



Score-balanced Loss for Multi-aspect Pronunciation Assessment

Heejin Do¹, Yunsu Kim^{1,2}, Gary Geunbae Lee^{1,2}

¹Graduate School of AI, POSTECH, Republic of Korea

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

{heejindo, yunsu.kim, gblee}@postech.ac.kr

Abstract

With rapid technological growth, automatic pronunciation assessment has transitioned toward systems that evaluate pronunciation in various aspects, such as fluency and stress. However, despite the highly imbalanced score labels within each aspect, existing studies have rarely tackled the data imbalance problem. In this paper, we suggest a novel loss function, score-balanced loss, to address the problem caused by uneven data, such as bias toward the majority scores. As a re-weighting approach, we assign higher costs when the predicted score is of the minority class, thus, guiding the model to gain positive feedback for sparse score prediction. Specifically, we design two weighting factors by leveraging the concept of an effective number of samples and using the ranks of scores. We evaluate our method on the speechocean762 dataset, which has noticeably imbalanced scores for several aspects. Improved results particularly on such uneven aspects prove the effectiveness of our method.

Index Terms: automated pronunciation assessment, imbalanced dataset, score-balanced loss

1. Introduction

Automatic pronunciation assessment supports non-native (L2) language learners in acquiring foreign spoken languages as part of a computer-assisted pronunciation training (CAPT) system [1, 2]. With the advantages of immediate feedback and convenience, CAPT has been actively studied mainly in assessing pronunciation with a phoneme-level score [3, 4, 5].

With increasing demands for detailed feedback, recent studies have evaluated pronunciation in several aspects of various granularity levels, such as stress, fluency, and prosody [6, 7, 8, 9]. As an alternative to multiple models that separately assess different aspects, current joint models [10, 11, 12] facilitate simultaneous prediction of multiple aspects using a single model considering the association between aspects.

However, notwithstanding the technical advances, extremely imbalanced score labels within each aspect that show high-score-biased distributions have rarely been studied on the pronunciation assessment task. Biased datasets of the CAPT system have been frequently reported [13, 5, 10, 11]; however, they are rarely optimized for the scoring task. Imbalanced datasets could cause the model to be overfitted toward the majority classes during the training [14, 15]; therefore, addressing the imbalance problem is crucial for qualified assessment. Furthermore, in the case of multi-aspect pronunciation scoring, uneven data distributions within certain aspects cause huge performance gaps between different aspects, which consequently hinders application to real-world educational situations. For practical use, accurately evaluating every aspect is important, without bias toward a specific aspect; in particular, good overall

quality is better than a part with exceptional quality.

In this paper, we introduce a novel loss function, score-balanced (SB) loss, to overcome the quality degradation caused by imbalanced data in multi-aspect pronunciation assessment. Motivated by the class-balanced loss suggested for visual classification [15], we introduce cost-sensitive re-weighting schemes for balancing the weight when scoring pronunciation. Unlike class-balanced loss, which is based on the ground-truth label class, the proposed SB_{num} loss directly targets the predicted scores and aims at score-labeled regression. Specifically, we design SB loss with two different factors: SB_{num} , which leverages the number of samples on categorized scores, and SB_{rank} , which exploits those ranks. By assigning high costs to the predicted score of the minority class, the model is directed to favorably predict sparse classes despite fewer training samples.

We evaluate the proposed loss function on the public speechocean762 dataset, which is widely used for multi-aspect pronunciation assessments [16]. We find highly imbalanced data distributions, which are densely distributed toward a high score, on certain aspects such as *Completeness* and *Stress*. Attributing this result to the particularly low assessment qualities on such aspects in existing studies [10, 11, 12], we train the open-source model, goodness of pronunciation feature-based transformer (GOPT) [10], with the proposed SB loss. Significantly improved results on the notably imbalanced aspects prove that the issue of unevenly distributed datasets has been successfully resolved, reducing the gap between different aspect assessment qualities. It is noteworthy that our enhancements are achieved without any augmentation or architecture modeling. Our codes are available on the GitHub link¹.

