

Illusion3D: 3D Multiview Illusion with 2D Diffusion Priors

Yue Feng¹ Vaibhav Sanjay¹ Spencer Lutz¹ Badour AlBahar² Songwei Ge¹ Jia-Bin Huang¹
¹University of Maryland, College Park ²Kuwait University
<https://3D-multiview-illusion.github.io>



Figure 1. **3D Multiview Illusion.** Our work expands the capability of existing multiview illusions (based on shadow, wire, or 2D plane) to 3D surfaces with perspective views. Different concepts can be observed when rendering our generated illusion from different perspectives. We showcase our generated 3D multiview illusions with different setups with text prompt inputs, including cubes, spheres, and reflective cylinders with 2D image grids and 3D shapes. For videos of the result gallery and real-world examples, please see [website](#).

Abstract

Automatically generating multiview illusions is a compelling challenge, where a single piece of visual content offers distinct interpretations from different viewing perspectives. Traditional methods, such as shadow art and wire art, create interesting 3D illusions but are limited to simple visual outputs (i.e., figure-ground or line drawing), restricting their artistic expressiveness and practical versatility. Recent diffusion-based illusion generation methods can generate more intricate designs but are confined to 2D images. In this work, we present a simple yet effective approach for creating 3D multiview illusions based on user-provided text prompts or images. Our method leverages a pre-trained text-to-image diffusion model to optimize the textures and geometry of neural 3D representations through differentiable rendering. When viewed from multiple angles, this produces different interpretations. We develop several techniques to improve the quality of the generated

3D multiview illusions. We demonstrate the effectiveness of our approach through extensive experiments and showcase illusion generation with diverse 3D forms.

1. Introduction

From one angle, it’s a campfire; from another, a butterfly—such is the magic of multiview illusions, where a single object shifts its interpretation with every change of perspective [18, 35, 48, 54, 58]. This is an exciting form of art where the viewer’s angles can affect their visual experiences. However, creating such an artwork is non-trivial. Shadow art is one of the popular forms [4, 23, 24, 45], using light and shadows to create different interpretations of a single 3D object. The 2D shadows cast by the 3D object can change based on the light’s position, altering the perceived visual output. Another form of multiview illusion is wire art [38, 57, 59], which involves creating 3D structures that re-

veal different line drawings when viewed from specific angles. Although these methods offer appealing 3D visual experiences, they are limited to producing simple images and require massive skill practice. Mirror and reflective cylinder art also present interesting visual effects (Figure 2d) but is confined to a *single* view. Motivated by the recent advances in using pretrained diffusion models to create detailed 2D illusion [3, 9], we aim to develop automatic methods for high-quality 3D multiview illusion generation.

We present a simple yet effective framework for creating 3D multiview illusions guided by text prompts and images. Our approach optimizes 3D neural representations through differentiable rendering to produce the desired visual interpretations. We achieve this with a pre-trained diffusion model using Variational Score Distillation [50]. Creating multiview 3D illusions induces additional challenges that are not in 2D settings. For example, the 3D ambiguity leads to substantial local minima that hinder optimization. To this end, we propose several techniques to improve the generation quality, including patch-wise denoising, scheduled camera jittering, and progressive render resolution scaling.

Existing diffusion models often function at a specific range of resolutions. We propose patch-wise denoising, which enables VSD-based methods to optimize effectively at higher resolutions. Consequently, our approach expands the capabilities of existing multiview illusions by generating detailed high-resolution 3D illusions at 1024×1024 .

When using the above method to generate 3D illusions, we observed a common issue where the described concepts appear multiple times within a single view, as shown in the third row of Figure 11. To mitigate this duplicate pattern issue, we introduce scheduled camera jittering. Specifically, we perturb the rendering camera with Gaussian noise during training, ensuring smooth and seamless transitions between views. We schedule the amount of Gaussian noise level during the denoising process and start with a large camera jittering to effectively reduce the duplicate pattern problem at early training stage.

We also develop a novel technique that progressively increases rendering resolution to improve the quality of the generated illusions. This technique enforces the fusion of image content across different views, while further reducing the duplicate pattern issue. Specifically, we progressively increase the rendering resolution throughout the optimization process. This approach ensures that the main object is optimized at the center of the view, resulting in clearer and more focused illusions when combined with camera jitter and patch-wise denoising.

With our illusion generation framework, we explore various 3D multiview illusions with different forms, including cubes and spheres with neural texture representations, reflective surfaces, and ambiguous 3D shapes. Our contributions include:

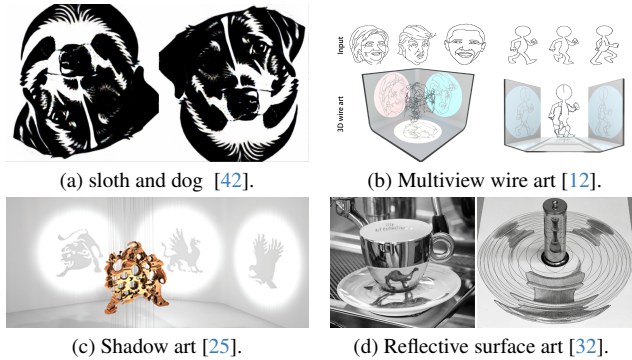


Figure 2. **2D and 3D Illusions. Each subfigure highlights different types of visual effects.** (a) 2D flip illusion [42]. Existing 3D illusions typically form line drawings [12] (b) or figure-ground images like shadows [24] (c). In contrast, our work expands the capability to generate 2D color images from different viewpoints. (d) When placed on a textured surface, a reflective surface can reveal new content [32]. However, it takes substantial time and effort for an artist to develop such a work. Our method simply describes the process and supports generating more than one single view.

- We explore a novel problem to generate 3D multiview illusions using pre-trained 2D diffusion models.
- We introduce several techniques and design choices to improve the quality of the 3D multiview illusions.
- We demonstrate that our approach can enhance and overcome the limitations of traditional artistic methods, with results applicable to real-world scenarios.

