PREDICTING MULTI-CODEBOOK VECTOR QUANTIZATION INDEXES FOR KNOWLEDGE DISTILLATION

Liyong Guo[*,1], *Xiaoyu Yang*[*,1], *Quandong Wang*[1], *Yuxiang Kong*[1], *Zengwei Yao*[1], *Fan Cui*[1],
Fangjun Kuang[1], *Wei Kang*[1], *Long Lin*[1], *Mingshuang Luo*[1], *Piotr Żelasko*[2], *Daniel Povey*[1]

[1] Xiaomi Corp., Beijing, China   [2]Meaning.Team Inc, USA

{guoliyong, xiaoyuyang6, dpovey}@xiaomi.com, pzelasko@meaning.team

ABSTRACT

Knowledge distillation (KD) is a common approach to improve model performance in automatic speech recognition (ASR), where a student model is trained to imitate the output behaviour of a teacher model. However, traditional KD methods suffer from teacher label storage issue, especially when the training corpora are large. Although on-the-fly teacher label generation tackles this issue, the training speed is significantly slower as the teacher model has to be evaluated every batch. In this paper, we reformulate the generation of teacher label as a codec problem. We propose a novel Multi-codebook Vector Quantization (MVQ) approach that compresses teacher embeddings to codebook indexes (CI). Based on this, a KD training framework (MVQ-KD) is proposed where a student model predicts the CI generated from the embeddings of a self-supervised pre-trained teacher model. Experiments on the LibriSpeech clean-100 hour show that MVQ-KD framework achieves comparable performance as traditional KD methods (l1, l2), while requiring 256 times less storage. When the full LibriSpeech dataset is used, MVQ-KD framework results in 13.8% and 8.2% relative word error rate reductions (WERRs) for non-streaming transducer on test-clean and test-other and 4.0% and 4.9% for streaming transducer. The implementation of this work is already released as a part of the open-source project icefall[1].

Index Terms— knowledge distillation, neural transducer, ASR

1. INTRODUCTION

In the field of speech processing, significant improvements have been witnessed in self-supervised pre-training in recent years [1–4]. After pre-training on a very large amount of unlabeled data, the model is then fine-tuned with task-specific labeled data for downstream tasks such as ASR[1, 3, 4], speaker verification[4, 5], emotion recognition[5, 6], etc.

To fully leverage the richness of unlabeled data, pre-trained models [1, 3, 4, 7, 8] usually have a large number of parameters, ranging from hundreds of millions to several billions.. Although these models achieves state-of-the-art performance, they are impractical to be used in real-life scenarios due to their large model size and footprint. To deal with this, efforts have been made to utilize pre-trained models for improving smaller model's performance. Knowledge distillation (KD)[9], also known as teacher-student training [10–13] is applied to transfer information from a pre-trained teacher model to a typically-smaller student model, where the student model

learns from labels generated from the teacher. Although the student model is typically of small size, there is an implicit problem being unnoticed: training efficiency. In traditional teacher-student training, the teacher labels are often float embeddings [11, 14, 15] extracted on-the-fly, which would slow down the training if the teacher model is an extremely large pre-trained model. In addition to that, the maximum training batch size has to be reduced, leading to potential performance degradation. Otherwise, one could save the float type embeddings to disk before training and load them during KD training. However, the training speed would be constrained by the I/O and a huge amount of storage space is needed, making training impractical if the training corpus is large.

Clustering or quantization is effective for representation learning [1, 3, 16, 17]. Wav2vec2.0 [1] takes vector quantization (VQ) for clustering and the codebook vector is taken into computing contrastive loss. BEST-RQ [18] takes the VQ indexes as the pre-training labels. Inspired by these and borrowing ideas of residual vector quantization [19] and the direct-sum codebooks [20, 21], we propose a trainable **M**ulti-codebook **V**ector **Q**uantization (MVQ) which compresses each embedding vector into a short sequence of 8-bit integer codebook indexes (CI)[2].

