

# DuRep: Dual-Mode Speech Representation Learning via ASR-Aware Distillation

*Prabash Reddy Male, Swayambhu Nath Ray, Harish Arsikere, Akshat Jaiswal, Prakhar Swarup, Prantik Sen, Debmalaya Chakrabarty, K V Vijay Girish, Nikhil Bhawe, Frederick Weber, Sambuddha Bhattacharya, Sri Garimella*

Amazon AGI, Bangalore, India

prabash, swayar, arsikere, akshatkj@amazon.com

## Abstract

Recent advancements in speech encoders have drawn attention due to their integration with Large Language Models for various speech tasks. While most research has focused on either causal or full-context speech encoders, there's limited exploration to effectively handle both streaming and non-streaming applications, while achieving state-of-the-art performance. We introduce DuRep, a Dual-mode Speech Representation learning setup, which enables a single speech encoder to function efficiently in both offline and online modes without additional parameters or mode-specific adjustments, across downstream tasks. DuRep-200M, our 200M parameter dual-mode encoder, achieves 12% and 11.6% improvements in streaming and non-streaming modes, over baseline encoders on Multilingual ASR. Scaling this approach to 2B parameters, DuRep-2B sets new performance benchmarks across ASR and non-ASR tasks. Our analysis reveals interesting trade-offs between acoustic and semantic information across encoder layers.

**Index Terms:** Speech encoder, dual-mode, best-in-class, HuBERT, Best-RQ, knowledge distillation.

## 1. Introduction

The design of speech encoders, which extract meaningful representations from raw speech data, often involves a trade-off between performance and latency. Full-context encoders [1, 2], leveraging entire speech utterances, offer superior accuracy but higher latency, making them ideal for offline use. Conversely, streaming encoders process speech frame-by-frame, reducing latency for real-time applications like smart speakers and on-device assistants, though with some potential performance compromise. Recent efforts have focused on unifying full-context and streaming models for automatic speech recognition (ASR) tasks [3–5], combining the benefits of both encoder types while reducing training time and deployment costs, making them a compelling solution for speech-based applications.

Dual-mode frameworks proposed for ASR systems, such as [3], unify streaming and non-streaming ASR models by jointly training both modes, assisted by an additional in-place distillation from full-context to streaming mode. However, this approach involves separate modules for convolution and normalization, leading to additional parameters. [4] showed that model with shared convolution and normalization layers performs as good as models with mode-specific components.

Instead of limiting the model's context to full-context and pure streaming during training, [5] proposed sampling the future context from a distribution, adapting the model to a wide range of latency scenarios. To address the limitations of fixed-window masking and unintended look-ahead accumulation, [4] divided input speech into non-overlapping chunks for attention.

[6] utilised variable attention mask to balance efficient modeling and reduced future frame dependency. Despite the focus on transducer [7] or CTC-based objectives, previous works have not explored distillation strategies for training dual-mode encoders, leaving a gap that our research aims to address.

In this paper, we present a zero-additional-parameter training framework for a dual-mode speech encoder adaptable to various latency scenarios. Our contributions include:

- A novel variable attention masking based distillation method to create a dual-mode encoder from a full-context teacher encoder.
- Analysis of the impact of varying teacher encoder and context sampling strategy on dual-mode encoder's performance.
- Comparative layer-wise analysis of semantic and acoustic feature representation using the SUPERB framework [8].
- Demonstrate the scalability of our method by successfully scaling it to a 2B encoder, resulting in significant performance in both ASR and non-ASR tasks.

