



UTMOS: UTokyo-SaruLab System for VoiceMOS Challenge 2022

Takaaki Saeki*, Detai Xin*, Wataru Nakata*,
Tomoki Koriyama, Shinnosuke Takamichi, and Hiroshi Saruwatari

The University of Tokyo, Japan.

{takaaki.saeki, detai.xin}@ipc.i.u-tokyo.ac.jp, nakata-wataru855@g.ecc.u-tokyo.ac.jp

Abstract

We present the UTokyo-SaruLab mean opinion score (MOS) prediction system submitted to VoiceMOS Challenge 2022. The challenge is to predict the MOS values of speech samples collected from previous Blizzard Challenges and Voice Conversion Challenges for two tracks: a main track for in-domain prediction and an out-of-domain (OOD) track for which there is less labeled data from different listening tests. Our system is based on ensemble learning of strong and weak learners. Strong learners incorporate several improvements to the previous fine-tuning models of self-supervised learning (SSL) models, while weak learners use basic machine-learning methods to predict scores from SSL features. In the Challenge, our system had the highest score on several metrics for both the main and OOD tracks. In addition, we conducted ablation studies to investigate the effectiveness of our proposed methods.

Index Terms: VoiceMOS Challenge 2022, mean opinion score prediction, self-supervised learning, ensemble learning

1. Introduction

Although subjective evaluation has been the gold standard in the field of speech synthesis [1], its high cost in terms of time and money motivates us to develop measures for automatically determining the performance. Although a number of neural network-based approaches for doing this have been proposed [2–4], there are still many challenges, such as developing a general-purpose prediction model.

The VoiceMOS Challenge [5], which was launched this year, provides the common database and baseline systems. The database contains synthetic speech samples and the corresponding mean opinion scores (MOS) on a five-point scale as assigned by human evaluators. The participants construct a prediction system and submit the system’s predicted MOS for the test data. There are two tracks in the challenge, the main and out-of-domain (OOD) tracks, and the system performance is evaluated using several metrics.

In this paper, we present our MOS prediction system, UTMOS (pronounced “u-t-mos”), which we submitted to VoiceMOS Challenge 2022. Our system is based on ensemble learning of strong and weak learners: the strong learners are obtained by fine-tuning models of self-supervised learning (SSL) models, and the weak learners predict scores from SSL features by using non-neural-network machine-learning methods from SSL features. The strong learner incorporates several improvement functions, including contrastive learning, listener dependency, and phoneme encoding. We also present the results of VoiceMOS Challenge 2022 and those of our ablation studies. Our implementation is publicly available¹. This paper makes three contributions in particular:

*Equal contribution.

¹<https://github.com/sarulab-speech/UTMOS22>

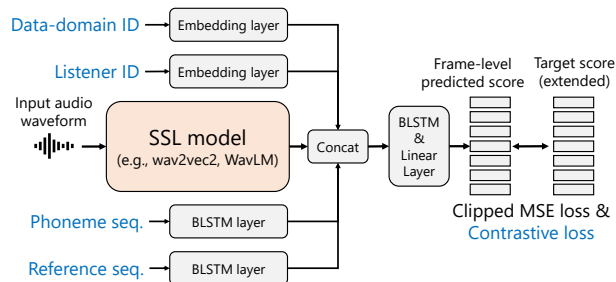


Figure 1: Architecture of the proposed strong learner.

- It describes and MOS prediction system that had the highest score on several metrics in the main and OOD tracks of VoiceMOS Challenge 2022.
- It presents proposed methods for predicting MOS that include contrastive learning and phoneme encoding.
- It presents the results of ablation studies demonstrating the effectiveness of listener-dependent learning and that of stacking by increasing the number of strong learners.

2. VoiceMOS Challenge 2022

The data used in the VoiceMOS Challenge 2022 were mainly synthetic speech samples from previous Blizzard Challenges and Voice Conversion Challenges. The VoiceMOS Challenge is divided into two tracks: the main track and the OOD track. The dataset statistics for both tracks are given in Table 1.

Main track. The main track uses the BVCC dataset [6], which contains data obtained by conducting large-scale listening tests on samples from 187 different systems from previous Blizzard Challenges, Voice Conversion Challenges, and ESPnet-TTS [7] published samples. The main track dataset consists of English synthetic speech samples. The test set contains samples obtained with the same listening test but from unseen systems, speakers, and listeners.

