest and smallest one, and another one sampled from the other sub-networks) in one updating iteration instead of all of them to speed up the training process. Also, we average all the test sets for clearer comparisons. Several trends can be found: **1)** All the sub-networks in the all-in-one model consistently produce signif-

Table 1: WER(↓) of all-in-one Conformer models against individually compressed Conformer systems through a single network attribute. Systems with the same Sys. Number share weights, where “ℓ” and “s” are for the largest and small sub-network, respectively. Blue cells: the performance is better than the baseline (Sys. 0) with a statistical significance (MAPSSWE [47], $\alpha=0.05$). Yellow cells: the performance has no significant difference from the baseline. * means it is significantly better than the individually compressed system of equal model size (Sys. 1, 4, 7, and 8).

Sys.	Attribute	All-in-one	+ KL	# Param.	Hub5'00			RT02			RT03		Avg.
					swbd	calhm	swbd1	swbd2	swbd3	swbd	fsk		
0	12-2048-32bit (baseline)				34m	7.4	15.2	8.9	13.0	15.9	10.6	16.6	12.86
1	Dep.	8	×	-	23m	7.8	15.9	9.7	13.4	16.5	11.2	17.5	13.49
2s		8	✓	×	share	7.8	15.6	9.2	13.7	16.4	10.7	17.0	13.25*
2l		12	✓	×	34m	7.4	15.0	9.0	13.1	15.8	10.3	16.3	12.71*
3s		8	✓	✓	share	7.8	15.8	9.4	13.3	16.5	10.8	16.9	13.23*
3l	12	✓	✓	34m	7.3	14.9	8.8	12.7	15.7	10.4	16.2	12.59*	
4	Wid.	1024	×	-	21m	7.7	15.5	9.0	13.3	16.2	10.8	16.9	13.06
5s		1024	✓	×	share	7.3	15.1	9.0	13.1	15.8	10.6	16.9	12.88*
5l		2048	✓	×	34m	7.1	14.8	8.7	12.6	15.5	10.3	16.5	12.54*
6s		1024	✓	✓	share	7.4	15.0	8.8	12.7	15.2	10.2	16.4	12.54*
6l	2048	✓	✓	34m	7.3	14.8	8.5	12.5	14.9	10.0	16.1	12.33*	
7	Prec.	4-bit	×	-	34m	7.6	15.6	8.9	13.3	16.1	10.6	17.0	13.06
8		8-bit	×	-	34m	7.4	15.2	8.9	13.0	15.8	10.6	16.7	12.86
9s		4-bit	✓	×	share	7.5	15.5	9.0	12.9	15.6	10.5	16.9	12.90*
9l		8-bit	✓	×	34m	7.4	15.2	8.8	12.6	15.5	10.4	16.6	12.70*
10s		4-bit	✓	✓	share	7.0	15.2	8.8	12.6	15.5	10.1	16.3	12.54*
10l		8-bit	✓	✓	34m	7.1	15.0	8.8	12.6	15.6	10.1	16.3	12.53*

icantly lower WER than that of individually compressed systems of equal complexity (i.e., Sys. 1A~8A vs. Sys. 1~8 in Tab. 2, respectively). **2)** Also, further performance improvements (or WER decrease) can be found when adding KL regularization, especially for small sub-networks (i.e., Sys. 1K~8K vs. Sys. 1A~8A in Tab. 2, respectively). **3)** With one-pass multitask training and KL regularization, a WER reduction up to 1.01% absolute (6.98% relative) is achieved for the smallest sub-network of the architecture 8-1024-4bit against the individually compressed system of the same architecture (i.e., Avg. of Sys. 8K vs. Sys. 8 in Tab. 2). Moreover, a maximum compression ratio of 12.8 times is obtained against the baseline Conformer system while incurring no statistically significant WER increase (i.e., Avg. of Sys. 4K vs. Sys. 0 in Tab. 2). **4)** An overall compression and training time speed-up ratio of 3.4 times is achieved (see Training time in Tab. 2, where the overall training time is measured by adding up the GPU hours needed for training all the sub-networks, namely 319 hours in total for training individual sub-networks separately). The all-in-one Conformer Encoder only contains 34 million parameters, while it can be re-configured to 8 different-sized systems that would contain 186m parameters in total if trained and stored separately (see #Param. in Tab. 2).

