

# FlashVGGT: Efficient and Scalable Visual Geometry Transformers with Compressed Descriptor Attention

Zipeng Wang    Dan Xu\*  
The Hong Kong University of Science and Technology  
zwang253@connect.ust.hk    danxu@cse.ust.hk

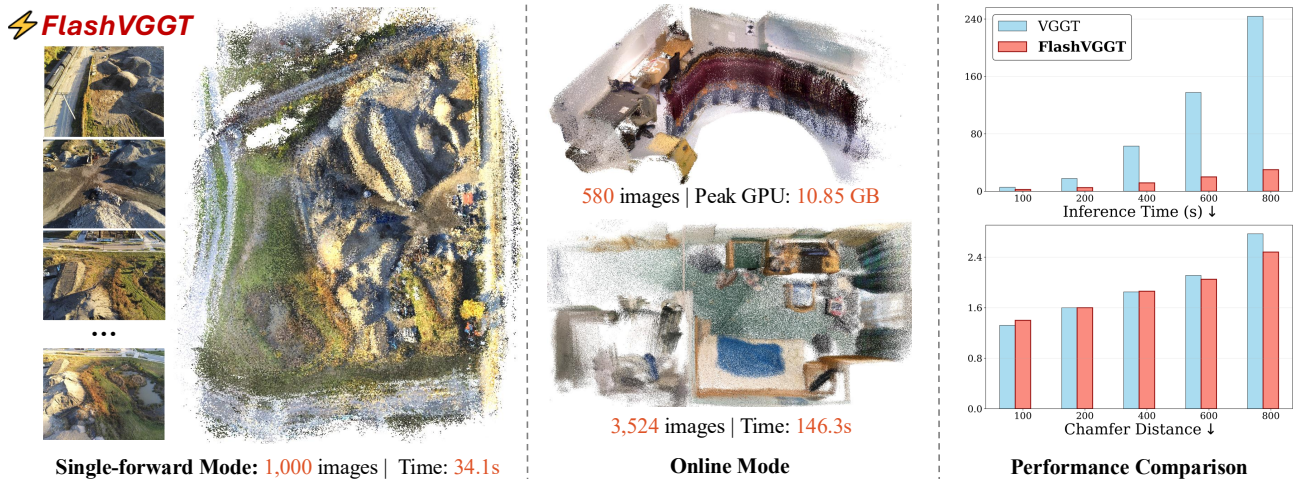


Figure 1. **FlashVGGT** achieves significant speedup and scales to larger inputs. Our method enables both fast single-forward inference on long sequences (left) and memory-efficient online inference (center) while maintaining competitive accuracy versus VGGT [58] (right).

## Abstract

3D reconstruction from multi-view images is a core challenge in computer vision. Recently, feed-forward methods have emerged as efficient and robust alternatives to traditional per-scene optimization techniques. Among them, state-of-the-art models like the Visual Geometry Grounding Transformer (VGGT) leverage full self-attention over all image tokens to capture global relationships. However, this approach suffers from poor scalability due to the quadratic complexity of self-attention and the large number of tokens generated in long image sequences. In this work, we introduce *FlashVGGT*, an efficient alternative that addresses this bottleneck through a descriptor-based attention mechanism. Instead of applying dense global attention across all tokens, *FlashVGGT* compresses spatial information from each frame into a compact set of **descriptor tokens**. Global attention is then computed as cross-attention between the full set of image tokens and this smaller descriptor set, significantly reducing computational overhead. Moreover, the compactness of the descriptors enables online inference

over long sequences via a chunk-recursive mechanism that reuses cached descriptors from previous chunks. Experimental results show that *FlashVGGT* achieves reconstruction accuracy competitive with VGGT while reducing inference time to just 9.3% of VGGT for 1,000 images, and scaling efficiently to sequences exceeding 3,000 images. Our project page is available at [https://wzpsscott.github.io/flashvggt\\_page/](https://wzpsscott.github.io/flashvggt_page/).

## 1. Introduction

Reconstructing 3D geometry from multi-view images is a fundamental problem in computer vision [21]. Traditional pipelines such as Structure-from-Motion (SfM) [26, 42, 46, 52] and Multi-View Stereo (MVS) [38, 47, 48] have dominated this task for decades. These methods rely on per-scene, iterative optimization pipelines that include feature detection, matching, triangulation, and bundle adjustment. While often accurate, these pipelines are computationally intensive and fragile, requiring extensive processing and careful tuning for each scene, especially under challenging

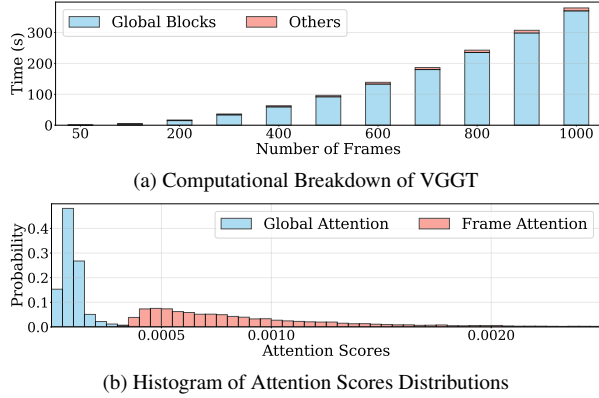


Figure 2. (a) The global attention block is the primary computational bottleneck in VGGT, dominating total inference time. (b) Global attention is highly sparse, with most scores concentrated near zero, suggesting that full self-attention is highly inefficient. In contrast, frame attention exhibits a more uniform distribution.

conditions [36, 45].

Recent advances have shifted toward learning-based approaches that bypass the complexities of traditional pipelines by directly predicting 3D structure using neural networks [27, 57, 60, 64]. These models are trained end-to-end on large-scale datasets [3, 12, 22, 40, 44, 65], enabling direct 3D prediction from multi-view input. This feed-forward paradigm eliminates the need for sequential post-processing and per-scene optimization. Moreover, they exhibit improved robustness by learning strong priors from diverse, large-scale data.

A recent milestone in this direction is the Visual Geometry Grounding Transformer (VGGT) [58], which performs high-fidelity 3D reconstruction from hundreds of views in a single forward pass. VGGT’s success stems from an alternating attention backbone that combines frame-wise and global attention blocks [16, 54], enabling effective aggregation of frame and global context. However, this architecture has a key limitation: the global attention block requires self-attention over all image tokens. As shown in Fig. 2a, this leads to quadratic complexity, creating a severe bottleneck and limiting scalability of VGGT to long sequences.

This work is driven by a central question: *Is full self-attention truly necessary for global reasoning in VGGT?* We base our approach on two key insights. First, classical methods show that accurate inter-frame associations can be inferred from sparse keypoints and descriptors [4, 35], suggesting that dense token-to-token attention may be unnecessary. Second, we observe that VGGT’s global attention maps are inherently sparse, with most scores near zero (Fig. 2b), implying that much of the computation is spent on irrelevant token pairs. These observations motivate our quest for a more efficient alternative that retains global reasoning while scaling to long input sequences.

Motivated by these insights, we propose **FlashVGGT**, an efficient architecture that overcomes the computation bottleneck of VGGT through compressed descriptor attentions. Our core innovation reformulates the global attention block by generating a compact set of descriptor tokens via spatial resampling, which encapsulate key information from each frame. Global attention is then approximated via cross-attention from image tokens to descriptors (*i.e.*, using image tokens as queries and descriptors as keys/values). This design reduces the computational complexity of global attention from  $\mathcal{O}(N^2)$  to  $\mathcal{O}(N^2/r^2)$ , where  $r$  is the spatial compression ratio. Empirically, FlashVGGT achieves over 90% inference speedup on 1,000-image sequences, with accuracy comparable to VGGT.

Furthermore, the compactness of our descriptor tokens enables scalable inference for very long sequences (e.g., 3,000 images) through an chunk-recursive scheme. When processing sequences that exceed memory limits, we divide the input into sequential chunks. By caching and reusing descriptor tokens from previous chunks, later chunks incorporate historical context while maintaining a global receptive field across the entire sequence. Crucially, unlike StreamVGGT [69], which caches full-resolution tokens from all transformer layers, our method only stores the compressed descriptors. This achieves an  $r^2$  reduction in peak memory usage, enabling scalable reconstruction for substantially larger inputs and resource-constrained scenarios.

The main contributions of this work are summarized as follows: (i) We propose FlashVGGT, an efficient framework that alleviates the quadratic complexity of global attention in VGGT. (ii) We design a chunk-recursive inference mechanism that enables online reconstruction of long sequences using cached descriptors. (iii) Extensive experiments demonstrate that FlashVGGT achieves competitive accuracy while reducing inference time by over 90% on 1,000-image inputs and scales to over 3,000 views.

