
Images that Sound: Composing Images and Sounds on a Single Canvas

Ziyang Chen Daniel Geng Andrew Owens

University of Michigan

<https://ificl.github.io/images-that-sound/>




Figure 1: **Images that sound.** We use diffusion models to generate visual spectrograms (second row) that look like natural images, which we call *images that sound*. These spectrograms can be converted into natural sounds (third row) using a pretrained vocoder, or colored to obtain more visually pleasing results (first row). Please refer to our [website](https://ificl.github.io/images-that-sound/) to listen to the sounds.

Abstract

Spectrograms are 2D representations of sound that look very different from the images found in our visual world. And natural images, when played as spectrograms, make unnatural sounds. In this paper, we show that it is possible to synthesize spectrograms that simultaneously look like natural images and sound like natural audio. We call these visual spectrograms *images that sound*. Our approach is simple and zero-shot, and it leverages pre-trained text-to-image and text-to-spectrogram diffusion models that operate in a shared latent space. During the reverse process, we denoise noisy latents with both the audio and image diffusion models in parallel, resulting in a sample that is likely under both models. Through quantitative evaluations and perceptual studies, we find that our method successfully generates spectrograms that align with a desired audio prompt while also taking the visual appearance of a desired image prompt. Please see our project page for video results: <https://ificl.github.io/images-that-sound/>

1 Introduction

The spectrogram is a ubiquitous low-dimensional representation for audio machine learning that plots the energy within different frequencies over time. But it is also widely used as a tool for converting




Figure 2: **Images vs. spectrograms.** We show grayscale images generated from Stable Diffusion [96] on the left, followed by log-mel spectrograms generated from Auffusion [118] in the middle, and our generated *images that sound* results on the right.

sound into a visual form that can be—at least partially—perceived by sight. For example, in this representation (Fig. 2), event onsets look to a human observer like lines, and speech looks like a sequence of waves and bands. This insight is commonly used within the audio community, which frequently repurposes pretrained visual networks for audio tasks, often with only relatively minor modifications [48, 90, 69, 68, 34, 118, 112].

We hypothesize that the success of spectrograms in these roles is due in part to the fact that they share many statistical properties with the distribution of natural images, providing visual structures like edges and textures that the human visual system can readily process. Given the statistical similarities between images and sounds, we ask whether it is possible to automatically generate examples that lie at the *intersection* of both modalities. We create *images that sound* (Fig. 1), 2D matrices that *look* semantically meaningful when viewed as images, but that also *sound* meaningful when played as a spectrogram. This generative modeling problem is challenging, because it requires modeling a distribution that is induced by two very different data sources, and no relevant paired data is available.

We are motivated by the “spectrogram art” that has been made by a variety of artists [10], most famously by musicians Aphex Twin [2], Venetian Snares [113] and Nine Inch Nails [85]. These artists manipulate their songs to display a desired image when they are visualized as spectrograms, such as by showing the artist’s face or album art. In current practice, there is a steep trade-off between the quality of the image and the sound, since it is difficult to simultaneously control the interpretation of a signal in both modalities. As a result, existing artwork often comes across to the listener as dissonant or as random noise, rather than as natural sounds.¹ By contrast, we aim to generate signals that are as natural as possible in both modalities, such as towers that simultaneously sound like ringing bells or images of tigers that make a roaring sound (Fig. 1).

In this work, we pose this problem as a multimodal compositional generation task and propose a simple, zero-shot method that composes off-the-shelf text-to-spectrogram and text-to-image diffusion models from different modalities. Inspired by prior work on compositionality in diffusion models [72, 29, 42, 41], we denoise using both a noise estimate from the spectrogram model and a noise estimate from the image model. This is possible because these two models perform diffusion in the same latent space. The result is a sample that is simultaneously likely under the (text-conditional) distribution of spectrograms and images. The spectrograms are then converted to waveforms using a pretrained vocoder. In addition, we show that these black-and-white images may be colorized, resulting in color images whose grayscale versions can be played as spectrograms.

Surprisingly, we find that off-the-shelf diffusion models *trained on different modalities* can be composed together to obtain samples that function as both an image and a sound. Often these examples reuse visual elements in unexpected ways (e.g., in Fig. 1, a line is both the onset of a bell chime and the contour of a bell tower). We provide qualitative results, as well as quantitative comparisons and human study results against baselines, indicating that our method produces spectrograms that better align with both the audio and image prompts. Our contributions are summarized as follows:

- We propose *images that sound*, a type of multimodal art that can be both understood as an image or played as a sound.
- We show that we can compose pretrained diffusion models from different modalities in a zero-shot fashion to produce examples at the intersection of image and spectrogram distributions.

¹We encourage the reader to listen to [popular examples](#) of spectrogram art [10].

- We propose alternative methods for generating *images that sound*, one based on score distillation sampling [11, 93] and another based on simply subtracting an image from a spectrogram.
- We find through qualitative and quantitative experiments that our method outperforms baseline approaches and generates high-quality samples.

2 Related Work

Diffusion models. Diffusion models [102, 54, 106, 27, 104] are a class of generative models that learn to reverse a forward process that iteratively corrupts data. Typically, this forward process adds Gaussian noise and the reverse process learns to denoise the data by predicting the added noise. Diffusion models have a variety of applications, including text-conditioned image generation [27, 96, 84, 26, 98], video generation [53, 56, 101, 6, 45, 115], image and video editing [94, 97, 79, 49, 9, 31, 40], audio generation [118, 70, 71, 32, 43, 76], 3D generation [74, 57, 12, 73, 8, 38], and camera pose estimation [121]. In this work, we use Stable Diffusion [96], a latent diffusion model trained for text-conditioned image generation, as well as Auffusion [118], a text-conditioned audio generation model trained to produce log-mel spectrograms. Auffusion is finetuned from Stable Diffusion, similar to Riffusion [34], and as a result, the two methods share a latent space. This is crucial for our technique, which jointly diffuses these shared latents.

Compositional generation. One property of diffusion models is that they admit a straightforward technique to compose concepts by summing noise estimates. This may be understood by viewing noise estimates as gradients of a conditional data distribution [105, 106], and the sum of these gradients as pointing in the direction that maximizes multiple conditional likelihoods. This approach has been applied to enable compositions of text prompts globally [72], spatially [7, 29], transformations of images [42], and image components [41]. We go beyond these works by showing that diffusion models *from two different modalities* can successfully be composed together.

Audio-visual learning. A variety of works have learned cross-modal associations between vision and sound. Some approaches establish *semantic correspondence*, *i.e.*, which sounds and visuals are commonly associated with one another [3, 108]. Previous work has used this cue to learn cross-modal representations [5, 83, 80, 46, 44, 69, 68] and audio-visual sound localization [4, 58, 81, 59, 100, 91, 75]. Some researchers focus on the *temporal correspondence* between audio and visual streams [64, 87, 33, 16, 107, 62] to study source separation [122, 1, 37], Foley sound generation [61, 28, 117, 78], and action recognition [39, 60, 86]. Others also explore the *spatial correspondence* between them [21, 35, 120, 19, 22, 77], including spatial sound generation [36, 82, 14, 67] and audio-visual acoustic learning [13, 103, 24, 18, 20]. Differing from the works above, our focus is to explore the intersection of the distributions between spectrograms and images, where we create spectrograms that can be understood as visual images and can also be played as sounds.

Audio steganography. Audio steganography is the practice of concealing information within an audio signal. Artists have explored it for creative expression [111, 110]. Aphex Twin embedded a visual of his face in the audio waveform of the track “Formula” [2]. Noam Oxman creates animal portraits made of musical notations [88]. Other work has proposed deep learning methods for steganography, such as hiding video content inside audio files with invertible generative models [119], hiding audio data inside an identity image [123], and audio watermarking [15, 89, 99]. Our approach can be viewed as a steganography method that hides an image within an audio track, and is only revealed when the track is converted to a spectrogram.