2. Related Work

2D illusion. Since 1930, psychologists have studied ambiguous and illusion figures [1, 2, 17, 20]. Human perception allows for multiple interpretations of the same image [29, 53]. Human perception is influenced by various factors such as the duration of stimulus [52], reaction time [36, 47], critical features [10], sensory memory [33], and predictive processing [34]. Researchers have developed algorithms in computer vision to select images and create customized hybrid [30] and mosaic images [56]. Recent works have extended the perception of visual illusions to convolutional neural networks (CNNs) [11]. Creating camouflage images by hiding objects in scenes has been achieved using various models [7, 31], including deep learning [60], GANs [55], and diffusion models [21]. Additionally, perceptual puzzles can be designed using differentiable algorithms [5]. Efforts have also been made to quantify ambiguities in illusion images [49] and understand perceptual illusions using generative classifiers [13]. However, models like CLIP [40] can still be fooled by optical illusions [28].

These illusions are time-consuming to make and require a high level of expertise. The development of generative AI, particularly Diffusion models [42], has simplified the

generation of high-quality images from text prompts. Recent works [3, 9] have utilized diffusion models to generate images with 2D illusion with different views corresponding to different text prompts (Fig 2a). These methods utilize diffusion models to generate illusion artwork in two main ways: 1) Simultaneously denoising the RGB image from different views [9] using a pixel-based diffusion model [16]. However, applying the inference pipeline to 3D generation presents significant challenges. This limitation arises because of the complexity introduced by the Gaussian property of noise in diffusion processes. While diffusion models excel in 2D settings, transitioning to 3D requires handling multiple overlapping views and maintaining coherence across different perspectives. 2) Using an optimization-based method to train an RGB image representation with Score Distillation Sampling (SDS) loss [44] or Variational Score Distillation (VSD) loss [50], which is most effective when training in the latent space [3]. We build upon an optimization-based approach to extend the capability to 3D illusion generation.

3D illusion. The perception of 3D objects may involve various forms of illusions. Different viewing angles can lead to varying interpretations of the same object [14, 46]. Additionally, the interpretation of 2D shadows cast by 3D surfaces can vary based on the light source’s position, affecting figure-ground perception [4, 23, 24, 45](see Fig 2b)). Similarly, the interpretation of line drawings of 3D wires are also illusions [12, 38, 57, 59] (see Fig 2c)). These approaches typically generate simple images, such as basic figure-ground representations, line drawings, or abstract forms. In contrast, our work creates colorful 3D illusions with complex, detailed interpretations, advancing the art of 3D illusion creation. Mirror and reflective cylinder art (Fig 2d)) present intriguing visual effects but are typically confined to a single view. Our approach generates multiple views, including top and side perspectives, and can even incorporate two reflective cylinders/mirrors to create three distinct views, which is beyond human ability.

3D generation using diffusion models. Diffusion models have demonstrated significant capability in generating photorealistic 2D images. Recently, they have been widely adopted for 3D generation, overcoming the need for large labeled 3D datasets [6, 19, 22, 37, 50?]. DreamFusion [37] leveraged text-to-image diffusion models for 3D synthesis by employing a probability density distillation (SDS) loss. However, the resulting images often have low resolution, excessive saturated colors, and over-smoothing issues. Building on this, Magic3D [19] achieved high-quality 3D mesh models through a two-stage optimization framework, significantly improving the quality of the generated 3D models. ProlificDreamer [50] further advanced the field by introducing Variational Score Distillation (VSD) loss to

address the saturation and smoothing problems associated with the SDS loss. This approach enhances the diversity and quality of the generated samples. We utilize Variational Score Distillation (VSD) as our baseline to generate our 3D multiview illusions using text-to-image diffusion models.

3. Method

Given multiple text prompts, we aim to create 3D illusions that reveal distinct interpretations from different viewing angles, with each interpretation corresponding to a specific text prompt. We create these illusions by changing the viewing angles of common 3D shapes or using reflective surfaces that generate reflections distinct from the original scene. Here, we present our approach for optimizing rendered views to accurately align with their corresponding text descriptions. Our method leverages the capabilities of pretrained diffusion models and incorporates several novel techniques and design choices to enhance the quality of these 3D illusions.

3.1. Background

Given a 3D representation G , we optimize the rendering of each viewpoint V_i to be aligned with a specific text prompt y_i . For example, we texture the surface of a 3D cube to generate content corresponding to multiple text prompts. We begin by selecting a corner of the cube and identifying the three adjacent faces connected to it. Each view V_i is configured to display any two adjacent faces, allowing us to associate three distinct text prompts with the corresponding views. This setup ensures that each face is visible in two different views, creating overlapping regions on the cube that are jointly optimized to satisfy two prompts.

We parameterize the texture field using a neural network with parameters θ of the 3D scene. This network employs a multiresolution hash encoding MLP from Instant-NGP [26] to embed the texture features. We denote the hashmap MLP as the function f , such that the queried texture RGB T is given by $T = f(\mathcal{E}(q); \theta)$, where q is the query coordinate and \mathcal{E} is the embedded texture map feature. To optimize our texturing module $f(\theta)$ for generating visual illusions, we utilize a pre-trained Stable Diffusion model [42]. While Score Distillation Sampling [37] is a common method for utilizing 2D diffusion priors, it often produces over-smoothing and color over-saturation artifacts. To mitigate these issues, we adopt the Variational Score Distillation (VSD) method [50].

