

IncEventGS: Pose-Free Gaussian Splatting from a Single Event Camera

Jian Huang^{1,2}Chengrui Dong^{1,2}Xuanhua Chen^{2,3}Peidong Liu^{2*}¹Zhejiang University²Westlake University³Northeastern University

{huangjian39, dongchengrui, chenxuanhua, liupeidong}@westlake.edu.cn

<https://github.com/WU-CVGL/IncEventGS>

Abstract

Implicit neural representation and explicit 3D Gaussian Splatting (3D-GS) for novel view synthesis have achieved remarkable progress with frame-based camera (e.g. RGB and RGB-D cameras) recently. Compared to frame-based camera, a novel type of bio-inspired visual sensor, i.e. event camera, has demonstrated advantages in high temporal resolution, high dynamic range, low power consumption, and low latency, which make it favored for many robotic applications. In this work, we present IncEventGS, an incremental 3D Gaussian Splatting reconstruction algorithm with a single event camera, without the assumption of known camera poses. To recover the 3D scene representation incrementally, we exploit the tracking and mapping paradigm of conventional SLAM pipelines for IncEventGS. Given the incoming event stream, the tracker first estimates an initial camera motion based on prior reconstructed 3D-GS scene representation. The mapper then jointly refines both the 3D scene representation and camera motion based on the previously estimated motion trajectory from the tracker. The experimental results demonstrate that IncEventGS delivers superior performance compared to prior NeRF-based methods and other related baselines, even if we do not have the ground-truth camera poses. Furthermore, our method can also deliver better performance compared to state-of-the-art event visual odometry methods in terms of camera motion estimation.

1. Introduction

Reconstructing accurate 3D scene representations from 2D images has been a long-standing challenge in computer vision and robotics, driving substantial efforts over the past few decades. Among those pioneering works, Neural Radiance Fields (NeRF) [19] and 3D Gaussian Splatting (3D-GS) [11], stand out for their utilization of differentiable rendering techniques, and have garnered significant attention due to their capability to recover high-quality 3D scene representation

from 2D images. Commonly used sensors for 3D scene reconstruction are usually frame-based cameras, such as the RGB and RGB-D cameras. They usually capture full-brightness intensity images within a short exposure time at a regular frequency. Due to the characteristic of this data-capturing process, they often suffer from motion blur or fail to capture accurate and informative intensity information under fast motion and low-light conditions, which would further affect the performance of downstream applications.

The event camera, a bio-inspired sensor, has gained significant attention in recent years for its potential to address the limitations of frame-based cameras under challenging conditions. Unlike conventional cameras, event cameras record brightness changes asynchronously at each pixel, emitting events when a predefined threshold is surpassed. This unique operation offers several advantages over conventional cameras, in terms of high temporal resolution, high dynamic range, low latency, and power consumption. Although event cameras have attractive characteristics for challenging environments, they cannot be directly integrated into existing frame-based 3D reconstruction algorithms that rely on processing dense 2D brightness intensity images.

Several pioneering works have been proposed to exploit event stream [3, 12, 23] to recover the motion trajectory and scene representation. While existing methods deliver impressive performance, they usually exploit 2.5D semi-dense depth maps to represent the 3D scene, and bundle adjustment (BA) is hardly performed, due to the asynchronous and sparse characteristics of event data stream. Klenk *et al.* [15] recently proposed to convert event stream into event voxel grids, and then adapt a previous frame-based deep visual odometry pipeline [30] for accurate camera motion estimation. As NeRF exhibited impressive scene representation capability recently, several works [8, 14, 17, 25] explore to recover the underlying dense 3D scene NeRF representation from event stream, by assuming ground-truth poses are available.

In contrast to those works, we propose IncEventGS, an incremental dense 3D scene reconstruction algorithm from a single event camera, by exploiting Gaussian Splatting as the underlying scene representation. Different from prior event-

*Corresponding author.

based NeRF reconstruction methods, *IncEventGS* does not require any ground-truth camera poses, which is more challenging and provides more flexibility for real-world robotic application scenarios. To overcome the challenges brought by unknown poses, *IncEventGS* adopts the tracking and mapping paradigm of conventional SLAM pipelines [20]. In particular, *IncEventGS* exploits prior explored and reconstructed 3D scenes for camera motion estimation of incoming event stream during the tracking stage. Both the 3D-GS scene representation and camera motions are then jointly optimized (*i.e.* event-based bundle adjustment) during the mapping stage, for more accurate scene representation and motion estimation. The 3D scene is progressively expanded and densified. The experimental results on both synthetic and real datasets demonstrate that *IncEventGS* can recover the underlying 3D scene representation and camera motion trajectory accurately. In particular, *IncEventGS* outperforms prior NeRF-based methods and other related baselines in terms of scene representation recovery, even without ground-truth poses. Furthermore, our method also delivers better camera motion estimation accuracy than the most recent state-of-the-art visual odometry algorithm, in terms of the Absolute Trajectory Error (ATE) metric. The recovered 3D scene representation can be further used to render novel brightness images. Our main contributions can be summarized as follows:

- We present an incremental 3D Gaussian Splatting reconstruction algorithm from a single event camera, without requiring the ground-truth camera poses.
- We propose a novel initialization strategy tailored to the event data stream, which is vital to the success of the algorithm.
- The experimental results on both the synthetic and real datasets demonstrate the superior performance of our method over prior NeRF-based methods and related baselines in terms of novel view synthesis and better performance over state-of-the-art event-based visual odometry algorithm in terms of camera motion estimation.

2. Related Works

We review two main areas of prior works: event-based neural radiance fields and 3D Gaussian Splatting, which are the most related to our work.

