
ProDyG: Progressive Dynamic Scene Reconstruction via Gaussian Splatting from Monocular Videos

Shi Chen
ETH Zürich

Erik Sandström
Google

Sandro Lombardi
Independent Researcher

Siyuan Li
ETH Zürich

Martin R. Oswald
University of Amsterdam

Abstract

Achieving truly practical dynamic 3D reconstruction requires online operation, global pose and map consistency, detailed appearance modeling, and the flexibility to handle both RGB and RGB-D inputs. However, existing SLAM methods typically merely remove the dynamic parts or require RGB-D input, while offline methods are not scalable to long video sequences, and current transformer-based feedforward methods lack global consistency and appearance details. To this end, we achieve online dynamic scene reconstruction by disentangling the static and dynamic parts within a SLAM system. The poses are tracked robustly with a novel motion masking strategy, and dynamic parts are reconstructed leveraging a progressive adaptation of a Motion Scaffolds graph. Our method yields novel view renderings competitive to offline methods and achieves on-par tracking with state-of-the-art dynamic SLAM methods. Project page: cs-vision.github.io/ProDyG.github.io

1 Introduction

Dynamic scene reconstruction is fundamental to problems like action recognition, scene understanding, autonomous driving, robotics and augmented reality, because it provides a temporally consistent spatial understanding of how objects and agents move and interact in their environment — an essential prerequisite for any system to perceive, predict, and act in the world. This problem has been tackled in various ways *e.g.* as an online SLAM task [4, 89, 34], an offline reconstruction task [29, 73, 79, 39, 56, 55], and lately as a feedforward task [86, 87, 11, 69]. However, most dynamic SLAM works ignore the dynamic parts [89, 4], only reconstructing the static world, only track rigid objects [85, 43] or are restricted to object-centric reconstruction [45, 58]. Offline methods typically separate the reconstruction task into pose estimation followed by reconstruction [55, 56], or are not scalable to long input videos due to their reliance on global optimization over all past frames [29, 73]. Feed-forward methods train large transformers for online dynamic scene reconstruction, but are yet to achieve global pose consistency and only produce point clouds [69, 11, 86].

Despite recent progress, existing methods fall short of at least one of the requirements for practical dynamic scene reconstruction: (1) **online operation**, tightly coupling pose estimation and dense map reconstruction for scalability; (2) **global pose and map consistency**; (3) **expressive representations** like 3DGS [23] for detailed appearance and geometry; and (4) **flexibility** to handle both **RGB and RGB-D** input. SLAM systems often ignore dynamics or lack detail; feedforward methods trade consistency and accuracy for speed; and offline methods are challenging to scale.

We propose **ProDyG**¹, a method for **Progressive, Dynamic** scene reconstruction with **Gaussians** from monocular input, that meets all four criteria. Our contributions are:

¹Pronounced “Prodigy”.



Figure 1: **ProDyG**: Given an online RGB stream, ProDyG robustly tracks the camera and progressively reconstructs the static background and the dynamic foreground using 3D Gaussian Splatting [23] (3DGS). ProDyG reconstructs dynamic scenes with high-quality novel view synthesis, by fusing observations from different timestamps into a consistent dynamic 3DGS [23] representation with Motion Scaffolds [29]. On the right, we show a rendering from a novel view overlaid with the estimated camera trajectory and dynamic Gaussian trajectories.

- A motion mask prediction strategy, using dynamic flow and refined with semantic point prompting [50]. This is integrated into the SLAM backend for robust pose estimation.
- An **online** dynamic reconstruction pipeline that uses Motion Scaffolds [29] to propagate 3DGS in space and time.
- Support for both RGB and RGB-D, with ProDyG as the **first online RGB-only** method.
- Competitive novel view synthesis results against state-of-the-art offline methods.

2 Related Work

Dynamic Dense SLAM. Dynamic dense SLAM methods mainly address two scenarios: filtering dynamic content to reconstruct only static environments [4, 13, 90, 89, 77, 19, 15, 81, 44, 5, 74, 59, 48, 57, 82, 37, 32, 88, 1, 24, 7, 64, 26, 72, 51], or explicitly tracking and reconstructing dynamic objects alongside static structures [85, 3, 43, 17, 52, 53, 78, 14, 45, 58, 8, 65, 34, 35, 31]. Basic strategies for handling dynamic elements include outlier filtering, robust loss functions [90, 44, 5], or covariance scaling [1], but these fail with extensive dynamic content. Instead, a motion mask is estimated with optical flow [88, 82], semantic segmentation [26, 14, 35, 57, 4, 33, 77, 51, 81, 74, 85, 59, 72, 53] or a combination of the two [34, 19, 15, 52], with uncertainty-guided loss functions [89, 37], with motion segmentation networks [57], point tracking [65], unsupervised clustering [17] or via conditional random fields [32]. In many real-world scenarios, reconstructing dynamic objects is critical. Some traditional works track rigid dynamic content during SLAM [85, 3, 43, 17, 52, 53, 78, 14] or perform object-centric non-rigid reconstruction [45, 58, 8]. While the above-mentioned methods can reconstruct dynamic point clouds [4, 85, 74, 59, 88, 17], surfels [52, 53], oct-trees [81] or signed distance functions [48, 45, 58, 8], they struggle to model photometric details and lighting effects needed for scene understanding and photorealistic rendering. In response, neural implicit methods [33, 77, 51, 19] have been proposed, but they are too slow for real-time SLAM. Dynamic 3DGS [23] SLAM has emerged as a solution [15, 34, 35, 31, 89, 32, 72, 26] with concurrent works such as DynaGSLAM [34] using dynamic scene flow from optical flow to update the Gaussian means, [65] adapts [79] to the online setting and [35] uses an MLP to deform the dynamic parts. These works cannot handle pure RGB input. For a recent survey on dynamic SLAM, we refer to [71].

Dynamic 3D Reconstruction. Dynamic 3D reconstruction shares many similarities with dynamic SLAM, but videos are processed offline [29, 73, 79, 9, 21, 80, 27, 28, 36, 40, 41, 49, 60, 84, 2, 16, 39, 25, 91, 18, 38, 68]. These works typically implement a strategy for deforming dynamic content over time. Some optimize a deformation field via an MLP [2, 60, 40, 80, 25, 73], use motion basis functions [16, 49, 68, 36, 29, 28, 27, 21, 39, 18], or extend 3D Gaussian Splatting [23] with a time attribute [79, 9]. Due to the ill-posed nature of dynamic reconstruction, strong priors are commonly used, such as point trackers, optical flow, and regularizers like as-rigid-as-possible (ARAP), constant velocity, and acceleration [29, 41, 56, 68, 39]. Inductive biases also help, such as MLP smoothness

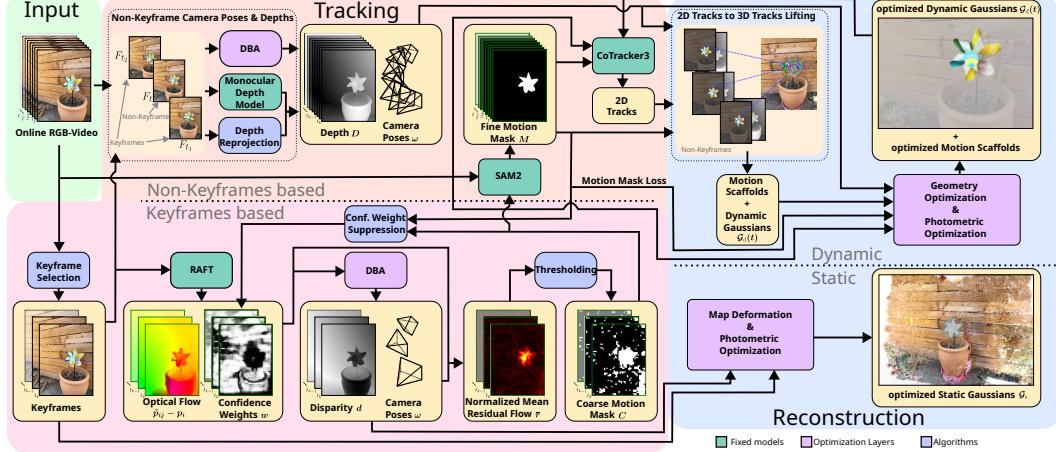


Figure 2: **ProDyG Architecture.** We achieve motion-agnostic online tracking by leveraging [54] to first create keyframe-based coarse motion masks, from which we seed prompts for SAM2 to distill per-frame fine-grained masks. ProDyG processes batches of frames incrementally, employing a keyframe selection similar to [54]. Static background is reconstructed by optimizing the static set of Gaussians with proxy depth maps [54]. For dynamic reconstruction, we attach Gaussians to Motion Scaffolds [29], which are initialized by lifting 2D tracks to 3D, to encode a dense motion field. Subsequent to a final geometric and photometric optimization, the Motion Scaffolds and dynamic Gaussians are extended temporally when a new batch of images arrives.

and continuity, or selecting a sparse set of motion bases. While some of these works use SLAM for pose estimation, reconstruction is a post-processing step [55, 56].

