

RETRIEVAL-AUGMENTED TEXT-TO-AUDIO GENERATION

Yi Yuan, Haohe Liu, Xubo Liu, Qiushi Huang, Mark D. Plumbley, Wenwu Wang

School of Computer Science and Electronic Engineering, University of Surrey, UK

ABSTRACT

Despite recent progress in text-to-audio (TTA) generation, we show that the state-of-the-art models, such as AudioLDM, trained on datasets with an imbalanced class distribution, such as AudioCaps, are biased in their generation performance. Specifically, they excel in generating common audio classes while underperforming in the rare ones, thus degrading the overall generation performance. We refer to this problem as long-tailed text-to-audio generation. To address this issue, we propose a simple retrieval-augmented approach for TTA models. Specifically, given an input text prompt, we first leverage a Contrastive Language Audio Pretraining (CLAP) model to retrieve relevant text-audio pairs. The features of the retrieved audio-text data are then used as additional conditions to guide the learning of TTA models. We enhance AudioLDM with our proposed approach and denote the resulting augmented system as Re-AudioLDM. On the AudioCaps dataset, Re-AudioLDM achieves a state-of-the-art Frchet Audio Distance (FAD) of 1.37, outperforming the existing approaches by a large margin. Furthermore, we show that Re-AudioLDM can generate realistic audio for complex scenes, rare audio classes, and even unseen audio types, indicating its potential in TTA tasks.

Index Terms— Audio generation, retrieval-information, diffusion model, deep learning, long tail problem

1. INTRODUCTION

The landscape of text-to-audio (TTA) generation has been revolutionized by advancements in diffusion-based generative modelling [1, 2, 3]. Leveraging powerful backbone models such as CLAP [1] and large language model (LLM) [4], these models are capable of extracting semantic information and enabling the creation of high-fidelity audio from textual descriptions.

In this work, we show that due to the scarcity and diversity of audio training data, bias appears in these state-of-the-art models, leading to performance degradation. Figure 1 (top) draws a statistical analysis conducted on the 327 labels of AudioCaps [5], one of the largest audio-text datasets, indicating a notable imbalance in data distribution. The bottom-left plot of Figure 1 shows a sample result of the state-of-the-art model trained with AudioCaps, when given the prompt “A man is talking then pops the champagne and laughs”, the model could only generate the content for “man talking”, but miss uncommon or complex events such as “champagne popped” then followed by “laugh”. Hence, an inherent limitation is seen due to the constrained scope and variability of the training dataset, where the quality of generated sounds seems largely correlated with their frequency of appearance during training. In this regard, these models can faithfully generate realistic audio clips for common sound events, but they may generate incorrect or irrelevant audio clips when encountering less frequent or unseen sound events.

We denote this the *long-tailed text-to-audio generation* problem, which influences the model performance in diversity and restricts

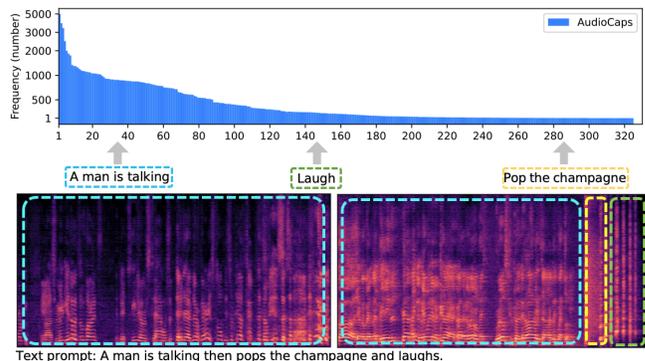


Fig. 1. The long-tailed problem in AudioCaps dataset (top). Example audio clips (bottom) generated with the baseline model (left) and Re-AudioLDM (right).

the applicability of these models, especially in real-world scenarios. Our motivation is to develop a robust TTA framework that breaks the barrier of imbalanced data and achieves realistic generation on diverse acoustic entities.

We propose a novel retrieval-augmented TTA framework to address the long-tailed generation issue. We enhance the state-of-the-art TTA model, AudioLDM [1], based on a latent diffusion model (LDM), with a retrieval module, dubbed Re-AudioLDM. Specifically, we first use the input text prompt to retrieve relevant references (e.g., text-audio pairs) from dataset (e.g., AudioCaps), and then use a pre-trained audio model and a language model to extract the acoustic and textual features, respectively. These extracted features are then further given to the cross-attention [6] module of the LDM to guide the generation process. The retrieved audio-text pairs serve as supplementary information that helps improve the modelling of low-frequency audio events in the training stage. In the inference stage, the retrieval-augmented strategy also provides references in relation to the text prompt, ensuring a more accurate and faithful audio generation result.

