

PRISM3D: Probabilistic Refinement and Robust Initialization for Physically Consistent Scene Modeling under Extreme Motion Blur

Gopi Raju Matta[✉], Reddypalli Trisha[✉], Vemunuri Divya Madhuri[✉], and Kaushik Mitra[✉]

Computational Imaging Lab, Dept. of Electrical Engineering, IIT Madras
{ee17d021, ee26s067, ee24d004}@smail.iitm.ac.in, kmitra@ee.iitm.ac.in

Abstract. We address the inverse problem of **blind 3D scene reconstruction** from extremely motion-blurred images, a scenario where traditional Structure-from-Motion (SfM) pipelines fail. Existing approaches typically circumvent this bottleneck by relying on impractical sharp-image supervision. In this work, we introduce **PRISM3D**, a unified framework enabling robust reconstruction directly from severely degraded inputs. To overcome the lack of a reliable starting point, we propose a **Robust Initialization** strategy utilizing deep dense tracking method (VGGSfM) to recover global topology where feature matching fails. To the best of our knowledge, we are the first to effectively leverage this paradigm to bootstrap 3D Gaussian Splatting from extreme motion blur. However, while robust, this initialization yields sparse and noisy geometry that causes deterministic optimization to diverge. To resolve this, we propose a coupled solution driven by probability and physics: we **adopt a probabilistic formulation** for geometric densification via Markov Chain Monte Carlo (MCMC) to robustly populate the sparse priors, while simultaneously modeling physical image formation via continuous **Bézier Trajectories**. Furthermore, while PRISM3D establishes a highly robust standalone pipeline, the availability of complementary event streams offers an opportunity to push the reconstruction fidelity further. To exploit this, we introduce **PRISM3D-E**, a multi-modal (RGB + Events) extension that seamlessly integrates high-temporal-resolution events as structural priors to maximize geometric recovery. Because existing datasets lack paired event streams under such severe degradation, we concurrently contribute the **PRISM3D-E Benchmark** to facilitate rigorous evaluation. Extensive experiments demonstrate that both our standalone RGB framework and its multi-modal extension establish new state-of-the-art performance. To facilitate future research in extreme motion deblurring, our project page, featuring extensive qualitative comparisons, video results, and the upcoming releases of our codebase and benchmark dataset, is available at <https://gopirajumatta.github.io/PRISM3D/>.

1 Introduction

Recovering high-fidelity 3D scenes from extremely motion-blurred images is a persistent inverse problem in computer vision. In real-world scenarios like robotics or handheld photography, rapid camera motion inevitably degrades image quality. Computational reconstruction requires solving a fundamental “chicken-and-egg” paradox: accurate deblurring requires precise camera poses,

yet estimating poses from blurred images requires a deblurred scene model to establish correspondences.

Despite the rapid evolution of Neural Radiance Fields (NeRF) [18] and 3D Gaussian Splatting (3DGS) [10], existing frameworks largely ignore the first half of this paradox. State-of-the-art methods typically rely on an “oracle” assumption: that accurate camera poses and point clouds are provided by Structure-from-Motion (SfM) pipelines like COLMAP [30] running on sharp images. This renders them impractical for blind deblurring. As demonstrated in our systematic analysis (Sec. 4.4), traditional SfM pipelines are brittle under extreme blur. Because blur acts as a low-pass filter, it obliterates the high-frequency features (e.g., corners, edges) required to establish reliable correspondences. Consequently, at moderate-to-extreme blur levels, standard initialization fails completely.

Recent advancements offer alternative paths to bypass this registration failure. Deep Structure-from-Motion (SfM) pipelines like VGGSfM [34] can successfully recover global topology directly from degraded inputs. Concurrently, feed-forward 3D Foundation Models (3DFMs) like VGGT [33] provide rapid scene initialization by directly predicting dense geometry. However, these dense but highly noisy point clouds trap initialization-sensitive 3DGS training in local minima [16]. While Deep SfM avoids these local minima by extracting a geometrically consistent, albeit sparse, topological skeleton, using it as a simple drop-in replacement for COLMAP remains insufficient. Its inherent geometric sparsity leads to inferior, artifact-prone reconstructions when paired with standard deterministic 3DGS optimization.

To bridge this **Initialization-Optimization Gap**, we introduce **PRISM3D**. We formulate blind deblurring as a joint optimization problem that explicitly addresses both the uncertainty of this deep blind initialization and the continuous dynamics of camera motion. Specifically, we tightly couple robust deep initialization with probabilistic scene densification (via MCMC) to gracefully populate the sparse geometry, while effectively capturing the physical blur image formation process using continuous Bézier trajectories. This yields a highly robust pipeline comprising three theoretically grounded components:

Robust Initialization via Deep Topology Recovery: We eliminate the dependency on brittle feature matching by leveraging VGGSfM [34]. Through a controlled evaluation of blur severity (1 to 11 frames), we demonstrate that VGGSfM maintains stable registration where COLMAP fails, providing the essential topological skeleton required to bootstrap optimization.

Probabilistic Scene Modeling (The MCMC Formulation): To handle the high variance of our deep initialization, we abandon heuristic densification in favor of a probabilistic formulation using 3DGS-MCMC [12]. We treat Gaussian primitives as samples from an underlying scene distribution rather than fixed parameters. By updating these samples via Stochastic Gradient Langevin Dynamics (SGLD), our framework robustly explores the solution space and “fills in” structural gaps, critically preventing local minima.

Physics-Consistent Trajectory Optimization: We solve the inverse image formation problem by modeling camera motion as a continuous $SE(3)$

Bézier trajectory. By integrating the scene radiance along this curved path, we mathematically formulate the blur as a differentiable function of the camera’s momentum. Jointly optimizing these trajectories with the probabilistic scene model ensures that the reconstruction is physically consistent with the exposure process [43].

