



# Exploring the influence of fine-tuning data on wav2vec 2.0 model for blind speech quality prediction

*Helard Becerra, Alessandro Ragano, Andrew Hines*

School of Computer Science, University of College Dublin, Dublin, Ireland

helard.becerra@ucd.ie, alessandro.ragano@ucdconnect.ie, andrew.hines@ucd.ie

## Abstract

Recent studies have shown how self-supervised models can produce accurate speech quality predictions. Speech representations generated by the pre-trained wav2vec 2.0 model allows constructing robust predicting models using small amounts of annotated data. This opens the possibility of developing strong models in scenarios where labelled data is scarce. It is known that fine-tuning improves the model's performance; however, it is unclear how the data (e.g., language, amount of samples) used for fine-tuning is influencing that performance. In this paper, we explore how using different speech corpus to fine-tune the wav2vec 2.0 can influence its performance. We took four speech datasets containing degradations found in common conferencing applications and fine-tuned wav2vec 2.0 targeting different languages and data size scenarios. The fine-tuned models were tested across all four conferencing datasets plus an additional dataset containing synthetic speech and they were compared against three external baseline models. Results showed that fine-tuned models were able to compete with baseline models. Larger fine-tune data guarantee better performance; meanwhile, diversity in language helped the models deal with specific languages. Further research is needed to evaluate other wav2vec 2.0 models pre-trained with multi-lingual datasets and to develop prediction models that are more resilient to language diversity.

**Index Terms:** wav2vec, speech quality, self-supervised models

## 1. Introduction

Speech quality assessment is key to monitoring and evaluating the performance of applications and services in which speech is an essential component. Their success depends heavily on assuring acceptable levels of perceived quality, thus the importance of measuring it. Accurate quality assessments can be gathered through subjective experiments in the form of mean opinion scores (MOS) [1] or using the Multiple Stimuli with Hidden Reference and Anchor method (MUSHRA) [2], yet, the resources required (e.g., time, human availability) and their nature make it hard to implement them over an automatic quality monitoring design. Computer algorithms (objective metrics) that can automatically estimate speech quality offer a practical solution to this problem. Within speech objective metrics, Full-reference (FR) metrics, which use both degraded and reference speech signals, have been shown to work very well at estimating the perceived speech quality [3, 4, 5, 6, 7]; however, they depend on the availability of the reference signal. On the other hand, No-reference metrics (NR), which rely solely on the degraded signal, can be deployed in real-time monitoring scenarios, but at the cost of accuracy [8].

Machine learning (ML) provides the potential for data-driven metrics rather than similarity or distance error-based approaches that have been shown to work well in FR

scenarios [5]. Self-supervised pre-trained speech models have shown to be very useful for different speech-related tasks, including speech quality prediction [9, 10]. The wav2vec 2.0 model, a self-supervised learning (SSL) framework composed of a multi-layer convolutional feature encoder and a Transformer, was successfully applied for MOS prediction using the fixed-size representations generated by the model [10, 9]. Models based on deep-learning techniques have also been developed to predict the speech quality with very accurate outcomes. NISQA is a model based on a convolutional neural network (CNN) for frame-wise feature computation, a self-attention block for time-dependency modelling, and an attention-pooling block for estimating the overall MOS. This model showed stable results when tested on unknown data and live phone calls [11]. Despite their promising results, data-driven models are exposed to bias depending on the type of data used to train them. Collecting data that does not bias the model is a challenge (specifically neutral to: speaker voice/accent, language, degradation) [12, 13, 14].

Several studies have pointed out the importance of the fine-tuning step to achieve better performance, especially over specific contexts where labelled data is scarce or hard to collect [15, 10]. For wav2vec 2.0, adding a fine-tune step has shown a noticeable accuracy improvement in MOS prediction over different speech datasets [9, 10]. Adding this step facilitates the exploration of different testing scenarios involving specific types of degradation, speech languages, or subjective collection methods (e.g., MUSHRA, MOS). However, further research is needed to understand the level of influence that the data (e.g., language, type of distortion, amount of samples) used for fine-tuning is exerting over the model's performance.

