

THE EFFECT OF SPOKEN LANGUAGE ON SPEECH ENHANCEMENT USING SELF-SUPERVISED SPEECH REPRESENTATION LOSS FUNCTIONS

George Close, Thomas Hain, and Stefan Goetze

Speech and Hearing Group, Department of Computer Science, University of Sheffield, United Kingdom

ABSTRACT

Recent work in the field of speech enhancement (SE) has involved the use of self-supervised speech representations (SSSRs) as feature transformations in loss functions. However, in prior work, very little attention has been paid to the relationship between the language of the audio used to train the self-supervised representation and that used to train the SE system. Enhancement models trained using a loss function which incorporates a self-supervised representation that shares exactly the language of the noisy data used to train the SE system show better performance than those which do not match exactly. This may lead to enhancement systems which are language specific and as such do not generalise well to unseen languages, unlike models trained using traditional spectrogram or time domain loss functions. In this work, SE models are trained and tested on a number of different languages, with self-supervised representations which themselves are trained using different language combinations and with differing network structures as loss function representations. These models are then tested across unseen languages and their performances are analysed. It is found that the training language of the self-supervised representation appears to have a minor effect on enhancement performance, the amount of training data of a particular language, however, greatly affects performance.

Index Terms— Speech enhancement, self-supervised speech representations, language, domain adaption, neural networks

1. INTRODUCTION

Speech enhancement (SE) is a longstanding, active area of speech research given its myriad applications in downstream tasks [1–4]. With the change in working patterns globally due to the COVID-19 pandemic, online remote meetings have become a mainstay in the working world. SE systems which are able to remove various kinds of background noise while maintaining the quality and intelligibility of speech are required features of online meeting software. Deep learning approaches utilising neural networks have shown better performance versus traditional signal-processing-based methods [5–8]. Another related topic of research is the development of sophisticated ways of measuring the degradation in a noisy signal and thus the effectiveness of speech enhancement systems [9–13]. These measures (or ‘metrics’) are often perceptually motivated, and aim in some way to behave similarly to human ratings of quality; some recent works propose neural networks which, given audio as input directly, predict human mean opinion score (MOS) quality ratings of the audio [14–16].

This work was supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation [grant number EP/S023062/1]. This work was also funded in part by TOSHIBA Cambridge Research Laboratory

The use of self-supervised speech representations (SSSRs) has become increasingly popular in various speech-related tasks, including speech enhancement [17–19]. SSSRs are particularly advantageous due to their ability to predict speech content’s contextual information within the input audio, thereby capturing patterns in spoken language.

Recent work [20–22] has utilised SSSRs in loss functions for the training of neural speech enhancement systems. Loss functions which use distances between SSSR representations have shown to be more strongly correlated with perceptual measures of quality and with human MOS scores compared to standard short time Fourier transform (STFT) spectrogram-based loss functions [20]. However, these works focus mainly on the use of different intermediate representations from the SSSR as features for loss functions and are not concerned with the relationship between the data used to train the SSSR and the data used to train the enhancement system.

This work investigates the relationship between the spoken language in the training data of the SSSR and the noisy speech data used to train the enhancement system. Hence, several monolingual noisy speech datasets are created, differing only in the language of the clean speech which is artificially corrupted. All other dataset parameters are kept as similar as possible. Subsequently, several speech enhancement models are trained using these datasets with SSSR-derived loss functions where the SSSRs have been trained on different combinations of languages, as well as with standard non-SSSR loss functions. The performance of these models is evaluated across testsets for each language, assessing their ability to adapt to enhancing speech in languages not observed in training.

The remainder of this paper is structured as follows. In Section 2, the SSSRs used in this work are introduced. In Section 3, the simulation process to generate the monolingual noisy speech datasets is detailed. In Section 4, an experiment involving the training of neural speech enhancement systems using the proposed datasets is carried out, and performance analysed. Section 5 concludes the paper.

