

GaINeR: Geometry-Aware Implicit Neural Representation for Image Editing

Weronika Jakubowska^{*,1}, Mikołaj Zieliński^{*,2}, Rafał Tobiasz^{*,3,4,5},
Krzysztof Byrski³, Maciej Zięba^{1,6}, Dominik Belter², and Przemysław
Spurek^{3,5}

¹ Wrocław University of Science and Technology, Poland

² Poznań University of Technology, Institute of Robotics and Machine Intelligence,
Poland

³ Faculty of Mathematics and Computer Science, Jagiellonian University, Poland

⁴ Doctoral School of Exact and Natural Sciences, Jagiellonian University, Poland

⁵ IDEAS Research Institute, Poland

⁶ Tooploox, Poland




Fig. 1: Given a single image, GaINeR reconstructs a manipulable representation that enables geometry-aware editing. Our method supports object-level edits (top left), 2D to 3D transformations (top right), editing 3D reconstruction (bottom left), and upscaling (bottom right) while preserving semantic structure and photometric coherence.

Abstract. Implicit Neural Representations (INRs) are widely used for modeling continuous 2D images, enabling high-fidelity reconstruction, super-resolution, and compression. Architectures such as SIREN, WIRE, and FINER demonstrate their ability to capture fine image details. However, conventional INRs lack explicit geometric structure, limiting local editing, and integration with physical simulation. To address these limitations, we propose GaINeR (Geometry-Aware Implicit Network Representation), a novel framework for 2D images that combines trainable Gaussian distributions with a neural network-based INR. For a given image coordinate, the model retrieves the K nearest Gaussians, aggregates distance-weighted embeddings, and predicts the RGB value via a neural network. This design enables continuous image representation,

interpretable geometric structure, and flexible local editing, providing a foundation for physically aware and interactive image manipulation. Our method supports geometry-consistent transformations, seamless super-resolution, and integration with physics-based simulations. Moreover, the Gaussian representation allows lifting a single 2D image into a geometry-aware 3D representation, enabling depth-guided editing. Experiments demonstrate that GaINeR achieves state-of-the-art reconstruction quality while maintaining flexible and physically consistent image editing.

Keywords: Implicit Neural Representations · Geometry-aware editing
· Gaussian-Based Representation

1 Introduction

Implicit Neural Representations (INRs) have transformed the modeling of continuous signals across domains such as computer vision and computational physics [12, 22, 24]. Instead of discrete grids tied to fixed spatial resolutions, INRs represent signals as continuous functions parameterized by neural networks, enabling resolution-independent representations with high expressiveness and compact memory usage. While architectures such as SIREN [22], GAUSS [19], WIRE [20], FINER [14], and positional encoding [25] improve the ability of INRs to capture high-frequency image details, they typically lack explicit geometric structure, limiting local editing and integration with physical simulations.

Recent work, MiRaGe [27], addresses image editability by representing 2D images by a Gaussian Splatting-based model, enabling intuitive modifications and coupling with physics engines. Each Gaussian acts as a discrete component, allowing local transformations and perspective changes. However, because the representation is discrete, MiRaGe suffers from producing artifacts when the physical engine significantly alters the geometry of objects in a 2D image. The problem can be partially solved by training very small Gaussians, but it requires dedicated parameters that must be manually tuned. Additionally a Gaussian Splatting-based model cannot provide a truly continuous image representation.

To overcome these limitations, we propose Geometry-Aware Implicit Network Representation (GaINeR), the first model that combines the continuous representation of INR for 2D images with trainable Gaussian embeddings. GaINeR enables editing 2D images, and 2D to 3D manipulations (see Fig. 1, Fig. 2). We achieve this by explicitly conditioning the implicit function on a geometric structure defined by a set of specific trainable Gaussian components, each parameterized by position, scale, and a unique learnable feature embedding. The core INR network then acts as a decoder. Given a query coordinate (x, y) , the network retrieves and aggregates the features of the k nearest Gaussians using a distance-weighted mechanism (see Fig. 3). The MLP then uses this aggregated geometric and feature information to synthesize the final, high-fidelity RGB colors. This design allows precise spatial manipulation of the image, as Gaussian positions and shapes can be adjusted to produce smooth, real-time deformations. Additionally, the Gaussians can act as discrete physical entities, enabling the entire 2D representation to interact naturally with standard physics engines




Fig. 2: GaINeR produces realistic image edits by preserving structural consistency under spatial transformations. Red arrows indicate the direction of the applied changes.

for dynamic simulation and animation. As a result, GaINeR achieves state-of-the-art performance in 2D image reconstruction, while maintaining fully continuous and editable representations. Importantly, due to the locality of Gaussian conditioning, our representation naturally extends beyond the 2D image plane. By treating the image as a continuous surface embedded in 3D space, GaINeR enables lifting a single 2D image into a geometry-aware 3D representation. This allows depth-guided Gaussian displacement and 3D-consistent editing. Furthermore, as a fully continuous representation, GaINeR inherently supports arbitrary-resolution image synthesis, enabling seamless super-resolution and upscaling without retraining.

