



# Enhancing Zero-shot Audio Classification using Sound Attribute Knowledge from Large Language Models

Xuenan Xu<sup>\*1</sup>, Pingyue Zhang<sup>\*1</sup>, Ming Yan<sup>2</sup>, Ji Zhang<sup>2</sup>, Mengyue Wu<sup>†1</sup>

<sup>1</sup>MoE Key Lab of Artificial Intelligence X-LANCE Lab, Shanghai Jiao Tong University, China

<sup>2</sup>Institute of Intelligent Computing, Alibaba Group, China

{wsntxxn, williamzhangsjtu, mengyuewu}@sjtu.edu.cn, {ym119608, zj122146}@alibaba-inc.com

## Abstract

Zero-shot audio classification aims to recognize and classify a sound class that the model has never seen during training. This paper presents a novel approach for zero-shot audio classification using automatically generated sound attribute descriptions. We propose a list of sound attributes and leverage large language model’s domain knowledge to generate detailed attribute descriptions for each class. In contrast to previous works that primarily relied on class labels or simple descriptions, our method focuses on multi-dimensional innate auditory attributes, capturing different characteristics of sound classes. Additionally, we incorporate a contrastive learning approach to enhance zero-shot learning from textual labels. We validate the effectiveness of our method on VGGSound and AudioSet<sup>1</sup>. Our results demonstrate a substantial improvement in zero-shot classification accuracy. Ablation results show robust performance enhancement, regardless of the model architecture.

**Index Terms:** zero-shot learning, audio classification, sound attribute, large language model, audio-text contrastive learning

## 1. Introduction

Supervised learning has shown promise in many fields however it is limited to pre-defined classes included during training. In audio classification, current datasets often only contain commonly-heard daily sounds [1] while rare sounds from specific domains are seldom included due to requirement of labor-intensive annotations. Hereby, such models can only predict the probability of pre-defined classes but the probability of unseen classes cannot be given, impeding its real-world applications.

To overcome the reliance on exhausting annotation of training data, zero-shot classification has drawn much attention recently [2, 3]. Most works learn the correspondence between audio features and semantic features of classes [4, 5]. Xie *et al.* used a bilinear model to calculate the audio-class similarity [6, 7]. The textual label is the class and a definition from wikipedia, e.g., “Sounds associated with the species of medium-to-large birds, *Corvus*” for the class `Crow`. Sims *et al.* added synonyms, semantic broadening and onomatopoeia as the auxiliary information [3]. However, labels and auxiliary information are often insufficient to discriminate sounds, in particular similar ones. Since the definition is about the sounding object instead of auditory attributes, they cannot reflect the difference between close-related classes, e.g., `train` and `underground`. Therefore, these approaches have not extensively exploited perceptual auditory knowledge to describe sound classes.

<sup>1</sup>The code is available at <https://www.github.com/wsntxxn/AttrEnhZsAc>.

<sup>\*</sup>Equal contribution.

<sup>†</sup>Corresponding author.

Auditory attributes, more representative of sounds’ inherent characteristics, can be a leading flag for sound descriptions. Some studies explored using sound attributes as either inputs or training objectives to facilitate learning: Lin *et al.* used attributes like pitch and material to characterize collision sounds [8] while Choi *et al.* took the occurrence of instruments as attributes to classify music.

Recently, audio foundation models from contrastive language-audio pre-training (CLAP) [9, 10] have drastically improved the generalizability of audio models by using natural language supervision. However, the success in effective zero-shot classification may be due to indirect exposure to testing sounds during training, rather than true zero-shot learning. For instance, CLAP’s ability to recognize crow cawing could be attributed to its training on paired data of crow cawing sounds and the text “A crow is cawing”, though the sound is not explicitly labelled as a pre-defined class. Thus, CLAP’s capability of classifying novel sound categories without prior exposure may be questioned. Moreover, CLAP predominantly focuses on scaling-up audio-text data while the text data are mostly limited to event labels. These textual descriptions may not adequately capture the distinct characteristics inherent in auditory signals, such as pitch, tone, and timbre. Therefore, we adopt the zero-shot setting that test audio is unseen during training and focus on improving textual descriptions in this work.