We perform extensive experiments on events with different frequencies of occurrence in the dataset. We show that Re-AudioLDM provides a stable performance among a variety of audio entities. It significantly improves the performance for tail classes over the baseline models, demonstrating that it can provide effective alleviation for long-tail TTA issues. Furthermore, as compared with the baseline models, Re-AudioLDM is capable of generating more realistic and complex audio clips, including rare, complex, or even unseen audio events. As the example with the same prompt shown in Figure 1 (bottom-left), Re-AudioLDM (bottom-right) can generate both uncommon entities “champagne popped” with a complex structure followed with the sound of “laugh”, achieving a better result than the baseline models with all the required entities and semantic orders correctly. In addition, Re-AudioLDM achieves an FAD score of 1.37, outperforming state-of-the-art TTA models by a large margin.

The remainder of this paper is organized as follows. Section 2

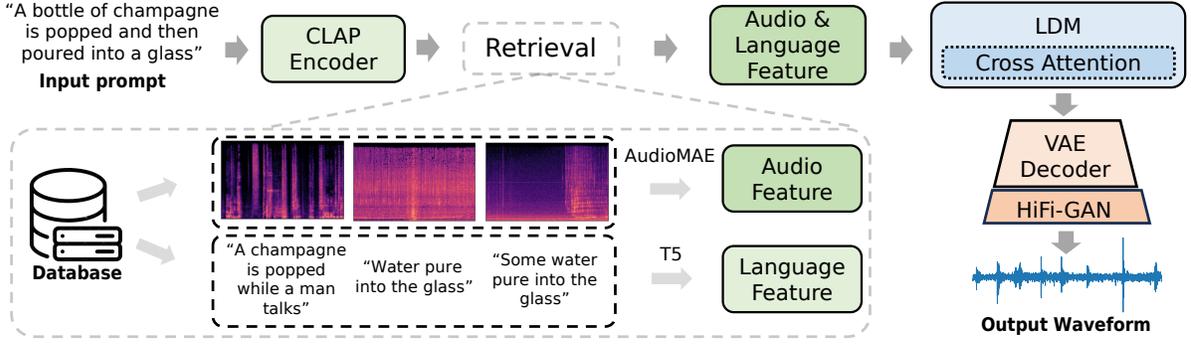


Fig. 2. The overview structure of Re-AudioLDM

introduces the related works of audio generation and retrieval-based models, followed by the details of Re-AudioLDM in Section 3. Section 4 presents the experimental setting and Section 5 shows the results and ablation studies. Conclusions are given in Section 6.

2. RELATED WORK

Our work relates to two major works, diffusion-based text-to-audio models and retrieval-based generative models. These two fields are briefly discussed in the following subsections.

2.1. Audio Generation

Recent works on audio generation follow an encoder-decoder framework [1, 7]. The model first uses an encoder to encode the information into a latent representation, which can be decompressed into a mel-spectrogram feature. The decoder then uses a variational autoencoder (VAE) and a generative adversarial network (GAN) vocoder to turn such features into waveforms. Liu *et al.* [8] used PixelSNAIL [9] as the encoder to represent labels while Iashin and Rahtu [10] apply a GPT2 [11] as the encoder to encode input images. Subsequently, diffusion-based models have been used for latent token generation. Yang *et al.* [12] replaced the transformer-based encoder with a diffusion-based encoder. Liu *et al.* [1] used the CLAP model [13] to obtain embeddings for the input data (audio or text), and used the LDM as the token generator. Ghosal *et al.* [4] then further improved this framework by replacing CLAP with LLM [14].

2.2. Retrieved Information Aggregation

Several studies in the field of image generation have considered leveraging retrieved information. Li *et al.* [15] extracted image features from a training set, and placed them in a memory bank which was then used as a parallel input condition for audio generation. Blattmann *et al.* [16] presented a nearest-neighbours strategy to select related image samples from a neighbourhood area. The KNN-Diffusion [17] used image features obtained from large-scale retrieval databases during the inference stage to perform new-domain image generation. Chen *et al.* [18] extended the image-only retrieval into image-text pair retrieval, augmenting both high-level semantics and low-level visual information for a diffusion model. In contrast, no similar works have been done for audio generation, and Re-AudioLDM is the first attempt to introduce retrieved information from a dataset to improve the text-to-audio generation performance.

