

# Mitigating Data Imbalance in Automated Speaking Assessment

Fong-Chun Tsai \*, Kuan-Tang Huang \*, Bi-Cheng Yan \*, Tien-Hong Lo \*, and Berlin Chen \*

\* National Taiwan Normal University

E-mail: {fongchun.tsai, 61347002s, 80847001s, teinhonglo, berlin}@ntnu.edu.tw

**Abstract**—Automated Speaking Assessment (ASA) plays a crucial role in evaluating second-language (L2) learners’ proficiency. However, ASA models often suffer from class imbalance, leading to biased predictions. To address this, we introduce a novel objective for training ASA models, dubbed the **Balancing Logit Variation (BLV) loss**, which perturbs model predictions to improve feature representation for minority classes without modifying the dataset. Evaluations on the ICNALE benchmark dataset show that integrating the BLV loss into a celebrated text-based (BERT) model significantly enhances classification accuracy and fairness, making automated speech evaluation more robust for diverse learners.

## I. INTRODUCTION

Driven by advances in automatic speech recognition (ASR) [1]–[3] and the worldwide growth of second-language learners (L2), Automated Speaking Assessment (ASA) has emerged as a practical and widely used tool for language acquisition [4]–[6], serving as a key component of computer-assisted language learning (CALL). ASA systems are developed to assess the spoken language proficiency of L2 learners in multiple dimensions of communicative competence, such as delivery, language use, and discourse coherence. Their applications primarily fall into two main areas: high-stakes exams and pedagogical tools. For example, organizations such as the ETS[7] leverage ASA systems to streamline scoring processes, reduce the workload of human experts, lower recruitment costs, and deliver consistent and objective evaluations of speaking proficiency[8]. On the other hand, ASA systems serve as a critical resource in classrooms, providing teachers with reliable and timely assessment feedback on student oral skills, in order to improve instructional quality and/or offer statistical data on student performance [6].

A de facto standard for ASA systems is typically instantiated in an open-response learning scenario, where an L2 learner is provided with text prompts and/or questions to elicit natural speech responses. Based on open-ended and spontaneous speech inputs, ASA systems manage to quantify communicative competence from multiple facets, including pronunciation, vocabulary, grammatical accuracy, and topical relevance. Existing studies frame the ASA task as a proficiency score classification problem, which comprises scoring, speech processing, and natural language processing modules. Among these modules, the speech processing module is a critical component of ASA, acting as the foundation for analyzing the spoken responses of learners in terms of specific aspects, including pronunciation, stress, fluency, and intonation[7]. There

are two primary approaches to feature extraction: hand-crafted features and self-supervised learning (SSL) features [9], [10]. To extract hand-crafted features, the spoken response is first transcribed into a sequence of linguistic units by an ASR system, and a set of corresponding features is extracted to portray the learner’s pronunciation quality through forced alignment with the recognized hypotheses. The commonly used hand-crafted features include, but are not limited to, silence duration, speaking rate, and repetition rate. In contrast, recent advances in self-supervised models like Wav2Vec 2.0 [11] and BERT [12] have introduced a new paradigm for feature extraction. These models leverage vast amounts of unlabeled data during pre-training to produce rich, high-dimensional representations of speech. While SSL models require fine-tuning to align with the specific goals of ASA and improve their interpretability, studies have shown that SSL features outperform traditional methods in ASA tasks [13]. This marks a significant step forward in enhancing the accuracy and robustness of speech assessments.

Most ASA data is derived from examination results, and the scores from these exams typically follow a normal distribution. This characteristic results in the ASA dataset often facing issues of imbalance, meaning that the number of data samples in each performance category or level within the dataset is uneven. Many studies in the past have pointed out that class imbalance can significantly affect the performance of classification tasks, leading to reduced accuracy and reliability [14], [15]. However, despite the well-recognized effects of data imbalance, there has not been much research or effort in previous ASA-related tasks to address or mitigate this issue effectively.

To address the issue of class imbalance, one common approach is data augmentation [16], [17]. This can involve techniques such as reducing the number of instances in majority classes or generating additional data through interpolation methods. However, because ASA datasets have very limited amounts of data, reducing the number of majority class examples may lead to model overfitting. Furthermore, since ASA data follows a Gaussian distribution, there is an inherent lack of examples at the high and low ends of the score range. This creates a problem when using interpolation-based data augmentation methods, as these methods rely on creating new synthetic examples by generating values between two data points. Therefore, this approach is also difficult to apply to ASA tasks.

