GS-LRM: Large Reconstruction Model for 3D Gaussian Splatting

Kai Zhang^{*1} Sai Bi^{*1} Hao Tan^{*1} Yuanbo Xiangli² Nanxuan Zhao¹ Kalyan Sunkavalli¹ Zexiang Xu¹

¹Adobe Research ²Cornell University

Abstract. We propose GS-LRM, a scalable large reconstruction model that can predict high-quality 3D Gaussian primitives from 2-4 posed sparse images in ~0.23 seconds on single A100 GPU. Our model features a very simple transformer-based architecture; we patchify input posed images, pass the concatenated multi-view image tokens through a sequence of transformer blocks, and decode final per-pixel Gaussian parameters directly from these tokens for differentiable rendering. In contrast to previous LRMs that can only reconstruct objects, by predicting per-pixel Gaussians, GS-LRM naturally handles scenes with large variations in scale and complexity. We show that our model can work on both object and scene captures by training it on Objaverse and RealEstate10K respectively. In both scenarios, the models outperform state-of-the-art baselines by a wide margin. We also demonstrate applications of our model in downstream 3D generation tasks. Our project webpage is available at: https://sai-bi.github.io/project/gs-lrm/.

Keywords: Large Reconstruction Models \cdot 3D Reconstruction \cdot Gaussian Splatting

1 Introduction

Reconstructing a 3D scene from image captures is both a central problem and a long-standing challenge in computer vision. Traditionally, high-quality 3D reconstruction relies on complex photogrammetry systems [23, 48, 50] and requires a dense set of multi-view images. Recent advancements in neural representations and differentiable rendering [9, 30, 40, 41] have shown superior reconstruction and rendering quality, by optimizing renderings on a per-scene basis. However, these methods are slow and still require a large number of input views. Recently, transformer-based 3D large reconstruction models (LRMs) have been proposed, learning general 3D reconstruction priors from vast collections of 3D objects and achieving sparse-view 3D reconstruction of unprecedented quality [27, 32, 61, 66]. However, these models adopt triplane NeRF [7, 43] as the scene representation, which suffers from a limited triplane resolution and expensive volume rendering. This leads to challenges in training and rendering speeds, preserving fine details, and scaling to large scenes beyond object-centric inputs.

^{*} Equal contribution.



Fig. 1: Novel-view renderings of our predicted Gaussians from object captures (top left), text-conditioned generated object images (top right), scene captures (bottom left) and text-conditioned generated scene images (bottom right, from Sora [42] with the prompt "*Tour of an art gallery with many beautiful works of art in different styles*"). The rendered depth for scenes is at the bottom. By predicting 3D Gaussians with our transformer-based *GS-LRM*, we can naturally handle objects and complex scenes.

Our goal is to build a general, scalable, and efficient 3D reconstruction model. To this end, we propose GS-LRM, a novel transformer-based large reconstruction model that predicts 3D Gaussian primitives [30] from sparse input images, enabling fast and high-quality rendering and reconstruction for both objects and scenes, as depicted in Fig. 1. The core of our approach is a simple and scalable transformer-based network architecture that predicts per-pixel Gaussians. Specifically, we patchify input posed images into patch tokens and process them through a series of transformer blocks comprising self-attention and MLP layers, and directly regress per-view per-pixel 3D Gaussian primitives from the contextualized multi-view tokens. Unlike previous LRMs that require careful designs of additional (triplane) NeRF tokens for reconstruction, we align input (2D images) and output (3D Gaussians) in the same pixel space, predicting one Gaussian per pixel along the ray. This alignment not only simplifies the transformer architecture but also facilitates 3D Gaussians to preserve the high-frequency details in the input images. Moreover, predicting per-pixel Gaussians allows our model to freely adapt to the input image resolution, exhibiting accurate scene details in high-resolution inputs, that previous LRMs with a fixed triplane resolution often struggle with.

Our transformer-based GS-LRM is highly scalable from multiple aspects, including model sizes, training data, and scene scales. We train two versions of our GS-LRM on two large-scale datasets including Objaverse [19] and Real-Estate10K [74], separately for object and scene reconstruction tasks, using the same transformer architecture with minimal domain-specific parameter changes. The results demonstrate that our GS-LRM (with 300M model parameters, up to 16K transformer token length) achieves high-quality sparse-view reconstruction for both object and scene scenarios. We also achieve state-of-the-art reconstruction quality and outperform previous methods by a large margin of 4dB PSNR for objects and 2.2dB PSNR for scenes.

