



Unsupervised Data Selection via Discrete Speech Representation for ASR

Zhiyun Lu, Yongqiang Wang, Yu Zhang, Wei Han, Zhehuai Chen, Parisa Haghani

Google Inc.

{zhiyunlu, yqw, ngyuzh, weihan, zhehuai, parisah}@google.com

Abstract

Self-supervised learning of speech representations has achieved impressive results in improving automatic speech recognition (ASR). In this paper, we show that data selection is important for self-supervised learning. We propose a simple and effective unsupervised data selection method which selects acoustically similar speech to a target domain. It takes the discrete speech representation available in common self-supervised learning frameworks as input, and applies a contrastive data selection method on the discrete tokens. Through extensive empirical studies we show that our proposed method reduces the amount of required pre-training data and improves the downstream ASR performance. Pre-training on a selected subset of 6% of the general data pool results in 11.8% relative improvements in LibriSpeech test-other compared to pre-training on the full set. On Multilingual LibriSpeech French, German, and Spanish test sets, selecting 6% data for pre-training reduces word error rate by more than 15% relatively compared to the full set, and achieves competitive results compared to current state-of-the-art performances.

Index Terms: speech recognition, data selection, self-supervised pre-training, discrete speech representation

1. Introduction

Self-supervised pre-training has demonstrated great success in learning representations from unlabeled data and improving the downstream automatic speech recognition (ASR) task [1–5]. The paradigm learns good representations from unlabeled audio through a proxy task which predicts the masked part of input from its visible parts. Popular proxy tasks include contrastive task [1, 6]; BERT-style masked language modeling task [5, 7]; and reconstruction task [8].

While the majority of research in the field focus on designing better context prediction tasks and self-supervision losses, an important question is left unanswered: what *data* should we use to do self-supervised pre-training? Is it always the more data the better? In supervised learning, it is well-known that learning from data of matched domain is important [9, 10]; in semi-supervised learning (teacher-student learning) [11–13], pseudo-label filtering and data weighting is carefully studied [14–16]. Few recent works study the effect of data selection on self-supervised learning [17–19]. [18] shows that domain shift hurts the self-supervised pre-training, and [19] shows that frame selection and data reweighting is helpful.

In this paper, we show that data selection is important for self-supervised learning. To this end, we propose a simple and flexible unsupervised data selection framework. As shown in Fig. 1, we first encode an utterance into a discrete speech representation as is commonly done in self-supervised learning methods [3, 5], and then apply a token-based data selection method on the discrete token sequences. The framework can work with different quantizers and data selection modules.

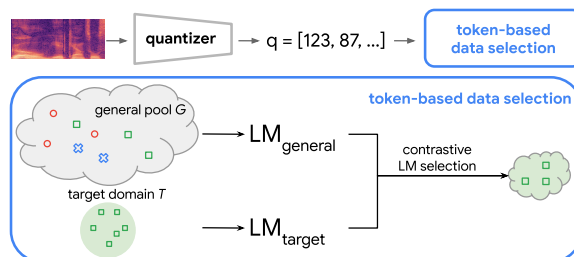


Figure 1: *Unsupervised data selection framework: the utterance is first quantized to a sequence of discrete tokens, and then we apply token-based data selection method to the discrete speech sequences. The quantizer is readily available from popular self-supervised learning methods. We use contrastive language model selection as the selection method.*

Inspired by [20], we use the discrete speech tokens as the features for data selection. For the data selection module, in order to select data that is close to the target domain, we adopt the language model based contrastive method introduced in [21]. Contrastive data selection is a well-known technique in the NLP and speech communities, and has been widely used in machine translation [22, 23], and ASR [24, 25]. It computes a domain relevance score for each utterance and filters by a threshold to keep those closest to the target domain.

Our proposed data selection framework has the following benefits: *i)* It improves the downstream ASR performance by exploiting similar non-domain-specific audio. *ii)* It is data efficient and greatly reduces computation in pre-training as it can select a small subset of data for pre-training. *iii)* It is *unsupervised*: it does not require labeled data which is appealing for low resource languages and domains. *iv)* It does not require extra feature learning as it uses available discrete tokens from self-supervised learning frameworks. *v)* Lastly, our experiment shows that it is not sensitive to the choice of quantizer nor other hyper-parameters thus almost tuning free.

