



# Speech Self-Supervised Representation Benchmarking: Are We Doing it Right?

Salah Zaiem<sup>1</sup>, Youcef Kemiche<sup>2,3</sup>, Titouan Parcollet<sup>4,5</sup>, Slim Essid<sup>1</sup>, Mirco Ravanelli<sup>6</sup>

<sup>1</sup>LTCI, Télécom Paris, Institut Polytechnique de Paris, France

<sup>2</sup>Hi! PARIS Engineering Team, France

<sup>3</sup>Capgemini, France

<sup>4</sup>Samsung AI Center, Cambridge, United-Kingdom

<sup>5</sup>University of Cambridge, United-Kingdom

<sup>6</sup>Mila-Quebec AI Institute, Université de Montréal, Concordia University, Canada

salah.zaiem@telecom-paris.fr

## Abstract

Self-supervised learning (SSL) has recently allowed leveraging large datasets of unlabeled speech signals to reach impressive performance on speech tasks using only small amounts of annotated data. The high number of proposed approaches fostered the need and rise of extended benchmarks that evaluate their performance on a set of downstream tasks exploring various aspects of the speech signal. However, and while the number of considered tasks has been growing, most rely upon a single decoding architecture that maps the frozen SSL representations to the downstream labels. This work investigates the robustness of such benchmarking results to changes in the decoder architecture. Interestingly, it appears that varying the architecture of the downstream decoder leads to significant variations in the leaderboards of most tasks. Concerningly, our study reveals that benchmarking using limited decoders may cause a counterproductive increase in the sizes of the developed SSL models.

**Index Terms:** self-supervised learning, representation learning

## 1. Introduction

Self-supervised learning (SSL) provides a powerful means of benefiting from large volumes of unlabeled data to achieve substantial performance gains across a range of downstream tasks, without relying on manual annotations. Several approaches have been proposed in the literature, including predictive coding [1], multi-task learning [2], and contrastive learning approaches [3, 4]. Self-supervised models have lately become a necessary tool for speech practitioners suffering from scarce annotations.

However, experimenting with large SSL models is a costly endeavor both in terms of time and computing. Thus, the proliferation of approaches for speech SSL [5] has highlighted the need for benchmarks that evaluate their performance across multiple downstream tasks. Ultimately, such benchmarks explore various aspects of the speech signal, helping practitioners to make informed choices adapted to their use cases. They also serve as a simple way to assess the effectiveness of different techniques and identify areas for improvement. As a result, an increasing number of extended benchmarks have emerged in recent years, providing standardized frameworks for evaluating the performance of speech SSL models and algorithms [6, 7, 8]. These benchmarks cover a large range of speech tasks and applications and, even among a single task such as automatic speech recognition (ASR), a large choice of linguistic, acoustic and prosodic settings [9].

In popular speech SSL benchmarks, for every considered task, the self-supervised representation is probed using a learned downstream decoder mapping the frozen representation

to the final labels. Simplicity and limited capacities are the main drivers of the choice of these downstream probes with, for instance, linear probing for classification tasks or shallow vanilla recurrent neural networks for speech recognition. We hypothesize that such a benchmarking policy may harm the development of novel SSL technologies in at least two ways. First, the popularity of the main benchmarks, *e.g.* SUPERB [6], has made the considered downstream probes the *de facto* evaluation setting for any new speech SSL model. This is also true for the metrics used in the benchmark, which will condition the optimization process during the development of these new approaches, possibly leading to decisions to drop models that perform poorly with the selected probes but may perform better with other downstream architectures. Second, as the simplicity of the probes contrasts with the growing complexity of the SSL encoders, testing with low-capacity probes may imply a transfer of complexity from the decoder (expected to be task-specific) to the encoder (supposedly more general), leading to superfluously large self-supervised parts. As an example, in computer vision, Dubois *et al.* [10] have shown that changing the probe family from linear to multi-layer perceptrons (MLP) leads to different optimal choices in the hyperparameters of the SSL models and enables smaller SSL representations.

