

GFlow: Recovering 4D World from Monocular Video

Shizun Wang, Xingyi Yang, Qiuhong Shen, Zhenxiang Jiang, Xinchao Wang*

National University of Singapore
shizun.wang@u.nus.edu, xinchao@nus.edu.sg

Abstract

Recovering 4D world from monocular video is a crucial yet challenging task. Conventional methods usually rely on the assumptions of multi-view videos, known camera parameters, or static scenes. In this paper, we relax all these constraints and tackle a highly ambitious but practical task: With only one monocular video without camera parameters, we aim to recover the dynamic 3D world alongside the camera poses. To solve this, we introduce **GFlow**, a new framework that utilizes only 2D priors (depth and optical flow) to lift a video to a 4D scene, as a flow of 3D Gaussians through space and time. GFlow starts by segmenting the video into still and moving parts, then alternates between optimizing camera poses and the dynamics of the 3D Gaussian points. This method ensures consistency among adjacent points and smooth transitions between frames. Since dynamic scenes always continually introduce new visual content, we present prior-driven initialization and pixel-wise densification strategy for Gaussian points to integrate new content. By combining all those techniques, GFlow transcends the boundaries of 4D recovery from causal videos; it naturally enables tracking of points and segmentation of moving objects across frames. Additionally, GFlow estimates the camera poses for each frame, enabling novel view synthesis by changing camera pose. This capability facilitates extensive scene-level or object-level editing, highlighting GFlow’s versatility and effectiveness.

Website — <https://littlepure2333.github.io/GFlow>

Extended version — <https://arxiv.org/abs/2405.18426>

1 Introduction

The quest for accurate reconstruction of 4D scene from video inputs stands at the forefront of contemporary research in computer vision and graphics. This endeavor is crucial for the advancement of virtual and augmented reality, video analysis, and multimedia applications. The main challenge lies in capturing the transient essence of dynamic scenes and the often absent camera pose information.

Traditional approaches are typically split between two types: the one relies on pre-calibrated camera parameters or multi-view video inputs to reconstruct dynamic scenes (Wu

et al. 2023; Luiten et al. 2023; Sun et al. 2024; Bansal et al. 2020; Cao and Johnson 2023; Fridovich-Keil et al. 2023; Li et al. 2022; Lin et al. 2023, 2021b; Pumarola et al. 2021), while the other estimates camera poses from static scenes using multi-view stereo techniques (Bian et al. 2023; Fu et al. 2023; Wang et al. 2023; Lin et al. 2021a; Wang et al. 2021; Xia et al. 2022; Schönberger et al. 2016; Schonberger and Frahm 2016; Tian, Du, and Duan 2023; Charatan et al. 2023). This division highlights a missing piece in this field: the challenge of reconstructing dynamic scenes using only a single monocular video without any camera parameters.

Addressing this challenge is particularly difficult due to the inherently *ill-posed* nature. From a single monocular video, multiple reconstructions might visually appear correct when projected onto the camera view. However, many of these reconstructions fail to conform to the physical laws of the real world. Although NeRF-based (Cao and Johnson 2023; Fridovich-Keil et al. 2023; Shao et al. 2023; Liu et al. 2023) methods attempt to solve this problem, they often yield poor results. This failure is primarily due to their implicit representation, which makes it challenging to accurately enforce physical constraints in the reconstruction.

Recent developments in 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) and its extensions (Wu et al. 2023; Luiten et al. 2023; Yang et al. 2023a,b) into dynamic scenes have emerged as promising alternatives. These techniques have shown promise in handling the complexities associated with the dynamic nature of real-world scenes, as well as the intricacies of camera movement and positioning. Yet, they still operate under the assumption of known camera poses (Schönberger et al. 2016; Schonberger and Frahm 2016).

To transcend these limitations and fully leverage the capabilities of 2D foundation models for dynamic scene reconstruction, we offer a novel insight:

Given 2D factors such as RGB, depth and optical flow from one video, we have enough clues to model the 4D (3D+t) world behind the video.

Leveraging this insight, we introduce **GFlow**, a novel framework that leverages 3D Gaussian Splatting (Kerbl et al. 2023) to reconstruct the video. It conceptualizes the video content as a dynamic flow of Gaussian points moving through space and time. We simultaneously optimize the flow and the camera poses together, to ensure that the projected video adheres to those 2D factors.