## 2. Methodology

### 2.1. Training stages

In this work, we propose a novel four-stage approach to develop a dual-mode encoder as demonstrated in Figure 1:

- **Full-context Pretraining (S1):** We begin by training a non-streaming (NS) encoder using BestRQ framework [9], which is a widely accepted pretraining method in the industry for speech encoders on unlabelled data [1, 10]. Pretraining helps prevent convergence issues for large model training.
- **Full-context Finetuning (S2):** The pretrained encoder from Stage 1 undergoes fine-tuning with an ASR-aware transducer objective, enhancing its performance for semantic tasks.
- **Dual-mode Distillation (S3):** A cross-entropy-based distillation process as outlined in Section 2.2, is employed to transfer knowledge from the non-streaming ASR-aware encoder to a dual-mode (DM) encoder, aka ASR-aware distillation. This stage utilizes variable attention masking, as detailed in Section 2.3, to achieve seamless performance across both streaming and non-streaming modes.
- **Dual-mode Finetuning (S4):** Similar to Stage 2, the dual-mode encoder from Stage 3 is fine-tuned using an ASR-aware transducer objective with variable attention masking. This final stage further optimizes the encoder for dual-mode operation, resulting in a robust and versatile encoder for downstream tasks.

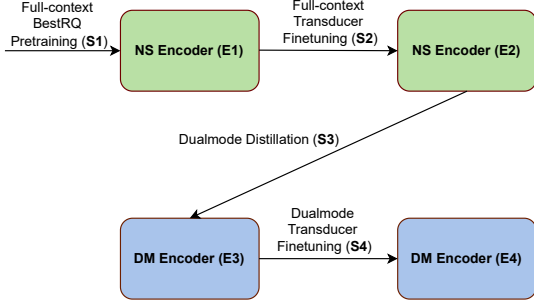


Figure 1: Dual-Mode encoder training framework

## 2.2. Knowledge Distillation

Following the pseudo-target prediction approach outlined in [11], during S3 (shown in Figure 1), we perform offline k-means clustering on the embeddings extracted from an intermediate conformer block of the full-context supervised encoder (E2). This block is chosen based on the semantic richness of its output representations (refer to section 5.2). The purpose of this clustering is to quantize the latent features of the full-context encoder, generating cluster IDs that serve as pseudo-labels for the dual-mode encoder (student) training. The dual-mode encoder is then trained to predict these cluster IDs. Our ablation studies showed that not applying masking to the input speech frames during distillation leads to a more effective student encoder.

## 2.3. Variable Attention Masking

In this work, we employ a variable attention masking strategy as discussed in [6] to optimize how the encoder processes contextual information. We introduce *variability* in the extent of the past and future contexts during the training process. Specifically, we define two uniform discrete distributions,  $L_{past}$  and  $L_{future}$ , which encompass all possible values for past and future context, respectively. For each input batch, the look-back (LB) and look-ahead (LA) values are independently sampled from these distributions. This approach exposes the encoder to a range of contextual scenarios, improving its ability to function effectively under different latency conditions.

To prevent the unintended accumulation of look-ahead information across multiple encoder layers, we deliberately avoid applying fixed-window masking to future context. Specifically, we use chunked attention for future context, while a fixed-window approach is applied to handle past context. This strategy is visually represented in Figure 2, which illustrates a look-back of 5 frames and a look-ahead of 2 frames. The impact of distribution’s sample space on encoder’s performance is examined in Section 4.3.

## 3. Experimental Details

We conducted all initial experiments and ablation studies on DuRep-200M, subsequently scaling our best-performing training strategy to DuRep-2B.

### 3.1. Data

We utilized a comprehensive speech dataset totalling 8.3M hours, which included both hand-transcribed speech data and a significant portion from untranscribed sources. Data consists

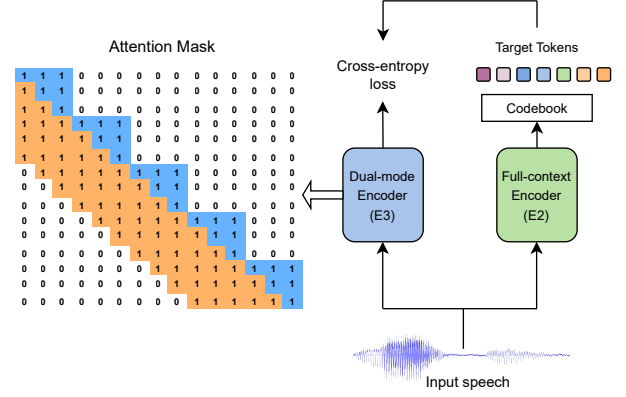


Figure 2: Cross entropy based distillation using variable attention masking - we use fixed-window for past and chunking for future context to prevent lookahead accumulation