Multi-Modal Degenerate Recovery: While PRISM3D is a highly robust standalone pipeline, the limits of reconstruction fidelity can be pushed further when complementary event sensor data is available. We introduce **PRISM3D-E**, a multi-modal (RGB + Events) extension designed for completely degenerate scenarios where photometric cues are nearly obliterated. PRISM3D-E leverages the microsecond resolution of event cameras via the Event-based Double Integral (EDI) model [21] to extract high-frequency structural priors during the exposure window. Routing these priors through our Deep SfM pipeline guarantees a strictly superior topological skeleton to bootstrap our MCMC densification. To facilitate rigorous evaluation of this extreme regime, we concurrently contribute the **PRISM3D-E Benchmark**.

Our contributions are as follows:

- We propose **PRISM3D**, a unified framework for solving blind deblurring by coupling **Robust Deep Initialization** (VGGSfM) and **Probabilistic Refinement** (MCMC) with continuous **Bézier Trajectories**. To our knowledge, this is the first work to effectively bootstrap 3DGS from extreme motion blur without sharp image supervision.
- We present a **systematic analysis of robustness** across varying blur levels, empirically validating the necessity of each core component. Specifically, we demonstrate that **Deep SfM** maintains reliable initialization where classical COLMAP fails, that our **Probabilistic MCMC** formulation is critical for resolving the geometric sparsity of deep tracking, and that continuous **Bézier trajectories** strictly outperform linear and cubic splines in modeling complex camera motion.
- We introduce the **PRISM3D-E Benchmark**, comprising a novel dataset that explicitly pairs extreme motion blur with complementary event streams, supported by standardized evaluation protocols and multi-modal baselines. Alongside this, we propose the **PRISM3D-E** multi-modal extension, demonstrating that routing event-derived priors through deep tracking can **further enhance reconstruction** in completely degenerate scenarios.
- Through extensive evaluations, we demonstrate that both our standalone framework (**PRISM3D**) and its multi-modal extension (**PRISM3D-E**) achieve state-of-the-art reconstruction quality, effectively bridging the gap between deep structure estimation and **physics-based differentiable rendering**.

2 Related Work

2.1 Robust Initialization in Degenerate Scenarios

Accurate pose estimation is a prerequisite for optimizing NeRF and 3D Gaussian Splatting (3DGS). However, standard SfM pipelines like COLMAP [30] and its

enhanced variants (GLOMAP [20], HLOC [28] with SuperGlue [29], Pixel-Perfect SfM [15]) rely on local feature correspondences, making them notoriously brittle under extreme motion blur.

Feed-forward 3D Foundation Models (3DFMs) like VGGT [33] and VGGT-X [16] offer rapid initialization, but scaling them to dense views yields noisy dense point clouds that trap initialization-sensitive 3DGS training in local minima. Conversely, deep dense tracking provides a geometrically rigorous alternative. By replacing discrete keypoint matching with differentiable tracking and bundle adjustment [9, 27, 31], VGGsFm [34] serves as a critical enabler for *blind* deblurring, recovering global topology where classical SfM and 3DFMs fail.

2.2 Probabilistic Scene Representation

While NeRF [18] and its implicit variants [1–3, 6, 8, 14, 19, 26, 32, 36, 38, 42] revolutionized view synthesis, they suffer from computationally expensive ray marching [17]. *3D Gaussian Splatting* (3DGS) [11] resolves this via efficient rasterization. However, standard 3DGS uses deterministic cloning and pruning heuristics assuming a dense, low-variance initialization (e.g., from COLMAP). Under the sparse initializations common in blind deblurring, these heuristics fail to meaningfully populate the scene, causing geometric collapse.

3DGS-MCMC [12] addresses this by reformulating Gaussians as probabilistic samples from a scene distribution, utilizing Markov Chain Monte Carlo for robust densification. We adopt this formulation to dynamically populate the sparse geometry from our deep SfM initialization, bridging the gap between noisy tracking and high-fidelity rendering.

2.3 Blind 3D Deblurring

Early NeRF-based deblurring methods like BAD-NeRF [35] converge slowly. While explicit 3DGS representations overcome this bottleneck, they remain limited by initialization strategies. ExBluRF [13] and Deblur-GS [4] assume impractical access to sharp reference images. Conversely, blind methods like BAD-Gaussians [43] attempt physical trajectory modeling but critically rely on standard COLMAP, failing on extremely blurred inputs.

To bypass photometric degradation, multi-modal approaches use supplementary hardware. Jang et al. [7] require active RGB-D sensors. Event-based NeRFs [23, 24] inherit implicit network bottlenecks, while recent event-assisted 3DGS methods (EvaGaussians [39], E2GS [5]) synthesize pseudo-sharp images purely to facilitate COLMAP initialization. Thus, they remain inapplicable in standard RGB-only scenarios and bottlenecked by classical SfM. In contrast, **PRISM3D** is an *event-agnostic*, standalone RGB solution capable of surviving extreme degradation. When complementary events *are* available, our multi-modal extension (**PRISM3D-E**) seamlessly integrates them to yield strictly superior geometry over SfM-dependent baselines.

3 Method

PRISM3D is a unified 3D Gaussian Splatting (3DGS) framework designed for robust novel view synthesis under extreme motion blur. To bridge the initialization-optimization gap, our architecture replaces brittle feature matching with deep




Fig. 1: Overview of the PRISM3D Framework. We address extreme motion blur through a unified inverse problem formulation. **(Top Left)** For the standalone **PRISM3D** pipeline, we input blurry images directly. For **PRISM3D-E**, we incorporate complementary event streams to generate high-frequency structural priors (EDI). **(Top Right)** These inputs feed into our **Robust Initialization** stage, which utilizes deep point tracking and bundle adjustment to recover a reliable sparse point cloud and camera poses, bypassing brittle classical feature matching. **(Bottom)** This sparse geometry bootstraps our **Probabilistic Refinement** module. We model the scene using 3D Gaussians updated via Stochastic Gradient Langevin Dynamics (SGLD flow, blue arrows) to robustly populate structural gaps. Concurrently, we ensure physical consistency by rendering latent sharp views along a continuous **Bézier Exposure Path**, averaging them to compute the blur modeling loss against the degraded input.

dense tracking via VGGSfM [34]. Because this deep initialization is inherently sparse, we couple it with a probabilistic MCMC-based densification strategy [12]. Finally, to ensure physical accuracy, we jointly optimize the scene geometry alongside continuous camera trajectories. Furthermore, we introduce **PRISM3D-E**, a multi-modal extension that leverages complementary event data, when available, to extract high-frequency structural priors and push reconstruction limits even further. Our complete framework is illustrated in Figure 1.

