

# TEXT-DRIVEN FOLEY SOUND GENERATION WITH LATENT DIFFUSION MODEL

Yi Yuan, Haohe Liu, Xubo Liu, Xiyuan Kang, Peipei Wu, Mark D. Plumbley, Wenwu Wang

Centre for Vision Speech and Signal Processing, University of Surrey, United Kingdom

## ABSTRACT

Foley sound generation aims to synthesise the background sound for multimedia content. Previous models usually employ a large development set with labels as input (e.g., single numbers or one-hot vector). In this work, we propose a diffusion model based system for Foley sound generation with text conditions. To alleviate the data scarcity issue, our model is initially pre-trained with large-scale datasets and fine-tuned to this task via transfer learning using the contrastive language-audio pretraining (CLAP) technique. We have observed that the feature embedding extracted by the text encoder can significantly affect the performance of the generation model. Hence, we introduce a trainable layer after the encoder to improve the text embedding produced by the encoder. In addition, we further refine the generated waveform by generating multiple candidate audio clips simultaneously and selecting the best one, which is determined in terms of the similarity score between the embedding of the candidate clips and the embedding of the target text label. Using the proposed method, our system ranks 1<sup>st</sup> among the systems submitted to DCASE Challenge 2023 Task 7. The results of the ablation studies illustrate that the proposed techniques significantly improve sound generation performance. The codes for implementing the proposed system are available at [https://github.com/yyua8222/Dcase2023\\_task7](https://github.com/yyua8222/Dcase2023_task7).

**Index Terms**— Sound generation, Diffusion model, Transfer learning, Language model

## 1. INTRODUCTION

The development of deep learning models has recently achieved remarkable breakthroughs in the field of sound generation [1, 2, 3, 4]. Among various application domains of sound, Foley sounds, the mimic of background sound, play a crucial role in enhancing the perceived acoustic properties of movies, music, videos and other multimedia content [5]. The development of an automatic Foley sound synthesis system holds immense potential in simplifying traditional sound generation, which often involves intensive labour work on sound recording and mixing.

Currently, many sound generation models [1, 6, 7] adopt an encoder-decoder architecture, showing remarkable performance. Liu et al. [7] utilize a convolutional neural network (CNN) encoder, a variational autoencoder (VAE) decoder and a generative adversarial network (GAN) vocoder. The encoder embeds the input feature (e.g., label) into latent variables and the decoder transforms this intermediate information into mel-spectrogram which is then converted to a waveform by the vocoder. DiffSound [6] utilizes text as input and obtains the semantic features by using a contrastive language image pre-training (CLIP) model [8]. AudioGen [4] further improves the performance by using a pre-trained Transfer Text-to-Text Transformer (T5) [9] to obtain text embedding, which is then used to generate the waveform directly without using a vocoder.

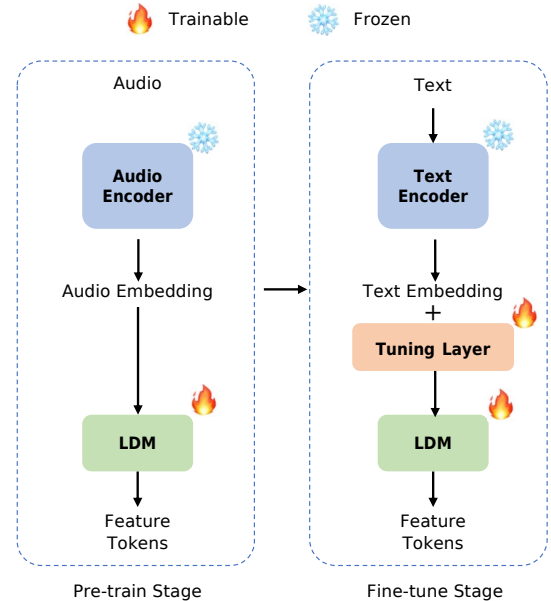


Figure 1: The training process of the LDM model. Audio Embedding are applied for pre-train the LDM on large dataset, while text embedding and an extra tuning layer are applied for fine-tune the LDM on target dataset

