

EVPGS: Enhanced View Prior Guidance for Splatting-based Extrapolated View Synthesis

Jiahe Li Feiyu Wang* Xiaochao Qu Chengjing Wu Luoqi Liu* Ting Liu*
 MT Lab, Meitu Inc., Beijing 100083, China
 {ljh15, wfy1, qxc, ethan, llq5, lt}@meitu.com
 *Corresponding Author

Abstract

Gaussian Splatting (GS)-based methods rely on sufficient training view coverage and perform synthesis on interpolated views. In this work, we tackle the more challenging and underexplored Extrapolated View Synthesis (EVS) task. Here we enable GS-based models trained with limited view coverage to generalize well to extrapolated views. To achieve our goal, we propose a view augmentation framework to guide training through a coarse-to-fine process. At the coarse stage, we reduce rendering artifacts due to insufficient view coverage by introducing a regularization strategy at both appearance and geometry levels. At the fine stage, we generate reliable view priors to provide further training guidance. To this end, we incorporate an occlusion awareness into the view prior generation process, and refine the view priors with the aid of coarse stage output. We call our framework Enhanced View Prior Guidance for Splatting (EVPGS). To comprehensively evaluate EVPGS on the EVS task, we collect a real-world dataset called Merchandise3D dedicated to the EVS scenario. Experiments on three datasets including both real and synthetic demonstrate EVPGS achieves state-of-the-art performance, while improving synthesis quality at extrapolated views for GS-based methods both qualitatively and quantitatively. Our code and dataset are available on the [EVPGS Homepage](#).

1. Introduction

Gaussian Splatting (GS) [15] and its variants [16, 19, 24, 42] have enabled high-fidelity and real-time novel view synthesis. Despite the achievements, a main drawback of the existing methods is that they rely on training images with sufficient view coverage, which is an interpolated view synthesis setting. In real-world scenarios, users often capture images by moving the camera horizontally around an object, yet may desire novel views from elevated pitch angles. Therefore, existing methods suffer from deteriorated




Figure 1. **Illustration of the EVS problem.** We visualize the camera pose at each view as a rectangle warped by the pose. **Left:** In conventional novel view synthesis, the training views sufficiently cover the scene. **Right:** In EVS, the training views have limited scene coverage (only horizontal and near-horizontal views as per the example) and overlap poorly with the testing views. Example scene from DTU [14].

synthesis quality in real-world applications due to both limited training view coverage and the resulting large disparity between the training and testing views. This gives rise to the extrapolated view synthesis (EVS) problem, which we tackle in this work. We illustrate EVS in Figure 1.

To remedy for the gap between the training and testing views, concurrent works in sparse view reconstruction [18, 36, 41] perform view augmentation by generating view priors at augmented views, which act as pseudo-labels during training. The view prior is typically generated by reprojecting a training-view image to an unseen augmented view via a reconstructed explicit 3D representation, e.g. mesh or depth map. While these sparse view reconstruction methods can produce novel-view renderings with reasonable realism, they essentially deal with interpolated view synthesis. We scrutinise the EVS problem and identify two challenges which have not been addressed by existing strategies for generating view priors, i.e. lack of view-dependent color information, and occlusions in complex scenes. We illustrate these challenges in Figure 2. These issues can cause rendering artifacts such as unrealistic appearance and corrupted scene parts, which call for more dedicated processing of the view priors.

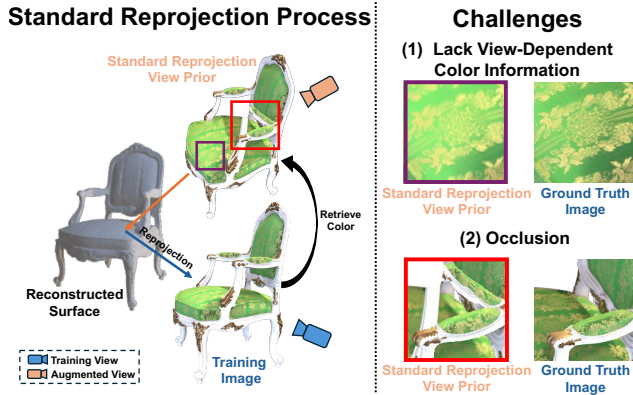


Figure 2. **Illustration of the standard reprojection process and its challenges.** **Left:** In the standard reprojection process, each pixel in the view prior of an augmented view retrieves color from the corresponding pixel in the nearest training view, using reprojection via the reconstructed mesh. **Right:** This technique risks losing view-dependent color information for the augmented views and may introduce occlusion artifacts, leading to corrupted view priors.

To address both challenges above, we present a comprehensive view augmentation framework for EVS called **Enhanced View Prior Guidance for Splatting (EVPGS)**. Our framework generates reliable view priors through a two-stage coarse-to-fine process. At the coarse stage, we leverage both the prior knowledge from the Denoising Diffusion Model (DDM) [11, 25] and the reconstructed mesh from pre-trained GS Model to reduce rendering artifacts at extrapolated views caused by insufficient view coverage. At the fine stage, we incorporate the occlusion-aware strategy into the view prior generation to eliminate the occlusion issue encountered by the standard reprojection technique (illustrated in Figure 2). We also refine the view priors by integrating the view-dependent color information from the regularized renderings at the coarse stage. Our fine stage can produce high-quality view priors as strong appearance guidance towards synthesizing realistic extrapolated views. We term the view priors generated by the EVPGS framework *Enhanced View Priors*.