In this paper, we integrate the advantages of semantic supervision and attribute supervision to improve zero-shot audio classification by learning the correspondence between audio and its description that focuses on sound attributes. Compared with previous works using sound attributes, we define a set of attributes applicable to general sound categories. Then we use ChatGPT to automatically generate attribute-focused descriptions for all categories. The capabilities of large language models (LLM) are exploited to provide accurate attribute descriptions without human annotation. Moreover, inspired by the success of self-supervised learning [11], we adapt a more advanced contrastive learning paradigm to zero-shot classification, thereby achieving improved alignment between audio and textual labels. Note that our approach is knowledge-driven, hence can be integrated to foundation models like CLAP by adding the sound attributes and improving textual labels during training. Experimental results demonstrate that our approach achieves significant improvement in zero-shot classification accuracy, regardless of the backbone model architecture. In summary, our contributions are three-fold:

- We define an inherent auditory attribute set for general sounds and incorporate ChatGPT [12] to automatically generate attribute-oriented descriptions, which can assist any audio model architecture training.
- We use an improved contrastive learning paradigm to learn

<b>Attributes</b>	Remarks: Take “Smoke detector, smoke alarm” as an example
<b>Sound class</b>	The name of the sound class is one of the attributes: <i>Smoke detector, smoke alarm.</i>
<b>Sound frequency (pitch)</b>	The frequency of the sound: <i>frequency is sharp, piercing, and high.</i>
<b>Timbre</b>	The timbre of the sound, that is, the unique quality or character of a sound: <i>timbre is shrill and alarming.</i>
<b>Onomatopoeia</b>	Words to imitate or resemble the sound: <i>sounds like beep, beep-beep, and loud screech</i>
<b>Simile</b>	Comparison to another common sounds: <i>like a persistent and urgent warning signal.</i>
<b>Emotion</b>	The emotions evoked by the sound: <i>evoking alertness and urgency.</i>
<b>Definition</b>	Only available for AudioSet, a wikipedia description of the sound: <i>Sounds emitted by a device that senses smoke, typically to warn of a fire.</i>

Table 1: Seven attributes defined to describe sound, using smoke detector as the example. To differentiate the “description” provided by AudioSet from our proposed attribute-oriented description, we use “definition” to denote the former.

better audio-text correspondence.

- Results on two datasets demonstrate that by leveraging the audio-related domain-specific knowledge in LLM mined from massive text data, the zero-shot classification performance is significantly improved.

## 2. Description of Sound Attributes

To describe the characteristics of a sound class, it is necessary to define a set of sound attributes. Although [13] defined several binary attributes such as duration and material, these attributes only cover a small portion of sound features and are somewhat only suitable for the involved collision sound dataset. Therefore, we try to determine general sound attributes. We ask ChatGPT which attributes can be used to describe a sound category. In addition to the summarized attributes, we manually add three attributes: 1) onomatopoeia: inspired by [3], we include onomatopoeia to supplement the description of sounds that are difficult to describe using attributes such as pitch and intensity; 2) timbre: a property that distinguishes different sounds, especially those with similar pitches; 3) simile: similar sounds. In zero-shot classification, the model can be extensively trained on seen categories. Therefore, when calculating the probability of an unseen class, information about what it is similar to among the seen classes is valuable.

Based on these attributes, we first ask ChatGPT<sup>2</sup> to generate descriptions for several sound classes. Subsequently, we take them as few-shot examples to guide ChatGPT to generate descriptions for the remaining classes. Such an in-context learning approach makes full use of the extensive acoustic knowledge learned by the large language model. Finally, descriptions in terms of 6 attributes are generated for each sound class, as shown in Table 1. For AudioSet, we also use its provided wikipedia description as an additional attribute.

## 3. Zero-Shot Classification by Audio-Text Contrastive Learning

### 3.1. Contrastive Learning Framework

Our zero-shot classification approach is similar to previous studies [3, 6], with a few modifications. The overall framework is shown in Figure 1. For a pair of an audio clip  $\mathcal{A}$  and a class label  $c$ , sound attributes  $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n$  of  $c$  are annotated by ChatGPT. An audio encoder transforms the input audio  $\mathcal{A}$  into an embedding vector  $\mathbf{a}$ :

$$\mathbf{a} = \text{Enc}_A(\mathcal{A}) \quad (1)$$

<sup>2</sup>We use *gpt-3.5-turbo*.

During training, the descriptive text for an audio sample is formed by sampling from  $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n$  and concatenation. Different sampling strategies are investigated and the comparison will be shown later in Section 5.2. The description is encoded by a text encoder accordingly:

$$\mathbf{t} = \text{Enc}_T(\text{Concat}(\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_s)) \quad (2)$$

where  $s$  is the number of sampled attributes.

