

DiveSound: LLM-Assisted Automatic Taxonomy Construction for Diverse Audio Generation

Baihan Li¹, Zeyu Xie¹, Xuenan Xu¹, Yiwei Guo¹, Ming Yan², Ji Zhang², Kai Yu^{*1}, Mengyue Wu^{*1}

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

²Institute of Intelligent Computing, Alibaba Group, China
{lbh0612, zeyu_xie, wsntxxn, mengyuewu}@sjtu.edu.cn

Abstract

Audio generation has attracted significant attention. Despite remarkable enhancement in audio quality, existing models overlook diversity evaluation. This is partially due to the lack of a systematic sound class diversity framework and a matching dataset. To address these issues, we propose DiveSound, a novel framework for constructing multimodal datasets with in-class diversified taxonomy, assisted by large language models. As both textual and visual information can be utilized to guide diverse generation, DiveSound leverages multimodal contrastive representations in data construction. Our framework is highly autonomous and can be easily scaled up. We provide a text-audio-image aligned diversity dataset whose sound event class tags have an average of 2.42 subcategories. Text-to-audio experiments on the constructed dataset show a substantial increase of diversity with the help of the guidance of visual information. Our samples are available at <https://divesounddemo.github.io>.

Index Terms: diverse audio generation, multimodal aligned dataset, diverse subcategory taxonomy

1. Introduction

Audio generation, especially text-conditioned audio generation, has been attracting more and more attention in recent years, with numerous applications in film dubbing, game production and virtual reality [1]. In this scope, the text-to-sound generation focuses on producing sound events beyond speech and music, such as those in the natural world. Many competitive audio generation methods have been proposed recently, such as DiffSound [2], AudioGen [3], Make-An-Audio [4] and AudioLDM [5], whose competence in different audio generation tasks has been well demonstrated.

However, most prevailing audio generation methods overlook the diversity of the generated sound, and diversity measurement for the same sound event is rarely mentioned. Meanwhile, the demand for diversity in audio generation continues to grow with the ever increasing need, because the natural world presents a wide array of sounds, even within the same sound event. The lack of diversity arises partly due to the long-tailed distribution of existing datasets, which causes models to predominantly learn the most frequent occurrences rather than diverse distributions. As shown in the Figure 1, different types of sound may belong to the same sound event, in accordance with different text descriptions or images. Nevertheless, Existing models can only generate monotonous sounds or a limited range of variation.

Some studies have attempted to enhance diversity in the generated audio [6], and they adopt unsupervised algorithms to

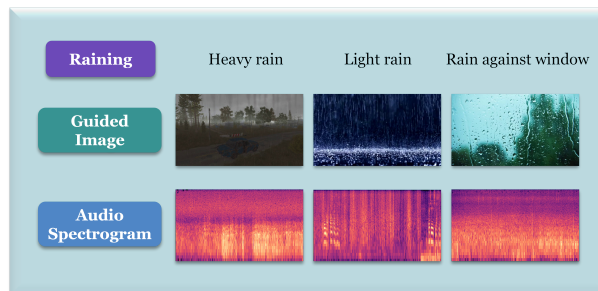


Figure 1: Multimodal text-audio-image dataset samples

cluster sound events and integrate visual information to guide generation. This line of research has achieved promising results on small datasets. However, unsupervised algorithms require researchers to explicitly assign labels to each clustered category and find matching images, heavily relying on the subjective interpretation of researchers and thus cannot be scaled up to large datasets. Hence, there is a need for an automated framework to construct diverse datasets, eliminating the necessity for manual intervention and enabling easy scalability. Due to the inherent alignment between images and audio, better control over audio can be achieved, which has been proved in [6, 7]. Consequently, the proposed subcategories should be able to be distinguished textually and visually.

In this paper, we propose a framework involving three processes: (1) clustering and (2) reasoning to obtain a taxonomy with diverse subcategories, and (3) a matching process to acquire high-quality text-audio-image paired data. Specifically, the proposed framework **DiveSound**, aided by large language models (LLMs) to cluster and inference, yields a taxonomy with diverse subcategories that are distinguishable both visually and textually. This takes into account the inherent alignment between audio and visual modalities. Leveraging the taxonomy descriptors as intermediaries, we construct an automatic data matching process to establish a dataset. DiveSound comprises diverse text-audio-image data, where models trained on it demonstrate outstanding diversity while ensuring the quality of audio. In contrast to prior efforts, DiveSound is automatically constructed, facilitating easy scalability without introducing bias from human annotators. The contributions of this work are summarized as follows:

- We propose DiveSound, a novel taxonomy to properly define sound diversity and sub-categorize within-class variety, accompanied by an automatic pipeline to align high-quality text-audio-image data.
- Both subjective and objective evaluations demonstrate that by incorporating such a taxonomy-based dataset can enhance

*Mengyue Wu and Kai Yu are the corresponding authors.

generated sound quality and diversity, where visual modality is most helpful in guiding audio generation.