2 Related Work

Multi-view 3D Reconstruction. 3D reconstruction has been extensively studied in computer vision and graphics for decades. To address this task, various traditional methods have been proposed, including structure-from-motion (SfM) [1, 47, 48, 54] for sparse reconstruction and calibration, and multi-view stereo (MVS) [23, 46, 50] for dense reconstruction. Recently, deep learning-based MVS methods have also been proposed [16, 25, 51, 67, 68], offering efficient highquality reconstruction in a feed-forward manner. Generally, these methods utilize 3D cost volumes — constructed by unprojecting 2D image features into plane sweeps — to achieve high-quality per-view depth estimation. In this work, we estimate pixel-aligned 3D Gaussian primitives, essentially achieving a per-pixel estimation of depth along with additional Gaussian properties. Instead of relying on 3D cost volumes, we adopt a multi-view transformer to directly regress Gaussians, using self-attention over all multi-view image patches, naturally allowing for multi-view correspondence reasoning. Our transformer model can effectively handle highly sparse input views, a big challenge for cost volume-based methods.

Radiance Field Reconstruction. Recently, a vast number of works have emerged to address scene reconstruction by optimizing radiance field representations [40] with differentiable rendering, bypassing traditional MVS pipelines. While NeRF [40] initially represents a radiance field as a coordinate-based MLP, this area has expanded with various models, including voxel-based [34, 57, 70], factorization-based [7, 9–11], hash grid-based [41], and point- (or primitive-) based [24, 30, 35, 65] representations among others [4–6, 72]. We leverage Gaussian Splatting [30], a state-of-the-art technique for radiance field modeling and rendering, allowing us to achieve real-time rendering and large-scale scene reconstruction. In contrast to these radiance field models optimized on a per-scene basis given dense images, we propose a generalizable model trained across scenes for feed-forward sparse-view reconstruction.

Feed-forward Reconstruction. Prior works have proposed various feed-forward 3D reconstruction and rendering methods. Early efforts adopt CNNs to estimate neural points [2,64,69] or multi-plane images (MPIs) [33,38,74], achieving

rendering via point splatting or alpha compositing. We utilize Gaussian Splatting, which can be seen as a generalization of point splatting and MPI compositing. Recently, generalizable radiance field-based methods have been proposed and achieved state-of-the-art quality [12, 56, 62, 70], leveraging NeRF-like volume rendering. These methods typically leverage 3D-to-2D geometric projection to sample per-view 2D features for 3D reconstruction, using model designs like epipolar lines [55, 56, 62] or plane-swept cost-volumes [12, 29, 36] similar to MVS. We instead leverage a clean large transformer model without such 3D inductive bias, and directly regress per-view pixel-aligned 3D Gaussian primitives. The dense self-attention layers in our model can effectively learn multi-view correspondence and general reconstruction priors, leading to significantly better rendering quality than previous epipolar-based methods (see Tab. 1).

Our model is inspired by the recent transformer-based 3D large reconstruction models (LRMs) [27, 32, 61, 66], which are based on triplane NeRF and focus on objects. We propose a novel LRM for Gaussian Splatting with a simplified network architecture and per-pixel Gaussian prediction mechanism, achieving better quality, faster rendering, and scaling up to handle large scenes.

PixelSplat [8] and LGM [59] are two concurrent works that are also based on pixel-aligned 3D Gaussian prediction. In particular, LGM leverages a U-Net architecture and only focuses on object generation; PixelSplat leverages epipolar line-based sampling and only tackles scene-level reconstruction. In contrast, our GS-LRM is a clean transformer model that is much simpler to implement and scale. We demonstrate that our model significantly outperform these two concurrent works in terms of both object and scene-level reconstructions.

3 Method

In this section, we present the technical details of our method, including the architecture of our transformer-based model (Sec. 3.1) and the loss functions (Sec. 3.2).

3.1 Transformer-based Model Architecture

As shown in Fig. 2, we train a transformer model to regress per-pixel 3D GS parameters from a set of images with known camera poses. We tokenize posed input images via a patchify operator [20]. Multi-view image tokens are then concatenated and passed through a sequence of transformer blocks consisting of self-attention and MLP layers. From each output token, we decode the attributes of pixel-aligned Gaussians in the corresponding patch with a linear layer.

Tokenizing posed images. The inputs to our model are N multi-view images $\{\mathbf{I}_i \in \mathbb{R}^{H \times W \times 3} | i = 1, 2, ..., N\}$ and their camera intrinsic and extrinsic parameters; here H and W are the height and width of the images. Following prior works [13, 66], we use the Plücker ray coordinates [44] of each image $\{\mathbf{P}_i \in \mathbb{R}^{H \times W \times 6}\}$ computed from the camera parameters for pose conditioning. Specifically, we concatenate the image RGBs and their Plücker coordinates



Fig. 2: Our simple transformer-based *GS-LRM* predicts 3D Gaussian parameters from sparse posed images. Images are patchified and the concatenated patch tokens are sent to the transformer blocks. By unpatchifying the transformer's output, each pixel is unprojected to a 3D Gaussian. The final output merges all 3D Gaussians. (Note that here we visualize the Gaussian centers and colors as point clouds for illustration; please refer to in Fig. 1 for the splatting-based renderings.)