OOD track. The OOD track uses a dataset consisting of Chinese synthetic speech collected using a listening test different from that used for the main track. The dataset provides a small amount of labeled data and a large amount of audio-only unlabeled data.

For both the main and OOD tracks, prediction performance was evaluated using four metrics; mean squared error (MSE), linear correlation coefficient (LCC), Spearman rank correlation coefficient (SRCC), and Kendall rank correlation coefficient (KTAU). Participants submit their predicted score for each speech utterance in the test set, and utterance-level and system-level metrics were calculated for each of the four metrics.

3. UTMOS

Our UTMOS method leverages ensemble learning by using multiple models, which consists of *strong* learners and *weak*

Table 1: *Datasets used in VoiceMOS Challenge 2022: “closed/open” indicates that system used to synthesize speech is included/excluded in training data; “labeled/unlabeled” indicates that corresponding MOS score is included/excluded.*

	Main			OOD		
	Train	Dev	Test	Train	Dev	Test
Evaluations	39,792	8,258	8,528	1,848	1,819	7,680
Audio clips	4,974	1,066 (open)	1,066 (open)	136 (labelled) + 540 (unlabelled)	136 (open)	540 (open)
Audio clips per system	12–36 (avg: 29.4)	1–37 (avg: 5.9)	1–38 (avg: 5.7)	4–10 (avg: 6.5)	4–46 (avg: 2.5)	6–52 (avg: 20.8)
Systems	175	175 (closed) + 6 (open)	175 (closed) + 12 (open)	21	21 (closed) + 3 (open)	21 (closed) + 5 (open)
Listeners	288	288 (closed) + 8 (open)	288 (closed) + 16 (open)	285	285 (closed) + 43 (open)	285 (closed) + 76 (open)

learners. The strong learner is an SSL-based neural network model that directly uses a speech waveform as input. The weak learners are basic machine learning models, such as ridge regression models and support vector machines with utterance-level SSL features as input.

3.1. Fine-tuned SSL model

3.1.1. Basic architecture

Fig. 1 illustrates the architecture of the strong learner. As in a previous study [8], we used a pretrained SSL model to extract features from input audio. First, the raw waveform of a speech utterance is input to the SSL model to obtain frame-level features. Unlike the previous study [8], our model does not average the frame-level output features; instead, it sends them to bidirectional long short-term memory (BLSTM) and linear layers to compute frame-level scores. During training, we extend the target score by the number of frames and define a frame-level loss function. We found that the frame-level loss achieves higher performance than the previous one using averaged features. During inference, the model predicts the score by averaging the frame-level scores. Using this model, we introduced several functions described in Sections 3.1.2 to 3.1.5 for the strong learner.

3.1.2. Contrastive loss

Contrastive learning is a self-supervised machine-learning method that can utilize unlabeled data by learning from intrinsic similarity relations between data. Contrastive learning is widely used in speech quality assessment, in which speech representations are learned from large-scale unlabeled speech data [9–11]. Given scores s_1 and s_2 of utterances x_1 and x_2 , the difference in the scores ($d_{x_1, x_2} = s_1 - s_2$) can be regarded as the difference in the two utterances in terms of speech quality. If the predicted scores for the two utterances are denoted as \hat{s}_1 and \hat{s}_2 , respectively, it is intuitive to assume that the predicted difference ($\hat{d}_{x_1, x_2} = \hat{s}_1 - \hat{s}_2$) is close to d_{x_1, x_2} . Therefore, we consider a contrastive loss defined as $\mathcal{L}_{x_1, x_2}^{\text{con}} = \max(0, |d_{x_1, x_2} - \hat{d}_{x_1, x_2}| - \alpha)$, where α is a hyperparameter greater than zero. We call α the margin since it is similar to the support vector machine margin, i.e. small errors are ignored by the model. One advantage of this loss function is that it penalizes the model when the signs of d_{x_1, x_2} and \hat{d}_{x_1, x_2} are opposite, e.g. the case in which x_1 is better than x_2 but the model predicts that x_1 is worse than x_2 , which makes the contrastive loss suitable for improving the metrics based on ranking accuracy, such as the SRCC used in the Challenge.