4.3. Results for wav2vec2.0 Systems

We then demonstrate the efficacy of the one-pass all-in-one model on SSL systems. For a fair comparison, we fine-tune individual and all-in-one models with the same iterations.

Single Attribute Compression: We first study single attribute compression for wav2vec2.0 systems, where the configuration of each attribute is set as {11, 12} for model depth, {2560, 3072} for model width, and {6-bit, 8-bit} for weight precision. Results are shown in Tab. 3. Similar trends can be found to those observed in Conformer systems. **1)** Sub-networks in all-in-one models have performance of no significant differences with (e.g., Sys. 3s and 2s vs. Sys. 1, etc.) or even better than (e.g., Sys. 5l vs. Sys. 0; Sys. 10s vs. Sys. 7, etc.) individual systems. **2)** They are further enhanced by KL regularization⁵

²ESPnet: egs/swbd/asr1/run.sh

³Huggingface: facebook/wav2vec2-base-100h

⁴For Conformer systems, we set $\lambda^1 = 0.8$ and $\lambda^2 = 0.2$ for depth, and $\lambda^1 = \lambda^2 = 1$ for other attributes.

⁵For SSL cases, we empirically set $\lambda^1 = 1$ and $\lambda^2 = 0.005$ for

Table 2: $WER(\downarrow)$ of all-in-one Conformer models against individually compressed Conformer systems through multiple network attributes. Systems with the same letter (A or K) share weights; systems with the same number (1~8) share the same network architecture. Blue cells: the performance is better than the baseline (Sys. 0) with a statistical significance (MAPSSWE [47], $\alpha=0.05$). Yellow cells: the performance has no significant difference from the baseline. * means it is significantly better than the individually compressed system of equal model size (Sys. 1~8). † means it has no significant difference with the corresponding individually-compressed system.

Sys.	Architecture	All-in-one	+KL	#Parameters Encoder	Training Time	Compression Ratio	Hub5'00		RT02			RT03		Avg.
							swbwd	calhm	swbwd1	swbwd2	swbwd3	swbwd	fsh	
0	12-2048-32bit (baseline)			34m	37hrs	1x	7.4	15.2	8.9	13.0	15.9	10.6	16.6	12.86
1	12-2048-8bit	x	-	34m	48hrs	8.0x	7.4	15.2	8.9	13.0	15.8	10.6	16.7	12.86
2	12-2048-4bit			34m	48hrs	4.0x	7.6	15.6	8.9	13.3	16.1	10.6	17.0	13.06
3	12-1024-8bit			21m	42hrs	6.4x	7.9	15.7	9.1	13.6	16.4	10.7	17.5	13.33
4	12-1024-4bit			21m	42hrs	12.8x	7.8	16.0	9.5	13.4	16.7	10.9	17.5	13.46
5	8-2048-8bit			23m	37hrs	5.8x	8.0	15.7	9.3	13.5	16.6	11.0	17.4	13.49
6	8-2048-4bit			23m	37hrs	11.7x	8.1	16.2	9.6	14.1	17.3	11.4	18.3	13.96
7	8-1024-8bit			15m	33hrs	8.2x	7.9	16.1	9.5	14.9	16.8	11.2	17.8	13.69
8	8-1024-4bit			15m	32hrs	18.4x	8.4	17.3	10.2	14.7	17.7	11.9	18.7	14.48
1A	12-2048-8bit	✓	x	share 34m	94hrs	4.0x	7.2	15.1	8.7	12.7	15.3	10.2	16.5	12.58*
2A	12-2048-4bit					8.0x	7.3	15.2	9.0	13.0	15.5	19.4	16.8	12.80*
3A	12-1024-8bit					6.4x	7.7	15.3	9.2	13.0	16.0	10.5	17.3	13.06*
4A	12-1024-4bit					12.8x	7.7	15.6	9.3	13.4	16.3	10.6	17.3	13.23*
5A	8-2048-8bit					5.8x	7.8	15.6	9.4	13.4	16.2	10.7	17.4	13.28*
6A	8-2048-4bit					11.7x	7.9	16.1	9.7	13.9	16.6	10.9	17.5	13.54*
7A	8-1024-8bit					8.2x	8.0	16.1	9.6	13.9	16.8	11.1	17.9	13.68†
8A	8-1024-4bit					18.4x	8.2	16.1	9.8	14.2	17.2	11.2	18.1	13.90*
1K	12-2048-8bit	✓	✓	share 34m	95hrs	4.0x	7.2	14.8	8.9	12.6	15.4	10.2	16.3	12.52*
2K	12-2048-4bit					8.0x	7.2	15.0	8.9	12.9	15.5	10.3	16.5	12.66*
3K	12-1024-8bit					6.4x	7.5	15.2	9.1	12.8	15.9	10.3	16.4	12.76*
4K	12-1024-4bit					12.8x	7.6	15.4	9.1	13.0	16.0	10.6	16.8	12.97*
5K	8-2048-8bit					5.8x	7.6	15.3	9.2	13.3	16.3	10.7	16.8	13.04*
6K	8-2048-4bit					11.7x	7.7	15.4	9.2	13.5	16.4	10.8	17.0	13.19*
7K	8-1024-8bit					8.2x	7.8	15.8	9.4	13.4	16.5	10.8	17.3	13.32*
8K	8-1024-4bit					18.4x	7.9	15.8	9.5	13.6	16.7	11.0	17.5	13.47*