A first solution to such limitations might come from considering novel downstream-agnostic evaluations [11], through intrinsic quality assessment metrics for speech embeddings [12], but the correlation of these metrics with downstream performances has not been clearly identified [13] yet. In the context of image classification, Garrido *et al.* [14] have indicated that the rank of the vision SSL representations correlates highly with the final downstream performance. However, this study has been conducted with linear downstream probes only. Conscious of these issues, SUPERB [6], the main speech SSL benchmark, proposes two tracks where the downstream probes can be chosen by the submitter, with or without capacity constraints on the probes but these two tracks remain empty of submissions.

Therefore, this work evaluates the robustness of speech SSL benchmarks to variations in the probing architecture. The contributions are three-fold: i) We benchmark a set of published state-of-the-art SSL models on various speech tasks, varying the downstream decoders (Section 2); ii) We show that, except for ASR on LibriSpeech, the rankings and relative performances of the models are severely impacted by a change of decoder (Section 3); iii) We release the code<sup>1</sup> for the benchmark, implemented in SpeechBrain [15], for replication and further research on SSL models and downstream tasks.

<sup>1</sup>[github.com/salah-zaiem/speechbrain-2/tree/develop/recipes/SSL\\_benchmark](https://github.com/salah-zaiem/speechbrain-2/tree/develop/recipes/SSL_benchmark)

## 2. Definition and Experimental Protocol

This section formally describes the limitation faced by current speech SSL benchmarks and also details the experiment protocol illustrating this issue.

### 2.1. Problem definition

Formally, a SSL pipeline consists of two models: a pre-trained encoder  $\phi$  and a downstream probe  $f$ .  $\phi$  is learned through solving a pretext task on unlabeled speech datasets (Libri-light [16] and LibriSpeech960 [17] have been popular choices in the literature), while  $f$  is learned on the annotated downstream dataset. In this framework, the SUPERB benchmark has chosen for every considered downstream task  $T$  a probing family  $\mathfrak{F}_T$  (i.e. a downstream architecture) and shows for every considered SSL encoder  $\phi$  a task error rate corresponding to:

$$\min_{f \in \mathfrak{F}_T} E_t(f \circ \phi) \quad (1)$$

with  $E_t(f \circ \phi)$  the test-set error rate of the full SSL pipeline.

However, proper benchmarking would rank the candidate models according to the minimum quantity defined above over all the possible downstream probing families, possibly with a model capacity constraint. This measure is computationally intractable, as it would require training a model with every possible known downstream architecture respecting the capacity constraints for every considered encoder and task. In this work, we will explore whether benchmarking according to the value obtained in Equation (1) leads to a proper and robust ranking. To this effect, for every considered downstream task, we will survey two (or more) probing families, and see whether the rankings and relative differences obtained using the first families are kept in the second experiment.

### 2.2. Self-supervised pretrained models

To perform this study, a subset of the models presented in the SUPERB benchmark is considered. In particular, we retrieved the best-performing and easiest-to-use SSL models. Hence, 9 models acting directly on the raw waveform have been selected: Wav2vec 2.0 [1], HuBERT [18], WavLM<sup>2</sup>[19], and Data2vec [20] both in their Base and Large versions. DistilHuBERT [21], a distilled version of Hubert Base with 4 times less transformer layers, is also added to the list. The considered models share the same output frequency generating a representation of size  $D$  every 20 ms of audio signal, with  $D = 1,024$  for Large versions and  $D = 768$  for Base ones and DistilHuBERT.

All the models considered are based on very similar Transformer-based architectures, but their pretraining pretext tasks differ. Wav2vec2.0 training is based on the contrastive predictive coding [22] (CPC) objective, maximizing the mutual information between a set of context features and predicted future samples. HuBERT and WavLM learn to map the unlabeled audio to sequences of pseudo-labels obtained through the clustering of previous representations. WavLM introduces distortions during the training to enforce noise-related invariances. Inspired by teacher-student approaches, Data2vec is trained by utilizing a masked input view to predict latent representations of the complete input data, in a self-distillation configuration. All the pre-trained checkpoints are obtained from their Hugging-Face (HF) official cards [23], with the exception of Wav2vec2.0

<sup>2</sup>We considered WavLM Base+, trained on 94k hours of speech data, for the Base WavLM version.