We demonstrate the effectiveness of the data selection method for pre-training on LibriSpeech and multilingual LibriSpeech (MLS) datasets. By selecting 60k hours of YouTube speech which is acoustically similar to LibriSpeech for pre-training, we get word error rate (WER) 1.7% and 3.0% on test-clean and test-other set, after fine-tuning on LibriSpeech 960h. Our method reduces WER by 11.8% relative on test-other compared to pre-training on the full set of 1 million hours. This performance is close to the state-of-the-art model pre-trained with the in-domain Libri-light 60k data [5]. On MLS French, German, and Spanish test sets, pre-training on selected 60k hours of YouTube speech gets WER of 3.7, 3.3, 3.2 respectively, which is better than the state-of-the-art WERs achieved by multilingual ASR model [26]. The relative WER reduction compared to pre-training on over 1 million hours is 18%, 15%, and 26%. Our selection method is also applicable in supervised learning, as shown through experiments on WSJ and CHiME-6.

2. Unsupervised Data Selection via Discrete Speech Representation

We first define the task of data selection for self-supervised learning and discuss related works in § 2.1, and then describe our unsupervised data selection method in § 2.2 and 2.3.

2.1. Data selection for self-supervised learning

Data selection improves ASR performance and data efficiency by identifying the most informative training examples. Concretely, the task of data selection for self-supervised learning is, given labeled data from a target domain T and *unlabeled* data from a general pool G , selecting the best subset of G for self-supervised pre-training, such that the model pre-trained on that subset and then fine-tuned on T achieves the best performance for the target domain. We assume that G is much larger than T .

Confidence filtering is a classic method in data selection for ASR [27, 28], and has been successfully applied to end-to-end models [12, 29, 30]. However, confidence methods focus on data of good *transcript quality*, which might not be useful for self-supervised pre-training. Moreover, training a confidence model requires supervised data, which is not often feasible. [19] proposes to do frame-level data selection and utterance-level reweighting, which is complementary to our task.

We propose to do data selection by applying contrastive data selection on discrete speech representations. It selects data that is acoustically similar to the target domain, and does not require any labeled data in the loop.

2.2. Discrete speech representation

An emergent trend in self-supervised learning for speech is to learn discrete tokens from the continuous speech signals [31, 32]. The discrete representations make it applicable to NLP methods that require discrete inputs, for example BERT-style pre-training algorithms. Empirically, it leads to better results compared to without quantization [3, 4].

A quantizer maps continuous features into discrete tokens from a learnt codebook. It can be placed at different depths of the network, which leads to latent codes of different semantic level. In wav2vec 2.0 family of models [3, 5], the quantization is immediately after the feature encoder and before any transformer/Conformer representation learning layers. Therefore the discrete tokens is of low semantic level, and could contain information like pitch, background noise, and other confounding details of the audio signal. On the other hand, in the generative VQ-VAE, the quantization is at the top of all representation layers, where the discrete code is more abstract. [31] shows the VQ-VAE quantized code is predictive of the phonetic content of the utterances. Moreover, the discrete token is used differently in the training objectives. W2v-BERT uses the token ID as the target label in the cross-entropy loss of the masked language model task; VQ-VAE uses the discrete token to look up a 1-of-K embedding vector, which is then used to reconstruct the spectrogram. It’s an open research question on how the quality of the quantization affects the self-supervised learning [20]. In this work, we experiment with both w2v-BERT and VQ-VAE discrete token sequences as input to the contrastive data selection module described next.

2.3. Contrastive data selection

Contrastive data selection selects examples from a source corpus that are matched to a target domain. It was originally pro-

posed on text data [21], and has since been successfully applied in other domains such as machine translation and ASR. We extend this method to work on discrete speech tokens.

More specifically, we train two language models (LM) on corpora of the target domain T and of general domain G using the discrete speech tokens as input. For each utterance, we compute the log probability difference [21, 25] between the two LMs normalized by the number of tokens in the utterance. Assume \mathbf{q} is the vector quantized representation of the utterance, the domain relevance score is $\frac{\log P_T(\mathbf{q}) - \log P_G(\mathbf{q})}{\text{length}(\mathbf{q})}$, where P_T and P_G are the probabilities of the target and general domain LM respectively. We select utterances with the top domain relevance scores. See Fig. 1 for an illustration.

