

LongStream: Long-Sequence Streaming Autoregressive Visual Geometry

Chong Cheng^{1,2} Xianda Chen⁴ Tao Xie^{2,3} Wei Yin²
 Weiqiang Ren² Qian Zhang² Xiaoyang Guo^{2,‡} Hao Wang^{1,†}

¹HKUST (GZ) ²Horizon Robotics ³ZJU ⁴CSU

[‡]Project Lead [†]Corresponding Author

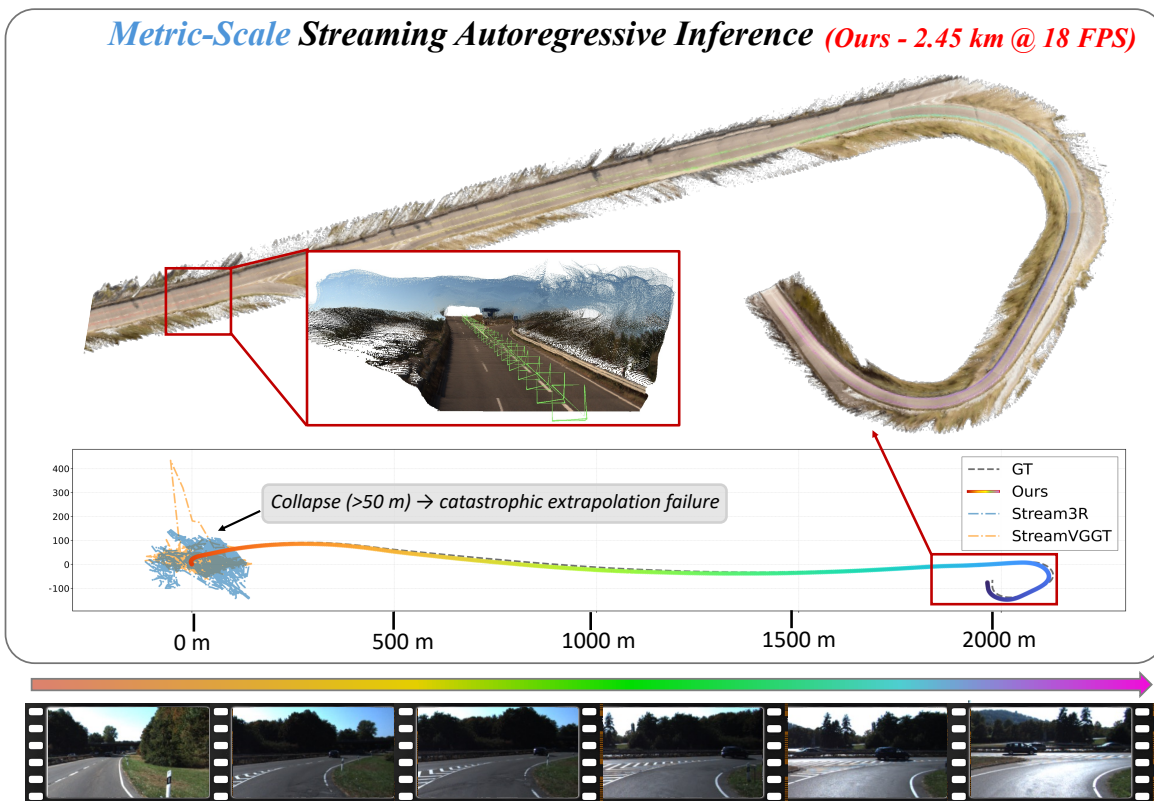


Figure 1. **Streaming Autoregressive Model Comparison for Metric-Scale Scene Reconstruction.** Existing streaming models (e.g., Stream3R [1], StreamVGGT [2]) collapse within tens of meters, suffering from extrapolation errors. In contrast, our proposed LongStream delivers stable, kilometer-scale reconstruction. Its gauge-decoupled formulation and cache-consistent training preserve metric accuracy and geometric stability, sustaining 18 FPS performance across multi-kilometer sequences.

Abstract

Long-sequence streaming 3D reconstruction remains a significant open challenge. Existing autoregressive models often fail when processing long sequences because they anchor poses to the first frame, leading to attention decay, scale drift, and extrapolation errors. We introduce LongStream, a novel gauge-decoupled streaming visual geometry model for metric-scale scene reconstruction across thousands of frames under a strictly online, future-invisible

setting. Our approach is threefold. First, we discard the first-frame anchor and predict keyframe-relative poses. This reformulates long-range extrapolation into a constant-difficulty local task. Second, we introduce orthogonal scale learning. This method fully disentangles geometry from scale estimation to suppress drift. Finally, we identify attention bias issues in Transformers, including attention-sink reliance and long-term KV-cache saturation. We propose cache-consistent training combined with periodic cache refresh. This approach suppresses attention biases and

contamination over ultra-long sequences and reduces the gap between training and inference. Experiments show that LongStream achieves state-of-the-art performance, enabling stable, metric-scale reconstruction over kilometer-scale sequences at 18 FPS. Project Page: <https://3dagentworld.github.io/longstream/>

1. Introduction

Geometry reconstruction, the joint estimation of camera poses and dense 3D structure from image sequences [3–8], is a cornerstone technology for applications like autonomous driving, AR/VR, and embodied robotics [9, 10]. These domains demand systems that can robustly process long-sequence video streams in real time.

Conventional pipelines [3, 4, 6, 11, 12] and recent Transformer-based approaches [13–18] achieve state-of-the-art accuracy, but are inherently offline. They require re-processing the entire sequence to integrate a new frame [1, 2], leading to massive computational redundancy and precluding real-time use. To address this, streaming models have been proposed. Recent works such as SStream3R [1] and StreamVGGT [2] employ causal Transformers and KV-caching to build reconstructions incrementally. However, these models suffer from catastrophic extrapolation failure when processing long sequences. As shown in Figures 1 and 2, existing streaming methods incrementally reconstruct scenes in linear time, but their trajectories collapse within tens of meters, leading to complete tracking failure.

We argue that this failure stems from the “**gauge-coupled**” design inherent in current models. They are anchored to the first-frame coordinate system and trained to regress absolute poses. This forces the model to learn a position-fixed mapping, making long-sequence prediction increasingly difficult. Due to hardware constraints, the model only sees small indices in a training batch, making it difficult to extrapolate reliably to large indices at inference. As a result, a domain gap emerges under the “train-short, test-long” bias [19, 20].

To this end, we propose LongStream, a “**gauge-decoupled**” streaming autoregressive geometry framework. The framework theoretically decouples the gauge freedoms of global $SE(3)$ coordinates and metric scale.

First, to handle coordinate gauge over long sequences, we remove the fixed first-frame anchor. Instead, we regress keyframe-relative poses. This reformulates the ill-posed long-range extrapolation problem into a constant-difficulty local estimation task. As a result, the predictions become invariant to the global coordinate choice.