Large whose Fairseq [24] checkpoint is used, as the HF one was underperforming compared to the results reported in SUPERB.

### 2.3. Downstream Tasks and Datasets

Speech SSL benchmarks have been attempting to evaluate universal speech representations by providing a large set of different tasks that assess different aspects of the speech signal. Similarly, we propose 7 tasks that tackle the phonetic, speaker-identity-related, emotive and semantic aspects.

**Speech Recognition Tasks.** Four speech recognition tasks are considered in this work. For the two first ones, LibriSpeech [17] *train-clean-100/dev-clean* subsets are used for training and validation while *test-clean* and *test-other* are used for testing. To test the ability of the models with fewer data and in a spontaneous context, the Buckeye dataset [25] is considered. The training, validation, and test splits used in our Buckeye experiments are available in the repository with the training set containing approximately 9.5 hours of audio and the test set 1.5 hours. For these two English ASR tasks, two results are shown depending on whether a language model (LM) is used or not during decoding. In the experiments labeled as *Without LM*, greedy decoding is applied. For the “With LM” experiments, the LibriSpeech official 4-gram language model is utilized. The language model is combined with shallow fusion to the acoustic model. Since low-resource languages are one of the main applications of SSL methods, two low-resource languages tasks, extracted from the CommonVoice 11.0 [26] release, are considered: Welsh (Cymraeg) and Basque (Euskera). The Word Error Rate (WER) is used as the error metric for all the ASR tasks. In all the ASR experiments, the probe is trained using the Connectionist Temporal Classification[27] loss at the character level.

**Automatic Speaker Verification.** Automatic Speaker Verification (ASV) involves a binary classification process to determine if the speakers in a pair of utterances are the same. As in the SUPERB benchmark, we use VoxCeleb1 [28] train and test splits for this task. It is important to note that the speakers in the testing set may not have been included in the training set. The evaluation metric is equal error rate (EER).

**Emotion Recognition (ER).** IEMOCAP [29], a dataset containing 10,039 from 10 speakers, is used for Emotion Recognition. The task consists of predicting an emotion class for the utterance among 4 candidates (neutral, happy, sad, angry). The performance shown is the mean of 10 runs done cross-validating on ten folds leaving each time one of the speakers’ data for testing.

**Intent Classification (IC).** Instead of Speech Commands (SC) [30], which is the dataset used in the SUPERB benchmarks for IC, we use the SLURP dataset [31], as the error rates with SC are already very low and SLURP is a more challenging task. The SLURP collection consists of approximately 72,000 audio recordings that capture single-turn user interactions with a home assistant. The IC task consists of classifying every utterance into one of the 18 considered scenarios. Examples of scenarios include “calendar”, “email” and “alarm”. The metric for emotion recognition and intent classification is accuracy.

### 2.4. Downstream Probes

This section provides a high-level description of the used downstream probes. The numerous hyperparameters and architectural details, allowing thorough replication of the experiments, are available in the provided repository.

Table 1: SSL benchmarking results for all tasks and downstream architectures. On most downstream tasks, rankings and relative performances are severely impacted by the change of the downstream architecture. The number of parameters of the SSL encoder and the probes is shown in the “Params” rows and columns.

