

WIPES: Wavelet-based Visual Primitives*

Wenhao Zhang^{1,†}, Hao Zhu^{1,†}, Delong Wu^{1,†}, Di Kang², Linchao Bao², Xun Cao^{1,*}, Zhan Ma¹

¹ Nanjing University, ² Tencent

[†] Equal contribution. * Corresponding author: caoxun@nju.edu.cn

Abstract


Pursuing a continuous visual representation that offers flexible frequency modulation and fast rendering speed has recently garnered increasing attention in the fields of 3D vision and graphics. However, existing representations often rely on frequency guidance or complex neural network decoding, leading to spectrum loss or slow rendering. To address these limitations, we propose **WIPES**, a universal **Wavelet-based vIsual PrimitivES** for representing multi-dimensional visual signals. Building on the spatial-frequency localization advantages of wavelets, WIPES effectively captures both the low-frequency “forest” and the high-frequency “trees.” Additionally, we develop a wavelet-based differentiable rasterizer to achieve fast visual rendering. Experimental results on various visual tasks, including 2D image representation, 5D static and 6D dynamic novel view synthesis, demonstrate that WIPES, as a visual primitive, offers higher rendering quality and faster inference than INR-based methods, and outperforms Gaussian-based representations in rendering quality.

1. Introduction


Representing visual signals as continuous functions has garnered increasing attention in recent years. Owing to the inherent advantages of end-to-end differentiability and scalable querying, these continuous representations can be seamlessly integrated with various imaging and rendering functions. This integration facilitates the optimization of inverse problems, whether the resolution is down-sampled or the dimensionality is reduced. Since these representations are optimized independently for different scenarios, pursuing a universal visual primitive, which offers complex frequency compatibility and fast rendering speed, is of paramount importance.

However, existing representations suffer from spectral loss or slow rendering. For instance, the representational capacity of popular implicit neural representations (INRs)


*The work was supported in part by NSFC Grant 62025108 and Tencent Rhino-Bird Joint Research Program RBFR2024009.



(a) Comparison of Gaussians and WIPES



(b) PSNR improvement of WIPES over Gaussian primitives [44]



(c) Performance comparison of WIPES and Gaussian primitives [44]

Figure 1. Spatial-Frequency Localization Benefits of WIPES. (a) shows the spatial and frequency distributions of Gaussian and WIPES primitives with various parameters. Adjusting σ only affects the Gaussian frequency bandwidth, unable to overcome its low-frequency constraint. In contrast, WIPES flexibly modulates frequency, making it advantageous for complex visual tasks. (b) displays image fitting on the Kodak dataset, where WIPES achieves higher PSNR with the same primitives count, and this advantage grows with more primitives. (c) illustrates that Gaussian primitives smooth complex textures due to low-frequency limitation, while WIPES accurately reconstructs rooftop patterns and railings through adaptable frequency properties.

[34] is often constrained by spectral bias [27], only the frequencies defined by hyperparameters are effectively learned [35, 41, 46], resulting in the loss of information in other frequency bands. Furthermore, the complex network architec-

tures of INRs require substantial computational resources [24] to decode signal attributes, thereby impeding processing speed across various visual tasks. Explicit representations, among which the 3D Gaussian Splatting (3DGS) [14] specifically models the visual world by employing multiple Gaussian functions. However, the inherent low-pass nature of Gaussian functions introduces inaccuracies in capturing high-frequency details. To mitigate these low-pass characteristics, frequency regularization [43] or modulation losses [11] are often employed, which necessitates extensive expertise to tune hyperparameters for different scenes.

These problems arise from a mismatch between the characteristics of the visual primitives used and the visual signals themselves. According to Marr’s vision theory, visual signals are perceived hierarchically, and intensity changes are detected at multiple scales. Wavelet, designed for multi-scale signal analysis, aligns well with the human visual system and can capture both the “forest” and the “trees” in visual signals [10]. Consequently, we developed WIPES, a universal primitive for representing visual signals.

WIPES leverages the Morlet wavelet—known for its continuous differentiability and robust spatial-frequency localization—to effectively capture signals exhibiting significant frequency variations. We extend the Morlet wavelet by developing a wavelet primitives adaptable to arbitrary dimensions, making it well-suited for various visual signals in real-world applications (*e.g.*, 2D images, 5D static and 6D dynamic radiance fields). Additionally, we develop a fast, differentiable rasterizer for WIPES by utilizing an efficient GPU sorting algorithm [18]. To mitigate the substantial computational overhead associated with automatic differentiation during training, we derive explicit gradients for all parameters in both forward and backward propagations. Our rasterizer effectively replaces Gaussian-based rasterizers, enabling seamless integration with different loss functions [5, 43] and subsequent visual tasks [38, 42, 44] built upon Gaussian splatting techniques.

In summary, we make the following contributions:

- We design a universal differentiable wavelet primitive that adaptively models various frequency features in real-world scenes without relying on handcrafted frequency guidance.
- We develop a fast differentiable rasterizer for our WIPES, enabling anisotropic splatting and efficient backpropagation across different dimensions to achieve general and effective visual representation.
- The proposed wavelet-based approach benefits superior reconstruction performance across multiple task dimensions compared to other visual representations.

2. Related Work

Implicit Neural Representations (INRs) play a pivotal role in modeling continuous functions of signals found in

images, scenes and videos [19]. They leverage MLPs to accurately map spatial coordinates to their corresponding values [24, 34]. This capability enables INRs to be applied in various domains, including novel view synthesis [24], free-hand 3D/4D echocardiography [32, 33], microscopy imaging [45], and nano-optics design [4].

