



DGMO: Training-Free Audio Source Separation through Diffusion-Guided Mask Optimization

Geonyoung Lee^{1,*}, Geonhee Han^{1,*}, Paul Hongsuck Seo¹

¹Department of Computer Science and Engineering, Korea University, Korea

gun0lee@korea.ac.kr, rtrt505@korea.ac.kr, phseo@korea.ac.kr

Abstract

Language-queried Audio Source Separation (LASS) enables open-vocabulary sound separation via natural language queries. While existing methods rely on task-specific training, we explore whether pretrained diffusion models, originally designed for audio generation, can inherently perform separation without further training. In this study, we introduce a training-free framework leveraging generative priors for zero-shot LASS. Analyzing naïve adaptations, we identify key limitations arising from modality-specific challenges. To address these issues, we propose Diffusion-Guided Mask Optimization (DGMO), a test-time optimization framework that refines spectrogram masks for precise, input-aligned separation. Our approach effectively repurposes pretrained diffusion models for source separation, achieving competitive performance without task-specific supervision. This work expands the application of diffusion models beyond generation, establishing a new paradigm for zero-shot audio separation.¹

Index Terms: target source separation, language-queried audio source separation (LASS), diffusion model

1. Introduction

Humans can focus on specific sounds in complex auditory environments, a phenomenon known as the cocktail party effect[1]. Computational models aim to replicate this ability through sound separation, isolating target sources from audio mixtures. Language-queried Audio Source Separation (LASS) has emerged as a flexible solution, allowing users to specify target sounds via natural language queries [2, 3, 4]. However, existing LASS models predominantly rely on task-specific training, where networks are explicitly trained for sound separation. Recent advances have explored generative models for LASS [5, 6], but these methods still require specialized training, limiting their flexibility and scalability across different domains.

In this study, we introduce a training-free framework that repurposes pretrained generative models for source separation. Diffusion models, which have demonstrated remarkable performance in audio generation [7, 8], remain largely unexplored for sound separation. Unlike prior LASS methods that require task-specific training, we investigate whether a pretrained generative model can inherently perform separation without further training for this task. Our approach leverages diffusion models' generalization ability, enabling zero-shot separation by extracting sound sources based on textual queries. To explore diffusion-based LASS, we first investigate naïve adaptations, such as in-

put mask optimization—an approach previously used in referring image segmentation [9], which is conceptually related to source separation. However, applying diffusion models to audio separation presents unique challenges due to the fundamental differences between audio and visual modalities including phase inconsistencies and the need for precise time alignment.

To overcome these challenges, we propose Diffusion-Guided Mask Optimization (DGMO), a test-time framework that integrates generative priors with explicit mask optimization. Rather than treating separation as a purely generative process, DGMO refines a learnable mask in the magnitude spectrogram domain, ensuring time alignment while leveraging diffusion-generated references in the mel spectrogram domain. This hybrid approach preserves the fidelity of separated audio, mitigating artifacts and inconsistencies seen in previous naïve generative methods [5, 6].

Our key contributions are as follows: (1) We establish a fully training-free framework by repurposing diffusion models for audio separation without additional training. (2) We identify limitations in naïve adaptations of diffusion models to LASS and propose Diffusion-Guided Mask Optimization (DGMO), a test-time optimization framework overcoming the unique challenges in the audio modality. (3) To the best of our knowledge, this is the first work to apply pretrained generative models to training-free, zero-shot source separation, expanding the role of diffusion models beyond generation.

2. Related Works

Language-queried Audio Source Separation Early sound separation models achieved success within predefined domains [10, 11, 12]. Research has since expanded to universal sound sources using vision [13], audio [14], label [3], and language queries. The language-based approach is appealing for its accessibility. LASS-Net [2] first introduced a BERT-based text encoder but required joint text-audio optimization. With multimodal learning advancements [15, 16, 17], methods aligning modalities in a shared space emerged, reducing alignment constraints [3, 4]. Moreover, generative approaches for LASS [5, 6] have been proposed to directly synthesize the target audio.