At each training timestep, we render an RGB image by querying the texture field, defined as $T = f(\mathcal{E}(q); \theta)$. The VSD gradient is computed using a pre-trained diffusion model $\phi_{\text{pretrained}}$ along with a trainable LoRA module ϕ :

$$L_{VSD}(\theta) = \mathbb{E}_{t, \epsilon} \left[w(t) (\epsilon_{\phi_{\text{pretrained}}} - \epsilon_{\phi}) \frac{\partial f(\theta)}{\partial \theta} \right], \quad (1)$$

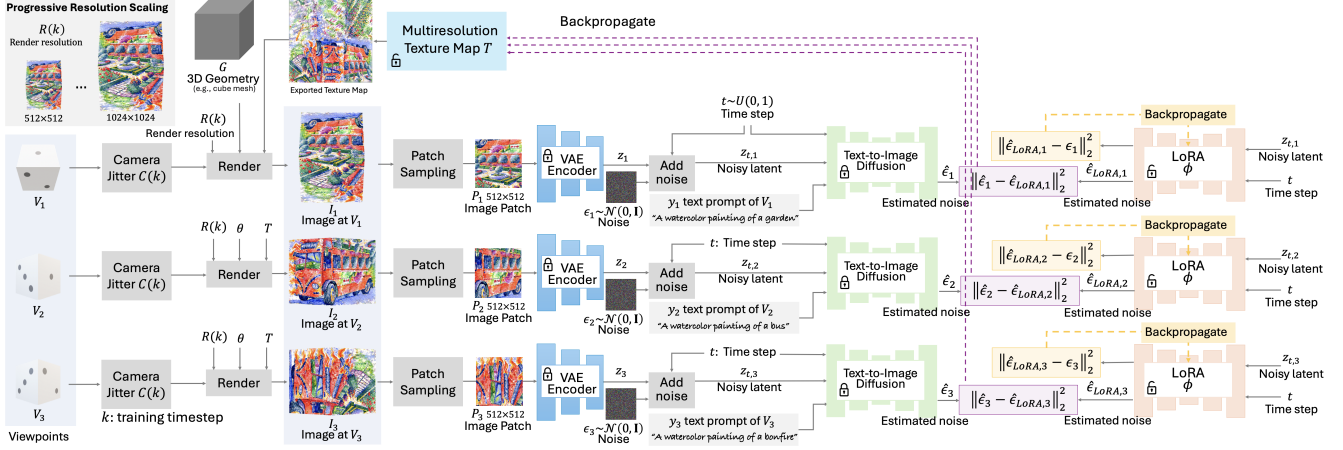


Figure 3. **Method overview.** We illustrate the process of generating a 3D multiview illusion from a cube with three varying interpretations from different viewpoints (V_1 , V_2 , and V_3) guided by text prompts (y_1 , y_2 , and y_3). First, we render the cube from the target viewpoints V_i applying scheduled camera jitter $C(k)$ and scheduled render size $R(k)$ with respect to the gradient flow time k . Camera jitter improves generation quality and render size scheduling helps reduce the duplicate pattern issue. We then utilize the multi-resolution texture field T to obtain the images (I_1 , I_2 , and I_3) at resolution $R(k) \in [512, 1024]$. To increase resolution during training, we extract a random patch P_i of size 512×512 from each rendered image I_i , which is then fed into a pre-trained VAE encoder. Given the 3D shape, we aim to optimize the parameters of the texture field. The generation of these viewpoints is optimized by leveraging a text-to-image diffusion model guided by the text prompts (y_1 , y_2 , and y_3). To avoid unnatural, saturated colors, we apply Variational Score Distillation (VSD) and LoRA model. We apply the same settings for spheres and scenes with reflective surfaces.

where $\epsilon_{\phi_{\text{pretrained}}} = \phi_{\text{pretrained}}(f(\theta); z, t)$ and $\epsilon_{\phi} = \phi(f(\theta); z, t)$. The time step t is randomly sampled from $t \sim \mathcal{U}(0.02, 0.98)$, and z represents the noisy input to the model with the injected noise following $\epsilon \sim \mathcal{N}(0, 1)$. The weighting function $w(t)$ adheres to the VSD configuration.

We alternate between updating the texture map parameters θ and the LoRA weights. The VSD gradient above is used to update the texture map parameters θ , while the LoRA module ϕ is optimized with the following objective by fine-tuning the diffusion model on current renderings:

$$L_{LoRA}(\phi) = \min_{\phi} \sum_{i=1}^n \mathbb{E}_{t, \epsilon} [\|\epsilon_{\phi}(f(\theta); z, t) - \epsilon\|_2^2]. \quad (2)$$

Generating 3D illusions using this baseline often fails, as shown in Figure 10, due to the under-constrained nature of 3D optimization compared to 2D illusion generation. In the following, we outline the key challenges faced and the strategies we developed to address them.

3.2. Camera jitter

Optimizing the texture representation with VSD loss from a fixed camera view can result in substantial artifacts. Interestingly, such artifacts are rarely observed in text-to-image generation results [50]. When computing the VSD gradient $\frac{\partial f(\theta)}{\partial \theta}$, we need to backpropagate the gradient to the texture map through the latent diffusion model encoder. We

hypothesize that these artifacts arise from the presence of a “blind spot” in the encoder. Specifically, the encoder’s latent representation does not change significantly with or without these artifacts. In contrast, text-to-3D approaches do not suffer from this issue due to dense camera views sampled during training. The artifacts in the “blind spot” of one view can be optimized in another.

Motivated by this observation, we introduce a random offset to the camera parameters during optimization. Without this adjustment, the output shows grid-like artifacts (see Figure 9). We find that incorporating a camera shift improves the smoothness of the generated images. However, applying this camera jitter technique from the start of training leads to duplicate pattern artifacts, particularly when patch-wise denoising is used to boost render resolution. To address this, we constrain the camera jitter and *gradually* increase its level as training progresses. We tested two scheduling strategies for increasing camera jitter, linear and sigmoid mapping, and found that a linear schedule generally performs better in practice.

