

GENERATION OR REPLICATION: AUSCULTATING AUDIO LATENT DIFFUSION MODELS

Bralios, Dimitrios; Wichern, Gordon; Germain, François G; Pan, Zexu; Khurana, Sameer; Hori, Chiori; Le Roux, Jonathan

TR2024-027 March 19, 2024

Abstract

The introduction of audio latent diffusion models possessing the ability to generate realistic sound clips on demand from a text description has the potential to revolutionize how we work with audio. In this work, we make an initial attempt at understanding the inner workings of audio latent diffusion models by investigating how their audio outputs compare with the training data, similar to how a doctor auscultates a patient by listening to the sounds of their organs. Using text-to-audio latent diffusion models trained on the AudioCaps dataset, we systematically analyze memorization behavior as a function of training set size. We also evaluate different retrieval metrics for evidence of training data memorization, finding the similarity between mel spectrograms to be more robust in detecting matches than learned embedding vectors. In the process of analyzing memorization in audio latent diffusion models, we also discover a large amount of duplicated audio clips within the AudioCaps database.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
2024*

GENERATION OR REPLICATION: AUSCULTATING AUDIO LATENT DIFFUSION MODELS

Dimitrios Bralios^{1,2}, Gordon Wichern¹, François G. Germain¹, Zexu Pan¹,
Sameer Khurana¹, Chiori Hori¹, Jonathan Le Roux¹

¹Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

²University of Illinois Urbana-Champaign, Urbana, IL, USA

ABSTRACT

The introduction of audio latent diffusion models possessing the ability to generate realistic sound clips on demand from a text description has the potential to revolutionize how we work with audio. In this work, we make an initial attempt at understanding the inner workings of audio latent diffusion models by investigating how their audio outputs compare with the training data, similar to how a doctor auscultates a patient by listening to the sounds of their organs. Using text-to-audio latent diffusion models trained on the AudioCaps dataset, we systematically analyze memorization behavior as a function of training set size. We also evaluate different retrieval metrics for evidence of training data memorization, finding the similarity between mel spectrograms to be more robust in detecting matches than learned embedding vectors. In the process of analyzing memorization in audio latent diffusion models, we also discover a large amount of duplicated audio clips within the AudioCaps database.

Index Terms— audio latent diffusion model, audio synthesis, memorization, acoustic similarity

1. INTRODUCTION

Diffusion models [1,2] have quickly become a powerful class of generative models, and are based on an elegant theoretical formulation that only requires a simple denoising network to iteratively generate complex data such as images and sound. Perhaps their most successful application is in training large scale models that generate images from text captions [3]. Text-to-image models based on latent diffusion [4], which perform the iterative diffusion process in a low-dimensional latent space, provide a good trade-off in terms of generation quality and computational cost. However, because text-to-image models are so easy to use, thanks to their control by natural language, and produce realistic images, concerns over the intellectual property rights of the data used to train these models has recently begun to emerge.

While many legal and ethical issues remain open, research attempting to answer the technical question of how diffusion models may memorize and/or copy their training data has begun to appear for the case of images [5–7]. While precisely defining replication in generative models is difficult, evidence that deep networks memorize training data has emerged through work related to the generalization and overfitting of deep neural networks [8, 9], along with research on membership inference attacks [10, 11], which determine whether a particular sample was part of a model’s training set. At the same time, we note that another somewhat counterintuitive by-product of that research is duplicate detection, as recent work found duplicated samples in the training set to be more likely to be replicated [12].

This work was performed while D. Bralios was an intern at MERL.

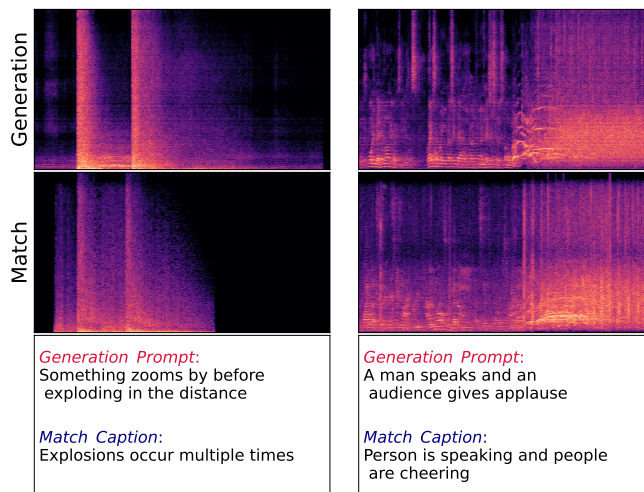


Fig. 1. Samples generated using TANGO [13] and their top matches in the AudioCaps [14] training set along with the corresponding prompts and captions. Found using our proposed method employing mel descriptors. Notice that despite not being identical to the generations, the matches contain semantically identical acoustic events that temporally coincide.