3. PROPOSED METHOD

Similar to previous audio generation works [1, 4, 19], Re-AudioLDM is a cascaded model including three parts: input embedding, diffusion-based feature generator, and a pipeline to reconstruct the waveform from the latent feature.

3.1. Text and Retrieval Embedding Encoder

Re-AudioLDM takes two paralleled inputs: a text input c_t as low-level semantic information, and a set of text-audio pairs as retrieval augmentation c_r for high-level semantic-audio information. The text embedding E^t is obtained as:

$$E^t = f_{\text{clap}}(c_t) \quad (1)$$

which $f_{\text{clap}}(\cdot)$ is the CLAP model [20] used for text encoding, as in AudioLDM [1]. The retrieved information $c_r = \langle \text{text}_1, \text{audio}_1 \rangle, \langle \text{text}_2, \text{audio}_2 \rangle, \dots, \langle \text{text}_k, \text{audio}_k \rangle$ are the top- k neighbours selected through the similarity comparison between the embedding of the target caption and those of the retrieval dataset. Here for each pair, the multi-modal embedding is divided into two groups of concatenation (CAT), presented as audio retrieval E^{ra} and text retrieval E^{rt} , encoded as:

$$E^{ra} = \text{CAT}[f_{\text{mae}}(\text{audio}_1), \dots, f_{\text{mae}}(\text{audio}_k)], \quad (2)$$

$$E^{rt} = \text{CAT}[f_{\text{t5}}(\text{text}_1), \dots, f_{\text{t5}}(\text{text}_k)] \quad (3)$$

where $f_{\text{t5}}(\cdot)$ is a pre-trained T5 model [14] for obtaining the text embedding, and $f_{\text{mae}}(\cdot)$ is a pre-trained AudioMAE model [21] for obtaining the embedding of the paired audio.

3.2. Retrieval-Augmented Diffusion Generator

Re-AudioLDM uses LDM as the generator to obtain the intermediate latent token of the target audio. The diffusion model involves two processes, a forward process to gradually add noise into the latent vectors and a reverse process to progressively predict the transition noise of the latent vector in each step. During the forward step, the latent representation z_0 is transformed into a standard Gaussian distribution z_n with a continuous noise injection:

$$q(z_n | z_{n-1}) = \mathcal{N}(z_n; \sqrt{1 - \beta_n} z_{n-1}, \beta_n \mathbf{I}), \quad (4)$$

$$q(z_n | z_0) = \mathcal{N}(z_n; \sqrt{\bar{\alpha}_n} z_0, (1 - \bar{\alpha}_n) \epsilon) \quad (5)$$

where ϵ denotes the Gaussian noise with $\alpha_n = 1 - \beta_n$ controlling the noise level. In the reverse process, LDM learns to estimate the

distribution of noise ϵ_θ in the latent space, given conditions from the text embedding \mathbf{E}^t , calculated with equation (1), and the retrieved embedding \mathbf{E}^{ra} and \mathbf{E}^{rt} , calculated with equation (2) and (3), respectively. The LDM model applies UNet as the general structure, where the input layer takes the noisy latent vector z_n , text embedding \mathbf{E}^t , and the time step n as the condition. Then the retrieved information of both text and audio is shared with all the cross-attention blocks within the remaining layers. Employing a re-weighted training objective [22], LDM is trained by:

$$L_n(\theta) = \mathbb{E}_{z_0, \epsilon, n} \left\| \epsilon - \epsilon_\theta(z_n, n, \mathbf{E}^t, \text{Attn}(\mathbf{E}^{ra}, \mathbf{E}^{rt})) \right\|_2^2 \quad (6)$$

3.3. VAE Decoder & HiFi-GAN Vocoder

Re-AudioLDM utilizes a combination of a VAE and a HiFi-GAN as the general pipeline for reconstructing waveform from the latent feature tokens. During the training stage, VAE learns to encode the mel-spectrogram into the intermediate representation and then decode it back to mel-spectrogram, while HiFi-GAN is trained to convert mel-spectrogram into waveform. For inference, Re-AudioLDM applies the VAE decoder for mel-spectrogram reconstruction and HiFi-GAN for waveform generation.

