


Image Editing with Gaussian Splatting

Originally published at <https://www.unite.ai/image-editing-with-gaussian-splatting/>



A new collaboration between researchers in Poland and the UK proposes the prospect of using [Gaussian Splatting](#) to edit images, by temporarily interpreting a selected part of the image into 3D space, allowing the user to modify and manipulate the 3D representation of the image, and then applying the transformation.



To change the orientation of the cat's head, the relevant section is moved into 3D space via Gaussian Splatting, and then manipulated by the user. The modification is then applied. The process is analogous to various modal techniques in Adobe software, that lock off the interface until a current complex process is completed. Source: <https://github.com/waczjoan/MiraGe/>

Since the Gaussian Splat element is temporarily represented by a mesh of triangles, and momentarily enters a 'CGI state', a physics engine integrated into the process can interpret natural movement, either to change the static state of an object, or to produce an animation.



A physics engine incorporated into the new MiraGe system can perform natural interpretations of physical movement, either for animations or static alterations to an image.

There is no generative AI involved in the process, meaning that no [Latent Diffusion Models](#) (LDMs) are involved, unlike Adobe's [Firefly system](#), which is trained on Adobe Stock (formerly Fotolia).

The system – called *MiraGe* – interprets selections into 3D space and infers geometry by creating a *mirror image* of the selection, and approximating 3D coordinates that can be embodied in a Splat, which then interprets the image into a mesh.




CLICK TO PLAY VIDEO ON SITE

Click to play. Further examples of elements that have been either altered manually by a user of the MiraGe system, or subject to physics-based deformation.

The authors compared the MiraGe system to former approaches, and found that it achieves state-of-the-art performance in the target task.

Users of the zBrush modeling system will be familiar with this process, since zBrush allows the user to essentially [‘flatten’ a 3D model](#) and add 2D detail, while preserving the underlying mesh, and interpreting the new detail into it – a ‘freeze’ that is the opposite of the MiraGe method, which operates more like Firefly or other Photoshop-style modal manipulations, such as warping or crude 3D interpretations.



Parametrized Gaussian Splats allow MiraGe to create high-quality reconstructions of selected areas of a 2D image, and apply soft-body physics to the temporarily-3D selection.

The paper states:

‘[We] introduce a model that encodes 2D images by simulating human interpretation. Specifically, our model perceives a 2D image as a human would view a photograph or a sheet of paper, treating it as a flat object within a 3D space.

‘This approach allows for intuitive and flexible image editing, capturing the nuances of human perception while enabling complex transformations.’


The [new paper](#) is titled *MiraGe: Editable 2D Images using Gaussian Splatting*, and comes from four authors across Jagiellonian University at Kraków, and the University of Cambridge. The full code for the system has been [released at GitHub](#).

Let’s take a look at how the researchers tackled the challenge.

Method

The MiraGe approach utilizes [Gaussian Mesh Splatting](#) (GaMeS) parametrization, a technique developed by a group that includes two of the authors of the new paper. GaMeS allows Gaussian Splats to be interpreted as traditional CGI meshes, and to become subject to the standard range of warping and modification techniques that the CGI community has developed over the last several decades.

MiraGe interprets ‘flat’ Gaussians, in a 2D space, and uses GaMeS to ‘pull’ content into GSplat-enabled 3D space, temporarily.



Each flat Gaussian is represented as three points in a cloud of triangles, called 'triangle soup', opening up the inferred image to manipulation.
Source: <https://arxiv.org/pdf/2410.01521>


We can see in the lower-left corner of the image above that MiraGe creates a 'mirror' image of the section of an image to be interpreted.

The authors state:

'[We] employ a novel approach utilizing two opposing cameras positioned along the Y axis, symmetrically aligned around the origin and directed towards one another. The first camera is tasked with reconstructing the original image, while the second models the mirror reflection.'

'The photograph is thus conceptualized as a translucent tracing paper sheet, embedded within the 3D spatial context. The reflection can be effectively represented by horizontally flipping the [image]. This mirror-camera setup enhances the fidelity of the generated reflections, providing a robust solution for accurately capturing visual elements.'

The paper notes that once this extraction has been achieved, perspective adjustments that would typically be challenging become accessible via direct editing in 3D. In the example below, we see a selection of an image of a woman that encompasses only her arm. In this instance, the user has tilted the hand downward in a plausible manner, which would be a challenging task by just pushing pixels around.



An example of the MiraGe editing technique.

Attempting this using the Firefly generative tools in Photoshop would usually mean that the hand becomes replaced by a synthesized, diffusion-imagined hand, breaking the authenticity of the edit. Even the more capable systems, such as the [ControlNet](#) ancillary system for Stable Diffusion and other Latent Diffusion Models, [such as Flux](#), struggle to achieve this kind of edit in an image-to-image pipeline.

This particular pursuit has been dominated by methods using Implicit Neural Representations (INRs), such as [SIREN](#) and [WIRE](#). The difference between an implicit and explicit representation method is that the coordinates of the model are not directly addressable in INRs, which use a [continuous function](#).

By contrast, Gaussian Splatting offers explicit and addressable X/Y/Z [Cartesian coordinates](#), even though it uses Gaussian ellipses rather than [voxels](#) or other methods of depicting content in a 3D space.