We define the camera jitter $C(k)$ as a gradual, time-dependent adjustment to the camera parameters during training. Specifically, $C(k)$ is a linear function of the current training timestep k , where C_{max} represents the maximum camera jitter range, and k_{total} is total number of training steps:

$$C(k) = C_{\text{max}} \cdot \frac{k}{k_{\text{total}}}. \quad (3)$$

The camera parameters, rotation R , translation T , and field of view FOV are updated as follows:

$$\begin{aligned} R &= R_{\text{initial}} + \mathcal{N}(0, C(k) \cdot \sigma_R), \\ T &= T_{\text{initial}} + \mathcal{N}(0, C(k) \cdot \sigma_T), \\ FOV &= FOV_{\text{initial}} + \mathcal{N}(0, C(k) \cdot \sigma_{FOV}), \end{aligned} \quad (4)$$

where R_{initial} , T_{initial} and FOV_{initial} are the initial camera intrinsic parameters, and σ_R , σ_T , and σ_{FOV} denote the standard deviations for rotation, translation, and field of view.

3.3. Resolution scaling with patch denoising

In 3D optimization, the problem is often under-constrained, leading to multiple local optima that degrade illusion quality. A common failure mode occurs when the optimization cheats by showing each prompt on only one face. Ideally, prompts should overlap across faces to create the illusion (e.g., for the cube example, prompt 1 \rightarrow faces 1 and 3, prompt 2 \rightarrow faces 1 and 2, prompt 3 \rightarrow faces 2 and 3). However, a local optimum can emerge where each prompt i only shows up on face i , preventing the desired overlap and disrupting the illusion effect.

Another local optimum occurs when each face contains multiple, separate illusion contents, leading to a duplicate pattern issue. For example, face 1 might display prompts 1 and 2, face 2 shows prompts 2 and 3, and face 3 shows prompts 1 and 3. This redundancy disrupts the intended multiview illusion, resulting in incoherent and visually unappealing outputs. We often observe this issue in the baseline method (see failure cases 4.3), significantly impacting the quality of the generated illusions.

Moreover, pre-trained Stable Diffusion models [42] are limited to single-stage resolutions of 512×512 or 768×768 . To achieve high-resolution 1024×1024 results, one way is to adopt a two-stage illusion generation process similar to pretrained high-resolution Stable Diffusion models [?]. However, this method is not feasible, even in 2D illusion cases. Existing work like [8] and [51] employ patch-based denoising techniques with diffusion models to generate high-resolution images. While effective for some tasks, using random patches for training introduces duplicate pattern artifacts, as the model starts generating primary content inconsistently across patches. A simple strategy to mitigate this involves starting patches at the center and gradually expanding their location outward. However, we found that both starting gradually from the center or sampling stochastically generate content and subsequently replicate it elsewhere resulting in duplicate pattern artifacts.

To address this, we propose a progressive renderer resolution scaling strategy. Specifically, we begin by denoising a 512×512 image to ensure that no repeated content is present in the final rendered RGB image I_i . We then progressively enlarge the renderer resolution and sample a random 512×512 image patch to refine the main content. This

process effectively mitigates the duplicate pattern issue and enhances the overall quality of the multiview illusion while achieving high-resolution results.

We tested two scheduling strategies to progressively scale the renderer resolution $R(k)$, linear and sigmoid mapping. We find both significantly reduce duplicate patterns in practice, while sigmoid mapping $\sigma(k)$ generally demonstrated improved performance across randomly sampled prompt pairs y_i . We define the render resolution $R(k)$ as follows, where a is the initial resolution of 512, b is the final resolution of 1024, and k is the current training step:

$$\sigma(k) = \frac{1}{1 + \exp\left(-10 \cdot \left(\frac{k}{k_{\text{total}}} - 0.5\right)\right)} \quad (5)$$

$$R(k) = a + \sigma(k) \cdot (b - a) \quad (6)$$

4. Results

We validate our method’s performance through various experiments, including ablation studies, to assess the impact of different components and hyperparameters.

Baselines. We call the 512 resolution setting described in Sec 3.1 baseline in the following. The inverse projection baseline uses Stable Diffusion to independently generate images of the three prompts and optimize the cube’s texture map by performing an inverse projection of the 3D shape. The latent blending baseline also uses Stable Diffusion to independently generate images of the three prompts and then blend these views by averaging their latent variables in the overlapping region. We call the result using Dream Target loss from [3] as Burgert *et al.*.

Datasets. Illusion on real prompts is hard to succeed with [3, 9]. We randomly select the painting style and primary object and compose 43 prompt pairs to test illusion generation on spheres and cubes. We have a total of 86 illusion examples evaluated on CLIP score [39], aesthetic score/artifact [27], and alignment and concealment score from [3, 9], on baselines and other design choices. Styles or adding negative prompts do not significantly affect the generation.

Table 1. Comparison with different baselines.

Method	CLIP \uparrow	Aes \uparrow	Aes artifact \downarrow	\mathcal{A} \uparrow	\mathcal{C} \uparrow
Inverse projection	0.271	2.303	3.505	0.229	0.385
Latent blending	0.193	2.528	3.413	0.160	0.272
Burgert <i>et al.</i> [3]	0.179	2.958	3.309	0.144	0.206
Ours	0.288	2.791	3.247	0.251	0.461

4.1. Quantitative Results

We quantitatively compare our method with the baselines above in Table 1. Burgert *et al.* [3] can have a higher score on Aesthetic score, but their visual results are no better than




Figure 4. **Comparison with basic baselines.** Inverse projection can blend the images but cannot generate an illusion. Blending in latent space also fails. Burget *et al.*'s [3] method can not generate an illusion of 3D shape. Our method can blend the primary content while generating an illusion.



Figure 5. **3D multiview illusion with reflective surfaces.** We demonstrate illusion generation on a reflective cylinder (left) and a curved mirror (right).

ours; we provide user study Table 3 and qualitative results Fig 4 to showcase it. Our method has the best score except for the aesthetic score among all the baselines.