3 Method

Our goal is to generate spectrograms that simultaneously represent both a sound and an image, each of which is specified by a text prompt. When the spectrogram is converted into a waveform, the sound matches the audio prompt, while when it is visually inspected, it should take the appearance of a given visual prompt (Fig. 1). To do this, we sample from the joint distribution of images and spectrograms, using off-the-shelf diffusion models trained on each modality independently.




Figure 3: **Composing audio and visual diffusion models.** We generate the visual spectrogram that can be visualized as an image or played as a sound. Given a noisy latent $\mathbf{z}_t$, we apply visual and audio diffusion models, each guided by a text prompt, to compute noise estimates $\epsilon_v^{(t)}$ and $\epsilon_a^{(t)}$ respectively. We obtain the multimodal noise estimate $\tilde{\epsilon}^{(t)}$ by a weighted average, then use it as part of the iterative denoising process. Finally, we decode the clean latent $\mathbf{z}_0$ to a spectrogram and convert it into a waveform using a pretrained vocoder (or by Griffin-Lim [47]).

3.1 Preliminaries

Diffusion models. Diffusion models [54, 106] iteratively denoise standard Gaussian noise, $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, to generate clean samples, $\mathbf{x}_0$, from some learned data distribution. At timestep t in the reverse diffusion process, the noise predictor, ϵ_θ , takes the intermediate noisy sample, $\mathbf{x}_t$, and the condition y , such as a text prompt embedding, to estimate the noise $\epsilon_\theta(\mathbf{x}_t; y, t)$. Following DDIM [104], we obtain the next, less noisy, sample $\mathbf{x}_{t-1}$ at the previous timestep via:

$$\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \cdot \hat{\epsilon}_\theta(\mathbf{x}_t; y, t)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \hat{\epsilon}_\theta(\mathbf{x}_t; y, t) + \sigma_t \epsilon_t, \quad (1)$$

where ϵ_t is independent Gaussian noise, α_t is a predefined coefficient, and σ_t controls the randomness level which we set to 0 for deterministic sampling. We may also optionally apply classifier-free guidance (CFG) [55] by modifying the noise estimate as:

$$\hat{\epsilon}_\theta(\mathbf{x}_t; y, t) = \epsilon_\theta(\mathbf{x}_t; \emptyset, t) + \gamma (\epsilon_\theta(\mathbf{x}_t; y, t) - \epsilon_\theta(\mathbf{x}_t; \emptyset, t)), \quad (2)$$

where γ denotes the strength of the conditional guidance and $\emptyset$ is the unconditional embedding of the empty string. This often results in much higher-quality samples.

Latent diffusion. Latent Diffusion Models (LDMs) [96] perform the diffusion process in a latent space rather than in pixel space. A pretrained encoder and decoder pair, $\mathcal{E}$ and $\mathcal{D}$, translates between pixel space and latent space. The latent space is typically much more compact and information-dense, which makes diffusion in this space more efficient. We use pretrained LDMs in our approach, due to the availability of audio and visual models with the same latent space.

3.2 Multimodal Denoising

Our goal is to generate an example $\mathbf{x} \in \mathbb{R}^{m \times n}$ that would be likely to appear under both visual and audio distributions, $p_a(\cdot)$ and $p_v(\cdot)$. We formulate this as sampling from a product of expert models² [52]: $p_{av}(\mathbf{x}) \propto p_a(\mathbf{x})p_v(\mathbf{x})$. We follow recent work on the compositional generation that

²Recent work has called this a *conjunction* [30], since conceptually the samples are roughly from the intersection of both distributions.

samples from this distribution using the score functions from pretrained diffusion models [30]. In contrast to these approaches, however, our two models are trained on *two different modalities*.

We create our spectrograms using two pretrained latent diffusion models. One trained to generate images, $\epsilon_{\phi,v}(\cdot, \cdot, \cdot)$, and the other to generate spectrograms, $\epsilon_{\phi,a}(\cdot, \cdot, \cdot)$, both operating in the same latent space. We show an overview of our method in Fig. 3. Given a noisy latent, $\mathbf{z}_t$, and text prompts y_v and y_a corresponding to the desired image and spectrogram prompt respectively, we compute two CFG noise estimates (Eq. (2)):

$$\epsilon_v^{(t)} = \epsilon_{\phi,v}(\mathbf{z}_t; \emptyset, t) + \gamma_v (\epsilon_{\phi,v}(\mathbf{z}_t; y_v, t) - \epsilon_{\phi,v}(\mathbf{z}_t; \emptyset, t)), \quad (3)$$

$$\epsilon_a^{(t)} = \epsilon_{\phi,a}(\mathbf{z}_t; \emptyset, t) + \gamma_a (\epsilon_{\phi,a}(\mathbf{z}_t; y_a, t) - \epsilon_{\phi,a}(\mathbf{z}_t; \emptyset, t)), \quad (4)$$

where γ_v and γ_a are the corresponding visual and audio guidance scales. We then combine the noise estimates from both modalities by applying weighted averaging, producing a multimodal noise estimate that steers the denoising process toward a sample that is likely under the distribution of both images and spectrograms:

$$\tilde{\epsilon}^{(t)} = \lambda_a^{(t)} \epsilon_a^{(t)} + \lambda_v^{(t)} \epsilon_v^{(t)}, \quad (5)$$

where $\lambda_a^{(t)}$ and $\lambda_v^{(t)}$ are the weights of the audio and visual noise estimates at timestep t respectively.

With this new noise estimate $\tilde{\epsilon}^{(t)}$, we perform a step of DDIM (Eq. (1)) to obtain a less noisy latent, $\mathbf{z}_{t-1}$. Repeating this process we obtain the clean latent $\mathbf{z}_0$, which is then decoded using the decoder $\mathcal{D}$ to obtain the spectrogram $\hat{\mathbf{x}} = \mathcal{D}(\mathbf{z}_0)$. This spectrogram can further be converted to a waveform using a pretrained vocoder or colorized to an RGB image whose grayscale version is the spectrogram.

Warm-starting. We find it useful to warm-start the denoising process. In Sec. 4.5, we experiment with warm-starting using only the spectrogram noise estimates or only the image noise estimates. This can be represented by using $w_a^{(t)}$ and $w_v^{(t)}$ as the relative weight on the audio and the visual noise estimates respectively. We let

$$\lambda_a^{(t)} = \frac{w_a^{(t)}}{w_a^{(t)} + w_v^{(t)}}, \quad \lambda_v^{(t)} = \frac{w_v^{(t)}}{w_a^{(t)} + w_v^{(t)}}, \quad (6)$$

with $w_a^{(t)} = H(t_a T - t)$ and $w_v^{(t)} = H(t_v T - t)$ being Heaviside step functions, and t_a and t_v indicating the *proportion* of the reverse process that has audio or visual denoising respectively. When $t_a < 1.0$ and $t_v = 1.0$, we warm-start with only image denoising, and vice-versa. The above ensures that the weights $\lambda_a^{(t)}$ and $\lambda_v^{(t)}$ sum to one, and are equally weighted after warm-starting.

Colorization. After we generate a spectrogram, $\hat{\mathbf{x}}$, we can optionally colorize it to create a more visually appealing result. Since our spectrograms fall outside the distribution of pre-trained colorization models, we use Factorized Diffusion [41] to colorize, which samples a diffusion model while projecting the noisy intermediate images such that they equal $\hat{\mathbf{x}}$ when turned into grayscale. In doing so, the denoising process synthesizes only the “color component” of the sampled image, while the “grayscale component” is constrained to equal the generated spectrogram. Note that this method is similar to prior work [63, 23, 106, 114]. We choose this particular method due to its simplicity.

4 Experiments

We evaluate our methods using quantitative metrics and human studies. We also present qualitative comparisons and an analysis of our method, and why it works.

