

Out-Of-Distribution Detection for Audio-visual Generalized Zero-Shot Learning: A General Framework

Liuyuan Wen
lywen@mail.ustc.edu.cn

School of Physical Sciences
University of Science and Technology
of China
Hefei, China

Abstract

Generalized Zero-Shot Learning (GZSL) is a challenging task requiring accurate classification of both seen and unseen classes. Within this domain, Audio-visual GZSL emerges as an extremely exciting yet difficult task, given the inclusion of both visual and acoustic features as multi-modal inputs. Existing efforts in this field mostly utilize either embedding-based or generative-based methods. However, generative training is difficult and unstable, while embedding-based methods often encounter domain shift problem. Thus, we find it promising to integrate both methods into a unified framework to leverage their advantages while mitigating their respective disadvantages. Our study introduces a general framework employing out-of-distribution (OOD) detection, aiming to harness the strengths of both approaches. We first employ generative adversarial networks to synthesize unseen features, enabling the training of an OOD detector alongside classifiers for seen and unseen classes. This detector determines whether a test feature belongs to seen or unseen classes, followed by classification utilizing separate classifiers for each feature type. We test our framework on three popular audio-visual datasets and observe a significant improvement comparing to existing state-of-the-art works. Codes can be found in <https://github.com/liuyuan-wen/AV-OOD-GZSL>.

1 Introduction

The fusion of acoustic and visual cues in human communication and scene comprehension enjoys widespread recognition, demonstrating notable advancements in various applications such as action recognition [1, 2], emotion recognition [3, 4], speech recognition [5, 6] and so on. Despite these advancements, gathering sufficient annotated data for task-specific audio-visual representations remains a daunting challenge. Zero-shot learning (ZSL) arises as a feasible solution, enabling the classification of instances from novel classes by leveraging knowledge acquired from known classes. The more intricate Generalized ZSL (GZSL) further extends this capability to classify test samples from both seen and unseen classes. In this work, we tackle Audio-Visual Generalized Zero-Shot Learning (AV-GZSL) classification task.

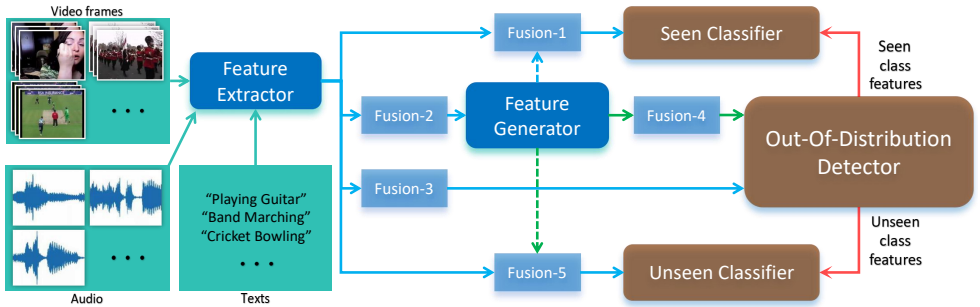


Figure 1: Our general framework. A generator is used to synthesize features, enabling the training of an OOD detector and separate classifiers for seen and unseen classes. OOD detector distinguishes seen and unseen features for separate classification. Different modality fusion methods (*Fusion 1-5*) can be implemented in different stages. (Blue and green arrows stands for training process using seen and synthesized unseen features respectively. Red arrows represents evaluation process. Dashed lines means “optional”).

Recently, several studies [17, 22] have delved into the field of AV-GZSL, innovating methods that project both modalities into a common embedding space and measure distances from class label text embeddings. Mercea *et al.* [19] introduced a integrated framework, employing a transformer-based cross-attention mechanism on temporally averaged audio and visual input features. Subsequently, numerous subsequent works have emerged, such as [10, 13, 14, 15, 18], including a partially generative approach [29] which synthesizes unseen features as negative samples in contrastive loss. However, all these approaches are essentially embedding-based and suffer significantly from the domain shift problem [23]. Although they employed calibrated stacking [6] to reduce the bias, it is merely a temporary solution with limited effectiveness.

From our perspective, on one hand, it is crucial to devise a comprehensive and promising strategy to address the domain shift problem. On the other hand, we also acknowledge the difficulty and instability of generative training, and the bias problem inherent in embedding methods. Therefore, we believe that integrating both approaches into a unified framework holds promise, allowing us to leverage the advantages of each method while mitigating their respective disadvantages.

