



Acoustic Feature Mixup for Balanced Multi-aspect Pronunciation Assessment

Heejin Do¹, Wonjun Lee², Gary Geunbae Lee^{1,2}

¹Graduate School of AI, POSTECH, South Korea

²Department of Computer Science and Engineering, POSTECH, South Korea

{heejindo, leeljun, gblee}@postech.ac.kr

Abstract

In automated pronunciation assessment, recent emphasis progressively lies on evaluating multiple aspects to provide enriched feedback. However, acquiring multi-aspect-score labeled data for non-native language learners' speech poses challenges; moreover, it often leads to score-imbalanced distributions. In this paper, we propose two *Acoustic Feature Mixup* strategies, linearly and non-linearly interpolating with the in-batch averaged feature, to address data scarcity and score-label imbalances. Primarily using goodness-of-pronunciation as an acoustic feature, we tailor mixup designs to suit pronunciation assessment. Further, we integrate fine-grained error-rate features by comparing speech recognition results with the original answer phonemes, giving direct hints for mispronunciation. Effective mixing of the acoustic features notably enhances overall scoring performances on the speechocean762 dataset, and detailed analysis highlights our potential to predict unseen distortions.

Index Terms: pronunciation assessment, multi-aspect pronunciation assessment, computer-assisted pronunciation training

1. Introduction

Assisting non-native (L2) language learners to acquire foreign speaking skills, automatic pronunciation assessment is pivotal for computer-assisted pronunciation training (CAPT) systems [1, 2]. Recently, moving beyond solely evaluating phone-level scores [3, 4, 5, 6], assessing pronunciation on multiple aspects and granularities has attracted increasing attention [7, 8, 9, 10]. To achieve multi-aspect pronunciation assessment via deep learning techniques, qualified data with labeled multi-aspect scores for learner utterances is required.

However, obtaining multi-dimensional score-labeled speech data poses challenges, and score labels are prone to have imbalanced distributions [11, 12], often failing to represent real-world minority cases. Such imbalanced training data skewed towards specific scores significantly degrades the model performance on samples with new or unseen score ranges [11]. For instance, a model trained on a biased dataset where most cases are labeled around the 2-point range may struggle to predict samples of other score ranges. Indeed, recent advancements in multi-aspect pronunciation assessment have yielded notable performance enhancements via meticulously crafted deep neural modeling [7, 8, 10] and extensive utilization of acoustic feature input [9]. However, a substantial gap persists between severely score-imbalanced aspects and others, exceeding fourfold.

In this paper, we propose two *Acoustic-feature Mixup* (AM) strategies to simulate distribution shifts toward scarce positions without original speech data, thereby guiding the balanced learning for multiple scoring dimensions. Mixup [13] is an ap-

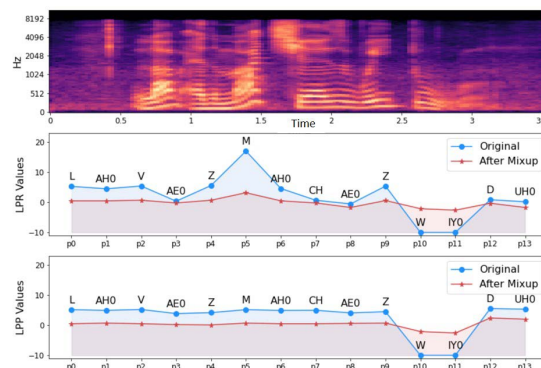


Figure 1: An example of GOP features, log phone posterior (LPP) and log posterior ratio (LPR), shift after applying dynamic Mixup.

proach that interpolates data samples to aid in model regularization and has primarily been applied for image classification tasks [14, 15, 16]. Distinct from its typical use, we suggest suitable methods for acoustic features and regression of continuous numeric labels for pronunciation assessment, where the utility is yet to be explored. In particular, we present two AM strategies: 1) static AM, which involves linear and simple combinations, and 2) dynamic AM, which integrates non-linear interpolations. Unlike existing approaches, where mixing policies are solely applied for two pairs, we consider all pairs within a batch by incorporating in-batch averaged values within the policy.