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

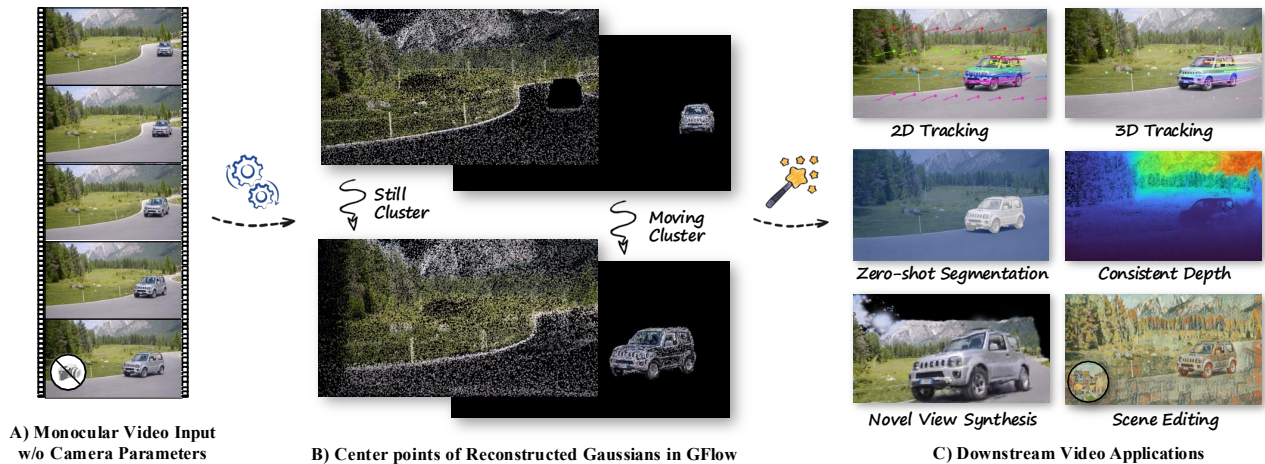


Figure 1: A) Given a monocular video in the wild, B) our proposed **GFlow** can reconstruct the underlying 4D world, *i.e.* the dynamic scene represented by 3D Gaussian splatting (Kerbl et al. 2023) and associated camera poses. Within GFlow, the Gaussians are split into still and moving clusters and are further densified. C) GFlow facilitates a range of applications, including tracking objects in 2D and 3D, segmenting video objects, synthesizing new views, estimating consistent depth and video editing. We encourage readers to visit the website for more video illustrations.

The key to GFlow lies in the alternating optimization of camera poses and dynamic 3D Gaussians. While directly estimating camera poses in dynamic scenes is considered highly challenging, we make it feasible by separating the scene into static and dynamic parts. For the static parts, we optimize the camera poses using reprojection error. In dynamic regions, Gaussian points are first reprojected using the optimized camera poses, then refined based on RGB, depth, and flow priors. This dual optimization ensures that each video frame is rendered accurately, capturing the dynamic nature of the original scene.

Apart from the optimization strategy, we propose two methods to effectively integrate new Gaussian points into the scene and accelerate convergence. The first method, *prior-driven initialization*, sets up initial Gaussian points in plausible 3D geometric positions, based on RGB and depth priors. The second method, *pixel-wise densification*, involves increasing the number of Gaussian points in regions with large pixel errors. Together, these strategies contribute to maintaining high fidelity in cross-frame rendering, also ensuring that transitions and movements between frames are smooth.

Beyond dynamic 3D scene recovery, GFlow can also serve as a powerful tool for video processing. It can track points across frames in 3D world coordinates without prior training and segment objects by propagating a given initial mask. Since it employs explicit representation, GFlow can render captivating new views of video scenes by easily changing camera poses and editing objects or entire scenes as desired, showcasing its versatility and power.

To conclude, our contributions are: **1)** A novel framework that recovers 4D scenes and associated camera poses from a monocular video. **2)** An alternating optimization process that ensures high fidelity and temporally smooth dynamics in 4D scenes. **3)** Two new strategies for initializing and den-

sifying Gaussian points in dynamic scenes. **4)** Enables new video processing capabilities, including tracking, segmentation, novel view rendering, and editing.

2 Related Works

3D Renderable Representations Static 3D scenes can be recovered as renderable representations from posed multi-view images through differentiable rendering, enabling novel view synthesis. Such 3D renderable representations can be categorized into implicit and explicit representations. Early works in 3D scene reconstruction primarily adopted implicit neural representations (Flynn et al. 2016). The most influential of these, Neural Radiance Fields (NeRFs) (Mildenhall et al. 2021), introduced importance sampling with volumetric ray-marching but relied on a deep multi-layer perceptron, significantly hindering rendering speed. Although follow-up works (Müller et al. 2022; Chen et al. 2022) adopted hash grids or structured tensors with smaller MLPs to represent density and appearance, their rendering speed is still constrained by the need to query substantial samples for single ray marching.

In contrast, the explicit category is dominated by differentiable point-based rendering techniques (Yifan et al. 2019; Kerbl et al. 2023). This approach eliminates the need to query samples from deep networks, instead directly fetching attributes from points, which enables a significant speedup compared to implicit neural-based methods. Recently, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) extends points to 3D Gaussians with opacity and spherical harmonics, and introduces tile-based rasterization to achieve real-time rendering speeds. Here, we choose 3DGS as our base representation, as its fast rendering speed and the explicit nature make the reconstructed scene flexible enough for content creation and editing.

4D Reconstruction 4D reconstruction from video, also known as dynamic 3D scene reconstruction. Many prior works extended NeRFs to handle dynamic scenes (Park et al. 2021a,b; Pumarola et al. 2021; Li et al. 2023), typically using grids, triplanes, voxels (Cao and Johnson 2023; Fridovich-Keil et al. 2023; Shao et al. 2023; Liu et al. 2023), or learning deformable fields to map a canonical template (Ouyang et al. 2023; Kasten et al. 2021). But the reconstruction quality is relatively low due to its implicit essence. Recent developments in 3DGS (Kerbl et al. 2023) have set new records in reconstruction quality and rendering speed. Extensions of 3DGS (Wu et al. 2023; Luiten et al. 2023; Yang et al. 2023a,b) have begun exploring dynamic scene reconstruction. However, they still operate under the assumption of a known camera sequence (Schönberger et al. 2016; Schönberger and Frahm 2016).