2. SELF SUPERVISED SPEECH REPRESENTATION MODELS

SSSR models are neural networks which output a *context* representation of the input speech audio waveform [17]. They are trained in a self-supervised manner using large corpora of speech data, typically with the objective to recreate some masked portion of the input audio. Structurally, they consist of two main stages. The first, denoted by the operator \mathcal{G}_{FE} in the following with subscript FE standing for *feature encoder*, is built from a number of 1D convolutional layers which convert the input time-domain speech signal $s[n]$ to a two-dimensional feature representation

$$\mathbf{S}_{\text{FE}} = \mathcal{G}_{\text{FE}}(s[n]), \quad (1)$$

with a feature dimension F (typically of size 512) and a time dimension T , i.e. the number of frames, which is dependent on the length of the input audio signal. The second stage, denoted by \mathcal{G}_{OL} , consists of a number of Transformer [23] layers and operates over the output of the feature encoder stage to result in the SSSR output

$$\mathbf{S}_{OL} = \mathcal{G}_{OL}(\mathcal{G}_{FE}(s[n])). \quad (2)$$

The output representation \mathbf{S}_{OL} shares the time dimension T with \mathbf{S}_{FE} but has a different, usually larger feature dimension F . Subscript OL stands for *output layer* of the SSSR.

Hidden Unit BERT (HuBERT) [18] is an SSSR model; during training it makes use of a BERT [24] inspired loss function. The output of its feature encoder \mathcal{G}_{FE} has a dimension of $F = 512$, while its final layer output after \mathcal{G}_{OL} has a feature dimension of 758. In this work, the HuBERT model used is trained on the 960 hour LibriSpeech [25] training set and is sourced from the fairseq Github repository¹. It is important to note that this dataset consists of English read speech only, so the model has only ever been exposed to English speech.

Multilingual HuBERT (mHuBERT) [26] is a variation on HuBERT which has been trained on multilingual speech data, specifically the English, French and Spanish language parts of the VoxPopuli [27] dataset, each containing 4.5k hours totalling 13.5k hours of speech. It has the same feature dimensions as HuBERT. It can be considered as a *middle point* between the monolingual HuBERT and the massively multilingual XLSR.

The **Cross-Lingual Speech Representation (XLSR)** [28] is an SSSR with the main distinguishing feature being that it is trained using audio containing a large number of languages. It is intended to act as a ‘universal’ model of speech, encoding latent speech representations which are shared across languages. It is trained on 436k hours of speech from 128 different languages from datasets including VoxPopuli, CommonVoice and BABEL², with the Wav2Vec2 [17] contrastive masking objective. Note that unlike the other two SSSRs used in this work, it is trained on potentially noisy data, notably CommonVoice (described below in Section 3.2) and BABEL which contains conversational telephone recordings. Its \mathcal{G}_{FE} representations have a feature dimension F of 512 while its \mathcal{G}_{OL} representations have an F of 1024.

3. COMMONVOICE-DEMAND: A MULTILINGUAL SPEECH ENHANCEMENT DATASET

This section details the creation process of the proposed *CommonVoice-DEMAND* speech enhancement dataset, which is intended to be a multilingual variation on the popular monolingual VoiceBank-DEMAND (VB-D) dataset.

3.1. VoiceBank-DEMAND Dataset

VoiceBank-DEMAND [29] is a widely used dataset for speech enhancement neural network training. It consists of clean English read speech, artificially corrupted with environmental noise from the DEMAND [30] noise dataset, as well as two additional noise types, speech-shaped noise (SSN) and babble. Note that the babble noise was created by randomly overlapping the clean speech audio. The clean speech files vary in length from around 3 to 10 seconds, while the DEMAND noise recordings are all 10 minutes long. The

noisy signal $x[n]$ is created by adding a random part of the noise recording $v[n]$ of the same length as the clean speech $s[n]$.

$$x[n] = s[n] + c \cdot v[n] \quad (3)$$

This scaling factor

$$c = \sqrt{\frac{P_s}{P_v \cdot 10^{\frac{SNR}{10}}}} \quad (4)$$

is computed using a given target mixing signal-to-noise ratio (SNR), and using the ITU-T P.56 method [31] for computing the active speaker power of the clean (speech) reference audio P_s and of the noise P_v .

The training set consists of 11572 pairs of clean and noisy speech, $s[n], x[n]$, from 28 different speakers (14 male, 14 female) with native British accents speaking English. The clean speech is mixed at 0, 5, 10 and 15 dB SNR with cafeteria, car, kitchen, meeting, metro, restaurant, station, and traffic noise from DEMAND as well as babble and speech-shaped noise. The test set consists of 824 $s[n], x[n]$ pairs from two additional speakers (one male, one female) who do not appear in the training set, mixed at 2.5, 7.5, 12.5 and 17.5 dB SNR with bus, cafe, living room, office and public square noise from DEMAND. All audio has a sample rate of 48000 Hz and is in WAV format, but the dataset is typically down-sampled to 16000 Hz for speech enhancement training.