In this study, we present a general framework named Audio-Visual Out-Of-Distribution (AV-OOD), as depicted in Figure 1. We first employ a generative model to synthesize unseen features, facilitating the training of an OOD detector alongside two separate classifiers for seen and unseen classes. Given the significance of modality fusion in multi-modal learning, different fusion methods (*Fusion 1-5*) can be implemented at various stages to fine-tune model design. During evaluation, the detector assesses whether a test feature belongs to seen or unseen classes, followed by classification using separate classifiers tailored to each feature type. In our experiment, we implement WGAN-GP [7, 8] as the feature generator, two Multilayer Perceptrons (MLPs) for both the OOD detector and the seen classifier. Lastly, an embedding-based model, adapted from [19], serves as the unseen classifier.

Furthermore, the generator can optionally synthesize novel seen or unseen features to aid in classifier training. However, in this work, we refrain from this approach due to our observation that the performance of our generator is not sufficiently satisfactory for synthe-

sizing discriminative trainable samples. Hence, we leave this option for future investigation. This decision also explains why we exclusively utilize the generator for training the OOD detector, as the precise manifold and distribution structure of unseen classes is less critical when only distinguishing between seen and unseen classes.

This framework is designed to be universal, as each component of the model operates relatively independently. Although we will introduce specific models for each part in Section 3, researchers can enhance overall performance by merely substituting any of them. For instance, models that excel in ZSL, such as spiking neural networks, can replace the unseen classifier, or a superior feature generative network can replace the WGAN-GP used in this work.

Contributions:

- (1) We are the first to introduce OOD detection to the field of AV-GZSL, alongside a general framework integrating both embedding-based and generative-based methods. This integration enables us to effectively harness the advantages of both approaches.
- (2) Each component of the framework operates relatively independently and each can be replaced with more effective alternatives, facilitating future researches in terms of structuration and overall performance.
- (3) We evaluate our framework on three widely-used audio-visual datasets and observe a significant improvement compared to existing state-of-the-art approaches.

2 Related Works

2.1 Audio-visual Generalized Zero-Shot Learning

In the realm of AV-GZSL, embedding-based models are proposed aiming to map video, audio, and text into a shared feature space for comparison. Initially proposed in [22], Coordinated Joint Multimodal Embedding (CJME) utilized triplet loss for proximity between modal features and class features. Audio-Visual Generalized Zero-shot Learning Network (AVGZSLNet) [17] introduced a module reconstructing text features from visual and audio inputs. Audio-Visual Cross-modal Attention (AVCA) [19] leveraged cross-modal attention mechanisms, effectively integrates information from both modalities. Temporal Cross-attention Framework (TCAF) [18] proposed a temporal cross-attention framework to enhance cross-attention across modalities and time. Audio-Visual Feature Synthesis (AVFS) [19] simulated unseen features with generative models, combining contrastive and discriminative losses. Variational Information Bottleneck for AV-GZSL (VIB-GZSL) [15] is a method based on variational information bottleneck. Audio-Visual Modality-fusion Spiking Transformer (AVMST) [13] and Motion-Decoupled Spiking Transformer (MDFT) [14] utilized spiking neural networks to process highly sparsity event data efficiently. Hypermultiple [11] used a hyperbolic transformation to achieve curvature-aware geometric learning.

All works above are essentially embedding-based methods including AVFS which only did a shallow addition and still suffered from bias problem. Our proposed framework is a fundamental solution for bias problem by implementing OOD detection, which integrates both embedding-based and generative-based methods.

