LookingGlass: Generative Anamorphoses via Laplacian Pyramid Warping

Pascal Chang*
ETH Zürich
Zürich, Switzerland
Disney Research Studios
Zürich, Switzerland
pascal.chang@disneyresearch.com

Sergio Sancho*
ETH Zürich
Zürich, Switzerland
Disney Research Studios
Zürich, Switzerland
sergio.sancho@disneyresearch.com

Jingwei Tang
Disney Research Studios
Zürich, Switzerland
jingwei.tang@disneyresearch.com

Vinicius Azevedo
Disney Research Studios
Zürich, Switzerland
vinicius.azevedo@disneyresearch.com

Ugaso Sheikh-Abdi
StudioLAB, The Walt Disney Company
Berkeley, USA
ugaso.sheikabdi@disney.com

Markus Gross
ETH Zürich
Zürich, Switzerland
Disney Research Studios
Zürich, Switzerland
grossm@inf.ethz.ch

Figure 1: LookingGlass is an interactive experience that leverages diffusion models to create ambiguous anamorphoses—images that reveal a hidden form when viewed through mirrors or lenses. Participants explore these illusions firsthand, experiment with different setups, and generate their own, which they can view later with a mini cylindrical mirror provided as a giveaway.

Abstract

We present *LookingGlass*, an interactive experience that gives a generative twist to classical anamorphoses—illusions that reveal their true form only from a specific viewpoint or through a mirror or lens. While these distortions date back to the 17th century [Niceron 1638], traditional anamorphic images lose meaning when viewed directly. Our method uses latent diffusion models to create ambiguous anamorphoses—images that remain interpretable in both distorted and undistorted forms. Central to this is Laplacian Pyramid Warping, a frequency-aware image warping technique for high-quality synthesis. Differing from previous works, we extend generative perceptual illusions to latent space models with broader spatial transformations while maintaining image quality. *LookingGlass* invites SIGGRAPH attendees to explore these illusions firsthand, blending geometry, optics, and AI creativity.

*Both authors contributed equally to this manuscript.

SIGGRAPH Emerging Technologies '25, Vancouver, BC, Canada

ACM ISBN 979-8-4007-1551-8/25/08
https://doi.org/10.1145/3721257.3734031

ACM Reference Format:
Pascal Chang, Sergio Sancho, Jingwei Tang, Vinicius Azevedo, Ugaso Sheikh-Abdi, and Markus Gross. 2025. LookingGlass: Generative Anamorphoses via Laplacian Pyramid Warping. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Emerging Technologies (SIGGRAPH Emerging Technologies '25), August 10–14, 2025, Vancouver, BC, Canada.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3721257.3734031

1 Introduction

Anamorphoses are optical illusions where images appear distorted until viewed from a specific angle or through a mirror or lens. Dating back to the 1600s, J.-F. Nicéron's *La Perspective Curieuse* laid the groundwork for their systematic construction. Traditional anamorphic images lose meaning when viewed directly; we explore a more challenging task—creating ambiguous anamorphoses that remain interpretable in both forms. Recent computational methods have pushed the boundaries of optical illusions in both 2D and 3D, from hybrid images that shift with viewing distance [Oliva et al. 2006] to 3D wire art that produces distinct projections [Hsiao et al. 2018]. More recently, generative models have been explored for creating optical illusions, with two notable works closely related to ours. Visual Anagrams [Geng et al. 2024] introduced a generative approach for ambiguous images by aligning diffusion paths across multiple views, enabling the creation of 2D images that change appearance under simple transformations like flipping or 90-degree rotation. However, their method is constrained to pixel-space diffusion and

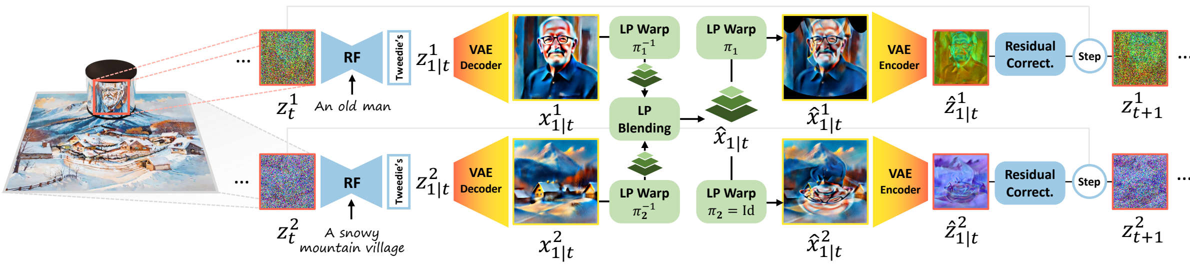

Figure 2: LookingGlass Pipeline Overview. At each denoising step, the predicted final image is derived from the network's noise estimate and decoded into image space. Warping and view aggregation are then applied in image space before re-encoding into latent space for the next diffusion step. Our novel Laplacian Pyramid Warping maintains high-frequency details.

limited to basic orthogonal transformations. In contrast, Diffusion Illusions [Burgert et al. 2024] can generate more complex illusions but suffer from artifacts common to Score Distillation Sampling methods, along with high memory and computational demands inherent to their optimization-based nature. In view of previous work, we create high-quality ambiguous anamorphic images using latent diffusion models in a single denoising pass. Our contributions are twofold: (1) We adapt latent diffusion and flow models for artifact-free anamorphic image synthesis, improving quality and accessibility. (2) We introduce Laplacian Pyramid Warping, a frequency-aware technique that enables complex spatial transformations while preserving high-frequency details. This allows us to generate intricate illusions involving reflective and refractive surfaces with minimal quality loss and improved expressiveness.