Feed-forward Methods. Feed-forward methods inspired by DUST3R [70] and MAST3R [30] have been proposed to solve dynamic 3D reconstruction. These works rely on large-scale training of transformer networks and can infer pointmaps from dynamic input image pairs [86, 87, 11, 13, 6, 69]. However, they typically entangle static and dynamic points [69, 86], meaning that the motion of dynamic points cannot be tracked over time. In concurrent work, [11, 13, 87] propose to enable 3D correspondence estimation of dynamic 3D points. DAS3R [75] learns to predict motion masks building on MonST3R [86], and only predicts the static world, using 3DGS [23]. Common to all feed-forward methods is that they process the video either in image pairs or in a sliding window, without guaranteeing global pose and map consistency, contrary to SLAM. In contrast, we perform globally consistent and online dynamic reconstruction using SLAM.

3 Method

ProDyG is an online dynamic dense mapping and tracking system that robustly tracks a monocular camera (sec. 3.1) while also reconstructing and disentangling the static and dynamic parts (sec. 3.3) with a 3D Gaussian Splatting representation (sec. 3.2). For an overview, see fig. 2.

3.1 Motion-Agnostic Online Camera Tracking

Flow-Based Robust Camera Tracking. We employ Splat-SLAM [54] as our tracking backend, *i.e.* we maintain a factor graph storing camera extrinsics ω_i , disparity estimates $d_i \in \mathbb{R}^{(H \times W) \times 1}$ per keyframe (node) i , and optical flow $\tilde{p}_{ij} - p_i \in \mathbb{R}^{(H \times W \times 2) \times 1}$ per edge (i, j) , where p_i is the flattened pixel grid from keyframe i , and \tilde{p}_{ij} represents the flattened predicted pixel coordinates when p_i is projected into keyframe j using optical flow. We also store the confidence $w_{ij} \in \mathbb{R}^{(H \times W) \times 1}$ associated with the optical flow. Tracking is achieved following Dense Bundle Adjustment (DBA) [66], where the keyframe poses ω and disparities d are optimized with a reprojection error:

$$\arg \min_{\omega, d} \sum_{(i, j) \in \mathcal{E}} \|\tilde{p}_{ij} - p_{ij}\|_{\Sigma'_{ij}}^2, \quad p_{ij} = K\omega_j^{-1}(\omega_i(1/d_i)K^{-1}[p_i, 1]^T), \quad \Sigma'_{ij} = \text{diag}(w_{ij}\bar{C}_i). \quad (1)$$

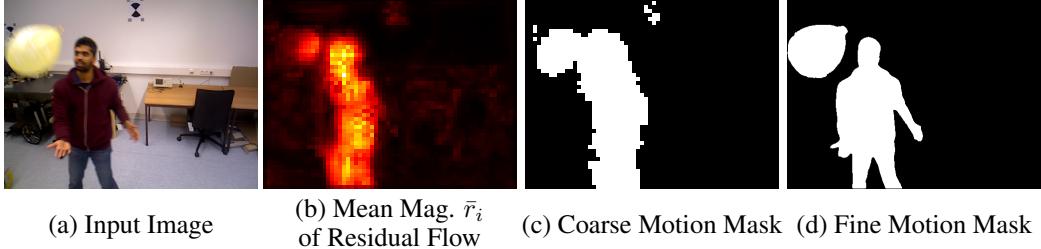


Figure 3: **Semantic-guided Motion Mask Refinement.** The flow magnitude \bar{r}_i (b) is thresholded to yield (c). Point prompts from (c) are used as input to SAM2 [50] to yield the fine motion masks (d).

Here, \mathcal{E} are the edges of a local factor graph, applied in a sliding window manner, K is the camera intrinsics, $\|\cdot\|_{\Sigma_{ij}}$ denotes the Mahalanobis distance with confidence weights suppressed by a coarse binary motion mask C_i , defined as 1 where motion is detected and \bar{C}_i stands for negation of C_i .

We generate the coarse motion masks $\{C_i\}$ as follows. After each DBA iteration, we compute the residual flow $\hat{r}_{ij} = (\tilde{p}_{ij} - p_i) - (p_{ij} - p_i) = \tilde{p}_{ij} - p_{ij}$ by subtracting the camera-induced flow $p_{ij} - p_i$ from the estimated optical flow $\tilde{p}_{ij} - p_i$. This residual is near zero in static regions and larger in dynamic ones. Dynamic areas are identified by evaluating the normalized mean magnitude $\bar{r}_i(x, y)$ of \hat{r}_{ij} over the connected target keyframes N_i for each keyframe i

$$\bar{r}_i(x, y) = \frac{1}{|N_i|} \sum_{j \in N_i} \frac{\|\hat{r}_{ij}(x, y)\|}{\frac{1}{HW} \sum_{x', y'} \|\tilde{p}_{ij}(x', y') - p_i(x', y')\|}. \quad (2)$$

For each keyframe, C_i is computed by thresholding $\bar{r}_i(x, y)$ at the top 20%, assigning zero weight to potentially dynamic regions during DBA. As $\bar{r}_i(x, y)$ is updated each DBA iteration, it is progressively refined, enabling robust pose estimation in dynamic environments.

Semantic-Guided Motion Mask Refinement. To improve the separation of dynamic and static regions for reconstruction (sec. 3.3), we use SAM2 [50] to generate fine-grained motion masks for every frame. We show an example visualization of the input image, the mean magnitude \bar{r}_i of the residual flow, the coarse motion mask, and the fine motion mask in fig. 3.

Our approach has two phases: initialization and incremental prediction. During initialization, we generate fine motion masks by prompting SAM2 at the centroids of connected regions in the median-filtered coarse masks C_i . During incremental prediction, we extend existing object-wise masks into new frames and add new segmentations based on the coarse masks. We validate each segmentation by counting prompt point candidates to filter false positives. Before this phase, we adjust the threshold (top 20% at initialization) based on the dynamic pixel ratio in the latest fine motion mask, adapting to sequence-specific motion. Detailed algorithms are in the supplemental material.

With the fine motion masks $\{M_i\}$, we further enhance our tracking robustness by replacing the coarse mask C_i with the fine during DBA (eq. (1)) as $\Sigma'_{ij} = \text{diag}(w_{ij}(1 - M_i))$. This semantic-guided approach provides more accurate exclusion of dynamic distractors for subsequent iterations of DBA, including local BA, global BA and loop BA [54].