3. Experimental Setup

While the proposed selection method is unsupervised, it can be applied to both unlabeled data for pre-training, and labeled data for supervised learning. We investigate both settings in the empirical study. We describe the datasets in § 3.1, the quantizer in § 3.2 and contrastive selection setup in § 3.3. We provide model architecture in § 3.4, and training hyper-parameters in § 3.5. In each subsections, we present pre-training and supervised learning experiments separately.

3.1. Data

In the data selection for pre-training experiments, we use LibriSpeech and MLS as the target dataset, and YT-U as the general pool. We first pre-train the encoder on YT-U and then fine-tune on the LibriSpeech or MLS labeled set. The quantizer for speech tokenization is trained on YT-U. In the data selection for supervised learning experiments, we use WSJ and CHiME-6 as the target datasets, and the People’s Speech dataset as the general pool. We add the selected People’s Speech data to the in-domain training set for supervised learning. We reuse a quantizer trained on SpeechStew for these experiments. See Table 1 summarizes the datasets in each experiment.

Table 1: Summary of datasets used in the experiments.

Experiment	Target domain T	General pool G	Quantizer data
Pre-train (§ 4.1.4.2)	LibriSpeech / MLS	YT-U	YT-U
Supervised (§ 4.3)	WSJ / CHiME-6	People’s Speech	SpeechStew

LibriSpeech [33] We consider both LibriSpeech 100 hours and 960 hours as the supervised fine-tuning data. We report word error rates (WERs) on the dev-clean, dev-other, test-clean, and test-other evaluation sets. We also use Libri-light 60k in pre-training as a baseline for comparison.

Multilingual LibriSpeech (MLS) [34] We use French, German, and Spanish sets from MLS corpus. We use MLS-full as the fine-tuning supervised data, and report WERs on test sets.

YouTube unsupervised data (YT-U) is collected from speech-heavy videos that cover many different domains, including lectures, news, and etc. The raw audio is segmented by a voice activity detector [35] to a length of 32 seconds per utterance. We prepare speech data of 1.0 million hours in English (YT-U-En), 1.2 million in French (YT-U-Fr), 1.1 million in German (YT-U-De), and 1.0 million in Spanish (YT-U-Es). We use YT-U to refer to one of YT-U-En, YT-U-Fr, YT-U-De, and YT-U-Es, when it is clear from the context.

Wall Street Journal (WSJ) (LDC93S6B, LDC94S13B) is 80 hours of read speech from Wall Street Journal news text corpus. We report WERs on the eval92 test set, and score our results

Table 2: Comparing WERs on Librispeech dev and test sets when the encoder is pre-trained on YouTube with and without data selection. The model is fine-tuned on 960h and 100h. Pre-training on 6% of selected subset reduces WER by 11.8% relative on test-other set.

Pre-train data	Pre-train # hours	Fine-tune 960h		Fine-tune 100h	
		Dev clean / other	Test clean / other	Dev clean / other	Test clean / other
Libri-light [5]	60k	1.5 / 2.9	1.5 / 2.9	2.4 / 4.4	2.5 / 4.6
YT-U-En	1,000k	1.7 / 3.3	1.8 / 3.4	3.0 / 5.8	3.1 / 5.9
YT-U-En select	60k	1.6 / 2.9	1.7 / 3.0	2.8 / 5.0	2.8 / 5.2
YT-U-En select + Libri-light	120k	1.5 / 2.7	1.6 / 2.8	2.4 / 3.9	2.5 / 4.4

with the Kaldi script [36].

CHiME-6 [37] is a set of approximately 40 hours of noisy distant microphone conversational speech in everyday home environments. We use the official front-end enhancement recipe to augment the dataset, and report WER on the test set where guided source separation with 12 channels enhancement is used.

People’s Speech [38] is a public speech dataset from diverse sources including movies, TV, local news, and etc. We use the 17.1k hours clean subset of People’s Speech.

SpeechStew [39] is an ensemble dataset combining 7 public speech corpora, including AMI, Common Voice, English Broadcast News, LibriSpeech, Switchboard/Fisher, TED-LIUM v3, and WSJ. We use the unlabeled speech in SpeechStew to train the quantizer.

