

DISTILMOS: LAYER-WISE SELF-DISTILLATION FOR SELF-SUPERVISED LEARNING MODEL-BASED MOS PREDICTION

Jianing Yang, Wataru Nakata, Yuki Saito, and Hiroshi Saruwatari

The University of Tokyo, Japan.

ABSTRACT

With the advancement of self-supervised learning (SSL), fine-tuning pretrained SSL models for mean opinion score (MOS) prediction has achieved state-of-the-art performance. However, during fine-tuning, these SSL-based MOS prediction models often suffer from catastrophic forgetting of the pretrained knowledge and tend to overfit the training set, resulting in poor generalization performance. In this study, we propose DistilMOS, a novel method that learns to predict not only MOS but also token IDs obtained by clustering the hidden representations of each layer in the pretrained SSL model. These layer-wise token targets serve as self-distillation signals that enables the MOS prediction model to extract rich internal knowledge from SSL models, enhancing both prediction accuracy and generalization capability. Experimental evaluations demonstrate that our method significantly outperforms standard SSL-based MOS prediction models on both in-domain and out-of-domain evaluations, verifying the effectiveness and practicality of the proposed method¹.

Index Terms— MOS prediction, SSL model, knowledge extraction, speech quality assessment, self-distillation

1. INTRODUCTION

Subjective speech quality assessment is a crucial means of evaluating the performance of text-to-speech (TTS) and voice conversion (VC) systems. Among such speech quality assessment, the mean opinion score (MOS) test evaluates speech naturalness based on the average of multiple human listeners' ratings. However, manual MOS testing is time-consuming and costly, motivating the development of automatic MOS prediction models based on deep neural networks (DNNs). For instance, MOSNet [1] adopts a CNN-BLSTM architecture to predict MOS from input speech at the utterance level and achieved high correlation with human ratings on the Voice Conversion Challenge 2018 dataset [2]. Subsequent works considered listener bias during MOS test. For example, MBNet [3] and LDNet [4] take individual listener ratings as additional inputs to model inter-listener variability, thereby improving prediction accuracy. Recently, with the emergence of large-scale self-supervised learning (SSL) models pretrained on speech [5, 6, 7], leveraging feature representations from hidden layers of these models for MOS prediction has become a mainstream method [8]. SSL-based models have demonstrated strong performance on public datasets, and in the VoiceMOS Challenge 2022 [9], most top-performing systems were built upon or extended from the SSL-MOS baseline [10, 11, 12, 13].

Despite achieving impressive accuracy, existing SSL-driven MOS prediction models still face several limitations. For instance, fine-tuning SSL models solely with MOS regression may lead to catastrophic forgetting [14] of previously learned knowledge, thus impairing the model's generalization capability [15]. It has been

reported that SSL-based MOS prediction models suffer a significant performance drop in zero-shot evaluations across different datasets [8]. Many studies have shown that different layers of SSL models capture distinct levels of speech information. That is, lower layers tend to encode acoustic features such as speaker and prosodic cues, while higher layers capture more abstract information like lexical and semantic content [7, 16, 17, 18]. These hierarchical features are all critical in human perception of speech quality [19, 20]. Hence, an urgent problem is how to effectively utilize the multi-level representations of SSL models to improve model robustness and generalization across domains in MOS prediction.

To address this issue, we propose a novel MOS prediction method based on layer-wise self-distillation. Our core idea is to incorporate a self-distillation mechanism, where the MOS prediction model learns not only to predict MOS but also to reconstruct the latent feature patterns from the internal layers of the pretrained SSL model. Concretely, we apply k -means clustering [21] to the frame-level outputs of each SSL layer to generate discrete token ID sequences. During model training, we introduce token predictors, which are supervised to classify the corresponding token sequences from each layer. This multi-task training enables the MOS prediction model to focus on both the acoustic and semantic information encoded at different levels of the SSL model without requiring additional data. Experiments on the BVCC dataset [8] demonstrate that our method achieves better performance compared to baseline. Furthermore, it also performs well in zero-shot prediction on the SOMOS dataset [22], confirming the effectiveness and robustness of our method for MOS prediction tasks.