The following constitutes a list of our key contributions:

- We introduce GaINeR, an implicit neural representation of 2D images based on Gaussian embeddings and MLPs, which provides superior reconstruction quality compared to prior INR-based approaches.
- We extend the INR framework by introducing a geometric component that allows direct and flexible 2D image modification and integrates with a physical simulation engine for interactive effects. Thanks to its fully continuous formulation, GaINeR also supports arbitrary-resolution image synthesis, enabling seamless super-resolution and upscaling without retraining.
- We demonstrate that our geometry-aware formulation enables single-image 3D-aware transformations via depth-guided Gaussian displacement, allowing consistent spatial manipulation beyond the 2D image plane.

2 Related work

Implicit Neural Representations (INRs) have received considerable attention in research [26] as neural models designed for signal representation. This paradigm has found broad application across various domains, including the representation of images [12], videos [6], and 3D shapes [18]. Significant uses span 3D shape reconstruction [16], robotics [13], and data compression [24].

INR for images INR architectures often remain simple, typically consisting of a single Multi-Layer Perceptron (MLP), with recent efforts focused on enhancing capacity through specialized embedding techniques [5, 17], unique activation functions [21], refined rendering methods [3], and regularization strategies [28].

Implicit Neural Representations (INRs), or coordinate-based MLPs, have fundamentally transformed the representation and processing of continuous signals across various domains. The main direction of INRs development for 2D images is to adapt various activation functions and training procedures. SIREN [22] uses sine activations alongside a frequency hyperparameter ω , enabling smooth fitting and stable training. While effective for continuous signals, its fixed-frequency multiplier limits versatility, resulting in lower accuracy when modeling details across diverse frequencies [9]. GAUSS [19] employs a Gaussian activation function providing smooth localized representations that are beneficial for signal reconstruction and denoising. However, its aperiodic nature hinders the effective modeling of sharp, high-frequency details. WIRE [20] utilizes the Gabor wavelet activation, combining sinusoidal and Gaussian components to capture localized frequencies. Although effective for modeling high-frequency details within specific spatial regions, its complex-valued formulation increases computational demands, slowing both training and inference. FINER [14] modifies SIREN with a frequency-variable activation, broadening the representable frequency spectrum while retaining SIREN’s benefits. However, FINER is sensitive to bias initialization and complicates the training process. Lastly, Positional Encoding (PE) [25] maps inputs into a higher-dimensional space using sinusoidal functions to mitigate the spectral bias common in MLPs.

Such a representation provides high-quality reconstruction but offers limited manual editing and integration with physical engines. In this paper, we aim to incorporate a geometric component into the INR framework to facilitate 2D image modification.

Explicit and Editable Representations Although INRs excel at compression and continuous representation, their inherent implicit nature severely limits direct interaction and editing. Conversely, explicit representations prioritize geometry and manipulability. 3D Gaussian Splatting (3DGS) [10] recently revolutionized novel view synthesis by employing a vast collection of explicit, 3D anisotropic Gaussians, achieving real-time rendering speeds. This approach naturally allows direct manipulation of Gaussian primitives, offering superior editability compared to traditional volumetric methods. Following this, Gaussian-Image [29] adapted the core GS concept to a 2D image representation, demonstrating competitive compression and quality. However, while these Gaussian-based methods introduce necessary explicit geometry, they still rely on simple color and opacity stored directly within each primitive. Consequently, they lose the continuous spatial coherence and superior color fidelity offered by a dedicated INR decoder. Models such as MiRaGe [27] attempted to address the editability challenge for 2D images by placing 2D Gaussian representations in a 3D space to enable physical interactions such as bending and folding.




Fig. 3: Overview of the GaINeR framework. (a) **Training:** For RGBA inputs, Gaussian components outside the alpha mask are removed from training. Each input coordinate $\mathbf{x}$ is encoded via Gaussian features $\mathbf{e}_i$ derived from the multi-resolution hashgrid $\mathcal{H}(\mu_i)$. The aggregated embedding $\mathbf{e}_{\text{KNN}}(\mathbf{x}, \mathcal{G})$ is obtained through radius-limited KNN interpolation and decoded by the MLP f_θ to reconstruct pixel color $\mathbf{c}$. (b) **Rendering and Editing:** After training, the optimized Gaussian set $\mathcal{G}$ is used to reconstruct the sampling mask (for RGBA cases) and to enable spatial transformations. Edited Gaussian means $\mu'_i \in \mathbb{R}^2$ are re-encoded to produce updated embeddings $\mathbf{e}'_{\text{KNN}}(\mathbf{x}, \mathcal{G}')$, which the decoder maps to modified image values, enabling geometry-consistent editing.