However, INRs are constrained by spectral bias, favoring low-frequency information over high-frequency details [27]. To mitigate this bias, strategies like positional encoding [35] and specifically activation functions [20, 28, 31, 34] have been developed. Despite these advancements, the spectral bias remains fundamentally unaddressed and is merely transformed into a “frequency”-specified spectral bias [46]. Additionally, INRs suffer from slow training and inference due to their reliance on computationally intensive MLPs. To enhance efficiency, recent approaches have integrated explicit storage with implicit decoding, including sparse voxel grids [9], hash tables [25, 47, 48], and multi-scale/grid INRs [29, 30, 49]. While these approaches offer improvements, they introduce additional complexity, making it challenging to balance memory consumption with real-time performance. Furthermore, INRs are highly sensitive to parameter changes. Modifications to the input or internal parameters can lead to significant variations in the output, complicating the process of making local adjustments [23]. Although strategies such as explicit encoding and latent disentanglement [17, 36] have been proposed to address these issues, achieving precise control over local features in INR-based approaches remains challenging.

Explicit Signal Representations with Gaussian Primitives, conversely, overcome several INR limitations and are growing in popularity in computer vision and graphics. They use multiple independent continuous Gaussian functions to render scenes, which offers marked benefits like enhanced expressiveness and editability, and superior fitting efficiency than conventional INRs.

These Gaussian-based approaches have achieved top-tier performance in rendering quality and speed [14]. Early research on splatting introduced Gaussian primitives for scene representation [50, 51]. Recent advancements such as FreGS [43], mip-splatting [40], and scaffold-GS [21] enhance frequency management, mitigate aliasing, and incorporate voxel-based initialization, thereby demonstrating the robustness and efficiency of Gaussian-based approaches.

However, Gaussian primitives are constrained by their inherent low-pass filtering characteristics, which are effective for smooth signals but hinder the capture of high-frequency details like sharp edges or intricate textures. Real-world scenes frequently contain abrupt changes, such as the complex textures of grass or nylon fabric. Accurately representing these details requires multiple small-scale Gaussians [11, 13], which significantly increase model complexity, thereby reducing scalability and efficiency.

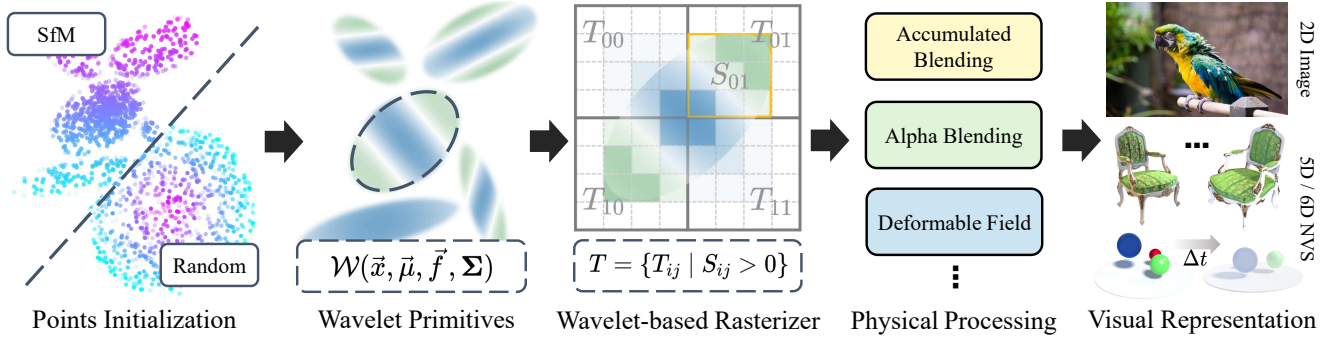


Figure 2. Pipeline of WIPES. WIPES can be seamlessly integrated into various visual reconstruction tasks, including 2D image representation, 5D static and 6D dynamic novel view synthesis (NVS).

Wavelet-based Signal Representations have been proven to be highly effective in analyzing complex signals within vision tasks. Traditional Fourier approaches use sine functions for their smoothness and clear characteristics in the Fourier domain. However, they decompose signals into global sine bases, which lack localization in both space and frequency [22]. This limitation makes Fourier approaches less suitable for visual data with local features and abrupt changes. In contrast, wavelets provide optimal localization in both domains, effectively capturing local high-frequency variations and sudden signal changes while maintaining a compact representation. Recent studies have leveraged wavelets to enhance vision models. MFN [8] employs wavelets as nonlinear filters to approximate complex functions with fewer parameters, improving frequency information representation. TriNeRFlet [15] integrates wavelets into Triplane’s multi-scale representations, enhancing cross-scale information sharing and high-frequency regularization for better 3D reconstructions. WavePlanes [2] uses wavelet compression on the spatiotemporal feature planes of dynamic NeRFs, reducing model size and enhancing high-frequency details.

However, recent wavelet-inspired representations are implicitly integrated into neural networks, which significantly reduce the rendering speed. On the contrary, the proposed wavelet-based primitives operate on the visual signal explicitly and directly, thus improving the efficiency while maintaining the advantages of spatial-frequency location.

3. Preliminaries

3.1. 3D Gaussian Splatting

3DGS [14] explicitly represents 3D scenes using Gaussian primitives, enabling efficient scene reconstruction. Specifically, each Gaussian primitive is parameterized by a spatial mean vector $\vec{\mu}$, a covariance matrix Σ (expressed as $\Sigma = \mathbf{R}\mathbf{S}\mathbf{S}^\top\mathbf{R}^\top$, where \mathbf{R} denotes rotation and \mathbf{S} denotes scaling), an opacity value α , and color attributes encoded by spherical harmonic coefficients c . In the world space,

the Gaussian primitive centered at $\vec{\mu}$ with Σ is defined as:

$$\mathcal{G}(\vec{x}, \vec{\mu}, \Sigma) = e^{-\frac{1}{2}(\vec{x}-\vec{\mu})^\top \Sigma^{-1}(\vec{x}-\vec{\mu})}, \quad (1)$$

where \vec{x} and $\vec{\mu}$ are column vectors in the world space, denoted as $(x, y, z)^\top$ and $(\mu_x, \mu_y, \mu_z)^\top$, respectively, the covariance matrix Σ is a 3×3 symmetric matrix.