Based on MVQ, a KD framework (MVQ-KD) is proposed by teaching a student model to predict CI generated from the embeddings at an intermediate layer of a teacher model. This could solve the computation or storage issue of traditional KD methods[22, 23]. For example, with the Hubert-Large model [3] whose dimension is 1024, it would cost 1976 gigabytes for storing 960 hours' of float-type teacher embeddings if 3-fold speed perturbation is used. However, only 7.72 gigabytes are needed for the corresponding CI in a 16-codebook MVQ setup, achieving a compression rate of 256. CI can be pre-computed and stored on disk at very low cost, which improves the training efficiency.

The key experimental findings of this paper are:

- MVQ-KD achieves comparable performance as using traditional l_1 or l_2 losses, while saving 256 times storage space, avoiding the need of on-the-fly teacher label generation.

- The performance of MVQ-KD can be further improved with more codebooks.

- MVQ-KD is effective both for streaming and non-streaming transducer models.

In the rest of this paper, Sec. 2 illustrates the details of the MVQ algorithm. Sec. 3 briefly reviews the self-supervised pre-trained Hu-BERT model and presents the MVQ-KD framework. In Sec. 4, the experimental setup and results are described. Finally, conclusions are drawn in Sec. 5.

---

* stands for equal contribution
[1]https://github.com/k2-fsa/icefall

[2]https://github.com/k2-fsa/multi_quantization

2. TRAINABLE MULTI-CODEBOOK QUANTIZER

Consider a quantization module $\mathcal{Q}$ encoding vectors $\boldsymbol{x} \in \Re^D$ from a known distribution, into a fixed-size sequence of N integers $0 \leq i_n < K$: let $\boldsymbol{i} = i_0, \ldots, i_{N-1}, \boldsymbol{i} \in \{0, ..., K-1\}^N$. $\mathcal{Q}$ should have

$$\boldsymbol{i} = \text{Encode}(\boldsymbol{x}), \tag{1}$$

$$\hat{\boldsymbol{x}} = \text{Decode}(\boldsymbol{i}). \tag{2}$$

We are interested in encodings that are as close as possible to optimal in an l_2^2-error sense, i.e. that minimizes $E[||\hat{\boldsymbol{x}} - \boldsymbol{x}||_2^2]$. To keep the encoding scheme practical, we consider *direct-sum* codebooks, i.e. a scheme where the $\text{Decode}(\cdot)$ function sums over the codebooks:

$$\text{Decode}(\boldsymbol{i}) = \sum_{n=0}^{N-1} \boldsymbol{c}_{i_n}^{(n)}. \tag{3}$$

This requires $K \times N$ codebook centers $c_k^{(n)} \in \Re^D$. The length of the integer sequence N can be referred to as the number of codebooks.

When encoding using direct-sum codebooks, it is impractical to enumerate all possible encodings as the encoding space is $\mathcal{O}(K^N)$. One straightforward way is to choose the codebook center sequentially that reduces the residual error the most while keeping other codebook centers unchanged. This heuristic is not guaranteed to yield the lowest residual error as the codebook centers from different codebooks are not jointly evaluated. Therefore, an iterative encoding scheme is proposed to improve the aforementioned heuristic, which efficiently searches for better encodings.

2.1. Iterative encoding scheme

Assuming the number of codebooks to be N, $\text{Encode}(\cdot)$ compresses a float vector x to N CI. The encoding function has N independent neural classifiers, which generate an intial estimate of CI. The encoding process $\text{Encode}(\cdot)$ is implemented as follows :

- Choose the initial codebook entries $\boldsymbol{i}$ as the arg-max of N independent logistic regression classifiers.

- For e.g. 3 iterations, refine the codebook entries: $\boldsymbol{i} \leftarrow \text{Refine}(\boldsymbol{x}, \boldsymbol{i})$, each time using the refined index generated from the previous call of $\text{Refine}(\cdot, \cdot)$.