## 2. Related Work

**Feed-forward 3D Reconstruction.** An emerging paradigm in 3D reconstruction is to directly predict 3D structures from images using deep-learning models trained on large datasets. This paradigm, often referred to as feed-forward 3D reconstruction, replaces per-scene iterative optimization with one or several forward passes of a neural network, offering greater efficiency and robustness than traditional methods. Early efforts in this line of research [27, 60, 66] predict pairwise 3D point maps and employ per-scene global alignment to reconstruct from multiple views. Subsequent works [24, 61, 64] sought to directly predict 3D structures from multiple images in a single forward pass. Notably, VGGT [58] introduced a transformer-based architecture with alternating frame-wise and global attention blocks. This approach enables reconstruction from hun-

reds of images in a single forward pass with high accuracy. However, VGGT’s global attention module requires full self-attention over all image tokens, causing computational cost to scale nearly quadratically with the number of input images. This results in significant overhead for large-scale inputs. Another line of research [11, 15, 59] performs 3D reconstruction in an online manner. While these methods offer improved memory efficiency, their reconstruction accuracy often lags behind that of offline counterparts.

**Efficient Vision Transformers.** Vision Transformers [1, 16, 32] have become a pivotal component in modern computer vision, powered by the self-attention mechanism. However, a major limitation of Vision Transformers is the quadratic computational complexity of self-attention [13, 29, 51, 63], which leads to prohibitive overhead for large-scale inputs. Many efforts [13, 30, 31, 53] have been made to mitigate the computational demands of self-attention. One line of research introduces sparsity [10, 62] or low-rank approximations [20, 23] into the attention computation. Another direction focuses on reducing the number of tokens involved in the attention operation by clustering [37], merging [5, 6] or selection [18, 56]. Despite these advancements, most existing efficient Transformer methods are designed for 2D tasks like image classification. Consequently, it remains an open challenge to design efficient Vision Transformers for 3D reconstruction, as it demands maintaining long-range, multi-view geometric consistency.

**Concurrent Work.** Concurrently with our work, several other approaches have been proposed to address the efficiency bottleneck in VGGT. FastVGGT [49] employs token merging to reduce the number of tokens fed into global attention. However, the process of identifying and merging similar tokens across the entire sequence introduces considerable computational overhead. FasterVGGT [55] introduces block sparsity into the attention matrix to reduce computation. While effective for moderate sparsity levels, this approach suffers from significant performance degradation when sparsity increases. StreamVGGT [69] enables incremental reconstruction from streaming input by caching intermediate tokens from all global attention blocks. This design, however, leads to substantial memory overhead, limiting its scalability to longer sequences beyond tens of images. In contrast, FlashVGGT provides a more principled approach to global context compression, enabling a superior balance among inference speed, memory consumption and accuracy compared to these concurrent approaches.

### 3. Method

#### 3.1. Visual Geometry Grounding Transformers

We build our approach upon the Visual Geometry Grounding Transformer (VGGT) [58]. For completeness, we briefly recap its architecture, which consists of three main

stages: image encoding, feature aggregation via alternating attention, and 3D reconstruction heads.

**Image Encoder.** Given a sequence of  $S$  input images  $\{I_i\}_{i=1}^S$ , a DINO [9] encoder extracts a set of feature tokens for each image. This results in a set of token sequences  $\mathbf{F} = \{\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_S\}$ , where each  $\mathbf{F}_i \in \mathbb{R}^{N \times C}$  represents the  $N$  tokens from the  $i$ -th image. Each sequence includes a special learnable camera token that stores camera information and several register tokens [14].

**Alternating Attention.** The core of VGGT is a transformer aggregator composed of  $L$  identical layers, each containing two attention blocks designed to capture both intra- and inter-frame relationships. For simplicity, we omit the layer index in the following formulations.

- **Frame Attention** processes tokens *within* each frame independently. For each frame  $i$ , self-attention is computed over its  $N$  tokens to refine local features:

$$\mathbf{F}'_i = \text{SelfAttn}(\mathbf{F}_i), \quad \mathbf{F}_i \in \mathbb{R}^{N \times C} \quad (1)$$

- **Global Attention** models interactions *across all frames*. The tokens are concatenated into a single global sequence  $\mathbf{G} \in \mathbb{R}^{K \times C}$ , where  $K = S \times N$ . Standard self-attention is then computed over all  $K$  tokens:

$$\mathbf{H} = \text{SelfAttn}(\mathbf{G}) \quad (2)$$

The output  $\mathbf{H} \in \mathbb{R}^{K \times C}$  is then reshaped back into  $S$  frame sequences for subsequent processing.

**Reconstruction Heads.** The final aggregated token features are used by two separate heads to predict 3D properties for each view: (i) a **camera head** predicts camera extrinsics and intrinsics, and (ii) a **DPT head** [43] predicts a depth map and an aleatoric uncertainty map [25].

While this architecture is powerful, its scalability is limited by the quadratic complexity  $\mathcal{O}(K^2) = \mathcal{O}(S^2 N^2)$  of the dense global self-attention, as identified in Fig. 2a.

#### 3.2. Descriptor-Based Global Attention

We introduce an efficient alternative to the dense global attention block that preserves global reasoning while reducing its complexity. Our approach replaces standard self-attention in global attention blocks with a descriptor-based cross-attention mechanism, while keeping the encoder, frame attention, and reconstruction heads unchanged due to their minimal computational overhead. Fig. 3 provides an overview of our framework.

**Spatially-Compressed Descriptor Tokens.** Given the input to a global block,  $\mathbf{G} = \text{Reshape}(\{\mathbf{F}'_i\}_{i=1}^S) \in \mathbb{R}^{K \times C}$ , we first restore its spatial structure by reshaping it to  $\mathbf{G} \in \mathbb{R}^{S \times H \times W \times C}$ , where  $H$  and  $W$  are the height and width of the 2D patch token grid for each frame. We then generate a compact set of descriptor tokens  $\mathbf{D}$  by applying spatial compression. Compared to pooling-based methods, interpolation better preserves local spatial information in the



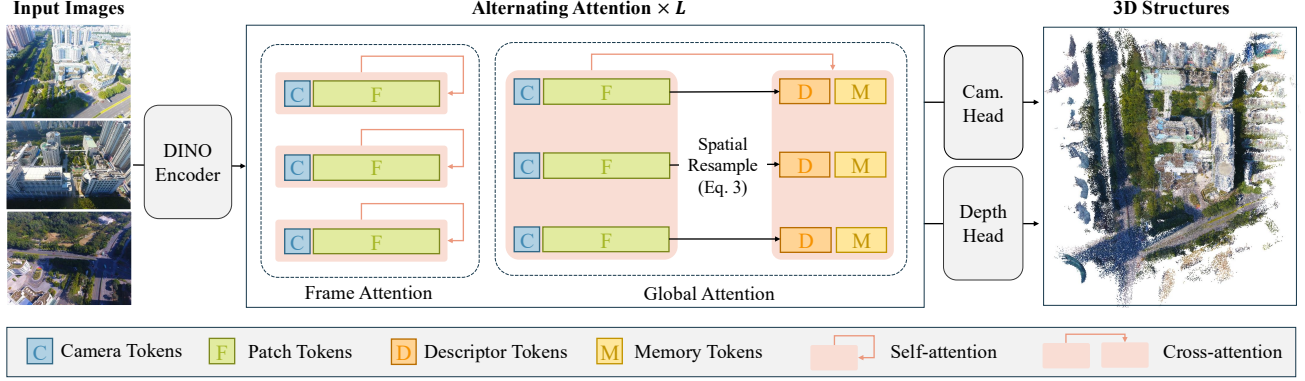


Figure 3. **Architecture Overview.** Our framework encodes input images into tokens using DINO [9] and processes them through alternating frame and global attention blocks. Unlike VGGT’s dense global attention over all tokens, FlashVGGT generates a compact set of descriptor tokens via spatial compression and computes efficient cross-attention from image tokens to these descriptors. The final aggregated tokens are fed into reconstruction heads to predict camera parameters and depth maps.

original features (see discussion in Sec. 4.4). We employ bilinear interpolation to resample each frame’s spatial dimensions  $(H, W)$  to a lower resolution  $(\lfloor H/r \rfloor, \lfloor W/r \rfloor)$ , where  $r$  is the compression factor:

$$\mathbf{D} = \text{Reshape}(\text{Interp}(\mathbf{G}, (\lfloor H/r \rfloor, \lfloor W/r \rfloor))) \quad (3)$$

The resulting descriptor tokens have a size of  $\mathbb{R}^{K_d \times C}$ , where  $K_d = S \times \lfloor H/r \rfloor \times \lfloor W/r \rfloor$ .