Our GaINeR framework distinguishes itself from these methods by creating a deeper synergy, where the explicit geometric primitives (Gaussians) are not renderers themselves, but rather rich feature providers that condition a dedicated Implicit Neural Representation decoder, thus ensuring both high-fidelity continuous color and native geometric editability and physical readiness. Examples of realistic edits enabled by our method are illustrated in Fig. 2.

3 Method

In this section, we present the GaINeR framework (see Fig. 3). We begin by revisiting the classical formulation of Implicit Neural Representations (INRs) to establish the foundational concepts. Next, we introduce GaINeR, which incorporates learnable Gaussian embeddings to achieve spatially consistent and editable 2D representations. Finally, we describe the training objective and optimization strategy used to jointly learn the embedding distribution and neural decoder.

While GaINeR is introduced in the context of 2D image reconstruction and editing, its formulation is inherently dimension-agnostic. Since Gaussian embeddings and KNN-based aggregation operate in continuous space, the same mechanism naturally extends to three-dimensional domains [17]. Consequently, GaINeR generalizes seamlessly from 2D image representations to 3D volumetric modeling.

Classical Implicit Neural Representations Implicit Neural Representations provide a continuous and compact way to encode signals such as images, shapes, or radiance fields using neural networks. Instead of storing discrete samples (e.g., pixels or voxels), an INR models the underlying signal as a function f_θ that directly maps input coordinates to output values. This formulation enables subpixel-level precision, smooth interpolation, and differentiability with respect to spatial coordinates.

Formally, an INR is typically realized as a multilayer perceptron (MLP):

$$f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^m,$$

where d denotes the dimensionality of the input domain (e.g., spatial coordinates), and m corresponds to the output space (e.g., color channels or density). For 2D images, the INR learns a mapping

$$f_\theta(x, y) = (r, g, b),$$

which associates each coordinate (x, y) with an RGB color value. In Neural Radiance Fields (NeRFs), the input extends to a 3D position and a viewing direction:

$$f_\theta(\mathbf{x}, \mathbf{d}) = (c, \sigma),$$

where c represents the emitted color and σ is the corresponding volume density.

This purely functional formulation captures all spatial information implicitly within the network parameters θ , without requiring discrete grids or explicit coordinate encoding. Classical INRs exhibit several key limitations. First, they lack explicit geometric structure, making it difficult to reason about spatial relationships or integrate with physical simulation engines. Second, the global nature of MLP parameterization restricts local controllability, complicating fine-grained image editing or deformation. Finally, reconstruction quality often depends on careful architectural tuning and positional encodings, which limit generalization across diverse image domains. These challenges motivate the development of geometry-aware extensions, such as our proposed GaINeR model, which bridges implicit representations with explicit geometric priors and physical consistency.

Geometry-Aware Implicit Network Representation Our method extends the classical INR formulation by introducing a geometry-consistent embedding mechanism based on a family of learnable Gaussian distributions. Each 2D coordinate in the image domain is represented not as a fixed input vector, but as a continuous Gaussian embedding that encodes its local spatial context. This design enables geometry-aware reasoning and smooth interpolation while preserving the functional nature of implicit representations.

Let the input domain be $\mathbf{x} = (x, y) \in \mathbb{R}^2$. We define a set of N Gaussian components, each parameterized by its mean $\mu_i \in \mathbb{R}^2$, covariance matrix $\Sigma_i \in \mathbb{R}^{2 \times 2}$, and a learnable embedding vector $\mathbf{e}_i \in \mathbb{R}^d$:

$$\mathcal{G} = \{(\mathcal{N}_i(\mu_i, \Sigma_i), \mathbf{e}_i)\}_{i=1}^N.$$

To enhance the representational capacity of the embeddings, we are using multi-resolution hashgrid features [17] as Gaussian features obtained from,

$\mathcal{H}(\mu) = \mathbf{e}$, where $\mathcal{H}$ is hashing function. These features provide fine-grained spatial detail during training but are omitted at inference, as the learned Gaussian parameters themselves encode the required multi-scale structure.