3.2 Static and Dynamic Map Representation

We represent the scene map $\mathcal{G} = \{\mathcal{G}_s, \mathcal{G}_d(t)\}$ as a set of 3D Gaussians [23], split into static $\mathcal{G}_s = \{g_i\}_{i=1}^{N_s}$ and time-varying dynamic components $\mathcal{G}_d(t) = \{g_i(t)\}_{i=1}^{N_d}$. Each Gaussian—whether static or dynamic—is parameterized by a mean $\boldsymbol{\mu}_i \in \mathbb{R}^3$, rotation matrix $R_i \in \mathbb{R}^{3 \times 3}$, scale $\mathbf{s}_i \in \mathbb{R}^3$, opacity $o_i \in [0, 1]$, and RGB color $\mathbf{c}_i \in \mathbb{R}^3$. Rendering is performed by first projecting the 3D Gaussians onto the image plane, approximated as 2D Gaussians. The 2D opacity at a pixel p is then

$$\alpha_i(p) = o_i \exp \left(-\frac{1}{2} (p - \boldsymbol{\mu}'_i)^\top \Sigma'_i (p - \boldsymbol{\mu}'_i) \right), \quad (3)$$

where $\boldsymbol{\mu}'_i = K\omega^{-1}\boldsymbol{\mu}_i$ with ω the camera-to-world pose, K the camera intrinsics, and $\Sigma'_i = P\Sigma_i P^\top = P(R_i S_i S_i^\top R_i^\top)P^\top$, with $S_i = \text{diag}(\mathbf{s}_i)$ and P the affine projection [92]. Given the pose

ω , we render a pixel p following the 3DGS [23] pipeline for RGB images C , RaDe-GS [83] for depth maps D^r , and additionally compute opacity maps O as

$$C(p) = \sum_{i \in \mathcal{K}} \mathbf{c}_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad D^r(p) = \sum_{i \in \mathcal{K}} \hat{d}_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad O(p) = \sum_{i \in \mathcal{K}} \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j). \quad (4)$$

Here, \mathcal{K} is the set of Gaussians projected to p . We use the unbiased RaDe-GS rasterizer, which replaces the z-buffer depth with \hat{d}_i , the ray–Gaussian intersection depth along the viewing ray.

3.3 Progressive Dynamic Scene Reconstruction

We propose a progressive approach to dynamic scene reconstruction, which enables online processing and allows the system to adapt to the continuously evolving scene geometry and camera motion. ProDyG uses the same keyframe selection strategy as in Splat-SLAM [54], and the addition of new keyframes triggers progressive reconstruction of both static and dynamic regions. To achieve global map consistency, we apply 3DGS map deformations for static regions followed by photometric optimization as in [54]. For dynamic Gaussians, we extend the Motion Scaffold representation [29] to handle incremental updates. Next, we review the necessary steps.

Non-Keyframe Camera Poses and Depths. Dynamic scene reconstruction benefits from temporally dense supervision to capture smooth motion, yet the SLAM backend [54] operates on sparse keyframes for efficiency. To bridge this gap, we derive the camera poses and depth maps for non-keyframes that exist between keyframes.

To obtain non-keyframe camera poses, we follow the practice of DROID-SLAM [66]. For each non-keyframe F_t positioned between two neighboring keyframes F_{t_1} and F_{t_2} , we construct a temporary local graph with edges connecting both keyframes to F_t . We then execute the recurrent update operator and the Dense Bundle Adjustment (DBA) layer on this graph, which optimizes the non-keyframe poses based on the already optimized keyframe poses.

For non-keyframe depth estimation, we leverage reprojection from neighboring keyframes. Given a non-keyframe F_t with its derived camera-to-world pose ω_t and neighboring keyframes F_{t_1} and F_{t_2} with poses ω_{t_1} and ω_{t_2} , we first back-project the pixels of both keyframes into 3D:

$$p_{t_1}^{3D}(p_{t_1}) = \omega_{t_1}(1/d_{t_1})K^{-1}[p_{t_1}, 1]^T, \quad p_{t_2}^{3D}(p_{t_2}) = \omega_{t_2}(1/d_{t_2})K^{-1}[p_{t_2}, 1]^T, \quad (5)$$

where K is the camera intrinsic matrix, and d_{t_1} and d_{t_2} are the disparities of keyframes F_{t_1} and F_{t_2} , respectively. We then reproject all back-projected 3D points $p_{t_1, t_2}^{3D} = p_{t_1}^{3D} \cup p_{t_2}^{3D}$ from both keyframes to the camera of the non-keyframe F_t and record their distance to the camera along the z-axis:

$$p_t^{3D} = \omega_t^{-1} p_{t_1, t_2}^{3D}, \quad p_t = K p_t^{3D}, \quad D_t^{\text{repro}}(p_t) = (p_t^{3D})_z. \quad (6)$$

We filter out points that are back-projected from dynamic pixels or using invalid disparity estimates (determined by a multi-view consistency check as in [54]). Using the remaining reprojected 2D coordinates and their corresponding depth values, we bilinearly interpolate to create a reprojection depth map D_t^{repro} . Finally, to get dense depth, we align a monocular depth estimate D_t^{mono} with D_t^{repro} by estimating scale θ_t and shift γ_t parameters through least squares fitting:

$$\theta_t, \gamma_t = \arg \min_{\theta, \gamma} \sum_{(u, v)} \left((\theta D_t^{\text{mono}}(u, v) + \gamma) - D_t^{\text{repro}}(u, v) \right)^2, \quad D_t^{\text{aligned}} = \theta_t D_t^{\text{mono}} + \gamma_t. \quad (7)$$

Motion Scaffolds. We adopt Motion Scaffolds [29] (MoSca) as our dynamic representation. MoSca a structured graph $(\mathcal{V}_d, \mathcal{E}_d)$ representing the underlying motion of the scene, where each node $v^{(m)} \in \mathcal{V}_d$ encodes the motion trajectory of a specific region. Each MoSca node $v^{(m)}$ is defined as

$$v^{(m)} = ([\mathbf{Q}_1^{(m)}, \mathbf{Q}_2^{(m)}, \dots, \mathbf{Q}_T^{(m)}], r^{(m)}), \quad (8)$$

where $\mathbf{Q}_t^{(m)} \in SE(3)$ represents the per-timestep rigid transformation at time t , and $r^{(m)}$ is an RBF radius parameterizing the node's influence. These nodes are initially anchored by lifting 2D pixel trajectories using the estimated camera poses ω_t and depth maps D_t^{aligned} computed in eq. (7).

Similar to the static map deformation [54], we also deform the dynamic Gaussians to reflect the continuously updated pose and depth estimates from the SLAM backend. Each dynamic Gaussian is defined relative to a MoSca node (not in world coordinates) and by re-anchoring the MoSca nodes, the dynamic Gaussians are updated accordingly. The MoSca nodes are re-anchored in 3D when new pose and depth updates are available. Thus, ProDyG achieves global dynamic map consistency prior to optimization. This ensures better convergence and consistency between the static and dynamic components of the scene.

Following MoSca [29], the set of dynamic Gaussians $\mathcal{G}_d(t)$ at any query timestamp t is formed by warping each individual Gaussian $g_i(t_i^{\text{ref}}) = (\mu_i, R_i, \mathbf{s}_i, o_i, \mathbf{c}_i; t_i^{\text{ref}}, \Delta\mathbf{w}_i)$ from its reference timestamp t_i^{ref} (the timestamp where it is initialized) to t

$$\mathcal{G}_d(t) = \{(\mathbf{T}_i(t)\mu_i, \mathbf{T}_i(t)R_i, \mathbf{s}_i, o_i, \mathbf{c}_i \mid \mathbf{T}_i(t) = \mathcal{W}(\mu_i, \mathbf{w}(\mu_i) + \Delta\mathbf{w}_i; t_i^{\text{ref}}, t)\}_{i=1}^{N_d} , \quad (9)$$

Here, $\mathbf{w}(\cdot)$ is the base RBF skinning weight parametrized by $\{r^{(m)}\}_{m \in \mathcal{E}_d(m^*)}$, where the neighborhood $\mathcal{E}_d(m^*)$ consists of the nearest MoSca node $v^{(m^*)}$ and all nodes connected to $v^{(m^*)}$, and $\Delta\mathbf{w}_i$ are learnable per-Gaussian skinning weight corrections. The warping function $\mathcal{W}(\cdot)$ is computed using Dual Quaternion Blending (DQB) [22]:

$$\mathcal{W}(\mathbf{x}, \mathbf{w}; t_{src}, t_{dst}) = \text{DQB} \left(\{w_m, \Delta\mathbf{Q}^{(m)}\}_{m \in \mathcal{E}_d(m^*)} \right) , \quad (10)$$

where $\Delta\mathbf{Q}^{(m)} = \mathbf{Q}_{t_{dst}}^{(m)}(\mathbf{Q}_{t_{src}}^{(m)})^{-1}$ is the relative transformation between t_{src} and t_{dst} for node m .

