

LVT: Large-Scale Scene Reconstruction via Local View Transformers

TOOBA IMTIAZ^{1,2*}, LUCY CHAI^{1*}, KATHRYN HEAL¹, XUAN LUO¹, JUNGYEON PARK¹, JENNIFER DY², and JOHN FLYNN¹, ¹Google, USA, ²Northeastern University, USA




Fig. 1. **Local View Transformer (LVT) for scalable scene reconstruction.** Instead of the quadratic self-attention operation typically used in transformers, LVT processes information within local neighborhoods relative to each view, based on the insight that spatially nearby views are generally more informative about the local scene structure than distant ones. Cascading these LVT blocks increases the effective receptive field. Our model takes a set of input views representing the entire scene and, in a single feed-forward pass, outputs a Gaussian splat scene representation which can then be rendered to various target cameras. LVT scales linearly, rather than quadratically, with the number of input images, and generalizes to varied sequence lengths and out-of-distribution camera trajectories.

Large transformer models are proving to be a powerful tool for 3D vision and novel view synthesis. However, the standard Transformer’s well-known quadratic complexity makes it difficult to scale these methods to large scenes. To address this challenge, we propose the Local View Transformer (LVT), a large-scale scene reconstruction and novel view synthesis architecture that circumvents the need for the quadratic attention operation. Motivated by the insight that spatially nearby views provide more useful signal about the local scene composition than distant views, our model processes all information in a local neighborhood around each view. To attend to tokens in nearby views, we leverage a novel positional encoding that conditions on the relative geometric transformation between the query and nearby views. We decode the output of our model into a 3D Gaussian Splat scene representation that includes both color and opacity view-dependence. Taken together, the Local View Transformer enables reconstruction of arbitrarily large, high-resolution scenes in a single forward pass. See our project page for results and interactive demos: <https://toobaimt.github.io/lvt/>.

CCS Concepts: • **Computing methodologies** → **Rendering; Computer vision; Neural networks.**

*denotes equal contribution

Authors’ Contact Information: Tooba Imtiaz^{1,2}, imtiaz.t@northeastern.edu; Lucy Chai¹, lucychai@google.com; Kathryn Heal¹, kheal@google.com; Xuan Luo¹, xuluo@google.com; Jungyeon Park¹, jungyeonp@google.com; Jennifer Dy², jdy@ece.neu.edu; John Flynn¹, jflynn@google.com. ¹Google, USA, ²Northeastern University, USA.



This work is licensed under a Creative Commons Attribution 4.0 International License. SA Conference Papers ’25, Hong Kong, Hong Kong
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2137-3/2025/12
<https://doi.org/10.1145/3757377.3763838>

Additional Key Words and Phrases: Neural rendering, novel view synthesis, 3D Gaussian Splatting, feed-forward models

ACM Reference Format:

Tooba Imtiaz^{1,2}, Lucy Chai¹, Kathryn Heal¹, Xuan Luo¹, Jungyeon Park¹, Jennifer Dy², and John Flynn¹. 2025. LVT: Large-Scale Scene Reconstruction via Local View Transformers. In *SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers ’25)*, December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3757377.3763838>

1 Introduction

Echoing their widespread impact across other domains [Carion et al. 2020; Caron et al. 2021; Devlin et al. 2019; Dosovitskiy et al. 2021; Radford et al. 2019; Ranftl et al. 2021], the recent application of Transformer [Vaswani et al. 2017]-based architectures to novel view synthesis (NVS) has achieved remarkable success [Hong et al. 2024; Zhang et al. 2024; Ziwen et al. 2025]. Unlike earlier techniques that relied more heavily on explicit 3D reasoning, these models instead opt for a primarily data-driven approach to infer scene structure from input images.

Despite their state-of-the-art performance, a significant drawback of transformers remains their high computational cost, primarily driven by the self-attention mechanism’s quadratic scaling with token count. Even with optimizations such as Flash Attention [Dao et al. 2022], global self-attention remains costly for long sequences. This computational burden poses a significant bottleneck to applying these models to large real-world scenes with many high resolution input views.

In contrast, per-scene optimization approaches, like 3DGS [Kerbl et al. 2023], are able to perform scene-scale reconstruction but require thousands of test-time gradient steps to do so. Other methods [Flynn et al. 2024, 2016] sidestep a persistent global scene representation and opt instead to reconstruct local geometry for each rendered view on-the-fly as the camera moves through the scene. Although this strategy is efficient, the lack of a global model results in noticeable flickering as the camera moves.

Recently, Long-LRM [Ziwen et al. 2025] addresses the computational cost of transformers by employing a hybrid architecture that integrates Mamba2 [Dao and Gu 2024a] and transformer blocks. The method allows scaling to a considerably larger number of input views than prior transformer-based scene-reconstruction methods. However, the approach requires several optimizations for efficiency, such as heuristic Gaussian merging and pruning. It also lacks natural generalization to varied-length sequences, and needs camera selection to maintain a fixed number of inputs. This highlights the need for more scalable and robust generalizable reconstruction techniques.

To this end, we propose the Local View Transformer (LVT), an architecture designed for efficient and high-fidelity 3D Gaussian splat reconstruction at the scene scale in a single inference pass. In contrast with prior global attention-based transformer models for 3D reconstruction, LVT exploits the inherent spatial locality in 3D scenes. We motivate our design by drawing a parallel to Convolutional Neural Networks [LeCun et al. 2015] (CNNs). CNNs largely replaced earlier neural network architectures based on dense connections with a local sliding window mechanism. The CNN layer brought several key advantages: enhanced efficiency, improved generalization via translation invariance, and the ability to scale to arbitrary-sized inputs. While an individual CNN’s receptive field is limited, stacking multiple CNN layers progressively widens the receptive field, enabling broader contextual understanding.

In LVT, we extend this localized-processing paradigm to the attention mechanism within a transformer. Specifically, our attention mechanism operates locally, limiting tokens from an input or query view to attend *only* to tokens within a small set of neighboring views. The definition of “neighboring” can vary; for a scene capture containing sequential images, this might involve attending to the previous and next views. For a more general collection of views, this could involve choosing views based on the distance between camera centers. Although each LVT transformer block limits attention to nearby views, stacking multiple blocks progressively expands the ‘view receptive field,’ in a similar manner to a stack of CNN layers.

Mirroring the translation invariance inherent in CNNs, our method is designed to be *transformation* invariant to the 3D coordinate frame of the query view. This means that tokens within each view encode splat geometry relative to that view’s local coordinate frame. Consequently, when attending from a query view to a neighboring view, the token features in the neighboring view must be transformed into the query view’s coordinate frame. To achieve this we incorporate a straightforward encoding of the relative poses between views, and found this sufficient to achieve the desired invariance. This invariance allows LVT to transfer priors learned from shorter sequences during training to arbitrarily long sequences during inference.

Finally, we introduce view-dependent opacity, allowing for a richer representation of specular highlights, dynamic lighting effects, and thin objects within scene reconstructions, in contrast to prior methods that solely encode view-dependent colors [Fridovich-Keil et al. 2022; Kerbl et al. 2023; Mildenhall et al. 2020].

Altogether, our LVT model surpasses prior methods to achieve state-of-the-art novel view synthesis on large scenes. Quantitatively, on the challenging DL3DV benchmark [Ling et al. 2024], LVT outperforms the previous state-of-the-art, Long-LRM [Ziwen et al. 2025], by a significant margin of +3.54dB PSNR and 3DGS, which is optimized per scene, by +2.11dB PSNR. Our model also exhibits zero-shot generalization capabilities on the Tanks&Temples dataset [Knapitsch et al. 2017], outperforming Long-LRM by +2.79dB PSNR, and per-scene optimized 3DGS by +0.72dB PSNR. In summary:

- We propose LVT, a transformer architecture for feed-forward novel view synthesis of large scenes that scales linearly to long and varied-length sequences.
- Our model introduces a local attention module in which image tokens attend only to tokens from nearby views thus eliminating the costly self-attention operation. To achieve this local processing we condition the tokens on relative camera poses,
- For further improvements, we incorporate view-dependent opacity prediction to enhance rendering fidelity, particularly in scenes with thin structures and complex lighting effects.
- Our extensive experiments and ablations demonstrate an effective, scalable approach for novel view synthesis of full scenes in a single inference pass.