This paper aims: 1) to better understand the role of fine-tuning in wav2vec 2.0 based models and how the content and size of the fine-tuning datasets impact the model performance over a range of speech datasets, 2) to compare wav2vec 2.0 based models and NISQA based models over a range of speech datasets. With that aim, two experiments were carried out using the fine-tuning system presented in [10]. The first experiment explores the content language aspect; for this experiment, three models were fine-tuned using only English, German, and Chinese speech samples. The experiment aims to better understand the influence of language on speech quality prediction on datasets with different language samples. The second experiment explores the size of the datasets used for fine-tuning; three models were fine-tuned using 1 k, 10 k, and 50 k mixed speech samples. The experiment seeks to understand the level of influence that fine-tuning data size has on speech quality predictions. Labelled speech datasets containing conferencing degradations (Tencent, IU Bloomington [16, 17], NISQA Corpus [11], and PSTN [18]) and synthesised speech (VoiceMOS [10]) were used to build the fine-tuning datasets and test the resulting models.

## 2. Methods

### 2.1. The wav2vec 2.0 pre-trained model

This model is based on the wav2vec 2.0 SSL framework, which consists of a convolutional feature encoder and a Transformer that takes raw waveforms as input and produces fixed-size speech representations [19]. The framework was successfully applied to speech-related tasks like speech recognition and language verification [15, 20]. Due to a large amount of data on which the model is pre-trained, only a small amount of data is required to apply the model over different tasks on which speech representation plays a key role. The promising results presented in prior studies, more specifically for speech quality prediction [10, 9], plus the release of annotated datasets for speech quality challenges, motivated the exploration of this framework and how fine-tuning can influence its performance in the speech quality prediction task. The wav2vec 2.0 model, pre-trained with 53.2 k hours of audio from the Librispeech dataset [21], served as the base to fine-tune a set of models targeting specific speech contexts as presented in Table 1.

### 2.2. Datasets

The availability of large multi-lingual labelled datasets, in the context of the VoiceMOS and ConferencingSpeech challenges, enabled the realization of this study. Such datasets include more than 200 hours of speech samples degraded with common degradations experienced over conferencing applications plus speech synthesis and voice conversion samples. This study considered the synthesised speech dataset VoiceMOS [10], plus four sets of datasets with speech conferencing distortions, namely: Tencent (2 datasets), NISQA (7 datasets) [11], IU-Bloomington (2 datasets)[16, 17] and PSTN [18]. These datasets were used to build subsets to fine-tune the wav2vec 2.0 pre-trained model targeting specific speech scenarios. Table 2 presents the main characteristics of the 13 datasets used in this study.

### 2.3. Fine-tune experiments

Two experiments were conducted to help understand the impact of fine-tuning data over the wav2vec 2.0 performance: fine-tuning models based on the datasets’ language and size. The language models were fine-tuned using samples containing speech in German, Chinese, and English. This experiment aims to understand how the language applied during the fine-tuning is affecting the models’ performance. For the second group of models, fine-tuning was done using different amounts of speech samples: 1 k, 10 k and 50 k. This experiment aims to observe how different datasets sizes used for fine-tuning the models influence their performance. Train and test subsets were randomly generated using an 80-20 percent ratio and maintaining a similar MOS distribution across all fine-tune datasets. The wav2vec 2.0 pre-trained model <sup>1</sup> and the fine-tune SSL system released for the VoiceMOS challenge <sup>2</sup> were used to fine-tune the models for the two experiments. In this system, fine-tuning is done by adding an output mean-pooling layer trained with an L1 loss function. The mean square error (MSE) between the predicted and target MOS is monitored on each training epoch, and training is stopped if no improvement is observed after 20 epochs. Details about the fine-tuned models and the subsets used to build them are presented in Table 1.