3.1 PRISM3D: Blind RGB Extreme Motion Deblurring

Standard 3DGS pipelines rely heavily on COLMAP [30] for initialization. However, under severe motion blur, the high-frequency features required for classical Structure-from-Motion (SfM) are obliterated, causing downstream deblurring methods (e.g., ExBluRF [13], Deblur-GS [4], and BAD-Gaussians [43]) to fail entirely. PRISM3D overcomes this catastrophic failure by explicitly coupling robust deep tracking with probabilistic geometric refinement.

Robust Deep Initialization and Probabilistic Refinement Accurate initialization is critical yet challenging when classical feature-based methods collapse.

We resolve this by coupling deep dense tracking for global registration with a probabilistic formulation for geometric densification.

Robust Initialization via VGGSfM: Unlike incremental SfM pipelines, VGGSfM [34] estimates global camera poses end-to-end. By leveraging a Transformer-based architecture for differentiable point tracking and bundle adjustment via the Theseus solver, it extracts reliable tracks even from severely degraded textures. As evaluated in Section 4.4, this yields valid initializations at extreme blur levels where traditional pipelines fail to establish correspondences.

Probabilistic Refinement via 3DGS-MCMC: While VGGSfM recovers global topology, the resulting point cloud is inherently sparse. Standard 3DGS heuristic cloning often fails to bridge these structural gaps. We therefore adopt a probabilistic formulation building upon 3DGS-MCMC [12], treating Gaussians as samples from an underlying scene distribution. Updated via Stochastic Gradient Langevin Dynamics (SGLD), these primitives dynamically relocate to high-likelihood regions, effectively populating the sparse deep priors without relying on rigid density control rules.

Synergistic Joint Optimization: The integration of deep tracking and MCMC creates a synergistic differentiable pipeline. While 3DGS-MCMC adaptively densifies the VGGSfM output, our joint optimization allows gradients to flow back through the continuous Bézier trajectories. This co-refinement ensures that both camera motion and scene geometry remain strictly consistent with the physical and photometric observations of the blurry input.

Physical Motion Blur Image Formation Model Motion blur arises from the accumulation of photons during exposure, expressed mathematically as a continuous integration over virtual latent sharp images $\mathbf{C}_t(\mathbf{u})$ (left), which we approximate using n discrete samples (right):

$$\mathbf{B}(\mathbf{u}) = \phi \int_0^\tau \mathbf{C}_t(\mathbf{u}) dt \quad (1) \quad \mathbf{B}(\mathbf{u}) \approx \frac{1}{n} \sum_{i=0}^{n-1} \mathbf{C}_i(\mathbf{u}) \quad (2)$$

Here, $\mathbf{B}(\mathbf{u})$ is the captured blurry image, ϕ is a normalization factor, and τ is the exposure time. Crucially, this discrete formulation ensures $\mathbf{B}(\mathbf{u})$ remains differentiable with respect to each virtual sharp image $\mathbf{C}_i(\mathbf{u})$, enabling joint optimization.

Continuous Camera Trajectory Modeling We parameterize camera poses in the $SE(3)$ Lie group. Because linear or cubic splines [43] are inadequate for modeling the complex physical momentum of severe camera shake, we adopt a continuous differentiable Bézier curve interpolation [4]. For a curve of degree M with control points T_j , the interpolated pose at time t is:

$$T_t = \prod_{j=0}^M \exp \left(\binom{M}{j} (1-u)^{M-j} u^j \cdot \log(T_j) \right) \quad (3)$$

where $u = t/\tau \in [0, 1]$. This mathematically ensures smooth, physics-consistent trajectory estimation for accurate deblurring.

Loss Functions and Joint Optimization Following [43], we jointly estimate the 3DGS parameters Θ (mean μ , covariance Σ , opacity o , color $\mathbf{c}$) and the Bézier control points $\mathbf{T}$ by minimizing a combined photometric loss against the captured blurry image $\mathbf{B}_k^{gt}(\mathbf{u})$:

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_1 + \lambda\mathcal{L}_{\text{D-SSIM}} \quad (4)$$

To optimize Θ and $\mathbf{T}$, we compute analytical Jacobians to ensure proper gradient flow [43] (for clarity, explicit dependence on the pixel coordinate $\mathbf{u}$ is omitted):

$$\frac{\partial \mathcal{L}}{\partial \Theta} = \sum_{k=0}^{K-1} \frac{\partial \mathcal{L}}{\partial \mathbf{B}_k} \cdot \frac{1}{n} \sum_{i=0}^{n-1} \frac{\partial \mathbf{B}_k}{\partial \mathbf{C}_i} \frac{\partial \mathbf{C}_i}{\partial \Theta} \quad (5) \quad \frac{\partial \mathcal{L}}{\partial \mathbf{T}} = \sum_{k=0}^{K-1} \frac{\partial \mathcal{L}}{\partial \mathbf{B}_k} \cdot \frac{1}{n} \sum_{i=0}^{n-1} \frac{\partial \mathbf{B}_k}{\partial \mathbf{C}_i} \frac{\partial \mathbf{C}_i}{\partial \Theta} \frac{\partial \Theta}{\partial \mathbf{T}} \quad (6)$$

Here, camera poses are represented as 6D vectors in their $SE(3)$ Lie algebra.

3.2 PRISM3D-E (Event-Assisted Extreme Motion 3D Deblurring)

While our standalone PRISM3D pipeline is highly capable, we introduce **PRISM3D-E** to achieve even higher reconstruction quality when complementary event data happens to be available. By leveraging these events in severely degenerate scenarios, we synthesize intermediate sharp images exclusively to bootstrap the VGGSfM initialization. Crucially, these event-derived images are never used for photometric supervision; the final probabilistic optimization remains strictly guided by the original blurry inputs and our physical blur model.