We mainly leverage the Goodness of Pronunciation (GOP) feature as the acoustic feature, which is determined by comparing the phone-level pronunciations of the learner and the correct answer. As GOP provides details on mispronounced phonemes, it has been widely used for pronunciation assessment. Our methods mix GOP features rather than the original speech data, allowing the generation of inputs that match the discriminative regions for grading without specific score-labeled speech data (Figure 1). Further, we introduce multi-granular error rate features obtained from the automatic speech recognition (ASR) system. Specifically, we measure the character- and token-level match error rate between ASR results and the correct phonemes of the utterance and concatenate it with the final representation vector, thus providing direct hints for mispronunciation. Mixing up these error-rate features in parallel with GOP features further assists the model training.

Extensive experiments on the publicly available speechocean762 dataset demonstrate the training assistance of two AM strategies on the multi-aspect pronunciation assessment framework. The original dataset exhibits severely imbalanced

score distributions for aspects such as *Stress* and *Completeness*, a major contributor to the low performance in these aspects [11]. Visualizing how the proposed mixup technique shifts the existing distribution demonstrates the ability to synthesize discriminative samples. Remarkably improved performance on imbalanced aspects further suggests that AM plays a complementary role in addressing vulnerabilities in unseen score samples; thus, it assists the system in achieving aspect-wise balanced scoring.

2. Related work

Although multi-aspect pronunciation assessment has achieved recent success [7, 8, 9, 10], this success has been limited to aspects where the score labels of the training data are evenly distributed. The inferior performance on a specific aspect might be attributed to its highly imbalanced score-label distributions, with the majority of samples having high scores [11, 10]. As scores in real-world scenarios are likely to be distributed diversely, addressing such imbalances is crucial. Recent related attempts focused on training optimization by either assigning balanced weights [17] or designing balanced loss functions [11]. However, there has been no direct research attempting data shift, and solely optimizing training with existing data may be susceptible to potential distortion encountered in practical use. We aim to achieve robustness even with unseen range data by synthesizing data in the latent space.

Mixup [13] is renowned for aiding model regularization by interpolating between data samples, particularly when labeled data is scarce or not representative [15, 18, 16]. Existing studies revealed that data distribution shift effectively enhances the robustness of DNNs against adversarial samples while reducing overconfident predictions [19, 20, 18]. Diverse shift policies on mixups have been extensively studied for visual classification tasks [21, 22, 23], but their use for pronunciation assessment has yet to be explored. Building upon these benefits, we suggest adopting a mixup for multi-aspect pronunciation assessment to overcome training difficulties induced by biased score labels.

3. Acoustic feature mixup

3.1. Mixup policy

To effectively shift the distribution of existing data skewed on specific score ranges (Figure 2) and synthesize corresponding pseudo acoustic features, we introduce two AM strategies, which are static (AM_{stat}) and dynamic (AM_{dyn}). Both methods employ the average feature values of the entire samples within a mini-batch for more stabilized training; however, static AM considers simple linear transformation, while dynamic AM further incorporates non-linearity.

3.1.1. Static AM

We intuitively explore a straightforward linear data transformation, which shifts the distribution in parallel. Given the i -th sample, where x_i denotes its acoustic feature and $y_i \in \mathbb{R}^m$ represents its corresponding score vector encompassing m distinct aspects, we compute the averaged acoustic feature $a_x = \frac{1}{b} \sum_{i=1}^b x_i$ and the averaged score label $a_y = \frac{1}{b} \sum_{i=1}^b y_i$ over a mini-batch of size b . AM_{stat} linearly interpolates x_i and y_i with a_x and a_y using a mixup ratio λ as follows:

$$\tilde{x} = x - \lambda \cdot a_x \quad (1)$$

$$\tilde{y} = y - \lambda \cdot a_y \quad (2)$$

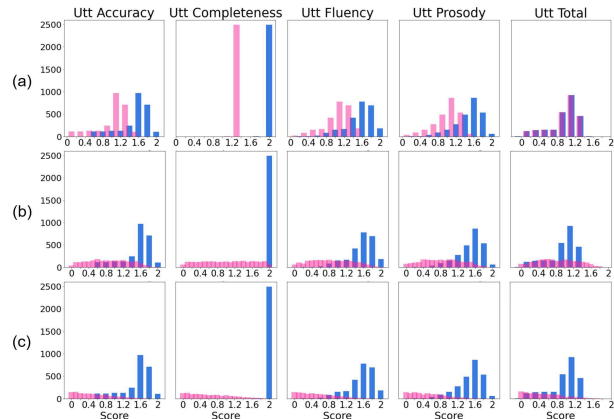


Figure 2: The utterance-level score-label distribution shift when AM_{stat} with fixed $\lambda=0.3$ (a), AM_{stat} with $\lambda \sim Beta(\alpha, \alpha)$ (b), and AM_{dyn} (c) are applied, respectively. blue and pink bars denote original and mixed-up distribution, respectively.