When employing the splatting method for rendering, the 3D reconstruction kernel must be integrated along the z -axis to evaluate its contribution to each pixel on the image plane. Since integrating a 3D Gaussian function along a given axis still results in a Gaussian function, we can effectively transform a 3D Gaussian $\mathcal{G}(\vec{x})$ into a 2D Gaussian $\mathcal{G}'(\vec{x}')$ on the image plane using the ray-coordinate transformation described in [50]:

$$\int_{\mathbb{R}} \mathcal{G}(\vec{x}, \vec{\mu}, \Sigma) dz = \mathcal{G}'(\vec{x}', \vec{\mu}', \Sigma'), \quad (2)$$

where $\vec{x}' = (x, y)^\top$ and $\vec{\mu}' = (\mu_x, \mu_y)^\top$ represent coordinates in the projected 2D image plane. To achieve accurate perspective projection, the original covariance matrix Σ undergoes transformation through the world-to-camera transformation matrix \mathbf{W} , followed by projection via the local affine Jacobian matrix \mathbf{J} . This yields a ray-space covariance matrix $\hat{\Sigma} = \mathbf{J}\mathbf{W}\Sigma\mathbf{W}^\top\mathbf{J}^\top$. Ultimately, the 2D covariance matrix Σ' in image space can be easily obtained by taking the upper left 2×2 submatrix of $\hat{\Sigma}$.

Finally, the pixel values on the 2D image are obtained by α -blending:

$$C(\vec{x}') = \sum_{i \in N} c_i \alpha_i \mathcal{G}'_i(\vec{x}', \vec{\mu}', \Sigma') \prod_{j=1}^{i-1} (1 - \alpha_j \mathcal{G}'_j(\vec{x}', \vec{\mu}', \Sigma')), \quad (3)$$

where \vec{x}' denotes the queried pixel position and N represents the set of sorted 2D Gaussians associated with \vec{x}' .

3.2. Expansion of 3D Gaussian Splatting

Gaussian-based representations have demonstrated exceptional performance not only in 3D scene signal modeling

Table 1. Quantitative comparison of image fitting on Kodak and DIV2K datasets. We color code each cell as **best** and **second best**.

Methods Metrics	Kodak				DIV2K			
	PSNR↑	SSIM↑	LPIPS↓	FPS↑	PSNR↑	SSIM↑	LPIPS↓	FPS↑
WIRE [31]	38.02	0.9541	0.0477	19.55	33.20	0.9565	0.1303	11.63
GSImage-RS [44]	43.09	0.9983	0.0216	2129.18	39.85	0.9978	0.0371	2190.40
GSImage-Cholesky	44.06	0.9985	0.0188	2061.99	39.53	0.9975	0.0430	2248.74
WIPES-RS	45.76	0.9988	0.0132	1708.28	40.34	0.9976	0.0360	1817.48
WIPES-Cholesky	45.87	0.9987	0.0120	1778.75	40.32	0.9978	0.0361	1830.10

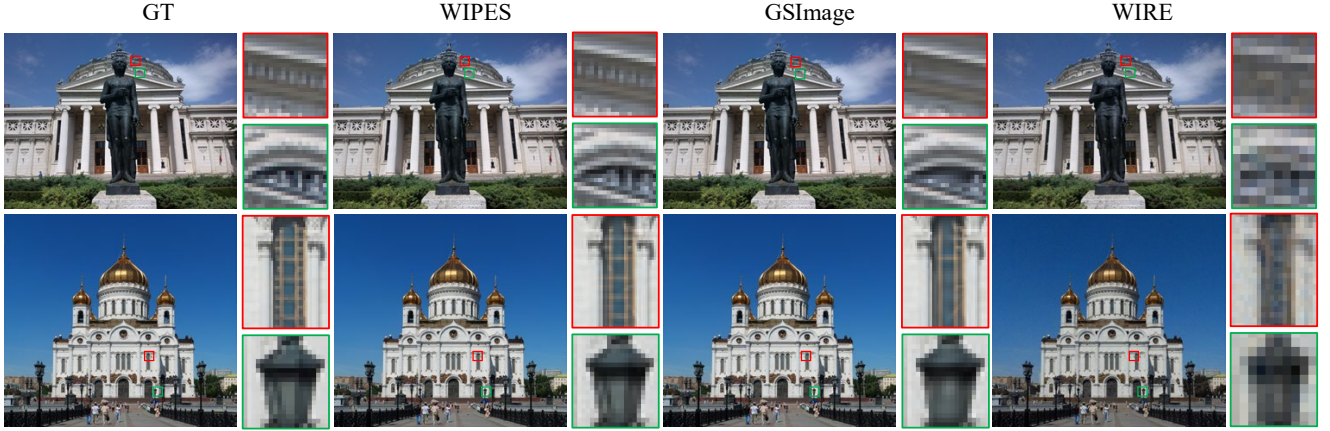


Figure 3. Qualitative comparison on Kodak and DIV2K datasets. Our approach effectively captures both low- and high-frequency details.

but also in various other domains. For instance, Gaussian Image [44] explicitly parameterizes 2D image signals using 2D Gaussian primitives defined directly on the image plane, replacing traditional α -blending with a simplified weighted summation mechanism:

$$C(\vec{x}') = \sum_{i \in N} c_i \alpha_i \mathcal{G}'_i(\vec{x}', \vec{\mu}', \Sigma'). \quad (4)$$

Moreover, to handle dynamic 3D content, D-3DGS [38] employs an MLP network \mathcal{F}_θ to model the deformation field of 3D Gaussian primitives. Thus, the primitives efficiently adapt to dynamic variations over time, yielding a high-fidelity representation of temporally-varying 3D scenes:

$$(\delta \vec{x}, \delta \Sigma) = \mathcal{F}_\theta(\gamma(\vec{x}), \gamma(t)), \quad (5)$$

where γ denotes positional encoding [24].