To comprehensively evaluate our EVPGS framework, we collect *Merchandise3D*, a real-world dataset specifically tailored for the EVS task. This dataset includes seven object instances of several categories with intricate structural and textural details, representing challenging real-world scenarios. We evaluate EVPGS on three different benchmark datasets, both real and synthetic, including DTU [14], Synthetic-NeRF [20], and our Merchandise3D. Experimental results demonstrate the superior performance of EVPGS over the different datasets.

In summary, our main contributions are as follows,

1. To address the challenges posed by EVS, we propose the view augmentation framework EVPGS based on GS

models to generate reliable view priors as high-quality augmented views.

2. We propose several complementary strategies to improve view prior quality through a coarse-to-fine process. At the coarse stage, we introduce a regularization strategy at both appearance and geometry levels to reduce rendering artifacts due to insufficient training view coverage. Then at the fine stage, we generate reliable view priors by both eliminating the occlusion issue in the standard reprojection technique, and refining the view prior with view specific visual information.
3. EVPGS can synthesize extrapolated views with realistic appearance and fine details. Experiments demonstrate EVPGS achieves state-of-the-art performance on three EVS benchmarks, including both public and our self-collected datasets. We release our Merchandise3D dataset to facilitate future research.

2. Related Work

2.1. Neural Scene Representations

Neural scene representations have been extensively studied for efficient and high-quality 3D reconstruction. The pioneering work Neural Radiance Fields (NeRF) [1, 20] utilizes a deep Multi-layer Perceptron to encode volume density and color across the continuous 3D space and employs the alpha-composition technique for rendering into 2D images. [30, 31, 39] introduces explicit surface representation for better reconstruction. Several other works [6, 7, 10, 21, 34, 40] adopt the voxel-grid representation to expedite training and rendering. To reduce memory cost, [3, 7] decompose the 3D volume into a set of 2D/1D planes/vectors, representing the color and opacity at each 3D point with features stored in the planes/vectors.

A more recent line of works [9, 12, 15, 16, 19, 24, 42–44] adopt the splatting technique [48] for rendering, achieving significant improvement in both efficiency and quality. The seminal work 3D Gaussian Splatting (3DGS) [15] proposes the Gaussian primitive representation which is flexible and parameter-efficient while being conducive to parallelized implementation. Mip-Splatting [42] introduces kernel filtering techniques to alleviate aliasing artifacts caused by resolution changes. By exploiting the characteristics of Gaussian primitives, [9, 12, 43, 44] design dedicated approaches to improve surface reconstruction. Our framework is generic and can be implemented with different GS-based methods as backbones.

2.2. View Augmentation

Sparse View Reconstruction: While sparse view reconstruction [2, 4, 17, 18, 22, 23, 29, 32, 36, 38, 41, 47] aims to improve view interpolation between a small set of training views, several works [18, 36, 41, 47] incorporate view

augmentation strategies. [18, 36] leverage the reprojection technique to generate view priors, which provide training supervision at additional views. [41] improves the geometric consistency across different augmented views by exploiting the point cloud generated with Structure-from-Motion (SfM) [26, 27]. [47] leverages augmented views to encourage the newly cloned Gaussians to be distributed around the more representative locations.

Extrapolated View Synthesis: Several other works utilize view augmentation strategies to better synthesize extrapolated views [13, 37, 45]. Inspired by the early method based on depth probability volume [5], [37] utilizes the reconstructed pseudo-depth map with additional view coverage information to produce priors at extrapolated views with reasonable fidelity. [45] leverages the reprojection technique to apply supervision on a set of augmented pixels based on reconstructed mesh vertices. Different from the methods using augmented views for pixel-wise supervision, [13] proposes an DDM-based [25] regularization strategy to remove artifacts caused by insufficient training view coverage at the image semantics level. In our framework, we encourage the backbone to extrapolate to unseen views by adopting a similar strategy to [13], while we additionally consider both problems of lacking view-dependent color information and occlusion issue during view prior generation for further improving fidelity of the view prior.

3. Method

3.1. The EVPGS Framework

In the EVS task, the goal is to synthesize extrapolated views by leveraging a training dataset with only limited view coverage. As shown in Figure 1, there can be a drastic discrepancy between the training and testing view coverages, which hampers the learned GS model from synthesizing high-quality images at testing views.

Our key idea is based on exploiting view priors as augmented views during training. We propose the EVPGS framework for generating reliable view priors which we term *Enhanced View Priors*. Our EVPGS framework conducts artifact removal and appearance refinement for GS models in a coarse-to-fine manner. We illustrate our overall framework in Figure 3.

We first pre-train a GS model using the training views with limited coverage. Then at the coarse stage, we propose the Appearance and Geometry Regularization (AGR) strategy to fine-tune the pre-trained model across the augmented views. The regularization objective is designed by leveraging both a pre-trained DDM [25] and the reconstructed mesh from the pre-trained GS model. Benefiting from the strong appearance priors in DDM, our appearance regularization can produce visually sanitized images with significantly reduced artifacts. To further reduce artifacts, we

additionally propose the geometry regularization strategy which indirectly improves appearance by smoothing the reconstructed surface at the augmented views. This is implemented by supervising the directly rendered depth map from pre-trained GS model through mesh-rasterized depth map from the reconstructed mesh.