Diffusion Models and Non-Generation Tasks Diffusion models excel in text-to-image [18, 19] and text-to-audio tasks. AudioLDM [7, 20] and Aufferion [8] leverage latent diffusion for realistic audio synthesis. Beyond generation, they enhance test-time optimization and editing. DreamFusion [21] applies score distillation sampling for 3D synthesis, while Peekaboo [9] refines segmentation via inference-time mask optimization. Furthermore, audio editing methods [22, 23] and image inversion techniques [24, 25] demonstrate diffusion models' versatility in refining and manipulating signals.

*These authors contributed equally.

¹The code is available at: <https://wlttschmrz.github.io/DGMO/>.




Figure 1: Training-free LASS framework using pre-trained diffusion model. It has two key processes: a Reference Generation and a Mask Optimization.

3. Method

3.1. Language-queried Audio Source Separation

Given an audio mixture x composed of multiple source signals $\{s_i\}$ and environmental noise e formulated as $x = \sum_i s_i + e$, LASS [2] aims to extract a target source s^* described by a natural language query q . Conventionally, this task is addressed by estimating a mask $M(x, q)$ and applying it to the mixture, such that $s^* = x \odot M$, where \odot denotes the element-wise multiplication, preventing additional artifacts that may arise from directly generating signals. By leveraging textual descriptions instead of predefined categories, LASS enables flexible and intuitive audio separation. However, this task requires learning cross-modal associations between natural language queries and audio sources, posing significant challenges in achieving precise text-audio alignment. This challenge has led prior work to train task-specific models for learning such associations. In contrast, we explore the capability of pretrained diffusion models [7, 8, 20] originally designed for audio generation, to perform source separation without any task-specific training, leveraging their inherent generative priors for zero-shot language-queried audio source separation.

3.2. Diffusion Models and Mask Optimization

Diffusion models [26] are generative models that iteratively refine noisy inputs by learning a data distribution through a forward noise-injection and reverse denoising process. The reverse process estimates the original data by predicting and removing noise, formulated as:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(x_t, q, t) \right) + \sigma_t z \quad (1)$$

where α_t is the noise scaling factor, $\epsilon_\theta(x_t, q, t)$ is the predicted noise conditioned on the noisy input x_t , language query q , and timestep t , and $z \sim \mathcal{N}(0, I)$ is a standard Gaussian noise term with σ_t controlling the variance of the stochastic update.

While diffusion models have demonstrated high-quality generation including in the audio domain, their potential for signal separation remains largely unexplored. A notable exception is [9], a prior approach in computer vision that performs test-time optimization using score distillation loss with a pretrained diffusion model for segmentation based on a language query—an approach analogous to sound separation, as both tasks aim to isolate distinct components from an input mixture. Specifically, x_0 is masked by M before the noise injection

and M is optimized to minimize the diffusion loss function²:

$$x_t = \sqrt{\alpha_t}(x_0 \odot M) + \sqrt{1 - \alpha_t}\epsilon \quad (2)$$

$$M^* = \operatorname{argmin}_M \mathbb{E}_{\epsilon, t} \left[w_t \cdot \|\hat{\epsilon}_\theta(x_t, q, t) - \epsilon\|_2^2 \right] \quad (3)$$

where w_t is a weighting term computed from noise schedule parameters that depends on timestep t . Through this optimization process, the optimal mask M^* learns to remove irrelevant regions of the input image x_0 , ensuring it best corresponds to the query q effectively achieving segmentation.

Given the similarity between image segmentation and LASS, one may think that we can directly apply above technique to LASS. However, unlike visual signals, which are non-additive due to occlusion—where objects can block and completely remove parts of other objects—audio signals are additive, meaning multiple sources mix without fully masking each other. Therefore, to separate audio signals through masking, we cannot simply apply a binary mask as in visual segmentation. Unlike in the visual domain, where occluded parts can be directly masked out, audio separation requires computing the remaining audio signals to be removed, making the process as challenging as directly generating the target sound. This poses a unique challenge in the audio domain, preventing the above mask optimization with diffusion models from succeeding in the same way it does for visual segmentation.

3.3. Separated Audio Generation

An alternative approach to constructing separated audio signals using a diffusion model is to generate the target sound s^* directly from the query q , conditioned on the input mixture x . This approach is inspired by inversion-based editing techniques [25], where a model refines an existing signal to align with a given target representation by q .