3.2. Quantizer

In the pre-training experiment, we use both the w2v-BERT [5] and the VQ-VAE [8] as the quantizer. In the supervised learning experiment, we use a VQ-VAE quantizer. In w2v-BERT quantizer, the codebook vocab size is 1024 and the codebook dimension is 1024. In VQ-VAE, the vocab size is 8192 and the dimension is 16. For both models, we apply Gumbel-Softmax technique [40] to the quantization operation, which makes the argmax differentiable through a temperature parameter in the backward pass. The temperature is gradually annealed towards a small but non-zero value during training.

3.3. Contrastive data selection

We use the same set of contrastive data selection parameters across all experiments. We build N -gram language model (LM) with $N = 5$ and back-off. We use the Kneser-Ney LM interpolation algorithm [41] to estimate LM probabilities. We use a few hundred to a few thousand hours of data as input. In the pre-training experiments, we use the in-domain training set for the target LM, and a random subsample of a few hundred hours YT-U for the general LM. In the supervised experiment, we use WSJ or CHiME-6 training set to build the target LM, and the People’s Speech clean set to build the general LM. As we will show in the experiment, the method is not sensitive to the amount of data in LM training.

3.4. Model architecture

In all our experiments, we use 80-dimensional log-mel filter bank coefficients as the acoustic inputs, computed with a 25ms window and shifted every 10ms. For transcript tokenization, we use a 1024-token WordPiece model [42] for all English experiments, and 4096-token WordPiece models for French, German, and Spanish experiments.

In the pre-training experiments, we follow the w2v-BERT XL setup in [5] which is a 24 layer Conformer. The feature encoder has two 2D-convolution layers with strides (2,2). We

Table 3: Comparing WERs on MLS test sets in French, German and Spanish when the pre-training data is YouTube with and without data selection. The model is fine-tuned on MLS-Full. Pre-training on the selected subset reduces WERs by more than 15% relative in all 3 languages. Our best WER is better than the state-of-the-art performance obtained by multilingual models.

Pre-train data	# hrs	French	German	Spanish
Prior work	multilingual model			
SoTA		4.9 [46]	4.1 [26]	3.7 [26]
Our work	monolingual model			
YT-U	1,000k	4.5	3.9	4.3
YT-U select	60k	3.7	3.3	3.2

use the 12-th Conformer block hidden features to compute the contrastive loss and the last (24-th) layer features to compute the masked prediction loss. We use a RNN-T model for fine-tuning. The decoder is a two-layer LSTM with a hidden dimension of 640. The model has 600 million parameters. In the supervised learning experiments, we use a 17 layer Conformer model of 119 million parameters, the same as ConformerL in [43].

3.5. Training details

In the pre-training experiments, we use w2v-BERT as the pre-training recipe. We use a masking length of 400ms with masking probability of 0.065. We use Adam optimizer [44], and a transformer learning rate schedule [45] with 1e-3 peak learning rate and 25,000 warm-up steps. The batch size is 4096. In the fine-tuning, we use different learning schedules for the pre-trained encoder and the decoder. The encoder has a peak learning rate of 3e-4 and 5,000 warm-up steps, while decoder has 1e-3 and 1,500 steps. The batch size is 256.

For the supervised learning experiment, we use Adam optimizer with a schedule of 0.002 peak learning rate and 10,000 warm-up steps. The batch size is 256 for WSJ experiment, and 2048 for CHiME-6 experiment. We find a larger batch size is essential for CHiME-6, otherwise the model may fail to train.

4. Experimental Results

We present key results of data selection for pre-training in § 4.1, followed by ablation studies in § 4.2. Then we show the selection method is also applicable to supervised learning in § 4.3.

4.1. Data selection for self-supervised pre-training

Table 2 compares WERs on LibriSpeech dev and test sets when the models are pre-trained on YT-U-En with and without data selection, and then fine-tuned on LibriSpeech 960h or 100h. We compare with state-of-the-art performance from [5] in the first row, where the pre-training data is the in-domain Libri-light 60k hours. For our models, we pre-train the encoder for up to 800k steps. Since there is no golden rule in picking pre-

Table 4: *Our unsupervised data selection method outperforms random sample and confidence filtering baselines. Both confidence and our method are better than no data selection.*

Selection	# hours	Dev clean / other	Test clean / other
no selection	1,000k	1.9 / 3.7	2.0 / 3.9
random	60k	1.9 / 4.1	2.1 / 4.1
confidence	60k	1.7 / 3.3	1.9 / 3.5
ours	60k	1.7 / 3.2	1.8 / 3.3