Diffusion models have also recently found many uses in audio generation for speech [15, 16], music [17, 18], and general sounds [13, 19–21]. The audio latent diffusion model (AudioLDM) [20], which generates mel spectrograms from a text description and then resynthesizes audio waveforms using a separately trained HiFi-GAN network [22], is a particularly promising approach. In a recent public challenge on Foley sound synthesis [23], a technique based on AudioLDM ranked first in terms of subjective quality. While recent works on text-to-music generation have uncovered evidence of training data memorization for transformer-based [24] models, and the novelty of generated samples from a diffusion-based model [18], they do not systematically analyze why memorization occurs, or compare the ability of different signal representations to detect replicated training data.

In this work, we make an attempt at quantifying training data replication for audio latent diffusion models trained on general sounds. We define a generated sound file as replicating training data when it contains nearly-identical complex spectro-temporal patterns. Certain sounds, such as stationary noise and constant tones, are easily replicated but excluded by this definition since they do not contain such complex patterns. While the term “nearly-identical”

lacks precision, we believe it is appropriate for this exploratory work, and we propose techniques based on both mel spectrograms and contrastive language audio pre-training (CLAP) [25, 26] descriptors for finding these potential replications. Our experiments utilize TANGO [13], a text-to-audio model descended from AudioLDM but using a different text encoder and with its diffusion core trained on AudioCaps [14] dataset. We find a small, but not insignificant number of replicated training samples (see Fig. 1 for examples), and validate our ability to find training data replications by decreasing the size of the training set to a small enough level that memorization is very likely to occur. We also discover a large amount of duplicated audio clips within the AudioCaps database. Audio examples of potentially replicated training data and identified AudioCaps duplicates are available on our demo page ¹.

2. SEARCHING FOR REPLICATIONS

In order to detect generations which potentially are replications of training samples, we frame the problem as a copy detection problem. Specifically, the generated samples are regarded as queries and the training samples as a reference set. For each pair of a query q and a reference r , a similarity score is computed, resulting in a pairwise similarity matrix. We can then follow a two stage semi-manual approach consisting of a *retrieval* stage and a *verification* stage. In the retrieval stage, we retrieve the queries q whose top-1 match similarity score is above a certain threshold τ . In the verification stage, we manually examine each retrieved query and reference pair.

The similarity score $s(q, r)$ between query q and reference r is computed based on a similarity score between descriptors $\mathbf{q}, \mathbf{r} \in \mathbb{R}^d$ extracted from q and r , in our case their cosine similarity:

$$s(q, r) = \text{sim}(\mathbf{q}, \mathbf{r}) = \frac{\mathbf{q}^\top \mathbf{r}}{\|\mathbf{q}\| \|\mathbf{r}\|}. \quad (1)$$

We investigate the use of a low-dimensional log mel spectrogram as a sample descriptor, a natural choice since most audio models employing latent diffusion extract the latent representation from a mel one. We also experiment with a CLAP [25] descriptor.

Additionally, since we threshold based on the similarity score in the retrieval stage, it is important to have comparable similarity scores across queries. Depending on the metric, some sounds may be inherently more similar to all other sounds, thus, we employ similarity normalization [27, 28]. For each query q , we discount its similarity with each reference r using a bias term based on the average similarity between q and its K nearest neighbors in a background set of other samples, resulting in the normalized similarity score:

$$\hat{s}(q, r) = s(q, r) - \beta \cdot \text{bias}(q), \quad (2)$$

$$\text{bias}(q) = \frac{1}{K} \sum_{k=1}^K s(q, b_k),$$

where b_k is the k -th nearest neighbour of q in the background set (based on the similarity of their descriptors), and β a scalar. Therefore, in the retrieval stage, we retrieve the following set of queries:

$$\{q : \max_r \hat{s}(q, r) \geq \tau\}, \quad (3)$$

which are then inspected along with their top-1 matches in the verification stage, based on the definition of Section 1.