Specifically, a denoised output x_0 can be generated from a noised input x_t , which is derived from the original mixture x . With an appropriately chosen t , the reconstructed x_0 serves as the separated audio, as it retains the essential content semantics of x while being regenerated under the condition q , effectively filtering out mismatching components. In this process, the choice of t is crucial: if too large, x_t may lose essential attributes from x , while if too small, it may not introduce enough noise for effective regeneration. A well-balanced t ensures that

²For notational simplicity, we present equations using regular diffusion models, though our experiments utilize latent diffusion models.

relevant information is preserved while allowing the model to refine the signal to align with the given query.

While this regeneration technique effectively generates sounds relevant to the query q and resembles the original source within x , the generated outputs often introduce artifacts or contain entirely new sounds that only superficially match the intended target, lacking true correspondence to the original signal. This highlights the need for an explicit constraint, similar to the mask optimization process, to ensure that the generated output remains faithful to the original source while effectively isolating the target sound.

3.4. Diffusion-Guided Mask Optimization

We propose a novel training-free LASS framework based on diffusion models, which overcomes the limitations of previous approaches by integrating both mask optimization and generative refinement into a unified process. This framework operates in two stages: reference generation and mask optimization.

Reference Generation In this stage, we generate separated audio given x and q following the procedure in Section 3.3, referring to the generated audio signals $\{s_i\}$ as references. As discussed, these references inherit attributes from x but often introduce sound elements that are not originally present in x due to the absence of explicit constraints, which are difficult to impose effectively within a diffusion model.

Mask Optimization Once the reference signals $\{s_i\}$ are generated, they encapsulate the knowledge embedded within the diffusion model regarding both the input mixture x and the query q . However, since there is no explicit constraint that ensures the separated sound s^* strictly belongs to the mixture x , we introduce a mask optimization process to enforce consistency with the input mixture. Specifically, rather than using the references directly as separated outputs, we use them as supervision signals to guide a mask M applied to the mixture x .

Since diffusion models operate in the mel spectrogram domain, we define the optimization loss by comparing the mel spectrograms of the masked mixture and the reference signal. However, applying the mask directly in the mel domain is infeasible due to the lossy, non-invertible mel transformation, which prohibits faithful waveform reconstruction. While vocoder-based reconstruction can be used to directly convert mel spectrograms back to waveforms, it typically induces temporal artifacts and alignment errors, as it generates phase through neural prediction instead of retaining the mixture’s true phase.

To mitigate these issues, we decouple the optimization and evaluation spaces: the mask is applied in the magnitude spectrogram domain for stable and interpretable reconstruction, while the loss is computed in the mel domain to maintain compatibility with the model’s conditioning. Formally, for each reference s_i , we define the objective as:

$$\mathcal{L}_i(M) = \|\text{mel}(x^{\text{spec}} \odot M) - s_i^{\text{mel}}\|_2^2 \quad (4)$$

where x^{spec} is the magnitude spectrogram of the input mixture, M is the mask, and s_i^{mel} is the mel-spectrogram of the corresponding reference s_i . This formulation enables effective gradient-based optimization while ensuring the extracted output remains both physically plausible and semantically aligned.

To improve robustness, we average the individual losses with multiple references $\{s_1, \dots, s_n\}$:

$$M^* = \underset{M}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(M) \quad (5)$$

Table 1: *Evaluation of Diffusion based LASS approaches. We present results on the AudioCaps dataset. Numbers reported are SI-SDR and SDRi values for our method and baselines.*

Models	SI-SDR	SDRi
Original Mixture (No Separation)	-0.07	0
Mask Optimization (Section 3.2)	-0.06	2.33
Separated Audio Generation (Section 3.3)	0.20	-0.24
Diffusion-Guided Mask Optimization (Ours)	1.99	3.57

Using multiple references mitigates high variance in mask optimization, as each reference captures different aspects of the target source. All components in this process are differentiable, allowing gradient-based optimization.

The estimated target waveform \hat{s}^* is then reconstructed using the optimized mask M^* and the original phase:

$$\hat{s}^* = \text{iSTFT}(x^{\text{phase}}, x^{\text{spec}} \odot M^*) \quad (6)$$

Here, x^{spec} and x^{phase} denote the magnitude and phase spectrograms of the mixture x , respectively.

