



The ZevoMOS entry to VoiceMOS Challenge 2022

Adriana STAN

Communications Department,
Technical University of Cluj-Napoca, Romania

adriana.stan@com.utcluj.ro

Abstract

This paper introduces the ZevoMOS entry to the main track of the VoiceMOS Challenge 2022. The ZevoMOS submission is based on a two-step finetuning of pretrained self-supervised learning (SSL) speech models. The first step uses a task of classifying natural versus synthetic speech, while the second step's task is to predict the MOS scores associated with each training sample. The results of the finetuning process are then combined with the confidence scores extracted from an automatic speech recognition model, as well as the raw embeddings of the training samples obtained from a wav2vec SSL speech model.

The team id assigned to the ZevoMOS system within the VoiceMOS Challenge is T01. The submission was placed on the 14th place with respect to the system-level SRCC, and on the 9th place with respect to the utterance-level MSE. The paper also introduces additional evaluations of the intermediate results.

Index Terms: MOS prediction, synthetic speech, self-supervised learning, ASR confidence scores, VoiceMOS Challenge.

1. Introduction

The evaluation of text-to-speech (TTS) synthesis systems still heavily relies on performing subjective listening tests. In these tests, several factors can affect the final judgement of the listeners with respect to an audio sample, such as: *exposure*—did the listener ever listen to synthetic samples before; *context*—what other audio systems and audio samples are present in the same listening test; or *failure-to-justify*—meaning that the listener cannot pinpoint exactly why a sample is bad, it is just that it does not sound natural.

This means that, aside from the additional and lengthy effort of preparing and conducting listening tests, their results are in some cases inconclusive or error prone. Another problem with the listening tests is that the listener only rates the samples at an utterance level, commonly on a scale of 1 to 5. This discrete scale, and high level evaluation does not provide sufficient information for the TTS developer on how to improve the system. And also, in the next iterations of the TTS development, the listening tests have to be conducted again.

The most common sections of a listening test pertain to naturalness, speaker similarity and intelligibility. In recent years, the latter two sections have seen some alternative objective measures being applied to their evaluation. For the speaker similarity, large speaker-verification networks can be involved [1]. The average embeddings, or equal error rates computed on the synthesised samples and the natural samples are beginning to show acceptability in the TTS community. On the other hand, for the intelligibility, highly-accurate automatic speech recognition systems and their word error rates (WER) can enable a measure of how well the TTS system pronounced the input text [2].

As a result, the last elusive objective measure of TTS quality lies within the most complex one, the naturalness. This is also the hardest measure to implement, as even natural samples are sometimes rated unnatural when minimum levels of noise and reverberation are present in the sample. Take for example the listening tests performed for Tacotron2 [3] and FastPitch [4] where there is an 0.5 difference between the average scoring of the natural samples. Yet, finding an objective measure of audio quality is not only useful in the final evaluation of TTS or voice conversion systems, but also, if adequate, could be used as loss measures in the training process of such systems. Nowadays, the loss of DNN-based speech systems uses L1 or L2 losses, which do not truly correlate with the perceived quality or naturalness of the output speech.

There are several works which attempt to solve the objective evaluation of the synthesised samples' issue. For example MOSNet [5] uses three DNN architectures, CNN, BLSTM, and CNNBLSTM to predict the MOS scores of the Voice Conversion Challenge 2018 using frame- and utterance-level predictions of the MOS scores. MBNet [6] introduces an additional sub-network which accounts for the individual listeners' evaluations of each utterance and aggregates these scores into the final MOS prediction. A similar approach is taken by LDNet [7] where individual listener scores, as well as utterance-averaged scores are predicted. LDNet also evaluates multiple NN architectures, such as convolutional 2D networks, MobileNetV2 [8] and MobileNetV3 [9]. The training dataset is also the Voice Conversion Challenge 2018 listening test's results. [10] looks into different embeddings of the audio samples, such as deep spectrum, x-vectors and acoustic embeddings in conjunction with the MOSNet, and with both speaker-level and system-level aggregation of the MOS scores. One interesting study is that of [11] where the authors look into several self-supervised models, such as wav2vec [12], CPC [13], APC [14] and Tera [15], and finetune them in an end-to-end manner to perform the naturalness score predictions on the Voice Conversion Challenge 2018 data. An extensive analysis of the MOS predictions is performed by the authors of [16], where they compare several previously published MOS prediction networks, and pretrained self-supervised models in a task of in-domain, as well as out-of-domain MOS prediction task.