While almost all previous methods either rely on known camera parameters or multi-view video inputs to reconstruct dynamic scenes (Sun et al. 2024; Bansal et al. 2020; Li et al. 2022; Lin et al. 2023, 2021b), or estimate camera poses from static scenes using multi-view stereo techniques (Bian et al. 2023; Fu et al. 2023; Wang et al. 2023; Lin et al. 2021a; Wang et al. 2021; Xia et al. 2022; Tian, Du, and Duan 2023; Charatan et al. 2023). The key difference between our GFlow and these approaches lies in our ability to recover dynamic scenes from a unposed monocular video. Additionally, some concurrent works (Wang et al. 2024; Stearns et al. 2024; Lei et al. 2024; Liu et al. 2024; Kong, Yang, and Wang 2025) also attempt to solve this problem.

3 Preliminaries

3.1 3D Gaussian Splatting

Recently, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) exhibits strong performance and efficiency in 3D representation. 3DGS fits a scene as a set of Gaussians $\{G_i\}$ from multi-view images $\{V_k\}$ and paired camera poses $\{P_k\}$ in an optimization pipeline. Adaptive densification and pruning of Gaussians are applied in this iterative optimization to control the total number of Gaussians. Generally, each Gaussian is composed of its center coordinate $\mu \in \mathbb{R}^3$, 3D scale $s \in \mathbb{R}^3$, opacity $\alpha \in \mathbb{R}$, rotation quaternion $q \in \mathbb{R}^4$, and associated view-dependent color represented as spherical harmonics $c \in \mathbb{R}^{3(d+1)^2}$, where d is the degree of spherical harmonics.

These parameters can be collectively denoted by G , with $G_i = \{\mu_i, s_i, \alpha_i, q_i, c_i\}$ denoting the parameters of the i -th Gaussian. The core of 3DGS is its tile-based differentiable rasterization pipeline to achieve real-time optimization and rendering. To render $\{G_i\}$ into a 2D image, each Gaussian is first projected into the camera coordinate frame given the camera pose P_i to determine the depth of each Gaussian. Then colors, depth, or other attributes in pixel space are rendered in parallel by alpha composition with the depth order of adjacent 3D Gaussians. Specifically, in our formulation, we do not consider view-dependent color variations for simplicity, thus the degree of spherical harmonics is set as $d = 0$, i.e., only the RGB color $c \in \mathbb{R}^3$.

3.2 Camera Model

To project the 3D point coordinates $\mu \in \mathbb{R}^3$ into the camera view, we use the pinhole camera model. The camera intrinsics is $K \in \mathbb{R}^{3 \times 3}$ and the camera extrinsics which define the world-to-camera transformation is $E = [R|t] \in \mathbb{R}^{3 \times 4}$. The camera-view 2D coordinates $x \in \mathbb{R}^2$ are calculated as $dh(x) = KEh(\mu)$, where $d \in \mathbb{R}$ is the depth, and $h(\cdot)$ represents the homogeneous coordinate mapping.

4 Methodology

Problem Definition We aim to address a highly challenging and ill-posed problem, which is commonly encountered in real-world scenarios though: Given a sequence of monocular video frames without known camera parameters, the objective is to model the dynamic 3D world and the associated camera poses to represent the video.

Overview To address this problem, we propose **GFlow**, a framework that represents videos through a flow of 3D Gaussians, as shown in Figure 2. We first preprocess the videos to derive several priors using advanced foundation models. The priors include depth (Leroy, Cabon, and Revaud 2024), optical flow (Xu et al. 2023), and camera intrinsics (Leroy, Cabon, and Revaud 2024), which we believe are the minimum necessary. These priors contribute to good initialization and regularization in the GFlow optimization process. Two novel strategies are devised to effectively deal with the Gaussian points initialization and densification in the dynamic scenes (Sec. 4.1). At the essence of proposed method, GFlow alternately optimizes the camera pose and Gaussian points for each frame in sequential order to reconstruct the 4D world, assisted by movement clustering of Gaussian points (Sec. 4.2).

4.1 Gaussian Points Allocation

This section introduces new strategies for initializing and densifying Gaussian points according to the video content.

Prior-driven Initialization of Gaussians The original 3D Gaussian Splatting (Kerbl et al. 2023) initializes Gaussian points using point clouds derived from Structure-from-Motion (SfM) (Schonberger and Frahm 2016; Schönberger et al. 2016), which are only viable for static scenes with dense views. However, our task involves dynamic scenes that change both spatially and temporally, making SfM infeasible.

To address this, we developed a new method called **prior-driven initialization** for single frames. This method fully utilizes the *texture* information and *depth* estimation obtained from the image to initialize the Gaussian points.

Intuitively, image areas with more edges usually indicate more complex textures, so more Gaussian points should be initialized in these areas. Given an image $I \in \mathbb{R}^{H \times W}$, we extract its texture map $T \in \mathbb{R}^{H \times W}$ the Sobel operator (Kanopoulos, Vasanthavada, and Baker 1988), an edge detection operator. We then normalize this texture map to create a probability map $P \in \mathbb{R}^{H \times W}$, from which we sample N points to obtain their 2D coordinates $\{x_i\}_{i=1}^N$.