2 Related Work

3D reconstruction and novel view synthesis have long been central challenges in computer vision. Early works used custom algorithms and heuristics to generate a scene representation from a set of input images [Agarwal et al. 2009; Chaurasia et al. 2013; Kutulakos and Seitz 2000; Penner and Zhang 2017; Seitz and Dyer 1999]. In contrast, [Kerbl et al. 2023; Mildenhall et al. 2020] leverage differentiable rendering to directly optimize a dedicated 3D representation for individual scenes. These techniques deliver high fidelity at the expense of requiring substantial test-time optimization, typically ranging from minutes to hours. Additionally, since these methods optimize on a per-scene basis, they do not learn generalizable priors applicable across different scenes. Consequently, to prevent overfitting, they require dense input views or the addition of ad-hoc regularization losses.

Feed-forward models. Generalizable methods avoid these drawbacks by predicting the entire scene in a single feed-forward pass though a deep learning pipeline [Aliiev et al. 2020; Flynn et al. 2016; Lin et al. 2020; Yifan et al. 2019; Zhou et al. 2018]. One approach for these methods is to leverage explicit 3D geometric structures, such as epipolar lines [Suhail et al. 2022; Wang et al. 2021] or plane-sweep cost volumes [Chen et al. 2021, 2024; Flynn et al. 2019, 2024, 2016; Johari et al. 2022; Liu et al. 2020; Mildenhall et al. 2019; Xu et al. 2025], to aggregate information from input views. While these methods can be effective for a small number of inputs views from forward-facing cameras, they lack a unified representation and




Fig. 2. **Overview of Local View Transformer architecture.** The pipeline takes input view sequences and the corresponding local ray maps, and patchifies them to obtain input tokens. Within each LVT block, the tokens from a window of w neighbor views (here $w = 3$ for illustration) relative to a query view are consolidated and updated with the corresponding relative transformation embeddings. The tokens from the *query* view then selectively attend to these neighbor tokens (instead of attending to all input tokens, as full self-attention does). This LVT block is repeated 24 times. The processed tokens are then unpatchified and decoded to pixel-aligned Gaussian splat parameters. The combined Gaussian splats are rendered using the 3DGS renderer and compared to ground truth during training.

typically resort to transitioning between local representations to process larger, general scenes.

Large reconstruction models. With the rise of transformers, view synthesis methods have moved away from explicit 3D inductive biases towards fully learned, data-driven geometric priors, culminating in Large Reconstruction Models (LRMs). Existing LRMs employ two primary output strategies: either decoding to an explicit 3D representation, such as triplane NeRF fields [Hong et al. 2024; Li et al. 2023; Wang et al. 2024b] or pixel-aligned 3-D Gaussian splats [Tang et al. 2024; Zhang et al. 2024], or directly decoding to target pixels, as exemplified by LVSM [Jin et al. 2025]. These models achieve impressive results from extremely sparse inputs—often one to four images with little or no view overlap—by leveraging the global context captured by transformers. However, the quadratic complexity of transformer attention poses a significant practical challenge to scaling these models to real-world scenes with numerous input views or high-resolution imagery. Thus, these prior models focus on short, low-resolution captures.

Scaling to longer input sequences. Recent work, such as Long-LRM [Ziwen et al. 2025] and Fast3R [Yang et al. 2025], explore scalable NVS architectures that can handle longer input sequences. Long-LRM [Ziwen et al. 2025] mitigates the computational cost of transformers by combining Mamba2 [Dao and Gu 2024b] and transformer blocks, pushing prior feed-forward LRMs [Jin et al. 2025; Zhang et al. 2024] from roughly 2–4 to 32 input views. However, the use of Mamba2 sacrifices the standard transformer model’s natural order invariance. This property is particularly beneficial for view synthesis, as input views often lack a definitive ordering. Fast3R [Yang et al. 2025] generalizes DUST3R [Wang et al. 2024a] from two input views to longer sequences for camera registration. However, both models at their core still adopt full self-attention,

which becomes computationally infeasible at for a large number of input views. Spann3R [Wang and Agapito 2025] and CUT3R [Wang* et al. 2025] build the scene incrementally using a memory bank or a recurrent state.

In contrast, our proposed Local View Transformer explores an alternative approach to handling long and variable-length input sequences, inspired by the spatial nature of 3D scenes. Unlike prior methods, LVT replaces *global* full self-attention with a more *localized* attention operation across tokens from neighboring views. We find that this strategy largely maintains the benefits of the attention module while scaling linearly with the number of input views.

Positional encoding. Transformers are by design permutation invariant, so positional encodings are used to provide the model with information about the order of input tokens. Approaches such as RoPE have explored using relative, rather than absolute, positional embeddings to improve scalability and generalization [Heo et al. 2024; Su et al. 2024]. Analogously, recent works have explored embedding camera information as a relative positional encoding for geometric tasks to allow for invariance towards the choice of a global reference frame [Kong et al. 2024; Li et al. 2025; Miyato et al. 2024]. Our Local View Transformer also adopts relative camera poses as a positional embedding signal; combined with our windowed attention mechanism, this maintains transformation invariance in the local neighborhoods for a more scalable transformer architecture over long input sequences.

View dependence in 3D reconstruction. Accurately modeling view-dependent effects, such as specular reflections and highlights, has been a consistent challenge in novel view synthesis. Ref-NeRF [Verbin et al. 2022] improves upon NeRF’s [Mildenhall et al. 2020] basic view dependence by factorizing color into diffuse and specular components, yielding more accurate reflections particularly those

with complex lighting and reflective materials. Drawing from their application in graphics, various approaches [Fridovich-Keil et al. 2022; Wizarwongsa et al. 2021; Yu et al. 2021a] employ spherical harmonics [Ramamoorthi and Hanrahan 2001] to model view-dependent effects. Spherical harmonics simplify the calculation of view-dependent color to a straightforward dot product between the viewing direction and the predicted coefficients.

While the use of view dependent color to model directional effects is widespread, the use of view-dependent opacity has been relatively limited in prior work, possibly due to overfitting risks for methods that employ per-scene optimization. Gaussian Opacity Fields (GOF) [Yu et al. 2024] perform heuristic merging of per-view opacity. VoD-3DGS [Nowak et al. 2025], uses a learnable per-Gaussian symmetric matrix that modulates opacity based on viewing direction. However, their view-dependent appearance model struggles to reconstruct scenes with low specular content, and suffers from degraded quantitative performance.

In our approach, we also include view-dependent opacity, finding that a straightforward application of spherical harmonics to opacity substantially improves rendered quality, especially on long input sequences.

3 Method

The Local View Transformer adapts the transformer architecture to process tokens in a locally-aware fashion. Below, we detail our model components and their differences from the standard Transformer.

3.1 Architecture Overview

As illustrated in Figure 2, given N_i input images $\mathbf{I}$ with known camera intrinsics and poses, i.e. $\mathbf{I} = \{I_j \in \mathbb{R}^{H \times W \times 3}, K_j, P_j\}, j \in \{1, \dots, N_i\}$, the Local View Transformer (LVT) predicts the parameters of input pixel-aligned 3D Gaussian splats. We train the network end-to-end by rendering from these splats to N_t target views $\mathbf{T} = \{T_k \in \mathbb{R}^{H \times W \times 3}, K_k, P_k\}, k \in \{1, \dots, N_t\}$ and comparing the rendered target images to the corresponding ground truth images. Our architecture mimics the structure of GS-LRM [Zhang et al. 2024], beginning with tokenization into image patches, followed by several transformer blocks, and ending with a decoding block that converts the tokens into per-pixel Gaussian Splat parameters [Kerbl et al. 2023]. However, we make modifications to ensure *transformation invariance* at each step, i.e. all processing is performed locally, relative to each input view, and independent of an arbitrary world coordinate system.

3.2 Tokenizing Images to Patches

Following prior works [Dosovitskiy et al. 2021], we first patchify the inputs into a series of non-overlapping tokens with patch-size p . The conventional practice in related LRM models [Jin et al. 2025; Zhang et al. 2024; Ziwen et al. 2025] concatenates input pixels with Plücker ray coordinates [Plücker 1865] to encode input camera parameters along the channel dimension, which are converted to tokens representing each $p \times p$ patch via a convolutional kernel.