4. Wavelet-based Visual Primitives

4.1. 3D Splatting with Wavelet Primitives

The wavelet transform decomposes a signal into a linear combination of translated and scaled short oscillatory pulses, effectively capturing localized variations in both spatial and frequency domains. Due to their compact support in these domains, wavelets achieve superior accuracy in signal approximation [6]. This property substantiates the feasibility of using differentiable wavelet basis functions in our methodology. Inspired by the Morlet wavelet, we obtain

the 3D wavelet primitive by applying cosine modulation to the 3D Gaussian primitive. A 3D wavelet primitive is described by:

$$\mathcal{W}(\vec{x}, \vec{\mu}, \vec{f}, \Sigma) = \frac{1}{2} [\cos(\vec{f} \cdot (\vec{x} - \vec{\mu})) + 1] \mathcal{G}(\vec{x}, \vec{\mu}, \Sigma), \quad (6)$$

where $\vec{f} = (f_x, f_y, f_z)^\top$ is the frequency of cosine modulation. Notably, the Gaussian primitive can be considered a subset of the proposed 3D wavelet primitive by setting the parameter \vec{f} to $\vec{0}$.

Integrating the 3D wavelet primitive along the z -axis produces results similar to those described in Eq. (2):

$$\int_{\mathbb{R}} \mathcal{W}(\vec{x}, \vec{\mu}, \vec{f}, \Sigma) dz = \mathcal{W}'(\vec{x}', \vec{\mu}', \vec{f}', \Sigma'), \quad (7)$$

where $\vec{f}' = (f_x, f_y)^\top$ and $\vec{x}' - \vec{\mu}'$ is the same as the 3D Gaussian primitive in Eq. (2).

The projected 2D covariance matrix Σ' for each wavelet primitive is computed via a similar spatial transformation used in projecting Gaussian primitives from 3D world space to the 2D image plane. Specifically, the covariance matrix Σ and modulation frequency vector \vec{f} of the wavelet primitive undergo the identical ray-space transformation defined by matrices \mathbf{W} and \mathbf{J} , resulting in the transformed frequency vector \vec{f}' and the projected covariance matrix Σ' .

For rasterization, we follow a similar pipeline as in 3D GS while substituting the primitives with our proposed 3D wavelets. In the end, the volumetric α -blending can be




Figure 4. Comparisons of Gaussian and wavelet primitives for representing complex texture patterns. Results indicate that wavelet primitives effectively overcome the low-frequency bias of Gaussian primitives and demonstrate excellent performance in high-frequency texture regions.

easily obtained by replacing the Gaussian primitive \mathcal{G}' in Eq. (3) with the wavelet primitive \mathcal{W}' .

4.2. Wavelet-Based Frequency Flexibility

Gaussian functions inherently act as low-pass filters, as illustrated in Fig. 1a. Specifically, in the frequency domain, Gaussians with large variance exhibit narrow bandwidths concentrated around zero frequency, while those with small variance produce wider bandwidths. This bell-shaped Gaussian spectrum rapidly attenuates high-frequency signals, thus limiting its capacity to represent signals rich in high-frequency content. Consequently, a large number of Gaussian basis functions with smaller variances are inevitably required to accurately capture intricate high-frequency details.

In contrast, wavelets provide significantly enhanced frequency flexibility. By employing adaptive frequency modulation, wavelet primitives simultaneously encode both low- and high-frequency components within a single basis function. This modulation introduces a controllable shift in the wavelet spectrum, positioning the Gaussian-like envelope around a desired center frequency. The modulation frequency defines this spectral center, and the Gaussian envelope controls the bandwidth, ensuring optimal coverage of the spatial-frequency domain.

We validate this advantage by comparing Gaussian and wavelet primitives on two images exhibiting complex moiré patterns (Fig. 4). Under identical conditions and equal numbers of primitives, Gaussian primitives consistently failed to represent detailed high-frequency structures, resulting in noticeable modeling voids. In contrast, the wavelet-based representation accurately captured intricate frequency transitions, effectively modeling sharp discontinuities and fine details. This result demonstrates that wavelets, through their inherent frequency adaptability, offer superior effi-

ciency and accuracy in representing signals characterized by rapid spatial variations.

4.3. Expansion of 3D Wavelet Splatting

We further leverage wavelet primitives to enhance representation quality, especially for signals exhibiting complex spatial-frequency content. Specifically, we validate the effectiveness of wavelet primitives in two key applications: image fitting and dynamic novel view synthesis.

For the image fitting task, we replace the original 2D Gaussian primitives \mathcal{G}' defined in Eq. (4) with frequency-adaptive wavelet primitives \mathcal{W}' . This substitution introduces modulation frequencies \vec{f}' , significantly expanding frequency coverage compared to Gaussian-based models. In dynamic novel view synthesis, we further generalize wavelet primitives by making them temporally adaptive. The modified formulation is given by:

$$(\delta\vec{x}, \delta\vec{f}, \delta\mathbf{\Sigma}) = \mathcal{F}_\theta(\gamma(\vec{x}), \gamma(t)). \quad (8)$$

Comprehensive experimental validation demonstrating the superiority of our wavelet-based approach is provided in Sec. 5, confirming their ability to efficiently represent high-frequency details and complex temporal variations.