3. EXPERIMENTAL FRAMEWORK

The model we focus on and experiment with is Tango [13], firstly because its code (including training scripts) and model checkpoints are available online. Secondly, it maintains near state-of-the-art performance while being trained on AudioCaps [14], which is a significantly smaller training set compared to similar models, containing around 45k audio and text pairs, thus making the search for replications more tractable. We note however that, despite the small training set size, level-weighted mixing augmentation of two sounds with concatenated captions is performed during training. Every sample generation is performed using 200 diffusion steps and classifier-free guidance [29] with a scale of 3.

All sounds have a sampling frequency of 16 kHz and are zero-padded to 10.242 s to match the length of the generated samples. The mel spectrogram is computed using 16 mel bins and a 128 ms long Hann window with 25% overlap. It is then normalized by dividing it by its maximum value, converted to decibel, and clipped with a lower value of -40 dB. By flattening, we get a 1712-dimensional descriptor. In order to compute CLAP [25] descriptors, we use the publicly available `laion/clap-htsat-unfused` model checkpoint on Hugging Face, resulting in 512-dimensional descriptors. For the calculation of the normalized similarity score, we use a background set consisting of 1000 samples randomly selected from the balanced train segment of the AudioSet dataset [30], which has no overlap with AudioCaps. We set $\beta = 0.5$ and use the top-5 nearest neighbors from the background set, i.e., $K = 5$.

4. RESULTS & DISCUSSION

4.1. Comparison of Descriptors

First, we compare the mel and CLAP descriptors in terms of detecting replicated samples. Thus, we train two additional Tango models on small, random subsets of the AudioCaps training set, while still performing the mixing augmentation. The first model was trained on 1000 audio and text pairs, and the second on 5000. Each model is trained for 62200 steps, the maximum number of training steps of the pre-trained checkpoint. With each model, we generate using the same training set prompts, resulting in as many generated samples as the size of each training set. We observe extreme overfitting, in terms of difference between the validation and training loss in the case of the first model, and slight overfitting in the case of the second.

As expected, we observe widespread replication in the generations of the first model, with over 90% of them being almost identical replications of training samples. Two examples of memorized generations are shown in the top of Fig. 2, together with their matches successfully identified by both mel and CLAP methods. Memorization can also be seen in the first row of Fig. 3, where we observe clear separation of the distributions of the top-1 normalized similarity scores of the generations to the training set, and of the training set to itself. We note that the mel method achieves better separation, which can translate into better precision in replication detection. Based on these histograms, we select the threshold value $\tau_{\text{mel}} = 0.5005$ for the following analyses of the two other models.

For the second model trained on 5000 training samples, we follow the methodology outlined in Section 2 using mel and CLAP descriptors. We set $\tau_{\text{mel}} = 0.5005$, resulting in 178 pairs, and we appropriately set τ_{CLAP} to get the same number of pairs. Out of those, following manual verification using the definition from Section 1, just 31 out of 178 are replicated in the case of mel descriptors and 28 out of 178 in the case of CLAP, with only 2 being in both sets.