In contrast, as our LVT model processes views *locally*, our input tokens do not need to capture where the view is situated in 3D space. Rather than Plücker rays which encode both the camera intrinsics

and extrinsics, we compute the set of normalized local ray maps $\mathbf{R} = \{R_j \in \mathbb{R}^{H \times W \times 3}, j = 1, \dots, N_i\}$ using just the intrinsics K_j . These local ray maps are concatenated to the image pixels and patchified into a set of tokens $\mathbf{M} \in \mathbb{R}^{N_i \times \frac{H}{p} \times \frac{W}{p} \times d}$ using a convolutional kernel:

$$\begin{aligned} M_j &= \text{Conv}(\text{Concat}([I_j, R_j])) \\ \mathbf{M} &= \{M_j | j \in 1, \dots, N_i\}, \end{aligned} \quad (1)$$

where d refers to the channel dimension. These tokens are then flattened to shape $N_i \times \frac{HW}{p^2} \times d$ and normalized prior to the transformer blocks.

3.3 Local View Transformer

The core idea of our Local View Transformer is that nearby image neighbors provide the most relevant information for a given view. Thus, rather than computing the full quadratic self-attention matrix in the standard Transformer architecture, we only need to compute attention between a query view and its neighbors.

In more detail, for a given image I_j associated with tokens M_j with shape $\frac{HW}{p^2} \times d$, we first define the set of neighbors $\mathcal{N}(j)$, where each image j is also in its own neighbor set. To provide a transformation invariant spatial conditioning signal, instead of directly using camera poses P_j , we use the relative poses $P_j P_{j'}^{-1}$ between the j -th and j' -th input views. We then apply a positional encoding [Vaswani et al. 2017] on the relative quaternions and translations. The positionally-encoded values are next processed by several residual MLP blocks to produce a set of conditioning features $C_{(j,j')} \in \mathbb{R}^d$ that capture the relative spatial information between neighboring views. Relatedly, rotary positional embeddings [Heo et al. 2024; Su et al. 2024] compute transformations on the key and query feature spaces to encode the relative rather than absolute positions of transformer tokens, whereas our conditioning features operate on the relative rigid-body transformation between neighboring camera poses. Several heuristics, including nearest w spatial neighbors, can be used to determine the set of neighbors $\mathcal{N}(j)$. In our experiments, we use the distance between camera centers to select the neighbors. We discuss alternative choices in our ablation study.

Next, we modify the transformer attention block to attend only to neighboring tokens and condition on the relative transformation feature. We use the notation $M_j^{(m,l)}$ to denote the m -th token at the l -th layer corresponding to input image j , where $m \in 1, \dots, \frac{HW}{p^2}$, and $C_{(j,j')}^{(l)}$ as the conditioning feature at that layer, transformed via a per-layer Dense projection. We compute:

$$\begin{aligned} \mathbf{Q}_j^{(m,l)} &= W_q M_j^{(m,l)} \\ \mathbf{K}_{j'}^{(m',l)} &= W_k \left(M_{j'}^{(m',l)} + C_{(j,j')}^{(l)} \right) \\ \mathbf{V}_{j'}^{(m',l)} &= W_v \left(M_{j'}^{(m',l)} + C_{(j,j')}^{(l)} \right) \end{aligned} \quad (2)$$

In essence, the keys and values contain not only information about the scene structure from the neighbors in $M_{j'}^{(m',l)}$ but also information about the relative positioning between neighbors in $C_{(j,j')}$. The

resulting feature becomes:

$$M_j^{(m, l+1)} = \underset{j' \in \mathcal{N}(j), m' \in 1, \dots, \frac{HWN_i}{p^2}}{\text{softmax}} \left(\frac{\mathbf{Q}_j^{(m, l)} \mathbf{K}_{j'}^{(m', l)T}}{\sqrt{d}} \right) \mathbf{V}_{j'}^{(m', l)} \quad (3)$$

The remainder of LVT follows the Vision Transformer architecture, interleaving our modified residual neighborhood attention blocks with standard residual feed-forward blocks. We use pre-layer normalization [Xiong et al. 2020] before each attention and feed-forward layer and use Flash Attention [Dao et al. 2022] within the attention blocks. We note that cascading the LVT blocks effectively grows the receptive field of the tokens, expanding the scope of attention to more views with each subsequent layer.

Following the transformer blocks, each processed token, $M_j^{(m, L)}$ at the final layer L is unpatchified using a transposed convolution layer into a patch size of $p \times p$ with d_G channels to yield pixel-aligned Gaussian parameters, where d_G is the total number of Gaussian parameters. At each pixel, we parameterize the corresponding 3D Gaussian by a 3-channel scale, a 4-channel quaternion, scalar depth along the pixel’s ray, d_c -channel color and d_o -dimensional opacity σ . The translation, rotation and scale are relative to the corresponding view’s 3D coordinate frame and are transformed into the world coordinate frame for rendering. During training, the decoded $N_i \times H \times W$ Gaussians are rendered to target views using the differentiable 3DGS rasterizer [Kerbl et al. 2023].

3.4 View-Dependent Splat Opacity

Previous works have used spherical harmonics to represent view-dependent effects in both 3D Gaussian Splat and other neural representations [Fridovich-Keil et al. 2022; Müller et al. 2022; Sloan et al. 2023]. We extend the application of spherical harmonics coefficients from color representation to also compute view-dependent opacity. Although view-dependent opacity has been less frequently utilized in prior work [Heitz et al. 2015; Nowak et al. 2025], potentially due to over-parameterization concerns, we note that LVT inherits the generalizability of feed-forward methods without overfitting. We find that this addition of view-dependent opacity better captures thin structures and sharp boundaries especially over long input sequences. This allows for a representation that works well locally and transitions smoothly as the camera changes. As a result, each view-dependent splat in our model is parameterized by $d_c = 3 \times (\text{deg_SH}_{\text{color}} + 1)^2$ and $d_o = (\text{deg_SH}_\sigma + 1)^2$ channels for color and opacity respectively. The spherical harmonics are computed relative to each view’s coordinate frame, consistent with our focus on transformation invariance, and are transformed to world space for rendering.

Regularization. To avoid spurious opacities when viewed from novel viewing angles, we introduce a regularization term that penalizes extreme magnitudes of splat opacity evaluated at randomly sampled viewing directions on the surface of a unit sphere. Let $\hat{\mathbf{n}}$ be a randomly sampled direction encoded with the spherical harmonic basis. We define the regularizer

$$R_\sigma = \frac{1}{HWN_i} \sum_{n=1}^{HWN_i} |\sigma_{(n)} \cdot \hat{\mathbf{n}}| \quad (4)$$

where $\sigma_{(n)}$ refers to the opacity coefficients corresponding to the n -th pixel. In our models that do not use view-dependent opacity, we simply regularize $R_\sigma = \frac{1}{HWN_i} \sum_{n=1}^{HWN_i} |\sigma_{(n)}|$ following LongLRM [Ziwen et al. 2025] where $\sigma_{(n)}$ is a per-pixel scalar opacity.

3.5 Training Objective

We render the predicted Gaussian splats using the target cameras and minimize the loss between the rendered images $\hat{\mathbf{T}}$ and the ground truth target views $\mathbf{T}$. Our reconstruction loss function $\mathcal{L}$ is a weighted combination of the Mean Squared Error (MSE) loss, Perceptual loss, and opacity regularization:

$$\mathcal{L} = \frac{1}{N_t} \sum_{k=1}^{N_t} \left(\text{MSE}(\hat{T}_k, T_k) + \lambda \cdot \text{Perceptual}(\hat{T}_k, T_k) \right) + \alpha \cdot R_\sigma, \quad (5)$$

where λ is the weighting factor for the perceptual loss and α is a weighting factor for the opacity regularizer. Following prior LRMs [Hong et al. 2024; Wang et al. 2024b; Xu et al. 2023], we adopt LPIPS [Zhang et al. 2018] as the perceptual loss, in contrast to [Zhang et al. 2024] which uses a VGG-based [Simonyan and Zisserman 2014] loss for training stability.