In this work, we introduce the results of our MOS prediction system, named ZevoMOS, entered into the 2022 VoiceMOS Challenge [17]. The system builds on top of the self-supervised speech learning (SSL) representations finetuned in a natural versus synthetic speech classification task, and aggregates the SSL-based predictions with the confidence scores of an ASR system. The paper is organised as follows: Section 2 introduces the general structure of the ZevoMOS system and its additional inputs. Section 3 describes the data, initial and VoiceMOS challenge results for our system, while Section 4 discusses the results and draws the conclusions.

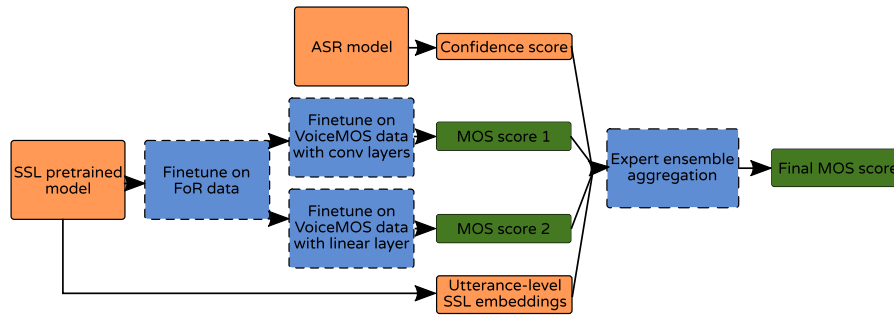


Figure 1: Overview of the ZevoMOS system’s architecture. The orange boxes mark the pretrained modules, and the blue boxes mark the modules trained for the VoiceMOS Challenge predictions.

2. ZevoMOS system description

The development of the ZevoMOS prediction system starts from the baseline networks provided by the VoiceMOS Challenge organisers. Three different systems were provided, and we selected the one described in [16] and its implementation.¹ In this baseline, the MOS predictions are learned by adding a final linear layer to a SSL model and finetuning the model in an end-to-end manner. The following subsections present the changes made to this baseline model, and the additional input and score aggregation method used in ZevoMOS.

2.1. Self-supervised learning speech models

The experiments presented in [16] use the wav2vec pretrained models available in the Fairseq [18] repository.² However, these models are trained on natural speech alone. We therefore started out by finetuning the model on the FoR dataset [19] in a natural versus synthetic speech classification task. The FoR dataset contains around 198,000 English utterances from some of the latest commercial deep-learning speech synthesizers (i.e. Amazon, Google and Baidu), as well as natural speech from the Arctic, LJSpeech and VoxForge datasets. Such that none of the synthetic samples overlap with the systems used to collect the VoiceMOS data. This was one of the rules of the challenge, which stated that the data from the ESP-Net TTS, Blizzard and Voice Conversion challenges should not be used to train the prediction models. However, the FoR dataset contains only high-quality synthetic speech samples, which are not as well correlated to, for example, the HMM-based TTS systems that are present in the Blizzard Challenge evaluations. Still, we considered that having the SSL models exposed to a limited set of artefacts which appear in the synthesised samples was important, and the numeric results support our initial hypothesis. As a result, the wav2vec hidden representations of the FoR samples averaged over the entire utterance were passed through a final linear layer to output a natural or synthetic binary classification label. The information from the classification task was enabled to flow back into the wav2vec network and update its weights. A similar finetuning step was performed for the Decoar 2.0 [20] pretrained model available in the S3PRL repository.³ The models were only trained for a few epochs (< 10) such that the bias of the classification task does not affect the learned audio

embeddings.

The resulting wav2vec and Decoar 2.0 models were then finetuned on the VoiceMOS training data. The models’ hidden representations were passed through either an average pooling layer followed by a linear layer, or a sequence of two 2D convolution layers, followed by three linear layers. The task of this finetuning step was to predict the utterance-level MOS scores. The models were trained with an early stopping criterion which looked into the MSE values computed over the dev set.