The idea of using GSplat in a 2D space has been most prominently presented, the authors note, in the 2024 Chinese academic collaboration [GaussianImage](#), which offered a 2D version of Gaussian Splatting, enabling inference frame rates of 1000fps. However, this model has no implementation related to image editing.

After GaMeS parametrization extracts the selected area into a Gaussian/mesh representation, the image is reconstructed using the Material Points Method (MPM) technique first outlined in a [2018 CSAIL paper](#).

In MiraGe, during the process of alteration, the Gaussian Splat exists as a guiding proxy for an equivalent mesh version, much as [3DMM CGI models](#) are [frequently used](#) as orchestration methods for implicit neural rendering techniques such as Neural Radiance Fields (NeRF).

In the process, two-dimensional objects are modeled in 3D space, and the parts of the image that are not being influenced are not visible to the end user, so that the contextual effect of the manipulations are not apparent until the process is concluded.

MiraGe can be integrated into the popular open source 3D program [Blender](#), which is now [frequently used](#) in AI-inclusive workflows, primarily for image-to-image purposes.



A workflow for MiraGe in Blender, involving the movement of the arm of a figure depicted in a 2D image.

The authors offer two versions of a deformation approach based on Gaussian Splatting – *Amorphous* and *Graphite*.

The Amorphous approach directly utilizes the GaMeS method, and allows the extracted 2D selection to move freely in 3D space, whereas the Graphite approach constrains the Gaussians to 2D space during initialization and training.

The researchers found that though the Amorphous approach might handle complex shapes better than Graphite, 'tears' or rift artefacts were more evident, where the edge of the deformation aligns with the unaffected portion of the image*.

Therefore, they developed the aforementioned 'mirror image' system:

[We] employ a novel approach utilizing two opposing cameras positioned along the Y axis, symmetrically aligned around the origin and directed towards one another.

The first camera is tasked with reconstructing the original image, while the second models the mirror reflection. The photograph is thus conceptualized as a translucent tracing paper sheet, embedded within the 3D spatial context. The reflection can be effectively represented by horizontally flipping the [image].

This mirror-camera setup enhances the fidelity of the generated reflections, providing a robust solution for accurately capturing visual elements.'

The paper notes that MiraGe can use external physics engines such as those [available in Blender](#), or in [Taichi Elements](#).

Data and Tests

For image quality assessments in tests carried out for MiraGe, the [Signal-to-Noise Ratio](#) (SNR) and [MS-SIM](#) metrics were used.

Datasets used were the [Kodak Lossless True Color Image Suite](#), and the [DIV2K validation](#) set. The resolutions of these datasets suited a comparison with the closest prior work, Gaussian Image. The other rival frameworks trialed were SIREN, WIRE, NVIDIA's  **NVDA 1.00%** [Instant Neural Graphics Primitives](#) (I-NGP), and [NeuRBF](#).

The experiments took place on a NVIDIA GEFORCE RTX 4070 laptop and on a NVIDIA RTX 2080.

	Kodak dataset		DIV2K dataset	
	PSNR ↑	MS-SSIM ↑	PSNR ↑	MS-SSIM ↑
WIRE (Saragadam et al. [2023a])	41.47	0.9939	35.64	0.9511
SIREN (Sitzmann et al. [2020a])	40.83	0.9960	39.08	0.9958
I-NGP (Müller et al. [2022])	43.88	0.9976	37.06	0.9950
NeuRBF (Chen et al. [2023])	43.78	0.9964	38.60	0.9913
3DGS (Kerbl et al. [2023])	43.69	0.9991	39.36	0.9979
GaussianImage-70k (Zhang et al. [2024])	44.08	0.9985	39.53	0.9975
GaussianImage-100k* (Zhang et al. [2024])	38.93	0.9948	41.48	0.9981
MiraGe-70k (our)	57.41	0.9998	53.22	0.9996
MiraGe-100k (our)	59.52	0.9999	54.54	0.9998

MiraGe offers state-of-the-art results against the chosen prior frameworks, according to the results featured in the new paper.

Of these results, the authors state:

'We see that our proposition outperforms the previous solutions on both datasets. The quality measured by both metrics shows significant improvement compared to all the previous approaches.'

Conclusion

MiraGe's adaptation of 2D Gaussian Splatting is clearly a nascent and tentative foray into what may prove to be a very interesting alternative to the vagaries and whims of using diffusion models to effect modifications to an image (i.e., via Firefly and other API-based diffusion methods, and via open source architectures such as Stable Diffusion and Flux).

Though there are many diffusion models that can effect minor changes in images, LDMs are limited by their semantic and often 'over-imaginative' approach to a text-based user request for a modification.

Therefore the ability to temporarily pull part of an image into 3D space, manipulate it and replace it back into the image, while using only the source image as a reference, seems a task that Gaussian Splatting may be well suited for in the future.

* There is some confusion in the paper, in that it cites 'Amorphous-Mirage' as the most effective and capable method, in spite of its tendency to produce unwanted Gaussians (artifacts), while arguing that 'Graphite-Mirage' is more flexible. It appears that Amorphous-Mirage obtains the best detail, and Graphite-Mirage the best flexibility. Since both methods are presented in the paper, with their diverse strengths and weaknesses, the authors' preference, if any, does not appear to be clear at this time.

First published Thursday, October 3, 2024

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More