Table 3: $WER(\downarrow)$ of all-in-one wav2vec2.0 models against individually compressed wav2vec2.0 systems through a single network attribute. Systems with the same Sys. Number share weights, where “l” and “s” are for the largest and small sub-network, respectively. Colors and marks are with the same meanings as the MAPSSWE [47] test in Table 1.

Sys.	Attribute	All-in-one	+KL	# Parm.	dev		test		
					clean	other	clean	other	
0	12-3072-32bit (baseline)			94m	5.78	13.65	5.89	13.32	
1	Depth	11	x	-	87m	5.88	13.96	5.94	13.54
2s		11	✓	x	share	5.88†	14.10†	5.98†	13.62†
2l		12	✓	x	94m	5.79†	13.87	5.82†	13.31†
3s		11	✓	✓	share	5.95†	14.00†	5.89†	13.59†
3l	12	✓	✓	94m	5.73†	13.81†	5.76*	13.31†	
4	Width	2560	x	-	85m	6.58	15.41	6.55	14.99
5s		2560	✓	x	share	6.47†	15.62	6.56†	15.27
5l		3072	✓	x	94m	5.81†	13.32*	5.77*	13.00*
6s		2560	✓	✓	share	6.37*	15.23†	6.32*	14.82†
6l	3072	✓	✓	94m	5.78†	13.60†	5.75*	13.12†	
7	Prec.	6-bit	x	-	94m	5.86	13.85	5.85	13.30
8		8-bit	x	-	94m	5.80	13.63	5.83	13.26
9s		6-bit	✓	x	share	5.84†	13.96†	5.88†	13.40†
9l		8-bit	✓	x	94m	5.76†	13.62†	5.75†	13.08*
10s		6-bit	✓	✓	share	5.72*	13.46*	5.75†	13.21†
10l		8-bit	✓	✓	94m	5.81†	13.63†	5.71†	13.19†

(e.g., Sys. 6s vs. Sys. 5s, Sys. 10s vs. Sys. 9s, etc.).

Multiple Attribute Compression: Experiments of multiple attribute compression of wav2vec2.0 systems also share trends with those in Conformer systems found in Tab. 2. The results are reported in Tab. 4. Particularly, **1)** a WER reduction up to 0.88% absolute (test other of Sys. 8K vs. Sys. 8 in Tab. 4) and 8.8% relative (dev clean of Sys. 8K vs. Sys. 8 in Tab. 4) is achieved for the wav2vec2.0 model of the compressed architecture 11-2560-6bit against the individually re-tuned system. **2)** A maximum compression ratio of 3.93 times is obtained while incurring no statistically significant WER increase (Sys. 6K vs. Sys. 0 in Tab. 4). **3)** Finally, an overall compression and training time speed-up ratio of 2.2 times is achieved (see Training Time in Tab. 4). In addition, the all-in-one model contains only 94m parameters against 690m in total for 8 individual systems.

depth; $\lambda^1 = 1$ and $\lambda^2 = 0.1$ for other attributes.

