

Factorized Diffusion: Perceptual Illusions by Noise Decomposition

Daniel Geng*, Inbum Park*, and Andrew Owens

University of Michigan
dgeng@umich.edu

https://dangeng.github.io/factorized_diffusion/

Abstract. Given a factorization of an image into a sum of linear components, we present a zero-shot method to control each individual component through diffusion model sampling. For example, we can decompose an image into low and high spatial frequencies and condition these components on different text prompts. This produces hybrid images, which change appearance depending on viewing distance. By decomposing an image into three frequency subbands, we can generate hybrid images with three prompts. We also use a decomposition into grayscale and color components to produce images whose appearance changes when they are viewed in grayscale, a phenomena that naturally occurs under dim lighting. And we explore a decomposition by a motion blur kernel, which produces images that change appearance under motion blurring. Our method works by denoising with a composite noise estimate, built from the components of noise estimates conditioned on different prompts. We also show that for certain decompositions, our method recovers prior approaches to compositional generation and spatial control. Finally, we show that we can extend our approach to generate hybrid images from real images. We do this by holding one component fixed and generating the remaining components, effectively solving an inverse problem.

Keywords: Diffusion models · Perceptual illusions · Hybrid images

1 Introduction

The visual world is full of phenomena that can be understood through image decompositions. For instance, objects look blurry when seen from a distance, while up close their details are highly salient—two distinct perspectives that can be captured by a decomposition in frequency space [46, 53]. During the day, we see in full color, while in dim light we perceive only luminance—an effect that can be appreciated with a color space decomposition.

We present a simple method for controlling the factors of such decompositions, allowing a user to generate images that are perceived differently under different viewing conditions, yet still globally coherent. We apply this approach to generating a variety of perceptual illusions (Fig. 1). (i) Inspired by the classic work of Oliva *et al.* [46] we generate *hybrid images*, whose interpretation changes

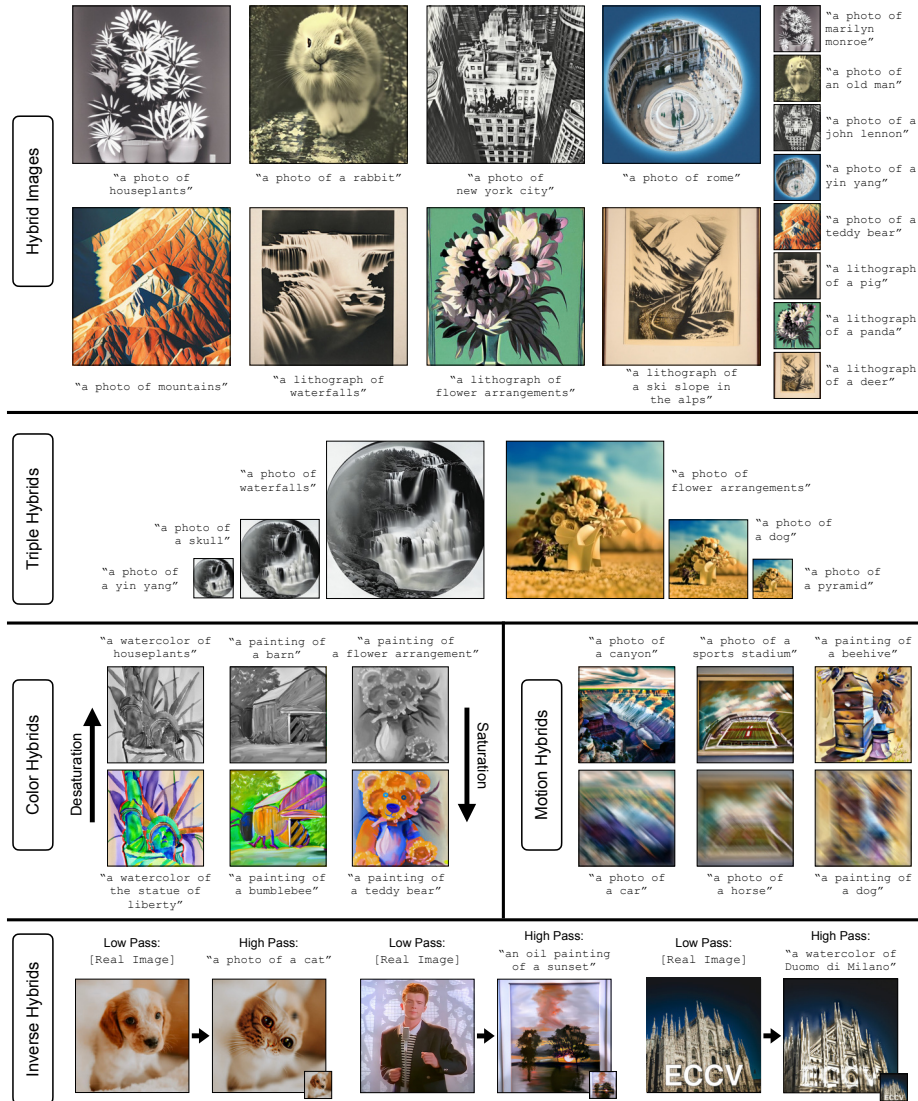


Fig. 1: Illusions by Factorized Diffusion. By conditioning the components of a generated image with different prompts, we can use off-the-shelf text-conditioned image diffusion models to synthesize hybrid images [46], hybrid images containing three objects, and new perceptual illusions which we refer to as *color hybrids* and *motion hybrids*, which change appearance when color is added or motion blur is induced. In addition, we can extract a component from an existing image and generate the missing components, allowing us to produce hybrid images from real images, which we term *inverse hybrids*. Examples shown are hand-picked. For random samples please see Fig. 9 and Fig. 18. For the hybrid images, we include insets to aid in visualization. However, perception of this effect depends on the resolution of the images, so *we highly encourage the reader to zoom so that an image fills the screen completely, or visit our webpage for easier viewing.*

with viewing distance, which we achieve by controlling the generated image’s low and high frequency components. A decomposition into three subbands allows us to produce hybrid images with three different prompts, which we refer to as *triple hybrids*. (ii) We generate color images whose appearance changes when they are viewed in grayscale, a phenomena that naturally occurs under dim lighting. We call these *color hybrids*, and make these by controlling image luminance separately from color. (iii) Finally, we produce images that change appearance under motion blur, by using a blur kernel to decompose an image. We refer to these as *motion hybrids*.

