Unblur-SLAM: Dense Neural SLAM for Blurry Inputs

Qi Zhang
University of Amsterdam
Amsterdam, Netherlands
q.zhang@uva.nl

Denis Rozumny
Meta Reality Labs
Zurich, Switzerland
rozumden@gmail.com

Francesco Girlanda
ETH Zürich
Zürich, Switzerland
fgirlanda@student.ethz.ch

Sezer Karaoglu
University of Amsterdam
Amsterdam, Netherlands
s.karaoglu@uva.nl

Marc Pollefeys
ETH Zürich
Zürich, Switzerland
marc.pollefeys@inf.ethz.ch

Theo Gevers
University of Amsterdam
Amsterdam, Netherlands
th.gevers@uva.nl

Martin R. Oswald
University of Amsterdam
Amsterdam, Netherlands
m.r.oswald@uva.nl

Figure 1. **Unblur-SLAM results.** Our SLAM-approach integrates both a feed-forward deblurring and rerendering-based test-time refinement effectively. The latter one estimates a local point spread function, which enables our method to handle multiple sources of blur, demonstrating excellent performance for both motion and defocus blur. While previous blur-aware SLAM approaches typically assume all input frames to be blurry and are thus significantly slower than regular SLAM methods, Unblur-SLAM detects the amount of blur in the input frame and skips the costly refinement for sharp frames.

Abstract

We propose Unblur-SLAM, a novel RGB SLAM pipeline for sharp 3D reconstruction from blurred image inputs. In contrast to previous work, our approach is able to handle different types of blur and demonstrates state-of-the-art performance in the presence of both motion blur and defocus blur. Moreover, we adjust the computation effort with the amount of blur in the input image. As a first stage, our method uses a feed-forward image deblurring model for which we propose a suitable training scheme that can improve both tracking and mapping modules. Frames that are successfully deblurred by the feed-forward network obtain refined poses and depth through local-global multi-view optimization and loop closure. Frames that fail the first stage deblurring are directly modeled through the global 3DGS representation and an additional blur network to model multiple blurred sub-frames and simulate the blur formation process in 3D space, thereby learning sharp details and refined sub-frame poses. Experiments on several real-world datasets demonstrate consistent improvements in both pose estimation and sharp reconstruction results of geometry and texture.[1]

[1]The code and dataset are open-sourced at https://github.com/SlamMate/Unblur-SLAM.git

1. Introduction

The estimation of high-quality 3D reconstructions from a sequence of images in an online setting using Simultaneous Localization and Mapping (SLAM) has been a long-studied topic and has numerous applications in robotics, construction, mixed reality, and autonomous driving.

The vast majority of SLAM methods only work well when a number of assumptions are met, *e.g*. static scenes, good lighting, sharp input frames, sufficient image overlap. In this paper, we study the SLAM problem for scenarios in which the input frames can be substantially blurred. Several traditional SLAM methods that rely on image features [4, 7, 36] fail on blurry inputs because corner-like features can no longer be found. To avoid the drawbacks of feature-based methods, a lot of subsequent research focused on denser approaches for both tracking and mapping [11, 12, 19]. Recent methods based on 3D Gaussian Splatting (3DGS) can produce even higher fidelity renderings [14, 16, 32, 44, 55, 58], and due to their flexibility in optimizing surface positions, compared to Neural Radiance Field-based methods [18, 33, 35], they are much faster in combination with slightly better accuracy. The recent online 3DGS-based SLAM approaches Splat-SLAM [44] and HI-SLAM2 [60] attempt to integrate the advantages of frame-by-frame approaches by incorporating loop closure detection, global bundle adjustment, and deformable 3D Gaussian maps. With the rise of 3D foundation models equipped with abundant 3D spatial prior knowledge [25, 47, 51], their SLAM variants effectively leverage this knowledge [28, 30, 37] to achieve robust but slightly less precise reconstructions.

While image blur is a common phenomenon degrading the image quality due to rapid camera motion or scene changes, none of these works explicitly handle blurry inputs, relying only on multi-view optimization and loop closure detection to reduce the impact. Recent deblur methods [3, 40, 46, 52] have achieved good results by explicitly modeling physical motion and defocus blur, and attempting to learn sharp details from blurry inputs. However, they have many limitations, *e.g*. the SLAM methods in these works [3, 52] mostly focus on simple motion blur, limiting the further development of their performance.

The proposed Unblur-SLAM method aims to further improve localization and high-fidelity reconstruction accuracy in more realistic blur scenarios, while addressing the limitations of the aforementioned methods. It combines explicit and implicit approaches to handle image blur. First, we preliminarily process real blur using the deblur foundation model through our proposed training method. Then, we reduce the impact of blurry frames on the system through multi-view optimization and loop closure detection, and estimate pixel-wise blur kernels through an MLP network and historical information. When the deblur network fails, we render multiple blurry sub-frames in 3D space by the MLP network and the global 3DGS [18] model. Sharp frames will be identified and jointly optimized with other frames for 3DGS reconstruction. As shown in Figure 1, our system exhibits excellent deblurring effects. We make the following contributions:

- We propose the first SLAM method capable of simultaneously handling defocus blur and motion blur. Using multi-scale blur kernels estimated via additional MLP layers, after processing by a deblurring neural network, we can handle various blur types.
- We develop an automatic method to detect the amount of blur in the input frame, treat sharp and blurry frames separately, and skip the costly refinement for sharp frames. This is in comparison to all previous blur-aware SLAM approaches that assume all input frames to be blurry.
- We integrate a single-image deblur feed-forward model into a high-fidelity reconstruction SLAM system, achieving 3D consistency through our proposed training method, thereby bringing continuous improvements to both SLAM tracking and reconstruction modules.
- We propose a robust online deblurring method based on 3D Gaussians. The method performs initial blur removal based on a feed-forward model and conducts further blur removal and detail enhancement using multi-view blur kernel learning. Meanwhile, the method incorporates loop closures and online Gaussian deformation to further enhance accuracy. The method achieves state-of-the-art performance on multiple datasets. Code will be released.

2. Related Work

Neural Network-Based Deblurring Methods. Single-frame image deblurring networks achieve superior restoration quality through problem-tailored architectures. For motion blur tasks, Transformer-based methods including Chen *et al*. [5] and LoFormer [31] enhance detail recovery by decomposing motion patterns. However, these 2D deblurring networks do not impose multi-view geometric constraints and are prone to generating 2D noise. Meanwhile, these networks have prohibitively long computation times and cannot perform online inference. In contrast, state space models such as XYScanNet [26] and EVSSM [20] adopt Mamba-based architectures to capture long-range dependencies with linear complexity, thereby enhancing real-time performance while maintaining good quality. These networks do not impose physical constraints during training. Furthermore, these networks have not been evaluated on images with high sensor noise [48] and heavy motion blur [45] as common in SLAM settings.

Deblurring methods based on 3D Gaussians. 3D Gaussian Splatting (3DGS) [17] can effectively utilize historical information to model input image blur in a 3D-consistent

manner based on its high-fidelity dense mapping capability. Similar to neural network-based deblurring work, these approaches can be divided into two categories. The first category only models motion blur. For example, BAD-Gaussians [62] uses historical models to render the physical imaging process of motion blur modeling, recovering rendered sub-frame trajectories while jointly optimizing Gaussian parameters. CoMoGaussian [24] refines the motion modeling by using neural ordinary differential equations to model continuous camera motion trajectories instead of the previous discretized interpolation, and uses per-pixel weights to perform weighted averaging of each sub-frame's contribution to the final rendered blurry image rather than directly averaging the sub-frames. These works often assume that motion is captured under small aperture camera settings with only motion blur. DoF-GS [53] represents another category of works which models defocus blur through a focus localization network. There are also methods that model both types of blur through blur kernels, such as Deblurring 3D Gaussian Splatting [21] and BAGS [40]. However, they are limited by kernel size and Gaussian count, failing to handle extreme blur robustly. In contrast, the blur in our system is preliminarily processed by a deblurring neural network.