2.2. ASR-based confidence scores

Part of the evaluation of the naturalness of a synthetic sample pertains to its intelligibility. The intelligibility is also commonly evaluated in subjective listening tests. Yet, recent studies showed that high-quality ASR systems can give a good indication of the intelligibility of synthetic samples, without the definite need for a subjective evaluation of this measure [2]. The intelligibility of synthetic samples is measured as the word error rate between the input/target text and the transcriptions of that sample performed by either the listeners, or an ASR system. However, in the VoiceMOS challenge, the transcripts of the training samples were not available. We therefore relied on the confidence scores exhibited by the Alpha Cephei generic US English model⁴. The self-reported WER on the LibriSpeech dataset is 5.69%. The model selection was made also because the Alpha Cephei ASR outputs word-level confidence scores. The average of these word-level confidence scores across the entire utterance was used as additional input into the final MOS predictions made by the ZevoMOS system. It is worth mentioning here that an initial confidence score estimation was performed with the Google Speech-to-Text API.⁵ Yet, for the very poor-quality samples, the API failed to output any transcript or confidence measure. This left around 100 samples in both the training and the dev sets without this measure, and we therefore dropped the use of this API.

2.3. Expert ensemble aggregation

The individual predictions of the finetuned SSL models or the ASR-based confidence scores exhibited sub-par results on the dev data. We present these individual performances in the Results section. However, we noticed that some of the predictions made by individual components were complementary. Such that we resorted to aggregate the individual predictions in a final expert ensemble module. The ensemble included the predic-

¹<https://github.com/nii-yamagishilab/mos-finetune-ssl>

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

³<https://github.com/s3prl/s3prl>

⁴<https://alphacephei.com/vosk/models>

⁵<https://cloud.google.com/speech-to-text>

tions of the SSL finetuned models, the ASR confidence scores, and augmented with the baseline wav2vec hidden representations of the audio samples. The aggregation was performed with either a LightGBM [21] method, or a three linear layer neural net. Both structures used equal input weights for each component. A diagram of the final ZevoMOS system is shown in Figure 1, and the source code to train the system, as well as the pretrained models are available at the following link: <https://github.com/adrianastan/ZevoMOS/>.

3. Evaluation

3.1. The VoiceMOS Challenge 2022 Speech Dataset

The VoiceMOS Challenge 2022 includes two tracks: a MAIN track in which all the samples were evaluated within the same listening test by the same listeners, and an out-of-domain track (OOD), in which the samples were evaluated in a separate listening test. Also, the main track comprises English samples, while the OOD track includes Chinese synthetic samples. The ZevoMOS entry made submissions only to the MAIN track of the challenge.

The VoiceMOS Challenge 2022 data for the MAIN track is composed of a subset of utterances gathered from the past Blizzard and Voice Conversion Challenges. These samples were reevaluated in a new large-scale listening test [22]. The listening test contained samples from 187 different systems covering various speech technologies, from unit selection to end-to-end neural systems. The data is sampled at 16kHz and was released as separate training, dev and test sets. The training set comprises 4974 samples from 175 systems. The dev set contains 1066 samples from 181 systems. 6 systems (amounting to 222 samples) from the dev set were not present within the training data, however their samples were evaluated in the same large-scale listening test. These two datasets were available in the training phase of the challenge, and could be used to evaluate the initial models in the challenge’s platform. A separate dataset, the test set, containing 1066 samples from 187 systems was used to perform the final evaluation of the submissions. The target MOS scores of the test set utterances were made available only after the completion of the evaluation phase. 12 systems present in the test set, with a total of 234 utterances were not present in the training set, however 6 of these systems were present in the dev set. But during the test phase, the dev set was not supposed to be used as additional training data.

Due to compute power limitations, we had to restrict the length of the models’ input files to 6 seconds in both the training and testing procedures. This, of course, might have led to sub-optimal predictions, as some artefacts might have only occurred later in the audio sample. The data released in the VoiceMOS Challenge also included listener information and individual listener ratings. Each sample was rated by 8 listeners, but the ZevoMOS entry did not take this information into account, and only used the sample-averaged MOS scores.