Second, to tackle Sim(3) scale drift, we introduce an orthogonal scale learning mechanism [21]. We decouple geometry learning from metric scale estimation at the objective level. The geometry branch optimizes shape in a scale-




Figure 2. **Memory and runtime comparison.** Our method keeps memory and latency stable, whereas VGGT and FastVGGT grow rapidly and hit OOM on long sequences.

invariant space, while a dedicated scale head independently predicts the global scale factor. This effectively reduces scale entanglement and ensures stable metric outputs.

Beyond this gauge-decoupled formulation, the streaming Transformer architecture itself presents challenges. We identify that attention sink [22], i.e., its biased dependence on the first-frame token, and long-term KV-cache contamination are primary causes of temporal degradation and drift. To address this, we propose a cache-consistent training scheme. It aligns training and inference contexts by explicitly passing and trimming the cache during training. We further introduce a periodic cache refresh approach, which marginalizes stale context, mitigates long-term memory saturation, and stabilizes geometry.

Experiments on both outdoor (KITTI, vKITTI, Waymo) and indoor (TUM-RGBD, ETH3D, 7Scenes) datasets show that LongStream achieves state-of-the-art streaming reconstruction. It enables **real-time (18 FPS)**, **metric-scale** AR reconstruction over kilometer-scale sequences. Our contributions are summarized as follows:

- We propose LongStream, a streaming geometry foundation model centered on a “gauge-decoupled” design. It predicts keyframe-relative poses and employs orthogonal scale decoupling. This design systematically eliminates first-frame anchor dependence and effectively mitigates failures in long-sequence extrapolation.
- We identify attention-sink reliance and KV-cache contamination as the primary causes of long-horizon degradation. Our cache-consistent training scheme and periodic cache refresh stabilize temporal attention and reduce geometric drift.
- Experiments across indoor and outdoor benchmarks show that LongStream achieves state-of-the-art performance under a strictly online, future-invisible setting, while maintaining real-time throughput and stable metric scale on long sequences.

2. Related Work

Classical SfM and MVS. Structure-from-Motion (SfM) and Multi-View Stereo (MVS) pipelines [3–6, 11, 23, 24] reconstruct 3D scenes by modeling geometric relations between image correspondences. SfM estimates camera poses and sparse structure via feature matching and bundle ad-

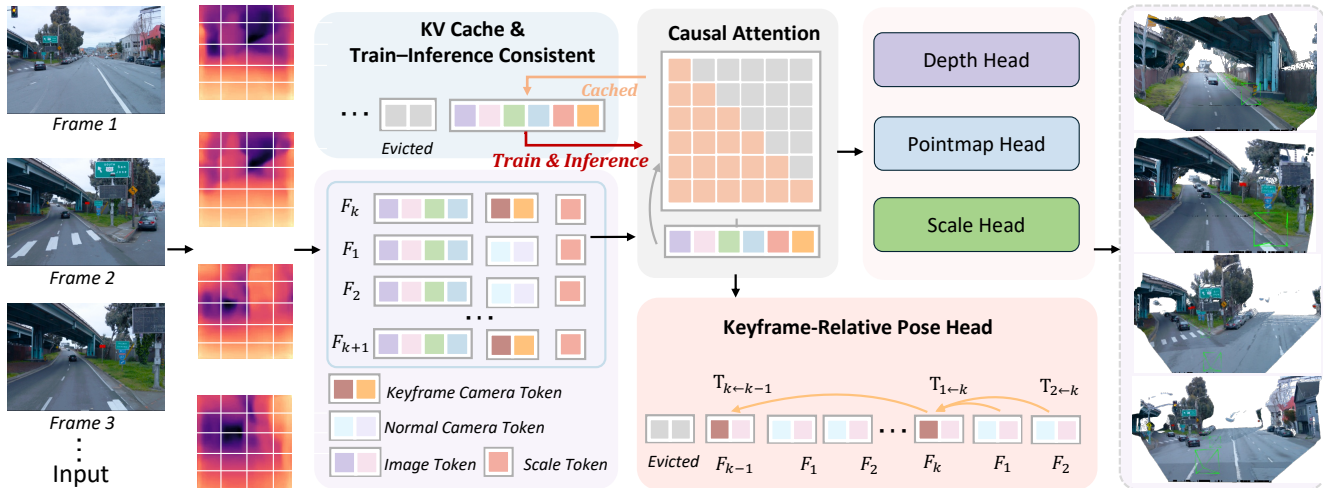


Figure 3. **Overview of our proposed LongStream.** Given streaming inputs, patch tokens are extracted by a ViT encoder and augmented with *keyframe*, *normal-frame*, and *scale tokens*. Tokens are fused via causal attention with a shared KV cache, which is consistently used in both training and inference for cache-consistent streaming modeling. The network predicts keyframe-relative poses $T_{i \leftarrow k}$, depth, pointmap, and global scale, enabling stable metric-scale reconstruction over long sequences.

justment [3, 4], while MVS densifies them with pixel-wise depth from plane sweeping or cost volumes [5, 6, 11, 23, 24]. Despite high accuracy and interpretability, these optimization-heavy pipelines with handcrafted features scale poorly to large or dynamic scenes and are difficult to deploy in real time.

Offline 3D reconstruction. Early end-to-end methods are mainly trained on image pairs [25–28]. Pointmaps offer efficient computation and real-time usage compared to voxel, mesh, or implicit fields [29–33], enabling SLAM and neural rendering [34–38]. DUST3R [13] regresses pointmaps and relative pose from two images without intrinsics, but remains pairwise and requires global alignment. MAST3R [14] adds dense features and reciprocal matching to improve calibration and matching robustness, yet operates on pairs and needs costly fusion for multi-view scenes. VGGT [15] predicts poses, depth, pointmaps, and tracks from multi-view inputs in a single feed-forward pass, but relies on a fixed reference frame and absolute pose supervision, causing reference and scale biases. π^3 [16] removes reference-view bias via permutation-equivariant design and predicts local pointmaps, but outputs remain ambiguous up to a global similarity transform and lack metric scale.

Streaming 3D reconstruction. Streaming methods update geometry frame by frame. Classical monocular SLAM and learning-based variants [39–44] incrementally recover structure and motion. CUT3R [45] maintains a recurrent state to output metric pointmaps online, but its RNN backbone struggles with long-term dependencies and degrades on long sequences. Stream3R [1] adopts a causal Transformer with a KV cache, scaling to long streams but suffering attention collapse as cached tokens dominate.

StreamVGGT [2] adds temporal causal attention and cache updates, with distillation improving consistency, yet long-horizon stability remains difficult due to cache contamination. Overall, existing streaming methods degrade noticeably as sequences grow and fail to generalize to much longer streams.