Deblurring SLAM Methods. In the SLAM field, advanced theories from offline deblurring 3D reconstruction methods have been leveraged to construct online SLAM approaches with deblurring capabilities. MBA-SLAM [52] adopts the motion blur modeling method based on BAD-Gaussians [62] for online operation and combines it with a motion blur-aware tracker. I2-SLAM [3] sets explicit variables for the image formation steps (white balance, exposure time, camera response function) to adapt to appearance changes in each frame, while also using sub-frames to model motion blur. I2-SLAM [3] attempted to use deblurring neural networks to enhance performance, but failed due to the 3D spatial inconsistency of single-image deblurring networks. Similar to I2-SLAM, Deblur-SLAM [13] further integrates local bundle adjustment and loop closures for improved consistency. DAGS [42] builds upon this by adding Gaussian-based deblurring awareness to output binary blur labels, unlike other works that assume all frames are blurred. All these works assume that the camera operates with a small aperture without any defocus blur and they cannot adapt their high computational effort to the amount of blur in the input images. Since in practice many SLAM datasets contain a mixture of motion blur and defocus blur, as well as various ratios of sharp vs. blurry images, we designed the first method to handle both types of blur and to adapt the computation effort to the estimated amount of blur in the input images for fast computation times.

3. Method

Unblur-SLAM assumes that real-world blur may have different sources: defocus blur occurs when light forms an out-of-focus image on the camera plane, while motion blur occurs through temporal integration of motion during exposure. Under this assumption, we propose a physics-constrained method to train a 3D-consistent single-image deblur model to enhance the blur robustness of the Droid-SLAM [49] tracker and the monocular depth estimation network [10].

We classify input frames into three categories based on blur detector scores: sharp frames, blurred frames with failed deblurring, and blurred frames with successful deblurring. Sharp frames are directly processed through Droid-SLAM for pose estimation and depth refinement. The 3D Gaussian representation is further refined using sharp frames through online deformation and rendering optimization guided by the Splat-SLAM [44] loss function. RGB frames with high blur scores where deblurring fails are skipped by the tracker. These frames are optimized by calculating photometric loss against motion-blurred frames formed from averaging multiple sub-frames that are blurred through the MLP layer. For successful deblurring cases, the frame is processed through Droid-SLAM, and the map representation is optimized using Gaussian deformation and photometric and depth errors of the rendered blurry frame modeled by the MLP layer. As shown in Figure 2, the blur models are divided into two categories: one category uses multiple blur sub-frames to optimize a photometric error with the input image, while the other category uses a single blurry rendered frame to compute the output image from the previous deblur model.

3.1. Blur Quantification and Recognition

As part of our pipeline, we dynamically detect which frames are sharp and which are blurry. To identify the optimal blur detector for this task, we constructed a comprehensive blur detection benchmark consisting of real-world datasets and two semi-synthetic blur datasets. The real-world data includes RealBlur [43] for motion blur, as well as DPDD [1] and RealDoF [23] for defocus blur.

To construct the first semi-synthetic dataset, we apply the Eden [61] video interpolation model to the ScanNet [9] dataset to generate intermediate frames. We physically model motion blur by averaging these interpolated frames in linear color space, utilizing the middle frame as the ground truth sharp reference. The second semi-synthetic dataset combines sequences from the RED [39], GoPro [38], and ReplicaBlurry [13] datasets, which are manually annotated to remove frames with moving objects or inherent blur.

Using this benchmark, we evaluated 39 different image quality metrics. We assessed them based on a custom proxy measurement that calculates a consistency score de-

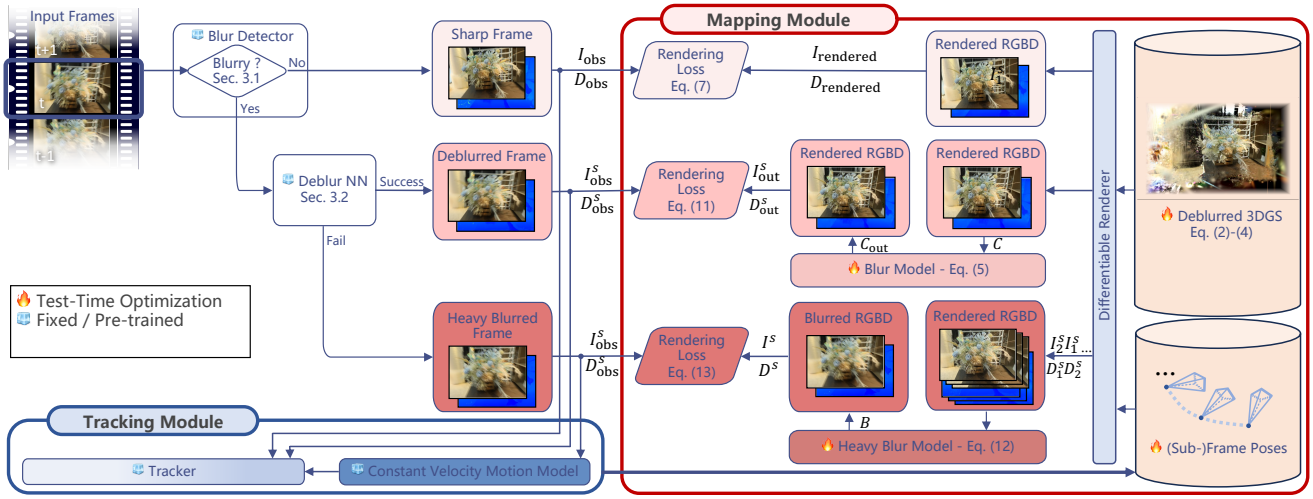

Figure 2. **Method overview.** Unblur-SLAM robustly handles varying blur by adaptively categorizing images into sharp, blurry, and heavily blurred levels (shown in different red shades). Since typically only a subset of frames is blurred, this improves both blur handling and average runtime. Both tracking and mapping modules optionally leverage the deblurring network. The red mapping module optimizes the 3DGS reconstruction using sliding-window (Eq. (14)) and global (Eq. (15)) losses, incorporating depth (Eq. (4)) and pose from the blue tracking module.

rived from the metric's accuracy and effect size. Based on our evaluation, ARNIQA [2] achieved the highest consistency score and was selected as our default blur detector. Detailed formulations regarding the blur synthesis and metric evaluations are provided in the supplementary material (Sec. 7.4).

3.2. Physics-constrained Deblurring Network

We impose additional physical constraints during the dataset creation process to guide the network in learning 3D-consistent deblurring capabilities. First, according to imaging principles, real-world motion blur is formed when light hits the camera plane and undergoes a certain degree of defocus based on focal length and then accumulates on the exposure plane. In this case, the image forms multiple trailing edges. That is, the image represents a video sequence over the exposure duration. The physical essence of deblurring lies in restoring the image corresponding to a specific moment during exposure. Real-world motion blur formation follows the physical process:

$$\mathrm{B} = \int_0^{\mathrm{T}} \Big(\mathrm{I(t)} \otimes \mathrm{K_{defocus}}\big(\mathrm{f}, \mathrm{d(t)}\big) \Big) \mathrm{d}t \tag{1}$$

where B is the blurred image, I(t) is the latent sharp image at time t, $\otimes$ denotes the convolution operation, $\mathrm{K_{defocus}}$ is the defocus kernel, and T is the exposure time.

We select the middle frame as the reconstruction target to avoid ambiguity in motion direction [41]. However, the inherent misalignment in real-world motion blur datasets [43] may lead the model to learn image sharpening instead of accurately extracting the mid-exposure frame. To address this, we adopt a two-stage training strategy: we first train on the RED [39], GoPro [38], and ReplicaBlurry [13] datasets (*i.e.* semi-synthetic datasets) to learn mid-exposure frame extraction. We use multiple video frames from real-world and synthetic environments to model motion blur, which allows the model to learn the accurate positions of intermediate frames. The combination of synthetic and real-world data enables it to achieve a certain degree of generalization across different camera models. Subsequently, we fine-tune on the DPDD [1] defocus dataset to enhance robustness against defocus blur. All datasets undergo inverse gamma correction to maintain color constancy.