4 Experiments

We provide training and implementation details of our method (Sec. 4.1), comparison against baselines (Sec. 4.2), and ablate the effect of various design choices (Sec. 4.3). We include additional experiments, including a model trained for 2K resolution, more ablations, and video results in our supplementary material.

4.1 Implementation Details

Datasets. We train our model on DL3DV [Ling et al. 2024], a large-scale dataset of 4K resolution videos captured from bounded and unbounded real-world scenes for benchmarking novel view synthesis. The training and test splits, DL3DV-10K and DL3DV-140, consist of 10,510 and 140 videos, respectively. The camera poses in each video are annotated using COLMAP [Schönberger et al. 2016]. For training, we sample a sequence of input images separated by a random stride between four and eight, and sample target images between the inputs. We construct the neighborhood graph by connecting each image $I(j)$ to its closest neighbors $\mathcal{N}(j)$ based on camera center distances.

We evaluate on the DL3DV test split and additionally perform zero-shot evaluation on Tanks&Temples [Knapitsch et al. 2017] (the train and truck scenes) and Mip-NeRF360 [Barron et al. 2022]. We use every eighth image in the sequence as targets. On DL3DV, we take every eighth image of the remaining images after removing the targets; as such, the number of input images varies with the length of the image sequence, and the distance of the targets to the nearest input also varies. On Tanks&Temples and Mip-NeRF360, we reduce our input stride to every fourth image after removing targets due to the larger camera spacing. We use temporal order for neighborhood graph on DL3DV and Tanks&Temples; for Mip-NeRF360 we compute neighbors based on camera center distances, as the cameras capture an upper hemisphere rather than a video

Model	Feed-Forward	DL3DV-140			Tanks&Temples			MipNeRF360		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
3DGS _{30k}	✗	25.53	0.849	0.210	20.45	0.780	0.251	24.81	0.745	0.246
Long-LRM	✓	24.10	0.783	0.254	18.38	0.601	0.363	–	–	–
LVT _{base}	✓	23.65	0.813	0.215	18.33	0.681	0.296	21.51	0.603	0.335
LVT _{SH-rgba}	✓	27.64	0.883	0.133	21.17	0.794	0.174	<u>24.53</u>	<u>0.698</u>	<u>0.251</u>

Table 1. **Comparison of LVT against baselines.** We evaluate at 960×540 resolution on the DL3DV-140 test set and perform zero-shot evaluation on Tanks&Temples and MipNeRF360. Our LVT_{SH-rgba} model outperforms both baselines across all metrics for DL3DV140, and perform competitively with 3DGS in a single feed-forward pass on the latter benchmark datasets. Our model is able to generalize to Tanks&Temples and MipNeRF360 datasets, despite differences in the sequence lengths and camera trajectories than those used for training. We obtain numbers of Long-LRM from their paper as the model is not publicly available, and note that Long-LRM is unable to generalize to Mip-NeRF360. We retrain 3DGS for each scene using the same input views as our model.

flight trajectory. In the supplementary (Sec. B.2), we also train and evaluate a model on RealEstate10k [Zhou et al. 2018].

Training details. We follow the basic transformer structure of GS-LRM [Zhang et al. 2024], using 24 transformer layers, with hidden dimensions of 1024 for the transformer and 4096 for the MLP. We use 16 heads for the attention layer. We scale the initialization of the final kernels of each residual block (in the attention and MLP modules) by $\sqrt{\frac{1}{N_r}}$ where N_r is the total number of residual connections, and otherwise use LeCun Normal initialization on the learnable parameters and omit biases. We use a peak learning rate of $2e-4$ for 150K steps, consisting of 2K steps linear warmup on the learning rate and cosine decay for the remaining steps. We first train at a maximum resolution of 480×270 , and then finetune at 960×540 for an additional 20K steps with peak learning rate $1e-5$ following 2K steps of linear warmup. We use loss weights $\lambda = 0.05$ and $\alpha = 0.001$, and batch size 256.

Multi-resolution, multi-length dataset mixing. As the computational costs of the transformer and rendering scale with the resolution and number of input images, we use a dataset mixing strategy to allow the model to generalize to longer sequences of varied length. That is, we train on three resolutions of the dataset: 240×135 , 480×270 , and 960×540 . We train on the longest sequences at the low 240×135 resolution, sampling between 40 and 52 input images. At the middle 480×270 resolution, we sample up to 24 input images, and at 960×540 resolution we sample 16 input images. Additional details are provided in the supplementary (Tab. 4).

4.2 Comparison Against Baselines

We train several versions of our LVT model with variations in the spherical harmonics coefficients. Our base model (LVT_{base}) does not have any view-dependence in the representation, while our full model (LVT_{SH-rgba}) uses view-dependence on both color and opacity. We compare our results to Long-LRM [Ziwen et al. 2025], which uses full self-attention and thus does not naturally generalize to variable-length input sequences and requires an input selection algorithm to retain a consistent number of inputs with optimal coverage of the scene. As the pre-trained model for Long-LRM is not available, we directly use the numbers reported by the authors. We also compare to 3DGS [Kerbl et al. 2023], which is an optimization-based approach that requires several minutes of optimization per scene. We generate a 3DGS representation using the same input

images as our model, simply selecting every eighth image after removing all target images.

As shown in Table 1, LVT_{SH-rgba} outperforms both Long-LRM and 3DGS in PSNR, SSIM, and LPIPS on DL3DV. Our LVT_{base} model, without spherical harmonics, offers a substantial boost in SSIM and LPIPS over Long-LRM, highlighting the effectiveness of variable-length input selection, as we do not require any heuristics to select a maximally informative fixed-length input set.

We also evaluate our model trained on DL3DV on the *train* and *truck* examples from the Tanks&Temples dataset, which we resize to 960×540 following Long-LRM (Figure 4). We run our model on 66 inputs for train and 55 for truck, longer than the sequence lengths that we use for training. Here, our LVT_{base} model surpasses Long-LRM in all three metrics. We hypothesize that our local attention operation also helps with generalization given this different distribution of camera trajectories, as local segments of the Tanks&Temples camera paths are more similar to those of DL3DV than the entire trajectory as a whole. Our full model achieves better scores than our base model and, despite not being trained on this data, surpasses 3DGS in LPIPS, PSNR and SSIM. We also evaluate our model on the MipNeRF360 dataset (Figure 5). Notably, Long-LRM reports inability to generalize to MipNeRF360. LVT outperforms 3DGS on this benchmark as well, and does so in a single inference pass. We benchmark the inference time of our model, which scales linearly with the number of inputs, in Supp. Sec. B.1.

4.3 Ablations

We conduct ablations to validate the design choices and the effectiveness of our proposed Local View Transformer (LVT). We train all ablation models at 480×270 resolution with batch size 64.

Generalization to varying sequence lengths. We compare our LVT local attention to a model with full self-attention, following GS-LRM, but with otherwise identical training and model design to ours (Table 2). We train both models on 16 inputs, and evaluate on 12, 16, 20, and 32 inputs to understand how sensitive the models are to varying numbers of inputs. Our model is both more robust to changes in the number of input views, while also being more scalable in the number of input views, with a linear rather than quadratic complexity attention module (see Sec. B.1). We hypothesize that this is because of the local, transformation-invariant design of LVT. By operating on *local* neighborhoods using only *relative* poses, any spatial reasoning learned is modular and independent of any global camera path. Relatedly, we hypothesize that the use

Number of Inputs	Complexity	PSNR $\uparrow$				SSIM $\uparrow$				LPIPS $\downarrow$			
		12	16	20	32	12	16	20	32	12	16	20	32
Global Attention	$O(N^2)$	25.24	25.01	24.39	22.83	0.837	0.830	0.817	0.779	0.133	0.141	0.152	0.182
Local Attention	$O(N)$	25.32	25.12	24.69	23.31	0.846	0.842	0.834	0.807	0.130	0.135	0.141	0.167

Table 2. **Comparison of local attention and global attention.** We train both models on 16 input views at 480×270 resolution. During inference, we change the number of input views and observe that LVT is more robust to differences in sequence length, while also exhibiting linear scaling behavior.