Table 4: $WER(\downarrow)$ of all-in-one wav2vec2.0 models against individually compressed wav2vec2.0 systems through multiple network attributes. Systems with the same letter (A or K) share weights, while systems with the same Sys. Number (1~8) share the same network architecture. Colors and marks are with the same meanings as the MAPSSWE [47] test in Table 2.

Sys.	Architecture	All-in-one	+KL	# Parm.	Train Time	Comp. Ratio	dev					
							clean	other	clean	other		
0	12-3072-32bit (baseline)			94m	3.8hrs	1x	5.78	13.65	5.89	13.32		
1	12-3072-8bit	x	-	share 94m	4.0hrs	3.08x	5.62*	13.43†	5.61*	12.94*		
2	12-3072-6bit					3.72x	5.65*	13.60*	5.69*	13.18†		
3	12-2560-8bit					85m	4.2hrs	3.34x	6.53	15.19	6.59	14.96
4	12-2560-6bit					85m	4.0hrs	4.00x	6.62	15.68	6.67	15.24
5	11-3072-8bit					87m	4.0hrs	3.27x	5.93	13.85	5.99	13.33
6	11-3072-6bit					87m	4.0hrs	3.93x	6.01	14.04	6.03	13.68
7	11-2560-8bit					79m	3.9hrs	3.53x	6.68	15.41	6.70	15.14
8	11-2560-6bit					79m	4.0hrs	4.22x	6.83	16.00	6.85	15.63
1A	12-3072-8bit	✓	x	share 94m	22.9hrs	3.08x	5.62*	13.43†	5.61*	12.94*		
2A	12-3072-6bit					3.72x	5.65*	13.60*	5.69*	13.18†		
3A	12-2560-8bit					3.34x	6.08*	14.98*	6.13*	14.48*		
4A	12-2560-6bit					4.00x	6.08*	15.21*	6.17*	14.70*		
5A	11-3072-8bit					3.27x	5.71*	13.66†	5.78*	13.20†		
6A	11-3072-6bit					3.93x	5.77*	13.86†	5.86*	13.46*		
7A	11-2560-8bit					3.53x	6.18*	15.21†	6.24*	14.68*		
8A	11-2560-6bit					4.22x	6.23*	15.43*	6.35*	14.99*		
1K	12-3072-8bit	✓	✓	share 94m	23.1hrs	3.08x	5.62*	13.33*	5.60*	13.12†		
2K	12-3072-6bit					3.72x	5.65*	13.50*	5.67*	13.16†		
3K	12-2560-8bit					3.34x	6.03*	14.86*	6.07*	14.44*		
4K	12-2560-6bit					4.00x	6.08*	14.95*	6.09*	14.54*		
5K	11-3072-8bit					3.27x	5.74*	13.57*	5.78*	13.35†		
6K	11-3072-6bit					3.93x	5.77*	13.71*	5.81*	13.42†		
7K	11-2560-8bit					3.53x	6.14*	15.06*	6.22*	14.73*		
8K	11-2560-6bit					4.22x	6.23*	15.15*	6.26*	14.75*		

5. Conclusion

We propose a novel one-pass multiple ASR systems joint compression and quantization approach using an all-in-one neural model, which enables the simultaneous construction of multiple systems with varying Encoder depths, widths, and quantization precision settings without the need to train and store individual target systems separately. Extensive results demonstrate that all-in-one models produce comparable or even better performance while costing less training time and memory storage against individually trained systems. Future work focuses on building more flexibly layer-wise re-configurable systems.

6. Acknowledgements

This research is supported by Hong Kong RGC GRF grant No. 14200220, 14200021, Innovation Technology Fund grant No. ITS/218/21.

