4DTAM: Non-Rigid Tracking and Mapping via Dynamic Surface Gaussians

Hidenobu Matsuki Gwangbin Bae Andrew J. Davison

Dyson Robotics Laboratory, Imperial College London

{h.matsuki20, g.bae, a.davison}@imperial.ac.uk

Website: https://muskie82.github.io/4dtam/ Video: https://youtu.be/MRGhggLmTF0?si=51bqfAe9pYQNWgf-/



Figure 1. 4DTAM jointly estimates camera-egomotion, appearance, geometry and scene dynamics without any template.

Abstract

We propose the first 4D tracking and mapping method that jointly performs camera localization and non-rigid surface reconstruction via differentiable rendering. Our approach captures 4D scenes from an online stream of color images with depth measurements or predictions by jointly optimizing scene geometry, appearance, dynamics, and camera ego-motion. Although natural environments exhibit complex non-rigid motions, 4D-SLAM remains relatively underexplored due to its inherent challenges; even with 2.5D signals, the problem is ill-posed because of the high dimensionality of the optimization space. To overcome these challenges, we first introduce a SLAM method based on Gaussian surface primitives that leverages depth signals more effectively than 3D Gaussians, thereby achieving accurate surface reconstruction. To further model non-rigid deformations, we employ a warp-field represented by a multi-layer perceptron (MLP) and introduce a novel camera pose estimation technique along with surface regularization terms that facilitate spatio-temporal reconstruction. In addition to these algorithmic challenges, a significant hurdle in 4D SLAM research is the lack of reliable ground truth and evaluation protocols, primarily due to the difficulty of 4D capture using commodity sensors. To address this, we present a novel open synthetic dataset of everyday objects with diverse motions, leveraging large-scale object models and animation modeling. In summary, we open up the modern 4D-SLAM research by introducing a novel method and evaluation protocols grounded in modern vision and rendering techniques.

1. Introduction

The world we live in has many moving elements. Rivers flow, trees sway, cookies crumble, and humans walk. Although Simultaneous Localization and Mapping (SLAM) methods which assume that most of the world is static are highly useful, embodied agents which aim to navigate and interact with their environments in the most general way should be able to operate in dynamic scenes. There are several ways to segment and ignore moving scene elements, and a SLAM system can be assembled by integrating these individual modules so that it can reconstruct the static parts of a scene and estimate camera ego-motion. However, in this work, we aim for a more comprehensive spatiotemporal (4D) reconstruction of scenes exhibiting significant dynamic motion. Our primary focus is on a unified framework that leverages intrinsic capabilities of the underlying scene representation without heavily relying on prior assumptions about moving elements. 4D-SLAM with general scene motion is difficult primarily because of the complex and high-dimensional nature of modeling non-rigid motions (and potential topological changes) while simultaneously optimizing the pose of a moving camera. There is much more redundancy than in rigid SLAM, and some prior assumptions are needed to combat this. Another challenge lies in the lack of datasets to train and/or evaluate techniques. Recent advances in computer vision and graphics make it a good time to revisit this problem. New 3D representations (e.g. neural fields and Gaussian splats) allow differentiable rendering of complex 3D scenes and optimization via 2D observations, and to model deformation fields smoothly without more specific assumptions. Also, the availability of high-quality 3D meshes on Internet and rendering software (e.g. Blender) gives the ability to render non-rigidly moving objects with ground truth.

We present 4DTAM, a novel approach for 4D Tracking And Mapping in dynamic scenes. We use Gaussian surface primitives to represent the scene and introduce a neural warp-field represented by a multi-layer perception (MLP) to model continuous temporal changes. We then utilize differentiable rendering to jointly optimize the scene geometry, appearance, dynamics, and camera ego-motion from an online stream of a single RGB-D camera. This enables accurate 3D reconstruction and real-time rendering, even in the presence of complex non-rigid deformations. To facilitate future research, we also introduce a new synthetic dataset of dynamic objects. Our focus in this dataset is realistic, complex motion of scenes that are not well represented by existing deformable object models. Animated 3D meshes are rendered and the ground truth depth, surface normals, and foreground masks are extracted together with the camera poses/intrinsics. This dataset provides challenging scenarios for 4D reconstruction methods. We also release the full rendering script to allow the generation of custom 4D datasets. Our experimental results demonstrate that 4DTAM achieves good performance in both camera tracking and scene reconstruction in the presence of dynamic objects. It can handle the complex motion of articulated objects (e.g., drawers) and non-rigid objects (e.g., curtains, flags, and animals), showcasing its potential for applications in robotics, augmented reality, and other fields requiring real-time dynamic scene understanding. We primarily use RGB-D sensor input, but also demonstrate an extension to monocular RGB streams by incorporating a monocular depth prediction network in the supplementary material.

In summary, the contributions of this paper are:

• **4DTAM**, the first 4D tracking and mapping method that uses differentiable rendering and Gaussian surface primitives for dynamic environments.

- The first 2DGS [17]-based SLAM method with analytic camera pose gradients, normal initialization, and regularization to fully exploit depth signals.
- An MLP-based warp-field for modeling non-rigid scene, complemented by a novel camera localization technique and rigidity regularization of surface Gaussians.
- A novel 4D-SLAM dataset with complex object motions, ground-truth camera trajectories, and dynamic object meshes, along with an evaluation protocol.
- Extensive evaluations demonstrating that the method achieves state-of-the-art performance.

2. Related Work

2.1. Visual SLAM

Visual SLAM has been an extensively researched field, with Dense SLAM specifically focusing on capturing detailed scene geometry [35] and semantics [30]. A central aspect of these methods lies in the choice of scene representation and the corresponding optimization framework. Dense SLAM methods based on traditional scene representations, such as volumetric Truncated Signed Distance Functions (TSDF) [22, 34, 60] or Surfels [44, 61], project 2D observations into 3D space and employ specific data fusion algorithms. While effective, these methods often fail to keep consistency between the model and sensor observations across multiple viewpoints, posing challenges for long-term operation.

However, recent advancements in graphics hardware have facilitated the adoption of differentiable rendering frameworks, which have revolutionized inverse rendering and scene reconstruction [23, 31, 33, 37]. Differentiable rendering ensures multi-view consistency through streamlined backpropagation, enhancing scene reconstruction accuracy. Notably, 3D Gaussian Splatting (3DGS) [25] has gained attention due to its flexible resource allocation and rapid forward rendering capabilities. Initially developed for photorealistic view synthesis, recent research has extended its application to surface reconstruction [15, 66]. Enhanced methods, such as 2D Gaussian Splatting (2DGS) [17], achieve superior geometry reconstruction by reducing the Gaussian dimension and explicitly defining surface normals. These differentiable rendering representations have been applied to visual SLAM, from coordinatebased MLPs [52] to explicit voxel grids [21, 55, 63, 67], points [43], and 3D Gaussians [24, 29, 62].

2.2. SLAM for 4D Scene Reconstruction

3D reconstruction of dynamic scenes has been extensively studied, with notable achievements using optimization methods, even for unknown non-rigid objects observed by a single moving RGB camera [14, 53]. However, these approaches typically require batch optimization and are limited to smaller scenes. In contrast, dynamic SLAM targets incremental, reconstruction and tracking of large, continuously moving scenes ideally in real-time. Most methods to date have relied on RGB-D data from moving depth cameras.