¹<https://www.merl.com/demos/auscultating-diffusion>

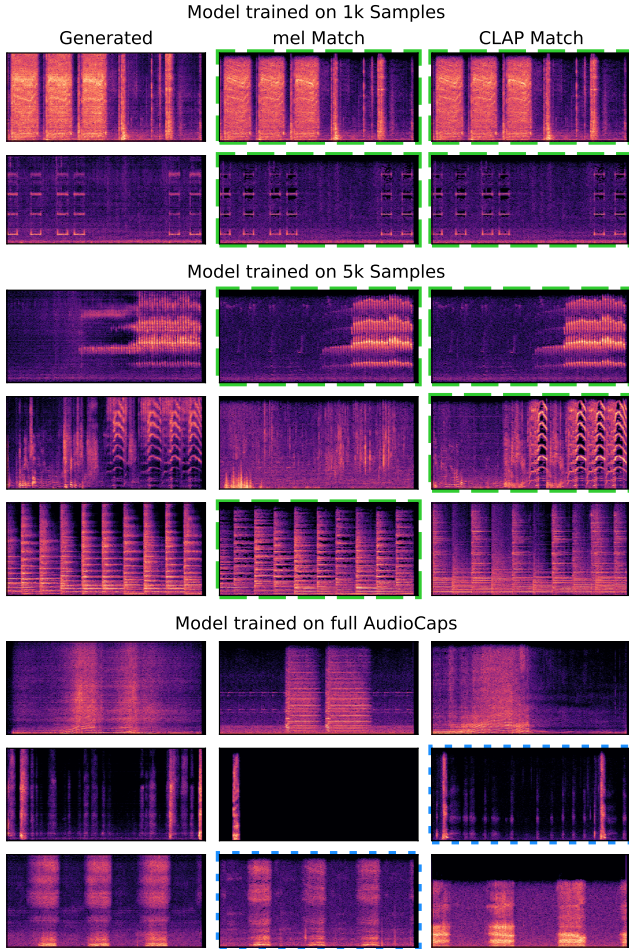


Fig. 2. Spectrograms of generated samples (left column) along with their top-1 matches using mel descriptors (middle column) and CLAP descriptors (right column) for each model. Green dashed outlines indicate successful match detection, while the blue dotted outlines indicate detection of a potential match.

We find that mel descriptors tend to pay attention to similar spectro-temporal energy presence, while CLAP descriptors focus mostly on similar semantic content. In Fig. 2 (middle), we see a failure case for each method. Also, in contrast to the previous model, replication pairs here are not always exact copies, but contain certain complex spectro-temporal patterns that are nearly identical. Finally, Fig. 3 (middle) confirms that replication is limited in comparison to the previous case, with the mel histograms more closely corresponding to our empirical observations (i.e., few replicated samples).

4.2. Investigation of Replication

We now shift our attention to the model trained on the whole AudioCaps dataset. As before, we generate around 45k samples using the text prompts of the training set and we then search for potential matches following the methodology described in Section 2. We use the threshold $\tau_{\text{mel}} = 0.5005$ that retrieves 2278 samples and the appropriate τ_{CLAP} that results in the same number of pairs. After manually inspecting those samples, the number of samples that po-

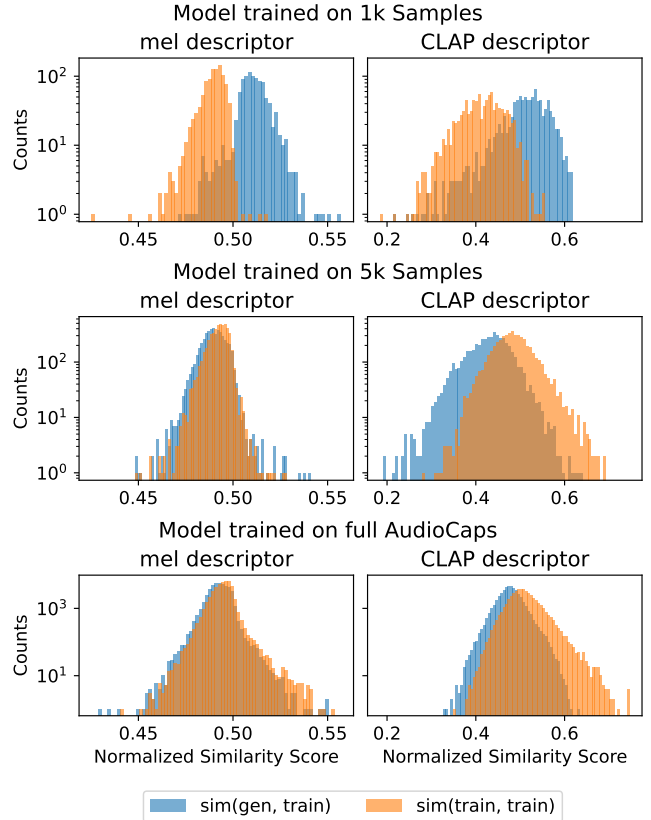


Fig. 3. Histograms of the top-1 normalized similarity scores, using both descriptors, between the generated samples and the corresponding training samples (blue), as well as between the training samples and themselves (orange) essentially describing the self-similarity of the training set. The vertical axis is in logarithmic scale.

tentially fit our replication definition is quite limited, around a dozen for each method with mel descriptors detecting slightly more.