7. References

- [1] A. Vaswani, N. Shazeer, N. Parmar *et al.*, “Attention is all you need,” in *NIPS*, 2017.
- [2] A. Gulati, J. Qin, C.-C. Chiu *et al.*, “Conformer: Convolution-augmented transformer for speech recognition,” in *Interspeech*, 2020.
- [3] Y. Peng, S. Dalmia, I. Lane *et al.*, “Branchformer: Parallel mlp-attention architectures to capture local and global context for speech recognition and understanding,” in *ICML*, 2022.
- [4] A. Baevski, Y. Zhou, A. Mohamed *et al.*, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *NeurIPS*, 2020.
- [5] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai *et al.*, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM T-ASLP*, vol. 29, pp. 3451–3460, 2021.
- [6] A. Baevski, W.-N. Hsu, Q. Xu *et al.*, “Data2vec: A general framework for self-supervised learning in speech, vision and language,” in *ICML*, 2022.
- [7] S. Chen, C. Wang, Z. Chen *et al.*, “WavLM: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE J-STSP*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [8] Z. Wu, D. Zhao, Q. Liang *et al.*, “Dynamic sparsity neural networks for automatic speech recognition,” in *ICASSP*, 2021.
- [9] J. Lee, J. Kang, and S. Watanabe, “Layer pruning on demand with intermediate ctc,” in *Interspeech*, 2021.
- [10] H. Jiang, L. L. Zhang, Y. Li *et al.*, “Accurate and structured pruning for efficient automatic speech recognition,” in *Interspeech*, 2023.
- [11] S. Li, M. Xu, and X.-L. Zhang, “Efficient conformer-based speech recognition with linear attention,” in *APSIPA ASC*, 2021.
- [12] Z. Li, T. Wang, J. Deng *et al.*, “Lossless 4-bit quantization of architecture compressed conformer asr systems on the 300-hr switchboard corpus,” in *Interspeech*, 2023.
- [13] J. Rathod, N. Dawlatabad, S. Singh *et al.*, “Multi-stage progressive compression of conformer transducer for on-device speech recognition,” in *Interspeech*, 2022.
- [14] J. Park, S. Jin, J. Park *et al.*, “Conformer-based on-device streaming speech recognition with KD compression and two-pass architecture,” in *IEEE SLT*, 2023.
- [15] S. Ding, P. Meadowlark, Y. He *et al.*, “4-bit conformer with native quantization aware training for speech recognition,” in *Interspeech*, 2022.
- [16] S. Kim, A. Gholami, Z. Yao *et al.*, “Integer-only zero-shot quantization for efficient speech recognition,” in *ICASSP*, 2022.
- [17] O. Rybakov, P. Meadowlark, S. Ding *et al.*, “2-bit conformer quantization for automatic speech recognition,” in *Interspeech*, 2023.
- [18] D. Qiu, S. Ding, and Y. He, “The role of feature correlation on quantized neural networks,” in *IEEE ASRU*, 2023.
- [19] J. Xu, S. Hu, J. Yu *et al.*, “Mixed precision quantization of transformer language models for speech recognition,” in *ICASSP*, 2021.
- [20] C.-I. J. Lai, Y. Zhang, A. H. Liu *et al.*, “PARP: Prune, adjust and re-prune for self-supervised speech recognition,” in *NeurIPS*, 2021.
- [21] Y. Peng, K. Kim, F. Wu *et al.*, “Structured pruning of self-supervised pre-trained models for speech recognition and understanding,” in *ICASSP*, 2023.
- [22] V. S. Lodagala, S. Ghosh, and S. Umesh, “PADA: Pruning assisted domain adaptation for self-supervised speech representations,” in *IEEE SLT*, 2023.
- [23] H.-J. Chang, S.-w. Yang, and H.-y. Lee, “DistilHuBERT: Speech representation learning by layer-wise distillation of hidden-unit bert,” in *ICASSP*, 2022.