Fig. 2: The PRISM3D-E Benchmark: We introduce a synthetic event dataset designed for extreme motion blur. While RGB frames suffer from severe degradation, the complementary event streams preserve fine structural cues.

Event Generation and Formulation Event cameras asynchronously trigger an event $\mathbf{e}(x, y, \tau, p)$ when the logarithmic brightness change exceeds a threshold Θ :

$$p_{x,y,\tau} = \begin{cases} -1, & \log(\mathcal{I}_{x,y,\tau}) - \log(\mathcal{I}_{x,y,\tau-\Delta\tau}) < -\Theta \\ +1, & \log(\mathcal{I}_{x,y,\tau}) - \log(\mathcal{I}_{x,y,\tau-\Delta\tau}) > \Theta \end{cases} \quad (7)$$

where p is the polarity. To facilitate processing and integration with the EDI model, we group these asynchronous events uniformly into b discrete bins over the exposure time t_{exp} :

$$B_k = \{\mathbf{e}_i(x_i, y_i, \tau_i, p_i)\}_{t_{k-1} < \tau_i \leq t_k}, \quad \text{where } t_k = t_{\text{start}} + \frac{k}{b}t_{\text{exp}} \quad (8)$$

4 Experiments

4.1 Experimental Setup

Datasets: For evaluation, we use the synthetic dataset from ExBluRF [13], comprising eight diverse outdoor scenes with challenging camera motion (20 to 40 blurry training views and 4 to 6 test views per scene). Utilizing the provided latent sharp image sequences, we process the frames via ESIM [25] to generate corresponding high-temporal-resolution event streams. This creates our complementary **PRISM3D-E Benchmark** (Figure 2), which we will release publicly to support event-based deblurring research in extreme motion scenarios. The complete event sequences are provided in the supplementary material.

Additionally, we validate on the real-world E2NeRF dataset [23], captured using a DAVIS346 color event camera. It features five challenging scenes (i.e., *letter*, *lego*, *camera*, *plant*, *toys*) with complex textures and irregular hand-held motion (100ms exposure), resulting in severe blur and complex trajectories.

Baseline Methods and Evaluation Metrics: We compare PRISM3D against three baseline categories: (1) **Deep Single-Image Deblurring:** MPRNet [41] and Restormer [40]; (2) **Event-Based Reconstruction:** EDI [21], E2NeRF [23], and EBADNeRF [24]; and (3) **Oracle Baselines (Sharp Supervision):** ExBluRF* [13], Deblur-GS* [4], and BAD-Gaussians* [43]. While BAD-Gaussians is theoretically a blind method designed for blurred inputs, its critical dependence on COLMAP initialization renders it an "Oracle" baseline in our extreme blur regime, as classical SfM fails to provide the necessary poses and point clouds from degraded frames. Consequently, ExBluRF*, Deblur-GS*, and BAD-Gaussians* rely on the oracle assumption that accurate poses and point clouds are obtainable from sharp images. Since this is impossible in real-world blind deblurring, there are no direct comparisons with these oracle methods on the real dataset (Figure 4). We retain event-based methods as valid baselines because EDI-deblurred images enable COLMAP initialization even under severe blur. We evaluate reconstruction quality using **PSNR**, **SSIM**, and **LPIPS**.

Implementation Details: Our method is implemented in PyTorch [22] within the **3DGS-MCMC** [12] framework using the *gsplat* [37] pipeline. We optimize Gaussians and camera poses in $SE(3)$ **Bézier space** (using 9 control points) with the Adam optimizer. Gaussian learning rates follow the original 3DGS [11], while the pose learning rate is 1×10^{-3} . We use 15 virtual camera samples (n in Eq. 2) to balance performance and efficiency. For PRISM3D-E, we use 13 event bins for event-based deblurring (EDI). All models are trained on an NVIDIA RTX 4090 GPU using 7k iterations with a data downsampling factor of 2.

4.2 Quantitative Results

Reconstruction Quality: Table 1 organizes methods into three categories for fair comparison: (1) **Blind RGB** (w/o Events), (2) **Event-Assisted** (w/ Events), and (3) **Oracle/Impractical** (w/ Sharp Supervision).

Operating solely on motion-blurred images, **PRISM3D** outperforms all competing blind methods and several event-based approaches, proving the efficacy

Table 1: Quantitative comparisons for sharp novel view synthesis on the PRISM3D-E Benchmark (Synthetic). The table is grouped by input modality: (1) **Blind RGB** (w/o Events), (2) **Event-Assisted** (w/ Events), and (3) **Oracle Baselines** (w/ Sharp Supervision). Group 3 methods (ExBluRF*, BAD-Gaussians*, Deblur-GS*) are included only for reference as they rely on impractical sharp-image initialization. Best and second-best results within ranked groups are highlighted in green and orange. *Note: For brevity, we display three representative scenes matching our qualitative comparisons (Figure 3); however, the **Average** row is computed over the entire 8-scene dataset. The full per-scene breakdown is provided in the supplementary material.*

Scene	Metric	Blind RGB (w/o Events)			Event-Assisted (w/ Events)				Oracle Baselines(*)		
		MPRNet	Restormer	PRISM3D	EDI+3DGS	E2NeRF	EBAD-NeRF	PRISM3D-E	ExBluRF*	BAD-Gaussians*	Deblur-GS*
Bench	PSNR↑	25.35	26.39	29.86	28.95	25.41	28.15	33.55	31.93	32.54	33.26
	SSIM↑	0.678	0.720	0.841	0.865	0.708	0.822	0.924	0.877	0.901	0.929
	LPIPS↓	0.425	0.356	0.118	0.201	0.438	0.172	0.063	0.111	0.046	0.116
Camellia	PSNR↑	24.84	25.14	28.56	22.46	28.07	24.33	29.47	28.02	28.83	29.14
	SSIM↑	0.669	0.690	0.821	0.762	0.721	0.743	0.873	0.715	0.815	0.874
	LPIPS↓	0.395	0.351	0.129	0.271	0.329	0.192	0.108	0.313	0.099	0.158
Stone Lantern	PSNR↑	24.97	26.68	28.29	26.48	30.47	26.29	29.43	28.24	28.29	28.59
	SSIM↑	0.785	0.831	0.849	0.825	0.836	0.802	0.894	0.765	0.843	0.863
	LPIPS↓	0.342	0.280	0.195	0.270	0.324	0.264	0.152	0.236	0.143	0.236
Average (All 8)	PSNR↑	26.19	26.83	29.49	27.55	29.35	28.36	31.95	31.15	30.95	31.43
	SSIM↑	0.733	0.756	0.840	0.830	0.771	0.833	0.908	0.834	0.863	0.895
	LPIPS↓	0.368	0.324	0.153	0.232	0.352	0.183	0.097	0.162	0.089	0.178