**Auxiliary Descriptor Tokens.** To maintain geometric consistency, we augment the compressed descriptors with three types of auxiliary tokens: (i) camera and register tokens from all frames; (ii) all tokens from the first image (which defines the world coordinate system) and (iii) all tokens from key-frames selected via k-means clustering [33] on average frame tokens. The key-frame selection is highly efficient, converging in under 2 seconds for 1,000 images on a single NVIDIA H800 GPU as it operates on per-frame averages rather than individual tokens. These auxiliary tokens act as geometric anchors, preserving high-fidelity information from camera parameters, the world coordinate frame, and representative views. This prevents the loss of critical details during descriptor compression, ensuring robust geometric reasoning across the entire sequence.

**Descriptor Attention.** We reformulate the global attention operation from Eq. (2) as a cross-attention layer. The original, full-resolution tokens  $\mathbf{G}$  are used as queries, while the descriptor tokens  $\mathbf{D}$  are used as the shared keys and values. This allows the full-resolution tokens to be updated by a compact set of descriptors representing the global context.

$$\mathbf{H} = \text{CrossAttn}(\mathbf{Q} = \mathbf{G}, \mathbf{KV} = \mathbf{D}). \quad (4)$$

Crucially, the operation maintains a *global receptive field*, preserving the model’s ability to capture long-range dependencies across all input images.

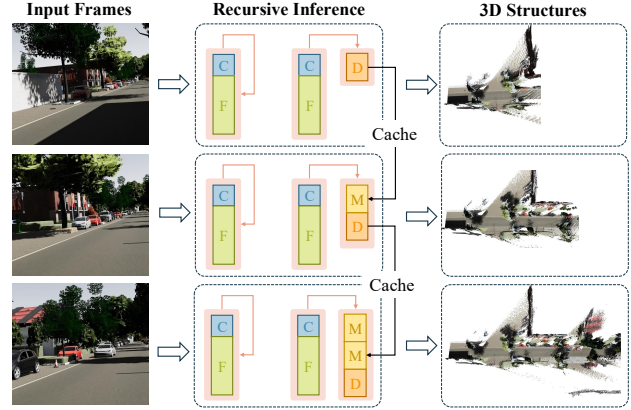


Figure 4. **Chunk Recursive Inference.** For long input sequences, we process them in a chunked manner while retaining global reception by caching descriptor tokens.

**Complexity.** Our design significantly reduces the computational complexity of the global block. The standard self-attention in Eq. (2) requires  $\mathcal{O}(K^2) = \mathcal{O}(S^2 N^2)$  operations. Our descriptor-based cross-attention reduces this to:  $\mathcal{O}(K \times K_d) = \mathcal{O}(S^2 N^2 / r^2)$ . With  $r = 4$  as in our experiments, the complexity reduction is about  $16 \times$ .

### 3.3. Chunk-Recursive Inference

To scale reconstruction to sequences that exceed GPU memory constraints, we propose a chunk-recursive inference scheme. This method processes long sequences sequentially while maintaining a global context across all previously seen chunks, as shown in Fig. 4.

**Problem Formulation.** Let the input sequence of  $S$  images be divided into  $T$  consecutive chunks,  $\{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_T\}$ . For the  $t$ -th chunk ( $1 \leq t \leq T$ ), let  $\mathbf{D}_t$  denote the descriptor tokens generated from chunk  $\mathcal{C}_t$  using the method described in Section 3.2. We maintain a set of memory tokens  $\mathbf{M}_t$  that

accumulates global information from all chunks processed up to step  $t$ , initialized as  $\mathbf{M}_0 = \emptyset$ .

**Descriptor Attention with Memory.** For chunk  $t$ , the global attention computation incorporates historical context through a memory mechanism that maintains information from all previously processed chunks. The queries remain the full-resolution image tokens  $\mathbf{G}_t$  from the current chunk  $\mathcal{C}_t$ . The keys and values are formed by concatenating the current chunk’s descriptors  $\mathbf{D}_t$  with the memory tokens from previous chunks  $\mathbf{M}_{t-1}$ :

$$\mathbf{H}_t = \text{CrossAttn}(\mathbf{Q} = \mathbf{G}_t, \mathbf{KV} = [\mathbf{M}_{t-1}, \mathbf{D}_t]), \quad (5)$$

where  $[\cdot, \cdot]$  denotes the concatenation operation along the sequence dimension. This design enables each token in the current chunk to attend to both the locally compressed context ( $\mathbf{D}_t$ ) and the globally accumulated history ( $\mathbf{M}_{t-1}$ ), effectively maintaining a global receptive field across the entire sequence while operating on individual chunks.

**Memory Update.** After processing chunk  $t$ , the memory is updated by appending the current chunk’s descriptors. To constrain memory growth during long sequences, we implement a dropping mechanism that retains only the descriptor tokens from every  $p$ -th frame. Let  $\mathbf{D}_t^{\text{retain}} = \mathbf{D}_t[:, p]$  denote the subset of descriptors from every  $p$ -th frame within the current chunk. The memory is then updated as:

$$\mathbf{M}_t = [\mathbf{M}_{t-1}, \mathbf{D}_t^{\text{retain}}] \quad (6)$$

This selective update rule ensures the memory  $\mathbf{M}_t$  compactly represents the entire sequence history while limiting its size to grow sublinearly with the number of frames.

**Complexity.** Our chunk-recursive scheme achieves substantial memory efficiency gains over the naive KV-caching in StreamVGGT [69]. While StreamVGGT’s memory usage is  $\mathcal{O}(KL)$ , where  $L$  is the number of global attention blocks, our approach reduces this to  $\mathcal{O}(KL/(pr^2))$  through descriptor compression ( $r$ ) and memory dropping ( $p$ ).

### 3.4. Model Training

**Training Strategy.** Our training follows a two-stage curriculum. The first stage trains the model on 2-24 randomly shuffled views following VGGT’s procedure. The second stage fine-tunes the model on ordered sequences to enable chunk-recursive inference, applying a causal mask to the global attention block that restricts each image to attend only to previous frames in the sequence. Unless otherwise specified, we use a spatial compression ratio of  $r = 4$ , a memory drop ratio of  $p = 5$ , and select a key frame every 200 images across all experiments.

**Training Data.** We train our model on seven datasets, which is a subset of VGGT’s training data, covering diverse scenarios including synthetic/real-world data, scene-level/object-centric configurations, and indoor/outdoor environments. Specifically, we use BlendedMVS [65],

CO3Dv2 [44], ScanNet [12], Mapillary [40], Arkitscenes [3], MVSSynth [22], and VirtualKitti [8].

## 4. Experiment

We evaluate FlashVGGT against state-of-the-art methods across three core tasks: (i) monocular and sparse-view reconstruction (Sec. 4.1), long-sequence dense 3D reconstruction (Sec. 4.2), and (iii) online dense 3D reconstruction (Sec. 4.3). We subsequently analyze the effectiveness of our key design choices in Sec. 4.4. All evaluations are conducted on a single NVIDIA H800 GPU.

### 4.1. Monocular and Sparse Reconstruction

**Camera Pose Estimation.** Following [58], we evaluate our method on the CO3Dv2 [44] and RealEstate10K [68] datasets for camera pose estimation on short sequences. For each scene, we randomly select 10 images and report standard metrics: RRA (Relative Rotation Accuracy) and RTA (Relative Translation Accuracy), and AUC, the area under the accuracy-threshold curve for the minimum of RRA and RTA. As shown in Tab. 1, FlashVGGT achieves highly competitive performance. On the out-of-distribution RealEstate10K dataset, our method closely matching VGGT’s metrics, while being better than the concurrent work FastVGGT [49]. On the CO3Dv2 dataset, FlashVGGT also significantly outperforms all other efficient methods like FastVGGT [49] and CUT3R [59], and remains within a narrow margin of the original VGGT.

**Monocular Depth Estimation.** We evaluate single-image depth prediction on the Sintel [7], Bonn [41], and NYU-v2 [39] datasets, reporting standard metrics: Absolute Relative Error (Abs Rel) and accuracy under threshold  $\tau < 1.25$ . As shown in Tab. 2, FlashVGGT demonstrates strong performance across all benchmarks. On Sintel, our method achieves competitive results that are comparable with VGGT and significantly outperform other efficient methods. For the Bonn and NYU-v2 datasets, FlashVGGT consistently ranks as the second-best method, closely following VGGT within a small gap while substantially outperforming FastVGGT and other concurrent approaches.