In practice the contrastive loss of all possible pairs in a mini-batch are considered: $\mathcal{L}^{\text{con}} = \sum_{i \neq j} \mathcal{L}_{x_i, x_j}^{\text{con}}$. It is worth noting that, as discussed in Section 4.3, our model can be trained using only the proposed contrastive loss function without using other regular loss functions like MSE or L1 and still achieve better performance than the baseline model.

In addition to the contrastive loss, we use the clipped MSE loss [4] for the regression loss between the discrete predicted and ground truth scores: $\mathcal{L}^{\text{reg}}(y, \hat{y}) = \mathbb{1}(|y - \hat{y}| > \tau)(y - \hat{y})^2$.

The final loss function is defined as:

$$\mathcal{L} = \beta \mathcal{L}^{\text{reg}} + \gamma \mathcal{L}^{\text{con}} \quad (1)$$

where β and γ are hyperparameters.

3.1.3. Data-domain and listener dependent learning

The previous MOS prediction model based on an SSL model [8] learns the utterance-level MOS as the target variable. Previous studies [4, 12] improved the prediction accuracy by making listener-dependent predictions instead of simply predicting utterance-level MOS; accuracy was improved because the distribution of evaluation scores is different for each listener.

We thus introduced a listener dependency function into the SSL-based MOS predictor. As shown in Fig. 1, listener-embedding is concatenated with the features extracted by the SSL model to predict the listener-dependent score. During training, we also include a “mean listener” for which the target score is the average value of all the listeners’ scores, as in a previous study [12] Since the listener information is unknown in inference, the mean-listener embedding is given to the model to predict the utterance-level MOS.

Furthermore, we need to consider the bias of each listening test instead of considering only the bias per listener within a given listening test. For example, different listening tests are conducted for the main and OOD datasets. To include data from multiple domains in our training, we use both listener and domain IDs as shown in Fig. 1. When we train our model on all of the main, OOD, and external datasets described in Section 3.2, for example, we assign different domain IDs to the respective datasets. In addition, we use the average score within each domain to obtain the score of the mean listener.

3.1.4. Phoneme encoding

In our preliminary studies, we observed that there was a strong correlation between the MOS ratings and clustering results of linguistic contents estimated with an automatic speech recognition (ASR) model. On the basis of this observation, we use the ASR results as auxiliary input of the strong learner to further improve prediction accuracy. To apply this method to multilingual synthetic speech samples, we use phoneme sequences as input instead of graphemes or character sequences.

Furthermore, intuitively, the larger the difference between the reference text used to generate synthetic speech and the text estimated with ASR, the lower the intelligibility and the expected MOS rating. Since the participants were not provided the actual texts used for synthesized speech, to estimate the reference text of each utterance, we perform clustering on the ASR results by using the DBSCAN algorithm [13] based on normalized Levenshtein distance and extract the median text corresponding to each cluster. We refer to this median text as the reference sequence. As shown in Fig. 1, phoneme and reference sequences are fed to the phoneme encoder which consists of BLSTM layers and the initial and last hidden states are concatenated. Finally it is replicated with the number of frames and concatenated to the output of the SSL model.

3.1.5. Data augmentation

Deep neural networks usually suffer from overfitting if the training data is limited. Given the limited data size of the challenge, especially for the OOD track, overfitting is likely to happen. We thus utilize data augmentation to alleviate this problem. We consider two augmentation methods: speaking-rate-changing and pitch-shifting, which can alter the utterances while maintaining the MOS. Speaking-rate-changing slows down or speeds up the audio by a factor f_t . Since a very large or small f_t will affect the MOS, we set f_t close to 1. Pitch-shifting changes the speaker identity of the utterances by raising or lowering the original pitch p to $p + f_p$. During training, the two parameters (f_t and f_p) are randomly selected from two ranges $[1 - F_t, 1 + F_t]$ and $[-F_p, F_p]$, respectively. We tune F_p so that the MOS of the augmented waveform has perceptually little difference from the original ones. We use WavAugment [14] to implement all data augmentation methods.