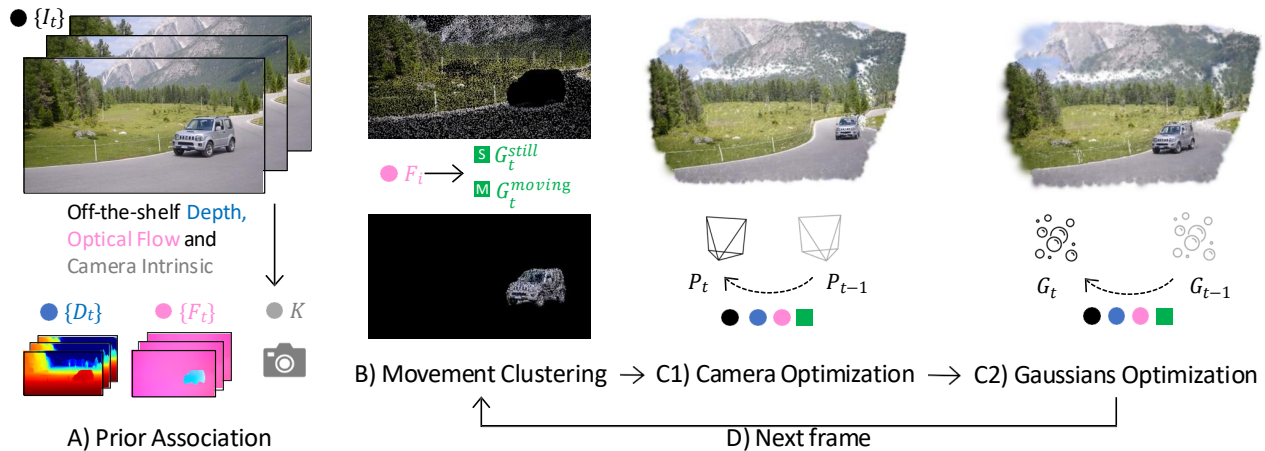


Figure 2: **Overview of GFlow.** **A)** Given a monocular video input consisting of image sequence $\{I_t\}$, the associated depth $\{D_t\}$, optical flow $\{F_t\}$ and camera intrinsic K are obtained using off-the-shelf prior. **B)** For each frame, GFlow first clustering the scene into still part $\{G_t^s\}$ and moving part $\{G_t^m\}$. Then optimization process in GFlow consists of two steps: **C1)** Only the camera pose P_t is optimized by aligning the appearance, depth and optical flow within the still cluster. **C2)** Under the optimized camera pose P_t^* , the Gaussian points $\{G_t\}$ are optimized and densified based on appearance, depth, optical flow and the two scene clusters. **D)** The same procedure of steps B, C1, and C2 loops for the next frame. The colorful marks under the dashed line represent the variables involved in the optimization.

To obtain their position in the 3D space, we use depth D estimated from off-the-shelf model (Leroy, Cabon, and Revaud 2024), as it can offer strong geometric information. The depth $\{d_i\}_{i=1}^N$ of sampled points can be retrieved from depth map D using 2D coordinates. The 3D center coordinate $\{\mu_i^s\}_{i=1}^N$ of Gaussian points is initialized by unprojecting depth $\{d_i\}_{i=1}^N$ and camera-view 2D coordinates $\{x_i\}_{i=1}^N$, according to the pinhole model. The scale $\{s_i\}_{i=1}^N$ and color $\{c_i\}_{i=1}^N$ of the Gaussian points are initialized based on the probability values and pixel colors retrieved from the image, respectively.

Pixel-wise Densification of Gaussians 3D Gaussian Splatting (Kerbl et al. 2023), designed for static scenes, uses gradient thresholding to densify Gaussian points; points exceeding a gradient threshold are cloned or split based on their scales. However, this method struggles in dynamic scenes, particularly when camera movements reveal new scene areas where no Gaussian points exist.

To address this, we introduce a new **pixel-wise densification** strategy that leverages image content, specifically targeting areas yet to be fully reconstructed. Our approach utilizes a pixel-wise photometric error map $E_{pho} \in \mathbb{R}^{H \times W}$ and a mask $M_e \in \mathbb{R}^{H \times W}$ as the basis for densification. This masked error map is then normalized into a probability map $P_e \in \mathbb{R}^{H \times W}$. To densify new Gaussian points, the same initialization method described in prior-driven initialization is adopted, with the exception of sampling from $P_e \odot M_e$. The number of new Gaussian points introduced is proportionate to the mask ratio, ensuring controlled expansion of the point set.

There are two scenarios for densification: 1) Before Gaussian optimization, the mask M_e only marks new content,

which is detected via a forward-backward consistency check using bidirectional flow from advanced optical flow estimators (Xu et al. 2022, 2023). And we set the P_e as uniform probability map, to fill new content emerged in a new frame. 2) During Gaussian optimization, P_e is unchanged, and the mask M_e is identified by thresholding the error map E_{pho} , ensuring densification occurs at details-lacking area.