3.3. 3D Gaussian Splatting Representation

Following Splat-SLAM [44], we represent the scene as a collection of 3D Gaussians $\mathrm{G} = \{\mathrm{g_i}\}$, where each Gaussian $\mathrm{g_i}$ is parameterized by:

$$\mathrm{g_{i(x)}} = \exp\left(-\frac{1}{2} (\mathrm{x} - \boldsymbol{\mu}_\mathrm{i})^\mathrm{T} \Sigma_\mathrm{i}^{-1} (\mathrm{x} - \boldsymbol{\mu}_\mathrm{i}) \right) \tag{2}$$

where $\boldsymbol{\mu}_\mathrm{i} \in \mathrm{R}^3$ is the mean position, $\Sigma_\mathrm{i} \in \mathrm{R}^{3\times3}$ is the covariance matrix decomposed as $\Sigma_\mathrm{i} = \mathrm{R_i S_i S_i^T R_i^T}$ with rotation $\mathrm{R_i}$ and scale $\mathrm{S_i}$, along with opacity $\mathrm{o_i} \in [0, 1]$ and color $\mathrm{c_i} \in \mathrm{R}^3$. The rendering projects 3D Gaussians onto the image plane and blends them using alpha compositing:

$$\mathrm{C} = \sum_{\mathrm{i}\in\mathrm{N}} \mathrm{c_i} \alpha_\mathrm{i} \prod_{\mathrm{j}=1}^{\mathrm{i}-1} (1-\alpha_\mathrm{j}), \quad \mathrm{D_r} = \sum_{\mathrm{i}\in\mathrm{N}} \mathrm{d_i} \alpha_\mathrm{i} \prod_{\mathrm{j}=1}^{\mathrm{i}-1} (1-\alpha_\mathrm{j}) \tag{3}$$

where $\alpha_\mathrm{i} = \mathrm{o_i} \cdot \mathrm{g_i^{2D}(x')}$ represents the combined opacity and 2D Gaussian density at pixel x', and $\mathrm{d_i}$ is the depth of

Gaussian center.

When Droid-SLAM performs multi-view depth optimization, we accordingly deform the Gaussian representation. Given a depth update from d to d' at pixel coordinate (u, v) in keyframe K, we update the mean position of affected Gaussians:

$$\boldsymbol{\mu}_{i'} = \boldsymbol{\mu}_i + \frac{d' - d}{d}(\boldsymbol{\mu}_i - t_K) \tag{4}$$

where t_K is the camera center of keyframe K. This deformation ensures the Gaussian map remains consistent with the optimized geometry.

3.4. Complex Blur Modeling

For successfully deblurred frames, we adopt the Blur Proposal Network (BPN) from BAGS [40] to model residual blur. BPN estimates per-pixel convolution kernels $h(x) \in R^{K\times K}$ and masks $m(x) \in [0, 1]$:

$$C_{out}(x) = \big(1 - m(x)\big)C(x) + m(x)\big(C(x) \otimes h(x)\big) \tag{5}$$

where C is the sharp rendered image. For frames with failed deblurring, we render multiple sub-frames and apply BPN to each, discretizing the temporal integration in Eq. (1):

$$B \approx \frac{1}{N_{sub}} \sum_{j=1}^{N_{sub}} \mathrm{BPN}(I_j, h_j, m_j) \tag{6}$$

3.5. Hybrid Bundle Constraint

We use deformable Gaussians to update the depth propagated from the front-end tracker during loop closure and multi-frame optimization, while simultaneously updating each camera's pose. Meanwhile, frames are categorized as sharp if the blur detection metric [2] is below a pre-defined threshold τ_{sharp}, otherwise as blurry. For blurry frames, we apply the physics-constrained deblurring network and evaluate success using both Laplacian ratio and blur quantification metric. Successful deblurring is determined when the combined metric exceeds the threshold $\tau_{success}$.

Sharp frame optimization. Sharp frames receive higher optimization weights w_{sharp} to provide strong priors:

$$\begin{aligned} L_{sharp} = w_{sharp}\Big[&\lambda_{rgb}(||I_{rendered} - I_{obs}||_1) \\ &+ \lambda_{depth}||D_{rendered} - D_{obs}||_1\Big] \end{aligned} \tag{7}$$

Multi-scale Blurry frame optimization. Successfully deblurred frames undergo multi-scale optimization without pose updates. We progressively optimize from coarse to fine scales s with exposure compensation, where $a^s_{exposure}$ and $b^s_{exposure}$ are learnable affine scale and shift parameters for the current scale s:

$$I^s_{adjusted}(x) = \exp(a^s_{exposure}) \cdot I^s_{rendered}(x) + b^s_{exposure} \tag{8}$$

The multi-scale loss combines RGB and depth terms with appropriate masking to handle boundary regions and invalid depth values. For each scale s, the Blur Proposal Network (BPN) from BAGS [40] learns depth-aware kernels to model residual blur and enhance details:

$$\begin{aligned} I^s_{out}(x) =&\big(1 - m^s(x)\big)I^s_{adjusted}(x) \\ &+ m^s(x)\big(I^s_{adjusted}(x) \otimes h^s(x)\big) \end{aligned} \tag{9}$$

where $h^s(x) \in R^{K_s\times K_s}$ is the per-pixel convolution kernel and $m^s(x) \in [0, 1]$ is the blending mask.

The kernel $h^s(x)$ adapts based on the rendered depth $D(x)$, effectively performing depth-dependent detail restoration:

$$h^s\big(x, D(x)\big) = h^s_{deblur}\big(D(x)\big) + \alpha(D(x)) \cdot h^s_{sharpen}\big(D(x)\big) \tag{10}$$

where h^s_{deblur} removes residual blur and $h^s_{sharpen}$ enhances edges, with $\alpha(D(x))$ being a learnable depth-dependent factor. For distant regions with larger depth values, the kernel provides stronger enhancement to compensate for angular resolution limitations and depth-dependent degradation. This enhancement remains physically plausible through multi-view consistency in our bundle adjustment.

Unlike BAGS [40], which requires a warm-up phase, we directly optimize each frame without initialization stages. Rather than employing Mip-Splatting [57] for multi-resolution consistency, we adopt direct image rescaling. The multi-scale loss combines RGB and depth terms with BPN modeling and appropriate masking to handle boundary regions and invalid depth values:

$$\begin{aligned} L^s_{deblur} =&\lambda_{rgb}\|I^s_{out} - I^s_{obs}\|_1 + \lambda_{depth}||D^s_{out} - D^s_{obs}||_1 \\ &+ \lambda_{sparse}\|m^s\|_1 \end{aligned} \tag{11}$$

where $\|m^s\|_1$ represents the L_1 norm of the mask over all valid pixels x. where the sparsity constraint on m^s ensures selective enhancement only where beneficial, preventing overfitting to sensor-specific noise patterns. The D^s_{out} and I^s_{out} denote the depth and image outputs from BPN (Eq. (6)) at scale s and the Different from offline reconstruction, we set the mask to 0 for insufficiently reconstructed regions in online reconstruction, preventing overly large or small masks from affecting the inference. This enhancement improves perceptual quality but may decrease pixel-wise metrics such as PSNR.

Heavy blurry frame optimization. For failed deblurring, we employ virtual camera trajectory optimization. We generate N_{sub} virtual cameras along the blur trajectory, each parameterized by $(R_k, t_k, \theta_k, \rho_k)$. The composite blur model is applied to each rendered sub-frame, and the averaged result is compared with the blurred observation. We evaluate the rendering errors for RGB I^s and depth D^s images at scale s.

