

SynSonic: Augmenting Sound Event Detection through Text-to-Audio Diffusion ControlNet and Effective Sample Filtering

Jiarui Hai, Mounya Elhilali

Department of Electrical and Computer Engineering
Johns Hopkins University, Maryland, USA

Abstract—Data synthesis and augmentation are essential for Sound Event Detection (SED) due to the scarcity of temporally labeled data. While augmentation methods like SpecAugment and Mix-up can enhance model performance, they remain constrained by the diversity of existing samples. Recent generative models offer new opportunities, yet their direct application to SED is challenging due to the lack of precise temporal annotations and the risk of introducing noise through unreliable filtering. To address these challenges and enable generative-based augmentation for SED, we propose SynSonic, a data augmentation method tailored for this task. SynSonic leverages text-to-audio diffusion models guided by an energy-envelope ControlNet to generate temporally coherent sound events. A joint score filtering strategy with dual classifiers ensures sample quality, and we explore its practical integration into training pipelines. Experimental results show that SynSonic improves Polyphonic Sound Detection Scores (PSDS1 and PSDS2), enhancing both temporal localization and sound class discrimination.

1. INTRODUCTION

Sound event detection (SED) is a technique to identify, classify, and temporally locate specific acoustic events within audio recordings [1]. It involves analyzing audio signals to detect occurrences of target sounds, such as speech, music, animal calls, or environmental sounds, along with their precise start and end times.

Recent advances in deep neural networks have led to significant progress in SED [1]–[3]. However, training SED models typically requires strongly labeled data, which provide precise annotations of the exact start and end times of each sound event. Acquiring such datasets is challenging and costly due to the labor-intensive annotation process. To address this limitation, synthetic soundscapes with precise annotations have been proposed as alternative training resources for SED systems [4], [5]. However, synthetic data introduces its own challenges. Although it enables the creation of numerous audio mixtures, the diversity of these mixtures is constrained by the range of available foreground and background samples. Furthermore, collecting high-quality, isolated foreground events remains a difficult task. When synthetic datasets are generated from limited sound samples, the resulting lack of diversity can lead to overfitting and impair the generalization capability of deep neural networks [1].

To mitigate these issues, data augmentation has emerged as a particularly effective solution. By artificially increasing the variability of training data, augmentation techniques can significantly enhance model robustness and improve generalization. Traditional data augmentation methods [6]–[9] have demonstrated strong performance in diversifying training signals without the need for additional manual annotations. However, these methods typically operate on existing data and are inherently limited by the diversity of the original samples, imposing an upper bound on the achievable variability.

Recently, generative models have been increasingly adopted for data augmentation in computer vision [10]–[12], thanks to their ability to automatically produce diverse and realistic data. In the audio domain, text-to-audio (T2A) models have shown promising capabilities in

generating sound clips containing single sound events, even when primarily trained on audio mixtures, and have been successfully used for data augmentation in sound extraction and classification [13], [14]. However, SED relies on precise temporal annotations, specifically, exact onset and offset times, to prepare strongly labeled data, which makes it challenging to directly apply the same augmentation strategies as sound separation and classification. Moreover, data selection is crucial when using generative models for data augmentation. Despite their ability to produce authentic data, generative models can still make mistakes, and indiscriminately incorporating generated samples may introduce noise and degrade performance. While models like CLAP [15] can assist with filtering, relying on a single scoring model may introduce selection bias, potentially limiting the quality and diversity of generated samples and impairing SED’s robustness.

To bridge this gap and bring the benefits of generative model-based augmentation to SED, we introduce SynSonic¹, which combines controllable foreground sound generation with an effective and reliable sample selection method. The main contributions of this work are summarized as follows:

(1) Controllable sound sample generation. To generate tightly segmented sound events with precise temporal structures, we integrate an energy envelope-based control signal into a text-to-audio diffusion model [16] using ControlNet [17]. This setup guides the generation process by enforcing target temporal patterns, enabling precise control over event timing during mixture synthesis.

(2) Robust sample selection. To automatically identify high-quality samples, we propose a dual-classifier joint ranking strategy that combines CLAP scores with classification logits from an AudioSet-based [18] audio classification model [19]. This method improves selection robustness, enhances sample quality, and reduces biases associated with single-score filtering.

