

Audio-visual Generalized Zero-shot Learning the Easy Way

Shentong Mo¹ and Pedro Morgado²

¹ Carnegie Mellon University
² University of Wisconsin-Madison

Abstract. Audio-visual generalized zero-shot learning is a rapidly advancing domain that seeks to understand the intricate relations between audio and visual cues within videos. The overarching goal is to leverage insights from seen classes to identify instances from previously unseen ones. Prior approaches primarily utilized synchronized auto-encoders to reconstruct audio-visual attributes, which were informed by cross-attention transformers and projected text embeddings. However, these methods fell short of effectively capturing the intricate relationship between cross-modal features and class-label embeddings inherent in pre-trained language-aligned embeddings. To circumvent these bottlenecks, we introduce a simple yet effective framework for Easy Audio-Visual Generalized Zero-shot Learning, named *EZ-AVGZL*, that aligns audio-visual embeddings with transformed text representations. It utilizes a single supervised text audio-visual contrastive loss to learn an alignment between audio-visual and textual modalities, moving away from the conventional approach of reconstructing cross-modal features and text embeddings. Our key insight is that while class name embeddings are well aligned with language-based audio-visual features, they don't provide sufficient class separation to be useful for zero-shot learning. To address this, our method leverages differential optimization to transform class embeddings into a more discriminative space while preserving the semantic structure of language representations. We conduct extensive experiments on VGGSound-GZSL, UCF-GZSL, and ActivityNet-GZSL benchmarks. Our results demonstrate that our *EZ-AVGZL* achieves state-of-the-art performance in audio-visual generalized zero-shot learning.

Keywords: Audio-Visual Learning · Zero-shot Learning · Audio-Visual Generalized Zero-shot Learning

1 Introduction

In the field of machine learning, the ability of models to recognize patterns and make predictions is crucial. Zero-shot learning exemplifies this capability, which leverages known information to understand the unknowns. To further explore its potential and applications, the principle of zero-shot learning has been applied in various contexts. One such context is Audio-visual generalized

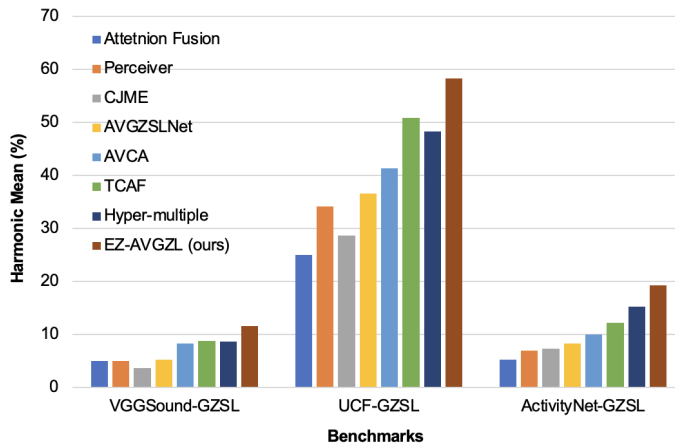


Fig. 1: Comparison of our *EZ-AVGZL* with state-of-the-art methods on VGGSound-GZSL, UCF-GZSL, ActivityNet-GZSL benchmarks in terms of Harmonic Mean (higher is better) for seen and unseen classes. Our method significantly outperforms previous baselines in terms of all datasets.

zero-shot learning (AVGZSL), an emerging topic of study that predicts unseen audio-visual categories by leveraging knowledge learned from seen classes.

Within the scope of AVGZSL, previous approaches [27, 28] have focused on designing autoencoders for reconstructing audio-visual attributes. For example, TCaF [27] adopted a temporal contrastive feature alignment strategy. This strategy aligns audio-visual cues with word2vec embeddings of class prompts by defining both cross-entropy and regression losses to reduce the distance between the output embedding for a sample and the corresponding projected word2vec embedding. These cross-modal and textual features were derived from cross-attention transformers and the projected text embeddings. However, the complexity of the network for AVGZSL, which is derived from the synergy between cross-attention transformers and the corresponding projected word2vec embeddings, makes it challenging to understand the deeper interplay between cross-modal features and class label embeddings.

In the context of audio-visual zero-shot learning, the main challenge lies in the absence of unseen classes during training. Without the supervision, the prediction of cross-modal class labels performs poorly on these audio-visual samples. To address the challenge, we propose a novel framework for Easy Audio-Visual Generalized Zero-shot Learning, referred to as *EZ-AVGZL*. This framework aligns audio-visual embeddings with transformed text representations in a simple yet effective manner. *EZ-AVGZL* learn well-separated text semantics across seen and unseen class names to achieve maximal separation while retaining the semantic relations between each other. During training, we leverage a simple supervised contrastive loss to learn the cross-modal alignment between audio-visual features

and transformed text embeddings. *EZ-AVGZL* enhances class separability while preserving semantics through the use of differential optimization, thereby generating more distinct text embeddings. Furthermore, it employs a single supervised text to audio-visual contrastive loss, establishing an alignment between the audio-visual and textual modalities. This approach effectively eliminates the need for reconstructing cross-modal features and text embeddings.

Empirical experiments on VGGSound-GZSL, UCF-GZSL, and ActivityNet-GZSL benchmarks demonstrate the state-of-the-art performance against previous audio-visual generalized zero-shot learning baselines, as illustrated in Fig. 1. Extensive ablation studies further confirm the importance of class embedding optimization and supervised audio-visual language contrastive objectives in learning language alignment with audio-visual representations. Additionally, we demonstrate the robustness of our approach across various zero-shot learning tasks and observe the benefit of our approach when applied to a variety of foundation models pre-trained from unimodal and cross-modal.

Our main contributions can be summarized as follows:

- We introduce the *EZ-AVGZL* framework, a paradigm shift in audio-visual generalized zero-shot learning. This framework simplifies the process with a supervised text-audio-visual contrastive loss, enhancing effectiveness.
- We propose a novel text transformation technique designed to construct distinct text embeddings, paving the way for better alignment and, by extension, superior recognition capabilities.
- We demonstrate that our method achieves state-of-the-art performance on extensive benchmarks in audio-visual generalized zero-shot learning.

2 Related Work

Audio-Visual Learning. Audio-visual learning problems have been addressed in many previous works [1, 2, 8, 10, 14, 17, 21, 32, 35, 36, 38, 41, 42, 44, 46, 47, 50, 60–62] to learn the audio-visual correlation between two distinct modalities from videos. Such cross-modal alignments are beneficial for a variety audio-visual tasks, including audio-event localization [22, 24, 30, 31, 33, 39, 40, 54, 58], audio-visual spatialization [4, 11, 42, 45], audio-visual continual learning [34, 49], audio-visual navigation [4–6] and audio-visual parsing [23, 37, 53, 57]. In this work, our primary focus is on learning audio-visual representations that align with text embeddings for generalized zero-shot learning, a task that presents more challenges than the aforementioned tasks.

