

# Balanced Audiovisual Dataset for Imbalance Analysis

Wenke Xia<sup>1,†</sup>, Xu Zhao<sup>2,†</sup>, Xincheng Pang<sup>1</sup>, Changqing Zhang<sup>3</sup>, Di Hu<sup>1,4,\*</sup>

<sup>1</sup>Gaoling School of Artificial Intelligence, Renmin University of China, Beijing

<sup>2</sup>School of Computer Science and Engineering, Northeastern University, Shenyang

<sup>3</sup>College of Intelligence and Computing, Tianjin University, Tianjin

<sup>4</sup>Beijing Key Laboratory of Big Data Management and Analysis Methods, Beijing

Project website: <https://gewu-lab.github.io/Balanced-Audiovisual-Dataset/>

## 1. Introduction

The issue of imbalance has garnered significant attention across various domains, including loss imbalance in multi-task learning, category imbalance in long-tailed datasets, etc. Imbalance problems are also prevalent in the area of multimodal learning, where models tend to rely on the dominant modality in the presence of discrepancies between modalities [7]. Recent studies [7, 9] have endeavored to address the issue of modality imbalance. Wang et al. [9] found the phenomenon that the performance of multimodal models would be inferior to the unimodal model due to the modality overfitting behaviors. Peng et al. [7] and Wu et al. [10] proposed to solve the modality imbalance problem with the perspective of the optimization process, while Han et al. [5] alleviated this issue by making reliable multi-view fusion. However, we observe that although existing imbalance methods exhibit superior performance on the overall testing set, they even fail to perform better than the unimodal model when inference on some modality-preferred subsets as illustrated in Figure 1. This phenomenon has not been adequately addressed in prior works but it may impact the reliability of multimodal models in certain scenarios, such as those involving modality noise.

To provide a comprehensive analysis of this phenomenon, we first introduce a metric to estimate the sample-wise modality discrepancy with unimodal confidence and divide existing datasets into different modality-preferred subsets. Concretely, we define a modality-preferred subset as a set of samples whose corresponding modality confidence is higher than the other modality. Further, we evaluate existing imbalance methods and find that: although existing imbalance methods achieve better performance than unimodal modality on the overall testing set, they consistently perform worse than the unimodal model on the visual preferred testing subset across different datasets as demonstrated in Figure 1(a-b).

However, given the data-driven nature of deep learning [10], it is inevitable that the serious modality bias present in existing datasets such as Kinetics-Sound [1] and

VGG-Sound [2] will cause multimodal models to learn this bias. Thus, to further investigate the effectiveness of imbalance methods, we build a balanced audiovisual dataset containing samples with various modality discrepancies, and the allocation of such modality discrepancy remains uniformly distributed over the dataset, as shown in Figure 1(c).

Based on our balanced dataset, we re-evaluate existing imbalance methods and find results consistent with the previous phenomenon: there still exists no imbalance method exceeding all unimodal models on every modality-preferred testing subset as shown in the histogram of Figure 1(c). We further select testing subsets with different modality discrepancies to evaluate the performance of existing imbalance methods and surprisingly find that the larger the modality discrepancy is, the worse the multimodal models perform compared with the unimodal one. This phenomenon indicates that existing imbalance methods have difficulty handling multimodal data with large modality discrepancy. Thus, the evaluation of different modalities discrepancies should also be considered to comprehensively analyze the performance of multimodal models.

To summarize, our contributions are as follows:

- We introduce an evaluation metric to estimate the modality discrepancy and point out that existing audiovisual datasets suffer from serious modality bias.
- We provide a balanced audiovisual dataset with uniformly distributed modality discrepancies to support further analysis for multimodal models.
- We find that existing multimodal models would exhibit inferior performance than the unimodal one in scenarios with serious modality discrepancy. We highlight the importance of addressing this issue in future research to ensure the reliability of multimodal models across various application scenarios.

## 2. Evaluation Methods

### 2.1. Modality discrepancy estimation

Inspired by some sample difficulty measure work in unimodal datasets [3], we provide an estimation method for

<sup>†</sup> Equal contribution. \* Corresponding author.