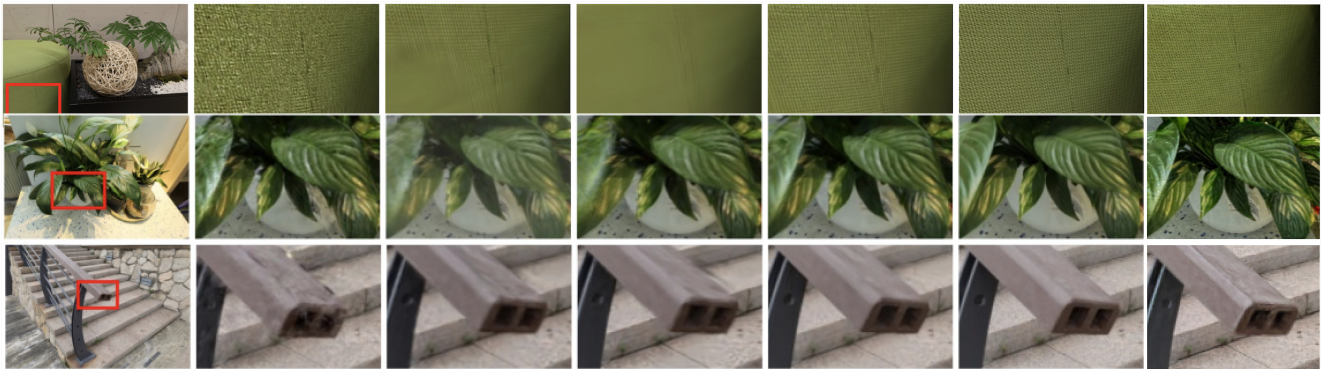

Figure 3. **Deblurring performance of our method** when compared to other state-of-the-art offline methods.

First, we synthesize the blurred approximation $\hat{Q}^s$:

$$\hat{Q}^s = \frac{1}{N_{sub}} \sum_{j=1}^{N_{sub}} \text{BPN}(Q_j^s, h_j^s, m_j^s) \tag{12}$$

The failure loss is then calculated as:

$$L_{fail} = \sum_{Q \in \{I^s, D^s\}} \lambda_Q \|\hat{Q}^s - Q_{obs}^s\|_1 + \lambda_{sparse} \|m^s\|_1 \tag{13}$$

where Q_j^s and Q_{obs}^s denote the rendered sub-frames (color or depth) and the observed input frame. Other terms are defined in Eq. (6) and (11).

Bundle adjustment. Joint bundle optimization is performed across all frame types with adaptive weighting:

$$L_{total} = \sum_{f \in F} w_f L_f \tag{14}$$

where L_f is the frame-specific loss term, which equals L_{sharp} for sharp frames, or L_{deblur}^s / L_{fail} for blurry frames. The $w_{sharp} > w_{deblur} = w_{fail}$ to prioritize sharp observations. Unlike Splat-SLAM's [44] selective keyframe filtering, we incorporate all tracker keyframes as backend optimization keyframes to fully leverage deblurring results.

3.6. Global Optimization and Loop Closure

Following Splat-SLAM [44], we maintain global consistency through the following three main paradigms.

▷ **Local bundle adjustment** optimizes within a sliding window of recent keyframes using the Disparity, Scale, and Pose Optimization (DSPO) layer.

▷ **Loop closure detection** identifies revisited locations through optical flow magnitude thresholding and temporal constraints.

▷ **Global bundle adjustment** periodically optimizes the entire pose graph with appropriately deformed Gaussians:

$$L_{global} = \sum_{f \in F} w_f L_f + \lambda_{reg} \sum_{i=1}^{N} \|s_i - \bar{s}\|_1 \tag{15}$$

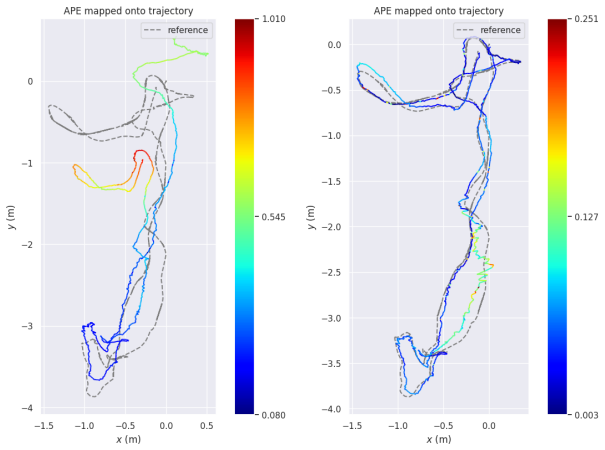

(a) Droid-SLAM (b) Ours

Figure 4. **Trajectory comparison with Droid-SLAM** on the indoor MCD dataset [45].

where w_f are frame-specific weights prioritizing sharp observations, and the regularization term with $\lambda_{reg} > 0$ prevents largely elongated Gaussians.

3.7. Final Refinement

Upon trajectory completion, we perform final optimization with multi-scale progressive refinement, re-estimating poses and exposure parameters for failed deblurring frames while maintaining global consistency for all other frames.

4. Experiments

Datasets. To evaluate our method under extreme motion blur conditions, we use a synthetically blurred ReplicaBlurry [13] dataset generated by (Sec. 3.1) and ArchViz dataset generated by the Unreal game engine[27]. We also test our method on real-world TUM-RGBD [48] and IndoorMCD [45] datasets to assess performance under different camera conditions with handheld devices and wheeled platforms. To measure the deblurring quality of our method, we further evaluate using the Deblur-NeRF [29] dataset, which is commonly used for offline deblurring methods.

Metrics. For experiments on synthetic datasets, we report PSNR between rendered keyframe images and the original sharp images to demonstrate the robustness of our method in extreme cases. For the Deblur-NeRF [29] dataset, we report PSNR, SSIM [54], and LPIPS [59] metrics with ground truth sharp images. For other real-world datasets (TUM-RGBD [48], IndoorMCD [45]), we only report ATE[m] due to the lack of sharp ground truth frames. Further, we evaluated PSNR on 3 sequences from the TUM dataset with manually annotated sharp frames provided by I2-SLAM.

Implementation details. We set three scales for mapping optimization at 3, 5, and 9. Under severe blur, we model motion blur using 3 virtual sub-frames. Experiments are conducted on an AMD EPYC-2 7282 processor with an NVIDIA RTX A6000 equipped with 48GB memory.

Evaluation. We first conduct a qualitative evaluation of the blur detector's performance on synthetic and real blurred datasets, as described in Sec. 3.1. We select the metric with the highest consistency score as our blur detector. We then evaluate pose estimation performance in real-world environments, as shown in Table 3. Ablation studies on different training datasets validate the correctness of our theory. To demonstrate whether our model learns 3D-consistent physical properties, we apply the deblurring model (without blur detector filtering) to the Droid-SLAM pipeline. To assess performance across diverse scenarios, we evaluate our system using 19 sequences from the TUM RGB-D dataset and 57 sequences from the MCD dataset. As shown in Figure 4, we get high ATE improvement.

To further evaluate our method's robustness against severe motion blur, we conduct an ablation study on the proposed fallback mechanism using the synthetic Replica-Blurry dataset. Please refer to Sec. 6.1 in the supplementary material for detailed experimental settings and quantitative results. We further tested our method on the synthetic dataset ArchViz, as shown in Table 1. We are more robust than MBA-SLAM on datasets where all frames are blurred.