- Return the refined indexes. They will be used as input for $\text{Decode}()$ and as label to train the classifiers (see Sec. 2.2)

The function $\text{Refine}(\boldsymbol{x}, \boldsymbol{i})$ is the most essential part in $\text{Encode}(\cdot)$. It is given initial codebook indexes $\boldsymbol{i} = i_n, 0 \leq n < N$, and returns possibly-improved codebook entries $\hat{\boldsymbol{i}} = \hat{i}_n, 0 \leq n < N$:

- For each codebook $0 \leq n < N$ and each index $0 \leq k < K$, compute modified residual $\hat{\boldsymbol{x}} - \boldsymbol{x}$ assuming we let the n'th codebook center be k but leaving all the other codebooks with their initial values (the ones in $\boldsymbol{i}$).

- For each codebook n, sort the residuals above and store indexes of the J smallest residuals.

- Construct a sub-problem that has $N/2$ codebooks, with each codebook being of size J^2, by summing pairs of J-best codebook centers, combining $n = 0$ with $n = 1$, $n = 2$ with $n = 3$, and so on.

- Recurse, call $\text{Refine}(\cdot, \cdot)$ on the smaller sub-problem.

- The result $\hat{\boldsymbol{i}}$ can be computed from the answer to the sub-problem and the indexes of the J-best entries for each codebook.

The recursion terminates when the sub-problem only has one codebook and the index resulting in the lowest residual is selected. The refined index $\hat{\boldsymbol{i}}$ can be returned recursively.

2.2. Training and Inference Procedure

The trainable parameters in the quantization module are codebook centers $c_k^{(n)}$ and N logistic-regression classifiers $\mathcal{C}_n$. For each float vector $\boldsymbol{x}$ and its encodings $\text{Encoder}(\boldsymbol{x}) = \boldsymbol{i}$, the training loss $\mathcal{L}$ consists of two parts:

$$\mathcal{L} = \mathcal{L}_{\text{residual}} + \mathcal{L}_{\text{prediction}}, \tag{4}$$

$$= ||\boldsymbol{x} - \text{Decode}(\boldsymbol{i})||_2^2 + \sum_{n=1}^{N} -\log \mathcal{C}_n(\boldsymbol{x})_{i_n} \tag{5}$$

where $\mathcal{C}_n(\boldsymbol{x})_{i_n}$ is the probability of predicting i_n in $\mathcal{C}_n$. The first term $\mathcal{L}_{\text{residual}}$ is the reconstruction loss, i.e l_2^2 residual and optimizes $\mathcal{C}_n$. The second term $\mathcal{L}_{\text{prediction}}$ is the prediction loss and encourages the neural classifiers to select the encoded indexes $\boldsymbol{i}$ improved by $\text{Refine}(\cdot, \cdot)$. By doing so, the initial estimate given by $\mathcal{C}_n$ is expected to be close to the refined CI. Training is performed based on gradient descent with Adam [24] optimizer. During inference, the encoding process described in Sec. 2.1 is repeated to generate CI for test data.

2.3. Analysis of reconstruction loss

Rate-distortion theory can be used to evaluate the reconstruction performance of the proposed quantization algorithm. For quantizers with various numbers of codebooks trained on HuBERT-L[3] with embedding dimension of 1024, the first column of Table 1 shows the relative reconstruction loss (RRL), defined as the mean of the squared reconstruction error $||\hat{\boldsymbol{x}} - \boldsymbol{x}||_2^2$ divided by the mean of $||\boldsymbol{x} - \mu_x||_2^2$, i.e. the average sum-square of $\boldsymbol{x}$ after mean normalization. The last column shows the best possible distortion assuming the 1024 dimensions were normally and independently distributed. This comes from the rate-distortion equation for a memoryless Gaussian source [25]: $R(D) = \frac{1}{2} \log_2(\sigma_x^2/D)$, with $\sigma_x^2 = 1$, the bit-rate per dimension R set to $\frac{8N}{1024}$ since there are N codebooks of 8 bits each, and solving for distortion D. Although this is a lower bound on the distortion, the values are actually higher than the second column as the HuBERT-L embeddings are not strictly independent and Gaussian.