(3) Balanced Real/Synthetic fusion. We conduct extensive experiments to establish practical guidelines for designing SynSonic, including analyses of foreground generation methods, data filtering settings, and mixture synthesis strategies. Additionally, we provide a usage recipe for effectively leveraging SynSonic, detailing configurations for Real/Synthetic data composition during training.

As a result, SynSonic consistently improves SED performance in both temporal localization and class discrimination metrics.

2. METHOD

In this section, we present the proposed data augmentation strategy, illustrated in Fig. 1, which comprises three key steps. First, foreground sound samples are generated using Energy-Envelope ControlNet to ensure both diversity and alignment with desired temporal characteristics. Second, the generated samples undergo filtering via a dual-classifier system to retain only high-quality and relevant examples. Finally, the selected samples are used to synthesize strongly labeled sound

This Work is supported by ONR N00014-23-1-2050 and N00014-23- 1-2086.

¹Source Code: <https://github.com/JHU-LCAP/SynSonic>

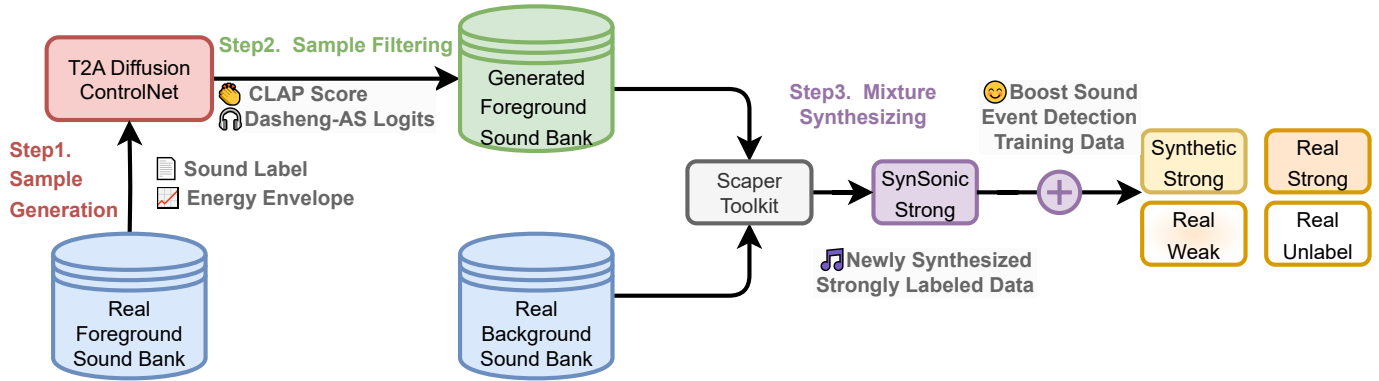


Fig. 1: The overall framework of SynSonic.

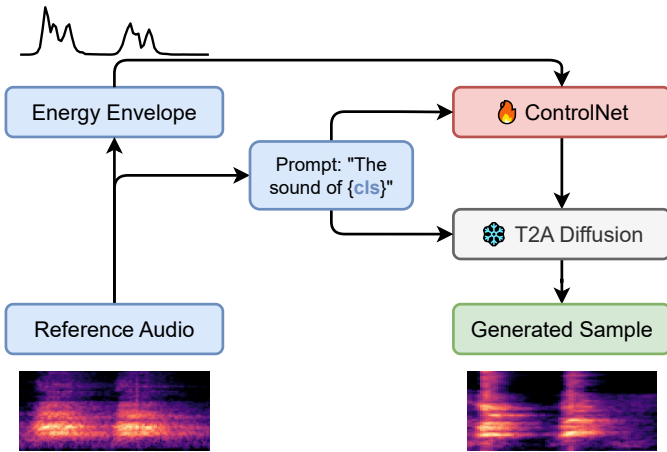


Fig. 2: Illustration of the energy-envelope ControlNet generating new foreground samples, using the sound class *dog* as an example.

mixtures, thereby enriching the dataset and improving the robustness of SED model training.