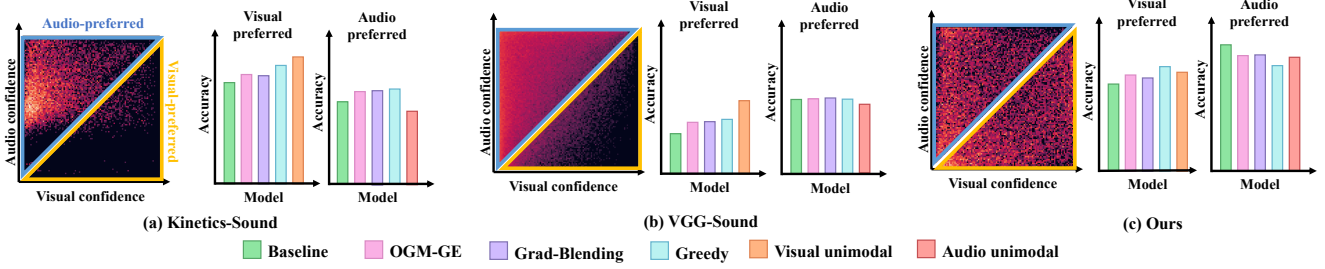


Figure 1: (a-b) show the heatmaps of modality discrepancy distribution and classification results on the existing datasets. In the heatmap, the color represents the proportion of samples, the brighter area indicates larger samples. (c) shows our balanced heatmap and the classification results.

measuring the sample-wise modality discrepancy, which utilizes the difference between unimodal confidence assigning to the true label during the training process.

Concretely, we consider a training dataset of size  $N$ ,  $S = \{(X_i, Y_i)\}_{i=1}^N$  where the  $i$ th sample contains the observation,  $X_i$  and its true label,  $Y_i$ . In multimodal datasets, we consider each observation  $X_i$  consists of  $M$  inputs from different modalities, referred as  $X_i = (x_i^{m_1}, \dots, x_i^{m_k}, \dots, x_i^{m_M})$ , where  $m_k$  refer to the  $k$ th modality respectively. We first train each unimodal model with an individual modality for  $E_{m_k}$  epochs, and utilize the average confidence of unimodal to represent the learning difficulty for the  $i$ th sample:

$$\hat{u}_i^{m_k} = \frac{1}{E_{m_k}} \sum_{e=1}^{E_{m_k}} p_\theta^{m_k}(e)(y_i|x_i^{m_k}),$$

where  $p_\theta^{m_k}(e)$  presents the confidence of unimodal model with parameters  $\theta$  at the  $e$ th epoch.

We divide such average unimodal confidence scores into 0.01 intervals, then count the samples of corresponding intervals and calculate the proportion, which could be regarded as a distribution of modality discrepancy. We present the resulting distribution as a heatmap in Figure 1, where the coordinates correspond to the average unimodal confidence scores and the color represents the proportion of samples within the corresponding interval. Our method thus provides an intuitive method to visualize the modality discrepancies in multimodal datasets.

## 2.2. Dataset split strategy

We use the average unimodal confidence  $\hat{u}_i^a$  and  $\hat{u}_i^v$ , to represent the unimodal confidence on audio and visual modalities. To investigate modality bias in audiovisual datasets, we split an existing audiovisual dataset into two subsets based on the unimodal confidence scores. We name these subsets the audio-preferred subset and the visual-preferred subset. If  $\hat{u}_i^a > \hat{u}_i^v$ , the sample would be split into the audio-preferred subset, while if  $\hat{u}_i^v > \hat{u}_i^a$ , the sample would be split into the visual-preferred subset instead.

With our split strategy, we divide the testing set of existing datasets into different modality-preferred subsets.

As depicted in Figure 1(a-b), our estimation results prove that there exists serious modality bias in these audiovisual datasets, which may not be adequate for comprehensive analysis of the existing modality imbalance methods.

## 3. Experiments on Existing Datasets

### 3.1. Experiments settings

We estimate the modality discrepancy of Kinetics-Sound [6] and VGG-Sound [2] datasets, and further, evaluate existing imbalance methods on these datasets. We use ResNet-18 as the backbone for both modalities. Following Peng et al. [7], for audio input, we extract the spectrogram with a sample rate of  $48k Hz$ , while we extract frames with 1fps and randomly select frames as the visual input.

For modality discrepancy estimation, we train each modality separately to evaluate the unimodal confidence. Given that the coverage speed of each modality varies in different datasets, we adopt an early-stop strategy during training to obtain reliable unimodal confidence.

### 3.2. Experiments results

In this section, we conduct a comprehensive evaluation of existing imbalance methods using different subsets, with a vanilla late-fusion method as the baseline. Specifically, we first evaluate the performance on the overall testing set, then compare the results on modality preferred subsets referring to the split strategy.

As presented in Table 1, we find some results consistent with the recent works [7]: all imbalance methods perform better than the unimodal model on the overall set, which proves the effectiveness of multimodal learning. Besides, compared with the baseline, all imbalance methods gain significant improvement across different datasets, which indicates that existing methods alleviate the modality imbalance problem to some extent.