3. Methodology

3.1. Overall

We propose LongStream, a *gauge-decoupled* streaming geometry framework that jointly predicts pose, depth, and scale within a unified spatiotemporal Transformer, as shown in Figure 3. A ViT encoder produces patch tokens augmented with *Camera*, *Register*, and *Scale* tokens to distinguish keyframe roles. These tokens are fused by a causal aggregator with a shared *KV cache*, enabling long-sequence streaming inference. For each frame, the model predicts a keyframe-relative pose $T_{i \leftarrow k}$, depth, a pointmap, and a global scale s . Training and inference use the same layout, ensuring gauge-decoupled and stable metric-scale reconstruction. We next detail the components of the design.

3.2. Gauge-Decoupled Formulation

We aim to overcome the limitations of existing streaming models that rely on gauge-coupled designs. We argue that a robust geometric learning system must remain theoretically invariant to the gauge freedoms, namely arbitrary global coordinates and metric scale. To this end, we propose a framework that systematically separates the $SE(3)$ and $Sim(3)$ degrees of freedom.

In the $SE(3)$ gauge, we discard absolute pose regression and redefine pose learning as gauge-invariant [46, 47] rela-

tive pose estimation. The learning objective becomes:

$$\mathbf{T}_{i \leftarrow k} = \mathbf{T}_i \circ \mathbf{T}_k^{-1}, \quad (1)$$

where \mathbf{T}_i and \mathbf{T}_k denote world-to-camera absolute poses and the k -th frame is the preceding keyframe. This formulation is mathematically gauge-invariant under any world-frame reparameterization $S \in SE(3)$. The decoupling design achieves two goals simultaneously: it transforms an out-of-distribution long-range extrapolation problem, caused by large indices [1, 2], into a constant-difficulty in-distribution local task with bounded index gap ($i - k$); and it mitigates the fixed-anchor bias that contributes to instability in first-frame-anchored causal models.

In the $Sim(3)$ gauge, we address chaotic scale entanglement with the scale-invariant (SI-Log) philosophy [21]. We use an orthogonal scale learning mechanism, where we separate geometry learning and metric scale estimation at both the architectural and objective levels. The geometry branch is normalized and supervised with scale-invariant principles [21]. A dedicated scale head then predicts the global scale factor s .

3.3. Network Architecture

Given input images I_i , the model predicts, for each frame i , its relative pose with respect to the reference keyframe k , $\mathbf{p}_{i \leftarrow k} = [\mathbf{t}_{i \leftarrow k}, \mathbf{q}_{i \leftarrow k}, f_{i \leftarrow k}]$, where \mathbf{t} is translation, \mathbf{q} a unit quaternion, and f a focal-length offset. The model also outputs the depth map D_i , the corresponding world-coordinate point cloud X_i , a global scale factor s , and the frame representation h_i :

$$\{h_i, \mathbf{p}_{i \leftarrow k}, D_i, X_i, s\} = F_\theta(I_i), \quad i = 1, \dots, S. \quad (2)$$

where h_i denotes the frame-level feature representation.

Following the architecture of VGGT-style streaming models [1, 2, 15], the architecture consists of a DINOv2-based [48] tokenizer, a causal Transformer aggregator, and task-specific heads for relative pose, depth, pointmap, and scale. Geometry is modeled in a keyframe-relative manner: each frame predicts its pose and pointmap with respect to the current keyframe k , enabling streaming reconstruction without reliance on a fixed first-frame coordinate frame.

For each frame I_i , the tokenizer extracts patch features $x_i^p \in \mathbb{R}^{P \times C}$ and augments them with camera token and register tokens, with distinct tokens for keyframes and non-keyframes. A shared Scale Token is added for $Sim(3)$ decoupling. All tokens are concatenated into $H^{(0)} \in \mathbb{R}^{B \times S \times (P+T) \times C}$ and processed by a stack of Transformer blocks with alternating intra-frame and global attention under a strictly causal mask:

$$H^{(l+1)} = \text{Block}^{(l)}(H^{(l)}, \text{AttnMask}), \quad (3)$$

with outputs fed into the pose, depth, and pointmap heads for iterative refinement.

Relative pose head. It takes both frame and keyframe features from the aggregator and explicitly predicts the relative transformation $\mathbf{T}_{i \leftarrow k}$ of the current frame i with respect to its reference keyframe k :

$$\mathbf{p}_{i \leftarrow k} = [\mathbf{t}_{i \leftarrow k}, \mathbf{q}_{i \leftarrow k}, f_{i \leftarrow k}], \quad (4)$$

where the translation $\mathbf{t}_{i \leftarrow k} \in \mathbb{R}^3$, rotation $\mathbf{q}_{i \leftarrow k} \in \mathbb{H}$ (unit quaternion), and focal offset $f_{i \leftarrow k} \in \mathbb{R}^2$ jointly form the relative camera pose:

$$\mathbf{T}_{i \leftarrow k} = \mathbf{T}_i \mathbf{T}_k^{-1}. \quad (5)$$

This definition remains invariant under any right-multiplicative transformation of the world coordinate frame, thereby removing the dependency on a fixed world anchor.

To ensure that the model learns only the relative relationship between (i, k) , we adopt a reference-aware attention scheme. For a non-keyframe i assigned to keyframe k , its tokens attend only to tokens from k and from frames between k and i under the causal or window mask. Keyframe tokens, in turn, attend only to tokens from the previous keyframe $k-1$ and from frames between $k-1$ and k , rather than to the entire history. After aggregation, we fuse the pose tokens of the current frame and its reference keyframe via concatenation and a linear projection:

$$\mathbf{h}_{\text{fused}} = \text{Proj}([\mathbf{h}_i, \mathbf{h}'_k]). \quad (6)$$

Finally, following the design of RAFT [49], the head employs an AdaLN-modulated Transformer to iteratively predict the relative pose $\mathbf{p}_{\text{rel}} = [\mathbf{t}, \mathbf{q}, f]$ through incremental updates $\mathbf{p}^{(t+1)} = \mathbf{p}^{(t)} + \Delta \mathbf{p}^{(t)}$.

Scale head. To achieve $Sim(3)$ gauge invariance, we design an explicit scale head that receives a dedicated *Scale Token*, which is similar to the scale token idea in concurrent work MapAnything [50]. It predicts an unconstrained log-scale variable $x_s \in \mathbb{R}$, which is then exponentiated to obtain a strictly positive scale factor:

$$s = \exp(\mathbf{w}^\top \mathbf{h}_{\text{scale}}), \quad (7)$$

where $\mathbf{h}_{\text{scale}}$ is the feature of the scale token at the last aggregator layer. The scale s affects only translation, depth, and pointmap outputs, while rotation and field of view remain unchanged. The scale head is trained only on datasets with available ground-truth metric scale.

Depth and pointmap heads. The depth head and pointmap head take the aggregated frame-level and patch-level features to predict, for each frame, a depth map $D_i \in \mathbb{R}^{H \times W}$ and the corresponding world-coordinate points $X_i \in \mathbb{R}^{H \times W \times 3}$, along with per-pixel confidence scores. These branches operate jointly with the scale head: geometry is optimized in a scale-invariant space, while the *ScaleToken* independently learns the global scaling factor, ensuring full $Sim(3)$ gauge decoupling.