2.1. Soundbank Generation with Diffusion ControlNet

To address the challenge that foreground samples generated directly by T2A models often lack precise control over temporal attributes—such as timing and duration—which are critical for synthesizing strongly labeled data, we integrate a ControlNet-based mechanism into the generative process. ControlNet provides a structured approach that enables the model to be guided by external control signals in addition to the text prompt.

Specifically, since only the temporal structure needs to be controlled when generating foreground sound for SED, we adopt the energy envelope as the guiding signal, inspired by [20], [21]. This simple yet effective representation captures the temporal dynamics of sound events, including their onset, duration, and intensity progression. As a result, the diffusion model can generate audio that not only aligns with the semantic content of the text but also faithfully follows the intended temporal pattern. Importantly, it results in a well-defined onset-offset boundary.

We build the diffusion-based ControlNet on top of EzAudio [16], a recently released T2A model based on diffusion transformers, which delivers superior audio quality and prompt coherence. Our implementation adopts a standard ControlNet architecture, in which the first half of the model’s blocks are duplicated. These ControlNet blocks are connected to the latter half of the main transformer through




Fig. 3: Illustration of the filtering methods applied to samples generated for the sound class *vacuum cleaner*.

Algorithm 1 Dual-Classifier Sample Filtering

Require: Generated samples S , CLAP scores $\{c_i\}$, classification logits $\{l_i\}$, weight w , top percentage $p\%$

- 1: Rank samples by CLAP score $\rightarrow R_{\text{CLAP}}$
- 2: Rank samples by classification logit $\rightarrow R_{\text{Cls}}$
- 3: **for** each sample i **do**
- 4: Compute weighted rank score: $s_i = w \cdot R_{\text{CLAP},i} + (1-w) \cdot R_{\text{Cls},i}$
- 5: **end for**
- 6: Sort samples by s_i and select top $p\%$
- 7: **return** Filtered samples

long skip connections. The energy-envelope is processed through a zero-initialized one-dimensional convolutional layer and subsequently added to the ControlNet input to provide temporal guidance.

Leveraging the generative capabilities of the diffusion ControlNet, as illustrated in Fig. 2, we observe that conditioning on energy references still yields diverse audio outputs. Although these outputs share similar temporal envelopes with the references, they exhibit substantial variation in frequency characteristics, underscoring the model’s ability to generate diverse and non-redundant samples from the reference samples. These generated samples, diverse in their frequency characteristics, can thus enhance data diversity and help improve the robustness of the SED model.

2.2. Effective Sample Filtering with Dual Classifiers

Although the T2A model with ControlNet generally produces realistic sound samples, it can occasionally generate incorrect sound classes or unnatural audio. To ensure the quality of samples used for synthesizing strongly labeled data and to avoid introducing errors during SED model training, we adopt a filtering strategy to select only high-quality outputs. A CLAP-based model [13], [15], [22] can be employed to assess the similarity between the text prompt and the generated sample, helping to filter out unsuitable synthetic examples. However, this approach may introduce biases inherent to a single CLAP model.

To address this issue, we incorporate an additional audio classification model, Dasheng-AS [19], an AudioSet-based [18] audio classification model, to provide a complementary score based on the predicted class logits. Since our focus is on foreground generation, and each foreground sample corresponds to a single class, this approach is well-suited. For each generated sample, both the CLAP similarity score and the classification logit are computed. As shown in Algorithm 1, to account for scale differences between these two metrics, samples within each class are independently ranked by each score. A weighted rank-based score is then calculated to determine the final ranking. Finally, only the top-ranked samples, selected according to a predefined percentage threshold, are retained for subsequent mixture creation.

As illustrated in Fig. 3, although the CLAP scores and Dasheng-AS logits exhibit an overall linear correlation, notable discrepancies between the two metrics remain. Some samples with high CLAP scores do not correspondingly achieve high logits, and vice versa, suggesting that relying on a single metric may introduce selection biases aligned with only one model’s perspective and lead to suboptimal sample quality. In contrast, the proposed joint ranking approach integrates both scores, mitigates such metric-specific biases, and ensures that only samples demonstrating consistently strong performance in both semantic similarity and classification confidence are selected.