Progressive Construction of Motion Scaffolds. After the initial bootstrapping phase of the SLAM backend, we initialize the Motion Scaffolds and dynamic Gaussians. First, we employ CoTracker3 [20] to generate dense long-term 2D pixel trajectories within the fine motion masks $\{M_i\}$ (sec. 3.1) and corresponding per-timestamp visibility labels. Similar to [29], these 2D trajectories are first lifted into 3D space at visible timestamps using the camera poses and depth maps estimated by the backend, while we linearly interpolate between nearby observations at invisible timestamps. Finally, we sample a subset of the lifted 3D tracks to serve as the initial positions for the MoSca nodes.

Subsequently, we carry out geometry optimization on the initialized MoSca nodes as described in [29] to infer rotations and positions of invisible nodes, minimizing the as-rigid-as-possible (ARAP) loss, velocity consistency loss and acceleration consistency loss.

After the geometry optimization of the Motion Scaffolds, we initialize the dynamic Gaussians at 3D positions obtained by back-projecting pixels within the fine motion masks. Finally, we perform a photometric optimization of both the Motion Scaffolds and the dynamic Gaussians using a combination of losses following [29]: an RGB loss \mathcal{L}_{rgb} that enforces color consistency, a depth loss $\mathcal{L}_{\text{depth}}$ that aligns rendered depth D^r with the aligned monocular depth D^{aligned} (eq. (7)), a track loss $\mathcal{L}_{\text{track}}$ that ensures consistent motion with the 2D trajectories, and the aforementioned ARAP, velocity, and acceleration losses for geometric regularization. In addition to [29], we introduce a novel motion mask loss that penalizes the rendered opacity of dynamic Gaussians at pixels identified as static by the motion masks

$$\mathcal{L}_{\text{mask}} = \frac{1}{|\mathcal{P}_t^{\text{static}}|} \sum_{p \in \mathcal{P}_t^{\text{static}}} O(p) , \quad (11)$$

where $\mathcal{P}_t^{\text{static}}$ is the set of static pixels at timestamp t given by negating the motion mask M_t , and $O(p)$ is the accumulated opacity of dynamic Gaussians at pixel p (eq. (4)). This loss effectively prevents dynamic Gaussians from overflowing into static regions, maintaining a clean separation between static and dynamic components of the scene.

As more frames become available, we extend our dynamic reconstruction. For each new frame batch, we run CoTracker3 [20] in a temporal window of the new frame batch and an 8-frame overlap with previously processed frames to ensure continuity in the reconstruction. First, we identify “recently visible” tracks as those marked as visible for at least 4 frames within the 8 overlapping frames and extend them into the new frame batch. These extended 2D tracks are then lifted into 3D using the latest camera pose and depth estimates. To identify newly visible dynamic regions, we back-project

all pixels within the fine motion masks of the new frames into 3D and perform a spherical search to determine whether each back-projected point has at least one lifted 3D track within a predefined radius r_{search} . Pixels without nearby 3D tracks are marked as “newly-seen”, representing previously unobserved portions of dynamic objects. To capture these newly-seen regions, we execute a second run of the point tracker, specifically querying newly-seen pixels. Finally, we run a third tracking pass within the entire dynamic regions of the new frames to replenish the density of visible 2D tracks. This multi-stage tracking strategy ensures dense coverage of all dynamic elements in the scene.

Within the new temporal window, all extended and newly added 2D tracks are lifted into 3D using the same procedure as during initialization. For temporal consistency, we warp new tracks to past timestamps with DQB (Eqn.10), using Motion Scaffolds from the previous update. This leverages the latest photometric optimization to guide newly-seen 3D tracks through past frames despite invisibility. We initialize dynamic Gaussians only at newly-seen pixels, then jointly optimize geometry and appearance over the expanded MoSca and dynamic Gaussians using the same losses as during initialization.

4 Experiments

4.1 Experimental Setup

Implementation Details. All experiments were conducted on a cluster with an AMD EPYC 7H12 or 7742 CPU and an NVIDIA A6000 GPU. The kernel size of the median filter used to denoise the coarse motion masks is 5×5 . The spherical search radius for “newly-seen” pixel identification is $r_{\text{search}} = 0.02m$. For geometry and photometric optimization, we keep our loss weights identical with those applied in MoSca [29], and set $\lambda_{\text{mask}} = 1$ as the weight of $\mathcal{L}_{\text{mask}}$. For more implementation details, we refer to the supplemental material.

Datasets. We evaluate our camera tracking on the Bonn RGB-D Dynamic Dataset [47] and the TUM RGB-D Dataset [62] (dynamic scenes). Since existing works report tracking results on different sets of sequences, we select four mostly used sequences from each dataset to evaluate our method. For rendering, we report novel view synthesis (NVS) results both qualitatively and quantitatively on the iPhone dataset [12]. To align with Shape of Motion [68], we evaluate our method and all baselines on the 5 sequences used in [68] with the 2x downsampled image resolution. For a fair comparison, we use the preprocessed motion masks given by [68] for all our experiments on the iPhone Dataset.

Baselines. For tracking, we compare with various works on RGB and RGB-D SLAM. The main baseline is Splat-SLAM [54] since we base our tracking pipeline on it. For rendering, the baseline methods are NVS-capable monocular Gaussian-based dynamic reconstruction methods including Shape of Motion [68], DynOMo [56], MoSca [29] and Gaussian Marbles [61].

Metrics. For tracking, we evaluate ATE RMSE [cm] [63] after aligning the estimated camera trajectory with the ground truth via Umeyama alignment [67]. For NVS, we report PSNR, SSIM and LPIPS evaluated within the covisibility masks provided by [12] and averaged over all novel views.

4.2 Tracking

In tab. 1, we evaluate tracking performance on the Bonn RGB-D Dynamic Dataset [47] and the TUM RGB-D Dataset [62]. For both datasets, ProDyG performs competitively among all the RGB-D and RGB SLAM works and shows a significant advantage over the main baseline Splat-SLAM [54]. This improvement validates the effectiveness of our motion-agnostic camera tracking method introduced in sec. 3.1. WildGS-SLAM [89] is the only baseline method to outperform ProDyG on both datasets. Since both methods build upon the tracking framework of Splat-SLAM [54], we attribute this performance gap to two key factors: (1) WildGS-SLAM employs a test-time-optimized MLP that produces soft uncertainty masks to suppress confidence weights. As it reconstructs only the static background, it can aggressively suppress regions beyond actual dynamic object boundaries (e.g. moving shadows) without degrading mapping quality. In contrast, ProDyG prioritizes accurate dynamic reconstruction and therefore requires motion masks with precise boundaries, which may miss some effective distractors. (2) WildGS-SLAM benefits from additional DINOv2 [46] features and on-the-fly training of the uncertainty MLP at the cost of higher computational complexity, while the primary computational overhead of our tracking approach comes from SAM2 [50] inference.