Our approach consists of a simple change to the sampling procedure of an off-the-shelf diffusion model. Given an image decomposition and a text prompt to control each component, in each step of the reverse diffusion process we estimate the noise multiple times: once for each component, conditioned on its corresponding text prompt. We then assemble a composite noise estimate by combining components from each individual noise estimate, obtained by applying the decomposition directly to the noise estimates (Fig. 2). Notably, our approach does not require finetuning [51, 72] or access to auxiliary networks, as in guidance based methods [3, 20, 24, 35, 35, 44].

We also show that we can take components from existing images, and generate the remaining components conditioned on text. This recovers a simple method to solve inverse problems, and is highly related to prior work on using diffusion models for solving inverse problems [2, 9, 11–13, 32, 38, 57, 58, 69]. We apply this technique to producing hybrid images from real images. Finally, we show that using certain decompositions with our method recovers prior techniques for spatial [4] and compositional [36] control over text prompts.

In summary, our contributions are as follows:

- Given a decomposition of an image into a sum of components, we propose a zero-shot adaptation of diffusion models to control these components during image generation.
- Using our method, we produce a variety of perceptual illusions, such as images that change appearance under different viewing distances (*hybrid images*), illumination conditions (*color hybrids*), and motion blurring (*motion hybrids*). Each of these illusions corresponds to a different image decomposition.
- We provide quantitative evaluations comparing our hybrid images to those produced by traditional methods, and show that our results are better.
- We give an analysis and intuition for how and why our method works.
- We show a simple extension of our method allows us to solve inverse problems, and we apply this approach to synthesizing hybrid images from real images.

2 Related Work

Diffusion models. Diffusion models [14, 28, 55, 56, 58] are trained to denoise data corrupted by added Gaussian noise. This is achieved by estimating the noise in noisy data, potentially with some additional conditioning, such as with

text embeddings. To sample data from a diffusion model, pure Gaussian noise is iteratively denoised until a clean image remains. Each denoising step consists of an update that removes a portion of the predicted noise from the noisy image, such as DDPM [28] or DDIM [56]. One noteworthy application of diffusion models is for text-conditional image generation [33, 44, 50, 52], which we build our method on top of.

Diffusion model control. Diffusion models are capable of both generating and editing images conditioned on text prompts. By modifying the reverse process [4, 21, 39, 68, 73], finetuning [51, 72], performing text inversion [30, 41, 56, 66, 70], swapping attention maps [17, 26, 64], supplying instructions [6], or using guidance [3, 17, 24, 35, 44, 48], modifying the style, location, and appearance of content in an image has become a relatively accessible task. Another line of work on compositional generation [4, 15, 37, 68] shows that diffusion models can generate images that conform to compositions of text prompts. Our work builds upon this, and shows that similar techniques can be applied to prompting individual components of an image to produce perceptual illusions. Our work is also similar to Wang *et al.* [68], in which a diffusion model is used to generate multiple images that blend seamlessly into a zooming video. However, we focus on generating only single images that can be understood at multiple resolutions.

Another line of work uses pretrained diffusion models to solve inverse problems, such as colorization or inpainting, in a zero-shot setting [2, 9, 11–13, 32, 38, 57, 58, 69]. We show that an extension of our method recovers these techniques, and we use it to generate hybrid images, which has not been considered before.

Computational optical illusions. Optical illusions are entertaining, but can also serve as windows into human and machine perception [16, 22, 23, 27, 31, 43, 61, 67]. As such, much work has gone into developing computational methods for generating optical illusions [7, 8, 10, 19, 21, 25, 46, 47, 63, 65]. In classic work, Oliva *et al.* introduced hybrid images [46], which are images that change their appearance depending on viewing distance or duration [53]. These images work by combining low frequencies from one image with high frequencies of another, exploiting the multiscale processing of human perception [45, 54]. By contrast, our approach generates hybrid images from scratch with a diffusion model. This avoids manual alignment steps and leads to higher quality illusions.

Artists and researchers have recently used text-conditioned image diffusion models to generate optical illusions. For example, a pseudonymous artist [65] adapted a QR code generation model [34, 72] to create images that subtly match a target template. While these are also images with multiple interpretations, they are restricted to binary mask templates and require a specialized finetuned model. Burgert *et al.* [7] use score distillation sampling [49] to generate images that match other prompts when viewed from different orientations or overlaid on top of each other. Other methods such as Tancik [63] and Geng *et al.* [21] use off-the-shelf diffusion models [33, 50] to generate multi-view optical illusions that change appearance upon transformations such as rotations, flips, permutations, skews, and color inversions. These methods work by transforming the noisy image

multiple ways during the reverse diffusion process, denoising each transformed version, then averaging the noise estimates together. However, many types of transformations, like the multiscale processing considered in hybrid images, fail because they perturb the noise distribution [21]. Like these approaches, our work also changes the reverse diffusion process to produce images that have multiple interpretations. However, our approach manipulates the *noise estimate* rather than the noisy image, enabling us to handle illusions that prior work cannot. Please see Appendix G for additional discussion and results.

3 Method

For a given decomposition of an image into components, our method allows for control of each of these components through text conditioning. We achieve this by modifying the sampling procedure of a text-to-image diffusion model.