4.1 Implementation Details

Models. We select a pair of off-the-shelf latent diffusion image and audio models that share the same latent space, encoder, and decoder. For the image model, we use Stable Diffusion v1.5³ [96]. For the audio model, we use Auffusion⁴ [118], which finetunes Stable Diffusion v1.5 on log-mel spectrograms. To synthesize audio from the log-mel spectrograms, we consider two options: following [118] and using off-the-shelf HiFi-GAN [66] vocoder, or the Griffin-Lim algorithm [47, 92]. We use HiFi-GAN for our main experiments. In Sec. 4.5, we evaluate the choice of vocoder and verify that our resultant waveforms do indeed encode to a visually interpretable spectrogram.

³Stable Diffusion v1-5 [hugging face model card](#)

⁴Auffusion [hugging face model card](#)

Table 1: **Quantitative evaluation on images that sound.** We report CLIP, CLAP, FID, and FAD metrics, along with 95% confidence intervals shown in gray. The best results are highlighted in **bold**.

Method	Modality	CLIP (%) $\uparrow$	CLAP (%) $\uparrow$	FID $\downarrow$	FAD $\downarrow$
Stable Diffusion [96]	$\mathcal{V}$	34.5 (± 0.1)	2.2 (± 0.2)	–	41.74
Auffusion [118]	$\mathcal{A}$	22.5 (± 0.1)	48.3 (± 0.6)	290.29	–
Imprint	$\mathcal{A} \& \mathcal{V}$	27.2 (± 0.2)	32.3 (± 1.0)	244.84	29.42
SDS	$\mathcal{A} \& \mathcal{V}$	25.4 (± 0.2)	23.4 (± 1.4)	273.03	32.57
Ours	$\mathcal{A} \& \mathcal{V}$	28.2 (± 0.1)	33.5 (± 0.9)	226.46	19.21

Hyperparameters. We begin the reverse process with random latent noise $\mathbf{z}_T \in \mathcal{R}^{4 \times 32 \times 128}$, the same shape that Auffusion was trained on. Despite the image model not being trained on this specific size, we found that it nevertheless produces visually appealing results. We set the classifier guidance scales γ_v and γ_a to be between 7.5 and 10 and denoise the latents for 100 inference steps with warm-start parameters of $t_a = 1.0, t_v = 0.9$ to preserve audio priors. We decode the latent variables into images of dimension $3 \times 256 \times 1024$. By averaging across each channel, we obtain spectrograms corresponding to 10 seconds of audio. We re-normalize the spectrograms for visualization.

Baselines. As there is no previous work in this domain, we propose two baseline approaches. The first, inspired by Diffusion Illusions of Burgert *et al.* [11], uses multimodal score distillation sampling (SDS). We optimize a single-channel image $\mathbf{x} = g(\theta)$, where g is an implicit function parameterized by θ , using two SDS losses: one from the image diffusion model ϕ_v and the other from the audio diffusion model ϕ_a . This results in a gradient of:

$$\nabla_{\theta} \mathcal{L}_{\text{SDS}}(\mathbf{x} = g(\theta)) = \lambda_{\text{SDS}} \mathbb{E}_{t, \epsilon} \left[\omega_v(t) \left(\epsilon_v^{(t)} - \epsilon \right) \frac{\partial \mathbf{x}}{\partial \theta} \right] + \mathbb{E}_{t, \epsilon} \left[\omega_a(t) \left(\epsilon_a^{(t)} - \epsilon \right) \frac{\partial \mathbf{x}}{\partial \theta} \right], \quad (7)$$

where λ_{SDS} is the weight of the image SDS gradient and ϵ is the noise added to the image or latents. We implement this with pixel-based diffusion model DeepFloyd IF [26], as we find it performs better than Stable Diffusion with the SDS loss, and Auffusion [118]. This model thus does not require a shared latent space between vision and audio. We refer to this baseline as the *SDS*.

The second baseline involves taking existing images and subtracting them from existing spectrograms, multiplied by some scaling factor, inspired by [25]. This works when the spectrograms have high power, as the subtraction does not significantly affect the audio but still imprints an image into the spectrogram. We obtain spectrograms and images for this baseline via Auffusion and Stable Diffusion. This approach, which we call *imprint*, is simple but can be surprisingly effective. All methods use the same vocoder and post-processing for fairness. Please see Appendix A.3 for more details.

4.2 Quantitative Evaluation

We start by quantitatively evaluating the quality of our generated *images that sound*, examining how well the generated examples match the provided text prompts for each modality.

Experimental setup. Following the evaluation of Visual Anagrams [42], we create two sets of text prompt pairs. We randomly select 5 discrete (onset-based) and 5 continuous sound category names from VGGSound Common [17] as audio prompts. We randomly chose 5 objects and 5 scene classes for image prompts, formatted as “a painting of [class], grayscale”. This yields a total of 100 prompt pairs. We report Stable Diffusion and Auffusion performances as single-modality

Table 2: **Human study.** We show win-rates of our spectrograms against those generated by the SDS and *imprint* baselines. The first row indicates which audiovisual prompt pair is evaluated, formatted as [audio prompt] / [visual prompt], with the last column being the average of all seven prompt pairs. Note that 50% win-rate is chance performance, and as such our method outperforms the baselines in the vast majority of cases. Also note that this is a *best-case* evaluation – please see Sec. 4.3 for details. All results reported are % win-rate against the baseline with a 95% confidence interval in gray ($N = 100$).

Baseline	Metric	bell/castle	bark/dog	birds/garden	meow/kitten	racecar/racecar	tiger/tiger	train/train	Average
SDS	audio quality	53.1 (± 4.9)	69.4 (± 4.2)	95.9 (± 0.8)	75.5 (± 3.7)	88.8 (± 2.0)	70.4 (± 4.1)	88.8 (± 2.0)	77.4 (± 0.7)
	visual quality	60.2 (± 4.7)	51.0 (± 4.9)	98.0 (± 0.4)	32.7 (± 4.4)	69.4 (± 4.2)	68.4 (± 4.3)	94.9 (± 1.0)	67.8 (± 0.8)
	alignment	58.2 (± 4.8)	63.3 (± 4.6)	93.9 (± 1.1)	62.2 (± 4.7)	82.7 (± 2.8)	59.2 (± 4.8)	91.8 (± 1.5)	73.0 (± 0.8)
Imprint	audio quality	82.1 (± 3.0)	73.7 (± 3.9)	53.7 (± 5.0)	54.7 (± 5.0)	86.3 (± 2.4)	85.3 (± 2.5)	85.3 (± 2.5)	74.4 (± 0.7)
	visual quality	92.6 (± 1.4)	86.3 (± 2.4)	66.3 (± 4.5)	68.4 (± 4.3)	66.3 (± 4.5)	77.9 (± 3.5)	56.8 (± 4.9)	73.5 (± 0.8)
	alignment	88.4 (± 2.1)	87.4 (± 2.2)	60.0 (± 4.8)	65.3 (± 4.6)	86.3 (± 2.4)	80.0 (± 3.2)	85.3 (± 2.5)	78.9 (± 0.6)




Figure 4: **Qualitative comparison.** We show our qualitative results along with the *imprint* and SDS baselines given visual (first) and audio (second) prompts. Please zoom in for better viewing.

benchmarks to establish upper and lower bounds. We generate 10 samples for each prompt pair, except for the SDS baseline, for which we generate 4 samples due to its slower speed.

Evaluation metric. Following [50], we use the CLIP [95] score to measure the alignment between spectrograms and image text prompts, and analogously we use the CLAP [116] score to evaluate the alignment of audio with audio text prompts. An ideal method should excel at both simultaneously. We also report FID [51] and FAD [65] to evaluate the quality of generated examples where we use the results of Stable diffusion and Auffusion as reference sets respectively.