Method	Type	Bonn RGB-D Dynamic Dataset [47]					TUM RGB-D Dataset [62]				
		Ball	Ball2	Pers	Pers2	Avg.	f3/ws	f3/wx	f3/wr	f3/whs	Avg.
<i>RGB-D Input</i>											
ORB-SLAM2 [44]	S	6.5	23.0	6.9	7.9	11.1	40.8	72.2	80.5	72.3	66.45
NICE-SLAM [90]	S	24.4	20.2	24.5	53.6	30.7	79.8	86.5	244.0	152.0	140.57
ReFusion [48]	R	17.5	25.4	28.9	46.3	29.5	1.7	9.9	40.6	10.4	15.7
DynaSLAM (N+G) [4]	R	3.0	2.9	6.1	7.8	5.0	0.6	1.5	3.5	2.5	2.03
DG-SLAM [76]	R	3.7	4.1	4.5	6.9	4.8	0.6	1.6	4.3	-	-
RoDyn-SLAM [19]	R	7.9	11.5	14.5	13.8	11.9	1.7	8.3	-	5.6	-
DDN-SLAM (RGB-D) [33]	R	1.8	4.1	4.3	3.8	3.5	1.0	1.4	3.9	2.3	2.15
<i>RGB Input</i>											
DSO [10]	S	7.3	21.8	30.6	26.5	21.6	1.5	12.9	13.8	40.7	17.23
DROID-SLAM [66]	S	7.5	4.1	4.3	5.4	5.3	1.2	1.6	4.0	2.2	2.25
MonoGS [42]	S	15.3	17.3	26.4	35.2	23.6	1.1	21.5	17.4	44.2	21.05
Splat-SLAM [54]	S	8.8	3.0	4.9	25.8	10.6	2.3	1.3	3.9	2.2	2.43
DDN-SLAM (RGB) [33]	R	-	-	-	-	-	2.5	2.8	8.9	4.1	4.58
MegaSaM [37]	R	3.7	2.6	4.1	4.0	3.6	0.6	1.5	2.6	1.8	1.63
WildGS-SLAM [89]	R	2.7	2.4	3.6	3.1	2.94	0.4	1.3	3.3	1.6	1.63
DynaMoN (MS) [55]	D	6.8	3.8	2.4	3.5	4.1	1.4	1.4	3.9	2.0	2.18
DynaMoN (MS&SS) [55]	D	2.8	2.7	14.8	2.2	5.6	0.7	1.4	3.9	1.9	1.98
D4DGS-SLAM* [65]	D	3.6	3.9	4.5	5.2	4.3	-	-	-	-	-
4D-GS SLAM* [35]	D	2.4	3.7	8.9	9.4	6.1	0.5	2.1	2.6	-	-
ProDyG (Ours)	D	2.7	2.6	4.9	2.9	3.29	1.6	1.2	3.0	1.7	1.89

Table 1: **Tracking Performance on Bonn RGB-D Dynamic Dataset [47] and TUM RGB-D Dataset [62].** (ATE RMSE \downarrow [cm]). Best results are highlighted as **first**, **second**, **third** and concurrent works with*. We take the numbers from [89] except [89, 65, 35]. We categorize each method into static reconstruction (S), robust against dynamics (R) and producing a globally consistent dynamic model (D). ProDyG is competitive with WildGS-SLAM [89] and MegaSaM [37] while explicitly reconstructing a consistent dynamic model, contrary to [89, 37].

4.3 Rendering

In tab. 2, we evaluate NVS performance quantitatively on the iPhone Dataset [12]. Notably, we differentiate the experiment setups of each methods using the checkboxes, whether the reconstruction or tracking is executed online, and whether the input modalities are RGB only or RGB-D. As shown in tab. 2, ProDyG outperforms the offline method Shape of Motion [68] in both PSNR and SSIM, while running both the tracking and dynamic scene reconstruction online. This is a significantly more difficult task. Compared to MoSca [29], which is an offline method that shares the same representation of the motion field with our framework, ProDyG shows only minimal disadvantages when optimized with RGB-D input. This demonstrates the effectiveness of our progressive dynamic reconstruction. DynOMo [56] is the only method capable of online reconstruction with precomputed camera poses, and ProDyG shows a significant advantage over DynOMo when tested under the same constraint. Furthermore, ProDyG still maintains reasonably good performance when only having access to RGB images, while none of the baseline methods is capable of operating under the same setting. Finally, we evaluate NVS when estimated poses from the SLAM backend are loaded offline to the online mapper. The performance is very similar in both RGB-D and RGB-only modes, showing that our method can find the offline solution, even when optimized in pure online mode.

In fig. 4, we show qualitative comparisons of novel view renderings on the iPhone Dataset [12]. Compared to other methods, DynOMo [56] exhibits worse quality in novel view renderings due to lack of motion constraints (e.g. 2D point tracking). Therefore, ProDyG is essentially the first method to support high-quality novel view synthesis of dynamic scenes through online reconstruction from monocular videos. Due to the motion mask loss (eq. (11)), dynamic objects reconstructed by our method tend to show more accurate silhouettes than Shape of Motion [68] and MoSca [29].

Limitations. Our dynamic representation struggles with objects that move predominantly outside the viewing frustum and later reappear, due to insufficient photometric constraints for Motion Scaffolds. This causes optimization to be dominated by regularization terms, often leading to undesired deformations. This issue is an inherent limitation also observed in MoSca [29]. Additionally, similar to most monocular dynamic view synthesis methods, our system requires per-scene test-time optimization, thereby limiting its applicability in real-time scenarios. Furthermore, ProDyG cannot effectively handle large changes in novel viewpoints, which would necessitate generative models or data-driven priors to hallucinate unseen regions. These challenges highlight key directions for future work.

Shape of Motion [68]	DynOMo [56]	MoSca [29]	Gaussian Marbles [61]	ProDyG (Ours)	ProDyG (Ours)	ProDyG (Ours)	ProDyG (Ours)
Online Reconstr.	✗	✓	✗	✗	✓	✓	✓
Online Tracking	✗	✗	✗	✗	✗	✓	✗
RGB-only	✗	✗	✗	✗	✗	✓	✓
PSNR↑	17.43	11.98	18.44	16.00	17.65	17.87	15.41
SSIM ↑	0.591	0.436	0.666	-	0.634	0.643	0.603
LPIPS↓	0.303	0.748	0.311	0.437	0.390	0.377	0.462

Table 2: **Novel View Synthesis Evaluation on iPhone Dataset [12]**. All results are averaged over the 5 sequences evaluated in [68], with the standard 2x downsampling. All methods except [68] are evaluated without ground truth camera poses. Best results are highlighted as **first**, **second**, **third**. Our method shows superior PSNR and SSIM over the offline Shape of Motion and falls short of the state-of-the-art offline method MoSca [29] by a small margin under the extra constraints of online reconstruction and tracking. When tested with precomputed camera poses, ProDyG outperforms the only online competitor DynOMo [56] by a significant advantage. Notably, our method still works reasonably well with RGB-only input while being online.

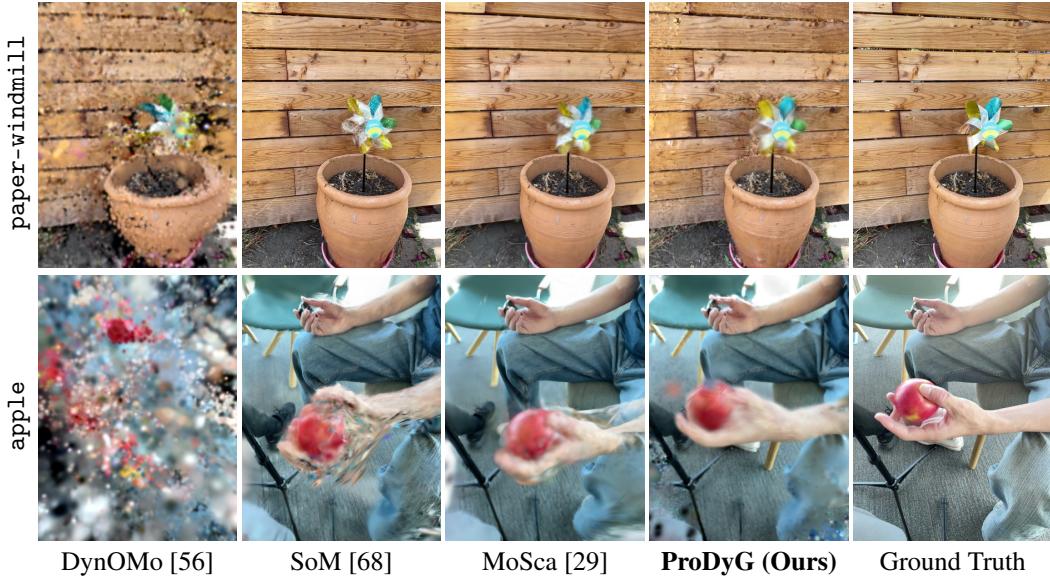


Figure 4: **Qualitative Novel View Synthesis Results on iPhone [12]**. All methods are trained with access to ground truth depth maps. Dynamic objects reconstructed by our method tend to show more accurate silhouettes than [68] and [29]. Note that ProDyG is the only one among the methods to perform both tracking and reconstruction online.