For a given input coordinate $\mathbf{x}$, we identify its K nearest Gaussian centers in Euclidean space using a radius-limited k -nearest neighbors (KNN) search, where candidates are considered only within a distance r for computational efficiency. The corresponding embedding vectors of these neighbors are then aggregated using Gaussian-weighted averaging:

$$\mathbf{e}_{\text{KNN}}(\mathbf{x}, \mathcal{G}) = \frac{\sum_{i \in \text{KNN}(\mathbf{x}, \mathcal{G})} \mathcal{N}_i(\mathbf{x}; \mu_i, \Sigma_i, \mathbf{e}_i) \mathbf{e}_i}{\sum_{i \in \text{KNN}(\mathbf{x}, \mathcal{G})} \mathcal{N}_i(\mathbf{x}; \mu_i, \Sigma_i, \mathbf{e}_i)},$$

where $\text{KNN}(\mathbf{x}, \mathcal{G})$ denotes the indices of the K nearest neighbors of $\mathbf{x}$. The resulting feature

$$\mathbf{e}(\mathbf{x}) = \mathbf{e}_{\text{KNN}}(\mathbf{x}, \mathcal{G}) \in \mathbb{R}^d$$

serves as a locally geometry-aware embedding that captures both spatial proximity and the density structure of the learned Gaussians.

The aggregated embedding $\mathbf{e}(\mathbf{x})$ is subsequently processed by a multilayer perceptron (MLP) decoder:

$$f_\theta(\mathbf{x}) = \text{MLP}_\theta(\mathbf{e}(\mathbf{x})),$$

where $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^3$ produces the RGB color value corresponding to the coordinate $\mathbf{x}$. This combination of continuous Gaussian embeddings and coordinate-based neural decoding allows the model to reconstruct and edit 2D images with smooth, geometry-consistent variations.

Unlike classical INRs, which operate directly on raw spatial coordinates, GaINeR embeds each position into a learned, spatially coherent feature space. This design enhances the model’s ability to capture local geometric relationships, improves reconstruction fidelity, and facilitates controllable 2D image editing through interpretable, geometry-aware embeddings.

Training Objective The training of GaINeR jointly optimizes the parameters of the Gaussian embedding $\{(\Sigma_i, \mathbf{e}_i)\}_{i=1}^N$, the multi-resolution hashgrid encoder $\mathcal{H}(\cdot)$, and the MLP decoder f_θ that reconstructs the target image.

Given a training set $\{(\mathbf{x}_j, \mathbf{c}_j)\}_{j=1}^M$, where $M \in N$ is a subset of Gaussians in training dataset, $\mathbf{x}_j = (x_j, y_j)$ denotes pixel coordinates, and $\mathbf{c}_j = (r_j, g_j, b_j)$ are the corresponding RGB values, the model is trained to minimize the Smooth L1 reconstruction loss:

$$\mathcal{L}_{\text{rec}} = \frac{1}{M} \sum_{j=1}^M \text{SmoothL1}(f_\theta(\mathbf{x}_j) - \mathbf{c}_j),$$

where

$$\text{SmoothL1}(\mathbf{r}) = \begin{cases} \frac{1}{2} \|\mathbf{r}\|_2^2 / \beta, & \text{if } \|\mathbf{r}\|_1 < \beta, \\ \|\mathbf{r}\|_1 - \frac{1}{2} \beta, & \text{otherwise,} \end{cases}$$




Fig. 4: Results of 2D-to-3D lifting with GaINeR. We elevate flat 2D images into fully volumetric 3D representations using monocular depth maps. The final column shows the reconstructed geometry rendered from a novel viewpoint. Our method generalizes well across complex paintings and masked objects.

and $\mathbf{r} = f_\theta(\mathbf{x}_j) - \mathbf{c}_j$ denotes the residual, while β controls the transition between the ℓ_2 and ℓ_1 regions.

Alpha-based Sampling and Mask Reconstruction. For images containing an alpha channel, i.e., RGBA inputs, we utilize the alpha map as a spatial sampling mask during training to improve computational efficiency. Specifically, only pixels with nonzero alpha values are included in the training set $\{(\mathbf{x}_j, \mathbf{c}_j)\}_{j=1}^M$, which focuses optimization on visible regions and avoids redundant computation in transparent areas.

During inference, the corresponding mask must also be reconstructed to enable consistent rendering and editing. To this end, we also train a dedicated mask model that learns to reproduce the alpha-based sampling mask. This model is subsequently animated in synchrony with the image model, ensuring temporally coherent and spatially accurate masks for all generated or edited frames.

Image Editing During inference, the hashgrid encoder $\mathcal{H}$ is no longer required, since the learned Gaussian embeddings $\mathbf{e}_i$ already encode multi-scale spatial information acquired during training. Thus, the model operates directly on the optimized Gaussian parameters $\mathcal{G} = \{\mathcal{N}(\mu_i, \Sigma_i, \mathbf{e}_i)\}_{i=1}^N$, enabling efficient and interpretable scene editing.