Results. We show our quantitative results in Tab. 1. Our method outperforms baselines across all metrics and performs comparably to single-modality models, which serve as rough upper bounds for each modality. This demonstrates our approach’s ability to generate meaningful *images that sound*, sampling from the intersection of natural image and spectrogram distributions. Stable Diffusion achieves a low CLAP score, indicating how poorly a randomly sampled natural image acts as a spectrogram. We observe that the SDS baseline often fails to optimize both modalities together. In contrast, our method achieves a higher success rate and generates more diverse results. Our method is significantly faster, generating one sample in 10 seconds compared to the SDS baseline’s 2-hour optimization time using NVIDIA L40s. The *imprint* baseline imprints the image onto the spectrogram, potentially degrading the sound pattern and leading to a lower CLAP score. Note that FID and FAD are distribution-based metrics, and as our task focuses on generating examples that lie in a small subset of the natural image and spectrogram distribution, higher FID scores, in general, are expected.

4.3 Human Studies

Experimental setup. We also perform two-alternative forced choice (2AFC) studies to evaluate our results. We construct seven paired text prompts by hand, ensuring semantic correlations between image and audio prompts, such as pairing a visual of dogs with the sound of dogs barking. Using these prompts, we generate samples using our method, the SDS baseline, and the *imprint* baseline, and hand-pick the best examples for evaluation. This *best-case evaluation* is useful as participants from MTurk are not expected to have prior knowledge about spectrograms, let alone domain expertise. Moreover, this evaluation matches the intended use case of our method, in which a user repeatedly queries the model for a result that they prefer based on artistic merit and quality. Participants, are presented with one sample from our method, and a corresponding sample from a baseline, and are




Figure 5: **Qualitative examples with colorization results.** We present 4 examples alongside their image prompts, audio prompts, and colorization prompts. Please refer to our [website](#) for video results.

asked to choose (1) the sample that looks most like the image prompt, (2) the sample that sounds most like the audio prompt, and (3) the sample in which the visual structure of spectrogram best aligns with that of image over time. The first two questions act as perceptual versions of CLIP and CLAP scores. The third question is designed to evaluate how well the visual structure of audio matches with the images as the spectrogram is played. Please see Appendix A.3 for further details and discussion.

Results. Win-rates between our method and baselines are presented in Tab. 2, broken down by prompt pair. We also include averaged win-rates over all prompt pairs in the final column. As can be seen, our method outperforms the two baselines in most cases. Human evaluators consistently rate our spectrograms as being higher in audio and visual quality, and as being better “visually-synced”; on average our method is 2-3 times as likely to be chosen as the better sample than baselines.

4.4 Qualitative Results

Results. We present qualitative results from our method as well as baselines in Fig. 4, with additional results from our method in Figs. 5, 8 and 9 in the appendix. Audio of all results can be found on our [website](#). As can be seen (and heard) our approach generates more visually appealing samples with better sound quality than compared to the baselines. The SDS baseline often focuses on one modality, to the detriment of the other, and in general, generates audio of lower quality. Moreover, the method suffers from the characteristic oversaturation of SDS-based results. The performance of the *imprint* baseline is highly dependent on the independently generated spectrogram and image and tends to fail when the two are misaligned, as in the castle example, or when the spectrogram has low energy, as the subtracted image is hard to see in already low energy regions of the spectrogram, such as with the kitten example. Interestingly, we found that our method often combines visual and acoustic elements. For example, the onsets of the bells ringing in Fig. 4 coincide with the towers of the castle, and the spectrogram patterns of meowing are hidden as stripes and edges on the kittens.

We show additional hand-picked results from our approach in Fig. 5 with colorization results, in which we can see more examples of our method blending acoustic and visual elements, such as the

Table 3: **Ablations.** We conduct the ablation study of the guidance scale (left) and the warm-starting (right). To evaluate overall performance, we normalize each score by boundaries in Tab. 1 and then sum them.

Method	Variation	CLIP (%) $\uparrow$	CLAP (%) $\uparrow$	Method	t_v	t_a	CLIP (%) $\uparrow$	CLAP (%) $\uparrow$	Overall $\uparrow$
Ours	$\gamma_v, \gamma_a = 5.0$	28.0	29.3	Ours	1.0	1.0	28.7	30.8	1.14
	$\gamma_v, \gamma_a = 7.5$	28.2	31.9		1.0	0.9	29.0	27.4	1.08
	$\gamma_v, \gamma_a = 10$	28.2	33.5		0.9	1.0	28.2	33.5	1.15
					0.8	1.0	27.4	35.9	1.14




Figure 6: **Cycle consistency check on the vocoders.** We show the original log mel-spectrogram decoded from latents and log mel-spectrograms obtained from waveforms synthesized by HiFi-GAN or Griffin-Lim.

water lilies corresponding to the frogs croaking, the corgis corresponding to the dogs barking, and the flowers corresponding to the birds chirping. Please see more results in Fig. 8 of Appendix A.2.

Multimodal compositionality. Prior work [72, 29, 7, 42, 41] shows that diffusion models may be “composed” to generate samples that are likely under two or more different probability distributions. Our method can be seen as extending this idea of compositionality to multiple modalities. On the face of it, the distribution of spectrograms and the distribution of natural images would seem to be completely disjoint. However, as our results show, perhaps surprisingly, some overlap exists. We believe that this is possible for two reasons. First, spectrograms and images are both fairly flexible, allowing for significant amounts of perturbation or changes in style before becoming unrecognizable. And second, images and spectrograms share certain low-level characteristics, such as edges, curves, and corners, indicating a certain amount of similarity. Please see Appendix A.1 for more analysis.

However, we find that not all compositions can be successful as shown in Fig. 9 of Appendix A.2. Moreover, careful selection of prompts is crucial to creating good results. For example, incorporating terms such as “lithograph style” or “black background” encourages the visual model to create areas of silence, which results in better quality, as shown in Fig. 5.

4.5 Ablations

Vocoder. To extract waveforms from our generated spectrograms we use HiFi-GAN [66], a neural vocoder. Given that our spectrograms are incredibly out of distribution, one concern with this setup is that the vocoder will ignore spectrograms and generate waveforms that do not match the inputs. To ameliorate this concern, we conduct a cycle consistency check by re-encoding the neural vocoder’s predicted waveform back into a spectrogram by performing an STFT. As can be seen in Fig. 6, the recomputed spectrograms are very similar to the original spectrogram, with only slightly less sharpness and some blurred textures⁵. This suggests we truly create visual spectrograms that look like images. We also experiment with using Griffin-Lim [92] as a vocoder, with similar results to HiFi-GAN as shown in Fig. 6. We opt to use HiFi-GAN as our default vocoder as it outperforms Griffin-Lim in audio quality, with Griffin-Lim attaining a CLAP score of **0.302**, compared to **0.335** obtained from HiFi-GAN. Please see more results in Fig. 7 of Appendix A.2.

Warm-starting. We also conduct an ablation study on our warm-starting strategy by varying which modality is warm-started and by how many steps. Results are presented in Tab. 3, where t_a and t_v are defined in Sec. 3.2. We find that warm-starting the denoising process with either image or audio diffusion yields higher scores in the corresponding modality, as that modality effectively gets free reign to set the high-level features of the final result. We find that allowing the audio diffusion model to denoise alone for the first 10% of the timesteps results in an attractive balance between CLIP and CLAP scores. Therefore, we adopt $t_v = 0.9$ and $t_a = 1.0$ for our main experiments.

⁵We note that perfect cycle consistency is not generally possible since vocoders are fundamentally lossy.

Guidance scale. We also explore different guidance scales γ_v and γ_a for our method. We present results in Tab. 3. We find that higher guidance scales generally yield better results on both modalities. We hypothesize that the higher guidance scales more strongly encourage the sample to come from the “intersection” of the conditional spectrogram and conditional image distributions, resulting in better alignment with both text prompts.