2. Automatic Construction of Diverse Subcategory Taxonomy

In the construction framework, we initially establish a taxonomy to characterize diverse subcategories for categorizing and organizing high-quality data. Subsequently, we employ an automatic matching process to construct the text-audio-image dataset.

2.1. Characterizing Sound Diversity Taxonomy with LLM

Many sound classes exhibit great within-class diversity however there currently lacks an investigation into what sound classes are mostly diversified and how many sub-categories can be identified. To cluster and reclassify sound event labels, we employ GPT-4 [8], a start-of-the-art LLM with outstanding performance in natural language processing and generation tasks. VGGSound comprises 300 labels, among which certain labels correspond to exceedingly specific sound events for which further diversification is unnecessary. For example, we rarely need to generate the sound of a particular specie of an unusual animal and some special sound events like ‘*smoke detector beeping*’ have no need for classification. At the same time, other labels correspond to sound events emitted by a singular entity and cannot be differentiated through visual stimuli alone (e.g., ‘*dog barking*’ and ‘*dog howling*’, which are included in ‘*dog*’). As a consequence, we need GPT-4 to select the significant labels and regroup them to certain meaningful sound event class, and then generate the subcategories.

2.1.1. Clustering VGGSound labels

VGGSound [9] dataset categorizes auditory labels into nine overarching categories, namely *animals*, *home*, *music*, *nature*, *people*, *sports*, *tools*, *vehicle* and *others*. To enhance the quality of the clustering, this study employed a consistent format of prompt to direct the GPT model to merge different categories separately, which means breaking down a substantial task into smaller components to ensure generation performance. Illustrative representation of the clustering methodology is presented in Figure 2. After clustering, a refined selection with 69 **new classes**, containing 83 distinct labels from the VGGSound.

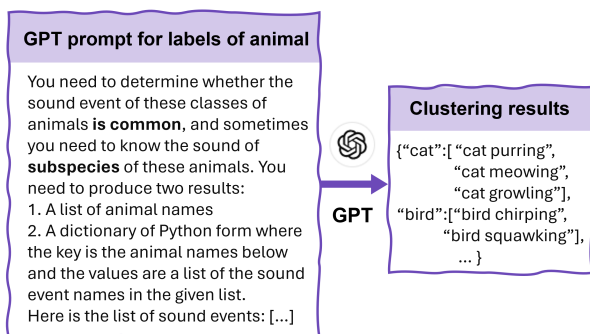


Figure 2: The process of regrouping the VGGSound labels into new sound event labels, in the example of animals. Sound event labels for other overarching categories are also classified using the same form of prompt.

2.1.2. Reasoning subcategories of each class

After obtaining the selected and summarized labels, we aim for GPT to reclassify these categories and generate more diverse labels. During this reclassification and generation process, we aim to leverage the visual information included in VGGSound. Therefore, we have incorporated into the classification criteria that each subcategory within each major class should be distinguishable both visually and auditorily. The specific form of the prompt and classification results is demonstrated in Figure 3. After this selection and generation work by GPT, we ultimately obtained 40 valid new sound event classes.

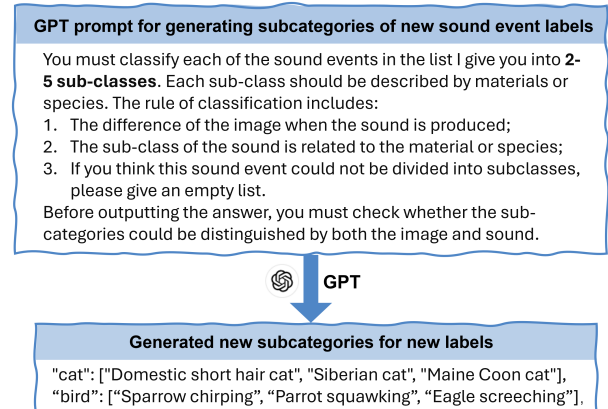


Figure 3: The prompt and examples of editing classifications and creating new subcategories. Three important rules are listed.