Our user study provides respondents with three results on the same prompt side-by-side. The result generated by our final method is compared with two other results generated through other methods. The user is then asked to select which one is the most visually appealing and which one aligns best with the set of prompts. This data is then aggregated, and we are able to compute the percentage of users who prefer our result over others, normalized by the number of times the result was presented to users.

4.2. Qualitative Results

We present qualitative results with basic baselines in Fig 4. Our method can fuse the content from different prompts and generate attractive results. We also show the setting of the reflective surface by placing reflective cylinders or curved mirrors on top of a plane in Fig 5 and Fig 6.




Figure 6. **3D multiview illusion with two reflective cylinders/mirrors.** Creating an illusion with a single reflective cylinder is already challenging for artists. For traditional methods, generating artwork on two cylinders simultaneously is even more difficult, if not impossible. Here, using 2D diffusion priors, we can create intriguing examples of illusion generation on two reflective cylinders. This showcases an extension of human capability in creating complex illusion artwork.




Figure 7. **Personalized image illusion generation with reflective surfaces.** Given an RGB image, we can supervise the generation of an image of one view by just using L2 loss and text to generate an image of another content. We reinvent the Finding Waldo game with a reflective cylinder. A real-world example is in the 3rd image.




Figure 8. **3D shape illusion.** A 3D shape generation model is trained with a view from the reflected cylinder. Columns 1 and 2 are different views of the generated 3D object. Column 3 is a view of the object from the reflective cylinder.

Table 2. **Ablation on various design choices on CLIP score, Aesthetic score, Alignment score and Concealment score.** Our method achieves the best among all the matrices

Method	Patch denoising	Camera jitter	Resolution scaling	CLIP \uparrow	Aes \uparrow	Aes artifact \downarrow	\mathcal{A} \uparrow	\mathcal{C} \uparrow
Baseline	×	×	×	0.184	2.526	3.481	0.143	0.195
Ablation-A (random patch denoising)	✓	×	×	0.128	2.700	3.784	0.083	0.172
Ablation-B (w/o resolution scaling, random camera jitter)	✓	random	×	0.163	2.682	3.594	0.125	0.169
Ablation-C (w/o resolution scaling, scheduled camera jitter)	✓	scheduled	×	0.155	2.520	3.693	0.119	0.193
Ablation-D (w/o camera jitter)	✓	×	✓	0.165	2.663	3.695	0.129	0.159
Ours	✓	scheduled	✓	0.288	2.791	3.247	0.251	0.461

Table 3. **Percentage comparison of visual appeal and prompt alignment for different models.** Numbers shown are the percentage of users who preferred a given method over alternatives in the user study.

	Burgert <i>et al.</i> [3]	Baseline	Ablation-A	Ablation-B	Ablation-C	Ablation-D	Ours
Visually Appealing?	10.14%	5.26%	28.15%	17.57%	29.50%	39.80%	60.61%
Aligns with Prompt?	4.73%	2.63%	30.37%	16.22%	18.78%	41.84%	63.17%



Figure 9. **Camera jitter can make the result smoother.** Adding camera jitter to the baseline method can make the transition between views more seamless but may also introduce two primary contents in the generation.

L2 loss. By supervising the cylinder view with an RGB image, taking the Waldo image as an example, we can generate a customized illusion with real example in Fig 7. We use a perfume cap placed on an iPad to fabricate the result.

4.3. 3D shape illusion

We also train a 3D shape illusion with reflective surfaces. Different content reveals when viewing a 3D object by looking at the reflective cylinder beside it; see Fig 8.

Ablation. Adding camera jitter can smooth the transition between views and may also introduce duplicate patterns if patch-wise training is included, see Fig 9. We present visual results under different design choices in Fig 10. Camera jitter and patch denoising can introduce duplicate pattern issues. Progressive resolution scaling with scheduled camera jitter can resolve this. We also shrink the view on the sphere

case in comparison with our orthogonal setting. In Fig 12 we can see how our method force the fusion between views.

Failure Cases. In Fig 11, row 1 is a typical failure case of random patch method, which happens on 70% of the randomly selected prompt pairs. Although our method solves most of the content fusion and duplicate pattern problems, there are still chances that these could happen, with less than 5% of the probability of generating it. Row 2 and 3 in Fig. 11 are failure cases of our methods. Row 2 has duplicate pattern artifacts; row 3 is a blending of image content.

5. Limitations

Our method may still suffer from duplicate pattern issues with a lower probability; see Fig 11. The illusion of an evenly shared surface is inherently ill-posed for the optimization process for each view. When optimizing two views that share the same area, for the cube case, once there is a principal component in that face that is not connected with the other face, we will have duplicate pattern issues (the same content on every single face of a cube). We resolved duplicate pattern issues by introducing progressive resolution scaling. However, there are still cases where the result is unsatisfactory or the training fails. This makes sense because we are experimenting with a random choice of prompt selections.

6. Conclusions

We introduce a novel approach for creating 3D multiview illusions guided by text prompts. Our work uses differentiable rendering and pre-trained diffusion models optimized with Variational Score Distillation (VSD) loss to generate detailed 3D illusions viewable from different angles. Key contributions include expanding multiview illusions to complex 3D shapes with detailed textures and introducing camera viewpoint jittering, patch denoising, and progressive resolution scaling, significantly improving quality and



Figure 10. **Comparison of different design choices.** Our method can have the best primary content fusion while obtaining visual quality. The baseline method suffers null space of VAE encoder without camera jitter. Random patch denoising can improve the resolution of results while introducing multiple duplicate primary content. Random camera jitter can make the transition smoother, but it still has duplicate pattern artifacts. The scheduled camera jitter resolved the duplicate pattern issue a bit but couldn't resolve it. Progressive render resolution scaling can also smooth the result, but the fusion of the primary content is less and has minor duplicate pattern issues on the horse. Our method combines progressive resolution scaling with scheduled camera jitter to produce results that center the primary content and diminish duplicate pattern artifacts.