3.2. Initial evaluation

Prior to submitting the ZevoMOS score predictions to the VoiceMOS Challenge scoreboard, a series of initial evaluations were performed. The objective evaluation of the VoiceMOS Challenge 2022 uses the Spearman Rank Correlation Coefficient (SRCC) and the Mean Squared Error (MSE) of the MOS scores for individual utterances, as well as the averaged scores across a system’s samples. Therefore, there are 4 objective measures to analyse: *utterance-level SRCC and MSE*, and *system-*

Table 1: *MOS scores’ prediction results on the dev dataset for the individual SSL models. The best results in each column are marked in boldface.*

Model	System		Utterance	
	SRCC	MSE	SRCC	MSE
w2v baseline	0.919	0.119	0.856	0.241
w2v-FoR-linear	0.927	0.110	0.865	0.253
w2v-FoR-conv	0.932	0.310	0.872	0.482
decoar2 baseline	0.890	0.214	0.824	0.352
decoar2-FoR-linear	0.927	0.120	0.856	0.278
decoar2-FoR-conv	0.935	0.264	0.863	0.350

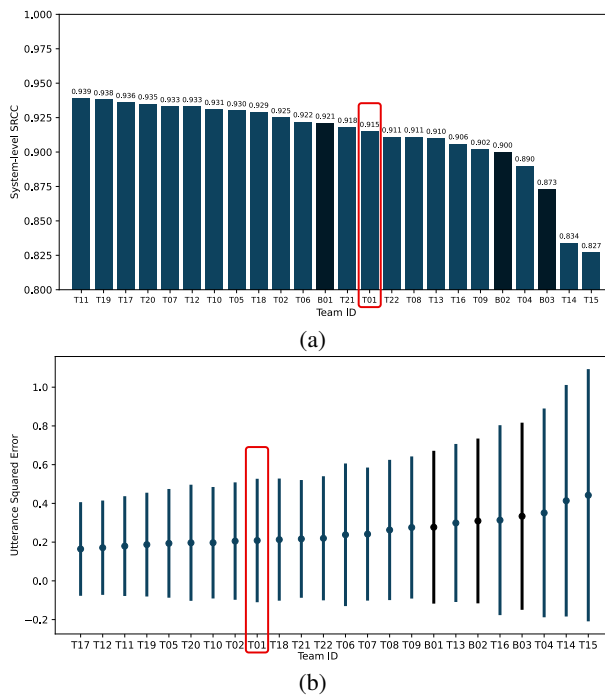


Figure 2: *VoiceMOS main track (a) system-level SRCC rankings and (b) utterance-level MSE rankings. The ZevoMOS system is under team id T01. The darker shades represent the baseline systems.*

level SRCC and MSE.

The first thing we wanted to look into were the system-level correlations between the VoiceMOS data subsets. There is an 0.917 SRCC between the training and dev sets and a 0.905 SRCC between the training and the test sets. In theory, without any additional information, and assuming similar data distributions across the datasets, these would be the top-most values that a system trained only on this data could attain. We also noticed that for part of the systems available only in the dev or test sets, there was only one sample available. In this case the system-level mean MOS value is equal to the rating of that particular sample.

We then proceeded to evaluate the intermediate models obtained in the training process of the ZevoMOS system. The results of this evaluation over the dev dataset are shown in Table 1. The models pertain to the following definitions: *w2v baseline* - the pretrained wav2vec model fine-

Table 2: Results on the *dev* and *test* data for various expert ensemble input components. The components’ numbering refer to: [A] ASR confidence scores, [B] wav2vec baseline embeddings, [C] wav2vec-FoR-linear, and [D] wav2-FoR-conv. The ZevoMOS submission includes all four components, but uses a neural network for ensembling, instead of LightGBM.