where λ is a randomly sampled weight from a $Beta(\alpha, \alpha)$ distribution. Figure 2 illustrates that selecting lambda from a beta distribution (b) instead of a fixed constant lambda, regardless of the sample (a), helps achieve more evenly distributed pseudo labels. The synthesized pseudo acoustic feature and label pairs, (\tilde{x}, \tilde{y}) , are then used for training along with the original data. Note that only mixed-up samples with labels within the range of 0 to 2 are utilized for training.

3.1.2. Dynamic AM

Emphasizing the importance of capturing intricate elements in distorted images, cutting-edge techniques for visual tasks applied dynamic mixup, which considers non-linearity existing between the samples [24, 15]. Motivated by their works and particularly tailoring for pronunciation assessment, we design a novel dynamic acoustic feature mixup policy. Specifically, we devise a non-linear interpolation between the given sample and the mini-batch mean value to shift them into a latent space. With two mixing weights λ_1 and λ_2 , which are separately and randomly derived from a $Beta(\alpha, \alpha)$ distribution, the AM_{dyn} is defined as follows:

$$\tilde{x} = \lambda_1 x - \lambda_2 a_x + \lambda_1 \lambda_2 (x - a_x) \quad (3)$$

$$\tilde{y} = \lambda_1 y - \lambda_2 a_y + \lambda_1 \lambda_2 (y - a_y) \quad (4)$$

where x_i , y_i , a_x , and a_y are defined same as AM_{stat} .

3.2. Acoustic features

As the primary acoustic feature, we adopt the GOP feature instead of the original speech data. We follow the process outlined in [25, 7] for GOP feature generation. Specifically, the speech audio and its canonical transcription are first given to the acoustic model, yielding a sequence of phonetic posterior probabilities. Subsequently, following phoneme-level force alignment, these probabilities are converted into 84-dimensional GOP features. The dimensionality 84 stems from the concatenation of log phone posterior (LPP) and log posterior ratio (LPR), each comprising 42 dimensions, calculated for each of the 42 source phones within the Librispeech acoustic model. The LPP of a phone φ and LPR of observing phone φ_j given phone φ_i are

Table 1: Averaged MSE (for phoneme level) and PCC scores (for all levels) with standard deviation across five runs. **Acc** and **Comp** are the Accuracy and Completeness, respectively. *GOPT-imp* is the result of our implemented version of GOPT. +ER denotes the addition of error-rate features. **Bold** and underline denote the best and the second-best performance in each column, respectively.

