

## Mirror Illusion Art

Xiaopei Zhu<sup>1\*</sup> Zeyuan Li<sup>2\*</sup> Jun Zhu<sup>1,3</sup> Xiaolin Hu<sup>1,3,4†</sup>

<sup>1</sup>Department of Computer Science and Technology, BNRist, Tsinghua University

<sup>2</sup>Huazhong University of Science and Technology

<sup>3</sup>IDG/McGovern Institute for Brain Research, Tsinghua University

<sup>4</sup>Chinese Institute for Brain Research (CIBR)




Figure 1. Physical visualization of Mirror Illusion Arts designed by our method. Interestingly, it seems that the front view of the 3D object and its mirror view “do not appear” to belong to the same object. See *Supplementary Video* for a full-angle visualization of these objects.

### Abstract

*Mirror Illusion Art is a novel reflection-conditioned 3D illusion where one object yields two target appearances (front and mirror). The task is formulated as inverse design from two target 2D images (front and mirror) to a printable 3D object with geometry and texture. Prior topology-driven and shadow-based approaches demand substantial manual effort, optimize shape only, and often yield non-smooth or incomplete geometry. To address these challenges, we propose AutoMIA, an automated Mirror Illusion Art design pipeline that jointly optimizes shape and color. To stabilize optimization and suppress artifacts, four mechanisms are introduced: (1) projection-alignment component (PAC) selection to reduce surface noise, (2) position-weighted adaptive (PWA) suppression for background noise, (3) internal*

*voxel preservation (IVP) to prevent internal fractures, and (4) shape-color decoupled (SCD) optimization that balance shape and color optimization. AutoMIA generate diverse smooth Mirror Illusion artworks successfully both in the digital and physical world, with only around 76s design time and 2.6 GB memory on average using a single RTX 3090, advancing inverse graphics and computational design. Our code is available at <https://github.com/zxp555/AutoMIA>.*

### 1. Introduction

Art and design are important domains for exploring the creativity of AI. In such fields, AI can reduce the amount of complex thinking required by human designers, and even create many unexpected new designs. Among various arts, optical illusion art stands out as both artistically valuable and scientifically insightful, as it exposes the limitations of

\*Equal contribution.

†Corresponding author.

the human visual system and offers opportunities for deeper understanding of human perception. Existing optical illusion art can be broadly divided into two categories: 2D illusion art and 3D illusion art. Well-known examples include Hybrid Images [13] and Visual Anagrams [7] in the 2D domain, as well as Shadow Art [12] and Multi-View Wire Art [10] in the 3D domain.

In this work, we explore an interesting 3D optical illusion art, termed Mirror Illusion Art. Unlike traditional visual illusions, Mirror Illusion Art leverages the unique properties of reflection: when a 3D object is placed in front of a mirror, its reflection appears as a completely different object, distinct from the original view (see Fig. 1). This phenomenon arises from the deliberate design of both the geometry and texture of the 3D object, thus effectively deceiving human perception. To automate the creation of such illusions, we formulate the following problem: given two 2D images, one representing the appearance in front of the mirror and the other specifying the desired reflection, we develop AI algorithms to generate a 3D illusion object whose real-world reflection achieves the target Mirror Illusion effect.

Sugihara [18] proposed a topology-based design method for Mirror Illusion Art. However, this approach has two key limitations. First, the design process heavily relies on human intuition and intricate mathematical calculations, making it challenging for beginners. Second, the method focuses solely on shape optimization and does not support the creation of designs involving colored patterns. Another approach is based on Shadow Art [12], however, this method also restricts optimization to shape and cannot handle color patterns. Moreover, since Shadow Art primarily focuses on the projected shadow rather than the geometry of the 3D object itself, it often results in 3D objects that are not smooth or complete, which in turn degrades the visual quality and physical realizability of Mirror Illusion Art.

To address these limitations, we propose an automated method named AutoMIA for designing 3D Mirror Illusion Art. Given any pair of user-specified 2D images, our approach optimizes a 3D model to produce the desired mirror illusion art. Furthermore, our method jointly optimizes both 3D shape and color, enabling the creation of smooth and color-rich 3D Mirror Illusions. To attain this goal, we found challenges to remove several annoying defects, including “surface noise”, “background noise”, “internal fracture” and “color-shape imbalance”.

To address these challenges, we introduce several targeted solutions. We propose a projection alignment-based component (PAC) selection method to mitigate the “surface noise”. To suppress “background noise”, we develop an position-weighted adaptive (PWA) noise suppression method. For resolving “internal fracture” inside the 3D object, we design an internal voxel preservation (IVP)

mechanism during optimization. Finally, to alleviate the “color-shape imbalance”, we develop a shape-color decoupled (SCD) optimization method.