The model is trained to minimize the distance between  $\mathbf{a}$  and its corresponding class’s  $\mathbf{t}$ , while repelling  $\mathbf{a}$  and  $\mathbf{t}$  of other classes, using a similarity function  $\text{sim}(\cdot, \cdot)$ . During inference, all attributes of the class  $c$  are utilized to calculate the similarity:

$$\mathbf{t}^c = \text{Enc}_T(\text{Concat}(\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n)) \quad (3)$$

$$y = \underset{c \in C}{\text{argmax}} \text{sim}(\mathbf{a}, \mathbf{t}^c) \quad (4)$$

where  $C$  is the set of zero-shot classes for testing which does not intersect with the training class set. Compared with the framework in [7, 8], our approach make two modifications:

**Similarity Function** The baseline in [7] calculates audio-class similarity by a bilinear model:

$$s = (W' \mathbf{a})' \mathbf{t} \quad (5)$$

We replace it with the cosine similarity between embeddings after fully-connected (FC) projection layers:

$$\mathbf{a}_P = \text{FC}_A(\mathbf{a}) \quad \mathbf{t}_P = \text{FC}_T(\mathbf{t})$$

$$s = \frac{\mathbf{a}_P \cdot \mathbf{t}_P}{\|\mathbf{a}_P\| \cdot \|\mathbf{t}_P\|} \quad (6)$$

**Loss Function** In [7] the model is trained by a weighted max margin ranking loss [14]. In light of the success of InfoNCE loss [15] in supervised contrastive learning (SupCon) [11], we adopt it in our improved learning framework. For a batch of  $B$  samples, the loss is calculated as:

$$\mathcal{L} = \sum_i \frac{-1}{N_{l_i}} \sum_{j, l_j=l_i} \log \frac{\exp(\text{sim}(\mathbf{a}_i, \mathbf{t}_j)/\tau)}{\sum_k \exp(\text{sim}(\mathbf{a}_i, \mathbf{t}_k)/\tau)} \quad (7)$$

where  $i, j, k \in \{1 \dots B\}$ ,  $l_i$  is the label of the audio sample  $i$  and  $\tau$  is the temperature.  $N_{l_i}$  is the number of samples in the batch that share the same label with  $i$ . Descriptions from these samples are treated as positive ones while other samples’ descriptions are negative. The InfoNCE loss trains the model to discriminate positive descriptions from negative ones.




Figure 1: The contrastive learning framework for zero-shot audio classification. The text is the description of sound attributes for the class generated by ChatGPT. During training, attributes are randomly selected to form the description while all attributes are used during inference.

To sum up, our primary approach involves two methods: **Baseline** and **SupCon**. In **Baseline**, we utilize bilinear similarity and employ the weighted max-margin ranking loss. This method is the same as [7], using all attributes as the textual label. **SupCon** combines cosine similarity with the InfoNCE loss as done in [11]. In this method, we randomly select attributes during training.

### 3.2. Model Architecture

To encode informative attribute descriptions into discriminative embeddings, we employ SentenceBERT [16] as the text encoder. It is capable of encoding sentences into discriminative embeddings, thus well-suited for distinguishing descriptions of different sound classes. We explore two variants of SentenceBERT: a standard `mpnet-base` with 12 layers and a hidden size of 768, and a smaller `MiniLM` with 6 layers and a hidden size of 384. Different text encoders are evaluated to investigate the correlation between the performance of our approach and the encoding ability of the text encoder. We validate the effectiveness of our method on 2 popular audio encoder backbones: a convolutional PANNs `CNN14` [17] and a Transformer-based `AST tiny224` [18]. Since the size of audio classification datasets is insufficient to fine-tune the pre-trained text encoder, we keep the text encoder frozen while the audio encoder is trained from scratch.

## 4. Experimental Setup

**Dataset** Although previous works on zero-shot audio classification mostly used ESC50 [19], we find results on ESC50 unstable due to limited data size. Therefore, we use two large datasets, VGGSound [20] and AudioSet [1]. VGGSound contains over 200k clips covering 309 classes. For AudioSet, we select only single-label audio samples from unbalanced set and exclude 21 abstract classes like “Animal” and “Music”. We also exclude classes with less than 100 samples. Following [7], we randomly select 1,500 samples for classes with more than 1,500 samples for a balanced distribution. Finally, 115k samples and 285 classes are selected for AudioSet.

**Fold Split** To ensure fully zero-shot setting, we split a dataset into 5 disjoint folds according to classes. The model is then trained on 4 of them and evaluated on the remaining one. For a balanced data distribution among folds, we sort classes in descending order based on their sample numbers. Starting with 5 empty folds, we select the next 5 unallocated classes in the sorted order, shuffle them, and allocate them to the 5 folds respectively. The splitting process yields approximately 35k samples in each fold for VGGSound and 23k samples for AudioSet.

