
Self-Supervised Learning with Random-Projection Quantizer for Speech Recognition

Chung-Cheng Chiu^{*1} James Qin^{*1} Yu Zhang¹ Jiahui Yu¹ Yonghui Wu¹

Abstract

We present a simple and effective self-supervised learning approach for speech recognition. The approach learns a model to predict the masked speech signals, in the form of discrete labels generated with a random-projection quantizer. In particular the quantizer projects speech inputs with a randomly initialized matrix, and does a nearest-neighbor lookup in a randomly-initialized codebook. Neither the matrix nor the codebook is updated during self-supervised learning. Since the random-projection quantizer is not trained and is separated from the speech recognition model, the design makes the approach flexible and is compatible with universal speech recognition architecture. On LibriSpeech our approach achieves similar word-error-rates as previous work using self-supervised learning with non-streaming models, and provides lower word-error-rates and latency than wav2vec 2.0 and w2v-BERT with streaming models. On multilingual tasks the approach also provides significant improvement over wav2vec 2.0 and w2v-BERT.

1. Introduction

Self-supervised learning has shown impressive improvement for the quality of the speech recognition models in recent years (Schneider et al., 2019; Baevski et al., 2020a; 2019; 2020b; Hsu et al., 2021; Zhang et al., 2020; Chung et al., 2021; Zhang et al., 2021). These learning approach enable the model to learn from unsupervised data and combine with supervised learning to improve the recognition accuracy. The capability of learning from unsupervised data is particularly beneficial when the supervised data is limited

^{*}Equal contribution ¹Google Research, Brain Team. Correspondence to: Chung-Cheng Chiu <chungchengc@google.com>, James Qin <jamesqin@google.com>, Yu Zhang <ngyuzh@google.com>.

and opens up new opportunities for low resource languages and domains.

One common design principle of self-supervised learning for speech recognition centers around learning representations. Inspired by the success of BERT (Devlin et al., 2018), one research trend in the speech community is to build BERT-inspired algorithms. One challenge in building BERT-style self-supervised learning for speech is to bridge the gap between continuous speech signals and the discrete text tokens, and a solution for addressing this issue is through learning speech representation (Schneider et al., 2019; Baevski et al., 2020b) or learning quantized representation (Baevski et al., 2020b;a; 2019; Hsu et al., 2021; Chung et al., 2021). Many previous works proposed effective algorithms for learning speech representations, and the quantized result of those learned representations showed encouraging correlation with the phoneme of the utterances.

While representation learning is a critical topic for the speech field, combining it with self-supervised learning leads to two limitations that can slow the research progress: (1) Model architecture limitation. The integration of representation learning and self-supervised learning often requires the model to act the role of providing speech representation while still being effective for the downstream tasks. An effective representation model, however, may not always be effective for the downstream tasks. For example, a good representation learning model may require accessing the future context of the utterance, while downstream tasks may require a low latency model which prohibits the access of the future context. (2) Increased complexity. The objectives of representation learning and self-supervised learning are not always aligned, and the complexity of designing both algorithms and finding their balance can impede the research development. This complexity can also motivate the field toward designing more complicated algorithms instead of finding a simple and effective alternative.

In this work we propose **BERT-based Speech pre-Training with Random-projection Quantizer (BEST-RQ)**, a simple and effective self-supervised learning algorithm for speech recognition. The algorithm masks speech signals and feeds them to the encoder part of the speech recognition model, and the encoder learns to predict the masked region based

on the unmasked speech signals where the learning targets are labels provided by a random-projection quantizer. The random projection quantizer projects speech signals to a randomly initialized matrix, and finds a nearest vector in a randomly initialized codebook. The index of that vector is the target label. Neither the projection matrix nor the codebook is updated throughout the learning process. The quantizer does not require representation learning, and its separation from the model removes the limitation on the architecture design of the model. Despite its simplicity, on LibriSpeech the algorithm achieves similar results as previous work with non-streaming models, and provides better improvement with streaming models compared with previous approaches. On multilingual tasks, the algorithm exhibits further gains compared to wav2vec 2.0 (Baevski et al., 2020b) and w2v-BERT (Chung et al., 2021).

We conduct further analysis on the relation between representation learning quality and the self-supervised learning quality, and demonstrate that the two objectives are not inherently aligned in Section 4.3 and Section 4.4. Such an observation is central to our design of self-supervised learning without representation learning, and opens up a new, less complicated research direction for self-supervised learning.