of probabilistic MCMC refinement for sparse initialization. When event data is available, **PRISM3D-E** achieves an average **2.5 dB PSNR gain** over one of the state-of-the-art event-based methods, E2NeRF. Notably, **PRISM3D-E delivers 0.5 - 0.8 dB PSNR improvement over sharp-supervised oracle baselines (ExBluRF*, BAD-Gaussians*, Deblur-GS*)** despite their privileged sharp image access. This highlights the immense strength of integrating event-based priors, VGGsFM initialization, and 3DGS-MCMC joint optimization.

Table 2: Performance comparison between PRISM3D and BAD-Gaussians (SOTA) across various blur levels.

Blur Level	PRISM3D (Ours)			BAD-Gaussians		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Sharp	41.42	0.975	0.016	41.02	0.974	0.017
3	37.35	0.949	0.038	37.14	0.946	0.036
5	36.70	0.942	0.032	33.72	0.901	0.069
7	35.33	0.928	0.037	30.99	0.842	0.111
9	35.09	0.924	0.035	x	x	x
11	32.22	0.897	0.093	x	x	x

Robustness vs. BAD-Gaussians:

As shown in Table 2, while BAD-Gaussians applies COLMAP directly to blurred inputs, their approach is primarily evaluated on mild motion blur. Relying on classical SfM becomes a bottleneck under extreme degradation. Because severe blur smears the high-frequency details (e.g., sharp corners and edges) required for reliable feature matching, COLMAP cannot establish correspondences and fails completely beyond 7 frames. In contrast, PRISM3D’s deep tracking initialization allows it to maintain robust reconstruction up to a severe 11-frame blur.

Why PRISM3D-E Outperforms Oracle Baselines: Our framework couples robust initialization with physics-consistent trajectory modeling. **MCMC refinement** is more resilient to pose errors than the deterministic heuristics used in BAD-Gaussians. Furthermore, our continuous **Bézier trajectory model** (Figure 8) better approximates real-world motion blur than linear or cubic splines. EDI-deblurred images also allow VGGsFM to estimate highly accurate camera poses (**mean APE of 0.0862**), which is crucial for reliable 3D reconstruc-

tion. Finally, unlike ExBluRF, whose voxel-based representation struggles with high-frequency content, our bundle-adjusted Gaussian representation maintains superior geometric fidelity, allowing PRISM3D-E to surpass oracle baselines without sharp supervision.

Training Time and GPU Memory Consumption:

Evaluated on the real dataset (Table 3), PRISM3D completes optimization in just **7 minutes** per scene. We exclude explicit baselines (e.g., BAD-Gaussians) here, as their COLMAP dependency causes initialization failure on these severely degraded sequences. Conversely, surviving event-based implicit baselines (EBAD-NeRF, E2NeRF) require 7–15 hours and ~ 15 GiB of GPU memory. In contrast, PRISM3D consumes only **~ 1.55 GiB**, confirming the extreme efficiency and scalability of our representation.

Table 3: System Efficiency (Real Dataset). Average training time and GPU memory consumption.

Method	Time (hh:mm:ss)↓	Memory (MiB)↓
EBAD-NeRF	06:44:00	14,836
E2NeRF	14:58:00	15,532
PRISM3D (Ours)	00:07:13	1,584

4.3 Qualitative Results

Qualitative comparisons on synthetic and real datasets are presented in Figures 3 and 4. PRISM3D and PRISM3D-E consistently produce sharper reconstructions with fewer artifacts than existing approaches. In synthetic scenes like *Bench*, *Camellia* and *Stone Lantern*, competing methods over-smooth fine structures or introduce distortions (**red arrows**, Figure 3), whereas PRISM3D and its multi-modal extension, PRISM3D-E faithfully restores object fine details.




Fig. 3: Results on the Synthetic Dataset: PRISM3D effectively resolves extreme motion blur, preserving fine structural details and minimizing artifacts compared to existing methods. Note that ExBluRF*, BAD-Gaussians*, and Deblur-GS* utilize an oracle initialization (sharp images for poses and point clouds), whereas our method operates blindly.

In real-world scenarios (Figure 4), our method achieves superior clarity in high-frequency regions, recovering legible text on the **CVPR 2023 poster** and sharp specular highlights. While prior event-based methods (E2NeRF, EBAD-NeRF) suffer from color distortions and residual blur, PRISM3D and its multi-modal




Fig. 4: Results on the Real Dataset: PRISM3D reconstructs sharp and high-quality images from severely motion-blurred real-world inputs. In contrast, existing methods struggle with artifacts, noise, loss of fine details, and text degradation. Our framework effectively restores textures and structural consistency, as evident in the insets.

extension preserve accurate color distributions and fine-grained textures without ghosting artifacts.

4.4 Robustness of Modules across Various Blur Levels

To comprehensively evaluate our framework, we assess the robustness of each module across progressive motion blur levels before conducting targeted analyses in extreme blur scenarios to isolate individual contributions.




Fig. 5: Synthetic Motion Blur Levels: Images generated by averaging 1 (sharp), 3, 5, 7, 9, and 11 consecutive frames from an 11-image burst at each viewpoint, illustrating the increasing severity of motion blur used for evaluation.