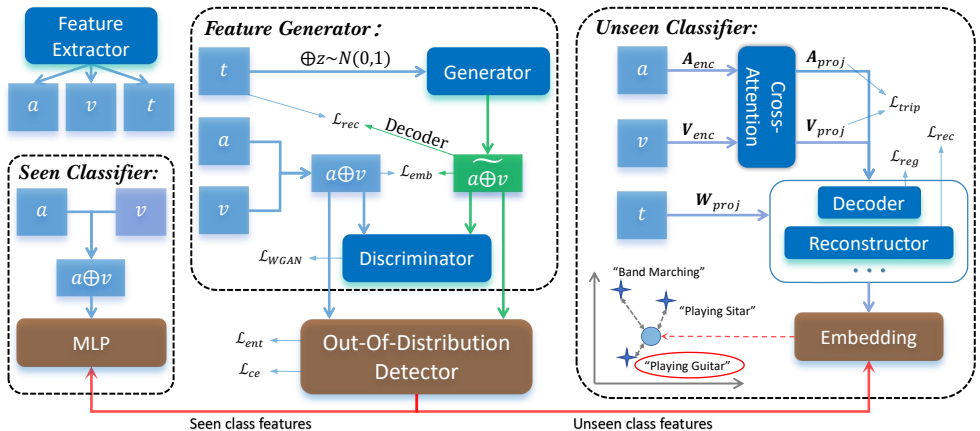


Figure 2: We employ MLPs as the seen classifier and OOD detector, an embedding-based model as the unseen classifier, and WGAN-GP for feature generation. Extracted features a , v , t stands for audio, visual, text respectively. Fusion methods involve simple concatenation \oplus except for *Fusion-5* in the unseen classifier, which utilizes cross-attention. The synthesized unseen features $\widetilde{a} \oplus v$ along with seen ones $a \oplus v$ are input for training of the OOD Detector, which distinguishes seen and unseen features for separate classification during test stage. (Blue and green arrows depict the training process using real seen and synthesized unseen features respectively. Red arrows represents evaluation process.)

2.2 Out-Of-Distribution Detection

Out-of-distribution (OOD) detection serves as a crucial method in GZSL, focusing on distinguishing between seen and unseen classes during inference to prevent the model from biased predictions towards the known categories. Socher *et al.* [25] pioneered the introduction of a binary novelty random variable for novelty detection. Atzmon *et al.* [9] proposed an adaptive confidence smoothing method, serving as a gating model for OOD detection. Subsequently, a variety of methods emerged, including entropy-based [16], probabilistic-based [27], and boundary-based [8], among others.

In this study, we primarily employ entropy-based methods [16] and integrate them into AV-GZSL alongside more expert classifiers. We evaluate two types of entropy losses and conduct a comparative analysis against the traditional calibrated stacking for the first time.

3 Model Architecture

In this section, we will introduce each component of the framework utilized in this study as Figure 2 shows. In general, we employ a conditional WGAN-GP [2, 8] as the feature generator, an MLP equipped with entropy-based loss as the OOD detector [16], another MLP serving as the seen classifier, and an embedding-based model adapted from AVCA [19] as the unseen classifier.

3.1 Feature Generator

We opt for a conditional WGAN-GP with an additional loss as our feature generator, inspired by its recent success in video classification [16]. On top of that, AV-GZSL requires an additional step of modality fusion, where we simply concatenate audio and visual features as $x = a \oplus v$. And the synthesized features are represented as $\tilde{x} = a \widetilde{\oplus} v$. The training process involves three distinct loss components: $\mathcal{L}_{\text{WGAN}}$, \mathcal{L}_{rec} and \mathcal{L}_{emb} .

Conditional WGAN-GP loss. A generator G and a discriminator D are employed in GAN. We denote the generated feature as $\tilde{x} = G(z, t)$ where $z \sim N(0, 1)$, the convex combination of x and \tilde{x} as \hat{x} , and the penalty coefficient as λ . Thus, conditioned on the textual embedding t , the expression is:

$$\mathcal{L}_{\text{WGAN}} = \mathbb{E}[D(x, t)] - \mathbb{E}[D(\tilde{x}, t)] - \lambda \mathbb{E} \left[(\|\nabla_{\hat{x}} D(\hat{x}, t)\|_2 - 1)^2 \right]. \quad (1)$$

Reconstruction loss. A decoder is utilized to reconstruct the class embedding t from the synthesized \tilde{x} for generated features to be discriminative. Let \mathcal{O}^{dec} represent the outputs of the decoder, and **MSE** represent the Mean Square Error metric, then $\mathcal{L}_{\text{rec}} = \mathbf{MSE}(\mathcal{O}^{\text{dec}}, t)$.

Embedding loss. We organize the real and synthesized features as matched (same classes) and unmatched (different classes) pairs. Then, a Cosine Embedding loss **CE** is employed to minimize and maximize the distance between the matched and unmatched features, respectively. Thus, $\mathcal{L}_{\text{emb}} = \mathbf{CE}(x, \tilde{x})$.

Thereby the total loss for feature generator is:

$$\mathcal{L}_{\text{gen}} = \min_G \max_D \mathcal{L}_{\text{WGAN}} + \alpha \mathcal{L}_{\text{rec}} + \beta \mathcal{L}_{\text{emb}}, \quad (2)$$

where α and β are weight coefficients.

3.2 Out-Of-Distribution Detector

Our OOD detector takes real seen and synthesized unseen features as positive and negative inputs, respectively. It outputs the probability of samples belonging to seen or unseen classes. In practice, we set a specific threshold, where outputs exceeding or falling below are categorized as seen or unseen accordingly.

We evaluate two types of OOD detectors, binary and entropy classifiers (See Section 4.4), both implemented as 3-layer MLPs denoted as O_{bin} and O_{ent} .