2. Related Work

Many of the previous work on self-supervised learning for speech recognition focus on learning speech representation. wav2vec (Schneider et al., 2019) applies contrastive learning to learn the future representation based on the past context. vq-wav2vec (Baevski et al., 2020a) uses wav2vec to learn the representations and quantizes them to discrete tokens, and performs BERT-style pre-training to further improve the representation learning. Discrete-BERT (Baevski et al., 2019) extends vq-wav2vec by fine-tuning the BERT-pre-trained model on the downstream tasks. wav2vec 2.0 (Baevski et al., 2020b) uses contrastive learning with both past and future context to predict the representation of the masked parts. HuBERT (Hsu et al., 2021) uses k-means to learn the initial quantizer that maps speech signals to discrete labels, and performs BERT-style pre-training where the inputs are masked speech signals and prediction targets are discrete labels. HuBERT further uses the pre-trained model as the new quantizer to train a new iteration of the model, and repeat the process to iteratively improve the pre-training results. w2v-BERT (Chung et al., 2021) uses a sub-network of the model to perform contrastive learning to learn speech representation, and use the rest of the network to perform BERT-style pre-training. w2v-BERT trains the representation learning and the BERT-style pre-training simultaneously. Our approach distinguishes from these work in avoiding the requirement of representation learning and




Figure 1. Overview of BEST-RQ. The approach applies random projections to project the input speech signals to a randomly initialized codebook, and map them to discrete labels through finding the nearest vector in the codebook. The pre-training objective is for the ASR encoder to take the masked input signals and predict the labels corresponding to the masked part provided by the random-projection quantizer.

separating the quantizer from the speech recognition model.

Our quantizer project input signals with a random matrix, which is similar to performing dimension reduction for the input signals. Using such quantization results as prediction target for self-supervised learning share a similar structure as the masked autoencoder (MAE) (He et al., 2021), which directly reconstruct the masked input signals. Another similar work in the computer vision community is BEiT (Bao et al., 2021), which trains a VQ-VAE (van den Oord et al., 2018) as the quantizer and use the VQ-VAE to perform BERT-style self-supervised learning. Different from these approaches, our algorithm does not require training the quantizer which further simplifies the training process.

3. Self-supervised Learning with Random-projection Quantizer

BEST-RQ applies a random-projection quantizer to map speech signals to discrete labels to enable BERT-style pre-training for ASR encoders. The quantizer randomly initializes a matrix and a codebook, and uses the matrix to project the input speech signals and the codebook to find the nearest vector where the index of the vector is the label. The pre-training process masks the speech signals and feeds them to the ASR encoder and trains the ASR encoder to predict labels of the masked part. Both the randomly initialized matrix and codebook are fixed during the pre-training process. The input data is normalized to have 0 mean and standard deviation of 1. The normalization is critical for pre-

venting the random projection to collapse to a small subset of codes. The framework is described in Figure 1. After the pre-training process the resulting ASR encoder is adopted to fine-tune on downstream ASR tasks.

The approach applies masks directly on the speech signal, where the masking strategy samples at every frame whether to apply masks with a fixed probability. Each mask spans from the starting frame with a fixed length. The masked parts are replaced with a noise sampled from a normal distribution with 0 mean and 0.1 standard deviation.

3.1. Random-projection Quantizer

Given an input vector x where x is a d -dimensional vector computed from speech signals, the random-projection quantizer maps x to discrete labels y through

$$y = \underset{i}{\operatorname{argmin}} \| \operatorname{norm}_{l_2}(c_i) - \operatorname{norm}_{l_2}(Ax) \|, \quad (1)$$

where A denotes a randomly initialized $h \times d$ matrix and $C = \{c_1, \dots, c_n\}$ is a set of randomly initialized h -dimensional vectors, $\operatorname{norm}_{l_2}()$ is a function that normalize the vector to have unit L_2 norm. The projection matrix A use Xavier initialization (Glorot & Bengio, 2010) and the codebook C use standard normal distribution for initialization, and the parameters are fixed during the pre-training process and therefore the quantizations are consistent during training.

3.2. Pre-training

The pre-training process adds a softmax layer on top of the ASR encoder to learn to predict the quantized speech labels. Since the random-projection quantizer is independent of the ASR encoder, the pre-training is flexible and can work with different architectures of the ASR encoder. We study the effectiveness of the algorithm on both non-streaming and streaming models, and in our experiments we use Conformer (Gulati et al., 2020) as the building block.

3.2.1. NON-STREAMING MODELS

Since the BERT-style pre-training is designed for the non-streaming models, training with this type of architecture is straightforward where the model uses both past and future context to learn to predict the quantized labels of the masked speech signals.

3.2.2. STREAMING MODELS

In addition to the non-streaming models, streaming architecture also plays a critical role for the speech recognition tasks as many of the applications require transcribing speakers' utterances with low-latency (Sainath et al., 2020). Streaming architecture however is less well-studied in the previous self-supervised learning work compared to the non-streaming ar-

chitecture. Moreover, many of the previous self-supervised learning approaches specify a pre-training setup that takes both the previous and future context, making it a question of how one can generalize the approaches to streaming models. We proposed two pre-training algorithms that are compatible with the streaming architecture:

Streaming pre-train. As our algorithm does not require learning quantization and focuses only on training the ASR encoder, this approach largely benefits the streaming models. Pre-training for streaming models follows the same setup as non-streaming models but the ASR encoder now learns to predict the quantized labels of the masked part based only on the past context.

Non-Streaming pre-train. Given that the neural network architecture like Transformer/Conformer allows switching from non-streaming to streaming behaviors by adding a mask for the future context within the same model, one can also perform pre-training with non-streaming setup for streaming models. Our algorithm provides benefits for streaming models with both non-streaming and streaming pre-training.