Experiments demonstrate that our method can efficiently generate various Mirror Illusion Arts, covering diverse categories such as letters, digits, Chinese characters, cartons, logos, emoji icons, and so on. Moreover, all of our designs are physically realizable. Figure 1 showcases real-world artworks produced using our method. On average, our approach requires only around 76 seconds of design time and 2.6 GB of memory on average using a single RTX 3090 GPU, highlighting its efficiency and lightweight advantage. This work shows the synergy of science and art, highlighting how AI can help create new kinds of art and also help us better understand how people see and interpret visual patterns.

## 2. Related Work

### 2.1. 2D Illusion Arts

Early computational approaches to 2D illusion art exploit specific properties of human vision. For example, hybrid images [13] utilize multi-scale perception by combining low- and high-frequency components from two images, producing a morphing effect with viewing distance. Motion illusions [5] leverage the visual system’s sensitivity to local phase changes, generating continuous motion perception without actual displacement. Camouflage arts [3, 8, 14, 19] hide targets by matching background statistics and adjusting luminance or feature cues, often using learning-based objectives for naturalness and difficulty control. However, these illusion arts are time-consuming to design, and usually require a high level of expertise.

Recent approaches frame 2D illusion synthesis using generative priors. Diffusion-based methods [1, 7] generate multi-view illusions by jointly denoising across pixel transformations, enabling concealment without explicit perceptual models. Analyses of GAN imagery [9] highlight “visual indeterminacy” as a source of illusion-like artifacts. However, these methods are computationally expensive and limited by the inherent randomness of generative models, making it difficult to produce stable illusion patterns.

### 2.2. 3D Illusion Arts

Compared to 2D illusion arts, 3D illusion arts involve more complex geometric relationships and must maintain illusion consistency across multiple viewpoints, distances, and lighting conditions, making both design and realization more challenging. Common types include shadow-arts [6, 11, 12, 17], which inverse-design 3D sculptures to match target shadows from specific light directions; wire-arts [10, 15], which construct connected 3D wire networks reproducing line drawings from multiple views; and warp-

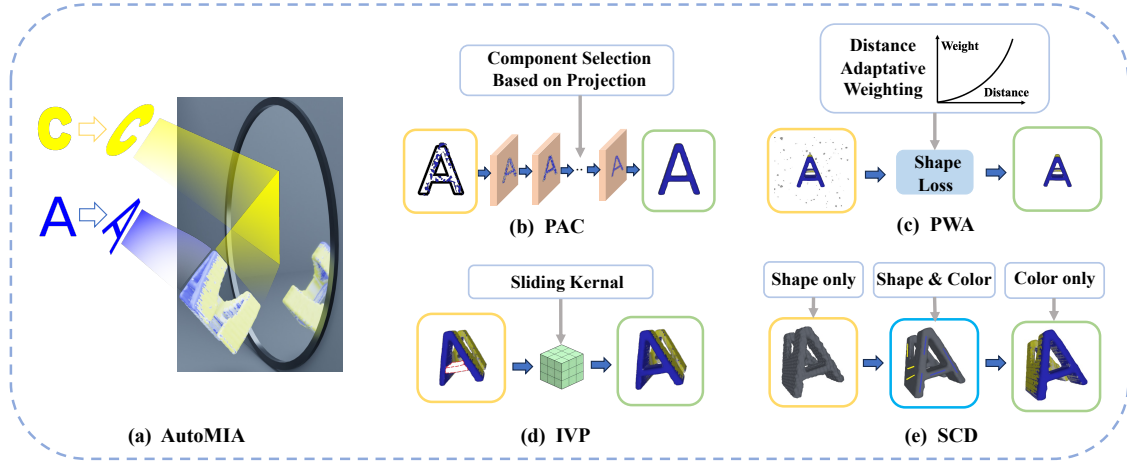


Figure 2. Pipeline of the proposed method. (a) Automatic design of Mirror Illusion Art (AutoMIA). (b) Projection Alignment-Based Component Selection (PAC). (c) Position-Weighted Adaptive Noise Suppression (PWA). (d) Internal Voxel Preservation Mechanism (IVP). (e) Shape-Color Decoupled Optimization (SCD).

arts [2, 4], which create multi-view or anamorphic illusions by learning viewpoint-dependent warps to reveal hidden patterns.

Mirror Illusion Art is an interesting 3D illusion art that optimizes a special 3D object capable of presenting completely different patterns in front of and inside a mirror. Sugihara [18] proposed a topology-based design method for Mirror Shadow Art. However, this approach heavily relies on human intuition and math, and it only supports shape optimization. Although some shadow art methods [12, 17] can be modified to achieve a certain level of effect, these methods also restrict optimization to shape and often result in 3D objects that are not smooth or complete.

