

FlowSep: Language-Queried Sound Separation with Rectified Flow Matching

Yi Yuan, Xubo Liu, Haohe Liu, Mark D. Plumbley, Wenwu Wang
Centre for Vision, Speech and Signal Processing (CVSSP)
University of Surrey, Guildford, UK

Abstract—Language-queried audio source separation (LASS) focuses on separating sounds using textual descriptions of the desired sources. Current methods mainly use discriminative approaches, such as time-frequency masking, to separate target sounds and minimize interference from other sources. However, these models face challenges when separating overlapping soundtracks, which may lead to artifacts such as spectral holes or incomplete separation. Rectified flow matching (RFM), a generative model that establishes linear relations between the distribution of data and noise, offers superior theoretical properties and simplicity, but has not yet been explored in sound separation. In this work, we introduce FlowSep, a new generative model based on RFM for LASS tasks. FlowSep learns linear flow trajectories from noise to target source features within the variational autoencoder (VAE) latent space. During inference, the RFM-generated latent features are reconstructed into a mel-spectrogram via the pre-trained VAE decoder, followed by a pre-trained vocoder to synthesize the waveform. Trained on 1,680 hours of audio data, FlowSep outperforms the state-of-the-art models across multiple benchmarks, as evaluated with subjective and objective metrics. Additionally, our results show that FlowSep surpasses a diffusion-based LASS model in both separation quality and inference efficiency, highlighting its strong potential for audio source separation tasks. Code, pre-trained models and demos can be found at: https://audio-agi.github.io/FlowSep_demo/.

Index Terms—Language-queried audio source separation, sound separation, rectified flow matching, multimodal learning

I. INTRODUCTION

Audio source separation systems aim to separate specific sound sources from audio mixtures. Previous research has made significant progress in various domains, including speech [1], [2], music [3], [4] and acoustic events [5], [6]. Recently, there has been growing interest in separating target audio sources using natural language queries as known as the Language-Queried Audio Source Separation (LASS) [7] task. LASS provides a useful tool for future source separation systems, allowing users to extract audio sources via natural language instructions. These systems could be useful in many applications, such as automatic audio editing [8], [9], multimedia content retrieval [10], and audio augmented listening [11].

The first attempt is the LASS-Net [7], which employs a BERT [12] network to encode textual queries, a ResUNet [3] module to predict the spectrogram masks of the target source. AudioSep [13] leverages contrastive multimodal pretraining models (e.g., CLAP [14]) and scales up the training data to 14,000 hours, which achieves state-of-the-art results and shows promising zero-shot separation performance. However, both LASS-Net [7] and AudioSep [13] primarily use masking-based

discriminative methods (e.g., spectrogram masking [2]) to separate target audio sources from mixtures. These systems may encounter challenges when dealing with overlapping sound events [15]. In particular, the masks generated by these models may be excessive or insufficiently selective, which leads to artifacts such as spectral holes or incomplete separation [16]. These limitations affect its effectiveness in real-world scenarios with diverse and dynamic acoustic environments.

Recently, researchers have been exploring separation source systems using non-discriminative models, such as generative models including Generative Adversarial Network (GANs) [17] and diffusion-based models [16]. Due to their generative nature, these models have the potential to enhance the perceptual quality of the separated sources and improve the overall subjective quality. Rectified Flow Matching (RFM) [18] is a recently proposed generative model. Similar to diffusion-based systems that learn to gradually denoise the output from general distribution, RFM models linear the relationships between data distributions and noise and provide superior theoretical properties and simplicity. However, it has not yet been explored in sound separation tasks.

In this work, we propose FlowSep, a LASS system using rectified flow matching. FlowSep employs a FLAN-T5 encoder [19] to embed the textual queries, followed by an RFM-based separation network. Specifically, FlowSep learns linear flow trajectories from noise to target features within the VAE latent space, conditioned on the query embeddings. During inference, the RFM-generated latent features are decoded into a mel-spectrogram using a pre-trained VAE decoder and then converted to a waveform with a pre-trained vocoder. We train FlowSep using 1.680 hours of audio data including AudioCaps [20], VGGSound [21] and WavCaps [22]. Experimental results across multiple datasets demonstrate that FlowSep significantly outperforms previous state-of-the-art models, such as AudioSep [13] in both objective and subjective metrics, showing promising separation performance in real-world scenarios. Additionally, comparisons between FlowSep and a diffusion-based LASS model indicate that RFM enhances both the output quality and inference efficiency, highlighting its strong potential for sound separation tasks.

