

BRIDGING THE GAP BETWEEN TEXT, AUDIO, IMAGE, AND ANY SEQUENCE: A NOVEL APPROACH USING GLOSS-BASED ANNOTATION

Sen Fang^{1,*}, Yaling Feng², Sizhou Chen³, Xiaofeng Zhang⁴, Teik Toe Teoh⁵

^{1,2}Victoria University, ³Chengdu University of Information Technology,
⁴Shanghai Jiao Tong University, ⁵Nanyang Technological University

ABSTRACT

This paper presents an innovative approach called **BGTAI** to simplify multimodal understanding by utilizing gloss-based annotation as an intermediate step in aligning Text and Audio with Images. While the dynamic temporal factors in textual and audio inputs contain various predicate adjectives that influence the meaning of the entire sentence, images, on the other hand, present static scenes. By representing text and audio as gloss notations that omit complex semantic nuances, a better alignment with images can potentially be achieved. This study explores the feasibility of this idea, specifically, we first propose the first **Langu2Gloss** model and then integrate it into the multimodal model UniBriVL for joint training. To strengthen the adaptability of gloss with text/audio and overcome the efficiency and instability issues in multimodal training, we propose a **DS-Net** (Data-Pair Selection Network), an **Result Filter** module, and a novel **SP-Loss** function. Our approach outperforms previous multimodal models in the main experiments, demonstrating its efficacy in enhancing multimodal representations and improving compatibility among text, audio, visual, and any sequence modalities.

Index Terms— Multimodal Learning, Data Enhancement, Neural Network, Audio-Visual, Image Generation

1. INTRODUCTION

In recent years, there has been a growing trend towards the integration of multiple modalities such as language, audio, and vision. Numerous studies [1, 2, 3, 4] have explored various methods to effectively fuse three or even up to twelve modalities into a unified representation. Capitalizing on the availability of large-scale datasets, these approaches primarily rely on massive pretraining to train models that can then be transferred to various downstream tasks, yielding impressive results. Most of these studies are based on transformer architectures or modified versions thereof. Depending on their specific usage and architectural details, they employ diverse fusion techniques to integrate the different modalities effectively. The prevailing approach [1, 2] involves jointly training based on specific architectures and then manually converting them into the desired task format. However, such parameters often fail to be effectively shared across modalities. The latest approach [3, 4] aims to achieve parameter sharing and structural generalization by employing a universalized architecture, such as sharing encoders or utilizing multi-path transformers. This facilitates task transferability and enhances the usability of the shared structure. Nevertheless, previous methods have not effectively achieved parameter sharing among multiple modalities, as they only store them within a single model.

* The work was done at MMLab@NTU during the internship of the first author, code will available at <https://www.mmlab-ntu.com/> after publication.

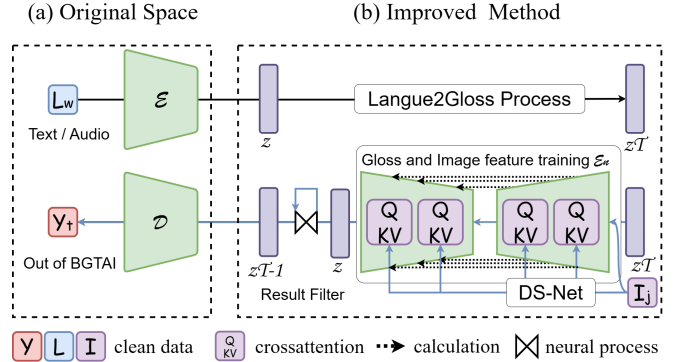


Fig. 1: (a) BGTAI Production Pipeline and (b) Structure Drawing. It reveals a structure that unifies all representation possibilities.

Gloss is an annotation method [5] applied in sign language, where certain postures or gestures (essentially sequences) can be represented as Gloss. In other words, it serves as an intermediary medium connecting natural language and complex sequences. In the field of sign language, such as sign language translation or sign language generation, utilize Gloss to enhance their model capabilities [6, 7]. Moreover, in reality, the application of Gloss has been employed for several decades to enhance language-gesture learning abilities for individuals with hearing impairments [5]. Therefore, we believe that incorporating real-world scenarios into multimodal learning can yield promising outcomes, based on the highly beneficial and validated methods. UniBriVL is a state-of-the-art multimodal model [8] for language-visual tasks that utilize a specialized encoder to jointly train text and audio, resulting in a universal language representation. This representation is then utilized to train the model for the task of matching with images. Our exploratory work focused on this model in the form of components.