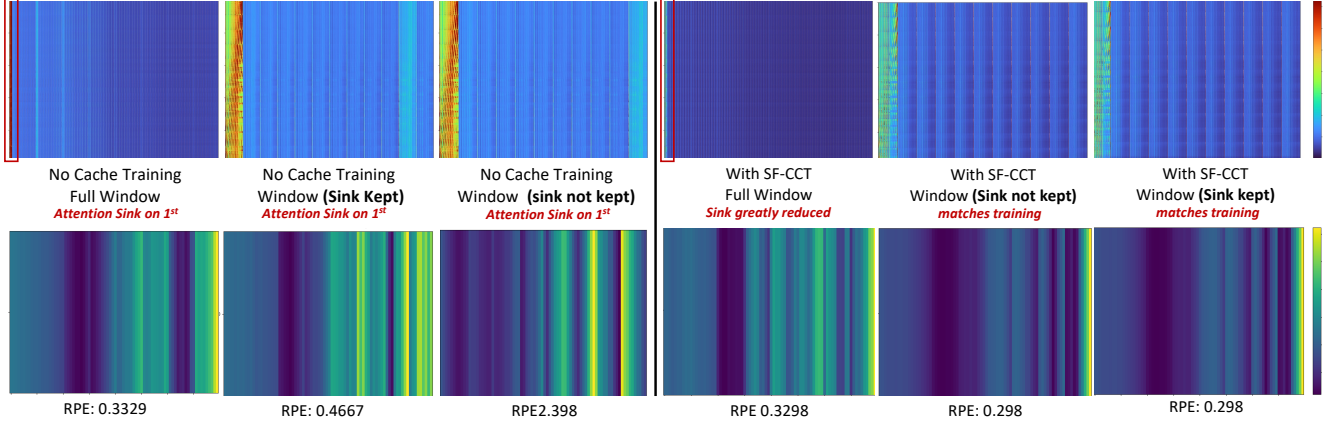


Figure 4. **Cache-consistent training (CCT)**. We show attention maps (top) and Relative Pose Error (RPE) heatmaps (bottom) under different training–inference settings. Without CCT (left), causal inference develops a strong attention sink; windowed inference either amplifies this sink when it is kept or collapses when it is removed. With CCT (right), the sink is strongly suppressed in causal mode and likewise suppressed in both windowed modes, yielding stable and best accuracy. Light blue denotes attention to the keyframe.

3.4. Probabilistic Framework and Loss Functions

To implement the gauge-decoupled design within a unified probabilistic framework, we formulate the overall objective as maximizing the joint likelihood of geometry, motion, and scale given the input sequence. Let I denote the image frames, X the 3D pointmap, D the depth, p the relative pose, and s the global scale. We minimize the Negative Log Posterior:

$$\mathcal{L} = \underbrace{\mathcal{L}_{\text{geom}} + \mathcal{L}_{\text{depth}}}_{\text{Geometry \& Depth Likelihood}} + \underbrace{\mathcal{L}_{\text{pose}}}_{\text{Pose Likelihood}} + \underbrace{\mathcal{L}_{\text{scale}}}_{\text{Scale Prior}}, \quad (8)$$

where each term corresponds to a conditional factor in the posterior decomposition

$$p(D, X, p, s | I) \propto p(D | X, I) \cdot p(X | p, s, I) \cdot p(p | I) \cdot p(s). \quad (9)$$

This factorization aligns the learning process with the gauge-invariant formulation of $SE(3)$ and $Sim(3)$ introduced in Sec. 3.2.

Relative pose loss. The relative pose loss $\mathcal{L}_{\text{pose}}$ corresponds to $p(p | I)$ and supervises the RelPoseHead output $\mathbf{p}_{\text{rel}} = [\mathbf{t}, \mathbf{q}, f]$ across iterative updates:

$$\mathcal{L}_{\text{pose}} = \sum_{t=1}^T \gamma^{t-1} \left(\ell(\hat{q}^{(t)}, q_{i \leftarrow k}) + \ell_t(\hat{t}^{(t)}, t_{i \leftarrow k}) + \ell(\hat{f}^{(t)}, f_{i \leftarrow k}) \right), \quad (10)$$

where ℓ denotes a L1 loss and $f_{i \leftarrow k}$ is the focal offset. To maintain gauge decoupling, the translation term ℓ_t is computed in a normalized coordinate space, ensuring that translation supervision does not implicitly encode global scale.

Geometry loss. Inspired by SI [21], the geometric loss $\mathcal{L}_{\text{geom}}$ (shape optimization) corresponds to $p(X | p, s, I)$

Algorithm 1 Cache-Consistent Training

- 1: **Input:** chunks $\{c_1, \dots, c_N\}$, initial cache $KV^{(0)} = \emptyset$
- 2: **for** $i = 1$ to N **do**
- 3: $(\text{out}_i, KV^{\text{new}}) = \text{model}(c_i, KV^{(i-1)})$
- 4: $KV^{(i)} = \text{trim}(KV^{\text{new}}, \text{window_size})$
- 5: **end for**

and operates in the normalized space:

$$\tilde{X}_{\text{pred}} = \frac{\hat{X}_{\text{raw}}}{\text{Norm}(\hat{X}_{\text{raw}})}, \quad \tilde{X}_{\text{gt}} = \frac{X}{\text{Norm}(X)}, \quad (11)$$

$$\mathcal{L}_{\text{geom}} = \|\tilde{X}_{\text{pred}} - \tilde{X}_{\text{gt}}\|_1.$$

Normalization removes explicit scale dependency, ensuring $\partial \mathcal{L} / \partial s = 0$. Hence, $\mathcal{L}_{\text{geom}}$ supervises only the backbone to learn correct 3D structure.

Scale loss. The scale loss $\mathcal{L}_{\text{scale}}$ regularizes the global scale s . Since scale is multiplicative, we compare predicted and ground-truth scale in log space to measure relative error and stabilize gradients:

$$\mathcal{L}_{\text{scale}} = \|\log \hat{s} - \log s_{\text{gt}}\|_1. \quad (12)$$

where s_{gt} is the metric scale computed from calibrated ground-truth depth. This loss is applied only to metric-calibrated samples; non-metric data are trained using geometry and depth losses alone.

3.5. KV Cache and Train–Inference Consistency

Prior work on streaming Transformers shows that models often rely on an attention sink [22, 52–54] to stabilize attention; without it, sliding past the first frame can trigger model collapse.