This paper is organized as follows. Section II introduces the proposed system, followed by the dataset and evaluation methods applied to evaluate the performance of FlowSep in Section III. Section IV presents the experimental results and conclusions are given in Section V.




Fig. 1. The architecture of FlowSep. FlowSep consists of four main components: (1) a FLAN-T5 encoder for text embedding; (2) a VAE for encoding and decoding mel-spectrograms; (3) an RFM module for generating audio features within the VAE latent space; (4) a BigVGAN vocoder to generate the waveform.

II. METHOD

We propose FlowSep, an RFM-based generative model for language-queried audio source separation. FlowSep consists of four main components: a FLAN-T5 encoder for text embedding, a VAE for encoding and decoding mel-spectrograms, an RFM-based latent feature generator for predicting audio features within the VAE latent space, and a GAN-based vocoder [23] to generate the waveform. The model architecture and the separation workflow are illustrated in Figure 1.

A. Text Encoder

Unlike AudioSep [13], which uses a contrastive language-audio pre-training (CLAP) encoder [14] for text query embedding, we use a FLAN-T5 encoder [19]. As motivated by advancements in audio synthesis [24], [25], where FLAN-T5 has demonstrated better performance compared to CLAP.

B. Latent Feature Generator

The latent feature generator is a UNet-based [26]–[28] network with cross-attention modules to process the T5-embedding. We first introduce the RFM approach, followed by a description of the channel-conditioned generation approach, where we adapt RFM to the LASS task.

1) *Rectified Flow Matching*: RFM aims to predict the vector field μ which maps a linear pathway between the noise distribution $z_0 \in N[0, I]$ and target feature z_1 . Specifically, RFM learns this UNet neural network $\mu(\cdot, \theta)$ to predict the optimal transformation path for any noisy audio feature z_t with a random flow step $t \in [0, 1]$ and parameter θ . During training, the RFM first computes the noisy version of each data z_1 as:

$$z_t = (1 - (1 - \sigma)t)z_0 + tz_1 \quad (1)$$

where σ is a relatively small value [29] and chosen empirically as 1×10^{-5} in FlowSep. Taking both mixture audio feature z^m and query embedding E as condition, the target of RFM is summarized as:

$$v = z_1 - (1 - \sigma)z_0 \quad (2)$$

$$L_{\text{RFM}}(\theta) = \mathbb{E}_{t, z_1, z_0} \|\mu(z_t, E, z^m) - v\|^2 \quad (3)$$




Fig. 2. The channel-concatenation conditioning mechanism

Then during the inference stage, the output latent vector \hat{z}_1 is sampled from a noise z_n sampled from Gaussian distribution. By leveraging the theoretical properties and simplicity of RFM (e.g., linear flow trajectories), the inference process can be completed in a relatively small number of steps (e.g., fewer than 10 steps), greatly improving the system’s efficiency.

2) *Channel-Conditioned Generation*: FlowSep aims to generate the target feature \hat{z}_1 based on the mixture waveform and text query. Hence, we propose channel-conditioned generation to guide the model by taking the mixed audio as input channel conditions. As shown in Figure 2, both standard Gaussian distribution z_t and the latent vectors of mixture mel-spectrogram z^m are concatenated into the input channel before being forwarded into the RFM module. In this way, the additional condition, z^m , is considered as extra information within the input so the feature can be processed in parallel with the target latent vector by the RFM model. The noise adding forward process does not affect the mixture channel.