At the fine stage, we guide EVS towards both better realism and finer details by introducing the Occlusion-Aware Reprojection and Refinement (OARR) strategy, which generates *Enhanced View Priors* through both occlusion removal and view prior refinement. To eliminate occlusion during view prior generation, we extend the standard reprojection technique [18, 36, 45] by introducing the reverse-reprojection process, and replace the occlusion-corrupted pixels with the coarse stage results by designing a selection criterion based on reverse-reprojection. As a result of reprojection, the view prior only contains visual information from its corresponding training image and lacks the color information specific to its view. To compensate for the view-specific visual information, we introduce a simple but effective refinement strategy to incorporate the complementary visual information from coarse stage renderings into the view priors.

While each stage of EVPGS can achieve significant improvement individually, we combine both stages into one comprehensive framework for the best performance. We present our proposed components in detail below.

3.2. Appearance and Geometry Regularization

Appearance Regularization: At the coarse stage, we fine-tune the pre-trained GS model augmented views to reduce artifacts due to insufficient training view coverage. To achieve this goal, we employ the estimated noise of Denoising Diffusion Model (DDM) [25] as the evaluation criterion for the augmented views.

For each rendered image $\hat{\mathcal{I}}$ at an augmented view, the DDM produces the noise estimate through a diffusion and denoising process. In the diffusion process, DDM gradually adds random Gaussian noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to $\hat{\mathcal{I}}$ following the process $\hat{\mathcal{I}}_\tau = \sqrt{\bar{\alpha}_\tau} \hat{\mathcal{I}} + (1 - \bar{\alpha}_\tau) \epsilon$, where $\hat{\mathcal{I}}_\tau$ is the noisy version of $\hat{\mathcal{I}}$, τ is the time step [11] and $\bar{\alpha}_\tau$ is the pre-defined noise variance schedule. Then in the denoising process, the noise estimation is output by the denoising model $\epsilon_\theta(\cdot, \tau)$, where θ is the model parameters. Previous works [11, 28] have shown the estimated noise of DDM is proportional to the score function, which is the negative gradient of the log-likelihood of $\hat{\mathcal{I}}$ in our scenario, i.e.,

$$\epsilon_\theta(\hat{\mathcal{I}}_\tau, \tau) \propto -\nabla_{\hat{\mathcal{I}}} \log p(\hat{\mathcal{I}}) \quad (1)$$

This equation suggests that by computing the estimated noise, we can obtain the direction of the image gradient $\nabla_{\hat{\mathcal{I}}} \log p(\hat{\mathcal{I}})$ which drives $\hat{\mathcal{I}}$ towards the data distribution modes captured by DDM. Therefore, as suggested in [13, 35], we formulate our appearance regularization term

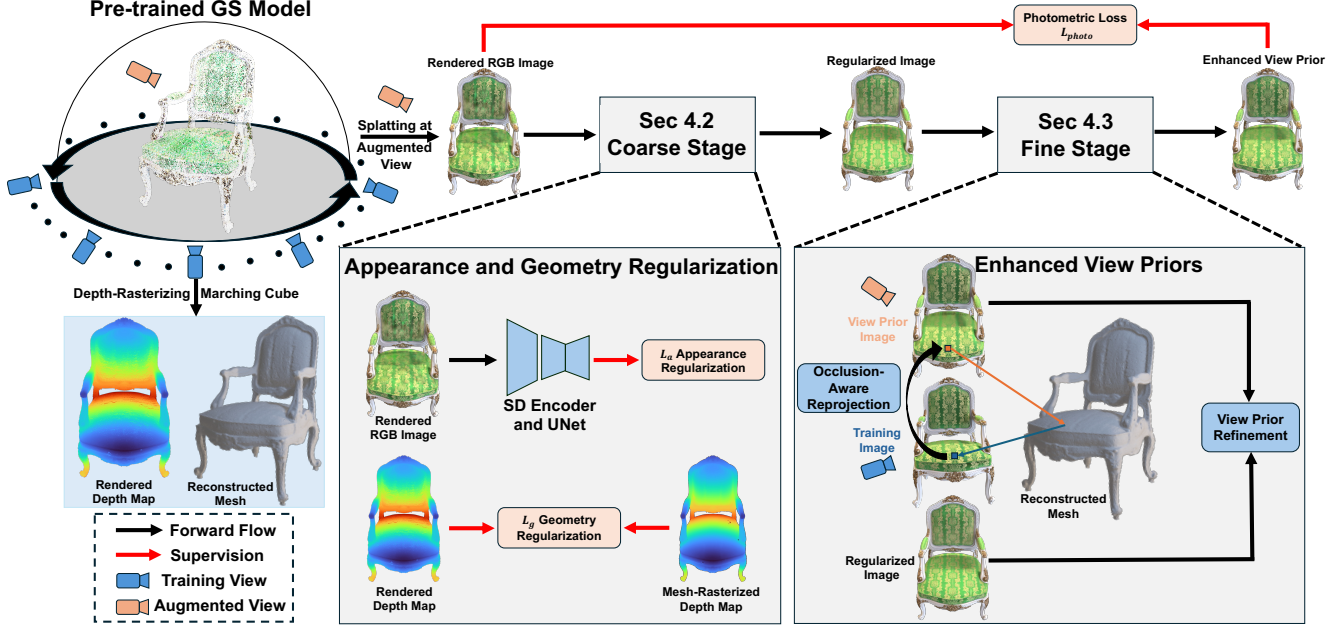


Figure 3. **Framework Overview.** In EVPGS, we first pre-train a GS model using the training set with limited view coverage (e.g. only horizontal views), then fine-tune the GS model at augmented views (e.g. obtained by elevating the training views) through a coarse-to-fine process. At the coarse stage, we propose the Appearance and Geometry Regularization (AGR) strategy, where we use the pre-trained GS model to produce synthesis at augmented views and reduce the artifacts at these views by leveraging the Denoising Diffusion Model [25]. We additionally generate depth maps rasterized from reconstructed mesh to supervise the depth maps directly rendered from the GS model. At the fine stage, we produce *Enhanced View Priors* at these augmented views as pseudo-labels, via our Occlusion-Aware Reprojection and Refinement (OARR) strategy. The OARR strategy comprises both the occlusion-aware reprojection technique to eliminate occlusion corruptions, and the view prior refinement strategy to incorporate the view-dependent colors obtained from the coarse stage.