channel-wise, enabling per-pixel pose conditioning and forming a per-view feature map with 9 channels. Similar to ViT [20], we patchify the inputs by dividing the per-view feature map into non-overlapping patches with a patch size of p. For each 2D patch, we flatten it into a 1D vector with a length of $p^2 \cdot 9$. Finally, we adopt a linear layer that maps the 1D vectors to image patch tokens of d dimensions, where d is the transformer width. Formally, this process can be written as:

$$\{\mathbf{T}_{ij}\}_{j=1,2,\dots,HW/p^2} = \text{Linear}(\text{Patchify}_p(\text{Concat}(\mathbf{I}_i, \mathbf{P}_i))), \tag{1}$$

where $\{\mathbf{T}_{ij} \in \mathbb{R}^d\}$ denotes the set of patch tokens for image *i*, and there are a total number of HW/p^2 such tokens (indexed by *j*) for each image. As Plücker coordinates vary across pixels and views, they naturally serve as spatial embeddings to distinguish different patches; hence we do not use additional positional embedding as in [20] or view embeddings [61].

Processing image tokens with transformer. Given the set of multi-view image tokens $\{\mathbf{T}_{ij}\}$, we concatenate and feed them through a chain of transformer blocks [60]:

$$\{\mathbf{T}_{ij}\}^0 = \{\mathbf{T}_{ij}\},\tag{2}$$

$$\{\mathbf{T}_{ij}\}^{l} = \operatorname{TransformerBlock}^{l}(\{\mathbf{T}_{ij}\}^{l-1}), l = 1, 2, \dots, L,$$
(3)

where L is the total number of transformer blocks. Each transformer block is equipped with residual connections [26] (i.e., the operator '+' in Fig. 2) and consists of Pre-LayerNorm [3], multi-head Self-Attention [60] and MLP.

Decoding output tokens to per-pixel Gaussians. With the output tokens $\{\mathbf{T}_{ij}\}^L$ from the transformer, we decode them into Gaussian parameters using a single linear layer:

$$\{\mathbf{G}_{ij}\} = \operatorname{Linear}(\{\mathbf{T}_{ij}\}^L),\tag{4}$$

6 K. Zhang, S. Bi, H. Tan et al.

where $\mathbf{G}_{ij} \in \mathbb{R}^{p^2 \cdot q}$ represents the 3D Gaussian and q is the number of parameters per Gaussian. We then unpatchify \mathbf{G}_{ij} into p^2 Gaussians. Since we use the same patch size p for patchifying and unpatchifying operations, we end up with HWGaussians for each view where each 2D pixel corresponds to one 3D Gaussian.

Similar to [30], the 3D Gaussian is parameterized by 3-channel RGB, 3channel scale, 4-channel rotation quaternion, 1-channel opacity, and 1-channel ray distance. Thus q=12 in our formulation. For splatting rendering, the location of the Gaussian center is obtained by the ray distance and the known camera parameters. Specifically, suppose that t, ray_o, ray_d are the ray distance, ray origin, and ray direction, respectively, the Gaussian center $xyz = ray_o + t \cdot ray_d$. We will discuss the detailed initialization of these parameters in the Appendix.

The final output of our model is simply the merge of 3D Gaussians from all N input views. Thus the model will output $N \cdot HW$ Gaussians in total. It's worth noting that the number of Gaussians scale up with increased input resolution, which is in contrast to the fixed-resolution triplane used in prior LRM works [27, 32, 61, 66]. This property allows us to better handle high-frequency details in the inputs and large-scale scene captures.

3.2 Loss Functions

During training, we render the images at the M supervision views using the predicted Gaussian splats, and minimize the image reconstruction loss. Let $\{\mathbf{I}_{i'}^*|i'=1,2,...,M\}$ be the set of groundtruth views, and $\{\hat{\mathbf{I}}_{i'}^*\}$ be the rendered images, our loss function is a combination of MSE (Mean Squared Error) loss and Perceptual loss:

$$\mathcal{L} = \frac{1}{M} \sum_{i'=1}^{M} \left(\text{MSE}\left(\hat{\mathbf{I}}_{i'}^{*}, \mathbf{I}_{i'}^{*} \right) + \lambda \cdot \text{Perceptual}\left(\hat{\mathbf{I}}_{i'}^{*}, \mathbf{I}_{i'}^{*} \right) \right),$$
(5)

where λ is the weight of the perceptual loss. We empirically find that the perceptual loss in [14] based on VGG-19 network [53] provides a more stable training than LPIPS [73] used in [27, 32, 61, 66], and we use it in this work.

4 Experiments

In this section, we first describe the training and testing datasets (Sec. 4.1), then introduce the implementation and training details (Sec. 4.2). We make both quantitative and qualitative comparisons (Sec. 4.3) against different baselines. Finally we show some downstream applications (Sec. 4.5) of our method. We refer the readers to our project page for video results and interactive visualizations.