5. Experiments

Our experiments demonstrate the robustness and versatility of WIPES in representing real-world scenes, effectively capturing both high-frequency details (*e.g.*, edge variations and intricate textures) and low-frequency contours. To validate WIPES as a generalized signal representation framework, we systematically evaluate its performance across three distinct tasks: 2D image fitting, 5D static novel view synthesis, and 6D dynamic novel view synthesis.

All experiments are consistently executed on a single NVIDIA A6000 GPU. To ensure a fair comparison, we adopt the same initialization schemes, optimization strategies, and hyperparameter configurations as those of the baseline approaches. Specifically, frequency modulation coefficients are randomly initialized in alignment with competing methods to exclude initialization bias. Additionally, our implementation builds upon the recent `gsplat` library [39], augmented by custom CUDA kernels that enable wavelet-based differentiable rasterization blending. This adaptation is critical for efficiently leveraging wavelet primitives in rendering tasks.

5.1. 2D Image Fitting

Configurations. In this section, we utilized the Kodak [7] and DIV2K [1] datasets. The DIV2K dataset was down-scaled by a factor of 2 using bicubic interpolation, resulting in images with resolutions ranging from 408×1020 to 1020×1020 . For our WIPES, we conducted experiments

Table 2. Quantitative comparison of static novel view synthesis on Mip-NeRF360, Tanks & Temples and DeepBlending datasets. Results marked with † were adopted from the 3DGS paper [14], all others come from our experiments.

Methods Metrics	Mip-NeRF360				Tanks & Temples				DeepBlending			
	PSNR↑	SSIM↑	LPIPS↓	FPS↑	PSNR↑	SSIM↑	LPIPS↓	FPS↑	PSNR↑	SSIM↑	LPIPS↓	FPS↑
Plenoxels [9]	23.32	0.646	0.451	6.79	21.08	0.719	0.379	13	23.06	0.795	0.510	11.2
INGP [25]	25.19	0.667	0.440	11.7	21.72	0.723	0.330	17.1	23.62	0.797	0.423	3.26
Mip-NeRF360† [3]	27.69	0.792	0.237	0.06	22.22	0.759	0.257	0.14	29.40	0.951	0.245	0.09
3DGS [14]	27.28	0.811	0.220	118.2	23.35	0.835	0.182	162.4	29.43	0.898	0.246	124.1
WIPES	27.55	0.816	0.215	95.7	23.78	0.852	0.171	126.3	29.82	0.907	0.238	92.6




Figure 5. Qualitative comparison of static novel view synthesis on Mip-NeRF360, Tanks & Temples and DeepBlending datasets.

using the two covariance construction methods proposed in GSImage, namely Cholesky and RS. The frequency coefficients were randomly initialized with a normal distribution, while all other learnable parameters followed the initialization strategy described in GSImage [44]. To ensure a fair comparison with WIRE [31] and GSImage [44], all hyperparameters were kept consistent with those used in GSImage.

Results. Tab. 1 presents a quantitative comparison between

our proposed WIPES and existing advanced INR-based and Gaussian-based approaches. Our approach achieves the best performance across all metrics, including PSNR, SSIM, and LPIPS. Additionally, as illustrated in Fig. 3, WIPES effectively captures both low-frequency features (*e.g.*, walls and gates) and high-frequency details (*e.g.*, tiles and streetlamps) within the scenes. Although GSImage demonstrates the highest rendering speed, our approach maintains a competitive frame rate, ensuring efficient per-

Table 3. Quantitative comparison on static novel view synthesis with GES [11], which incorporates frequency guidance, across the Mip-NeRF360, Tanks & Temples and DeepBlending datasets. GES(3e-4) uses the default parameters from the original GES paper, while GES(2e-4) adopts the same settings as 3DGS [14].

(a) Mip-NeRF360				
Methods Metrics	PSNR↑	SSIM↑	LPIPS↓	Points↓
GES(3e-4)	26.98	0.7208	0.3257	1.521M
WIPES(3e-4)	27.11	0.7958	0.2518	1.401M
GES(2e-4)	27.36	0.8148	0.2172	3.055M
WIPES(2e-4)	27.55	0.8164	0.2149	2.787M

(b) Tanks & Temples				
Methods Metrics	PSNR↑	SSIM↑	LPIPS↓	Points↓
GES(3e-4)	23.44	0.8394	0.1954	0.913M
WIPES(3e-4)	23.76	0.8433	0.1914	0.931M
GES(2e-4)	23.59	0.8481	0.1772	1.601M
WIPES(2e-4)	23.78	0.8523	0.1709	1.591M

(c) DeepBlending				
Methods Metrics	PSNR↑	SSIM↑	LPIPS↓	Points↓
GES(3e-4)	29.59	0.9053	0.2492	1.243M
WIPES(3e-4)	29.72	0.9062	0.2501	1.137M
GES(2e-4)	29.55	0.9041	0.2434	2.059M
WIPES(2e-4)	29.82	0.9069	0.2383	1.933M

formance without sacrificing quality. This balance between high-quality image reconstruction and rendering efficiency underscores the practical effectiveness of our approach.

5.2. 5D Static Novel View Synthesis

Configurations. In this section, we evaluated WIPES on 13 scenes sourced from the Mip-NeRF360 [3], Tanks & Temples [16], and DeepBlending [12] datasets. Specifically, we set the learning rate for the frequency coefficients to 2.5×10^{-3} in WIPES. Differentiable rasterization was implemented via custom CUDA kernels, while all other learnable parameters and experimental settings remained aligned with those of 3DGS [14]. This configuration facilitated a direct comparison with state-of-the-art approaches including Plenoxels [9], INGP [25], Mip-NeRF360, and 3DGS.