- [24] Y. Lee, K. Jang, J. Goo *et al.*, “FitHuBERT: Going thinner and deeper for knowledge distillation of speech self-supervised learning,” in *Interspeech*, 2022.
- [25] T. Ashihara, T. Moriya, K. Matsuura *et al.*, “Deep versus wide: An analysis of student architectures for task-agnostic knowledge distillation of self-supervised speech models,” in *Interspeech*, 2022.
- [26] Y. Fu, Y. Kang, S. Cao *et al.*, “DistillW2V2: A small and streaming wav2vec 2.0 based asr model,” *arXiv preprint arXiv:2303.09278*, 2023.
- [27] H.-J. Chang, N. Dong, R. Mavlyutov *et al.*, “CoLLD: Contrastive layer-to-layer distillation for compressing multilingual pre-trained speech encoders,” in *ICASSP*, 2024.
- [28] C. J. Cho, A. Mohamed, S.-W. Li *et al.*, “SD-HuBERT: Self-distillation induces syllabic organization in HuBERT,” in *ICASSP*, 2024.
- [29] D. de Oliveira and T. Gerkmann, “Distilling HuBERT with LSTMs via decoupled knowledge distillation,” in *ICASSP*, 2024.
- [30] R. Wang, Q. Bai, J. Ao *et al.*, “LightHuBERT: Lightweight and configurable speech representation learning with once-for-all hidden-unit bert,” in *Interspeech*, 2022.
- [31] Y. Peng, Y. Sudo, S. Muhammad *et al.*, “DPHuBERT: Joint distillation and pruning of self-supervised speech models,” in *Interspeech*, 2023.
- [32] N. Wang, C.-C. C. Liu, S. Venkataramani *et al.*, “Deep compression of pre-trained transformer models,” in *NeurIPS*, 2022.
- [33] E. Fish, U. Michieli, and M. Ozay, “A model for every user and budget: Label-free and personalized mixed-precision quantization,” in *Interspeech*, 2023.
- [34] S. Ding, Q. David, D. Rim *et al.*, “USM-Lite: Quantization and sparsity aware fine-tuning for speech recognition with universal speech models,” in *ICASSP*, 2024.
- [35] H. Cai, C. Gan, T. Wang *et al.*, “Once for all: Train one network and specialize it for efficient deployment,” in *ICLR*, 2020.
- [36] J. Yu and T. S. Huang, “Universally slimmable networks and improved training techniques,” in *ICCV*, 2019.
- [37] L. Hou, Z. Huang, L. Shang *et al.*, “DynaBERT: Dynamic bert with adaptive width and depth,” in *NeurIPS*, 2020.
- [38] A. Radford, J. W. Kim, T. Xu *et al.*, “Robust speech recognition via large-scale weak supervision,” in *ICML*, 2023.
- [39] S. Watanabe, T. Hori, S. Kim *et al.*, “Hybrid CTC/attention architecture for end-to-end speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [40] H.-J. Chen, Y. Meng, and H.-y. Lee, “Once-for-all sequence compression for self-supervised speech models,” in *ICASSP*, 2023.
- [41] J. Yu, P. Jin, H. Liu *et al.*, “BigNAS: Scaling up neural architecture search with big single-stage models,” in *ECCV*, 2020.
- [42] S. Hu, R. Wang, L. Hong *et al.*, “Generalizing few-shot NAS with gradient matching,” in *ICLR*, 2022.
- [43] S. Watanabe, T. Hori, S. Karita *et al.*, “ESPnet: End-to-end speech processing toolkit,” *arXiv preprint arXiv:1804.00015*, 2018.
- [44] J. J. Godfrey, E. C. Holliman, and J. McDaniel, “SWITCHBOARD: Telephone speech corpus for research and development,” in *ICASSP*, 1992.
- [45] V. Panayotov, G. Chen, D. Povey *et al.*, “LibriSpeech: an asr corpus based on public domain audio books,” in *ICASSP*, 2015.
- [46] M. Nagel, M. Fournarakis, R. A. Amjad *et al.*, “A white paper on neural network quantization,” *arXiv preprint arXiv:2106.08295*, 2021.
- [47] L. Gillick and S. J. Cox, “Some statistical issues in the comparison of speech recognition algorithms,” in *ICASSP*, 1989.