4.1 Datasets

Table 1: Comparison against baselines on object-level (left) and scene-level (right) reconstructions. We matched the baseline settings by comparing with Instant3D's Triplane-LRM [32] and LGM [59] at 512 and 256 resolution respectively for both input and rendering. All scene-level results were performed at 256 resolution for fair comparisons. We outperform relevant baselines by a large margin in both scenarios.

								RealEstate10k		
		GSO	I DIDG I	DOM: N	ABO	I DIDG I		$\mathrm{PSNR}\uparrow$	\uparrow SSIM \uparrow	LPIPS ↓
1	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS \downarrow	nivelNeBF [70]	20.43	0 589	0.550
Triplane-LRM [32]	26.54	0.893	0.064	27.50	0.896	0.093	GPNR [55]	24.11	0.793	0.255
Ours (Res-512)	30.32	0.952	0.030	29.09	0.925	0.085	Du et. al [22]	24.78	0.820	0.213
LGM [59]	21.44	0.832	0.122	20.79	0.813	0.158 0.074	pixelSplat [8]	25.89	0.858	0.142
Ours (nes-250)	20.00	0.344	0.001	20.00	0.320		Ours	28.10	0.892	0.114

Object-level dataset. We use the Objaverse dataset [19] to train our objectlevel reconstruction model. We only leverage the multi-view renderings of the objects without accessing explicit 3D information (such as depths). Following [27], we center and scale each 3D object to a bounding box of $[-1, 1]^3$, and render 32 views randomly placed around the object with a random distance in the range of [1.5, 2.8]. Each image is rendered at a resolution of 512×512 under uniform lighting. We render a total of 730K objects. We evaluate our model on two 3D object datasets including the full Google Scanned Objects (GSO) [21] that contains 1009 objects and the Amazon Berkeley Objects (ABO) [17] dataset from which we sample 1000 objects. We follow Instant3D [32] and render 4 structured input views evenly placed around the objects with an elevation of 20° to ensure a good coverage, and randomly select another 10 views for testing.

Scene-level dataset. We use the RealEstate10K [74] dataset to train our scene-level model. It sources from real estate video footages and has both indoor and outdoor scenes. The dataset contains 80K video clips curated from 10K YouTube videos. Each clip has a set of extracted frames with known camera poses estimated by SfM [49]. We follow the standard training/testing split for the dataset, which is also used in pixelSplat [8].

4.2 Implementation Details

We have two models trained independently in this paper: object-level GS-LRM and scene-level GS-LRM. The two models share the same model architecture and take almost the same training recipe. The differences are in the training data (Sec. 4.1) and view selection and normalization (see details below). We also made necessary changes for fair comparisons with baseline methods (Sec. 4.3).

Model details. We use a patch size of 8×8 for the image tokenizer. Our transformer consists of 24 layers, and the hidden dimension of the transformer is 1024. Each transformer block consists of a multi-head self-attention layer with 16 heads, and a two-layered MLP with GeLU activation. The hidden dimension of the MLP is 4096. Both layers are equipped with Pre-Layer Normalization

8



Fig. 3: Visual comparisons to Instant3D's Triplane-LRM [32]. The 4-view input images are shown in the leftmost column, and we compare novel view renderings on the right. The Triplane-LRM cannot reconstruct high-frequency details (top left and top right) and thin structures (bottom left) well. It also suffers from texture distortion artifacts (bottom right), possibly due to a lack of representation capability and the non-pixel-aligned prediction of triplanes. In contrast, our *GS-LRM* works significantly better in these cases. PSNRs are shown under each image.

(LN) and residual connections. Besides the above Pre-LNs, Layer Normalization is used after the patchifying Linear layer and before the unpatchifying Linear layer to stabilize the training.

Training details. To enable efficient training and inference, we adopt Flash-Attention-v2 [18] in the xFormers [31] library, gradient checkpointing [15], and mixed-precision training [37] with BF16 data type. We also apply deferred backpropagation [71] for rendering the GS to save GPU memory. We pre-train the model with a resolution of 256×256 and fine-tune the trained model with a resolution of 512×512 for a few epochs. The fine-tuning shares the same model architecture and initializes the model with the pre-trained weights, but processes more tokens than the pre-training. At each training step, for object-level, we sample a set of 8 images (from 32 renderings) as a data point, and from which we randomly select 4 input views and 4 supervision views independently. This sampling strategy encourages more overlap between input views and rendering views than directly sampling from 32 rendering views, which helps the model's convergence. For scene-level, we adopt two input views for a fair comparison with pixelSplat [8]. Following pixelSplat [8], we select 2 random input views and then randomly sample supervision views between them; we use 6 supervision views for each batch. We normalize the camera poses for scene-level input images following common practices in previous forward-facing reconstructions as done in [9, 39]. We further fine-tune a model that takes 2-4 input images of 512×512 for generating visual results. For both models, we use 64 A100 (40G VRAM) GPUs to train our models; 256-res pretraining takes 2 days, while 512-res finetuning takes 1 additional day. For more details, please refer to the Appendix.