O_{bin} is a straightforward classifier using Binary Cross Entropy loss

$$\mathcal{L}_{\text{bin}} = \mathbf{BCE}(O_{\text{bin}}(x), 1) + \mathbf{BCE}(O_{\text{bin}}(\tilde{x}), 0). \quad (3)$$

However, O_{bin} proves inadequate due to the complex boundaries between seen and unseen classes (See Section 4.4.1). Therefore, we adopt an information entropy-based classifier with loss

$$\mathcal{L}_{\text{ent}} = \mathbf{H}(O_{\text{ent}}(x)) - \mathbf{H}(O_{\text{ent}}(\tilde{x})), \quad (4)$$

where $\mathbf{H}(p) = -\sum_i p_i \log(p_i)$ represents the entropy of p . Entropy effectively quantifies the uncertainty of features; the lower a feature’s entropy, the more likely it belongs to seen classes. Additionally, we include a Cross Entropy loss to accelerate convergence:

$$\mathcal{L}_{\text{ce}} = \mathbf{CE}(O_{\text{ent}}(x), y(x)), \quad (5)$$

where $y(x)$ is the corresponding label of sample x . Thus, the total loss for the OOD detector becomes $\mathcal{L}_{\text{ood}} = \mathcal{L}_{\text{ent}} + \mathcal{L}_{\text{ce}}$.

3.3 Seen Classifier

We adopt a 3-layer MLP for our seen classifier, finding it more stable and efficient compared to overly complex model structures (See Section 4.4.2). It takes $x = a \oplus v$ as input and employs Cross Entropy loss $\mathcal{L}_{sc} = \mathbf{CE}(x, y(x))$, $y(x)$ is the corresponding label of sample x .

3.4 Unseen Classifier

We enhance the AVCA [19] framework by introducing an additional negative loss component and take it as our unseen classifier. Instead of simple concatenation, we employ a transformer-based cross-attention for modality fusion.

As shown on the right side of Figure 2, in each training session, features a and v undergo a cross-attention block after being embedded into the same dimension as t by encoders \mathbf{A}_{enc} and \mathbf{V}_{enc} . This block facilitates the exchange of semantics among different modalities, capitalizing on shared cross-modal information. Subsequently, projectors \mathbf{V}_{proj} , \mathbf{A}_{proj} , and \mathbf{W}_{proj} project all modalities of samples into a common embedding space. Following this, the remaining procedures, including decoding, reconstruction, and backward process, are conducted for loss calculation. During evaluation, all three modalities are re-projected into the same joint embedding space as described above for classification. The closest textual label embedding of the current sample is then selected as the classified target.

To ensure effective feature representation learning, the loss function \mathcal{L}_{uc} consists of three components: triplet loss \mathcal{L}_{trip} , reconstruction loss \mathcal{L}_{rec} , and regularization loss \mathcal{L}_{reg} . While a contrastive learning strategy typically randomly selects a negative sample from a different class for every feature in the baseline, it’s important to note that previous studies often apply these losses solely for positive samples, neglecting losses from negative samples. To promote more robust training, we treat them equally and compute all three types of losses for both positive and negative samples.

Triplet loss. Four types of triplet loss functions cooperate to align contrastive audio, visual, and textual label embeddings:

$$\begin{aligned} \mathcal{L}_{trip}^{\pm} = & \mathbf{T}\left(\mathcal{O}_{av}^{proj,\pm}, \mathcal{O}_t^{proj,\pm}, \mathcal{O}_{av}^{proj,\mp}\right) + \mathbf{T}\left(\mathcal{O}_t^{proj,\pm}, \mathcal{O}_{av}^{proj,\pm}, \mathcal{O}_t^{proj,\mp}\right) \\ & + \mathbf{T}\left(\mathcal{O}_t^{proj,\pm}, \mathcal{O}_{av}^{proj,\pm}, \mathcal{O}_{av}^{proj,\mp}\right) + \mathbf{T}\left(\mathcal{O}_{av}^{proj,\pm}, \mathcal{O}_t^{proj,\pm}, \mathcal{O}_t^{proj,\mp}\right), \end{aligned} \quad (6)$$

where $\mathbf{T}(\cdot)$ represents the triplet loss function $\mathbf{T}(x, y, z) = \max\{\|x - y\|_2 - \|x - z\|_2 + \mu, 0\}$. Here, x , y , and z are anchor embeddings, and μ denotes the margin. $\mathcal{O}_m^{proj,\pm}$ represents the output from the projector in the joint space, where $m \in \{a, v, t\}$ denotes three modalities. Superscripts $\{+, -\}$ represent positive and negative samples, respectively. These four types of triplet loss functions cooperate to align contrastive audio, visual, and textual label embeddings.