2.2. Automatic Text-Audio-Image Data Matching Process

To construct text-audio-image data pair, we use text descriptions serving as a bridge for connecting visual and audio modalities, thus CLIP [10] and CLAP [11] are utilized to create a shared space for text-image and text-audio respectively. The pipeline of this auto-matching process is illustrated in Figure 4, which could be summarized by 3 steps.

Step 1: Connect text and image features We use text as an intermediate pivot to bridge visual and auditory information. In step 1, we focus on matching text with corresponding images. After obtaining the reclassification result, we utilized CLIP to extract features from randomly selected frames. Simultaneously, we extracted features from the subcategories’ text of each new class generated by GPT. This serves two purposes:

- Calculate the probability of images being classified as various texts, identifying the potential sound events depicted in each audio’s visual frame.
- Find the image with the highest similarity to each class of text and image, serving as representative image information for that audio class.

Step 2: Connect text and audio features A similar approach is utilized to align text with audio features. We employ CLAP to identify the closest subclass within each new class for every audio. To enhance classification accuracy, we incorporated 2-4 adjectives generated by GPT-4 for each subclass during similarity calculations. The Top diagram in Figure 4 shows that the text description of each subcategory on the left is concatenated with the generated adjectives on the right.




Figure 4: Top: the automated matching process of text-audio data pairs. The example here uses the new class dog to demonstrate how an audio clip is matched with its corresponding text data pair. Bottom: the statistics of the newly selected dataset, including the 35 class labels with an average of 2.42 subcategories.

By computing the similarity between audio and text data, we determine the subclass text that best matches each audio within the respective new class.

Step 3: Match the text-audio-image data After obtaining text descriptions with the highest similarity scores to the audio and its associated visual frame respectively, we proceed with audio matching and filtering. When both image and audio features are most similar to the same text description, we establish a text-audio data pair, assigning audio data to this subclass. The new class should have enough audio data, as a result if the number of audio data corresponding to a subclass is less than 20, we discard that subclass.

End Product This process resulted in 35 classes, with an average of 2.42 subcategories per class and over 10K clips. The details of each class and the class labels are shown in Figure 4. These subcategories are not only audibly diversified from each other, but also visually distinguishable and can be described with varying text. For instance the different species of dogs could be distinguished both by their appearance and their barking sound. Hereby, we identify sound classes that exhibit distinguishable inner-diversity and provide corresponding textual descriptions as well as visual supplements.

3. Diverse Audio Generation Model

In this section, we introduce the generation model, which is similar to Xie *et al.* [6] where latent diffusion model is used for generation. Text and visual features are embedded into training with a fusion module as diversity guidance.

3.1. Modality Fusion

The information from different modalities is fused into the model in the training process. Initially, the new class labels C_i are encoded into embedding $\mathcal{E}_{\text{label}}$ using a lookup table. Then, the audio data from each subcategory $C_{i,j}$ of each class obtained in Section 2.2 is augmented with the descriptive text v_{text}

and representative image v_{image} features extracted by CLAP and CLIP respectively. The embedding of fusion inputs concatenated with visual and textual information are as follows:

$$\begin{aligned}\mathcal{E}_{\text{base}} &= \mathcal{E}_{\text{label}} \\ \mathcal{E}_{\text{text}} &= \text{concat}(\mathcal{E}_{\text{label}}, \text{CLAP}(v_{\text{text}})) \\ \mathcal{E}_{\text{image}} &= \text{concat}(\mathcal{E}_{\text{label}}, \text{CLIP}(v_{\text{image}}))\end{aligned}$$

3.2. Latent Diffusion Model

We utilized the Latent Diffusion Model (LDM) as the backbone structure, comprising Variational Autoencoder (VAE), diffusion, and vocoder models, which have demonstrated excellent performance in audio generation tasks [12, 5]. The VAE model is employed to extract representations of the audio, aiming to reduce computational complexity. It comprises an encoder, which compresses the Mel spectrogram into the latent space, and a decoder, which reconstructs the spectrogram based on samples. Finally, the waveform is reconstructed by the vocoder. A diffusion model is employed to predict the latent representation, relying on conditional inputs $\mathcal{E}_{\text{base}}/\mathcal{E}_{\text{text}}/\mathcal{E}_{\text{image}}$. In detail, a forward process continuously introduces noise into the latent representation, resulting in Gaussian noise. A reverse process gradually eliminates noise from the Gaussian noise.