3.2. CommonVoice Dataset

To create a multilingual dataset which is as similar as possible to the VB-D dataset described in Section 3.1, speech is sourced from the Mozilla CommonVoice [32] dataset, which consists of recordings of read speech in 108 languages, with corresponding text prompt sentences. The recordings are crowd-sourced using the CommonVoice website³. Validation that the recordings properly represent the prompt sentence is also crowd-sourced. In addition to the audio recording and prompt sentence text, some additional metadata is sometimes available such as gender and accent of the speaker. In this work, we make use of the English, Spanish, and Welsh portions of the dataset. These portions contain 3209, 2152 and 152 hours of audio, respectively.

3.3. Candidate Selection

Due to their crowd-sourced nature, the quality of the recordings in the CommonVoice dataset varies considerably. For the creation of the multilingual CommonVoice-DEMAND dataset, the aim is to select the cleanest possible audio for use as reference signals. Additionally, certain signal enhancement metrics require the input audio to have a minimum length. The process for the selection of candidate reference signals for a given language to create the CommonVoice-DEMAND datasets can be summarised as follows: Firstly, only recordings which have been validated by the crowd-sourced validation process are selected. This is to ensure that the audio does contain the prompt sentence and is not a failed recording or too noisy for the speech to be intelligible.

Secondly, recordings of less than 2 seconds length and those which contain a single-word utterance are excluded since it was found that some speech enhancement metrics have difficulties assessing such recordings.

Finally, the quality of the remaining audio recordings is assessed. A

¹<https://github.com/facebookresearch/fairseq>

²<https://catalog.ldc.upenn.edu/byyear>

³<https://commonvoice.mozilla.org>

voice activity detector (VAD) is used to segregate frames of length L of the signal $x[n]$ into disjoint sets \mathcal{A} and \mathcal{B} , for which the signal fulfils either the hypothesis that speech is present \mathcal{H}_1 , or that speech is absent \mathcal{H}_0 , respectively. An SNR estimate is obtained by

$$\widehat{\text{SNR}}(x[n]) = 10 \log_{10} \left(\frac{\frac{1}{|\mathcal{A}|} \sum_{\ell \in \mathcal{A}} \sum_{n=0}^{K-1} x^2[\ell K + n]}{\frac{1}{|\mathcal{B}|} \sum_{\ell \in \mathcal{B}} \sum_{n=0}^{K-1} x^2[\ell K + n]} \right) \quad (5)$$

for frame index ℓ with $|\mathcal{A}|$ and $|\mathcal{B}|$ denoting the respective cardinalities of the sets \mathcal{A} and \mathcal{B} . For simplicity, the Google WebRTC-VAD⁴ is used. A minimum threshold of 50 dB estimated SNR by the above formulation is used to select the candidate recordings. Note that while this is a somewhat crude estimator in that it does not account for the noise which is present in the speech-active frames in \mathcal{A} , it was found empirically that this approach works sufficiently well to select high-quality recordings containing little to no background noise, with a low computation overhead. For each language, 20000 of such candidate recordings are selected in this way. During this process, the candidate recordings are converted from the MP3 format to the WAV format.

3.4. Dataset Creation

In order to ensure that the proposed CommonVoice-DEMAND datasets are comparable to the original VoiceBank-DEMAND dataset, the log files describing VoiceBank-DEMAND are used. These consist of two lists (one for the training set, one for the test set) for clean audio file $s[n]$, the name of the noise file for $v[n]$, from which a random section of the same length as $s[n]$ is obtained and the desired mixing SNR value. The process for the creation of the CommonVoice-DEMAND datasets is as follows:

- A candidate CommonVoice recording is selected which is closest to the length in seconds to the clean audio recording from the original VoiceBank dataset. This candidate recording is either padded with zero values or truncated such that it is the same number of samples exactly as the original VoiceBank recording.
- The mixing process as described in Section 3.1 is carried out, using the selected CommonVoice recording as $s[n]$, and creating a corresponding noisy version $x[n]$ according to (3) and (4). Since CommonVoice data has a sampling rate of 32 kHz, the resulting sampling rate of CommonVoice-DEMAND is also 32 kHz and thus lower than VoiceBank-DEMAND's 48 kHz. Note, that for SE, data is usually anyhow downsampled to 16 kHz.
- The selected CommonVoice recording is then removed from the list of available candidate recordings, ensuring that uniqueness for each $s[n], x[n]$ pair in the resultant dataset.