4. EXPERIMENTS

4.1. Datasets

We use the AudioCaps dataset [25] for the experiments, which comprises 46,000 ten-second audio clips, each paired with a human-annotated caption. We follow the official training-testing split, where each training audio clip is assigned a single caption, while in the testing split, each audio clip is annotated with five captions. During the inference stage, we employ the first caption of each audio clip that appears in the test split as the text input. The remaining four captions are used only for the ablation study in Section 5.

4.2. Experiment Setup

Data Preparation. For a retrieval-based AudioCaps dataset, we apply a CLAP-score based retrieval function to find the top-50 nearest neighbours of the target text embedding. The waveform and the text from each neighbour are stored as a text-audio pair. It is noted that for both training and testing samples, the target audio clips are excluded from the retrieval information, which can avoid any access to the target data during both the training and inference stages.

Implementation Detail. As a cascaded model, the encoder and decoder parts of Re-AudioLDM are trained separately with audio clips sampled at 16 kHz. For the target, we use the short-time Fourier transform (STFT) with a window of 1024 samples and a hop size of 160 samples, resulting in a mel-spectrogram with 64 mel-filterbanks. Then, a VAE model is applied to compress the spectrogram with a ratio of 4, resulting in a feature vector with a frequency dimension of 16. For the information provided by the retrieval strategy, the text feature is directly extracted by a pre-trained T5-medium model, presenting a fixed sequence length of 50. The audio feature, on the other hand, is first converted into filter banks with 128 mel-bins and then processed by a pre-trained AudioMAE model, leading to a vector of dimension 32.

Training Detail. The LDM is optimized with a learning rate of 5.0×10^{-5} . Re-AudioLDM is trained for up to 80 epochs with a batch size of 4 and the evaluation is carried out every 100,000 steps. Re-AudioLDM-S applies a UNet architecture consisting of

128 channels, while we enlarge the model into Re-AudioLDM-L with 196 channels for experiments on more complex models.

Evaluation Metrics. Following Liu *et al.* [1], we use the Inception Score (IS), Fréchet Audio Distance (FAD), and Kullback–Leibler (KL) divergence to evaluate the performance of Re-AudioLDM. A higher IS score indicates a larger variety in the generated audio, while lower KL and FAD scores indicate better audio quality. For the semantic-level evaluation, we calculate the cosine similarity between the output audio embedding and the target text embedding calculated by the CLAP encoders, which demonstrates the correlation between audio and text.

5. RESULTS

5.1. Evaluation Results

The experiments are carried out on AudioCaps evaluation set. We compare the performance with several state-of-the-art frameworks, including AudioGen [23], AudioLDM [1] and Tango [4]. Selecting only the first caption of each audio clip as the text description, each framework infers 975 10-second audio clips with a sampling rate of 16 kHz. Table 1 compares the metrics achieved with different text-to-audio models, where Re-AudioLDM outperforms other methods by a large margin. It is noted that without the additional information provided by retrieval, Re-AudioLDM does not exhibit any advantage and is generally inferior to Tango, the current state-of-the-art model on AudioCaps. However, upon incorporating retrieval information, Re-AudioLDM successfully outperformed the baseline models in all four evaluation metrics. By enlarging the size of the hidden layer in LDM structure, the Re-AudioLDM-L using 10 retrieved pairs further decreases the FAD score to below 1.4, which is a significant improvement over the baseline frameworks.

5.2. Ablation Study

Retrieval Type. Experiments in Table 1 show the results on different retrieval information, e.g. audio, text, or both. With the audio features extracted by AudioMAE, only a slight improvement is achieved by Re-AudioLDM, mainly because Re-AudioLDM misses the relationship between sound events, although it captures the features of related sound events. By adding the paired text information of each retrieved audio clip, Re-AudioLDM learns the relationship between audio features and high-level semantic information, contributing a significant improvement in capturing highly related semantic features for the sound events.

Number of Retrieved Pairs. Several experiments are carried out to assess the impact of the number of retrieved audio-text pairs on audio generation performance. As depicted in Figure 3, the incorporation of retrieval information improves the performance, as the number of retrieved pairs increases, while such improvement slows down after the number reaches five and it becomes flattened at around ten. Therefore, in order to strike a balance between training costs and model performance, the number of retrieved pairs is chosen empirically to be in the range of 3 to 5 for this data.