2. DISTILMOS

2.1. Overall Architecture

The overall architecture of the proposed MOS prediction model are illustrated in Fig. 1(b). As in prior work, we first extract frame-level representations from all Transformer blocks of an SSL model initialized with pretrained weights and fine-tuned during training, and then aggregate them by a learnable layer-wise weighted sum [23]. These features are then passed through a *Feature Processor*, which consists of three stacked blocks, each consisting of a linear layer, a 1D convolution (Conv1D), batch normalization, and a GELU activation [24]. This module is designed to capture local contextual information along the sequence. Because the SSL representations are high-dimensional, we insert a linear projector between the SSL encoder and the Feature Processor to reduce the dimensionality to 256.

The output of the Feature Processor is fed into a *CNN-BLSTM* module. The CNN-BLSTM structure begins with a Conv1D layer, followed by a BLSTM [25] to model temporal dependencies. The forward and backward sequence embeddings are concatenated and projected back to the original dimension via a linear layer, then

¹Code: <https://github.com/BaleYang/DistilMOS>

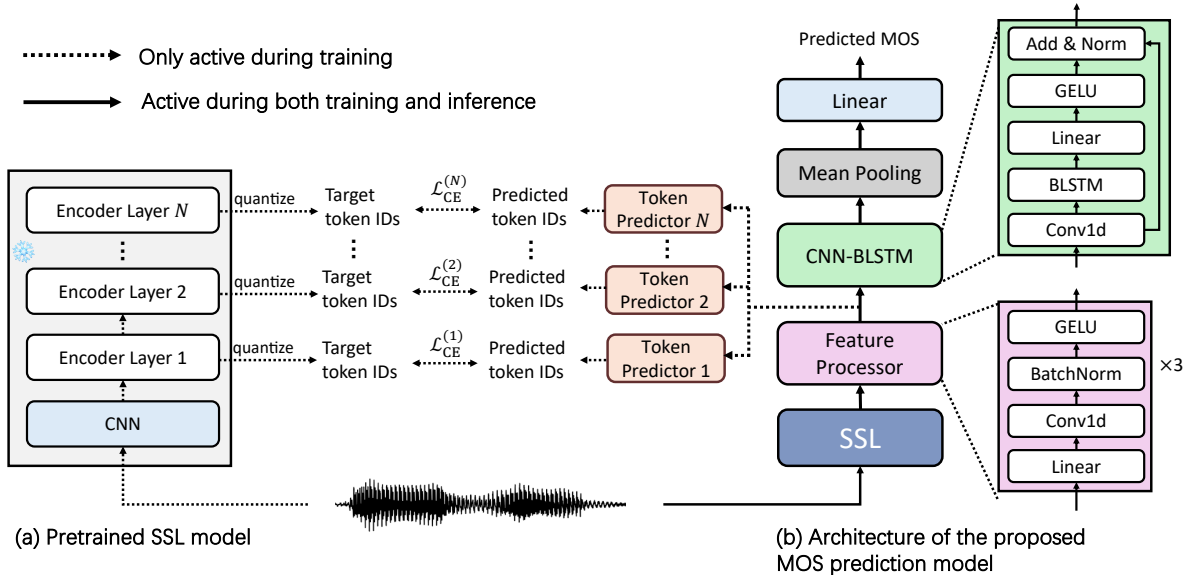


Fig. 1: Overview of the proposed method. (a) Pretrained SSL model: frame-level representations from each encoder layer are quantized via k -means clustering to generate discrete token IDs. (b) Proposed MOS prediction model architecture: during training, the model is jointly optimized with two objectives, MOS regression and token prediction for self-distillation, where token predictors learn to reconstruct the quantized token IDs from each SSL layer.

activated by GELU. Finally, a residual connection with the initial Conv1D output is applied, followed by a layer normalization [26] operation, yielding the final output of the CNN-BLSTM.