These results demonstrate that our efficient descriptor-based attention preserves the strong geometric reasoning capabilities of VGGT, while our subsequent experiments will show it does so at a fraction of the computational cost.

### 4.2. Long-sequence Dense 3D Reconstruction

This section presents our main results on long-sequence dense 3D reconstruction. We conduct a thorough evaluation of state-of-the-art offline methods capable of processing long sequences, comparing against Fast3R [64], VGGT [58], and FastVGGT [49]. For VGGT, we adopt the memory-efficient implementation from [49], which maintains the original accuracy while enabling inference on

Table 1. **Camera Pose Estimation on RealEstate10K [68] and Co3Dv2 [44].**

Method	RealEstate10K (unseen)			Co3Dv2		
	RRA@30↑	RTA@30↑	AUC@30↑	RRA@30↑	RTA@30↑	AUC@30↑
Fast3R [64]	99.05	81.86	61.68	97.49	91.11	73.43
CUT3R [59]	99.82	95.10	81.47	96.19	92.69	75.82
FLARE [67]	99.69	95.23	80.01	96.38	93.76	73.99
VGGT [58]	<b>99.97</b>	<b>96.22</b>	<b>85.32</b>	<b>98.96</b>	<b>97.13</b>	<b>88.59</b>
FastVGGT [49]	<u>99.92</u>	94.76	84.37	97.51	96.01	86.55
<b>FlashVGGT</b>	99.92	<u>95.61</u>	<u>85.30</u>	<u>98.23</u>	<u>96.75</u>	<u>86.88</u>

Table 2. **Monocular Depth Estimation on Sintel [7], Bonn [41] and NYU-v2 [39].**

Method	Sintel		Bonn		NYU-v2	
	Abs Rel↓	$\tau < 1.25 \uparrow$	Abs Rel↓	$\tau < 1.25 \uparrow$	Abs Rel↓	$\tau < 1.25 \uparrow$
Fast3R [64]	0.544	0.509	0.169	0.796	0.093	0.898
CUT3R [59]	0.418	0.520	0.058	0.967	0.081	0.914
FLARE [67]	0.606	0.402	0.130	0.836	0.089	0.898
VGGT [58]	<b>0.335</b>	<b>0.599</b>	<b>0.053</b>	<b>0.970</b>	<b>0.056</b>	<b>0.951</b>
FastVGGT [49]	<u>0.337</u>	0.582	0.056	0.952	0.058	0.943
<b>FlashVGGT</b>	0.346	<u>0.586</u>	<u>0.054</u>	<u>0.957</u>	<u>0.058</u>	<u>0.947</u>

Table 3. **Large-Scale Dense 3D Reconstruction.** Evaluation across 100, 500, and 1,000 image sequences, with results averaged over N-RGBD [2], 7-Scenes [50], and ScanNet [12]. Point cloud and camera pose metrics are multiplied by 100 for better readability.

Frames	Method	Depth		Point				Camera				Resource	
		Abs Rel↓	$\tau < 1.25 \uparrow$	Acc↓	Comp↓	CD↓	NC↑	APE↓	ARE↓	RPE-Trans↓	RPE-Rot↓	Time (s)↓	Mem. (GB)↓
100	Fast3R [64]	0.038	0.951	1.164	1.900	1.532	62.10	2.654	3.123	0.494	0.756	4.40	13.94
	VGGT [58]	0.029	0.983	0.962	1.162	1.062	<b>72.48</b>	<b>1.537</b>	2.935	<b>0.353</b>	<b>0.493</b>	4.93	12.26
	FastVGGT [49]	<u>0.029</u>	<u>0.984</u>	<u>0.988</u>	<b>1.092</b>	<b>1.040</b>	68.34	1.663	3.011	0.507	0.702	<u>2.74</u>	<u>12.68</u>
	<b>FlashVGGT</b>	<b>0.028</b>	<b>0.990</b>	<b>0.897</b>	<u>1.142</u>	1.019	<u>70.14</u>	<u>1.648</u>	<b>2.834</b>	<u>0.447</u>	<u>0.621</u>	<b>1.54</b>	<b>12.07</b>
500	Fast3R [64]	0.045	0.962	1.432	1.590	1.511	58.8	6.784	8.570	2.343	2.120	62.40	33.30
	VGGT [58]	0.035	0.967	1.484	<b>1.209</b>	1.347	<b>71.15</b>	<u>4.414</u>	<b>6.855</b>	<b>1.453</b>	<b>1.558</b>	90.97	37.22
	FastVGGT [49]	<u>0.034</u>	<u>0.967</u>	<u>1.388</u>	<u>1.241</u>	<u>1.314</u>	66.70	4.561	7.064	1.722	1.952	<u>29.04</u>	<u>39.33</u>
	<b>FlashVGGT</b>	<b>0.034</b>	<b>0.969</b>	<b>1.314</b>	1.283	<b>1.298</b>	<u>70.18</u>	<b>4.298</b>	<u>6.950</u>	<u>1.474</u>	<u>1.576</u>	<b>12.54</b>	<b>33.39</b>
1000	Fast3R [64]	0.122	0.855	3.076	1.457	2.267	52.5	12.67	22.36	9.530	11.34	224.10	61.95
	VGGT [58]	0.048	0.951	2.039	1.004	1.521	<u>68.65</u>	6.519	15.80	<u>2.222</u>	7.029	372.80	68.40
	FastVGGT [49]	<u>0.034</u>	<u>0.986</u>	<u>1.322</u>	<u>1.089</u>	<u>1.206</u>	66.05	<u>5.651</u>	<u>8.400</u>	2.553	<u>2.898</u>	<u>78.22</u>	<u>72.60</u>
	<b>FlashVGGT</b>	<b>0.032</b>	<b>0.991</b>	<b>1.160</b>	<b>1.096</b>	<b>1.128</b>	<b>69.63</b>	<b>5.237</b>	<b>8.242</b>	<b>2.067</b>	<b>2.802</b>	<b>35.32</b>	<b>60.74</b>

sequences of over 1,000 images. Following [49], we benchmark performance across totally 107 scenes from N-RGBD [2], 7-Scenes [50], and a subset of ScanNet [12] on 100, 500 and 1,000 images respectively. We evaluate each method comprehensively across 12 metrics covering depth estimation, point cloud reconstruction and camera pose estimation. For depth estimation, we report Absolute Relative Error (Abs Rel) and  $\tau < 1.25$  ratio. For point cloud reconstruction, we report Accuracy (Acc), Completeness (Comp), Chamfer Distance (CD), and Normal Consistency (NC). For camera estimation, we report Absolute Translation Error (APE), Absolute Rotation Error (ARE), Relative Translation Error (RPE-Trans), and Relative Rotation Error (RPE-Rot). Additionally, we report inference time and maximum GPU memory usage to provide a complete picture of each method’s practical utility.

The comprehensive results in Tab. 3 demonstrate FlashVGGT’s superior scalability and efficiency while maintaining competitive reconstruction quality across varying sequence lengths. For 100-image sequences, FlashVGGT achieves performance comparable to VGGT while being over  $3\times$  faster. At 500 images, it remains highly competitive with VGGT while achieving an over  $8\times$  speedup. For 1,000-image sequences, VGGT suffers from noticeable performance degradation due to attention dilution across excessive tokens, whereas FlashVGGT maintains high accuracy with over  $10\times$  faster inference.

Across all sequence lengths, FlashVGGT consistently outperforms other efficient methods like FastVGGT [49] and Fast3R [64], establishing a superior balance between efficiency and reconstruction fidelity. Fig. 5 presents a qualitative comparison of FlashVGGT against other methods. The results demonstrate that FlashVGGT produces more complete and robust reconstructions from long input sequences while achieving significantly faster inference. Notably, VGGT exhibits substantial performance degradation as sequence length increases, as seen in the room reconstruction from 1,000 images. We attribute this failure to noisy and redundant interactions over extremely long input (over 1M tokens for 1,000 images). In contrast, FlashVGGT avoids this pitfall by learning a compact, stable set of descriptor tokens that distill essential information, maintaining consistent performance across long sequences.