C. VAE Decoder and GAN Vocoder

FlowSep leverages a combination of a VAE and a GAN network for reconstructing the target waveform. The VAE encodes the mel-spectrogram into an intermediate representation (latent space feature), which is then decoded back to its original form. Following this, a GAN-based vocoder is trained to convert the decoded output into the waveform. In FlowSep, we use the state-of-the-art universal vocoder BigVGAN [23].

III. DATASET AND EVALUATION METRICS

A. Dataset

1) *Training Set*: we use 1,680 hours of audio from AudioCaps [20], VGGsSound [21] and WavCaps [22] for training.

TABLE I
OBJECTIVE EVALUATION ON LASS, WHERE AC, VGG AND ESC ARE SHORT FOR AUDIOCAPS [20], VGG SOUND [21] AND ESC50 [30] RESPECTIVELY. FAD DOES NOT APPLY ON UNPROCESSED DATA AS IT CALCULATE THE DIFFERENCE BETWEEN TWO GROUPS OF AUDIO, WHILE THE TARGET AND MIXED AUDIO SHARE THE SAME AUDIO EVENTS.

Model	FAD ↓				CLAPScore ↑					CLAPScore _A ↑			
	AC	DE-S	VGG	ESC	AC	DE-S	DE-R	VGG	ESC	AC	DE-S	VGG	ESC
Unprocessed	—	—	—	—	11.9	23.2	22.7	13.6	19.1	64.9	71.3	66.7	71.3
LASS-Net [7]	5.09	1.83	3.09	3.28	14.4	24.4	25.3	17.4	20.5	70.2	76.6	69.5	79.6
AudioSep [13]	4.38	1.21	2.30	1.93	13.6	26.1	29.7	19.0	21.2	69.6	78.9	72.4	80.5
FlowSep	2.86	0.90	2.06	1.49	21.9	26.9	31.3	19.5	22.7	81.7	80.1	73.2	80.7

AudioCaps [20] is the largest publicly available audio dataset consisting of 10-second audio clips paired with human-annotated captions. As the subset of AudioSet [31], AudioCaps contains 49,837 training clips and 957 testing samples.

VGGSound [21] is a large-scale audio dataset consisting of approximately 200,000 audio clips sourced from YouTube. Each audio clip in VGGSound has a duration of 10 seconds, while the audio clips in VGGSound are annotated with a set of 309 audio event textual labels, rather than detailed captions. **WavCaps [22]** is an audio dataset with machine captions generated using Large Language Models (LLM). We used audio samples shorter than 10 seconds and selected a total of 400,000 audio clips for training.

2) *Test Set*: We evaluate performance using five different benchmarks: the VGGSound and ESC50 [30] test datasets, as applied in the evaluation benchmark proposed by AudioSep [13], the AudioCaps testing set, and the two official evaluation datasets from DCASE2024 Task 9¹. During evaluation, it is ensured that target and noise sources in each mixture do not share any overlapping sound classes.

VGGSound testing set selected 200 clean and distinct audio samples as the target audio and 10 audio samples from the remaining testing set. The testing set consists of 2,000 mixtures by mixing each target sample and noise sample with random LUFS loudness between -35 and -25 dB.

ESC50 [30] evaluation set provides a total of 2,000 mixtures, where each clip is mixed with a sample from the ESC50 with a signal-to-noise ratio (SNR) at 0 dB.

AudioCaps testing set consisting of 928 samples, where we take each audio clip as the target source and mix it with a noise source selected from the testing set under random SNR rate between -15 and 15 dB. We select the first caption of the target source as the query for separation.

DCASE2024 Task 9 provides two evaluation sets which can be directly applied for evaluation: DCASE-Synth (DE-S) and DCASE-Real (DE-R). In detail, DE-S includes 3,000 synthetic mixtures mixed from 1,000 audio clips with an SNR rate ranging from -15 to 15 dB. DE-R consists of 100 audio clips collected from real-world scenarios. Each audio clip contains at least two overlapping sound sources. Each audio clip was annotated their component sources using text descriptions, so that each clip can be used as a mixture from

which to extract one or more of the component sources based on a text query. Each audio clip in DE-R was labeled with two such text queries.