While many methods detect and exclude dynamic objects to focus on static scene reconstruction [45], full spatiotemporal reconstruction (which we refer to as *4D-SLAM*) requires more advanced solutions. For instance, tracking and reconstructing rigid moving objects separately [42] or employing parametric shape models for known semantic classes like humans or animals [26] are effective strategies. Specialized domains, such as endoscopic imaging, have utilized scene-specific priors or deformation models to handle non-rigid dynamics [28, 41].

An incremental 4D-SLAM for general dynamic scenes has remained more challenging, but has been addressed based on various regularizing assumptions and representations. DynamicFusion [36] pioneered a line of work [13, 19, 47, 48] which captures temporal evolution in the scene geometry by jointly optimizing a canonical volumetric representation (e.g., TSDF volume [36]) and a deformation field. As the solution space is extremely high-dimensional, additional constraints are often introduced to regularize the motion field [47, 48] or to align visual features [4, 19]. Recent advances in 3D representations, such as neural fields and Gaussian primitives, have opened new possibilities for dynamic scene reconstruction. Canonical radiance and motion fields can be jointly optimized via differentiable rendering, as demonstrated with NeRF [38, 40, 54] and SDF [7, 56]. For 3D Gaussians, which can explicitly represent points, motion can be estimated either through perprimitive trajectories [27] or learnable motion bases [57]. However, warp-field-based motion representation offers inherent smoothness regularization, leveraging the properties of neural fields [10, 18, 64, 65]. Most existing methods, however, rely on known camera poses or multi-camera setups to capture dense spatiotemporal observations. While DyNoMo [46] supports camera pose optimization, its 3D Gaussian representation is not suited for geometrically accurate reconstruction. In contrast, our 4DTAM framework enables 4D reconstruction using a single RGB-D camera, jointly optimizing camera poses, appearance, geometry, and dynamics, making it practical for most embodied agents.

2.3. Datasets for 4D Reconstruction

4D reconstruction has been studied extensively for the case of the human body. Datasets like Human3.6M [20], Deep-Cap [16], and ZJU-MoCap [39] capture diverse human motions under a multi-camera setup. The cameras are fixed, synchronized, and calibrated to reduce the difficulty in establishing dense multi-view correspondences. Only a small number of datasets provide single-stream RGB- D sequences captured from a moving camera [5, 12, 47]. Recovering the camera poses is not trivial for such realworld captures, and additional post-processing (e.g. robust depth map alignment [56]) is required. Another challenge lies in ground truth acquisition. Besides the depth measurements, other ground truths (e.g., scene flow, object mask) often require manual labeling. On the contrary, synthetic datasets [6, 59] provide perfect ground truths. Recent advances in open-source datasets [9] and rendering software [8] also close the synthetic-to-real domain gap significantly. To this end, we introduce a new high-quality synthetic dataset tailored for 4D reconstruction and camera pose estimation.

3. Method

3.1. 2D Gaussian Splatting

Our geometric scene representation is based on 2D Gaussian Splatting (2DGS) [17]. Unlike 3D Gaussian Splatting (3DGS), which uses blob-like splats, 2DGS functions as a stretchable surfel with explicitly defined surface normal directions. This property makes 2DGS particularly well-suited for non-rigid scene reconstruction with a single camera, where effectively handling 2.5D input signals is critical.

Each 2D Gaussian \mathcal{G} is represented by its 3D mean position \mathbf{P}_{μ} , rotation $\mathbf{R} \in SO(3)$, color **c**, opacity *o*, and a scaling vector $\mathbf{S} \in \mathbb{R}^2$. The rotation matrix **R** is decomposed as $\mathbf{R} = [\mathbf{t}_u, \mathbf{t}_v, \mathbf{t}_w]$, where \mathbf{t}_u and \mathbf{t}_v represent two principal tangential vectors, and \mathbf{t}_w is the normal vector, defined as $\mathbf{t}_w = \mathbf{t}_u \times \mathbf{t}_v$. For simplicity, spherical harmonics are omitted in this work.

The 2D Gaussian function is parameterized on the local tangent plane in world space as:

$$P(u,v) = \mathbf{P}_{\mu} + s_u \mathbf{t}_u u + s_v \mathbf{t}_v v = \mathbf{H}(u,v,1,1)^{\mathrm{T}} \quad (1)$$

where
$$\mathbf{H} = \begin{bmatrix} s_u \mathbf{t}_u & s_v \mathbf{t}_v & \mathbf{0} & \mathbf{p}_k \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \mathbf{RS} & \mathbf{p}_k \\ \mathbf{0} & 1 \end{bmatrix}$$
 (2)

For a point $\mathbf{u} = (u, v)$ in the tangential plane of 2D Gaussian (uv space), its projection onto the image plane is given by

$$\mathbf{x} = (xz, yz, z, 1)^{\mathrm{T}} = \mathbf{WP}(u, v) = \mathbf{WH}(u, v, 1, 1)^{\mathrm{T}}$$
(3)

where $\mathbf{W} \in \mathbb{R}^{4 \times 4}$ is the transformation matrix from world space to screen space.

To avoid numerically unstable matrix inversion of $\mathbf{M} = (\mathbf{W}\mathbf{H})^{-1}$, 2DGS applies ray-splat intersection by finding the intersection of non-parallel planes (x-plane and yplane). The ray $\mathbf{x} = (x, y)$ is determined by the intersection of the x-plane \mathbf{h}_x and the y-plane \mathbf{h}_y , represented as $\mathbf{h}_x = (-1, 0, 0, x)^{\mathrm{T}}$ and $\mathbf{h}_y = (0, -1, 0, y)^{\mathrm{T}}$, respectively. In the uv coordinates of the 2D Gaussian, this is expressed



Figure 2. Method overview of 4DTAM.

as:

$$\mathbf{h}_u = (\mathbf{W}\mathbf{H})^{\mathrm{T}}\mathbf{h}_x$$
 and $\mathbf{h}_v = (\mathbf{W}\mathbf{H})^{\mathrm{T}}\mathbf{h}_y$ (4)

The intersection point meets the following condition,

$$\mathbf{h}_{u} \cdot (u, v, 1, 1)^{\mathrm{T}} = \mathbf{h}_{v} \cdot (u, v, 1, 1)^{\mathrm{T}} = 0$$
 (5)

This leads to an solution for the intersection point $\mathbf{u}(\mathbf{x})$:

$$u(\mathbf{x}) = \frac{\mathbf{h}_{u}^{2}\mathbf{h}_{v}^{4} - \mathbf{h}_{u}^{4}\mathbf{h}_{v}^{2}}{\mathbf{h}_{u}^{1}\mathbf{h}_{v}^{2} - \mathbf{h}_{u}^{2}\mathbf{h}_{v}^{1}} \qquad v(\mathbf{x}) = \frac{\mathbf{h}_{u}^{4}\mathbf{h}_{v}^{1} - \mathbf{h}_{u}^{1}\mathbf{h}_{v}^{4}}{\mathbf{h}_{u}^{1}\mathbf{h}_{v}^{2} - \mathbf{h}_{u}^{2}\mathbf{h}_{v}^{1}} \tag{6}$$

where \mathbf{h}_{u}^{i} , \mathbf{h}_{v}^{i} are the *i*-th parameter of the 4D homogeneous plane parameters.

The 2D Gaussian at (u, v) is evaluated as:

$$\mathcal{G}(\mathbf{u}) = \exp\left(-\frac{u^2 + v^2}{2}\right) \tag{7}$$

The 2D Gaussians are sorted along the camera ray by their center depth and organized into image tiles. Per-pixel color is rendered via volumetric alpha blending:

$$c(\mathbf{x}) = \sum_{i=1} \mathbf{c}_i \alpha_i \mathcal{G}_i(\mathbf{u}(x)) \prod_{j=1}^{i-1} (1 - \alpha_j \mathcal{G}_j(\mathbf{u}(x)))$$
(8)

where depth and normal can be rendered similarly.