Event-based Neural Radiance Fields. Prior works [8, 14, 25] propose to exploit event stream to recover the neural radiance fields with known camera motion trajectory. Low *et al.* [17] further improves the reconstruction algorithm to handle sparse and noisy events under non-uniform motion. The recovered neural radiance fields can then be used to render novel view brightness images. The ground-truth poses are usually computed from corresponding brightness images via COLMAP [26] or provided by the indoor motion-capturing system. Recently, Qu *et al.* [22] proposed to integrate event measurements into an RGB-D implicit neural SLAM framework and achieved robust performance in mo-

tion blur scenarios. Li *et al.* [16] also proposed exploiting event measurements and a single blurry image to recover the underlying neural 3D scene representation. In contrast to those works, *IncEventGS* conducts incremental 3D scene reconstruction without requiring any prior ground-truth poses, which is more challenging and provides more flexibility for practical robotic application scenarios. The method further exploits 3D Gaussian Splatting as the underlying scene representation, which demonstrates better image rendering quality and efficiency, compared to the NeRF-based representation.

3D Gaussian Splatting. 3D Gaussian Splatting [11] proposes a novel explicit 3D representation to further improve both the training and rendering efficiency compared to Neural Radiance Fields. Due to its impressive efficient scene representation capability, several pioneering works have been proposed to exploit 3D-GS for incremental 3D reconstruction. For example, Keetha *et al.* [10] propose an RGBD-based 3D-GS SLAM, employing an online tracking and mapping system tailored to the underlying Gaussian representation. Yan *et al.* [36] implement a coarse-to-fine camera tracking approach based on the sparse selection of Gaussians. Matsuki *et al.* [18] propose to apply 3D Gaussian Splatting to do incremental 3D reconstruction using a single moving monocular or RGB-D camera. Huang *et al.* [7] exploit ORB-SLAM3 to compute accurate camera poses and feed it into a 3D-GS algorithm for dense mapping. Fu *et al.* [2] use monocular depth estimation with 3D-GS. Yugay *et al.* [37] combine DROID-SLAM [29] based camera tracking with active and inactive 3D-GS sub-maps. Wang *et al.* [33] integrate bundle-adjustment and 3DGS [38] to estimated camera trajectory within the exposure time. Hu *et al.* [6] propose a novel depth uncertainty model to ensure the selection of valuable Gaussian primitives during optimization. While those methods deliver impressive performance in terms of 3D scene recovery and motion estimation, they usually assume the usage of frame-based images (*i.e.* either RGB or RGB-D data). On the contrary, we propose to exploit pure event measurements for incremental 3D-GS reconstruction. Several concurrent studies have recently explored the use of 3D Gaussians for event-based reconstruction, including EvGGS [31], Event3DGS [35], and E2GS [1]. EvGGS and Event3DGS rely solely on event data, whereas E2GS incorporates both event data and blurry images. However, the key difference between these methods and ours is that they all rely on ground-truth poses.

3. Method

The overview of our *IncEventGS* is shown in Fig. 1. Given only a single event camera, *IncEventGS* incrementally performs tracking and dense mapping under the framework of 3D Gaussian Splatting, to recover both the camera motion trajectory and 3D scene representation simultaneously. Since event data are asynchronous, they cannot be directly integrated with the 3D-GS representation. We therefore process the event data stream into chunks according to a fixed time window.

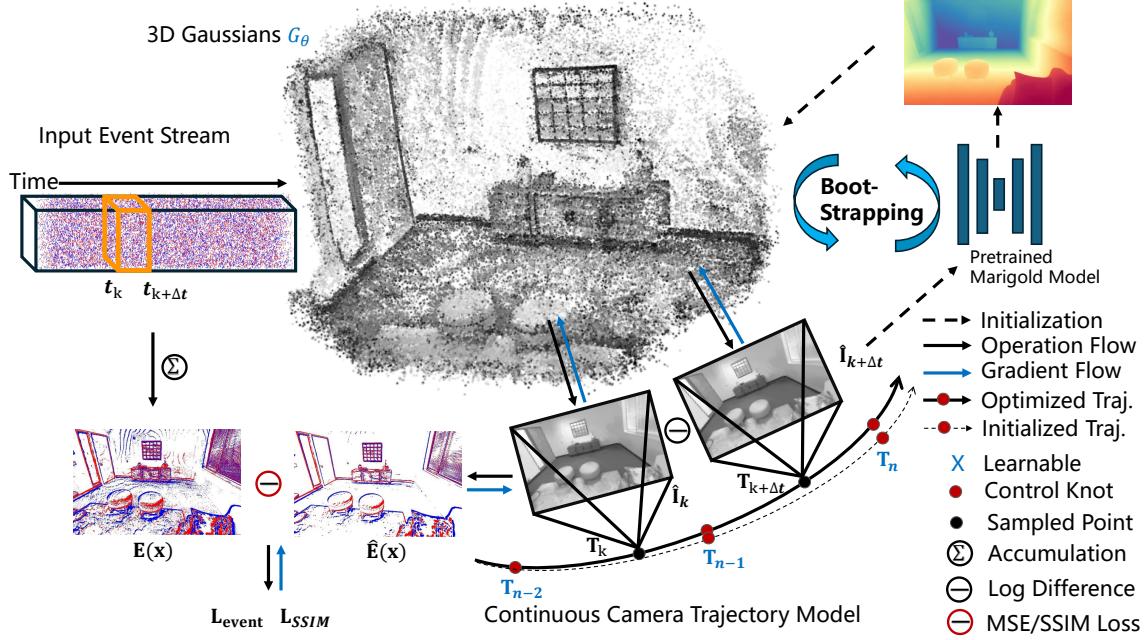


Figure 1. **The pipeline of IncEventGS.** IncEventGS processes incoming event stream by dividing it into chunks and representing the camera trajectory as a continuous model. It randomly samples two close consecutive timestamps to integrate the corresponding event streams. Two brightness images are rendered from 3D-GS at the corresponding poses, and we minimize the photometric loss between the synthesized and measured brightness change. During initialization, a pre-trained depth estimation model estimates depth from the rendered images to bootstrap the system.