Zero-shot Learning. Zero-shot learning aims to leverage semantic relationships, typically utilizing attributes or word embeddings associated with class labels to make informed predictions about previously unseen categories. This is often achieved by drawing parallels to the semantic information of known categories. Early audio-visual zero-shot learning methods, such as those proposed by Parida et al. [48] and Mazumder et al. [26], attempted to map video, audio, and text into the same feature space and compute the similarity across each other for comparison. For example, Audio-Visual Generalized Zero-shot Learning Network [26]

introduced a reconstruction module to predict text representations from audio and visual features. More recently, ImageBind [12] proposed learning a joint embedding across six different modalities, especially on audio, text, and video. This approach demonstrated the ability to recognize or retrieve unseen categories without explicit training. Unlike these methods, our approach does not require large-scale pre-training. Instead, we present a simple yet effective approach based on unimodal or crossmodal weights pre-trained on large-scale data to achieve audio-visual generalized zero-shot learning.

Audio-Visual Generalized Zero-shot Learning. Audio-visual generalized zero-shot learning extends the zero-shot learning paradigm to scenarios where both seen and unseen classes are present at test time. The seminal work, AVCA [28] proposed an attention mechanism to selectively focus on discriminative audio and visual cues. By harnessing the combined strength of these cues, AVCA not only achieved superior performance on seen classes but also exhibited proficiency in recognizing unseen classes.

Following up, TCaF [27] introduced a temporal contrastive feature alignment strategy. This strategy aligns audio-visual cues across time and applies two separate decoders to reconstruct the audio-visual embedding and the projected word2vec embedding to be close to the original word2vec embedding of class labels. More recently, Hong *et al.* [16] incorporated cross-modality alignment between video and audio features in the hyperbolic space. They used multiple adaptive curvatures for hyperbolic projections to achieve curvature-aware geometric learning. In contrast, we develop a novel framework to aggregate the alignment between audio-visual representations and text embeddings with explicit class embedding optimization. To the best of our knowledge, this work is the first to leverage an explicit class embedding optimization mechanism for audio-visual generalized zero-shot learning. Our experiments in Section 4.2 also demonstrate the effectiveness of our method in classifying both seen and unseen categories.

3 Method

We aim to learn a language-aligned representation for a video clip, which includes both the visual frame sequence and corresponding audio, to enable open-dictionary recognition of audio-visual events. To achieve this, we propose a simple yet effective approach, named Easy Audio-Visual Generalized Zero-shot Learning (*EZ-AVGZL*). This approach seeks to 1) construct multi-modal representations of the data and 2) align them with language representations using a supervised contrastive learning objective while leveraging strong foundation models. The proposed approach is based on two components: class embedding optimization described in Section 3.2 and supervised alignment between language and audio-visual representations described in Section 3.3. Despite its simplicity, *EZ-AVGZL* outperforms recent state-of-the-art methods on Audio-Visual Generalized Zero-shot Learning (AVGZSL) on a variety of datasets.

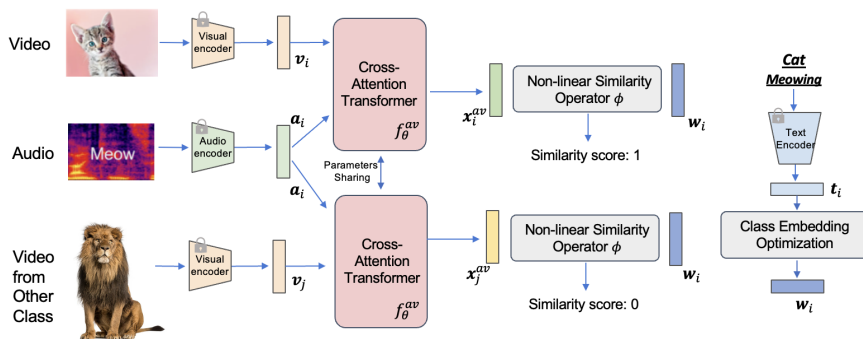


Fig. 2: Illustration of the proposed Easy Audio-Visual Generalized Zero-shot Learning (*EZ-AVGZL*) framework. The initial class embeddings t_i from a frozen text encoder are optimized with maximal separability and preserved semantics to generate the new embeddings w_i . Then the cross-attention transformer $f_{\theta}^{av}(\cdot, \cdot)$ takes visual and audio features (v_i, a_i) from the unimodal encoder to generate the multi-modal representations x_i^{av} . Finally, a non-linear similarity function is applied to align representations x_i^{av} with the corresponding class embeddings w_{y_i} by minimizing the cross-entropy loss between the predicted similarity score and the target score as 1 in the target class entry, while the target is 0 for multi-modal representations x_j^{av} given visual features v_j of videos from other classes.

3.1 Preliminaries: Audio-Visual Generalized Zero-shot Learning (AVGZSL)

Problem Statement. The objective of AVGZSL is to utilize natural language descriptions of audio-visual events t_i and the semantic relationships to enable a model to recognize events from a set of audio-visual event classes $\mathcal{T} = \mathcal{T}_s \cup \mathcal{T}_u$. This should be possible regardless of whether the model has seen them ($t_i \in \mathcal{T}_s$) or not seen them ($t_i \in \mathcal{T}_u$) during training.

To accomplish this, a model $f_{\theta} : \mathcal{V} \times \mathcal{A} \rightarrow \mathcal{T}$ is trained on the seen dataset $\mathcal{D}_s = \{(v_i, a_i, t_i)\}$ for audio-visual event recognition, where $v_i \in \mathcal{V}$ denotes the i^{th} visual clip, typically multiple frames spanning a short time window, $a_i \in \mathcal{A}$ denotes the corresponding audio, and $t_i \in \mathcal{T}_s$ denotes the (human-readable) class name of the audio-visual event depicted in video (v_i, a_i) . For effective zero-shot learning, a model needs to learn a mapping that can generalize beyond \mathcal{T}_s by inferring from the semantic relationships between seen and unseen classes. In other words, the AVGZSL setting evaluates models based on their classification performance on videos (v_i, a_i) belonging to both seen \mathcal{T}_s and unseen classes \mathcal{T}_u .