Components				Dev dataset				Test dataset			
				System		Utterance		System		Utterance	
[A]	[B]	[C]	[D]	SRCC	MSE	SRCC	MSE	SRCC	MSE	SRCC	MSE
✓				0.522	0.557	0.296	0.767	0.482	0.539	0.394	0.717
	✓			0.905	0.130	0.825	0.261	0.894	0.150	0.835	0.260
✓	✓			0.906	0.130	0.825	0.261	0.907	0.140	0.842	0.250
✓			✓	0.929	0.096	0.865	0.208	0.912	0.126	0.868	0.214
✓		✓		0.931	0.090	0.860	0.216	0.912	0.127	0.853	0.248
	✓	✓		0.935	0.091	0.873	0.198	0.911	0.124	0.870	0.219
		✓	✓	0.941	0.083	0.876	0.194	0.913	0.123	0.869	0.220
	✓	✓	✓	0.941	0.078	0.883	0.181	0.916	0.120	0.874	0.210
✓	✓	✓	✓	0.942	0.070	0.883	0.180	0.917	0.119	0.873	0.211
ZevoMOS submission				0.951	0.067	0.883	0.182	0.915	0.122	0.878	0.208

tuned on the MOS scores using a single linear output layer; *w2v-FoR-linear* - the wav2vec model finetuned on the FoR dataset and finetuned again to predict the MOS scores using a single linear output layer; *w2vFoR-conv* - the wav2vec model finetuned on the FoR dataset and finetuned again to predict the MOS scores using the sequence of 2D convolution layers followed by the three linear output layers. The same notations apply for the Decoar 2.0 model.⁶ It can be noticed that the simple finetuning of SSL models using the natural versus synthetic classification task improves the MOS score predictions’ accuracy for both the wav2vec and Decoar 2.0 models. The choice of the sequence of convolution plus linear output layers also add a marginal improvement to the SRCC measures. However, compared to the single linear output layer, the convolution sequence drastically affects the MSE values. Overall, the wav2vec model performs better using the baseline model, yet after the finetuning processes, the results of wav2vec and Decoar 2.0 models are very similar. We eventually selected the wav2vec model for the final submission, as we focused more on the utterance-level SRCC measure.

Starting from these initial results, we evaluated the different expert ensembles. The two structures, the LightGBM and linear NN aggregation methods exhibited very similar results, with only some light improvements on the NN side. Table 2 presents the results of various component combinations using the LightGBM method. However, the ZevoMOS submission (marked as such in Table 2) used the NN-based ensemble, due to its slightly better results. It can be noticed that, although no major improvement is obtained by any of the individual components, their combined knowledge is incrementally beneficial.

3.3. VoiceMOS Challenge evaluation

Figure 2 shows the results of all the submissions made to the challenge. The ZevoMOS system is under the T01 team id. With respect to utterance-level MSE, ZevoMOS was ranked 9th, and on the 14th place with respect to the system-level SRCC. The best performing systems obtained a system SRCC of 0.939 and an utterance MSE of 0.165. The results do not

⁶We should note here that although we started from the B01 baseline provided by the challenge organisers, we retrained the MOS prediction model and also used the large version of the wav2vec. Therefore, our results for this baseline differ slightly.

pertain to the same submission. It is interesting to notice that among the first 10 ranked systems w.r.t. system-level SRCC there is only a 0.01 difference. The same variation applies for the first 15 ranked systems in terms of utterance-level MSE.

The complete results of the ZevoMOS entry for both the dev and the test datasets are shown in Table 2. In the training phase, ZevoMOS was placed on the 6th place w.r.t. system-level SRCC and on the 5th place w.r.t. utterance-level MSE.⁷

4. Discussions and Conclusions

This paper introduced the ZevoMOS entry to the VoiceMOS Challenge 2022. The system uses an expert ensemble to aggregate individual MOS predictions made by self-supervised learning models. The SSL models were previously finetuned in a classification task of synthesised versus natural speech prediction. The expert ensemble also includes the confidence scores of a high-quality ASR system, as well as the baseline SSL model embeddings of the training/testing audio samples. The results placed the ZevoMOS system on the 14th place in terms of system-level SRCC and on the 9th place in terms of utterance-level MSE.

Although the results of the ZevoMOS submission are behind the best performing systems, it still supports the idea that models pretrained on large amounts of natural speech samples are informative in terms of estimating the quality of a synthetic sample. And therefore, their study w.r.t. naturalness estimation should be further explored. It is true, however, that an ideal MOS predictor would be a simple signal-based differentiable measure which would ensure its use within the training process of synthetic speech generation architectures. As future work, we plan to investigate such signal-based measures which might be indicative of quality degradations of speech without a reference signal, but also to investigate other model architectures pretrained on large amounts of natural speech data, and which might have learned abstract features pertaining to the naturalness of the output signal.