	Model	Phoneme Score		Word Score (PCC)			Utterance Score (PCC)				
		Acc(MSE ↓)	Acc(PCC ↑)	Acc ↑	Stress ↑	Total ↑	Acc ↑	Comp ↑	Fluency ↑	Prosody ↑	Total ↑
Baseline	LSTM	0.089 ±0.002	0.587 ±0.014	0.511 ±0.014	0.297 ±0.012	0.524 ±0.011	0.717 ±0.004	0.123 ±0.143	0.741 ±0.01	0.744 ±0.006	0.743 ±0.006
	GOPT	0.085 ±0.001	0.612 ±0.003	0.533 ±0.004	0.291 ±0.030	0.549 ±0.002	0.714 ±0.004	0.155 ±0.039	0.753 ±0.008	0.760 ±0.006	0.742 ±0.005
	GOPT-imp	0.086 ±0.001	0.608 ±0.004	0.529 ±0.005	0.292 ±0.036	0.544 ±0.006	0.712 ±0.005	0.217 ±0.091	0.755 ±0.003	0.756 ±0.003	0.737 ±0.005
Ours	AM_{stat}	0.085 ±0.001	0.611 ±0.007	0.532 ±0.009	0.347 ±0.008	0.551 ±0.006	0.723 ±0.007	0.281 ±0.090	0.769 ±0.004	0.766 ±0.003	0.752 ±0.007
	+ER	0.085 ±0.001	<u>0.614</u> ±0.005	<u>0.538</u> ±0.005	0.306 ±0.009	0.558 ±0.005	<u>0.735</u> ±0.001	<u>0.402</u> ±0.085	<u>0.780</u> ±0.002	<u>0.779</u> ±0.003	<u>0.764</u> ±0.005
	AM_{dyn}	0.086 ±0.001	0.609 ±0.007	0.531 ±0.009	<u>0.332</u> ±0.022	0.547 ±0.009	0.726 ±0.003	0.403 ±0.130	0.769 ±0.004	0.765 ±0.004	0.753 ±0.003
	+ER	0.084 ±0.001	0.617 ±0.004	0.539 ±0.003	0.317 ±0.027	<u>0.557</u> ±0.004	0.738 ±0.002	0.392 ±0.182	0.782 ±0.002	0.780 ±0.001	0.768 ±0.003

defined as follows [7]:

$$LPP(\varphi) \approx \frac{1}{t_e - t_s + 1} \sum_{t_s}^{t_e} \log p(\varphi|o_t) \quad (5)$$

$$LPR(\varphi_j|\varphi_i) = \log p(\varphi_j|o; t_s, t_e) - \log p(\varphi_i|o; t_s, t_e) \quad (6)$$

where o_t is the input observation of the frame t , and the start and end frame indexes are t_s and t_e , respectively.

In addition, we incorporate fine-grained error rate features to provide the model with direct information about mispronunciations. Considering that correct phonemes for the utterances learners need to mimic are provided, we compare the learner’s ASR-hypothesized phonemes to the reference answer phonemes to extract the error rate. Specifically, we use the character error rate (CER) and the match error rate (MER). CER is measured by dividing the number of missed characters by the number of characters in the reference. MER is calculated by dividing the number of missed tokens (phonemes in our work) by the total number of tokens in the union of the hypothesis and reference. While CER focuses on individual character errors, MER focuses on correct phoneme matches. The extracted error rates are concatenated with the model representation before passing to the final linear layer for each aspect score prediction.

3.3. Loss function

For training, we employ the mean squared error (MSE) loss, a widely utilized function for the pronunciation assessment task [7, 8, 9]. The overall loss is determined by aggregating the individual losses at each granularity level, where each loss represents the multi-aspect-averaged value within that level. The total loss is defined as follows:

$$MSE_{total} = \sum \frac{1}{N} \sum^N MSE_{mn} \quad (7)$$

given the M granularity levels and N aspects. In this work, 3 levels of granularity and 9 aspects are applied.

4. Experiments

We evaluate our AM methods on the open-source speechocean762 ([26]) dataset, which includes the speech data of non-native language learners and the corresponding labeled

multi-aspect scores. While its multifaceted labeled scores on multi-granular levels provide diverse opportunities for the multi-aspect pronunciation assessment, they have severely imbalanced labels, particularly for specific aspects. The dataset comprises 2500 utterances of training and test sets, respectively. We employ the fundamental framework, the GOPT [7] model, for training to explore the sole effects of the mixup itself without supplementary modeling techniques. GOPT is based on a Transformer [27] encoder and utilizes the 84-dimensional GOP features obtained with the process described in Section 3.2. The GOP features are first projected to 24 dimensions by a projection layer and combined with canonical phoneme and positional embedding. Then, the combined input is fed into a three-layer transformer encoder with 24 embedding dimensions.

To ensure a fair comparison, we kept all settings except those related to the proposed method and GPU identical to the GOPT. Specifically, using the Adam optimizer, we set the learning rate as 1e-3 and batch size as 25 on 100 epoch training. For the acoustic model¹ to obtain GOP features, we used the LibriSpeech [28] 960-hour data-trained model. α for beta distribution is set as 1 to create even likelihoods for mixing coefficients. To acquire error-rate features, we employed a wav2vec2.0 with 315 million parameters [29] as the ASR model. For phoneme transcription and evaluation, we aligned the ASR model’s vocabulary with the speechocean762 dataset and trained the ASR model with the CTC head [30]. GTX 2080Ti GPU is used, and the averaged PCC results of five distinct runs are reported with the standard deviation. Following prior studies, MSE is also used to measure phoneme-level accuracy.