Fig. 4: Visual comparisons to LGM [59]. The LGM renderings have obvious distorted textures (top) and broken geometries (bottom) and are inferior in recovering accurate surface opacity (top left; bottom left; bottom right). Our *GS-LRM* renderings recover the high-frequency details and are visually much closer to the ground truth. PSNRs are shown under each image.

4.3 Evaluation against Baselines

Object-level. We compare our object-level GS-LRM with the Triplane-LRM in Instant3D [32]. We outperform this baseline by a large margin across all view synthesis metrics; for example, as shown in Tab. 1, we improve the PSNR for novel-view rendering by 3.98dB on GSO data, and by 1.59dB on ABO data. This is also reflected by our much sharper renderings in Fig. 3; our GS-LRM manages to faithfully reproduce the high-frequency details, e.g., texts, in the input images, while Triplane-LRM tends to blur out the details. We attribute this to our pixel-aligned Gaussian prediction scheme which creates a shortcut for learning accurate per-Gaussian colors from input RGB images; this is in contrast to the non-pixel-aligned prediction of triplanes in Instant3D's LRM where the relationship between input pixel colors and triplane features is less straightforward, and might be more challenging to learn for the network. Another advantage of our GS-LRM is that our predicted Gaussians are much faster to render than the predicted NeRF from Triplane-LRM, making it easier to deploy in downstream applications. We also tried to compare against another baseline SparseNeuS [36]; however, we found that it failed to produce plausible reconstructions given 4 highly sparse inputs; this was also observed in the prior Instant3D work (they had to use 8 views to run SparseNeuS, which is not a fair comparison).

In addition, we compare with the concurrent work LGM [59] which is also based on Gaussian Splatting [30]. The official LGM is trained with a special setting using 256×256 resolution input and 512×512 resolution output supervision. Since their model only accepts 256×256 input, we compare with LGM using our low-res model, trained with 256×256 images only from our 256-res pre-training stage. We evaluate both models with 256×256 renderings for comparison. As



Fig. 5: We compare scene-level *GS-LRM* with the best-performing baseline pixel-Splat [8]. We can observe that our model is better in sharpness (leftmost column), has fewer floaters (mid-right and rightmost), and is more faithful to the original scenes (mid-left). Our superiority in visual quality aligns with the significant quantitative metric improvement in Tab. 1. PSNRs are shown at the corner of each image.

seen in Tab. 1, our approach significantly outperforms LGM, achieving a notable 8dB higher PSNR on both GSO and ABO testing data. The improvement can be visualized in Fig. 4. It's worth noting that this is an almost equal-compute comparison: LGM is trained on 32 A100 (80G VRAM) for 4 days, while our low-res base model is trained using 64 A100 (40G VRAM) for 2 days. This further highlights the method-wise advantage of our GS-LRM — a transformer model predicting per-pixel Gaussians that scales up easily with data and compute.

Scene-level. We compare our scene-level GS-LRM against previous generalizable neural rendering techniques [22, 55, 70] and the state-of-the-art GS-based concurrent work pixelSplat [8]. Since pixelSplat and GPNR are trained at 256×256 image resolution, we use our low-res model in these comparisons, and follow exactly the same evaluation setting as in pixelSplat. We directly take the reported quantitative results for all baselines from pixelSplat [8]. As shown in Tab. 1, our approach achieves the best quantitative results on the RealEstate10k

testing set, substantially surpassing all baselines for all metrics. In particular, when compared to pixelSplat, the top-performing baseline, our model leads to significant improvements of 2.2db in PSNR, 0.034 in SSIM, and 0.028 in LPIPS. These metric improvements align with the visual comparisons in Fig. 5, where our results are sharper and have fewer floaters. Note that pixelSplat and other baseline methods all leverage more complex model designs such as epipolar line-based feature aggregation. In contrast, we utilize a straightforward transformer model with self-attention layers. These self-attention layers effectively learn to aggregate context information from all the intra-view and inter-view pixels (as opposed to a subset of them on epipolar lines) for accurate per-pixel Gaussian prediction, when trained on large amount of data.

4.4 High-resolution Qualitative Results

We showcase some high-res reconstructions of our *GS-LRM* in Fig. 6. For the object captures, texts on the top-left product box remain readable in our rendered novel views, even when the inputs are captured from oblique viewpoints; we also manage to reconstruct the challenging thin structures and transparent glasses in the top-right table example. For the scene captures, we are able to handle large outdoor depth variations and faithfully reproduce complex structures, e.g., trees, in the presented examples. Please refer to our project page for videos and interactive rendering results.

4.5 Applications in 3D Generation

Following the Instant3D [32] work, we can also chain our GS-LRM with a textconditioned or image-conditioned multi-view generator to achieve text-to-3D or image-to-3D. We qualitatively show some results to demonstrate such a workflow for applying our models in this downstream 3D generation task.