However, this anchoring creates a fragile reliance on the first frame. Over long sequences, it leads to geometric

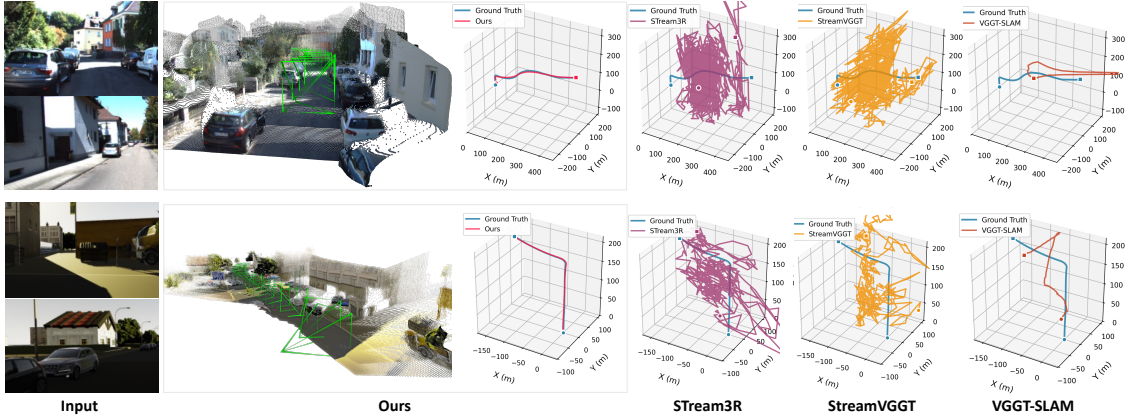


Figure 5. **Qualitative comparison on long-sequence pose estimation.** We compare LongStream against streaming and SLAM baselines on KITTI and vKITTI sequences spanning several hundred meters. Stream3R and StreamVGGT accumulate drift over long trajectories, and VGGT-SLAM runs out of memory on the second vKITTI sequence. LongStream preserves stable and coherent poses across all scenes, maintaining trajectory continuity even under large loop motions.

Methods	KITTI [51] (ATE ↓)											Avg.
	00 4542x, 3.7km	01 1101x, 2.5km	02 4661x, 5.1km	03 801x, 0.6km	04 271x, 0.4km	05 2761x, 2.2km	06 1101x, 1.2km	07 1101x, 0.7km	08 4071x, 3.2km	09 1591x, 1.7km	10 1201x, 0.9km	
FastVGGT	*	705.39	*	62.38	10.27	157.74	124.43	69.27	*	190.10	194.75	189.29
MAS3R-SLAM	*	530.37	*	18.87	88.98	159.430	92.00	*	*	*	*	177.93
VGGT-SLAM	*	607.16	*	169.83	13.12	*	*	*	*	*	*	263.37
CUT3R	185.89	651.52	296.98	148.06	22.17	155.61	132.54	77.03	238.39	205.94	193.39	209.78
TTT3R	190.93	546.84	218.77	105.28	11.62	153.12	132.94	70.95	180.57	211.01	133.00	177.73
SStream3R	190.98	681.95	301.40	158.25	102.73	159.85	135.03	90.37	261.15	216.31	207.49	227.77
StreamVGGT	191.93	653.06	303.35	157.50	108.24	160.46	133.71	89.00	263.95	216.69	209.80	226.15
Ours	92.55	46.01	134.70	3.81	1.95	84.69	23.12	14.93	62.07	85.61	21.48	51.90

Table 1. **Quantitative comparison on the KITTI dataset in terms of ATE.** The upper block lists optimization-based baselines, and the lower block reports streaming methods. Our approach achieves the best accuracy across all sequences.

saturation [53], manifested as unstable attention, keyframe jumps, and growing pose errors. These issues persist even with relative-pose supervision, indicating that sink anchoring itself imposes an asymmetric positional bias.

We argue that “short-horizon collapse” is not caused directly by removing the sink, but is a symptom of train–inference mismatch. We therefore introduce Cache-Consistent Training (CCT), introduced in Algorithm 1, which explicitly passes and trims the KV cache between chunks to align training and inference visibility. During training, we remove the constant sink token and use purely causal masking with a sliding window, while *explicitly passing and trimming* the KV cache between chunks so that cache visibility mirrors inference.

As shown in Figure 4, CCT makes the attention pattern mathematically equivalent between chunked training and frame-by-frame inference, forcing the model to operate in a pure sliding window without a persistent anchor and thereby removing sink dependence.

For ultra-long sequences, accumulated KV still yields long-term memory saturation and geometric drift. We thus adopt a periodic cache refresh that hard-marginalizes stale

context by resetting the sink frame and KV cache every N keyframes, akin to state marginalization in SLAM.

This retains geometric continuity while clearing degraded features, allowing periodic memory resets with no extra compute; because the entire model operates in a keyframe-relative coordinate system, the cache can be refreshed at any keyframe without breaking consistency or degrading accuracy.

Combining CCT with periodic cache refresh yields stable, generalizable streaming over thousands of frames, maintaining consistent geometric accuracy and well-behaved attention distributions.

4. Experiments

4.1. Implementation Details

Model configurations. We initialize LongStream from VGGT and retain its 24-layer backbone with alternating global and frame-level attention, containing roughly $1.3B$ parameters. During training, we use a fixed visibility layout, a consistent sliding window, and a keyframe interval of ten so that each batch contains multiple keyframe transitions.

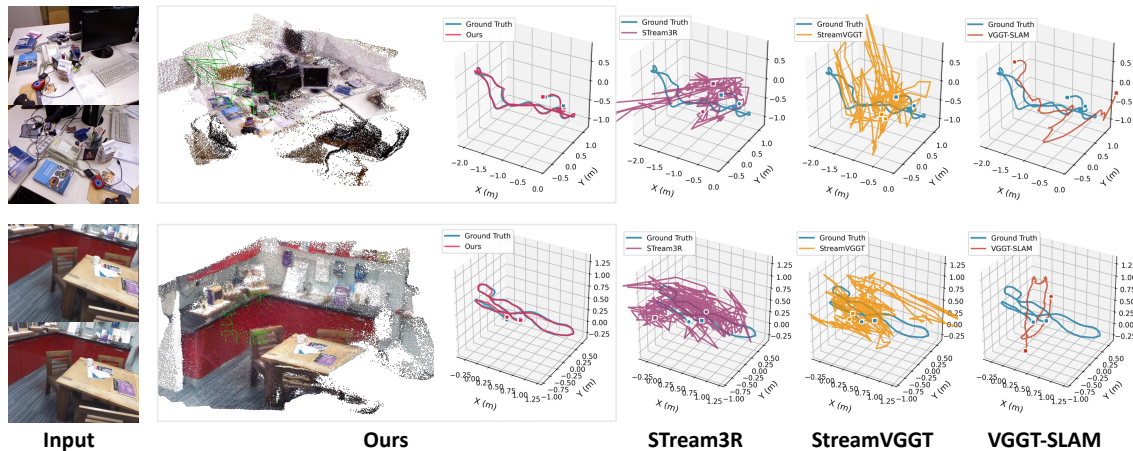


Figure 6. **Qualitative comparison on indoor sequences.** We evaluate challenging scenes with strong viewpoint changes, occlusions, and repeated back-tracking. While Stream3R, StreamVGGT, and VGGT-SLAM drift on these highly folded trajectories, LongStream maintains stable poses and consistent 3D structure throughout the sequence. The symbol * denotes runs with out-of-memory errors or more than three tracking failures.