5. Results and discussion

5.1. Main result

The main results presented in Table 1 highlight the effectiveness of both our AM_{stat} and AM_{dyn} methods in improving the training of the DNN-based model across multiple aspects at the phoneme, word, and utterance levels. Particularly noteworthy is the approximately 25% enhancement in assessment performance for the previously weakest aspect, *Completeness*, indicating a more balanced outcome across various aspects. Also, improvements are observed for the *Stress*, another highly imbalanced aspect, but the extents are not as significant. Notably,

¹<https://kaldi-asr.org/models/m13>

Table 2: Comparison of results between using a fixed lambda value of 0.3 in AM_{stat} (fix) and using random weights following a beta distribution (beta). +ER denotes the addition of error-rate features.

Model	Phoneme Score		Word Score (PCC)			Utterance Score (PCC)				
	Acc(MSE ↓)	Acc(PCC ↑)	Acc ↑	Stress ↑	Total ↑	Acc ↑	Comp ↑	Fluency ↑	Prosody ↑	Total ↑
AM_{stat} (fix) +ER	0.085 ±0.001	0.614 ±0.004	0.537 ±0.004	0.324 ±0.025	0.555 ±0.003	0.736 ±0.007	0.302 ±0.054	0.780 ±0.004	0.780 ±0.003	0.766 ±0.007
AM_{stat} (beta) +ER	0.085 ±0.001	0.614 ±0.005	0.538 ±0.005	0.306 ±0.009	0.558 ±0.005	0.735 ±0.001	0.402 ±0.085	0.780 ±0.002	0.779 ±0.003	0.764 ±0.005

Table 3: Ablation results in error-rate features. The multi-aspect averaged performances within each level are reported.

Model	Phoneme		Word	Utterance
	MSE ↓	PCC ↑	Avg PCC ↑	Avg PCC ↑
GOPT	0.085	0.612	0.458	0.625
CER	0.086	0.610	0.453	0.637
MER	0.085	0.613	0.456	0.650
CER + MER	0.086	0.612	0.462	0.656
+ AM_{stat}	0.085	0.614	0.467	0.692
+ AM_{dyn}	0.084	0.617	0.471	0.692

Completeness is scored on a continuous scale from 0 to 10, while *Stress* is scored on a scale of either 5 or 10. Therefore, our method of smoothly shifting the distribution to achieve evenness might be more suitable for the former.

Overall, AM_{dyn} +ER exhibits the highest performance tendency, followed closely by AM_{stat} +ER. While pseudo labels generated by static mixup span the entire score spectrum (Figure 2; b), those from dynamic mixup tend to be distributed more on rare or lower scores (Figure 2; c); thus, higher results on AM_{dyn} +ER imply its potential guidance for more adversarial synthesis. A noticeable point is made for severely imbalanced and inferior aspects such as *Completeness* and *Stress*: excluding error-rate features in static and dynamic mixups yields better performance. This discrepancy could be attributed to ER’s reliance on ASR model results, which may propagate ASR errors during the mixing process, unlike fixed and reliable human-annotated score labels.

5.2. Mixup weight choices

We investigate the impact of the choice of mixture ratio in static AM, whether to set it to a fixed value or follow a random beta distribution. When weights are fixed at a static value of 0.3, the shifted distribution of labels appears quite rigid (Figure 2; a). However, the superior performance of the fixed AM_{stat} in word-level *Stress* as shown in Table 2 suggests that such rigidity might be advantageous in discrete aspects. Conversely, the contrasting trend observed in *Completeness* indicates that a smoother shift could be beneficial for aspects requiring continuous predictions.