of a combination of proprietary and publicly available datasets. Public datasets included Librispeech [12], Multilingual Librispeech [13], Commonvoice [14], PeopleSpeech [15], Fleurs [16], VoxPopuli [17], Librivox [18] and others. For stages S1 and S3, which do not require transcriptions, we utilized the complete dataset. For stages S2 and S4 (Transducer training), we used existing transcriptions where available and machine-generated transcriptions for the untranscribed partitions. These machine-generated transcriptions were produced using multiple proprietary models.

## 3.2. Encoder Architecture

Our proposed dual-mode encoder is based on the Conformer architecture [19], featuring a convolutional subsampling layer followed by a series of Conformer blocks. The input audio is processed into 128-dimensional log filter bank energy features, computed over a 25ms window with a 10ms shift. Each feature vector is stacked with 3 preceding frames and then downsampled to a 40ms frame rate, resulting in an overall input feature dimension of 512. The DuRep-200M model uses 1 convolutional front-end filter and 18 Conformer blocks, with each block having a model dimension (d\_model) of 512 and a feed-forward dimension (d\_ff) of 2048. Similarly, the DuRep-2B model also features 1 convolutional front-end filter and 20 Conformer blocks, with each block having a d\_model of 2048 and a d\_ff of 8192. All convolution operations are causal to prevent unintended look-ahead, and no positional embeddings are used, simplifying the encoder’s overall design. In contrast to the original conformer paper, our conformer block places the convolution module before the multi-head self-attention module. This arrangement implicitly learns positional information without the need for additional positional embeddings.

After conducting a SUPERB layer analysis (5.2) on E2, we selected the output from the 14th block for DuRep-200M and the 15th block for DuRep-2B for offline quantization during S3 (2.2). K-means clustering was performed on the extracted embeddings, with 2,500 clusters for DuRep-200M and 4,500 clusters for the DuRep-2B encoder. During transducer-finetuning (S2 and S4 of Figure 1), the prediction network consisted of a 2-layer LSTM model [20] with 1,280 hidden units, and the joint network contained 1,280 units. SpecAug [21] is applied during all supervised training phases to enhance generalization. Dur-

ing variable-attention masking of S3 and S4, we sample past and future context (in seconds) uniformly from  $L_{past}=[\infty, 5.4, 4.6, 3.6]$  and  $L_{future}=[0, 1, 1.8, \infty]$ . These values were chosen based on our internal encoder usage scenarios.

### 3.3. Baselines

We adopt the well-established BestRQ framework for pre-training our 200M full-context and streaming baseline encoders. The pre-trained encoders are then individually fine-tuned using ASR transducer objective in their corresponding modes. Both models are trained until the convergence of Word Error Rate (WER) and the resulting fine-tuned encoders serve as single-mode baselines. The 200M baseline encoders are architecturally similar to the DuRep-200M encoder, with the exception that the non-streaming baseline encoder utilizes positional encoding and non-causal convolutions. During BestRQ pre-training, we mask 320ms of audio with a starting probability of 0.02 for the streaming baseline, and mask 400ms of audio with the same starting probability for non-streaming baseline.

### 3.4. Evaluation Setup

SUPERB [8, 22] is a benchmarking framework designed to evaluate speech encoders and is used for relative comparison of encoders. During SUPERB evaluation, the candidate encoder is frozen, and a lightweight decoder (CTC for ASR and MLP for classification tasks) is trained. We evaluate our proposed encoders and all open-source encoders using an internal implementation of this framework that closely resembles the open-source SUPERB implementation. Major differences in the setup include: 1) we use a 4000 vocabulary wordpiece tokenizer for monolingual ASR and 15000 vocabulary wordpiece tokenizer for multilingual ASR tasks respectively instead of character based tokenizer, 2) instead of beam search we use greedy search for simplicity while evaluating CTC models and 3) no external LM is used in any evaluation. Due to these variations, absolute WERs reported after our SUPERB benchmarking are higher than those reported in the original works. But since the setup is kept constant across all runs, we can still use the WER numbers for relative comparisons across encoder models. We evaluate the encoders on Monolingual ASR (MonoASR), Multilingual ASR (MulASR), Language Identification (LID) and Emotion Recognition (ER) tasks. For simplicity, we report results on only ASR tasks for our 200M ablations.