We associate each chunk with a continuous time trajectory parameterization in the $\mathfrak{se}(3)$ space. Two close consecutive timestamps (i.e., t_k and $t_{k+\Delta t}$, where Δt is a small time interval) can be randomly sampled, and the measured brightness change $E(x)$ can be computed from the corresponding event stream. Based on the parameterized trajectory, the corresponding camera poses (i.e., T_k , $T_{k+\Delta t}$) can be determined and the brightness images (i.e., \hat{I}_k , $\hat{I}_{k+\Delta t}$) can be further rendered from the 3D-GS. The synthesized brightness change $\hat{E}(x)$ can be computed for event loss computation.

During the tracking stage, we optimize only the camera motion trajectory of the newly accumulated event chunk and exploit the recovered trajectory to initialize the dense bundle adjustment (BA) algorithm for the mapping stage. During the mapping stage, we continuously densify 3D Gaussians for newly explored areas and prune transparent 3D Gaussians. For computational efficiency, we exploit a sliding window of the latest chunks and perform BA only within this window for both 3D-GS reconstruction and motion trajectory estimation. We will detail each component as follows.

3.1. 3D Scene Representation

Following 3D-GS [11], the scene is represented by a set of 3D Gaussian primitives, each of which contains mean position $\mu \in \mathbb{R}^3$ in the world coordinate, 3D covariance $\Sigma \in \mathbb{R}^{3 \times 3}$, opacity $\mathbf{o} \in \mathbb{R}$, and color $\mathbf{c} \in \mathbb{R}^3$. To ensure that the covariance matrix remains positive semi-definite throughout the gradient descent, the covariance Σ is parameterized using a

scale vector $\mathbf{s} \in \mathbb{R}^3$ and rotation matrix $\mathbf{R} \in \mathbb{R}^{3 \times 3}$:

$$\Sigma = \mathbf{R} \mathbf{S} \mathbf{s}^T \mathbf{R}^T, \quad (1)$$

where scale matrix $\mathbf{S} = \text{diag}([\mathbf{s}])$ is derived from the scale vector $\mathbf{s} \in \mathbb{R}^3$.

In order to enable rendering, 3D-GS projects 3D Gaussian primitives to the 2D image plane from a given camera pose $\mathbf{T}_c = \{\mathbf{R}_c \in \mathbb{R}^{3 \times 3}, \mathbf{t}_c \in \mathbb{R}^3\}$ using following equation:

$$\Sigma' = \mathbf{J} \mathbf{R}_c \Sigma \mathbf{R}_c^T \mathbf{J}^T, \quad (2)$$

where $\Sigma' \in \mathbb{R}^{2 \times 2}$ is the 2D covariance matrix, $\mathbf{J} \in \mathbb{R}^{2 \times 3}$ is the Jacobian of the affine approximation of the projective transformation. After projecting 3D Gaussians onto the image plane, the color of each pixel is determined by sorting the Gaussians according to their depth and then applying near-to-far α -blending rendering via the following equation:

$$\mathbf{I} = \sum_i^N \mathbf{c}_i \alpha_i \prod_j^{i-1} (1 - \alpha_j), \quad (3)$$

where \mathbf{c}_i is the learnable color of each Gaussian, and α_i is the alpha value computed by evaluating the 2D covariance Σ' multiplied with the learned Gaussian opacity \mathbf{o} :

$$\alpha_i = \mathbf{o}_i \cdot \exp(-\sigma_i), \quad \sigma_i = \frac{1}{2} \Delta_i^T \Sigma'^{-1} \Delta_i, \quad (4)$$

where $\Delta_i \in \mathbb{R}^2$ is the offset between the pixel center and the

2D Gaussian center. Depth is rendered by:

$$\mathbf{D} = \sum_i^N \mathbf{d}_i \alpha_i \prod_j^{i-1} (1 - \alpha_j), \quad (5)$$

where d_i denotes the z-depth of the center of the i -th 3D Gaussian to the camera. We also render alpha map to determine visibility:

$$\mathbf{V} = \sum_i^N \alpha_i \prod_j^{i-1} (1 - \alpha_j), \quad (6)$$

The derivations presented above demonstrate that the rendered pixel color, denoted as \mathbf{I} in Eq. (3), is a function that is differentiable with respect to the learnable attributes of 3D-GS, and the camera poses \mathbf{T}_c . This facilitates our bundle adjustment formulation, accommodating a set of event chunks and inaccurate camera motion trajectories within the framework of 3D-GS.

3.2. Event Data Formation Model

An event camera records changes in brightness as a stream of events asynchronously. To relate 3D-GS representation with the event stream, we sample two close consecutive timestamps (i.e., t_k and $t_{k+\Delta t}$, where Δt is a small time interval), and the measured brightness change between Δt is:

$$\mathbf{E}(\mathbf{x}) = C \{e_i(\mathbf{x}, t_i, p_i)\}_{t_k < t_i < t_k + \Delta t}, \quad (7)$$

where $e(\mathbf{x}, t_i, p_i)$ is the i^{th} event within the defined time interval corresponding to pixel \mathbf{x} and C is the fixed contrast threshold. The corresponding camera poses \mathbf{T}_k and $\mathbf{T}_{k+\Delta t}$ can be interpolated from the camera motion trajectory parameterization, and two corresponding brightness images (i.e. $\hat{\mathbf{I}}_k$ and $\hat{\mathbf{I}}_{k+\Delta t}$) can be rendered from 3D-GS. The synthesized brightness change $\hat{\mathbf{E}}$ is modelled as:

$$\hat{\mathbf{E}}(\mathbf{x}) = \log(\hat{\mathbf{I}}_{k+\Delta t}(\mathbf{x})) - \log(\hat{\mathbf{I}}_k(\mathbf{x})), \quad (8)$$

where $\hat{\mathbf{E}}(\mathbf{x})$ depends on the parameters of both the motion trajectory parameters and 3D-GS, and is differentiable with respect to them.

Both in tracking and mapping, inspired by the work of [25], we segment the current event chunks into n_{seg} equal segments according to the number of events, obtaining n_{seg} timestamps that correspond to the end of each segment. We then randomly select one timestamp from these n_{seg} timestamps to serve as $t_{k+\Delta t}$, and we randomly sample an integer n_{win} between the integer bounds n_{low} and n_{up} . The index of t_k is equal to the index of $t_{k+\Delta t}$ subtract n_{win} . n_{seg} , n_{low} and n_{up} are hyperparameters. This sampling strategy enables the model to capture both local and global information.