The goal of this process is to create a dataset which is as similar as possible to the original VoiceBank-DEMAND dataset, but using clean speech with a different language or source. These datasets differ from the original VoiceBank-DEMAND in the (usually) greater number of speakers in the training and test sets and the exact portion of the noise file from which $v[n]$ is created due to the random sampling of the noise file. The created CommonVoice-DEMAND data ensures reproducibility by fixed seeds in the random number generator. The code for this process is available on GitHub⁵ to allow for the recreation and reuse of the generated corpora.

⁴<https://github.com/wiseman/py-webrtcvad>

⁵<https://github.com/letol9/CommonVoice-DEMAND>

4. SPEECH ENHANCEMENT EXPERIMENTS

4.1. Experiment Setup

Masking-based speech enhancement networks are trained using SSSR derived loss functions using the proposed CommonVoice-DEMAND datasets. The SpeechBrain [33] toolkit is used to facilitate the training and testing of the models. The models are trained for 50 epochs with the Adam [34] optimiser, with a learning rate of 0.001. At test time, the model with the highest Perceptual Evaluation of Speech Quality (PESQ) [9] score on the validation set is loaded and evaluated.

The model used for all training consists of two bidirectional long short-term memory (BLSTM) layers and two linear layers. The first linear layer has LeakyReLU activation and the second has Sigmoid activation. The model takes a spectrogram of the noisy audio $x[n]$ as input and produces a *mask* which is multiplied with the input spectrogram to generate an enhanced spectrogram. Using the original noisy phase of $x[n]$, the enhanced time domain signal $\hat{s}[n]$ is produced via the overlap-add resynthesis method. Despite its simplicity, this model's training is stable, efficient, and has demonstrated state-of-the-art performance [5, 35].

4.2. Datasets

CommonVoice-DEMAND training and test sets were generated using the process described in Section 3 for English and Spanish. These languages were chosen as they match languages used to train the mHuBERT model, and the CommonVoice corpus component for each is sufficiently large. A testset for Welsh was also created as a language which was not used to train HuBERT or mHuBERT but which was used as one of XLSRs 128 languages. The CommonVoice-DEMAND datasets have the same size training and test sets as the original VoiceBank-DEMAND with 11572 and 824 $s[n], x[n]$ pairs, respectively. A validation set of size 770 is created from each CommonVoice-DEMAND training set. All audio is at 16 kHz sample rate.

4.3. SSSR Signal Enhancement Loss Function

The SSSR loss function as defined in [20] is used, which is based on the mean squared error (MSE) distance between the output feature encoder representations of the enhanced signal $\hat{s}[n]$ and the reference signal $s[n]$. The SSSR loss function is given by

$$L_{\text{FE}}(\mathbf{S}_{\text{FE}}, \hat{\mathbf{S}}_{\text{FE}}) = \frac{1}{TF} \sum_{t,f} (\mathbf{S}_{\text{FE}}[t, f] - \hat{\mathbf{S}}_{\text{FE}}[t, f])^2, \quad (6)$$

where \mathbf{S}_{FE} and $\hat{\mathbf{S}}_{\text{FE}}$ are the feature encoder output representations of $s[n]$ and $\hat{s}[n]$, respectively. F and T denote the feature and time dimensions of the representations, with F being 512 for all models used in this work and T depending on the length in samples of the time domain audio. Models are trained using HuBERT, mHuBERT, and XLSR to obtain the representations in (6). In addition, the spectrogram MSE loss, short-time objective intelligibility (STOI) [10] loss, and scale-invariant signal-to-distortion ratio (SISDR) [11] loss are used as baselines. These baselines are popular loss functions for speech enhancement training and are language-independent. Spectrograms are created using a STFT with a FFT length of 512, a window length of 32 ms, a hop length of 16 ms, and a hamming window.

Table 1: Performance of models trained on CommonVoice-DEMAND *English*; tested on English, Spanish and Welsh testsets.