Ablation	Mixed-Res. Train	Spatial Neighbor	Tokenize Cond.	Attention Cond.	SH	σ Reg.	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Initial Config.	$\times$	$\times$	Intrinsics	Relative Pose	$\times$	$\times$	22.02	0.774	0.196
Mixed-res. Train	$\checkmark$	$\times$	Intrinsics	Relative Pose	$\times$	$\times$	22.57	0.786	0.186
Add Spatial	$\times$	$\checkmark$	Intrinsics	Relative Pose	$\times$	$\times$	22.40	0.783	0.190
Conditioning Variants	$\times$	$\times$	Plücker	None	$\times$	$\times$	20.59	0.696	0.260
	$\times$	$\times$	Intrinsics	World Pose	$\times$	$\times$	21.44	0.745	0.220
	$\times$	$\times$	Intrinsics	Relative Plücker	$\times$	$\times$	21.81	0.767	0.202
SH Variants	$\times$	$\times$	Intrinsics	Relative Pose	$\checkmark$	$\times$	23.89	0.828	0.152
	$\times$	$\times$	Intrinsics	Relative Pose	$\checkmark$	$\checkmark$	<u>24.69</u>	<u>0.844</u>	<u>0.136</u>
Full Model	$\checkmark$	$\checkmark$	Intrinsics	Relative Pose	$\checkmark$	$\checkmark$	27.156	0.882	0.102

Table 3. **Design ablations.** We train all configurations at 480×270 resolution. Our full model adds mixed-resolution training and regularized spherical harmonics to improve over our initial configuration, whereas we find that alternative conditioning variations lead to sub-optimal metrics.

of relative rather than absolute pose in our network simplifies the network’s learning task. Using absolute poses requires learning over the combinatorial space of all key and query poses, whereas our approach constrains this problem by learning over the much smaller space of relative transformations between two nearby views. However, we note that this version of our model is not fully invariant to the number of inputs - empirically, we find that while rendering to nearby cameras is accurate, incorrectly-placed distant Gaussians can end up interfering with other cameras in the sequence, so performance still degrades as sequence length increases. This motivates our integration of multi-resolution and multi-length dataset mixing, which enables the accurate learning of 3D scene geometry from long sequences without requiring full-resolution training across all sequence lengths.

Mixed-resolution training. Our evaluations show that training with a mixed-resolution strategy leads to superior performance when evaluated across entire input sequences (whose lengths are scene-dependent), compared to models trained without this strategy. While both models train on 16 inputs at 480×270 resolution the mixed-resolution model is also trained with longer sequences at lower 240×135 resolution. This enables it to develop a more comprehensive understanding of scene structure, that it then can transfer to higher resolutions. We report a qualitative comparison of the two configurations in the supplementary (Fig. 6).

Spatial neighbor selection. Since our training dataset consists of ordered image sequences, a simple way to construct the neighborhood graph is to take the images in sequential order, which we use for our initial configuration. However, we find that a spatial neighbor selection strategy, in which the neighboring images are selected based on camera center distance, outperforms the simple sequential image order strategy, and has the additional benefit of being agnostic to the order that input images are taken.

Relative transformation conditioning. We conduct ablations on the method used to encode the relative transformations between the central view and its neighboring views within LVT. Table 3 compares different conditioning strategies. We find that using Plücker rays in the tokenization step, which encodes both camera intrinsics and extrinsics, but removing the conditioning to the attention module and thus removing information about relative transformations, degrades metrics. Encoding the local rays only during tokenization but directly using camera extrinsics relative to a world frame performs better than the previous variant but is still sub-optimal to our base model, which explicitly computes the relative transformations. Interestingly, we find that a straightforward, although more computationally expensive, encoding of the relative Plücker rays, e.g. applying $P_j P_{j'}^{-1} K_{j'}^{-1}$ to the pixel grid and adding to the attention conditioning, performs slightly worse than our base model, which directly encodes the relative quaternion and translation once per view and broadcasts that encoding across view tokens. These experiments confirm the benefits of transformation-invariant relative embeddings, enabling generalization from shorter training sequences to longer inference sequences.

View-dependent splat parameters. As observed in Table 1, adding spherical harmonics on the color and opacity improves results over a model without view-dependent effects. We also ablate our opacity regularizer, which encourages the computed opacity from randomly sampled viewing directions to be zero. Adding this regularizer improves metrics, as the model without this regularization is prone to producing spurious opacity when seen from different viewing directions at inference time. Our best model combines these design choices, achieving the best metrics among the ablation variations.

4.4 Limitations

Our method achieves state-of-the-art performance on multiple benchmarks, and generalizes well to very long sequences with high fidelity, with inference time scaling linearly with the input sequence length. However, it suffers from some limitations, which present an opportunity for future investigation.

When input sequences contain many redundant, overlapping views, the model’s inherently local attention mechanism may restrict communication between more distant, yet potentially relevant, views. Selecting a more spatially diverse subset of these input views, thereby increasing the typical separation between selected viewpoints, could potentially alleviate this issue. More complex neighbor selection strategies, that included both nearby and distant neighbors, could also allow for longer range communication. Alternatively, a hybrid approach could be explored, where LVT layers are interleaved with standard transformer layers.

Our model’s per-pixel splat generation results in a dense, redundant representation when processing many high-resolution input views and we find that splat rasterization can become a significant bottleneck. Our use of view dependent opacity complicates the simple approach of culling almost-transparent splats, however, several previous works [Hanson et al. 2025; Milef et al. 2025; Niemeyer et al. 2025], have introduced methods to efficiently render large 3DGS scene that could likely be applied to our model’s outputs.

Finally, the use of view-dependent opacity can occasionally produce popping artifacts in areas with extreme textures (Supp. Fig 12). Increasing the regularization weight α can mitigate such artifacts, but this often comes at the cost of reduced rendering quality, indicating a trade-off.

5 Conclusion

We introduce the Local View Transformer (LVT), a new architecture for large-scale scene reconstruction and novel view synthesis. LVT’s local attention mechanism allows scaling to many more input views than prior transformer-based methods. Additionally, through the use of a relative, rather than global, camera pose conditioning mechanism, LVT improves generalization across long and varied camera trajectories. We further propose to decode to 3D Gaussian splats with both view-dependent color *and* opacity, and demonstrate that this extension improves fidelity on thin structures and complex lighting effects.

Our experiments show that LVT achieves state-of-the-art results on a variety of benchmarks and demonstrates strong generalization capabilities. Furthermore, we perform comprehensive ablation studies confirming the efficacy of LVT’s local attention, robustness to varying input sequence lengths, and benefits from mixed-resolution training, relative transformation conditioning, and view-dependent opacity with regularization. In summary, LVT offers an effective, scalable, and generalizable solution for high-fidelity, large-scale 3D scene reconstruction in a single forward pass.

Acknowledgments

The authors would like to thank Stephen Lombardi, Clement Godard, Tiancheng Sun, and Yiming Wang for their invaluable feedback and discussions throughout the course of this research. We are

also thankful to Peter Hedman, Daniel Duckworth, Ryan Overbeck, and Jason Lawrence for their valuable feedback and guidance for preparing this manuscript. Tooba Imtiaz was partly supported by NIH Graduate Research Fellowship under Grant No. 5U24CA264369-03.