Results. Tab. 2 presents a quantitative comparison between WIPES and leading 5D novel view synthesis approaches, demonstrating that our approach outperforms existing methods in most scenarios. Notably, WIPES exhibits significant advantages in environments with complex frequency components, such as large-scale outdoor scenes. Additionally, Fig. 5 provides qualitative comparisons for selected indoor and outdoor scenes from the dataset. The results indicate that WIPES excels at reconstructing reflective environments (e.g., mirror frames and truck windshields) and effectively captures high-frequency details (e.g., potted plants and distant transmission towers).

Comparison with frequency-constrained approaches.

We compare WIPES with Gaussian-based approaches that utilize frequency priors or guidance. Since FreGS [43] has not released its source code, we present quantitative results against GES [11]. GES employed a density gradient threshold of $3e^{-4}$, whereas the original 3DGS [14] used $2e^{-4}$. To ensure fairness, we evaluated GES with both thresholds. Experiments on the Mip-NeRF360, Tanks & Temples, and DeepBlending datasets (Tab. 3) demonstrate that WIPES achieves superior reconstruction without requiring frequency guidance. Additionally, when using identical density gradient thresholds, WIPES matches GES in parameter count, highlighting its ability to deliver enhanced results without relying on frequency priors or guidance while maintaining comparable parameter efficiency.

5.3. 6D Dynamic Novel View Synthesis

Configurations. In this section, we employed the D-NeRF [26] and NeRF-DS [37] datasets to assess performance under varying dynamic conditions. WIPES utilized a frequency coefficient learning rate of 2.5×10^{-3} , consistent with the settings of D-3DGS [38]. All other learnable parameters were maintained in alignment with D-3DGS to ensure a fair comparison. Differentiable rasterization was achieved through the integration of custom CUDA kernels, mirroring the configurations used in the 6D tasks.

Results. Tab. 4 presents the quantitative experimental results of WIPES in comparison with INR-based approaches and Gaussian-based approaches referenced in our work. WIPES exhibits enhanced expressive capabilities for various frequency components in dynamic scenes, enabling similar reconstruction performance to D-3DGS with fewer primitives. This indicates that the deformable INR in WIPES more effectively learns the dynamic features of primitives. Consequently, WIPES achieves superior performance in both reconstruction quality and rendering speed. Fig. 6 showcases qualitative results of WIPES across different dynamic datasets. It is evident that WIPES delivers more detailed and accurate results for both low-frequency elements, such as cloth, and high-frequency textures.

6. Conclusion

In this paper, we have proposed WIPES, a universal wavelet-based visual primitive for representing multi-dimensional visual signals. We have enhanced the Morlet function by designing a universally applicable and continuously differentiable wavelet primitive. Without relying on any frequency guidance or priors, WIPES has adaptively captured various real-world visual signals with rich frequency components, achieving impressive results in reconstruction tasks for 2D image representation, 5D static, and 6D dynamic novel view synthesis. To facilitate WIPES, we have developed a fast differentiable wavelet rasterizer that

Table 4. Quantitative comparison of dynamic novel view synthesis on D-NeRF and NeRF-DS datasets.

Methods Metrics	D-NeRF				NeRF-DS			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FPS \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FPS \uparrow
D-NeRF [26]	30.43	0.9572	0.0724	0.117	—	—	—	—
NeRF-DS [37]	—	—	—	—	23.58	0.8582	0.2496	0.008
D-3DGS [38]	39.16	0.9895	0.0132	71.7	23.69	0.8423	0.1976	38.36
WIPES	39.52	0.9899	0.0127	84.0	23.95	0.8527	0.1762	42.45