\mathcal{L}_a as follows¹,

$$\mathcal{L}_a = \sum (\epsilon_\theta(\hat{\mathcal{I}}_\tau, \tau) \odot \hat{\mathcal{I}}_\tau) \quad (2)$$

where \odot is the Hadamard product. Unlike previous works [13, 35] which fine-tune the DDM model to adapt to their scenarios, we avoid fine-tuning the DDM model, thereby reducing computational costs and mitigating the risk of catastrophic forgetting.

Geometry Regularization: In the coarse stage of our framework, we observe the depth maps rendered from the pre-trained GS model (i.e. denoted as \bar{D}) at augmented views are usually inaccurate, which deteriorates the rendering results at the extrapolated views. As a result, we use the reconstructed mesh from the pre-trained model to smooth the rendered depth map, further reducing severe artifacts. To smooth the rendered depth map, for each augmented view, we use the pseudo-depth map (i.e. denoted as \hat{D}) rasterized from the reconstructed mesh to enforce geometry regularization, which we formulate as supervision on the rendered depth map \bar{D} , i.e.,

$$\mathcal{L}_g = \frac{1}{H} \frac{1}{W} \sum_{i=1}^H \sum_{j=1}^W |\hat{D}_{ij} - \bar{D}_{ij}| \quad (3)$$

¹Following the derivation of [35], the term $\epsilon_\theta(\hat{\mathcal{I}}_\tau, \tau)$ in Eq. 2 does not receive gradient when optimizing the GS model.

where H and W are the image height and width.

3.3. Enhanced View Priors

At the fine stage, we generate reliable view priors to further improve realism at augmented views while retaining both the view-dependent effects inherit from the coarse stage and the high-frequency details which can be lost after the coarse stage. This is achieved by our joint strategy to both remove occlusion corruption and refine the view priors. We term this strategy Occlusion-Aware Reprojection and Refinement (OARR) strategy.

Occlusion-Aware Reprojection: We first generate view priors based on our improved reprojection technique. For each augmented view v , we rasterize the depth map \bar{D}_v at this augmented view and the depth map \bar{D}_t at the nearest training view t using reconstructed mesh from the pre-trained GS model. For each pixel coordinate $\mathbf{p}_v \in \mathbb{R}^2$ in the augmented view, the reprojection technique calculates its correspondence pixel $\mathbf{p}_{v \rightarrow t} \in \mathbb{R}^2$ in the nearest training view, i.e.,

$$d_{\mathbf{p}_{v \rightarrow t}} \begin{pmatrix} \mathbf{p}_{v \rightarrow t} \\ 1 \end{pmatrix} = \mathbf{K}_t (d_{\mathbf{p}_v} \mathbf{R} \mathbf{K}_v^{-1} \begin{pmatrix} \mathbf{p}_v \\ 1 \end{pmatrix} + \mathbf{t}) \quad (4)$$

where \mathbf{K}_v and \mathbf{K}_t are the camera intrinsic matrices from the views v and t , \mathbf{R} and \mathbf{t} are the rotation matrix and translation vector both mapping from the camera space of the

augmented view to the camera space of the training view, $d_{\mathbf{p}_v}$ and $d_{\mathbf{p}_{v \rightarrow t}}$ are the depths of \mathbf{p}_v and $\mathbf{p}_{v \rightarrow t}$ from \bar{D}_v and \bar{D}_t , respectively. Once the correspondence $\mathbf{p}_{v \rightarrow t}$ is determined, the color at \mathbf{p}_v in the augmented view can be obtained by passing the color at $\mathbf{p}_{v \rightarrow t}$ in the training view to \mathbf{p}_v . Since $\mathbf{p}_{v \rightarrow t}$ is usually a sub-pixel coordinate, we calculate its color using bilinear interpolation based on its four nearest pixels.

To avoid potential occlusion when generating view priors, we extend the standard reprojection technique by introducing the reverse-reprojection technique which eliminates occlusion corruptions. Our technique finds the reverse-correspondence in the augmented view for each $\mathbf{p}_{v \rightarrow t}$ following a reverse manner to Eq. 4, i.e.,

$$d_{\mathbf{p}_{t \rightarrow v}} \begin{pmatrix} \mathbf{P}_{t \rightarrow v} \\ 1 \end{pmatrix} = \mathbf{K}_v(d_{\mathbf{p}_{v \rightarrow t}} \mathbf{R}^T \mathbf{K}_t^{-1} \begin{pmatrix} \mathbf{P}_{v \rightarrow t} \\ 1 \end{pmatrix} - \mathbf{R}^T \mathbf{t}) \quad (5)$$

where $\mathbf{p}_{t \rightarrow v}$ is the reverse-correspondence of $\mathbf{p}_{v \rightarrow t}$ and $d_{\mathbf{p}_{t \rightarrow v}}$ is the depth of $\mathbf{p}_{t \rightarrow v}$. We observe the pixel distance $e = \|\mathbf{p}_v - \mathbf{p}_{t \rightarrow v}\|_2$, which we call reverse-reprojection error, can be used as a criterion for detecting occlusion, i.e. the larger the distance, the more likely \mathbf{p}_v is corrupted by occlusion. Therefore, we set an error threshold ξ on e to filter out the pixels that are likely corrupted. When $e < \xi$, we pass the color of $\mathbf{p}_{v \rightarrow t}$ to \mathbf{p}_v ; otherwise \mathbf{p}_v adopts the corresponding pixel color of the regularized renderings from the coarse stage.