VGGSfM Robustness: To evaluate VGGSfM’s robustness, we simulated increasing blur severity by averaging 1, 3, 5, 7, 9, and 11 consecutive frames from sharp burst sequences (Figure 5). While COLMAP fails to register images beyond moderate blur, VGGSfM consistently extracts stable poses and dense point clouds even under extreme degradation (Table 4). Translating this to quantitative accuracy, VGGSfM maintains highly reliable camera poses compared to ground truth (Table 5), demonstrating its superiority.


MCMC Robustness: We conducted two experiments to validate 3DGS-MCMC. First, we evaluated its resilience to noisy blur point cloud initializations by feeding it blur-degraded point clouds from VGGSfM (e.g., blur level 3, blur level 5), while giving the poses estimated from ground truth sharp images. As shown in Figure 6, standard 3DGS suffers significant performance drops as blur increases, whereas 3DGS-MCMC maintains stable reconstruction quality, proving the robustness of its probabilistic sampling framework towards varying blur level point clouds.

Table 4: SfM comparison: Performance across various blur levels.

Blur	VGGSfM		COLMAP	
	#Pts	#Imgs	#Pts	#Imgs
Sharp	23498	20	4987	20
3	22873	20	1257	20
5	21660	20	606	20
7	18518	20	246	15
9	21351	20	x	x
11	20283	20	x	x

**Fig. 6:** MCMC robustness to various blur-corrupted point cloud initializations: PSNR comparison of 3DGS-MCMC with 3DGS across various blur point cloud initializations obtained from VGGSfM.**Table 5:** VGGSfM Pose translation error statistics (in meters) with respect to ground truth.

Blur	RMSE ↓	Mean ↓	Med. ↓	Std ↓	Max ↓
3	0.18	0.16	0.13	0.09	0.35
5	0.32	0.27	0.20	0.18	0.65
7	0.40	0.36	0.29	0.17	0.75
9	0.61	0.49	0.36	0.36	1.43
11	0.72	0.52	0.35	0.50	2.24

**Fig. 7:** MCMC robustness to various blur levels: PSNR comparison of PRISM3D for deblurring with and without MCMC across various blur levels, SfM initialization are obtained from VGGSfM.

Second, we assessed end-to-end deblurring performance within PRISM3D across varying blur levels. As shown in Figure 7, PRISM3D equipped with MCMC consistently achieves superior PSNR. Furthermore, it successfully suppresses the noticeable artifacts and structural inconsistencies present in the non-MCMC variant (see supplementary material for visual comparisons), highlighting its practical value for reliable scene recovery.

4.5 Virtual Camera Count & Trajectory Representation

As illustrated in Fig. 8, we analyze the number of interpolated virtual camera poses ($n \in \{5, 10, 15, 20\}$) alongside various trajectory models. Across all configurations, increasing n from 5 to 15 yields substantial improvements, beyond which performance saturates. We thus adopt $n = 15$ to balance photometric accuracy with computational efficiency.

Comparing trajectory models, naive linear interpolation consistently underperforms, confirming the non-linear nature of real-world motion blur. While cubic B-splines provide a strong baseline, our continuous $SE(3)$ Bézier formulation consistently achieves the highest performance. This superiority is particularly pronounced in the multi-modal **PRISM3D-E** setting, where the Bézier trajectory outperforms splines by 0.59 dB in PSNR at $n = 15$. This validates that modeling camera momentum via continuous Bézier curves is a fundamental requirement for maximizing geometric recovery across sensor modalities.




Fig. 8: Impact of trajectory representations and virtual camera count. We evaluate linear, spline, and Bézier interpolations across varying discrete trajectory steps (n). **Left:** Standalone PRISM3D. **Right:** Multi-modal PRISM3D-E. In both regimes, reconstruction quality optimizes at $n = 15$. Linear trajectories fail to capture complex intra-exposure motion, while our continuous $SE(3)$ Bézier formulation consistently outperforms cubic splines, proving its necessity for physics-consistent deblurring.

4.6 Ablations

Component-wise Analysis for Extreme Motion Blur To quantify the contribution of each module in PRISM3D and PRISM3D-E, we performed ablation studies on the synthetic dataset (Table 6). Our analysis (Table 4) shows that classical COLMAP fails entirely under severe motion blur. While the deep feature matcher HLOC successfully runs, it lags behind VGGsFM by an average of 1.45 dB in PSNR. Furthermore, even recent 3D foundation models like VGGT-X struggle in this extreme regime, underperforming our initialization by 1.67 dB. This firmly highlights VGGsFM’s superior capability in extracting robust global topology from severe degradation. The inclusion of MCMC-based probabilistic refinement contributes an additional 0.47 dB in PSNR. We also validate our physics-consistent blur formation model. Assuming a naive linear camera trajectory results in a steep 1.55 dB drop. Our continuous $SE(3)$ Bézier formulation outperforms standard spline interpolation by 0.24 dB in the standalone setting. Crucially, when high-temporal-resolution event priors (EDI) are integrated, the necessity of accurate trajectory modeling becomes even more pronounced: our Bézier formulation outperforms linear and spline interpolations by 2.18 dB and 0.59 dB, respectively. Overall, integrating these event-based priors provides the most significant boost, yielding a 2.46 dB PSNR improvement over the standalone PRISM3D.

Why Deep SfM over 3D Foundation Models?. As highlighted by recent studies on 3D Foundation Models (3DFMs) like VGGT [33], scaling feed-forward models to dense views yields imperfect outputs with higher noise levels [16]. For initialization-sensitive representations like 3DGS, this noise undermines the learning of 3D primitives and traps the highly non-convex optimization process in local minima, leading to significant degradation in rendering quality [16]. Our empirical analysis directly confirms this fundamental limitation under extreme blur degradation. As shown in Table 6, substituting the robust tracking initialization of PRISM3D with VGGT-X causes a severe drop in downstream deblurring quality (from 29.49 dB down to 27.82 dB). This 1.67 dB performance penalty reinforces that the feed-forward geometry of 3DFMs is currently insufficient for

blind deblurring. Conversely, because VGGsFm is a dedicated deep SfM pipeline, its intrinsic bundle adjustment and strict multi-view constraints successfully extract the clean, geometrically consistent topological skeleton strictly necessary to guide our probabilistic densification.