**Hyper-parameters** Following [7], we apply a two-stage training paradigm: 1) training the audio encoder from scratch; 2) performing **Baseline** or **SupCon** audio-text alignment. The model is trained with a batch size of 64 using SGD optimizer with an initial learning rate of 0.1 and the cosine learning rate scheduler in the first stage. Note that in the first stage, the model is also trained on 4 folds and has no access to the leave out fold. In the second stage, we use Adam optimizer with a learning rate of  $1 \times 10^{-4}$  and a batch size of 256.

## 5. Results and Analysis

### 5.1. Comparison with Baselines

Results on VGGSound and AudioSet are presented in Table 2. To achieve fair comparison between **Baseline** and **SupCon**, we report results of **SupCon** with all attributes used during training, which is consistent with **Baseline**. We provide accuracy averaged over 5 random seeds. The standard deviation is omitted since it is minor. Results of t-tests demonstrate that in all cases where the proposed method outperforms the baseline, the improvement is statistically significant with p-values less than 0.05. Note that “definition” in Table 1 is unavailable for VGGSound. Therefore, in terms of the textual label, we solely use class as the baseline for VGGSound, while class and definition for AudioSet. Results demonstrate the efficacy of using descriptions focusing on sound attributes we define for supervision. Whatever the textual labels and dataset used, **SupCon** approach achieves consistent improvement over **Baseline**. The comparison validates the effectiveness of our improved contrastive learning paradigm.

### 5.2. Ablation Study

**Attribute Sampling Strategy** During training, attributes can be sampled with different strategies to form the description text. We explore three strategies: *Deterministic*, *Random*, and *With-class*.

- *Deterministic*: all attributes are selected.
- *Random*: attributes are randomly selected.
- *With-class*: attributes are also randomly selected but we ensure the class name is always sampled.

By default we use the *Deterministic* strategy. Results in Table 3 showcase the superiority of the *With-class* strategy. This indicates that class label is a vital and clean attribute hence it should be always included. Other attributes are generated automatically with inevitable noise. Therefore, randomness in attributes other than class should be incorporated into the training process to enhance the model’s robustness against noisy attributes.

Table 2: Zero-shot classification performance using CNN14 backbone. “Definition” here refers to wikipedia description only available for AudioSet. All attributes are used during training for both approaches.

		VGGSound					AudioSet				
Approach	Text	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5
Baseline [6]	Class (& Definition)	20.5	19.9	23.7	20.8	21.4	23.7	17.7	24.2	26.1	23.6
	Attributes	24.1	22.6	27.1	<b>25.6</b>	22.5	25.2	23.5	26.6	28.9	26.8
SupCon	Class (& Definition)	23.2	21.7	23.5	21.6	24.2	25.8	21.6	27.0	27.9	27.8
	Attributes	<b>27.0</b>	<b>25.3</b>	<b>27.3</b>	25.3	<b>26.6</b>	<b>26.1</b>	<b>26.0</b>	<b>29.0</b>	<b>31.8</b>	<b>32.5</b>

Table 3: VGGSound zero-shot classification performance averaged on 5 folds using different attribute sampling strategies.

Strategy	Deterministic	Random	With-class
Accuracy	26.3	25.3	<b>26.6</b>

**Model Architecture** To investigate the performance of our method under different model architectures, we change the audio / text encoder backbone. Here we use the *With-class* strategy during training. Results are presented in Table 4. Although AST performs worse than CNN14, the supervision of attribute descriptions does bring improvement over the textual label of class, indicating that our method is robust to audio backbones. In addition, by replacing the MiniLM text encoder with mpnet which performs better in text matching, further performance boost is achieved. Our approach is capable of leveraging advanced text encoders to achieve better zero-shot classification performance.

Table 4: VGGSound zero-shot classification performance averaged on 5 folds using different audio / text backbones.

Text	CNN14-MiniLM	AST-MiniLM	CNN14-mpnet
Class	22.8	17.9	24.1
Attributes	<b>26.6</b>	<b>21.6</b>	<b>28.1</b>

### 5.3. Analysis on Attributes




Figure 2: VGGSound zero-shot classification performance with each attribute included during training.

To analyze the improvement brought by attribute descriptions, we include each attribute during training to explore its influence. The result is shown in Figure 2. Whatever attribute is included, the performance is consistently enhanced. Com-

pared with low-level auditory attributes like pitch and timbre, high-level semantic attributes such as onomatopoeia and simile bring more significant improvement. This indicates that high-level semantic attributes generated by LLMs are more accurate and representative for audio classification.