When applied to features that are normally distributed and independent, the algorithm achieves RRL (see second column in Table 1) that is within 10% of the Shannon lower bound, so the performance in terms of reconstruction loss does not have that much potential for further improvement as far as our purposes are concerned.

Table 1: Relative Reconstruction Loss (RRL) and Shannon distortion if embedding dimensions were i.i.d. normal

N	RRL(HuBERT)	RRL(Gaussian)	Shannon distortion
1	0.517	0.992	0.989
4	0.356	0.969	0.958
8	0.278	0.938	0.917
16	0.225	0.876	0.841
32	0.206	0.760	0.707

3. PROPOSED DISTILLATION FRAMEWORK

3.1. Self-supervised pre-trained HuBERT

Recently, self-supervised pre-training has shown promising results [1, 3] in ASR. Among these methods, HuBERT[3] is one of the most effective frameworks. HuBERT model comprises of three parts: convolutional neural network (CNN) encoder, transformer and acoustic unit discovery system. The CNN encoder processes

raw speech waveform w and generates embedding $X = x_{1:T}$. The acoustic unit discovery system then produces the hidden unit target z_t for each x_t using k-means clustering. Before feeding embedding X to the transformer to generate contextualized representations, a set of randomly selected timestamps are masked. The self-supervised pre-training objective is to predict the correct hidden unit z_t for both masked and unmasked timestamps with $\mathcal{L} = \alpha\mathcal{L}_m + (1 - \alpha)\mathcal{L}_u$, where $\mathcal{L}_m$ and $\mathcal{L}_u$ is the CrossEntropy loss for masked and unmasked timestamps and α is a tunable coefficient. During training, z_t is refined to improve the clustering quality. After pre-training, HuBERT can be fine-tuned with labeled speech for ASR tasks[3, 5].

3.2. Traditional KD Methods for Neural Transducers

Neural transducers is a powerful modelling framwork E2E ASR. It has gained more popularity recently due to its natural support for streaming and superior performance. To further improve the performance of neural transducer, knowledge distillation (KD), or teacher-student training, is common used. During KD, a student model is trained to imitate the output of a teacher model. Depending on the teacher's output, different loss functions can be applied for KD training. Kullback-Leibler (KL) divergence is commonly used if teacher labels are distributions whereas l_1 or l_2 are more appropriate for continuous feature. As a transducer generates a 3-D distribution lattice, directly applying KL-divergence is computational intractable. [26] used a collapsed version of the distribution lattice to reduce computation, whereas [22] approximated the distribution lattice with its one-best alignment. Both methods[22, 26] pre-computed the teacher labels and stored them to disk, which could be problematic for large training corpora. Instead of utilizing the output distribution, embedding features are another straightforward teacher label. [23] uses the l_2 loss between the encoder embeddings of teacher and student model for KD. Let teacher embedding $\mathbf{TE}^{l_{th}} = \mathbf{TE}_1^{l_{th}}, ..., \mathbf{TE}_T^{l_{th}}$ be the embedding extracted from the l_{th}-th layer of the teacher model and $\mathrm{SE}^{l_{st}} = \mathbf{SE}_1^{l_{st}}, ..., \mathbf{SE}_T^{l_{st}}$ be the embedding at l_{st}-th layer of the student model, the KD loss function is:

$$\mathcal{L}_{embedding} = \sum_{t=1}^{T} \mathbf{Dist}(\mathrm{TE}_t^{l_{th}}, \mathbf{LossNet}(\mathrm{SE}_t^{l_{st}})), \quad (6)$$

where $\mathbf{Dist}$ is any function (e.g l_1, l_2) that measures the distance between two vectors and $\mathbf{LossNet}$ is usually a linear layer that maps $\mathbf{SE}^{l_{st}}$ to the same dimension as $\mathbf{TE}^{l_{th}}$. $\mathbf{TE}^{l_{th}}$ are generated on-the-fly since teacher and student are jointly trained[23]. However, this will inevitably affect batch size or training speech, which could affect the performance of the student model.