3.3. Camera Motion Trajectory Modeling

Since each event chunk usually contains too many events, we sample a portion of them according to the total number of

events during optimization. Following [32, 38], we formulate the corresponding poses (i.e. \mathbf{T}_k and $\mathbf{T}_{k+\Delta t}$) at the beginning and end of the sampled event portion within each chunk, by employing a camera motion trajectory. The trajectory is represented through linear interpolation between two camera poses, one at the beginning of the chunk $\mathbf{T}_{\text{start}} \in \mathbf{SE}(3)$ and the other at the end $\mathbf{T}_{\text{end}} \in \mathbf{SE}(3)$. The camera pose at time t_k can thus be expressed as follows:

$$\mathbf{T}_k = \mathbf{T}_{\text{start}} \cdot \exp\left(\frac{t_k - t_{\text{start}}}{t_{\text{end}} - t_{\text{start}}} \cdot \log(\mathbf{T}_{\text{start}}^{-1} \cdot \mathbf{T}_{\text{end}})\right), \quad (9)$$

where t_{start} and t_{end} represent the timestamps corresponding to the boundary of the event chunk. It follows that \mathbf{T}_k is differentiable with respect to both $\mathbf{T}_{\text{start}}$ and \mathbf{T}_{end} . The objective of *IncEventGS* is thus to estimate both $\mathbf{T}_{\text{start}}$ and \mathbf{T}_{end} for each event chunk, along with the learnable parameters of 3D-GS \mathbf{G}_θ .

3.4. Incremental Tracking and Mapping

For both tracking and mapping, we aim to minimize the difference between the synthesized and measured brightness changes. In particular, we compute the loss of the latest event chunk only for the tracking stage and minimize the following loss function:

$$\mathbf{T}_{\text{start}}^*, \mathbf{T}_{\text{end}}^* = \underset{\mathbf{T}_{\text{start}}, \mathbf{T}_{\text{end}}}{\operatorname{argmin}} \left\| \mathbf{E}(\mathbf{x}) - \hat{\mathbf{E}}(\mathbf{x}) \right\|_2, \quad (10)$$

where $\hat{\mathbf{E}}(\mathbf{x})$ and $\mathbf{E}(\mathbf{x})$ are the synthesized and measured brightness changes respectively, corresponding to a randomly sampled event portion within the latest event chunk.

Once the tracking is done, we insert the latest event chunk to the mapper and exploit the estimated $\mathbf{T}_{\text{start}}^*$ and $\mathbf{T}_{\text{end}}^*$ as the initial value of the chunk to perform dense bundle adjustment. For computational consideration, we exploit a sliding window BA of the latest n_w chunks, and n_w is a hyperparameter. In particular, we optimize both the motion trajectories and the 3D-GS jointly by minimizing the following loss functions:

$$\mathcal{L} = (1 - \lambda) \mathcal{L}_{\text{event}} + \lambda \mathcal{L}_{\text{ssim}}, \quad (11)$$

$$\mathcal{L}_{\text{event}} = \left\| \mathbf{E}_i(\mathbf{x}) - \hat{\mathbf{E}}_i(\mathbf{x}) \right\|_2, \quad (12)$$

$$\mathcal{L}_{\text{ssim}} = \text{SSIM}(\mathbf{E}_i(\mathbf{x}), \hat{\mathbf{E}}_i(\mathbf{x})), \quad (13)$$

where λ is a hyperparameter, SSIM is the structural dissimilarity loss [34]. As the event data streams in, we alternatively perform tracking and mapping.

3.5. System Initialization and Boot-strapping

Conventional frame-based 3D-GS methods usually require a good initial point cloud and camera poses computed via COLMAP [27] for initialization. However, they are usually not easy to obtain if we are only given a single event camera. We therefore initialize the 3D-GS by sampling point cloud randomly within a bounding box. The first m event chunks

(where m is a hyperparameter) are selected for initialization, and all corresponding camera poses are randomly initialized to be near the identity matrix. We then minimize the loss computed by Eq. (11) with respect to the attributes of 3D-GS and the parameters of camera motion trajectories jointly.

Through experiments, we found that the above initialization procedure consistently produces satisfactory brightness images. However, the 3D structure remains of low quality due to the short baselines of the moving event camera. We further find that it could potentially affect the performance of the whole pipeline without a good initial 3D structure as more event data is received. Therefore, we utilize a monocular depth estimation network [9] to predict a dense depth map from the rendered brightness image after the pipeline is trained for certain iterations. This depth map is then used to re-initialize the centers of the 3D Gaussians by unprojecting the pixel depths, after which we repeat the minimization of Eq. (11) for system bootstrapping. More details about system re-initialization can be found in our supplementary material.

3D-GS Incrementally Growing. As the camera moves, new Gaussians are periodically introduced to cover newly explored regions. After tracking, we obtain an accurate camera pose estimate for each new event chunk. The centers of new Gaussians are determined by:

$$\mathbf{p} = \mathbf{T} \cdot \pi^{-1}(\mathbf{u}, d_u) \quad (14)$$

where $\mathbf{u} \in \mathbb{R}^2$ is pixel coordinate in the image plane, d_u is depth of the 3D point \mathbf{p} projecting onto the image plane, which is rendered by Eq. (5), π^{-1} denotes camera inverse projection, \mathbf{T} is the camera pose from tracking. To ensure that new Gaussians are only added in previously unmapped areas, a visibility mask is computed to guide the expansion of the Gaussian splatting process, as following:

$$\mathbf{M}(p) = V < \lambda_V, \quad (15)$$

where \mathbf{V} is the rendered alpha map and λ_V is a hyperparameter.