2. Related work

A dataset is an essential factor for both supervised classification and regression. Therefore, the data imbalance problem has been actively discussed for decades, particularly around visual and text classification tasks [17, 18, 14, 19]. Generally, the related studies are divided according to two main approaches: re-sampling and re-weighting. Re-sampling over-samples or under-samples the data by repeating or deleting the existing data examples [20, 21], while re-weighting adjusts the loss function by assigning more costs to under-represented-class samples [14, 15]. Recently, pointing out the lack of research on imbalanced regression which handles continuous targets, few studies dealing with non-categorical imbalanced data have been proposed, mostly for the visual dataset [22, 23]. Emphasizing the need for a balancing strategy for challenging imbalanced regression, particularly for speech pronunciation assessment, we propose a novel loss optimized for score-labeled datasets.

¹https://github.com/doheejin/SB_loss_PA




Figure 1: Data distribution histogram across score labels for each aspect of the utterance (1st row) and word (2nd row).

Real-word assessments for educational use have suffered from naturally imbalanced datasets [24], which tend to have biased grades. To improve the pronunciation error detection of CAPT by addressing problems caused by extremely imbalanced class datasets, machine learning-based re-weighting methods have been investigated [13]. Most recently, to compensate for the score imbalance in pronunciation assessment datasets, a transfer learning-based study [5] applied simple balanced weights with the inverse of the number of frames. Considering the score labels that are biased toward high values, GOPT [10] set the Pearson correlation coefficient (PCC) as the major evaluation metric. However, their focus is not on data inequality, and an intensive investigation for tackling score imbalance is lacking. We propose a novel solution focusing on scoring data imbalance in multi-aspect pronunciation assessment.

3. Methods

3.1. Score-label imbalance dataset

We investigate the score label distribution of the pronunciation assessment data by each aspect before presenting our method. We observe the publicly available speechocean762 dataset [16], which provides rich score labels for various aspects and has been mainly used for multi-aspect pronunciation assessment tasks [10, 11, 12]. For each non-native speaker utterance, the dataset comprises phoneme-level *Accuracy* score; word-level *Accuracy*, *Stress*, and *Total* scores; and utterance-level *Accuracy*, *Completeness*, *Fluency*, *Prosody*, and *Total* scores. The phoneme-level aspect has a score between 0-2, and the word-level and utterance-level aspects have a score between 0-10.

Although the dataset promotes research on multi-aspect pronunciation assessment, the provided score labels are imbalanced, showing high-score-biased distributions, particularly in utterance *Completeness* and word *Stress* aspects (Figure 1). Note that *Completeness* actually has six labels but are not shown due to extremely huge samples of score 2. We attribute this imbalance to the notably inferior quality of scoring tasks of those two aspects in existing models [10, 11, 12]. Figure 2 shows that the mean squared error (MSE) loss, which is a generally used loss function for pronunciation assessment, is insufficient for scoring the high-score-biased *Completeness* aspect. To overcome the limitations of the existing learning, we suggest SB loss, which re-weights the MSE loss and aids model training.

3.2. Score-balanced loss

To effectively re-weight the loss when training with imbalanced data, we introduce the SB factor. Our work is motivated by the




Figure 2: MSE and PCC score gap between different aspect-assessment tasks on the GOPT [10] model. MSE is used as the loss function; however, for the *Completeness* aspect, it indeed does not contribute to the training.

class-balanced loss [15], which adds the weighting factor to the loss using the inverse of a demonstrated effective-number-of-samples term for an imbalanced visual classification task.

Unlike its application to the cross-entropy loss for classification, we extend the concept of the effective number of samples to the MSE loss for regression. Specifically, we re-scale the range of word and utterance aspect scores from 0-10 to 0-2 and regard scores at intervals of 0.2 as different classes, $\{0, 0.2, \dots, 2\}$, defining two forms of score-balanced loss with different weighting factors: SB_{num} and SB_{rank} .