To predict the utterance-level MOS, we apply mean pooling to the sequential output from the CNN-BLSTM, followed by a linear projection. The MOS prediction is optimized with a mean squared error (MSE) loss between the ground-truth MOS s and predicted MOS \hat{s} :

$$\mathcal{L}_{\text{MOS}} = (\hat{s} - s)^2. \quad (1)$$

2.2. Self-Distillation from Discretized SSL Features

To incorporate richer internal knowledge from the pretrained SSL model, we adopt a token prediction strategy. As illustrated in Fig. 1(a), for each layer of the pretrained SSL model, we apply k -means clustering to the frame-level representations to generate discrete token sequences for each utterance. These cluster assignments serve as pseudo labels representing difference speech features of each layer.

During training, the output from the Feature Processor is fed into N parallel *Token Predictors* (one for each SSL layer), where each predictor consists of a 3-layer multi-layer perceptron (MLP) with GELU activations. These modules are trained to classify the token IDs obtained from the corresponding SSL layer.

The token classification for the n -th layer is optimized to minimize a cross-entropy loss:

$$\mathcal{L}_{\text{CE}}^{(n)} = -\frac{1}{L} \sum_{i=1}^L \log p_i^{(n)}, \quad (2)$$

where L denotes the sequence length and $p_i^{(n)}$ is the predicted probability of the correct token at position i for layer n .

The motivation behind this self-distillation strategy is that human MOS perception depends on various levels of speech features, such as prosody, speaking rate, and intelligibility [19]. The token

prediction task implicitly encourages the MOS prediction model to learn such acoustic and semantic features distributed across different SSL layers [16]. When these representations are fused into the MOS prediction head, they provide a stronger basis for subjective quality estimation.

Moreover, since the token labels are derived from the pretrained SSL model, no additional human-annotated data is required. This makes our method efficient and scalable. The token prediction branch can be discarded during inference, incurring no extra computational cost. This design also mitigates the risk of in-domain overfitting caused by direct MOS-only optimization and improves generalization to unseen domains.

Unlike conventional distillation methods such as DistilHuBERT [27], which aim to predict the exact embeddings of each SSL layer, our method instead focuses on learning essential representations for MOS prediction while discarding irrelevant details. This design prioritizes task-relevant abstraction rather than full feature preservation.

2.3. Training Objective

The final training objective combines MOS regression and token classification across N SSL layers:

$$\mathcal{L} = \mathcal{L}_{\text{MOS}} + \alpha \cdot \frac{1}{N} \sum_{n=1}^N \mathcal{L}_{\text{CE}}^{(n)}, \quad (3)$$

where α is a weighting coefficient that balances the contribution of the token-level self-distillation loss.

3. EXPERIMENTAL EVALUATION

3.1. Datasets

Our experiments were conducted on the BVCC dataset [8], the official corpus for the main track of the VoiceMOS Challenge 2022 [9].

Table 1: Experimental results on BVCC (in-domain) and SOMOS (zero-shot) datasets with different SSL backbones. **Bold** value indicates the best result for each SSL backbone.

SSL backbone	Model	BVCC								SOMOS			
		Utterance-level				System-level				Utterance-level			
		LCC	SRCC	KTAU	MSE	LCC	SRCC	KTAU	MSE	LCC	SRCC	KTAU	MSE
Wav2Vec2 Base	SSL-MOS	0.869	0.867	0.687	0.472	0.926	0.923	0.769	0.301	0.320	0.310	0.210	0.654
	DistilMOS (ours)	0.884	0.883	0.711	0.186	0.934	0.934	0.782	0.098	0.344	0.318	0.217	0.960
	w/o token prediction	0.872	0.871	0.695	0.231	0.928	0.927	0.774	0.143	0.316	0.305	0.207	1.403
	MSE distillation	0.884	0.881	0.709	0.187	0.932	0.930	0.777	0.089	0.277	0.255	0.173	0.940
WavLM Base	SSL-MOS	0.869	0.870	0.694	0.354	0.917	0.916	0.763	0.210	0.353	0.336	0.230	0.973
	DistilMOS (ours)	0.886	0.884	0.712	0.190	0.931	0.928	0.778	0.087	0.396	0.375	0.257	1.051
	w/o token prediction	0.881	0.882	0.711	0.190	0.925	0.922	0.768	0.094	0.312	0.297	0.203	0.796
	MSE distillation	0.874	0.871	0.695	0.214	0.923	0.919	0.763	0.098	0.383	0.362	0.248	0.935
WavLM Base+	SSL-MOS	0.862	0.870	0.691	0.281	0.917	0.927	0.779	0.142	0.315	0.287	0.196	0.317
	DistilMOS (ours)	0.885	0.885	0.713	0.191	0.936	0.932	0.782	0.083	0.396	0.370	0.253	0.998
	w/o token prediction	0.879	0.878	0.705	0.193	0.931	0.927	0.775	0.090	0.320	0.297	0.202	0.748
	MSE distillation	0.882	0.881	0.709	0.193	0.927	0.922	0.773	0.091	0.277	0.360	0.225	0.724
WavLM Large	SSL-MOS	0.884	0.881	0.708	0.195	0.934	0.930	0.784	0.104	0.197	0.188	0.127	1.829
	DistilMOS (ours)	0.891	0.889	0.720	0.188	0.939	0.936	0.791	0.105	0.404	0.387	0.264	0.965
	w/o token prediction	0.888	0.887	0.718	0.180	0.934	0.932	0.782	0.088	0.393	0.375	0.256	0.823
	MSE distillation	0.887	0.886	0.715	0.184	0.933	0.930	0.779	0.097	0.387	0.370	0.252	1.200