Architecture Following TCaF [27], we tackle AVGZSL using a variety of pre-trained models to encode all input signals. We experiment with a wider variety of pre-trained models than in [27], in all cases, we train our system on a series of unimodal pre-extracted feature representations for the visual frames $v_i = f_v(v_i)$, audio $a_i = f_a(a_i)$ and class names $t_i = f_t(t_i)$. As our goal is multi-modal GZSL,

we further encode the unimodal audio and visual representations into a joint multi-modal representation using the fusion model $\mathbf{x}_i^{av} = f_{av}(\mathbf{v}_i, \mathbf{a}_i)$ proposed in [27], which consists of a sequence of L transformer layers with cross-attention between audio and visual modalities.

The quality of representations obtained from foundation models is crucial for effective zero-shot learning. However, the focus of this work is the language alignment step. While prior work often relies on elaborate alignment strategies, we demonstrate that a simple supervised contrastive objective between text and audio-visual representations can achieve state-of-the-art performance. This is possible as long as the class name embeddings for seen and unseen events are maximally separated while retaining the semantic relations between each other. An overview of the proposed approach is shown in Fig. 2.

Apart from the quality of representations derived from foundation models, language alignment, which is the focus of this work, is a critical component for effective zero-shot learning. While prior work often relies on complex alignment strategies, we demonstrate that a *simple supervised contrastive objective* between text and audio-visual representations can achieve state-of-the-art performance, as long as the class name embeddings for seen and unseen events are *maximally separated* while retaining the *semantic relations* between each other. An overview of the proposed approach is shown in Fig. 2.

3.2 Class Embedding Optimization

Language-aligned classifiers can achieve zero-shot generalization as class embeddings $\mathbf{t}_1, \dots, \mathbf{t}_{|\mathcal{T}|}$ obtained from large-scale models like LLaMa or CLIP, encode semantic relationships between classes. However, while these class embeddings can effectively encode semantics, they are sub-optimal for classification due to a lack of class separability. To obtain well-separated semantic embeddings, we optimize a new set of class embeddings $\mathbf{w}_1^*, \dots, \mathbf{w}_{|\mathcal{T}|}^*$ to satisfy two competing constraints: class separability and semantic preservation.

Class Separability. Class embeddings that are maximally separated from each other are better suited for recognition of seen classes [43]. Thus, to achieve class embedding separation, we minimize the negative distance of each class embedding to its nearest neighbor

$$\mathcal{L}_{sep} = - \sum_{c=1}^{|\mathcal{T}|} \min_{k \neq c} \|\mathbf{w}_c - \mathbf{w}_k\|_2. \quad (1)$$

However, maximizing separability alone would destroy the semantics of the pre-trained embeddings. In fact, the minimum of Eq. 1 (if norm constrained) would be achieved for a set of embeddings $\mathbf{w}_1, \dots, \mathbf{w}_{|\mathcal{T}|}$ following a "Tammes problem" configuration [52]. Note that the embeddings are normalized. When the embedding dimension d is larger than the number of classes in the dictionary \mathcal{T} , the Tammes configuration provides a solution where all embeddings are equally spaced from each other, *i.e.* $\|\mathbf{w}_i - \mathbf{w}_j\|_2 = \text{Cte} \quad \forall i \neq j$. For example, the distance between the class "a person playing guitar" and the class "a person

playing the violin" would be the same as the distance to the class "a dog barking". Hence, since effective zero-shot recognition heavily relies on the semantic relations between seen and unseen classes, class separability must be balanced with the preservation of the semantics available in pre-trained embeddings.

Semantic preservation To preserve the semantics of the initial embeddings \mathbf{t}_i , a simple approach would be to use \mathbf{t}_i as proxies and constrain the new embeddings \mathbf{w}_i to be in their proximity

$$\mathcal{L}_{sem}^{prox} = \sum_{c=1}^{|\mathcal{T}|} \|\mathbf{w}_c - \mathbf{t}_c\|_2. \quad (2)$$

However, this constraint would significantly limit the separability of class embeddings \mathbf{w} , given their direct ties to \mathbf{t} . Instead, we seek new class embeddings \mathbf{w} that retain a similar ranking of similarities to other classes. This is accomplished using a margin ranking loss of the form

$$\mathcal{L}_{sem}^{rank} = \sum_{c=1}^{|\mathcal{T}|} \sum_{\substack{i=1 \\ c \neq i}}^{|\mathcal{T}|} \sum_{j=1}^{|\mathcal{T}|} \max \{0, m - \text{sign}(d_{ci}^{\mathbf{t}} - d_{cj}^{\mathbf{t}}) (d_{ci}^{\mathbf{w}} - d_{cj}^{\mathbf{w}})\} \quad (3)$$

where m is a margin coefficient, and $d_{ci}^{\mathbf{t}} = \|\mathbf{t}_c - \mathbf{t}_i\|_2$ and $d_{ci}^{\mathbf{w}} = \|\mathbf{w}_c - \mathbf{w}_i\|_2$ are the ℓ_2 distances between pairs of class embeddings, using either the pre-trained embeddings \mathbf{t} or the new optimized embeddings \mathbf{w} , respectively. Intuitively, this ranking loss will seek to organize the new embeddings \mathbf{w} such that if \mathbf{t}_c is closer to \mathbf{t}_i than to \mathbf{t}_j , then \mathbf{w}_c will also be closer to \mathbf{w}_i than to \mathbf{w}_j . Note that Eq. 3's complexity grows cubically with the number of classes. While manageable for the tested datasets, it can be impractical for large vocabularies. A potential solution is to sample a large (but fixed) number of class triplets per iteration, and replace Eq. 3 with an expectation over these triplets.

To balance separability and semantic preservation, we obtain the class embeddings by minimizing the joint loss

$$\arg \min_{\mathbf{w}_1, \dots, \mathbf{w}_{|\mathcal{T}|}} (1 - \alpha) \mathcal{L}_{sem} + \alpha \mathcal{L}_{sep}, \quad (4)$$

where α is set to 0.5. Given that the objective is differentiable, we optimize the new embeddings using gradient descent, starting from the initialization $\mathbf{w}_i = \mathbf{t}_i \forall i$.