Although the imbalance methods achieve great performance on the overall subset, we observe that existing methods consistently perform worse than the unimodal model in the weak-modality preferred subset. Concretely, consider-

Methods	Kinetics-Sound			VGG-Sound		
	Overall	Audio preferred	Visual preferred	Overall	Audio preferred	Visual preferred
Baseline	61.73	60.96	69.09	48.65	52.38	27.50
OGM-GE	64.57	63.20	74.44	50.89	53.43	35.63
Grad-Blending	66.00	64.34	73.82	<b>51.47</b>	<b>53.73</b>	35.87
Greedy	<b>66.35</b>	<b>64.51</b>	80.75	50.32	52.68	37.37
Audio unimodal	51.52	54.32	31.23	41.25	45.39	10.86
Visual unimodal	47.67	42.28	<b>86.43</b>	32.44	28.65	<b>51.67</b>

Table 1: Evaluation results on the existing audiovisual dataset.

ing the VGG-Sound [2] dataset is a curated sound-oriented dataset, where the audio-preferred subset contains a significantly larger number of samples than the visual-preferred subset. The visual modality is weaker in comparison to the audio modality. However, Table 1 shows that although the visual unimodal performance is worse than existing imbalance methods on the overall set, the visual unimodal model exceeds other multimodal models with a large margin on the visual-preferred subset, which hints that existing imbalance methods still fail to fully utilize the weak modality.

However, considering the greedy nature of deep learning [10], the multimodal models would inevitably learn the modality biases inherent in these biased audiovisual datasets, which would potentially impact the evaluation of these imbalance methods.

#### 4. The Balanced Audiovisual Dataset

To alleviate the influence of modality bias and further study the effectiveness of imbalance methods, we build a balanced audiovisual dataset with full consideration of the modality discrepancy. As shown in Figure 1(c), our balanced dataset contains samples with various modality discrepancy, and the allocation of such modality discrepancy is uniformly distributed over the dataset.

To automatically collect videos with corresponding labels, we filter the crawled videos with the off-the-shelf pretrained unimodal models (SlowFast [4] pretrained on Kinetics-400 [6] for visual modality and VGGish [8] pretrained on VGG-Sound [2] for audio modality). Constrained by the pretrained datasets, we select several categories with the same semantic concept existing in both datasets, then integrate them into 30 categories and crawl videos from YouTube for corresponding topics.

Then we extract the full videos into 10-seconds clips and evaluate unimodal confidence for each clip with pretrained model. Referring to the unimodal confidence, we split the clips into three types, called the high-correspondence clips, audio-correct clips, and visual-correct clips to build our dataset, where the high-correspondence clips mean both unimodal models predict true label, audio-correct clips mean only audio modality model predict the true label while the visual-correct clips are on the opposite. With such clip

selection strategy, we could preliminarily build a nearly balanced audiovisual dataset, which contains audiovisual clips of three types with similar proportion.

However, the preliminary dataset suffers from long-tailed distribution and slight modality bias due to the limitation of crawling from the website, such as the instruments videos are easy to collect while some human action categories are difficult to obtain, etc. Thus, to deal with such problem, we screen extra samples from VGG-Sound [2] and Kinetics-400 [6] and incorporate them into our dataset. In the end, we estimate the modality discrepancy of our dataset and recalibrate by removing noise and some data from dominated modality referring to such estimation. Finally, our final overall dataset retains 7,929 samples from the Kinetics-400 dataset, 14,498 samples from the VGG-Sound dataset, and 13,424 samples collected from YouTube. Our dataset remains a uniformly distributed modality discrepancy over the dataset, as shown in Figure 1(c).

### 5. Experiments on Our Dataset

#### 5.1. Testing on modality-preferred subset

As Table 2 shows, when trained on our balanced audiovisual dataset, all the imbalance methods exhibit superior performance than the baseline and the unimodal model on the overall set, which proves the effectiveness of existing imbalance methods.

However, there still exists no method that exceeds all unimodal models on every modality-preferred subset. The Baseline, OGM-GE [7] and Grad-Blending [9] methods gain significant improvement on the audio-preferred subset, but fail to perform better than the visual unimodal model in the visual-preferred subset. While the Greedy [10] method performs better than the visual unimodal model but less than the audio unimodal model in the audio-preferred subset.

The results of Table 2 indicate that although existing imbalance methods have achieved notable enhancements, they still have difficulty handling samples with modality discrepancy. To analyze the influence of the modality discrepancy on multimodal models, we propose to evaluate these methods with various modality discrepancy subsets.