Table 5: *The effect of quantizer on fine-tuning WERs. The pre-training data is YT-U selected with w2v-BERT or VQ-VAE quantizers. Both quantizers perform similarly.*

Quantizer	Dev clean / other	Test clean / other
w2v-BERT	1.7 / 3.2	1.8 / 3.3
VQ-VAE	1.7 / 3.3	1.8 / 3.4

training checkpoints, we fine-tune at every 100k checkpoints. We pick the best one in terms of fine-tuning WERs on dev sets for each method to report. Comparing the YT-U with YT-U select row, we reduce WER by 5% and 11% relative on the test clean and other set respectively, and reduce the amount of pre-training data to 6%. In the last row, by combining the selected YT-U data with Libri-light, we are able to slightly improve the state-of-the-art WER on the test-other set.

Table 3 shows WERs on MLS French, German, Spanish test sets. For our monolingual models, the pre-training data is YT-U and the fine-tuning data is MLS-full. To the best of our knowledge, the state-of-the-art (SoTA) WERs on these test sets are achieved by multilingual models in [26, 46], as detailed in the first row. We pre-train for 100k steps for fast experimentation. Comparing YT-U with YT-U select, the WER reduction is 18%, 15%, and 26% for French, German and Spanish. And the amount of data is less than 6%.

In Table 4 we compare the proposed selection method with random subsample and confidence filtering. We provide WERs with no data selection for reference. For the confidence filtering method, we use a strong RNN-T model of 183 million parameters trained from around 150k hours of labeled YouTube data. We interpret RNN-T loss as the confidence measure following [29]. All methods pre-train for 100k steps and fine-tune on 960h. Both confidence filtering and our method are better than no data selection. Note the confidence method requires labeled YouTube data. In contrast, our method does not require any paired data, yet still outperforms the confidence method.

4.2. Ablation studies

We conduct ablation studies on the choice of the quantizer, and hyper-parameters of the contrastive selection method. For all variants in this section, the pre-training data is 60k select YT-U-En, and the fine-tuning data is LibriSpeech 960 hours. We pre-train for 100k steps, as the relative comparison stays the same for more steps.

Table 5 compares WERs when different quantizers are used to extract the discrete tokens for data selection. Both w2v-BERT and VQ-VAE quantizers perform similarly. The proposed method is not sensitive to the choice of the quantizer.

Table 6 compares WERs when different amount of data is used for LM training in the contrastive selection module. The quantizer is w2v-BERT. Our method is not sensitive to the amount of data used in LM training, once in a reasonable range.

Table 6: *The effect of LM training data on fine-tuning WERs. The pre-training data is YT-U selected with LMs trained from different data. All variants perform similarly.*

LM data # hours	Dev clean / other	Test clean / other
960h / 100h	1.7 / 3.2	1.8 / 3.3
100h / 100h	1.6 / 3.2	1.8 / 3.4
960h / 1 million h	1.7 / 3.1	1.7 / 3.4

Table 7: *WERs on WSJ and Chime6. We select 0.8k hours of data from People’s Speech using different methods and add it to the training set. Our method outperforms other baselines. By using only 7% of the selected subset, our method achieves close WER compared to the full set.*

Training set	+ # hours	WSJ	CHiME-6
in-domain wsj 80h / chime 40h [39]	0	28.2	66.7
+ random sample	0.8k	3.1	51.2
in-domain + text contrastive	0.8k	3.1	49.7
+ ours	0.8k	2.7	48.0
in-domain + full People’s Speech	11.8k	2.6	47.9

4.3. Data selection for supervised learning

While our work is mainly motivated to perform data selection for pre-training, the proposed method is also applicable to labeled data. We can select transcribed speech close to the target domain and add it to the training set to improve the ASR performance. This is useful for low resource domains.

To demonstrate the performance in this setting, we apply data selection to the People’s Speech dataset to improve ASR performance on WSJ and CHiME-6. We select around 800 hours of data from People’s Speech and add that to the in-domain training data for supervised learning. In Table 7, we compare three data selection methods: random sample, text contrastive LM selection, and our unsupervised method. We do not include confidence method because the ground-truth transcript is already given. We also present WERs when we use the original training set only from [39], and when we add the full 11.8k hours People’s Speech data.