4.2 Alternating Gaussian-Camera Optimization

Once the first frame has been initialized and optimized, for subsequent frames, we adopt a alternating optimization strategy for the camera poses $\{P_i\}$ and the Gaussians $\{G_i\}$.

Movement Clustering of Gaussian Points In constructing dynamic scenes that include both camera and object movements, treating these scenes as static can lead to inaccuracies and loss of crucial temporal information. To better manage this, we propose a method to cluster the scene into still and moving parts, which will be incorporated in the optimization process.

We calculate the epipolar error map based on the optical flow estimated by UniMatch (Xu et al. 2023, 2022). The moving mask M_t at time t is identified by thresholding the epipolar error map. When Gaussian points are initialized or densified, those within M_t are considered as moving points $\{G_t^m\}_t \subseteq \{G_i\}_t$, while others are considered as still points $\{G_t^s\}_t \subseteq \{G_i\}_t$, whose center coordinate $\{\mu_i^s\}$ will stop updating. This simple yet effective movement clustering method enables GFlow to model and track both rigid and non-rigid movement, whether it occurs on objects or other elements like water.

Camera Optimization The camera intrinsic K is estimated by MAST3R (Leroy, Cabon, and Revaud 2024). Be-

tween two consecutive frames, only camera extrinsic E is optimizable, all Gaussian points are frozen. The camera extrinsic $E = [R|t]$ consists of a rotation $R \in \mathbf{SO}(3)$ and a translation $t \in \mathbb{R}^3$.

For a given frame at time t , we optimize the camera extrinsic E_t by minimizing the errors in its photometric appearance, depth and optical flow.

$$E_t^* = \arg \min_{E_t} (\lambda_p \mathcal{L}_{pho} + \lambda_d \mathcal{L}_{dep} + \lambda_f \mathcal{L}_{flo}), \quad (1)$$

where λ_p , λ_d and λ_f are weighting factors. During this optimization, since the moving part is not contribute to camera pose estimation, we will mask out the moving area according to current and previous moving mask M_t, M_{t-1} .

Here, $\mathcal{R}(\cdot)$ denotes the 3D Gaussian splatting rendering process for desired output. The *photometric loss* combines MSE and SSIM (Wang et al. 2004) loss between the rendered image $\hat{I}_t = \mathcal{R}(\{G_i\}_t)$ and the actual frame image I_t .

$$\mathcal{L}_{pho} = \mathcal{L}_{mse}(\hat{I}_t, I_t) + \mathcal{L}_{ssim}(\hat{I}_t, I_t) \quad (2)$$

The *depth loss* is calculated using the L1 loss between the rendered depth $\hat{D}_t = \mathcal{R}(\{G_i\}_t)$ and prior depth D_t . To address discrepancies in scale and shift between the rendered and prior depths, we apply a scale and shift-invariant term on the loss, where a and b are optimizable.

$$\mathcal{L}_{dep} = |(a * \hat{D}_t + b) - D_t| \quad (3)$$

The *optical flow loss* is calculated using the mean squared error (MSE) loss between the movements of Gaussian points in camera view \hat{F}_t and the prior optical flow F_t , to ensure the temporal smoothness of the Gaussian points' trajectories.

$$\mathcal{L}_{flo} = \mathcal{L}_{mse}(\hat{F}_t, F_t) \quad (4)$$

Gaussians Optimization Once the optimized camera extrinsics E_t^* is obtained, we first conduct **pre-optimization gaussians relocation** for those moving Gaussian points $\{G_i^m\}_t$. Initially, we retrieve the 2D coordinates of moving Gaussian points from the previous frame $\{x(G_{i,t-1}^m)\}$. Using these coordinates, we calculate their movement based on the previous frame's optical flow map $\{F_{t-1}(x(G_{i,t-1}^m))\}$ and update their current position: $x(G_{i,t}^m) = x(G_{i,t-1}^m) + F_{t-1}(x(G_{i,t-1}^m))$. With the updated coordinates, we then extract the depth from the current frame's depth map $\{D_t(x(G_{i,t}^m))\}$, and project these points from the camera view to world coordinates using E_t^* . This step ensures that the moving Gaussian points are accurately positioned and serves as a good initialization for subsequent optimization.

Then, the total optimization objectives of Gaussian points contains photometric loss, depth loss, optical flow loss, and additional isotropic loss.

$$\{G_i^*\}_t = \arg \min_{\{G_i\}_t} (\lambda_p \mathcal{L}_{pho} + \lambda_d \mathcal{L}_{dep} + \lambda_f \mathcal{L}_{flo} + \lambda_i \mathcal{L}_{iso}) \quad (5)$$

Where λ_i is a weight factor. The *isotropic loss* is calculated as the mean of the standard deviation of the Gaussian points' 3D scales $\{s_i\}$. In this monocular setting with sparse views, the Gaussians tend to elongate along the view ray direction due to the lack of sufficient multi-view constraints. Applying isotropic loss will encourage the Gaussians to be isotropic, helping to reduce needle-like artifacts.

$$\mathcal{L}_{iso} = \frac{1}{N} \sum_{i=1}^N \text{std}(s_i) \quad (6)$$

The photometric and isotropic loss is applied to all Gaussian points, whereas the depth and optical flow losses focus specifically on the moving cluster $\{G_i^m\}_t$. This approach ensures tailored optimization that balances the dynamics and stability of Gaussian points in the scene.