Text/Image-to-object. For the text-to-3D application, we use the finetuned SDXL [45] model in Instant3D as the text-to-multi-views generator. Since the above generator generates four structured views with known camera parameters, we directly feed these posed images into our object-level GS-LRM to get 3D GS instantly. The results are visualized in Fig. 7 (top two rows). We provide both novel view renderings and point clouds for illustrating the appearance and geometry, respectively. For the image-to-3D application, we use the image-conditioned multi-view diffusion model in Zero123++ [52], which generates 6 structured views at fixed viewpoints. Though being trained with 4 input views, our transformer-based GS-LRM can take a variable number of images, e.g., 6 images, during inference; hence we simply input the 6 generated images (along with their cameras) into our object-level GS-LRM to predict the 3D Gaussians. We also show our novel-view renderings and a point cloud visualization of 3D Gaussians in Fig. 7 (bottom two rows).

Text-to-scene. We adopt the recently proposed Sora video generation model [42] as the multi-view scene generator from texts. As the Sora model has not publicly



Fig. 6: We show high-res novel-view renderings from our predicted GS given high-res input images (4 512×512 images for objects, and 2 512×904 images for a scene; rendering resolution is the same to the input), to demonstrate our *GS-LRM*'s capability to represent fine-grained details, e.g., readable texts (top left), translucent and thin structures (top right, bottom). Image sources are from GSO (top left), ABO (top right), and RealEstate10K (bottom).

released, we use the generated videos published in [42]. Our current model is limited to static scenes only, and we thus pick the generated videos from relevant



Fig. 7: Text-to-3D (top two rows) and image-to-3D (bottom two rows) results by chaining Instant3D's [32] text-conditioned and Zero123++'s [52] image-conditioned multi-view generators to our GS-LRM reconstructor. For each result, we show novel view renderings and a visualization of the point cloud with point positions and colors extracted from the predicted Gaussians.

text prompt guidance. We use COLMAP [49] to register the video frames, then feed a selected sparse subset of posed frames into our scene-level GS-LRM to reconstruct the 3D scene. The visualization is in Fig. 8.

4.6 Limitations

Although our method shows high-quality reconstruction results from posed sparse images, there are still a few limitations to be addressed in future work. Firstly, the highest resolution our model can currently operate on is about 512×904 ; it is of significant interest to extend the model to work on 1K, even 2K resolution images for best visual quality. Secondly, our model requires known camera parameters; this assumption may not be practical in certain application scenarios (e.g., a user only captures 4 views around an object, making it extremely challenging for SfM to work). To make our model more practical, it is interesting to explore ways to get rid of the camera parameters from the input end [28,61,63]. Thirdly, the pixel-aligned representation only explicitly models the surface inside



Fig. 8: Text-to-scene result. We take the generated video from the Sora text-to-video model [42] (Prompt: drone view of waves crashing against the rugged cliffs along Big Sur's garay point beach). We select 4 frames and reconstruct the 3D scene from them using our scene-level GS-LRM. We show novel-view RGB/depth rendering from our predicted GS (top) and 4 input images (bottom). Please refer to our project page for the video and interactive rendering results.

the view frustum, which means that unseen regions cannot be reconstructed. We found that the model has a certain ability to re-purposing points for hallucinating unseen parts (also observed in [58]), but this capacity is limited and not guaranteed. We leave it to future work to improve the unseen regions.

5 Conclusion

In this work, we present a simple and scalable transformer-based large reconstruction model for Gaussian Splatting (GS) representation. Our method enables fast feed-forward high-res GS prediction from a sparse set of posed images in ~ 0.23 seconds on a single A100 GPU. Our model can work on both object-level and scene-level captures, and achieves state-of-the-art performance in both scenarios when trained on large amount of data. We hope that our work can inspire more future work in the space of data-driven feed-forward 3D reconstruction.

Acknowledgement

We thank Nathan Carr and Duygu Ceylan for useful discussions.

References

- Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Curless, B., Seitz, S.M., Szeliski, R.: Building rome in a day. Communications of the ACM 54(10), 105–112 (2011) 3
- Aliev, K.A., Sevastopolsky, A., Kolos, M., Ulyanov, D., Lempitsky, V.: Neural point-based graphics. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16. pp. 696–712. Springer (2020) 3
- Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016) 5
- 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) 3
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mipnerf 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5470– 5479 (2022) 3
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. arXiv preprint arXiv:2304.06706 (2023) 3
- Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16123–16133 (2022) 1, 3
- Charatan, D., Li, S., Tagliasacchi, A., Sitzmann, V.: pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. arXiv preprint arXiv:2312.12337 (2023) 4, 7, 8, 10
- 9. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: European Conference on Computer Vision (ECCV) (2022) 1, 3, 8
- Chen, A., Xu, Z., Wei, X., Tang, S., Su, H., Geiger, A.: Dictionary fields: Learning a neural basis decomposition. ACM Transactions on Graphics (TOG) 42(4), 1–12 (2023) 3
- 11. Chen, A., Xu, Z., Wei, X., Tang, S., Su, H., Geiger, A.: Factor fields: A unified framework for neural fields and beyond. arXiv preprint arXiv:2302.01226 (2023) 3
- Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., Su, H.: Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14124–14133 (2021) 4
- Chen, E.M., Holalkere, S., Yan, R., Zhang, K., Davis, A.: Ray conditioning: Trading photo-consistency for photo-realism in multi-view image generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 23242–23251 (October 2023) 4
- Chen, Q., Koltun, V.: Photographic image synthesis with cascaded refinement networks. In: Proceedings of the IEEE international conference on computer vision. pp. 1511–1520 (2017) 6
- Chen, T., Xu, B., Zhang, C., Guestrin, C.: Training deep nets with sublinear memory cost. arXiv preprint arXiv:1604.06174 (2016) 8