To address the issue of class imbalance in ASA, we introduce Balancing Logit Variation (BLV) loss from the semantic segmentation domain [18], which does not require any modifications to the data. This method is particularly well-suited for ASA tasks that suffer from a lack of high- and low-score data and limited sample sizes. During training, this approach perturbs the model’s predictions, causing each sample to be projected into a small range of feature space rather than being mapped to a single fixed feature point.

By projecting the data from minority classes into larger intervals, we effectively balance the feature space utilized by the model, thereby mitigating the data imbalance issue. This innovative strategy improves the model’s ability to generalize across classes while addressing the challenges associated with imbalanced datasets in ASA tasks.

Our main contributions are threefold:

- 1) **Introduction of BLV loss for ASA:** The paper proposes a novel loss function that perturbs the logits based on class frequency, enabling more effective learning from imbalanced datasets in ASA tasks.
- 2) **Integration into an ASA pipeline using Whisper and BERT:** The method is integrated into a practical ASA pipeline using Whisper [3] for speech transcription and BERT for proficiency classification, showing its compatibility with state-of-the-art models.
- 3) **Empirical Validation with Extensive Experiments on Long-Tailed Speech Assessment Dataset:** Through extensive experiments, the paper demonstrates that BLV loss improves classification accuracy and fairness across both head and tail classes, outperforming baseline loss functions in real ASA scenarios.

## II. RELATED WORK

### A. Feature Extraction

ASA systems evaluate communicative competence by extracting features related to pronunciation, vocabulary, grammatical accuracy, and topical relevance. Traditional approaches relied on hand-crafted features to capture various aspects of language proficiency, including pronunciation, stress, fluency, and intonation [19], [20]. However, these manual feature engineering methods often incorporate implicit assumptions that may inadvertently omit crucial proficiency indicators [5].

Recent advancements in SSL have introduced more sophisticated methods for spoken response representation. Transformer models, trained through contrastive learning, have demonstrated remarkable capability in capturing multiple dimensions of speech, including semantics, acoustic properties, fluency, and pronunciation features [21]. These models have achieved state-of-the-art performance across various downstream tasks. In the context of L2 learner assessment, several studies have successfully applied SSL features to ASA tasks [5], [6], [13]. Notably, Bannò et al. demonstrated that Wav2Vec2.0-based systems surpass BERT-based text-only approaches in proficiency assessment, particularly in evaluating speech-related components [13]. The superior performance of SSL features

in ASA tasks can be attributed to their comprehensive representation of both semantic content and prosodic elements.

### B. Data Imbalance

Data imbalance is a common classification problem that has been discussed in various fields. It leads to models that focus on the majority class, causing the performance of minority classes to degrade.

The previous literature has proposed several methods to address data imbalance, which can be mainly divided into two approaches. The first approach involves data augmentation techniques, such as oversampling, undersampling, or mixing. However, data augmentation can generate overly ambiguous samples, which, in tasks such as ASA with limited training data, could have a severe negative impact on model performance.

The second approach involves adjusting the loss function, such as using focal loss [22] or BLV loss. These methods are independent of the data and suitable for tasks such as ASA. In this work, we introduce the BLV loss from image segmentation. This loss introduces perturbations to the model predictions, allowing the features of minority classes to be projected onto a larger area, thus balancing the model performance.

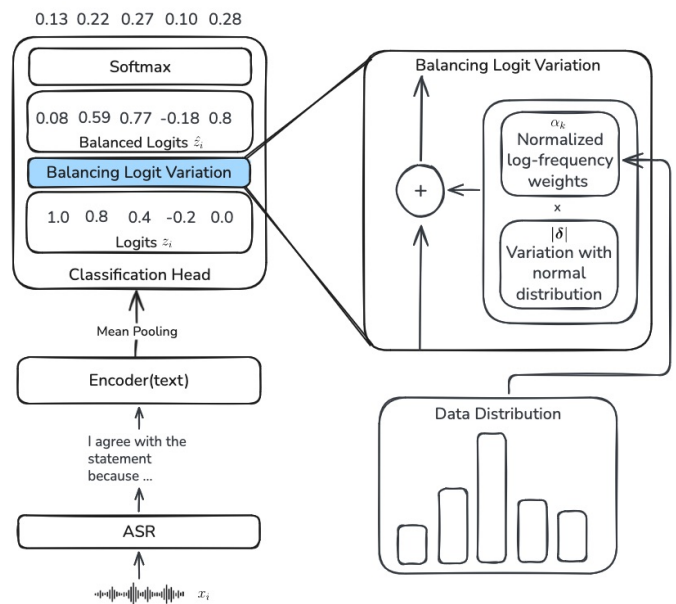


Fig. 1: Model Architecture

## III. METHODOLOGY

### A. Model Overview

Our model utilizes a pre-trained BERT encoder[12] to generate representations for classification tasks. We use this model for simplicity to demonstrate the effectiveness of our proposed loss. Initially, the audio input is transcribed into text using the Whisper ASR system, which may introduce