DDIM Inversion A naive approach for reference generation injects random Gaussian noise into the input mixture. However, such arbitrary noise overwrites the structure and source-related signals, resulting in outputs that deviate from the original mixture content. While reducing the noise level might help retain more structure, it hampers the removal of non-target components. To address this, we adopt DDIM inversion [24, 25], a deterministic alternative that transforms the input mixture x_0 into a noisy x_t without randomness. Unlike random noise injection, DDIM inversion preserves the content structure of x_0 and maintains semantic fidelity throughout the reference generation process. This improvement ensures reliable reference signals, facilitating effective mask optimization.

4. Experiments

4.1. Evaluation Benchmarks

For evaluation, we use four publicly available text-aligned audio datasets and construct artificial mixtures following prior research in LASS [3, 4]. All datasets include both training and test sets. However, as our method is entirely training-free, we exclusively utilize the test set for evaluation. separation models. **VGGSound** [27] We adopt the evaluation setup of [4], where 100 clean target audio samples are each mixed with 10 randomly selected background samples from the test set. Loudness is uniformly sampled between -35 dB and -25 dB LUFS, and mixtures are normalized to 0.9 if clipping occurs, resulting in 1,000 mixtures with an average SNR of 0 dB.

AudioCaps [28] We follow [4], where the AudioCaps test set of 957 audio clips, each with five captions, is used to construct 4,785 mixtures for LASS. Each target source is mixed with five randomly selected background sources with different sound event tags. Mixtures are generated at 0 dB SNR, ensuring equal energy levels between the target and background sounds.

MUSIC [13] MUSIC contains 536 high-quality videos of 11 musical instruments sourced from YouTube. Following [3], 5,004 test examples for sound source separation constructed from 46 test videos from MUSIC by mixing randomly selected segments from different instrument classes at an SNR of 0 dB.

ESC-50 [29] While the dataset contains 2,000 audio clips across 50 classes, mixtures are created by pairing clips from different classes at 0 dB SNR. Constructing 40 mixtures per class, it contains 2,000 evaluation pairs.

Table 2: Benchmark evaluation results of DGMO and comparison with state-of-the-art LASS systems. For CLAP scores, except for our model, the results are sourced from [5].

Training Type	Models	VGGSound			AudioCaps			MUSIC			ESC-50		
		SI-SDR	SDRi	CLAP _{Score}	SI-SDR	SDRi	CLAP _{Score}	SI-SDR	SDRi	CLAP _{Score}	SI-SDR	SDRi	CLAP _{Score}
Supervised training	LASSNet [2]	-4.50	1.17	17.40	-0.96	3.32	14.40	-13.55	0.13	-	-2.11	3.69	20.50
	CLIPSep [3]	1.22	3.18	-	-0.09	2.95	-	-0.37	2.50	-	-0.68	2.64	-
	AudioSep [4]	9.04	9.14	19.00	7.19	8.22	13.60	9.43	10.51	-	8.81	10.04	21.20
Train-free	Ours	1.80	2.65	18.70	1.89	3.62	18.60	0.56	2.82	24.60	1.98	3.27	22.00

Table 3: DGMO with Various Diffusion Models. It presents the performance of DGMO applied to different models. Results are evaluated on the AudioCaps test set with 100 samples. Metrics reported are SI-SDR and SDRi. Additionally, we present FAD of these models, which are taken from the [8], where lower values indicate better generation performance by measuring the distance between generated and real audio distributions.

Audio Diffusion Model	FAD (Generation)	SI-SDR	SDRi
AudioLDM [7]	4.40	1.10	3.12
AudioLDM2 [20]	2.19	1.58	2.89
Auffusion [8]	1.63	1.99	3.57

4.2. Evaluation Metrics

We evaluate the performance of our methods using three widely adopted metrics: scale-invariant source-to-distortion ratio [30] (SI-SDR), signal-to-distortion ratio improvement [14] (SDRi), and CLAP Score [31]. SI-SDR measures the quality of separated signals by assessing residual distortion and interference, independent of signal scale. SDRi quantifies the improvement in separation quality relative to the original mixture, providing a comparative measure of enhancement. CLAP Score, a reference-free metric, evaluates the semantic alignment between the separated audio and the text prompt, reflecting how well the output matches the intended content. Higher values across all metrics indicate better separation performance.