Figure 11. **Failure cases.** We observe duplicate object generation and blending of objects among each view. Row 1, duplicate red panda and mice in random patch method. Row 2, duplicate monkey faces in the third view. Row 3, blending on primary objects.



Figure 12. **Shrink view on sphere.** We shrink the view from an orthogonal view to a more compact view on the sphere case. Intermediate views are presented in 2nd row and 4th row.

consistency. The camera viewpoint jitter improved the robustness and visual quality between the viewpoints. Patch denoising boosts the resolution of illusion results. Progressive resolution scaling worked best for illusion generation, ensuring more evident, focused illusions. We also identified challenges in generating illusions for highly overlapping and minimally distorted shapes, suggesting future research directions. Our work advances multiview illusions by combining text-to-image generation with 3D rendering, creating dynamic visual experiences. Future research can explore more complex shapes, optimize efficiency, and enhance realism, further unlocking the potential of 3D multiview illusions.

References

- [1] Michael Bach. Object perception: When our brain is impressed but we do not notice it. 2009. 2
- [2] Edwin G. Boring. A new ambiguous figure. *American Journal of Psychology*, 42:444, 1930. 2
- [3] Ryan Burgert, Xiang Li, Abe Leite, Kanchana Ranasinghe, and Michael Ryoo. Diffusion illusions: Hiding images in plain sight. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 2, 3, 5, 6, 7, 1
- [4] R. Chaine, Z. Deng, M. H. Kim, Aalok Gangopadhyay, Prajwal Singh, Ashish Tiwari, and Shanmuganathan Raman. Hand shadow art: A differentiable rendering perspective. 1, 3
- [5] Kartik Chandra, Tzu-Mao Li, Joshua B. Tenenbaum, and Jonathan Ragan-Kelley. Designing perceptual puzzles by differentiating probabilistic programs. *ACM SIGGRAPH Conference Proceedings*, 2022. 2
- [6] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In *ICCV*, 2023. 3
- [7] Hung-Kuo Chu, Wei-Hsin Hsu, Niloy J. Mitra, Daniel Cohen-Or, T. Wong, and Tong-Yee Lee. Camouflage images. *ACM TOG*, 2010. 2
- [8] Zheng Ding, Mengqi Zhang, Jiajun Wu, and Zhuowen Tu. Patched denoising diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2023. 5
- [9] Daniel Geng, Inbum Park, and Andrew Owens. Visual anagrams: Generating multi-view optical illusions with diffusion models. *ArXiv*, abs/2311.17919, 2023. 2, 3, 5
- [10] Michael S. Georgiades and John Harris. Biasing effects in ambiguous figures: Removal or fixation of critical features can affect perception. *Visual Cognition*, 4:383–408, 1997. 2
- [11] Aaron Hertzmann. Visual indeterminacy in gan art. In *ACM SIGGRAPH 2020 Art Gallery*, pages 424–428, 2020. 2
- [12] Kai-Wen Hsiao, Jia-Bin Huang, and Hung-Kuo Chu. Multi-view wire art. *ACM TOG*, 37:1 – 11, 2018. 2, 3
- [13] Priyank Jaini, Kevin Clark, and Robert Geirhos. Intriguing properties of generative classifiers. *ArXiv*, abs/2309.16779, 2023. 2
- [14] JJA Keiren, Freek van Walderveen, and Alexander Wolff. Constructability of trip-lets. In *Abstracts 25th European Workshop on Computational Geometry (EuroCG’09, Brussels, Belgium, March 16-18, 2009)*, pages 251–254, 2009. 3
- [15] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. 1
- [16] Mikhail Konstantinov, Alex Shonenkov, Daria Bakshandaeva, and Ksenia Ivanova. If by deepfloyd lab at stabilityai. *GitHub Repo*, 2023. 3
- [17] Jürgen Kornmeier and Michael Bach. Ambiguous figures – what happens in the brain when perception changes but not the stimulus. *Frontiers in Human Neuroscience*, 6, 2011. 2
- [18] Thi Ngoc Hanh Le, Sheng-Yi Yao, Chih-Kuo Yeh, Sheng-Jie Wang, and Tong-Yee Lee. Optimized binarization for eggshell carving art. *SIGGRAPH Asia 2021 Posters*, 2021. 1
- [19] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, 2023. 3
- [20] Gerald M. Long and Thomas C. Toppino. Enduring interest in perceptual ambiguity: alternating views of reversible figures. *Psychological bulletin*, 130 5:748–68, 2004. 2
- [21] Xuejiao Luo, Shuo Wang, Zongwei Wu, Christos Sakaridis, Yun Cheng, Deng-Ping Fan, and Luc Van Gool. Camdiff: Camouflage image augmentation via diffusion model. *ArXiv*, abs/2304.05469, 2023. 2
- [22] Gal Metzger, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *CVPR*, 2023. 3
- [23] Sehee Min, Jaedong Lee, Jungdam Won, and Jehee Lee. Soft shadow art. In *International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*, 2017. 1, 3
- [24] Niloy Jyoti Mitra and Mark Pauly. Shadow art. *ACM SIGGRAPH Asia 2009 papers*, 2009. 1, 2, 3
- [25] Niloy J. Mitra and Mark Pauly. Shadow art. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 28(5): 156:1–156:7, 2009. 2
- [26] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 41:1 – 15, 2022. 3
- [27] Naila Murray, Luca Marchesotti, and Florent Perronnin. Ava: A large-scale database for aesthetic visual analysis. *CVPR*, pages 2408–2415, 2012. 5
- [28] Jerry Ngo, Swami Sankaranarayanan, and Phillip Isola. Is clip fooled by optical illusions? In *Tiny Papers @ ICLR*, 2023. 2
- [29] Michael E. R. Nicholls, Owen Churches, and Tobias Loetscher. Perception of an ambiguous figure is affected by own-age social biases. *Scientific Reports*, 8, 2018. 2
- [30] Aude Oliva, Antonio Torralba, and Philippe G. Schyns. Hybrid images. *ACM SIGGRAPH 2006 Papers*, 2006. 2
- [31] Andrew Owens, Connelly Barnes, Alex Flint, Hanumant Singh, and William T. Freeman. Camouflaging an object from many viewpoints. *CVPR*, pages 2782–2789, 2014. 2
- [32] Bored Panda. These anamorphic illusions reveal their true form only when viewed in a mirror, 2016. Accessed: 2024-05-19. 2
- [33] Joel Pearson and Jan W. Brascamp. Sensory memory for ambiguous vision. *Trends in Cognitive Sciences*, 12:334–341, 2008. 2
- [34] Robert Pepperell. Being alive to the world: an artist’s perspective on predictive processing. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 379, 2023. 2
- [35] Maxine Perroni-Scharf and Szymon Rusinkiewicz. Constructing printable surfaces with view-dependent appearance. *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. 1
- [36] Thorsten Plewan, Ralph Weidner, and Gereon Rudolf Fink. The influence of stimulus duration on visual illusions and simple reaction time. *Experimental Brain Research*, 223:367 – 375, 2012. 2