View Prior Refinement: As a result of reprojection, each view prior is only a warped version of its corresponding training image and lacks visual information specific to its view. This may lead to unrealistic appearance at the extrapolated views. To incorporate view-dependent information into the view priors, we refine the view priors by blending them with regularized renderings from the coarse stage in a complementary way. We observe that the view priors can better retain high-frequency details, while the regularized renderings from the coarse stage exhibit more view-dependent coloring. Therefore, for each view prior $\bar{\mathcal{I}}$, we obtain the regularized rendering $\tilde{\mathcal{I}}$ from the coarse stage at the same augmented view, then blend the high-frequency component of $\bar{\mathcal{I}}$ with the low-frequency component of $\tilde{\mathcal{I}}$ through weighted summation in the frequency domain. We update $\bar{\mathcal{I}}$ to *Enhanced View Prior* as follows,

$$\bar{\mathcal{I}} \leftarrow \mathcal{F}^{-1}(\mathbf{M} \odot \mathcal{F}(\bar{\mathcal{I}}) + (1 - \mathbf{M}) \odot \mathcal{F}(\tilde{\mathcal{I}})) \quad (6)$$

where \mathbf{M} is the weighting map in the frequency domain, \mathcal{F} and \mathcal{F}^{-1} are the Fast Fourier Transform [8] and its inverse transform, respectively. The weight map \mathbf{M} assigns a large weight w_h to the center pixel, which corresponds to the high-frequency area, and a small weight w_l to the outer edges, which correspond to the low-frequency area. The weight at each coordinate on \mathbf{M} decreases linearly from w_h to w_l as the distance between the coordinate and the center of \mathbf{M} (i.e. the highest frequency) increases.

3.4. Training Strategies

Pre-train Stage: During the pre-train stage, we use the photometric loss employed in GS-based methods, which is a combination of L_1 and SSIM[33] loss, i.e.,

$$\mathcal{L}_{photo} = \lambda \mathcal{L}_1(\hat{\mathcal{I}}, \mathcal{I}) + (1 - \lambda) SSIM(\hat{\mathcal{I}}, \mathcal{I}) \quad (7)$$

where $\hat{\mathcal{I}}$ is the rendered image and \mathcal{I} is the ground truth image. Our EVPGS framework is generic and can be implemented with various GS-based methods as backbones. For each backbone, we also incorporate the specific strategies proposed in their respective papers.

Coarse Stage: We sample the augmented views by elevating the camera perspective at the training views for ω degrees. We formulate the coarse-stage loss function \mathcal{L}_{coarse} by summing the photometric loss in Eq. 7 and our AGR regularization loss terms in Eq. 2 and 3. During optimization, we update the GS model by taking gradients collectively on all loss terms in \mathcal{L}_{coarse} , i.e.,

$$\mathcal{L}_{coarse} = \mathcal{L}_{photo}^{train} + \lambda_a \mathcal{L}_a + \lambda_g \mathcal{L}_g \quad (8)$$

where λ_a and λ_g are the loss balancing weights, $\mathcal{L}_{photo}^{train}$ is the photometric loss applied to the training images. We apply \mathcal{L}_a and \mathcal{L}_g only to the augmented views.

Fine Stage: At the fine stage, we use *Enhanced View Priors* as pseudo-labels to fine-tune GS model on the augmented views. We employ the photometric loss $\mathcal{L}_{photo}^{aug}$ with *Enhanced View Prior* $\bar{\mathcal{I}}$ as the supervision signal at the augmented views, i.e.,

$$\mathcal{L}_{photo}^{aug} = \lambda \mathcal{L}_1(\hat{\mathcal{I}}, \bar{\mathcal{I}}) + (1 - \lambda) SSIM(\hat{\mathcal{I}}, \bar{\mathcal{I}}) \quad (9)$$

where $\hat{\mathcal{I}}$ is the directly rendered images from the GS model. We formulate the fine-stage loss \mathcal{L}_{fine} by incorporating the training view loss from Eq. 7 into Eq. 9, i.e.,

$$\mathcal{L}_{fine} = \mathcal{L}_{photo}^{train} + \mathcal{L}_{photo}^{aug} \quad (10)$$

4. Experiments

4.1. Experimental Setup

Datasets: We conduct experiments on three datasets with training/testing splits organized following the EVS scenario. These include the two public datasets DTU [14] and Synthetic-NeRF [20], and our custom-captured real-world dataset called Merchandise3D which closely reflects the EVS scenario. Merchandise3D includes seven common products in real life with a significant difference in view coverage between training views and testing views. More details of our dataset are in supplementary material.

- 1) **DTU[14]:** We select 14 scanned scenes in our EVS task. Each scan contains either 49 or 64 views. For each scene, we use the 21 or 31 views with the smallest pitch angles for training and the remaining for testing.
- 2) **Merchandise3D:** For each object, we use 50 images captured from horizontal viewpoints as training views




Figure 4. **Qualitative results on the real-world datasets DTU [14] (top two rows) and our Merchandise3D (bottom two rows).** On both datasets, the baseline methods (the first and the third rows) exhibit rendering artifacts and lack of details due to limited training view coverage, while all five variants of EVPGS can effectively deal with these issues.