We show in Fig. 2 (bottom) a few generations and their closest mel and CLAP matches, demonstrating how both methods operate. In the first case, both methods do not return a match, with the mel method returning a sample with similar spectro-temporal distribution of energy and the CLAP method returning a semantically similar sample. In the second case, the CLAP method detects a training set sample containing a sneeze sound that temporally coincides with the first one in the generation, while the mel method returns a burp sound occurring at that same timestep. In the third and final case, we have a snoring sound, where the mel method finds a very close albeit not identical sample, while the CLAP method returns a semantically similar sample also consisting of snoring.

In Fig. 4, we showcase some detected potential replication matches using the mel and CLAP methods. In the case of mel, we have very similar spectro-temporal patterns coinciding, especially in the temporal dimension. In the case of CLAP, we have semantic matches whose acoustic components are very similarly structured. Minor variations between the matches could be attributed to parts of the models (e.g., the vocoder) that are pre-trained on other datasets. Additionally, the presence of additional acoustic events in the generations might be a result of the augmentation used during training.

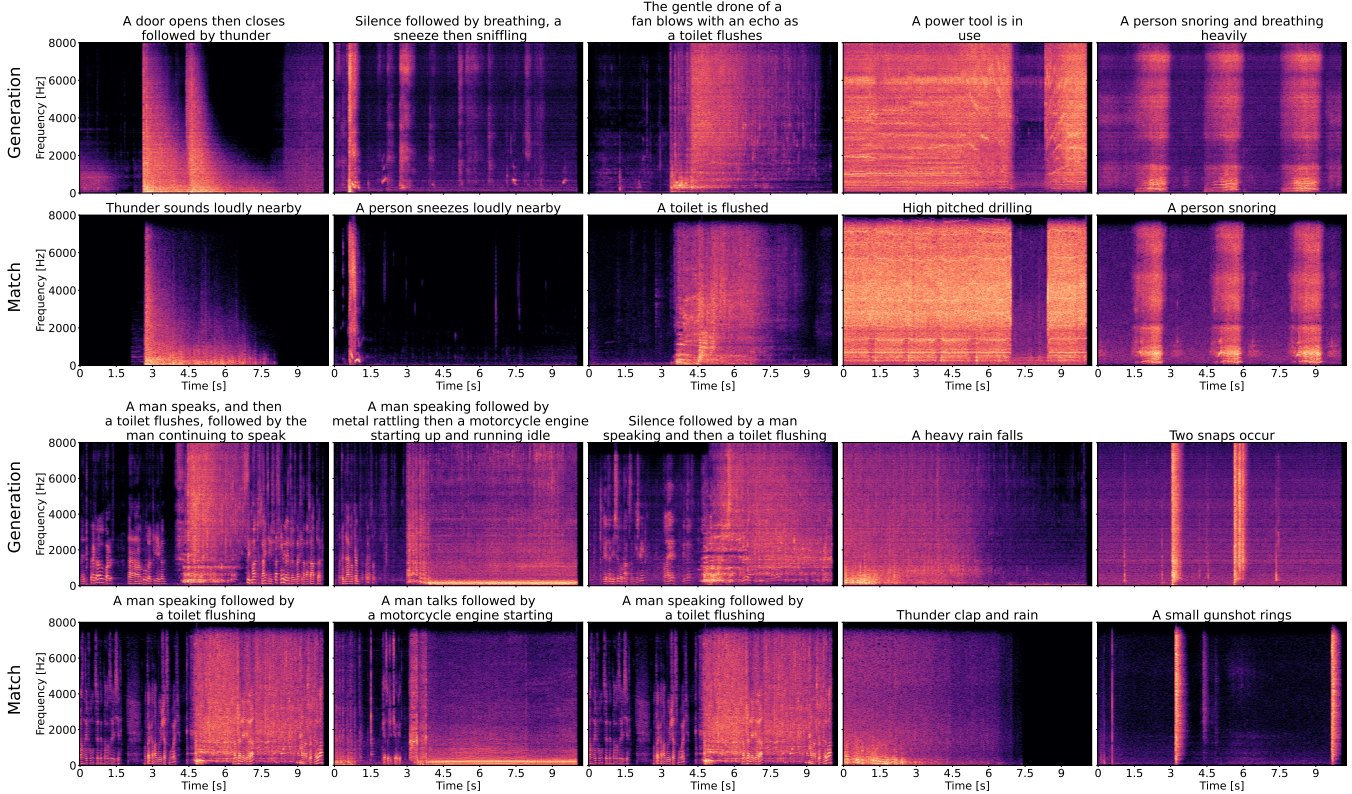


Fig. 4. Matches found using mel descriptors in the top two rows. Matches found using CLAP descriptors in the bottom two rows.