References

- Sameer Agarwal, Noah Snavely, Ian Simon, Steven M. Seitz, and Richard Szeliski. 2009. Building Rome in a Day. In *Int. Conf. Comput. Vis.* 72–79.
- Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. 2020. Neural point-based graphics. In *Eur. Conf. Comput. Vis.*
- Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. 2022. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *Eur. Conf. Comput. Vis.* Springer, 213–229.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. Emerging properties in self-supervised vision transformers. In *Int. Conf. Comput. Vis.* 9650–9660.
- David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. 2024. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.* 19457–19467.
- Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. 2013. Depth synthesis and local warps for plausible image-based navigation. *ACM Trans. Graph.* 32, 3 (2013), 1–12.
- Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. 2021. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Int. Conf. Comput. Vis.* 14124–14133.
- Yuedong Chen, Haoifei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. 2024. MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images. In *Eur. Conf. Comput. Vis.*
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. In *Adv. Neural Inform. Process. Syst.*, Vol. 35. 16344–16359.
- Tri Dao and Albert Gu. 2024a. Transformers are SSMS: Generalized Models and Efficient Algorithms Through Structured State Space Duality. In *Int. Conf. Machine Learning.*
- Tri Dao and Albert Gu. 2024b. Transformers are SSMS: Generalized Models and Efficient Algorithms Through Structured State Space Duality. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 235)*. PMLR, 10041–10071. <https://proceedings.mlr.press/v235/dao24a.html>
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Minneapolis, Minnesota, 4171–4186.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Int. Conf. Learn. Represent.*
- Yilun Du, Cameron Smith, Ayush Tewari, and Vincent Sitzmann. 2023. Learning to render novel views from wide-baseline stereo pairs. In *IEEE Conf. Comput. Vis. Pattern Recog.* 4970–4980.
- John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. 2019. Deepview: View synthesis with learned gradient descent. In *IEEE Conf. Comput. Vis. Pattern Recog.* 2367–2376.
- John Flynn, Michael Broxton, Lukas Murmann, Lucy Chai, Matthew DuVall, Clément Godard, Kathryn Heal, Srinivas Kaza, Stephen Lombardi, Xuan Luo, et al. 2024. Quark: Real-time, High-resolution, and General Neural View Synthesis. *ACM Trans. Graph.* 43, 6 (2024), 1–20.
- John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. 2016. Deepstereo: Learning to predict new views from the world’s imagery. In *IEEE Conf. Comput. Vis. Pattern Recog.* 5515–5524.
- Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. 2022. Plenoxels: Radiance fields without neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.* 5501–5510.
- Alex Hanson, Allen Tu, Geng Lin, Vasu Singla, Matthias Zwicker, and Tom Goldstein. 2025. Speedy-Splat: Fast 3D Gaussian Splatting with Sparse Pixels and Sparse Primitives. In *IEEE Conf. Comput. Vis. Pattern Recog.*
- Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX microflake distribution. *ACM Trans. Graph.* 34, 4 (2015), 1–11.
- Byeongho Heo, Song Park, Dongyoon Han, and Sangdoo Yun. 2024. Rotary position embedding for vision transformer. In *Eur. Conf. Comput. Vis.* Springer, 289–305.
- Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. 2024. Lrm: Large reconstruction model for

- single image to 3d. *Int. Conf. Learn. Represent.* (2024).
- Haian Jin, Hanwen Jiang, Hao Tan, Kai Zhang, Sai Bi, Tianyuan Zhang, Fujun Luan, Noah Snavely, and Zexiang Xu. 2025. LVSM: A Large View Synthesis Model with Minimal 3D Inductive Bias. In *Int. Conf. Learn. Represent.*
- Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. 2022. Geonerf: Generalizing nerf with geometry priors. In *IEEE Conf. Comput. Vis. Pattern Recog.* 18365–18375.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* 42, 4 (2023), 139–1.
- Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. 2017. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Trans. Graph.* 36, 4 (2017), 1–13.
- Xin Kong, Shikun Liu, Xiaoyang Lyu, Marwan Taher, Xiaojuan Qi, and Andrew J Davison. 2024. Eschernet: A generative model for scalable view synthesis. In *IEEE Conf. Comput. Vis. Pattern Recog.* 9503–9513.
- Kiriakos N. Kutulakos and Steven M. Seitz. 2000. A Theory of Shape by Space Carving. *Int. J. Comput. Vis.* 38, 3 (2000), 199–218. <https://doi.org/10.1023/A:1008191222954>
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. 2023. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. *arXiv preprint arXiv:2311.06214* (2023).
- Ruilong Li, Brent Yi, Junchen Liu, Hang Gao, Yi Ma, and Angjoo Kanazawa. 2025. Cameras as Relative Positional Encoding. *arXiv preprint arXiv:2507.10496* (2025).
- Kai-En Lin, Zexiang Xu, Ben Mildenhall, Pratul P Srinivasan, Yannick Hold-Geoffroy, Stephen DiVerdi, Qi Sun, Kalyan Sunkavalli, and Ravi Ramamoorthi. 2020. Deep multi depth panoramas for view synthesis. In *Eur. Conf. Comput. Vis.* Springer, 328–344.
- Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. 2024. DL3dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In *IEEE Conf. Comput. Vis. Pattern Recog.* 22160–22169.
- Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. 2020. Neural sparse voxel fields. In *Adv. Neural Inform. Process. Syst.*, Vol. 33. 15651–15663.
- Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. 2019. Local Light Field Fusion: Practical View Synthesis with Prescriptive Sampling Guidelines. *ACM Trans. Graph.* (2019).
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2020. Nerf: Representing scenes as neural radiance fields for view synthesis. In *Eur. Conf. Comput. Vis.*
- N. Milef, D. Seyb, T. Keeler, T. Nguyen-Phuoc, A. Božić, S. Kondguli, and C. Marshall. 2025. Learning Fast 3D Gaussian Splatting Rendering using Continuous Level of Detail. *Computer Graphics Forum* (04 2025). <https://doi.org/10.1111/cgf.70069>
- Takeru Miyato, Bernhard Jaeger, Max Welling, and Andreas Geiger. 2024. Gta: A geometry-aware attention mechanism for multi-view transformers. In *Int. Conf. Learn. Represent.*
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* 41, 4 (2022), 1–15.
- Michael Niemeyer, Fabian Manhardt, Marie-Julie Rakotosaona, Michael Oechsle, Daniel Duckworth, Rama Gosula, Keisuke Tateno, John Bates, Dominik Kaeser, and Federico Tombari. 2025. RadSplat: Radiance Field-Informed Gaussian Splatting for Robust Real-Time Rendering with 900+ FPS. In *Int. Conf. on 3D Vision*.
- Mateusz Nowak, Wojciech Jarosz, and Peter Chin. 2025. VoD-3DGS: View-opacity-Dependent 3D Gaussian Splatting. *arXiv preprint arXiv:2501.17978* (2025).
- Eric Penner and Li Zhang. 2017. Soft 3d reconstruction for view synthesis. *ACM Trans. Graph.* 36, 6 (2017), 1–11.
- Julius Plücker. 1865. Xvii. on a new geometry of space. *Philosophical Transactions of the Royal Society of London* 155 (1865), 725–791.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. (2019).
- Ravi Ramamoorthi and Pat Hanrahan. 2001. An efficient representation for irradiance environment maps. In *Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. Association for Computing Machinery, New York, NY, USA, 497–500. <https://doi.org/10.1145/383259.383317>
- René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. 2021. Vision transformers for dense prediction. In *Int. Conf. Comput. Vis.* 12179–12188.
- Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. 2016. Pixelwise view selection for unstructured multi-view stereo. In *Eur. Conf. Comput. Vis.* Springer, 501–518.
- Steven M. Seitz and Charles R. Dyer. 1999. Photorealistic Scene Reconstruction by Voxel Coloring. *Int. J. Comput. Vis.* 35, 2 (1999), 151–173. <https://doi.org/10.1023/A:1008176507526>
- Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- Peter-Pike Sloan, Jan Kautz, and John Snyder. 2023. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 339–348.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* 568 (2024), 127063.
- Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. 2022. Generalizable patch-based neural rendering. In *Eur. Conf. Comput. Vis.* Springer, 156–174.
- Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. 2024. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. In *Eur. Conf. Comput. Vis.* Springer, 1–18.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Adv. Neural Inform. Process. Syst.* 30 (2017).
- Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. 2022. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *IEEE Conf. Comput. Vis. Pattern Recog.* IEEE, 5481–5490.
- Hengyi Wang and Lourdes Agapito. 2025. 3D Reconstruction with Spatial Memory. In *Int. Conf. on 3D Vision*.
- Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexiang Xu, and Kai Zhang. 2024b. Pf-lrm: Pose-free large reconstruction model for joint pose and shape prediction. In *Int. Conf. Learn. Represent.*
- Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. 2021. Ibrnet: Learning multi-view image-based rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.* 4690–4699.
- Qianqian Wang*, Yifei Zhang*, Aleksander Holynski, Alexei A. Efros, and Angjoo Kanazawa. 2025. Continuous 3D Perception Model with Persistent State. In *IEEE Conf. Comput. Vis. Pattern Recog.*
- Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. 2024a. Dust3r: Geometric 3d vision made easy. In *IEEE Conf. Comput. Vis. Pattern Recog.* 20697–20709.
- Suttisak Wizatwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. 2021. Nex: Real-time view synthesis with neural basis expansion. In *IEEE Conf. Comput. Vis. Pattern Recog.* 8534–8543.
- Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. 2020. On layer normalization in the transformer architecture. In *Int. Conf. Machine Learning (ICML'20)*. JMLR.org, Article 975, 10 pages.
- Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. 2025. DepthSplat: Connecting Gaussian Splatting and Depth. In *IEEE Conf. Comput. Vis. Pattern Recog.*
- Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, et al. 2023. Dmv3d: Denoising multi-view diffusion using 3d large reconstruction model. *arXiv preprint arXiv:2311.09217* (2023).
- Jianing Yang, Alexander Sax, Kevin J Liang, Mikael Henaff, Hao Tang, Ang Cao, Joyce Chai, Franziska Meier, and Matt Feiszli. 2025. Fast3R: Towards 3D Reconstruction of 1000+ Images in One Forward Pass. In *IEEE Conf. Comput. Vis. Pattern Recog.*
- Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. 2019. Differentiable surface splatting for point-based geometry processing. *ACM Trans. Graph.* 38, 6 (2019), 1–14.
- Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. 2021a. Plenotrees for real-time rendering of neural radiance fields. In *Int. Conf. Comput. Vis.* 5752–5761.
- Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. 2021b. pixelnerf: Neural radiance fields from one or few images. In *IEEE Conf. Comput. Vis. Pattern Recog.* 4578–4587.
- Zehao Yu, Torsten Sattler, and Andreas Geiger. 2024. Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Trans. Graph.* 43, 6 (2024), 1–13.
- Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. 2024. Gs-lrm: Large reconstruction model for 3d gaussian splatting. In *Eur. Conf. Comput. Vis.* Springer, 1–19.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conf. Comput. Vis. Pattern Recog.* 586–595.
- Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. 2018. Stereo magnification: Learning view synthesis using multiplane images. *ACM Trans. Graph.* (2018).
- Chen Zhiwen, Hao Tan, Kai Zhang, Sai Bi, Fujun Luan, Yicong Hong, Li Fuxin, and Zexiang Xu. 2025. Long-LRM: Long-sequence Large Reconstruction Model for Wide-coverage Gaussian Splats. In *Int. Conf. Comput. Vis.*