3.2. Analytic Camera Pose Jacobian

One major advantage of Gaussian Splatting is its analytical formulation of gradient flow for model parameters, enabling real-time full-resolution rendering. However, it assumes posed images as input and does not provide gradients for camera poses. To accelerate optimization, we derive the analytic Jacobian of the camera pose for 2D Gaussian Splatting and implement it using a CUDA kernel. This formulation has potential applications for a wide range of tasks involving pose estimation in surface-based Gaussian Splatting.

We use Lie algebra to derive the minimal Jacobians for the camera pose matrix from the world coordinate system to the camera's local coordinate system, defining $T_{CW} \in$ SE(3) and $\tau \in \mathfrak{se}(3)$. Since 2DGS backpropagates gradients to $\mathbf{M}^T = \mathbf{W}\mathbf{H}$ during the optimization of the 3D mean, we require the partial derivative $\frac{\partial \mathbf{M}^T}{\partial \tau}$. Let $\mathbf{K} \in \mathbb{R}^{4\times 4}$ represent the camera projection matrix. Then, equation 3 is rewritten as:

$$\mathbf{x} = \mathbf{M}^T(u, v, 1, 1)^{\mathrm{T}} = \mathbf{K} \boldsymbol{T}_{CW} \mathbf{H}(u, v, 1, 1)^{\mathrm{T}}$$
(9)

Using the chain rule, the partial derivatives are computed as:

$$\frac{\partial \mathbf{M}^T}{\partial \boldsymbol{\tau}} = \frac{\partial \mathbf{M}^T}{\partial \mathbf{W}} \frac{\partial \mathbf{W}}{\partial \boldsymbol{T}_{CW}} \frac{\partial \boldsymbol{T}_{CW}}{\partial \boldsymbol{\tau}}, \qquad (10)$$

$$\frac{\partial \boldsymbol{T}_{CW}}{\partial \boldsymbol{\tau}} = \begin{bmatrix} \boldsymbol{0} & -\boldsymbol{R}_{CW}_{;1}^{\times} \\ \boldsymbol{0} & -\boldsymbol{R}_{CW}_{;2}^{\times} \\ \boldsymbol{0} & -\boldsymbol{R}_{CW}_{;3}^{\times} \\ \boldsymbol{I} & -\boldsymbol{t}_{CW}^{\times} \end{bmatrix}$$
(11)

where $\mathbf{R}_{CW} \in \mathbf{SO}(3)$ and $\mathbf{t}_{CW} \in \mathbb{R}^3$ denote the rotation and translation parts of \mathbf{T}_{CW} , respectively. The notation \times represents the skew-symmetric matrix of a 3D vector, and $\mathbf{R}_{CW,i}$ denotes the *i*th column of \mathbf{R}_{CW} .

2DGS also renders a normal map, which can be supervised using the loss computed from the rendered normals. Let n_c denote the camera-space normal. The normal of a 2D Gaussian in the camera's local coordinate system is defined as:

$$\mathbf{n}_c = \boldsymbol{T}_{CW} \mathbf{t}_w \tag{12}$$

where \mathbf{t}_w is the surface normal in the world coordinate system.

Borrowing the notation of the left Jacobian for Lie groups from [49], the partial derivative is given by:

$$\frac{\partial \mathbf{n}_c}{\partial \tau} = \frac{\mathcal{D} \mathbf{n}_c}{\mathcal{D} \boldsymbol{T}_{CW}} = \begin{bmatrix} \boldsymbol{I} & -\mathbf{n}_c^{\times} \end{bmatrix}$$
(13)

Further details of the derivation are provided in the supplementary material.

3.3. Warp Field

To model time-varying deformations, we use a warp-field represented by a coordinate-based network [64, 65]. In our hand-held single-camera setup, the limited view coverage of dynamic objects necessitates structural priors in the motion representation. For this, we employ a compact MLP as the warp-field to estimate transitions from the canonical Gaussians following [64].

Given time t and center position x of 2D Gaussians in canonical space as inputs, the deformation MLP f_{θ} produces offsets, which subsequently transform the canonical 2D Gaussians to the deformed space:

$$(\delta \boldsymbol{x}, \delta \boldsymbol{r}, \delta \boldsymbol{s}) = \mathbf{f}_{\theta}(\gamma_1(\boldsymbol{x}), \gamma_2(t))$$
(14)

where $\delta x \in \mathbb{R}^3$, $\delta r \in SO(3)$, $\delta s \in \mathbb{R}^2$ denotes the offsets of 2D Gaussian's mean position, rotation and scale respectively, γ denotes the frequency-based positional encoding [31]. For deformable SLAM applications, we leverage a CUDA-optimized MLP implementation [33] to enable fast, interactive reconstruction.

3.4. Tracking and Mapping Framework

Our SLAM method follows the standard tracking and mapping architecture, where the tracking module is in charge of fast online camera pose estimation while the mapping performs a relatively more involved joint opimitization of the camera poses, geometry and motion of selected keyframes. Further details of the hyperparameters are available in the supplementary material.

3.4.1. Tracking

The tracking module estimates the coarse camera pose for the latest incoming frame. This is achieved by minimizing the photometric and depth rendering errors between the sensor observation and the rendering from the deformable Gaussian model. Unlike static 3DGS SLAM methods, we estimate the camera pose relative to the warped Gaussians at the latest keyframes timestamp t_{kf} , assuming the deformed scene structure at t_{kf} is closest to the current state. We define photometric rendering loss as:

$$L_p = \left\| I(\mathcal{G}_{cano}, \boldsymbol{T}_{CW}, t_{kf}) - \bar{I} \right\|_1$$
(15)



Figure 3. 2D Gaussian's Surface Normal Rendering based on Different Initialization. Left: Random initialization. Right: Our initialization aligned with sensor measurement.

Here $I(\mathcal{G}, \mathbf{T}_{CW})$ denotes a rendered color image from the canonical Gaussians \mathcal{G}_{cano} , timestamp of the latest keyframe t_{kf} and camera pose \mathbf{T}_{CW} , and \bar{I} is an observed image. Similarly, we also minimize geometric depth error:

$$L_g = \left\| D(\mathcal{G}_{cano}, \boldsymbol{T}_{CW}, t_{kf}) - D \right\|_1 \tag{16}$$

Following MonoGS [29], we further optimize affine brightness parameters. Keyframes are selected every N-th frame and sent to the mapping process for further refinement.

3.4.2. Mapping

The mapping module performs joint optimization of the camera pose, canonical Gaussians, and the warp field within a sliding window.

Gaussian Management When a new keyframe is registered, we add new Gaussians to the canonical Gaussians \mathcal{G}_{cano} , based on the back-projected point cloud from the RGB-D observations. Unlike 3DGS, 2DGS explicitly encodes surface normal information in its rotation vector, making it beneficial to initialize using surface normals estimated from sensor depth measurements. To achieve this, we compute the surface normals of the current depth observation by taking the finite difference of neighboring backprojected depth points and assign them as the normal vectors of the 2D Gaussianss \mathbf{t}_w . This is formulated as:

$$\mathbf{t}_{w} = \frac{\nabla_{x} \mathbf{p}_{d} \times \nabla_{y} \mathbf{p}_{d}}{|\nabla_{x} \mathbf{p}_{d} \times \nabla_{y} \mathbf{p}_{d}|}$$
(17)

where \mathbf{p}_d denotes points back-projected by the current sensor depth observation. We store the computed normal information as a 2D image \mathbf{N}_{sensor} for normal supervision. Pruning and densification parameters follow MonoGS, which effectively prunes the wrongly inserted Gaussians in the canonical space due to the object movement.