4.3. Duplication in the AudioCaps Training Set

The presence of training samples that form a tail of high self-similarity in the histogram in Fig. 3 (bottom row) prompts us to investigate duplication in the AudioCaps dataset. First, we compute the pairwise normalized self-similarity matrix S between the audio data of the AudioCaps training set, using mel descriptors. Second, we zero out the diagonal of S . Next, we compute the binary matrix $A_{ij} = \mathbb{1}[S_{ij} > \tau \text{ and } S_{ji} > \tau]$ where $\tau = 0.5025$ and $\mathbb{1}[\cdot]$ represents the indicator function. The resulting symmetric matrix A can be viewed as a graph adjacency matrix and we can find its connected components. Finally, we manually examine each found connected component and select ones that correspond to clusters of duplicated samples in the training set.

We find 276 connected components where 88 of them are indeed duplicated clusters containing 257 samples in total. We present these findings in Fig. 5. Examples include audio files extracted from different YouTube videos containing the same intro audio effects. However, despite the presence of duplicate audio files in the training set, we did not find generations replicating them.

5. CONCLUSIONS AND FUTURE WORK

We performed an initial analysis of training data replication in TANGO, a recently proposed state-of-the-art text-to-audio latent diffusion model. While we only manually verified a small number of replications on the model trained on the full dataset (approximately 5-10 per every 10,000 training samples), we have provided some evidence of training data memorization. We believe studies such as this




Fig. 5. Clusters of audio duplicates found in the AudioCaps training set.

one are important to help technically inform many of the outstanding legal issues related to generative modeling. While we compared mel and CLAP descriptors for finding replications, there are countless additional techniques to be explored, e.g., mel+CLAP ensembles, audio fingerprinting [31, 32], self-supervised learning [33], etc. We also surprisingly found duplicate sound files in AudioCaps, a curated subset of AudioSet. It has recently been observed that de-duplicated training data can improve performance in language models [12, 34], and a similar analysis for audio is an interesting area for further study.