2.3. Mixture Synthesis

After obtaining the generated and filtered foreground sound samples, we use the Scaper toolkit [23], a soundscape synthesis library designed for creating labeled audio mixtures from foreground and background elements, to generate strongly labeled synthetic data. By default, we exclusively use the generated foreground sounds without mixing them with real foreground samples, while background and non-target sounds are sourced from real recordings [5], [24], [25]. Alternative settings, such as incorporating synthetic background sounds, are further explored in Section 4.3.1.

This process involves two stages of synthesis, foreground generation via the diffusion transformer and mixture creation via Scaper, which differs from the synthetic strong subset used in DCASE dataset [5], where real foreground sounds are employed. To reflect this distinction, we refer to the dataset generated by our method as the SynSonic Strong subset.

3. EXPERIMENTAL SETUPS

3.1. SED Dataset

We use the Domestic Environment Sound Event Detection (DESED) dataset [5] to evaluate the effectiveness of the proposed method. This dataset focuses on 10 domestic sound classes and provides a comprehensive collection of recordings for training and evaluation. The training set includes real weakly labeled data (1,578 clips), real strongly labeled data (3,470 clips) [26], synthetic strongly labeled data (10,000 clips), and real unlabeled data (14,412 clips). Additionally, a development set comprising 1,168 real clips is available for model development and validation. The SynSonic strong subset is used to augment the training data; specifically, we incorporate 10,000 samples generated using the method described in Section 2.

3.2. Implementation Details

We build the energy-envelope ControlNet on top of **EzAudio-L**² and train it using the AudioCaps dataset [27]. For the energy envelope, we compute waveform energy, ensuring that the window and hop sizes are aligned with those of the audio latent representation. The ControlNet

²<https://huggingface.co/OpenSound/EzAudio>

is trained with a batch size of 16, a learning rate of 1×10^{-5} , and for 10 epochs. During inference, we use a classifier-free guidance (CFG) scale of 3.5 and perform sampling with 50 diffusion steps.

For sample filtering, we employ the official checkpoints of Laion CLAP³ and Dasheng-AS⁴, the latter being based on Dasheng-base and fine-tuned on AudioSet. Parameter selections are detailed in Section 4.2.

For SED experiments, we adopt the FDY-CRNN model [2] and follow the official training protocols, including Mean Teacher [28], Mixup [8], and SpecAugment [7]. Most hyperparameters, such as audio sample rate, mel-spectrogram settings, optimizer, learning rate, and loss weighting, remain consistent with the official implementation⁵, except for the batch size, which is adjusted to accommodate the additional synthesized data. The default batch composition is described and discussed in Section 4.3.2.

3.3. Evaluation Metrics

For evaluation, we adopt the PSDS1 and PSDS2 metrics [29], [30] to assess model performance, where the former emphasizes temporal localization and the latter focuses on event class discrimination. Following the DCASE baseline⁶, we apply a fixed median filter with a window size of 7 to smooth the model predictions, without further tuning this parameter in our experiments. To ensure robust and reliable results, we report **the average PSDS scores across three trials** with different random seeds.

4. RESULTS AND DISCUSSION

4.1. Comparison of Foreground Generation Methods

We compare the ControlNet-based foreground generator with two alternative strategies: (1) simple T2A diffusion, as used in data augmentation for target sound extraction [13], and (2) audio-to-audio (A2A) diffusion, a simple method that aims to achieve controllable generation by conditioning on audio inputs. In the simple T2A diffusion approach, the generation length is determined based on reference samples, and silence at the beginning and end of the generated audio is trimmed to achieve the desired duration. However, because the generation process is not explicitly guided by temporal cues, the resulting samples often exhibit loose temporal structure, with unsmooth onsets and offsets as well as unintended pauses. In the A2A diffusion approach, inspired by the Img2Img framework [31], noise is added to the reference samples at a specific diffusion step (we use 75% of the total diffusion steps). This encourages the generation of audio that preserves the temporal structure of the reference, while partially retaining its frequency characteristics, as they are not fully disentangled — unlike in the energy-envelope ControlNet. For fair comparison, we apply the same filtering and mixture synthesis strategies to both alternative methods as used in the proposed approach. Furthermore, we provide a model trained without any additional data for reference.