### 3. Method

#### 3.1. Problem Formulation and Overview

Our goal is to automatically reconstruct a 3D object  $V$  that exhibits a “mirror illusion” effect, supervised by two given colored images,  $A$  and  $B$ . Specifically, we require the 3D object’s front view ( $C_{\text{direct}}$ ) in front of a mirror appears as  $A$ , while its mirrored view ( $C_{\text{mirror}}$ ) resembles  $B$ . We denote the 3D renderer as  $R$ , and the viewing angles in front of and within the mirror as  $\theta_{\text{direct}}$  and  $\theta_{\text{mirror}}$ , respectively. The similarity function  $\Phi$  measures the resemblance between the rendered view and the target image, incorporating both shape and color similarity. The optimization objective is thus given by:

$$\min_V (\Phi(C_{\text{direct}}, A) + \Phi(C_{\text{mirror}}, B)) = \min_V (\Phi(R(V, \theta_{\text{direct}}), A) + \Phi(R(V, \theta_{\text{mirror}}), B)). \quad (1)$$

The function  $\Phi$  consists of the shape loss function  $L_{\text{shape}}$  and the color loss function  $L_{\text{color}}$ .  $L_{\text{shape}}$  measures the Binary Cross Entropy (BCE) loss between the target pattern mask and the projected mask, while  $L_{\text{color}}$  measures the  $L_1$  loss between the target pattern’s color values and the projected pattern’s color values. Specifically, suppose there are  $N$  pixels on the current projection plane  $F$ . For the  $i$ -th pixel ( $i \in [1, N]$ ), let  $m_i^t$  and  $c_i^t$  denote the mask value and color value of the target pattern, respectively, and  $m_i^r$  and  $c_i^r$  denote the corresponding mask value and color value of the rendered projection. Then

$$L_{\text{shape}} = -\frac{1}{N} \sum_{i=1}^N (m_i^t \cdot \log(m_i^r) + (1 - m_i^t) \cdot \log(1 - m_i^r)), \quad (2)$$

and

$$L_{\text{color}} = -\frac{1}{N} \sum_{i=1}^N |c_i^t - c_i^r|. \quad (3)$$

To facilitate the optimization in 3D space, we build a 3D voxel model of the object  $V$ , represented as

$$V = \{(x_i, \rho_i, c_i)\}, \quad i = 1, 2, \dots, L, \quad (4)$$

where  $L$  denotes the total number of voxels in the 3D object  $V$ ,  $x_i$  is the 3D spatial coordinate of each voxel,  $\rho_i \in [0, 1]$  indicates the density (opacity) of each voxel, and  $c_i$  represents the color of each voxel. We jointly optimize  $x_i$ ,  $\rho_i$ , and  $c_i$  to achieve our goal (Equation 1).

However, during optimization, we identified four main challenges: “surface noise”, “background noise”, “internal fracture”, and “color-shape imbalance”. We address these with four corresponding methods: Projection-Alignment Component (PAC) selection (Section 3.2),




Figure 3. The “surface noise” defect (a) and its mitigation after applying our PAC method (b).

Position-Weighted Adaptive (PWA) suppression (Section 3.3), Internal Voxel Preservation (IVP) (Section 3.4), and Shape-Color Decoupling (SCD) optimization (Section 3.5).

After that, we convert the optimized 3D voxel model into a smoother 3D mesh model and fabricate it using 3D printing (detailed in *Supplementary Material (SM)*). Figure 1 presents the overall pipeline of our method.

### 3.2. Projection Alignment-Based Component Selection

Since the 3D object is constrained by two different view patterns during the reconstruction process, the supervision signal from one angle may introduce noise into the reconstruction from the other angle. This often results in surface voxel noise on the 3D object, which in turn affects the quality of the projected views. We refer to this as the “surface noise”. Figure 3(a) shows an example.