We also tested our method's performance in mapping. After communicating with the authors of I2-SLAM [3], we obtained their manually annotated sharp frames on the TUM dataset and the keyframes used by their system. Under the same keyframes, we compared our system with I2-SLAM. As shown in Table 5, our system achieved state-of-the-art performance in regions with a high proportion of blurred frames. As show in Figure 5, we also have a better result than MBA-SLAM. I2-SLAM also attempted to use the single-image deblurring model to improve their method's effectiveness, but it failed due to the model's 3D inconsistency. Our method, after further incorporation of Gaussian-based deblurring, demonstrated detail enhancement that surpasses annotated sharp images [3, 29] as shown in Figures 6 and 7. This enhancement disrupts the evaluation of reference-based mapping metrics that assume a perfect ground truth. This phenomenon often occurs in large scenes when the camera moves forward and backward, causing the same object to occupy inconsistent pixel sizes across multiple views. In small scenes, such as the Deblur-NeRF motion blur dataset, our complete system shows significant improvement over the system without Gaussian-based deblurring, as shown in Table 4 and Figure 3. In addition, we conducted experiments to observe the performance of using our method in the Deblur-NeRF defocus blur dataset, as shown in Table 2. We compare to offline methods on the two subsets of the Deblur-NeRF dataset.

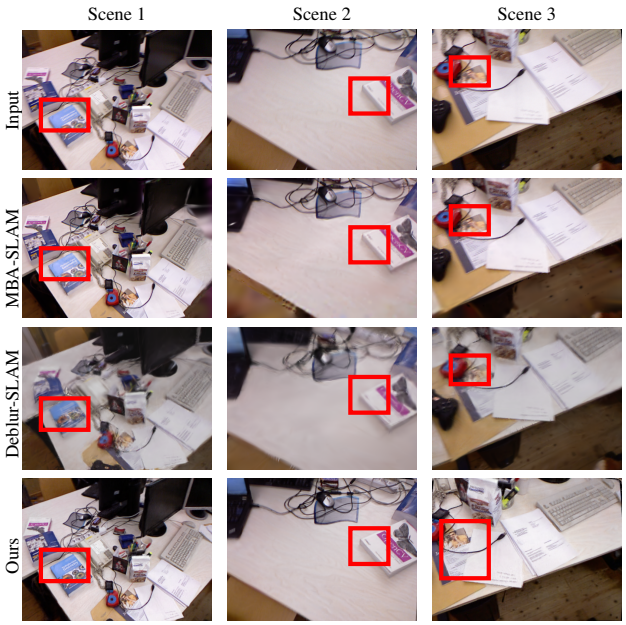

Figure 5. **Qualitative comparison with MBA-SLAM** [52] on the TUM dataset [48]. Our method yields sharper reconstruction results in many scene parts. The qualitative experimental results for the shown fr1_desk sequence were obtained through communication with the MBA-SLAM authors.

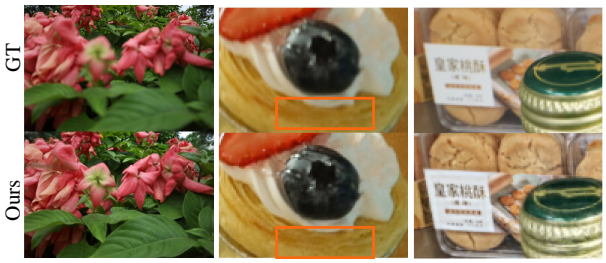

Figure 6. **Qualitative comparison of our method** with ground truth sharp frames (top row) from the Deblur-NeRF defocus dataset [29].

Runtime. We used three sequences from the TUM dataset: fr1_desk, fr2_xyz, and fr3_office to perform runtime analysis of our method under online settings in Table 6. We also ablated another version of our method without the Gaussian refinement-based approach using only the single-image deblurring neural network. It can be observed that our method

Method	ArchViz-1		ArchViz-2		ArchViz-3		Avg. ATE
	ATE ↓	PSNR ↑	ATE ↓	PSNR ↑	ATE ↓	PSNR ↑	
MBA-SLAM [52]	**0.0075**	28.45	0.0036	30.16	0.0141	27.85	0.0084
Ours	**0.0075**	**28.76**	**0.0027**	**32.71**	**0.0067**	**30.09**	**0.0056**

Table 1. **Qualitative experiments on the ArchViz** [52] synthetic dataset to demonstrate the performance comparison between our method and MBA-SLAM [52]. This dataset contains extreme blur cases, and we achieve the best results using ATE RMSE [m] and PSNR metrics.

Method	Deblur-NeRF [29]	DP-NeRF [22]	BAGS [40]	Deblur-GS [21]	DOF-GS [53]	Ours
PSNR ↑	23.98	24.12	24.17	24.21	24.12	**27.45**

Table 2. **Comparison on the Deblur-NeRF defocus blur** dataset [29]. We report PSNR (dB) ↑ metric and compare to several state-of-the-art methods. Our method outpeforms other methods by a wide margin, both quantitatively and qualitatively.

Dataset	Droid-SLAM [49]	Real_J	Ours*	Ours
TUM	0.380	0.400	0.352	**0.336**
MCD	0.138	0.139	0.155	**0.128**

Table 3. **Comparison of trajectory errors** ATE RMSE [m] ↓. Our method yields the best tracking results. Real_J differs from Ours by replacing the DPDD dataset with the RealBlur_J dataset. 'Ours' differs from 'Ours*' by training on five ReplicaBlurry [13] sequences instead of one.

Method	PSNR↑	SSIM↑	LPIPS↓
Mip-Splatting [57]	21.87	0.6270	0.3066
Deblur-NeRF [29]	25.63	0.7645	0.1820
BAD-Gaussians [62]	21.69	0.6471	0.1262
Deblurring 3DGS [21]	26.61	0.8224	0.1096
BAGS [40]	26.70	0.8237	0.0956
CoMoGaussian [24]	27.85	0.8431	0.0822
Ours w/o ref.	28.22	0.9007	0.1053
Ours	**29.49**	**0.9213**	**0.0728**

Table 4. **Comparison against offline deblurring methods** on the Deblur-NeRF dataset [29]. The comparison on both the motion blur and the defocus blur dataset of this established benchmark for offline deblurring methods demonstrates that Unblur-SLAM is able to outperform current state-of-the-art despite solving the harder online problem in which only partial data is available.

	fr1_desk	fr2_xyz	fr3_office
Sharp ratio	16%	100%	37%
I^2-SLAM [3]	27.23	**32.06**	28.91
Deblur-SLAM [13]	19.93	27.21	21.67
Ours	**28.03**	31.14	**29.22**

Table 5. **Rendering comparison (PSNR)** on TUM-RGBD dataset [48]. For fair comparison, we only use the same keyframes as in I2-SLAM [3] to obtain results, which are then compared against the sharp frames of the keyframes marked by I2-SLAM. The sharp ratio indicates the percentage of sharp keyframes in the input sequence.

achieves excellent real-time performance, and the limitation in real-time performance is primarily due to the deblur neural network which performs inference at 600ms per frame on our device.

Figure 7. **Qualitative comparison results on sharp frames manually labeled in I2-SLAM[3]**. The first row is the ground truth image, the second row is the experimental results of our system.

Method	Splat-SLAM	I^2-SLAM	Ours	Ours w/o ref.
FPS↑	1.24	0.095	0.74	**0.85**

Table 6. **Runtime comparison in frames per second (FPS)**. Ours w/o ref. indicates that we only use the deblurring model without refinement, while "Ours" refers to our complete pipeline. We compute the average FPS of our system on the TUM dataset sequences: fr1_desk, fr2_xyz, and fr3_office. Note that Splat-SLAM is not a deblurring method and is only listed for reference.

5. Conclusions

We propose Unblur-SLAM, a novel online deblurring approach that is able to handle both motion and defocus blur to reconstruct sharp 3D scenes from blurred images. By applying middle frame constraints and assumptions about the blur process, we enable single-image deblurring large models to learn 3D-consistent deblurring priors. Based on the physical assumptions of the blur formation, we construct a Gaussian-based system to estimate the blur kernels of the input images. To prevent system failure in extreme cases, we propose directly rendering multiple blurred sub-frames in 3D to model severe blur scenarios. Experiments on several real-world datasets show that Unblur-SLAM surpasses state-of-the-art methods in 3D scene deblurring. Unfortunately, our current system remains a trade-off between real-time efficiency and accuracy, requiring further improvements in runtime and memory consumption.