4.3. Implementation Details

We use the pre-trained text-to-audio diffusion model, Auffusion [8], following the original diffusion model’s preprocessing. Audio is sampled at 16 kHz, padded to 10.24 s, then centered and normalized. We apply STFT with 256 mel filter banks, a window length of 1024, an FFT size of 2048, and a hop length of 160. For reference generation, DDIM inversion is performed in 25 steps with a noising step ratio of 0.7 and null text. We sample references in batches of 4 and optimize masks for 300 epochs per iteration, over 2 iterations.

4.4. Results

Comparisons to Naïve Approaches Table 1 compares the proposed method with the naïve approaches described in Sections 3.2 and 3.3. The naïve mask optimization method completely fails to find separation masks resulting in even lower scores than the original mixture x due to the complexity of the task. The separated audio generation technique improves scores but its effectiveness is limited, as the generated audio often contains signals not originally present in x . In contrast, the proposed diffusion-guided mask optimization successfully separates the target sound using only a pretrained diffusion model without any task-specific training.

Comparisons to Supervised Methods We compare our method with other supervised methods. LASS-Net [2] uses a

Table 4: Effect of Noising Step Ratio on DGMO Performance. Performance of DGMO with varying noising step ratios, evaluated on the AudioCaps test set (100 samples). The results demonstrate how the inversion ratio influences audio source separation quality. Metrics reported are SI-SDR and SDRi.

Method	Metric	Noising Step Ratio (t/T)				
		0.1	0.3	0.5	0.7	0.9
Random	SI-SDR	-0.59	-0.79	-0.80	-0.86	-1.05
Noise Injection	SDRi	2.28	2.57	2.68	1.60	2.48
DDIM Inversion	SI-SDR	-0.57	-0.34	0.62	1.99	2.04
	SDRi	2.71	2.81	3.15	3.57	3.64

pre-trained BERT [32] and ResUNet [33]. CLIPSep [3] employs CLIP [17] and SOP [13]. Both models operate in the frequency domain and reconstruct waveforms using noisy phase information. AudioSep [4] also employs the CLAP and trained with captioning data [28, 34]. We report the evaluation results as provided in prior work [5, 31], where models were assessed on the same dataset using predefined metrics

Ablations with Various Diffusion Models We evaluate our framework using multiple audio diffusion models. As shown in Table 3, our framework performs consistently well across different models, demonstrating its robustness. Additionally, the zero-shot separation performance generally aligns with the audio generation quality of each model (*e.g.*, SI-SDR vs. FAD for generation), indicating a strong correlation between a model’s generative capability and its effectiveness in source separation.

Effects of DDIM Inversion and Noising Step t Table 4 shows performance variations across different noising steps t . With random noise injection, too small a ratio t/T introduces insufficient noise, degrading separation quality. As the ratio increases, the injected noise dominates, reducing the correlation between the original input and the resulting signal. In contrast, DDIM inversion shows stable and superior performance across all noise scales. By leveraging structured, content-aware noise injection, it consistently mitigates the trade-off observed in random noise injection. These results highlight the robustness and effectiveness of DDIM inversion across different noise scales, reinforcing its suitability for source separation tasks.

5. Conclusion

We explored the feasibility of training-free LASS by leveraging pretrained diffusion models, originally designed for audio generation, for zero-shot source separation. We analyzed naïve adaptations of diffusion models to LASS and identified key limitations. To address these challenges, we introduced Diffusion-Guided Mask Optimization, a test-time optimization framework that refines spectrogram masks for accurate, input-aligned separation. Our results demonstrate that pretrained generative models can be effectively repurposed for source separation without task-specific training, achieving competitive performance.

6. Acknowledgements

This research was supported by IITP grants (IITP-2025-RS-2020-II201819, IITP-2025-RS-2024-00436857, IITP-2025-RS-2024-00398115, IITP-2025-RS-2025-02263754, IITP-2025-RS-2025-02304828), and the KOCCA grant (RS-2024-00345025) funded by the Korea government (MSIT, MOE and MSCT).