<sup>1</sup><https://github.com/pytorch/fairseq/tree/main/examples/wav2vec>

<sup>2</sup><https://github.com/nii-yamagishilab/mos-finetune-ssl>

Table 1: *Fine-tune description for Experiments 1 and 2. all(L): German, Chinese and English. all(D): NISQA TestLiveTalk, Tencent, and PSTN. PT: pre-trained. FT: fine-tuned.*

Experiment	Models		Language	Fine-tune Datasets		Source Dataset
	FT Model	PT Model		Train samples	Val samples	
Exp. 1	w2v_ger	wav2vec 2.0	German	185	47	Nisqa TestLiveTalk
Exp. 1	w2v_cn	wav2vec 2.0	Chinese	2424	606	Tencent
Exp. 1	w2v_en	wav2vec 2.0	English	7515	1879	PSTN
Exp. 1	w2v_all	wav2vec 2.0	all(L)	11429	2858	all(D)
Exp. 2	w2v_1k	wav2vec 2.0	English	1884	472	PSTN/Tencent
Exp. 2	w2v_10k	wav2vec 2.0	English	8995	2248	PSTN/Tencent
Exp. 2	w2v_50k	wav2vec 2.0	English	56217	14005	PSTN/Tencent
Baseline	w2v_base	wav2vec 2.0	English	4974	1066	VoiceMOS
Baseline	nq_base1		English	70764	9713	Nisqa/Tencent/PSTN
Baseline	nq_base2		English	70764	9713	Nisqa/Tencent/PSTN

## 3. Results

This section presents the results of the fine-tuned models across all 13 datasets included in this study. These models are compared against wav2vec 2.0 fine-tuned with the VoiceMOS dataset (w2v\_base) and two NISQA based models (nq\_base1 and nq\_base2) trained both over NISQA, PSTN and Tencent datasets with different architecture settings. The selection of w2v\_base as a baseline was led by its good performance over synthesised speech data reported in [10]. The performance of these models was evaluated using the MSE, the Spearman’s rank correlation coefficient (SRCC) and the linear correlation coefficient (LCC). Results for all fine-tuned and baseline models were grouped for the two experiments, and they are presented in Figure 1. The first row presents the performance of the models for the language fine-tuning; meanwhile, the second row shows the results for the size fine-tuning aspect. Evaluated datasets are listed in the x axis.

### 3.1. Experiment 1: Language analysis

Among the fine-tuned models generated for this experiment, w2v\_en and w2v\_all showed good performance across most NISQA and PSTN corpora (English speech samples). The w2v\_en model, fine-tuned exclusively with English samples, outperformed the other models in terms of SRCC and LCC at the NISQA ValLive dataset. Meanwhile, the w2v\_all model shows its best individual SRCC and LCC scores on the NISQA TestLiveTalk and the Tencent w/o Reverberation datasets, German and Chinese speech datasets, respectively. The w2v\_ger model surpassed all other models when tested against the NISQA TestLive corpus, a dataset containing only German speech. As for w2v\_cn, its best individual performance was reported for the Tencent w/o Reverberation corpus (Chinese speech samples). Despite the good performance observed for the fine-tuned models, baseline models still surpass them at several datasets. nq\_base2 reported the best correlation scores across several NISQA, PSTN and Tencent corpora. As for the w2v\_base model, its best performance was observed for the IU-Bloomington corpora and the VoiceMOS dataset; yet, it surprises its low performance over the Tencent w Reverberation dataset, which was the worst correlation performance across all models. Finally, the nq\_base1 model showed a poor performance across most datasets, and it reported its worst individual correlation performance for the NISQA TestLiveTalk dataset.

From an MSE perspective, IU-Bloomington Cosine and Voices datasets showed the highest values across all models. This is explained in the methodology adopted for the subjective experiment (MUSHRA method) and its following re-scaling procedure. MUSHRA is 100 point scale (IUVo), 10 point scale (IUCo) compared to other datasets where MOS is 5 point scale. Aside from those datasets, most models reported MSE

Table 2: Description of the speech datasets.