	DTU			Merchandise3D			Synthetic-NeRF		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
3DGS[15]	25.188	0.8599	0.0865	23.946	0.9125	0.0618	24.479	0.8857	0.0804
Mip-Splatting[42]	25.443	0.8668	0.0851	24.276	0.9139	0.0607	27.507	0.9134	0.0517
2DGS[12]	25.589	0.8843	0.0762	22.181	0.9037	0.0746	27.058	0.9137	0.0573
GOF[43]	25.874	0.8633	0.0756	22.903	0.9169	0.0639	27.528	0.9170	0.0530
RaDe-GS[44]	25.778	0.8882	0.0753	23.559	0.9134	0.0645	27.531	0.9190	0.0530
VEGS*[13]	24.927	0.8674	0.0985	24.039	0.9115	0.0654	24.509	0.8935	0.0944
EVPGS(3DGS)	26.101	0.8784	0.0773	24.789	0.9204	0.0563	27.270	0.9173	0.0540
EVPGS(MS)	26.107	0.8777	0.0802	24.736	0.9207	0.0563	27.735	0.9206	0.0507
EVPGS(2DGS)	26.028	0.8936	0.0715	24.554	0.9194	0.0574	27.266	0.9183	0.0565
EVPGS(GOF)	26.343	0.8959	0.0702	23.565	0.9239	0.0577	27.666	0.9221	0.0526
EVPGS(RaDe-GS)	26.488	0.8991	0.0670	25.136	0.9267	0.0496	27.849	0.9243	0.0498

Table 1. **Quantitative evaluation of our EVPGS framework.** We compare the EVPGS variants based on different backbones with the baselines on the DTU [14], Merchandise3D, and Synthetic-NeRF [20] datasets. For each dataset, we report the average metrics across all scenes. Each EVPGS variant substantially improves performance over the corresponding baselines.

and the remaining 50 images from elevated viewpoints as testing views.

- 3) **Synthetic-NeRF[20]:** We sort the camera locations of the original training/testing split in ascending order along the z-axis to select the lowest 100 views as training views, leaving the remaining 200 views for testing.

We quantify the view coverage disparity between the

training and testing splits by calculating the pitch angle difference between each test view and its closest training view, averaged across all testing views for each dataset, as shown in Table 2.

Baselines and metrics: We use six GS-based methods as baselines, i.e. 3DGS [15], Mip-Splatting [42], 2DGS [12], GOF [43], RaDe-GS [44], and VEGS [13]. These baselines




Figure 5. **Qualitative results on the synthetic dataset Synthetic-NeRF [20].** When using different backbones, our EVPGS produces synthesis closer to the ground truth than the baselines at extrapolated views.

Dataset	Average Training/Testing Pitch Angle Difference (Degree)	
	Original Split	Our EVS-specific split
DTU [14]	11.576	24.497
Synthetic-NeRF [20]	8.463	26.491
Merchandise3D	-	39.580

Table 2. **Average pitch angle difference between the original training/testing split and our EVS-specific split.** We measure the difference for the three datasets DTU, Synthetic-NeRF, and Merchandise3D. Our EVS splits exhibit a larger view disparity on each dataset, posing a more challenging scenario.

are trained only on the training views for each scene in each dataset. To evaluate the flexibility of our EVPGS framework, we integrate it with each GS-based backbone except VEGS, as VEGS is specifically designed to tackle the EVS problem. Based on the results, we select RaDe-GS as the primary backbone for EVPGS, as it achieves the best performance in our experiments. In the following figures and tables, we use *EVPGS(Backbone)* to denote each backbone in the corresponding experiment. Following the baseline methods, we adopt PSNR, SSIM [33], and LPIPS [46] as quantitative metrics.

Implementation Details: We first pre-train the backbone model for 20k iterations. Then at the coarse stage, we employ the Stable-Diffusion-2.1 model [25] trained on billion-scale real-world images to provide strong knowledge priors for our AGR strategy. We train for 1k iterations in this stage. At the fine stage, we set the reverse-reprojection error threshold $\xi = 1$ and train for 9k iterations. We will explain more implementation details in supplementary.

4.2. Comparison to baselines

We compare our EVPGS framework with 3DGS [15] and four of its variants, i.e. Mip-Splatting [42], 2DGS [12], GOF [43] and RaDe-GS [44]. To further evaluate the generality of EVPGS for the EVS task, we adopt these five GS methods as alternative backbones. Additionally, we compare EVPGS with VEGS [13], a GS-based method specifically designed for the EVS problem in urban scenes. To

	DTU			Synthetic-NeRF		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
<i>baseline</i>	25.778	0.8882	0.0753	27.531	0.9190	0.0530
<i>only coarse</i>	26.219	0.8927	0.0699	27.683	0.9222	0.0507
<i>only fine</i>	26.165	0.8978	0.0703	27.708	0.9224	0.0523
<i>full w/o occ.</i>	26.027	0.8791	0.0714	27.606	0.9217	0.0519
<i>full w/o ref.</i>	26.155	0.8951	0.0714	27.457	0.9221	0.0514
<i>full w/o mesh</i>	26.337	0.8956	0.0680	27.829	0.9237	0.0501
<i>full</i>	26.488	0.8991	0.0670	27.849	0.9243	0.0498

Table 3. **Ablation study on DTU [14] and Synthetic-NeRF [20].** We use RaDe-GS [44] as the baseline.

adapt VEGS to our scenario, we remove the LiDAR initialization and object detection, referring to this variant as VEGS* in the following.