To address this challenge, we propose a projection alignment-based component (PAC) selection method. A component is defined as an independent subset of spatially connected voxels. Two voxels are considered connected if they share a common face. We use a depth-first traversal algorithm to identify all  $K$  components in the current 3D space, denoted as  $S_k, k = 1, 2, \dots, K$ . Next, for any given component  $S_k$ , we render its projections under  $\theta_{\text{direct}}$  and  $\theta_{\text{mirror}}$  to obtain the projection masks  $M_k^{\text{direct}}$  and  $M_k^{\text{mirror}}$ , respectively. We then compute their alignment scores with the supervision pattern masks  $M_A$  and  $M_B$  using the following formula:

$$l_k = \alpha \text{IoU}(M_k^{\text{direct}}, M_A) + \beta \text{IoU}(M_k^{\text{mirror}}, M_B) - \gamma (O(M_k^{\text{direct}}, M_A) + O(M_k^{\text{mirror}}, M_B)). \quad (5)$$

Here,  $\text{IoU}(X, Y) = \frac{|X \cap Y|}{|X \cup Y|}$  measures the overlap between each component’s projection mask and the corresponding supervision mask, while  $O(X, Y) = \frac{|X \cap (1-Y)|}{|X|}$  quantifies the proportion of pixels in  $X$  that fall outside the supervision mask  $Y$ .  $\alpha, \beta$ , and  $\gamma$  are weighting coefficients. We set a threshold  $\tau$  for the alignment score  $l_k$ , and retain only those components with  $l_k \geq \tau$ , removing all others.




Figure 4. The “background noise” defect (a) and its mitigation after applying our PWA method (b).

Through this approach, we iteratively remove surface noise from the 3D object  $V$  while preserving its effective main structure. This results in a smoother object surface. Figure 3 (b) presents an example after denoising.

### 3.3. Position-Weighted Adaptive Noise Suppression

Since the supervision images  $A$  and  $B$  mainly constrain the surface shape and color of the 3D object  $V$ , their ability to suppress external noise voxels far from the object’s surface is limited. As a result, noisy voxels may appear at background locations away from the surface of  $V$ , as illustrated in Figure 4(a). We refer to this defect as the “background noise.”

To address this issue, we propose a position-weighted adaptive (PWA) noise suppression method. The core idea is to assign a distance-adaptive weight to each pixel on the projection plane, thereby reducing the loss contribution from pixels that are far from the target mask. In other words, projections that are farther from the target mask are penalized more heavily during optimization.

Specifically, let  $F$  denote the projection plane in the viewing direction. Both the projection mask  $C$  of the 3D object  $V$  and the shape mask  $M$  of the target image are defined on this plane. For any pixel  $u$  on  $F$ , we let  $d_{\text{max}}(u)$  denote the maximum Euclidean distance from  $u$  to the farthest pixel in  $M$ . Similarly,  $d_{\text{min}}(u)$  denotes the minimum Euclidean distance from  $u$  to the nearest pixel in  $M$ . We then define the distance-adaptive weight  $w(u)$  as follows:

$$w(u) = 1 + (w_{\text{max}} - 1) \cdot \left( \frac{d_{\text{min}}(u)}{d_{\text{max}}(u)} \right)^q \quad (6)$$

Here,  $q$  is a hyperparameter that controls the distance gain, and the maximum value of  $w(u)$  is set to  $w_{\text{max}}$ . Based on Equation 2, the shape optimization loss  $L_{\text{shape}}$  after incorporating the distance-adaptive weighting is given by:

$$L_{\text{shape}} = -\frac{1}{N} \sum_{i=1}^N w(u_i) \cdot (m_i^t \cdot \log(m_i^r)) + (1 - m_i^t) \cdot \log(1 - m_i^r). \quad (7)$$




Figure 5. The “internal fracture” defect (a) and its mitigation after applying our IVP method (b).

In this way, we can effectively penalize projection pixels that are in the background areas, with the penalty increasing as the distance grows. This approach helps eliminate distant external noise. Figure 4(b) presents an example after applying this method.

### 3.4. Internal Voxel Preservation Mechanism

Since the two supervision images A and B mainly constrain the optimization of surface voxels of the 3D object  $V$ , the voxels located inside  $V$  lack direct supervision signals. We observe that some internal voxels may have their densities reduced to zero (i.e., they disappear) during optimization and are difficult to recover. This leads to internal fractures within the 3D object  $V$ , preventing it from forming a complete and connected structure—a defect we refer to as “internal fracture”, as illustrated in Figure 5(a).

To address this challenge, we propose an internal voxel preservation (IVP) mechanism. Specifically, we first define the “solid voxel” as voxels whose density  $\rho(x)$  exceeds a threshold  $\gamma$ . The identification rule is expressed as  $o(x) = 1[\rho(x) > \gamma]$ , where  $o(x) = 1$  indicates that  $x$  is a “solid voxel”.