3.1 Preliminaries: Diffusion Models

Diffusion models sample from a distribution by iteratively denoising noisy data. Over T timesteps, they denoise pure random Gaussian noise, \mathbf{x}_T , until a clean image, \mathbf{x}_0 , is produced at the final step. At intermediate timesteps, a variance schedule is followed such that the noisy image at timestep t is of the form

$$\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1 - \alpha_t}\epsilon, \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is a sample from a standard Gaussian distribution, and α_t is a predetermined variance schedule. To sample \mathbf{x}_{t-1} from \mathbf{x}_t the diffusion model, $\epsilon_\theta(\cdot, \cdot, \cdot)$, predicts the noise in \mathbf{x}_t , conditioned on the timestep t and optionally on context y , such as a text prompt embedding. Afterwards, an update step, $\text{update}(\cdot, \cdot)$, is applied which removes a portion of the estimated noise, $\epsilon_\theta := \epsilon_\theta(\mathbf{x}_t, y, t)$, from the noisy image \mathbf{x}_t . The exact implementation of this step depends on the specifics of the method used, but it is—critically for our method—often a linear combination of \mathbf{x}_t and ϵ_θ (and possibly noise, $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$). For example, DDIM [56] (with $\sigma_t = 0$) performs the update as:

$$\mathbf{x}_{t-1} = \text{update}(\mathbf{x}_t, \epsilon_\theta) = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1 - \alpha_t}\epsilon_\theta}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1}}\epsilon_\theta. \quad (2)$$

3.2 Factorized Diffusion

An overview of our method can be found in Fig. 2. Our method works by manipulating the noise estimate during the reverse diffusion process such that different components of the estimate are conditioned on different prompts. Given a decomposition of an image, $\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$, into the sum of N components,

$$\mathbf{x} = \sum_i^N f_i(\mathbf{x}), \quad (3)$$

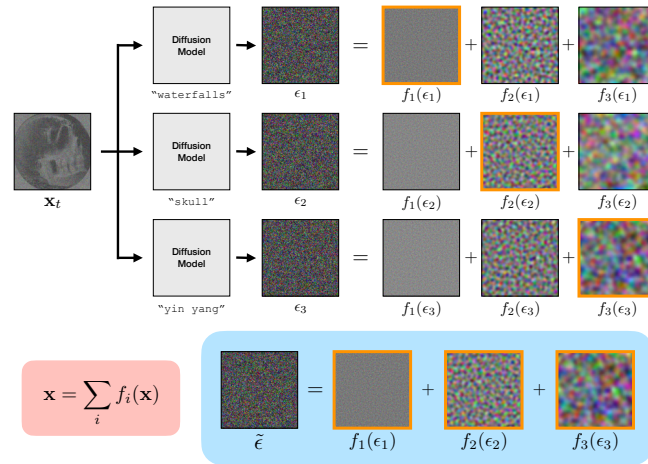


Fig. 2: Factorized Diffusion. Given an **image decomposition**, we control components of the decomposition through text conditioning during image generation. To do this, we modify the sampling procedure of a pretrained diffusion model. Specifically, at each denoising step, t , we construct a **new noise estimate**, $\tilde{\epsilon}$, to use for denoising, whose components **come from components** of ϵ_i , which are noise estimates conditioned on different prompts. Here, we show a decomposition into three frequency subbands, used for creating triple hybrid images, but we consider a number of other decompositions.

where each $f_i(\mathbf{x})$ is a component, we can correspond to each component a different text prompt y_i . At each step of the reverse diffusion process, instead of computing a single noise estimate we compute N —one conditioned on each y_i —which we denote by $\epsilon_i = \epsilon_\theta(\mathbf{x}_t, y_i, t)$. We then construct a composite noise estimate $\tilde{\epsilon}$ made up of components from each ϵ_i :

$$\tilde{\epsilon} = \sum f_i(\epsilon_i). \quad (4)$$

This new noise estimate, $\tilde{\epsilon}$, is used to perform the diffusion update step. In effect, each component of the image is denoised while being conditioned by a different text prompt, resulting in a clean image whose components are conditioned on the different prompts. We refer to this technique as *factorized diffusion*.

As noted in Sec. 2, our method is similar to recent work by Tancik [63] and Geng *et al.* [21] in that we modify noise estimates with the aim of generating visual illusions. However, our method differs in that we modify only the *noise estimate*, and not the input to the diffusion model, \mathbf{x}_t . As a result, our method produces a different class of perceptual illusions from prior work. Please see Appendix G for additional discussion and results.

3.3 Analysis of Factorized Diffusion

To give intuition for why our method works, we suppose that our update function $\text{update}(\cdot, \cdot)$ is a linear combination of the noisy image \mathbf{x}_t and the noise estimate

ϵ_θ , as is commonly the case¹ [28, 56]. The update function also depends on t , which we omit for brevity. We may then decompose the update step as

$$\mathbf{x}_{t-1} = \text{update}(\mathbf{x}_t, \epsilon_\theta) \quad (5)$$

$$= \text{update}\left(\sum f_i(\mathbf{x}_t), \sum f_i(\epsilon_\theta)\right) \quad (6)$$

$$= \sum_i \text{update}(f_i(\mathbf{x}_t), f_i(\epsilon_\theta)) \quad (7)$$

where the first equality is by definition of the update step, the second equality is by applying the image decomposition, and the third equality is by linearity of the update function. Eq. (7) tells us that an update step on \mathbf{x}_t , with ϵ_θ , can be interpreted as the *sum* of updates on *the components* of \mathbf{x}_t and of ϵ_θ . Our method can be understood as using different conditioning on each of these components. Written explicitly, the update our method uses is

$$\mathbf{x}_{t-1} = \sum_i^N \text{update}(f_i(\mathbf{x}_t), f_i(\epsilon_\theta(\mathbf{x}_t, y_i, t))). \quad (8)$$

Moreover, let us write out the update step explicitly as

$$\mathbf{x}_{t-1} = \text{update}(\mathbf{x}_t, \epsilon_\theta) = \omega_t \mathbf{x}_t + \gamma_t \epsilon_\theta, \quad (9)$$

for ω_t and γ_t determined by the variance schedule and scheduler. Then if the f_i 's are linear, we have

$$f_i(\mathbf{x}_{t-1}) = f_i(\text{update}(\mathbf{x}_t, \epsilon_\theta)) \quad (10)$$

$$= f_i(\omega_t \mathbf{x}_t + \gamma_t \epsilon_\theta) \quad (11)$$

$$= \omega_t f_i(\mathbf{x}_t) + \gamma_t f_i(\epsilon_\theta) \quad (12)$$

$$= \text{update}(f_i(\mathbf{x}_t), f_i(\epsilon_\theta)) \quad (13)$$

meaning that updating the i th component of \mathbf{x}_t with the i th component of ϵ_θ will only affect the i th component of \mathbf{x}_{t-1} .