3.2.1. Re-weighting by the number of samples: SB_{num}

According to a proven theoretical framework [15], the effective number of samples can be defined as $(1 - \beta^n)/(1 - \beta)$, where n is the number of samples and β is the hyperparameter defined with $(N - 1)/N$. In contrast to class-balanced loss, which is weighted according to the number of samples n for ground-truth class y and assigned a fixed term irrespective of the direction of learning, we define SB_{num} according to the number of samples of the *predicted* score. Thus, the higher weights are assigned to the smaller number of samples of the predicted score class. By giving positive feedback for *predicting* rare classes during training, the model is induced to predict more confidently regarding sparsely observed scores. Specifically, we consider predicted score \hat{y} in the interval of $[s, s + 0.2)$ as class s and set the number of samples of score class s (n_s) as $n_{\hat{y}}$. The SB_{num} loss weight factor for the m -th aspect, α^m , is defined as follows:

$$\alpha^m = \begin{cases} (1 - \beta)/(1 - \beta^{n_{\hat{y}}^m}) & , \text{if } n_{\hat{y}}^m \neq 0 \\ 1 & , \text{otherwise} \end{cases} \quad (1)$$

where $n_{\hat{y}}^m$ denotes the number of samples for predicted score \hat{y} on the m -th aspect. Then, SB_{num} loss, where the MSE loss is multiplied by the SB factor can be defined as the following:

$$SB_{num} = \sum_{m=1}^M \alpha^m \cdot (L_{MSE})^m \quad (2)$$

where M is the number of all aspects at all phoneme, word, and utterance granularity levels; and α^m and $(L_{MSE})^m$ are the balancing factor and MSE loss of the m -th aspect, respectively.

3.2.2. Re-weighting by the ranking: SB_{rank}

Considering that the variance of the number of samples is large depending on the classes, we further design SB_{rank} to assign weights according to ranks rather than the actual sample number per class. As there are only 11 classes (from 0 to 2), the rank gap between classes is not significant compared with up to 2500

Table 1: *Experimental results of average MSE (only for phoneme-level) and PCC scores with standard deviation for five different runs. Acc is Accuracy and Comp is Completeness. -baseline denotes our implemented results of GOPT. Bold text denotes remarkably outperforming results over the baseline.*

Model	Phoneme Score		Word Score (PCC)			Utterance Score (PCC)				
	Acc(MSE ↓)	Acc(PCC ↑)	Acc ↑	Stress ↑	Total ↑	Acc ↑	Comp ↑	Fluency ↑	Prosody ↑	Total ↑
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
Gong et al.	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
*Gong et al.-baseline	0.086 ±0.001	0.609 ±0.003	0.526 ±0.004	0.272 ±0.048	0.543 ±0.005	0.717 ±0.006	0.134 ±0.197	0.755 ±0.006	0.756 ±0.003	0.739 ±0.004
+ SB_{num}	0.086 ±0.001	0.605 ±0.006	0.531 ±0.005	0.386 ±0.015	0.547 ±0.005	0.722 ±0.004	0.427 ±0.101	0.750 ±0.013	0.752 ±0.007	0.747 ±0.002
+ SB_{rank}	0.086 ±0.001	0.605 ±0.003	0.529 ±0.004	0.341 ±0.036	0.544 ±0.002	0.717 ±0.007	0.335 ±0.235	0.750 ±0.005	0.745 ±0.015	0.743 ±0.007

sample gaps; therefore, replacing balancing criteria with rank enables smooth learning.