As shown in Table 1, simple T2A diffusion yields a clear improvement in the PSDS2 score, which emphasizes classification accuracy over temporal localization. However, it does not significantly improve PSDS1, which focuses on timing precision. This reflects the limitation discussed above: T2A diffusion lacks strong control over the temporal structure of generated samples. In contrast, A2A generation achieves better improvement in PSDS1 by producing samples with the desired temporal structure. However, it offers limited gains in

³<https://huggingface.co/laion/clap-htsat-fused>

⁴<https://github.com/XiaoMi/dasheng>

⁵<https://github.com/frednam93/FDY-SED>

⁶https://github.com/DCASE-REPO/DESED_task/tree/master

Table 1: Comparison of different foreground augmentation methods.

Foreground Generator	PSDS1↑	PSDS2↑
None (Baseline)	0.4417	0.6639
T2A Diffusion	0.4507	0.6859
A2A Diffusion	0.4556	0.6799
Energy ControlNet	0.4641	0.6868

Table 2: Ablation study of sample filtering using different CLAP and AS weighting schemes.

Filter Ratio	CLAP Weight	AS Weight	PSDS1↑	PSDS2↑
Top 50%	1.0	0.0	0.4565	0.6750
Top 50%	0.7	0.3	0.4590	0.6748
Top 50%	0.5	0.5	0.4641	0.6868
Top 50%	0.3	0.7	0.4573	0.6874
Top 50%	0.0	1.0	0.4574	0.6784
None	/	/	0.4436	0.6690
Top 75%	0.5	0.5	0.4500	0.6776
Top 25%	0.5	0.5	0.4589	0.6802

PSDS2, as it tends to preserve certain frequency characteristics from the original audio, which reduces the diversity of generated samples. Compared to the two baselines, temporal ControlNet improves both PSDS1 and PSDS2, demonstrating its ability to balance precise temporal control and sample diversity, thereby enhancing model performance in both class discrimination and temporal localization.

4.2. Comparison of Sample Filtering Settings

To investigate the effectiveness of the proposed sample filtering strategy, we conduct experiments using different weightings between the CLAP score and the audio classification score, ranging from purely CLAP-based to purely classifier-based filtering. In addition, we examine the impact of the filtering ratio, varying from no filtering to extreme filtering, where only the top 25% of samples are retained.

As shown in Table 2, relying on a single filtering model leads to suboptimal performance in both PSDS1 and PSDS2. Specifically, using only CLAP (1.0/0.0) or only the audio classifier (0.0/1.0) results in noticeable degradation. In contrast, combining the two models improves performance, with equal weighting (0.5/0.5) achieving the best results. This suggests that integrating semantic relevance and audio classification confidence offers complementary benefits, enabling the selection of higher-quality, less noisy samples and thereby improving the SED model.

Meanwhile, we investigate the impact of the filtering threshold. Omitting the threshold results in almost no improvement over the baseline without newly synthesized data, highlighting the importance of data filtering. Furthermore, a lower threshold (top 75%) introduces noisy samples, while a higher threshold (top 25%) reduces data diversity, both leading to inferior improvements. By contrast, a threshold of 50% achieves the best performance by effectively balancing quality and diversity.

4.3. Discussion

In this section, we present additional configurations of the proposed augmentation strategy explored during method development.

4.3.1. Soundbank Generation and Mixture Synthesis: As shown in Table 3, we compare two alternative strategies: (1) including real foreground audio when creating new strongly labeled mixtures, and (2) adding generated background or non-target sounds using the same generation strategy as for foreground samples. Incorporating real foreground audio leads to performance degradation, as these foreground samples have already been used in the original synthetic

Table 3: Comparison of different foreground mixing strategies.

Mixing Strategy	PSDS1↑	PSDS2↑
Generated Foreground Only	0.4641	0.6868
+ Real Foreground	0.4484	0.6826
+ Generated Background	0.4613	0.6866

Table 4: Ablation study on training batch composition. The composition ratio indicates the proportion of: SynSonic strong / synthetic strong / real strong / real weak / real unlabeled samples in each batch.