5 Discussion and Conclusion

In this work, we demonstrate that, perhaps surprisingly, there is a non-trivial overlap between the distribution of natural images and the distribution of natural spectrograms. We show this by sampling from the intersection of these two distributions, resulting in spectrograms that look like real images but also sound like real sounds. The method we proposed is simple and zero-shot, and leverages the compositional nature of diffusion with cross-modal models. We see our work as advancing multimodal compositional generation and opening up new possibilities for multimodal art.

Limitations. One limitation of our method is that it cannot generate examples that have both high-fidelity audio and image. We show failure cases, which occur for many prompts, in Fig. 9. Some of these failures may be due to the strict constraints of the problem, since realistic examples may not always exist at the intersection of both distributions. Our method is also limited by the quality of the audio diffusion model, whose performance lags behind that of visual models.

Potential negative societal impacts. The image and audio generation models that our method leverages are becoming progressively more powerful, and care must be taken in their deployment. Moreover, our method could potentially be used for steganography, secretly embedding images within audio. This capability may be used for deception, and we believe it deserves further consideration.

Acknowledgements. We thank Ang Cao, Linyi Jin, Jeongsoo Park, Chris Donahue, Alexei Efros, Prem Seetharaman, Justin Salamon, Julie Zhu, and John Granzow for their helpful discussions. This project is supported by the Sony Research Award and Cisco Systems. Daniel is supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1841052.

References

- [1] T. Afouras, A. Owens, J. S. Chung, and A. Zisserman. Self-supervised learning of audio-visual objects from video. *European Conference on Computer Vision (ECCV)*, 2020. 3
- [2] Aphex Twin. Formula, 1994. audio track. 2, 3
- [3] R. Arandjelovic and A. Zisserman. Look, listen and learn. In *Proceedings of the IEEE international conference on computer vision*, pages 609–617, 2017. 3
- [4] R. Arandjelovic and A. Zisserman. Objects that sound. In *Proceedings of the European conference on computer vision (ECCV)*, pages 435–451, 2018. 3
- [5] Y. Asano, M. Patrick, C. Rupprecht, and A. Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. *Advances in Neural Information Processing Systems*, 33:4660–4671, 2020. 3
- [6] O. Bar-Tal, H. Chefer, O. Tov, C. Herrmann, R. Paiss, S. Zada, A. Ephrat, J. Hur, Y. Li, T. Michaeli, et al. Lumiere: A space-time diffusion model for video generation. *arXiv preprint arXiv:2401.12945*, 2024. 3
- [7] O. Bar-Tal, L. Yariv, Y. Lipman, and T. Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation. *arXiv preprint arXiv:2302.08113*, 2023. 3, 9
- [8] R. Bensadoun, T. Monnier, Y. Kleiman, F. Kokkinos, Y. Siddiqui, M. Kariya, O. Harosh, R. Shapovalov, B. Graham, E. Garreau, et al. Meta 3d gen. *arXiv preprint arXiv:2407.02599*, 2024. 3
- [9] T. Brooks, A. Holynski, and A. A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 3
- [10] B. Buckle. Spectrogram art: A short history of musicians hiding visuals inside their tracks. Available from: <https://mixmag.net/feature/spectrogram-art-music-aphex-twin>, 2022. Mix-mag article. 2

- [11] R. Burgert, X. Li, A. Leite, K. Ranasinghe, and M. S. Ryoo. Diffusion illusions: Hiding images in plain sight. *arXiv preprint arXiv:2312.03817*, 2023. 3, 6, 18
- [12] A. Cao, J. Johnson, A. Vedaldi, and D. Novotny. Lightplane: Highly-scalable components for neural 3d fields. *arXiv preprint arXiv:2404.19760*, 2024. 3
- [13] C. Chen, R. Gao, P. Calamia, and K. Grauman. Visual acoustic matching. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [14] C. Chen, A. Richard, R. Shapovalov, V. K. Ithapu, N. Neverova, K. Grauman, and A. Vedaldi. Novel-view acoustic synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6409–6419, 2023. 3
- [15] G. Chen, Y. Wu, S. Liu, T. Liu, X. Du, and F. Wei. Wavmark: Watermarking for audio generation. *arXiv preprint arXiv:2308.12770*, 2023. 3
- [16] H. Chen, W. Xie, T. Afouras, A. Nagrani, A. Vedaldi, and A. Zisserman. Audio-visual synchronisation in the wild. *arXiv preprint arXiv:2112.04432*, 2021. 3
- [17] 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)*, pages 721–725. IEEE, 2020. 6
- [18] M. Chen, K. Su, and E. Shlizerman. Be everywhere-hear everything (bee): Audio scene reconstruction by sparse audio-visual samples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7853–7862, 2023. 3
- [19] Z. Chen, D. F. Fouhey, and A. Owens. Sound localization by self-supervised time delay estimation. *European Conference on Computer Vision (ECCV)*, 2022. 3
- [20] Z. Chen, I. D. Gebru, C. Richardt, A. Kumar, W. Laney, A. Owens, and A. Richard. Real acoustic fields: An audio-visual room acoustics dataset and benchmark. In *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*, 2024. 3
- [21] Z. Chen, X. Hu, and A. Owens. Structure from silence: Learning scene structure from ambient sound. In *5th Annual Conference on Robot Learning*, 2021. 3
- [22] Z. Chen, S. Qian, and A. Owens. Sound localization from motion: Jointly learning sound direction and camera rotation. In *International Conference on Computer Vision (ICCV)*, 2023. 3
- [23] J. Choi, S. Kim, Y. Jeong, Y. Gwon, and S. Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021. 5, 17
- [24] S. Chowdhury, S. Ghosh, S. Dasgupta, A. Ratnarajah, U. Tyagi, and D. Manocha. Adverb: Visually guided audio dereverberation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7884–7896, 2023. 3
- [25] Classical Music Reimagined. Fun with spectrograms! how to make an image using sound and music. Available from: <https://www.youtube.com/watch?v=N2DQFfID6eY>, 2017. Youtube video. 6
- [26] DeepFloyd Lab at StabilityAI. DeepFloyd IF: a novel state-of-the-art open-source text-to-image model with a high degree of photorealism and language understanding. <https://www.deepfloyd.ai/deepfloyd-if>, 2023. 3, 6, 17
- [27] P. Dhariwal and A. Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 3
- [28] Y. Du, Z. Chen, J. Salamon, B. Russell, and A. Owens. Conditional generation of audio from video via foley analogies. *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [29] Y. Du, C. Durkan, R. Strudel, J. B. Tenenbaum, S. Dieleman, R. Fergus, J. Sohl-Dickstein, A. Doucet, and W. S. Grathwohl. Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc. In *International Conference on Machine Learning*, pages 8489–8510. PMLR, 2023. 2, 3, 9
- [30] Y. Du, S. Li, and I. Mordatch. Compositional visual generation with energy based models. *Advances in Neural Information Processing Systems*, 33:6637–6647, 2020. 4, 5
- [31] D. Epstein, A. Jabri, B. Poole, A. Efros, and A. Holynski. Diffusion self-guidance for controllable image generation. *Advances in Neural Information Processing Systems*, 36:16222–16239, 2023. 3