Models /Tasks	SSL Params.	LibriSpeech train-100 ASR				Buckeye ASR		Welsh	Basque	ASV	ER	IC
Evaluation Metrics		WER ↓				WER ↓		WER ↓	WER ↓	EER ↓	Acc. ↑	Acc. ↑
First downstream architectures		LSTM				LSTM		LSTM	LSTM	Xvectors	Pool + Lin.	Pool + Lin.
		Clean	Other	Clean LM	Other LM	w/o LM	with LM	Welsh	Basque	ASV	ER	IC
DistilHuBERT	23.5M	13.99	34.91	9.96	28.26	35.59	28.29	53.20	46.78	9.1	65	46.6
Wav2vec 2.0 Base	95M	6.23	14.93	4.86	11.97	24.87	19.48	54.45	51.21	5.29	66.4	59.0
Wav2vec 2.0 Large	317.4M	3.72	9.25	3.13	7.48	<b>20.72</b>	16.11	45.42	<b>37.98</b>	5.69	69.3	66
HuBERT Base	94.7M	6.24	15.03	5.03	12.31	45.53	26.51	52.92	46.91	4.50	67.5	53.8
HuBERT Large	316.6M	3.57	8.12	2.90	6.59	51.30	33.10	51.21	46.15	5.20	71.3	69.9
WavLM Base+	94.7M	5.96	14.33	4.84	11.72	42.21	24.41	51.31	46.40	3.74	67.1	57.9
WavLM Large	316.6M	3.48	7.37	2.87	5.96	27.31	<b>14.27</b>	48.92	41.89	<b>2.98</b>	<b>75.3</b>	<b>78.8</b>
Data2vec Base	93.8M	5.30	13.79	4.03	10.97	37.26	30.50	54.00	46.37	5.43	63.0	56.9
Data2vec Large	314.3M	<b>3.10</b>	<b>6.50</b>	<b>2.58</b>	<b>5.38</b>	22.63	18.63	<b>44.32</b>	38.23	4.89	64.1	69.8
<b>Probe size and inference metrics</b>												
Downstream Parameters Base		39.9M				39.9M		40.3M	40.3M	7.0M	13.8k	3.1k
Downstream Parameters Large		42M				42M		42.4M	42.4M	7.7M	18.4k	4.1k
Second downstream architectures		Conformer				ContextNet		Lin.	Lin	ECAPA	ECAPA	LSTM + Lin.
		Clean	Other	Clean LM	Other LM	w/o LM	with LM	Welsh	Basque	ASV	ER	IC
DistilHuBERT	23.5M	14.97	36.51	11.54	31.41	58.56	43.61	80.78	77.04	2.85	72.4	74.9
Wav2vec 2.0 Base	95M	6.91	15.39	5.09	12.29	30.04	23.04	74.31	71.76	2.82	73.2	77.7
Wav2vec 2.0 Large	317.4M	4.32	9.25	3.58	7.03	23.92	18.68	75.45	78.48	3.17	68.4	79.0
HuBERT Base	94.7M	6.88	15.68	5.23	12.63	30.44	23.11	77.39	73.40	2.40	<b>78.2</b>	79.4
HuBERT Large	316.6M	3.96	8.60	<b>3.10</b>	6.88	39.39	31.57	71.58	60.24	3.84	71.5	80.1
WavLM Base+	94.7M	6.55	14.93	4.98	11.80	27.73	21.69	75.87	69.43	<b>1.76</b>	72.6	81.2
WavLM Large	316.6M	4.08	8.10	3.13	<b>6.31</b>	<b>15.61</b>	<b>12.1</b>	<b>68.73</b>	<b>56.32</b>	1.77	77.4	<b>85.8</b>
Data2vec Base	93.8M	5.85	14.32	4.53	12.52	40.53	33.45	77.49	75.26	3.75	72.0	73.4
Data2vec Large	314.3M	<b>3.43</b>	<b>6.82</b>	3.27	6.58	25.26	21.5	69.09	63.31	2.67	71.3	79.9
<b>Probe size and inference metrics</b>												
Downstream Parameters Base		11.2M				32.4M		1.9M	1.9M	9.2M	7.3M	42M
Downstream Parameters Large		11.2M				32.5M		2.3M	2.3M	9.8M	7.9M	44.1M

**Global settings.** During fine-tuning, the SSL encoder weights are frozen and only the downstream decoder weights are learned. Similarly to what is done in SUPERB, we observed that the last-layer representation may not always be optimal and, thus, collected all the hidden layers representations from the pretrained model. These hidden states are then weighted and summed together to create the representation fed to the decoder. The layer weights are trained during the fine-tuning. To check that the benchmark is properly set, we reproduced the SUPERB downstream architectures in the first set of experiments. Second, the probes are changed with simpler or more complex alternatives inspired from the relevant literature of every task.