3.3. Fine-tuning

After the pre-training, the approach initializes the encoder of the downstream ASR from the pre-trained model, and fine-tunes on the supervised set. The softmax layer added on top of the encoder during the pre-training process is not used in fine-tuning. We focus on end-to-end models with RNN transducers (Graves, 2012), where the decoder uses LSTMs for the prediction network. On constructing the encoder, an additional projection layer is added on top of the pre-trained encoder to help it adapt to the downstream ASR task. The training process also updates the encoder during the supervised fine-tuning.

3.4. Understanding the Effectiveness of the Random-projection Quantizer

Our algorithm use a random-projection quantizer for self-supervised learning, and such a design raise two questions: how good is the resulting quantization quality with this quantizer and how much does the quantization quality affect the effectiveness of the self-supervised learning? We address these two questions through comparing our quantizer with VQ-VAEs. Using random-projections for quantizing speech signals shares some similarity as VQ-VAEs. The random projection performs dimension reduction for the speech signals while the random codebook provides an approximated discrete representation of the speech data distribution. VQ-VAEs also provide a discrete representation for the speech signals, but do so by learning a representation in the latent space that best preserves the speech data. Thus, comparing with VQ-VAEs gives us insight on the quantization quality

Table 1. LibriSpeech results with non-streaming models. The LM used in our experiment is a Transformer LM with model size 0.1B.

Method	Size (B)	No LM				With LM			
		dev	dev-other	test	test-other	dev	dev-other	test	test-other
wav2vec 2.0 (Baevski et al., 2020b)	0.3	2.1	4.5	2.2	4.5	1.6	3.0	1.8	3.3
HuBERT Large (Hsu et al., 2021)	0.3	–	–	–	–	1.5	3.0	1.9	3.3
HuBERT X-Large (Hsu et al., 2021)	1.0	–	–	–	–	1.5	2.5	1.8	2.9
w2v-Conformer XL (Zhang et al., 2020)	0.6	1.7	3.5	1.7	3.5	1.6	3.2	1.5	3.2
w2v-BERT XL (Chung et al., 2021)	0.6	1.5	2.9	1.5	2.9	1.4	2.8	1.5	2.8
BEST-RQ (Ours)	0.6	1.5	2.8	1.6	2.9	1.4	2.6	1.5	2.7

of our quantizer and the effect of representation learning for self-supervised learning.

We demonstrate that the quantization quality of the random-projection quantizer is not ideal but yet effective for self-supervised learning by comparing it with VQ-VAE-based quantizations in Section 4.3. We also show that the gap in the quantization quality is less an issue with the increase of the unsupervised data in Section 4.4. The main objective of self-supervised learning for speech recognition is to train the model to learn contextual information. The random-projection quantizer preserve the distribution of the speech data, and in order for the model to learn to predict the quantized token based on unmasked signals, the model needs to learn to process the raw signals and infer the contextual information among speech data. Such a criterion allows the model to perform effective self-supervised learning with a random-projection quantizer.

4. Experiments

We perform self-supervised learning experiments on LibriSpeech with non-streaming and streaming models, and assess the approach on multilingual tasks with non-streaming models. We study the quantization quality of the random-projection quantizer by comparing it with the quantizer learned with VQ-VAEs. The implementation use Lingvo (Shen et al., 2019) library.

4.1. LibriSpeech

Following (Zhang et al., 2020), we conduct experiments on the LibriLight dataset (Kahn et al., 2020) for pre-training, and fine-tune on the LibriSpeech training set which contains 960 hours of data. The input speech signals are 80-dimensional log-mel filter bank coefficients, and each frame has the stride of 10ms. In the fine-tuning phase, the decoder has a vocab size 1024 and uses a 1024-token WordPiece model (Schuster & Nakajima, 2012) for tokenizations that is constructed from the transcripts of the LibriSpeech training set.

4.1.1. NON-STREAMING MODELS

We use the same architectures reported in (Zhang et al., 2020) for fair comparisons. The model has two convolution layers at the bottom which provide 4 times temporal-dimension reduction for the input sequences. The rest of the layers are a stack of Conformer models. We explore 0.6B model size which is extensively studied in the previous works. The model contains 24 layers of Conformer models.

Pre-train. The pre-training uses mask length 400ms with masking probability of 0.01. The learning rate schedule uses a transformer learning rate schedule (Vaswani et al., 2017). The training of the model uses Adam optimizer (Kingma & Ba, 2015) with 0.004 peak learning rate and 25000 warm-up steps. The batch size is 2048. Since the encoder has 4 times temporal-dimension reduction, the quantization with random projections stacks every 4 frames for projections. The vocab size of the codebook is 8192 and the dimension is 16.

Fine-tune. The fine-tuning model also follow the same architecture as in (Zhang et al., 2020) and use RNN Transducer (RNN-T) (Graves, 2012) for decoder with 2 layers of unidirectional LSTMs, where the hidden dimension of LSTMs are 1280. The fine-tuning process use the Transformer learning rate schedule. Since the encoder is initialized from a pre-trained model, the fine-tuning process use a lower learning rate for the encoder than the decoder. The encoder use 0.0003 peak learning rate and 5000 warmup steps, while the decoder use 0.001 peak learning rate and 1500 warmup steps.