- [32] Z. Evans, C. Carr, J. Taylor, S. H. Hawley, and J. Pons. Fast timing-conditioned latent audio diffusion. *arXiv preprint arXiv:2402.04825*, 2024. 3
- [33] C. Feng, Z. Chen, and A. Owens. Self-supervised video forensics by audio-visual anomaly detection. *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [34] S. Forsgren and H. Martiros. Riffusion - Stable diffusion for real-time music generation, 2022. 2, 3
- [35] R. Gao, C. Chen, Z. Al-Halah, C. Schissler, and K. Grauman. Visualechoes: Spatial visual representation learning through echolocation. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [36] R. Gao and K. Grauman. 2.5d visual sound. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [37] R. Gao and K. Grauman. Visualvoice: Audio-visual speech separation with cross-modal consistency. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [38] R. Gao, A. Holynski, P. Henzler, A. Brussee, R. Martin-Brualla, P. Srinivasan, J. T. Barron, and B. Poole. Cat3d: Create anything in 3d with multi-view diffusion models. *arXiv preprint arXiv:2405.10314*, 2024. 3
- [39] R. Gao, T.-H. Oh, K. Grauman, and L. Torresani. Listen to look: Action recognition by previewing audio. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10457–10467, 2020. 3
- [40] D. Geng and A. Owens. Motion guidance: Diffusion-based image editing with differentiable motion estimators. *arXiv preprint arXiv:2401.18085*, 2024. 3
- [41] D. Geng, I. Park, and A. Owens. Factorized diffusion: Perceptual illusions by noise decomposition. *arXiv:2404.11615*, April 2024. 2, 3, 5, 9, 17
- [42] D. Geng, I. Park, and A. Owens. Visual anagrams: Generating multi-view optical illusions with diffusion models. In *CVPR*, 2024. 2, 3, 6, 9
- [43] 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. 3
- [44] R. Girdhar, A. El-Nouby, Z. Liu, M. Singh, K. V. Alwala, A. Joulin, and I. Misra. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190, 2023. 3
- [45] R. Girdhar, M. Singh, A. Brown, Q. Duval, S. Azadi, S. S. Rambhatla, A. Shah, X. Yin, D. Parikh, and I. Misra. Emu video: Factorizing text-to-video generation by explicit image conditioning. *arXiv preprint arXiv:2311.10709*, 2023. 3
- [46] Y. Gong, A. Rouditchenko, A. H. Liu, D. Harwath, L. Karlinsky, H. Kuehne, and J. Glass. Contrastive audio-visual masked autoencoder. *arXiv preprint arXiv:2210.07839*, 2022. 3
- [47] D. Griffin and J. Lim. Signal estimation from modified short-time fourier transform. *IEEE Transactions on acoustics, speech, and signal processing*, 32(2):236–243, 1984. 4, 5
- [48] G. Gwardys and D. Grzywczak. Deep image features in music information retrieval. *International Journal of Electronics and Telecommunications*, 60:321–326, 2014. 2
- [49] A. Hertz, R. Mokady, J. Tenenbaum, K. Aberman, Y. Pritch, and D. Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 3
- [50] J. Hessel, A. Holtzman, M. Forbes, R. L. Bras, and Y. Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021. 7
- [51] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 7
- [52] G. E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 2002. 4
- [53] J. Ho, W. Chan, C. Saharia, J. Whang, R. Gao, A. Gritsenko, D. P. Kingma, B. Poole, M. Norouzi, D. J. Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 3

- [54] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3, 4
- [55] J. Ho and T. Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 4
- [56] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022. 3
- [57] L. Höllein, A. Cao, A. Owens, J. Johnson, and M. Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7909–7920, October 2023. 3
- [58] X. Hu, Z. Chen, and A. Owens. Mix and localize: Localizing sound sources in mixtures. *Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [59] C. Huang, Y. Tian, A. Kumar, and C. Xu. Egocentric audio-visual object localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22910–22921, 2023. 3
- [60] J. Huh, J. Chalk, E. Kazakos, D. Damen, and A. Zisserman. Epic-sounds: A large-scale dataset of actions that sound. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 3
- [61] V. Iashin and E. Rahtu. Taming visually guided sound generation. *arXiv preprint arXiv:2110.08791*, 2021. 3
- [62] V. Iashin, W. Xie, E. Rahtu, and A. Zisserman. Synchformer: Efficient synchronization from sparse cues. *arXiv preprint arXiv:2401.16423*, 2024. 3
- [63] B. Kawar, M. Elad, S. Ermon, and J. Song. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022. 5, 17
- [64] E. Kidron, Y. Y. Schechner, and M. Elad. Pixels that sound. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 88–95. IEEE, 2005. 3
- [65] K. Kilgour, M. Zuluaga, D. Roblek, and M. Sharifi. Fr\`echet audio distance: A metric for evaluating music enhancement algorithms. *arXiv preprint arXiv:1812.08466*, 2018. 7
- [66] J. Kong, J. Kim, and J. Bae. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in neural information processing systems*, 33:17022–17033, 2020. 5, 9
- [67] S. Liang, C. Huang, Y. Tian, A. Kumar, and C. Xu. Av-nerf: Learning neural fields for real-world audio-visual scene synthesis. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [68] Y.-B. Lin and G. Bertasius. Siamese vision transformers are scalable audio-visual learners. *arXiv preprint arXiv:2403.19638*, 2024. 2, 3
- [69] Y.-B. Lin, Y.-L. Sung, J. Lei, M. Bansal, and G. Bertasius. Vision transformers are parameter-efficient audio-visual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2299–2309, 2023. 2, 3
- [70] 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. *arXiv preprint arXiv:2301.12503*, 2023. 3
- [71] H. Liu, Q. Tian, Y. Yuan, X. Liu, X. Mei, Q. Kong, Y. Wang, W. Wang, Y. Wang, and M. D. Plumbley. Audioldm 2: Learning holistic audio generation with self-supervised pretraining. *arXiv preprint arXiv:2308.05734*, 2023. 3
- [72] N. Liu, S. Li, Y. Du, A. Torralba, and J. B. Tenenbaum. Compositional visual generation with composable diffusion models. In *European Conference on Computer Vision*, pages 423–439. Springer, 2022. 2, 3, 9
- [73] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9298–9309, 2023. 3
- [74] S. Luo and W. Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2837–2845, 2021. 3
- [75] T. Mahmud, Y. Tian, and D. Marculescu. T-vsl: Text-guided visual sound source localization in mixtures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 3

- [76] 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 preprint arXiv:2404.09956*, 2024. 3
- [77] S. Majumder, Z. Al-Halah, and K. Grauman. Learning spatial features from audio-visual correspondence in egocentric videos. *arXiv preprint arXiv:2307.04760*, 2023. 3
- [78] X. Mei, V. Nagaraja, G. L. Lan, Z. Ni, E. Chang, Y. Shi, and V. Chandra. Foleygen: Visually-guided audio generation. *arXiv preprint arXiv:2309.10537*, 2023. 3
- [79] 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. 3
- [80] H. Mittal, P. Morgado, U. Jain, and A. Gupta. Learning state-aware visual representations from audible interactions. *Advances in Neural Information Processing Systems*, 35:23765–23779, 2022. 3
- [81] S. Mo and P. Morgado. Localizing visual sounds the easy way. In *European Conference on Computer Vision*, pages 218–234. Springer, 2022. 3
- [82] P. Morgado, N. Vasconcelos, T. Langlois, and O. Wang. Self-supervised generation of spatial audio for 360 video. *Advances in neural information processing systems*, 31, 2018. 3
- [83] P. Morgado, N. Vasconcelos, and I. Misra. Audio-visual instance discrimination with cross-modal agreement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12475–12486, 2021. 3
- [84] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models, 2021. 3
- [85] Nine Inch Nails. Year zero, 2007. Music Album. 2
- [86] M. A. Nugroho, S. Woo, S. Lee, and C. Kim. Audio-visual glance network for efficient video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10150–10159, 2023. 3
- [87] A. Owens and A. A. Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 631–648, 2018. 3
- [88] N. Oxman. Sympawnies: animal portraits made of musical notations. Available from: <https://www.youtube.com/@Sympawnies>, 2023. Youtube Channel. 3
- [89] P. O'Reilly, Z. Jin, J. Su, and B. Pardo. Maskmark: Robust neuralwatermarking for real and synthetic speech. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4650–4654. IEEE, 2024. 3
- [90] K. Palanisamy, D. Singhania, and A. Yao. Rethinking cnn models for audio classification. *arXiv preprint arXiv:2007.11154*, 2020. 2
- [91] S. Park, A. Senocak, and J. S. Chung. Can clip help sound source localization? In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5711–5720, 2024. 3
- [92] N. Perraudin, P. Balazs, and P. L. Søndergaard. A fast griffin-lim algorithm. In *2013 IEEE workshop on applications of signal processing to audio and acoustics*, pages 1–4. IEEE, 2013. 5, 9
- [93] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *ICLR*, 2023. 3
- [94] C. Qi, X. Cun, Y. Zhang, C. Lei, X. Wang, Y. Shan, and Q. Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15932–15942, 2023. 3
- [95] 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*, pages 8748–8763. PMLR, 2021. 7
- [96] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 3, 4, 5, 6