**ASR tasks.** In the first set of experiments, reproducing the SUPERB conditions, a vanilla 2-layer 1, 024-units BiLSTM followed by a linear layer mapping the audio to the characters are used. For the second downstream architectures, we consider an encoder-decoder Conformer [32] downstream architecture for LibriSpeech, with 12 encoder layers, 4 decoder ones and 4 attention heads. For the Buckeye task, we use a ContextNet [33] with unitary strides to keep the frame rate of the SSL models. For Welsh and Basque, a two-layered dense neural network is used to map every frame representation to the probabilities of the characters. In a second time, experiments using ContextNet with LibriSpeech are also performed. ContextNet and Conformer performances, which are close to state-of-the-art on LibriSpeech, motivated their selection as downstream probes.

**Automatic Speaker Verification.** First, the X-vector [34] architecture is tested with an AM-Softmax [35] loss training of the speaker embeddings. For verification, we use the cosine similarity between speaker representations. In the second experiment, we replace the X-vectors with an ECAPA-TDNN neural network [36]. ECAPA-TDNN combines time-delay neural networks with parallel attention mechanisms to capture temporal dependencies and the long-range context in sequential data and

achieves state-of-the-art results in speaker verification.

**Classification tasks.** Similarly to SUPERB, we perform linear probing for classification tasks (*i.e.* intent classification and emotion recognition) in the first set of experiments. The representations are first average-pooled along the time axis before being fed to a classification linear layer. For the second downstream architecture, and inspired by state-of-the-art approaches [37], ECAPA-TDNN is chosen for emotion recognition. For intent classification, following published architectures [38], two layers of BiLSTM with hidden size 1, 024 followed by a linear classifier are used. In contrast to using time-pooled features, this allows for considering the order of frame representations.

### 3. Benchmarking Results and Discussion

Table 1 shows the full benchmarking results of the various SSL models. The upper and lower parts of the table display the obtained performance for the first and second sets of downstream architectures respectively. The number of neural parameters are also reported both for the SSL encoder and for the downstream decoders. For the latter, only two numbers are provided per task as it only depends on the dimension of the encoder output vector, *i.e.* “Base” or “Large”. In the first set of experiments, the settings of two tasks (LibriSpeech and ASV) are exactly similar to the SUPERB benchmark conditions: the Pearson correlation between our results and those shown in the SUPERB leaderboard reaches respectively 0.99 and 0.97, validating our reproduction of its experimental setting.

To study the impact of a decoder change on the final performances, we compute, for every task, the Pearson and Spearman correlations between the performances obtained with the first downstream architectures and those obtained with the second ones, and collect them in Table 2. The Pearson correlation evaluates the linear relationship between the two sets of per-

Table 2: Pearson and Spearman correlations between the performances obtained with first and second downstream probes, for every considered task. The “Diff” column shows the relative difference in performance between the two architectures.

Task	Pearson	Spearman	Mean DS1	Mean DS2	Diff (%)
LibriSpeech 1-2	0.99	0.97	5.8	6.48	-11.7
LibriSpeech 1-3	0.99	0.98	5.8	7.03	-21.2
Buckeye	0.42	0.56	34.16	32.39	5.2
Welsh	0.59	0.62	50.64	74.52	-47.2
Basque	0.19	0.15	44.66	69.47	-55.6
VoxCeleb	0.47	0.75	5.2	2.78	46.5
Iemocap	0.22	0.34	67.66	73	7.9
Slurp	0.75	0.66	62.1	79.04	27.3

performances, while the Spearman one assesses the strength and direction of their monotonic relationship. Correlation metrics close to 1 imply proportional performances and similar rankings between the SSL models used with different probes, making the benchmark robust to the considered downstream change.