- 16 K. Zhang, S. Bi, H. Tan et al.
- Cheng, S., Xu, Z., Zhu, S., Li, Z., Li, L.E., Ramamoorthi, R., Su, H.: Deep stereo using adaptive thin volume representation with uncertainty awareness. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2524–2534 (2020) 3
- Collins, J., Goel, S., Deng, K., Luthra, A., Xu, L., Gundogdu, E., Zhang, X., Vicente, T.F.Y., Dideriksen, T., Arora, H., et al.: Abo: Dataset and benchmarks for real-world 3d object understanding. In: CVPR. pp. 21126–21136 (2022) 7
- Dao, T.: Flashattention-2: Faster attention with better parallelism and work partitioning. arXiv preprint arXiv:2307.08691 (2023) 8
- Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: CVPR. pp. 13142–13153 (2023) 3, 7
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020) 4, 5
- Downs, L., Francis, A., Koenig, N., Kinman, B., Hickman, R., Reymann, K., McHugh, T.B., Vanhoucke, V.: Google scanned objects: A high-quality dataset of 3d scanned household items. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 2553–2560. IEEE (2022) 7
- Du, Y., Smith, C., Tewari, A., Sitzmann, V.: Learning to render novel views from wide-baseline stereo pairs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4970–4980 (2023) 7, 10
- Furukawa, Y., Ponce, J.: Accurate, dense, and robust multiview stereopsis. IEEE transactions on pattern analysis and machine intelligence **32**(8), 1362–1376 (2009) 1, 3
- Gao, Q., Xu, Q., Su, H., Neumann, U., Xu, Z.: Strivec: Sparse tri-vector radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17569–17579 (2023) 3
- Gu, X., Fan, Z., Zhu, S., Dai, Z., Tan, F., Tan, P.: Cascade cost volume for high-resolution multi-view stereo and stereo matching. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2495–2504 (2020) 3
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) 5
- Hong, Y., Zhang, K., Gu, J., Bi, S., Zhou, Y., Liu, D., Liu, F., Sunkavalli, K., Bui, T., Tan, H.: Lrm: Large reconstruction model for single image to 3d (2023) 1, 4, 6, 7
- Jiang, H., Jiang, Z., Zhao, Y., Huang, Q.: Leap: Liberate sparse-view 3d modeling from camera poses. ArXiv 2310.01410 (2023) 13
- Johari, M.M., Lepoittevin, Y., Fleuret, F.: Geonerf: Generalizing nerf with geometry priors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18365–18375 (2022) 4
- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics 42(4) (2023) 1, 2, 3, 6, 9
- Lefaudeux, B., Massa, F., Liskovich, D., Xiong, W., Caggiano, V., Naren, S., Xu, M., Hu, J., Tintore, M., Zhang, S., Labatut, P., Haziza, D., Wehrstedt, L., Reizenstein, J., Sizov, G.: xformers: A modular and hackable transformer modelling library. https://github.com/facebookresearch/xformers (2022) 8