We utilize the VAE introduced by AudioLDM [5]. Our LDM shares a similar architecture to that of Ghosal D *et al.* [13] but with a reduced parameter count, employing attention dimensions of $\{4, 8, 16, 16\}$ and block channels of $\{128, 256, 512, 512\}$. The LDM is trained using the AdamW optimizer for 80 epochs with a linear decay scheduler. We set the learning rate to 3×10^{-5} and the classifier free guidance scale to 3.

3.3. Evaluation Metrics

We use both objective and subjective metrics to assess the quality and diversity of the generated audio, thus comparing the im-

Table 1: *Objective and subjective evaluation metrics, tested on baseline, text-guided and image-guided systems. Real audios are considered as ground truth. p -value for the t -test comparing MSD results of the image-guided system with baseline and text-guided systems are 0.0055 and 0.0032 respectively. * indicates significance < 0.05 and ** for < 0.005*

System	Objective		Subjective		
	Quality (FAD) ↓	Diversity (MSD) ↑	Accuracy ↑	Diversity ↑	Naturalness ↑
Ground truth	-	-	4.00	4.07	3.87
Baseline	15.20 ± 0.93	10.92 ± 0.17	3.59	3.06	3.43
Text-guided	14.60 ± 1.02	11.30 ± 0.21	3.69	3.40	3.46
Image-guided	13.13 ± 0.79	12.62 ± 0.13	3.69	3.77	3.53

pect of different modalities on audio generation.

Objective-Quality Fréchet Audio Distance (FAD) [14] is an objective metric widely used in the field of text audio generation to measure the quality of generated audio. It is applied to measure the model performance in terms of generation quality by measuring the distance between generated audio and real audio. We chose it as our metric for objective-quality evaluation to measure model performance by calculating the average value of the model over 35 classes of generated audio.

Objective-Diversity Having ensured that the quality of the generated audio clips are comparable, Mean Squared Distance (MSD) is used as a measure of the diversity of the generated audio [7]. To calculate MSD between generated audio clips, the self-supervised BEATs [15] model is used to extract audio features that are not specific to categories. After the features have been extracted, we calculate the distance between any two audio clips in each class and use the average distance to represent an objective measure of the diversity of the audio in this class.

Subjective-Evaluation The Mean Opinion Score (MOS) is evaluated based on three criteria: event accuracy, generation diversity, and audio naturalness. We randomly selected 20% of the classes for subjective evaluation. 10 normally-hearing participants aged 22-24 who are familiar with audio/speech subjective evaluation tasks are selected as evaluators. They were asked to listen to concatenated audio segments, each 30 seconds long, consisting of generated audio or real audio inferred by the same system. This method is inspired by the generated image diversity evaluation [16] and allows evaluators to compare the diversity between audios generated by the same system. They then scored the entire audio segment with a marking scheme from 1 to 5. In terms of the diversity, they need to compare the difference between the . The final score was calculated as the average of the ratings from 10 evaluators.

3.4. Results

We compared vanilla system (baseline) results with text- and audio-guided generation models, with all experiments repeated using 3 random seeds. We calculate the objective metrics for all the systems. In Table 1, we present the average and standard deviation of FAD and MSD and the results of the MOS evaluation.

Diversity Enhancement Both objective and subjective metrics in Table 1 shows an enhancement in the diversity of generated audios with the guidance of both textual and visual information.

In the MSD evaluation, it is observed that the diversity of text-guided system is slightly higher than that of the baseline system, while the diversity metric of the image-guided system is notably elevated the previous two systems. In subjective evaluations, despite the diversity scores of the generated audio be-

ing lower than those of real audio, the image-guided system still outperforms the other two approaches. Low standard deviations show that the image-guided system has the strongest stability. The p -values indicates the guidance of image has a significant impact on the enhancement of the diversity of audio generation.

The subjective results indicate that the in-class diversity of the proposed dataset is guaranteed, which means that our framework is able to construct a high-quality dataset. According to both objective and subjective metrics, the improvement of generation diversity from baseline to systems with different modalities shows the effectiveness of training with model fusion. Furthermore, the visual information have a better performance than the textual information in terms of guiding sound generation.

Overall Quality In Table 1, FAD and subjective scores in terms of accuracy and naturalness indicate that incorporating image or text information has a positive effect on enhancing the quality of generated audio. Furthermore, the image-guided system performs the best, followed by the text-guided system, both slightly outperforming the baseline.