All the models considered converge to competitive performances on every downstream task and with all the decoding architectures. With the notable exception of LibriSpeech, all the downstream tasks exhibit very different behavior when probes change. In Table 2, the last three columns show the mean performance of the SSL candidates with the first and second downstream decoders. We observe, first, high sensitivity to the chosen decoder, with relative improvements reaching 46.5% and 27.3% respectively for ASV and IC when decoding with the second set of probes. Second, and more interestingly towards answering our main question, the low correlation values (while positive) testify to the large variations in relative performances and rankings between the performances with DS1 and those with DS2. The Spearman correlation, for instance, only reaches 0.34 and 0.66 for ER and IC respectively. It is worth noting that while LibriSpeech performance assessment seems robust to changes in the decoder, this does not generalize to the other ASR tasks, starting with the spontaneous English Buckeye corpus with 0.56 Spearman and 0.42 Pearson correlation, or the Basque task with 0.19 Pearson and 0.15 Spearman correlations. The case of Buckeye ASR is particularly interesting as changing the decoder, from BiLSTM to ContextNet, improves the results for a few models and harms them for others. The best performing one, WavLM Large using the second decoder, is only the fourth best performing one with the SUPERB settings.

However, we observe that contrarily to the other downstream tasks, the rankings and performances of the considered SSL encoders on the ASR task using LibriSpeech *train-clean-100*, shown in Table 2 vary only very lightly with the change of downstream decoder. To cross-validate this finding, experiments using a third downstream decoder, ContextNet, are produced for this task. The results of this additional experiment are reported in Table 3. Similarly, no major differences are witnessed in the ranking of the SSL candidates. For instance, in the three setups, without LM decoding, Data2vec Large is always the best-performing representation, whether it is on the *test-clean* or the *test-other* split. Similarly, DistilHuBERT is systematically the worst performing one and the “Large” versions of the considered candidates perform consistently better than their “Base” counterparts. These results are confirmed in Table 2 where Spearman and Pearson correlations between the performances using the different downstream decoders are over 0.97 for LibriSpeech while the highest correlation value is 0.75 for the other tasks. With LibriSpeech being one of the main ASR benchmarking dataset, and as it is systematically part of the pretraining dataset, these observations suggest that the SSL encoders may be overfitting on this task. These results al-

Table 3: WER results of LibriSpeech experiments on the two considered test splits with a third downstream probe

Tasks \ Models	SSL Params	Clean	Other	Clean LM	Other LM
DistilHuBERT	23.5M	20.52	43.27	10.44	29.17
Wav2vec 2.0 Base	95M	7.24	15.66	4.73	11.21
Wav2vec 2.0 Large	317.4M	4.35	8.68	03.03	6.86
HuBERT Base	94.7M	7.31	16.00	4.60	11.11
HuBERT Large	316.6M	4.04	8.63	2.98	6.45
WavLM Base+	94.7M	6.73	15.33	4.52	10.84
WavLM Large	316.6M	4.09	8.43	2.94	6.15
Data2vec Base	93.8M	5.46	13.34	3.76	10.04
Data2vec Large	314.3M	<b>3.50</b>	<b>6.94</b>	<b>2.56</b>	<b>5.36</b>
Probe size and inference metrics					
Downstream Parameters Base				32.4M	
Downstream Parameters Large				32.5M	

low us to conclude that, except for LibriSpeech ASR, current SSL benchmarking is not robust to the choice of the downstream probe. Furthermore, the reported results highlight an interesting consequence of using low-capacity decoders. With the first set of downstream architectures, “Large” versions of SSL models are almost always better performing than their “Base” counterparts. However, this is not the case with higher-capacity decoders in the second set of probes. For instance, the best ASV and ER performances are reached, respectively, with WavLM Base+ and HuBERT Base. For intent classification, when changing the downstream decoder from linear to BiLSTMs, the mean absolute difference between the “Base” and “Large” versions’ performance drops from 14.23 to 3.28. For emotion recognition, while all four “Large” versions perform better than the “Base” ones when probed linearly, enhancing the capacity of the decoder reverses this order for all of them except WavLM. Furthermore, concerning the ASV results, DistilHuBERT performs better with an ECAPA decoder than the best model (WavLM Large) with an x-vector-based head, while containing more than 13 times fewer parameters. Since the number of parameters does not present a full picture of the involved computations, the THOP library<sup>3</sup> is used to compute the number of Multiply-accumulate operations (MACs). During inference, running DistilHuBERT-ECAPA on VoxCeleb1 enroll and test sets involves computations reaching a mean 249.4G MACs per batch of 8 samples, for a result of 2.85 EER, while WavLM Large-Xvectors requires a mean batch computation of 1.2T MACs to reach 2.98 EER. All these observations lead to the conclusion that keeping very small-capacity decoders may lead to an inflation in model sizes.