5 Conclusion

We proposed ProDyG, a progressive dynamic 3D reconstruction framework that meets four key criteria for practical deployment: online operation, global pose and map consistency, detailed appearance and geometry modeling through 3D Gaussian Splatting, and flexibility to operate with either RGB or RGB-D input. Our novel flow-based motion-mask prediction integrated with the SLAM backend enables robust camera tracking in dynamic environments, while our online dynamic reconstruction pipeline updates and optimizes Motion Scaffolds and dynamic Gaussians in a progressive manner. Our experiments demonstrate that ProDyG achieves competitive performance in both tracking and novel view synthesis.

Acknowledgments

We thank the Computer Vision Lab at ETH Zurich for providing the GPU compute resources used in this research, with approval from Ender Konukoglu. Erik Sandström conducted this work as part of his employment at Google Zurich, which supported his participation in this research; Google had no direct role in the project's design, execution, or analysis. Shi Chen acknowledges support from the NeurIPS 2025 Scholar Award and thanks Stefan Leutenegger for additional travel support.

References

- [1] Agarwal, P., Tipaldi, G.D., Spinello, L., Stachniss, C., Burgard, W.: Robust map optimization using dynamic covariance scaling. In: 2013 IEEE international conference on robotics and automation. pp. 62–69. Ieee (2013)
- [2] Bae, J., Kim, S., Yun, Y., Lee, H., Bang, G., Uh, Y.: Per-gaussian embedding-based deformation for deformable 3d gaussian splatting. In: European Conference on Computer Vision. pp. 321–335. Springer (2024)
- [3] Bescos, B., Campos, C., Tardós, J.D., Neira, J.: Dynaslam ii: Tightly-coupled multi-object tracking and slam. IEEE robotics and automation letters **6**(3), 5191–5198 (2021)
- [4] Bescos, B., Fácil, J.M., Civera, J., Neira, J.: Dynaslam: Tracking, mapping, and inpainting in dynamic scenes. IEEE robotics and automation letters **3**(4), 4076–4083 (2018)
- [5] Campos, C., Elvira, R., Rodríguez, J.J.G., Montiel, J.M., Tardós, J.D.: Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap slam. IEEE transactions on robotics **37**(6), 1874–1890 (2021)
- [6] Chen, X., Chen, Y., Xiu, Y., Geiger, A., Chen, A.: Easi3r: Estimating disentangled motion from dust3r without training. arXiv preprint arXiv:2503.24391 (2025)
- [7] Dai, W., Zhang, Y., Li, P., Fang, Z., Scherer, S.: Rgb-d slam in dynamic environments using point correlations. IEEE transactions on pattern analysis and machine intelligence **44**(1), 373–389 (2020)
- [8] Dou, M., Khamis, S., Degtyarev, Y., Davidson, P., Fanello, S.R., Kowdle, A., Escolano, S.O., Rhemann, C., Kim, D., Taylor, J., et al.: Fusion4d: Real-time performance capture of challenging scenes. ACM Transactions on Graphics (ToG) **35**(4), 1–13 (2016)
- [9] Duan, Y., Wei, F., Dai, Q., He, Y., Chen, W., Chen, B.: 4d-rotor gaussian splatting: towards efficient novel view synthesis for dynamic scenes. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11 (2024)
- [10] Engel, J., Koltun, V., Cremers, D.: Direct sparse odometry. IEEE transactions on pattern analysis and machine intelligence **40**(3), 611–625 (2017)
- [11] Feng, H., Zhang, J., Wang, Q., Ye, Y., Yu, P., Black, M.J., Darrell, T., Kanazawa, A.: St4rtrack: Simultaneous 4d reconstruction and tracking in the world. arXiv preprint arXiv:2504.13152 (2025)
- [12] Gao, H., Li, R., Tulsiani, S., Russell, B., Kanazawa, A.: Dynamic novel-view synthesis: A reality check. In: NeurIPS (2022)
- [13] Han, J., An, H., Jung, J., Narihira, T., Seo, J., Fukuda, K., Kim, C., Hong, S., Mitsuji, Y., Kim, S.: D²ust3r: Enhancing 3d reconstruction with 4d pointmaps for dynamic scenes. arXiv preprint arXiv:2504.06264 (2025)
- [14] Henein, M., Zhang, J., Mahony, R., Ila, V.: Dynamic slam: The need for speed. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 2123–2129. IEEE (2020)
- [15] Hou, C., Yeo, Q.X., Guo, M., Su, Y., Li, Y., Lee, G.H.: Mvgsr: Multi-view consistency gaussian splatting for robust surface reconstruction. arXiv preprint arXiv:2503.08093 (2025)
- [16] Hu, B., Li, Y., Xie, R., Xu, B., Dong, H., Yao, J., Lee, G.H.: Learnable infinite taylor gaussian for dynamic view rendering. arXiv preprint arXiv:2412.04282 (2024)
- [17] Huang, J., Yang, S., Zhao, Z., Lai, Y.K., Hu, S.M.: Clusterslam: A slam backend for simultaneous rigid body clustering and motion estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5875–5884 (2019)

[18] Jeong, Y., Lee, J., Choi, H., Cho, M.: Rodygs: Robust dynamic gaussian splatting for casual videos. arXiv preprint arXiv:2412.03077 (2024)

[19] Jiang, H., Xu, Y., Li, K., Feng, J., Zhang, L.: Rodyn-slam: Robust dynamic dense rgb-d slam with neural radiance fields. IEEE Robotics and Automation Letters (2024)

[20] Karaev, N., Makarov, I., Wang, J., Neverova, N., Vedaldi, A., Rupprecht, C.: Cotracker3: Simpler and better point tracking by pseudo-labelling real videos. arXiv preprint arXiv:2410.11831 (2024)

[21] Katsumata, K., Vo, D.M., Nakayama, H.: A compact dynamic 3d gaussian representation for real-time dynamic view synthesis. In: European Conference on Computer Vision. pp. 394–412. Springer (2024)

[22] Kavan, L., Collins, S., Žára, J., O’Sullivan, C.: Skinning with dual quaternions. In: Proceedings of the 2007 symposium on Interactive 3D graphics and games. pp. 39–46 (2007)

[23] Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. **42**(4), 139–1 (2023)

[24] Kim, D.H., Kim, J.H.: Effective background model-based rgb-d dense visual odometry in a dynamic environment. IEEE Transactions on Robotics **32**(6), 1565–1573 (2016)

[25] Kong, H., Yang, X., Wang, X.: Efficient gaussian splatting for monocular dynamic scene rendering via sparse time-variant attribute modeling. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 4374–4382 (2025)

[26] Kong, M., Lee, J., Lee, S., Kim, E.: Dgs-slam: Gaussian splatting slam in dynamic environment. arXiv preprint arXiv:2411.10722 (2024)

[27] Kratimenos, A., Lei, J., Daniilidis, K.: Dynmf: Neural motion factorization for real-time dynamic view synthesis with 3d gaussian splatting. In: European Conference on Computer Vision. pp. 252–269. Springer (2024)