Image editing in GaINeR is performed by modifying the spatial configuration of the Gaussian means μ_i . Because the feature associated with any query coordinate $\mathbf{x}$ is computed through a radius-limited KNN aggregation, $\mathbf{e}_{\text{KNN}}(\mathbf{x}, \mathcal{G})$, shifting a subset of Gaussian centers $\mu_i \mapsto \mu'_i$ effectively alters the local neighborhood structure and, consequently, the interpolated feature $\mathbf{e}_{\text{KNN}}(\mathbf{x}, \mathcal{G})$. The MLP decoder f_θ then produces a modified color field corresponding to the updated embedding distribution:

$$f'_\theta(\mathbf{x}) = \text{MLP}_\theta(\mathbf{e}_{\text{KNN}}(\mathbf{x}, \mathcal{G}')).$$

This formulation allows geometric transformations such as translation, rotation, bending, or stretching to be applied directly in the Gaussian space while

maintaining photometric consistency in the output. Since the MLP decoder is conditioned only on the aggregated embeddings, not on absolute Gaussian positions, the model exhibits a form of spatial equivariance: moving the Gaussians induces coherent deformations in the reconstructed image without introducing artifacts.




Fig. 5: Overview of 3D reconstruction with GaIneR. Input images are segmented with SAM [11] and lifted to 3D using depth estimated by Depth-Pro [4]. The Gaussian means are projected into 3D while the decoder predicts color c and density σ , enabling NeRF-style volumetric rendering and novel-view synthesis, and direct 3D scene modifications.

2D To 3D Transformations The proposed framework (Fig. 5) naturally generalizes beyond 2D image reconstruction. By extending the MLP decoder to additionally predict a density term σ ,

$$f_{\theta}(\mathbf{x}) = (\mathbf{c}, \sigma),$$

the model integrates with a volumetric rendering formulation analogous to NeRF [16]. In this case, the rendering equation is expressed as

$$\hat{\mathbf{C}}(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t)) dt,$$

where,

$$T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right),$$

$\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ denotes a camera ray with origin $\mathbf{o}$ and direction $\mathbf{d}$.

Each Gaussian component is then lifted into three-dimensional space. The local embedding for any query position $\mathbf{x}$ is computed through the Gaussian-weighted KNN aggregation $\mathbf{e}_{\text{KNN}}(\mathbf{x}, \mathcal{G})$. However the KNN utilizes Mahalanobis distance instead of Euclidean distance for the search for nearest neighbors. As in the 2D formulation, shifting the Gaussian means $\mu_i \mapsto \mu'_i$ directly modifies the interpolated embedding field $\mathbf{e}_{\text{KNN}}(\mathbf{x}, \mathcal{G})$, thereby altering the reconstructed volumetric scene while preserving photometric coherence.

Since the decoder f_{θ} depends solely on the aggregated embeddings rather than absolute coordinates, these transformations yield physically consistent and




Fig. 6: Visual comparison of pixel-wise absolute errors between ground truth and reconstructed images produced by our method and competing approaches. Error magnitudes are contrast-enhanced using gamma correction ($\gamma = 0.2$) for better visibility.

spatially smooth changes in appearance. Consequently, GaINeR provides a unified framework for continuous, differentiable editing across both 2D and 3D domains, enabling 2D-to-3D lifting followed by direct 3D scene editing.

4 Experiments

We evaluated our proposed method across three tasks: image reconstruction, image editing, and 2D-to-3D lifting. The image editing task encompassed both user-driven transformations and physically simulated deformations. For reconstruction, we employed two standard benchmarks: Kodak [8] and DIV2K [2]. Additional ablation studies, runtime analysis, and implementation details are provided in the supplementary material. The Kodak dataset contains 24 natural images of resolution 768×512 , covering diverse scenes such as portraits, landscapes, and textures.

The DIV2K validation set comprises 100 high-resolution images subjected to 2x bicubic downscaling, with spatial resolutions ranging from 408×1020 to 1020×1020 , and is widely used for evaluating image restoration and super-resolution methods.

Model	Kodak		DIV2K	
	PSNR $\uparrow$	MS-SSIM $\uparrow$	PSNR $\uparrow$	MS-SSIM $\uparrow$
WIRE [20]	41.47	0.9939	35.64	0.9511
SIREN [22]	40.83	0.9960	36.02	0.9568
I-NGP [17]	43.88	0.9976	37.06	0.9894
NeuRBF [7]	43.78	0.9984	37.06	0.9901
3DGS [10]	44.09	0.9991	39.12	0.9980
GaussianImage [29]	38.93	0.9984	41.48	0.9981
MiraGe [27]	59.52	0.9999	54.54	0.9998
GaINeR (ours)	77.09	0.9999	62.20	0.9999

Table 1: Quantitative comparison on the Kodak and DIV2K datasets. Our method achieves the highest reconstruction quality across both datasets in terms of PSNR and MS-SSIM.