In this study, we propose the **BGTAI** (Bridging the Gap between Text, Audio, Image, and any Sequence) framework aimed at addressing the disparities among text, audio, images, and any other sequences. The proposed framework comprises several components including the groundbreaking Langu2Gloss model, the DS-Net model, a Result Filter module, and a novel SP-Loss function, which are designed to solve some of the problems encountered in model training and application. We conducted a comprehensive evaluation involving various qualitative and quantitative assessments, such as baseline comparisons (including audio retrieval, audio classification, audio captioning, image retrieval, etc.), ablation experiments, and training efficiency tests. The experimental results show our method's superior effectiveness, surpassing previous similar approaches in multimodal tasks, and highlighting its significance for future multimodal model development and application.

2. PROPOSED METHOD

In this chapter, we provide a detailed introduction to our motivation and propose several components including Langu2Gloss, DS-Net, Result Filter, and SP-Loss as part of our method.

2.1. Motivation

The motivation behind our research is the need to address the limitations of existing multimodal training approaches in aligning image (A kind of sequence) with text/audio. To overcome these limitations, we propose the DS-Net, a Data-Pair Selection Network, which enhances alignment by selecting suitable training data pairs. We also introduce an improvement module to address efficiency and stability issues in multimodal training. Furthermore, we present a novel SP-Loss function that optimizes the training process and improves overall model performance. Our work also integrates the Langu2Gloss model into the UniBriVL for joint training, enhancing the adaptability of the image with langu. Through the incorporation of the DS-Net, Result Filter module, and SP-Loss function, we aim to provide more efficient and effective solutions for multimodal alignment.

2.2. Our components

2.2.1. Langu2Gloss

Given an audio/text sentence, $L^N = \{w_1, w_2, \dots, w_N\}$, use SpeechLM [9] encoder (It can unify audio and text into a single representation) maps the sequence into a latent representation as follows:

$$o_{1:N}, h_L = \text{Encoder}(L^N) \quad (1)$$

Here, $o_{1:N}$ represents the output of the encoder for each word w , and h_L is the hidden representation of the provided sentence. The hidden representation and the encoder outputs are then passed to our decoder. The decoder utilizes an attention mechanism to generate sign gloss sequences, one gloss at a time:

$$\text{gloss}_m = \text{Decoder}(h_L, \alpha(o_{1:N}), \text{gloss}_{m-1}) \quad (2)$$

In the above equation, $\alpha(\cdot)$ denotes the attention function, and gloss_m is the gloss produced at time step m . We ingeniously leverage the original text/audio inputs and the textually summarized/summarized audio outputs from the How2 dataset [10] as our training set, with Gloss serving as an intermediary.

2.2.2. DS-Net

To co-articulate between inputs and gloss, we propose a Data-Pair Selection Network (DS-Net) that learns to predict the temporal alignment between the original text/audio input and the summarized text/summary audio output. The DS-Net predicts a discrete sparse monotonic temporal alignment path, represented by a binary decision matrix \hat{A} of size $Q \times T$, where Q is the length of the summarized sequence and T is the length of the original sequence.

This alignment path determines which data-pairs to select and which data-pairs to skip in order to generate a co-articulated continuous sequence. T is the length of the original sequence. Formally, the DS-Net predicts the alignment path as follows:

$$\hat{A} = \text{DS-Net}(R, h_{1:W}) \quad (3)$$

$$\hat{Y} = Y \times \hat{A} \quad (4)$$

This operation allows us to map sequences of varying lengths, with the end of the sequence determined by the alignment selection of the final data-pair, which is built into the Langu2Gloss model.

2.2.3. Result Filter

DS-Net can optimize output for any sequence if glosses are replaced with images or any sequence by using the Dynamic Time Warping [11] (DTW) algorithm as a supervision signal, at this point we call this module the Result Filter module. We pre-compute the DTW path, denoted as \mathcal{A}^* , between the interpolated text/audio dictionary sequence \mathcal{I} and the target image continuous sequence \mathcal{Y} . To train the Result Filter, we compute a cross entropy loss \mathcal{L}_{CE} between the predicted 1D temporal alignment, $\hat{A} \in \mathbb{R}^Q$, and the ground truth DTW alignment, $A^* \in \mathbb{R}^{Q \times 1}$, as:

$$\mathcal{L}_{CE}(\hat{A}, A^*) = -\frac{1}{Q} \sum_{q=1}^Q A_q^* \cdot \log(\hat{A}_q) \quad (5)$$

The final continuous image sequence, $\hat{Y} = (y_1, \dots, y_T)$, is produced as shown in Eq. 1. The first two components are also in it.