7. References

- [1] S. Haykin and Z. Chen, “The cocktail party problem,” *Neural computation*, vol. 17, no. 9, pp. 1875–1902, 2005.
- [2] 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,” in *Proceedings of Interspeech 2022*, 2022, pp. 1801–1805.
- [3] H.-W. Dong, N. Takahashi, Y. Mitsufuji, J. McAuley, and T. Berg-Kirkpatrick, “Clipsep: Learning text-queried sound separation with noisy unlabeled videos,” *arXiv preprint arXiv:2212.07065*, 2022.
- [4] 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,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [5] Y. Yuan, X. Liu, H. Liu, M. D. Plumbley, and W. Wang, “Flowsep: Language-queried sound separation with rectified flow matching,” *arXiv preprint arXiv:2409.07614*, 2024.
- [6] H. Wang, J. Hai, Y.-J. Lu, K. Thakkar, M. Elhilali, and N. Dehak, “Soloaudio: Target sound extraction with language-oriented audio diffusion transformer,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
- [7] 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,” *Proceedings of the International Conference on Machine Learning*, pp. 21 450–21 474, 2023.
- [8] J. Xue, Y. Deng, Y. Gao, and Y. Li, “Auffusion: Leveraging the power of diffusion and large language models for text-to-audio generation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 4700–4712, 2024.
- [9] R. Burgert, K. Ranasinghe, X. Li, and M. S. Ryo, “Peekaboo: Text to image diffusion models are zero-shot segmentors,” *ArXiv*, vol. abs/2211.13224, 2022.
- [10] Q. Kong, Y. Cao, H. Liu, K. Choi, and Y. Wang, “Decoupling magnitude and phase estimation with deep resunet for music source separation,” in *22nd International Conference on Music Information Retrieval, ISMIR 2021*. International Society for Music Information Retrieval, 2021, pp. 342–349.
- [11] 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.
- [12] Y. Liu, X. Liu, Y. Zhao, Y. Wang, R. Xia, P. Tain, and Y. Wang, “Audio prompt tuning for universal sound separation,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 1446–1450.
- [13] H. Zhao, C. Gan, A. Rouditchenko, C. Vondrick, J. McDermott, and A. Torralba, “The sound of pixels,” in *The European Conference on Computer Vision (ECCV)*, September 2018.
- [14] Q. Kong, K. Chen, H. Liu, X. Du, T. Berg-Kirkpatrick, S. Dubnov, and M. D. Plumbley, “Universal source separation with weakly labelled data,” *arXiv preprint arXiv:2305.07447*, 2023.
- [15] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, “Clap learning audio concepts from natural language supervision,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [16] 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.
- [17] 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.
- [18] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” 2021.
- [19] P. Esser, S. Kulal, A. Blattmann, R. Entezari, J. Müller, H. Saini, Y. Levi, D. Lorenz, A. Sauer, F. Boesel, D. Podell, T. Dockhorn, Z. English, K. Lacey, A. Goodwin, Y. Marek, and R. Rombach, “Scaling rectified flow transformers for high-resolution image synthesis,” 2024.
- [20] 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*, vol. 32, pp. 2871–2883, 2024.
- [21] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall, “Dreamfusion: Text-to-3d using 2d diffusion,” *arXiv preprint arXiv:2209.14988*, 2022.
- [22] 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.
- [23] M. Xu, C. Li, D. Su, W. Liang, D. Yu *et al.*, “Prompt-guided precise audio editing with diffusion models,” *arXiv preprint arXiv:2406.04350*, 2024.
- [24] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” *arXiv preprint arXiv:2010.02502*, 2020.
- [25] R. Mokady, A. Hertz, K. Aberman, Y. Pritch, and D. Cohen-Or, “Null-text inversion for editing real images using guided diffusion models,” *arXiv preprint arXiv:2211.09794*, 2022.
- [26] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [27] 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.
- [28] C. D. Kim, B. Kim, H. Lee, and G. Kim, “Audiocaps: Generating captions for audios in the wild,” in *NAACL-HLT*, 2019.
- [29] K. J. Piczak, “Esc: Dataset for environmental sound classification,” in *Proceedings of the 23rd ACM international conference on Multimedia*, 2015, pp. 1015–1018.
- [30] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “Sdr-half-baked or well done?” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.
- [31] 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,” *CoRR*, 2024.
- [32] J. Devlin, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [33] F. I. Diakogiannis, F. Waldner, P. Caccetta, and C. Wu, “Resuneta: A deep learning framework for semantic segmentation of remotely sensed data,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 162, pp. 94–114, 2020.
- [34] K. Drossos, S. Lipping, and T. Virtanen, “Clotho: an audio captioning dataset,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 736–740.