### 4.3. Online Dense 3D Reconstruction

We evaluate FlashVGGT in an online inference setting with a chunk size of 10 images against three recent online reconstruction methods: CUT3R [59], TTT3R [11], and StreamVGGT [69]. Experiments are conducted on N-RGBD [2] using 500-image sequences from each scene. As shown in Tab. 4, our method significantly outperforms previous approaches across all metrics. FlashVGGT achieves the best reconstruction quality while being over  $3.3\times$  faster than the fastest competitor CUT3R. Notably, we achieve



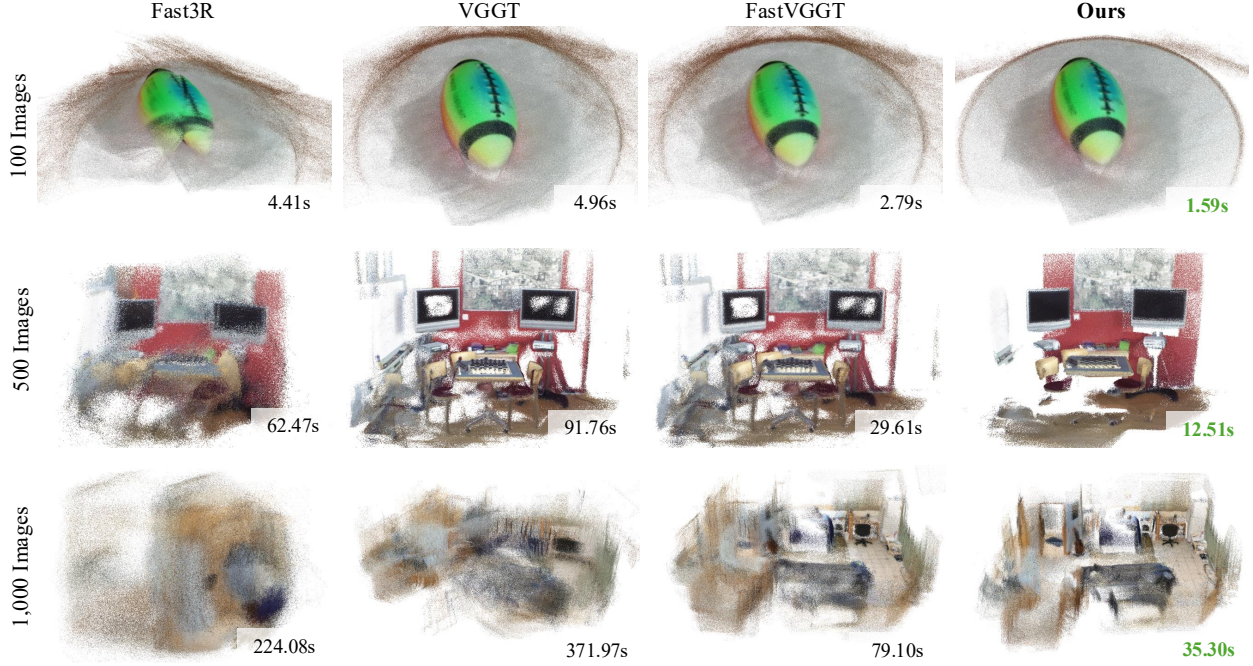


Figure 5. **Qualitative comparison of long-sequence dense 3D reconstruction.** FlashVGGT produces more robust reconstructions over long input sequences while being significantly faster. Only points with top 90% confidence are shown for better visualization.

Table 4. **Online Dense 3D Reconstruction on N-RGBD [2].**

Method	Abs Rel↓	Acc↓	Comp↓	APE↓	Time (s)	Mem (GB)
CUR3R [59]	0.375	4.890	3.426	23.456	34.19	<b>6.16</b>
TTT3R [11]	0.134	3.567	1.954	16.434	35.67	<b>6.16</b>
StreamVGGT [69]	0.086	2.456	1.235	6.543	209.50	70.70
<b>FlashVGGT</b>	<b>0.047</b>	<b>1.912</b>	<b>0.625</b>	<b>4.792</b>	<b>12.52</b>	<u>13.10</u>

this while using less than a quarter of the memory required by StreamVGGT. Qualitative results in Fig. 6 further demonstrate our approach’s superiority. FlashVGGT successfully reconstructs complete room geometry and fine details (e.g., tiny objects on the table), while CUT3R and TTT3R suffer from accumulated errors and fail to recover meaningful structures. Although StreamVGGT produces reasonable geometry, it requires over  $20\times$  more time and still fails to capture fine details effectively.

#### 4.4. Model Analysis and Discussion

**Spatial Compression Methods.** We evaluated five strategies for producing descriptor tokens: average pooling, top-k selection based on token norm, nearest-neighbor interpolation, bilinear interpolation, and a lightweight learnable compressor consisting of a depth-wise convolution followed by a point-wise linear layer (Table 5). Our analysis reveals that interpolation-based methods consistently outperform other approaches. We argue that their edge stems from locality preservation: as the DINO encoder outputs tokens corresponding to  $14\times 14$  pixel patches, aggressive aggregation methods such as pooling merge information from

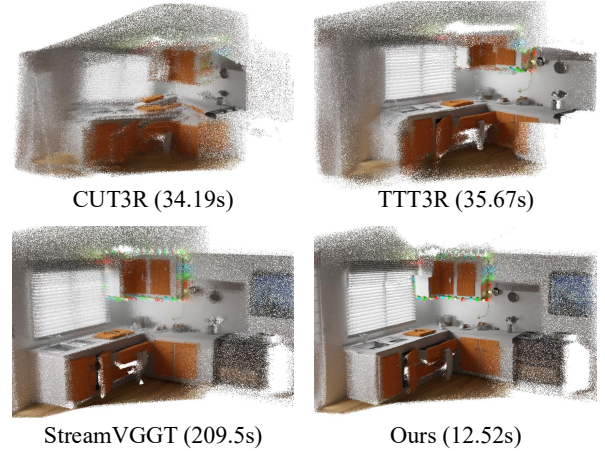


Figure 6. **Qualitative comparison of online 3D reconstruction.** All methods are evaluated on a sequence of 500 images.

many distant patches and wash out fine-grained cues. Interpolation, instead, blends only a handful of spatially adjacent tokens with distance-aware weights, retaining high-frequency detail that downstream tasks find useful. While top-k selection preserves original token values, it relies on the assumption that larger token norms correlate with informative local descriptors, which may not always hold. Finally, the learnable compressor neither improves quality nor stability; its limited capacity appears insufficient to capture the rich spatial patterns of the descriptors.

**Spatial Compression Ratio.** Fig. 7 illustrates the trade-off

Table 5. **Comparison of different spatial compression techniques.** Evaluated on N-RGBD [2] with 100 input images.

	Abs Rel↓	Acc↓	Comp↓	NC↑	APE↓	ARE↓
Pooling	0.019	0.560	0.301	75.68	2.256	4.008
Top-k	0.019	0.569	0.331	75.13	2.234	4.516
Learned	0.023	0.643	0.675	68.33	2.658	5.183
Nearest	0.014	0.441	0.273	76.96	1.902	3.456
Bilinear	<b>0.014</b>	<b>0.436</b>	<b>0.272</b>	<b>77.75</b>	<b>1.890</b>	<b>3.438</b>

between reconstruction accuracy (Chamfer Distance, left axis) and inference speed (right axis) across different compression ratios  $r$ . While larger  $r$  values yield faster inference, they also lead to a progressive decline in accuracy as fine-grained spatial information is lost. The ratio  $r = 4$  provides an optimal balance, offering a significant speedup with minimal loss in reconstruction quality. Beyond this point, the rate of performance degradation increases substantially for diminishing gains in speed.




Figure 7. **The impact of different compression ratios.** The ratio of 4 provides a balanced choice between quality and speed.

**Auxiliary Descriptor Tokens.** We identify that including the auxiliary descriptor tokens as described in Sec. 3.2 is crucial for maintaining reconstruction quality, as they augment the fine-grained information losses in compression. As shown in Fig. 8, omitting these auxiliary tokens degrades global geometric consistency. This effect is particularly pronounced in long sequences with low inter-frame overlap, *e.g.*, autonomous driving scenarios.




Figure 8. **The impact of auxiliary descriptor tokens.** Including auxiliary tokens improves geometric quality.

**Confidence Maps.** Fig. 9 shows a comparison of the confidence maps predicted by VGGT and our method. VGGT




Figure 9. **Analysis of confidence maps.** Ours method produces spatially coherent confidence scores that better preserve planar structures (*e.g.*, the computer screen) while filtering noise.

tends to produce over-confident predictions, assigning disproportionately low confidence scores to homogeneous regions such as walls and computer screens. This often results in gaps and holes in the final reconstruction after filtering out low-confidence points. In contrast, our method generates a more calibrated and spatially coherent confidence map. This allows for the effective preservation of structural details while robustly filtering out noise, leading to more complete and reliable 3D reconstructions.