2.2.4. SP-Loss.

Specifically, the \mathcal{A}^* and \hat{A} obtained through the DTW algorithm can be utilized for calculating the SP-Loss. To optimize the SP-Loss, we introduce a normalization step before computing the loss. Firstly, we normalize \mathcal{A}^* and \hat{A} to ensure that their values fall within the range of $[0, 1]$. Then, we use the normalized \mathcal{A}^* and \hat{A} to calculate the SP-Loss. The SP-Loss is computed as follows:

$$\text{SP-Loss} = \frac{1}{Q} \sum_{q=1}^Q \left| \mathcal{A}_q^* - \hat{A}_q \right| \quad (6)$$

Here, Q represents the length of the sequence when calculating. The aforementioned process describes the computation of the optimized SP-Loss based on the input provided by the DTW algorithm.

3. EXPERIMENTS

In this section, we assess the performance and effectiveness of our model, including baseline comparison, training efficiency evaluation, ablation study, qualitative evaluation, and zero-shot comparison with previous works.

3.1. Experimental Setup

Table 1 shows the datasets used in our primary downstream tasks and their corresponding details, encompassing classification tasks such as multi-class (MC), multi-label (ML), and zero-shot (ZS), retrieval tasks including audio retrieval (AR) and cross-modal retrieval (CMR), as well as the audio captioning (AC) task.

Dataset	Task	# Clip (Split)	# Class	Metric
ESC-50 [12]	MC/ZS	2k (5 folds)	50	ACC
UrbanSound8K [13]	MC/ZS	8k (10 folds)	10	ACC
FSD50K [14]	ML/ZS	50k	200	mAP
VGGSound [15]	MC/ZS	185k	309	mAP
TAU Audio [16]	MC/ZS	190k	10	ACC
DESED [17]	AR	2.5k (valid)	10	F1
VGGSound [15]	CMR	15k (test)	309	MRR
Clotho [18]	AC	5k (evaluation)		COCO ¹ [19]

Table 1: Datasets for Downstream Tasks and Evolution Content.

¹<https://github.com/tylin/coco-caption>. Other datasets are mentioned later.

Model	Classification					Retrieval		
	ESC-50	UrbanSound8K	FSD50K	VGGSound	TAU Audio Only	DESED (AR)	VGGSound (CMR)	
	ACC	ACC	mAP	mAP	ACC	F1	A→I (MRR)	I→A (MRR)
Supervise	0.5300	0.6289	0.3212	0.4512	0.3827			
OpenL3 [20]	0.723	0.7681	0.4053	0.3405	0.651 [16]	0.1176	0.0171	0.0160
YamNet	0.8505	0.7832	0.5039	0.4541	0.5667	0.3420		
Wav2CLIP [21]	0.8595	0.8101	0.4308	0.4663	0.6302	0.3955	0.0566	0.0678
WavBriVL [22]	0.9117	0.8832	-	0.4741	-	0.3720	0.0611	0.0608
UniBriVL	0.9307	0.8722	-	0.4885	-	0.4111	0.0641	0.0612
BGTAI	0.9455	0.8911	0.4748	0.5130	0.6932	0.4351	0.0622	0.0745
BGTAI (ZS)	0.5015	0.4896	0.0365	0.1214	0.2089			
SOTA	0.959 [23]	0.8949 [23]	0.5671 [24]	0.544 [25]	0.687 [26]			

Table 2: The downstream classification and retrieval tasks’ results are compared with supervised training from scratch, OpenL3, YamNet, Wav2CLIP, WavBriVL, UniBriVL and current state-of-the-art (SOTA) models. Zero-shot classification uses the BGTAI model.

Model	B1	B4	M	RL	Cr	S	Sr
Baseline ⁸ [18]	0.389	0.015	0.084	0.262	0.074	0.033	0.054
Wav2CLIP	0.393	0.054	0.104	0.271	0.100	0.045	0.073
WavBriVL	0.412	0.92	0.108	0.268	0.126	-	-
UniBriVL	0.434	0.107	0.115	0.300	0.11	0.050	0.080
BGTAI	0.476	0.185	0.126	0.328	0.121	0.054	0.088

Table 3: The audio captioning results are compared to the baseline, omitting Bleu2/3 and including Bleu1/4 (B1/4), METEOR (M), ROUGEL (RL), CIDEr (Cr), SPICE (S), and SPIDer (Sr) metrics due to space limitations.