References

[1] Abdullah Abuolaim and Michael S Brown. Defocus deblurring using dual-pixel data. In *European Conference on Computer Vision*, pages 111–126. Springer, 2020. 3, 4

[2] Lorenzo Agnolucci, Leonardo Galteri, Marco Bertini, and Alberto Del Bimbo. Arniqa: Learning distortion manifold for image quality assessment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 189–198, 2024. 4, 5, 2

[3] Gwangtak Bae, Changwoon Choi, Hyeongjun Heo, Sang Min Kim, and Young Min Kim. I2-slam: Inverting imaging process for robust photorealistic dense slam. *arXiv preprint arXiv:2407.11347*, 2024. 2, 3, 7, 8

[4] Carlos Campos, Richard Elvira, Juan J.Gomez Rodriguez, Jose M.M. Montiel, and Juan D. Tardos. ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual-Inertial, and Multimap SLAM. *IEEE Transactions on Robotics*, 37 (6):1874–1890, 2021. 2

[5] Duosheng Chen, Shihao Zhou, Jinshan Pan, Jinglei Shi, Lishen Qu, and Jufeng Yang. A polarization-aided transformer for image deblurring via motion vector decomposition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 28061–28070, 2025. 2

[6] Jerred Chen and Ronald Clark. Image as an imu: Estimating camera motion from a single motion-blurred image. In *International Conference of Computer Vision (ICCV)*, 2025. 2

[7] Chi-Ming Chung, Yang-Che Tseng, Ya-Ching Hsu, Xiang-Qian Shi, Yun-Hung Hua, Jia-Fong Yeh, Wen-Chin Chen, Yi-Ting Chen, and Winston H Hsu. Orbeez-slam: A real-time monocular visual slam with orb features and nerf-realized mapping. *arXiv preprint arXiv:2209.13274*, 2022. 2

[8] Jacob Cohen. *Statistical Power Analysis for the Behavioral Sciences*. Lawrence Erlbaum Associates, Hillsdale, NJ, 2nd edition, 1988. 2

[9] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, 2017. 3, 2

[10] Ainaz Eftekhar, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10786–10796, 2021. 3, 2

[11] Jakob Engel, Thomas Schöps, and Daniel Cremers. Lsd-slam: Large-scale direct monocular slam. In *European conference on computer vision*, pages 834–849. Springer, 2014. 2

[12] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):611–625, 2017. 2

[13] Francesco Girlanda, Denys Rozumnyi, Marc Pollefeys, and Martin R Oswald. Deblur gaussian splatting slam. *arXiv preprint arXiv:2503.12572*, 2025. 3, 4, 6, 8, 1, 2

[14] Huajian Huang, Longwei Li, Hui Cheng, and Sai-Kit Yeung. Photo-slam: Real-time simultaneous localization and photo-realistic mapping for monocular, stereo, and rgb-d cameras. *arXiv preprint arXiv:2311.16728*, 2023. 2

[15] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5148–5157, 2021. 2

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

[17] 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. 2

[18] 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. 2

[19] Christian Kerl, Jürgen Sturm, and Daniel Cremers. Dense visual slam for rgb-d cameras. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2100–2106. IEEE, 2013. 2

[20] Lingshun Kong, Jiangxin Dong, Jinhui Tang, Ming-Hsuan Yang, and Jinshan Pan. Efficient visual state space model for image deblurring. *arXiv preprint arXiv:2405.14343*, 2025. 2, 1

[21] Byeonghyeon Lee, Howoong Lee, Xiangyu Sun, Usman Ali, and Eunbyung Park. Deblurring 3d gaussian splatting. *arXiv preprint arXiv:2401.00834*, 2024. 3, 8

[22] Dogyoon Lee, Minhyeok Lee, Chajin Shin, and Sangyoun Lee. Dp-nerf: Deblurred neural radiance field with physical scene priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12386–12396, 2023. 8

[23] Junyong Lee, Hyeongseok Son, Jaesung Rim, Sunghyun Cho, and Seungyong Lee. Iterative filter adaptive network for single image defocus deblurring. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2034–2042, 2021. 3

[24] Jungho Lee, Donghyeong Kim, Dogyoon Lee, Suhwan Cho, Minhyeok Lee, Wonjoon Lee, Taeoh Kim, Dongyoon Wee, and Sangyoun Lee. Comogaussian: Continuous motion-aware gaussian splatting from motion-blurred images. *arXiv preprint arXiv:2503.05332*, 2025. 3, 8

[25] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, pages 71–91. Springer, 2024. 2

[26] Hanzhou Liu, Chengkai Liu, Jiacong Xu, Peng Jiang, and Mi Lu. Xyscannet: A state space model for single image deblurring. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 779–789, 2025. 2

[27] Peidong Liu, Xingxing Zuo, Viktor Larsson, and Marc Pollefeys. Mba-vo: Motion blur aware visual odometry. In

Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 5550–5559, 2021. 6

[28] Yuzheng Liu, Siyan Dong, Shuzhe Wang, Yingda Yin, Yanchao Yang, Qingnan Fan, and Baoquan Chen. Slam3r: Real-time dense scene reconstruction from monocular rgb videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16651–16662, 2025. 2

[29] Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12861–12870, 2022. 6, 7, 8, 2

[30] Dominic Maggio, Hyungtae Lim, and Luca Carlone. Vggt-slam: Dense rgb slam optimized on the sl (4) manifold. *arXiv preprint arXiv:2505.12549*, 2025. 2

[31] Xintian Mao, Jiansheng Wang, Xingran Xie, Qingli Li, and Yan Wang. Loformer: Local frequency transformer for image deblurring. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 10382–10391, 2024. 2

[32] 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

[33] 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. In *European Conference on Computer Vision (ECCV)*. CVF, 2020. 2

[34] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21 (12):4695–4708, 2012. 2

[35] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. 2

[36] Raul Mur-Artal and Juan D. Tardos. ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras. *IEEE Transactions on Robotics*, 33(5): 1255–1262, 2017. 2

[37] Riku Murai, Eric Dexheimer, and Andrew J Davison. Mast3r-slam: Real-time dense slam with 3d reconstruction priors. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16695–16705, 2025. 2

[38] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*, 2017. 3, 4, 1, 2

[39] Seungjun Nah, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, Radu Timofte, and Kyoung Mu Lee. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *CVPR Workshops*, 2019. 3, 4, 1, 2

[40] Cheng Peng, Yutao Tang, Yifan Zhou, Nengyu Wang, Xijun Liu, Deming Li, and Rama Chellappa. Bags: Blur agnostic gaussian splatting through multi-scale kernel modeling. In *European Conference on Computer Vision*, pages 293–310. Springer, 2024. 2, 3, 5, 8

[41] Bang-Dang Pham, Phong Tran, Anh Tran, Cuong Pham, Rang Nguyen, and Minh Hoai. Hypercut: Video sequence from a single blurry image using unsupervised ordering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9843–9852, 2023. 4

[42] Xiaofei Qin, Haoying Ye, Changxiang He, and Xuedian Zhang. Deblur-aware gaussian splatting simultaneous localization and mapping. *Knowledge-Based Systems*, 316: 113366, 2025. 3

[43] Jaesung Rim, Haeyun Lee, Jucheol Won, and Sunghyun Cho. Real-world blur dataset for learning and benchmarking deblurring algorithms. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 3, 4

[44] Erik Sandström, Keisuke Tateno, Michael Oechsle, Michael Niemeyer, Luc Van Gool, Martin R Oswald, and Federico Tombari. Splat-slam: Globally optimized rgb-only slam with 3d gaussians. *arXiv preprint arXiv:2405.16544*, 2024. 2, 3, 4, 6