Fig. 3. **Qualitative comparison of LVT and baselines on DL3DV.** We observe that 3DGS is prone to floater artifacts in reconstructed scenes. Long-LRM is less prone to floaters, but struggles to reconstruct thin structures (middle example) and specularities. LVT achieves high-fidelity reconstructions, and reconstructs specular reflections and thin structures particularly well when view-dependent opacity is introduced.




Fig. 4. **Comparison of LVT models (zero-shot) on Tanks&Temples dataset.** LVT trained on DL3DV generalizes well when tested zero-shot, with the full model achieving high-fidelity detail and structure reconstruction. Qualitatively, our LVT_{SH-rgba} model preserves similar levels of detail as 3DGS, despite no per-scene optimization.




Fig. 5. **Comparison of LVT (zero-shot) on Mip-NeRF360 dataset.** The camera trajectories in Mip-NeRF360 differ from the drone flight sequences used for training, yet our model is able to generalize to this dataset as it aggregates information in local neighborhoods. We note that 3DGS requires several minutes of optimization while our model performs an inference pass in seconds.

A Additional Training Details

A.1 Dataset Mixing and Finetuning Strategy

We employ a dataset mixing strategy that randomly samples between longer sequences at lower resolution, and shorter sequences at higher resolution. This encourages the model to learn both high resolution details and scene priors on long scenes in a computationally feasible manner. We further detail our dataset mixing approach in Tab. 4. We initially train our model to a maximum resolution of 480×270 and then finetune to 960×540 . We show qualitative examples of the model output with and without this dataset mixing in Fig. 6. To manage the memory footprint during finetuning, we cull 40% and 20% of splats with the smallest opacity for LVT_{base} and $LVT_{SH-rgba}$ respectively. We found that LVT_{base} is more prone to running out of memory during training, so we additionally finetune the model while randomly cropping target images to size of 672×378 and empirically reduce the LPIPS loss weight to 0.02 to balance the two reconstruction losses with this change in target image size.

Model Resolution	Weight	Resolution	# Inputs	# Targets
480×270	0.125	240x135	40	24
	0.125	240x135	44	24
	0.125	240x135	48	24
	0.125	240x135	52	24
	0.5	480x270	16	8
960×540	0.125	240x135	40	24
	0.125	240x135	44	24
	0.125	240x135	48	16
	0.125	240x135	52	16
	0.25	480x270	24	2
	0.25	960x540	16	1

Table 4. **Dataset composition.** For mixed-resolution training, we train with longer sequences at lower resolution, and shorter sequences at higher resolution to learn both the necessary structure and details for scene reconstruction.




Fig. 6. **Ablation for mixed-resolution training on DL3DV.** Our model trained with the mixed-resolution training strategy surpasses the model trained with a single input resolution.




Fig. 7. **Feature fusion for 2K outputs.** To produce higher resolution outputs, we use a larger patch size and then fuse in features from the input images to mitigate quality loss due to the increased patch size.

A.2 Model Training Time

For DL3DV, pretraining takes 4 days at 480×270 resolution (with a batch size of 256, 1 per H100 GPU), followed by 2 days of finetuning at 960×540 . We believe further optimizations can reduce this computational cost.

A.3 Towards Higher Resolution with LVT-Flash

Our model’s computational cost scales linearly with the number of input images but quadratically with the number of pixels per image. This is because increasing the input resolution, while keeping the patch size constant, proportionally increases the number of tokens. For instance, doubling image width and height quadruples the tokens, leading to a 16-fold increase in transformer complexity, quickly becoming impractical at high resolutions.

While larger patch sizes (e.g., using 12×12 instead of 8×8 patches) would reduce token count, this approach compromises quality as individual tokens lack the capacity to adequately encode larger input and output patches. To mitigate this quality loss, we introduce a post-processing step. After decoding transformer tokens to pixels, we fuse them with features extracted directly from the input images (Fig. 7). This allows high-resolution information from the input to enhance the decoded features. We call this faster, reduced quality model *LVT-Flash*, and with it we can achieve 2k resolution in the transformer output.

B Additional Experiments

B.1 Model Inference Times

We compare the inference time of our model, with local attention windows, to a model with fully connected self attention, following the implementation of GS-LRM. Besides the tokenization and attention operations as described in Sec. 3, all other components and training of these two models are identical. We evaluate these two models, running in 16-bit floating point, at 480×270 resolution and benchmark on a single A100 40G GPU. In addition to finding that our LVT model is more robust to changes in sequence length at inference time, as reported in Tab. 2, the inference time our model




Fig. 8. **Runtime comparison of local and global attention.** We evaluate all models on a single A100 40G GPU at 480x270 resolution. Our LVT model uses local attention and scales roughly linearly ($O(N)$) with the number of input views N . We implement a model with full self-attention, following GS-LRM, which scales quadratically with the number of inputs ($O(N^2)$). This difference is also observed empirically in the model inference times.




Fig. 9. **Inference time comparison to baselines.** We evaluate all models on a single A100 40G GPU at 960x540 resolution. LVT_{SH-rgba} scales roughly linearly ($O(N)$) with the number of input views N . Long-LRM, despite a lower constant factor from its Mamba-attention hybrid, still includes global self-attention operation. The optimization time of 3DGS depends on the number of optimization steps.

also scales linearly with sequence length, whereas the self attention model exhibits dramatically longer inference time as sequence length increases (Fig. 8).

We also benchmark our full model trained at 540x960 resolution in a similar fashion. 3DGS requires optimization per scene, and its runtime is dominated by the number of optimization steps rather than the number of inputs. While our model is not as fast as the times reported by Long-LRM, we note their model is a hybrid of Mamba and transformer blocks, with the transformer block still operating at quadratic complexity. In contrast, our model uses transformer blocks only with local attention, and for model blocks with memory-intensive intermediate states, operations can be computed sequentially on a per-image basis rather than processing all images at once. This incremental approach effectively reduces memory consumption by trading off some computational speed. Again, we observe a linear trend between input sequence length and runtime

for our model, significantly faster than optimization-based 3DGS (Fig. 9).