3.2. Evaluation for Performance

3.2.1. Baseline Comparison

Table 2 compares the results of different models on classification and retrieval tasks. Models like OpenL3, YamNet, Wav2CLIP, WavBriVL, and UniBriVL achieve varying levels of accuracy on different datasets. The BGTAI model performs well in most tasks, with the highest accuracy on ESC-50 and UrbanSound8K datasets. In retrieval tasks, DESED (AR) evaluates based on the F1 score, where BGTAI achieves the highest score. For audio-visual cross-modal retrieval, UniBriVL achieves the highest MRR for A→I and BGTAI for I→A. The results presented in Table 2 demonstrate the effectiveness of different models on downstream classification and retrieval tasks. The BGTAI model consistently achieves competitive performance, outperforming other models on certain tasks.

Table 3 compares the results of different models for audio captioning using multiple evaluation metrics. The proposed models outperform the baseline. Wav2CLIP shows improved Bleu scores and higher METEOR and ROUGEL scores. WavBriVL performed well in most experiments and significantly improved METEOR. The UniBriVL model also performs well, particularly in terms of METEOR and other metrics. Finally, the BGTAI model outperforms all others, achieving the highest scores across all metrics. Most of the data came from the original paper, and we tried to keep the Settings consistent with the original paper in all tests. It’s worth noting that some previous models have been fine-tuned on specific tasks that we haven’t, which may result in our advantage not being as obvious as it should be. Overall, the proposed method and models demonstrate superior performance and potential in audio captioning tasks.




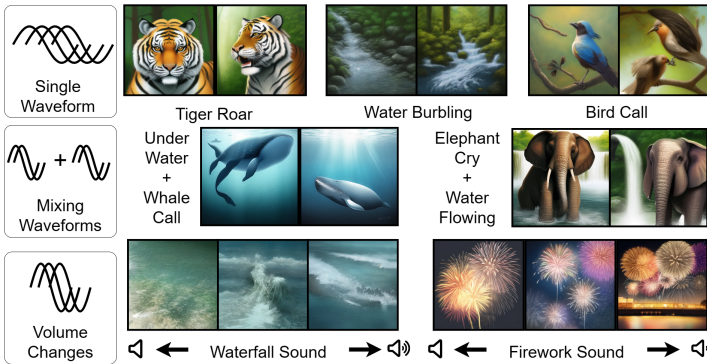
Fig. 2: Training efficiency evaluation. We tested the performance of models under different settings on the ESC-50 dataset. We sampled the performance of different periods, determined by epoch.

3.2.2. Training Efficiency Study

According to the Fig 2, the Base model exhibits the lowest ACC values across all periods, indicating its relatively lower efficiency in training. On the other hand, the set to "DS-Net" and "Result Filter" models show similar ACC values, indicating their comparable training efficiency. Interestingly, the set to "DS&RF" model demonstrates consistently higher ACC values across all periods, suggesting its superior training efficiency compared to the other models. In summary, the comparison of ACC values in different periods clearly indicates that the set to "DS&RF" model outperforms the other model configs in terms of training efficiency.

We found the reason based on the results of the observation experiment, as shown in the bar chart, the "Result Filter" performance is not particularly impressive due to its comparatively long training time requirement. On the other hand, "DS-Net" is capable of eliminating poor data, resulting in a shorter training time requirement and consequently better performance within the same period. When the two are combined, they produce a better result than if set alone. Based on our speculation, theoretically, they perform better because they complement each other’s shortcomings when combined: when used at the same time, the training time remains the same, while the overall efficiency performance becomes better than the individual.

Generated Images from Waveform Manipulation



Generated Images from Latent Manipulation




Fig. 3: Sound-to-image generation. Our model introduces a groundbreaking approach for synthesizing natural scene images from sound. It is trained solely on paired audio-visual data, eliminating the need for labels or language supervision. Importantly, our model offers remarkable controllability by manipulating input waveforms (left) and controlling the generated images in the latent space of diffusion models [27] (right), the I_j indicates image embedding. This innovative methodology provides enhanced flexibility and control over the model’s outputs.