Then, some cases are selected to show the effectiveness of adding specific attributes, shown in Figure 3. Numbers on the green and red bars show numbers of correctly and mistakenly classified samples, respectively. These classes are easily misclassified for they share similar acoustic characteristics or their labels are semantically close. Additional attributes describe unique acoustic characteristics to distinguish similar classes. For example, onomatopoeia is utilized to imitate the sound of dog growling and barking to highlight the difference.




Figure 3: Performance enhancement brought by adding attributes. Green and red denote ground truth and misclassified classes. The corresponding attribute description is also presented. Ono.= onomatopoeia, simi.=simile, emo.=emotion.

## 6. Conclusion

In this paper, we propose a zero-shot audio classification approach using sound attribute descriptions. The approach combines the advantages of semantic supervision and attribute supervision by making use of rich knowledge from ChatGPT. We define a set of sound attributes and use ChatGPT to automatically describe a sound category in terms of each attribute. Moreover, we propose an improved contrastive learning paradigm to augment the capability of our model to learn effectively from textual labels. The zero-shot accuracy on VGGSound and AudioSet is significantly boosted, indicating the effectiveness of our approach. The ablation study on backbones shows consistent improvement over the baseline, validating that our method is agnostic to model architectures. The limitation is that the attributes described by ChatGPT are occasionally misaligned with their respective attribute names. Additionally, the exploration of utilizing different LLMs for generation is left for future investigation.

## 7. Acknowledgements

This work was supported by National Natural Science Foundation of China (Grant No.92048205), the Key Research and Development Program of Jiangsu Province (No.BE2022059), Guangxi major science and technology project (No. AA23062062) and Alibaba Innovative Research.

## 8. References

- [1] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [2] P. Primus and G. Widmer, "Improved zero-shot audio tagging & classification with patchout spectrogram transformers," in *European Signal Processing Conference (EUSIPCO)*. IEEE, 2022, pp. 410–413.
- [3] Y. Sims, A. Mendes, and S. Chalup, "Enhanced embeddings in zero-shot learning for environmental audio," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [4] G. K. Gudur and S. K. Perepu, "Zero-shot federated learning with new classes for audio classification," *arXiv preprint arXiv:2106.10019*, 2021.
- [5] J. Choi and J. Nam, "Zero-shot learning for audio-based music classification and tagging," in *Proceedings of International Society for Music Information Retrieval Conference*, 2019, pp. 67–74.
- [6] H. Xie, O. Räsänen, and T. Virtanen, "Zero-shot audio classification with factored linear and nonlinear acoustic-semantic projections," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 326–330.
- [7] H. Xie and T. Virtanen, "Zero-shot audio classification via semantic embeddings," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 29, pp. 1233–1242, 2021.
- [8] Y.-H. Lin, X. Chen, R. Takashima, and T. Takiguchi, "Zero-shot sound event classification using a sound attribute vector with global and local feature learning," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [9] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, "Clap learning audio concepts from natural language supervision," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [10] Y. Wu, K. Chen, T. Zhang, Y. Hui, T. Berg-Kirkpatrick, and S. Dubnov, "Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [11] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," *Advances in neural information processing systems*, vol. 33, pp. 18 661–18 673, 2020.
- [12] OpenAI, "Introducing chatgpt," <https://openai.com/blog/chatgpt>, 2022.
- [13] Y. Lin, X. Chen, R. Takashima, and T. Takiguchi, "Binary attribute embeddings for zero-shot sound event classification," in *Global Conference on Consumer Electronics (GCCE)*, 2022, pp. 13–14.
- [14] J. Weston, S. Bengio, and N. Usunier, "Wsabie: scaling up to large vocabulary image annotation," in *Proceedings of International Joint Conference on Artificial Intelligence*, 2011, pp. 2764–2770.
- [15] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [16] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," in *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 2019.
- [17] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, "Panns: Large-scale pretrained audio neural networks for audio pattern recognition," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [18] Y. Gong, Y.-A. Chung, and J. Glass, "AST: Audio Spectrogram Transformer," in *Proceedings of Conference of the International Speech Communication Association*, 2021, pp. 571–575.
- [19] K. J. Piczak, "Esc: Dataset for environmental sound classification," in *Proceedings of ACM International Conference on Multimedia*, 2015, pp. 1015–1018.
- [20] H. Chen, W. Xie, A. Vedaldi, and A. Zisserman, "Vggsound: A large-scale audio-visual dataset," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 721–725.