4.3 Overall Pipeline

To conclude, the overall pipeline can be summarized as follows: Given an image sequence $\{I_t\}_{t=0}^T$ of monocular video input, we first utilizes off-the-shelf algorithms (Xu et al. 2023, 2022; Leroy, Cabon, and Revaud 2024) to derive the corresponding depth $\{D_t\}_{t=0}^T$, optical flow $\{F_t\}_{t=0}^T$ and camera intrinsic K . The initialization of the Gaussians is performed using the prior-driven initialization. Then for each frame I_t at time t , GFlow first divides the Gaussian points $\{G_i\}_t$ into still cluster $\{G_i^s\}_t$ and moving cluster $\{G_i^m\}_t$ according to the optical flow. The optimization process then proceeds in two steps. In the first step, only the camera extrinsics E_t is optimized. This is achieved by aligning the Gaussian points within the still part with the appearance I_t , depth D_t and optical flow F_t . Following that, under the optimized camera extrinsics E_t^* , the Gaussian points G_t are further refined using constraints from the RGB I_t , depth D_t , optical flow F_t , and isotropic loss \mathcal{L}_{iso} . Additionally, the Gaussian points are densified using our proposed pixel-wise strategy to incorporate newly visible scene content. After optimizing the current frame, the procedure — movement clustering, camera optimization, and Gaussian point optimization — is repeated for subsequent frames. It is worth noting that in practice, GFlow is highly efficient, taking only 10 to 20 minutes to optimize a video — significantly faster than other methods that typically require more than an hour.

5 Experiments

Dataset and Metrics We conduct experiments on a challenging video dataset, **DAVIS** (Perazzi et al. 2016; Pont-Tuset et al. 2017) dataset, which contains real-world videos of about 30 ~ 100 frames with various scenarios and motion dynamics. We report the reconstruction quality results on 30 DAVIS 2017 evaluation videos. For quantitative evaluation of reconstruction quality, we report standard PSNR, SSIM and LPIPS (Zhang et al. 2018) metrics.

Implementation Details All image sequences are resized to the shortest side is 480 pixels. The initial number of Gaussian points is set to 50,000. The camera intrinsics K are estimated by MAST3R (Leroy, Cabon, and Revaud 2024). The learning rate for Gaussian optimization is set to $4e-3$,



Figure 3: Our GFlow can explicitly model the dynamic 3D scene in the video. Here we show some rendered examples of videos from DAVIS (Perazzi et al. 2016; Pont-Tuset et al. 2017) dataset in the 3D world space.

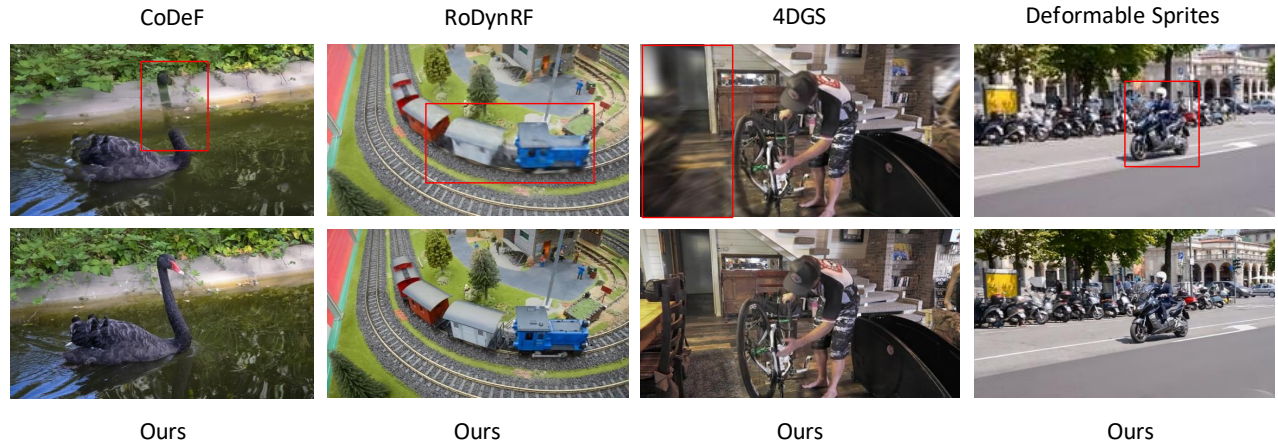


Figure 4: Visual comparison of reconstruction quality on the DAVIS (Perazzi et al. 2016; Pont-Tuset et al. 2017) dataset: CoDeF (Ouyang et al. 2023), RoDynRF (Liu et al. 2023), 4DGS (Yang et al. 2024), and Deformable Sprites (Ye et al. 2022) and Ours.

and for camera optimization, it is set to $1e-3$. The Adam optimizer is used with 500 iterations for Gaussian optimization in the first frame, 150 iterations for camera optimization, and 300 iterations for Gaussian optimization in subsequent frames. The gradient of color is set to zero, enforcing Gaussian points to move rather than lazily changing color. We balance the loss term by setting $\lambda_p = 1$, $\lambda_d = 0.1$, $\lambda_f = 0.01$, and $\lambda_i = 50$. Denoising is conducted at the 150 and 300 steps in the first frame optimization. For subsequent frames, the denoising occurs at the first step with a new content mask applied, and also occurs at the 100 and 200 steps with error-thresholding mask applied. The error threshold in denoising is set to 0.01. All experiments are conducted on a single NVIDIA RTX A5000 GPU. Note that the dynamics within each video could be distinct, so for better reconstruction, the hyperparameters could be tuned in practice.