Specifically, we rank the number of samples of each class in descending order and use the inverse rank as an SB term; therefore, the majority label is weighted less. For the same values, the average-rank policy was applied (e.g., rank 1.5 is assigned to two classes of the same smallest number of samples). For the obtained ranking of the number of samples of the predicted class \hat{y} , normalization is applied. The weighting factor of SB_{rank} loss when predicting the m -th aspect is defined as

$$\gamma^m = (1/r_{\hat{y}}^m) \quad (3)$$

where $r_{\hat{y}}^m$ is the normalized ranking of predicted class size on the m -th aspect. Then, the total SB_{rank} loss across all M aspects is represented as follows:

$$SB_{rank} = \sum_{m=1}^M \gamma^m \cdot (L_{MSE})^m \quad (4)$$

4. Experiments

We evaluate our SB loss on the publicly available speechocean762 [16] dataset, which has imbalanced labels and is well-built for the multi-aspect pronunciation assessment task. We use the public GOPT² model [10], which predicts multiple aspect scores in parallel. GOPT is based on a transformer [25] architecture and uses the goodness of pronunciation (GOP) features. In detail, with the audio and corresponding canonical transcription input, the acoustic model outputs frame-level phonetic posterior probabilities, which are then transformed into GOP features with 84 dimensions. Then, GOP features projected to 24 dimensions, canonical phoneme embedding, and positional embedding are added and input to the three-layer transformer encoder with 24 embedding dimensions.

We follow the same settings described for GOPT experiments, except for the loss function and graphics processing unit (GPU). In detail, the Adam optimizer, an initial $1e-3$ learning rate, 25 batch size, and 100 epochs are set for training. All models have 26.577k parameters. The automatic speech recognition model³ trained with LibriSpeech [26] 960-hour data is used. Training and test sets, each comprising 2500 utterances, are used. NVIDIA RTX A5000 GPU is used for all experiments. Five different experiments are conducted by different random seeds, and their mean and standard deviation are reported. PCC

²<https://github.com/YuanGongND/gopt>

³<https://kaldi-asr.org/models/m13>




Figure 3: *Test PCC scores of utterance Completeness and word Stress aspect according to training epochs.*

is used for the evaluation metric; additionally, MSE is used for phoneme-level accuracy. We set hyperparameter β in S_{num} as 0.9 after conducting experiments with diverse values.

5. Results and discussion

To compare the results in the same computational environment, we re-implement the GOPT model and use for comparison. Table 1 reveals that both models trained with SB_{num} and SB_{rank} loss exhibit exceptional improvements regarding aspects with extremely unbalanced score distributions, indicating that our approach effectively tackles the negative impact of imbalanced data. Specifically, after training with SB_{num} loss, absolute improvements of 11.4% and 29.3% are observed for the word *Stress* and utterance *Completeness*, respectively. This further reduces the performance gap between different tasks in a multi-aspect learning setting, as the difference between the most superior (*Prosody*) and inferior (*Completeness*) aspect is reduced from 5 to 1.7 times. Meanwhile, there are no significant PCC score differences over the baseline in aspects where samples are somewhat present even in minority score classes. Similar trends are observed for the model trained with SB_{rank} loss, which shows absolute improvements of 6.9% and 20.1% in *Stress* and *Completeness*, respectively. Overall evaluation results for aspects indicate that specifying the number of samples in each scoring class allows more precise re-weighting than abstracting them with the rank. The changes in the test PCC score by the learning process (Figure 3) reveal that SB_{num} and SB_{rank} converge to a specific value at the latter and higher point than the baseline, implying that the training validity continued longer.

5.1. Ablation study

Apart from the proposed SB_{num} loss, which is based on the predicted value, we additionally experiment with SB loss using the ground-truth label of each sample. Instead of the number of

Table 2: Comparison results between SB_{num} with the sample number of predicted \hat{y} (SB_{num}) and SB_{num} with the sample number of ground-truth y (SB_{num} (no pred)). **Bold** text denotes the particularly improved results.