MonoASR setup uses LibriSpeech corpus while MultiASR setup uses Multilingual LibriSpeech, CommonVoice, PeopleSpeech and VoxPopuli datasets during training and uses the test partitions of en, de, es, fr, it and pt of CommonVoice corpus for evaluation. For LID and ER, we use FLEURS and MSP Podcast [23] datasets respectively. Encoders are evaluated in Full-Context (LB, LA = ( $\infty, \infty$ )), Streaming (LB, LA = (5.4s, 0s)) and partial Look-Ahead (LB, LA = (5.4s, \*)) modes.

## 4. DuRep-200M

### 4.1. Results

Table 1 presents a comparative analysis of the dual-mode encoder and single-mode baseline encoders across full-context and streaming scenarios, in ASR tasks. In the full-context setting, Dual-mode encoder outperformed the baseline encoder, achieving an average improvement of 9.78% and 11.6% across monolingual and multilingual testsets. When evaluated in streaming mode, Dual-mode encoder shows an average perfor-

Table 1: Comparing DuRep-200M with single-mode baseline encoders using SUPERB on Monlingual and Multilingual ASR

Encoders	Inference Setting		MonoASR		MulASR
	LB	LA	test-clean	test-other	test
Dualmode (E4)	$\infty$	$\infty$	<b>4.61</b>	<b>10.16</b>	<b>21.25</b>
	5.4s	1s	5.12	11.30	23.35
	5.4s	0.6s	5.25	11.91	24.43
	5.4s	0s	<b>5.81</b>	<b>13.63</b>	<b>27.31</b>
Full-Context Baseline	$\infty$	$\infty$	5.07	11.36	24.05
Streaming Baseline	5.4s	0s	6.75	15.70	31.06

Table 2: Comparing ASR performance of DuRep-200M encoders obtained using different pre-training strategies

Pre-Training Strategy	Inference Setting		MonoASR		MulASR
	LB	LA	test-clean	test-other	test
BestRQ DM	$\infty$	$\infty$	6.024	13.13	25.29
Training	5.4s	0s	6.851	15.76	31.69
DM Distillation	$\infty$	$\infty$	5.49	11.80	23.18
from E1	5.4s	0s	6.82	15.06	29.02
DM Distillation	$\infty$	$\infty$	<b>4.61</b>	<b>10.16</b>	<b>21.25</b>
from E2	5.4s	0s	<b>5.81</b>	<b>13.63</b>	<b>27.31</b>

Table 3: Effect of context’s distribution during distillation (S3) on 200M dual-mode student encoder’s (E3) performance

Sample Space Configurations	Inference Setting (seconds)		MonoASR		MulASR
	LB	LA	test-clean	test-other	test
<b>T1</b>	$\infty$	$\infty$	6.02	13.42	26.75
	5.4	1	6.39	14.54	29.38
	5.4	0.6	6.64	15.12	30.28
	5.4	0	7.65	17.36	34.56
<b>T2</b>	$\infty$	$\infty$	<b>5.88</b>	<b>13.20</b>	<b>26.57</b>
	5.4	1	6.77	15.61	30.91
	5.4	0.6	6.99	16.07	31.47
	5.4	0	<b>7.33</b>	<b>16.63</b>	<b>33.6</b>
<b>T3</b>	$\infty$	$\infty$	6.14	13.45	26.63
	5.4	1	6.38	14.85	29.14
	5.4	0.6	6.64	15.03	30.06
	5.4	0	7.62	17.4	34.87

mance gain of 15.14% and 12% on monolingual and multilingual testsets.