5.3. Error rate ablation studies

We conduct extensive ablation studies to examine the individual and combined effects of each error rate on model training. The results in Table 3 indicate that, when used individually, MER has a greater impact than CER. Particularly at the utterance level, MER proves beneficial, likely due to its measurement method focusing on phonemes across the entire utterance. Notably, while neither individually aids at the word level, their combined usage shows performance improvement, indicating a synergistic effect between the two error factors. Moreover, the




Figure 3: Score-label distribution shift when AM_{dyn} is applied with the *original* and the *reverse* directions (left), and PCC performance and standard deviation of PCC of aspects within each granularity level (right).

inclusion of + AM_{stat} and + AM_{dyn} , which incorporate original and mixed-up error rates into the final model vector, remarkably improves the PCC across all levels, highlighting the effectiveness of auxiliary combining ER features.

5.4. Mixup direction matters

We further analyze whether our hypothesized shift toward underserved areas is indeed beneficial compared to the opposite direction. In particular, we adjust our formula from the original ($\tilde{x} = \lambda_1 x - \lambda_2 a_x + \lambda_1 \lambda_2 (x - a_x)$) to the following ($\tilde{x} = \lambda_1 x + \lambda_2 a_x + \lambda_1 \lambda_2 (x - a_x)$), aiming to move in a direction proportional to the average score, inspired by [15]. We call this as *reversed* AM_{dyn} . In the left part of Figure 3, we observe that the *reversed* AM_{dyn} indeed induces shifts in the opposite direction as intended. This suggests that while the original AM_{dyn} generates minority samples more frequently, the reversed AM_{dyn} favorably synthesizes majority samples. An interesting finding is that AM_{dyn} outperforms reversed AM_{dyn} across all granularity levels (Figure 3; bar charts), with even the decreasing PCC standard deviation among the aspects within each level. The result reveals that our approach not only contributes to achieving competitive performance but also facilitates balanced learning across overall aspects as we intended.

6. Conclusion

In this work, we propose two *Acoustic Feature Mixup* strategies, AM_{stat} and AM_{dyn} , which consider linear and non-linear interpolation between the samples and in-batch averaged feature, respectively. Primarily leveraging the GOP features but additionally introducing the error rate features, we design effective mixup policies. To evaluate our method on the DNN-based model, we use the foundational system for the multi-aspect pronunciation assessment task. Experiments with the highly imbalanced speechocean762 dataset exhibit overall performance improvement across all aspects, demonstrating our assistance in balanced scoring. Extensive analysis further demonstrates the potential for our smoother shift with AM to enhance prediction for adversarial or unseen samples.

7. Acknowledgements

This research was partly supported by the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2024-2020-0-01789) supervised by the IITP (Institute for Information & Communications Technology Planning & Evaluation) and by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2022-0-00223, Development of digital therapeutics to improve communication ability of autism spectrum disorder patients).