3.3. MVQ-based KD for Neural Transducers

To alleviate the aforementioned issues, we propose to apply MVQ on the embedding extracted from an intermediate layer of teacher $\mathbf{TE}_t^{l_{th}}$ and compress it to CI. Then, instead of regressing $\mathbf{TE}_t^{l_{th}}$ using l_1 or l_2 loss, the student model is trained to predict its CI. Let N be the number of codebooks, MVQ compresses $\mathbf{TE}_t^{l_{th}}$ to $\mathbf{i}_t = (i_{t,1}, ..., i_{t,N})$, with $i_{t,n}$ representing which entry in the n-th codebook is chosen at t-th frame, the loss function is:

$$\mathcal{L}_{cb} = \sum_{t=1}^{T} \sum_{n=1}^{N} \mathbf{CrossEntropy}(\mathbf{i}_t, \mathbf{LossNet}(\mathrm{SE}_t^{l_{st}})), \quad (7)$$

where $\mathbf{LossNet}$ is a module consisting of a linear layer with softmax activation that transforms the student embedding to probabilities. At each timestamp t, MVQ-KD performs N independent classification. If the number of codebook centers in each codebook is 256, each en-
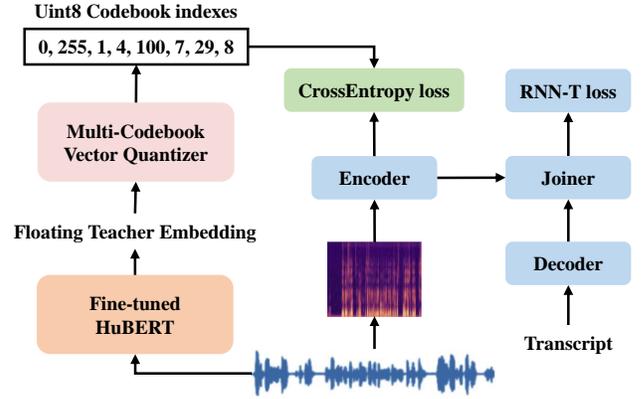
Fig. 1: MVQ based teacher-student learning.

try in $\mathbf{i}_t$ can be represented by an 8-bit integer. Therefore, they can be pre-computed and store on disk at very low cost. With $N=16$ and teacher embeddings of 1024 dimensional, MVQ achieves a compression ratio of 256, significantly increasing KD training's scalability compared to traditional KD methods. Fig. 1 illustrated the proposed KD framework for MVQ-based KD.

Different from non-streaming transducer models, streaming transducer model only has limited access to future context and tends to emit symbols later. Therefore, applying Eqn. (7) directly on a streaming transducer could be problematic as it may force the student to guess into the future. Inspired by[22], a time-shift variable δ is introduced to address the temporal mismatch between teacher and student model. This leads to a modified version of $\mathcal{L}_{cb}$:

$$\mathcal{L}_{cb} = \sum_{t=1}^{T-\delta} \sum_{n=1}^{N} \mathbf{CrossEntropy}(\mathbf{i}_t, \mathbf{LossNet}(\mathrm{SE}_{t+\delta}^{l_{st}})), \quad (8)$$

The codebook loss $\mathcal{L}_{cb}$ will be used as an auxiliary loss to the original transducer loss:

$$\mathcal{L}_{total} = \mathcal{L}_{transducer} + \lambda\mathcal{L}_{cb}, \quad (9)$$

where λ is a tunable scale of the auxiliary loss.