[28] Lee, J., Won, C., Jung, H., Bae, I., Jeon, H.G.: Fully explicit dynamic gaussian splatting. Advances in Neural Information Processing Systems **37**, 5384–5409 (2024)

[29] Lei, J., Weng, Y., Harley, A., Guibas, L., Daniilidis, K.: Mosca: Dynamic gaussian fusion from casual videos via 4d motion scaffolds. arXiv preprint arXiv:2405.17421 (2024)

[30] Leroy, V., Cabon, Y., Revaud, J.: Grounding image matching in 3d with mast3r. In: European Conference on Computer Vision. pp. 71–91. Springer (2024)

[31] Li, H., Meng, X., Zuo, X., Liu, Z., Wang, H., Cremers, D.: Pg-slam: Photo-realistic and geometry-aware rgb-d slam in dynamic environments. arXiv preprint arXiv:2411.15800 (2024)

[32] Li, M., Chen, W., Cheng, N., Xu, J., Li, D., Wang, H.: Garad-slam: 3d gaussian splatting for real-time anti dynamic slam. arXiv preprint arXiv:2502.03228 (2025)

[33] Li, M., Zhou, Y., Jiang, G., Deng, T., Wang, Y., Wang, H.: Ddn-slam: Real-time dense dynamic neural implicit slam. arXiv preprint arXiv:2401.01545 (2024)

[34] Li, R.B., Shaghaghi, M., Suzuki, K., Liu, X., Moparthi, V., Du, B., Curtis, W., Renschler, M., Lee, K.M.B., Atanasov, N., et al.: Dynagslam: Real-time gaussian-splatting slam for online rendering, tracking, motion predictions of moving objects in dynamic scenes. arXiv preprint arXiv:2503.11979 (2025)

[35] Li, Y., Fang, Y., Zhu, Z., Li, K., Ding, Y., Tombari, F.: 4d gaussian splatting slam. arXiv preprint arXiv:2503.16710 (2025)

[36] Li, Z., Chen, Z., Li, Z., Xu, Y.: Spacetime gaussian feature splatting for real-time dynamic view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8508–8520 (2024)

[37] Li, Z., Tucker, R., Cole, F., Wang, Q., Jin, L., Ye, V., Kanazawa, A., Holynski, A., Snavely, N.: Megasam: Accurate, fast, and robust structure and motion from casual dynamic videos. arXiv preprint arXiv:2412.04463 (2024)

[38] Liang, H., Ren, J., Mirzaei, A., Torralba, A., Liu, Z., Gilitschenski, I., Fidler, S., Oztireli, C., Ling, H., Gojcic, Z., et al.: Feed-forward bullet-time reconstruction of dynamic scenes from monocular videos. arXiv preprint arXiv:2412.03526 (2024)

[39] Liang, Y., Xu, T., Kikuchi, Y.: *Himor: Monocular deformable gaussian reconstruction with hierarchical motion representation*. arXiv preprint arXiv:2504.06210 (2025)

[40] Liang, Y., Khan, N., Li, Z., Nguyen-Phuoc, T., Lanman, D., Tompkin, J., Xiao, L.: *Gaufre: Gaussian deformation fields for real-time dynamic novel view synthesis*. In: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 2642–2652. IEEE (2025)

[41] Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: *Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis*. In: 2024 International Conference on 3D Vision (3DV). pp. 800–809. IEEE (2024)

[42] Matsuki, H., Murai, R., Kelly, P.H., Davison, A.J.: *Gaussian splatting slam*. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18039–18048 (2024)

[43] Morris, J., Wang, Y., Kliniewski, M., Ila, V.: *Dynosam: Open-source smoothing and mapping framework for dynamic slam*. arXiv preprint arXiv:2501.11893 (2025)

[44] Mur-Artal, R., Tardós, J.D.: *Orb-slam2: An open-source slam system for monocular, stereo, and rgbd cameras*. IEEE transactions on robotics **33**(5), 1255–1262 (2017)

[45] Newcombe, R.A., Fox, D., Seitz, S.M.: *Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time*. In: 2015 Proceedings of the IEEE conference on computer vision and pattern recognition. IEEE (2015)

[46] Oquab, M., Darzet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.Y., Li, S.W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: *Dinov2: Learning robust visual features without supervision* (2024), <https://arxiv.org/abs/2304.07193>

[47] Palazzolo, E., Behley, J., Lottes, P., Giguère, P., Stachniss, C.: *ReFusion: 3D Reconstruction in Dynamic Environments for RGB-D Cameras Exploiting Residuals* (2019), <https://www.ipb.uni-bonn.de/pdfs/palazzolo2019iros.pdf>

[48] Palazzolo, E., Behley, J., Lottes, P., Giguere, P., Stachniss, C.: *Refusion: 3d reconstruction in dynamic environments for rgbd cameras exploiting residuals*. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 7855–7862. IEEE (2019)

[49] Park, J., Bui, M.Q.V., Bello, J.L.G., Moon, J., Oh, J., Kim, M.: *Splinesg: Robust motion-adaptive spline for real-time dynamic 3d gaussians from monocular video*. arXiv preprint arXiv:2412.09982 (2024)

[50] Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al.: *Sam 2: Segment anything in images and videos*. arXiv preprint arXiv:2408.00714 (2024)

[51] Ruan, C., Zang, Q., Zhang, K., Huang, K.: *Dn-slam: A visual slam with orb features and nerf mapping in dynamic environments*. IEEE Sensors Journal **24**(4), 5279–5287 (2023)

[52] Rünz, M., Agapito, L.: *Co-fusion: Real-time segmentation, tracking and fusion of multiple objects*. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 4471–4478. IEEE (2017)

[53] Runz, M., Buffier, M., Agapito, L.: *Maskfusion: Real-time recognition, tracking and reconstruction of multiple moving objects*. In: 2018 IEEE international symposium on mixed and augmented reality (ISMAR). pp. 10–20. IEEE (2018)

[54] Sandström, E., Tateno, K., Oechsle, M., Niemeyer, M., Van Gool, L., Oswald, M.R., Tombari, F.: *Splat-slam: Globally optimized rgbd-only slam with 3d gaussians*. arXiv preprint arXiv:2405.16544 (2024)

[55] Schischka, N., Schieber, H., Karaoglu, M.A., Gorgulu, M., Grötzner, F., Ladikos, A., Navab, N., Roth, D., Busam, B.: *Dynamon: Motion-aware fast and robust camera localization for dynamic neural radiance fields*. IEEE Robotics and Automation Letters (2024)

[56] Seidenschwarz, J., Zhou, Q., Duisterhof, B., Ramanan, D., Leal-Taixé, L.: *Dynomo: Online point tracking by dynamic online monocular gaussian reconstruction*. arXiv preprint arXiv:2409.02104 (2024)

[57] Shen, S., Cai, Y., Wang, W., Scherer, S.: *Dytanvo: Joint refinement of visual odometry and motion segmentation in dynamic environments*. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 4048–4055. IEEE (2023)

[58] Slavcheva, M., Baust, M., Cremers, D., Ilic, S.: *Killingfusion: Non-rigid 3d reconstruction without correspondences*. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1386–1395 (2017)

[59] Soares, J.C.V., Gattass, M., Meggiolaro, M.A.: Crowd-slam: visual slam towards crowded environments using object detection. *Journal of Intelligent & Robotic Systems* **102**(2), 50 (2021)

[60] Somraj, N., Choudhary, K., Mupparaju, S.H., Soundararajan, R.: Factorized motion fields for fast sparse input dynamic view synthesis. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–12 (2024)

[61] Stearns, C., Harley, A., Uy, M., Dubost, F., Tombari, F., Wetzstein, G., Guibas, L.: Dynamic gaussian marbles for novel view synthesis of casual monocular videos. In: SIGGRAPH Asia 2024 Conference Papers. pp. 1–11 (2024)

[62] Sturm, J., Engelhard, N., Endres, F., Burgard, W., Cremers, D.: A benchmark for the evaluation of rgb-d slam systems. In: Proc. of the International Conference on Intelligent Robot Systems (IROS) (Oct 2012)