After that, we design a kernel  $\Omega$  of size  $k \times k \times k$  to determine whether a voxel belongs to the internal region. The core idea is to slide the kernel  $\Omega$  over all voxels in the 3D object. For each voxel at the center of  $\Omega$ , it is classified as an “internal voxel” only if all other voxels within the kernel are “solid voxel”. This rule can be formulated as follows:

$$m(x) = \mathbf{1} \left[ \frac{1}{|\Omega|} \sum_{y \in \Omega(x)} o(y) = 1 \right]. \quad (8)$$

When  $m(x) = 1$ ,  $x$  is identified as an “internal voxel”. Based on this mechanism, we impose a lower bound constraint on the density of all “internal voxel”, requiring their density to be no less than  $\rho_{\min}$ . This approach ensures that the density of “internal voxels” does not decrease to zero, thereby preventing the “internal fracture” defect and maintaining the connectivity and integrity of the 3D object  $V$ . This is also beneficial for subsequent 3D printing. Figure 5(b) presents an example after applying this method.

### 3.5. Shape-Color Decoupled Optimization




Figure 6. The “color-shape imbalance” defect (a) and its mitigation after applying our SCD method (b). The target view is a yellow letter C. In sub-figure (a), although the shape meets the requirement, red color from the other view “leaks” into the current view, resulting in a color inconsistency.

Since our approach jointly optimizes both the surface shape and color of the 3D object, balancing these two objectives is a critical challenge. We observe that the optimization process often exhibits a “color-shape imbalance” defect, where the optimized 3D object  $V$  may satisfy only the shape requirement or only the color requirement, but not both, as illustrated in Figure 6(a).

To address this issue, we propose a shape-color decoupled (SCD) optimization method. Specifically, we divide the optimization timeline  $[0, T]$  into three stages. In the first stage ( $0 \leq t < t_1$ ), SCD optimizes only the shape of the 3D object, ensuring a stable initial structure. In the second stage ( $t_1 \leq t < t_2$ ), SCD jointly optimizes both the shape and color for refinement. In the third stage ( $t_2 \leq t \leq T$ ), SCD optimizes only the color to fine-tune the appearance. The overall loss can thus be expressed as

$$L_{\text{total}} = w_s(t) \cdot L_{\text{shape}} + w_c(t) \cdot L_{\text{color}}, \quad (9)$$

where  $w_s(t)$  and  $w_c(t)$  are the weights that balance shape and color optimization, respectively, and are defined as follows:

$$w_s(t) = \begin{cases} 1, & 0 \leq t < t_2, \\ 0, & t_2 \leq t \leq T. \end{cases} \quad (10)$$

and

$$w_c(t) = \begin{cases} 0, & 0 \leq t < t_1, \\ \lambda, & t_1 \leq t < t_2, \\ 1, & t_2 \leq t \leq T. \end{cases} \quad (11)$$

where  $\lambda \in (0, 1)$  is the weighting factor that balances  $L_{\text{shape}}$  and  $L_{\text{color}}$ . Through this staged optimization approach, we are able to effectively balance the shape and color optimization of the 3D object, resulting in mirror illusion art that meets both shape and color requirements. Figure 6(b) presents an example after applying this method.

## 4. Experiments

### 4.1. Datasets

To show the generality of our method, we collected a wide variety of 2D images as input for our models. These images include English letters, Arabic numerals, Chinese characters, geometric patterns, emoji icons, cartoon designs, and commercial logos, etc. For each category, we randomly sample 200 images, resulting in a total of 1200 images, similar to the settings used in previous works [2, 4, 7, 12, 17]. The images were either generated using Python scripts or collected from the internet, and all internet-sourced images comply with their respective copyright agreements and are used solely for academic research. We name this dataset *Mirror-2D* and will release it soon.

### 4.2. Evaluation Metrics

We employ the following evaluation metrics to systematically assess the 3D generation quality of different methods. Shape Score (SS) evaluates the shape consistency between the 3D object’s projections and the supervision images, ranging from 0 to 1, with higher values indicating greater shape similarity. Color Score (CS) assesses color consistency, ranging from 0 to 1, with lower values indicating greater color similarity. Noise Level (NL) measures the intensity of surface noise in the 3D voxel model, ranging from 0 to 1, with lower values indicating less noise. Smooth Level (SL) evaluates the surface smoothness of the 3D mesh model, ranging from 0 to 1, with higher values indicating a smoother surface. The formula definitions and computation details are provided in the *SM*.

### 4.3. Experimental Settings

To ensure fair comparison, we adopt the same experimental settings for all experiments. The resolution of the 3D voxel model is set to  $128 \times 128 \times 128$ . Rendering is performed using PyTorch3D [16]. See *SM* for more details.