3.2. External data collection

In the OOD track, the size of the labeled training data (136 utterances) is not sufficient to train a robust MOS prediction model. Therefore we utilized the 540 unlabeled utterances and collecting the corresponding MOS as external data. To this end, we first selected the system with the highest MOS (BC2019-A) and regarded all utterances of this system as natural speech after double-checking the utterances with a Chinese native speaker. We then conducted a standard 5-point-scale MOS test for all 540 unlabeled utterances and 249 labeled utterances. A total of 32 Chinese listeners participated in this test; each listener rated 55 utterances, so each utterance had 2 answers on average. The utterance-level SRCC between the ground truth scores and collected scores for the labeled data was 0.757, which indicates a strong correlation. Although the distribution of these collected scores was not exactly the same as that of the original utterances, since they were all evaluated by Chinese, we think it is appropriate to utilize the external data along with the original data. Using the external data substantially improved performance, as discussed in Section 4.

3.3. Ensemble learning with strong and weak learners

We use an ensemble of models for prediction robustness. Specifically, we use the stacking method [15, 16] illustrated in Fig. 2 as the ensemble method. We use not only fine-tuned SSL models but also simple regression models using utterance-level features. We refer to the former and latter models as “strong learners” and “weak learners,” respectively.

The weak learners are a combination of feature extractions and regression methods. We propose using pretrained-SSL-model-based mean embeddings for feature extraction. Specifically, we extract embeddings of input utterances and compute the mean for all frames, taking as inspiration the structure of SSL-MOS [8]. Although this process might be too rough for obtaining utterance-level characteristics, we assume that only the mean embeddings have efficient information for MOS prediction. For the simple regression models, we use basic ones such as linear regression, decision-tree-based methods, and kernel methods. In general, model diversity is important for prediction performance in ensemble training [17]. Hence, we use multiple pretrained SSL models for feature extraction to increase the number of models. Moreover, we enhance the diversity of weak learners by using different data domains, i.e. languages and MOS test environments, for the OOD track.

The stacking method comprises stages 0 to 3. After feature extraction, we train strong and weak learners individually and




Figure 2: Flow of stacking with strong and weak learners.

Table 2: Results of ablation study.

(a) Main								
	Utterance-level				System-level			
	MSE	LCC	SRCC	KTAU	MSE	LCC	SRCC	KTAU
UTMOS strong	0.276	0.883	0.881	0.708	0.148	0.930	0.925	0.774
w/o contrastive loss	0.241	0.881	0.879	0.706	0.114	0.932	0.930	0.781
w/o listener ID	<u>0.307</u>	<u>0.880</u>	<u>0.878</u>	<u>0.704</u>	<u>0.160</u>	0.935	0.933	0.784
w/o phoneme encoder	0.249	0.881	0.882	0.709	0.119	0.935	0.936	0.790
w/o data augmentation	0.226	0.885	0.882	0.710	0.103	0.936	0.933	0.784
w/o MSE loss	0.219	0.882	0.880	0.707	0.114	0.932	0.929	0.778
SSL-MOS	0.380	0.869	0.871	0.695	0.223	0.920	0.918	0.758

(b) OOD								
	Utterance-level				System-level			
	MSE	LCC	SRCC	KTAU	MSE	LCC	SRCC	KTAU
UTMOS strong	0.378	0.891	0.871	0.690	0.248	0.970	0.972	0.879
w/o contrastive loss	0.407	0.870	0.862	0.676	0.272	0.945	0.957	0.841
w/o listener ID	<u>0.636</u>	<u>0.847</u>	<u>0.825</u>	<u>0.638</u>	<u>0.490</u>	<u>0.931</u>	<u>0.944</u>	<u>0.820</u>
w/o phoneme encoder	0.390	0.893	0.881	0.702	0.258	0.966	0.967	0.868
w/o data augmentation	0.322	0.887	0.872	0.691	0.191	0.960	0.967	0.872
w/o external data	0.412	0.883	0.868	0.684	0.253	0.960	0.961	0.861
SSL-MOS	0.676	0.872	0.842	0.654	0.500	0.957	0.964	0.862

predict scores using cross validation. We then train meta learners using the first stage scores. Finally, we train the third stage model with the second stage scores and obtain the final score.