Dataset	ID	Language	Degradation Types	Samples	Speakers	Systems	Methodology	Ratings per sample
VoiceMOS	VoMOS	English	text-to-speech, voice conversion, natural speech	7106	27	187	MOS	8
IU-Bloomington Cosine	IUCo	English	Background noise	18000			MUSHRA	
IU-Bloomington Voices	IUVo	English	Reverberation	18000			MUSHRA	
Tencent w/o Reverberation	TcwoR	Chinese	white noise, background noise, frequency filtering, amplitude clipping, codecs	8366			MOS	
Tencent w Reverberation	TCwR	Chinese	Reverberation	3197			MOS	24
NISQA Corpus TestFor	NTsF	English	Codecs, background noises, packet-loss, clipping, live conditions with WhatsApp, Zoom, and Discord	240	80	60	MOS	30
NISQA Corpus TestLiveTalk	NTsLT	German	Natural environment conditions	232	8	58	MOS	24
NISQA Corpus TestP501	NTsP501	English	Codecs, background noises, packet-loss, clipping, live conditions with Skype, Zoom, WhatsApp, and mobile network recordings	240	4	60	MOS	28
NISQA Corpus TrainLive	NTL	English	live telephone and Skype calls	1020	486	1020	MOS	5
NISQA Corpus TrainSim	NTs	English	white gaussian noise, MNRU noise, noise clips, frequency filters, amplitude clipping, speech level changes, codecs, packet-loss	10000	2322	10000	MOS	5
NISQA Corpus ValLive	NVaL	English	live telephone and Skype calls	200	102	200	MOS	5
NISQA Corpus ValSim	NVaS	English	white gaussian noise, MNRU noise, noise clips, frequency filters, amplitude clipping, speech level changes, codecs, packet-loss	2500	2500	2500	MOS	5
PSTN	PSTN	English	Background noise	58709	2150		MOS	5




Figure 1: SRCC, LCC and MSE results for all fine-tuned and baseline models across all speech datasets. First row: Experiment 1 - language influence. Second row: Experiment 2 - size influence.

scores below 1 for most datasets. For both Tencent datasets, the w2v\_base model showed MSE values up to 4, which evidence the model’s difficulty in dealing with non-English speech samples. This experiment showed that fine-tuning with different languages helped the model generalize better across other datasets. This is evidenced by the enormous gap between the w2v\_all and w2v\_base models, especially over non-English datasets.

### 3.2. Experiment 2: Size analysis

For the second experiment, three models were fine-tuned using different amounts of labelled data containing English and Chinese speech samples. Among these models, it is observed that the model w2v\_50K consistently reports better performance than the models fine-tuned with 10 k and 1 k samples. However, the performance gap between these three models is not very pronounced; they all show very similar correlation scores across all 13 datasets. The w2v\_50K model surpasses all other models on at least 3 NISQA datasets (TestFor, TestLiveTalk and ValLive), the Tencent w/o Reverberation and the PSTN dataset. Compared to the baseline models, nq\_base2 outperforms w2v\_50K

at 3 NISQA corpora (TrainLive, TrainSim and ValSim) and the Tencent w Reverberation. These models showed a very narrow performance gap, except for the NISQA TrainLive and TrainSim and the Tencent w Reverberation datasets. As for the w2v\_base model, its best performance was reported for the IU-Bloomington datasets (Cosine and Voices) and the VoiceMOS corpus. As in the previous experiment, MSE performance reported values around 1 for all datasets and models, except for the IU-Bloomington datasets.

The results from this experiment showed that increasing the size of the subset to fine-tune a model will certainly benefit the model’s performance, even at the point of surpassing strong baseline models at some datasets. However, it was also observed that fine-tuning, even with smaller datasets (1 k and 10 k), results in a relatively small cost in performance. This can be beneficial for a context in which labelled data is scarce, and a slight drop in prediction accuracy is accepted.

### 3.3. Overall Analysis

In order to understand how wav2vec 2.0 fine-tuned models behave compared to NISQA based models, performance results




Figure 2: Aggregated performance results across datasets.

from all 13 datasets were grouped and analysed in this section. Individual SRCC, LCC and MSE scores for each fine-tuned and baseline model across 13 datasets were averaged, and its distribution is depicted on the boxplots presented in Figure 2. The figure presents the distribution of the SRCC and LCC coefficients (first row) and the MSE scores (second row) across all 13 datasets for each evaluated model (x axis). It is observed that most of the fine-tuned models, except for the w2v\_ger model, report a low variance among their correlation scores. Fine-tuned models exhibit a more stable performance across different datasets despite being individually surpassed at several datasets. Compared to the baseline models, w2v\_all and w2v\_50K reported almost equivalent median values than the nq\_base2 baseline model but with lower variance across datasets. In terms of MSE, it is clear that NISQA based models produce closer predictions to the MOS ground truth data. Among the fine-tuned models, it is observed how the size of the dataset helps the model at making predictions with fewer errors (narrower variance towards the w2v\_50K model).