### 4.2. Importance of 4-stage framework

We conducted a series of experiments to understand the contribution of the first three stages in our proposed four-stage framework. In the first experiment, we pretrained a dual-mode encoder using the BestRQ framework. In the second experiment, we pre-trained a dual-mode encoder through dual-mode distillation from the unsupervised full-context encoder, E1. The third dual-mode encoder (E3), was prepared by distilling from the supervised full-context encoder (E2). All 3 encoders underwent dual-mode transducer finetuning (S4).

Results (Table 2), indicate that E4 outperforms the other encoders in both MonoASR and MultiASR tasks in both inference settings. This finding underscores the importance of dis-

Table 4: Comparison of DuRep-2B with open source models on ASR and Non-ASR tasks evaluated using SUPERB

Encoder (Inference Mode)	MonoASR		MulASR	Emotion	Language Id
	test-clean(WER)	test-other(WER)	test(WER)	MSP(Acc)	Flours(Acc)
DuRep-2B (FC)	<b>3.32</b>	<b>6.46</b>	<b>10.89</b>	<b>53.39</b>	<b>90.83</b>
DuRep-2B (S)	4.32	9.27	15.51	52.73	88.74
Whisper-large-v3-encoder (FC)	4.02	8.19	14.84	49.34	90.33
Wav2Vec2 1B (FC)	6.05	11.97	35.24	49.66	36.85
Wav2Vec2 2B (FC)	5.49	10.73	33.63	50.11	41.18
mHUBERT-147 95M (FC)	5.22	10.77	27.85	50.31	87.9
WavLM 316M (FC)	9.12	19.09	45.74	52.95	69.52

tillation from a supervised full-context encoder in producing a high-quality pre-trained dual-mode encoder, which, when further fine-tuned, led to a superior dual-mode encoder.

#### 4.3. Role of context’s sampling space during distillation

Table 3 summarizes SUPERB-ASR results for three different past and future context sampling configurations used during dual-mode distillation (**S3**); **T1** = ( $L_{past}=[\infty, 5.4, 4.6, 3.6]$ ,  $L_{future}=[0, 1, 1.8, \infty]$ ), **T2** = ( $L_{past}=[\infty, 5.4]$ ,  $L_{future}=[0, \infty]$ ), **T3** = ( $L_{past}=[\infty, 5.8, 5.6, \dots, 3.6]$ ,  $L_{future}=[0, 0.2, \dots, 1.8, \infty]$ ). The context values were chosen based on our internal encoder usage scenarios. T1 represents the ablation’s base configuration. Restricting the sample space of encoder’s context to either causal or full-context modes during training (T2) improved inference performance in full-context and causal modes by an average of 1.55% and 3.72%. However, it significantly degraded inference performance in partial look-ahead scenarios (by 6.17% and 5.16%). Expanding the sample space of encoder’s context (T3) did not yield noticeable improvements for any inference scenarios. We attribute this to the inherent diversity of audio duration in the training data, which implicitly facilitates the model to attend to different contexts. Similar findings were observed during the dual-mode finetuning phase (**S4**). As we aim to train an encoder that has reliable performance across deployment scenarios, we selected the base configuration (T1) for variable attention masking in S3 and S4.

## 5. DuRep-2B

### 5.1. Results

Table 4 compares the performance of the DuRep-2B encoder in two inference modes with open-source encoders on ASR and non-ASR tasks, evaluated using SUPERB framework (3.4). In full-context mode, the DuRep-2B encoder outperforms all encoders [24–26] across all tasks. It exceeds the performance of the whisper-large-v3 encoder by 19.26% and 26.61% on MonoASR and MulASR. Notably, the DuRep-2B encoder operating in purely causal mode, outperforms most open-source encoders operating in full-context mode across tasks.

### 5.2. Conformer Layer Analysis Using SUPERB

We conducted layer-wise analysis on encoder E2 and E4 using the SUPERB benchmarking setup to investigate the variations in acoustic and semantic information across the encoder layers. Unlike the standard SUPERB setup, which utilizes embeddings from all layers of the encoder, this analysis focused on embeddings extracted from individual encoder layers. We analyzed the encoder’s semantic capabilities using Monolingual ASR and Speaking Rate Prediction, while Speech Intensity Contour Pre-

diction and Pitch Contour Prediction were used to assess its acoustic capabilities using Dynamic Time Warping (DTW) with correlation distance measure.