The results of pre-training with LibriLight and fine-tuning on LibriSpeech, along with comparisons with previous works, are shown in Table 1. Our results with LM use shallow fusion to incorporate the LM. The LM is a 0.1B Transformer model trained on the LibriSpeech LM corpus, and the model has 8 layers, 1024 model dimension, and 4096 feed-forward network dimension. By using the same architecture and similar optimization strategy as (Zhang et al., 2020), our approach shows similar WERs as previous best results on LibriSpeech both with and without LM.

4.1.2. STREAMING MODELS

The architecture we use for the streaming experiments follows a similar design as previous work for building streaming ASRs (Yu et al., 2021). We scale the model size to be also 0.6B to be consistent with the non-streaming experiments. The architecture has 3 Conformer layers at the bottom, followed by a stacking layer with 2 times temporal-dimension reduction and 20 Conformer layers on top of the stacking layer. The Conformer has 1024 hidden dimension for the self-attention layer and 4096 for the feed-forward layers. The self-attention layer attends to the current and the previous 64 frames, and the convolution has a kernel that covers the current and the past 3 frames.

The training setup is mostly the same as the 0.6B model in the non-streaming experiments, with some changes on the masking ratio for different pre-training approaches.

Streaming pre-train. The streaming pre-training uses the same setup as the original architecture, and the mask length is 300ms and the masking probability is 0.02. The random-projection quantizer stacks every 2 frames for projections.

Non-streaming pre-train. The non-streaming pre-training extends the original architecture to have access for future context by having the convolution kernel within the Conformer layer to have access for the future 3 frames. The self-attention is still limited to having access only for the previous context. We also explored having future context access for the self-attention, but this setup tends to be less stable. The masking length is 400ms and the masking probability is 0.02.

Fine-tune. The fine-tuning ASR model uses RNN-T for decoder with a layer of unidirectional LSTM, where the hidden dimension of the LSTM is 640. The training setup is the same as the fine-tuning config for the 0.6B model in the non-streaming experiments. When initializing from a non-streaming pre-trained model, the convolution only uses the kernel weight that accesses the previous context to keep the model streaming.

Latency measurement. A streaming model can learn to delay its prediction to access the future context and improve the prediction accuracy, and therefore it is critical to measure the latency of the streaming models to see whether the model maintains similar latency. This assessment helps us identify whether the underlying approach provides real improvement instead of trading off latency for prediction accuracy. Our latency comparison process first calculates the starting time and ending time of every word for each hypothesis generated by the two models. The comparison then aligns the hypotheses from the two models, finds the matching words and calculates the difference of their starting and ending time. The relative latency measurement is the average word timing difference of all matched words

between the two models among all utterances. Specifically the relative latency is calculated with

$$\sum_{i,j} \frac{s'_{ij} - s_{ij} + e'_{ij} - e_{ij}}{2N}, \quad (2)$$

where i denotes the index of the matched words between the two hypotheses, j is the utterance index, s_{ij} and e_{ij} correspond to the starting and ending time of the word from the baseline model, s'_{ij} and e'_{ij} correspond to the starting and ending time of the word from the compared model, and N is the total number of matched word among all utterances. A negative relative latency means the compared model has lower latency than the baseline model.

The word-error-rates and the relative latency are shown in Table 2. In this comparison experiment both wav2vec 2.0 and w2v-BERT use the same architecture, same masking and training setup as BEST-RQ. Using the conventional masking setup for wav2vec 2.0 and w2v-BERT gives worse performance. Since there is no convolution layers at the bottom, the contrastive learning use speech signals as targets. The w2v-BERT model use 12 layers for the contrastive module and 12 layers for the masked prediction module, to be consistent with the non-streaming setup (Chung et al., 2021). Our algorithm outperforms wav2vec 2.0 and w2v-BERT for both streaming and non-streaming pre-training. In particular our algorithm performs well with both pre-training, while wav2vec 2.0 and w2v-BERT favors more with non-streaming pre-training. This is likely due to the fact that the representation learning of both approaches is more compatible with non-streaming architectures. Increasing the model size from 0.1B to 0.6B results a slight increase in latency, but models trained with self-supervised learning algorithms has lower latency with streaming pre-training giving the most significant latency reduction. This indicates that the self-supervised learning preserve the low-latency property while providing quality gain.

4.2. Multilingual Tasks

We present multilingual results in this section. We use the same model setup as the LibriSpeech non-streaming experiment for these tasks.

4.2.1. DATA

Multilingual LibriSpeech(MLS-10hrs) The Multilingual LibriSpeech dataset (Pratap et al., 2020) is a large corpus derived from read audiobooks of Librivox and consists of 8 languages: *Dutch (du)*, *English (en)*, *French (fr)*, *German (de)*, *Italian (it)*, *Polish (pl)*, *Portuguese (pt)*, *Spanish (es)*. The latest version of this corpus contains around 50k hours including 44k hours in English. We use the official 10 hours split of training data to evaluate few-shot learning capabilities.