4D Map optimization We perform joint optimization of the camera ego-motion, appearance, geometry and scene dynamics. In a single-camera setup, the lack of spatiotemporally dense observations makes fully capturing dynamic

scenes challenging, as complete spatial (xyz) coverage over time (t) is only feasible with multi-camera systems. To address this, we introduce regularization terms for both shape and motion.

In addition to photometric and depth losses, we apply a normal regularization based on sensor measurements to better align 2D Gaussians. Unlike the original 2DGS methods, which compute normals by finite differences of rendered depth during every optimization step—leading to high computational costs—we instead propose to use normals precomputed from depth input as supervision. This reduces computational overhead, as normals are calculated only when a new keyframe is inserted:

$$L_n = \sum_{i \in h \times w} (1 - \mathbf{n}_i^{\mathrm{T}} \mathbf{N}_{sensor,i})$$
(18)

To constrain motion in unobserved regions, we apply an as-rigid-as possible regularization loss L_{ARAP} from [27] to the Gaussian means. Additionally, we introduce a novel surface normal rigidity loss, constraining the 2D Gaussians' surface normals to stay similar between timesteps t_1 and t_2 , preserving local surface rigidity:

$$L_{ARAP_{n}} = w_{i,j} \left\| (\mathbf{t}_{w})_{i,t_{1}}^{T} (\mathbf{t}_{w})_{j,t_{1}} - (\mathbf{t}_{w})_{i,t_{2}}^{T} (\mathbf{t}_{w})_{j,t_{2}} \right\|_{1}$$
(19)

where $w_{i,j}$ is a distance-based weighting factor like L_{ARAP} . We apply ARAP regularizers between the oldest and latest keyframe in the current window.

Together with the isotropic loss L_{iso} proposed in [29], we minimize the following total cost function:

$$L_{total} = \lambda_p L_p + \lambda_g L_g + \lambda_n L_n + \lambda_{iso} L_{iso} + L_{ARAP} + L_{ARAP.n}$$
(20)

The optimization is based on the sliding window heuristics in [11], with two additional keyframes randomly selected from the history.

Global Optimization Sliding window-based optimization prioritizes the latest frame, causing past keyframe information to degrade over time. After tracking, if required we can perform global optimization to finalize the map, which takes less than 1 minute on an RTX 4090. During this step, the poses and number of Gaussians are fixed, and one keyframe is randomly selected per iteration. The process uses the normal consistency loss of 2DGS, ensuring global consistency despite being relatively slow.

3.5. Dataset Generation

We introduce **Sim4D**, a new synthetic dataset for 4D reconstruction. Recently, a large number of photo-realistic, animated 3D meshes have become available [2, 9]. Combined with open-source graphics software [3, 8], such meshes provide a scalable way of generating datasets for non-rigid 4D reconstruction. The data generation pipeline is illustrated in Fig. 4.

Meshes and background. We collected 50 high-quality, animated 3D meshes from Objaverse [9] and Sketchfab [2], all of which are under CC-BY license. The collected meshes exhibit a wide variety of motions, including nonrigid deformation and topological changes. We then place the object inside a cube and randomize the background texture. Texture maps are collected from Poly Haven [1] and are all under CC0 license.

Rendering. We render 240 to 540 frames for each object. The camera trajectories are defined along arcs of 20 degrees, and test viewpoints are defined outside of these arcs to evaluate the performance of novel-view synthesis and to quantify the accuracy of the reconstructed geometry. At each timestamp, the RGB image, ground truth depth, surface normals, and foreground mask are rendered and the camera intrinsics/extrinsics saved. Please refer to the supplementary material for additional details.

4. Evaluation

4.1. Experimental Setup

We extensively evaluate our non-rigid SLAM method on both synthetic and real-world datasets. Previous non-rigid RGB-D SLAM work has primarily focused on qualitative demonstrations using limited datasets, showcasing the early-stage potential of the field. To advance research, we introduce a quantitative evaluation protocol with the new Sim4D dataset. Our evaluation covers camera pose accuracy, as well as the appearance and geometric quality of the reconstructed models. Additionally, we demonstrate realworld performance using a self-captured dataset.

While designed primarily for dynamic scenes, our method is the first to leverage surface Gaussian splatting for both static SLAM and non-rigid RGB-D reconstruction. To further validate our approach, we perform a detailed quantitative component-wise ablation analysis.

Metrics and Datasets For our main Non-Rigid SLAM evaluation, we evaluate our method on 8 sequences from the Sim4D dataset. We first report ATE RMSE for trajectory evaluation. To assess SLAM map quality, we report depth rendering error (L1 error) for geometry and PSNR, SSIM, and LPIPS for appearance evaluation. For Sim4D, metrics are calculated from test views (extrapolated positions across different timestamps). The estimated and ground truth trajectories are aligned on the first frame, and test view positions are queried in the ground truth trajectory's coordinate system. Details about the test viewpoints are in the supplementary material. Since SurfelWarp [13] requires explicit foreground segmentation, we collect its results only on pixels with valid reconstruction. For Static



Figure 4. **Sim4D dataset.** We create a new dataset for 4D reconstruction by rendering animated 3D meshes.

SLAM ablation, we report ATE RMSE, rendering performance, and TSDF-fused mesh metrics, following the protocol in [43]. We evaluate our method on the Replica [50] dataset and the TUM RGB-D dataset [51]. To isolate the impact of scene representation from system differences, we replaced MonoGS's representation with 2DGS while keeping all other system configurations identical. For Offline Non-Rigid Reconstruction ablation, we report the average geometry and appearance rendering metrics on subsets of the DeepDeform [5], KillingFusion [47], and iPhone datasets [12], which are used in [56]. Numerical quantities for each sequence is available in supplementary material. Since [56] primarily focuses on object shape completion, metrics are calculated only within the given segmentation mask. The camera pose is provided by the dataset, and pose optimization is disabled to focus solely on reconstruction performance. We perform 30000 iteration for training, which takes approximately 30 mins.

Baseline Methods For quantitative non-rigid SLAM evaluation, we compare our method with SurfelWarp [13], the only non-rigid RGB-D SLAM method with publicly available code. For component-wise ablation analysis, we compare against MonoGS [29] for static SLAM evaluation and Morpheus [56] for offline reconstruction.

Implementation Details Our SLAM system runs on a desktop equipped with an Intel Core i9-12900K (3.50GHz) processor and a single NVIDIA GeForce RTX 4090 GPU. The camera pose jacobian for 2DGS, described in Section 3.2, is implemented using a CUDA rasterizer, similar to other gradients in Gaussian Splatting. For real-world data capture, we used the Realsense D455.

4.2. Quantitative Evaluation

Table 1 compares our method with SurfelWarp [13]. Our method outperforms SurfelWarp across all metrics. To analyze this further, Fig. 5 provides qualitative visualizations and trajectory plots for the modular_vehicle sequence.