4. Experiments

4.1. Experimental Setups.

Implementation Details. All experiments are conducted on a desktop PC equipped with a 5.73GHz AMD Ryzen 9 7900x CPU and an NVIDIA RTX 3090 GPU. The first $m = 3$ event chunks are used for initialization. During the mapping stage, a sliding window size of $n_w = 20$ is employed for the bundle adjustment algorithm. The hyperparameters are set as follows: $\lambda = 0.05$, $\lambda_V = 0.8$, and $n_{seg} = 100$. For the synthetic dataset, $n_{low} = 400k$ and $n_{up} = 500k$, while for the real dataset, $n_{low} = 60k$ and $n_{up} = 80k$. Each event chunk has a time interval of 50 ms. The learning rate of the camera poses is set to 1e-4 and that for the attributes of 3D-GS are set the same as the original 3D-GS work. The number

of optimization steps for initialization is 4500, and that for tracking and mapping are set to 200 and 1500 respectively. The contrast threshold C of the event camera is set to 0.1 for synthetic datasets and 0.2 for real datasets empirically.

Baselines and Evaluation Metrics. To the best of our knowledge, there are no existing event-only NeRF or 3D-GS methods that do not rely on ground-truth poses, making direct comparisons challenging. Therefore, we conduct a thorough comparison of our method with several event-based NeRF approaches, including E-NeRF [14], EventNeRF [25], and Robust e-NeRF [17], as well as our custom implemented two-stage method (*i.e.* E2VID [24] + COLMAP [27] + 3DGS [11]). E-NeRF, EventNeRF, and Robust e-NeRF leverage implicit neural radiance fields for 3D scene representation, requiring ground-truth camera poses for accurate NeRF reconstruction. For E2VID + COLMAP + 3DGS, event data is first converted into brightness images using E2VID. The camera poses are then estimated from these images using COLMAP, and 3D-GS is trained with the generated images and poses. Both the quantitative and qualitative comparisons are performed on the synthetic dataset. Since there are no paired ground truth images for the real dataset, we only perform qualitative comparisons on the real dataset. In terms of motion trajectory evaluations, we use the publicly available state-of-the-art event-only visual odometry methods, *i.e.* DEVO (mono) [15] and ESVO2(stereo) [21], for comparison.

The metrics used for novel view synthesis (NVS) include the commonly employed PSNR, SSIM, and LPIPS. For motion trajectory evaluations, we utilize Absolute Trajectory Error (ATE). To ensure fair comparisons, we employ the evaluation code provided by EventNeRF to compute the NVS metrics, which applies a linear color transformation between predictions and ground truth. Additionally, we use the public EVO toolbox [5] to compute the trajectory metrics.

Benchmark Datasets. To properly evaluate the performance of NVS and motion trajectory estimation, we synthesized event data using the 3D scene models from the Replica dataset [28]. In particular, we exploit the *room0*, *room2*, *office0*, *office2*, and *office3* scenes. We render high frame rate RGB images at 1000 Hz with a resolution of 768x480 pixels. These images are then converted to grayscale, and the event data is generated via the events simulator [4]. The contrast threshold is set to 0.1. To simulate real-world camera motions, we exploit the same motion trajectories as that of NICE-SLAM [39] for data generation.

We use the event dataset provided by TUM-VIE [13] for real data evaluations, which is also used by E-NeRF and Robust e-NeRF. TUM-VIE captures the event datasets by a pair of Prophesee Gen4 HD event cameras with a resolution of 1280x720 pixels. We only use the left-event camera data for our experiment.



Figure 2. Qualitative evaluation of novel view image synthesis on the Replica dataset. The experimental results demonstrate that our method renders higher-quality images with fewer artifacts compared to event-based NeRF and two-stage approaches.

	room0			room2			office0			office2			office3		
	PSNR↑	SSIM↑	LPIPS↓												
E-NeRF	13.99	0.58	0.51	15.56	0.47	0.58	18.91	0.51	0.57	13.05	0.65	0.44	14.01	0.62	0.48
EventNeRF	17.29	0.62	0.39	16.02	0.54	0.64	18.90	0.43	0.62	15.18	0.66	0.45	16.77	0.73	0.33
Robust e-NeRF	17.26	0.84	0.18	16.43	0.50	0.52	18.93	0.52	0.56	16.81	0.81	0.25	19.22	0.84	0.18
E2VID+ COLMAP+3DGS	14.45	0.44	0.52	15.74	0.51	0.55	18.91	0.31	0.68	14.03	0.57	0.48	13.25	0.47	0.53
Ours	24.31	0.85	0.17	23.75	0.79	0.23	25.64	0.54	0.30	21.74	0.82	0.23	21.18	0.88	0.13

Table 1. NVS performance comparison on Replica dataset. All results are calculated using the same code and ground-truth images. The results demonstrate that our method outperforms NeRF-based and two-stage methods.

4.2. Ablation Study

We conduct ablation studies to confirm our design choices. In particular, we study the effect of a monocular depth estimation network for system bootstrapping and event slicing hyperparameters n_{low} , n_{up} . The experiments are conducted with the Replica dataset, and the results are shown in Table 3 and Table 4, respectively.

We found that depth initialization significantly impacts

pose estimation, reducing the Average Trajectory Error (ATE) from 1.534 cm to 0.064 cm. Additionally, this improvement in pose estimation leads to a slight enhancement in Novel View Synthesis (NVS) performance. These results verify the importance of using depth initialization during the bootstrapping stage.

We compare several combinations of hyperparameters n_{low} and n_{up} , which refer to the range of event slicing window

	room0	room2	office0	office2	office3	1d	3d	6dof	desk	desk2
DEVO	0.289	0.266	0.138	0.281	0.156	0.147	0.303	2.93	0.732	0.201
E2VID+COLMAP	17.93	59.96	105.19	18.414	17.28	4.268	16.90	9.88	21.57	10.13
ESVO2	-	-	-	-	-	0.337	1.066	0.587	1.147	2.506
Ours	0.046	0.067	0.045	0.046	0.054	0.115	0.298	0.251	0.231	0.129

Table 2. Pose accuracy (ATE, cm) on Replica and TUM-VIE datasets. The results demonstrate that our method delivers better performance in terms of camera motion estimation. Since ESVO2 is a stereo event camera method, it cannot run on the Replica dataset.