According to the objective metrics, text-guided and image-guided systems also show their better controllability of generated audio. That’s because while the textual information is augmented in the embedding and provides more information for each subcategory. Furthermore, image contains more aligned information to audio than text, leading to lower FAD between ground truth and image-guided system.

4. Conclusion

The aim of this study is to propose a method for automatically matching text-audio-image data pairs and creating multimodal datasets, with the further objective of enhancing the diversity of text-to-audio generation. We introduce the DiveSound framework, composed of a 3-stage approach using a large language model and pretrained CLIP and CLAP models. Based on the re-classification of VGGSound, a dataset comprising multiple categories including subcategories and multimodal information, is introduced. Furthermore, by training a diffusion-based text-to-audio model on the proposed dataset, we demonstrate the positive impact of incorporating different modality information on the quality and diversity of generated audio and image-guided generation outperforms other methods. However, the quality of subcategories’ text description is constrained by the knowledge of LLM and the accuracy of data matching is also limited by the representation ability of CLIP and CLAP. In our future study, we will apply our DiveSound framework on more data and explore different method for augmenting multimodal information, in order to furtherly increase the sound generation diversity.

5. 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.

6. References

- [1] Q. Kong, Y. Xu, T. Iqbal, Y. Cao, W. Wang, and M. D. Plumbley, "Acoustic scene generation with conditional sampler-nn," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 925–929.
- [2] D. Yang, J. Yu, H. Wang, W. Wang, C. Weng, Y. Zou, and D. Yu, "DiffSound: Discrete diffusion model for text-to-sound generation," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 31, pp. 1720–1733, 2023.
- [3] F. Kreuk, G. Synnaeve, A. Polyak, U. Singer, A. Défossez, J. Copet, D. Parikh, Y. Taigman, and Y. Adi, "AudioGen: Textually guided audio generation," in *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023.
- [4] R. Huang, J. Huang, D. Yang, Y. Ren, L. Liu, M. Li, Z. Ye, J. Liu, X. Yin, and Z. Zhao, "Make-An-Audio: Text-to-audio generation with prompt-enhanced diffusion models," in *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, ser. Proceedings of Machine Learning Research, vol. 202. PMLR, 2023, pp. 13 916–13 932.
- [5] H. Liu, Z. Chen, Y. Yuan, X. Mei, X. Liu, D. P. Mandic, W. Wang, and M. D. Plumbley, "AudioLDM: Text-to-audio generation with latent diffusion models," in *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, ser. Proceedings of Machine Learning Research, vol. 202. PMLR, 2023, pp. 21 450–21 474.
- [6] Z. Xie, B. Li, X. Xu, M. Wu, and K. Yu, "Enhancing audio generation diversity with visual information," *arXiv preprint arXiv:2403.01278*, 2024.
- [7] H. Ohnaka, S. Takamichi, K. Imoto, Y. Okamoto, K. Fujii, and H. Saruwatari, "Visual onoma-to-wave: environmental sound synthesis from visual onomatopoeias and sound-source images," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [8] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "GPT-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [9] H. Chen, W. Xie, A. Vedaldi, and A. Zisserman, "VGGSound: A large-scale audio-visual dataset," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 721–725.
- [10] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [11] 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 *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [12] H. Liu, Q. Tian, Y. Yuan, X. Liu, X. Mei, Q. Kong, Y. Wang, W. Wang, Y. Wang, and M. D. Plumbley, "AudioLDM 2: Learning holistic audio generation with self-supervised pretraining," *CoRR*, vol. abs/2308.05734, 2023.
- [13] D. Ghosal, N. Majumder, A. Mehrish, and S. Poria, "Text-to-audio generation using instruction-tuned LLM and latent diffusion model," *arXiv preprint arXiv:2304.13731*, 2023.
- [14] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi, "Fréchet audio distance: A reference-free metric for evaluating music enhancement algorithms," in *INTERSPEECH*, 2019, pp. 2350–2354.
- [15] S. Chen, Y. Wu, C. Wang, S. Liu, D. Tompkins, Z. Chen, W. Che, X. Yu, and F. Wei, "BEATs: Audio pre-training with acoustic tokenizers," in *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, ser. Proceedings of Machine Learning Research, vol. 202. PMLR, 2023, pp. 5178–5193.
- [16] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with clip latents," *arXiv preprint arXiv:2204.06125*, vol. 1, no. 2, p. 3, 2022.