## 5. Conclusion

In this work, we introduced FlashVGGT, a novel framework that overcomes the scalability bottleneck of global attention in feed-forward 3D reconstruction. We identified that the full global self-attention in models like VGGT is computationally prohibitive and largely unnecessary. Our solution centers on a compressed descriptor attention paradigm that replaces dense attention with efficient cross-attention from all image tokens to a compact set of learned descriptor tokens. Furthermore, the compact nature of these descriptors enables a chunk-recursive inference scheme, allowing FlashVGGT to process very long sequence with a manageable memory footprint. Through extensive experiments, we demonstrated that FlashVGGT achieves a superior balance between efficiency and accuracy. It matches the reconstruction quality of VGGT while achieving over 90% inference speedup for 1,000-image sequences and scaling effectively to over 3,000 images. By making high-fidelity, large-scale 3D reconstruction both fast and practical, FlashVGGT opens the door for more demanding real-world applications.

**Limitations:** While highly efficient for long sequences, our method exhibits a slight performance degradation on shorter sequences. Furthermore, the design space for descriptor attention remains largely unexplored. Please refer to the supplementary material for a detailed discussion.



## Acknowledgments

The research is supported in part by Early Career Scheme of the Research Grants Council (RGC) of the Hong Kong SAR under grant No. 26202321, ITF PRP/046/24FX, Science & Technology Cooperation Program of Shandong under grant No. SDST26EG01, SAIL Research Project, HKUST-Zeekr Coolaborative Research Fund, Westwell Project, and Tencent Rhino-Bird Focused Research Program.

## References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, 2021. 3
- [2] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural rgb-d surface reconstruction. In *CVPR*, 2022. 6, 7, 8
- [3] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, et al. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data. *arXiv*, 2021. 2, 5
- [4] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer vision and image understanding*, 2008. 2
- [5] Daniel Bolya and Judy Hoffman. Token merging for fast stable diffusion. In *CVPR*, 2023. 3
- [6] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. *arXiv*, 2022. 3
- [7] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 5, 6
- [8] Johann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2. *arXiv*, 2020. 5
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 3, 4
- [10] Tianlong Chen, Yu Cheng, Zhe Gan, Lu Yuan, Lei Zhang, and Zhangyang Wang. Chasing sparsity in vision transformers: An end-to-end exploration. *NeurIPS*, 2021. 3
- [11] Xingyu Chen, Yue Chen, Yuliang Xiu, Andreas Geiger, and Anpei Chen. Ttt3r: 3d reconstruction as test-time training. *arXiv*, 2025. 3, 6, 7
- [12] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. 2, 5, 6
- [13] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *NeurIPS*, 2022. 3
- [14] Timothe Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv*, 2023. 3
- [15] Kai Deng, Zexin Ti, Jiawei Xu, Jian Yang, and Jin Xie. Vggt-long: Chunk it, loop it, align it—pushing vggt’s limits on kilometer-scale long rgb sequences. *arXiv*, 2025. 3
- [16] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv*, 2020. 2, 3
- [17] Xianze Fang, Jingnan Gao, Zhe Wang, Zhuo Chen, Xingyu Ren, Jiangjing Lyu, Qiaomu Ren, Zhonglei Yang, Xiaokang Yang, Yichao Yan, and Chengfei Lyu. Dens3r: A foundation model for 3d geometry prediction. *arXiv preprint arXiv:2507.16290*, 2025. 3
- [18] Mohsen Fayyaz, Soroush Abbasi Koohpayegani, Farnoush Rezaei Jafari, Sunando Sengupta, Hamid Reza Vaezi Joze, Eric Sommerlade, Hamed Pirsiavash, and Jürgen Gall. Adaptive token sampling for efficient vision transformers. In *ECCV*, 2022. 3
- [19] Jingnan Gao, Zhe Wang, Xianze Fang, Xingyu Ren, Zhuo Chen, Shengqi Liu, Yuhao Cheng, Jiangjing Lyu, Xiaokang Yang, and Yichao Yan. More: 3d visual geometry reconstruction meets mixture-of-experts. In *CVPR*, 2026. 3
- [20] Dongchen Han, Tianzhu Ye, Yizeng Han, Zhuofan Xia, Siyuan Pan, Pengfei Wan, Shiji Song, and Gao Huang. Agent attention: On the integration of softmax and linear attention. In *ECCV*, 2024. 3
- [21] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 1
- [22] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. In *CVPR*, 2018. 2, 5
- [23] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *ICML*, 2021. 3, 2
- [24] Nikhil Keetha, Norman Müller, Johannes Schönberger, Lorenzo Porzi, Yuchen Zhang, Tobias Fischer, Arno Knapitsch, Duncan Zauss, Ethan Weber, Nelson Antunes, et al. Mapanything: Universal feed-forward metric 3d reconstruction. *arXiv*, 2025. 2, 3
- [25] Alex Kendall and Roberto Cipolla. Modelling uncertainty in deep learning for camera relocation. In *ICRA*, 2016. 3
- [26] Jan J Koenderink and Andrea J Van Doorn. Affine structure from motion. *Journal of the Optical Society of America A*, 8 (2):377–385, 1991. 1
- [27] Vincent Leroy, Johann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In *ECCV*, 2024. 2
- [28] Hao Li, Zhengyu Zou, Fangfu Liu, Xuanyang Zhang, Fangzhou Hong, Yukang Cao, Yushi Lan, Manyuan Zhang, Gang Yu, Dingwen Zhang, et al. Iggt: Instance-grounded geometry transformer for semantic 3d reconstruction. *arXiv*, 2025. 3
- [29] Xinyang Li, Muyang Li, Tianle Cai, Haocheng Xi, Shuo Yang, Yujun Lin, Lvmin Zhang, Songlin Yang, Jinbo Hu, Kelly Peng, et al. Radial attention: O(nlogn) sparse attention with energy decay for long video generation. *arXiv preprint arXiv:2506.19852*, 2025. 3
- [30] Yanyu Li, Geng Yuan, Yang Wen, Ju Hu, Georgios Evangelidis, Sergey Tulyakov, Yanzhi Wang, and Jian Ren. Efficientformer: Vision transformers at mobilenet speed. *NeurIPS*, 2022. 3

- [31] Xinyu Liu, Houwen Peng, Ningxin Zheng, Yuqing Yang, Han Hu, and Yixuan Yuan. Efficientvit: Memory efficient vision transformer with cascaded group attention. In *CVPR*, 2023. 3
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 3
- [33] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982. 4
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv*, 2017. 1
- [35] David G Lowe. Object recognition from local scale-invariant features. In *ICCV*, 1999. 2
- [36] David G Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004. 2
- [37] Dmitrii Marin, Jen-Hao Rick Chang, Anurag Ranjan, Anish Prabhu, Mohammad Rastegari, and Oncel Tuzel. Token pooling in vision transformers. *arXiv*, 2021. 3
- [38] Pierre Moulon, Pascal Monasse, Romuald Perrot, and Renaud Marlet. OpenMVG: Open multiple view geometry. In *International Workshop on Reproducible Research in Pattern Recognition*, 2016. 1
- [39] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 5, 6
- [40] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulo, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017. 2, 5
- [41] E. Palazzolo, J. Behley, P. Lottes, P. Giguère, and C. Stachniss. ReFusion: 3D Reconstruction in Dynamic Environments for RGB-D Cameras Exploiting Residuals. In *IROS*, 2019. 5, 6
- [42] Linfei Pan, Dániel Baráth, Marc Pollefeys, and Johannes L Schönberger. Global structure-from-motion revisited. In *ECCV*, 2024. 1
- [43] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. 3
- [44] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *ICCV*, 2021. 2, 5, 6
- [45] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *CVPR*, 2020. 2
- [46] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 1
- [47] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, 2016. 1
- [48] Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *CVPR*, 2006. 1
- [49] You Shen, Zhipeng Zhang, Yansong Qu, and Liujuan Cao. Fastvggt: Training-free acceleration of visual geometry transformer. *arXiv*, 2025. 3, 5, 6, 2
- [50] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in rgb-d images. In *CVPR*, 2013. 6, 1
- [51] Jiaming Tang, Yilong Zhao, Kan Zhu, Guangxuan Xiao, Baris Kasikci, and Song Han. Quest: Query-aware sparsity for efficient long-context llm inference. *arXiv preprint arXiv:2406.10774*, 2024. 3
- [52] Shimon Ullman. The interpretation of structure from motion. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 203(1153):405–426, 1979. 1
- [53] Pavan Kumar Anasosalu Vasu, James Gabriel, Jeff Zhu, Oncel Tuzel, and Anurag Ranjan. Fastvit: A fast hybrid vision transformer using structural reparameterization. In *ICCV*, 2023. 3
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 2
- [55] Chung-Shien Brian Wang, Christian Schmidt, Jens Piekenbrinck, and Bastian Leibe. Faster vgggt with block-sparse global attention. *arXiv*, 2025. 3
- [56] Junke Wang, Xitong Yang, Hengduo Li, Li Liu, Zuxuan Wu, and Yu-Gang Jiang. Efficient video transformers with spatial-temporal token selection. In *ECCV*, 2022. 3
- [57] Jianyuan Wang, Nikita Karaev, Christian Rupprecht, and David Novotny. Vggsfm: Visual geometry grounded deep structure from motion. In *CVPR*, 2024. 2
- [58] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *CVPR*, 2025. 1, 2, 3, 5, 6
- [59] Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *CVPR*, 2025. 3, 5, 6, 7
- [60] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *CVPR*, 2024. 2
- [61] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He. Permutation-equivariant visual geometry learning. *arXiv*, 2025. 2, 3
- [62] Cong Wei, Brendan Duke, Ruowei Jiang, Parham Aarabi, Graham W Taylor, and Florian Shkurti. Sparsifiner: Learning sparse instance-dependent attention for efficient vision transformers. In *CVPR*, 2023. 3
- [63] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*, 2023. 3
- [64] Jianing Yang, Alexander Sax, Kevin J Liang, Mikael Henaff, Hao Tang, Ang Cao, Joyce Chai, Franziska Meier, and Matt Feiszli. Fast3r: Towards 3d reconstruction of 1000+ images in one forward pass. In *CVPR*, 2025. 2, 5, 6
- [65] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *CVPR*, 2020. 2, 5