Figure 3 presents the development evaluation results for semantic and acoustic tasks across the layers of DuRep-2B. Our findings show that embeddings from the deeper Conformer blocks yield the best performance for semantically-oriented tasks like ASR and Speaking Rate Prediction. In contrast, for tasks that capture acoustic properties like Intensity and Pitch Contour Prediction, embeddings from the initial blocks of the encoder achieve optimal performance. [27] conducted a similar study on prosody reconstruction and related tasks using various self-supervised learning models, showing that dominant prosodic information is mainly present in the initial layers.

Interestingly, our results reveal a unique trend for semantic tasks. The final Conformer block’s embeddings provide the best performance for encoder E4, whereas embeddings from slightly earlier block yields the highest performance for encoder E2. We attribute the shift in peak performance towards deeper layers to the ASR-aware distillation stage (S3), which likely enhances model capacity utilization for semantic tasks.




Figure 3: SUPERB style evaluation performed on semantic and acoustic tasks for 2B scaled versions of encoders E2 and E4

## 6. Conclusions

We introduce DuRep, a dual-mode speech representation learning approach using variable attention based distillation, to train a general-purpose speech encoder that is optimized for multiple look-ahead inference scenarios. In ASR-SUPERB evaluations, on average, the DuRep-200M encoder outperforms baselines by 13.06% in streaming and 10.42% in non-streaming mode. The scaled-up DuRep-2B encoder shows average 21.71% improvement over Whisper-v3 encoder in non-streaming mode. Despite not being fine-tuned for non-ASR tasks, DuRep-2B encoder outperforms open-source encoders across ASR and non-ASR tasks. Our analysis highlights the trade-off between semantic and acoustic information across encoder layers.

## 7. Acknowledgements

We would additionally like to acknowledge Jahn Heymann, Ankish Bansal, Harish Mallidi, Phani Nidadavolu, Milind Rao, Bharat Padi, Venkata Kishore Nandury, Tuan Dinh, and Andreas Schwarz for their valuable contribution to data preparation efforts.