Table 1 shows quantitative results on DTU [14], Merchandise3D, and Synthetic-NeRF [20]. EVPGS(RaDe-GS) consistently outperforms the baseline methods across all metrics, including PSNR, SSIM, and LPIPS, on all three datasets. Integrating the five baseline models into our EVPGS framework brings substantial performance improvements over the corresponding baselines.

Figure 4 shows the qualitative results on the real-world datasets DTU and Merchandise3D. In comparison with the baselines, our EVPGS framework can synthesize extrapolated images with fewer and milder artifacts, while preserving high-frequency details (e.g. text on the merchandise and tiles on the roof). These results highlight the ability of EVPGS to recover fine-grained details and produce realistic extrapolated views. Moreover, we show in Figure 5 our qualitative improvements on the synthetic dataset Synthetic-NeRF. The results on Synthetic-NeRF also demonstrate the efficacy of EVPGS.

4.3. Ablation Study

Ablation on both coarse and fine stages: We investigate the impact of each proposed strategies on EVS performance. Following the results in Table 3, we use RaDe-GS [44] as the backbone in our ablation study, since EVPGS(RaDe-GS) brings the best performance. We eval-




Figure 6. **Qualitative results from the ablation studies.** The first row demonstrates the effectiveness of the coarse stage in reducing prominent artifacts. The second row illustrates the impact of the view prior refinement strategy, which effectively combines view-dependent color information from the coarse stage while preserving fine-grained details.

uate the effectiveness of both coarse and fine stages in our framework, as well as the two strategies of the fine stage, i.e. occlusion-aware reprojection and view prior refinement. For the coarse and fine stages, we conduct two separate experiments where each experiment has only one of the stages applied. We denote these two simplified versions of EVPGS as *only coarse* and *only fine*, respectively. For the two strategies of the fine stage, we conduct another two experiments by removing either strategy from the complete EVPGS framework, resulting in two more simplified versions of EVPGS which we denote as *full w/o occ.* and *full w/o ref.*, respectively. As shown in Table 3, all of the proposed components contribute significantly to EVPGS.

We also provide qualitative comparison in Figure 6 to demonstrate the effectiveness of both the coarse stage and the view prior refinement strategy. The coarse stage removes severe artifacts, therefore benefiting the subsequent fine stage. View prior refinement helps preserve more view-specific visual information in the final rendered images.

Ablation on initial reconstructed mesh: We also investigate how the usage of the reconstructed mesh benefits both the coarse and fine stages. In our pipeline, we reconstruct an initial mesh after pre-training the GS backbone to enhance depth map quality, addressing the cases where GS-rendered depth maps are overall relatively accurate but incomplete in some augmented views. To account for scenarios where the initial reconstructed mesh is inaccurate or unavailable (e.g., complex objects or outdoor scenes), we conduct experiments on a variant of our EVPGS which does not rely on the initial mesh. Specifically, we remove the geometry regularization in AGR module (Sec. 3.2) and rely the GS-rendered depth map instead of the mesh-rendered depth map for occlusion detection (Sec. 3.3). We denote this version of EVPGS as *full w/o mesh*. As shown in Table 3, *full w/o mesh* has moderate performance drop compared to the full EVPGS, while still outperforming the baseline as well as the other simplified variants. Incorporating the reconstructed mesh in both the coarse and fine stages contribute to improved performance.

5. Limitations

There are two limitations to address in future work. While we handle view-dependent color at augmented views using the OARR strategy (Sec. 3.3), our method does not account for the part of view dependence caused by complex lighting effects as we do not explicitly model lighting. Despite the high-fidelity renderings at extrapolated views produced without considering lighting, there remains a noticeable color discrepancy compared to the ground truth, which can affect the overall visual experience. The other limitation of our method is its reliance on the quality of the reconstructed mesh. While our ablation study in Sec. 4.3 shows that EVPGS still performs competitively without an initial mesh, we aim to explore mesh-free approaches in future work that are better applicable to scenarios where meshes are inaccessible.

6. Conclusion

In this work, we proposed the view augmentation framework called EVPGS based on Gaussian Splatting (GS) to tackle the challenging Extrapolated View Synthesis (EVS) problem. EVPGS generates enhanced view priors to fine-tune GS models through a coarse-to-fine process. At the coarse stage, we proposed the Appearance and Geometry Regularization (AGR) strategy to effectively reduce severe rendering artifacts. At the fine stage, we proposed the Occlusion-Aware Reprojection and Refinement (OARR) strategy to generate enhanced view priors, which serves as strong appearance guidance for synthesis at extrapolated views. To comprehensively evaluate our EVPGS framework, we conducted experiments on two public real-world and synthetic datasets, i.e. DTU, Synthetic-NeRF, and our self-collected real-world dataset Merchandise3D. Both quantitative and qualitative results on all three datasets demonstrate that EVPGS significantly improves the synthesis quality at extrapolated views compared to the baseline methods. Our EVPGS framework is generic and achieves consistent improvements across different backbones.