This paper proposes a latent diffusion model (LDM) based method for Foley sound generation. Our model follows the structure of AudioLDM [1], an audio generation model that comprises a diffusion model based encoder, a VAE based module for learning audio prior, and a HiFi-GAN vocoder for waveform generation. Due to the lack of training data for the sound generation task, we follow the idea of pre-training [10, 11], by initially training all three models on large-scale datasets such as AudioSet [12], AudioCaps [13] and Freesound<sup>1</sup>, and then transferring them onto the target development set. For inputs, the category labels are initially wrapped into relevant texts (e.g., turning the label “Keyboard” into text “Someone using keyboard”) before they are passed into the contrastive language-audio pre-training (CLAP) [14] for generating the text embeddings. To learn the most suitable semantic features of each sound, an embedding tuning layer is then added to text embedding for finding the optimal embedding during the fine-tuning stage. As shown in Figure 1, we first use audio embeddings to pre-train the LDM model. Then a tuning layer is introduced into the system, which is updated via transfer learning along with the LDM module. For outputs, the cosine-similarity score obtained in terms of the outputs and target labels is used to select the best-related sounds from a pool of candi-

<sup>1</sup><https://freesound.org>

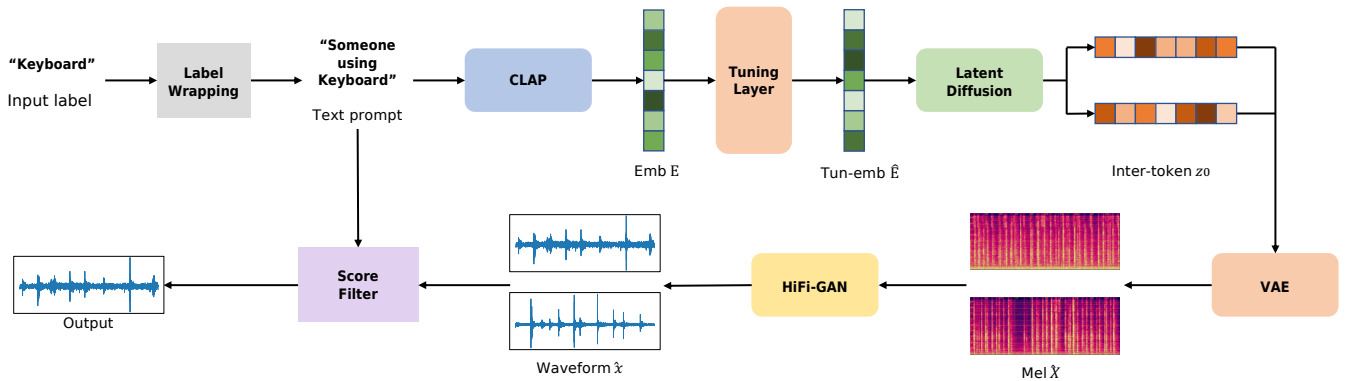


Figure 2: The overview of the system for both fine-tune stage and inference stage

date sound clips, which can improve the overall quality of the final result. Through experiments and ablation studies, we observe that the proposed techniques in our system can significantly improve the model performance in both the relevance of generated sound and the stability of the overall quality. Our system achieves a Fréchet audio distance (FAD) score of 4.52 on the DCASE task 7 validation set, significantly better than the baseline model with a FAD of 9.7.

The remaining sections of this paper are organised as follows. Section 2 describes the overview of the proposed system. The methodology of the network is explained in Section 3. Section 4 introduces the experimental setup. Results are shown in Section 5. Section 6 summarizes this work and draws the conclusion.

## 2. SYSTEM OVERVIEW

Our proposed system is based on the widely used structure on sound generation, which consists of an encoder, a generator, a decoder and a vocoder. The system adopts the same structure as AudioLDM [1], which used the CLAP [14] as the encoder and a latent diffusion model as the generator.