Long-Tailed Situations. Re-AudioLDM aims to tackle the long-tailed generation problem and generate more realistic audio clips on uncommon or unseen sound events. In order to evaluate the accuracy of each generated audio clip, we applied the CLAP score [20] to show the relationship between the audio clip and text description. We first calculate the frequency of the occurrence of each sound event by counting the label of each audio clip and then illustrate the model performance by averaging the CLAP score of each sound class for

Model	Dataset	Retrieval Info	Retrieval Number	KL ↓	IS ↑	FAD ↓	CLAP _{score} (%)↑
AudioGen [23]	AC+AS+8 others	✗	✗	1.69	5.13	2.15	23.44
AudioLDM [1]	AC+AS+2 others	✗	✗	1.66	6.51	2.08	25.39
Tango [4]	AudioCaps	✗	✗	1.32	6.45	1.68	29.28
Re-AudioLDM-S	AudioCaps	✗	✗	1.63	6.48	2.31	26.75
	AudioCaps	Audio	3	1.54	6.88	1.95	31.05
	AudioCaps	Audio & Text	3	1.27	7.31	1.48	37.07
	AudioCaps	Audio & Text	10	1.23	7.33	1.40	37.15
Re-AudioLDM-L	AudioCaps	Audio & Text	10	1.20	7.39	1.37	37.12

Table 1. The comparison between different frameworks, with and without retrieval information. AC and AS are short for AudioCaps [5] and AudioSet [24] respectively.

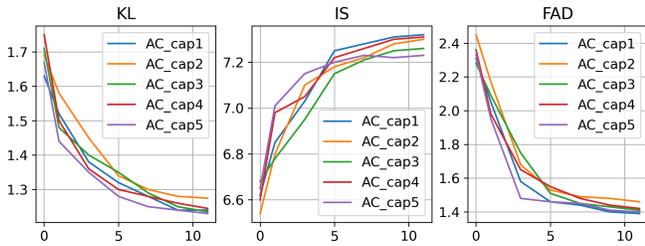


Fig. 3. Performance comparison on numbers of retrieved information, where AC_cap 1-5 refers to the caption groups of the testing set.

the AudioCaps testing set. The bar chart on the left side of Figure 4 presents a statistical analysis of the quantities of all 327 sound event classes in the AudioCaps training set. Similar to Figure 1 (top), tail classes constitute a significant portion, especially in the label group of 1 and 10. Figure 4 (right) shows the performance of each model on the events with different frequencies of event occurrence within the training set. Despite the initial gap in highly frequent audio events between Re-AudioLDM and baseline models, the baseline models perform worse when handling tailed entities. However, Re-AudioLDM has demonstrated more stable results, with a decrease of less than 3 in the CLAP score as the frequency of event occurrence is reduced in training data. Hence, Re-AudioLDM can reduce the degradation of output quality when generating tailed sound events, enhancing the overall model performance.

Zero-Shot Generation. For experiments on unseen entities, we evaluate several scenarios with events that are excluded during training. In Figure 4 (right), we found that baseline models show performance degradation on generating unseen audio (zero frequency occurrence). This may be because the model has not learned the features of unseen entities, while Re-AudioLDM can still achieve realistic results by providing related audio and semantic information. Hence, with essential retrieval information, Re-AudioLDM has the potential to generate sounds which are excluded from training data. The retrieval-based generation may significantly enhance the robustness of zero-shot tasks, which is one of the directions we will explore in the future.

Comparison with Mixup Strategy. Another way to address the class imbalance problem is to use a mixup strategy [26]. While mixup can increase the occurrence frequency for the tail entities, it also introduces more complex audio examples, as well as the synthetic audio data that may not align with real-world distributions. The results in [1] have shown that the mixup strategy leads to degradation

in overall performance. In contrast to the mixup method, our proposed retrieval augmentation strategy reduces the complexity of the training processes, resulting in an overall performance improvement.

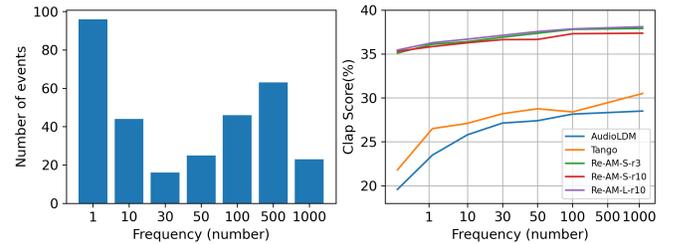


Fig. 4. Performance on different frequency entities, where S and L indicate model size and r refers to the number of retrieved clips.