2 Method

Our method differs from previous work in two key aspects. First, unlike Visual Anagrams, which is limited to orthogonal transformations, our approach supports arbitrary warpings, including those induced by curved mirrors. Second, while Visual Anagrams only work in pixel space, our method can operate in latent space, making it compatible with models like Stable Diffusion 3.5, which improves generation quality and versatility. Compared to Diffusion Illusions, our approach avoids Score Distillation Sampling, reducing artifacts while being faster and more memory-efficient, as it does not require optimization. Our pipeline is illustrated in Figure 2.

2.1 Extending to Latent Space

Tancik [2023] generates multi-view illusions with Stable Diffusion 1.5 [Rombach et al. 2022] using a similar pipeline to Geng et al. [2024]. Images are transformed from a secondary view (through the lens or mirror in our case) to the main view space in latent space, and decoded to the clean image after the denoising process. The results contain visual artifacts due to the fact that VAEs for latent diffusion and flow models are generally not trained to be equivariant: when the latent images are deformed, the corresponding decoded image does not necessarily share the same transformation. To tackle this issue, we propose to keep all image transformations in image space. This is done by decoding the latent predicted clean image to pixel space before doing the image transformation, and encoding to latent again after transforming back. An additional correction term corrects any reconstruction error caused by the VAE.

2.2 Laplacian Pyramid Warping

While the modifications above enable latent diffusion models to generate ambiguous images with arbitrary 2D transformations, anamorphosis poses additional challenges. First, the anamorphic view often reveals only part of the image, leading to visible seams. Second, transforms like stretching degrade high-frequency details.

The key issue lies in using averaging as the aggregation method for different views. Transforming multiple views to a canonical frame can induce mismatched frequency components, some might map from a small subset of pixels, while others undergo extreme stretching. Averaging these inconsistently scaled views fails to account for frequency discrepancies. To address this, we propose blending views using Laplacian pyramids [Burt and Adelson 1987].

In addition, we introduce Laplacian Pyramid Warping (LPW), a novel image warping technique. Rather than warping all pixels uniformly, LPW adapts pixel placement based on local stretching in the mapping, assigning each pixel to an appropriate level in the Laplacian pyramid after warping. Lower-resolution levels retain less detail but influence a larger image region, enabling adaptive blending between the inverse-warped mirror view and the main view by aligning frequency bands accordingly.

3 Experience

At our interactive booth, attendees will experience generative illusions firsthand. They will begin by exploring cylindrical and conic mirror anamorphoses, using physical mirrors to reveal hidden images by positioning them correctly on a series of images. Next, they will discover the Nicéron lens illusion by looking through a tube with a specialized lens. Finally, an interactive demo will allow them to generate their own illusions by prompting the model, which takes 1–2 minutes on an NVIDIA RTX 4090 GPU. The resulting images will be downloadable to their phones via a QR code. As a giveaway, attendees will receive mini cylindrical mirrors to view their creations at home.

4 Conclusion & Future Work

We present LookingGlass, a training-free approach for generating high-quality ambiguous anamorphoses using latent diffusion models, improving upon previous work by enhancing efficiency and image quality. Future directions could explore automating the selection of effective prompt pairs for greater visual coherence.

References

Ryan Burgert, Xiang Li, Abe Leite, Kanchana Ranasinghe, and Michael Ryoo. 2024. Diffusion Illusions: Hiding Images in Plain Sight. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) *(SIGGRAPH '24)*. Association for Computing Machinery, New York, NY, USA, Article 131, 11 pages. doi:10.1145/3641519.3657500

Peter J Burt and Edward H Adelson. 1987. The Laplacian pyramid as a compact image code. In *Readings in computer vision*. Elsevier, 671–679.

Daniel Geng, Inbum Park, and Andrew Owens. 2024. Visual Anagrams: Generating Multi-View Optical Illusions with Diffusion Models. In *Conference on Computer Vision and Pattern Recognition (CVPR).* https://arxiv.org/abs/2311.17919

Kai-Wen Hsiao, Jia-Bin Huang, and Hung-Kuo Chu. 2018. Multi-view wire art. *ACM Trans. Graph.* 37, 6, Article 242 (Dec. 2018), 11 pages. doi:10.1145/3272127.3275070

Jean François Niceron. 1638. *La perspective curieuse.* P. Billaine, Paris.

Aude Oliva, Antonio Torralba, and Philippe G. Schyns. 2006. Hybrid images. In *ACM SIGGRAPH 2006 Papers* (Boston, Massachusetts) *(SIGGRAPH '06)*. Association for Computing Machinery, New York, NY, USA, 527–532. doi:10.1145/1179352.1141919

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10684–10695.

Matthew Tancik. 2023. Illusion Diffusion: optical illusions using stable diffusion. https://github.com/tancik/Illusion-Diffusion.