Fig. 7: Comparison of MiraGe and GaINeR under manually applied geometric modifications in Blender. GaINeR maintains stable and realistic results under bending, cutting and stretching, whereas MiraGe exhibits visible artifacts.

Quantitative and Qualitative Evaluation Table 1 reports the average PSNR and MS-SSIM values obtained on both datasets. We compared our approach against WIRE [20], SIREN [22], I-NGP [17], NeuRBF [7], 3DGS [10], GaussianImage [29] and MiraGe [27]. In both cases, our method demonstrates a significant improvement in PSNR compared to previous approaches. Specifically, compared to the strongest existing method (MiraGe), our approach improves the average PSNR from 59.52 to 77.09 on the Kodak dataset, and from 54.54 to 62.20 on DIV2K, while maintaining near-perfect MS-SSIM values. This significant gain indicates that our representation captures image details with much higher fidelity and preserves structural consistency more effectively than prior implicit representations.

To ensure consistency across Gaussian-based representations, we evaluate MiraGe, GaussianImage, and our proposed method under an identical training setup. Each model is initialized with 500k Gaussian primitives, which remain fixed throughout optimization (no densification or pruning), and is trained for 30k iterations. As reported in Table 2 our method achieves the highest accuracy among all evaluated approaches.

Since the PSNR values across all methods are exceptionally high, the visual differences are barely perceptible to the human eye. Therefore, we compute the *absolute error* for each pixel and visualize the results as error maps, shown in Fig. 6. These error visualizations reveal that GaINeR method produces almost no noticeable errors compared to other approaches, demonstrating its superior reconstruction accuracy and stability.

Editing comparisons We conducted a series of experiments on images from the DIV2K dataset, exploring various types of spatial modifications. For consis-

Model	Kodak		DIV2K	
	PSNR $\uparrow$	MS-SSIM $\uparrow$	PSNR $\uparrow$	MS-SSIM $\uparrow$
GaussianImage [29]	30.46	0.9635	28.58	0.9565
MiraGe [27]	62.49	0.9999	58.45	0.9999
GaINeR (ours)	77.09	0.9999	62.20	0.9999

Table 2: Comparison with Gaussian-based approaches on the Kodak and DIV2K datasets. Our method achieves the highest reconstruction quality in terms of PSNR and MS-SSIM.




Fig. 8: Qualitative comparison of physics-based animations. **Elastic Flower:** MiraGe’s render exhibits transparencies, while ours remains opaque. **Elastic Beans:** MiraGe produces spiky boundary artifacts, which are absent in our result. **Sand Apple:** As the object collapses, MiraGe’s rendering shows structural artifacts, unlike GaINeR.

tency across methods, we evaluate MiraGe using two camera configurations: one and mirror. The one camera setup corresponds to the same viewpoint configuration employed in our GaINeR model, ensuring direct comparability. The mirror camera provides a reflective viewpoint that captures mirrored scene geometry and appearance. As illustrated in Fig. 7, we evaluate the visual effects of bending, stretching, cutting and displacing Gaussian components. MiraGe exhibits severe distortions under these transformations, primarily due to its reliance on Gaussian primitives for direct rendering. In contrast, our GaINeR framework produces visually consistent and realistic results across all conditions, as it utilizes Gaussian components solely for encoding rather than for rendering.

Physics comparisons To demonstrate applicability in dynamic scenarios, we treat the centers of our learned Gaussians as material points in a Material Point Method (MPM) simulation driven by the Taichi Elements [23] engine. Fig. 8 compares our method against MiraGe, with full animations provided in the supplementary material. While MiraGe’s discrete representation produces geometric artifacts and spiky boundaries under stress, GaINeR renders coherent surfaces even under significant deformation. Since our MLP decoder is conditioned on the geometric structure, it can robustly interpolate visual information




Fig. 9: Physical simulation of complex multi-material interactions using GaINeR. By treating learned Gaussian centers as material points in a Material Point Method engine, we simulate granular dynamics (sand) and fluid mechanics (water) simultaneously. Our continuous MLP decoder maintains cohesive surfaces, fluid-like blending, and high-frequency textures even under severe topological deformations over time.

from sparse or distorted particle distributions, avoiding the visual discontinuities inherent to discrete representations. We further showcase this robustness in Fig. 9, where GaINeR successfully handles complex, multi-material interactions, such as fluid mechanics and granular dynamics, maintaining continuous, high-fidelity textures despite extreme topological changes over time. Tab. 3 reports perceptual metrics for physics-based editing results. GaINeR achieves the highest MUSIQ and CLIPQA scores while remaining competitive in FID.

Table 3: Perceptual evaluation of physics-based deformation and editing results.