BVCC consists of 7,106 English synthetic speech utterances generated from 187 different systems across multiple Blizzard Challenges [28] and Voice Conversion Challenges [2]. Each utterance was rated by 8 listeners for naturalness on a 1–5 scale. We followed the official split, using 4,974 utterances for training, 1,066 for validation, and 1,066 for testing.

To further examine the robustness of our MOS prediction model, we adopted a zero-shot evaluation setting on the SOMOS dataset (without fine-tuning) [22]. SOMOS is a large-scale open-source corpus for subjective naturalness assessment of neural TTS systems. It contains 20,000 synthetic utterances generated from 200 acoustic models trained on the LJSpeech dataset [29] (mainly Tacotron-based [30]), along with 100 natural utterances. Each utterance was rated by at least 17 different listeners on a 1–5 scale.

3.2. Evaluation Metrics

For in-domain experiments on the BVCC dataset, we report results at both the utterance level and the system level using four metrics: MSE, linear correlation coefficient (LCC), Spearman rank correlation coefficient (SRCC), and Kendall’s τ (KTAU). Higher correlation coefficients indicate better alignment between predicted and human ratings, while lower MSE indicates smaller prediction errors. For out-of-domain zero-shot evaluation on SOMOS, we report only utterance-level metrics, which impose stricter requirements on generalization to unseen utterances and ranking consistency.

3.3. Implementation Details

To verify the general applicability of our method, we conducted experiments using four different SSL models: Wav2Vec2 Base [5], WavLM Base, WavLM Base+, and WavLM Large [7]. Wav2Vec2 Base and WavLM Base are of similar size and trained on the same datasets, but differ in pretraining objectives: Wav2Vec2 uses contrastive learning, whereas WavLM Base follows HuBERT-style masked token prediction [6]. WavLM Base and WavLM Base+ share the same architecture, but WavLM Base+ was pretrained on a larger dataset, thus possessing richer knowledge. WavLM Base+ and WavLM Large were trained on the same dataset, but differ in model size.

All linear, Conv1D, and BLSTM modules within the Feature Processor and CNN-BLSTM were set to 256 dimensions. For clustering, we applied k -means clustering ($k = 200$) to each SSL layer’s outputs on the BVCC training set to obtain token IDs for each utterance. To accelerate centroid estimation, we used MiniBatchK-Means [31] with the online partial fit routine, updating clusters with batch size 64. We empirically set the token distillation weight to $\alpha = 0.1$. Training was performed for 10,000 steps with a batch size of 32 on a single NVIDIA GH200 (120 GB) GPU. Checkpoints were saved every 1,000 steps, and the model with the highest utterance-level SRCC on the validation set was selected for evaluation on the test set.

We used the AdamW optimizer [32] with a OneCycleLR scheduler [33]. Hyperparameters were set as: learning rate = $1e-4$, betas = (0.9, 0.98), weight decay = $1e-4$, and gradient clipping threshold = 10.0.