3.3 Supervised Audio-Visual Language Alignment

Given the enhanced class embeddings $\mathbf{w}_1, \dots, \mathbf{w}_{|\mathcal{T}|}$ (with improved class separability), we empirically show that a contrastive objective is sufficient to learn state-of-the-art models for generalized zero-shot learning of audio-visual events. Specifically, let $\mathbf{x}_i^{av} = f_{\theta}^{av}(\mathbf{v}_i, \mathbf{a}_i)$ be the multi-modal representation of an audio-visual sample (v_i, a_i) belonging to class y_i . Then, we align representations \mathbf{x}_i^{av} with the corresponding class embeddings \mathbf{w}_{y_i} by minimizing the contrastive loss

$$\mathcal{L}_{AVLA} = -\log \frac{\exp(\text{sim}(\mathbf{x}_i^{av}, \mathbf{w}_{y_i}))}{\sum_{k=1}^B \exp(\text{sim}(\mathbf{x}_i^{av}, \mathbf{w}_{y_k}))} \quad (5)$$

where $\text{sim}(\cdot, \cdot)$ is a similarity function used to contrast the class embedding \mathbf{w}_{y_i} against others in the current batch of data \mathbf{w}_{y_k} .

Similarity. While contrastive objectives typically rely on cosine similarity metrics for $\text{sim}(\cdot, \cdot) : \mathcal{X} \times \mathcal{T} \rightarrow \mathfrak{R}$, we found that non-linear similarity functions can better capture the relationship between audio-visual representations and class embeddings. Specifically, we draw inspiration from transformer attention modules and use a multi-head cross-attention layer. In this layer, class embeddings \mathbf{w}_c are used as queries and audio-visual representations \mathbf{x}_i^{av} as key-value pairs. After cross-attention, class similarities are computed through a linear projection \mathbf{p} . Optimizing Eq. 5 pushes the learned embeddings closer to the text embeddings of the same category. During inference, we use the cross-attention transformer layer to output C similarity scores corresponding to each category, including unseen ones. We then classify the audio-visual samples according to the class with highest similarity.

4 Experiments

In this section, we evaluate the proposed *EZ-AVGZL* on three benchmarks for audio-visual generalized zero-shot learning, including VGGSound-GZSL [28], UCF-GZSL [28], and ActivityNet-GZSL [28].

4.1 Experimental Setup

Datasets. VGGSound-GZSL [28] contains 93,752 videos across 276 classes, selected from VGGSound [7], a large audio-visual dataset of 309 classes over 200k videos. We follow the prior work [27, 28] and use the same split for train/val(U)/test(U) over 138/69/69 classes. UCF-GZSL [28] includes 6,816 videos from 51 classes, chosen from UCF101 [51], the original video action recognition dataset, which consists of over 13k videos in 101 classes. We use the same train/val(U)/test(U) split over 30/12/9 classes in previous work [27, 28]. ActivityNet-GZSL [28] contains 20k videos in 200 classes of varying duration for action recognition [13]. We use the same train/val(U)/test(U) split over 99/51/50 classes as in the prior work [27, 28].

Evaluation Metrics. In line with prior work [27, 28], we use the mean class accuracy to evaluate all models. For ZSL performance, we evaluate only on the subset of test samples from unseen classes. For the GZSL performance, we evaluate on the full test set, which includes both seen and unseen classes. We report results for the subsets of seen (S) and unseen (U) classes, as well as their harmonic mean.

Implementation. We follow previous work [27] and use 8 heads with a dimension of 64 per head for all attention blocks in the cross-attention transformer. The model is trained for 50 epochs with a batch size of 64, using the Adam optimizer [20] with a learning rate of 1e-4, running average coefficients $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a weight decay of 1e-5. Following the prior work [27, 28], we use the word2vec [29] embeddings for text, and VGGish [15] embeddings for audio, and C3D [56] features for video frames.

Table 1: Quantitative results of VGGSound-GZSL benchmark.

Method	Seen	Unseen	Harmonic Mean	ZSL
AttentionFusion [9]	14.13	3.00	4.95	3.37
Perceiver [19]	13.25	3.03	4.93	3.44
CJME [48]	10.86	2.22	3.68	3.72
AVGZSLNet [26]	15.02	3.19	5.26	4.81
AVCA [28]	12.63	6.19	8.31	6.91
TCaF [27]	12.63	6.72	8.77	7.41
Hyper-alignment [16]	12.50	6.44	8.50	7.25
Hyper-single [16]	12.56	5.03	7.18	5.47
Hyper-multiple [16]	15.62	6.00	8.67	7.31
EZ-AVGZL (ours)	17.26	8.68	11.55	8.92

Table 2: Quantitative results of UCF-GZSL benchmark.

Method	Seen	Unseen	Harmonic Mean	ZSL
AttentionFusion [9]	39.34	18.29	24.97	20.21
Perceiver [19]	46.85	26.82	34.11	28.12
CJME [48]	33.89	24.82	28.65	29.01
AVGZSLNet [26]	74.79	24.15	36.51	31.51
AVCA [28]	63.15	30.72	41.34	37.72
TCaF [27]	67.14	40.83	50.78	44.64
Hyper-alignment [16]	57.13	33.86	42.52	39.80
Hyper-single [16]	63.47	34.85	44.99	39.86
Hyper-multiple [16]	74.26	35.79	48.30	52.11
EZ-AVGZL (ours)	78.32	46.35	58.24	48.27

4.2 Comparison to Prior Work

In this work, we propose a novel and effective framework for audio-visual generalized zero-shot learning. In order to demonstrate the effectiveness of the proposed *EZ-AVGZL*, we comprehensively compare it to previous audio-visual baselines [9, 19] adapted to our problem and current state-of-the-art audio-visual generalized zero-shot learning approaches [16, 26–28, 48].

For the VGGSound-GZSL dataset, we report the quantitative comparison results in Table 1. As shown, we achieve the best results in all metrics for both seen and unseen classes compared to previous audio-visual generalized zero-shot learning approaches. In particular, the proposed *EZ-AVGZL* outperforms Hyper-multiple [16], the current state-of-the-art audio-visual generalized zero-shot learning baseline, by 1.64@Seen, 2.68@Unseen, 2.88@Harmonic Mean, and 1.61@ZSL. Moreover, we also perform better than TCaF [27], a strong audio-visual generalized zero-shot learning baseline, by 4.63@Seen, 1.96@Unseen, 2.78@Harmonic Mean, and 1.51@ZSL. Furthermore, we achieve significant performance gains compared to AVCA [28], the first baseline using cross-modal attention and textual label embeddings. This indicates the importance of optimized text transformation and non-linear similarity functions in learning better transferrable knowledge for generalized zero-shot learning.

Table 3: Quantitative results of ActivityNet-GZSL benchmark.