Methods	TUM [55]	Oxford Spires [56]	Waymo [57]
	ATE ↓	ATE ↓	ATE ↓
FastVGGT	0.418	36.577	1.281
MASt3R-SLAM	0.082	37.728	7.625
VGGT-SLAM	0.123 (0.053 [†])	31.003	7.431
CUT3R	0.542	32.440	9.396
TTT3R	0.308	36.214	3.486
SStream3R	0.633	37.569	42.203
StreamVGGT	0.627	37.255	45.101
Ours	0.076	19.815	0.737

Table 2. **Quantitative comparison on TUM [55], Oxford Spires [56], and Waymo [57].** Top: optimization-based methods; Bottom: streaming methods. Our method demonstrates robustness on these small-scale trajectories, achieving the best performance across all online benchmarks. [†] Reported in [58].

Methods	vKITTI [59] (ATE ↓)					Avg.
	Scene 01 447×, 332m	Scene 02 223×, 113m	Scene 06 270×, 51m	Scene 18 339×, 254m	Scene 20 837×, 711m	
FastVGGT	3.435	0.311	0.120	2.050	101.667	31.427
MASt3R-SLAM	83.771	20.206	3.840	68.875	231.064	98.714
VGGT-SLAM	25.128	0.237	0.281	1.641	68.840	23.667
CUT3R	50.968	29.913	0.820	29.012	127.583	55.276
TTT3R	29.877	11.785	0.598	7.445	71.208	28.099
SStream3R	68.280	26.450	8.185	43.597	198.279	82.815
StreamVGGT	71.616	15.349	10.274	23.900	221.407	83.916
Ours	1.422	0.185	0.303	0.683	4.030	1.610

Table 3. **Quantitative comparison on vKITTI.** Top: optimization-based methods; Bottom: streaming methods. Our method achieves the best accuracy across all sequences.

We optimize with AdamW and cosine decay, using a peak learning rate of 4×10^{-6} , and a warmup of $1k$ steps. All images, depths, and pointmaps are resized to a maximum long side of 518 pixels, with aspect-ratio jittering, interval sampling, and cross-block shuffling.

Training method. We train LongStream with a two-

stage schedule: the first stage performs batch-independent training with batch size 22 for $50k$ iterations over three days on 32 A100 GPUs. The second stage applies KV-cache-consistent training that matches streaming inference by sampling sequence lengths between 10 and 80 with a cache window of 10. Metric scale supervision is used only when calibrated ground truth is available. At inference time, LongStream reaches 18 FPS on a single GPU. Further implementation details are available in the Appendix.

Training data. LongStream is trained on a multi-domain dataset collection, including Kubric [60], WildRGB [61], ScanNet [62], HyperSim [63], Mapillary [64], Replica [65], MVS-Synth [66], PointOdyssey [67], Virtual KITTI [59], Aria Synthetic Environments [68], Aria Digital Twin [68] Objaverse [69], Spring [70], and Waymo Open [57]. BlendedMVS [71], Co3Dv2 [72], MegaDepth [73], and DL3DV [74] do not provide metric ground truth and are therefore excluded from scale training, while all other datasets used in training provide metric supervision.

Baselines. We compare LongStream with offline transformers, streaming models, and SLAM-style systems. Because both VGGT [15] and π^3 [16] exceed memory limits on long clips, we report FastVGGT [75] as a practical replacement. Streaming baselines include CUT3R [45], TTT3R, SStream3R [1], and StreamVGGT [2]. MASt3R-SLAM [34] and VGGT-SLAM [58] represent optimization-based SLAM baselines. VGGT-SLAM performs chunk-wise inference over 16/32 frames and thus accesses future frames within each chunk. All methods are evaluated with official default settings under a unified protocol.

4.2. Quantitative Results

Camera pose estimation. We evaluate ATE on vKITTI [59] (training), Waymo [57] (held-out), and the un-

Methods	7Scenes		TUM	
	CD ↓	F1@0.25 ↑	CD ↓	F1@0.25 ↑
FastVGGT	6.373	0.710	<u>0.104</u>	<u>0.926</u>
MASt3R-SLAM	5.987	0.691	0.057	0.954
VGGT-SLAM	6.306	<u>0.696</u>	1.993	0.633
CUT3R	6.281	0.274	0.474	0.533
TTT3R	6.231	0.260	0.249	0.792
STream3R	<u>6.353</u>	0.479	1.126	0.444
StreamVGGT	6.630	<u>0.483</u>	0.680	0.402
Ours	2.260	0.641	0.225	0.673

Table 4. **Quantitative comparison on 7Scenes and TUM.** CD (lower) and F1@0.25 (higher) are adopted for evaluation. Best numbers are in **bold**; second best are underlined.

RelPose	Scale Head	CCT	Cache Refresh	ATE ↓	RPE ↓	Scale Err. ↓
×	×	×	×	8.043	2.207	-
✓	×	×	×	2.819	0.750	-
✓	✓	×	×	2.645	0.484	0.010
✓	✓	✓	×	0.984	0.454	0.032
✓	✓	✓	✓	0.115	0.126	0.035

Table 5. **Ablation study on RelPose, Scale head, CCT, and cache refresh.** Green indicates enabled, red indicates disabled. Rows 2 and 3 ATE gap is caused by a few large trajectory outliers. Scale Error reports absolute scale deviation; lower is better.

seen KITTI [35], TUM-RGBD [55], and Oxford Spires [56] datasets. As shown in Tables 1–3, streaming baselines exhibit nonlinear error growth as sequence length increases, likely due to long-horizon history saturation and contamination, whereas LongStream remains stable. Offline models frequently encounter out-of-memory (OOM) and tracking loss on long clips. Despite operating fully streaming, LongStream achieves state-of-the-art accuracy at 18 FPS and generalizes robustly across environments.

3D reconstruction. We evaluate full-sequence reconstruction on 7Scenes and TUM on test sequences without subsampling, reporting Chamfer Distance and F1@0.25. As shown in Table 4 and Figure 6, LongStream performs competitively with both offline approaches on both benchmarks. The small spatial extent and dense frame coverage of these datasets lead to largely saturated metrics. For 7Scenes, failure cases of some methods under all-frame evaluation produce CD values several orders of magnitude larger on certain scenes, leading to an inflated mean CD.

Scale estimation. We evaluate the accuracy of the recovered metric scale on vKITTI. LongStream produces a stable scale estimate across the entire sequence, achieving a scale ratio of 0.9905 with respect to ground truth.

4.3. Qualitative Results

Figures 5 and 6 visualize trajectories on kilometer-level and room-scale sequences, confirming stable pose predic-

tion under both large and small spatial extents. In outdoor settings (Figure 5), existing streaming methods such as STream3R and StreamVGGT suffer from accumulated drift over long trajectories, while the optimization-based VGGT-SLAM encounters memory limitations and tracking loss on longer sequences. In contrast, LongStream preserves trajectory continuity and metric accuracy across several hundred meters, successfully closing large loops without explicit loop closure modules.