Setting	PSNR↑	SSIM↑	LPIPS↓	ATE
full	21.74	0.82	0.23	0.046
w/o	17.80	0.76	0.26	1.534

Table 3. Ablation Study about Depth Initialization. The unit of ATE is cm. The experimental results demonstrate the effectiveness of the initialization strategy. It not only improves the quality of rendered images, but also improves the accuracy of the camera motion estimation significantly.

Setting	PSNR↑	SSIM↑	LPIPS↓	ATE
1k-10k	16.07	0.64	0.46	0.167
10k-50k	18.41	0.72	0.33	0.079
80k-200k	20.99	0.79	0.25	0.079
400k-500k	21.74	0.82	0.23	0.046
500k-600k	20.95	0.79	0.23	0.050
600k-700k	18.06	0.75	0.28	0.214

Table 4. Ablation Study on Event Slice Window Size (Hyperparameters n_{low} and n_{up}). The unit of ATE is cm.

size. Table 4 demonstrates that both too small and too large window sizes negatively impact the performance of Novel View Synthesis (NVS) and pose estimation. Consequently, we select $n_{low} = 400k$ and $n_{up} = 500k$ for our experiments on the Replica dataset.

4.3. Quantitative Evaluations.

We conduct quantitative evaluations against event NeRF methods (E-NeRF, EventNeRF, and Robust e-NeRF) and our custom implemented two-stage method (*i.e.* E2VID + COLMAP + 3DGS) in terms of the quality of NVS and pose estimation performance.

The NVS performance is evaluated on Replica-dataset and the results are presented in Table 1. It is important to note that the metrics are lower than those typically observed in standard NeRF/3D-GS methods for RGB images, primarily due to the lack of absolute brightness supervision. Even though NeRF-based methods use ground truth poses for training, *IncEventGS* still significantly outperforms them, highlighting the advantages of our approach utilizing a 3D Gaussian representation. Additionally, our method greatly surpasses two-stage method that also employs 3D Gaussian representation, demonstrating superior pose estimation and the effectiveness of our bundle adjustment technique.

We evaluate pose estimation performance using the ATE

metric on both synthetic and real datasets, comparing our method with DEVO, ESVO2 and E2VID + COLMAP. The results, presented in Table 2, show that our method outperforms both baselines, validating the effectiveness of our incremental tracking and mapping technique.

4.4. Qualitative Evaluations.

We evaluate our method against event NeRF methods and two-stage method qualitatively in terms of novel view image synthesis, both on synthetic and real data. The results are presented in both Fig. 2 and Fig. 3. It demonstrates that our method can deliver better novel view images, while event NeRF methods and two-stage method render images with additional artifacts. Compared to NeRF-based methods, our approach demonstrates the advantage of *IncEventGS* by leveraging 3D Gaussian Splatting as the underlying scene representation. In contrast to two-stage method, our dense bundle adjustment optimizes both 3D Gaussian Splatting and camera pose using event data, whereas two-stage approaches tend to accumulate errors over time, as confirmed by the experimental results. We also provide representative visualization of ATE error mapped onto trajectories in Fig. 4, both on synthetic and real dataset. It demonstrates that *IncEventGS* is able to recover more accurate motion trajectories.

5. Conclusion

We present the first pose-free 3D Gaussian splatting reconstruction model based on event camera, *i.e.* *IncEventGS*, and is useful for real world applications. We adopt the tracking and mapping paradigm in conventional SLAM pipeline to do incremental motion estimation and 3D scene reconstruction simultaneously. To handle the continuous and asynchronous characteristics of event stream, we exploit a continuous trajectory model to model the event data formation process. The experimental results on both synthetic and real datasets demonstrate the superior performance of *IncEventGS* over prior state-of-the-art methods in terms of high-quality novel image synthesis and camera pose estimation.

Acknowledgements. This work was supported in part by NSFC under Grant 62202389, in part by a grant from the Westlake University-Muyuan Joint Research Institute, and in part by the Westlake Education Foundation.