Reconstruction loss. A decoder reconstructs and aligns the initial and final states of textual features. We define

$$\mathcal{L}_{rec}^{\pm} = \sum_{m \in \{a, v, t\}} \mathbf{MSE}\left(\mathcal{O}_m^{dec,\pm}, t^{\pm}\right), \quad (7)$$

where $\mathcal{O}_m^{dec,\pm}$ denotes the output tensor from decoder operators, considered as the final state, and t^{\pm} represents the input textual features, considered as the initial state.

Regularization loss. The regularization loss encourages the alignment of audio and visual embeddings with text embeddings while preserving information from their respective input

modalities. It is expressed as

$$\mathcal{L}_{reg}^{\pm} = \sum_{m \in \{a, v\}} \text{MSE}(\mathcal{O}_m^{rec, \pm}, \mathcal{O}_m^{enc, \pm}), \quad (8)$$

where $\mathcal{O}_m^{rec, \pm}$ and $\mathcal{O}_m^{enc, \pm}$ represent output samples from reconstructors and encoders, respectively.

Thus, the total loss of the unseen classifier is formulated as $\mathcal{L}_{uc} = \mathcal{L}^+ + \mathcal{L}^-$, where $\mathcal{L}^{\pm} = \mathcal{L}_{trip}^{\pm} + \mathcal{L}_{rec}^{\pm} + \mathcal{L}_{reg}^{\pm}$.

4 Experiments

4.1 Experimental Setup

We adopt AVCA [19] as our baseline, along with a commonly used benchmark proposed by them on datasets curated from VGGSound [6], UCF101 [26], and ActivityNet [9]. Notably, AVCA employs a two-stage training and evaluation protocol where epochs and calibrated stacking of the second stage are determined by the best performance on the validation set in the first stage. Since we have replaced calibrated stacking with OOD detection, we only run the second stage with all epochs fixed, which is equivalent.

For evaluation on the test dataset, we report mean class accuracy for both the Seen (S) and Unseen (U) subsets, with the Harmonic Mean $\mathbf{HM} = \frac{2US}{U+S}$. The Zero-Shot Learning (ZSL) performances are determined solely by evaluating on the test unseen subset.

4.2 Implementation Details

Feature Extractor: Following [19], we extract audio features a and visual features v , and average them second-wise using self-supervised SeLaVi [8] framework pretrained on VGGSound[6]. Textual label embeddings t are acquired through word2vec network pretrained on Wikipedia [20].

Feature Generator: We employ WGAN-GP [4, 5] with the loss function in Equation 2. In our experiments, we set α to 0.1 and β to 0.01. Training is conducted over 5 epochs with a learning rate of 0.0001, using the Adam optimizer.

Out-Of-Distribution Detector: Our detector is a 3-layer MLP with hidden layers of dimensions 512 and 128. For training, we generate 50, 50, and 1000 unseen samples from VGGSound, UCF, and ActivityNet, respectively. Training spans 80 epochs with learning rates and batch sizes of 0.001/6900 for VGGSound, 0.009/16 for UCF, and 0.005/64 for ActivityNet.

Seen Classifier: The seen classifier is also a 3-layer MLP, featuring hidden layers of dimensions 512 and 256. Training is conducted over 200 epochs with learning rates and batch sizes of 0.008/1024 for VGGSound, 0.0006/32 for UCF, and 0.008/128 for ActivityNet.

Unseen Classifier: We utilize triplet loss with a margin μ of 1. Dropout rates for the encoder/projector/decoder are set at 0.3/0.1/0.2 for VGGSound, 0.5/0.4/0.4 for UCF, and 0.2/0.3/0.2 for ActivityNet. Training spans 50 epochs with learning rates and batch sizes of 0.0005/256 for VGGSound, 0.0024/112 for UCF, and 0.0005/256 for ActivityNet.