TABLE I: Comparison of Models on ICNALE Dataset

Method	$\sigma$	PCC $\uparrow$	RMSE $\downarrow$		Accuracy (%) $\uparrow$		Adjacent Accuracy (%) $\uparrow$		F1-score $\uparrow$
			Standard	Macro	Standard	Macro	Standard	Macro	
Bert (baseline)	-	0.6690	0.8614	0.9449	55.30	47.96	91.71	85.79	0.4672
Bert + focal loss	-	0.6466	0.8877	0.9646	54.83	48.13	90.32	83.67	0.4828
Bert + BLV	2	0.6719	0.8587	0.9598	<b>57.14</b>	48.60	91.24	86.45	<b>0.5000</b>
Bert + BLV	6	<b>0.6748</b>	<b>0.8479</b>	<b>0.9281</b>	56.22	<b>49.29</b>	<b>92.17</b>	<b>87.12</b>	0.4952

transcription errors. The resulting text is tokenized using a pre-trained BERT tokenizer.

Subsequently, the tokenized text is processed by the pre-trained BERT model, which serves as a feature extractor. The model obtains a fixed-length representation of the input sequence by applying mean pooling over the BERT embeddings. This representation is then fed into a feed-forward classification head that outputs logits for each class. The model architecture is illustrated in Fig. 1.

In the classification phase, the contextualized representation is passed through a mean pooling layer to aggregate features, followed by a classifier head.

*B. BLV Loss*

When training on a long-tailed dataset  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$  of spoken utterances  $\mathbf{x}_i$  with labels  $y_i \in \{1, \dots, C\}$ , standard cross-entropy tends to favor majority classes. To mitigate this, we adopt the *Balancing Logit Variation* (BLV) strategy [18], which injects class-dependent perturbations into the logits before computing the loss, where the perturbation for minority classes is larger and for majority classes is smaller.

Let  $\mathbf{z}_i \in \mathbb{R}^C$  be the original C-dimensional logit vector for instance  $i$ , and let  $q_k$  denote the total count of category  $k$  with  $N = \sum_{j=1}^C q_j$ . Compute the normalized log-frequency weight

$$\alpha_k = \frac{\log(N/q_k)}{\max_{j \in [C]} (\log(N/q_j))} \tag{1}$$

At each optimization step, sample  $\delta \sim \mathcal{N}(0, \sigma^2)$  and form the perturbed logit.

$$\hat{\mathbf{z}}_i^k = \mathbf{z}_i^k + \alpha_k |\delta| \tag{2}$$

Here,  $\sigma > 0$  is the only BLV loss hyperparameter. Next, compute class probabilities via softmax on  $\{\hat{\mathbf{z}}_i\}$  and evaluate the cross-entropy loss.

$$\mathcal{L}_{BLV} = -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^C \log \frac{\exp(\hat{\mathbf{z}}_i^k)}{\sum_{j=1}^C \exp(\hat{\mathbf{z}}_i^j)} \tag{3}$$

IV. EXPERIMENTAL SETUPS

*A. Dataset*

To validate the generalizability of our approach, we conducted experiments on the International Corpus Network of Asian Learners of English (ICNALE) corpus, a publicly




Fig. 2: Distribution of ICNALE Dataset on CEFR LEVEL

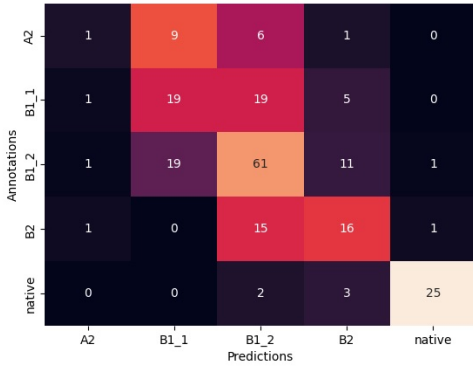

available dataset. We used the Spoken Monologues section of ICNALE, which contains a total of 4,332 samples. Each sample includes a topic, along with the user’s spoken and written responses. The users consist of both native speakers and L2 learners whose CEFR (Common European Framework of Reference for Languages) levels[23] range from A2 to B2. The CEFR levels are determined based on the learners’ vocabulary size and proficiency scores in English proficiency tests such as IELTS and TOEFL. As shown in Fig. 2, there is an imbalance in the distribution of user levels within the dataset.