B. Evaluation Metrics

Unlike discriminative networks used in previous systems [13], FlowSep does not separate audio events by masking the mixture input. The separated results are not strictly aligned with the target audio samples in the temporal dimension. Hence, traditional sample-level objective metrics for source separation tasks, such as source-to-distortion ratio (SDR), are not well-suited for evaluating the proposed system. Instead, we assessed its performance using five different metrics. For objective evaluation, we first use frechet audio distance (FAD) [32], a common metric for evaluating generation systems. Next, we apply two metrics based on CLAP [14] including CLAPScore [33], a reference-free metric measuring similarity between the output audio and the text query, and CLAPScore_A, which evaluates the similarity between the output audio and the target audio. For human evaluation, we follow the official methods used in the DCASE2024 Task 9 Challenge, which include assessing the relevance between the target audio and the language query (REL) and the overall sound quality (OVL). Both OVL and REL metrics are rated on a Likert scale from 1 to 5. All five datasets were evaluated using objective metrics, and we further conducted subjective evaluations on the AudioCaps, DE-S, and DE-R datasets, using a subset of 50 samples randomly selected from each. The subjective evaluations were performed on ten different listeners, with six researchers in audio and speech and four from other fields.

IV. EXPERIMENTS AND RESULTS

A. Data Processing

For training data processing, we first apply the PANNs [34], an off-the-shelf audio tagging model, to label the audio clips to ensure that every two audio clips used for creating synthetic mixture do not share any overlapping sound source classes. Then, the audio segments are padded or cropped to 10 seconds with 16 kHz sampling rate. We create synthetic training mixtures with a random SNR between -15 and 15 dB. The mixed waveform is then calculated through STFT under a frame of 1024 and a hop size of 160 to obtain the mel-spectrogram. For the language query, all the textual captions are converted into lower cases and punctuation is removed.

¹<https://dcase.community/challenge2024/task-language-queried-audio-source-separation>




Fig. 3. A case study of separation results on DE-S test set, as compared with the ground truth. More results can be found online from online.

TABLE II
SUBJECTIVE EVALUATION RESULTS ON LASS, WHERE AC, DE-S AND DE-R ARE SHORT FOR AUDIOCAPS [20], DCASE-SYNTH AND DCASE-REAL RESPECTIVELY.

Model	REL \uparrow			OVL \uparrow		
	AC	DE-S	DE-R	AC	DE-S	DE-R
LASS-Net [7]	3.12	2.96	3.59	2.16	2.84	3.88
AudioSep [13]	3.66	3.24	3.93	2.69	3.53	4.02
FlowSep	4.08	3.62	4.11	3.98	3.72	4.26

B. Experimental Details

We first train a 16 kHz BigVGAN [23] on the training datasets. For the encoder and decoder of the FlowSep, we apply the FLAN-T5-large [19] checkpoint and the pre-trained VAE model from AudioLDM [27] model. We freeze all the other components and the RFM model is trained for 1M steps with a batch size of 8 and a base learning rate of 5×10^{-5} . We use the publicly released LASS-Net [7] and AudioSep [13] models as baselines. For the diffusion-based network, denoted as DiffusionSep, we replace the RFM with general diffusion-based loss [27] and train the system using the same parameters.

C. Experimental Results

We evaluated FlowSep against baseline models using 10 steps with the ODE solver across five datasets, with the AudioCaps and VGGSound testing set being the non-zero-shot test dataset for FlowSep. Results in Table II and Table I show that FlowSep significantly outperforms the baselines across all datasets. Results on DE-R highlight its enhanced capability in real-world scenarios, results achieved on ESC50 and two DCASE2024 testing sets show its effectiveness in zero-shot cases compared to state-of-the-art models. FlowSep also achieved a higher average CLAPScore_A of 79 across four evaluation sets, showing better alignment and content accuracy with the ground truth. Furthermore, subjective metrics REL and OVL illustrate that FlowSep not only generates relevant audio events from natural language descriptions but also delivers higher overall perceptual quality.