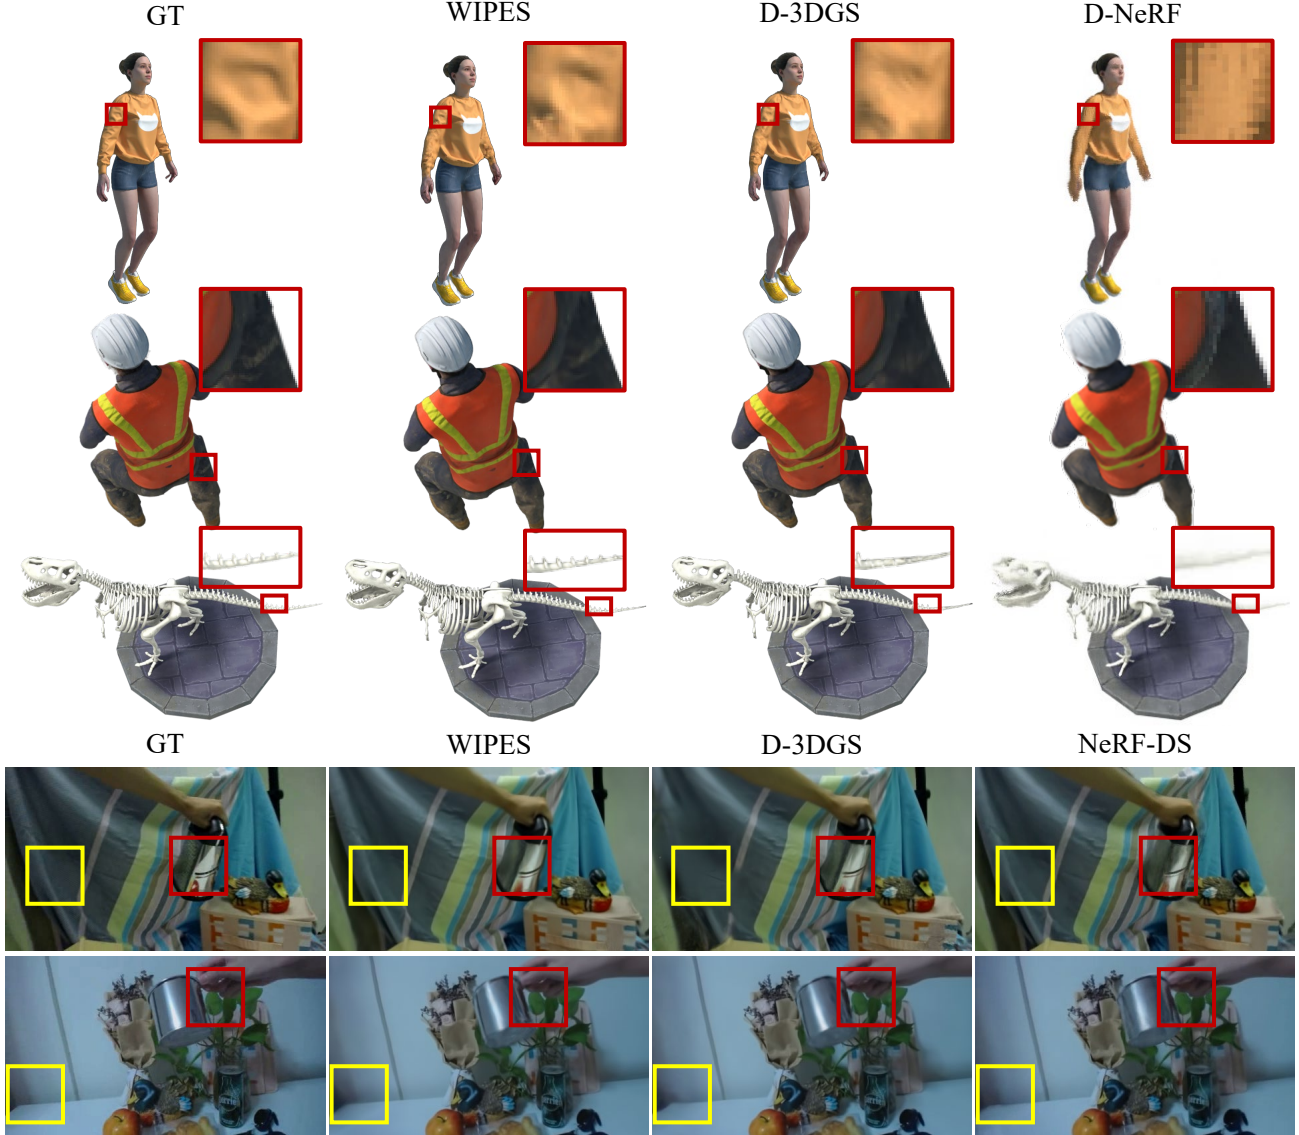


Figure 6. Qualitative comparison on D-NeRF and NeRF-DS datasets. Our approach demonstrates more accurate registration, delivering outstanding performance in both static and dynamic regions.

seamlessly integrates into Gaussian-based pipelines, enhancing rendering speed and efficiency during reconstruction. We have conducted extensive experiments on several challenging datasets across different dimensions, demonstrating the generality and effectiveness of our approach through both qualitative and quantitative evaluations.

However, our framework has limitations. The current Gaussian-based training pipeline has restricted scene

representation due to fixed hyperparameters. Additionally, gradient-based densification in Gaussian approaches has caused instability and performance variability. Future work will aim to develop a more robust and efficient wavelet-based training and rendering framework to address these issues, fully leveraging WIPES’s potential and advancing continuous signal representation in computer vision.

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 126–135, 2017. 5
- [2] Adrian Azzarelli, Nantheera Anantprasirichai, and David R. Bull. Waveplanes: A compact wavelet representation for dynamic neural radiance fields. *arXiv preprint arXiv:2312.02218*, 2023. 3
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5460–5469, 2022. 6, 7
- [4] Yuyao Chen, Lu Lu, George Em Karniadakis, and Luca Dal Negro. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Opt. Express*, 28(8): 11618–11633, 2020. 2
- [5] Jaeyoung Chung, Jeongtaek Oh, and Kyoung Mu Lee. Depth-regularized optimization for 3d gaussian splatting in few-shot images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 811–820, 2024. 2
- [6] Ronald A. DeVore. Nonlinear approximation. *Acta Numerica*, 7:51–150, 1998. 4
- [7] Eastman Kodak Company. Kodak lossless true color image suite. <https://r0k.us/graphics/kodak/>, 1999. 5
- [8] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J. Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations (ICLR)*, 2021. 3
- [9] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 2, 6, 7
- [10] A. Graps. An introduction to wavelets. *IEEE Computational Science and Engineering*, 2(2):50–61, 1995. 2
- [11] Abdullah Hamdi, Luke Melas-Kyriazi, Jinjie Mai, Guocheng Qian, Ruoshi Liu, Carl Vondrick, Bernard Ghanem, and Andrea Vedaldi. Ges: Generalized exponential splatting for efficient radiance field rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19812–19822, 2024. 2, 7
- [12] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics*, 37(6):257, 2018. 7
- [13] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Papers '24*. ACM, 2024. 2
- [14] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)*, 42(4), 2023. 2, 3, 6, 7
- [15] Rajaei Khatib. Trinerflet: A wavelet-based triplane nerf representation. In *ECCV 2024 Workshop on Wild 3D: 3D Modeling, Reconstruction, and Generation in the Wild*, 2024. 3
- [16] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *Association for Computing Machinery*, 36(4):78, 2017. 7
- [17] Amit Pal Singh Kohli, Vincent Sitzmann, and Gordon Wetzstein. Semantic implicit neural scene representations with semi-supervised training. In *2020 International Conference on 3D Vision (3DV)*, pages 423–433, 2020. 2
- [18] Christoph Lassner and Michael Zollhöfer. Pulsar: Efficient sphere-based neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1440–1449, 2021. 2
- [19] David B. Lindell, Dave Van Veen, Jeong Joon Park, and Gordon Wetzstein. Bacon: Band-limited coordinate networks for multiscale scene representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [20] Zhen Liu, Hao Zhu, Qi Zhang, Jingde Fu, Weibing Deng, Zhan Ma, Yanwen Guo, and Xun Cao. 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*, pages 2713–2722, 2024. 2
- [21] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20654–20664, 2024. 2
- [22] Stéphane Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999. 3
- [23] Julien N. P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. *ACM Transactions on Graphics (TOG)*, 40(4), 2021. 2
- [24] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020. 2, 4
- [25] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):1–15, 2022. 2, 6, 7
- [26] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 7, 8
- [27] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 5301–5310. PMLR, 2019. 1, 2