3.4. Baseline and Ablation Studies

Although our proposed method can be applied to any MOS prediction models that use pretrained SSL models for feature extraction, we adopted SSL-MOS [8] as the baseline for simplicity. To eliminate potential bias introduced by differences in SSL backbones, we fine-tuned each of the four SSL models described in Section 3.3 and compared their performance against our proposed method initialized with the same backbones.

To further validate the effectiveness of our method, we conducted ablation studies. Specifically, for each configuration, we trained a variant without token prediction, i.e., using only the proposed architecture optimized with the MSE loss (Eq. (1)). In addition, to demonstrate the effectiveness of predicting discrete tokens, we performed another variant where, instead of token prediction, the model was equipped with N prediction heads (one per SSL layer) to predict the original SSL embeddings directly, and optimized with an MSE loss.

We further examine the sensitivity of our method to the number of clusters used for SSL layer discretization. Using WavLM Base as the backbone, we vary $k \in \{50, 100, 150, 200, 250, 300\}$ and report the corresponding SRCC on both BVCC (in-domain) and SOMOS




Fig. 2: Effect of the number of clusters k on SRCC performance. Left: utterance-level SRCC on BVCC (in-domain) dataset. Middle: system-level SRCC on BVCC dataset. Right: utterance-level SRCC on SOMOS (zero-shot) dataset.

(zero-shot) dataset.

3.5. Results

The experimental results in Table 1 show that our method surpasses SSL-MOS across all backbones and metrics on BVCC, with the large gains on system-level and utterance-level SRCC. On SOMOS (zero-shot), it improves all correlation metrics, indicating better in-domain performance and stronger robustness to unseen conditions. The relatively high MSE on SOMOS is likely caused by cross-dataset score calibration differences, under which correlation-based metrics better reflect perceptual consistency than absolute error.

The results of the ablation studies further validate the contribution of token prediction. Removing token prediction and training the model solely with the MOS regression loss (w/o token prediction) led to noticeable performance degradation. Moreover, replacing token prediction with prediction heads that regress the SSL embeddings directly (MSE distillation) also resulted in performance declines, with all correlation-based metrics falling significantly behind the token prediction method. We speculate that this is because MSE-based distillation makes the MOS prediction model overfit to the pretrained SSL features, leading to poor performance.

The results of varying the number of clusters are shown in Fig. 2. We observe that the cluster size does not significantly affect performance, as consistently strong results are maintained across all values of k . Notably, in the zero-shot SOMOS evaluation, all cluster configurations surpass both the baseline and ablation variants, further demonstrating the robustness and generalization capability of our proposed method.

3.6. Analysis of Learned Representations

To examine whether our token-prediction self-distillation method enables the MOS prediction model to capture rich speech features from the pretrained SSL backbone, we conducted canonical correlation analysis (CCA) [34]. Specifically, for both MOS prediction model of the proposed method and the ablation baselines, we computed the CCA between the average-pooled outputs of the Feature Processor and the average-pooled representations from each Transformer layer of the pretrained SSL model. We additionally performed the same analysis on a randomly initialized MOS prediction model of the proposed model to assess whether architectural biases affect the results. For the SSL-MOS baseline, we instead computed the CCA between the average-pooled output of its final Transformer layer with the average-pooled representations from each Transformer layer of the pretrained SSL model. All experiments were conducted using WavLM Base as the backbone. We




Fig. 3: CCA curves between MOS prediction model representations and pretrained SSL layer outputs.

computed the CCA following [35]² on BVCC test set.

As shown in Fig. 3, both self-distillation method (DistilMOS and MSE distillation) maintain consistently high CCA values across layers, indicating that self-distillation successfully encourages the MOS prediction model to learn rich speech features present from the pretrained SSL model. In contrast, the MOS-only fine-tuned model (SSL-MOS and w/o token prediction) exhibits correlations close to random initialization, suggesting severe catastrophic forgetting. This finding provides further evidence that catastrophic forgetting undermines generalization in zero-shot scenarios.

4. CONCLUSION

In this paper, we proposed a novel method that enables the model to predict token IDs derived from clustering its own intermediate representations, thereby extracting and leveraging its internal knowledge. This self-distillation strategy effectively enhances both performance and generalization capability. Our method achieves superior results compared to baselines on both BVCC (in-domain) and SOMOS (zero-shot) datasets, and the ablation studies further validate its feasibility and robustness.