[45] Marco Sewtz, Yunis Fanger, Xiaozhou Luo, Tim Bodenmüller, and Rudolph Triebel. Indoormcd: A benchmark for low-cost multi-camera slam in indoor environments. *IEEE Robotics and Automation Letters*, 8(3):1707–1714, 2023. 2, 6, 7, 1

[46] Liao Shen, Tianqi Liu, Huiqiang Sun, Jiaqi Li, Zhiguo Cao, Wei Li, and Chen Change Loy. Dof-gaussian: Controllable depth-of-field for 3d gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26462–26471, 2025. 2

[47] Brandon Smart, Chuanxia Zheng, Iro Laina, and Victor Adrian Prisacariu. Splatt3r: Zero-shot gaussian splatting from uncalibrated image pairs. *arXiv preprint arXiv:2408.13912*, 2024. 2

[48] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A benchmark for the evaluation of rgb-d slam systems. In *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, 2012. 2, 6, 7, 8, 1

[49] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34:16558–16569, 2021. 3, 8, 1

[50] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, pages 2555–2563, 2023. 2

[51] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025. 2

[52] Peng Wang, Lingzhe Zhao, Yin Zhang, Shiyu Zhao, and Peidong Liu. Mba-slam: Motion blur aware dense visual slam with radiance fields representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2, 3, 7, 8, 1

[53] Yujie Wang, Praneeth Chakravarthula, and Baoquan Chen. Dof-gs: Adjustable depth-of-field 3d gaussian splatting for post-capture refocusing, defocus rendering and blur removal.

In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 21297–21306, 2025. 3, 8

[54] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. 7

[55] Chi Yan, Delin Qu, Dong Wang, Dan Xu, Zhigang Wang, Bin Zhao, and Xuelong Li. Gs-slam: Dense visual slam with 3d gaussian splatting. *arXiv preprint arXiv:2311.11700*, 2023. 2

[56] Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1191–1200, 2022. 2

[57] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19447–19456, 2023. 5, 8

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

[59] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 7

[60] Wei Zhang, Qing Cheng, David Skuddis, Niclas Zeller, Daniel Cremers, and Norbert Haala. Hi-slam2: Geometry-aware gaussian slam for fast monocular scene reconstruction. *arXiv preprint arXiv:2411.17982*, 2024. 2

[61] Zihao Zhang, Haoran Chen, Haoyu Zhao, Guansong Lu, Yanwei Fu, Hang Xu, and Zuxuan Wu. Eden: Enhanced diffusion for high-quality large-motion video frame interpolation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2105–2115, 2025. 3, 2

[62] Lingzhe Zhao, Peng Wang, and Peidong Liu. Bad-gaussians: Bundle adjusted deblur gaussian splatting. In *European Conference on Computer Vision*, pages 233–250. Springer, 2024. 3, 8

Unblur-SLAM: Dense Neural SLAM for Blurry Inputs

Supplementary Material

This supplementary material provides additional experimental results and further details on the experiments and methodology. The accompanying video further showcases experimental results.

6. Experiment Details

6.1. Ablation of the Fallback Mechanism

To demonstrate the effectiveness and necessity of our fallback mechanism under severe motion blur, we conduct an ablation study on the ReplicaBlurry dataset. This dataset simulates severe motion blur by averaging 36 Replica frames. We specifically select the sequence with the largest variation in motion blur degree for this evaluation.

As shown in Table 7, the inclusion of the fallback mechanism improves the reconstruction quality, raising the PSNR from 29.38 dB to 29.94 dB, demonstrating its critical role in robustly handling extreme blur cases.

Scenario	PSNR [dB]
With Fallback Mechanism	**29.94**
Without Fallback Mechanism	29.38

Table 7. **Ablation of the fallback mechanism.** ReplicaBlurry dataset simulates severe motion blur by averaging 36 Replica frames. We select the sequence with the largest variation in motion blur degree to demonstrate the effectiveness and necessity of our fallback mechanism.

6.2. Additional Information

In the qualitative experiments in Fig. 5, we run Deblur Gaussian SLAM [13] to compare the performance of its system with our system in TUM dataset [48]. It can be seen that, compared to methods based on multiple sub-frames, it is not precise for mild blurring and easily falls into local minima. Especially in the online setting, the rendered images inherently have a certain degree of blurriness due to insufficient observations.

The synthetic dataset ArchViz [52] in Table 1 was obtained through communication with the MBA-SLAM authors. For fair comparison and reasonable blur modeling, we use a deblurring neural network that was not trained on defocus datasets for preliminary blur removal. Meanwhile, to avoid introducing overfitting to synthetic data, we only used one sequence from the synthetic dataset. Specifically, this neural network was trained only on the Red [39], GoPro [38], and Office3 sequences (from the ReplicaBlurry [13] dataset) after manually filtering out frames containing moving objects.

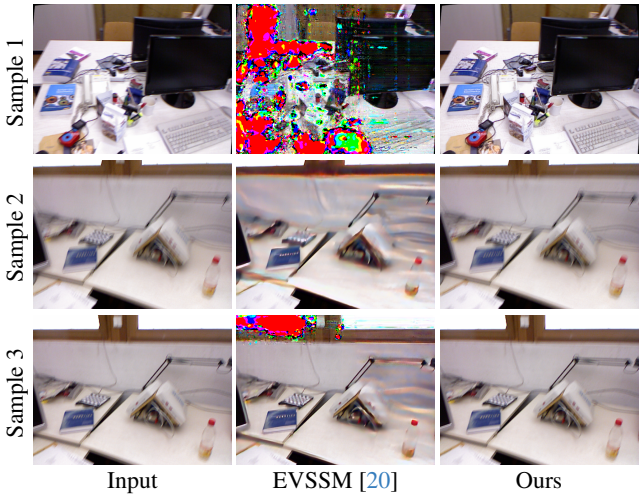

Figure 8. Comparison of input images and deblurred results

In Figure 4, we demonstrate the ATE performance on the MCD [45] s2r16 sequence to show that under our proposed training method, the single-image deblurring neural network can achieve good 3D consistency, thereby significantly improving multi-view pose estimation performance. We present the trajectory plots of Droid-SLAM [49] and our method combined with Droid-SLAM.

Table 3 shows that insufficient training data prevents the network from learning rigorous intermediate frame extraction, causing the deblurring neural network to acquire 3D-inconsistent sharpening abilities. In complex scenes, Our* augmented solely with one synthetic blur sequence cannot extract reasonable intermediate frames or less-blurred intermediate frames in more general scenarios, thus resulting in potential 3D inconsistencies.

6.3. Comparison to EVSSM

We show additional results in comparison to the recent offline method EVSSM [20]. We used the publicly available weights of EVSSM which was trained on the GoPro dataset [38]. The results are shown in Figure 8. EVSSM struggles with motion blur from moving objects along their edges, producing strange color patches and noise as camera motion and object motion occur jointly. In contrast, Unblur-SLAM decouples camera motion and object motion, and their separate processing makes our approach more robust.

7. Design Details

7.1. Tracking Module:

Tracking Module Algorithm 1 shows the structure of our tracking module. Here, $\Phi_{\text{mono}}(\cdot)$ denotes the monocular

depth estimator [10] and $b_{\text{fail}} \in \{0, 1\}$ represents the failed flag of the deblur neural network (Sec. 3.2).

7.2. Thresholds:

Thresholds are an effective solution for handling various blur levels and are fixed for all experiments (Sec. 4), based on the statistical distribution of values in the large-scale dataset (Sec. 7.4). During mapping, we employ a maximum iteration limit and robustly handle 7-12 consecutive blurry frames on the Deblur-NeRF dataset[29]. The deblurring network (Sec. 3.2) could identify the blurred frame implicitly with more computation time.

Method	PSNR [dB]	FPS
Ours_no_threshold	27.18	0.41
Ours	**29.46**	**0.74**

The table compares the average metrics to a variant without threshold on the fr1_desk, fr2_xyz and fr3_office sequences, which relies on the deblurring network for judgment.