As a cascade model, the decoder and vocoder are trained separately and then built into the overall system with the trained parameters frozen when training the LDM model based generator. Instead of directly using labels as the input, we employ a wrapping strategy to generate text descriptions for each label as the initial mechanism to enhance the semantic information of the input. For example, we turn the label “Keyboard” into text “Someone using keyboard”. Then, we introduce an embedding tuning layer after the encoder in order to produce a more suitable embedding for each sound.

During the generating stage, with the text input, the system extracts the text embedding using the CLAP model, and the LDM model then generates the intermediate representation of the sound feature, using the text embedding as a condition. Subsequently, the mel-spectrogram can be decoded from the tokens by the VAE decoder, which is then transformed into waveform by the GAN vocoder. This system is then further improved with several techniques:

- Transfer learning is introduced to boost the performance by pre-training the model on larger datasets.
- A tuning layer is applied during the fine-tuning stage to find the optimal embedding.
- Similarity score between the embedding of the generated out-

put and the target embedding is applied to select the best match results among a group of waveform clips generated by the system.

Detailed explanations of these methods are provided in following section. The overall structure of the system is shown in Fig. 2.

## 3. PROPOSED METHOD

### 3.1. System structure

#### 3.1.1. CLAP based encoder

We use the CLAP model to obtain the embedding of the input. CLAP consists of a text encoder  $f_{text}$  that turns a text description  $y$  into text embedding  $\mathbf{E}^y$  and an audio encoder  $f_{audio}$  that computes an audio embedding  $\mathbf{E}^x$  from audio samples  $x$ . The two encoders are trained with cross-entropy loss, resulting in an aligned latent space with the same dimension  $D_e$  for both audio and text embedding. Since most large audio datasets (e.g., AudioSet) only provide audio-label pairs, we leverage the cross-modal information provided by two encoders. Specifically, the system is pre-trained on larger datasets with audio embedding and fine-tuned with text embedding on the task development set. During the fine-tuning process, the text embedding  $\mathbf{E}^y$  is passed through a trainable linear layer to find the optimal embedding feature for each class of sound. Details of this mechanism are presented in Section 3.2

#### 3.1.2. LDM based generator

Our system uses an LDM[15] to generate the intermediate latent tokens, with the feature embedding ( $\mathbf{E}^y$  or  $\mathbf{E}^x$ ) as the condition. These tokens are then used by the VAE decoder to generate the mel-spectrogram. During training, the LDM involves two processes: 1) A forward process where the latent vector  $\mathbf{z}_0$  is gradually turned into a standard Gaussian distribution  $\mathbf{z}_N$  in  $N$  steps, with noise  $\epsilon$  added in each step. 2) A reverse process for the model to predict the transition probabilities  $\epsilon_\theta$  of each step  $n$ , for reconstructing the data  $\mathbf{z}_0$  by removing the noise  $\mathbf{z}_N$ . The model is trained with a re-weighted objective [16] as:

$$L_n(\theta) = E_{\mathbf{z}_0, \epsilon, n} \|\epsilon - \epsilon_\theta(\mathbf{z}_n, n, \mathbf{E})\|_2^2 \quad (1)$$

where  $\epsilon_\theta$  is the Gaussian distribution predicted by LDM with current state  $\mathbf{z}_n$ , current step  $n$ , and current condition  $\mathbf{E}$ . During sam-

pling, the model first generates random Gaussian noise as  $z_N$ , and then applies the denoising process by predicting the reverse transition probability and taking the  $E$  from CLAP as the condition.

### 3.1.3. VAE decoder & HiFi-GAN vocoder

We utilize a combination of a VAE decoder and a HiFi-GAN vocoder to transform latent feature tokens into waveforms. Our approach involves training a VAE [17] to decode the latent feature tokens into mel-spectrograms, and a HiFi-GAN [18] to generate the corresponding waveforms. To achieve this, we initially convert the waveforms into mel-spectrograms using the Short-Time Fourier Transformation (STFT). The VAE is trained to compress the mel-spectrograms,  $X$ , into a latent space vector  $z_0$ , then reconstruct the mel-spectrograms  $\hat{X}$  from the compressed representation. In parallel, we employ a HiFi-GAN to convert the mel-spectrograms  $\hat{X}$  into the corresponding waveform representations, denoted as  $\hat{x}$ .