Condition	Composition Ratio	PSDS1↑	PSDS2↑
No Gen. (Baseline)	6/6/0/6/12	0.4417	0.6639
+Gen. / -Real Strong	3/6/3/6/12	0.4498	0.6861
+Gen. / -Synth. Strong	6/3/3/6/12	0.4606	0.6877
+Gen. / -Real Weak	6/6/3/3/12	0.4477	0.6817
+Gen. / -Unlabeled	6/6/3/6/9	0.4519	0.6890
+Gen.	6/6/6/6/12	0.4641	0.6868

strong subset. Reusing them in newly synthesized mixtures increases the risk of overfitting, ultimately impairing the model’s generalization to unseen data during evaluation. Meanwhile, since the SED model focuses on detecting target foreground sounds and the real background soundbank is already large and diverse, adding generated background samples does not offer an additional performance benefit.

4.3.2. Batch Composition: As we introduce new strongly labeled subset during training, we explore different batch composition strategies to integrate it effectively. To better accommodate the new samples and understand their impact, we first reduce the proportion of each existing subset and fill the freed space with generated data. As shown in Table 4, incorporating generated samples significantly improves performance when replacing synthetic strongly labeled data, which tends to lack foreground diversity. In contrast, replacing real data, whether it is strongly labeled, weakly labeled, or unlabeled, yields similar improvement on PSDS2, but minor improvements in PSDS1, suggesting the essential role of real data and the necessity of retaining it during training. In the final configuration, we directly add the generated samples on top of the original batch without substantially reducing the real data proportion. This strategy, which slightly lowers the relative weight of other subsets, achieves the best PSDS1 performance by increasing the diversity of strongly labeled data while maintaining sufficient real data.

5. CONCLUSION AND FUTURE WORK

In this work, we explore SynSonic, a data augmentation strategy for sound event detection that harnesses the diffusion ControlNet’s ability to generate single-event sound samples with precise control over temporal structures, enabling the synthesis of strongly labeled mixtures. To ensure the quality of the generated samples, we investigate a joint rank filtering strategy designed to mitigate the selection bias inherent in single-model filtering. Furthermore, we offer practical guidelines for effective integration of SynSonic, including mixing strategies and training batch composition. Experimental results show that this approach improves both PSDS1 and PSDS2 scores, reflecting gains in event localization and class discrimination.

For future work, we plan to extend this approach to more complex sound event detection scenarios and explore the direct end-to-end generation of strongly labeled sound mixtures using generative models.