3.4 Decompositions Considered

We present details for the decompositions that we consider in this paper. Results for all decompositions are presented and discussed in Sec. 4.

Spatial frequencies. We consider factorizing an image into frequency subbands, and conditioning the subbands on different prompts, with the goal of producing hybrid images [46]. First, we consider a decomposition into two components:

$$\mathbf{x} = \underbrace{\mathbf{x} - G_\sigma(\mathbf{x})}_{f_{\text{high}}(\mathbf{x})} + \underbrace{G_\sigma(\mathbf{x})}_{f_{\text{low}}(\mathbf{x})}, \quad (14)$$

¹ The update may also include adding random noise, $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$, in which case our analysis still holds with a modification to the argument, discussed in Appendix H.

where G_σ is a low pass filter implemented as a Gaussian blur with standard deviation σ , and $\mathbf{x} - G_\sigma(\mathbf{x})$ acts as a high pass of \mathbf{x} . For a decomposition into three subbands, to make the triple hybrid images in Fig. 1, components are levels of a Laplacian pyramid which we define as

$$\mathbf{x} = \underbrace{\mathbf{x} - G_{\sigma_1}(\mathbf{x})}_{f_{\text{high}}(\mathbf{x})} + \underbrace{G_{\sigma_1}(\mathbf{x}) - G_{\sigma_2}(G_{\sigma_1}(\mathbf{x}))}_{f_{\text{med}}(\mathbf{x})} + \underbrace{G_{\sigma_2}(G_{\sigma_1}(\mathbf{x}))}_{f_{\text{low}}(\mathbf{x})} \quad (15)$$

where σ_1 and σ_2 roughly define cutoffs for the low, medium, and high passes.

Color spaces. We also consider decomposition by color space, with the goal of creating *color hybrids*—images with different interpretations when seen in grayscale or color. Similarly to the CIELAB color space, we decompose an image into a lightness component, L , and a chromaticity component ab . CIELAB seeks to represent colors in a perceptually uniform space, and therefore requires nonlinear transformations of RGB values. Instead, we use a simple linear decomposition. Our L component is a channel-wise average of all the pixels

$$f_{\text{gray}}(\mathbf{x}) = \frac{1}{3} \sum_{c \in \{R, G, B\}} \mathbf{x}_c, \quad (16)$$

where \mathbf{x}_c are the color channels of the image, \mathbf{x} , and the resultant $f_{\text{gray}}(\mathbf{x})$ has the same shape as \mathbf{x} . We define the color component as the residual:

$$f_{\text{color}}(\mathbf{x}) = \mathbf{x} - f_{\text{gray}}(\mathbf{x}). \quad (17)$$

Motion blurring. Motion blur may be modeled as a convolution with a blur kernel \mathbf{K} [5, 18, 40, 42, 62, 71]. To produce images that change appearance when blurred, what we call *motion hybrids*, we study the following decomposition:

$$\mathbf{x} = \underbrace{\mathbf{K} * \mathbf{x}}_{f_{\text{motion}}(\mathbf{x})} + \underbrace{\mathbf{x} - \mathbf{K} * \mathbf{x}}_{f_{\text{res}}(\mathbf{x})}, \quad (18)$$

where we have split an image into a motion blurred component and a residual component. We specifically study simple constant velocity motions, in which \mathbf{K} may be modeled as a matrix of zeros with a line of non-zero values. This may also be thought of as decomposing an image into an oriented low frequency component, and a residual component.

Spatial decomposition. While our primary focus is on perceptual illusions, we also consider spatial masking as a decomposition. Given binary spatial masks \mathbf{m}_i whose disjoint union covers the entire image, we can use the decomposition

$$\mathbf{x} = \sum_i \underbrace{\mathbf{m}_i \odot \mathbf{x}}_{f_i(\mathbf{x})}, \quad (19)$$

where \odot denotes element-wise multiplication and each $\mathbf{m}_i \odot \mathbf{x}$ is a component. The effect of this decomposition is to enable control of the prompts spatially. This is a special case of MultiDiffusion [4]. We discuss this connection in Appendix E.

Scaling. A final interesting decomposition is of the form $\mathbf{x} = \sum_i^N a_i \mathbf{x}_i$, for $\sum_i^N a_i = 1$. Taking $a_i = \frac{1}{N}$ recovers the compositional diffusion method of Liu *et al.* [37], in which noise estimates are averaged to sample from conjunctions of multiple prompts. Taking $a_0 = 1 - \gamma$ and $a_1 = \gamma$, gives us CFG [29] and negative prompting [1].

3.5 Inverse Problems

If we know what one of the components must be in our generated image, perhaps extracted from some reference image \mathbf{x}_{ref} , we can then fix this component while generating all other components with our method. This enables us to produce hybrid images from real images (see Figs. [1] and [8]). Without loss of generality, suppose we want to fix the first component. To do this, we can project \mathbf{x}_t after every reverse process step:

$$\mathbf{x}_t \leftarrow f_1(\sqrt{\alpha_t} \mathbf{x}_{\text{ref}} + \sqrt{1 - \alpha_t} \epsilon) + \sum_{i=2}^N f_i(\mathbf{x}_t) \quad (20)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$, and α_t is determined by the variance schedule. The argument of f_1 is a sample from the forward process, given the reference image—that is, a noisy version of \mathbf{x}_{ref} with the correct amount of noise for timestep t . Essentially, we project \mathbf{x}_t such that its first component matches that of \mathbf{x}_{ref} . This amounts to solving a (noiseless) inverse problem characterized by $\mathbf{y} = f_1(\mathbf{x})$. Much work has gone into developing methods to solve inverse problems using diffusion models as priors, and this extension of our method can be viewed as a simplified version of prior work [2, 9, 11, 32, 38, 57, 69].