Model	Phoneme Score		Word Score (PCC)			Utterance Score (PCC)				
	Acc(MSE ↓)	Acc(PCC ↑)	Acc ↑	Stress ↑	Total ↑	Acc ↑	Comp ↑	Fluency ↑	Prosody ↑	Total ↑
+ SB_{num} (no pred)	0.085 ±0.000	0.611 ±0.002	0.534 ±0.006	0.293 ±0.031	0.549 ±0.004	0.717 ±0.006	0.218 ±0.097	0.760 ±0.007	0.758 ±0.011	0.745 ±0.006
+ SB_{num}	0.086 ±0.001	0.605 ±0.006	0.531 ±0.005	0.386 ±0.015	0.547 ±0.005	0.722 ±0.004	0.427 ±0.101	0.750 ±0.013	0.752 ±0.007	0.747 ±0.002




Figure 4: Change in PCC scores of the model with SB_{num} according to change in hyperparameter β value for each aspect.

Table 3: Averaged $(1 - \beta)/(1 - \beta^{n_y})$ values across 11 classes for different β in the case of the phoneme-level Accuracy aspect.

β	0.5	0.8	0.9	0.99	0.999
Avg $(1 - \beta)/(1 - \beta^{n_y})$	0.542	0.267	0.175	0.095	0.088

samples in the class of the predicted \hat{y} , that of ground-truth y is used; therefore, the weighting factor for the m -th aspect, α^m , is defined as $(1 - \beta)/(1 - \beta^{n_y^m})$. Note that cases are not needed here, because all samples satisfy $n_y^m \geq 1$.

Table 2 shows almost similar PCC scores across most aspects, except for *Stress* and *Completeness*. For these two aspects, our prediction-based approach achieved higher performance than the label-based application. This implies that dynamically re-weighting the loss according to predicted values for each sample during training is more effective than assigning fixed weights from the beginning.

5.2. Effects of different β values

We observe the effects of different hyperparameter β values when training with SB_{num} loss. Figure 4 shows the best performance at $\beta = 0.9$, and thereafter, the PCC values sharply decrease with an increase in β . This can be explained by the variations on the scale of weight factors. When the value of β is greater, the scale of the weighting factor α becomes smaller, and vice versa (Table 3). Thus, increasing the β value hinders training in the case of the proposed SB loss.

5.3. Qualitative evaluation

To examine the actual mechanism behind our model aiding in pronunciation evaluation, we conduct a qualitative analysis focusing on *Completeness* and *Stress*. For each aspect, we compare the predicted score distribution of the baseline model and that of the SB_{num} -loss-trained model with the actual ground-truth score on the test set. Figure 5 reveals that the model with our approach predicts a wider range of scores while approximating closer to the predicted score distribution, which is densely

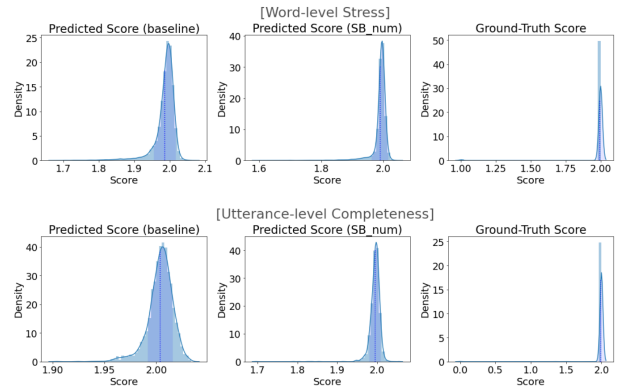


Figure 5: Comparison of predicted score distributions (of baseline and our method) and test set ground-truth scores; vertical lines represent mean, darker area shows standard deviations.

distributed around the score of 2. This comparison indicates that our strategy not only helps predict minority samples but also supports accurate predictions for majority samples.

In particular, for the word-level *Stress*, the actual target scores are distributed in the range 1-2. The lowest score predicted by the baseline is 1.67, while that by our method is 1.58, which is closer to the actual score of 1. For utterance-level *Completeness*, the lowest predicted score by the baseline is 1.90, while that with ours is 1.69, both of which are the samples of the ground-truth 0 score. Samples with *Completeness* score of 0 do not exist in the training set, thus our closer prediction implies the assistance on missing data at continuous target values.