Method	Seen	Unseen	Harmonic Mean	ZSL
AttentionFusion [9]	11.15	3.37	5.18	4.88
Perceiver [19]	18.25	4.27	6.92	4.47
CJME [48]	10.75	5.55	7.32	6.29
AVGZSLNet [26]	13.70	5.96	8.30	6.39
AVCA [28]	16.77	7.04	9.92	7.58
TCaF [27]	30.12	7.65	12.20	7.96
Hyper-alignment [16]	29.77	8.77	13.55	9.13
Hyper-single [16]	24.61	10.10	14.32	10.37
Hyper-multiple [16]	36.98	9.60	15.25	10.39
EZ-AVGZL (ours)	38.52	12.85	19.27	13.01

Table 4: Ablation studies on Class Embedding Optimization and Audio-Visual-Language Alignment.

CEP	AVLA	Seen	Unseen	Harmonic Mean	ZSL
\times	\times	12.63	6.72	8.77	7.41
\checkmark	\times	15.35	7.92	10.45	8.36
\times	\checkmark	14.67	7.23	9.69	7.91
\checkmark	\checkmark	17.26	8.68	11.55	8.92

Table 5: Ablation studies on optimization objectives in Class Embedding Optimization.

\mathcal{L}_{sem}	\mathcal{L}_{sep}	Seen	Unseen	Harmonic Mean	ZSL
\times	\times	14.67	7.23	9.32	7.91
\mathcal{L}_{sem}^{prox}	\times	15.75	7.83	10.46	8.26
\times	\checkmark	16.13	8.17	10.85	8.45
\mathcal{L}_{sem}^{prox}	\checkmark	17.26	8.68	11.55	8.92
\mathcal{L}_{sem}^{rank}	\checkmark	18.05	10.28	13.10	10.36

Additionally, we observe significant gains in the UCF-GZSL and ActivityNet-GZSL benchmarks, as shown in Table 2 and Table 3. Compared to TCaF [27], a strong baseline using three manually-designed objectives, we achieve result gains of 11.18@Seen, 5.52@Unseen, 7.46@Harmonic Mean, and 3.63@ZSL on UCF-GZSL dataset. When evaluated on the challenging ActivityNet-GZSL benchmark, our method still outperforms TCaF [27] by 8.40@Seen, 5.20@Unseen, 7.07@Harmonic Mean, and 5.05@ZSL. We also achieve better results than AVCA [28], a cross-modal baseline based on cross-modal attention between audio and visual embeddings. These results demonstrate the effectiveness of our approach in learning optimized textual label embeddings and non-linear similarity functions for generalized audio-visual zero-shot recognition.

4.3 Experimental Analysis

In this section, we conduct ablation studies to demonstrate the benefit of introducing the Class Embedding Optimization (CEO) and Audio-Visual Language Alignment (AVLA) modules. We also carry out extensive experiments to investigate optimization objectives in class embedding optimization, non-linear similarity functions, and generalization across diverse cross-modal settings.

Class Embedding Optimization (CEO) & Audio-Visual Language Alignment (AVLA). In order to demonstrate the effectiveness of CEO and AVLA, we conduct an ablation study and report the quantitative results on VGGSound-GZSL dataset in Table 4.

As observed, adding CEO to the vanilla baseline improves the results by 2.72@Seen, 1.20@Unseen, 1.68@Harmonic Mean, and 0.95@ZSL, which validates the benefit of optimizing class embeddings so as to obtain separated embeddings that preserve the semantic relationships between classes. On the other hand, introducing AVLA in the baseline increases the zero-shot performance across all metrics. More importantly, incorporating both CEP and AVLA into the baseline significantly raises the performance by 4.63@Seen, 1.96@Unseen, 2.78@Harmonic Mean, and 1.51@ZSL. These improvements validate the importance of text transformation and supervised text audio-visual contrastive in extracting separated textual label embeddings for AVGZSL.

Objectives in Class Embedding Optimization. Optimization objectives in class embedding transformation, including global semantic preservation and class separability, are essential for learning optimized textual class embeddings for zero-shot learning. To explore the effects of each objective more comprehensively, we ablate each text optimization objective and report the quantitative results in Table 5. We observe that adding the semantic preservation loss to the vanilla baseline improves the performance by 1.08@Seen, 0.60@Unseen, 1.14@Harmonic Mean, and 0.35@ZSL. On the other hand, introducing the class separability loss in the baseline enhances the zero-shot performance across all metrics. Furthermore, combining the semantic preservation and class separability objectives for the text transformer significantly increases the performance by 2.59@Seen, 3.05@Unseen, 3.78@Harmonic Mean, and 2.45@ZSL. This result validates the importance of class separability loss in balancing the preservation of the semantics available in pre-trained representations for zero-shot learning. Replacing the Euclidean-based semantic preservation loss \mathcal{L}_{sem}^{prox} with a ranking based loss \mathcal{L}_{sem}^{rank} further achieves performance gains of 0.79@Seen, 1.60@Unseen, 1.55@Harmonic Mean, and 1.44@ZSL. These results demonstrate the effectiveness of both semantic preservation and class separability objectives in class embedding optimization to generate meaningful text representations.

Non-linear Similarity Functions. The non-linear similarity operator used in supervised audio-visual language alignment impacts the extracted cross-modal representations. To explore these effects more comprehensively, we varied the functions from {Cosine-similarity, Linear, MLP, Cross-Attention}. For Cosine-Similarity, we directly compute the dot product between audio-visual embeddings \mathbf{x}_i^{av} and text embeddings \mathbf{t}_i . For Linear, we apply a linear layer to the concatenation of both embeddings to predict the similarity score. For MLP, we replace the linear layer with a multi-layer perceptron to generate the score. For Cross-Attention, we use text embeddings \mathbf{t}_i as the queries and audio-visual representations \mathbf{x}_i^{av} as key-value pairs in a transformer attention module to output the score. The comparison results of zero-shot performance on VGGSound-GZSL benchmark are reported in Table 6. When using Cosine-Similarity as the function, we achieve the worst results across all metrics. Replacing the cosine-similarity operator with both Linear increases the seen results by 2.55@Seen, but decreases the performance by 0.07@Unseen and 0.12@ZSL. This might be because adding complexity to the function causes the model to overfit the seen classes in the

Table 6: Exploration studies on non-linear similarity functions $\phi(\cdot)$.

$\phi(\cdot)$	Seen	Unseen	Harmonic Mean	ZSL
Cosine-Similarity	13.97	8.59	10.64	8.67
Linear	16.52	8.52	11.24	8.55
MLP	17.03	8.46	11.30	8.49
Cross-Attention	17.26	8.68	11.55	8.92

Table 7: Exploration studies on diverse cross-modal settings. CEP denotes Class Embedding Optimization.