8. References

- [1] M. Eskenazi, "An overview of spoken language technology for education," *Speech Communication*, vol. 51, no. 10, pp. 832–844, 2009.
- [2] H. Franco, L. Neumeier, Y. Kim, and O. Ronen, "Automatic pronunciation scoring for language instruction," in *1997 IEEE international conference on acoustics, speech, and signal processing*, vol. 2. IEEE, 1997, pp. 1471–1474.
- [3] S. M. Witt and S. J. Young, "Phone-level pronunciation scoring and assessment for interactive language learning," *Speech communication*, vol. 30, no. 2-3, pp. 95–108, 2000.
- [4] D. Luo, Y. Qiao, N. Minematsu, Y. Yamauchi, and K. Hirose, "Analysis and utilization of mlr speaker adaptation technique for learners' pronunciation evaluation," in *Tenth annual conference of the international speech communication association*, 2009.
- [5] Y.-B. Wang and L.-S. Lee, "Improved approaches of modeling and detecting error patterns with empirical analysis for computer-aided pronunciation training," in *2012 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2012, pp. 5049–5052.
- [6] J. Shi, N. Huo, and Q. Jin, "Context-Aware Goodness of Pronunciation for Computer-Assisted Pronunciation Training," in *Proc. Interspeech 2020*, 2020, pp. 3057–3061.
- [7] Y. Gong, Z. Chen, I.-H. Chu, P. Chang, and J. Glass, "Transformer-based multi-aspect multi-granularity non-native english speaker pronunciation assessment," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7262–7266.
- [8] H. Do, Y. Kim, and G. G. Lee, "Hierarchical pronunciation assessment with multi-aspect attention," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [9] F.-A. Chao, T.-H. Lo, T.-I. Wu, Y.-T. Sung, and B. Chen, "3m: An effective multi-view, multi-granularity, and multi-aspect modeling approach to english pronunciation assessment," in *2022 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2022, pp. 575–582.
- [10] F.-A. Chao, T.-H. Lo, and et al., "A Hierarchical Context-aware Modeling Approach for Multi-aspect and Multi-granular Pronunciation Assessment," in *Proc. INTERSPEECH 2023*, 2023, pp. 974–978.
- [11] H. Do, Y. Kim, and G. G. Lee, "Score-balanced Loss for Multi-aspect Pronunciation Assessment," in *Proc. INTERSPEECH 2023*, 2023, pp. 4998–5002.
- [12] Y. Basuki, "The use of drilling method in teaching phonetic transcription and word stress of pronunciation class," *Karya Ilmiah Dosen*, vol. 1, no. 1, 2018.
- [13] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=r1Ddp1-Rb>
- [14] A. F. M. S. Uddin, M. S. Monira, W. Shin, T. Chung, and S.-H. Bae, "Saliencymix: A saliency guided data augmentation strategy for better regularization," in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=M0QkvBGTTq>
- [15] Sanskriti, S. Jang, D. Kim, S. Cha, D. Kim, and K. Kim, "A dynamic mixup approach towards improved robustness of classifiers," 2024. [Online]. Available: <https://openreview.net/forum?id=YMHDDeDTWbE>
- [16] Z. Liu, S. Li, G. Wang, L. Wu, C. Tan, and S. Z. Li, "Harnessing hard mixed samples with decoupled regularizer," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [17] M. Sancinetti, J. Vidal, C. Bonomi, and L. Ferrer, "A transfer learning approach for pronunciation scoring," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6812–6816.
- [18] S. Venkataramanan, E. Kijak, Y. Avrithis et al., "Embedding space interpolation beyond mini-batch, beyond pairs and beyond examples," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [19] A. Chakraborty, M. Alam, V. Dey, A. Chattopadhyay, and D. Mukhopadhyay, "A survey on adversarial attacks and defenses," *CAAI Transactions on Intelligence Technology*, vol. 6, no. 1, pp. 25–45, 2021.
- [20] J. Liu, Z. Shen, Y. He, X. Zhang, R. Xu, H. Yu, and P. Cui, "Towards out-of-distribution generalization: A survey," *arXiv preprint arXiv:2108.13624*, 2021.
- [21] J.-H. Kim, W. Choo, and H. O. Song, "Puzzle mix: Exploiting saliency and local statistics for optimal mixup," in *International Conference on Machine Learning*. PMLR, 2020, pp. 5275–5285.
- [22] S. Li, Z. Wang, Z. Liu, D. Wu, C. Tan, and S. Z. Li, "Openmixup: A comprehensive mixup benchmark for visual classification," *ArXiv*, vol. abs/2209.04851, 2022.
- [23] J. Liu, B. Liu, H. Zhou, H. Li, and Y. Liu, "Tokenmix: Rethinking image mixing for data augmentation in vision transformers," in *European Conference on Computer Vision*. Springer, 2022, pp. 455–471.
- [24] J.-H. Kim, W. Choo, H. Jeong, and H. O. Song, "Co-mixup: Saliency guided joint mixup with supermodular diversity," *arXiv preprint arXiv:2102.03065*, 2021.
- [25] W. Hu, Y. Qian, F. K. Soong, and Y. Wang, "Improved mispronunciation detection with deep neural network trained acoustic models and transfer learning based logistic regression classifiers," *Speech Communication*, vol. 67, pp. 154–166, 2015.
- [26] J. Zhang, Z. Zhang, Y. Wang, Z. Yan, Q. Song, Y. Huang, K. Li, D. Povey, and Y. Wang, "speechocan762: An Open-Source Non-Native English Speech Corpus for Pronunciation Assessment," in *Proc. Interspeech 2021*, 2021, pp. 3710–3714.
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [28] 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.
- [29] 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, p. 12449–12460, 2020.
- [30] A. Graves and A. Graves, "Connectionist temporal classification," *Supervised sequence labelling with recurrent neural networks*, pp. 61–93, 2012.