4 Results

We provide results organized by decomposition, followed by results on inverse problems, and then random samples. Additional implementation details can be found in Appendix [A], and additional results can be found in Appendix [K].

4.1 Hybrid Images

We show qualitative results in Fig. [1], Fig. [3], and Fig. [4], as well as in Fig. [16] in the appendix. As can be seen, our method produces high quality hybrid images. Interestingly, we were also able to produce hybrid images with three different prompts (Figs. [1] and [14]) by using the Laplacian pyramid decomposition (Eq. [15]). While prior work [59] has attempted to generate these *triple hybrids* with traditional methods, our method far exceeds their results in terms of quality and recognizability (for details, see Appendix [C]).

Effect of blur kernel. In Fig. [3] we show how the strength of the Gaussian blur, σ , affects results. A lower σ value corresponds to a higher cut-off frequency on the low-pass filter, and results in the low-pass prompt being more prominently featured. Interpolating between σ values gives hybrid images.

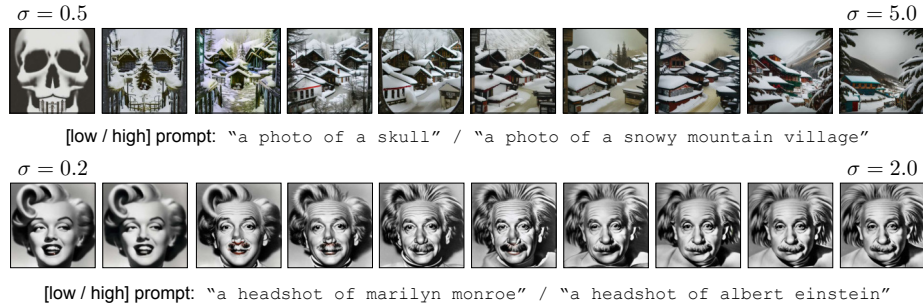


Fig. 3: Effect of σ . We show a linear sweep over the σ value used in our hybrid decomposition. A lower σ results in the low pass prompt being more prominent, and vice-versa. In between lies hybrid images. *Best viewed digitally, with zoom.*

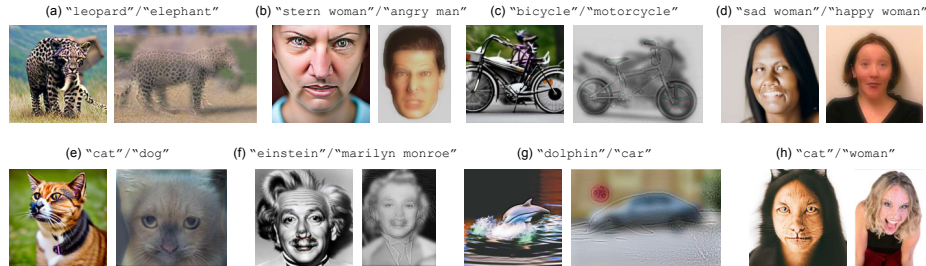


Fig. 4: Comparison to Oliva *et al.* [46]. We take hybrid images from Oliva *et al.* [46], and generate our own versions. Left is from our method, and right is from Oliva *et al.*'s. As can be seen, our method produces much more realistic images while still containing both subjects. *Best viewed digitally, with zoom.*

Comparisons to Oliva *et al.* [46]. In Fig. 4 we qualitatively compare our method to samples from Oliva *et al.* [46]. We directly take samples from [46], and manually create prompts to generate corresponding hybrids using our method. As can be seen, our hybrids are considerably more realistic, while containing the desired prompts at different viewing distances. One advantage our technique has is that the low frequency and high frequency components are generated with knowledge of each other, as the diffusion model is given the entire image. This is in contrast to the hybrid images of Oliva *et al.*, in which frequency components are extracted from two independent images and combined. Moreover, these two images must be found and made to align manually, whereas our method simply generates low and high frequency components that align well.

We also provide quantitative comparisons between our hybrid images and those of Oliva *et al.* [46]. In Tab. 1 we show results of a two-alternative forced choice (2AFC) study, in which human participants are asked to choose between our hybrid images or Oliva *et al.*'s. Participants of the study were asked which image better contained the prompts, and which of the images were of higher overall quality. For details, please see Appendix B. We find that participants consistently choose our images as being both higher in quality and better containing the prompts.

Table 1: Human Studies. We compare our hybrid images and Oliva *et al.*'s with a two-alternative forced choice test. Participants were shown results from Fig. 4 and were asked which images better contained the prompts, and which were of higher overall quality. Percentages denote the proportion that chose our method. Please see Appendix B for additional details. We find that our method is rated as both higher in quality and better aligned with the prompts. ($N = 77$)

	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	Average
High Prompt	70.1%	81.8%	63.6%	51.9%	61.0%	53.2%	70.1%	54.5%	63.3%
Low Prompt	83.1%	84.4%	74.0%	87.0%	93.5%	87.0%	75.3%	81.8%	83.3%
Quality	92.2%	87.0%	83.1%	77.9%	85.7%	79.2%	92.2%	33.8%	78.9%

Table 2: Hybrid Image CLIP Evaluation. We evaluate hybrid images by reporting the maximum clip score over different amounts of blurring. We report the max to compensate for the fact that different hybrid images may be best viewed at different resolutions. Please see Fig. 4 for the referenced hybrid images, and Appendix D for metric implementation details.