Since BVCC is a relatively small dataset, the clustering results may exhibit mismatches when applied to other datasets. As part of future work, we plan to extend this method to larger or combined datasets to further refine the method and develop a MOS prediction model with stronger generalization ability.

Acknowledgements: The work was supported by JST Moonshot Grant Number JPMJMS2011 and based on results obtained from a project outsourced by the New Energy and Industrial Technology Development Organization (NEDO).

²<https://github.com/ankitapasad/layerwise-analysis/>

5. REFERENCES

- [1] Chen-Chou Lo, Szu-Wei Fu, Wen-Chin Huang, Xin Wang, Junichi Yamagishi, Yu Tsao, and Hsin-Min Wang, "MOSNet: Deep Learning based Objective Assessment for Voice Conversion," in *Proc. Interspeech*, 2019, pp. 1541–1545.
- [2] Jaime Lorenzo-Trueba, Junichi Yamagishi, Tomoki Toda, Daisuke Saito, Fernando Villavicencio, Tomi Kinnunen, and Zhenhua Ling, "The Voice Conversion Challenge 2018: Promoting development of parallel and nonparallel methods," in *Proc. Speaker Odyssey*, 2018, pp. 195–202.
- [3] Yichong Leng, Xu Tan, Sheng Zhao, Frank Soong, Xiang-Yang Li, and Tao Qin, "MBNet: MOS prediction for synthesized speech with mean-bias network," in *Proc. ICASSP*, 2021, pp. 391–395.
- [4] Wen-Chin Huang, Erica Cooper, Junichi Yamagishi, and Tomoki Toda, "LDNet: Unified listener dependent modeling in MOS prediction for synthetic speech," in *Proc. ICASSP*, 2022, pp. 896–900.
- [5] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proc. NeurIPS*, 2020, pp. 12449–12460.
- [6] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 29, pp. 3451–3460, 2021.
- [7] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al., "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [8] Erica Cooper, Wen-Chin Huang, Tomoki Toda, and Junichi Yamagishi, "Generalization ability of MOS prediction networks," in *Proc. ICASSP*, 2022, pp. 8442–8446.
- [9] Wen-Chin Huang, Erica Cooper, Yu Tsao, Hsin-Min Wang, Tomoki Toda, and Junichi Yamagishi, "The VoiceMOS Challenge 2022," in *Proc. Interspeech*, 2022, pp. 4536–4540.
- [10] Takaaki Saeki, Detai Xin, Wataru Nakata, Tomoki Koriyama, Shinnosuke Takamichi, and Hiroshi Saruwatari, "UTMOS: UTokyo-SaruLab System for VoiceMOS Challenge 2022," in *Proc. Interspeech*, 2022, pp. 4521–4525.
- [11] Wei-Cheng Tseng, Wei-Tsung Kao, and Hung-yi Lee, "DDOS: A MOS prediction framework utilizing domain adaptive pre-training and distribution of opinion scores," in *Proc. Interspeech*, 2022, pp. 4541–4545.
- [12] Xiaohai Tian, Kaiqi Fu, Shaojun Gao, Yiwei Gu, Kai Wang, Wei Li, and Zejun Ma, "A transfer and multi-task learning based approach for MOS prediction," in *Proc. Interspeech*, 2022, pp. 5438–5442.
- [13] Zhengdong Yang, Wangjin Zhou, Chenhui Chu, Sheng Li, Raj Dabre, Raphael Rubino, and Yi Zhao, "Fusion of self-supervised learned models for MOS prediction," in *Proc. Interspeech*, 2022, pp. 5443–5447.
- [14] Zhiyuan Chen and Bing Liu, "Continual learning and catastrophic forgetting," in *Lifelong Machine Learning*, pp. 55–75. Springer, 2022.
- [15] Xiaohui Zhang, Jiangyan Yi, Jianhua Tao, Chenglong Wang, and Chu Yuan Zhang, "Do you remember? Overcoming catastrophic forgetting for fake audio detection," in *Proc. ICML*, 2023, pp. 41819–41831.
- [16] Ankita Pasad, Ju-Chieh Chou, and Karen Livescu, "Layer-wise analysis of a self-supervised speech representation model," in *Proc. ASRU*, 2021, pp. 914–921.