- [97] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 conference proceedings*, pages 1–10, 2022. 3
- [98] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes, T. Salimans, J. Ho, D. J. Fleet, and M. Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022. 3
- [99] R. San Roman, P. Fernandez, H. Elshar, A. Défossez, T. Furon, and T. Tran. Proactive detection of voice cloning with localized watermarking. In *Forty-first International Conference on Machine Learning*, 2024. 3
- [100] A. Senocak, H. Ryu, J. Kim, T.-H. Oh, H. Pfister, and J. S. Chung. Sound source localization is all about cross-modal alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7777–7787, 2023. 3
- [101] U. Singer, A. Polyak, T. Hayes, X. Yin, J. An, S. Zhang, Q. Hu, H. Yang, O. Ashual, O. Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 3
- [102] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In F. Bach and D. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. 3
- [103] A. Somayazulu, C. Chen, and K. Grauman. Self-supervised visual acoustic matching. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [104] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3, 4
- [105] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. 3
- [106] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 3, 4, 5, 17
- [107] J. Sun, L. Deng, T. Afouras, A. Owens, and A. Davis. Eventfulness for interactive video alignment. *ACM Transactions on Graphics (TOG)*, 42(4):1–10, 2023. 3
- [108] K. Sung-Bin, A. Senocak, H. Ha, A. Owens, and T.-H. Oh. Sound to visual scene generation by audio-to-visual latent alignment. *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [109] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33:7537–7547, 2020. 18
- [110] The Beatles. Lucy in the sky with diamonds, 1967. 3
- [111] Tool. 10,000 days, 2006. Volcano Entertainment. 3
- [112] D. Ulyanov. Audio texture synthesis and style transfer. <https://dmitryulyanov.github.io/audio-texture-synthesis-and-style-transfer>, 2016. 2
- [113] Venetian Snares. Songs about my cats, 2001. Music Album. 2
- [114] Y. Wang, J. Yu, and J. Zhang. Zero-shot image restoration using denoising diffusion null-space model. *The Eleventh International Conference on Learning Representations*, 2023. 5, 17
- [115] J. Z. Wu, Y. Ge, X. Wang, S. W. Lei, Y. Gu, Y. Shi, W. Hsu, Y. Shan, X. Qie, and M. Z. Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023. 3
- [116] 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)*, pages 1–5. IEEE, 2023. 7
- [117] Z. Xie, S. Yu, M. Li, Q. He, C. Chen, and Y.-G. Jiang. Sonicvisionlm: Playing sound with vision language models. *arXiv preprint arXiv:2401.04394*, 2024. 3

- [118] J. Xue, Y. Deng, Y. Gao, and Y. Li. Auffusion: Leveraging the power of diffusion and large language models for text-to-audio generation. *arXiv preprint arXiv:2401.01044*, 2024. 2, 3, 5, 6, 18
- [119] H. Yang, H. Ouyang, V. Koltun, and Q. Chen. Hiding video in audio via reversible generative models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1100–1109, 2019. 3
- [120] K. Yang, B. Russell, and J. Salamon. Telling left from right: Learning spatial correspondence of sight and sound. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9932–9941, 2020. 3
- [121] J. Y. Zhang, A. Lin, M. Kumar, T.-H. Yang, D. Ramanan, and S. Tulsiani. Cameras as rays: Pose estimation via ray diffusion. In *International Conference on Learning Representations (ICLR)*, 2024. 3
- [122] H. Zhao, C. Gan, W.-C. Ma, and A. Torralba. The sound of motions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1735–1744, 2019. 3
- [123] L. Zhao, H. Li, X. Ning, and X. Jiang. Thinimg: Cross-modal steganography for presenting talking heads in images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5553–5562, 2024. 3

A.1 Multimodal Compositionality Analysis

Model capabilities. Through our experiments, we observed that our method generally performs well with “continuous” sound events (*e.g.*, racing cars or train whistles) and simple visual prompts. Continuous sounds typically produce spectrograms with high energy distributed across time and frequencies, resulting in “white” spectrograms. This allows our model to effectively reduce sound energy, creating visual patterns that align with the audio.

Simple visual prompts with object or scene nouns provide the diffusion models with more flexibility during denoising, enabling sampling from the overlapping distributions of images and spectrograms. However, more complex prompts could push the models into highly constrained distributions where images that sound are less likely to be generated.

Generating discrete sounds (*e.g.*, dog barking or birds chirping) is more challenging due to their sparse energy distribution. In these cases, the models are more constrained, making it difficult to produce visual content with clear edges and structures aligned with sound onsets, leading to less satisfactory results sometimes.

Additionally, we emphasize that some prompt pairs may not have overlapping image and spectrogram distributions, making it impossible to create meaningful examples. For instance, combining the visual prompt starry night with the audio prompt playing guitar leads to a conflict, where the image modality tends toward a dark image, while the audio suggests a brighter one.

Style words. Visual diffusion models are capable of generating RGB images, but spectrograms are only one channel. We therefore use style words like “grayscale” or “black background” to nudge the image denoising process toward a distribution that matches the spectrograms. As suggested by the reviewer, we conducted ablation

experiments by removing the grayscale style word. The results are shown in Tab. 4. The model produces similar results, but (as expected) the image quality slightly decreases while the audio quality slightly improves.

Table 4: Ablation on style words.

Method	CLIP ↑	CLAP ↑	FID ↓	FAD ↓
w/o “grayscale”	28.1	33.7	237.12	18.09
w/ “grayscale”	28.2	33.5	226.46	19.21

A.2 Qualitative results

More qualitative results. We show more qualitative results from our method with different prompts in Fig. 8. Please see our [website](#) for video results. We also provide random examples with random prompt pairs in Fig. 9 with the last two rows as failure cases. For the failure cases, we can see that they either have good audio quality but lose clear visual patterns (mountains/fireworks) or have clear visual appearances but noisy audio (dogs/trains).

Vocoder analysis. We show more examples of the vocoder cycle consistency experiment in Fig. 7. As can be seen, the spectrograms from HiFi-GAN are quite similar to the original ones decoded from latents, indicating that our method does not find adversarial examples against the vocoder, but truly does find spectrograms that look like images.

A.3 Implementation Details

Colorization. We use DeepFloyd IF⁶ [26] following Factorized Diffusion [41] for colorizing spectrograms. This technique colorizes a grayscale image by using a pretrained diffusion model zero-shot to generate the color component, and is similar to prior work such as [63, 23, 106, 114]. We use it due to its simplicity. We colorize spectrograms of size $1 \times 256 \times 1024$ by directly feeding these into the diffusion model, which we found produced reasonable results despite the fact that the model was not trained for this size. We use prompts of the form “a colorful photo of [image prompt]” and denoise for 30 steps with a guidance scale of 10. Additionally, we found that starting the denoising at step 7 of 30 gave better results, which we hypothesize works because it gives the model a stronger prior for what the structure of the image is than starting from pure noise.