Acknowledgements. This work was funded by Zevo Technology SRL. The author would like to thank Lucian Georgescu for setting up the ASR system.

⁷<https://codalab.lisn.upsaclay.fr/competitions/695#results>

5. References

- [1] W. Xie, A. Nagrani, J. S. Chung, and A. Zisserman, "Utterance-level aggregation for speaker recognition in the wild," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 5791–5795.
- [2] J. Taylor and K. Richmond, "Confidence Intervals for ASR-based TTS Evaluation," in *Proceedings of Interspeech*, 2021.
- [3] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu, "Natural TTS Synthesis by Conditioning Wavenet on MEL Spectrogram Predictions," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4779–4783.
- [4] A. Łańcucki, "Fastpitch: Parallel text-to-speech with pitch prediction," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6588–6592.
- [5] C. Lo, S. Fu, W. Huang, X. Wang, J. Yamagishi, Y. Tsao, and H. Wang, "MOSNet: Deep Learning based Objective Assessment for Voice Conversion," in *Interspeech*, 2019.
- [6] Y. Leng, X. Tan, S. Zhao, F. Soong, X.-Y. Li, and T. Qin, "MB-NET: MOS Prediction for Synthesized Speech with Mean-Bias Network," in *ICASSP*, 06 2021, pp. 391–395.
- [7] W. Huang, E. Cooper, J. Yamagishi, and T. Toda, "LDNet: Unified Listener Dependent Modeling in MOS Prediction for Synthetic Speech," *CoRR*, vol. abs/2110.09103, 2021. [Online]. Available: <https://arxiv.org/abs/2110.09103>
- [8] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4510–4520.
- [9] A. Howard, M. Sandler, B. Chen, W. Wang, L.-C. Chen, M. Tan, G. Chu, V. Vasudevan, Y. Zhu, R. Pang, H. Adam, and Q. Le, "Searching for MobileNetV3," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 1314–1324.
- [10] J. Williams, J. Rownicka, P. Oplustil, and S. King, "Comparison of speech representations for automatic quality estimation in multi-speaker text-to-speech synthesis," in *Proceedings of SpeakerOdyssey*, 2020.
- [11] W.-C. Tseng, C.-y. Huang, W.-T. Kao, Y. Y. Lin, and H.-y. Lee, "Utilizing self-supervised representations for mos prediction," *arXiv preprint arXiv:2104.03017*, 2021.
- [12] 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.
- [13] A. Van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv e-prints*, pp. arXiv–1807, 2018.
- [14] Y.-A. Chung, W.-N. Hsu, H. Tang, and J. Glass, "An unsupervised autoregressive model for speech representation learning," in *Proc. of Interspeech*, 2019.
- [15] A. T. Liu, S.-W. Li, and H.-y. Lee, "Tera: Self-supervised learning of transformer encoder representation for speech," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2351–2366, 2021.
- [16] E. Cooper, W.-C. Huang, T. Toda, and J. Yamagishi, "Generalization Ability of MOS Prediction Networks," in *Proceedings of ICASSP 2022*, 2022.
- [17] W.-C. Huang, E. Cooper, Y. Tsao, H.-M. Wang, T. Toda, and J. Yamagishi, "The VoiceMOS Challenge 2022," in *arXiv Preprint, arXiv:2203.11389*, 2022.
- [18] 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," in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- [19] R. Reimao and V. Tzerpos, "FoR: A Dataset for Synthetic Speech Detection," in *2019 International Conference on Speech Technology and Human-Computer Dialogue (SpeD)*, 2019, pp. 1–10.
- [20] S. Ling and Y. Liu, "Decoar 2.0: Deep contextualized acoustic representations with vector quantization," *arXiv preprint arXiv:2012.06659*, 2020.
- [21] 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 *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017.
- [22] E. Cooper and J. Yamagishi, "How do Voices from Past Speech Synthesis Challenges Compare Today?" in *Speech Synthesis Workshop*, 2021.