- [66] Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the presence of motion. *arXiv*, 2024. [2](#)
- [67] Shangzhan Zhang, Jianyuan Wang, Yinghao Xu, Nan Xue, Christian Rupprecht, Xiaowei Zhou, Yujun Shen, and Gordon Wetzstein. Flare: Feed-forward geometry, appearance and camera estimation from uncalibrated sparse views. In *CVPR*, 2025. [6](#)
- [68] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv*, 2018. [5](#), [6](#)
- [69] Dong Zhuo, Wenzhao Zheng, Jiahe Guo, Yuqi Wu, Jie Zhou, and Jiwen Lu. Streaming 4d visual geometry transformer. *arXiv*, 2025. [2](#), [3](#), [5](#), [6](#), [7](#)



# FlashVGGT: Efficient and Scalable Visual Geometry Transformers with Compressed Descriptor Attention

## Supplementary Material

**Overview.** The supplementary material is organized as follows: Sec. A elaborates on training specifics and hyperparameters. Sec. B presents extended experimental analysis and ablation studies. Finally, Sec. C discusses the limitations of our current work and promising directions for future research.

### A. Implementation Details

**Model and Optimization.** We initialize FlashVGGT from a pre-trained VGGT [58] checkpoint. During training, we freeze the image encoder and reconstruction heads, optimizing only the alternating attention aggregator which comprises approximately 50% of the total parameters while employing the original VGGT loss functions. We use the Adam-W optimizer [34] with an initial learning rate of  $4 \times 10^{-6}$ , linear warmup, and cosine decay. Both training stages run for 10,000 iterations on 4 H800 GPUs, completing in approximately 16 hours.

**Training Protocol.** We adopt VGGT’s dynamic batching scheme, randomly sampling 2 to 24 frames per iteration. Input images are pre-processed by resizing the longer side to 518 pixels while randomizing the aspect ratio between 0.33 and 1.0. We apply standard data augmentation including color jittering and random grayscale conversion. For training stability, we employ gradient norm clipping with a threshold of 1.0, and leverage both bfloat16 precision and gradient checkpointing as in the original VGGT.

### B. Additional Analysis

**Ablation of auxiliary descriptor tokens.** As shown in Tab. 6, we analyze the contribution of each auxiliary token component using 500-image sequences from 7Scenes [50]. The full model achieves the best overall performance, validating the importance of all auxiliary tokens. The most critical component is the reference frame tokens, whose removal causes the most severe degradation, particularly in pose estimation (APE increases by 96% and ARE by 68%). This confirms that preserving the full coordinate frame is essential for global geometric consistency. Camera tokens also prove vital, as their absence leads to noticeable deterioration in reconstruction quality while providing minimal memory savings. This demonstrates that explicit camera parameter representation significantly aids the network’s geometric reasoning. While key frame tokens offer the most modest improvements, they still enhance both reconstruction (CD) and camera pose (APE) metrics with negli-

ble computational overhead. This suggests that distributing full-resolution information across the sequence helps maintain local detail preservation. Notably, all auxiliary components contribute to performance with minimal impact on efficiency, confirming our design provides an effective accuracy-efficiency trade-off.

Table 6. **Detailed ablations of auxiliary descriptor tokens.** Evaluated on 7Scenes [50] with 500 input images.

	Abs Rel↓	CD↓	NC↑	APE↓	ARE↓	Time (s)	Mem (GB)
w/o Cam tokens	0.066	2.849	64.01	3.908	8.115	12.99	33.40
w/o First frame	0.067	2.866	58.91	7.660	13.608	12.67	33.39
w/o Key frames	0.067	2.859	63.68	4.183	8.123	<b>12.04</b>	<b>33.33</b>
Full model	<b>0.066</b>	<b>2.748</b>	<b>64.12</b>	<b>3.904</b>	<b>8.115</b>	12.99	33.40

**Key-frame selection methods.** As shown in Tab. 7, we compare different strategies for selecting key frames. The clustering-based approach consistently outperforms both random and fixed-stride selection across all accuracy metrics. Our proposed clustering method achieves the best performance, demonstrating its ability to select representative frames that better capture the scene’s geometric diversity. This comes at a minimal computational cost, adding only 0.29 seconds compared to the fastest method. While fixed-stride selection offers slightly better efficiency, it shows clear performance limitations, particularly in pose estimation. This suggests that uniformly distributed frames may miss critical viewpoints needed for optimal geometric reconstruction. Random selection performs similarly to fixed-stride but with more variability across metrics, confirming that naive approaches cannot reliably identify the most informative frames for 3D reconstruction. The results validate that our clustering-based key frame selection effectively identifies geometrically representative views, providing better reconstruction quality with negligible overhead compared to simpler alternatives.

Table 7. **Comparison of different key-frame selection methods.**

	Abs Rel↓	CD↓	NC↑	APE↓	ARE↓	Time (s)	Mem (GB)
Random	0.067	2.789	63.92	4.108	8.123	12.74	33.39
Fix stride	0.067	2.784	64.02	4.096	8.189	<b>12.70</b>	<b>33.39</b>
Cluster	<b>0.066</b>	<b>2.748</b>	<b>64.12</b>	<b>3.904</b>	<b>8.115</b>	12.99	33.40

**Memory retain rate  $p$ .** We analyze the trade-off between efficiency and accuracy by varying the memory retain rate  $p$ , which controls how many historical descriptor tokens are preserved during chunk-recursive inference. As shown in Fig. 10, lower values of  $p$  (more aggressive mem-

ory dropping) yield better efficiency but gradually degrade reconstruction quality. Notably, the performance drop from  $p = 1$  to  $p = 5$  is minimal compared to the substantial efficiency gains. This suggests that carefully selected memory dropping can eliminate redundant historical information with negligible impact on reconstruction quality.




Figure 10. **The impact of different memory retain rate  $p$ .** The rate of 5 provides a balanced choice between quality and speed.

**Chunk size.** We analyze the impact of chunk size on online reconstruction performance in Tab. 8. The results demonstrate that chunk size has minimal effect on reconstruction quality across all metrics. However, chunk size significantly impacts computational efficiency: increasing from 1 to 100 frames per chunk provides a  $2.3\times$  speedup at the cost of 49% higher memory usage. This reveals a flexible trade-off between speed and memory constraints. Applications prioritizing throughput can use larger chunks (*e.g.*, 100) for faster processing, while memory-constrained environments can employ smaller chunks (*e.g.*, 10-50). The stability of reconstruction metrics across chunk sizes confirms the robustness of our chunk-recursive scheme, making it adaptable to diverse deployment scenarios.

Table 8. **The impact of different chunk size in online setting.**

Chunk	Abs Rel↓	CD↓	NC↑	APE↓	ARE↓	Time (s)	Mem (GB)
1	0.086	2.342	59.16	5.165	1.031	25.50	11.42
10	0.087	2.391	59.74	5.042	0.980	14.58	13.33
50	0.087	2.362	60.70	5.224	1.003	12.58	14.95
100	0.087	2.330	60.97	5.173	1.034	10.90	16.98

**Resources consumption.** As shown in Tab. 9, FlashVGGT demonstrates substantial efficiency gains across all resource metrics while scaling to longer sequences than competing methods. **(i) Computational Efficiency:** FlashVGGT achieves remarkable reductions in both time and FLOPs. For 1,000 images, our method is  $10.1\times$  faster than VGGT [58] and requires  $15.8\times$  fewer FLOPs. Even compared to FastVGGT [49], FlashVGGT provides a  $2.1\times$  speedup and  $2.1\times$  FLOP reduction, demonstrating the superiority of our descriptor-based approach over token merging. **(ii) Memory Efficiency:** FlashVGGT maintains the lowest memory footprint across all sequence lengths. At

Table 9. **Comparison of resources consumption across different input images.** ‘-’ denotes model running out memory.