### 4.4. Baselines

We select the methods of Shadow Art (SA) [12] and Shadow Art Revisited (SAR) [17] as baselines. Since both SA and SAR support only shape optimization and do not allow for color optimization, we convert the input 2D color images to 2D masks during optimization, with other settings following Section 4.3. We use only SS, NL, and SL as evaluation metrics when comparing our method with these baselines for the same reason. Since we primarily focus on automated design methods and the results of manual design are highly influenced by individual designer differences, which makes quantitative evaluation challenging, therefore, we do not include the manual design approach [18] as baselines in this work.

Table 1. Comparison with baseline methods

Methods	SL $\uparrow$	NL $\downarrow$	SS $\uparrow$	Time $\downarrow$	Memory $\downarrow$
SA [12]	0.827	0.507	0.499	50s	2.5 GB
SAR [17]	0.834	0.120	0.668	140s	3.3 GB
Ours	<b>0.989</b>	<b>0.049</b>	<b>0.931</b>	76s	2.6 GB

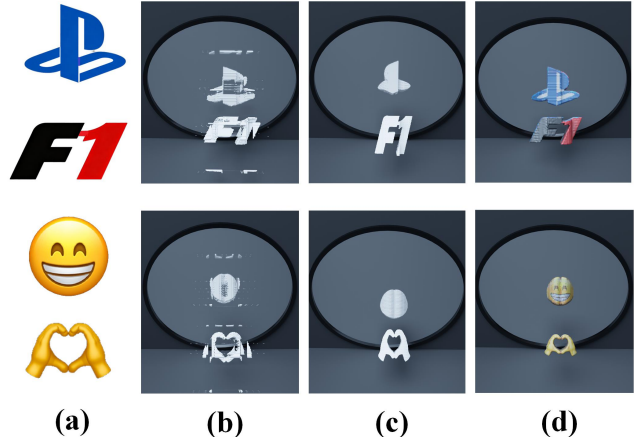


Figure 7. Examples of 3D objects generated by different methods. (a) Input 2D images. (b) Shadow Art (SA). (c) Shadow Art Revisited (SAR). (d) AutoMIA (Ours).

### 4.5. Comparison with Baseline Methods

We evaluated the performance of our method and the baseline methods (SA [12] and SAR [17]) on the *Mirror-2D* dataset. The reconstruction quality was assessed using SS, NL, and SL. The results are shown in Table 1. The results indicate that our method achieved the best performance across all three metrics. A set of visualization examples is presented in Figure 7. As can be seen, the SA and SAR methods not only achieved lower consistency with the target patterns, but also exhibited defects such as surface noise, background noise, and internal fractures. In contrast, our approach produced 3D objects that are closer to the target patterns, with smoother surfaces and less noise, thus achieving the desired *Mirror Illusion Art* effect. In addition to shape reconstruction, our method can also accurately recover the color patterns of the 2D images, whereas the two baseline methods are limited to shape reconstruction only. This highlights a unique advantage of our approach.

The main reasons are as follows. First, SA and SAR are mainly designed for orthogonal settings (e.g. reconstructing 3D objects from 3 orthogonal images) and struggle with non-orthogonal scenarios like *Mirror Illusion Art* (reconstructing 3D objects from 2 non-orthogonal images), where limited supervision leads to more surface noise and unwanted structures. In contrast, our proposed PAC and PWA effectively suppress such noise under non-orthogonal



Figure 8. Digital visualization of Mirror Illusion Arts designed by our AutoMIA method. See *SM* for more examples.

settings, demonstrating their unique advantages. Second, both SA and SAR do not account for internal connectivity, often resulting in internal fractures. In contrast, our IVP method effectively prevents the occurrence of “internal fractures”, ensuring the integrity of the optimized 3D patterns and making them more suitable for physical fabrication. Third, compared to SA and SAR, we introduce color constraints and propose the SCD method to balance shape and color optimization, enabling simultaneous reconstruction of both shape and color.

We further evaluated the average computational time and memory consumption on the same RTX 3090 GPU. The results (Table 1) show that, despite incorporating 4 core modules to enforce physical constraints, AutoMIA achieves similar computational speed and resource usage to SA and is significantly more efficient than SAR, while delivering substantially better reconstruction quality than both baselines. This demonstrates the combined advantages of our method in both efficiency and reconstruction quality.

#### 4.6. Digital Visualization

We constructed a virtual environment for digital visualization of Mirror Illusion Art, based on Blender 4.5. See *SM* for the details. Figure 8 presents several examples of digital Mirror Illusion Art visualizations, with additional examples provided in the *SM*. Interestingly, when viewing these de-

signs for the first time, the mirrored pattern often appears quite different from the pattern seen in front of the mirror, making it seem as if they are two distinct objects rather than a single object and its mirror view. We provide an in-depth discussion of the underlying neuroscience and psychological mechanisms behind this illusion phenomenon in the *SM*.