	Model	PESQ	STOI	CSIG	CBAK	COVL
English	Noisy	2.19	0.95	3.27	2.40	2.67
	Spec Loss	2.64	0.96	3.66	2.64	3.14
	STOI Loss	2.46	0.96	3.47	2.35	2.93
	SISDR Loss	2.73	0.96	3.45	2.74	3.07
	HuBERT L_{FE}	2.75	0.95	3.78	2.71	3.25
	mHuBERT L_{FE}	2.79	0.96	3.70	2.76	3.23
	XLSR L_{FE}	2.48	0.93	3.30	2.92	2.86
Spanish	Noisy	2.12	0.95	2.98	2.23	2.46
	Spec Loss	2.57	0.96	3.37	2.62	2.94
	STOI Loss	2.39	0.96	3.19	2.31	2.73
	SISDR Loss	2.68	0.96	3.15	2.74	2.88
	HuBERT L_{FE}	2.72	0.95	3.57	2.71	3.11
	mHuBERT L_{FE}	2.75	0.96	3.48	2.75	3.08
	XLSR L_{FE}	2.51	0.93	2.93	2.81	2.65
Welsh	Noisy	2.12	0.96	3.06	2.18	2.50
	Spec Loss	2.61	0.96	3.39	2.56	2.97
	STOI Loss	2.45	0.97	3.33	2.27	2.83
	SISDR Loss	2.72	0.97	3.15	2.69	2.89
	HuBERT L_{FE}	2.71	0.96	3.56	2.61	3.09
	mHuBERT L_{FE}	2.78	0.96	3.48	2.68	3.09
	XLSR L_{FE}	2.44	0.93	2.93	2.73	2.62

4.4. Result

Tables 1 and 2 display results on CommonVoice-DEMAND testsets for models trained on English and Spanish, respectively. PESQ [9], STOI [10], and the components of the Composite [36] intrusive metrics are reported, where higher values are better. The scores for the best performing model on each testset are highlighted in bold. The models perform better on the respective testset matching their language of training; the drop in performance on the non-matching testsets is consistent across all the models trained, including the baseline systems. Performance on the proposed English CommonVoice-DEMAND dataset is similar to that of models trained on the original VoiceBank-DEMAND in [20]. The best performing models, in terms of CBAK score are those trained with XLSR L_{FE} loss, except for one case. This is again consistent with the findings in [20]. HuBERT and mHuBERT perform similarly, with mHuBERT slightly outperforming in most cases.

Interestingly, all SSSR loss function models trained using Spanish audio perform better on the Welsh testset than those trained on the English audio. This is despite the fact that Welsh is lexically more similar to English than Spanish [37].

The quantity of data used for training the SSSR appears to be more important than language, as mHuBERT is trained with more English audio than HuBERT, however XLSR, trained with the most English speech data, performs worse. To further investigate this, an additional model was trained utilising the WavLM Base+ [38] SSSR, training and testing on the English CommonVoice-DEMAND dataset. WavLM Base+ has a parameter count comparable to HuBERT and mHuBERT and is trained with a similar objective and but with an additional speech denoising task. It is trained on 96k hours of English only audio. These results are shown in Table 3; the new model performs similarly in terms of PESQ score but somewhat better than all others in terms of CSIG.

Table 2: Performance of models trained on CommonVoice-DEMAND *Spanish*; tested on English, Spanish and Welsh testsets.

	Model	PESQ	STOI	CSIG	CBAK	COVL
English	Noisy	2.19	0.95	3.27	2.40	2.67
	Spec Loss	2.50	0.95	3.54	2.59	3.00
	STOI Loss	2.37	0.96	3.32	2.35	2.80
	SISDR Loss	2.44	0.95	3.25	2.58	2.82
	HuBERT L_{FE}	2.60	0.95	3.58	2.62	3.06
	mHuBERT L_{FE}	2.70	0.95	3.62	2.69	3.14
	XLSR L_{FE}	2.44	0.93	3.13	2.75	2.72
Spanish	Noisy	2.12	0.95	2.98	2.23	2.46
	Spec Loss	2.64	0.96	3.50	2.68	3.04
	STOI Loss	2.40	0.96	3.11	2.36	2.70
	SISDR Loss	2.61	0.96	3.18	2.72	2.86
	HuBERT L_{FE}	2.81	0.96	3.75	2.76	3.25
	mHuBERT L_{FE}	2.89	0.96	3.73	2.81	3.29
	XLSR L_{FE}	2.63	0.95	3.06	2.83	2.77
Welsh	Noisy	2.12	0.96	3.06	2.18	2.50
	Spec Loss	2.63	0.96	3.49	2.59	3.03
	STOI Loss	2.42	0.97	3.22	2.29	2.77
	SISDR Loss	2.60	0.97	3.18	2.63	2.85
	HuBERT L_{FE}	2.72	0.96	3.62	2.62	3.13
	mHuBERT L_{FE}	2.83	0.96	3.64	2.69	3.21
	XLSR L_{FE}	2.52	0.94	2.94	2.64	2.65

