Supplementary Material for Mirror Illusion Art

Xiaopei Zhu[1 *] Zeyuan Li[2 *] Jun Zhu[1,3] Xiaolin Hu[1,3,4 †]
[1]Department of Computer Science and Technology, BNRist, Tsinghua University
[2]Huazhong University of Science and Technology
[3]IDG/McGovern Institute for Brain Research, Tsinghua University
[4]Chinese Institute for Brain Research (CIBR)

1. Supplementary Video

See "Demo Video.mp4" for a full-angle visualization of mirror illusion 3D objects.

2. From Voxel to Mesh

Although voxel representation is convenient for 3D optimization, we convert the optimized 3D voxel model V into a 3D mesh model to obtain a smoother representation suitable for 3D printing. To achieve this, we first use the classic Marching Cubes [4] algorithm, which extracts an isosurface from the voxel density field to generate the mesh. However, applying this method alone often leaves jagged artifacts on the mesh surface. To address this, we further apply the Taubin smoothing [6] technique. For any vertex y on the mesh, let $\bar{y}$ denote the mean position of all its neighboring vertices; then, the smoothed vertex coordinate y' after Taubin smoothing is given by:

$$y' = (1 - \xi) \cdot y + \xi \cdot \bar{y}. \tag{1}$$

Here, ξ is a hyperparameter that controls the degree of smoothing. Through this approach, we obtain a much smoother 3D mesh suitable for 3D printing.

3. From Digital to Physical

To transform our design into a real-world artwork, we employ 3D printing technology to fabricate the optimized 3D mesh model. We use the Mimaki 3DUJ 553 full-color 3D printer, which utilizes UV-cured inkjet printing technology. This device forms objects layer by layer by jetting microdroplets of liquid photosensitive resin and instantly curing them with ultraviolet light, directly producing physical objects with textures and colorsm (no post-processing is required). The printing material is a water-soluble resin, which can be washed away after printing, facilitating the formation of complex structures. The layer thickness is set to approximately 20 μm, and the average model height is about 7 cm per piece. For each batch, we print five 3D models simultaneously, with an average printing time of about 11 hours.

4. Evaluation Metrics

We employ the following evaluation metrics to systematically assess the 3D generation quality of Mirror Shadow Art, covering shape consistency, color consistency, surface noise intensity, and surface smoothness.

Shape Score (SS) is used to evaluate the shape consistency between the 3D object's projections and the supervision images. We denote M^{direct} and M^{mirror} as the projection masks of the 3D object in the direct and mirror directions, respectively, and M_A and M_B as the corresponding supervision pattern masks. The Shape Score is defined as follows:

$$SS = \frac{\text{IoU}(M^{\text{direct}}, M_A) + \text{IoU}(M^{\text{direct}}, M_A)}{2}. \tag{2}$$

Color Score (CS) is used to evaluate the color consistency between the 3D object's projections and the supervision images. Based on Equation 3 of the main paper, we compute the average L1 loss between the projected color and the reference image color for the front view (C_{direct}) and the mirrored view (C_{mirror}), denoted as $L_{\text{color}}^{\text{direct}}$ and $L_{\text{color}}^{\text{mirror}}$, respectively. The Color Score can then be calculated as follows:

$$CS = \frac{L_{\text{color}}^{\text{direct}} + L_{\text{color}}^{\text{mirror}}}{2}. \tag{3}$$

Noise Level (NL) is used to evaluate the intensity of surface noise in the 3D voxel model. Let ρ_{origin} denote the original density distribution of the 3D object. We apply a 3D Gaussian filter G_σ with kernel radius σ to obtain the smoothed density distribution $\rho_{\text{Gauss}} = G_\sigma(\rho_{\text{origin}})$. The surface noise distribution R is then calculated as the difference between the original and smoothed densities, i.e.,

*Equal contribution.
†Corresponding author.