Table 1: Comparison of AV-GZSL performances

Model	VGGSound-GZSL				UCF-GZSL				ActivityNet-GZSL			
	S	U	HM	ZSL	S	U	HM	ZSL	S	U	HM	ZSL
CJME[24]	8.69	4.78	6.17	5.16	26.04	8.21	12.48	8.29	5.55	4.75	5.12	5.84
AVGZSLNet[10]	18.15	3.48	5.83	5.28	52.52	10.90	18.05	13.65	8.93	5.04	6.44	5.40
AVCA[19]	14.90	4.00	6.31	6.00	51.53	18.43	27.15	20.01	24.86	8.02	12.13	9.13
TCAF[23]	9.64	5.91	7.33	6.06	58.60	21.74	31.72	24.81	18.70	7.50	10.71	7.91
VIB-GZSL[15]	<u>18.42</u>	6.00	9.05	6.41	90.35	21.41	<u>34.62</u>	22.49	22.12	8.94	12.73	9.29
AVFS[29]	15.20	5.13	7.67	6.20	54.87	16.49	25.36	22.37	<u>29.00</u>	9.13	<u>13.89</u>	11.18
AVMST[18]	14.14	5.28	7.68	6.61	44.08	22.63	29.91	28.19	17.75	<u>9.90</u>	12.71	10.37
MDFT[21]	16.14	5.97	8.72	7.13	48.79	<u>23.11</u>	31.36	31.53	18.32	10.55	13.39	12.55
Hyper-multiple[16]	15.02	<u>6.75</u>	<u>9.32</u>	<u>7.97</u>	63.08	19.10	29.32	22.24	23.38	8.67	12.65	9.50
AV-OOD (ours)	21.25	7.57	11.16	8.78	<u>65.78</u>	26.61	37.89	<u>28.21</u>	31.9	9.28	14.38	<u>11.49</u>

4.3 Results Analysis

Compared methods. We compare our AV-OOD with current state-of-the-art AV-GZSL models as introduced in Section 2.1, with methods range from embedding-based ones to generative-based ones. The results are presented in Table 1.

Comparison of performances. From Table 1, it is evident that AV-OOD stands out as the top performer across all three datasets. Specifically, AV-OOD achieves a remarkable **HM** score of 11.16% on the VGGSound dataset, marking a substantial 77% enhancement over the baseline AVCA score of 6.31%. Notably, AV-OOD also exhibits outstanding performance on UCF and ActivityNet, with **HM** scores of 37.89% and 14.38% respectively, surpassing all competing models. Moreover, the **S** scores of AV-OOD demonstrate notable stability, particularly evident on VGGSound and ActivityNet, recording 21.25% and 31.9% respectively. This commendable performance can be attributed not only to the implementation of a more efficient MLP for replacing the seen classifier but also to the replacement of calibrated stacking, which significantly curtails the model’s ability on seen classes. However, it is noteworthy that while AV-OOD excels in **HM** scores, its **ZSL** scores do not achieve the same level of prominence. This underscores the primary objective of our framework, which prioritizes enhancing the model’s discrimination between seen and unseen classes rather than focusing solely on **ZSL** skills. Nonetheless, despite not reaching the top position in **ZSL**, AV-OOD remains notable for its superiority over the baseline AVCA [[19](#)] across all metrics. This can be attributed to our incorporation of losses from negative samples, which evidently enhance the model’s robustness and learning efficiency.

4.4 Ablation Study

4.4.1 Effect of Different Bias Reduction Methods

As introduced in Section 1, domain shift problem arises due to discrepancies between the seen and unseen data distributions across different semantic domains. Previous works on AV-GZSL all adopted Calibrated Stacking [[5](#)] as the bias reduction method, which we replaced with OOD detection. We tested three bias reduction methods while preserving the performances of seen and unseen classifiers: *Calibrated Stacking*, *OOD-binary*, and *OOD-entropy*, to compare their capability of distinguishing between seen and unseen classes. Seen classes are regarded as positive and in-distribution, while unseen classes are considered neg-




Figure 3: Comparison between ROC curves of three bias reduction methods on all three datasets. Obviously, our proposed “OOD-entropy” is the most stable one and gets the best AUC on two of the three datasets.

Table 2: Comparison of different bias reduction methods

Method	VGGSound			UCF			ActivityNet		
	AUC	FPR	HM	AUC	FPR	HM	AUC	FPR	HM
Calibrated Stacking	66.04	36.43	9.32	60.06	39.78	27.57	69.58	30.82	12.04
OOD-binary	53.66	54.9	7.82	47.09	63.38	17.9	56.34	50.94	10.44
OOD-entropy (ours)	72.59	26.64	10.11	87.66	8.68	37.03	67.09	33.87	11.56

ative and out-of-distribution.