*B. Configurations*

We adopt a two-stage pipeline for ASA:

- 1) **Transcription Stage:** Spoken responses are first transcribed using Whisper-large-v2, a state-of-the-art ASR model known for its robustness across accents and noise conditions.
- 2) **Classification Stage:** The transcribed text is then fed into a BERT-base-uncased model for proficiency classification. Token-level embeddings are aggregated using mean pooling, followed by a dropout layer and a linear classification head.

To ensure stable and effective learning, we configured training with a batch size of 4 and gradient accumulation over 2 steps (yielding an effective batch size of 8), trained for 20 epochs using a learning rate of  $1 \times 10^{-5}$ , and optimized with either the standard cross-entropy loss, focal loss, or the BLV loss. Where applicable, we employ mixed-precision (FP16)

(a) BLV  $\sigma = 6$ 

(b) Focal Loss

Fig. 3: Confusion Matrix on ICNALE Testing Set

training and apply a 0.1 dropout rate to the final layer of the feed-forward network.

### C. Evaluation Metrics

In order to validate whether our method effectively addresses the data imbalance issue in the ASA task, we employed a set of complementary evaluation metrics that capture different aspects of prediction quality.

To validate the effectiveness of our method on the ASA task, we employed Pearson Correlation Coefficient (PCC), Root Mean Square Error (RMSE), Accuracy, and Adjacent Accuracy as primary evaluation metrics. To further investigate the impact of data imbalance, we utilized macro-averaged variants of Accuracy and RMSE, which give equal weight to all classes regardless of their frequency. Additionally, the F1-score was included to balance precision and recall, particularly important for assessing performance on minority classes in imbalanced datasets. Collectively, these metrics provide a comprehensive framework to evaluate both overall prediction quality and fairness across proficiency levels.

## V. RESULTS

### A. Main Result on ICNALE

The results in Table I demonstrate the effectiveness of our BLV variance injection method in improving the model’s overall performance. We observe consistent gains in PCC, RMSE, standard accuracy, and adjacent accuracy, which collectively reflect better prediction accuracy and stronger correlation between predicted and true proficiency levels.

To specifically verify the method’s impact on data imbalance, improvements in macro RMSE and macro accuracy indicate a more balanced performance across all classes, mitigating bias toward majority classes. Additionally, the increase in F1-score highlights enhanced detection of minority classes by effectively balancing precision and recall, which is critical in imbalanced classification scenarios. Together, these metrics provide a comprehensive validation of our method’s ability to

improve both overall accuracy and fairness across imbalanced proficiency levels.

Among the BLV variants tested,  $\sigma = 6$  outperforms  $\sigma = 2$ , indicating that a larger variance injection plays a critical role in enhancing the model’s predictive capability. This suggests that increasing the variance allows the model to better address data imbalance and learn more discriminative representations. Importantly, both  $\sigma = 2$  and  $\sigma = 6$  settings deliver clear improvements over the baseline model, highlighting the robustness and effectiveness of our BLV approach across a range of parameter values. These findings underscore the importance of carefully tuning the variance parameter to fully realize the benefits of variance injection in mitigating imbalance-related biases.

We also evaluated the BLV loss against focal loss, a widely adopted objective for addressing class imbalance. Although focal loss leads to higher F1 scores and improved macro-accuracy, this gain comes at the cost of reduced overall accuracy and PCC. The confusion matrices in Fig. 3 highlight how these two methods differ in their error distributions.

Overall, these grouped evaluation metrics provide comprehensive evidence that our method effectively mitigates the impact of data imbalance, leading to a more equitable and accurate predictive model.

### B. Visualization of logits on ICNALE with BLV loss Injection

We further investigate how variance injection via BLV loss affects the model’s learned representations by comparing the output logits of the baseline BERT and our BERT + BLV  $\sigma = 6$  variant through t-distributed stochastic neighbor embedding (t-SNE) [24]. We applied t-SNE with a perplexity of 30 over 5,000 iterations to project each model’s high-dimensional logit vectors into two dimensions. As illustrated in Fig. 4, the BERT + BLV embeddings exhibit substantially larger inter-cluster separations and markedly tighter intra-cluster groupings than those of the baseline. This pronounced clustering behaviour confirms that injecting higher variance into minority-class




Fig. 4: t-SNE Visualization of the logits on ICNALE Training Set

logits enhances the model’s ability to discriminate between adjacent proficiency levels.