⁶<https://huggingface.co/DeepFloyd>

SDS baseline. We follow Diffusion Illusions [11] to implement our SDS baseline with an implicit image representation. We use Fourier Features Networks [109] with a learnable MLP to generate images of size $1 \times 256 \times 1024$. We use stage I of DeepFloyd IF-M to perform image score distillation sampling. We randomly make eight overlapping 256×256 crops and resize them to 64×64 to compute the averaged image SDS loss with a guidance scale of 80. For the audio modality, we use Auffusion [118]. As Auffusion is a latent diffusion model, we encode the images into $4 \times 32 \times 128$ latents and perform the audio SDS loss with a guidance scale of 10, which we found gave the best performance in the audio-only generation. We set the weight of the image SDS loss λ_{sds} to 0.4 to ensure balanced optimization for both modalities. We use the AdamW optimizer with a learning rate of 10^{-4} and weight decay of 10^{-3} , and optimize the Fourier Feature Network for 40,000 steps. We also apply the warm-start strategy to this method by optimizing the audio SDS loss only for the first 5,000 steps by setting λ_{sds} to zero. We note this method does not require a shared latent codebook between image and audio diffusion models.

Imprint baseline. We begin by generating images, $\mathbf{x}_{\text{img}}$, and spectrograms, $\mathbf{x}_{\text{spec}}$, of size 256×1024 using Stable diffusion and Auffusion, respectively, both with a guidance scale of 7.5. Next, we use the generated images as masks by converting them into inverse grayscale images and scaling them by a factor ρ . This mask is then applied to the generated spectrogram to obtain the final result, given by $\mathbf{x}_{\text{spec}}(1 - \rho \text{gray}(1 - \mathbf{x}_{\text{img}}))$. The hyperparameter ρ controls the strength of energy reduction: larger values yield clearer visual patterns but poorer audio quality, and vice versa. To strike a good balance, we set $\rho = 0.5$. The imprint baseline takes 10 seconds to generate a sample on NVIDIA L40s.

Algorithm 1 Pseudocode in a PyTorch-like style for the *imprint* baseline.

```
# get images and specs from LDMs
img = stable_diffusion(text_v)
spec = auffusion(text_a)
# reduce the energy give image masks
mask = 1 - rho * (1.0 - img.mean(0))
spec = mask * spec
audio = vocoder(spec)
```

Prompt selection. We present the image and audio prompts used for the quantitative evaluation in Tab. 5. We use 10 prompts for each modality, for a total of 100 prompt pairs.

Table 5: Text prompts for the quantitative evaluation.

Image prompts	Audio prompts
a painting of castles, grayscale	dog barking
a painting of dogs, grayscale	cat meowing
a painting of kittens, grayscale	bird chirping, tweeting
a painting of tigers, grayscale	tiger growling
a painting of auto racing game, grayscale	church bell ringing
a painting of mountains, grayscale	race car, auto racing
a painting of a garden, grayscale	train whistling
a painting of a forest, grayscale	fireworks banging
a painting of a farm, grayscale	people cheering
a painting of a beach, grayscale	playing acoustic guitar

A.4 Human Studies

Participants for the human study were recruited from Amazon Mechanical Turk (MTurk), and were paid 0.50 USD for a task lasting less than 5 min. We use a total of seven prompt pairs and compare them against two baselines: the SDS baseline and the *imprint* baseline. For each method and prompt pair, we hand-selected two high-quality samples for a total of 84 videos. Each video is about 10 seconds long, and includes a vertical line moving from left to right, indicating the current temporal position in the spectrogram. All participants were shown 14 pairs of videos—seven pairs comparing our method to the SDS baseline, and seven pairs comparing our method to the *imprint* baseline, all randomly selected and blinded. The participants are then asked to answer three questions:

1. Which video LOOKS most like a [visual prompt]?
2. Which video SOUNDS most like a [audio prompt]?
3. In the video, we play the image as a sound, from left to right. In which video does the [visual prompt] better align with the [audio prompt] sounds?




Figure 7: **More results on the vocoder cycle consistency check.** We show the original log mel-spectrogram decoded from latents and log mel-spectrograms obtained from waveforms synthesized by HiFi-GAN and Griffin-Lim.

The first two questions are designed to evaluate the quality of the audio and image generated by the methods, and their alignment with the respective prompts. The third question seeks to understand how well the visual structure or texture of the generated image and spectrogram align. However, note we were not able to guarantee that the participants had prior experience with spectrograms. To mitigate this to an extent, we include the description as a preamble to the third question. Also, note that we use abbreviated versions of the audio and visual prompts to avoid excessively long questions. We provide the prompt pairs we used for human studies in Tab. 6 for reference, and screenshots of our survey including the title block as well as the first video pair and associated questions in Fig. 10.

Table 6: **Text prompts for the human studies.** We note that the prompts are paired.

Image prompts	Audio prompts
a painting of castle towers, grayscale	bell ringing
a painting of cute dogs, grayscale	dog barking
a painting of a blooming garden with many birds, grayscale	birds singing sweetly
a painting of furry kittens, grayscale	a kitten meowing for attention
a painting of auto racing game, grayscale	a race car passing by and disappearing
a painting of trains, grayscale	train whistling
a painting of tigers, grayscale	tiger growling




Figure 8: **More qualitative results.** We show more qualitative results of our approach. Please zoom in for better viewing.




Figure 9: **Random results.** We show random results from our approach, using random audio and visual text prompt pairs. We provide failure cases in the last two rows. Please zoom in for better viewing.

Survey on Audiovisual Perception v2

[SOUND REQUIRED]

PLEASE TURN ON YOUR SOUND FOR THIS SURVEY

Thank you for participating in this survey! We are creating images that can also be read as a sound, and need your help evaluating their quality. We will show pairs of videos, and you will answer questions about each pair.

anonymous@gmail.com [Switch account](#)

Not shared

* Indicates required question

VIDEO PAIR 1

Video 1

bell

Image prompt: a painting of castle towers, grayscale

Audio prompt: bell ringing

Watch on

Video 2

bell

Image prompt: a painting of castle towers, grayscale

Audio prompt: bell ringing

Watch on

Which video LOOKS most like a "castle tower"? *

☐ Video 1

☐ Video 2

Which video SOUNDS most like "bells ringing"? *

☐ Video 1

☐ Video 2

In the video, we play the image as a sound, from left to right. *

In which video does the "castle" better align with the "bells ringing" sounds?

☐ Video 1

☐ Video 2

Figure 10: **Human study screenshots.** We show screenshots from our human study survey. Here we show the title block, as well as the first pair of videos. The full survey contains 14 video pairs.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We confirm that the claims in the abstract and introduction (Sec. 1) accurately reflect our contribution.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We provide a discussion about the limitations of our approach in Sec. 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our paper does not include theoretical results or proofs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide all implementation details in Sec. 4.1 and Appendix A.3. We also include all details of the human study experiments in Sec. 4.3 and Appendix A.3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide all code in our supplemental material, including code for baselines. We will release all code on acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide all the experimental setup details for quantitative evaluation in Sec. 4.2 and Appendix A.3 and for human study in Sec. 4.3 and Appendix A.3.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in the appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report performance with 95% confidence intervals for our main experiments in Tabs. 1 and 2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We specify that we use NVIDIA L40s for all the experiments including baselines in Sec. 4.2, where we also specify that our method takes about 10 seconds to generate a sample.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We confirm that the research conducted conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss potential societal impacts in Sec. 5.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We use pretrained diffusion models with well-known risks. We refer readers to the model cards of these diffusion models for further discussion of safeguards. Our work does not introduce significant risk.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We use Stable Diffusion v1.5, DeepFloyd IF, and Auffusion. We cite all these models, and provide links to model weights and licenses in Sec. 4.1 and Appendix A.3.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Our paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [Yes]

Justification: We provide all details about the human study in Sec. 4.3 and Appendix A.3, including full text and screenshots.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: The organization where this research was conducted has IRB approval for perceptual studies conducted over Mechanical Turk.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.