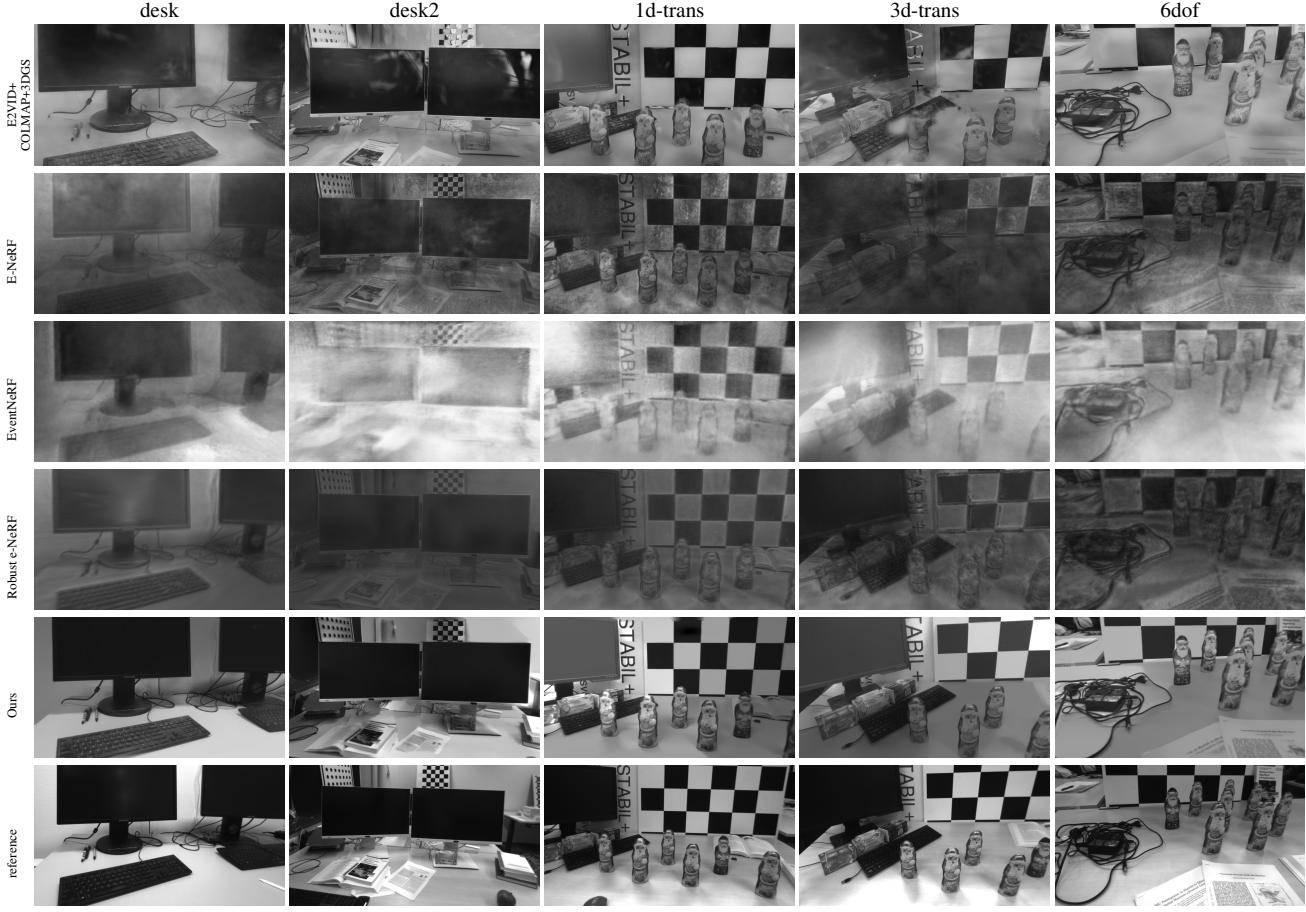


Figure 3. Qualitative evaluation for novel view image synthesis on real dataset. It demonstrates that our method is able to render better images with fewer artifacts than event NeRF methods and two-stage methods. Note that there are no GT images aligned with the event camera, and we choose the closest images from the RGB camera and crop them to the same size as the rendered images for visual comparisons.

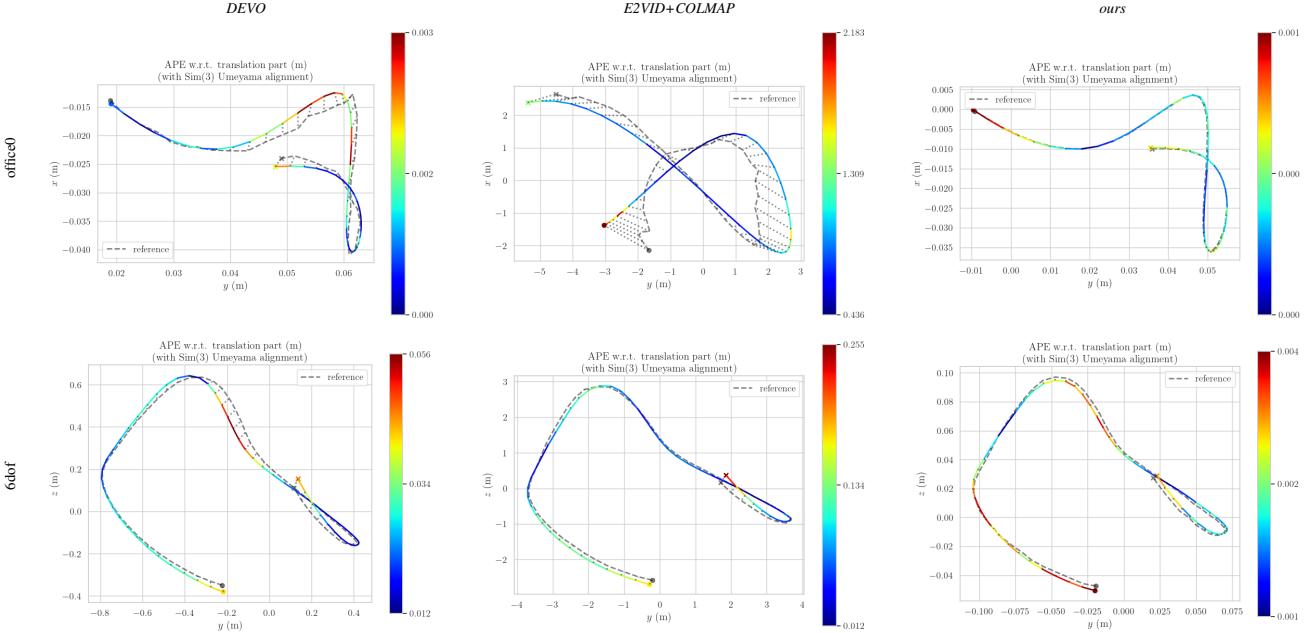


Figure 4. Representative visualization of ATE error mapped onto trajectories for the synthetic (office0) and real (6dof) datasets, generated by the EVO toolbox using the same ground truth poses, demonstrating the superior performance of our method in pose estimation.