Table 2. LibriSpeech results compared with previous works with the same streaming architecture, and use LibriLight set for pre-training and LibriSpeech 960h set for fine-tuning. The relative latency (the lower the better) is the average difference of the word prediction time when comparing with the baseline Conformer 0.1B model. Our algorithm outperforms wav2vec 2.0 and w2v-BERT on both WERs and latency.

Method	Size (B)	dev	dev-other	test	test-other	Relative latency (ms)
Conformer 0.1B	0.1	4.1	10.3	4.5	9.8	0
Conformer 0.6B	0.6	3.9	9.8	4.4	9.4	15.3
Non-Streaming pre-train						
wav2vec 2.0	0.6	2.6	7.3	3.0	7.2	-10.1
w2v-BERT	0.6	2.8	7.2	3.3	6.9	-0.7
BEST-RQ (Ours)	0.6	2.5	6.9	2.8	6.6	-16.3
Streaming pre-train						
wav2vec 2.0	0.6	2.7	8.0	2.9	7.9	-130.6
w2v-BERT	0.6	2.7	8.4	3.0	8.1	-117.1
BEST-RQ (Ours)	0.6	2.5	6.9	2.8	6.6	-130.9

Multilingual Voice Search (VS-1000hrs) Our high resource finetune datasets is multilingual Voice Search dataset (Li et al., 2021). We sample random 1000 hour subsets (VS-1000h) across 15 languages, including English (US), English (IN), Spanish (US), Portuguese (BR), Spanish (ES), Arabic (GULF), Arabic (EG), Hindi (IN), Marathi (IN), Bengali (BD), Chinese (TW), Russian (RU), Turkish (TR), Hungarian (HU), and Malay (MY). The test set for each language contains around 3–19K utterances.

XLS-R unsupervised data (XLS-R -U) Our public unlabeled speech data follows the pre-training data used for XLS-R (Babu et al., 2021) with one major difference: we do not use any data from VoxLingua-107 due to license constraint. In total, we utilize approximately 429k hours of unlabeled speech data in 51¹ languages. As a consequence our model is pre-trained on speech from 51 languages as compared to 128 for XLS-R, and our pre-training set is smaller by 6.6k hours. We use this pretrain data on MLS-10hrs to compare with published results.

Youtube unsupervised data (YT-U) Following (Zhang et al., 2021), we collected a multilingual Youtube dataset for pretraining. For each language we prepare an unlabeled YouTube dataset segmented using voice activation detection (VAD (Zazo Candil et al., 2016)). The number of hours per languages are: English (800k hrs), Spanish (800k hrs), Marathi (600k hrs), Portuguese (800k hrs), Russian (800k), Arabic (800k), Hindi (800k), Chinese (800k), Malay (250k), Turkish (800k), Bengali (800k), Hungarian (300k). In practice, we found this data performs much better than XLS-R -U on VS-1000hrs. Thus, we use this pretrain data on VS-1000hrs to compare the performance of different pretrain

¹Counting languages with more than 1 hour of speech data.

methods.

4.2.2. RESULTS ON MLS-10HRS

We conduct our multilingual low resource finetune experiments on MLS-10hrs. We use XLS-R -U as pretraining data and finetune it on MLS-10hrs. As shown in Table 3, our baseline w2v-BERT already outperform previous strong model from XLS-R(2B) (Babu et al., 2021). The average WER further bring down by 3% relative by using the proposed BEST-RQ. This demonstrate a simple random-projection quantizer is also effective for multilingual pre-training. We also report finetune results on the MLS full supervised data. Interestingly, with more finetune data, BEST-RQ perform even better than w2v-BERT, especially for pt and pl. Our results also comparable with previously state-of-the-art results in (Bai et al., 2021) which conduct joint training for multilingual ASR.

While fine-tuning with MLS-full and MLS-10hrs both exhibit improvement compared to existing approaches, fine-tuning with MLS-full provides more relative improvement. This likely implies that pre-training with random-projection quantizers is more effective when there is more fine-tuning data.

4.2.3. RESULTS ON VOICE SEARCH

To understand how the proposed model work for high resource (1000hrs per language), we pretrain our model on YT-U and finetune it on VS-1000hrs. We can see with more finetune data, the relative improvement is smaller compared with no pretrain baseline. However, our proposed BEST-RQ consistently outperform w2v-BERT and wav2vec 2.0 by 9% and 5% relatively. Compare to w2v-BERT, our proposed method outperform on all the languages. Among the 15

Table 3. Test set WER (%) comparisons on the MLS full and 10hrs set.