D. Case Studies

We visualized the separated spectrograms of AudioSep and FlowSep using an example from the DE-S testing set. As shown in Figure 3, the spectrogram generated by FlowSep is closely matching the ground truth. In contrast, AudioSep’s results show incomplete separation with noticeable spectral gaps. The results highlight that, unlike previous discriminative approaches, FlowSep, as a generative network, demonstrates

TABLE III
THE EFFICIENCY ANALYSIS OF FLOWSEP AS COMPARED WITH THE BASELINE MODELS. THE VAE DECODER AND VOCODER INFERENCE TIME IS SHOWN AS SUPERSCRIPTS.

Model	Infer-step	Time(s)	FAD \downarrow	CLAPScore \uparrow
AudioSep	–	0.06	4.38	13.6
DiffusionSep	50	4.9 ^{+0.12}	4.52	10.4
DiffusionSep	100	9.4 ^{+0.12}	3.46	12.3
DiffusionSep	200	18.1 ^{+0.12}	2.76	18.8
FlowSep	10	0.58 ^{+0.12}	2.86	21.9
FlowSep	100	5.1 ^{+0.12}	2.75	22.8
FlowSep	200	9.0 ^{+0.12}	2.74	23.1

promising capabilities in source separation tasks. More case studies can be found online².

E. Efficiency Analysis

We conduct an efficiency analysis of FlowSep in comparison with AudioSep and a diffusion-based LASS model we implemented, with the performance results presented in Table III. The inference experiments were conducted on the AudioCaps dataset using a single A100 GPU with 80 GB of memory. As compared with AudioSep, our generative approach does not present better efficiency than discriminative approaches, while FlowSep generates results with better quality. Furthermore, FlowSep surpasses the diffusion-based LASS model in both FAD and CLAPScore metrics with the same number of inference steps (e.g., 200 steps). FlowSep-RFM maintains stable performance when the inference steps are reduced (e.g., only minor degradation when reducing the steps from 200 to 100) and achieves acceptable performance with even fewer steps (e.g., 10 steps). These findings underscore that RFM outperforms diffusion-based methods in separation tasks, and our proposed system can effectively reducing inference time while enhancing separation performance.

V. CONCLUSION

We introduced FlowSep, an RFM-based generative model for LASS, designed to leverage the strengths of RFM while overcoming the limitations of previous discriminative models for source separation. Experiments across various datasets demonstrate that FlowSep outperforms baseline models in both objective and subjective quality assessments. Furthermore, our results show that FlowSep present better performance on separation tasks than the diffusion-based model in both quality and the efficiency of inference. These findings highlight the significant potential of RFM for source separation tasks.