4. Experimental evaluations

4.1. Experimental conditions

For the strong learners, we trained models with six different configurations for the main track and six different configurations for the OOD track. The main configuration was **UTMOS strong**, the strong learner used for submission. We also trained the strong learner without the functions described from Sections 3.1.2 to 3.2 and the use of MSE loss for the main track and the use of external data for the OOD track. For the preprocessing, we downsampled all speech samples to 16 kHz and normalized the volume. During training, each MOS rating was normalized to the range $[-1, 1]$ by applying linear projection. For the SSL models of strong learners, we used the published wav2vec2.0 [18] base model² pretrained on Librispeech [19]. For phoneme transcription in the phoneme encoder, we used the ASR model proposed by Xu et. al [20]. This model is xlsr-53 [21] fine-tuned on phonetic annotations from word transcriptions obtained using ESpeak³ and speech samples of Common-Voice [22]. For the phoneme encoder, we used 3-layer BLSTM with a hidden size of 256. For domain and listener embedding, we used an embedding dimension of 128. For the main track, we only used main track dataset. For the OOD track, we used the OOD track dataset, and external dataset we collected for training except when trained for “w/o external data”. For “w/o external data”, we used only the dataset from OOD track for training. For the hyperparameters of the loss function defined in Eq. (1), we set $\beta = 1$, and $\gamma = 0.5$ except for “w/o contrastive

²<https://github.com/pytorch/fairseq/blob/main/examples/wav2vec>

³<https://github.com/espeak-ng/espeak-ng>

Table 3: Results of stacking. “Strong” and “Weak” are the number of strong and weak learners used for stacking, except the case that Strong is 1 and Weak is 0, which means a single SSL model is used. With regard to the numbers of weak learners at OOD track, 48, 96, and 144 corresponds to 1 (OOD), 2 (OOD, external), and 3 (OOD, external, main) domains, respectively.

		(a) Main							
Strong	Weak	Utterance-level				System-level			
		MSE	LCC	SRCC	KTAU	MSE	LCC	SRCC	KTAU
1	0	0.216	0.894	0.890	0.720	0.105	0.937	0.934	0.792
17	0	0.169	0.896	0.893	0.725	0.088	0.939	0.936	0.792
0	48	0.186	0.887	0.885	0.714	0.108	0.928	0.927	0.777
1	48	0.172	0.896	0.894	0.726	0.098	0.935	0.933	0.789
5	48	0.169	0.898	0.895	0.728	0.095	0.938	0.936	0.793
12	48	0.169	0.898	0.895	0.728	0.094	0.938	0.935	0.792
17	48	0.165	0.899	0.896	0.730	0.090	0.939	0.936	0.795

		(b) OOD							
Strong	Weak	Utterance-level				System-level			
		MSE	LCC	SRCC	KTAU	MSE	LCC	SRCC	KTAU
1	-	0.280	0.905	0.885	0.704	0.160	0.972	0.965	0.858
6	0	0.155	0.920	0.896	0.720	0.029	0.988	0.975	0.886
0	48	0.204	0.893	0.858	0.674	0.033	0.985	0.963	0.860
0	96	0.179	0.907	0.877	0.696	0.030	0.988	0.975	0.890
0	144	0.176	0.909	0.882	0.702	0.033	0.987	0.974	0.888
1	144	0.174	0.910	0.883	0.704	0.033	0.986	0.976	0.894
6	144	0.162	0.917	0.892	0.715	0.028	0.989	0.977	0.900

loss” and “w/o MSE loss.” For “w/o contrastive loss” and “w/o MSE loss,” we used $\beta = 1, \gamma = 0$, and $\beta = 0, \gamma = 1$, respectively. For α and τ , we set $\alpha = 0.5, \tau = 0.25$ except for “w/o listener ID.” For “w/o listener ID,” we set $\alpha = 0.1, \tau = 0.1$. For data augmentation, we set $F_t = 0.1$, and $F_p = 300$ cents except for “w/o data augmentation.” For “w/o data augmentation,” no data augmentation was performed. For the optimizer, we used Adam [23] ($\beta_1 = 0.9, \beta_2 = 0.99$) with linear warmup and linear decay learning rate scheduling. Learning rate warmup was performed for 4000 steps, and the total number of training steps was 15,000. The batch size was 12, and gradient accumulation was performed every 2 steps. The best model checkpoint was selected on the basis of the highest system-level SRCC calculated from the development set. For the ablation study of strong learners and stacking, training was performed for five times and the results were calculated by averaging scores for each metric as model performance varies depending on the random seed.