Overall, these results suggest that the BERT style training objective HuBERT, mHuBERT and WavLM might make them better suited as loss function feature representations when signal quality is the main concern as shown by the high PESQ and CSIG scores while the contrastive feature encoding masking objective of XLSR makes it more suitable if the objective of the enhancement is background noise reduction at the cost of additional speech distortion as the higher CBAK scores of the XLSR models demonstrates.

Table 3: Performance of L_{FE} Loss models trained on CommonVoice-DEMAND English and tested on English testset.

	Model	PESQ	STOI	CSIG	CBAK	COVL
	Noisy	2.19	0.95	3.27	2.40	2.67
	HuBERT L_{FE}	2.75	0.95	3.78	2.71	3.25
	mHuBERT L_{FE}	2.79	0.96	3.70	2.76	3.23
	XLSR L_{FE}	2.48	0.93	3.30	2.92	2.86
	WavLM L_{FE}	2.76	0.96	3.84	2.71	3.28

5. CONCLUSION

In this work, a system to create noisy speech datasets for a number of languages are proposed. These noisy speech datasets are used to train and test neural speech enhancement models using SSSR based loss functions. It is found that the language of the audio used to train the representations has a minimal impact on their performance when used in this manner, and that training objective and amount of training data has a greater effect. However, questions remain as the degree that these factors effect the performance of these loss functions; future work will aim to investigate this.