6. REFERENCES

- [1] Y. Song and S. Ermon, “Generative modeling by estimating gradients of the data distribution,” in *Proc. NeurIPS*, 2019.
- [2] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” in *Proc. NeurIPS*, 2020.
- [3] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, “Hierarchical text-conditional image generation with CLIP latents,” *arXiv preprint arXiv:2204.06125*, 2022.
- [4] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proc. CVPR*, 2022, pp. 10 684–10 695.
- [5] N. Carlini, J. Hayes, M. Nasr, M. Jagielski, V. Shwag, F. Tramer, B. Balle, D. Ippolito, and E. Wallace, “Extracting training data from diffusion models,” in *Proc. USENIX Secur. Symp.*, 2023.
- [6] G. Somepalli, V. Singla, M. Goldblum, J. Geiping, and T. Goldstein, “Diffusion art or digital forgery? Investigating data replication in diffusion models,” in *Proc. CVPR*, 2023, pp. 6048–6058.
- [7] —, “Understanding and mitigating copying in diffusion models,” in *Proc. NeurIPS*, 2023.
- [8] V. Feldman and C. Zhang, “What neural networks memorize and why: Discovering the long tail via influence estimation,” in *Proc. NeurIPS*, 2020.
- [9] N. Carlini, M. Jagielski, C. Zhang, N. Papernot, A. Terzis, and F. Tramer, “The privacy onion effect: Memorization is relative,” in *Proc. NeurIPS*, 2022.
- [10] N. Carlini, S. Chien, M. Nasr, S. Song, A. Terzis, and F. Tramer, “Membership inference attacks from first principles,” in *Proc. SP*, 2022, pp. 1897–1914.
- [11] W.-C. Tseng, W.-T. Kao, and H.-y. Lee, “Membership inference attacks against self-supervised speech models,” in *Proc. Interspeech*, 2021, pp. 5040–5044.
- [12] N. Kandpal, E. Wallace, and C. Raffel, “Deduplicating training data mitigates privacy risks in language models,” in *Proc. ICML*, 2022, pp. 10 697–10 707.
- [13] D. Ghosal, N. Majumder, A. Mehrish, and S. Poria, “Text-to-audio generation using instruction tuned LLM and latent diffusion model,” *arXiv preprint arXiv:2304.13731*, 2023.
- [14] C. D. Kim, B. Kim, H. Lee, and G. Kim, “AudioCaps: Generating captions for audios in the wild,” in *Proc. NAACL-HLT*, 2019, pp. 119–132.
- [15] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, “DiffWave: A versatile diffusion model for audio synthesis,” in *Proc. ICLR*, 2020.
- [16] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov, “Grad-TTS: A diffusion probabilistic model for text-to-speech,” in *Proc. ICML*, 2021, pp. 8599–8608.
- [17] Q. Huang, D. S. Park, T. Wang, T. I. Denk, A. Ly, N. Chen, Z. Zhang, Z. Zhang, J. Yu, C. Frank *et al.*, “Noise2Music: Text-conditioned music generation with diffusion models,” *arXiv preprint arXiv:2302.03917*, 2023.
- [18] K. Chen, Y. Wu, H. Liu, M. Nezhurina, T. Berg-Kirkpatrick, and S. Dubnov, “MusicLDM: Enhancing novelty in text-to-music generation using beat-synchronous mixup strategies,” *arXiv preprint arXiv:2308.01546*, 2023.
- [19] D. Yang, J. Yu, H. Wang, W. Wang, C. Weng, Y. Zou, and D. Yu, “DiffSound: Discrete diffusion model for text-to-sound generation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 2023.
- [20] 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 *Proc. ICML*, 2023.
- [21] R. Huang, J. Huang, D. Yang, Y. Ren, L. Liu, M. Li, Z. Ye, J. Liu, X. Yin, and Z. Zhao, “Make-An-Audio: Text-to-audio generation with prompt-enhanced diffusion models,” in *Proc. ICML*, 2023.
- [22] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” in *Proc. NeurIPS*, 2020.
- [23] 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,” *arXiv preprint arXiv:2304.12521*, 2023.
- [24] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzett, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi *et al.*, “MusicLM: Generating music from text,” *arXiv preprint arXiv:2301.11325*, 2023.
- [25] 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 *Proc. ICASSP*, 2023, pp. 1–5.
- [26] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, “Clap learning audio concepts from natural language supervision,” in *Proc. ICASSP*, 2023, pp. 1–5.
- [27] E. Pizzi, S. D. Roy, S. N. Ravindra, P. Goyal, and M. Douze, “A self-supervised descriptor for image copy detection,” in *Proc. CVPR*, 2022, pp. 14 532–14 542.
- [28] M. Douze, G. Tolias, E. Pizzi, Z. Papakipos, L. Chanussot, F. Radenovic, T. Jenicek, M. Maximov, L. Leal-Taixé, I. Elezi *et al.*, “The 2021 image similarity dataset and challenge,” *arXiv preprint arXiv:2106.09672*, 2021.
- [29] J. Ho and T. Salimans, “Classifier-free diffusion guidance,” *arXiv preprint arXiv:2207.12598*, 2022.
- [30] 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 *Proc. ICASSP*, 2017, pp. 776–780.
- [31] P. Cano, E. Batle, T. Kalker, and J. Haitsma, “A review of algorithms for audio fingerprinting,” in *Proc. MMSP*, 2002, pp. 169–173.
- [32] S. Chang, D. Lee, J. Park, H. Lim, K. Lee, K. Ko, and Y. Han, “Neural audio fingerprint for high-specific audio retrieval based on contrastive learning,” in *Proc. ICASSP*, 2021, pp. 3025–3029.
- [33] Y. Gong, C.-I. Lai, Y.-A. Chung, and J. Glass, “SSAST: Self-supervised audio spectrogram transformer,” in *Proc. AAAI*, vol. 36, no. 10, 2022, pp. 10 699–10 709.
- [34] K. Lee, D. Ippolito, A. Nystrom, C. Zhang, D. Eck, C. Callison-Burch, and N. Carlini, “Deduplicating training data makes language models better,” in *Proc. ACL*, 2022, pp. 8424–8445.