²https://audio-agi.github.io/FlowSep_demo

VI. 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] Y. Luo and N. Mesgarani, “Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM Transactions on Audio, Speech, and Language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [2] D. Wang and J. Chen, “Supervised speech separation based on deep learning: An overview,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [3] Q. Kong, Y. Cao, H. Liu, K. Choi, and Y. Wang, “Decoupling magnitude and phase estimation with deep resunet for music source separation,” *International Society for Music Information Retrieval Conference*, pp. 342–349, 2021.
- [4] H. Liu, L. Xie, J. Wu, and G. Yang, “Channel-wise subband input for better voice and accompaniment separation on high resolution music,” *INTERSPEECH*, pp. 1241–1245, 2020.
- [5] T. Ochiai, M. Delcroix, Y. Koizumi, H. Ito, K. Kinoshita, and S. Araki, “Listen to what you want: Neural network-based universal sound selector,” *arXiv:2006.05712*, 2020.
- [6] Y. Liu, X. Liu, Y. Zhao, Y. Wang, R. Xia, P. Tain, and Y. Wang, “Audio prompt tuning for universal sound separation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 1446–1450.
- [7] X. Liu, H. Liu, Q. Kong, X. Mei, J. Zhao, Q. Huang, M. D. Plumbley, and W. Wang, “Separate what you describe: Language-queried audio source separation,” *arXiv:2203.15147*, 2022.
- [8] Y. Wang, Z. Ju, X. Tan, L. He, Z. Wu, J. Bian *et al.*, “Audit: Audio editing by following instructions with latent diffusion models,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 71 340–71 357, 2023.
- [9] J. Liang, H. Zhang, H. Liu, Y. Cao, Q. Kong, X. Liu, W. Wang, M. D. Plumbley, H. Phan, and E. Benetos, “WavCraft: Audio Editing and Generation with Large Language Models,” in *ICLR 2024 Workshop on Large Language Model (LLM) Agents*, 2024.
- [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*, 2021, pp. 8748–8763.
- [11] F. Z. Kaghat, A. Azough, M. Fakhour, and M. Mekkassi, “A new audio augmented reality interaction and adaptation model for museum visits,” *Computers and Electrical Engineering*, vol. 84, p. 106606, 2020.
- [12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” *arXiv e-prints*, p. arXiv:1810.04805, Oct. 2018.
- [13] X. Liu, Q. Kong, Y. Zhao, H. Liu, Y. Yuan, Y. Liu, R. Xia, Y. Wang, M. D. Plumbley, and W. Wang, “Separate anything you describe,” *arXiv:2308.05037*, 2023.
- [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] H. Wang, D. Yang, C. Weng, J. Yu, and Y. Zou, “Improving target sound extraction with timestamp information,” *arXiv:2204.00821*, 2022.
- [16] J. Hai, H. Wang, D. Yang, K. Thakkar, N. Dehak, and M. Elhilali, “Dpm-tse: A diffusion probabilistic model for target sound extraction,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2024, pp. 1196–1200.
- [17] K. Chen, J. Su, and Z. Jin, “Mdx-gan: Enhancing perceptual quality in multi-class source separation via adversarial training,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 741–745.
- [18] X. Liu, C. Gong *et al.*, “Flow straight and fast: Learning to generate and transfer data with rectified flow,” in *The Eleventh International Conference on Learning Representations*, 2022.
- [19] 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.
- [20] 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.
- [21] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” *arXiv:1409.1556*, 2014.
- [22] X. Mei, C. Meng, H. Liu, Q. Kong, T. Ko, C. Zhao, M. D. Plumbley, Y. Zou, and W. Wang, “WavCaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 3339–3354, 2024.
- [23] S.-g. Lee, W. Ping, B. Ginsburg, B. Catanzaro, and S. Yoon, “BigVGAN: A universal neural vocoder with large-scale training,” in *International Conference on Learning Representations*, 2023.
- [24] Y. Yuan, H. Liu, X. Liu, Q. Huang, M. D. Plumbley, and W. Wang, “Retrieval-augmented text-to-audio generation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2024, pp. 581–585.
- [25] N. Majumder, C.-Y. Hung, D. Ghosal, W.-N. Hsu, R. Mihalcea, and S. Poria, “Tango 2: Aligning diffusion-based text-to-audio generations through direct preference optimization,” *arXiv:2404.09956*, 2024.
- [26] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.
- [27] 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.
- [28] H. Liu, Y. Yuan, X. Liu, X. Mei, Q. Kong, Q. Tian, Y. Wang, W. Wang, Y. Wang, and M. D. Plumbley, “Audioldm 2: Learning holistic audio generation with self-supervised pretraining,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [29] A. Vyas, B. Shi, M. Le, A. Tjandra, Y.-C. Wu, B. Guo, J. Zhang, X. Zhang, R. Adkins, W. Ngan *et al.*, “Audiobox: Unified audio generation with natural language prompts,” *arXiv preprint arXiv:2312.15821*, 2023.
- [30] K. J. Piczak, “ESC: dataset for environmental sound classification,” in *Proceedings of the ACM International Conference on Multimedia*, 2015.
- [31] 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.
- [32] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi, “Fréchet Audio Distance: A Metric for Evaluating Music Enhancement Algorithms,” *arXiv:1812.08466*, 2018.
- [33] F. Xiao, J. Guan, Q. Zhu, X. Liu, W. Wang, S. Qi, K. Zhang, J. Sun, and W. Wang, “A reference-free metric for language-queried audio source separation using contrastive language-audio pretraining,” *Workshop on Detection and Classification of Acoustic Scenes and Events*, 2024.
- [34] 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.