## 8. References

- [1] Y. Zhang, W. Han, J. Qin, Y. Wang, A. Bapna, Z. Chen, N. Chen, B. Li, V. Axelrod, G. Wang, Z. Meng, K. Hu, A. Rosenberg, R. Prabhavalkar, D. S. Park, P. Haghighi, J. Riesa, G. Perng, H. Soltau, T. Strohman, B. Ramabhadran, T. Sainath, P. Moreno, C.-C. Chiu, J. Schalkwyk, F. Beaufays, and Y. Wu, "Google usm: Scaling automatic speech recognition beyond 100 languages," 2023. [Online]. Available: <https://arxiv.org/abs/2303.01037>
- [2] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," 2022. [Online]. Available: <https://arxiv.org/abs/2212.04356>
- [3] J. Yu, W. Han, A. Gulati, C.-C. Chiu, B. Li, T. N. Sainath, Y. Wu, and R. Pang, "Dual-mode asr: Unify and improve streaming asr with full-context modeling," 2021. [Online]. Available: <https://arxiv.org/abs/2010.06030>
- [4] F. Weninger, M. Gaudesi, M. A. Haidar, N. Ferri, J. Andrés-Ferrer, and P. Zhan, "Conformer with dual-mode chunked attention for joint online and offline asr," 2022. [Online]. Available: <https://arxiv.org/abs/2206.11157>
- [5] K. Kim, F. Wu, P. Sridhar, K. J. Han, and S. Watanabe, "Multi-mode transformer transducer with stochastic future context," 2021. [Online]. Available: <https://arxiv.org/abs/2106.09760>
- [6] P. Swietojanski, S. Braun, D. Can, T. F. da Silva, A. Ghoshal, T. Hori, R. Hsiao, H. Mason, E. McDermott, H. Silovsky, R. Travadi, and X. Zhuang, "Variable attention masking for configurable transformer transducer speech recognition," 2023. [Online]. Available: <https://arxiv.org/abs/2211.01438>
- [7] A. Graves, "Sequence transduction with recurrent neural networks," 2012. [Online]. Available: <https://arxiv.org/abs/1211.3711>
- [8] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin, T.-H. Huang, W.-C. Tseng, K. tik Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. Mohamed, and H. yi Lee, "Superb: Speech processing universal performance benchmark," 2021. [Online]. Available: <https://arxiv.org/abs/2105.01051>
- [9] C.-C. Chiu, J. Qin, Y. Zhang, J. Yu, and Y. Wu, "Self-supervised learning with random-projection quantizer for speech recognition," 2022. [Online]. Available: <https://arxiv.org/abs/2202.01855>
- [10] AssemblyAI, "Universal-1: Robust and accurate multilingual speech-to-text," 2024. [Online]. Available: <https://www.assemblyai.com/research/universal-1>
- [11] A. Mohamed, H.-y. Lee, L. Borgholt, J. D. Havtorn, J. Edin, C. Igel, K. Kirchhoff, S.-W. Li, K. Livescu, L. Maaløe, T. N. Sainath, and S. Watanabe, "Self-supervised speech representation learning: A review," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, p. 1179–1210, Oct. 2022. [Online]. Available: <http://dx.doi.org/10.1109/JSTSP.2022.3207050>
- [12] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [13] V. Pratap *et al.*, "Mls: A large-scale multilingual dataset for speech research," in *Interspeech 2020*. ISCA, Oct. 2020. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2020-2826>
- [14] R. Ardila *et al.*, "Common voice: A massively-multilingual speech corpus," 2020. [Online]. Available: <https://arxiv.org/abs/1912.06670>
- [15] D. Galvez *et al.*, "The people's speech: A large-scale diverse english speech recognition dataset for commercial usage," 2021. [Online]. Available: <https://arxiv.org/abs/2111.09344>
- [16] A. Conneau *et al.*, "Fleurs: Few-shot learning evaluation of universal representations of speech," *2022 IEEE Spoken Language Technology Workshop (SLT)*, pp. 798–805, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:249062909>
- [17] C. Wang *et al.*, "Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," 2021. [Online]. Available: <https://arxiv.org/abs/2101.00390>
- [18] LibriVox, "free public domain audiobooks," 2024. [Online]. Available: <https://librivox.org/>
- [19] A. Gulati, J. Qin, C.-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," 2020. [Online]. Available: <https://arxiv.org/abs/2005.08100>
- [20] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [21] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Interspeech*, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:121321299>
- [22] J. Shi, D. Berrebbi, W. Chen, H.-L. Chung, E.-P. Hu, W. P. Huang, X. Chang, S.-W. Li, A. Mohamed, H. yi Lee, and S. Watanabe, "MI-superb: Multilingual speech universal performance benchmark," 2023. [Online]. Available: <https://arxiv.org/abs/2305.10615>
- [23] "Msp-podcast corpus: A large naturalistic speech emotional dataset," <https://ecs.utdallas.edu/research/researchlabs/msp-lab/MSP-Podcast.html>, 2024, accessed: 29-08-2024.
- [24] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," 2020. [Online]. Available: <https://arxiv.org/abs/2006.11477>
- [25] M. Z. Boito, V. Iyer, N. Lagos, L. Besacier, and I. Calapodescu, "mhubert-147: A compact multilingual hubert model," 2024. [Online]. Available: <https://arxiv.org/abs/2406.06371>
- [26] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, p. 1505–1518, Oct. 2022. [Online]. Available: <http://dx.doi.org/10.1109/JSTSP.2022.3188113>
- [27] G.-T. Lin, C.-L. Feng, W.-P. Huang, Y. Tseng, T.-H. Lin, C.-A. Li, H.-y. Lee, and N. G. Ward, "On the utility of self-supervised models for prosody-related tasks," in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 1104–1111.