B.2 RealEstate-10k

We also train a version of our model on RealEstate-10K [Zhou et al. 2018] to benchmark our model against GS-LRM [Zhang et al. 2024]. Following prior work, we train at 256×256 resolution. We use loss weights $\lambda = 0.25$ and $\alpha = 0$ and the same learning rate schedule as our DL3DV model.

Two-view model. We train a version of our LVT_{base} model on two input views. In this case the neighbor set only contains two items per view, itself and the other input view. While this is not the intended use-case of our model, we conduct this experiment to ensure that our conditioning and attention modules do not degrade performance compared to the standard transformer architecture, as used in GS-LRM. We use the standard training/test split (comprising $\sim 90\%/10\%$ clips respectively) and input and target indices as adopted in PixelSplat [Charatan et al. 2024]. We compare to previous generalizable methods [Charatan et al. 2024; Du et al. 2023; Suhail et al. 2022; Yu et al. 2021b; Zhang et al. 2024; Ziwen et al. 2025] in Tab. 5. Our model obtains competitive performance with GS-LRM and Long-LRM on this simplified setting, ensuring that our changes to the model do not hinder the transformer. We show qualitative results from our model in Figure 10.

Model	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
pixelNeRF	20.43	0.589	0.550
GPNR	24.11	0.793	0.255
Du et al.	24.78	0.820	0.213
pixelSplat	25.89	0.858	0.142
GS-LRM	28.10	0.892	0.114
Long-LRM (no token merge)	28.54	0.895	0.109
LVT _{base} (Ours)	28.75	0.915	0.072

Table 5. **RealEstate10K evaluations.** We compare LVT against baseline methods at a resolution of 256×256 . We use the evaluation split from pixelSplat [Charatan et al. 2024], which takes two images as input.

Sequence model. Additionally, to demonstrate the generalization ability of LVT to varying sequence lengths, we train our LVT_{base} model with a window size of three neighbors and a full self-attention model. The full self-attention model is based on GS-LRM, which concatenates world-space Plücker rays to the input images and uses standard transformer blocks. All other components of the models are identical, and we train both using 8 input views. As reported in Table 6 and Figure 11, we evaluate both models on scenes by feeding 4, 8, 12 and 16 input views. Notably, the performance difference between the two models widens as the discrepancy between the number of training and evaluation input views increases, showing that our model generalizes better to varied sequence lengths.

B.3 Additional Ablations

We investigate additional design decisions of our LVT architecture via ablation experiments, computed with batch size 64 at 480x270 resolution.



Fig. 10. **Qualitative results on RealEstate10K.** We show qualitative results of our LVT model trained on 2 input frames of RealEstate10K dataset at 256×256 resolution.

Model	# of inputs	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Full SA	4	28.02	0.929	0.064
	8	27.81	0.922	0.079
	12	28.15	0.932	0.071
	16	26.72	0.922	0.080
LVT _{base} (Ours)	4	31.67	0.951	0.046
	8	27.80	0.921	0.078
	12	29.00	0.932	0.070
	16	29.46	0.936	0.066

Table 6. **RealEstate10K evaluations with variable sequence lengths.** We train a full self-attention model, following GS-LRM and LVT on 8 input views, but evaluate on varying number to inputs to assess their generalization ability for changing input sequence lengths. We use a resolution of 256×256 throughout.

Window Size	Dilation	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
3	1	21.75	0.765	0.205
5	1	22.02	0.774	0.196
5	2	21.26	0.746	0.221
7	1	21.91	0.772	0.198

Table 7. **Ablation for neighborhood window sizes and dilations in the local view transformer.** We evaluate various window sizes with and without dilation in neighbor selection, where a dilation of 1 corresponds to no dilation/selecting directly adjacent neighbors. A neighborhood of window size 5 with no dilation is the best-performing variant.

Neighborhood size and dilation. On DL3DV, for 32 input frames at 480×270 resolution, our LVT model uses a neighborhood size of




Fig. 11. **Ablation for generalization to variable sequence lengths on RealEstate10K.** We train both models on 8 input frames and test. LVT effectively consolidates information from more input views.

five neighbors. Here, we investigate the impact of this parameter, as well as the dilation factor. A dilation of 1 indicates that only the N closest neighbors are attended to, whereas a dilation of 2 takes every other of the $2N$ closest neighbors. Results are shown in Tab. 7. We find that reducing the window size to 3 or increasing it to 7 with no dilation leads to a marginal decrease. Introducing a dilation factor of 2 with a window size of 5 further degrades performance, implying that considering non-adjacent neighbors within the local window is detrimental to the reconstruction quality.

Ablation	Relative SH	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
LVT _{base}	-	22.02	0.774	0.196
+ SH color	$\times$	20.08	0.707	0.243
	$\checkmark$	22.48	0.795	0.181
+ SH opacity	$\times$	19.56	0.676	0.291
	$\checkmark$	23.89	0.828	0.152

Table 8. **Ablation for view-dependent splat color and opacity.** Incorporating view dependence into both color and opacity consistently boosts reconstruction quality compared to the baseline. The SH need to be computed relative to the input cameras (instead of world space, which can be arbitrarily defined).

Model	Input Stride	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
LVT _{base}	6	23.18	0.810	0.221
	8	23.65	0.813	0.215
	10	23.67	0.807	0.218
LVT _{SH-rgba}	6	27.95	0.891	0.127
	8	27.64	0.883	0.133
	10	27.13	0.873	0.141

Table 9. **Varying input image selection.** We vary the stride of the input images, selecting every sixth, eighth, or tenth image in the sequence as model inputs after all target images have been removed. Our models are relatively robust to the input selection, which in turn changes the overall total number of input images per scene.




Fig. 12. **Limitations of view-dependent opacity.** Using the view-dependent opacity occasionally produces ‘popping’ artifacts in regions with extreme (highly dense or smooth) textures.

Moreover, with no dilation the inference times are 0.95s, 1.27s, and 1.62s for neighborhood window sizes 3, 5 and 7, respectively. As such, we default to a neighborhood window size of 5 without dilation for our main model.

View-dependent splat parameters. We present results on incorporating view dependence into the color and opacity predictions of our model in Tab. 8 and Fig. 13. We find that maintaining transformation invariance by predicting spherical harmonics coefficients relative to each input camera’s local coordinate system (Relative SH in Tab. 8) is critical. Without it, the model cannot generalize to different sequence lengths at inference time.

Adding spherical harmonics on color improves over our LVT_{base} model, and adding view-dependent opacity offers further improvements. However, we occasionally observe some failure modes in which the view-dependent opacity can cause ‘popping’ artifacts, as shown in Fig. 12.

Input Camera Selection. For our experiments, we default to selecting every eighth image in the full image sequence as input images to the model, after all target images have been removed. This varies the distance of the selected input images to the target cameras, without requiring more advanced heuristics that maintain a fixed sequence length. In Table 9, we test our 960x540 resolution models while varying the input selection stride at inference time, taking every sixth,

eighth, or tenth image as input. Our models are relatively robust to this selection, and we choose every eighth as a good middle ground.

Higher resolution with LVT-Flash. In Fig. 14 we show outputs from our LVT-Flash model. Qualitatively we find that using a feature fusion module can recover some of the quality loss incurred from increasing the patch size and remains efficient enough to render at 2K resolution compared to using a smaller patch size.




Fig. 13. **Impact of view dependence.** We show qualitative examples of our model adding spherical harmonics on color, and on both color and opacity. LVT_{SH-rgba} attains the best metrics and is better able to represent complex structures and lighting effects.




Fig. 14. **LVT-Flash - DL3DV-140 at 2K resolution.** LVT-Flash enables efficient processing of high-resolution images by employing larger patch sizes, which reduces the number of tokens to be processed. To counteract potential quality degradation from these larger patches, our proposed feature fusion module integrates high-resolution details from the input image directly with the decoded transformer tokens, thereby preserving output detail and enabling high-resolution results.