7.3. Notion Table:

In brief, λ is the weight for the loss terms of a single frame, excluding λ_{reg} (Eq. (15)). Here is a notion overview:

Symbol	Definition	Context
x	Pixel coordinates x = (u, v) on image plane	All
f	Focal length used in defocus kernel calculation	Eq. (1)
d(t)	Camera depth at time t during exposure	Eq. (1)
σ_Δ	Standard Deviation of the Score Differences	Sec. 7.4
C, D_r	Rendered 2D RGB and depth image	Eq. (3)
d, d'	Original Gaussian depth(d) and updated tracker's depth(d')	Eq. (4)
R_k, t_k	Rotation and translation of the k-th virtual sub-frame	Sec. 3.4
θ_k, ρ_k	Lie algebra parameters for trajectory interpolation	Sec. 3.4
$w_{\text{sharp/deblur/fail/f}}$	Weights for sharp, deblurred, and failed frames	Eq. (14)-(15)
F, L_f	Union and Loss term of sharp, deblurred, and failed frames	Eq. (14)-(15)
a_{exposure}, b_{exposure}	Learnable affine Scale and Shift parameters for exposure	Eq. (8)

7.4. Blur Quantification and Recognition Details

This section provides further details on the dataset construction and metric evaluation introduced in Sec. 3.1 of the main paper.

Algorithm 1 Robust Blur-Aware Tracking Pipeline

Input: Current RGB image I_t, Pose history $H = \{T_{t-1}, T_{t-2}, \dots\}$, Deblurring failure flag b_{fail}
Output: Current camera pose T_t, Dense depth map D_t

1. Initial Depth Estimation
$D_{\text{mono}} \leftarrow \Phi_{\text{mono}}(I_t)$
2. Blur Evaluation and State Estimation
if b_{fail} **then**
  $V_{\text{rel}} \leftarrow T_{t-1} \cdot T_{t-2}^{-1}$
  $T_t \leftarrow T_{t-1} \cdot V_{\text{rel}}$
  $D_t \leftarrow D_{\text{mono}}$
else
  $(T_t, D_{\text{ref}}) \leftarrow \text{Droid-SLAM}(I_t, D_{\text{mono}}, T_{t-1})$
  $D_t \leftarrow D_{\text{ref}}$
end if
return T_t, D_t

Dataset construction. Our comprehensive blur detection benchmark incorporates real-world datasets alongside two custom semi-synthetic datasets. Following "Image as an IMU" [6], we first apply the Eden [61] video interpolation model to the ScanNet [9] dataset to generate intermediate frames. To physically model motion blur, we convert images from sRGB to linear space using inverse gamma correction, synthesize motion blur by averaging N interpolated frames, and then convert back to sRGB space:

$$I_{\text{blur}} = \gamma^{-1}\left(\frac{1}{N}\sum_{i=1}^{N}\gamma(I_{\text{linear}}^{(i)})\right) \tag{16}$$

where $\gamma(\cdot)$ and $\gamma^{-1}(\cdot)$ represent the gamma correction and its inverse. The middle interpolated frame serves as the ground truth sharp frame to ensure physical correctness.

The second semi-synthetic dataset combines sequences from the RED [39], GoPro [38], and ReplicaBlurry [13] datasets, with manual annotation to remove frames containing moving objects or inherent blur.

Blurry frames detection. To evaluate the 39 blur detection metrics, we designed a proxy measurement on a selected dataset with paired sharp and blurry images. The improvement score $\overline{\Delta}$ quantifies the magnitude of the quality difference between sharp and blurred images:

$$\overline{\Delta} = \frac{1}{M}\sum_{i=1}^{M}|s_{\text{sharp}}^{(i)} - q_{\text{blur}}^{(i)}| \tag{17}$$

where q denotes quality scores, and M is the total number of sharp/blurry pairs.

Effect size uses Cohen's d [8] to measure standardized separation: $d = \overline{\Delta}/\sigma_\Delta$. Additionally, accuracy a measures the percentage of correctly ranked blurry images. The final consistency score combines accuracy and effect size by measuring $(a \times |d|)/100$. We select the blur detection metric that achieves the highest consistency score as our blur detector, which is ARNIQA [2] (Table 8).

Metric	Acc. [%]	Effect size	Consistency score
ARNIQA-CSIQ [2]	92.93	**37.01**	**34.39**
MANIQA-KADID [56]	96.55	6.35	6.13
CLIPIQA+ [50]	96.58	2.26	2.18
MUSIQ [15]	95.13	2.14	2.03
BRISQUE [34]	79.78	1.33	1.06

Table 8. Blur detection metric. We compared the accuracy and consistency scores of 39 blur metrics and selected 5 representative ones for reporting. We chose the metric with the highest consistency score among them as our blur detector (Sec. 3.1).

8. Additional qualitative results

We conducted additional qualitative experiments on various datasets to demonstrate the robustness of our method. As

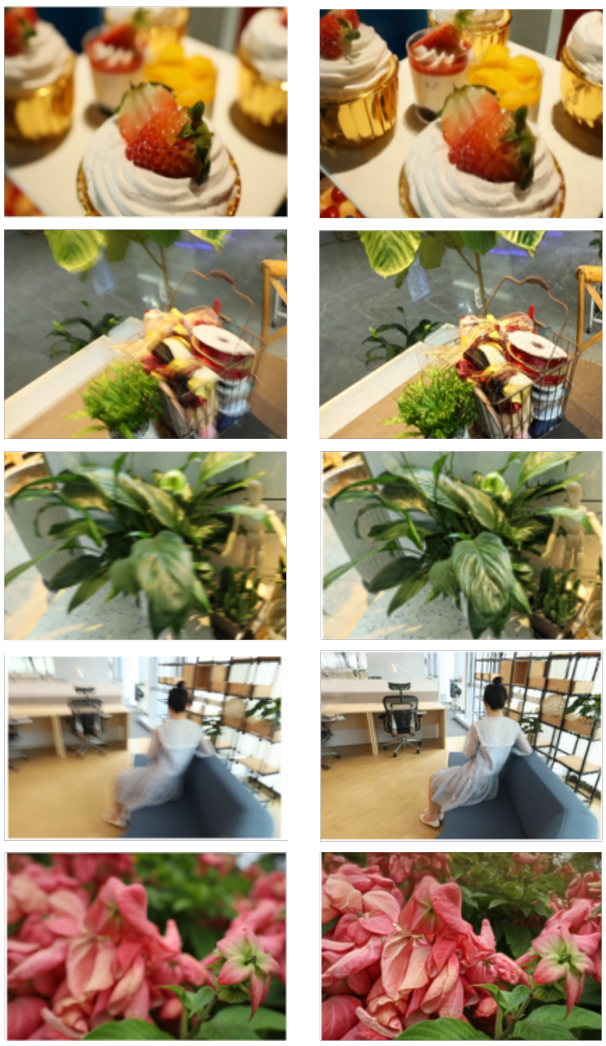

Input Ours

Figure 9. **Qualitative results on various datasets.** The left column shows the input images, while the right column displays the rendered images of our reconstructed 3D models.

shown in Figure 9, our approach exhibits improved detail recovery in various scenarios.

9. Limitations and Future Work

The physically constrained deblurring neural network training method proposed in Sec. 3.2 is unsuitable for images where the linear RGB color space cannot be accurately restored. These images typically undergo complex optimization algorithms by mobile phone pipelines, rendering them nearly impossible to reverse-engineer. Furthermore, while the proposed fallback mechanism for severe blur modeling renders multiple virtual sub-frames, the memory footprint grows explosively as the number of sub-frames increases. This memory constraint inherently limits its further expansion and scalability within online methodologies.

Nevertheless, this work proposes a reasonable engineering pipeline to implement a theoretical framework aimed at promoting progress within the community. Given the rapid development in the SLAM field, pure neural network-based approaches represent a highly promising direction for the further development of this theory.