Table 6: Ablation study. Average metrics across all 8 synthetic scenes. The inclusion of each module (VGGsFm, MCMC, EDI) and the use of Bézier trajectories yield synergistic improvements.

Variant	Configuration				Metrics		
	SfM	MCMC	Traj.	EDI	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
w/ HLOC	HLOC	✓	Bézier	×	28.04	0.761	0.148
w/ VGGT-X	VGGT-X	✓	Bézier	×	27.82	0.749	0.181
w/o MCMC	VGGsFm	×	Bézier	×	29.02	0.806	0.133
w/ Linear	VGGsFm	✓	Linear	×	27.94	0.790	0.190
w/ Spline	VGGsFm	✓	Spline	×	29.25	0.840	0.150
PRISM3D	VGGsFm	✓	Bézier	×	29.49	0.840	0.153
w/ COLMAP	COLMAP	✓	Bézier	✓	31.28	0.886	0.139
w/o MCMC	VGGsFm	×	Bézier	✓	31.06	0.869	0.083
w/ Linear	VGGsFm	✓	Linear	✓	29.77	0.860	0.150
w/ Spline	VGGsFm	✓	Spline	✓	31.36	0.900	0.100
PRISM3D-E	VGGsFm	✓	Bézier	✓	31.95	0.908	0.097

Event Modality vs. View Sparsity Analysis To comprehensively evaluate the utility of the event modality beyond initialization, we investigate integrating continuous event supervision directly into our core optimization loop via an explicit event reconstruction loss. As summarized in Table 7, utilizing event priors during the initialization phase (PRISM3D-E Init. Only) yields a substantial +2.46 dB PSNR gain over the blind RGB baseline. Conversely, appending an explicit event constraint to the downstream optimization loss (PRISM3D-E Init. + Event Loss) provides no additional performance improvement and marginally increases rendering artifacts under normal view densities.




Fig. 9: Novel view synthesis LPIPS ($\downarrow$) of PRISM3D-E with and without continuous event loss under varying sparsity regimes. Left: absolute training views. Right: relative percentage.

To further untangle this behavior, we analyze the performance of both variants across progressive degrees of view sparsity. As illustrated in Fig. 9, continuous event

Table 7: Impact of event utilization during initialization vs. continuous optimization.

Method Configuration	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
PRISM3D (RGB Only)	29.49	0.840	0.153
PRISM3D-E (Init. + Event Loss)	31.86	0.909	0.098
PRISM3D-E (Init. Only) (Ours)	31.95	0.908	0.097

supervision during optimization becomes significantly beneficial under extremely sparse-view settings (e.g., $N \leq 4$) where spatial and geometric constraints from the sparse RGB images are severely limited. However, as the density of the viewpoints increases, our coupled probabilistic MCMC refinement and continuous $SE(3)$ Bézier physical trajectory modeling prove sufficiently robust to recover high-fidelity scene geometry, resulting in diminishing returns for continuous multi-modal constraints. This observation justifies our lightweight front-end architectural design choice for PRISM3D-E under standard multi-view settings.

5 Discussion

In this work, we presented **PRISM3D**, a unified framework that solves the inverse problem of blind 3D scene reconstruction from extreme motion blur. We identified that the primary barrier to robust deblurring is the “Initialization-Optimization Gap”: traditional SfM fails to register blurred images, while deep trackers yield geometry too sparse for deterministic optimization.

PRISM3D bridges this gap by coupling **Robust Deep Initialization** (via VGGSfM) with **Probabilistic Scene Modeling** (via MCMC). This synergy is critical. The deep tracker recovers the global topology that feature-based methods miss, while the probabilistic formulation enables Gaussian primitives to adaptively explore and populate the sparse geometry, preventing local minima that plague deterministic baselines. By jointly optimizing these probabilistic priors with continuous **Bézier Trajectories**, we ensure the reconstruction respects the physical integration of photons during the exposure.

Furthermore, beyond our robust standalone pipeline, we demonstrated that complementary sensor data can further extend reconstruction capabilities in severely degraded conditions. Our multi-modal extension, **PRISM3D-E**, utilizes event streams purely to secure a reliable VGGSfM initialization via intermediate pseudo-sharp image synthesis. Importantly, we preserve physical accuracy by isolating these synthesized priors from the training phase; the actual scene densification is driven entirely by the raw blurred captures and our continuous trajectory formulation. Validated on our newly introduced **PRISM3D-E Benchmark**, both configurations consistently achieve state-of-the-art performance, providing a highly versatile, event-agnostic framework for tackling extreme real-world motion blur.

Limitations and Future Work. While robust, PRISM3D currently assumes static scenes; extending it to dynamic environments via 4D Gaussian Splatting deformation fields is a natural next step. Additionally, fine-tuning the VGGSfM pipeline directly on extremely blurred datasets offers a clear path for immediate robustness gains. Finally, although feed-forward 3D Foundation Models (e.g., VGGT) currently introduce optimization-breaking noise, adapting them for severe degradation remains a promising long-term direction to achieve instantaneous, feed-forward scene initialization.

Acknowledgements

We would like to acknowledge partial support from the Institute of Eminence (IoE) Research Center on VR and Haptics. KM would like to acknowledge support from the Qualcomm Faculty Award 2024.