## VI. CONCLUSION

In conclusion, this study addresses the challenge of data imbalance in prediction models by introducing variance injection via BLV loss. Our method is rigorously validated across multiple evaluation metrics, providing strong evidence of its effectiveness in reducing bias and improving model robustness. Additionally, visualization of the learned representations confirms that variance injection enhances the model’s discriminative ability. These findings demonstrate the potential of our approach to develop fairer and more reliable models in imbalanced data scenarios.

## REFERENCES

- [1] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd international conference on Machine learning*, 2006.
- [2] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds., Curran Associates, Inc., 2015.
- [3] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. Mcleavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *Proceedings of the 40th International Conference on Machine Learning*, 2023, pp. 28 492–28 518.
- [4] Y. Wang, M. Gales, K. Knill, *et al.*, “Towards automatic assessment of spontaneous spoken english,” *Speech Communication*, pp. 47–56, 2018.
- [5] S. Bannò, K. M. Knill, M. Matassoni, V. Raina, and M. J. F. Gales, *L2 proficiency assessment using self-supervised speech representations*, Arxiv preprint arXiv:2211.08849, 2022.
- [6] T.-H. Lo, F.-A. Chao, T.-i. Wu, Y.-T. Sung, and B. Chen, “An Effective Automated Speaking Assessment Approach to Mitigating Data Scarcity and Imbalanced Distribution,” in *North American Chapter of the Association for Computational Linguistics*, 2024.
- [7] K. Zechner and K. Evanini, *Automated Speaking Assessment: Using Language Technologies to Score Spontaneous Speech*. 2019, Google-Books-ID: 88bAD-wAAQBAJ.
- [8] J. Xu, E. Jones, V. Laxton, and E. Galaczi, “Assessing L2 English speaking using automated scoring technology: Examining automarker reliability,” *Assessment in Education: Principles, Policy & Practice*, pp. 411–436, 2021.
- [9] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM transactions on audio, speech, and language processing*, pp. 3451–3460, 2021.
- [10] A. Baeovski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, “Data2vec: A general framework for self-supervised learning in speech, vision and language,” in *International conference on machine learning*, 2022.
- [11] A. Baeovski, Y. Zhou, A. Mohamed, and M. Auli, “Wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Advances in Neural Information Processing Systems*, 2020.
- [12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, *BERT: Pre-training of deep bidirectional transformers for language understanding*, Arxiv preprint arXiv:1810.04805, 2019.

- [13] S. Bannò and M. Matassoni, *Proficiency assessment of L2 spoken English using wav2vec 2.0*, Arxiv preprint arXiv:2210.13168, 2022.
- [14] M. Buda, A. Maki, and M. A. Mazurowski, “A systematic study of the class imbalance problem in convolutional neural networks,” *Neural networks*, pp. 249–259, 2018.
- [15] J. M. Johnson and T. M. Khoshgoftaar, “Survey on deep learning with class imbalance,” *Journal of Big Data*, 2019.
- [16] D. S. Park, W. Chan, Y. Zhang, *et al.*, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Interspeech*, 2019.
- [17] C. Shorten and T. M. Khoshgoftaar, “A survey on image data augmentation for deep learning,” *Journal of big data*, pp. 1–48, 2019.
- [18] Y. Wang, J. Fei, H. Wang, *et al.*, “Balancing logit variation for long-tailed semantic segmentation,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [19] E. Coutinho, F. Höning, Y. Zhang, *et al.*, “Assessing the prosody of non-native speakers of English: Measures and feature sets,” in *Proceedings of the Tenth International Conference on Language Resources and Evaluation*, 2016.
- [20] S. Bhat and S.-Y. Yoon, “Automatic assessment of syntactic complexity for spontaneous speech scoring,” *Speech Communication*, pp. 42–57, 2015.
- [21] Y. K. Singla, J. Shah, C. Chen, and R. R. Shah, “What do audio transformers hear? probing their representations for language delivery & structure,” in *IEEE International Conference on Data Mining Workshops (ICDMW)*, 2022.
- [22] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [23] C. o. E. C. f. C. C.-o. E. C. M. L. Division, *Common European Framework of Reference for Languages: Learning, Teaching, Assessment*. 2001, Google-Books-ID: PygQ8Gk4k4YC.
- [24] L. Van der Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, 2008.