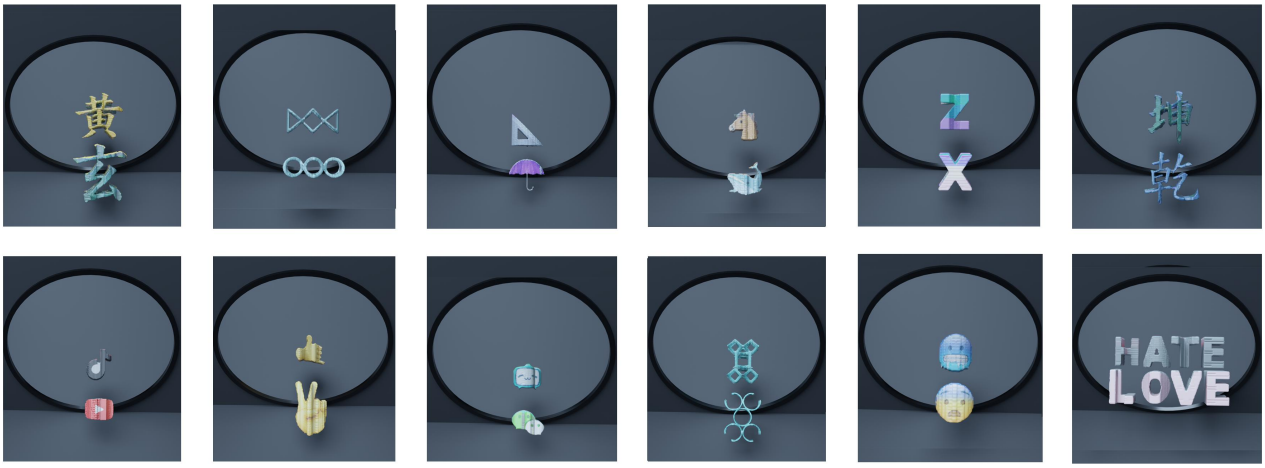

Figure S1. More Digital visualization Examples of Mirror Illusion Arts designed by our AutoMIA method.

$R = \rho_{\text{origin}} - \rho_{\text{Gauss}}$. The Noise Level is computed as follows:

$$NL = \frac{\|R\|_2^2}{\|\rho_{origin}\|_2^2}. \tag{4}$$

Smooth Level (SL) is used to evaluate the surface smoothness of the 3D mesh model. Suppose the normal vectors of any two adjacent faces on the 3D mesh surface are n_1 and n_2, respectively, then the cosine of the angle between them is given by $n_1 \cdot n_2$. The Smooth Level can be calculated as follows:

$$SL = \text{Mean}(\frac{n_1 \cdot n_2 + 1}{2}). \tag{5}$$

Here, Mean denotes the average taken over all pairs of adjacent faces.

SS ranges from 0 to 1, with higher values indicating greater shape similarity. CS ranges from 0 to 1, with lower values indicating greater color similarity. NL ranges from 0 to 1, with lower values reflecting less surface noise. SL ranges from 0 to 1, with higher values indicating a smoother mesh surface.

5. 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 [5], with the light source configured in direct illumination mode and the Grid Raysampler used as the ray sampler, where each ray samples 150 points. During rendering, the virtual viewing angle is randomly sampled between 22.5° and 60°, which aligns with the real-world physical setup. The Adam [2] optimizer is used for optimization, with 1000 epochs and a learning rate of 0.05.

6. Digital Visualization

We constructed a virtual environment for digital visualization of Mirror Illusion Art, based on Blender 4.5. First, we created a virtual circular mirror with length, width, and height ratios of 12.5:12.5:1.3. The mirror material was set to a metallic value of 1, roughness of 0.04, refractive index of 1.5, and opacity of 1.0. Next, the 3D objects optimized by the AutoMIA method were imported into Blender and positioned in front of the virtual mirror. The virtual camera's viewing angle was then adjusted to observe the Mirror Illusion effect, with specific requirements for camera angles discussed in Section 4.10. We used Blender's default renderer, setting the denoising filter threshold to 0.1 and the maximum number of samples to 2048.

Figure S1 provides more visualization examples based on this digital environment. All of the examples are generated by our AutoMIA method.