Exp.	Languages								Avg.
	en	de	nl	fr	es	it	pt	pl	
MLS-full									
wav2vec 2.0 from XLSR-53 (Conneau et al., 2020)	-	7.0	10.8	7.6	6.3	10.4	14.7	17.2	10.6
w2v-BERT from JUST (Bai et al., 2021)	6.6	4.3	9.9	5.0	3.8	9.1	14.6	8.1	7.8
JUST (Bai et al., 2021) (co-train)	6.5	4.1	9.5	5.2	3.7	8.8	8.0	6.6	6.5
w2v-BERT (0.6B)	5.5	4.3	10.9	5.6	4.5	10.1	13.4	11.2	8.2
BEST-RQ (Ours, 0.6B)	6.8	4.1	9.7	5.0	4.9	7.4	9.4	5.2	6.6
MLS-10hrs									
XLSR-53 (Conneau et al., 2020)	14.6	8.4	12.8	12.5	8.9	13.4	18.2	21.2	13.8
XLS-R(0.3B) (Babu et al., 2021)	15.9	9.0	13.5	12.4	8.1	13.1	17.0	13.9	12.8
XLS-R(1B) (Babu et al., 2021)	12.9	7.4	11.6	10.2	7.1	12.0	15.8	10.5	10.9
XLS-R(2B) (Babu et al., 2021)	14.0	7.6	11.8	10.0	6.9	12.1	15.6	9.8	11.0
w2v-BERT (0.6B)	12.7	7.0	12.6	8.9	5.9	10.3	14.6	6.9	9.9
BEST-RQ (Ours, 0.6B)	12.8	7.4	12.7	9.6	5.4	9.9	12.1	7.1	9.6

Table 4. Test set WER (%) comparisons using YT-U for pretrain and VS-1000hrs for finetune, across 15 languages.

Exp.	Avg. on 15 langs (VS)
Baseline (0.6B)	12.6
wav2vec 2.0 (0.6B)	12.0
w2v-bert (0.6B)	11.5
BEST-RQ (Ours) (0.6B)	10.9

languages, English, Portuguese, Russian and Turkish, are improved more than 10%, relatively. Indic languages (Hindi, Marathi and English (IN)) are only slightly improved, all smaller than 3% relatively.

4.3. Analyzing Quantization Quality

As our self-supervised learning algorithm eliminates the requirement of representation learning through applying a random-projection quantizer, it is crucial to understand the representation quality of this quantizer and how the quality of the quantization affect the self-supervised learning. We analyze the quality of quantizers by training ASR models feeding labels generated by quantizing utterances as input. The performance of the resulting ASR provides us insights on the quality of the quantizer. The ASR model embeds quantized labels and feeds the embedding to a stack of Conformer layers, followed by a CTC decoder. 16 Conformer layer has feature dim 256, local self attention with 8 heads and 128 context length and kernel size 5 for lightweight convolution, in total the model size is 25M. We study the effect of representation learning through comparing with

quantizers trained with the VQ-VAE. We compare 3 types of quantizers: a) a random-projection quantizer b) a quantizer trained with VQ-VAE where the encoder has the same architecture as the random-projection quantizer and the decoder contains only a projection layer c) a trained VQ-VAE whose encoder/decoder are Transformer models. For trained quantizers, we train on the whole LibriSpeech 960 hours audio-only data, with a constant learning rate of $1e-4$ and train for 400k steps with batch size 256. For all quantizers, the input frames are stacked with 3 frames on each’s left, resulting in 4x input length reduction. We also use the quantizers for self-supervised learning with the LibriSpeech 0.6B non-streaming setup to compare their performance.

Table 5 shows the WER on LibriSpeech 960h. Both the random-projection quantizer and the projection-based VQ-VAE quantizer lead to poor ASR performance, while the Transformer-based VQ-VAE quantizer provides a significantly better performance. This implies that the Transformer-based VQ-VAE quantizer learns a better representation. On the other hand, when using these quantizers for the purpose of self-supervised learning, all quantizers lead to similar WERs. This indicates that the quantizer quality does not translate to self-supervised learning quality.

4.4. Analyzing the Effect of Pre-training Data Size

One potential explanation for the above observation, that a sub-optimal quantization can work well for self-supervised learning, is that the self-supervised learning algorithm can learn to mitigate the quality gap given sufficient amounts of pre-training data. We investigate whether a quantizer with a better quantization quality performs better when the

Table 5. Quantizer quality’s impact on ASR tasks. Although the Transformer-based quantizer gets much better performance when used as input directly, the random-projection quantizer is equally effective for self-supervised learning. The model used in the direct ASR task has size 25M. The self-supervised learning tasks use the same setup as the LibriSpeech non-streaming experiment, which use LibriLight for pre-training and LibriSpeech for fine-tuning and has 0.6B model size.

Configuration	Quantizer size (M)	Direct ASR WER				Pretrain-finetune WER			
		dev	dev-other	test	test-other	dev	dev-other	test	test-other
Random quantizer	1	58.8	78.8	57.9	72.8	1.5	2.8	1.6	2.9
Projection VQ-VAE	1	61.4	74.8	60.9	75.2	1.5	2.8	1.6	2.9
Transformer VQ-VAE	10	17.8	35.8	17.6	36.1	1.4	2.9	1.6	3.1