Methods	Overall	Audio preferred	Visual preferred
Baseline	70.47	78.62	59.37
OGM-GE	<b>72.27</b>	78.92	63.59
Grad-Blending	72.04	<b>79.49</b>	61.89
Greedy	72.01	72.46	<b>71.41</b>
Audio unimodal	62.96	77.76	44.79
Visual unimodal	49.82	36.79	67.58

Table 2: Evaluation results on our dataset.

## 5.2. Testing with different modality discrepancy

Given we have obtained the average unimodal confidence  $\hat{u}_i^m$  of each sample, we propose a threshold  $T$  to split the testing subset. Concretely, if  $\hat{u}_i^a - \hat{u}_i^v > T$ , the sample would be split into the audio-dominated subset, while if  $\hat{u}_i^v - \hat{u}_i^a > T$ , the sample would be split into the visual-dominated subset. When threshold  $T$  increases, the modality discrepancy in the corresponding subset would be larger. In experiments, we set the threshold  $T$  as 0, 0.2, and 0.4 to select the subset with different modality discrepancies.

We provide the results with different thresholds  $T$  on the audio-dominated subset and the visual-dominated subset in Figure 2. When we only consider the  $T = 0$  scenarios, there exist some multimodal models exceeding the unimodal model on the corresponding subset (e.g., the OGM-GE in the audio-dominated subset and the Greedy in the visual-dominated subset). However, when the threshold  $T$  increases, only the samples with larger modality discrepancy would remain in the modality-dominated subset, and the unimodal model performs better than other multimodal models on the corresponding modality-dominated subset (The audio model performs better than other imbalance methods on the audio dominated subset when  $T = 0.4$ ).

The multimodal models are expected to perform better than unimodal models, however, the extra information instead hinders the performance on multimodal data of large modality discrepancy. Considering the samples with larger modality discrepancy would have more modality noise, these results indicate that the modality noise would affect the performance of multimodal models.

## 6. Conclusion

In this work, we first introduce a modality discrepancy estimation metric and find that existing audiovisual datasets suffer from serious modality bias. To support the comprehensive analysis of the modality imbalance methods, we build a balanced audiovisual dataset with uniformly distributed modality discrepancies. Further, we evaluate existing imbalance methods with different modality discrepancy subsets and find that multimodal models would exhibit inferior performance than the unimodal one in some scenarios with serious modality discrepancy. This compre-

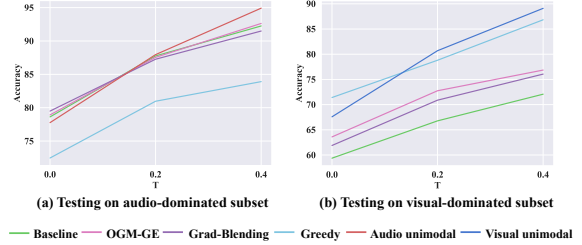


Figure 2: Evaluation results on modality-dominated subsets with different threshold  $T$ .

hensive evaluation method with different modality discrepancies should be considered to make sure the reliability of multimodal models in different application scenarios.

## 7. Acknowledgement

This research was supported by National Natural Science Foundation of China (NO.62106272), the Young Elite Scientists Sponsorship Program by CAST (2021QNR001), and Public Computing Cloud, Renmin University of China.

## References

- [1] R. Arandjelovic and A. Zisserman. Look, listen and learn. In *ICCV*, pages 609–617, 2017. 1
- [2] H. Chen, W. Xie, A. Vedaldi, and A. Zisserman. Vggsound: A large-scale audio-visual dataset. In *ICASSP*, pages 721–725. IEEE, 2020. 1, 2, 3
- [3] K. Ethayarajh, Y. Choi, and S. Swayamdipta. Information-theoretic measures of dataset difficulty. *arXiv preprint arXiv:2110.08420*, 2021. 1
- [4] H. Fan, Y. Li, B. Xiong, W.-Y. Lo, and C. Feichtenhofer. Pyslowfast. <https://github.com/facebookresearch/slowfast>, 2020. 3
- [5] Z. Han, C. Zhang, H. Fu, and J. T. Zhou. Trusted multi-view classification with dynamic evidential fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2022. 1
- [6] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 2, 3
- [7] X. Peng, Y. Wei, A. Deng, D. Wang, and D. Hu. Balanced multimodal learning via on-the-fly gradient modulation. In *CVPR*, pages 8238–8247, 2022. 1, 2, 3
- [8] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [9] W. Wang, D. Tran, and M. Feiszli. What makes training multi-modal classification networks hard? In *CVPR*, pages 12695–12705, 2020. 1, 3
- [10] N. Wu, S. Jastrzebski, K. Cho, and K. J. Geras. Characterizing and overcoming the greedy nature of learning in multimodal deep neural networks. In *ICML*, pages 24043–24055. PMLR, 2022. 1, 3