	Method	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	Average
Low Pass	Oliva <i>et al.</i> [46]	0.268	0.258	0.316	0.250	0.237	0.264	0.257	0.241	0.261
	Ours	0.286	0.252	0.307	0.273	0.275	0.260	0.244	0.269	0.271
High Pass	Oliva <i>et al.</i> [46]	0.297	0.230	0.306	0.272	0.276	0.306	0.260	0.231	0.272
	Ours	0.321	0.242	0.301	0.258	0.292	0.324	0.320	0.277	0.292

Finally, we present CLIP alignment scores in Tab. 2. To account for the fact that the hybrid images are best viewed at many different resolutions, we report the maximum CLIP score between the prompt and the image blurred by different amounts. Please see Appendix D for metric implementation details. We find that our method generates hybrid images with better alignment to the prompts.

4.2 Other Decompositions

Color hybrids. We provide qualitative color hybrid results in Fig. 1 and Fig. 5, as well as in Fig. 17 in the appendix. As can be seen, the grayscale image aligns with one prompt, while the color image aligns with another. For example, in the "rabbit"/"volcano" image from Fig. 5 the ears of the rabbit are repurposed as plumes of lava in the grayscale image. Note that it is not sufficient to simply add arbitrary amounts of color to a grayscale image to achieve this effect, as the colors added must not change the alignment of the grayscale image with its prompt. One interesting application of this technique is to produce images that appear different under bright lighting versus dim lighting, where human vision has a much harder time discerning color.

Motion hybrids We provide qualitative motion hybrid results in Fig. 1 and Fig. 6, as well as in Fig. 15 in the appendix. These are images that change appearance when motion blurred. For all motion hybrids in the paper, we use a blur kernel of $\mathbf{K} = \frac{1}{k}\mathbf{I} \in \mathbb{R}^{k \times k}$, with $k = 29$, corresponding to a diagonal motion from upper left to bottom right.

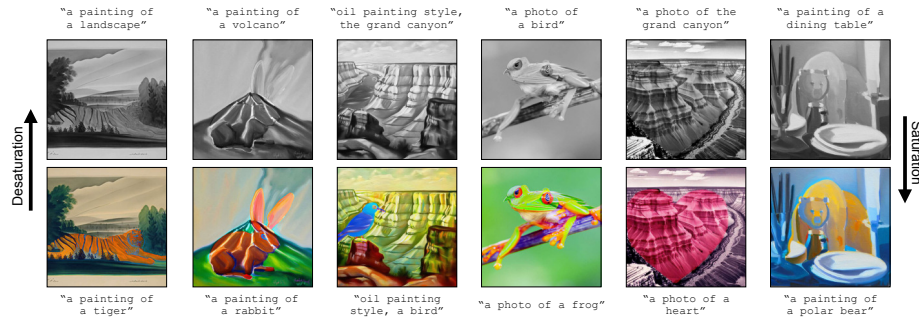


Fig. 5: Color Hybrids. We show additional *color hybrid* results. These are images that change appearance when color is added or subtracted away. These images change appearance when moved from bright to dim lighting, in which color is harder to see.




Fig. 6: Motion Hybrids. We show additional *motion hybrid* results. These are images that change appearance when motion blurred. Here, the motion is from upper left to bottom right.

Spatial decomposition By decomposing an image into disjoint spatial regions and applying our method, we can recover a technique that is a special case of MultiDiffusion [4]. Using this method, we can effect fine-grained control over where the text prompts act spatially, as shown in Fig. 7. For additional discussion, please see Appendix E.

Scaling decomposition By using the scaling decomposition with $a_i = \frac{1}{N}$, our method reduces exactly to prior work on compositionality in diffusion models, by Liu *et al.* [37]. Specifically, we recover the conjunction operator proposed by Liu *et al.* We demonstrate this in Fig. 7, but refer the reader to [37] for more examples.

4.3 Inverse Problems

As discussed in Sec. 3.5, we can modify our approach to solve inverse problems, resulting in a technique highly similar to prior work [2, 9, 11, 13, 32, 38, 58, 69].



Fig. 7: Spatial and Scaling Decompositions. Special cases of our method reduce to prior work. **(Left)** Decomposing images into spatial regions recovers a special case of MultiDiffusion [4], and allows us to assign prompts to spatial regions. **(Right)** Decomposing an image by scaling allows us to compose concepts, and recovers the method of Liu *et al.* [37].




Fig. 8: Hybrids from Real Images. We show hybrid images generated from real images. We take low or high passes of real images, and use our method to fill in the missing component, conditioned on a prompt. *Best viewed digitally, with zoom.*

While previous work investigates using diffusion priors for solving problems such as colorization, inpainting, super-resolution, or phase retrieval, we apply the idea towards generating hybrid images from real images. Specifically, we take low or high frequency components from a real image and use our method to fill in the missing components, conditioned on a prompt. Results are shown in Fig. [I] and Fig. [8]. We also provide colorization results in Appendix [J].

4.4 Limitations, Random Examples, and Societal Impacts

While our method can produce decent images fairly consistently, very high quality images are rarer. This can be seen in Fig. [9] and Fig. [18], in which we visualize random samples for hybrid images, color hybrids, and motion hybrids. We attribute this to the fact that our method produces images that are highly out-of-distribution for the diffusion model. Another failure case of our method is that the prompt for one component may dominate the generated image. Empirically, the success rate of our method can be improved by carefully choosing prompt pairs (see Appendix [I] for additional discussion), or by manually tuning decomposition parameters. Finally, our method relies on pixel diffusion models, as the image decompositions must be done in pixel space (see Appendices [A] and [F] for additional discussion and results).

The ability to better control image synthesis opens numerous societal and ethical considerations. We apply our method to generating illusions, which in a sense seeks to deceive perception, possibly leading to applications in misinformation. We believe this and other concerns deserve further study.




Fig. 9: Random Samples. We provide random samples for selected prompts and decompositions. As can be seen, most random results are of passable quality, with some catastrophic failures, and some very high quality illusions. More random samples are shown in Fig. 18.