Figure 5. **Qualitative comparison to SurfelWarp.** Left: Rendered image, Middle: Rendered normal map, Right: Estimated camera trajectory

Since SurfelWarp relies on a foreground mask, its reconstruction lacks scene completeness. In contrast, our method reconstructs the entire scene within a joint optimization framework, providing more comprehensive coverage. Additionally, compared to SurfelWarp's back-projection and Surfel fusion scheme, our differentiable rendering-based optimization enforces multi-view consistency over time, resulting in superior camera tracking and consistent 3D reconstruction. Our method achieves camera pose estimation at approximately 1.5 fps and completes the final global optimization in 1 minute.

4.3. Qualitative Evaluation

Fig. 6 presents qualitative reconstruction results on realworld dynamic scenes. Our method successfully reconstructs dynamic scenes with non-rigid deformations, whereas MonoGS fails to handle such complexities.

4.4. Ablation Study

Static SLAM Table 2 provides the camera ATE and 3D reconstruction evaluation results. Our 2DGS-based implementation shows competitive performance and achieves the best result in 6 out of 8 sequences for camera ATE, and consistently better result on rendering and 3D reconstruction metrics. The reconstruction is visualized in Fig 7 which

Method	Category	Metric	curtain	flag	mercedes	modular_vehicle	rhino	shoe_rack	water_effect	wave_toy
	Trajectory	ATE RMSE[cm]↓	6.10	31.9	5.21	4.21	2.81	2.16	2.60	1.45
	Geometry	L1 Depth[cm]↓	49.1	50.8	5.2	10.3	44.9	4.25	46.7	47.5
Method SurfelWarp [13] Ours		PSNR [dB] ↑	15.78	11.04	25.7	19.45	16.76	26.3	17.3	16.4
	Appearance	SSIM ↑	0.468	0.343	0.779	0.362	0.188	0.795	0.325	0.364
		LPIPS ↓	0.56	0.659	0.483	0.638	0.665	0.397	0.587	0.555
Method SurfelWarp [13] Ours	Trajectory	ATE RMSE[cm]↓	0.25	1.00	0.18	0.31	0.25	0.18	0.29	0.32
	Geometry	L1 Depth[cm]↓	0.96	3.58	0.62	1.44	1.85	0.99	3.20	3.43
Ours		PSNR [dB] ↑	28.01	21.01	32.13	30.59	24.13	31.7	27.12	27.10
SurfelWarp [13]	Appearance	SSIM ↑	0.787	0.601	0.894	0.801	0.742	0.901	0.795	0.794
		LPIPS ↓	0.096	0.150	0.138	0.210	0.260	0.12	0.908	0.097

Table 1. Non-rigid SLAM Evaluation on Sim4D Dataset.



Figure 6. **Qualitative Results on Real-World Datset.** Our method effectively handles dynamic objects compared to MonoGS.



Figure 7. **3D Reconstruction Result on Replica Office4.** Left: MonoGS. Right: Ours (MonoGS-2D). Our surface Gaussianbased approach yields more accurate geometric reconstructions.

	Metric	r0	r1	r2	00	01	o2	03	o4	avg
	ATE RMSE[cm]↓	0.44	0.32	0.31	0.44	0.52	0.23	0.17	2.25	0.59
	Depth L1[cm]↓	3.00	3.47	4.66	3.10	6.08	6.15	4.77	4.94	4.52
MonoGS	Precision[%]↑	39.0	28.8	28.9	39.7	15.8	28.0	32.5	25.5	29.7
	Recall[%]↑	44.2	34.5	32.8	47.6	24.3	30.0	35.4	28.5	34.6
	Metric r0 r1 r2 00 o1 o2 o3 ATE RMSE[cm]↓ 0.44 0.32 0.31 0.44 0.52 0.23 0.1 Depth L1[cm]↓ 3.00 3.47 4.66 3.10 6.08 6.15 4.7 Precision[%]↑ 3.00 28.8 28.9 39.7 15.8 28.0 32. Recall[%]↑ 44.2 34.5 32.8 47.6 24.3 30.0 3.5 F1[%]↑ 41.5 31.4 30.7 43.3 19.1 29.0 3.5 F1[%]↑ 0.42 0.45 0.28 0.57 0.37 0.59 0.85 0.6 MonoGS Precision[%]↑ 0.42 0.43 0.37 0.59 0.85 0.6 MonoGS-2D Precision[%]↑ 97.0 97.0 97.0 97.1 97.9 95.8 94. Recall[%]↑ 90.9 91.3 90.5 93.1 91.1 88.2 87. </td <td>33.9</td> <td>26.9</td> <td>31.9</td>	33.9	26.9	31.9						
	ATE RMSE[cm]↓	0.42	0.43	0.35	0.19	0.19	0.22	0.27	0.80	0.36
	Depth L1[cm]↓	0.45	0.28	0.57	0.37	0.59	0.85	0.62	0.63	0.54
MonoGS-2D	Precision[%]↑	97.0	97.0	97.0	97.1	97.9	95.8	94.8	83.9	95.0
	Recall[%]↑	85.5	86.0	84.8	89.4	85.1	81.8	81.5	72.4	83.3
	F1[%]↑	90.9	91.3	90.5	93.1	91.1	88.2	87.6	77.7	88.8

Table 2. Static SLAM Ablation on Replica.

Method	Metric	fr1/desk	fr2/xyz	fr3/office	avg.
Method MonoGS MonoGS-2D	ATE RMSE[cm]↓	1.50	1.44	1.49	1.47
Monous	Depth L1[cm]↓	6.2	13.0	fr3/office 1.49 13.0 1.83 4.30	10.7
ManaGS 2D	MetheInfidesATE RMSE[cm] \downarrow 1.50Depth L1[cm] \downarrow 6.2ATE RMSE[cm] \downarrow 1.58Depth L1[cm] \downarrow 3.00	1.58	1.20	1.83	1.57
W011003-2D	Depth L1[cm]↓	3.00	2.30	z fr3/office 1.49 13.0 1.83 4.30	3.2

Table 3. Static SLAM Ablation on TUM

shows the comparison of the mesh generated by TSDF Fusion between MonoGS and MonoGS-2D. Table 3 provides the camera ATE and rendering metrics evaluation on TUM

Method	Metric	KillingFusion	DeepDeform	iPhone
$\begin{array}{c c} Method & Metric \\ & Depth L1[cm] \downarrow \\ Morpheus [56] & PSNR [dB] \uparrow \\ SSIM \uparrow \\ LPIPS \downarrow \\ Depth L1[cm] \downarrow \\ PSNR [dB] \uparrow \\ SSIM \uparrow \end{array}$	3.2	1.9	2.4	
	PSNR [dB] ↑	27.02	26.81	25.28
	SSIM \uparrow	0.77	0.81	0.46
	$\begin{tabular}{ c c c c c c } \hline Metric & KillingFusion & DeepDeform & i \\ \hline Depth L1[cm] \downarrow & 3.2 & 1.9 \\ \hline PSNR [dB] \uparrow & 27.02 & 26.81 \\ \hline SSIM \uparrow & 0.77 & 0.81 \\ \hline LPIPS \downarrow & 0.40 & 0.38 \\ \hline Depth L1[cm] \downarrow & 4.9 & 1.1 \\ \hline PSNR [dB] \uparrow & 31.13 & 24.15 \\ \hline SSIM \uparrow & 0.93 & 0.90 \\ \hline LPIPS \downarrow & 0.13 & 0.27 \\ \hline \end{tabular}$	0.63		
	Depth L1[cm]↓	4.9	1.1	0.57
Ours	PSNR [dB] ↑	31.13	24.15	27.54
Ours	SSIM ↑	0.93	0.90	0.79
$\begin{array}{c c c c c c c c c c c c c c c c c c c $	0.26			

Table 4.	Offline	Non-Rigid	Reconstruction	Ablation:	Render-
ing Error	Metrics	on Real-wo	orld Dataset.		