References

1. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: CVPR (2021) 4
2. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: CVPR (2022) 4
3. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: ECCV (2022) 4
4. Chen, W., Liu, L.: Deblur-gs: 3d gaussian splatting from camera motion blurred images. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 7(1), 1–15 (2024) 4, 5, 6, 8
5. Deguchi, H., Masuda, M., Nakabayashi, T., Saito, H.: E2gs: Event enhanced gaussian splatting. In: 2024 IEEE International Conference on Image Processing (ICIP). pp. 1676–1682. IEEE (2024) 4
6. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: CVPR (2022) 4
7. Jang, H., Choi, D., Kim, D., Kang, W., Kim, M.H.: Splat-based 3d scene reconstruction with extreme motion-blur. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 26425–26434 (2025) 4
8. Jiang, Y., Hedman, P., Mildenhall, B., Xu, D., Barron, J.T., Wang, Z., Xue, T.: Alignerf: High-fidelity neural radiance fields via alignment-aware training. CVPR (2023) 4
9. Jin, Y., Mishkin, D., Mishchuk, A., Matas, J., Fua, P., Yi, K.M., Trulls, E.: Image matching across wide baselines: From paper to practice. *International Journal of Computer Vision* 129(2), 517–547 (2021) 4
10. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM TOG* 42(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/> 2
11. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *TOG* (2023) 4, 8
12. Kheradmand, S., Rebain, D., Sharma, G., Sun, W., Tseng, Y.C., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: 3d gaussian splatting as markov chain monte carlo. *Advances in Neural Information Processing Systems* 37, 80965–80986 (2025) 2, 4, 5, 6, 8
13. Lee, D., Oh, J., Rim, J., Cho, S., Lee, K.M.: Exblurf: Efficient radiance fields for extreme motion blurred images. In: ICCV (2023) 4, 5, 8
14. Lindell, D.B., Martel, J.N., Wetzstein, G.: Autoint: Automatic integration for fast neural volume rendering. In: CVPR (2021) 4
15. Lindenberger, P., Sarlin, P.E., Larsson, V., Pollefeys, M.: Pixel-perfect structure-from-motion with featuremetric refinement. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 5987–5997 (2021) 4
16. Liu, Y., Luo, C., Tang, Z., Peng, J., Zhang, Z.: Vggt-x: When vggt meets dense novel view synthesis. *arXiv preprint arXiv:2509.25191* (2025) 2, 4, 13
17. Max, N.: Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* pp. 99–108 (1995). <https://doi.org/10.1109/2945.468400> 4
18. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis (Aug 2020). <https://doi.org/10.48550/arXiv.2003.08934> 2, 4

19. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *TOG* (2022) 4
20. Pan, L., Baráth, D., Pollefeys, M., Schönberger, J.L.: Global structure-from-motion revisited. In: *European Conference on Computer Vision*. pp. 58–77. Springer (2024) 4
21. Pan, L., Scheerlinck, C., Yu, X., Hartley, R., Liu, M., Dai, Y.: Bringing a Blurry Frame Alive at High Frame-Rate With an Event Camera. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6820–6829 (2019) 3, 8
22. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An Imperative Style, High-performance Deep Learning Library. *Advances in neural information processing systems* **32** (2019), <https://pytorch.org/> 8
23. Qi, Y., Zhu, L., Zhang, Y., Li, J.: E2NeRF: Event Enhanced Neural Radiance Fields from Blurry Images. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 13254–13264 (2023) 4, 8
24. Qi, Y., Zhu, L., Zhao, Y., Bao, N., Li, J.: Deblurring neural radiance fields with event-driven bundle adjustment. In: *Proceedings of the 32nd ACM International Conference on Multimedia*. pp. 9262–9270 (2024) 4, 8
25. Rebecq, H., Gehrig, D., Scaramuzza, D.: ESIM: An Open Event Camera Simulator. In: *Proceedings of The 2nd Conference on Robot Learning*. pp. 969–982. PMLR (Oct 2018) 8
26. Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: *ICCV* (2021) 4
27. Reizenstein, J., Shapovalov, R., Henzler, P., Sbordone, L., Labatut, P., Novotny, D.: Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 10901–10911 (2021) 4
28. Sarlin, P.E., Cadena, C., Siegwart, R., Dymczyk, M.: From coarse to fine: Robust hierarchical localization at large scale. In: *CVPR* (2019) 4
29. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: SuperGlue: Learning feature matching with graph neural networks. In: *CVPR* (2020) 4
30. Schonberger, J.L., Frahm, J.M.: Structure-from-motion Revisited. In: *CVPR* (2016), <https://github.com/colmap/colmap> 2, 3, 5
31. Schops, T., Schonberger, J.L., Galliani, S., Sattler, T., Schindler, K., Pollefeys, M., Geiger, A.: A multi-view stereo benchmark with high-resolution images and multi-camera videos. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3260–3269 (2017) 4
32. Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. *CVPR* (2022) 4
33. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Ruppert, C., Novotny, D.: Vggt: Visual geometry grounded transformer. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 5294–5306 (2025) 2, 4, 13
34. Wang, J., Karaev, N., Ruppert, C., Novotny, D.: Vggsfm: Visual geometry grounded deep structure from motion. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 21686–21697 (2024) 2, 4, 5, 6
35. Wang, P., Zhao, L., Ma, R., Liu, P.: BAD-NeRF: Bundle Adjusted Deblur Neural Radiance Fields. In: *CVPR* (2023), <https://wangpeng000.github.io/BAD-NeRF/> 4

36. Wu, L., Lee, J.Y., Bhattad, A., Wang, Y., Forsyth, D.: Diver: Real-time and accurate neural radiance fields with deterministic integration for volume rendering (2022) [4](#)
37. Ye, V., Li, R., Kerr, J., Turkulainen, M., Yi, B., Pan, Z., Seiskari, O., Ye, J., Hu, J., Tancik, M., et al.: gsplat: An open-source library for gaussian splatting. *Journal of Machine Learning Research* **26**(34), 1–17 (2025) [8](#)
38. Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenotrees for real-time rendering of neural radiance fields. In: *ICCV* (2021) [4](#)
39. Yu, W., Feng, C., Li, J., Tang, J., Yang, J., Tang, Z., Cao, M., Jia, X., Yang, Y., Yuan, L., et al.: Evagaussians: Event stream assisted gaussian splatting from blurry images. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 24780–24790 (2025) [4](#)
40. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: *CVPR* (2022) [8](#)
41. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Multi-stage Progressive Image Restoration. In: *CVPR* (2021), <https://github.com/swz30/MPRNet> [8](#)
42. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: NeRF++: Analyzing and improving neural radiance fields. *arXiv:2010.07492* (2020) [4](#)
43. Zhao, L., Wang, P., Liu, P.: BAD-Gaussians: Bundle Adjusted Deblur Gaussian Splatting. In: *ECCV*. Springer (2024) [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)