5.1 Evaluation of Reconstruction Quality

Quantitative Results Reconstructing the 4D world, particularly with camera and content movement, is an extremely challenging task. We choose several methods clos-

est to tackle this task as our baseline. Deformable Sprites (Ye et al. 2022) decomposes the video into layers of persistent motion groups, which are then composed to reconstruct the video. RoDynRF (Liu et al. 2023) uses neural voxel radiance fields to model the dynamic scene and camera poses simultaneously. CoDeF (Ouyang et al. 2023) employs implicit content deformation fields to learn a canonical template for modeling monocular video, but it lacks physical interpretability, such as the ability to estimate camera poses. 4DGS (Yang et al. 2024) models the space and time dimensions for dynamic scenes by formulating unbiased 4D Gaussian primitives, though it requires camera poses as input. We use the camera poses estimated by MAST3R as its input. As shown in Table 1, our GFlow demonstrates significant advantages in reconstruction quality. This improvement stems from its explicit representation and tailored optimization process design, which can adapt Gaussian points over time while maintaining visual content coherence.

Qualitative Results The visual comparison in Figure 4 shows that CoDeF struggles to reconstruct videos with significant movement due to its reliance on representing a video as a canonical template. RoDynRF has difficulty re-

Method	DAVIS		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Deformable Sprites	22.83	0.6983	0.3014
RoDynRF	24.79	0.7230	0.3940
CoDeF	24.89	0.7703	0.2932
4DGS	24.60	0.7315	0.3710
GFlow (Ours)	29.74	0.9205	0.1237
GFlow*	29.21	0.9162	0.1320
w/o \mathcal{L}_{dep}	29.30	0.9086	0.1444
w/o \mathcal{L}_{iso}	26.32	0.8664	0.2003

Table 1: Reconstruction quality results on DAVIS(Perazzi et al. 2016; Pont-Tuset et al. 2017) dataset. Average PSNR, SSIM and LPIPS scores on all videos are reported.

constructing high-quality moving objects. Even with camera pose inputs, 4DGS falls short in reconstructing the entire frame image. Additionally, Deformable Sprites can only reconstruct videos at a very low resolution. In contrast, our GFlow can faithfully reconstruct both static and moving content in high quality. The visual illustration in Figure 3 demonstrates the dynamic 3D scene recovered from monocular videos, showcasing the effectiveness of our approach.

5.2 Ablation Study

Effect of Isotropic Loss Since the monocular video input only provides sparse and underconstrained views, traditional 3DGS, which relies on dense multi-view constraints, struggles to achieve good results. The sparse views will result in needle-like artifacts along the camera view ray direction. As shown in Table 1, the isotropic loss helps to overcome these drawbacks and improves the reconstruction quality.

Effect of Depth Loss Depth loss is used for ensuring a consistent and reasonable 3D geometry structure. Table 1 shows the reconstruction quality will decrease without depth loss.

Effect of Optimizing Camera Pose When preprocessing the monocular video using MAST3R (Leroy, Cabon, and Revaud 2024), it can also estimate camera poses. The results labeled as ‘GFlow*’ in Table 1 show the effect of directly using the camera poses estimated by MAST3R instead of optimizing them. A decrease in quality is observed, demonstrating the necessity of optimizing camera poses. Additionally, GFlow is capable of optimizing camera poses even when most areas in the video are in motion, where MAST3R fails.

5.3 Downstream Video Applications

Various downstream applications can be extended from our GFlow framework, as it is an explicit representation. We encourage readers to view the videos in the website for more intuitionistic illustration.

Point Tracking Due to the nature of GFlow, all Gaussian points can serve as query tracking points, enabling tracking in both 2D and 3D space. The tracking trajectories are

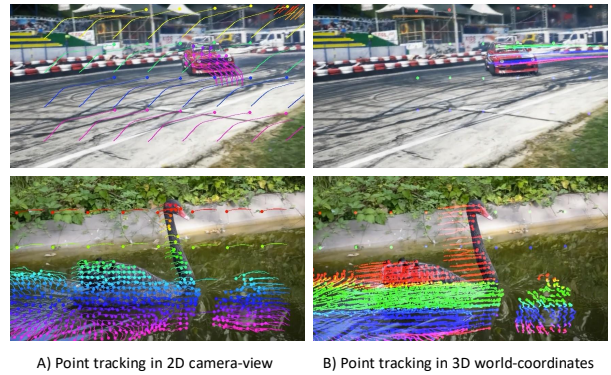


Figure 5: Point tracking visualization on DAVIS dataset. A) tracking in the 2D camera-view which contains joint motion of camera and content movement. B) tracking in the 3D world-coordinates which only present content movement.

illustrated in Figure 5. In conventional **2D tracking**, tracking occurs in the camera view, which includes the combined motion of both the camera and the content. In contrast, the Gaussian points in GFlow reside in 3D world coordinates, representing only content movement. As a result, some **3D tracking** trajectories tend to remain in their original locations, as shown in Figure 5B), because they belong to the static background. These results demonstrate that GFlow can achieve accurate tracking even on water ripples and remains reliable for fast-moving objects and scenes.