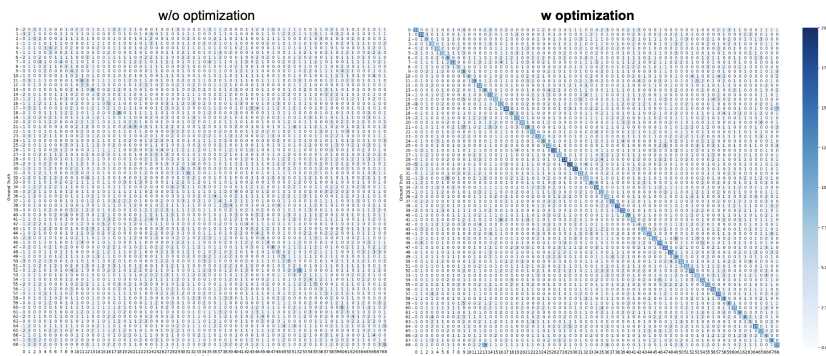
Text Encoder	Audio Encoder	Visual Encoder	CEP	Seen	Unseen	Harmonic Mean	ZSL
word2vec	VGGish	C3D	✗	12.63	6.72	8.77	7.41
			✓	17.26	8.68	11.55	8.92
LLaMa	AudioMAE	DINO	✗	18.35	10.63	13.46	12.26
			✓	23.95	13.32	17.12	13.73
CLAP	CLAP	X-CLIP	✗	27.79	14.01	18.63	17.39
			✓	33.28	20.15	25.10	21.06
ImageBind	ImageBind	ImageBind	✗	28.32	14.68	19.34	19.87
			✓	35.56	21.63	26.90	22.35

training data, leading to poorer generalization to unseen categories. Additionally, using MLP as the function operator significantly raises the performance by 3.06@Seen, but slightly deteriorates the results on unseen classes. Our method achieves the best results across all metrics by using the cross-attention operator to predict the similarity score. These improved results demonstrate the effectiveness of the cross-attention operator in balancing the trade-off between linear and non-linear similarity functions.

Generalization to Diverse Cross-modal Settings. To validate the generalizability of the proposed method across a flexible number of feature encoders, we experiment with different settings: unimodal, paired (video-text, audio-text), and unified embedding space (image-audio-text). In the unimodal setting, we use DINO [3] for the visual encoder, AudioMAE [18] for the audio encoder, and LLaMa [55] for the text encoder. In the paired setting, we use X-CLIP [25] pre-trained with video-text objectives for the visual encoder, and CLAP [59] pre-trained with audio-text objectives for the audio and text encoder, respectively. In the unified embedding setting, we apply ImageBind [12] pre-trained with image-audio-text alignment in the same embedding space as the encoder for all three modalities. The quantitative results on the VGGSound-GZSL benchmark are compared in Table 7. As observed, the unimodal setting, which uses transformers with strong unimodal features, significantly outperforms the vanilla baseline that uses convolution networks, where *EZ-AVGZL* achieves performance gains of 6.69@Seen, 4.64@Unseen, 5.77@Harmonic Mean, and 4.81@ZSL using class embedding optimization. In the paired setting, the proposed method further increases the results by 9.33@Seen, 6.93@Unseen, 7.98@Harmonic Mean, and 7.331@ZSL, benefiting from the text-audio correspondence and text-video alignment pre-trained on large-scale data. In the unified embedding setting using ImageBind [12] embeddings, we achieve the best results across all metrics

Table 8: Comparison of class embeddings w/o and w optimization.

base class	NN (w/o opt)	D (w/o opt)	NN (w opt)	D (w opt)
alligators, crocodiles hissing	church bell ringing	0.455	playing marimba, xylophone	0.498
barn swallow calling	black capped chickadee calling	0.307	black capped chickadee calling	0.465
basketball bounce	playing volleyball	0.324	bouncing on trampoline	0.394
bee, wasp, etc. buzzing	motorboat, speedboat acceleration	0.322	fly, housefly buzzing	0.333
bird squawking	pheasant crowing	0.400	duck quacking	0.496
black capped chickadee calling	barn swallow calling	0.307	magpie calling	0.436
bouncing on trampoline	basketball bounce	0.416	basketball bounce	0.494
bowling impact	playing table tennis	0.410	striking bowling	0.467
bull bellowing	eagle screaming	0.449	dog growling	0.532
canary calling	whale calling	0.403	black capped chickadee calling	0.513

**Fig. 3:** Confusion matrices of zero-shot predictions on VGGSound-GZSL using model with and without class embedding optimization.

compared to other settings. This might be because text embeddings with unified training on the other two modalities strengthen the alignment for audio-visual zero-shot learning. Furthermore, when class embedding optimization is added to all settings, we improve the baseline based on the initial text embeddings, further demonstrating the effectiveness of the proposed class embedding optimization in generating text representations for AVZSL.

Qualitative Comparisons with and without Optimization. To provide insights into why the proposed loss terms improve the separability of class embeddings while keeping semantics, we conduct qualitative studies. These include comparing the visualization of the class embedding space with and without optimization. The comparisons of class embeddings with and without optimization are shown in Table 8. We report the base class and identify 1) the nearest neighbor (NN) class and 2) the distance (D) to the nearest neighbor. We observe larger distances and similar NN classes. The confusion matrices for class-wise performances of the model are reported in Fig. 3.

Impact of Hyper-parameters α and Margin m . To investigate the impact of α and m on the performance thoroughly, we set α from $\{0.1, 0.5, 0.9\}$ and m from $\{0, 0.5, 1\}$. The quantitative comparisons on VGGSound-GZSL are reported in Table 9. As can be seen, with the increase of α from 0.1 to 0.5, the proposed approach achieves better results regarding all metrics. However, when α is increased to 0.9, the performance decreases a lot, which might be caused by not

Table 9: Study of α and margin m on VGGSound-GZSL.

α	m	Seen	Unseen	Harmonic Mean	ZSL
0.5		18.05	10.28	13.10	10.36
0.1	1	16.35	8.39	11.09	8.58
0.9		16.26	8.73	11.36	8.92
0.5	0	17.82	9.71	12.57	9.67
	0.5	17.67	9.85	12.65	9.95

Table 10: Study of distance functions on VGGSound-GZSL.

Distance	Seen	Unseen	Harmonic Mean	ZSL
cosine	18.05	10.28	13.10	10.36
ℓ_2	17.58	9.32	12.18	9.27
ℓ_1	17.35	8.78	11.66	8.85

much semantics preservation during training. Meanwhile, with the decrease of m , we observe a clearly dropping trend in our metrics.