We plot Receiver Operating Characteristic (ROC) curves with True Positive Rate (**TPR**) and False Positive Rate (**FPR**) by sweeping over classification thresholds. Figure 3 displays all the ROC curves along with their Area Under Curve (**AUC**). **FPR** and **HM** in Table 2 correspond to the threshold that yields a 60% **TPR** for detecting in-distribution samples. We can observe that *OOD-entropy* is the most stable method and achieves the best AUC on two of the three datasets. Taking UCF as an example, *OOD-entropy* obtains an AUC of 87.66%, significantly surpassing 60.06% of *Calibrated Stacking* and 47.09% of *OOD-binary*. Meanwhile, *OOD-entropy* achieves the lowest **FPR** of 8.68%, compared to 39.78% of *Calibrated Stacking* and 63.38% of *OOD-binary*. This contributes to the highest **HM** of *OOD-entropy*, reaching 37.03%.

We observed that *Calibrated Stacking* and *OOD-binary* fall below the random guess line in certain areas. This is because UCF contains too few samples for the generator and classifiers to learn a comprehensive representation, which can be well supplemented using *OOD-entropy*. Moreover, *OOD-binary* performs the worst across all three datasets, confirming our judgment in Section 3.2 that a binary classifier is too simplified to learn the complex boundaries between seen and synthesized unseen samples whose precision is also limited by the generator.

Lastly, although *Calibrated Stacking* outperforms *OOD-entropy* on ActivityNet, its performance declines quickly when **TPR** (or **FPR**) is relatively low, as seen in Figure 3, where the best **HM** often appears. Thus, *OOD-entropy* still holds the first place with the final best **HM** of 14.38% (See Table 1).

Table 3: Comparison of different classifier models

Classifiers (Seen & Unseen)	VGGSound			UCF			ActivityNet		
	SC _{acc}	UC _{acc}	HM	SC _{acc}	UC _{acc}	HM	SC _{acc}	UC _{acc}	HM
MLP & MLP	58.27	3.39	5.51	96.12	17.7	26.86	62.33	3.71	5.81
Embedding & Embedding	49.42	8.78	10.97	64.97	28.21	30.26	33.99	11.49	12.14
MLP & Embedding (ours)	58.27	8.78	11.16	96.12	28.21	37.89	62.33	11.49	14.38

4.4.2 Effect of Different Classifier Models

We conducted tests using different choices of seen and unseen classifiers: *MLP* or *Embedding* methods, and the results are presented in Table 3. **SC_{acc}** and **UC_{acc}** represent the accuracy of the seen and unseen classifiers, respectively. **HM** is obtained with the assistance of OOD-entropy.

It can be concluded that *MLP* is the better choice for seen classes, while *Embedding* is better for unseen classes. For example, on VGGSound, *MLP* achieves a **SC_{acc}** of 58.27% compared to *Embedding*’s 49.42%. However, the **UC_{acc}** is only 3.39% compared to *Embedding*’s 8.78%. The same trend applies to the other two datasets.

This difference stems from two main aspects. Firstly, our embedding model is originally designed for ZSL and may be too complex for familiar class classification with the risk of overfitting, whereas a simple *MLP* can perform this task proficiently. Secondly, the performance of **UC_{acc}** of *MLP* is limited because, unlike *Embedding*, it can only learn from synthesized unseen features due to its fixed output dimension, without including the more crucial real seen features. However, there is considerable room for improvement in terms of both **SC_{acc}** and **UC_{acc}** across all datasets. This deficiency could potentially be addressed by employing more powerful specialized classification models in future works.

5 Conclusion

In this study, we introduce out-of-distribution (OOD) detection into the realm of Audio-Visual Generalized Zero-Shot Learning (AV-GZSL) for the first time. We establish a comprehensive framework that forthcoming researches can adopt to achieve more structural, organized and promising results. We first employ generative adversarial networks to synthesize unseen features, thereby enabling the training of an OOD detector alongside classifiers for both seen and unseen classes. This detector determines whether a test feature belongs to seen or unseen classes, followed by classification utilizing separate classifiers for each feature type.

Our work demonstrates superiority in two key aspects. Firstly, we integrate both embedding-based and generative-based methods, effectively leveraging the advantages of each approach. Secondly, each component of our framework operates relatively independently and can be substituted with more effective alternatives. This facilitates future exploration into structured models and enhances overall performances.