Method	FID ↓	MUSIQ ↑	CLIPQA ↑
GaINeR	376.34	62.62	0.530
MiraGe	372.41	61.03	0.528


2D-to-3D Lifting. Our geometry-aware formulation enables single-image 3D-aware transformations. As illustrated in Fig. 5, input images are processed with SAM [11] and lifted to 3D using Depth-Pro [4]. To isolate GaINeR’s representational capacity from the artifacts of monocular depth estimation, we evaluate single-image Novel View Synthesis (10° – 15° baseline) on Hotdog and Scan 114 (DTU [1]) using both ground-truth depth and Depth-Pro (AbsRel error of 2.3% and 4.4%, respectively).

As shown in Tab. 4, GaINeR significantly outperforms an Explicit Point Cloud (EPC) reprojection baseline given identical depth sources, confirming that our continuous INR decoder effectively interpolates sparse geometry to prevent rendering holes. Furthermore, rendering errors (PSNR drops) are confined strictly to inevitable disocclusions and baked-in specularities (Fig. 10), while the underlying geometry remains structurally coherent. Additional qualitative results (Fig. 4) demonstrate that GaINeR generalizes well across complex datasets (e.g., LLFF [15]) and paintings, producing realistic 3D structures from flat 2D inputs.

Upscaling Image upscaling aims to reconstruct high-resolution images from low-resolution inputs while preserving fine details and structural consistency, a task that becomes particularly challenging at high magnification factors where artifacts and over-smoothing often appear. Fig. 11 shows a qualitative comparison of GaINeR with 3D Gaussian Splatting and MiraGe at $\times 8$, $\times 10$, and

Table 4: Single-image NVS evaluation. GaINeR significantly outperforms the EPC baseline using identical depth sources.

Scene	Depth	EPC	GaINeR (Ours)		
		PSNR	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Hotdog (NeRF Syn.)	GT Depth	8.57	20.58	0.874	0.186
	Depth-Pro	8.57	19.50	0.865	0.193
Scan 114 (DTU)	GT Depth	13.51	18.09	0.724	0.237
	Depth-Pro	13.51	17.74	0.712	0.245

**Fig. 10:** Single-image NVS. GaINeR preserves structural geometry; errors localize at disocclusions and baked-in specularities.

$\times 12$ scales. As an implicit neural representation (INR), GaINeR generates high-resolution outputs without relying on fixed discrete grids. As illustrated in the figure, it preserves fine textures and sharp edges even at extreme magnifications, while 3D Gaussian Splatting and MiraGe exhibit visible degradation artifacts, including Gaussian-like patterns and over-smoothed regions. All experiments were conducted on the DIV2K dataset, where the input images were obtained by bicubic downscaling of the original images by a factor of 8. Quantitatively, on DIV2K at $\times 8$ super-resolution, GaINeR achieves 23.22dB PSNR, compared to 17.33dB for 3D Gaussian Splatting and 16.61dB for MiraGe.

5 Conclusion

In this work, we introduced GaINeR, a geometry-aware implicit representation that combines the continuous flexibility of neural networks with the structured expressiveness of learnable Gaussian embeddings. By encoding spatial coordinates through Gaussian-weighted KNN aggregation, GaINeR achieves smooth, spatially coherent reconstructions while enabling 2D image manipulations.

Our framework bridges coordinate-based MLPs with explicit geometric priors, yielding a unified model capable of high-fidelity reconstruction and photometrically consistent transformations. Comprehensive experiments demonstrate that GaINeR surpasses prior implicit representations such as SIREN, WIRE, and MiraGe in both quantitative accuracy and visual quality, while offering greater interpretability through its Gaussian-based encoding. In addition, the continuous formulation of GaINeR naturally supports arbitrary-resolution image synthesis, enabling seamless super-resolution and upscaling without retraining.




Fig. 11: Visual comparison of GaNeR with 3D Gaussian Splatting and MiraGe at $\times 8$, $\times 10$, and $\times 12$ scaling factors. Our method maintains higher perceptual quality, whereas competing approaches introduce Gaussian artifacts.

Furthermore, we extend the formulation to 2D-to-3D lifting, where the learned Gaussian embeddings are projected into three-dimensional space to reconstruct depth-aware geometry and novel viewpoints from a single image. This generalization highlights the potential of Gaussian-encoded neural fields for bridging 2D perception and 3D reasoning, paving the way for future research in continuous, geometry-aware scene representations.

Limitations While GaNeR enables 3D lifting, its geometry is bounded by the quality of the off-the-shelf monocular depth estimator. Errors in depth prediction can lead to geometric artifacts.

6 Acknowledgements

M. Zieliński and D. Belter were supported by the National Science Centre, Poland, under research project no UMO-2023/51/B/ST6/01646. The scientific work was carried out using the infrastructure of the Poznan Supercomputing and Networking Center.