- [37] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2022. 3
- [38] Zhiyu Qu, Lan Yang, Honggang Zhang, Tao Xiang, Kaiyue Pang, and Yi-Zhe Song. Wired perspectives: Multi-view wire art embraces generative ai. *ArXiv*, abs/2311.15421, 2023. 1, 3
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021. 5
- [40] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv*, abs/2204.06125, 2022. 2
- [41] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 1
- [42] Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *CVPR*, pages 10674–10685, 2021. 2, 3, 5
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 1
- [44] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 3
- [45] Kaustubh Sadekar, Ashish Tiwari, and Shanmuganathan Raman. Shadow art revisited: A differentiable rendering based approach. *CVPR*, pages 628–636, 2021. 1, 3
- [46] Guy Sela and Gershon Elber. Generation of view dependent models using free form deformation. *The Visual Computer*, 23(3):219–229, 2007. 3
- [47] Irene Sperandio, Silvia Savazzi, and Carlo A. Marzi. Is simple reaction time affected by visual illusions? *Experimental Brain Research*, 201:345–350, 2010. 2
- [48] Yingtao Tian. Evolving three dimension (3d) abstract art: Fitting concepts by language. *ArXiv*, abs/2304.12932, 2023. 1
- [49] Xi Wang, Zoya Bylinskii, Aaron Hertzmann, and Robert Pepperell. Toward quantifying ambiguities in artistic images. *ACM Transactions on Applied Perception (TAP)*, 17:1 – 10, 2020. 2
- [50] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In *NeurIPS*, 2023. 2, 3, 4, 1
- [51] Zhendong Wang, Yifan Jiang, Huangjie Zheng, Peihao Wang, Pengcheng He, Zhangyang Wang, Weizhu Chen, Mingyuan Zhou, et al. Patch diffusion: Faster and more data-efficient training of diffusion models. *NeurIPS*, 36, 2024. 5
- [52] Richard N. Wilton. The recency effect in the perception of ambiguous figures. *Perception*, 14:53 – 61, 1985. 2
- [53] Marina C Wimmer. Children’s perception and understanding of ambiguous figures. 2007. 2
- [54] Kang Wu, Xiaoming Fu, Renjie Chen, and Ligang Liu. Survey on computational 3d visual optical art design. *Visual Computing for Industry, Biomedicine, and Art*, 5, 2022. 1
- [55] Rui xiong Guo, Jasmine Collins, Oscar de Lima, and Andrew Owens. Ganmouflage: 3d object nondetection with texture fields. *CVPR*, pages 4702–4712, 2022. 2
- [56] Pengfei Xu, Jianqiang Ding, Hao Zhang, and Hui Huang. Discernible image mosaic with edge-aware adaptive tiles. *Computational Visual Media*, 5:45 – 58, 2019. 2
- [57] Pengfei Xu, Zhijin Yang, Hongbo Fu, and Hui Huang. Model-guided explorative design of abstract wire art. 2021. 1, 3
- [58] Jingru Yang, Shan He, and Lin Lu. Binary image carving for 3d printing. *Comput. Aided Des.*, 114:191–201, 2019. 1
- [59] Chih-Kuo Yeh, Thi Ngoc Hanh Le, Zhiling Hou, and Tong-Yee Lee. Generating virtual wire sculptural art from 3d models. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 18:1 – 23, 2022. 1, 3
- [60] Qing Zhang, Gelin Yin, Yongwei Nie, and Weishi Zheng. Deep camouflage images. In *AAAI Conference on Artificial Intelligence*, 2020. 2

Illusion3D: 3D Multiview Illusion with 2D Diffusion Priors

Supplementary Material

Our results can be best viewed as videos; please see [website](#). We provide more examples and more ablation studies aside from the main paper in the website.

A. Implementation details

We implement the differentiable rendering and reflective render using Pytorch3D [41]. Our method runs on an NVIDIA RTX A6000 GPU with 48GB RAM. We set the training time step for 2000, with timestep annealing to $t \sim \mathcal{U}(0.02, 0.5)$ after step 1000. We set the standard deviations in equation (4) to 1, and C_{\max} to 0.3. Adam optimizer [15] was used, with learning rates 1×10^{-3} for VSD loss and 1×10^{-4} for the LoRA loss. Each training takes 2 hours to converge. We use Stable Diffusion v2-1-base [43] as a pre-trained diffusion model.

B. Ablation

We show more examples of ablation among different methods in Fig 13 in addition to Fig 10. To ablate the ability to generate illusion, we increase the number of views to a more challenging setting in Fig 14. We look at every corner of the cube, and for each view, we look at three faces. In total, we have eight views of a cube, and each face is shared by four corners, i.e., four prompts. Videos can be found at the ablation part in [website](#).