#### 4.7. Ablation Study

PAC, PWA, IVP, and SCD are the four core modules of our method. We conducted an ablation study to demonstrate their individual contributions, with the results summarized in Table 2. The findings indicate that removing any one of these modules leads to a noticeable decline in 3D reconstruction quality. Specifically, PAC has a significant impact on NL, SS, and CS; PWA mainly affects SL, NL, and SS; IVP primarily influences SS; and SCD shows a clear effect on SL, SS, and CS. Interestingly, adding any module improves all four metrics, even if a module is mainly designed for specific aspects. For example, while PWA primarily targets SL, NL, and SS, it also brings positive effects for CS. This may be because our method effectively balances shape and color optimization, and smoother surfaces with less noise are beneficial for further color refinement. Figures 3-6 provide visual examples to intuitively illustrate the effects of each module.

#### 4.8. Volume Size and Ray Sampling Density

During optimization, the volume size and ray sampling density are two important hyperparameters that directly affect both the quality of the rendered 3D object and the computational cost. We provide a detailed quantitative analysis of these parameters in *SM*.

#### 4.9. Physical Visualization

We selected six representative 3D objects from different categories and fabricated them using 3D printing technology

Table 2. Ablation Study

Methods	SL $\uparrow$	NL $\downarrow$	SS $\uparrow$	CS $\downarrow$
Ours	<b>0.989</b>	<b>0.049</b>	<b>0.931</b>	<b>0.018</b>
– PAC	0.910	0.248	0.629	0.323
– PWA	0.822	0.373	0.790	0.034
– IVP	0.979	0.050	0.748	0.021
– SCD	0.755	0.101	0.548	0.050




Figure 9. Mathematical model for Mirror Illusion Art. (a) Boundary case 1. (b) Boundary case 2.

(See *SM* for details.) The printed objects were placed in front of a circular mirror with a diameter of 16 cm, and the physical demonstration of Mirror Illusion Art follows similar settings of the simulation environment (Section 4.6). The light source was standard indoor lighting, and all photos were taken using an iPhone 13.

Figure 1 shows several physical-world visualization examples. Because our optimization process accounts for various physical constraints (removing surface noise, external noise, and internal fractures), our designs can transition smoothly from digital models to physical realizations. In contrast, many 3D patterns generated by SA and SAR cannot be physically fabricated, as these methods do not consider such physical constraints. Especially for 3D objects with intricate line patterns and hollow interiors, SA and SAR tend to produce internal fractures and noise when handling such complex structures, making 3D printing infeasible. In contrast, our method can successfully handle these cases.

#### 4.10. Discussion

Mirror Illusion Art requires specific viewing angles for the illusion to be perceived. An interesting question is: under what conditions can we observe the Mirror Illusion Art effect? To explore this, we constructed a simple mathematical model, as illustrated in Figure 9. In the figure, A and B represent the front and back of the object, while  $P_0$ ,  $P_1$ , and  $P_2$  denote three vertices of the object.

Next, we consider two boundary cases. In the first case, as shown in Figure 9(a), a light ray emitted from  $P_2$  is reflected by the mirror, passes through  $P_0$ , and reaches the observer’s eye. In this scenario, the observer can exactly see the entire back pattern B via the mirror. Let  $\Theta$  denote the observer’s viewing angle and  $\theta_1$  represent the viewing angle for this first case. If  $\Theta$  is less than  $\theta_1$ , the observer will not be able to see the complete back pattern B via mirror reflection. In the second case, as illustrated in Figure 9(b), a light ray from  $P_2$  travels along the surface of the back pattern B, passes through  $P_0$ , and reaches the eye di-




Figure 10. Typical failure cases.

rectly without reflection. In this situation, the observer can directly see the back pattern B, causing the illusion effect to disappear. Let  $\theta_2$  denote the viewing angle for this second case. To observe the ideal mirror illusion art, the observer’s viewing angle  $\Theta$  should be less than  $\theta_2$ . In summary, the observer’s viewing angle should satisfy:

$$\theta_1 < \Theta < \theta_2 \quad (12)$$

where  $\theta_1$  and  $\theta_2$  are determined by the shape of the object as well as the relative distances between the observer, the object, and the mirror. When this condition is met, the ideal mirror illusion art can be perceived.

#### 4.11. Limitations

We present AutoMIA for generating diverse and compelling mirror illusion art with low time and computational cost. However, there are two limitations. First, supervision images A and B should have similar widths to avoid conflicting signals and incomplete patterns (Fig. 10(a)), which can be easily resolved by scaling the images to the same width. Second, the image resolution should not exceed the voxel model’s resolution, otherwise, high-frequency noise and visual artifacts may occur (Fig. 10(b)). This can be solved by using higher-resolution voxel models (with increased cost) or down-sampling images.