Figure 8. **Non-rigid Reconstruction Results.** Our method flexibly models non-rigid deformations without requiring any shape templates or foreground/background separation.

dataset. Our method shows on par camera ATE but shows the increased geometric reconstruction quality.

Offline Non-rigid RGB-D Surface Reconstruction Table 4 reports offline reconstruction results, where camera poses are given. Our 2DGS+MLP deformation model shows competitive rendering performance compared to NeRF based methods. Note that Gaussian Splatting has the additional advantage of its rendering speed. We further provide qualitative visualizations in Fig. 8.

5. Conclusion

We presented the first tracking and mapping method for non-rigid surface reconstruction using Surface Gaussian Splatting. Our approach integrates a 2DGS + MLP warpfield SLAM framework with camera pose estimation and regularization, leveraging RGB-D input. To support further research, we also introduced a novel dataset for dynamic scene reconstruction with reliable ground truth. Experimental results demonstrate that our method outperforms traditional non-rigid SLAM approaches.

Limitations: Our method has primarily been tested on small-scale scenes; extending it to complex real-world scenarios may require 2D priors like point tracking or optical flow. The current implementation runs at 1.5 fps, limiting real-time use. Developing interactive dynamic scene scanning remains important future work.

6. Acknowledgement

Research presented in this paper has been supported by Dyson Technology Ltd. We are very grateful to members of the Dyson Robotics Lab for their advice and insightful discussions.

4DTAM: Non-Rigid Tracking and Mapping via Dynamic Surface Gaussians

Supplementary Material

We encourage readers to watch the supplementary video for additional details and qualitative results.

7. Implementation Details

7.1. System Details and Hyper parameters

Non-Rigid SLAM: We set the learning weights as follows: $\lambda_p = 0.9$, $\lambda_g = 0.1$, $\lambda_{iso} = 10.0$ and $\lambda_n = 0.002$. For the ARAP regularization [27], we use a nearest neighbor count of 20, a radius of 0.05, and an exponential decay weight of 500. Keyframes are selected with N = 1. For the MLP, we use an 8-layer architecture with 256 neurons per layer. Frequency encoding is set to 1 for time and 4 for position. MLP is implemented with CUDA-optimized CutlassMLP in tiny-cuda-nn [32] for the fast optimization.

Static SLAM Ablation: We followed the same hyperparameters as MonoGS [29], but we use normal loss L_n with the weight $\lambda_n = 0.01$ for the entire mapping process and $\lambda_g = 0.5$ for the final refinement. For the Replica 3D reconstruction evaluation, we have used the script introduced in [43].

Offline Non-rigid RGB-D Reconstruction Ablation: Camera poses are provided by the dataset and remain fixed during training. For the MLP, we adopt the same architecture described in [64], consisting of an 8-layer network with 256 dimensions per layer, where a concatenated feature vector is input to the fourth layer. The positional encoding frequencies are set to 6 for time and 10 for position. Following the approach in [7, 56], we evaluate the geometric and appearance metrics against the input views and report the average values.

8. Camera Pose Jacobian

We provide the detail of the derivation of camera pose jacobian of 2D Gaussian Splatting in 3.2.

We use the notation from [49]. Let $T \in SE(3)$ and $\tau = (\rho, \theta) \in \mathfrak{se}(3)$, the left-side partial derivative on the manifold is defined as:

$$\frac{\mathcal{D}f(\boldsymbol{T})}{\mathcal{D}\boldsymbol{T}} \triangleq \lim_{\boldsymbol{\tau} \to 0} \frac{\operatorname{Log}(f(\operatorname{Exp}(\boldsymbol{\tau}) \circ \boldsymbol{T}) \circ f(\boldsymbol{T})^{-1})}{\boldsymbol{\tau}} \quad (21)$$

Eq 11:

$$T = \operatorname{Exp}(\tau) = \exp(\tau^{\wedge})$$
$$= \exp\left(\sum_{j=1}^{6} \mathbf{E}_{j}\tau_{j}\right), \quad j = 1, \dots, 6, \quad \tau \in \mathbb{R}^{6}.$$
(22)

where the matrices $\mathbf{E}_j \in \mathbb{R}^{4 \times 4}$ are the SE(3) group generators and form a basis for $\mathfrak{se}(3)$:

Since the meaningful elements of the camera T is 12 number variables, we stack the elements for 12×6 matrix and we obtain

$$\frac{\partial \boldsymbol{T}}{\partial \boldsymbol{\tau}}\Big|_{\boldsymbol{\tau}=0} = \begin{bmatrix} \mathbf{0} & -\mathbf{R}_{:,1}^{\times} \\ \mathbf{0} & -\mathbf{R}_{:,2}^{\times} \\ \mathbf{0} & -\mathbf{R}_{:,3}^{\times} \\ \boldsymbol{I} & -\mathbf{t}^{\times} \end{bmatrix}.$$
 (26)

(25)

where $\mathbf{R} \in SO(3)$ and $\mathbf{t} \in \mathbb{R}^3$ denote the rotation and translation parts of T.

Eq 13:

$$\left. \frac{\partial \mathbf{n}_c}{\partial \tau} \right|_{\tau=0} = \frac{\mathcal{D} \mathbf{n}_c}{\mathcal{D} \mathbf{T}_{CW}} = \lim_{\tau \to 0} \frac{\mathrm{Exp}(\tau) \mathbf{n}_c - \mathbf{n}_c}{\tau}$$
(27)

$$=\lim_{\tau\to 0}\frac{(\boldsymbol{I}+\boldsymbol{\tau}^{\wedge})\cdot\mathbf{n}_{c}-\mathbf{n}_{c}}{\tau}$$
 (28)

$$=\lim_{\tau\to 0}\frac{\tau\wedge\mathbf{n}_c}{\tau} \tag{29}$$

$$= \lim_{\tau \to 0} \frac{\boldsymbol{\theta}^* \mathbf{n}_c + \boldsymbol{\rho}}{\tau}$$
(30)

$$= \lim_{\tau \to 0} \frac{-\mathbf{n}_c^{\times} \boldsymbol{\theta} + \boldsymbol{\rho}}{\tau}$$
(31)

$$= \begin{bmatrix} \mathbf{I} & -\mathbf{n}_c^{\times} \end{bmatrix}$$
(32)

9. Sim4D Training/Test Views

We define the training and test views on a sphere, with its center representing the target object. In spherical coordinates (r, θ, ϕ) , we set r = 2.0. The training view is sampled from two arcs on the sphere's surface, defined by $\theta \in [-10^{\circ}, 10^{\circ}]$ and $\phi \in [-10^{\circ}, 10^{\circ}]$. The test views are sampled from a circle on the sphere's surface that pass through four key points: $(\theta, \phi) = (5^{\circ}, 0^{\circ}), (0^{\circ}, 5^{\circ}), (-5^{\circ}, 0^{\circ}), and <math>(0^{\circ}, -5^{\circ})$. These points are chosen to ensure uniform sampling around the target object while maintaining a clear separation between the training and test views.



Figure 9. **Training and Test Views on the Sim4D Dataset: Blue** indicates training views, and **Red** indicates test views. Views are sampled (top right) from an arc on an object-centered sphere (top left) for dynamic scene reconstruction (bottom).