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5855–5864, 2021. 2
- [2] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction, 2024. 2
- [3] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields, 2022. 2
- [4] Yuedong Chen, Haoifei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *ECCV*, 2024. 2
- [5] Inchang Choi, Orazio Gallo, Alejandro Troccoli, Min H Kim, and Jan Kautz. Extreme view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7781–7790, 2019. 3
- [6] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 2
- [7] Sara Fridovich-Keil, Giacomo Meanti, Frederik Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance, 2023. 2
- [8] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing (3rd Edition)*. Prentice Hall, 2008. 5
- [9] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5354–5363, 2024. 2
- [10] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *ICCV*, 2021. 2
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. 2, 3
- [12] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 2, 6, 7
- [13] Sungwon Hwang, Min-Jung Kim, Taewoong Kang, Jayeon Kang, and Jaegul Choo. Vegs: View extrapolation of urban scenes in 3d gaussian splatting using learned priors, 2024. 3, 4, 6, 7
- [14] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014. 1, 2, 5, 6, 7
- [15] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering, 2023. 1, 2, 6, 7
- [16] Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and George Drettakis. A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM Transactions on Graphics (TOG)*, 43(4):1–15, 2024. 1, 2
- [17] Jiahe Li, Jiawei Zhang, Xiao Bai, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20775–20785, 2024. 2
- [18] Xinhang Liu, Jiaben Chen, Shiu-Hong Kao, Yu-Wing Tai, and Chi-Keung Tang. Deceptive-nerf/3dgs: Diffusion-generated pseudo-observations for high-quality sparse-view reconstruction. In *ECCV*, 2024. 1, 2, 3
- [19] Saswat Subhajyoti Mallick, Rahul Goel, Bernhard Kerbl, Francisco Vicente Carrasco, Markus Steinberger, and Fernando De La Torre. Taming 3dgs: High-quality radiance fields with limited resources. *arXiv preprint arXiv:2406.15643*, 2024. 1, 2
- [20] 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, 2020. 2, 5, 6, 7
- [21] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. 2
- [22] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs, 2021. 2
- [23] Avinash Paliwal, Wei Ye, Jinhui Xiong, Dmytro Kotovenko, Rakesh Ranjan, Vikas Chandra, and Nima Khademi Kalantari. Coherentgs: Sparse novel view synthesis with coherent 3d gaussians. In *European Conference on Computer Vision*, pages 19–37. Springer, 2025. 2
- [24] Lukas Radl, Michael Steiner, Mathias Parger, Alexander Weinrauch, Bernhard Kerbl, and Markus Steinberger. Stopthepop: Sorted gaussian splatting for view-consistent real-time rendering. *ACM Transactions on Graphics (TOG)*, 43(4):1–17, 2024. 1, 2
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022. 2, 3, 4, 7
- [26] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *CVPR*, 2019. 3
- [27] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 3
- [28] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011. 3
- [29] Guangcong Wang, Zhaoxi Chen, Chen Change Loy, and Ziwei Liu. Sparsenerf: Distilling depth ranking for few-shot novel view synthesis, 2023. 2
- [30] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit

- surfaces by volume rendering for multi-view reconstruction. In *NIPS*, 2021. [2](#)
- [31] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [2](#)
- [32] Yunsong Wang, Tianxin Huang, Hanlin Chen, and Gim Hee Lee. Freesplat: Generalizable 3d gaussian splatting towards free-view synthesis of indoor scenes. *arXiv preprint arXiv:2405.17958*, 2024. [2](#)
- [33] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. [5](#), [7](#)
- [34] Tong Wu, Jiaqi Wang, Xingang Pan, Xudong Xu, Christian Theobalt, Ziwei Liu, and Dahua Lin. Voxurf: Voxel-based efficient and accurate neural surface reconstruction. In *International Conference on Learning Representations (ICLR)*, 2023. [2](#)
- [35] Jamie Wynn and Daniyar Turmukhambetov. Diffusernerf: Regularizing neural radiance fields with denoising diffusion models, 2023. [3](#), [4](#)
- [36] Wangze Xu, Huachen Gao, Shihe Shen, Rui Peng, Jianbo Jiao, and Ronggang Wang. Mvpgs: Excavating multi-view priors for gaussian splatting from sparse input views, 2024. [1](#), [2](#), [3](#)
- [37] Chen Yang, Peihao Li, Zanwei Zhou, Shanxin Yuan, Bingbing Liu, Xiaokang Yang, Weichao Qiu, and Wei Shen. Nerfvs: Neural radiance fields for free view synthesis via geometry scaffolds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16549–16558, 2023. [3](#)
- [38] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization, 2023. [2](#)
- [39] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. [2](#)
- [40] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. [2](#)
- [41] Hanyang Yu, Xiaoxiao Long, and Ping Tan. Lm-gaussian: Boost sparse-view 3d gaussian splatting with large model priors, 2024. [1](#), [2](#), [3](#)
- [42] 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. [1](#), [2](#), [6](#), [7](#)
- [43] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Transactions on Graphics*, 2024. [2](#), [6](#), [7](#)
- [44] Baowen Zhang, Chuan Fang, Rakesh Shrestha, Yixun Liang, Xiaoxiao Long, and Ping Tan. Rade-gs: Rasterizing depth in gaussian splatting, 2024. [2](#), [6](#), [7](#)
- [45] Jian Zhang, Yuanqing Zhang, Huan Fu, Xiaowei Zhou, Bowen Cai, Jinchi Huang, Rongfei Jia, Binqing Zhao, and Xing Tang. Ray priors through reprojection: Improving neural radiance fields for novel view extrapolation. In *CVPR*, pages 18376–18386, 2022. [3](#)
- [46] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [7](#)
- [47] Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. Fsgs: Real-time few-shot view synthesis using gaussian splatting. In *European Conference on Computer Vision*, pages 145–163. Springer, 2025. [2](#), [3](#)
- [48] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa volume splatting. In *Visualization, 2001. VIS 01. Proceedings*, pages 29–538. IEEE, 2001. [2](#)