5 Conclusion

We present a zero-shot method that enables control over different components of an image through diffusion model sampling and apply it to the task of creating perceptual illusions. Using our method, we synthesize hybrid images, hybrid images with three prompts, and new classes of illusions such as *color hybrids* and *motion hybrids*. We give an analysis and provide intuition for why our method works. For certain image decompositions, we show that our method reduces to prior work on compositional generation and spatial control of diffusion models. Finally, we make a connection to inverse problems, and use this insight to generate hybrid images from real images.

Acknowledgements. We thank Patrick Chao, Aleksander Holynski, Richard Zhang, Trenton Chang, Utkarsh Singhal, Huijie Zhang, Bowen Song, Jeongsoo Park, Jeong Joon Park, Jeffrey Fessler, Liyue Shen, Qing Qu, Antonio Torralba, and Alexei Efros for helpful discussions. We also thank Walter Scheirer, Luba Elliott, and Nicole Finn for reaching out and giving us the (amazing) opportunity to create an illusion for the CVPR 2024 T-shirt as part of the AI Art Gallery (see Appendix D). Daniel is supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1841052.

References

1. AUTOMATIC1111: Negative prompt. <https://github.com/AUTOMATIC1111/stable-diffusion-webui/wiki/Negative-prompt> (2022), <https://github.com/AUTOMATIC1111/stable-diffusion-webui/wiki/Negative-prompt>, accessed: November 7, 2023
2. Avrahami, O., Lischinski, D., Fried, O.: Blended diffusion for text-driven editing of natural images. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 18208–18218 (2022)
3. Bansal, A., Chu, H.M., Schwarzschild, A., Sengupta, S., Goldblum, M., Geiping, J., Goldstein, T.: Universal guidance for diffusion models (2023)
4. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: Multidiffusion: Fusing diffusion paths for controlled image generation. arXiv preprint arXiv:2302.08113 (2023)
5. Brooks, T., Barron, J.T.: Learning to synthesize motion blur. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6840–6848 (2019)
6. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: CVPR (2023)
7. Burgert, R., Ranasinghe, K., Li, X., Ryoo, M.: Diffusion illusions: Hiding images in plain sight. <https://ryanndagreat.github.io/Diffusion-Illusions> (Mar 2023)
8. Chandra, K., Li, T.M., Tenenbaum, J., Ragan-Kelley, J.: Designing perceptual puzzles by differentiating probabilistic programs. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–9 (2022)
9. Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: Ilvr: Conditioning method for denoising diffusion probabilistic models. arXiv preprint arXiv:2108.02938 (2021)
10. Chu, H.K., Hsu, W.H., Mitra, N.J., Cohen-Or, D., Wong, T.T., Lee, T.Y.: Camouflage images. ACM Trans. Graph. **29**(4), 51–1 (2010)
11. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. arXiv preprint arXiv:2209.14687 (2022)
12. Chung, H., Sim, B., Ryu, D., Ye, J.C.: Improving diffusion models for inverse problems using manifold constraints. Advances in Neural Information Processing Systems **35**, 25683–25696 (2022)
13. Chung, H., Sim, B., Ye, J.C.: Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12413–12422 (2022)
14. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021)
15. Du, Y., Li, S., Mordatch, I.: Compositional visual generation with energy based models. Advances in Neural Information Processing Systems **33**, 6637–6647 (2020)
16. Elsayed, G., Shankar, S., Cheung, B., Papernot, N., Kurakin, A., Goodfellow, I., Sohl-Dickstein, J.: Adversarial examples that fool both computer vision and time-limited humans. Advances in neural information processing systems **31** (2018)
17. Epstein, D., Jabri, A., Poole, B., Efros, A.A., Holynski, A.: Diffusion self-guidance for controllable image generation (2023)
18. Fergus, R., Singh, B., Hertzmann, A., Roweis, S.T., Freeman, W.T.: Removing camera shake from a single photograph. In: Acm Siggraph 2006 Papers, pp. 787–794 (2006)
19. Freeman, W.T., Adelson, E.H., Heeger, D.J.: Motion without movement. ACM Siggraph Computer Graphics **25**(4), 27–30 (1991)

20. Geng, D., Owens, A.: Motion guidance: Diffusion-based image editing with differentiable motion estimators. *International Conference on Learning Representations (2024)*
21. Geng, D., Park, I., Owens, A.: Visual anagrams: Generating multi-view optical illusions with diffusion models. *Computer Vision and Pattern Recognition (CVPR) 2024 (2024)*
22. Gomez-Villa, A., Martin, A., Vazquez-Corral, J., Bertalmío, M.: Convolutional neural networks can be deceived by visual illusions. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12309–12317 (2019)
23. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572 (2014)*
24. Gu, Z., Davis, A.: Filtered-guided diffusion: Fast filter guidance for black-box diffusion models. *arXiv preprint arXiv:2306.17141 (2023)*
25. Guo, R., Collins, J., de Lima, O., Owens, A.: Ganmouflage: 3d object nondetection with texture fields. *Computer Vision and Pattern Recognition (CVPR) (2023)*
26. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control (2022)
27. Hertzmann, A.: Visual indeterminacy in gan art. In: *ACM SIGGRAPH 2020 Art Gallery*, pp. 424–428 (2020)
28. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *arXiv preprint arxiv:2006.11239 (2020)*
29. Ho, J., Salimans, T.: Classifier-free diffusion guidance (2022)
30. Huberman-Spiegelglas, I., Kulikov, V., Michaeli, T.: An edit friendly ddpm noise space: Inversion and manipulations. *arXiv preprint arXiv:2304.06140 (2023)*
31. Jaini, P., Clark, K., Geirhos, R.: Intriguing properties of generative classifiers. *arXiv preprint arXiv:2309.16779 (2023)*
32. Kawar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. *Advances in Neural Information Processing Systems* **35**, 23593–23606 (2022)
33. Konstantinov, M., Shonenkov, A., Bakshandaeva, D., Ivanova, K.: If by deepfloyd lab at stabilityai (2023), <https://github.com/deep-floyd/IF/>, gitHub repository
34. Labs, M.: Controlnet qr code monster v2 for sd-1.5 (July 2023), https://huggingface.co/monster-labs/control_v1p_sd15_qrcode_monster
35. Lee, Y., Kim, K., Kim, H., Sung, M.: Syncdiffusion: Coherent montage via synchronized joint diffusions. In: *Thirty-seventh Conference on Neural Information Processing Systems (2023)*
36. Liu, N., Li, S., Du, Y., Tenenbaum, J., Torralba, A.: Learning to compose visual relations. *Advances in Neural Information Processing Systems* **34**, 23166–23178 (2021)
37. Liu, N., Li, S., Du, Y., Torralba, A., Tenenbaum, J.B.: Compositional visual generation with composable diffusion models. In: *European Conference on Computer Vision*. pp. 423–439. Springer (2022)
38. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Repaint: Inpainting using denoising diffusion probabilistic models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11461–11471 (2022)
39. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: SDEdit: Guided image synthesis and editing with stochastic differential equations. In: *International Conference on Learning Representations (2022)*
40. Mildenhall, B., Barron, J.T., Chen, J., Sharlet, D., Ng, R., Carroll, R.: Burst denoising with kernel prediction networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2502–2510 (2018)

41. Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models (2022)
42. Nayar, S.K., Ben-Ezra, M.: Motion-based motion deblurring. *IEEE transactions on pattern analysis and machine intelligence* **26**(6), 689–698 (2004)
43. Ngo, J., Sankaranarayanan, S., Isola, P.: Is clip fooled by optical illusions? (2023)
44. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models (2021)
45. Oliva, A., Schyns, P.G.: Coarse blobs or fine edges? evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive psychology* **34** (1997)
46. Oliva, A., Torralba, A., Schyns, P.G.: Hybrid images. *ACM Trans. Graph.* **25**(3), 527–532 (jul 2006). <https://doi.org/10.1145/1141911.1141919>, <https://doi.org/10.1145/1141911.1141919>
47. Owens, A., Barnes, C., Flint, A., Singh, H., Freeman, W.: Camouflaging an object from many viewpoints (2014)
48. Parmar, G., Singh, K.K., Zhang, R., Li, Y., Lu, J., Zhu, J.Y.: Zero-shot image-to-image translation (2023)
49. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. *arXiv* (2022)
50. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), <https://github.com/CompVis/latent-diffusion><https://arxiv.org/abs/2112.10752>
51. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation (2022)
52. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Ayan, B.K., Mahdavi, S.S., Lopes, R.G., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding (2022)
53. Schyns, P.G., Oliva, A.: From blobs to boundary edges: Evidence for time- and spatial-scale-dependent scene recognition. *Psychological science* **5**(4), 195–200 (1994)
54. Schyns, P.G., Oliva, A.: Dr. angry and mr. smile: when categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognitive* **69** (1999)
55. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *Bach, F., Blei, D. (eds.) Proceedings of the 32nd International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 37, pp. 2256–2265. PMLR, Lille, France (07–09 Jul 2015), <https://proceedings.mlr.press/v37/sohl-dickstein15.html>
56. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv:2010.02502* (October 2020), <https://arxiv.org/abs/2010.02502>
57. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005* (2021)
58. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: *International Conference on Learning Representations* (2021), <https://openreview.net/forum?id=PXTIG12RRHS>

59. Sripian, P., Yamaguchi, Y.: Hybrid image of three contents. *Visual computing for industry, biomedicine, and art* **3**(1), 1–8 (2020)
60. Svantner, P.: City skyline under blue sky during daytime (2024), <https://unsplash.com/photos/city-skyline-under-blue-sky-during-daytime-bM6LEOXfjP8>
61. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013)
62. Takeda, H., Milanfar, P.: Removing motion blur with space–time processing. *IEEE Transactions on Image Processing* **20**(10), 2990–3000 (2011)
63. Tancik, M.: Illusion diffusion. <https://github.com/tancik/Illusion-Diffusion> (Feb 2023)
64. Tumanyan, N., Geyer, M., Bagon, S., Dekel, T.: Plug-and-play diffusion features for text-driven image-to-image translation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1921–1930 (June 2023)
65. Ugleh: Spiral town - different approach to qr monster. https://www.reddit.com/r/StableDiffusion/comments/16ew9fz/spiral_town_different_approach_to_qr_monster/ (September 2023), https://www.reddit.com/r/StableDiffusion/comments/16ew9fz/spiral_town_different_approach_to_qr_monster/
66. Wallace, B., Gokul, A., Naik, N.: Edict: Exact diffusion inversion via coupled transformations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 22532–22541 (2023)
67. Wang, X., Bylinskii, Z., Hertzmann, A., Pepperell, R.: Toward quantifying ambiguities in artistic images. *ACM Transactions on Applied Perception (TAP)* **17**(4), 1–10 (2020)
68. Wang, X., Kontkanen, J., Curless, B., Seitz, S., Kemelmacher, I., Mildenhall, B., Srinivasan, P., Verbin, D., Holynski, A.: Generative powers of ten. *arXiv preprint arXiv:2312.02149* (2023)
69. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490* (2022)
70. Wu, C.H., De la Torre, F.: Unifying diffusion models’ latent space, with applications to cyclediffusion and guidance. *arXiv preprint arXiv:2210.05559* (2022)
71. Yitzhaky, Y., Mor, I., Lantzman, A., Kopeika, N.S.: Direct method for restoration of motion-blurred images. *JOSA A* **15**(6), 1512–1519 (1998)
72. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models (2023)
73. Zhang, Q., Song, J., Huang, X., Chen, Y., Liu, M.Y.: Diffcollage: Parallel generation of large content with diffusion models. *arXiv preprint arXiv:2303.17076* (2023)