- 32. Li, J., Tan, H., Zhang, K., Xu, Z., Luan, F., Xu, Y., Hong, Y., Sunkavalli, K., Shakhnarovich, G., Bi, S.: Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model (2023) 1, 4, 6, 7, 8, 9, 11, 13
- Lin, K.E., Xu, Z., Mildenhall, B., Srinivasan, P.P., Hold-Geoffroy, Y., DiVerdi, S., Sun, Q., Sunkavalli, K., Ramamoorthi, R.: Deep multi depth panoramas for view synthesis. In: European Conference on Computer Vision. pp. 328–344. Springer (2020) 3
- Liu, L., Gu, J., Zaw Lin, K., Chua, T.S., Theobalt, C.: Neural sparse voxel fields. Advances in Neural Information Processing Systems 33, 15651–15663 (2020) 3
- Lombardi, S., Simon, T., Schwartz, G., Zollhoefer, M., Sheikh, Y., Saragih, J.: Mixture of volumetric primitives for efficient neural rendering. ACM Transactions on Graphics (ToG) 40(4), 1–13 (2021) 3
- Long, X., Lin, C., Wang, P., Komura, T., Wang, W.: Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In: European Conference on Computer Vision. pp. 210–227. Springer (2022) 4, 9
- Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., Ginsburg, B., Houston, M., Kuchaiev, O., Venkatesh, G., et al.: Mixed precision training. In: International Conference on Learning Representations (2018) 8
- 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) 38(4), 1–14 (2019) 3
- 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) 8
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020) 1, 3
- 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) 1, 3
- OpenAI: Creating video from text (February 2024), https://openai.com/sora 2, 11, 12, 14
- Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16. pp. 523–540. Springer (2020) 1
- Plücker, J.: Xvii. on a new geometry of space. Philosophical Transactions of the Royal Society of London (155), 725–791 (1865)
- Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. In: The Twelfth International Conference on Learning Representations (2023) 11
- Pollefeys, M., Nistér, D., Frahm, J.M., Akbarzadeh, A., Mordohai, P., Clipp, B., Engels, C., Gallup, D., Kim, S.J., Merrell, P., et al.: Detailed real-time urban 3d reconstruction from video. International Journal of Computer Vision 78, 143–167 (2008) 3
- 47. Pollefeys, M., Van Gool, L., Vergauwen, M., Verbiest, F., Cornelis, K., Tops, J., Koch, R.: Visual modeling with a hand-held camera. International Journal of Computer Vision 59, 207–232 (2004) 3

- 18 K. Zhang, S. Bi, H. Tan et al.
- Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4104–4113 (2016) 1, 3
- Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016) 7, 13
- Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise view selection for unstructured multi-view stereo. In: European Conference on Computer Vision (ECCV) (2016) 1, 3
- Shen, Z., Dai, Y., Rao, Z.: Cfnet: Cascade and fused cost volume for robust stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13906–13915 (2021) 3
- Shi, R., Chen, H., Zhang, Z., Liu, M., Xu, C., Wei, X., Chen, L., Zeng, C., Su, H.: Zero123++: a single image to consistent multi-view diffusion base model (2023) 11, 13
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) 6
- Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3d. In: ACM siggraph 2006 papers, pp. 835–846 (2006) 3
- 55. Suhail, M., Esteves, C., Sigal, L., Makadia, A.: Generalizable patch-based neural rendering. In: European Conference on Computer Vision. Springer (2022) 4, 7, 10
- Suhail, M., Esteves, C., Sigal, L., Makadia, A.: Light field neural rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8269–8279 (2022) 4
- 57. Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5459–5469 (2022) 3
- Szymanowicz, S., Rupprecht, C., Vedaldi, A.: Splatter image: Ultra-fast single-view 3d reconstruction. arXiv: (2023) 14
- 59. Tang, J., Chen, Z., Chen, X., Wang, T., Zeng, G., Liu, Z.: Lgm: Large multiview gaussian model for high-resolution 3d content creation. arXiv preprint arXiv:2402.05054 (2024) 4, 7, 9
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017) 5
- 61. Wang, P., Tan, H., Bi, S., Xu, Y., Luan, F., Sunkavalli, K., Wang, W., Xu, Z., Zhang, K.: Pf-lrm: Pose-free large reconstruction model for joint pose and shape prediction. arXiv preprint arXiv:2311.12024 (2023) 1, 4, 5, 6, 13
- Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snavely, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4690–4699 (2021) 4
- Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: Dust3r: Geometric 3d vision made easy. arXiv preprint arXiv:2312.14132 (2023) 13
- Wiles, O., Gkioxari, G., Szeliski, R., Johnson, J.: Synsin: End-to-end view synthesis from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7467–7477 (2020) 3
- Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5438–5448 (2022) 3

- Xu, Y., Tan, H., Luan, F., Bi, S., Wang, P., Li, J., Shi, Z., Sunkavalli, K., Wetzstein, G., Xu, Z., Zhang, K.: Dmv3d: Denoising multi-view diffusion using 3d large reconstruction model (2023) 1, 4, 6
- Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: Proceedings of the European conference on computer vision (ECCV). pp. 767–783 (2018) 3
- Yao, Y., Luo, Z., Li, S., Shen, T., Fang, T., Quan, L.: Recurrent mysnet for highresolution multi-view stereo depth inference. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5525–5534 (2019) 3
- Yifan, W., Serena, F., Wu, S., Öztireli, C., Sorkine-Hornung, O.: Differentiable surface splatting for point-based geometry processing. ACM Transactions on Graphics (proceedings of ACM SIGGRAPH ASIA) 38(6) (2019) 3
- Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4578–4587 (2021) 3, 4, 7, 10
- Zhang, K., Kolkin, N., Bi, S., Luan, F., Xu, Z., Shechtman, E., Snavely, N.: Arf: Artistic radiance fields (2022) 8
- 72. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492 (2020) 3
- 73. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) 6
- Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. arXiv preprint arXiv:1805.09817 (2018) 3, 7