10. Further Ablation Analysis

10.1. Normal Rigidity Loss

Table 5 presents the quantitative results demonstrating the effect of the normal rigidity loss defined in Equation 19. The normal rigidity loss improves the overall geometric metrics, such as camera ATE and L1 Depth, for the benchmark sequences by preserving the local geometric consistency of 2D Gaussians.

	ATE RMSE	L1 Depth	PSNR	SSIM	LPIPS
Ours full	0.28	1.71	28.47	0.820	0.12
w/o L _{ARAP_n}	0.52	2.00	29.04	0.853	0.13

Table 5. Ablation Study on $L_{ARAP.n}$. We report the average number of Sim4D dataset.

10.2. Monocular Depth Prior

While our method was primarily tested with RGB-D camera input, we conducted an ablation study using depth input from the state-of-the-art monocular prediction network [58], as shown in Table 9. The results demonstrate performance competitive with SurfelWarp, highlighting the potential for purely monocular non-rigid SLAM.

10.3. Static SLAM Ablation Analysis

Replica: Table 8 shows the photometric rendering performance analysis on the Replica dataset. The results demonstrate that the 2DGS-based SLAM approach offers an advantage in achieving accurate appearance reconstruction.

TUM: Table 6 presents the full ablation analysis on the TUM dataset. The 2DGS-based approach maintains competitive ATE and appearance metrics while achieving significantly better geometric rendering accuracy, as reflected in the Depth L1 error.

Method	Metric	fr1	fr2	fr3
MonoGS	ATE RMSE [cm]↓	1.50	1.44	1.49
MOHOUS	Depth L1 [cm] ↓	6.2	13.0	13.0
	PSNR [dB] ↑	23.5	24.65	25.09
	SSIM ↑	0.775	0.785	0.842
	LPIPS \downarrow	0.26 1	0.201	0.200
ManaCE 2D	ATE RMSE [cm]↓	1.58	1.2	1.83
M01003-2D	Depth L1 [cm] ↓	3.0	2.3	4.3
	PSNR [dB] ↑	23.63	24.47	24.05
	SSIM ↑	0.782	0.79	0.826
	LPIPS \downarrow	0.251	0.228	0.223

Table 6. **Static SLAM Ablation on TUM Dataset.** Comparison of ATE RMSE, Depth L1, and Rendering Performance Metrics.

Memory Analysis Table 7 presents the average memory usage on the TUM dataset sequences. Due to the geometrically accurate alignment, 2D Gaussians require fewer primitives to represent the scene, resulting in reduced memory consumption.

Memory Usage [MB]							
MonoGS-2D	MonoGS						
2.73MB	3.97MB						

Table 7. Memory Analysis on TUM RGB-D dataset.

10.4. Offline Non-Rigid RGB-D Reconstruction Ablation

Table 10 provides the full evaluation details of the offline non-rigid RGB-D reconstruction ablation analysis.

	Metric	room0	room1	room2	office0	office1	office2	office3	avg
	PSNR [dB] ↑	34.83	36.43	37.49	39.95	42.09	36.24	36.70	37.50
MonoGS	SSIM \uparrow	0.954	0.959	0.9665	0.971	0.977	0.964	0.963	0.96
	LPIPS \downarrow	0.068	0.076	0.075	0.072	0.055	0.078	0.065	0.07
	PSNR [dB]↑	36.21	37.81	38.7	43.45	43.8	37.48	37.43	39.14
MonoGS-2D	SSIM \uparrow	0.966	0.969	0.9737	0.985	0.984	0.972	0.971	0.975
MonoGS MonoGS-2D	LPIPS \downarrow	0.04	0.042	0.044	0.025	0.029	0.04	0.039	0.038

Table 8. Static SLAM Ablation: Rendering Performance Metrics [43] on Replica Dataset

Method	Category	Metric	curtain	flag	mercedes	modular_vehicle	rhino	shoe_rack	water_effect	wave_toy
Ours (Monocular)	Trajectory	ATE RMSE[cm]↓	6.23	16.29	4.90	1.86	3.17	8.02	5.52	7.21
	Geometry	L1 Depth[cm]↓	74.2	155	59.2	38.0	37.7	89.8	72.4	80.8
		PSNR [dB] ↑	17.73	16.22	20.72	26.28	21.48	17.49	18.86	17.98
	Appearance	SSIM ↑	0.461	0.455	0.636	0.578	0253	0.448	0.390	0.441
		LPIPS ↓	0.297	0.517	0.282	0.380	0.339	0.391	0.258	0.281

Table 9. Non-rigid SLAM Evaluation on Sim4D Dataset with Monocular Depth Prior.

		K	illingFus	sion	D	DeepDefor	m	iPhone		
		frog	duck	snoopy	seq002	seq004	seq028	teddy	mochi	haru
Morpheus [56]	Depth L1 [cm]	4.37	3.01	2.30	2.08	1.24	2.26	5.40	0.31	1.63
	PSNR [dB] ↑	27.2	28.17	25.73	27.21	26.94	26.30	23.40	28.12	24.34
	SSIM ↑	0.802	0.716	0.779	0.809	0.823	0.795	0.237	0.623	0.510
	LPIPS ↓	0.31	0.419	0.483	0.301	0.428	0.397	0.776	iPhone mochi 0.31 28.12 0.623 0.55 0.22 36.15 0.926 0.131	0.564
	Depth L1 [cm]	0.65	1.91	12.1	0.78	1.07	1.30	0.32	0.22	0.12
Ours	PSNR [dB] ↑	33.72	32.75	26.95	24.36	24.13	24.02	23.89	36.15	22.60
Ours	SSIM ↑	0.941	0.949	0.899	0.897	0.897	0.902	0.739	0.926	0.690
	$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	0.391								

Table 10. Offline RGB-D Reconstruction Results

References

- [1] Poly haven. https://polyhaven.com/textures/ fabric. Accessed: 2024-11-01. 6
- [2] Sketchfab. https://sketchfab.com/. Accessed: 2024-11-01.6
- [3] Oliver Boyne. Blendersynth. https://ollieboyne. github.io/BlenderSynth, 2023. 6
- [4] Aljaž Božič, Pablo Palafox, Michael Zollhöfer, Justus Thies, Angela Dai, and Matthias Nießner. Neural deformation graphs for globally-consistent non-rigid reconstruction. arXiv preprint arXiv:2012.01451, 2020. 3
- [5] Aljaž Božič, Michael Zollhöfer, Christian Theobalt, and Matthias Nießner. Deepdeform: Learning non-rigid rgb-d reconstruction with semi-supervised data. 2020. 3, 7
- [6] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012. 3
- [7] Hongrui Cai, Wanquan Feng, Xuetao Feng, Yan Wang, and Juyong Zhang. Neural surface reconstruction of dynamic scenes with monocular rgb-d camera. In *Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS)*, 2022. 3, 1
- [8] Blender Online Community. Blender a 3d modelling and rendering package, 2018. 3, 6
- [9] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. 3, 6
- [10] Bardienus P Duisterhof, Zhao Mandi, Yunchao Yao, Jia-Wei Liu, Jenny Seidenschwarz, Mike Zheng Shou, Ramanan Deva, Shuran Song, Stan Birchfield, Bowen Wen, and Jeffrey Ichnowski. DeformGS: Scene flow in highly deformable scenes for deformable object manipulation. WAFR, 2024. 3
- [11] J. Engel, V. Koltun, and D. Cremers. Direct sparse odometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2017. 6
- [12] Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. In *NeurIPS*, 2022. 3, 7
- [13] Wei Gao and Russ Tedrake. Surfelwarp: Efficient nonvolumetric single view dynamic reconstruction. In Proceedings of Robotics: Science and Systems (RSS), 2018. 3, 6, 7, 8
- [14] R. Garg, A. Roussos, and L. Agapito. Dense variational reconstruction of non-rigid surfaces from monocular video. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013. 2
- [15] Antoine Guédon and Vincent Lepetit. Sugar: Surfacealigned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. 2024. 2
- [16] Marc Habermann, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular human performance capture using weak supervision. In *Pro-*

ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5052–5063, 2020. 3

- [17] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *Proceedings of SIGGRAPH*, 2024. 2, 3
- [18] Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [19] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. VolumeDeform: Real-time Volumetric Non-rigid Reconstruction. In Proceedings of the European Conference on Computer Vision (ECCV), 2016. 3
- [20] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 3
- [21] M. M. Johari, C. Carta, and F. Fleuret. ESLAM: Efficient dense slam system based on hybrid representation of signed distance fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [22] Olaf Kähler, Victor Adrian Prisacariu, and David W. Murray. Real-time large-scale dense 3d reconstruction with loop closure. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 2
- [23] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3D mesh renderer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3907–3916, 2018. 2
- [24] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track map 3d gaussians for dense rgb-d slam. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024.
- [25] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics (TOG), 2023. 2
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIGGRAPH Asia), 34(6):248:1–248:16, 2015. 3
- [27] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *3DV*, 2024. 3, 6, 1
- [28] Ruibin Ma, Rui Wang, Yubo Zhang, Stephen Pizer, Sarah K McGill, Julian Rosenman, and Jan-Michael Frahm. Rnnslam: Reconstructing the 3d colon to visualize missing regions during a colonoscopy. *Medical image analysis*, 72: 102100, 2021. 3
- [29] Hidenobu Matsuki, Riku Murai, Paul H. J. Kelly, and Andrew J. Davison. Gaussian Splatting SLAM. 2024. 2, 5, 6, 7, 1

- [30] J. McCormac, A. Handa, A. J. Davison, and S. Leutenegger. SemanticFusion: Dense 3D semantic mapping with convolutional neural networks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017. 2
- [31] 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 Proceedings of the European Conference on Computer Vision (ECCV), 2020. 2, 5
- [32] Thomas Müller. tiny-cuda-nn, 2021. 1
- [33] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics (TOG), 2022. 2, 5
- [34] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon. KinectFusion: Real-Time Dense Surface Mapping and Tracking. In *Proceedings of the International Symposium on Mixed and Augmented Reality (ISMAR)*, 2011. 2
- [35] R. A. Newcombe, S. Lovegrove, and A. J. Davison. DTAM: Dense Tracking and Mapping in Real-Time. In *Proceedings* of the International Conference on Computer Vision (ICCV), 2011. 2
- [36] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2015. 3
- [37] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2
- [38] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. 3
- [39] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9054–9063, 2021. 3
- [40] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. 3
- [41] Juan J. Gomez Rodriguez, J. M. M Montiel, and Juan D. Tardos. Nr-slam: Non-rigid monocular slam. *IEEE Transactions on Robotics (T-RO)*, 2023. 3
- [42] Martin Rünz and Lourdes Agapito. Co-fusion: Realtime segmentation, tracking and fusion of multiple objects. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2017. 3
- [43] Erik Sandström, Yue Li, Luc Van Gool, and Martin R. Oswald. Point-slam: Dense neural point cloud-based slam. In Proceedings of the International Conference on Computer Vision (ICCV), 2023. 2, 7, 1, 3

- [44] Thomas Schöps, Torsten Sattler, and Marc Pollefeys. Bad slam: Bundle adjusted direct rgb-d slam. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019. 2
- [45] Raluca Scona, Mariano Jaimez, Yvan R Petillot, Maurice Fallon, and Daniel Cremers. StaticFusion: Background reconstruction for dense rgb-d slam in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2018. 3
- [46] Jenny Seidenschwarz, Qunjie Zhou, Bardienus Duisterhof, Deva Ramanan, and Laura Leal-Taixé. Dynomo: Online point tracking by dynamic online monocular gaussian reconstruction, 2024. 3
- [47] Miroslava Slavcheva, Maximilian Baust, Daniel Cremers, and Slobodan Ilic. Killingfusion: Non-rigid 3d reconstruction without correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2017. 3, 7
- [48] Miroslava Slavcheva, Maximilian Baust, and Slobodan Ilic. Sobolevfusion: 3d reconstruction of scenes undergoing free non-rigid motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [49] J. Solà, J. Deray, and D. Atchuthan. A micro Lie theory for state estimation in robotics. arXiv:1812.01537, 2018. 5, 1
- [50] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. arXiv preprint arXiv:1906.05797, 2019. 7
- [51] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. A Benchmark for the Evaluation of RGB-D SLAM Systems. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, 2012. 7
- [52] E. Sucar, S. Liu, J. Ortiz, and A. J. Davison. iMAP: Implicit mapping and positioning in real-time. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [53] L. Torresani, A. Hertzmann, and C. Chris Bregler. Nonrigid structure-from-motion: Estimating shape and motion with hierarchical priors. *IEEE Transactions on Pattern Analysis* and Machine Intelligence (PAMI), 30(5), 2008. 2
- [54] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Nonrigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. 2021. 3
- [55] Hengyi Wang, Jingwen Wang, and Lourdes Agapito. Coslam: Joint coordinate and sparse parametric encodings for neural real-time slam. *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2023.
- [56] Hengyi Wang, Jingwen Wang, and Lourdes Agapito. Morpheus: Neural dynamic 360deg surface reconstruction from

monocular rgb-d video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20965–20976, 2024. 3, 7, 8, 1

- [57] Qianqian Wang, Vickie Ye, Hang Gao, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4d reconstruction from a single video. 2024. 3
- [58] Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang. Moge: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision, 2024. 2
- [59] Wenshan Wang, Delong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. Tartanair: A dataset to push the limits of visual slam. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 4909–4916. IEEE, 2020. 3
- [60] T. Whelan, M. Kaess, H. Johannsson, M. F. Fallon, J. J. Leonard, and J. B. McDonald. Real-time large scale dense RGB-D SLAM with volumetric fusion. *International Journal of Robotics Research (IJRR)*, 34(4-5):598–626, 2015. 2
- [61] T. Whelan, S. Leutenegger, R. F. Salas-Moreno, B. Glocker, and A. J. Davison. ElasticFusion: Dense SLAM without a pose graph. In *Proceedings of Robotics: Science and Systems (RSS)*, 2015. 2
- [62] Chi Yan, Delin Qu, Dan Xu, Bin Zhao, Zhigang Wang, Dong Wang, and Xuelong Li. Gs-slam: Dense visual slam with 3d gaussian splatting. In CVPR, 2024. 2
- [63] Xingrui Yang, Hai Li, Hongjia Zhai, Yuhang Ming, Yuqian Liu, and Guofeng Zhang. Vox-fusion: Dense tracking and mapping with voxel-based neural implicit representation. In Proceedings of the International Symposium on Mixed and Augmented Reality (ISMAR), 2022. 2
- [64] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for highfidelity monocular dynamic scene reconstruction. 2024. 3, 5, 1
- [65] Zeyu Yang, Hongye Yang, Zijie Pan, Xiatian Zhu, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. *Proceedings* of the International Conference on Learning Representations (ICLR), 2024. 3, 5
- [66] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. ACM Transactions on Graphics (TOG), 2024. 2
- [67] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R. Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2022. 2