Loss	D	Duration	VGGSound (50 classes)			
			R@1	R@5	FID (\downarrow)	IS (\uparrow)
(A) L_2	✓	10 sec.	14.57	37.35	19.24	7.98
(B) $L_{SP-Loss}$	✓	10 sec.	30.05	62.74	25.70	12.27
(C) L_{total}		10 sec.	36.10	70.89	20.52	17.00
(D) L_{total}	✓	1 sec.	34.78	73.62	19.37	19.25
(E) L_{total}	✓	5 sec.	39.00	79.51	21.56	20.16
(F) L_{total}	✓	10 sec.	37.89	81.22	16.88	20.96

Table 4: Ablation studies of our proposed method. We compare the different configurations of our method by changing the loss functions, frame selection method, and duration of the audio. D denotes the Data-Pair Selection Network.

3.2.3. Ablation Study

Table 4 presents ablation studies of our proposed method, evaluating different configurations by varying loss functions, frame selection, and audio duration. Configuration (B) improves the basic setup by incorporating the Self-Paced Learning loss, achieving higher performance. Removing the frame selection network and focusing on the total loss in configuration (C) further improves results. Surprisingly, reducing audio duration to 1 and 5 seconds in configurations (D) and (E) maintains high performance. Configuration (F) combines the advantages of previous setups and achieves the best performance, generating high-quality audio-visual embeddings.

3.2.4. Zero-Shot Comparison with Previous Works

Table 5 presents a zero-shot study on emergent audio retrieval and classification. The table compares methods based on their use of audio and text supervision or loss. In the first category, AudioCLIP achieves the highest Top-1 accuracy on the ESC dataset. AVFIC, in the second category, performs reasonably on Clotho and AudioCaps, but corresponding results for the ESC dataset are unavailable. In the third category, our proposed method, BGTAI, exhibits emergent capabilities and outperforms previous work on all datasets. In addition, we also provide performance comparisons with Imagebind, a relatively new multimodal large-scale model. Despite BGTAI having significantly fewer parameters, it achieves comparable performance to imagebind. Overall, BGTAI demonstrates superior performance in emergent zero-shot audio retrieval and classification, surpassing existing methods with and without audio and text supervision.

Method	Emergent	Clotho		AudioCaps		ESC
		R@1	R@10	R@1	R@10	Top-1
<i>Uses audio and text supervision</i>						
AudioCLIP [23]	×	-	-	-	-	68.6
<i>Uses audio and text loss</i>						
AVFIC [28]	×	3.0	17.5	8.7	37.7	-
<i>No audio and text supervision</i>						
BGTAI (Ours)	✓	4.2	24.9	9.5	39.5	62.3
Imagebind	✓	6.0	28.4	9.3	42.3	66.9

Table 5: Zero-Shot Study. Emergent zero-shot audio retrieval and classification, tests performed on Clotho [18], AudioCaps [29] and ESC [12] datasets, models are all default Settings.

3.3. Evaluation for Qualitative

As shown in Figure 3, we conducted a qualitative evaluation to assess the capability of BGTAI in sound-guided image generation. It includes two major categories of demonstrations: direct manipulation of audio attributes such as volume and frequency, and audio embedding affecting image embedding. The evaluation demonstrates that our model successfully captures the intricate relationship between sound and visual content, translating auditory information into meaningful visual representations. Through these qualitative assessments, it is evident that our model achieves impressive performance in generating visual content from sound, offering exciting prospects for applications in creative industries and beyond.

4. CONCLUSION

In this paper, we propose the BGTAI framework, which aims to bridge the gap between text, audio, images, and any other sequences in multimodal learning. The key component of our framework is the Langue2Gloss model, which represents text and audio inputs as gloss notations to better align them with images. We also introduce the DS-Net model, a Result Filter module, and a novel SP-Loss function to enhance the adaptability of gloss and improve training efficiency and stability. Our method outperforms previous multimodal models in various tasks, including audio retrieval, audio classification, audio captioning, and image retrieval. These findings suggest that the BGTAI framework has great potential for advancing multimodal learning and its application in real-world scenarios.