Regarding the effect of the degradation type on the models' performance, no unexpected patterns were observed across datasets. However, it was observed that IU-Bloomington Cosine, a dataset containing samples degraded by background noise, received the lowest scores across all models. A similar condition was observed for VoiceMOS dataset, where most models, except for w2v\_base, reported low correlation scores and high error rates. Which was expected, since w2v\_base was fine-tuned using synthetic speech.

#### 4. Discussion

In this paper, two sets of models were generated by fine-tuning the wav2vec 2.0 pre-trained model targeting different languages and dataset sizes. These models were tested against one w2v\_base model and two NISQA baseline models across 13 different datasets. Results showed that nq\_base2 model performs best at most datasets; however, w2v\_all and w2v\_50K reported similar performance with lower variance across several datasets.

Findings from Experiment 1, where the language aspect of the fine-tuning step was evaluated, showed that performance from wav2vec 2.0 based models could be benefit from includ-

ing speech samples from different languages, even in small amounts. These results align with the findings presented in [10], where wav2vec 2.0 models were evaluated over datasets containing Chinese and Japanese speech samples. In their study, the xlsr model, a wav2vec 2.0 based model pre-trained with almost 56 k hours of speech containing 53 different languages [22], showed good generalization ability. Multi-language pre-trained wav2vec 2.0 models can lead to the development of more robust models that could deal with language bias more efficiently [13]. In [22], a cluster analysis of the speech representations produced by the xlsr pre-train model showed that languages that share a similar root, like English and German or Italian and Spanish, tended to cluster together. This can benefit languages that have scarce annotated material as they can rely on other similar languages to fine-tune and obtain higher prediction performance. High-performance techniques that require less annotated data are highly needed, and they demand further exploration [23]. Moreover, studies targeting the pre-training step would reveal how accurate a model can get using a language-dedicated pre-trained model.

Experiment 2 aimed to explore how the dataset size used to fine-tune the models affected the overall performance. The findings from this experiment showed that, although a larger amount of samples will guarantee higher correlation and lower error rates, the performance gap is not necessarily large. This behaviour, where wav2vec 2.0 models report stable and moderate correlations, was also observed in [10]. In contrast, the performance of NISQA baseline models, trained with different architecture settings, showed noteworthy differences across the evaluated datasets. These baseline models show the best and worst performances for several datasets. A similar fine-tuning experiment using NISQA would allow us to explore and understand how is the size aspect affects NISQA.

#### 5. Conclusions

This study explored the wav2vec 2.0 pre-trained model and the influence of fine-tuning it with different languages and data sizes. Based on the findings of this study, it is possible to conclude that fine-tuning a wav2vec 2.0 model using 1) a mid-size labelled dataset and 2) with a good language representation in the fine-tuning dataset will produce competitive models with steady performance. Better performance can be achieved by increasing the size of the fine-tune dataset with minor variance. Such models might be capable of competing against strong models trained with larger datasets like NISQA based models evaluated in this study. Similar wav2vec 2.0 models, pre-trained with multi-lingual datasets, need to be explored to build more resilient models across different languages.

Overall, this study showed the feasibility of tailoring wav2vec 2.0 for specific scenarios to perform well with limited annotated data. Within Quality of Experience, a research area where labelled data is a critical limitation, wav2vec stands as a promising alternative to deal with setups involving audio or even explore its application for other media types.

#### 6. Acknowledgements

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number 17/RC/2289\_P2 and 17/RC-PhD/3483. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