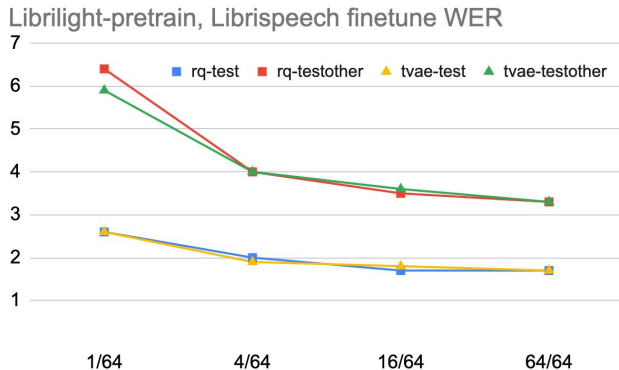


Figure 2. Comparing the self-supervised learning quality of the random-projection quantizer (rq) and the Transformer-based VQ-VAE quantizer (tvae) with different pre-training data size. Starting from low amount of pre-train data, the random-projection quantizer is behind the trained Transformer VQ-VAE quantizer. As the amount of pre-train data increases, the random-projection quantizer catches up.

amount of the pre-training data is limited, and whether increasing the amount of the pre-training data alleviate the discrepancy when compared to a random-projection quantizer. In this study, we compare self-supervised learning quality between a random-projection quantizer (rq) and a trained transformer-based VQ-VAE quantizer (tvae) with different pre-training data sizes. The random quantizer is untrained, and 4 Transformer VQ-VAE quantizers are trained with $\{1/64, 4/64, 16/64, 64/64\}$ LibriLight data, respectively. Then 4 identical random-projection quantizers and the above 4 transformer VAE quantizers are pre-trained separately with the same distinct percentages of LibriLight data as above for 100k steps with global batch size 2048. The pre-trained models fine-tune on LibriSpeech 960h. The result in Figure 2 shows that a quantizer with better representation quality (Transformer-based VQ-VAE) performs better when pre-training data is limited, but the gap disappears as the pre-training data increase.

5. Conclusions and Discussions

We proposed BEST-RQ to perform self-supervised learning for speech recognition models. BEST-RQ uses a random-projection quantizer to quantize speech signals to discrete labels. The pre-training process masks the speech signals and trains the model to predict labels corresponding to the masked parts. This approach shows similar WERs as the existing state-of-the-art results on LibriSpeech with non-streaming models, and outperform wav2vec 2.0 and w2v-BERT on LibriSpeech with streaming models and on multilingual tasks with non-streaming models. Further analysis showed that despite the fact that the random-projection quantizer provides a poorer representation compared to a trained VQ-VAE quantizer, it is effective for the purpose of self-supervised learning.

Our algorithm untangle the quantizer from the speech recognition model and also eliminates the requirement of representation learning. This simpler framework makes it easier to find a good recipe for the target task. The improvement on streaming models shows that the separation of the quantizer from the model makes the algorithm more effective for architectures that can be less effective for representation learning. The improvement on multilingual tasks shows that complicated tasks can benefit more from a simpler framework where finding a good recipe becomes more challenging. The quantization quality analysis implies that representation learning is not necessarily critical for self-supervised learning.

Codebook utilization. One of the most critical factors for pre-training quality is the percentage of the codebook that is used during training. In particular, at each training step a higher percentage of the codebook being used in each batch correlates strongly with a good pre-training quality. When the distribution of the codebook utilization is skewed toward a smaller subset of codes, this usually makes the pre-training task easier and provides less effective pre-training. The l_2 normalizations on the projected vector and the codebook are critical for providing more uniform codebook utilization. On the other hand, using randomly initialized codebook and projection matrix can introduce different codebook utiliza-

tions with different random seeds, which impact the pre-training quality across different runs with same experiment configurations. This variance impacts quality more when training with smaller pre-training and fine-tuning datasets. How to reduce this reproducibility issue caused by random initialization is an important next step for improving random-projection quantizations.

Hyperparameters. The pre-training quality is not very sensitive to the codebook vocab size and the codebook dimension, and is more sensitive to the masking probability and the mask length. The role of the projection layer in the random-projection quantizer is to allow using different codebook dimensions, and one can achieve similar results without the projection and set the codebook dimension to be the same as the input dimension. Due to the variance coming from the random initialization, the impact of a hyperparameter usually requires multiple runs of experiments to verify the result.

Longer convergence time for non-streaming models. One observation we have is that the algorithm takes more steps to converge with non-streaming models. We still observe improvement compared to wav2vec 2.0 and w2v-BERT at the same training step on multilingual tasks, though the final convergence usually takes 50% more steps. On the other hand, our training setup follows (Zhang et al., 2020), and it is unclear to us whether further hyperparameter tuning can help the model to converge faster. We did not observe the longer convergence property with streaming models.

Initialization. The quantizer uses random initialization and does not update the parameters, and therefore the initialization algorithm can play an important role on the results. In this paper we showed results with Xavier initialization for the projection matrix and the standard normal distribution for the codebook, and further comparisons on different initialization algorithms can be conducted in the future work.

6. Acknowledgements

We thank Wei Han and Johan Schalkwyk for helpful discussions, and Rohit Prabhavalkar, Izhak Shafran, and Hagen Soltau for insightful feedback. We also want to thank Bo Li for the help on multilingual tasks.