Adding extra labeled data from People’s Speech greatly improves WERs on WSJ and CHiME-6. Our selection method outperforms random sample and text contrastive selection method. It reduce the amount of labeled data to 7%, and the WER is 0.1% worse compared to using the full People’s Speech data. This result shows that we can apply the proposed method to select unlabeled data to send for transcription, and achieves good WER with small amount of annotation.

Lastly, we compare the contrastive score ranking from quantized codes with contrastive score ranking from text. There is almost zero rank correlation between the two ordinal variables, measured by Kendall τ coefficient. It suggests that the two data selection methods can be complementary, and we leave empirical studies on this as future work.

5. Conclusion

We show that data selection is important for self-supervised pre-training. We propose a simple and effective unsupervised data selection method, which applies contrastive data selection on discrete speech representations. Our experimental results demonstrate the effectiveness of the proposed approach in improving the ASR performance in both unsupervised and supervised settings.

6. References

- [1] A. Van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv e-prints*, pp. arXiv:1807.2018, 2018.
- [2] S. Pascual, M. Ravanelli, J. Serra, A. Bonafonte, and Y. Bengio, "Learning problem-agnostic speech representations from multiple self-supervised tasks," *arXiv preprint arXiv:1904.03416*, 2019.
- [3] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proc. NeurIPS*, 2020.
- [4] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," in *Proc. ICLR*, 2019.
- [5] Y.-A. Chung, Y. Zhang, W. Han, C.-C. Chiu, J. Qin, R. Pang, and Y. Wu, "W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training," in *Proc. ASRU*, 2021.
- [6] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," in *Proc. Interspeech*, 2019.
- [7] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," in *IEEE T-ASLP*, 2021.
- [8] A. Van Den Oord, O. Vinyals *et al.*, "Neural discrete representation learning," in *Proc. NeurIPS*, 2017.
- [9] M. L. Seltzer, D. Yu, and Y. Wang, "An investigation of deep neural networks for noise robust speech recognition," in *proc. ICASSP*, 2013.
- [10] T. Likhomanenko, Q. Xu, V. Pratap, P. Tomasello, J. Kahn, G. Avidov, R. Collobert, and G. Synnaeve, "Rethinking evaluation in asr: Are our models robust enough?" *arXiv preprint arXiv:2010.11745*, 2020.
- [11] J. Kahn, A. Lee, and A. Hannun, "Self-training for end-to-end speech recognition," in *Proc. ICASSP*, 2020.
- [12] D. S. Park, Y. Zhang, Y. Jia, W. Han, C.-C. Chiu, B. Li, Y. Wu, and Q. V. Le, "Improved noisy student training for automatic speech recognition," *Proc. Interspeech*, 2020.
- [13] T. Doutré, W. Han, M. Ma, Z. Lu, C.-C. Chiu, R. Pang, A. Narayanan, A. Misra, Y. Zhang, and L. Cao, "Improving streaming automatic speech recognition with non-streaming model distillation on unsupervised data," in *Proc. ICASSP*, 2021.
- [14] D. Charlet, "Confidence-measure-driven unsupervised incremental adaptation for hmm-based speech recognition," in *Proc. ICASSP*, 2001.
- [15] F. Wessel and H. Ney, "Unsupervised training of acoustic models for large vocabulary continuous speech recognition," *IEEE T-SAP*, 2004.
- [16] K. Vesely, L. Burget, and J. Cernocký, "Semi-supervised dnn training with word selection for asr," in *Proc. Interspeech*, 2017.
- [17] K. Kawakami, L. Wang, C. Dyer, P. Blunsom, and A. v. d. Oord, "Learning robust and multilingual speech representations," in *Proc. EMNLP*, 2020.
- [18] W.-N. Hsu, A. Sriram, A. Baevski, T. Likhomanenko, Q. Xu, V. Pratap, J. Kahn, A. Lee, R. Collobert, G. Synnaeve *et al.*, "Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training," in *Proc. Interspeech*, 2021.
- [19] M. K. Baskar, A. Rosenberg, B. Ramabhadran, Y. Zhang, and P. Moreno, "Ask2mask: Guided data selection for masked speech modeling," *arXiv preprint arXiv:2202.12719*, 2022.