Regarding the conditions for stacking and weak learners, the strong learners for stacking were chosen from the candidates during hyperparameter tuning by using the Optuna [24] based on the system-level SRCC of development set. We used a maximum of 17 and 6 strong learners for the main and OOD tracks, respectively. For the pretrained SSL features for weak learners, we used four wav2vec 2.0 [18], two HuBERT [25], and two WavLM [26] models, which differed from each other in model size, database, and training method. The simple regression methods for the weak and meta learners were two linear regressions (ridge regression and linear support vector regression (SVR)), two tree-based models (random forests and LightGBM [27]), and two kernel methods (kernel SVR and Gaussian process regression). By combining pretrained SSL models and simple regression methods, we obtained 48 weak learners.

The training of the weak and meta learners for the main track was performed using only the main track data. For the OOD track, we trained weak models for respective domains, which were main, OOD, and external ones, and integrated the results at the second stage. Hence, the number of weak learners

for the OOD track was 144. The meta learners for the OOD track were trained using the OOD track data.

4.2. VoiceMOS2022 results

In the both tracks, utterance-level (Utt.) and system-level (Sys.) metrics were calculated as described in Section 2. Three baseline methods and the 21 teams participated in the Main track. For the OOD track, scores of three baseline methods and 15 teams were submitted. Our team ID is “T17.”

Part of our results in the Main track were Utt. MSE = 0.165 (1), Utt. SRCC = 0.897 (1), Sys. MSE = 0.090 (1), Sys. SRCC = 0.936 (3), where the numbers in parentheses mean the rankings. The results in the OOD track were Utt. MSE = 0.162 (1), Utt. SRCC = 0.893 (2), Sys. MSE = 0.030 (1), Sys. SRCC = 0.988 (1).

4.3. Ablation study on SSL-based models

We conducted ablation studies for each of the methods described in Section 3.1. We denote a strong learner using all the methods in Section 3.1 as “UTMOS strong.” A method based on fine-tuning of the SSL model [8], which is a baseline method of the challenge, was designated as “SSL-MOS.” Table 2 lists the results. The best results are shown in bold, while the worst ones are underlined except for UTMOS strong and SSL-MOS.

We can see that all of our methods outperformed SSL-MOS in almost all indices. Furthermore, in the main track, the method that excluded data augmentation or phoneme encoder from UTMOS strong showed better results, which may be due to the larger amount of data of the main track than the OOD track. In the OOD track, UTMOS strong showed the best results in several indices including test system SRCC. This suggests that all of the proposed methods have effectiveness in cases with smaller amounts of data. For both the main and OOD tracks, the performance of the methods without the listener ID significantly degraded in many cases, indicating the effectiveness of listener dependency.

4.4. Evaluation on stacking

To investigate the effectiveness of strong and weak learners at stacking, we computed the prediction accuracy scores by changing the number of strong and weak learners. The results are shown in Table 3. The 1, 5, and 12 strong learners were chosen greedily based on system-level SRCC of development set.

We can see that even single strong layer gave high SRCCs although the MSEs were still large. By using the stacking ensemble with multiple strong learners, MSEs were reduced while SRCCs were kept high. The stacking using only weak models even had high SRCCs although the extracted features were simpler than fine-tuned SSL models. We also see that the increase of the number of strong and weak learners tended to improve prediction accuracy, which indicates it is promising to increase the number of models by using multiple hyperparameter values and multiple domains.

5. Conclusion

We presented the system we submitted to VoiceMOS Challenge 2022. Our system is based on ensemble learning of strong learners, which are obtained by fine-tuning SSL models, and weak learners that predict scores from SSL features. Future work includes constructing a larger-scale general-purpose MOS prediction model by collecting a wider variety of data.

Acknowledgements: Part of this work was supported by JSPS KAKENHI Grant Number 21H04900, 21K11955, and JST SPRING, Grant Number JPMJSP2108 (for implementation) and JST Moonshot R&D Grant Number JPMJPS2011 (for evaluation).