Ablating Distance Functions for Eq. 1-2. Distance functions in Eq. 1-2 is crucial for us to maximize separation while preserving semantics. We experimented with different distance functions, ℓ_2 and ℓ_1 distances for Eq. 1-2. Note that while, Eq. 1-2 show ℓ_2 distances in the paper, the embeddings are normalized. The quantitative results are shown in Table 10. We can observe that using the cosine distance function performed the best regarding all metrics. However, when optimizing ℓ_1 distance functions, we achieve the worst results in terms of all metrics. These ablation results highlight the importance of designing appropriate distance functions in maximizing separation while preserving the semantics of text class embeddings.

5 Conclusion

In this work, we present a novel and effective framework, *EZ-AVGZL*, that aligns audio-visual representations with optimized text embeddings to achieve audio-visual generalized zero-shot learning. We leverage differential optimization to learn better-separated textual representations with maximal separability and semantic preservation. Additionally, we introduce supervised audio-visual language alignment to learn the correspondence between audio-visual features and textual embeddings, which captures the cross-modal dynamics with class labels as prompts. Experimental results on VGGSound-GZSL, UCF-GZSL, and ActivityNet-GZSL datasets demonstrate our method’s superiority against previous baselines. Extensive ablation studies also validate the importance of class embedding optimization and supervised audio-visual language alignment for audio-visual generalized zero-shot learning. We also conduct comprehensive experiments to demonstrate the effectiveness of non-linear similarity functions and generalization to diverse cross-modal settings.

Broader Impact. This work makes substantial progress in audio-visual generalized zero-shot learning. Our method is novel in its simplicity, demonstrating that a simple supervised contrastive objective, without "bells and whistles", outperforms complex alignment strategies. Without our work, progress in this area would likely be driven by increasingly complex methods, which are difficult to implement, reproduce and compare. Beyond the strong results and simplicity of our method, the class embedding optimization procedure is also technically novel to maximize semantic separation while keeping semantics preserved.

References

1. Arandjelovic, R., Zisserman, A.: Look, listen and learn. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 609–617 (2017) [3](#)
2. Aytar, Y., Vondrick, C., Torralba, A.: Soundnet: Learning sound representations from unlabeled video. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2016) [3](#)
3. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021) [12](#)
4. Chen, C., Jain, U., Schissler, C., Garí, S.V.A., Al-Halah, Z., Ithapu, V.K., Robinson, P., Grauman, K.: Soundspaces: Audio-visual navigation in 3d environments. In: Proceedings of European Conference on Computer Vision (ECCV). pp. 17–36 (2020) [3](#)
5. Chen, C., Majumder, S., Ziad, A.H., Gao, R., Kumar Ramakrishnan, S., Grauman, K.: Learning to set waypoints for audio-visual navigation. In: Proceedings of International Conference on Learning Representations (ICLR) (2021) [3](#)
6. Chen, C., Schissler, C., Garg, S., Kobernik, P., Clegg, A., Calamia, P., Batra, D., Robinson, P.W., Grauman, K.: Soundspaces 2.0: A simulation platform for visual-acoustic learning. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track (2022) [3](#)
7. Chen, H., Xie, W., Vedaldi, A., Zisserman, A.: Vggsound: A large-scale audio-visual dataset. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 721–725. IEEE (2020) [8](#)
8. Ephrat, A., Mosseri, I., Lang, O., Dekel, T., Wilson, K., Hassidim, A., Freeman, W.T., Rubinstein, M.: Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. arXiv preprint arXiv:1804.03619 (2018) [3](#)
9. Fayek, H.M., Kumar, A.: Large scale audiovisual learning of sounds with weakly labeled data. In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (2021) [9](#), [10](#)
10. Gan, C., Huang, D., Zhao, H., Tenenbaum, J.B., Torralba, A.: Music gesture for visual sound separation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10478–10487 (2020) [3](#)
11. Gao, R., Grauman, K.: 2.5d visual sound. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 324–333 (2019) [3](#)
12. Girdhar, R., El-Nouby, A., Liu, Z., Singh, M., Alwala, K.V., Joulin, A., Misra, I.: Imagebind: One embedding space to bind them all. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023) [4](#), [12](#)
13. Heilbron, F.C., Escorcia, V., Ghanem, B., Niebles, J.C.: Activitynet: A large-scale video benchmark for human activity understanding. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 961–970 (2015) [8](#)
14. Hershey, J., Casey, M.: Audio-visual sound separation via hidden markov models. Advances in Neural Information Processing Systems **14** (2001) [3](#)
15. Hershey, S., Chaudhuri, S., Ellis, D.P.W., Gemmeke, J.F., Jansen, A., Moore, R.C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., Slaney, M., Weiss, R.J., Wilson, K.: Cnn architectures for large-scale audio classification. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (2017) [8](#)

16. Hong, J., Hayder, Z., Han, J., Fang, P., Harandi, M., Petersson, L.: Hyperbolic audio-visual zero-shot learning. In: Proceedings of the International Conference on Computer Vision (ICCV) (2023) [4](#), [9](#), [10](#)
17. Hu, D., Nie, F., Li, X.: Deep multimodal clustering for unsupervised audiovisual learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9248–9257 (2019) [3](#)
18. Huang, P.Y., Xu, H., Li, J., Baevski, A., Auli, M., Galuba, W., Metze, F., Feichtenhofer, C.: Masked autoencoders that listen. In: Proceedings of Advances In Neural Information Processing Systems (NeurIPS) (2022) [12](#)
19. Jaegle, A., Gimeno, F., Brock, A., Zisserman, A., Vinyals, O., Carreira, J.: Perceiver: General perception with iterative attention. In: Proceedings of the International Conference on Machine Learning (ICML) (2021) [9](#), [10](#)
20. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [8](#)
21. Korbar, B., Tran, D., Torresani, L.: Cooperative learning of audio and video models from self-supervised synchronization. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2018) [3](#)
22. Lin, Y.B., Li, Y.J., Wang, Y.C.F.: Dual-modality seq2seq network for audio-visual event localization. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 2002–2006 (2019) [3](#)
23. Lin, Y.B., Tseng, H.Y., Lee, H.Y., Lin, Y.Y., Yang, M.H.: Exploring cross-video and cross-modality signals for weakly-supervised audio-visual video parsing. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2021) [3](#)
24. Lin, Y.B., Wang, Y.C.F.: Audiovisual transformer with instance attention for audio-visual event localization. In: Proceedings of the Asian Conference on Computer Vision (ACCV) (2020) [3](#)
25. Ma, Y., Xu, G., Sun, X., Yan, M., Zhang, J., Ji, R.: X-CLIP:: End-to-end multi-grained contrastive learning for video-text retrieval. arXiv preprint arXiv:2207.07285 (2022) [12](#)
26. Mazumder, P., Sing, P., Kumar Parida, K., Namboodiri, V.P.: Avgzslnet: Audio-visual generalized zero-shot learning by reconstructing label features from multi-modal embeddings. In: Proceedings of 2021 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 3089–3098 (2021) [3](#), [9](#), [10](#)
27. Mercea, O.B., Hummel, T., Koepke, A.S., Akata, Z.: Temporal and cross-modal attention for audio-visual zero-shot learning. In: Proceedings of European Conference on Computer Vision (ECCV) (2022) [2](#), [4](#), [5](#), [6](#), [8](#), [9](#), [10](#)
28. Mercea, O.B., Riesch, L., Koepke, A.S., Akata, Z.: Audio-visual generalised zero-shot learning with cross-modal attention and language. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10553–10563 (2022) [2](#), [4](#), [8](#), [9](#), [10](#)
29. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. In: Proceedings of International Conference on Learning Representations (ICLR) (2013) [8](#)
30. Mo, S., Morgado, P.: Benchmarking weakly-supervised audio-visual sound localization. In: European Conference on Computer Vision (ECCV) Workshop (2022) [3](#)
31. Mo, S., Morgado, P.: A closer look at weakly-supervised audio-visual source localization. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2022) [3](#)