4. EXPERIMENTS

4.1. Datasets and Model

The LibriSpeech ASR corpus [27] was used for all experiments. The full dataset contains 960h hours of transcribed audio. Among these, the "train clean 10" subset was used for comparison with other baseline models and hyper-parameter tuning. During training, SpecAug[28] and speed perturbation with rate 0.9 and 1.1 are used for data-augmentation. MUSAN[29] is used for noise-augmentation. The output vocabulary has 500 subword units and WERs are reported on test-clean and test-other sets using beam search.

The large version of HuBERT[3] is adopted to initialise the encoder of the teacher transducer model and finetuned on full LibriSpeech. The student model is also a transducer model with a reworked version of Conformer[3][30] as encoder. In streaming experiments, we apply causal convolutions and blockwise-limited right context in attention and train the student model with dynamic chunk size [31]. Pruned RNNT loss[32] is used for computing $\mathcal{L}_{transducer}$.

---

[3]https://github.com/k2-fsa/icefall/tree/master/egs/librispeech/ASR

Details of teacher and student model are listed in Table 2. The embeddings of the 18-th transformer block are extracted for CI generation as we believe this layer contains rich information while being less difficult to learn from. Consequently, KD is carried out on the 9-th layer of the student model to share the same relative position as in the teacher model. A randomly sampled subset of 1000 audios are used to generate embeddings and train the quantizer $\mathcal{Q}$. Then, CI are generated by feeding the whole training set to $\mathcal{Q}$ and stored to disk. To compare with l_1 and l_2 losses, the 18-th layer's float embeddings are also stored.

Table 2: Details of teacher and student models

	Teacher model	Student model
Encoder	HuBERT-Large	Conformer
Encoder dim	1024	512
Encoder layer	24	12
Num Params	318M	78M

4.2. Impact of the Number of Codebooks

Table 3 demonstrates the impact of codebook numbers (N) on the MVQ distillation, from 1 to 32. With $N=1$, the quantization degenerates into the traditional VQ. For different N, the scale of $\mathcal{L}_{cb}$ is tuned individually for optimal performance. Using only one codebook, the student model already outperforms the baseline model. The student model consistently improves as N increases. This also accords with the fact that MVQ achieves lower reconstruction error if more codebooks are used (see Table 1). As $N=16$ achieves similar WERs compared to $N=32$ while doubling the compression rate, $N=16$ is selected in future experiments.

Table 3: WERs with different number of codebooks

N	compression-rate	test-clean	test-other
baseline, 0	-	6.83	18.19
1	4096	5.67	15.77
2	2048	5.58	15.27
4	1024	5.39	14.68
8	512	5.14	14.51
16	256	5.01	13.80
32	128	**4.99**	**13.68**

4.3. Teacher-student learning with different losses

Table 4 compares the WERs of student model trained using MVQ KD with other traditional l_1 and l_2 losses. To ensure the fairness of comparison, on-the-fly teacher label generation is not adopted for l_1 and l_2 as this will limit the batch size. Experiments are only carried out on clean-100 subset as storing teacher labels of 960h audio for l_1 and l_2 loss computation is impractical. The scales of all auxiliary losses are tuned individually and only the setup with the lowest WERs are reported. The following observations can be made from Table 4. First, the proposed KD framework successfully improves the performance of the student model. Both MVQ KD and traditional l_1 and l_2 method are able to reduce the WERs of the student model, indicating the effectiveness of using an intermediate layer for KD. Second, although KD with l_2 loss results in the lowest WERs, MVQ still achieves comparable performance, while being able to be flexibly applied in larger scale of experiment. During the experiments, it is found that the embedding values of HuBERT model are unstable, sometimes ranging from -2000 to +3000. Applying l_1 and

l_2 loss requires special design such as clamping the embedding values, whereas MVQ-KD is less sensitive to this.