## 4. Acknowledgements

This work has benefited from funding from l’Agence de l’Innovation de Défense, and was performed using HPC resources from GENCI-IDRIS (Grant 2021-AD011012801R1).

## 5. Conclusion

In this work, various experiments have been described to evaluate the robustness of speech SSL benchmarks to changes in the downstream probes. The results obtained show a high sensitivity of the leaderboard rankings and relative performances to the choice of the downstream architectures. Furthermore, it is shown that selecting low-capacity decoders may lead to oversized SSL encoders. We hope this diagnosis will help the community to design new benchmarking approaches and encourage submissions to the SUPERB “Constrained” track.

<sup>3</sup>[github.com/Lyken17/pytorch-OpCounter](https://github.com/Lyken17/pytorch-OpCounter)

## 6. References

- [1] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Neurips*, 2020.
- [2] S. Zaiem, T. Parcollet, S. Essid, and A. Heba, “Pretext tasks selection for multitask self-supervised audio representation learning,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1439–1453, 2022.
- [3] A. Saeed, D. Grangier, and N. Zeghidour, “Contrastive learning of general-purpose audio representations,” 2020.
- [4] S. Zaiem, T. Parcollet, and S. Essid, “Automatic data augmentation selection and parametrization in contrastive self-supervised speech representation learning,” *Interspeech*, 2022.
- [5] A.-R. Mohamed, H. yi 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 JSTSP*, 2022.
- [6] S. 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.
- [7] T.-h. Feng, A. Dong, C.-F. Yeh, S.-w. Yang, T.-Q. Lin, J. Shi, K.-W. Chang, Z. Huang, H. Wu, X. Chang, S. Watanabe, A. Mohamed, S.-W. Li, and H.-y. Lee, “Superb @ slt 2022: Challenge on generalization and efficiency of self-supervised speech representation learning,” 01 2023, pp. 1096–1103.
- [8] S. Evain, H. Nguyen, H. Le, M. Z. Boito, S. Mdhaffar, S. Alisamir, Z. Tong, N. Tomashenko, M. Dinarelli, T. Parcollet, A. Al-lauzen, Y. Esteve, B. Lecouteux, F. Portet, S. Rossato, F. Ringeval, D. Schwab, and L. Besacier, “Lebenchmark: A reproducible framework for assessing self-supervised representation learning from speech,” 2021.
- [9] H.-S. Tsai, H.-J. Chang, W.-C. Huang, Z. Huang, K. Lakhotia, S.-w. Yang, S. Dong, A. Liu, C.-I. Lai, J. Shi, X. Chang, P. Hall, H.-J. Chen, S.-W. Li, S. Watanabe, A. Mohamed, and H.-y. Lee, “SUPERB-SG: Enhanced speech processing universal PERFORMANCE benchmark for semantic and generative capabilities.” Dublin, Ireland: ACL, 2022.
- [10] Y. Dubois, T. Hashimoto, S. Ermon, and P. Liang, “Improving self-supervised learning by characterizing idealized representations,” *arXiv preprint arXiv:2209.06235*, 2022.
- [11] D. Álvarez, S. Pascual, and A. Bonafonte, “Problem-agnostic speech embeddings for multi-speaker text-to-speech with sampler,” 10th ISCA Speech Synthesis Workshop, 2019.
- [12] T. Schatz, V. Peddinti, F. Bach, A. Jansen, H. Hermansky, and E. Dupoux, “Evaluating speech features with the Minimal-Pair ABX task: Analysis of the classical MFC/PLP pipeline,” in *INTERSPEECH 2013*, Lyon, France, Aug. 2013, pp. 1–5.
- [13] R. Algayres, M. S. Zaiem, B. Sagot, and E. Dupoux, “Evaluating the reliability of acoustic speech embeddings,” *INTERSPEECH*, 2020.
- [14] Q. Garrido, R. Balestrieri, L. Najman, and Y. Lecun, “Rankme: Assessing the downstream performance of pretrained self-supervised representations by their rank,” *arXiv preprint*, 2022.