32. Mo, S., Morgado, P.: A unified audio-visual learning framework for localization, separation, and recognition. In: Proceedings of the International Conference on Machine Learning (ICML) (2023) [3](#)
33. Mo, S., Morgado, P.: Unveiling the power of audio-visual early fusion transformers with dense interactions through masked modeling. arXiv preprint arXiv:2312.01017 (2023) [3](#)
34. Mo, S., Pian, W., Tian, Y.: Class-incremental grouping network for continual audio-visual learning. arXiv preprint arXiv:2309.05281 (2023) [3](#)
35. Mo, S., Shi, J., Tian, Y.: DiffAVA: Personalized text-to-audio generation with visual alignment. arXiv preprint arXiv:2305.12903 (2023) [3](#)
36. Mo, S., Shi, J., Tian, Y.: Text-to-audio generation synchronized with videos. arXiv preprint arXiv:2403.07938 (2024) [3](#)
37. Mo, S., Tian, Y.: Multi-modal grouping network for weakly-supervised audio-visual video parsing. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2022) [3](#)
38. Mo, S., Tian, Y.: Semantic-aware multi-modal grouping for weakly-supervised audio-visual video parsing. In: European Conference on Computer Vision (ECCV) Workshop (2022) [3](#)
39. Mo, S., Tian, Y.: Audio-visual grouping network for sound localization from mixtures. arXiv preprint arXiv:2303.17056 (2023) [3](#)
40. Mo, S., Tian, Y.: AV-SAM: Segment anything model meets audio-visual localization and segmentation. arXiv preprint arXiv:2305.01836 (2023) [3](#)
41. Mo, S., Tian, Y.: Semantic grouping network for audio source separation. arXiv preprint arXiv:2407.03736 (2024) [3](#)
42. Morgado, P., Li, Y., Vasconcelos, N.: Learning representations from audio-visual spatial alignment. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS). pp. 4733–4744 (2020) [3](#)
43. Morgado, P., Li, Y., Costa Pereira, J., Saberian, M., Vasconcelos, N.: Deep hashing with hash-consistent large margin proxy embeddings. International Journal of Computer Vision (2020) [6](#)
44. Morgado, P., Misra, I., Vasconcelos, N.: Robust audio-visual instance discrimination. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12934–12945 (2021) [3](#)
45. Morgado, P., Vasconcelos, N., Langlois, T., Wang, O.: Self-supervised generation of spatial audio for 360 video. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2018) [3](#)
46. Morgado, P., Vasconcelos, N., Misra, I.: Audio-visual instance discrimination with cross-modal agreement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12475–12486 (June 2021) [3](#)
47. Owens, A., Wu, J., McDermott, J.H., Freeman, W.T., Torralba, A.: Ambient sound provides supervision for visual learning. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 801–816 (2016) [3](#)
48. Parida, K.K., Matiyali, N., Guha, T., Sharma, G.: Coordinated joint multimodal embeddings for generalized audio-visual zero-shot classification and retrieval of videos. In: Proceedings of 2020 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 3240–3249 (2020) [3](#), [9](#), [10](#)
49. Pian, W., Mo, S., Guo, Y., Tian, Y.: Audio-visual class-incremental learning. arXiv preprint arXiv:2308.11073 (2023) [3](#)
50. Senocak, A., Oh, T.H., Kim, J., Yang, M.H., Kweon, I.S.: Learning to localize sound source in visual scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4358–4366 (2018) [3](#)

51. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012) [8](#)
52. Tammes, P.M.L.: On the origin of number and arrangement of the places of exit on the surface of pollen-grains. *Recueil des travaux botaniques néerlandais* **27**(1), 1–84 (1930) [6](#)
53. Tian, Y., Li, D., Xu, C.: Unified multisensory perception: Weakly-supervised audio-visual video parsing. In: *Proceedings of European Conference on Computer Vision (ECCV)*. p. 436–454 (2020) [3](#)
54. Tian, Y., Shi, J., Li, B., Duan, Z., Xu, C.: Audio-visual event localization in unconstrained videos. In: *Proceedings of European Conference on Computer Vision (ECCV)* (2018) [3](#)
55. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C.C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P.S., Lachaux, M.A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E.M., Subramanian, R., Tan, X.E., Tang, B., Taylor, R., Williams, A., Kuan, J.X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., Scialom, T.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023) [12](#)
56. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (2015) [8](#)
57. Wu, Y., Yang, Y.: Exploring heterogeneous clues for weakly-supervised audio-visual video parsing. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1326–1335 (2021) [3](#)
58. Wu, Y., Zhu, L., Yan, Y., Yang, Y.: Dual attention matching for audio-visual event localization. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 6291–6299 (2019) [3](#)
59. Wu, Y., Chen, K., Zhang, T., Hui, Y., Berg-Kirkpatrick, T., Dubnov, S.: Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In: *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP* (2023) [12](#)
60. Zhang, L., Mo, S., Zhang, Y., Morgado, P.: Audio-synchronized visual animation. arXiv preprint arXiv:2403.05659 (2024) [3](#)
61. Zhao, H., Gan, C., Ma, W.C., Torralba, A.: The sound of motions. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 1735–1744 (2019) [3](#)
62. Zhao, H., Gan, C., Rouditchenko, A., Vondrick, C., McDermott, J., Torralba, A.: The sound of pixels. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 570–586 (2018) [3](#)