- [20] C.-C. Chiu, J. Qin, Y. Zhang, J. Yu, and Y. Wu, "Self-supervised learning with random-projection quantizer for speech recognition," *arXiv preprint arXiv:2202.01855*, 2022.
- [21] R. C. Moore and W. Lewis, "Intelligent selection of language model training data," in *Proc. ACL*, 2010.
- [22] A. Axelrod, X. He, and J. Gao, "Domain adaptation via pseudo in-domain data selection," in *Proc. EMNLP*, 2011.
- [23] M. van der Wees, A. Bisazza, and C. Monz, "Dynamic data selection for neural machine translation," in *EMNLP*, 2017.
- [24] F. Mezzoudj, D. Langlois, D. Juvet, and A. Benyettou, "Textual data selection for language modelling in the scope of automatic speech recognition," *Procedia Computer Science*, 2018.
- [25] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, G. Wang, and P. Moreno, "Injecting text in self-supervised speech pretraining," in *Proc. ASRU*, 2021.
- [26] J. Bai, B. Li, Y. Zhang, A. Bapna, N. Siddhartha, K. C. Sim, and T. N. Sainath, "Joint unsupervised and supervised training for multilingual asr," in *Proc. ICASSP*, 2022.
- [27] G. Zavaliagkos and T. Colthurst, "Utilizing untranscribed training data to improve performance," in *DARPA Broadcast News Transcription and Understanding Workshop*. Citeseer, 1998.
- [28] H. Y. Chan and P. Woodland, "Improving broadcast news transcription by lightly supervised discriminative training," in *Proc. ICASSP*, 2004.
- [29] Y. Zhang, D. S. Park, W. Han, J. Qin, A. Gulati, J. Shor, A. Jansen, Y. Xu, Y. Huang, S. Wang *et al.*, "Bigssl: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition," *arXiv preprint arXiv:2109.13226*, 2021.
- [30] D. Hwang, A. Misra, Z. Huo, N. Siddhartha, S. Garg, D. Qiu, K. Chai Sim, T. Strohmaier, F. Beaufays, and Y. He, "Large-scale asr domain adaptation by self-and semi-supervised learning," in *Proc. ICASSP*, 2022.
- [31] J. Chorowski, R. J. Weiss, S. Bengio, and A. Van Den Oord, "Unsupervised speech representation learning using wavenet autoencoders," in *T-ASLP*, 2019.
- [32] A. H. Liu, T. Tu, H.-y. Lee, and L.-s. Lee, "Towards unsupervised speech recognition and synthesis with quantized speech representation learning," in *Proc. ICASSP*, 2020.
- [33] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *Proc. ICASSP*, 2015.
- [34] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "Mls: A large-scale multilingual dataset for speech research," in *Proc. Interspeech*, 2020.
- [35] R. Zazo Candil, T. N. Sainath, G. Simko, and C. Parada, "Feature learning with raw-waveform cldnns for voice activity detection," in *Proc. Interspeech*, 2016.
- [36] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldi speech recognition toolkit," in *Proc. ASRU*, 2011.
- [37] M. M. Watanabe, Shinji *et al.*, "Chime-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings," in *ChIME 2020*.
- [38] D. Galvez, G. Damos, J. Ciro, J. F. Cerón, K. Achorn, A. Gopi, D. Kanter, M. Lam, M. Mazumder, and V. J. Reddi, "The people's speech: A large-scale diverse english speech recognition dataset for commercial usage," *arXiv preprint arXiv:2111.09344*, 2021.
- [39] W. Chan, D. Park, C. Lee, Y. Zhang, Q. Le, and M. Norouzi, "Speechstew: Simply mix all available speech recognition data to train one large neural network," *arXiv preprint arXiv:2104.02133*, 2021.
- [40] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," in *Proc. ICLR*, 2017.
- [41] K. Heafield, I. Pouzyrevsky, J. H. Clark, and P. Koehn, "Scalable modified Kneser-Ney language model estimation," in *Proc. ACL*, 2013.
- [42] M. Schuster and K. Nakajima, "Japanese and korean voice search," in *Proc. ICASSP*, 2012.
- [43] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *Proc. Interspeech*, 2020.
- [44] K. Diederik, B. Jimmy *et al.*, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2014.
- [45] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NeurIPS*, 2017.
- [46] B. Li, R. Pang, T. N. Sainath, A. Gulati, Y. Zhang, J. Qin, P. Haghani, W. R. Huang, M. Ma, and J. Bai, "Scaling end-to-end models for large-scale multilingual asr," in *Proc. ASRU*, 2021.