REFERENCES

- [1] A. Mesaros, T. Heittola, T. Virtanen, and M. D. Plumbley, "Sound event detection: A tutorial," *IEEE Signal Processing Magazine*, vol. 38, no. 5, pp. 67–83, 2021.
- [2] H. Nam, S.-H. Kim, B.-Y. Ko, and Y.-H. Park, "Frequency dynamic convolution: Frequency-adaptive pattern recognition for sound event detection," *arXiv preprint arXiv:2203.15296*, 2022.
- [3] K. Miyazaki, T. Komatsu, T. Hayashi, S. Watanabe, T. Toda, and K. Takeda, "Conformer-based sound event detection with semi-supervised learning and data augmentation," *dim*, vol. 1, no. 4, 2020.
- [4] R. Serizel, N. Turpault, A. Shah, and J. Salamon, "Sound event detection in synthetic domestic environments," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 86–90.
- [5] N. Turpault, R. Serizel, A. P. Shah, and J. Salamon, "Sound event detection in domestic environments with weakly labeled data and soundscape synthesis," in *Workshop on Detection and Classification of Acoustic Scenes and Events*, 2019.
- [6] W.-C. Lin, L. Bondi, and S. Ghaffarzadegan, "Background domain switch: A novel data augmentation technique for robust sound event detection," in *Interspeech 2023*, 2023, pp. 326–330.
- [7] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *Interspeech 2019*, p. 2613, 2019.
- [8] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017.
- [9] Q. Xie, Z. Dai, E. Hovy, T. Luong, and Q. Le, "Unsupervised data augmentation for consistency training," *Advances in neural information processing systems*, vol. 33, pp. 6256–6268, 2020.
- [10] W. Wu, T. Dai, X. Huang, F. Ma, and J. Xiao, "Image augmentation with controlled diffusion for weakly-supervised semantic segmentation," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 6175–6179.
- [11] B. Li, X. Xu, X. Wang, Y. Hou, Y. Feng, F. Wang, X. Zhang, Q. Zhu, and W. Che, "Semantic-guided generative image augmentation method with diffusion models for image classification," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 4, 2024, pp. 3018–3027.
- [12] B. Trabucco, K. Doherty, M. Gurinas, and R. Salakhutdinov, "Effective data augmentation with diffusion models," *arXiv preprint arXiv:2302.07944*, 2023.
- [13] 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.
- [14] S. Ghosh, S. Kumar, Z. Kong, R. Valle, B. Catanzaro, and D. Manocha, "Synthio: Augmenting small-scale audio classification datasets with synthetic data," *arXiv preprint arXiv:2410.02056*, 2024.
- [15] 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.
- [16] J. Hai, Y. Xu, H. Zhang, C. Li, H. Wang, M. Elhilali, and D. Yu, "Ezaudio: Enhancing text-to-audio generation with efficient diffusion transformer," *arXiv preprint arXiv:2409.10819*, 2024.
- [17] L. Zhang, A. Rao, and M. Agrawala, "Adding conditional control to text-to-image diffusion models," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 3836–3847.
- [18] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [19] H. Dinkel, Z. Yan, Y. Wang, J. Zhang, Y. Wang, and B. Wang, "Scaling up masked audio encoder learning for general audio classification," *arXiv preprint arXiv:2406.06992*, 2024.
- [20] S.-L. Wu, C. Donahue, S. Watanabe, and N. J. Bryan, "Music control-net: Multiple time-varying controls for music generation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 2692–2703, 2024.
- [21] H. F. García, O. Nieto, J. Salamon, B. Pardo, and P. Seetharaman, "Sketch2sound: Controllable audio generation via time-varying signals and sonic imitations," in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
- [22] Z. Kong, S.-g. Lee, D. Ghosal, N. Majumder, A. Mehrish, R. Valle, S. Poria, and B. Catanzaro, "Improving text-to-audio models with synthetic captions," *arXiv preprint arXiv:2406.15487*, 2024.
- [23] J. Salamon, D. MacConnell, M. Cartwright, P. Li, and J. P. Bello, "Scaper: A library for soundscape synthesis and augmentation," in *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2017, pp. 344–348.
- [24] E. Fonseca, J. Pons, X. Favory, F. Font, D. Bogdanov, A. Ferraro, S. Oramas, A. Porter, and X. Serra, "Freesound datasets: A platform for the creation of open audio datasets," in *ISMIR*, 2017, pp. 486–493.
- [25] A. Mesaros, T. Heittola, and T. Virtanen, "A multi-device dataset for urban acoustic scene classification," *arXiv preprint arXiv:1807.09840*, 2018.
- [26] S. Hershey, D. P. Ellis, E. Fonseca, A. Jansen, C. Liu, R. C. Moore, and M. Plakal, "The benefit of temporally-strong labels in audio event classification," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 366–370.
- [27] C. D. Kim, B. Kim, H. Lee, and G. Kim, "Audiocaps: Generating captions for audios in the wild," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 119–132.
- [28] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," *Advances in neural information processing systems*, vol. 30, 2017.
- [29] Ç. Bilen, G. Ferroni, F. Tuveri, J. Azcarreta, and S. Krstulović, "A framework for the robust evaluation of sound event detection," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 61–65.
- [30] A. Mesaros, T. Heittola, and T. Virtanen, "Metrics for polyphonic sound event detection," *Applied Sciences*, vol. 6, no. 6, p. 162, 2016.
- [31] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon, "Sdedit: Guided image synthesis and editing with stochastic differential equations," *arXiv preprint arXiv:2108.01073*, 2021.