C. Random sample, failure case and real-world example

We provide random samples of our result in Fig 16 and 17. All the prompts are randomly paired. A single reflective surface example is in Fig 5. We also give more examples of failure cases videos and analysis at [website](#). The human brain has objectified human faces in some sense, so it is harder to hide a human face from another view in the illusion generation process. Real-world examples videos can also be found at [website](#). We use an iPad and a perfume cap/reflective paper card to make the example. We just grab what we have to make the examples and haven't text other materials. We note that the surface of the perfume cap and paper card is not smooth. If the reflective material is more specular, the result will be more pretty.

D. 3D shape illusion

We provide more examples of 3D shape illusion generation at [website](#).

E. Texture map representation

We tested various texture map representations: plain RGB images, MLP-represented images, and multiresolution texture maps. Plain RGB images failed to produce consistent 2D illusions. MLP representations [3] generated 3D illusions but were limited in resolution due to GPU memory constraints. Multiresolution texture maps successfully produced high-quality 3D illusions with higher render resolutions, highlighting the importance of appropriate texture representation.

F. Discussion

Negative prompt of other views. To make the content of one view hide better from another view, i.e., improving the concealment of content. One may also propose that we add the text prompt of view 1 to the negative prompt of view 2, for example. We explored this method but figured it would not only ruin the generation but could also remove the content from view 2.

Gradient mask. Applying a gradient mask during optimization focusing on the central region can also enhance the fusion of input views, especially when the object is off-center. We explored this method on a 512-resolution case. We didn't include it in the 1024-resolution version because the render resolution scaling can do the job.

Prompt selection. Illusion generation struggles to converge without style constraints, yielding lower success rates. Style-constrained prompts like “*an oil painting of*” and “*a lithograph of*” perform better due to higher pixel consistency, aligning with findings in Diffusion Illusion [3].

Cube and sphere. The cube scenario was more challenging due to minimal surface distortion, which complicates the concealment of the content, which can be observed in Fig 16. In contrast, scenes with more significant view distortion (sphere, cylinder, and curved mirror) did not have this issue, indicating that high overlap and low distortion make content hiding more difficult.

Dream target loss. Dream target loss from Burgert *et al.* [3] uses a target image to help generate the illusion; this is useful in 2D cases when the view of the images doesn't change. In the 3D case, this method fails as it only supervises pixels within a fixed space that a VAE encoder can see. If we move the camera, the training will fail as the L2 loss part of Dream Target loss only supervises images pixel-wise with the target image. VSD loss [50] doesn't




Figure 13. **Comparison of different design choices on sphere case.** Our method can have the best primary content fusion while obtaining visual quality. The baseline method suffers null space of VAE encoder without camera jitter. Random patch denoising can improve the resolution of results while introducing multiple duplicate primary content. Random camera jitter can make the transition smoother, but it still has duplicate pattern artifacts. The scheduled camera jitter resolved the duplicate pattern issue a bit but couldn't resolve it. Progressive render resolution scaling can also smooth the result, but the fusion of the primary content is less and has minor duplicate pattern issues on the horse. Our method combines progressive resolution scaling with scheduled camera jitter to produce results that center the primary content and diminish duplicate pattern artifacts.




Figure 14. **Ablation on more views.** We extend the multiview illusion to 8 corners of a cube, each view containing three faces of a cube, and each adjacent corner has 180° flip viewing directions. We demonstrate our ability to generate illusions on more views on cube and sphere cases.

have this problem because it supervises the content based on the whole input image and the text prompt. So, if we jitter the view a bit, the training will still succeed, and we train part of the null space of a VAE encoder.

L2 loss. We also use L2 loss in the Waldo case in Fig. 7, supervising one view to a ground truth image pixel-wise. Because of the resolution mismatch of the low sampling

frequency of the texture map, we will have stripe-like artifacts on the texture map around the Waldo area on the texture map. To alleviate this problem, we query a patch of the pixel around that area where the reflected ray from the cylinder hit the plane. It can reduce the artifact, but cannot fully solve it. This makes sense because, technically, we are always under-sampling the multiresolution texture map.

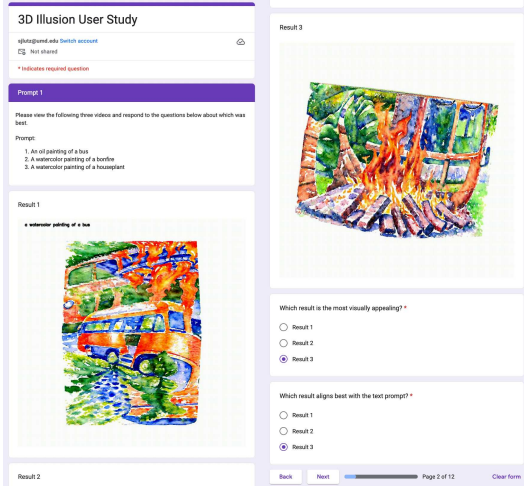


Figure 15. **Preview of User Study.** The user study was distributed over Google Forms. The user is provided with some context and the relevant prompts, followed by three animated GIFs, one for each method. They are then prompted with two multiple choice questions to evaluate the quality of the results.

G. User Study

We present details of user study. We distributed three surveys, each containing 11 sections of method comparisons to 39 users. Each section presented the user with three results generated on the same prompt but with different methods. The results are animated GIFs, allowing the user to evaluate all views of each result. The user is then prompted to answer two questions: “Which result is the most visually appealing?” and “Which result aligns best with the text prompt?”. A screenshot of user study is in Fig 15

To compute the result, we divide the number of times each method was selected, T_{selected} , by the number of times it was shown to the user, denoted as T_{shown} :

$$\text{Score} = \frac{T_{\text{selected}}}{T_{\text{shown}}} \quad (7)$$

This gives us a metric of the probability a user prefers a method’s result over others, given that the result is displayed.



Figure 16. **Random samples.** We show random samples with the same prompt pairs on cube and sphere cases.



Figure 17. **Random samples.** We present more random examples on cube and sphere case.