Similarly, in indoor environments (Figure 6), the model demonstrates robustness against highly folded camera trajectories characterized by strong viewpoint changes, occlusions, and repeated back-tracking. Where baselines tend to produce unstable or drift-prone poses under these erratic motion patterns, LongStream maintains a coherent global trajectory.

4.4. Ablation Study

We conduct ablations on a single vKITTI sequence to validate LongStream’s four core components: the keyframe-relative pose head, the scale branch, cache-consistent training (CCT), and periodic cache refresh. As shown in Table 5, combining all modules reduces ATE from 8.043 to 0.115, nearly two orders of magnitude.

Gauge-decoupled pose and scale. Switching from absolute pose regression to our gauge-decoupled formulation provides the largest gain (Row 1 → Row 2), confirming that separating local geometry from global coordinates is essential for generalizing beyond the training window. The scale branch is required for metric consistency; removing it prevents the model from producing a stable global scale across the sequence.

Temporal cache consistency. The periodic cache refresh prevents long-term memory saturation in infinite streams, while the dedicated scale branch is essential for maintaining metric consistency.

Combining all components reduces the final ATE by nearly two orders of magnitude relative to the baseline. More detailed analyses are provided in the Appendix.

5. Conclusion

LongStream delivers stable, metric-scale reconstruction over ultra-long sequences, overcoming the drift and extrapolation failures of existing methods. Its gauge-decoupled pose design and cache-consistent training preserve consistent geometry and scale across thousands of frames, achieving strong accuracy and real-time performance on diverse indoor and outdoor datasets.

Limitations. The model still assumes a largely static world, relies on a heuristic keyframe schedule, and shows mild degradation in pointmap consistency over very long windows. These limitations suggest clear directions for improving robustness and generality in future work.

6. Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62406267), Guangdong Provincial Project (No. 2024QN11X072), Guangzhou-HKUST(GZ) Joint Funding Program (No. 2025A03J3956) and Guangzhou Municipal Education Project (No. 2024312122).

References

- [1] Yushi Lan, Yihang Luo, Fangzhou Hong, Shangchen Zhou, Honghua Chen, Zhaoyang Lyu, Shuai Yang, Bo Dai, Chen Change Loy, and Xingang Pan. Stream3r: Scalable sequential 3d reconstruction with causal transformer. *arXiv preprint arXiv:2508.10893*, 2025. 1, 2, 3, 4, 7
- [2] Dong Zhuo, Wenzhao Zheng, Jiahe Guo, Yuqi Wu, Jie Zhou, and Jiwen Lu. Streaming 4d visual geometry transformer. *arXiv preprint arXiv:2507.11539*, 2025. 1, 2, 3, 4, 7
- [3] Johannes L. Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2, 3
- [4] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 2, 3
- [5] Chong Cheng, Gaochao Song, Yiyang Yao, Qinzheng Zhou, Gangjian Zhang, and Hao Wang. Graph-guided scene reconstruction from images with 3d gaussian splatting, 2025. 3
- [6] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 785–801, Cham, 2018. Springer International Publishing. 2, 3
- [7] Gaochao Song, Chong Cheng, and Hao Wang. Gvxf: Gaussian voxel kernel functions for highly efficient surface reconstruction in open scenes, 2024.
- [8] Jiyuan Wang, Chunyu Lin, Lang Nie, Kang Liao, Shuwei Shao, and Yao Zhao. Digging into contrastive learning for robust depth estimation with diffusion models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, MM '24, page 4129–4137, New York, NY, USA, 2024. Association for Computing Machinery. 2
- [9] Wenhao Shen, Wanqi Yin, Hao Wang, Chen Wei, Zhonggang Cai, Lei Yang, and Guosheng Lin. Hmr-adapter: A lightweight adapter with dual-path cross augmentation for expressive human mesh recovery. In *Proceedings of the 32nd ACM International Conference on Multimedia*, MM '24, page 6093–6102, New York, NY, USA, 2024. Association for Computing Machinery. 2
- [10] Wenhao Shen, Wanqi Yin, Xiaofeng Yang, Cheng Chen, Chaoyue Song, Zhonggang Cai, Lei Yang, Hao Wang, and Guosheng Lin. Adhmr: Aligning diffusion-based human mesh recovery via direct preference optimization, 2025. 2
- [11] Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent mvsnet for high-resolution multi-view stereo depth inference. *Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 3
- [12] Jiyuan Wang, Chunyu Lin, Lang Nie, Shujun Huang, Yao Zhao, Xing Pan, and Rui Ai. Weatherdepth: Curriculum contrastive learning for self-supervised depth estimation under adverse weather conditions. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, page 4976–4982. IEEE, May 2024. 2
- [13] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 2, 3
- [14] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, pages 71–91. Springer, 2024. 3
- [15] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. 3, 4, 7
- [16] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He. π^3 : Permutation-equivariant visual geometry learning. *arXiv preprint arXiv:2507.13347*, 2025. 3, 7
- [17] Yu Hu, Chong Cheng, Sicheng Yu, Xiaoyang Guo, and Hao Wang. Vggt4d: Mining motion cues in visual geometry transformers for 4d scene reconstruction, 2025.
- [18] Jiyuan Wang, Chunyu Lin, Cheng Guan, Lang Nie, Jing He, Haodong Li, Kang Liao, and Yao Zhao. Jasmine: Harnessing diffusion prior for self-supervised depth estimation. *arXiv preprint arXiv:2503.15905*, 2025. 2
- [19] Ofir Press, Noah A. Smith, and Mike Lewis. Train short, test long: Attention with linear biases enables input length extrapolation, 2022. 2
- [20] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context, 2019. 2
- [21] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data, 2024. 2, 4, 5
- [22] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks, 2024. 2, 5
- [23] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015. 2, 3
- [24] Qiancheng Fu, Qingshan Xu, Yew-Soon Ong, and Wenbing Tao. Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction, 2022. 2, 3
- [25] David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19457–19467, 2024. 3