- [28] Sameera Ramasinghe and Simon Lucey. Beyond periodicity: Towards a unifying framework for activations in coordinate-mlps. In *European Conference on Computer Vision*, pages 142–158. Springer, 2022. 2
- [29] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14335–14345, 2021. 2
- [30] Vishwanath Saragadam, Jasper Tan, Guha Balakrishnan, Richard G. Baraniuk, and Ashok Veeraraghavan. Miner: Multiscale implicit neural representation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–333. Springer, 2022. 2
- [31] Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G. Baraniuk. Wire: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18507–18516, 2023. 2, 4, 6
- [32] Chengkang Shen, Hao Zhu, You Zhou, Yu Liu, Si Yi, Lili Dong, Weipeng Zhao, David J Brady, Xun Cao, Zhan Ma, et al. Cardiacfield: computational echocardiography for automated heart function estimation using two-dimensional echocardiography probes. *European Heart Journal-Digital Health*, page ztae072, 2024. 2
- [33] Chengkang Shen, Hao Zhu, You Zhou, Yu Liu, Si Yi, Lili Dong, Weipeng Zhao, David J Brady, Xun Cao, Zhan Ma, et al. Continuous 3d myocardial motion tracking via echocardiography. *IEEE Transactions on Medical Imaging*, 2024. 2
- [34] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, pages 7462–7473, 2020. 1, 2
- [35] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:7537–7547, 2020. 1, 2
- [36] Delong Wu, Hao Zhu, Qi Zhang, You Li, Zhan Ma, and Xun Cao. Neural poisson solver: A universal and continuous framework for natural signal blending. In *Computer Vision – ECCV 2024*, pages 259–275, Cham, 2025. Springer Nature Switzerland. 2
- [37] Zhiwen Yan, Chen Li, and Gim Hee Lee. Nerf-ds: Neural radiance fields for dynamic specular objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8285–8295, 2023. 7, 8
- [38] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20331–20341, 2024. 2, 4, 7, 8
- [39] Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for gaussian splatting. *arXiv preprint arXiv:2409.06765*, 2024. 5
- [40] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19447–19456, 2024. 2
- [41] Gizem Yüce, Guillermo Ortiz-Jiménez, Beril Besbinar, and Pascal Frossard. A structured dictionary perspective on implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19228–19238, 2022. 1
- [42] Ruyi Zha, Tao Jun Lin, Yuanhao Cai, Jiwen Cao, Yanhao Zhang, and Hongdong Li. R²-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2
- [43] Jiahui Zhang, Fangneng Zhan, Muyu Xu, Shijian Lu, and Eric Xing. Fregs: 3d gaussian splatting with progressive frequency regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21424–21433, 2024. 2, 7
- [44] Xinjie Zhang, Xingtong Ge, Tongda Xu, Dailan He, Yan Wang, Hongwei Qin, Guo Lu, Jing Geng, and Jun Zhang. Gaussianimage: 1000 fps image representation and compression by 2d gaussian splatting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. 1, 2, 4, 6
- [45] Hao Zhu, Zhen Liu, You Zhou, Zhan Ma, and Xun Cao. Dnf: Diffractive neural field for lensless microscopic imaging. *Optics Express*, 30(11):18168–18178, 2022. 2
- [46] Hao Zhu, Zhen Liu, Qi Zhang, Jingde Fu, Weibing Deng, Zhan Ma, Yanwen Guo, and Xun Cao. Finer++: Building a family of variable-periodic functions for activating implicit neural representation. *arXiv preprint arXiv:2407.19434*, 2024. 1, 2
- [47] Hao Zhu, Shaowen Xie, Zhen Liu, Fengyi Liu, Qi Zhang, You Zhou, Yi Lin, Zhan Ma, and Xun Cao. Disorder-invariant implicit neural representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 2
- [48] Hao Zhu, Fengyi Liu, Qi Zhang, Zhan Ma, and Xun Cao. Rhino: Regularizing the hash-based implicit neural representation. *SCIENCE CHINA Information Sciences*, pages –, 2025. 2
- [49] Junyu Zhu, Hao Zhu, Qi Zhang, Fang Zhu, Zhan Ma, and Xun Cao. Pyramid nerf: Frequency guided fast radiance field optimization. *International Journal of Computer Vision*, pages 1–16, 2023. 2
- [50] Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. Ewa volume splatting. In *Proceedings Visualization, 2001. VIS '01.*, pages 29–538, 2001. 2, 3
- [51] Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. Surface splatting. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, pages 371–378, New York, NY, USA, 2001. Association for Computing Machinery. 2