5. REFERENCES

- [1] Nanyi Fei, Zhiwu Lu, Yizhao Gao, Guoxing Yang, Yuqi Huo, Jingyuan Wen, Haoyu Lu, Ruihua Song, Xin Gao, Tao Xiang, et al., “Towards artificial general intelligence via a multimodal foundation model,” *Nature Communications*, vol. 13, no. 1, pp. 1–13, 2022.
- [2] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al., “Image as a foreign language: Beit pretraining for all vision and vision-language tasks,” *arXiv preprint arXiv:2208.10442*, 2022.
- [3] Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra, “Imagebind: One embedding space to bind them all,” 2023.
- [4] Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue, “Meta-transformer: A unified framework for multimodal learning,” 2023.
- [5] Samuel J. Supalla, Jody H. Cripps, and Andrew P. J. Byrne, “Why american sign language gloss must matter,” *American Annals of the Deaf*, vol. 161, no. 5, pp. 540–551, 2017.
- [6] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden, “Sign language transformers: Joint end-to-end sign language recognition and translation,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [7] Ben Saunders, Necati Cihan Camgoz, and Richard Bowden, “Signing at scale: Learning to co-articulate signs for large-scale photo-realistic sign language production,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 5141–5151.
- [8] Sen Fang, Bowen Gao, Yangjian Wu, Jingwen Cai, and Teik Toe Teoh, “Unibrivl: Robust universal representation and generation of audio driven diffusion models,” 2023.
- [9] Ziqiang Zhang, Sanyuan Chen, Long Zhou, Yu Wu, Shuo Ren, Shujie Liu, Zhuoyuan Yao, Xun Gong, Lirong Dai, Jinyu Li, and Furu Wei, “Speechlm: Enhanced speech pre-training with unpaired textual data,” 2023.
- [10] Roshan Sharma, Shruti Palaskar, Alan W Black, and Florian Metze, “End-to-end speech summarization using restricted self-attention,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 8072–8076.
- [11] Donald J Berndt and James Clifford, “Using Dynamic Time Warping to Find Patterns in Time Series,” in *AAAI-94 Workshop on Knowledge Discovery in Databases*, 1994.
- [12] Karol J. Piczak, “ESC: Dataset for Environmental Sound Classification,” in *ACM Multimedia*. 2015, p. 1015, ACM Press.
- [13] J. Salamon, C. Jacoby, and J. P. Bello, “A dataset and taxonomy for urban sound research,” in *ACM Multimedia*, Orlando, FL, USA, Nov. 2014, pp. 1041–1044.
- [14] Eduardo Fonseca, Xavier Favory, Jordi Pons, Frederic Font, and Xavier Serra, “Fsd50k: an open dataset of human-labeled sound events,” *arXiv preprint arXiv:2010.00475*, 2020.
- [15] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman, “Vggsound: A large-scale audio-visual dataset,” in *ICASSP*. IEEE, 2020, pp. 721–725.
- [16] Shanshan Wang, Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen, “A curated dataset of urban scenes for audio-visual scene analysis,” in *ICASSP*, 2021, pp. 626–630.
- [17] Nicolas Turpault, Romain Serizel, Ankit Parag Shah, and Justin Salamon, “Sound event detection in domestic environments with weakly labeled data and soundscape synthesis,” in *DCASE*, New York City, United States, October 2019.
- [18] Konstantinos Drossos, Samuel Lipping, and Tuomas Virtanen, “Clotho: An audio captioning dataset,” in *ICASSP*, May 2020.
- [19] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, and C. Lawrence Zitnick, “Microsoft coco captions: Data collection and evaluation server,” 2015.
- [20] Relja Arandjelovic and Andrew Zisserman, “Look, listen and learn,” in *ICCV*. IEEE, 2017, pp. 609–617.
- [21] Ho-Hsiang Wu, Prem Seetharaman, Kundan Kumar, and Juan Pablo Bello, “Wav2clip: Learning robust audio representations from clip,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 4563–4567.
- [22] Sen Fang, Yangjian Wu, Bowen Gao, Jingwen Cai, and Teik Toe Teoh, “Exploring efficient-tuned learning audio representation method from brivl,” 2023.
- [23] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel, “Audioclip: Extending clip to image, text and audio,” *arXiv preprint arXiv:2106.13043*, 2021.
- [24] Yuan Gong, Yu-An Chung, and James Glass, “Psla: Improving audio tagging with pretraining, sampling, labeling, and aggregation,” *arXiv preprint arXiv:2102.01243*, 2021.
- [25] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen, “Slow-fast auditory streams for audio recognition,” in *ICASSP*, 2021, pp. 855–859.
- [26] Dennis Fedorishin, Nishant Sankaran, Deen Mohan, Justas Birgiolas, Philip Schneider, Srirangaraj Setlur, and Venu Govindaraju, “Investigating waveform and spectrogram feature fusion for audio classification,” Tech. Rep., DCASE2021 Challenge, June 2021.
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 10684–10695.
- [28] Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun, “Attention bottlenecks for multimodal fusion,” in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, Eds. 2021, vol. 34, pp. 14200–14213, Curran Associates, Inc.
- [29] Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim, “Audiocaps: Generating captions for audios in the wild,” in *NAACL-HLT*, 2019.