- [17] Anton de la Fuente and Dan Jurafsky, "A layer-wise analysis of Mandarin and English suprasegmentals in SSL speech models," in *Proc. Interspeech*, 2024, pp. 1290–1294.
- [18] Kwanghee Choi, Ankita Pasad, Tomohiko Nakamura, Satoru Fukayama, Karen Livescu, and Shinji Watanabe, "Self-Supervised Speech Representations are More Phonetic than Semantic," in *Proc. Interspeech*, 2024, pp. 4578–4582.
- [19] Erica Cooper, Wen-Chin Huang, Yu Tsao, Hsin-Min Wang, Tomoki Toda, and Junichi Yamagishi, "A review on subjective and objective evaluation of synthetic speech," *Acoustical Science and Technology*, vol. 45, no. 4, pp. 161–183, 2024.
- [20] Alexandra Vioni, Georgia Maniati, Nikolaos Ellinas, June Sig Sung, Inchul Hwang, Aimilios Chalamandaris, and Pirros Tsiakoulis, "Investigating content-aware neural text-to-speech MOS prediction using prosodic and linguistic features," in *Proc. ICASSP*, 2023.
- [21] John A Hartigan and Manchek A Wong, "Algorithm AS 136: A k-means clustering algorithm," *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol. 28, no. 1, pp. 100–108, 1979.
- [22] Georgia Maniati, Alexandra Vioni, Nikolaos Ellinas, Karolos Nikitaras, Konstantinos Klapsas, June Sig Sung, Gunu Jho, Aimilios Chalamandaris, and Pirros Tsiakoulis, "SOMOS: The Samsung Open MOS Dataset for the Evaluation of Neural Text-to-Speech Synthesis," in *Proc. Interspeech*, 2022, pp. 2388–2392.
- [23] Kaito Baba, Wataru Nakata, Yuki Saito, and Hiroshi Saruwatari, "The T05 system for the VoiceMOS Challenge 2024: Transfer learning from deep image classifier to naturalness MOS prediction of high-quality synthetic speech," in *Proc. SLT*, 2024, pp. 818–824.
- [24] Dan Hendrycks and Kevin Gimpel, "Gaussian error linear units (GELUs)," *arXiv preprint arXiv:1606.08415*, 2016.
- [25] Sepp Hochreiter and Jürgen Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [26] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton, "Layer Normalization," *CoRR*, vol. abs/1607.06450, 2016.
- [27] Heng-Jui Chang, Shu-wen Yang, and Hung-yi Lee, "DistilHuBERT: Speech representation learning by layer-wise distillation of hidden-unit BERT," in *Proc. ICASSP*, 2022, pp. 7087–7091.
- [28] Simon King and Vasilis Karaiskos, "The Blizzard Challenge 2009," Blizzard Challenge workshop, University of Edinburgh, 2009.
- [29] Keith Ito and Linda Johnson, "The LJ Speech Dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [30] Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al., "Tacotron: Towards end-to-end speech synthesis," in *Proc. Interspeech*, 2017, pp. 4006–4010.
- [31] David Sculley, "Web-scale k-means clustering," in *Proc. WWW*, 2010, pp. 1177–1178.
- [32] Ilya Loshchilov and Frank Hutter, "Decoupled Weight Decay Regularization," in *Proc. ICLR*, 2017.
- [33] Leslie N Smith and Nicholay Topin, "Super-convergence: Very fast training of neural networks using large learning rates," in *Artificial intelligence and machine learning for multi-domain operations applications*, 2019, vol. 11006, pp. 369–386.
- [34] Harold Hotelling, "Relations between two sets of variates," in *Breakthroughs in statistics: methodology and distribution*, pp. 162–190. Springer, 1992.
- [35] Ankita Pasad, Bowen Shi, and Karen Livescu, "Comparative layer-wise analysis of self-supervised speech models," in *Proc. ICASSP. IEEE*, 2023, pp. 1–5.