References

- [1] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman. Deep audio-visual speech recognition. *TPAMI*, 44(12):8717–8727, 2018.
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *ICML*, volume 70, pages 214–223, 2017.
- [3] Y. Asano, M. Patrick, C. Rupprecht, and A. Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. In *NeurIPS*, volume 33, pages 4660–4671, 2020.
- [4] Y. Atzmon and G. Chechik. Adaptive confidence smoothing for generalized zero-shot learning. In *CVPR*, pages 11671–11680, 2019.
- [5] W. L. Chao, S. Changpinyo, B. Gong, and F. Sha. An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *ECCV*, volume 14, pages 52–68. Springer International Publishing, 2016.
- [6] H. Chen, W. Xie, A. Vedaldi, and A. Zisserman. Vggsound: A large-scale audio-visual dataset. In *ICASSP*, pages 721–725, May 2020.
- [7] X. Chen, X. Lan, F. Sun, and N. Zheng. A boundary based out-of-distribution classifier for generalized zero-shot learning. In *ECCV*, pages 572–588, August 2020.
- [8] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. In *NeurIPS*, volume 30, 2017.
- [9] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015.
- [10] Jingwei Hong, Zulqarnain Hayder, Junting Han, Ping Fang, Mehrtash Harandi, , and Lars Petersson. Hyperbolic audio-visual zero-shot learning. In *ICCV*, pages 7873–7883, 2023.
- [11] A. Kerrigan, K. Duarte, Y. Rawat, and M. Shah. Reformulating zero-shot action recognition for multi-label actions. In *NeurIPS*, volume 34, pages 25566–25577, 2021.
- [12] J. T. Lee, M. Jain, H. Park, and S. Yun. Cross-attentional audio-visual fusion for weakly-supervised action localization. In *ICLR*, October 2020.
- [13] Wenrui Li, Zhengyu Ma, Liang-Jian Deng, Hengyu Man, and Xiaopeng Fan. Modality-fusion spiking transformer network for audio-visual zero-shot learning. In *ICME*, pages 426–431, 2023.
- [14] Wenrui Li, Xi-Le Zhao, Zhengyu Ma, Xingtao Wang, Xiaopeng Fan, , and Yonghong Tian. Motion-decoupled spiking transformer for audio-visual zero-shot learning. In *ACM MM*, pages 3994–4002, 2023.
- [15] Yapeng Li, Yong Luo, and Bo Du. Audio-visual generalized zero-shot learning based on variational information bottleneck. In *ICME*, pages 450–455, 2023.

- [16] D. Mandal, S. Narayan, S. K. Dwivedi, V. Gupta, S. Ahmed, F. S. Khan, and L. Shao. Out-of-distribution detection for generalized zero-shot action recognition. In *CVPR*, pages 9985–9993, 2019.
- [17] P. Mazumder, P. Singh, K. K. Parida, and V. P. Namboodiri. Avgzslnet: Audio-visual generalized zero-shot learning by reconstructing label features from multi-modal embeddings. In *WACV*, pages 3090–3099, 2021.
- [18] O. B. Mercea, T. Hummel, A. S. Koepke, and Z. Akata. Temporal and cross-modal attention for audio-visual zero-shot learning. In *ECCV*, pages 488–505, 2022.
- [19] O. B. Mercea, L. Riesch, A. Koepke, and Z. Akata. Audio-visual generalised zero-shot learning with cross-modal attention and language. In *CVPR*, pages 10553–10563, 2022.
- [20] T. Mikolov, E. Grave, P. Bojanowski, C. Puhrsch, and A. Joulin. Advances in pre-training distributed word representations. *arXiv preprint*, 2017.
- [21] Y. Mroueh, E. Marcheret, and V. Goel. Deep multimodal learning for audio-visual speech recognition. In *ICASSP*, pages 2130–2134, April 2015.
- [22] K. Parida, N. Matiyali, T. Guha, and G. Sharma. Coordinated joint multimodal embeddings for generalized audio-visual zero-shot classification and retrieval of videos. In *WACV*, pages 3251–3260, 2020.
- [23] F. Pourpanah, M. Abdar, Y. Luo, X. Zhou, R. Wang, C. P. Lim, and Q. J. Wu. A review of generalized zero-shot learning methods. *TPAMI*, 45(4):4051–4070, 2022.
- [24] L. Schoneveld, A. Othmani, and H. Abdelkawy. Leveraging recent advances in deep learning for audio-visual emotion recognition. *Pattern Recognition Letters*, 146:1–7, 2021.
- [25] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. In *NeurIPS*, volume 26, 2013.
- [26] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint*, 2012.
- [27] X. Wang, S. Pang, and J. Zhu. Domain segmentation and adjustment for generalized zero-shot learning. *arXiv preprint*, 2020.
- [28] S. Zhang, S. Zhang, T. Huang, and W. Gao. Multimodal deep convolutional neural network for audio-visual emotion recognition. In *ICMR*, pages 281–284, June 2016.
- [29] Q. Zheng, J. Hong, and M. Farazi. A generative approach to audio-visual generalized zero-shot learning: Combining contrastive and discriminative techniques. In *IJCNN*, pages 1–8, 2023.