References

- Babu, A., Wang, C., Tjandra, A., Lakhota, K., Xu, Q., Goyal, N., Singh, K., von Platen, P., Saraf, Y., Pino, J., et al. Xls-r: Self-supervised cross-lingual speech representation learning at scale. *arXiv preprint arXiv:2111.09296*, 2021.
- Baevski, A., Auli, M., and Mohamed, A. Effectiveness of self-supervised pre-training for speech recognition. *arXiv preprint arXiv:1911.03912*, 2019.
- Baevski, A., Schneider, S., and Auli, M. vq-wav2vec: Self-supervised learning of discrete speech representations. In *ICLR*, 2020a.
- Baevski, A., Zhou, H., Mohamed, A., and Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations. *arXiv preprint arXiv:2006.11477*, 2020b.
- Bai, J., Li, B., Zhang, Y., Bapna, A., Siddhartha, N., Sim, K. C., and Sainath, T. N. Joint unsupervised and supervised training for multilingual asr. *arXiv preprint arXiv:2111.08137*, 2021.
- Bao, H., Dong, L., and Wei, F. Beit: Bert pre-training of image transformers, 2021.
- Chung, Y.-A., Zhang, Y., Han, W., Chiu, C.-C., Qin, J., Pang, R., and Wu, Y. W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training. *arXiv preprint arXiv:2108.06209*, 2021.
- Conneau, A., Baevski, A., Collobert, R., Mohamed, A., and Auli, M. Unsupervised cross-lingual representation learning for speech recognition. *arXiv preprint arXiv:2006.13979*, 2020.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 249–256. PMLR, 13–15 May 2010.
- Graves, A. Sequence transduction with recurrent neural networks. *CoRR*, abs/1211.3711, 2012.
- Gulati, A., Qin, J., Chiu, C.-C., Parmar, N., Zhang, Y., Yu, J., Han, W., Wang, S., Zhang, Z., Wu, Y., and Pang, R. Conformer: Convolution-augmented transformer for speech recognition, 2020.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021.
- Hsu, W.-N., Bolte, B., Tsai, Y.-H. H., Lakhota, K., Salakhutdinov, R., and Mohamed, A. HuBERT: Self-supervised speech representation learning by masked prediction of hidden units. *arXiv preprint arXiv:2106.07447*, 2021.

- Kahn, J., Rivière, M., Zheng, W., Kharitonov, E., Xu, Q., Mazaré, P.-E., Karadayi, J., Liptchinsky, V., Collobert, R., Fuegen, C., Likhomanenko, T., Synnaeve, G., Joulin, A., Mohamed, A., and Dupoux, E. Libri-light: A benchmark for ASR with limited or no supervision. In *ICASSP*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Li, B., Pang, R., Sainath, T. N., Gulati, A., Zhang, Y., Qin, J., Haghani, P., Huang, W. R., and Ma, M. Scaling end-to-end models for large-scale multilingual asr. *arXiv preprint arXiv:2104.14830*, 2021.
- Pratap, V., Xu, Q., Sriram, A., Synnaeve, G., and Collobert, R. Mls: A large-scale multilingual dataset for speech research. In *INTERSPEECH*, 2020.
- Sainath, T. N., He, Y., Li, B., Narayanan, A., Pang, R., Bruguier, A., Chang, S.-y., Li, W., Alvarez, R., Chen, Z., and et al. A streaming on-device end-to-end model surpassing server-side conventional model quality and latency. In *ICASSP*, 2020.
- Schneider, S., Baeovski, A., Collobert, R., and Auli, M. wav2vec: Unsupervised pre-training for speech recognition. *arXiv preprint arXiv:1904.05862*, 2019.
- Schuster, M. and Nakajima, K. Japanese and Korean voice search. *2012 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2012.
- Shen, J., Nguyen, P., Wu, Y., Chen, Z., and et al. Lingvo: a modular and scalable framework for sequence-to-sequence modeling, 2019.
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural discrete representation learning, 2018.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention Is All You Need. *CoRR*, abs/1706.03762, 2017. URL <http://arxiv.org/abs/1706.03762>.
- Yu, J., Chiu, C.-C., Li, B., et al. FastEmit: Low-latency Streaming ASR with Sequence-level Emission Regularization. In *Proc. ICASSP*, 2021.
- Zazo Candil, R., Sainath, T. N., Simko, G., and Parada, C. Feature learning with raw-waveform cldnns for voice activity detection. In *Interspeech 2016*, 2016.
- Zhang, Y., Qin, J., Park, D. S., Han, W., Chiu, C.-C., Pang, R., Le, Q. V., and Wu, Y. Pushing the limits of semi-supervised learning for automatic speech recognition. *arXiv preprint arXiv:2010.10504*, 2020.
- Zhang, Y., Daniel Park, S., Han, W., Qin, J., Gulati, A., Shor, J., Jansen, A., Xu, Y., Huang, Y., Wang, S., Zhou, Z., Li, B., Ma, M., Chan, W., Yu, J., Wang, Y., Cao, L., Sim, K. C., Ramabhadran, B., Sainath, T. N., Beaufays, F., Chen, Z., Le, Q. V., Chiu, C.-C., Pang, R., and Wu, Y. Bigssl: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition. *arXiv preprint arXiv:2109.13226*, 2021.