Table 4: WER for baseline and distillation with different losses

config	test-clean	test-other
baseline	6.83	18.19
l_1	5.1	13.69
l_2^2	4.99	13.39
MVQ, $N=16$	5.01	13.80
MVQ, $N=32$	4.99	13.68

4.4. Training with full LibriSpeech

To further demonstrate the effectiveness and robustness of MVQ, experiments are scaled up to the full LibriSpeech for both non-streaming and streaming student transducer models. For non-streaming models, relative WER reductions (WERRs) of 13.8% and 8.2% are achieved on test-clean and test-other. For streaming models, both WERs and latency are shown for different δ. The latency is measured against the word-level alignment obtained from a hidden Markov model using [33]. To get a reasonable estimate for δ, the locations of posterior peaks in the lattice of teacher and student models are compared. The following three key observations can be made. First, MVQ-KD improves the accuracy of streaming model if a sensible δ is selected. The model trained with $\delta = 0$ has higher WERs than the baseline model, while the model trained with larger δ outperforms the baseline model, achieving WERRs of 4.0% and 4.9% with $\delta = 5$. Second, MVQ-KD reduces the latency of streaming models. Setting δ to 4 or 5 not only improves the model accuracy, but also encourages the model to emit faster, achieving a latency reduction of 0.1 seconds compared to the baseline model. Third, as δ increases, the latency also increases while the WER decrease, suggesting that δ controls the trade-off between model latency and model accuracy.

Table 5: WER of models trained with full LibriSpeech

	test-clean	test-other	latency (s)
Reference models			
Teacher, HuBERT-L	1.9	3.94	-
Baseline, non-streaming	2.69	6.11	-
Baseline, streaming	3.03	7.98	0.335
MVQ-KD trained model			
Non-streaming	**2.32**	**5.61**	-
Streaming, $\delta = 0$	3.13	7.9	0.165
Streaming, $\delta = 4$	2.99	7.64	**0.235**
Streaming, $\delta = 5$	**2.91**	**7.59**	0.259

5. CONCLUSIONS

In this paper, we present an efficient and effective knowledge distillation (KD) framework for neural transducers based on a novel Multi-codebook Vector Quantization (MVQ) algorithm. With a fine-tuned self-supervised pre-trained model, we show that our framework achieves comparable performance as the traditional l_1 and l_2 losses, while being much faster or requiring hundreds of times less storage. We also demonstrate that the proposed KD framework is effective both for non-streaming and streaming student model. In future works, we would like to incorporate multiple teacher layers for KD to further improve the student model. Since MVQ is a general quantization algorithm, we would also like to explore the feasibility of applying MVQ-KD on other speech processing tasks.

6. REFERENCES

[1] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS*, Vancouver, 2020.

[2] Y. Zhang, James Qin, Daniel S. Park, Wei Han, C. Chiu, et al., "Pushing the limits of semi-supervised learning for automatic speech recognition," in *NeurIPS SAS Workshop*, Vancouver, 2020.

[3] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, et al., "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2021.

[4] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al., "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, 2022.

[5] Yingzhi Wang, Abdelmoumene Boumadane, and Abdelwahab Heba, "A fine-tuned wav2vec 2.0/hubert benchmark for speech emotion recognition, speaker verification and spoken language understanding," *arXiv preprint arXiv:2111.02735*, 2021.

[6] Leonardo Pepino, Pablo Riera, and Luciana Ferrer, "Emotion recognition from speech using wav2vec 2.0 embeddings," in *Interspeech*, Brno, 2021.

[7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *NAACL*, Minneapolis, 2019.

[8] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli, "wav2vec: Unsupervised pre-training for speech recognition," in *Interspeech*, Graz, 2019.

[9] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean, "Distilling the knowledge in a neural network," in *NIPS Deep Learning Workshop*, Montreal, 2014.

[10] Zhong Meng, Jinyu Li, Yashesh Gaur, and Yifan Gong, "Domain adaptation via teacher-student learning for end-to-end speech recognition," in *ASRU*, Sentosa, 2019.

[11] Vimal Manohar, Pegah Ghahremani, Daniel Povey, and Sanjeev Khudanpur, "A teacher-student learning approach for unsupervised domain adaptation of sequence-trained asr models," in *SLT*, Athens, 2018.