## 3.2. Practical issues

**Transfer learning** To deal with the issue of data scarcity, our system takes advantage of a pre-trained model [10] by initially training all three models on extensive audio datasets, followed by fine-tuning them on our development dataset. Specifically, the LDM model undergoes its initial training phase using large-scale datasets with audio embeddings as inputs, while the model is then trained on the development dataset utilizing text embeddings.

**Embedding tuning** To first initialize the text embedding with more semantic features, we apply some hand-picked text by extending the label with some adjunct word (e.g., dogbark into a dog bark). We then apply a tuning strategy to determine the optimal embedding of each sound class. To implement this, we introduce a linear layer  $L(x)$  with trainable parameters to fine-tune the text embedding before passing it to the LDM model. To guide this trainable layer with only minor updates on the embedding, the parameters are initialized with an identity matrix as weight, along with a Gaussian noise as bias  $l_b$ . Hence, the initial  $L(x)$  serves as an adding function of input  $x$  and bias  $l_b$  at the beginning of the training process. Then, the system learns to update the parameter of both weight and bias for optimal embedding during training. The embedding updated by this linear layer is also used as the target embedding for the score-selecting function discussed in the following section.

**Score-based selection** To improve the overall generation quality and robustness, a scoring mechanism is applied to determine the best matches among sampling results. Leveraging the fact that CLAP provides embeddings in the same latent space for audio and text, we utilize the cosine similarity between the output audio and the target text. By comparing the FAD score of different groups of output clips with different score-selecting thresholds, specific thresholds are established for each class, allowing the system to only select the results surpassing these thresholds.

## 4. EXPERIMENTAL SETUP

### 4.1. Dataset

**DCASE2023-T7** consists of a training set and an official evaluation set with seven different classes of fully labelled urban sounds. Each class has around 600 to 800 4-second sound clips in the training set and exactly 100 clips in the evaluation set. We randomly partitioned the training dataset into two subsets, with a ratio of 9 : 1 for training

and validation purposes, while the evaluation set was exclusively used during the evaluation phase.

**AudioSet** is a large-scale dataset for audio research, which consists of a wide range of sounds. In detail, Audioset provides around 2.1 million 10-second audio with 527 classes of labels. Our system uses AudioSet during the pre-training stage.

**Freesound** is a similar audio dataset with labels but with a non-fixed length, ranging from one second to several minutes. To unify the output length, all the sounds in Freesound are padded into a 10-second-long clip to match the data in Audioset.

By combining AudioSet and Freesound, we collected around 2.2M sounds in 22.05Khz for pre-training the LDM, VAE and GAN models. By using the audio-embedding and mel-spectrogram as input conditions, we only utilize the audio features to pre-train the models, while label features are then used during fine-tuning stage with the official training dataset.

## 4.2. Evaluation metrics

We apply the FAD [19] score as main evaluation metric. In detail, FAD calculates the Fréchet distance  $F$  between a group of target sound audio clip  $t$  and a group of generated sound audio clip  $r$ , formed:

$$F = \|\mu_r - \mu_t\|^2 + \text{tr}(\Sigma_r + \Sigma_t - 2\sqrt{\Sigma_r \Sigma_t}) \quad (2)$$

where  $\mu$  and  $\Sigma$  are the mean and covariance of Gaussian of the embedding vector from each group extracted by VGGish [20].

## 4.3. Parameter setting

Both the decoder and vocoder are trained separately, then they are integrated into the overall system with parameters fixed when training the LDM model. Initially, all three models are pre-trained using AudioSet and Freesound from scratch and then fine-tuned with the development set.