References

[1] Hiroyuki Deguchi, Mana Masuda, Takuya Nakabayashi, and Hideo Saito. E2gs: Event enhanced gaussian splatting. In *2024 IEEE International Conference on Image Processing (ICIP)*, pages 1676–1682. IEEE, 2024. [2](#)

[2] Yang Fu, Sifei Liu, Amey Kulkarni, Jan Kautz, Alexei A Efros, and Xiaolong Wang. Colmap-free 3d gaussian splatting. *arXiv preprint arXiv:2312.07504*, 2023. [2](#)

[3] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth and optical flow estimation. In *CVPR*, 2018. [1](#)

[4] Daniel Gehrig, Mathias Gehrig, Javier Hidalgo-Carrió, and Davide Scaramuzza. Video to events: Recycling video datasets for event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3586–3595, 2020. [5](#)

[5] Michael Grupp. evo: Python package for the evaluation of odometry and slam. <https://github.com/MichaelGrupp/evo>, 2017. [5](#)

[6] Jiarui Hu, Xianhao Chen, Boyin Feng, Guanglin Li, Liangjing Yang, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Cg-slam: Efficient dense rgb-d slam in a consistent uncertainty-aware 3d gaussian field. *arXiv preprint arXiv:2403.16095*, 2024. [2](#)

[7] Huajian Huang, Longwei Li, Hui Cheng, and Sai-Kit Yeung. Photo-slam: Real-time simultaneous localization and photorealistic mapping for monocular, stereo, and rgb-d cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [2](#)

[8] Inwoo Hwang, Junho Kim, and Young Min Kim. Ev-nerf: Event based neural radiance field. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 837–847, 2023. [1, 2](#)

[9] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [5](#)

[10] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [2](#)

[11] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. [1, 2, 3, 5](#)

[12] Hanme Kim, Stefan Leutenegger, and Andrew J. Davison. Real-time 3d reconstruction and 6-dof tracking with an event camera. In *ECCV*, 2016. [1](#)

[13] Simon Klenk, Jason Chui, Nikolaus Demmel, and Daniel Cremers. Tum-vie: The tum stereo visual-inertial event dataset. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8601–8608. IEEE, 2021. [5](#)

[14] Simon Klenk, Lukas Koestler, Davide Scaramuzza, and Daniel Cremers. E-nerf: Neural radiance fields from a moving event camera. *IEEE Robotics and Automation Letters*, 8(3):1587–1594, 2023. [1, 2, 5](#)

[15] Simon Klenk, Marvin Motzett, Lukas Koestler, and Daniel Cremers. Deep event visual odometry. In *2024 International Conference on 3D Vision (3DV)*, pages 739–749. IEEE, 2024. [1, 5](#)

[16] Wenpu Li, Pian Wan, Peng Wang, Jinghang Li, Yi Zhou, and Peidong Liu. Beneref: Neural radiance fields from a single blurry image and event stream. In *European Conference on Computer Vision (ECCV)*, 2024. [2](#)

[17] Weng Fei Low and Gim Hee Lee. Robust e-nerf: Nerf from sparse and noisy events under non-uniform motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [1, 2, 5](#)

[18] Hidenobu Matsuki, Riku Murai, Paul HJ Kelly, and Andrew J Davison. Gaussian splatting slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18039–18048, 2024. [2](#)

[19] 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. [1](#)

[20] Raul Mur-Artal and Juan D Tardós. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE Transactions on Robotics*, 33(5):1255–1262, 2017. [2](#)

[21] Junkai Niu, Sheng Zhong, Xiuyuan Lu, Shaojie Shen, Guillermo Gallego, and Yi Zhou. Esvo2: Direct visual-inertial odometry with stereo event cameras. *IEEE Transactions on Robotics*, 2025. [5](#)

[22] Delin Qu, Chi Yan, Dong Wang, Jie Yin, Dan Xu, Bin Zhao, and Xuelong Li. Implicit event-rgb-d neural slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [2](#)

[23] Henri Rebecq, Timo Horstschaefer, Guillermo Gallego, and Davide Scaramuzza. EVO: A geometric approach to event-based 6-dof parallel tracking and mapping in real time. *IEEE Robotics and Automation Letters*, 2017. [1](#)

[24] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. High speed and high dynamic range video with an event camera. *IEEE transactions on pattern analysis and machine intelligence*, 43(6):1964–1980, 2019. [5](#)

[25] Viktor Rudnev, Mohamed Elgharib, Christian Theobalt, and Vladislav Golyanik. Eventnerf: Neural radiance fields from a single colour event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4992–5002, 2023. [1, 2, 4, 5](#)

[26] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion Revisited. In *CVPR*, pages 4104–4113, 2016. [2](#)

[27] Johannes Lutz Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [4, 5](#)

[28] 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. *arXiv preprint arXiv:1906.05797*, 2019. 5

[29] Zachary Teed and Jia Deng. DROID-SLAM: Deep Visual SLAM for Monocular, Stereo, and RGB-D Cameras. *Advances in neural information processing systems*, 2021. 2

[30] Zachary Teed, Lahav Lipson, and Jia Deng. Deep patch visual odometry. *Advances in Neural Information Processing Systems*, 2023. 1

[31] Jiaxu Wang, Junhao He, Ziyi Zhang, Mingyuan Sun, Jingkai Sun, and Renjing Xu. Evggs: A collaborative learning framework for event-based generalizable gaussian splatting. *arXiv preprint arXiv:2405.14959*, 2024. 2

[32] Peng Wang, Lingzhe Zhao, Ruijie Ma, and Peidong Liu. BAD-NeRF: Bundle Adjusted Deblur Neural Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4170–4179, 2023. 4

[33] Peng Wang, Lingzhe Zhao, Yin Zhang, Shiyu Zhao, and Peidong Liu. Mba-slam: Motion blur aware dense visual slam with radiance fields representation. *arXiv preprint arXiv:2411.08279*, 2024. 2

[34] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 4

[35] Tianyi Xiong, Jiayi Wu, Botao He, Cornelia Fermuller, Yiannis Aloimonos, Heng Huang, and Christopher A Metzler. Event3dgs: Event-based 3d gaussian splatting for fast ego-motion. *arXiv preprint arXiv:2406.02972*, 2024. 2

[36] Chi Yan, Delin Qu, Dong Wang, Dan Xu, Zhigang Wang, Bin Zhao, and Xuelong Li. Gs-slam: Dense visual slam with 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2

[37] Vladimir Yugay, Yue Li, Theo Gevers, and Martin R Oswald. Gaussian-slam: Photo-realistic dense slam with gaussian splatting. *arXiv preprint arXiv:2312.10070*, 2023. 2

[38] Lingzhe Zhao, Peng Wang, and Peidong Liu. Bad-gaussians: Bundle adjusted deblur gaussian splatting. In *European Conference on Computer Vision (ECCV)*, 2024. 2, 4

[39] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022. 5