[63] Sturm, J., Engelhard, N., Endres, F., Burgard, W., Cremers, D.: A benchmark for the evaluation of RGB-D SLAM systems. In: International Conference on Intelligent Robots and Systems (IROS). IEEE/RSJ (2012). <https://doi.org/10.1109/IROS.2012.6385773>, <http://ieeexplore.ieee.org/document/6385773/>

[64] Sun, Y., Liu, M., Meng, M.Q.H.: Improving rgb-d slam in dynamic environments: A motion removal approach. *Robotics and Autonomous Systems* **89**, 110–122 (2017)

[65] Sun, Z., Lo, J., Hu, J.: Embracing dynamics: Dynamics-aware 4d gaussian splatting slam. arXiv preprint arXiv:2504.04844 (2025)

[66] Teed, Z., Deng, J.: Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems* **34**, 16558–16569 (2021)

[67] Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence* **13**(04), 376–380 (1991)

[68] Wang, Q., Ye, V., Gao, H., Austin, J., Li, Z., Kanazawa, A.: Shape of motion: 4d reconstruction from a single video. arXiv preprint arXiv:2407.13764 (2024)

[69] Wang, Q., Zhang, Y., Holynski, A., Efros, A.A., Kanazawa, A.: Continuous 3d perception model with persistent state. arXiv preprint arXiv:2501.12387 (2025)

[70] Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: Dust3r: Geometric 3d vision made easy. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20697–20709 (2024)

[71] Wang, Y., Tian, Y., Chen, J., Xu, K., Ding, X.: A survey of visual slam in dynamic environment: The evolution from geometric to semantic approaches. *IEEE Transactions on Instrumentation and Measurement* (2024)

[72] Wen, L., Li, S., Zhang, Y., Huang, Y., Lin, J., Pan, F., Bing, Z., Knoll, A.: Cassidy: Gaussian splatting slam in dynamic environments. arXiv preprint arXiv:2411.15476 (2024)

[73] Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20310–20320 (2024)

[74] Xiao, L., Wang, J., Qiu, X., Rong, Z., Zou, X.: Dynamic-slam: Semantic monocular visual localization and mapping based on deep learning in dynamic environment. *Robotics and Autonomous Systems* **117**, 1–16 (2019)

[75] Xu, K., Tse, T.H.E., Peng, J., Yao, A.: Das3r: Dynamics-aware gaussian splatting for static scene reconstruction. arXiv preprint arXiv:2412.19584 (2024)

[76] Xu, Y., Jiang, H., Xiao, Z., Feng, J., Zhang, L.: Dg-slam: Robust dynamic gaussian splatting slam with hybrid pose optimization. arXiv preprint arXiv:2411.08373 (2024)

[77] Xu, Z., Niu, J., Li, Q., Ren, T., Chen, C.: Nid-slam: Neural implicit representation-based rgb-d slam in dynamic environments. In: 2024 IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6. IEEE (2024)

[78] Yang, S., Scherer, S.: Cubeslam: Monocular 3-d object slam. *IEEE Transactions on Robotics* **35**(4), 925–938 (2019)

[79] Yang, Z., Yang, H., Pan, Z., Zhang, L.: Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. arXiv preprint arXiv:2310.10642 (2023)

[80] Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20331–20341 (2024)

[81] Yu, C., Liu, Z., Liu, X.J., Xie, F., Yang, Y., Wei, Q., Fei, Q.: Ds-slam: A semantic visual slam towards dynamic environments. In: 2018 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 1168–1174. IEEE (2018)

[82] Yu, X., Ye, W., Guo, X., Ming, Y., Li, J., Bao, H., Cui, Z., Zhang, G.: D³flowslam: Self-supervised dynamic slam with flow motion decomposition and dino guidance. arXiv preprint arXiv:2207.08794 (2022)

[83] Zhang, B., Fang, C., Shrestha, R., Liang, Y., Long, X., Tan, P.: Rade-gs: Rasterizing depth in gaussian splatting. arXiv preprint arXiv:2406.01467 (2024)

[84] Zhang, D., Li, G., Li, J., Bressieux, M., Hilliges, O., Pollefeys, M., Van Gool, L., Wang, X.: Egogaussian: Dynamic scene understanding from egocentric video with 3d gaussian splatting. arXiv preprint arXiv:2406.19811 (2024)

[85] Zhang, J., Henein, M., Mahony, R., Ila, V.: Vdo-slam: A visual dynamic object-aware slam system. arXiv preprint arXiv:2005.11052 (2020)

[86] Zhang, J., Herrmann, C., Hur, J., Jampani, V., Darrell, T., Cole, F., Sun, D., Yang, M.H.: Monst3r: A simple approach for estimating geometry in the presence of motion. arXiv preprint arXiv:2410.03825 (2024)

[87] Zhang, S., Ge, Y., Tian, J., Xu, G., Chen, H., Lv, C., Shen, C.: Pomato: Marrying pointmap matching with temporal motion for dynamic 3d reconstruction. arXiv preprint arXiv:2504.05692 (2025)

[88] Zhang, T., Zhang, H., Li, Y., Nakamura, Y., Zhang, L.: Flowfusion: Dynamic dense rgb-d slam based on optical flow. In: 2020 IEEE international conference on robotics and automation (ICRA). pp. 7322–7328. IEEE (2020)

[89] Zheng, J., Zhu, Z., Bieri, V., Pollefeys, M., Peng, S., Armeni, I.: Wildgs-slam: Monocular gaussian splatting slam in dynamic environments. arXiv preprint arXiv:2504.03886 (2025)

[90] Zhu, Z., Peng, S., Larsson, V., Xu, W., Bao, H., Cui, Z., Oswald, M.R., Pollefeys, M.: Nice-slam: Neural implicit scalable encoding for slam. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12786–12796 (2022)

[91] Zou, Z., Qu, Z., Peng, X., Boominathan, V., Pediredla, A., Chakravarthula, P.: High-speed dynamic 3d imaging with sensor fusion splatting. arXiv preprint arXiv:2502.04630 (2025)

[92] Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Surface splatting. In: Proceedings of the 28th annual conference on Computer graphics and interactive techniques. pp. 371–378 (2001)

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: **[Yes]**

Justification: We cover our contributions in the abstract and as bullet points in the introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss limitations in a separate paragraph towards the end of the experiment section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We do not derive new theoretical results in this paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [No]

Justification: There are details related to hyperparameters that are not covered in the paper, but disclosed in the works which we reference. All information will, however, be made available once the source code is published.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We wished to have provided the code at submission time, but to improve readability of the code, we choose to release this after submission time, but as soon as possible.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.

- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: **[No]**

Justification: We provide all information necessary to understand the results, but as mentioned before, some hyperparamters that are the same as in Splat-SLAM [54] and MoSca [29] are not disclosed.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: **[No]**

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We include the CPU and GPU specs used to benchmark the runtime of our system, but believe that small differences in the setup of that hardware may still result in differences in runtime across devices.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We provide a framework for dynamic dense SLAM whose primary objective is AR/VR applications and household robotics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: We target dense dynamic scene reconstruction with the primary objective of household AR/VR and robotics.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.

- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: **[No]**

Justification: We will not release data or models that have a high risk of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: **[Yes]**

Justification: We credit the authors where we are inspired by their works, and we respect their licenses, but we do not describe their licenses etc. This will be properly done when releasing the code publicly.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: **[No]**

Justification: We do not release assets (code) yet, but hope to do so as soon as possible. Then we will document this.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: **[NA]**

Justification: All data we use is publicly available covering scenes where appropriate licenses and compensation have been taken care of for potential participants.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: **[No]**

Justification: All data we use is publicly available covering scenes where appropriate licenses and compensation have been taken care of for potential participants.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorosity, or originality of the research, declaration is not required.

Answer: **[No]**

Justification: We do not use LLMs for our methodology in any way.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.