For the mel-spectrogram of 22.05kHz sounds, we set the window length as 1024 samples, the hop size as 256 and the number of mel-filterbank as 80. The VAE is trained with a compression level of 4, which encodes the mel-spectrogram into a latent vector of 20 in the frequency dimension and 86 in the time dimension. The length of the audio embedding  $E^x$  and text embedding  $E^y$  from the CLAP encoder in Section 3.1 is 512. All the models are optimized with Adam optimizer under an initial learning rate of  $3.0 \times 10^{-5}$ , with 3 epochs on the pre-training dataset and up to 1000 epochs on the training set. We test the model (LDM-S) performance by generating 100 clips per class and calculate the FAD score as the evaluation metric for every 100,000 step.

To investigate the influence of the input embedding features, we employ a diverse set of labels and texts as embeddings for training the model. In the case of embedding tuning, we begin by selecting a specific set of text for providing the initial text embedding. Then, we train the tuning layer along with the LDM model, guiding the embedding towards the optimal value.

To further investigate the potential of the model on sound generation, we trained a larger LDM with a bigger CLAP model (LDM-L) with the same training configurations. To balance the computing complexity and the output quality, we trained this model on 16kHz sounds and upsample it to 22.05kHz before output the results. For the 16kHz mel-spectrogram, the hop size is decreased to 160 with a mel-bin dimension of 64. The results of both models are shown in the following section.

System	Dog Bark	Footstep	Gun Shot	Keyboard	Moving Motor Vehicle	Rain	Sneeze Cough
Baseline [7]	13.41	8.11	7.95	5.23	16.11	13.34	3.77
LDM-S	4.41	7.44	7.46	3.13	16.97	12.62	3.02
LDM-S+Pre	4.17	6.86	7.25	3.15	15.68	12.95	2.85
LDM-S+Pre+Text	3.84	5.66	6.66	3.48	14.35	12.62	2.12
LDM-S+Pre+Text+Filter	3.53	5.04	5.65	2.80	15.29	9.76	<b>1.92</b>
LDM-S+Pre+Text+Filter+Tuned	<b>3.36</b>	<b>4.77</b>	<b>5.19</b>	<b>2.69</b>	14.83	10.00	1.98
LDM-L+Pre+Text+Filter+Tuned	6.04	5.05	6.44	3.07	<b>11.08</b>	<b>4.74</b>	2.93

Table 1: The best results of each system on the DCASE2023-T7 evaluation set. LDM-S: model trained from scratch. Pre: model with pre-training on large datasets. Text: using label-related text as input. Filter: applying the score-selecting function. Tuned: model with a fine-tuned embedding. The score selection for motor sound is used with the text embedding of “ A moving motor ”.

## 5. RESULTS AND ANALYSIS

The performance of our system on DCASE2023-T7 validation set is reported in Table 1. Most of our models outperform the baseline [7] by a large margin in terms of FAD. The results obtained from different sizes of LDM highlight distinct strengths: LDM-S is better at generating clear and distinct sounds like dog barks, footsteps, and gunshots, whereas the larger model (LDM-L) demonstrates superior performance in handling complex sounds such as motor sounds and rain sounds.

Ablation studies are also conducted to investigate the effects of each proposed technique. The experimental results in Table 1 demonstrate that transfer learning generally improves the system performance in most cases with respect to the evaluation metrics. Applying the embedding tuning strategy enables the system to optimize the embedding value for each class, which further improves the performance. To validate the effectiveness of this embedding tuning mechanism, we conduct several experiments with models trained with different frozen text embeddings. All the models are sampled under the same configuration for up to 20 times and the results of motor sound are presented in the box chart of Fig. ???. It can be observed that training without embedding tuning may yield results with varying quality, ranging from the best FAD of around 12.5 to the largest score of up to 18. On the contrary, generating a well-trained embedding value can contribute to more stable results. This might be because the updated embedding during training can provide more semantic information for both LDM denoising and the waveform tuning process.

Embedding	Moving Motor Vehicle
Label	16.97
Motor	13.14
A moving motor	12.12
Sound of motor	12.87
Driving/motor/car	12.07
Tuned embedding	11.08
Audio embedding	<b>8.88</b>

Table 2: The best results on LDM-L with FAD on motor sounds between different score-selection. Embeddings indicate the text/label value for training and similarity calculation.