6. CONCLUSION

In this paper, we have presented a retrieval-augmented model, Re-AudioLDM, to tackle the long-tailed problem in AudioCaps. The comparisons with current state-of-the-art models (i.e., AudioLDM and Tango) using several performance metrics (i.e., FAD, and CLAP-score) demonstrate that Re-AudioLDM can significantly enhance the performance of TTA models in generating high-fidelity audio clips. By integrating retrieved features, Re-AudioLDM not only achieves improvements in overall performance, but enables the generation of rare or unseen sound entities. In future work, we will investigate the model with external large datasets and explore the potential of the model in downstream tasks, such as zero-shot generation.

7. ACKNOWLEDGMENT

This research was partly supported by a research scholarship from the China Scholarship Council (CSC), funded by British Broadcasting Corporation Research and Development (BBC R&D), Engineering and Physical Sciences Research Council (EPSRC) Grant EP/T019751/1 “AI for Sound”, and a PhD scholarship from the Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising.

References

- [1] H. Liu, Z. Chen, Y. Yuan, X. Mei, X. Liu, D. Mandic, W. Wang, and M. D. Plumbley, "AudioLDM: Text-to-Audio generation with latent diffusion models," in *International Conference on Machine Learning*, 2023.
- [2] 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*, 2023.
- [3] X. Liu, Z. Zhu, H. Liu, Y. Yuan, M. Cui, Q. Huang, J. Liang, Y. Cao, Q. Kong, M. D. Plumbley, and W. Wang, "WavJourney: Compositional audio creation with large language models," *arXiv preprint arXiv:2307.14335*, 2023.
- [4] 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.
- [5] C. D. Kim, B. Kim, H. Lee, and G. Kim, "AudioCaps: Generating captions for audios in the wild," in *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2019.
- [6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [7] Y. Yuan, H. Liu, J. Liang, X. Liu, M. D. Plumbley, and W. Wang, "Leveraging pre-trained AudioLDM for sound generation: A benchmark study," in *European Signal Processing Conference*, 2023.
- [8] X. Liu, T. Iqbal, J. Zhao, Q. Huang, M. Plumbley, and W. Wang, "Conditional sound generation using neural discrete time-frequency representation learning," *IEEE International Workshop on Machine Learning for Signal Processing*, 2021.
- [9] X. Chen, N. Mishra, M. Rohaninejad, and P. Abbeel, "PixelSNAIL: An improved autoregressive generative model," in *International Conference on Machine Learning*, 2018, pp. 864–872.
- [10] V. Iashin and E. Rahtu, "Taming visually guided sound generation," in *British Machine Vision Conference*, 2021.
- [11] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [12] 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 Transactions on Audio, Speech, and Language Processing*, vol. 31, pp.1720-1733, 2023.
- [13] 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*, 2021, pp. 8748–8763.
- [14] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
- [15] B. Li, P. H. Torr, and T. Lukasiewicz, "Memory-driven text-to-image generation," in *British Machine Vision Conference*, 2022.
- [16] A. Blattmann, R. Rombach, K. Oktay, J. Müller, and B. Ommer, "Retrieval-augmented diffusion models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 15 309–15 324, 2022.
- [17] S. Sheynin, O. Ashual, A. Polyak, U. Singer, O. Gafni, E. Nachmani, and Y. Taigman, "KNN-diffusion: Image generation via large-scale retrieval," in *International Conference on Learning Representations*, 2023.
- [18] W. Chen, H. Hu, C. Saharia, and W. W. Cohen, "Re-Imagen: Retrieval-augmented text-to-image generator," in *International Conference on Learning Representations*, 2023.
- [19] Y. Yuan, H. Liu, X. Kang, P. Wu, M. D. Plumbley, and W. Wang, "Text-driven Foley sound generation with latent diffusion model," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events Workshop*, 2023, pp. 231–235.
- [20] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*, 2023.
- [21] H. Xu, J. Li, A. Baevski, M. Auli, W. Galuba, F. Metze, C. Feichtenhofer *et al.*, "Masked autoencoders that listen," *arXiv preprint:2207.06405*, 2022.
- [22] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Neural Information Processing Systems*, 2020.
- [23] 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 *International Conference on Learning Representations*, 2023.
- [24] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "AudioSet: An ontology and human-labeled dataset for audio events," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017, pp. 776–780.
- [25] T. Heittola, A. Mesaros, and T. Virtanen, "TAU Urban Acoustic Scenes 2019, Development dataset," Mar. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.2589280>
- [26] 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.