7. Underlying Mechanisms of Mirror Illusion Art

Let's delve into the underlying mechanisms from a neuroscience and psychology perspective. According to the Bayesian theory of perception [3], the human visual system operates with prior experiences; in other words, our

Table S1. The effect of volume size

Volume Size	64	128	256
SL ↑	0.979	0.989	0.977
NL ↓	0.059	0.050	0.021
SS ↑	0.929	0.931	0.972
CS ↑	0.019	0.018	0.018
Time	61s	76s	185s
Memory	1978MB	2672MB	6952MB

Table S2. The effect of ray sampling density

Volume Size	80	115	150	185	220
SL ↑	0.935	0.984	0.989	0.991	0.974
NL ↓	0.066	0.057	0.050	0.060	0.050
SS ↑	0.917	0.915	0.931	0.908	0.684
CS ↑	0.102	0.020	0.018	0.018	0.267
Time	71s	72s	76s	78s	80s
Memory	2300MB	2520MB	2672MB	3036MB	3904MB

brain has an automatic completion mechanism. We first predict what we expect to see based on past experience and then update our perception using new sensory input. When a contour from a particular viewpoint resembles a familiar pattern, we tend to "fill in" the expected object in our mind. The Biased Competition Principle [1] suggests that our brain generates multiple possible interpretations for the same scene, while attention selectively enhances certain candidates, allowing them to dominate perception. Then what is the most likely interpretation? According to the Law of Prägnanz [7], among all possible interpretations after this mental completion, our brain prefers the simplest, most symmetrical, and most stable patterns, similar to the principle of Occam's Razor. This explains why we naturally expect the pattern in the mirror to be symmetrical and consistent with the one in front of the mirror. However, Mirror Illusion Art deliberately violates this intuitive expectation, creating a perceptual effect that surprises and challenges our usual ways of interpreting mirrored scenes.

8. Volume Size

During the optimization process, the volume size of the voxel model determines the level of detail in the reconstructed patterns, but also affects optimization time and memory usage. To quantitatively assess the impact of volume size on both reconstruction quality and computational cost, we tested three typical volume sizes: 64, 128, and 256. Following the experimental settings in Section 6, the results are shown in Table S1. The findings indicate that lower resolutions ($64 \times 64 \times 64$) reduce optimization time and memory consumption, but at the expense of reconstruction quality, while higher resolutions ($256 \times 256 \times 256$) offer only marginal improvements in quality with significantly increased computational costs. Therefore, in our experiments, the volume size is set to 128 to balance reconstruction quality and optimization efficiency.

9. Ray Sampling Density

An important hyperparameter in the 3D rendering process is the ray sampling density, which directly affects both the quality of the rendered 3D object and the computational cost. To further quantify its impact, we selected five representative ray sampling densities: 80, 115, 150, 185, and 220. All other experimental settings followed Section 6. The results are shown in Table S2. Overall, increasing the sampling density improves the quality of the rendered 3D object, but also increases computation time and memory consumption. We also observed that excessively high sampling density can even reduce some reconstruction metrics, such as SL and SS. Therefore, to balance reconstruction quality and computational efficiency, we set the sampling density to 150 in our experiments.

(a) **(b)**

Figure S2. 3-fold (non-orthogonal) mirror illusion art generated by (a) AutoMIA (Ours) and (b) Shadow Art (SA). Angle between two mirrors is 110°. Note that the SA-generated object has many artifacts (noise) on the surface and in the background.

10. Three-Fold Mirror Illusion Art

We further show that our method supports 3D reconstruction from three non-orthogonal views (3-fold images) (Figure S2), producing richer colors, smoother surfaces, and less noise than Shadow Art—beneficial for high-fidelity reconstruction and physical printing. This further suggests that AutoMIA opens the door to a broader range of innovative applications beyond its current scope.

References

[1] Diane M Beck and Sabine Kastner. Top-down and bottom-up mechanisms in biasing competition in the human brain. *Vision research*, 49(10):1154–1165, 2009. 3

[2] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 2

[3] David C Knill and Whitman Richards. *Perception as Bayesian inference*. Cambridge University Press, 1996. 2

[4] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. 1

[5] 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. 2

[6] Gabriel Taubin. Curve and surface smoothing without shrinkage. In *Proceedings of IEEE international conference on computer vision*, pages 852–857. IEEE, 1995. 1

[7] Eline Van Geert and Johan Wagemans. Prägnanz in visual perception. *Psychonomic Bulletin & Review*, 31(2):541–567, 2024. 3