6. References

- [1] A. W. Black and K. Tokuda, "The Blizzard Challenge-2005: Evaluating corpus-based speech synthesis on common datasets," in *Proc. INTERSPEECH*, Lisbon, Portugal, Sep. 2005.
- [2] B. Patton, Y. Agiomyrgiannakis, M. Terry, K. W. Wilson, R. A. Saurous, and D. Sculley, "AutoMOS: Learning a non-intrusive assessor of naturalness-of-speech," *arXiv preprint arXiv:1611.09207*, 2016.
- [3] 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," *Proc. Interspeech*, pp. 1541–1545, 2019.
- [4] Y. Leng, X. Tan, S. Zhao, F. K. Soong, X.-Y. Li, and T. Qin, "MB-NET: MOS prediction for synthesized speech with mean-bias network," *Proc. ICASSP*, pp. 391–395, 2021.
- [5] W.-C. Huang, E. Cooper, Y. Tsao, H.-M. Wang, T. Toda, and J. Yamagishi, "The VoiceMOS Challenge 2022," *arXiv preprint arXiv:2203.11389*, 2022.
- [6] E. Cooper and J. Yamagishi, "How do voices from past speech synthesis challenges compare today?" in *Proc. SSW*, 2021, pp. 183–188.
- [7] T. Hayashi, R. Yamamoto, K. Inoue, T. Yoshimura, S. Watanabe, T. Toda, K. Takeda, Y. Zhang, and X. Tan, "ESPnet-TTS: Unified, reproducible, and integratable open source end-to-end text-to-speech toolkit," *Proc. ICASSP*, pp. 7654–7658, 2020.
- [8] E. Cooper, W.-C. Huang, T. Toda, and J. Yamagishi, "Generalization ability of MOS prediction networks," *arXiv preprint arXiv:2110.02635*, 2021.
- [9] J. Serrà, J. Pons, and S. Pascual, "SESQA: semi-supervised learning for speech quality assessment," in *Proc. ICASSP*. IEEE, 2021, pp. 381–385.
- [10] P. Manocha, Z. Jin, R. Zhang, and A. Finkelstein, "CDPAM: Contrastive learning for perceptual audio similarity," in *Proc. ICASSP*. IEEE, 2021, pp. 196–200.
- [11] P. Manocha, B. Xu, and A. Kumar, "NORESQA: A framework for speech quality assessment using non-matching references," *Proc. NeurIPS*, vol. 34, 2021.
- [12] W.-C. Huang, E. Cooper, J. Yamagishi, and T. Toda, "LDNet: Unified listener dependent modeling in MOS prediction for synthetic speech," *arXiv preprint arXiv:2110.09103*, 2021.
- [13] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*. AAAI Press, 1996, p. 226–231.
- [14] E. Kharitonov, M. Rivière, G. Synnaeve, L. Wolf, P.-E. Mazaré, M. Douze, and E. Dupoux, "Data augmenting contrastive learning of speech representations in the time domain," *arXiv preprint arXiv:2007.00991*, 2020.
- [15] D. H. Wolpert, "Stacked generalization," *Neural Networks*, vol. 5, no. 2, pp. 241–259, 1992.
- [16] L. Breiman, "Stacked regressions," *Machine learning*, vol. 24, pp. 49–64, 1996.
- [17] Z.-H. Zhou, *Ensemble methods: foundations and algorithms*. CRC press, 2012.
- [18] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *arXiv preprint arXiv:2006.11477*, 2020.
- [19] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. ICASSP*, South Brisbane, Australia, Apr. 2015, pp. 5206–5210.
- [20] Q. Xu, A. Baevski, and M. Auli, "Simple and Effective Zero-shot Cross-lingual Phoneme Recognition," *arXiv preprint arXiv:2109.11680*, 2021.
- [21] 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.
- [22] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proc. LREC 2020*, 2020, pp. 4218–4222.
- [23] D. Kingma and B. Jimmy, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [24] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A next-generation hyperparameter optimization framework," in *Proc. KDD*, 2019.
- [25] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *arXiv preprint arXiv:2106.07447*, 2021.
- [26] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *arXiv preprint arXiv:2110.13900*, 2021.
- [27] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "LightGBM: A highly efficient gradient boosting decision tree," in *Proc. NIPS*, vol. 30, 2017.