- [26] Jiale Xu, Shenghua Gao, and Ying Shan. Freesplatter: Pose-free gaussian splatting for sparse-view 3d reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 25442–25452, 2025.
- [27] 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, 2025.
- [28] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-lrm: Large reconstruction model for 3d gaussian splatting. In *European Conference on Computer Vision*, pages 1–19. Springer, 2024. 3
- [29] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019. 3
- [30] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh r-cnn. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9785–9795, 2019.
- [31] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019.
- [32] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [33] Chong Cheng, Yu Hu, Sicheng Yu, Beizhen Zhao, Zijian Wang, and Hao Wang. Reggs: Unposed sparse views gaussian splatting with 3dgs registration, 2025. 3
- [34] Riku Murai, Eric Dexeimer, and Andrew J Davison. Mast3r-slam: Real-time dense slam with 3d reconstruction priors. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16695–16705, 2025. 3, 7
- [35] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 8
- [36] Chong Cheng, Zijian Wang, Sicheng Yu, Yu Hu, Nanjie Yao, and Hao Wang. Unposed 3dgs reconstruction with probabilistic procrustes mapping, 2025.
- [37] Chong Cheng, Sicheng Yu, Zijian Wang, Yifan Zhou, and Hao Wang. Outdoor monocular slam with global scale-consistent 3d gaussian pointmaps, 2025.
- [38] Sicheng Yu, Chong Cheng, Yifan Zhou, Xiaojun Yang, and Hao Wang. Rgb-only gaussian splatting slam for unbounded outdoor scenes, 2025. 3
- [39] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *IEEE transactions on pattern analysis and machine intelligence*, 29(6):1052–1067, 2007. 3
- [40] Yuzheng Liu, Siyan Dong, Shuzhe Wang, Yingda Yin, Yanchao Yang, Qingnan Fan, and Baoquan Chen. Slam3r: Real-time dense scene reconstruction from monocular rgb videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16651–16662, 2025.
- [41] Zihan Zhu, Songyou Peng, Viktor Larsson, Zhaopeng Cui, Martin R Oswald, Andreas Geiger, and Marc Pollefeys. Nicer-slam: Neural implicit scene encoding for rgb slam. In *2024 International Conference on 3D Vision (3DV)*, pages 42–52. IEEE, 2024.
- [42] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*, pages 628–644. Springer, 2016.
- [43] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4578–4587, 2021.
- [44] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2021. 3
- [45] Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 10510–10522, 2025. 3, 7
- [46] Bill Triggs, Philip F. McLauchlan, Richard I. Hartley, and Andrew W. Fitzgibbon. Bundle adjustment — a modern synthesis. In Bill Triggs, Andrew Zisserman, and Richard Szeliski, editors, *Vision Algorithms: Theory and Practice*, pages 298–372, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg. 3
- [47] Zichao Zhang, Guillermo Gallego, and Davide Scaramuzza. On the comparison of gauge freedom handling in optimization-based visual-inertial state estimation. *IEEE Robotics and Automation Letters*, 3(3):2710–2717, 2018. 3
- [48] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2024. 4
- [49] Zachary Teed and Jia Deng. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *European Conference on Computer Vision (ECCV)*, 2020. 4
- [50] Nikhil Keetha, Norman Müller, Johannes Schönberger, Lorenzo Porzi, Yuchen Zhang, Tobias Fischer, Arno Knapitsch, Duncan Zauss, Ethan Weber, Nelson Antunes, Jonathon Luiten, Manuel Lopez-Antequera, Samuel Rota Bulò, Christian Richardt, Deva Ramanan, Sebastian Scherer, and Peter Kotschieder. Mapanything: Universal feed-forward metric 3d reconstruction, 2026. 4
- [51] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012. 6

- [52] Ruyi Xu, Guangxuan Xiao, Yukang Chen, Liuning He, Kelly Peng, Yao Lu, and Song Han. Streamingvlm: Real-time understanding for infinite video streams, 2025. 5
- [53] Shuai Yang, Wei Huang, Ruihang Chu, Yicheng Xiao, Yuyang Zhao, Xianbang Wang, Muyang Li, Enze Xie, Yingcong Chen, Yao Lu, Song Han, and Yukang Chen. Longlive: Real-time interactive long video generation, 2025. 6
- [54] Joonghyuk Shin, Zhengqi Li, Richard Zhang, Jun-Yan Zhu, Jaesik Park, Eli Shechtman, and Xun Huang. Motionstream: Real-time video generation with interactive motion controls, 2026. 5
- [55] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012. 7, 8
- [56] Yifu Tao, Miguel Ángel Muñoz-Bañón, Lintong Zhang, Jiahao Wang, Lanke Frank Tarimo Fu, and Maurice Fallon. The oxford spires dataset: Benchmarking large-scale lidar-visual localisation, reconstruction and radiance field methods, 2025. 7, 8
- [57] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Sheng Zhao, Shuyang Cheng, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset, 2020. 7
- [58] Dominic Maggio, Hyungtae Lim, and Luca Carlone. Vggt-slam: Dense rgb slam optimized on the sl(4) manifold, 2025. 7
- [59] Johann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2, 2020. 7
- [60] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanaprasam, Florian Golemo, Charles Herrmann, Thomas Kipf, Abhijit Kundu, Dmitry Lagun, Issam Laradji, Hsueh-Ti (Derek) Liu, Henning Meyer, Yishu Miao, Derek Nowrouzezahrai, Cengiz Oztireli, Etienne Pot, Noha Radwan, Daniel Rebain, Sara Sabour, Mehdi S. M. Sajjadi, Matan Sela, Vincent Sitzmann, Austin Stone, Deqing Sun, Suhani Vora, Ziyu Wang, Tianhao Wu, Kwang Moo Yi, Fangcheng Zhong, and Andrea Tagliasacchi. Kubric: a scalable dataset generator. In *CVPR*, 2022. 7
- [61] Hongchi Xia, Yang Fu, Sifei Liu, and Xiaolong Wang. Rgb-d objects in the wild: Scaling real-world 3d object learning from rgb-d videos, 2024. 7
- [62] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 7
- [63] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M. Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *International Conference on Computer Vision (ICCV) 2021*, 2021. 7
- [64] Manuel Lopez-Antequera, Pau Gargallo, Markus Hofinger, Samuel Rota Buló, Yubin Kuang, and Peter Kotschieder. Mapillary planet-scale depth dataset. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 7
- [65] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The replica dataset: A digital replica of indoor spaces, 2019. 7
- [66] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. Deepmvs: Learning multi-view stereopsis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 7
- [67] Yang Zheng, Adam W. Harley, Bokui Shen, Gordon Wetstein, and Leonidas J. Guibas. Pointodyssey: A large-scale synthetic dataset for long-term point tracking. In *ICCV*, 2023. 7
- [68] Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar Parkhi, Richard Newcombe, and Yuheng (Carl) Ren. Aria digital twin: A new benchmark dataset for egocentric 3d machine perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20133–20143, October 2023. 7
- [69] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. 7
- [70] Lukas Mehl, Jenny Schmalfuss, Azin Jahedi, Yaroslava Nalivayko, and Andrés Bruhn. Spring: A high-resolution high-detail dataset and benchmark for scene flow, optical flow and stereo, 2023. 7
- [71] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1790–1799, 2020. 7
- [72] Jeremy Reizenstein, Roman Shapovalov, Philipp Hertzler, 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. 7
- [73] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2041–2050, 2018. 7
- [74] Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. D3dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22160–22169, 2024. 7

- [75] You Shen, Zhipeng Zhang, Yansong Qu, Xiawu Zheng, Jiayi Ji, Shengchuan Zhang, and Liujuan Cao. Fastvggt: Training-free acceleration of visual geometry transformer, 2025. [7](#)