The work of W. Jakubowska, R. Tobiasz, K. Byrski, P. Spurek, was supported by the National Centre of Science (Poland) Grant No. 2023/50/E/ST6/00068.

We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Centers: ACK Cyfronet AGH, WCSS, CI TASK) for providing computer facilities and support within computational grant no. PLG/2026/019272.

Some of the computations presented in this work were carried out using the infrastructure of the Poznań Supercomputing and Networking Center (PSNC).

References

1. Aanaes, H., Dahl, A., Steenstrup Pedersen, K.: Interesting interest points. *International Journal of Computer Vision* **97**, 18–35 (2012)
2. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (July 2017)
3. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 5855–5864 (2021)
4. Bochkovskiy, A., Delaunoy, A., Germain, H., Santos, M., Zhou, Y., Richter, S.R., Koltun, V.: Depth pro: Sharp monocular metric depth in less than a second. In: *ICLR. OpenReview.net* (2025)
5. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: *European conference on computer vision*. pp. 333–350. Springer (2022)
6. Chen, Z., Chen, Y., Liu, J., Xu, X., Goel, V., Wang, Z., Shi, H., Wang, X.: Videoinr: Learning video implicit neural representation for continuous space-time super-resolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2047–2057 (2022)
7. Chen, Z., Li, Z., Song, L., Chen, L., Yu, J., Yuan, J., Xu, Y.: Neurbf: A neural fields representation with adaptive radial basis functions (2023), <https://arxiv.org/abs/2309.15426>, accessed: April 28, 2026
8. Company, E.K.: Kodak lossless true color image suite. Web page (1991), available at <https://r0k.us/graphics/kodak/>
9. Kania, A., Mihajlovic, M., Prokudin, S., Tabor, J., Spurek, P.: Fresh: Frequency shifting for accelerated neural representation learning. In: *The Thirteenth International Conference on Learning Representations* (2025)
10. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
11. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.B.: Segment anything. In: *ICCV*. pp. 3992–4003. IEEE (2023)
12. Kloczek, S., Maziarka, Ł., Wołczyk, M., Tabor, J., Nowak, J., Śmieja, M.: Hyper-network functional image representation. In: *International Conference on Artificial Neural Networks*. pp. 496–510. Springer (2019)
13. Lin, C.H., Ma, W.C., Torralba, A., Lucey, S.: Barf: Bundle-adjusting neural radiance fields. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 5741–5751 (2021)
14. Liu, Z., Zhu, H., Zhang, Q., Fu, J., Deng, W., Ma, Z., Guo, Y., Cao, X.: Finer: Flexible spectral-bias tuning in implicit neural representation by variable-periodic activation functions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2713–2722 (2024)
15. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* (2019)
16. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)

17. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)* **41**(4), 1–15 (2022)
18. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 165–174 (2019)
19. Ramasinghe, S., Lucey, S.: Beyond periodicity: Towards a unifying framework for activations in coordinate-mlps. In: *European Conference on Computer Vision*. pp. 142–158. Springer (2022)
20. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.G.: Wire: Wavelet implicit neural representations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18507–18516 (2023)
21. Sitzmann, V., Chan, E., Tucker, R., Snavely, N., Wetzstein, G.: Metasdf: Meta-learning signed distance functions. *Advances in Neural Information Processing Systems* **33**, 10136–10147 (2020)
22. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Advances in neural information processing systems* **33**, 7462–7473 (2020)
23. Taichi Graphics: Taichi elements: A high-performance multi-material continuum physics engine. *GitHub repository* (2020), https://github.com/taichi-dev/taichi_elements
24. Takikawa, T., Saito, S., Tompkin, J., Sitzmann, V., Sridhar, S., Litany, O., Yu, A.: Neural fields for visual computing: Siggraph 2023 course. In: *ACM SIGGRAPH 2023 Courses*, pp. 1–227 (2023)
25. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems* **33**, 7537–7547 (2020)
26. Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., Yifan, W., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., et al.: Advances in neural rendering. In: *Computer Graphics Forum*. vol. 41, pp. 703–735. Wiley Online Library (2022)
27. Waczynska, J., Szczepanik, T., Borycki, P., Tadeja, S., Bohné, T., Spurek, P.: Mirage: Editable 2d images using gaussian splatting. In: *Forty-second International Conference on Machine Learning* (2025)
28. Yang, J., Pavone, M., Wang, Y.: Freenerf: Improving few-shot neural rendering with free frequency regularization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8254–8263 (2023)
29. Zhang, X., Ge, X., Xu, T., He, D., Wang, Y., Qin, H., Lu, G., Geng, J., Zhang, J.: Gaussianimage: 1000 fps image representation and compression by 2d gaussian splatting. In: *European Conference on Computer Vision*. pp. 327–345. Springer (2024)