[12] Gakuto Kurata and George Saon, "Knowledge distillation from offline to streaming rnn transducer for end-to-end speech recognition.," in *Interspeech*, Shanghai, 2020.

[13] Thibault Doutre, Wei Han, Min Ma, Zhiyun Lu, Chung-Cheng Chiu, et al., "Improving streaming automatic speech recognition with non-streaming model distillation on unsupervised data," in *ICASSP*, Toronto, 2021.

[14] Yoshua Bengio, "Deep learning of representations for unsupervised and transfer learning," in *Proceedings of ICML workshop on unsupervised and transfer learning*, 2012.

[15] Naoyuki Kanda, Yusuke Fujita, and Kenji Nagamatsu, "Sequence distillation for purely sequence trained acoustic models," in *ICASSP*, Calgary, 2018.

[16] Alexei Baevski, Steffen Schneider, and Michael Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," *arXiv preprint arXiv:1910.05453*, 2019.

[17] Hangbo Bao, Li Dong, and Furu Wei, "Beit: Bert pre-training of image transformers," *arXiv preprint arXiv:2106.08254*, 2021.

[18] Chung-Cheng Chiu, James Qin, Yu Zhang, Jiahui Yu, and Yonghui Wu, "Self-supervised learning with random-projection quantizer for speech recognition," *arXiv preprint arXiv:2202.01855*, 2022.

[19] Christopher F Barnes, Syed A Rizvi, and Nasser M Nasrabadi, "Advances in residual vector quantization: A review," *IEEE transactions on image processing*, vol. 5, 1996.

[20] Christopher F Barnes and Richard L Frost, "Vector quantizers with direct sum codebooks," *IEEE Trans. on information theory*, vol. 39, 1993.

[21] Christopher F Barnes and John P Watkins, "Embedded wavelet zerotree coding with direct sum quantization structures," in *Proceedings DCC'95 Data Compression Conference*, Snowbird, 1995.

[22] Xiaoyu Yang, Qiujia Li, and Philip C Woodland, "Knowledge distillation for neural transducers from large self-supervised pre-trained models," in *ICASSP*, Singapore, 2022.

[23] Rupak Vignesh Swaminathan, Brian King, et al., "Codert: Distilling encoder representations with co-learning for transducer-based speech recognition," in *Interspeech*, Brno, 2021.

[24] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *Computer Science*, 2014.

[25] MTCAJ Thomas and A Thomas Joy, *Elements of information theory*, Wiley-Interscience, 2006.

[26] Sankaran Panchapagesan, Daniel S Park, Chung-Cheng Chiu, Yuan Shangguan, Qiao Liang, and Alexander Gruenstein, "Efficient knowledge distillation for rnn-transducer models," in *ICASSP*, Toronto, 2021.

[27] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *ICASSP*, Brisbane, 2015.

[28] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, et al., "Specaugment: A simple data augmentation method for automatic speech recognition," in *Interspeech*, Graz, 2019.

[29] David Snyder, Guoguo Chen, and Daniel Povey, "MUSAN: A Music, Speech, and Noise Corpus," 2015, arXiv:1510.08484v1.

[30] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, et al., "Conformer: Convolution-augmented transformer for speech recognition," in *Interspeech*, Shanghai, 2020.

[31] Binbin Zhang, Di Wu, Zhuoyuan Yao, Xiong Wang, Fan Yu, Chao Yang, Liyong Guo, Yaguang Hu, Lei Xie, and Xin Lei, "Unified streaming and non-streaming two-pass end-to-end model for speech recognition," *arXiv preprint arXiv:2012.05481*, 2020.

[32] Fangjun Kuang, Liyong Guo, Wei Kang, Long Lin, Mingshuang Luo, Zengwei Yao, and Daniel Povey, "Pruned rnn-t for fast, memory-efficient asr training," in *Interspeech*, Incheon, 2022.

[33] Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldi.," in *Interspeech*, Stockholm, 2017.