From Table 1, the utilization of the similarity score function significantly enhances the overall performance, leading to improved output quality in most scenarios. However, despite the improvements observed in the majority of classes, we noticed that the generation quality of motor sounds did not exhibit a significant decrease in FAD (best achieved 11.08). By operating several subjective evaluations (human evaluation), we find out that this might be because

most motor sounds consist of noise-like sounds and sound events with distinct differences (e.g., driving sounds and engine sounds), making it challenging for CLAP to identify and extract a single embedding that aligns perfectly with all the target clips. To address this issue and improve the correlation of the score function, we introduced a multi-target-selection approach to replace the single embedding score-selection. Specifically, we collected a set of audio embeddings that demonstrated top feature correlation with the training dataset and randomly selected an audio embedding for the score-selection during each iteration. As the result presents in Table 2, our system with multiple audio-embedding filters achieves a notable FAD score of 8.88 for motor sounds.

## 6. CONCLUSION

This paper proposes a framework for small-domain Foley sound generation. Our system leverages a diffusion-based model and applied several methods to enhance performance. On the input feature, our experiment shows that the input embedding can significantly affect the overall quality. To alleviate this distinct gap between label and sound alignment, we proposed a trainable embedding for tuning the embedding value. Our result indicates that an improved embedding can further improve the quality and stability of the model. For output, a score-selection strategy is utilized to select the best clip along with CLAP score similarity. The experimental result shows that our system can significantly improve over the baseline network by a large margin. In the future, we will explore more efficient and end-to-end methods for audio feature extraction and fine-tuning.

## 7. ACKNOWLEDGMENT

This research was partly supported by a research scholarship from the China Scholarship Council (CSC) No.202208060240, funded from 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.

## 8. 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] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Sharifi, D. Roblek, O. Teboul, D. Grangier, M. Tagliasacchi, *et al.*, “Audiolm: a language modeling approach to audio generation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [3] 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,” *arXiv preprint arXiv:2301.12661*, 2023.
- [4] 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.
- [5] K. Choi, J. Im, L. Heller, B. McFee, K. Imoto, Y. Okamoto, M. Lagrange, and S. Takamichi, “Foley sound synthesis at the dcase 2023 challenge,” *In arXiv e-prints: 2304.12521*, 2023.
- [6] D. Yang, J. Yu, H. Wang, W. Wang, C. Weng, Y. Zou, and D. Yu, “Diffsound: Discrete Diffusion Model for Text-to-sound Generation,” *arXiv preprint arXiv:2207.09983*, 2022.
- [7] 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*, pp. 1–6, 2021.
- [8] 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.
- [9] 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,” *arXiv preprint arXiv:1910.10683*, 2019.
- [10] Y. Yuan, H. Liu, J. Liang, X. Liu, M. D. Plumbley, and W. Wang, “Leveraging pre-trained audiolm for sound generation: A benchmark study,” *arXiv preprint arXiv:2303.03857*, 2023.
- [11] D. Hendrycks, K. Lee, and M. Mazeika, “Using pre-training can improve model robustness and uncertainty,” in *International Conference on Machine Learning*, vol. 97, 2019, pp. 2712–2721.
- [12] 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.
- [13] C. D. Kim, B. Kim, H. Lee, and G. Kim, “AudioCaps: Generating captions for audios in the wild,” in *NAACL-HLT*, 2019.
- [14] 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.
- [15] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 2022, pp. 10 674–10 685.
- [16] J. Ho, A. Jain, and P. Abbeel, “Denosing diffusion probabilistic models,” in *Neural Information Processing Systems*, 2020.
- [17] V. Iashin and E. Rahtu, “Taming visually guided sound generation,” in *British Machine Vision Conference (BMVC)*, 2021.
- [18] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: generative adversarial networks for efficient and high fidelity speech synthesis,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020, pp. 17 022–17 033.
- [19] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi, “Fréchet Audio Distance: A Metric for Evaluating Music Enhancement Algorithms,” *arXiv preprint arXiv:1812.08466*, 2018.
- [20] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” *arXiv preprint arXiv:1409.1556*, 2014.