6. Conclusions

In this paper, we propose a simple yet effective score-balanced loss function for multi-aspect pronunciation assessment. Assuming continuous scores as the categorical classes, we design two forms of SB loss, assigning high costs when the model predicts minor classes. Both re-weighting by exploiting the number of samples in the predicted value and by the ranks of sample size effectively balance the biased training. The experimental results showing remarkable improvement on the highly imbalanced aspects prove that SB loss overcomes the limitation of unevenly distributed datasets.

Acknowledgements: This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-01906, Artificial Intelligence Graduate School Program (POSTECH)) and MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2023-2020-0-01789) supervised by the IITP.

7. 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] J. Shi, N. Huo, and Q. Jin, "Context-aware goodness of pronunciation for computer-assisted pronunciation training," *arXiv preprint arXiv:2008.08647*, 2020.
- [5] 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.
- [6] J. Tepperman and S. Narayanan, "Automatic syllable stress detection using prosodic features for pronunciation evaluation of language learners," in *Proceedings (ICASSP'05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.*, vol. 1. IEEE, 2005, pp. I-937.
- [7] C. Cucchiari, H. Strik, and L. Boves, "Quantitative assessment of second language learners' fluency by means of automatic speech recognition technology," *The Journal of the Acoustical Society of America*, vol. 107, no. 2, pp. 989–999, 2000.
- [8] J. P. Arias, N. B. Yoma, and H. Vivanco, "Automatic intonation assessment for computer aided language learning," *Speech communication*, vol. 52, no. 3, pp. 254–267, 2010.
- [9] K. Li, X. Wu, and H. Meng, "Intonation classification for 12 english speech using multi-distribution deep neural networks," *Computer Speech & Language*, vol. 43, pp. 18–33, 2017.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] X. Yang, A. Loukina, and K. Evanini, "Machine learning approaches to improving pronunciation error detection on an imbalanced corpus," in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 300–305.
- [14] Y.-X. Wang, D. Ramanan, and M. Hebert, "Learning to model the tail," *Advances in neural information processing systems*, vol. 30, 2017.
- [15] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, "Class-balanced loss based on effective number of samples," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 9268–9277.
- [16] J. Zhang, Z. Zhang, Y. Wang, Z. Yan, Q. Song, Y. Huang, K. Li, D. Povey, and Y. Wang, "speechocean762: An open-source non-native english speech corpus for pronunciation assessment," *arXiv preprint arXiv:2104.01378*, 2021.
- [17] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Transactions on knowledge and data engineering*, vol. 21, no. 9, pp. 1263–1284, 2009.
- [18] N. Japkowicz and S. Stephen, "The class imbalance problem: A systematic study," *Intelligent data analysis*, vol. 6, no. 5, pp. 429–449, 2002.
- [19] C. Padurariu and M. E. Breaban, "Dealing with data imbalance in text classification," *Procedia Computer Science*, vol. 159, pp. 736–745, 2019.
- [20] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: synthetic minority over-sampling technique," *Journal of artificial intelligence research*, vol. 16, pp. 321–357, 2002.
- [21] A. Estabrooks, T. Jo, and N. Japkowicz, "A multiple resampling method for learning from imbalanced data sets," *Computational intelligence*, vol. 20, no. 1, pp. 18–36, 2004.
- [22] J. Ren, M. Zhang, C. Yu, and Z. Liu, "Balanced mse for imbalanced visual regression," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 7926–7935.
- [23] Y. Yang, K. Zha, Y. Chen, H. Wang, and D. Katabi, "Delving into deep imbalanced regression," in *International Conference on Machine Learning*. PMLR, 2021, pp. 11 842–11 851.
- [24] 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.
- [25] 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.
- [26] 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.