6. REFERENCES

- [1] P. C. Loizou, *Speech Enhancement: Theory and Practice*, 2nd ed. USA: CRC Press, Inc., 2013.
- [2] S. Doclo, W. Kellermann, S. Makino, and S. E. Nordholm, "Multi-channel signal enhancement algorithms for assisted listening devices: Exploiting spatial diversity using multiple microphones," *IEEE Signal Processing Magazine*, vol. 32, no. 2, 2015.
- [3] N. Moritz, K. Adiloğlu, J. Anemüller, S. Goetze, and B. Kollmeier, "Multi-channel speech enhancement and amplitude modulation analysis for noise robust automatic speech recognition," *Computer Speech & Language*, vol. 46, November 2017.
- [4] R. Haeb-Umbach, J. Heymann, L. Drude, S. Watanabe, M. Delcroix, and T. Nakatani, "Far-field automatic speech recognition," *Proceedings of the IEEE*, vol. 109, no. 2, 2021.
- [5] S.-W. Fu, C. Yu, T.-A. Hsieh, P. Plantinga, M. Ravanelli, X. Lu, and Y. Tsao, "Metricgan+: An improved version of metricgan for speech enhancement," 2021.
- [6] M. Maciejewski, G. Wichern, and J. Le Roux, "WHAMR!: Noisy and reverberant single-channel speech separation," in *ICASSP 2020*, May 2020.
- [7] X. Chang, W. Zhang, Y. Qian, J. L. Roux, and S. Watanabe, "MIMO-SPEECH: End-to-End Multi-Channel Multi-Speaker Speech Recognition," *ASRU 2019*, October 2019.
- [8] G. Close, T. Hain, and S. Goetze, "PAMGAN+/-: Improving Phase-Aware Speech Enhancement Performance via Expanded Discriminator Training," in *AES Convention Europe 2023*, 2023.
- [9] A. Rix, J. Beerends, M. Hollier, and A. Hekstra, "Perceptual evaluation of speech quality (PESQ)-a new method for speech quality assessment of telephone networks and codecs," in *ICASSP 2001*, vol. 2, 2001.
- [10] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time-frequency weighted noisy speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, 2011.
- [11] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR – Half-baked or Well Done?" in *ICASSP 2019*, May 2019.
- [12] A. Avila, B. Cauchi, S. Goetze, S. Doclo, and T. Falk, "Performance comparison of intrusive and non-intrusive instrumental quality measures for enhanced speech," in *2016 IEEE International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2016.
- [13] G. Close, S. Hollands, T. Hain, and S. Goetze, "Non-intrusive Speech Intelligibility Metric Prediction for Hearing Impaired Individuals for the Clarity Prediction Challenge 1," in *Proc. Interspeech 2022*, 2022.
- [14] B. Cauchi, K. Siedenbarg, J. F. Santos, T. H. Falk, S. Doclo, and S. Goetze, "Non-Intrusive Speech Quality Prediction Using Modulation Energies and LSTM-Network," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 7, 2019.
- [15] C. K. A. Reddy, V. Gopal, and R. Cutler, "Dnsmos: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," 2021.
- [16] C.-C. Lo, S.-W. Fu, W.-C. Huang, X. Wang, J. Yamagishi, Y. Tsao, and H.-M. Wang, "MOSNet: Deep Learning-Based Objective Assessment for Voice Conversion," in *Proc. Interspeech 2019*, 2019.
- [17] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [18] W.-N. Hsu, B. Bolte, Y.-H. Hubert Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," 2021.
- [19] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhota, 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," in *Proc. Interspeech 2021*, 2021.
- [20] G. Close, W. Ravenscroft, T. Hain, and S. Goetze, "Perceive and predict: Self-supervised speech representation based loss functions for speech enhancement," in *ICASSP 2023*, 2023.
- [21] O. Tal, M. Mandel, F. Kreuk, and Y. Adi, "A Systematic Comparison of Phonetic Aware Techniques for Speech Enhancement," in *Proc. Interspeech 2022*, 2022.
- [22] T.-A. Hsieh, C. Yu, S.-W. Fu, X. Lu, and Y. Tsao, "Improving perceptual quality by phone-fortified perceptual loss using wasserstein distance for speech enhancement," in *Interspeech '21*, 2021.
- [23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [24] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. of ACL 2019*, Minneapolis, 2019.
- [25] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *ICASSP 2015*, 2015.
- [26] A. Lee, H. Gong, P.-A. Duquenne, H. Schwenk, P.-J. Chen, C. Wang, S. Popuri, Y. Adi, J. Pino, J. Gu, and W.-N. Hsu, "Textless speech-to-speech translation on real data," in *ACL*, Seattle, United States, July 2022.
- [27] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, "VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *ACL Proceedings*. Association for Computational Linguistics, Aug. 2021.
- [28] A. Babu, C. Wang, A. Tjandra, K. Lakhota, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, "XLS-R: Self-supervised Cross-lingual Speech Representation Learning at Scale," in *Proc. Interspeech 2022*, 2022.
- [29] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Investigating RNN-based speech enhancement methods for noise-robust Text-to-Speech," in *ISCA Speech Synthesis Workshop*, 2016.
- [30] J. Thiemann, N. Ito, and E. Vincent, "DEMAND: a collection of multi-channel recordings of acoustic noise in diverse environments," June 2013.
- [31] *Objective measurement of active speech level*, International Telecoms Union (ITU) Std. P.56, 1993.
- [32] 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," in *Proc. Conf. on Language Resources and Evaluation*, 2020.
- [33] 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, arXiv:2106.04624.
- [34] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [35] G. Close, T. Hain, and S. Goetze, "MetricGAN+/-: Increasing Robustness of Noise Reduction on Unseen Data," in *EUSIPCO 2022*, Belgrade, Serbia, Aug. 2022.
- [36] Z. Lin, L. Zhou, and X. Qiu, "A composite objective measure on subjective evaluation of speech enhancement algorithms," *Applied Acoustics*, vol. 145, 02 2019.
- [37] G. Bella, K. Batsuren, and F. Giunchiglia, "A database and visualization of the similarity of contemporary lexicons," in *Text, Speech, and Dialogue*, K. Ekstein, F. Pártl, and M. Konopík, Eds. Cham: Springer International Publishing, 2021.
- [38] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, M. Zeng, X. Yu, and F. Wei, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE J. Sel. Topics in Signal Processing*, vol. 16, 10 2022.