#### 5. Conclusion

In this work, we propose AutoMIA, an automated Mirror Illusion Art design pipeline that jointly optimizes both shape and color to generate compelling mirror illusion artworks. To stabilize the optimization and suppress artifacts, we developed the PAC method for reducing surface noise, PWA suppression for background noise, the IVP mechanism to prevent internal fractures, and SCD optimization to balance shape and color optimization. Experimental results demonstrate that AutoMIA successfully produces a wide variety of mirror illusion art in both digital and physical forms, with superior reconstruction quality, fast generation speed, and low computational resource requirements.

## 6. Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. U2341228 and No. 62576187).

## References

- [1] 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](#)
- [2] Pascal Chang, Sergio Sancho, Jingwei Tang, Markus Gross, and Vinicius Azevedo. Lookingglass: Generative anamorphoses via laplacian pyramid warping. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24–33, 2025. [3](#), [6](#)
- [3] Hung-Kuo Chu, Wei-Hsin Hsu, Niloy J Mitra, Daniel Cohen-Or, Tien-Tsin Wong, and Tong-Yee Lee. Camouflage images. *ACM Trans. Graph.*, 29(4):51–1, 2010. [2](#)
- [4] Yue Feng, Vaibhav Sanjay, Spencer Lutz, Badour AlBahar, Songwei Ge, and Jia-Bin Huang. Illusion3d: 3d multiview illusion with 2d diffusion priors. *arXiv preprint arXiv:2412.09625*, 2024. [3](#), [6](#)
- [5] William T Freeman, Edward H Adelson, and David J Heeger. Motion without movement. *ACM Siggraph Computer Graphics*, 25(4):27–30, 1991. [2](#)
- [6] Aalok Gangopadhyay, Prajwal Singh, Ashish Tiwari, and Shanmuganathan Raman. Hand shadow art: A differentiable rendering perspective. *arXiv preprint arXiv:2505.21252*, 2025. [2](#)
- [7] Daniel Geng, Inbum Park, and Andrew Owens. Visual anagrams: Generating multi-view optical illusions with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24154–24163, 2024. [2](#), [6](#)
- [8] Rui Guo, Jasmine Collins, Oscar de Lima, and Andrew Owens. Ganmouflage: 3d object nondetection with texture fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4702–4712, 2023. [2](#)
- [9] Aaron Hertzmann. Visual indeterminacy in gan art. In *ACM SIGGRAPH 2020 Art Gallery*, pages 424–428, 2020. [2](#)
- [10] Kai-Wen Hsiao, Jia-Bin Huang, and Hung-Kuo Chu. Multi-view wire art. *ACM Trans. Graph.*, 37(6):242, 2018. [2](#)
- [11] Sehee Min, Jaedong Lee, Jungdam Won, and Jehee Lee. Soft shadow art. In *Proceedings of the symposium on Computational Aesthetics*, pages 1–9, 2017. [2](#)
- [12] Niloy J Mitra and Mark Pauly. Shadow art. *ACM Trans. Graph.*, 28(5):156, 2009. [2](#), [3](#), [6](#)
- [13] Aude Oliva, Antonio Torralba, and Philippe G Schyns. Hybrid images. *ACM Transactions on Graphics (TOG)*, 25(3): 527–532, 2006. [2](#)
- [14] Andrew Owens, Connelly Barnes, Alex Flint, Hanumant Singh, and William Freeman. Camouflaging an object from many viewpoints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2782–2789, 2014. [2](#)
- [15] Zhiyu Qu, Lan Yang, Honggang Zhang, Tao Xiang, Kaiyue Pang, and Yi-Zhe Song. Wired perspectives: Multi-view wire art embraces generative ai. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6149–6158, 2024. [2](#)
- [16] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020. [6](#)
- [17] Kaustubh Sadekar, Ashish Tiwari, and Shanmuganathan Raman. Shadow art revisited: a differentiable rendering based approach. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 29–37, 2022. [2](#), [3](#), [6](#)
- [18] Kokichi Sugihara. Topology-disturbing objects: A new class of 3d optical illusion. *Journal of Mathematics and the Arts*, 12(1):2–18, 2018. [2](#), [3](#), [6](#)
- [19] Qing Zhang, Gelin Yin, Yongwei Nie, and Wei-Shi Zheng. Deep camouflage images. In *Proceedings of the AAAI conference on artificial intelligence*, pages 12845–12852, 2020. [2](#)