## 7. References

- [1] I. Rec, “P. 800: Methods for subjective determination of transmission quality,” *International Telecommunication Union, Geneva*, vol. 22, 1996.
- [2] I. Recommendation, “Method for the subjective assessment of intermediate sound quality (mushra),” *ITU, BS*, pp. 1543–1, 2001.
- [3] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, “Perceptual evaluation of speech quality (pesq)—a new method for speech quality assessment of telephone networks and codecs,” in *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)*, vol. 2. IEEE, 2001, pp. 749–752.
- [4] R. Huber and B. Kollmeier, “Pemo-q—a new method for objective audio quality assessment using a model of auditory perception,” *IEEE Transactions on audio, speech, and language processing*, vol. 14, no. 6, pp. 1902–1911, 2006.
- [5] A. Hines and N. Harte, “Speech intelligibility prediction using a neurogram similarity index measure,” *Speech Communication*, vol. 54, no. 2, pp. 306–320, 2012.
- [6] J. G. Beerends, C. Schmidmer, J. Berger, M. Obermann, R. Ullmann, J. Pomy, and M. Keyhl, “Perceptual objective listening quality assessment (polqa), the third generation itu-t standard for end-to-end speech quality measurement part i—temporal alignment,” *Journal of the Audio Engineering Society*, vol. 61, no. 6, pp. 366–384, 2013.
- [7] M. Narbutt, J. Skoglund, A. Allen, M. Chinen, D. Barry, and A. Hines, “Ambiquat: Towards a quality metric for headphone rendered compressed ambisonic spatial audio,” *Applied Sciences*, vol. 10, no. 9, p. 3188, 2020.
- [8] L. Malfait, J. Berger, and M. Kastner, “P. 563—the itu-t standard for single-ended speech quality assessment,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 6, pp. 1924–1934, 2006.
- [9] W.-C. Tseng, C. Yu Huang, W.-T. Kao, Y. Y. Lin, and H. Yi Lee, “Utilizing Self-Supervised Representations for MOS Prediction,” in *Proc. Interspeech 2021*, 2021, pp. 2781–2785.
- [10] E. Cooper, W.-C. Huang, T. Toda, and J. Yamagishi, “Generalization ability of mos prediction networks,” *arXiv preprint arXiv:2110.02635*, 2021.
- [11] G. Mittag, B. Naderi, A. Chehadi, and S. Möller, “NISQA: A Deep CNN-Self-Attention Model for Multidimensional Speech Quality Prediction with Crowdsourced Datasets,” in *Proc. Interspeech 2021*, 2021, pp. 2127–2131.
- [12] R. Z. Jiménez, G. Mittag, and S. Möller, “Removing the bias in speech quality scores collected in noisy crowdsourcing environments,” in *2021 13th International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 2021, pp. 49–54.
- [13] R. Z. Jiménez, B. Naderi, and S. Möller, “Influence of language differences in crowdsourcing speech quality assessment studies,” in *13th International Conference on Quality of Multimedia Experience, QoMEX 2021, Montreal, Canada, June 14-17, 2021*. IEEE, 2021, pp. 43–48. [Online]. Available: <https://doi.org/10.1109/QoMEX51781.2021.9465437>
- [14] M. Chinen, “Marginal effects of language and individual raters on speech quality models,” *IEEE Access*, vol. 9, pp. 127 320–127 334, 2021.
- [15] Z. Fan, M. Li, S. Zhou, and B. Xu, “Exploring wav2vec 2.0 on Speaker Verification and Language Identification,” in *Proc. Interspeech 2021*, 2021, pp. 1509–1513.
- [16] A. Stupakov, E. Hanusa, J. Bilmes, and D. Fox, “Cosine—a corpus of multi-party conversational speech in noisy environments,” in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2009, pp. 4153–4156.
- [17] C. Richey, M. A. Barrios, Z. Armstrong, C. Bartels, H. Franco, M. Graciarena, A. Lawson, M. K. Nandwana, A. Stauffer, J. van Hout, P. Gamble, J. Hetherly, C. Stephenson, and K. Ni, “Voices Obscured in Complex Environmental Settings (VOICES) Corpus,” in *Proc. Interspeech 2018*, 2018, pp. 1566–1570.
- [18] G. Mittag, R. Cutler, Y. Hosseinkashi, M. Revow, S. Srinivasan, N. Chande, and R. Aichner, “DNN No-Reference PSTN Speech Quality Prediction,” in *Proc. Interspeech 2020*, 2020, pp. 2867–2871.
- [19] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [20] A. Baevski and A. Mohamed, “Effectiveness of self-supervised pre-training for asr,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7694–7698.
- [21] 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)*. IEEE, 2015, pp. 5206–5210.
- [22] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, “Unsupervised Cross-Lingual Representation Learning for Speech Recognition,” in *Proc. Interspeech 2021*, 2021, pp. 2426–2430.
- [23] A. Ragano, E. Benetos, and A. Hines, “More for less: Non-intrusive speech quality assessment with limited annotations,” in *2021 13th International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 2021, pp. 103–108.