	Methods	200	400	600	800	1000	1200
Time (s)	VGGT	17.01	61.82	137.84	245.47	386.07	-
	FastVGGT	6.45	16.63	32.19	52.01	79.31	-
	<b>FlashVGGT</b>	<b>4.05</b>	<b>9.84</b>	<b>17.25</b>	<b>26.44</b>	<b>38.1</b>	<b>51.25</b>
PFLOPs	VGGT	4.24	16.92	38.04	67.61	105.61	-
	FastVGGT	0.59	2.23	5.19	9.20	14.34	-
	<b>FlashVGGT</b>	<b>0.29</b>	<b>1.10</b>	<b>2.43</b>	<b>4.30</b>	<b>6.70</b>	<b>9.62</b>
Mem (GB)	VGGT	18.50	30.98	43.45	55.93	68.40	-
	FastVGGT	19.34	32.66	45.97	59.29	72.60	-
	<b>FlashVGGT</b>	<b>16.97</b>	<b>27.92</b>	<b>38.83</b>	<b>49.76</b>	<b>60.68</b>	<b>71.61</b>

1,000 images, it uses 11% less memory than VGGT and 16% less than FastVGGT. This memory advantage enables FlashVGGT to successfully process 1,200-image sequences where both baselines fail due to out-of-memory errors. Note that FastVGGT takes more memory as it needs to compute token similarity across all tokens. **(iii) Scalability:** The computational advantages become more pronounced with longer sequences. While VGGT and FastVGGT cannot process beyond 1,000 images, FlashVGGT maintains efficient operation at 1,200 images with only 51.25 seconds and 71.61GB memory, demonstrating robust scalability for large-scale reconstruction tasks. These results confirm that our descriptor-based attention mechanism provides fundamental improvements in computational efficiency without compromising on reconstruction quality, enabling practical processing of very long image sequences.

**Comparison with latent cross-attention.** While our method shares the high-level intuition of reducing computation via asymmetric attention with latent cross-attention methods like Perceiver [23], our approach differs fundamentally. Unlike Perceiver-style methods that use randomly initialized learnable tokens as queries to aggregate information from the input, we use the original input tokens as queries and a spatially compressed version as keys and values. This design preserves the original input resolution, making it more suitable for dense prediction tasks like 3D reconstruction. Furthermore, our compressed descriptors carry strong data-dependent priors through spatial resampling, maintaining the input’s structural distribution rather than learning a generic latent representation. To validate our approach, we compare against a Perceiver-style alternative that uses additional learnable latent tokens per frame. During frame attention, these latent tokens interact with tokens in the same frame, while in global attention, we compute cross-attention from frame tokens to these latent tokens. As shown in Tab. 10, our method significantly outperforms the Perceiver-style approach across all metrics. These results demonstrate that leveraging data-dependent compression is crucial for high-quality 3D reconstruction.

**In-the-wild Benchmark.** We evaluated our model on the

Table 10. Comparison between our compressed descriptor attention and Perceiver-style latent cross-attention.

	Abs Rel↓	CD↓	NC↑	APE↓	ARE↓	Time (s)	Mem (GB)
Perceiver-style	0.097	5.645	34.02	14.573	12.564	13.56	34.56
<b>Ours</b>	<b>0.066</b>	<b>2.748</b>	<b>64.12</b>	<b>3.904</b>	<b>8.115</b>	<b>12.99</b>	<b>33.40</b>

IMC PhotoTourism benchmark. Although our architecture is primarily optimized for long sequences, it exhibits only a minor accuracy gap compared to VGGT on these shorter in-the-wild sequences (5-25 frames) and consistently outperforms the concurrent efficient alternative, FastVGGT.

Table 11. Evaluation on the IMC PhotoTourism benchmark.

	AUC@3	AUC@5	AUC@10	Time (s)
VGGT	<b>39.23</b>	<b>52.74</b>	<b>71.26</b>	0.37
FastVGGT	38.58	51.43	70.12	<u>0.35</u>
FlashVGGT	<u>38.62</u>	<u>51.87</u>	<u>70.49</u>	<b>0.26</b>

**Stricter Camera Pose Metrics.** We evaluated our model with stricter thresholds of 5 and 10 degrees on RealEstate10K (10 frames per sequence). Although our architecture is primarily optimized for long sequences, it closely matches VGGT’s metrics on these short sequences, while performing better than the concurrent work, FastVGGT [49].

Table 12. Camera pose evaluation with stricter thresholds.

	Racc@5	Tacc@5	Auc@5	Racc@10	Tacc@10	Auc@10	Times (s)
VGGT	<b>97.06</b>	<b>60.61</b>	<b>35.46</b>	<b>99.40</b>	<b>80.20</b>	<b>54.76</b>	0.22
FastVGGT	96.52	58.32	34.63	98.93	78.69	53.11	<u>0.20</u>
FlashVGGT	<u>96.67</u>	<u>58.44</u>	<u>34.75</u>	<u>99.11</u>	<u>78.98</u>	<u>53.78</u>	<b>0.15</b>

**Downsampling Alternatives.** To justify our descriptor-based attention, we compared it against three alternative strategies with a similar computational budget and the same training scheme: **(a) Input Downsampling:** Resizing input images before encoding. **(b) Feature Downsampling:** Resizing DINO features immediately after the encoder. **(c) Global Bottleneck:** Downsampling tokens before the global attention block and upsampling them afterward to restore resolution. As shown in the table below, our descriptor-based attention outperforms all variants. The reason is that all other variants irreversibly destroy high-frequency details through downsampling, while our method maintains the encoder, frame attention, heads, and global queries at full resolution. By only compressing the global keys and values, we model inter-frame correspondences efficiently without sacrificing intra-frame details, and therefore achieve better reconstruction quality.

**Inference Breakdown.** We provide a detailed breakdown of the inference time across different components for 1,000-image sequences, measured in seconds.

Table 13. Comparison with downsampling alternatives.

Method	Abs Rel↓	Acc↓	Comp↓	APE↓	Time (s)	Mem (GB)
(a) Input Downsampling	0.332	1.306	0.936	6.843	<b>9.72</b>	<b>33.87</b>
(b) Feature Downsampling	0.312	0.954	0.923	5.560	10.87	34.23
(c) Global Bottleneck	<u>0.241</u>	<u>0.810</u>	<u>0.884</u>	<u>4.630</u>	11.03	36.87
<b>Descriptor (Ours)</b>	<b>0.161</b>	<b>0.512</b>	<b>0.310</b>	<b>2.733</b>	<u>12.08</u>	<u>37.21</u>

Table 14. Breakdown of the inference time.

	Encoder	Frame Blocks	Global Blocks	Reconstruction Heads	Total
VGGT	2.25	5.15	368.16	2.04	377.60
FlashVGGT	<b>2.24</b>	<b>5.15</b>	<b>25.93</b>	<b>2.04</b>	<b>35.37</b>

## C. Discussions

**Limitations.** While FlashVGGT achieves substantial efficiency improvements over VGGT [58], several limitations merit discussion. First, our method exhibits a slight performance gap on short sequences, as evidenced in Tab. 1. However, this gap diminishes with longer sequences where our architectural advantages become more pronounced. Second, similar to the original VGGT, our model’s performance can degrade under challenging conditions involving large deformations or extreme lighting variations. This limitation, however, could potentially be mitigated through fine-tuning on domain-specific data without requiring architectural changes. These aspects represent promising directions for future work in enhancing robustness and applicability.

**Design Space.** While our framework achieves substantial efficiency improvements, its design space offers rich opportunities for future exploration. For instance, although our experiments in Tab. 5 demonstrate that a simple convolutional compressor underperforms interpolation, more sophisticated learnable architectures for token compression and selection remain unexplored. Investigating adaptive mechanisms that can dynamically adjust compression strategies based on input characteristics could potentially yield further performance gains.

**Integration with other architectures.** While our primary evaluation is based on VGGT [58], our descriptor-based attention mechanism is a general module that can be readily integrated into other architectures employing the alternating attention backbone [17, 19, 24, 61]. This portability paves the way for developing more efficient variants of state-of-the-art models across various tasks, including metric reconstruction [24] and semantic 3D understanding [28].