- [15] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, “Speechbrain: A general-purpose speech toolkit,” 2021.
- [16] J. Kahn, M. Riviere, W. Zheng, E. Kharitonov, Q. Xu, P. Mazare, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, “Libri-light: A benchmark for ASR with limited or no supervision,” in *ICASSP*, 2020.
- [17] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *2015 (ICASSP)*, 2015, pp. 5206–5210.
- [18] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 29, oct 2021.
- [19] 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 JSTSP*, vol. 16, no. 6, pp. 1505–1518, oct 2021.
- [20] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, “Data2vec: A general framework for self-supervised learning in speech, vision and language,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 1298–1312.
- [21] H.-J. Chang, S.-w. Yang, and H.-y. Lee, “DistilHuBERT: Speech Representation Learning by Layer-wise Distillation of Hidden-unit BERT,” *ICASSP*, 2021.
- [22] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *Neurips* 2018.
- [23] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, “Transformers: State-of-the-art natural language processing,” in *EMNLP*, 2020.
- [24] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, “fairseq: A fast, extensible toolkit for sequence modeling,” *arXiv preprint arXiv:1904.01038*, 2019.
- [25] M. Pitt, K. Johnson, E. Hume, S. Kiesling, and W. Raymond, “The buckeye corpus of conversational speech: Labeling conventions and a test of transcriber reliability,” *Speech Communication*, vol. 45, pp. 89–95, 01 2005.
- [26] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, “Common voice: A massively-multilingual speech corpus,” 2020.
- [27] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” ser. ICML ’06, 2006.
- [28] A. Nagrani, J. S. Chung, and A. Zisserman, “Voxceleb: A large-scale speaker identification dataset,” *Interspeech 2017*, Aug 2017.
- [29] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, “Iemocap: Interactive emotional dyadic motion capture database,” *Language resources and evaluation*, vol. 42, pp. 335–359, 2008.
- [30] P. Warden, “Speech commands: A dataset for limited-vocabulary speech recognition,” *arXiv preprint arXiv:1804.03209*, 2018.
- [31] E. Bastianelli, A. Vanzo, P. Swietojanski, and V. Rieser, “Slurp: A spoken language understanding resource package,” *EMNLP 2020*.
- [32] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, “Conformer: Convolution-augmented transformer for speech recognition,” *Interspeech*, 2020.
- [33] W. Han, Z. Zhang, Y. Zhang, J. Yu, C.-C. Chiu, J. Qin, A. Gulati, R. Pang, and Y. Wu, “Contextnet: Improving convolutional neural networks for automatic speech recognition with global context,” *Interspeech*, 2020.
- [34] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust dnn embeddings for speaker recognition,” in *(ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [35] F. Wang, J. Cheng, W. Liu, and H. Liu, “Additive margin softmax for face verification,” *IEEE Signal Processing Letters*, 2018.
- [36] B. Desplanques, J. Thienpondt, and K. Demuynck, “Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” *Interspeech*, 2020.
- [37] Y. Wang, A. Boumadane, and A. Heba, “A fine-tuned wav2vec 2.0/hubert benchmark for speech emotion recognition, speaker verification and spoken language understanding,” *arXiv preprint arXiv:2111.02735*, 2021.
- [38] L. Lugosch, M. Ravanelli, P. Ignoto, V. S. Tomar, and Y. Bengio, “Speech model pre-training for end-to-end spoken language understanding,” *Interspeech*, 2019.