Video Object Segmentation Since GFlow drives the Gaussian points to follow the movement of the visual content, given an initial mask, all Gaussian points within this mask can propagate to subsequent frames. This propagation forms a new mask as a concave hull (Park and Oh 2012) around these points. Notably, this capability is a by-product of GFlow, achieved without extra intended training.

Video Editing Since explicit representation can be easily edited: **Camera-level manipulation:** Changing the camera extrinsics can render novel views of dynamic scenes. When combined with camera intrinsics, it can create visual effects like dolly zoom. **Object-level editing:** With the cluster labels of moving Gaussian points, we can add, remove, resize, or stylize these points, allowing for precise object-level editing. **Scene-level editing:** Editing can also be applied to the entire scene, enabling the application of visual effects globally, as illustrated in Figure 1.

6 Conclusion

We have presented ‘GFlow’, a novel framework designed to address the challenging task of reconstructing the 4D world from casual monocular videos. Through Gaussian points allocation and alternating camera-Gaussian optimization, GFlow enables the recovery of dynamic scenes alongside camera poses across frames. Further capabilities such as tracking, segmentation, editing, and novel view rendering, highlighting GFlow’s potential to revolutionize video understanding and manipulation.

Acknowledgments

This project is supported by the National Research Foundation, Singapore, under its Medium Sized Center for Advanced Robotics Technology Innovation, and the Singapore Ministry of Education Academic Research Fund Tier 1 (WBS: A-0009440-01-00).

References

- Bansal, A.; Vo, M.; Sheikh, Y.; Ramanan, D.; and Narasimhan, S. 2020. 4d visualization of dynamic events from unconstrained multi-view videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5366–5375.
- Bian, W.; Wang, Z.; Li, K.; Bian, J.-W.; and Prisacariu, V. A. 2023. Nope-nerf: Optimising neural radiance field with no pose prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4160–4169.
- Cao, A.; and Johnson, J. 2023. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 130–141.
- Charatan, D.; Li, S.; Tagliasacchi, A.; and Sitzmann, V. 2023. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. *arXiv preprint arXiv:2312.12337*.
- Chen, A.; Xu, Z.; Geiger, A.; Yu, J.; and Su, H. 2022. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, 333–350. Springer.
- Flynn, J.; Neulander, I.; Philbin, J.; and Snavely, N. 2016. Deepstereo: Learning to predict new views from the world’s imagery. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5515–5524.
- Fridovich-Keil, S.; Meanti, G.; Warburg, F. R.; Recht, B.; and Kanazawa, A. 2023. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12479–12488.
- Fu, Y.; Liu, S.; Kulkarni, A.; Kautz, J.; Efros, A. A.; and Wang, X. 2023. Colmap-free 3d gaussian splatting. *arXiv preprint arXiv:2312.07504*.
- Kanopoulos, N.; Vasanthavada, N.; and Baker, R. L. 1988. Design of an image edge detection filter using the Sobel operator. *IEEE Journal of solid-state circuits*, 23(2): 358–367.
- Kasten, Y.; Ofri, D.; Wang, O.; and Dekel, T. 2021. Layered neural atlases for consistent video editing. *ACM Transactions on Graphics (TOG)*, 40(6): 1–12.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4): 1–14.
- Kong, H.; Yang, X.; and Wang, X. 2025. Efficient Gaussian Splatting for Monocular Dynamic Scene Rendering via Sparse Time-Variant Attribute Modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Lei, J.; Weng, Y.; Harley, A.; Guibas, L.; and Daniilidis, K. 2024. MoSca: Dynamic Gaussian Fusion from Casual Videos via 4D Motion Scaffolds. *arXiv preprint arXiv:2405.17421*.
- Leroy, V.; Cabon, Y.; and Revaud, J. 2024. Grounding Image Matching in 3D with MAST3R. *arXiv preprint arXiv:2406.09756*.
- Li, T.; Slavcheva, M.; Zollhoefer, M.; Green, S.; Lassner, C.; Kim, C.; Schmidt, T.; Lovegrove, S.; Goesele, M.; Newcombe, R.; et al. 2022. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5521–5531.
- Li, Z.; Wang, Q.; Cole, F.; Tucker, R.; and Snavely, N. 2023. Dynibar: Neural dynamic image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4273–4284.
- Lin, C.-H.; Ma, W.-C.; Torralba, A.; and Lucey, S. 2021a. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5741–5751.
- Lin, H.; Peng, S.; Xu, Z.; Xie, T.; He, X.; Bao, H.; and Zhou, X. 2023. High-fidelity and real-time novel view synthesis for dynamic scenes. In *SIGGRAPH Asia 2023 Conference Papers*, 1–9.
- Lin, K.-E.; Xiao, L.; Liu, F.; Yang, G.; and Ramamoorthi, R. 2021b. Deep 3d mask volume for view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1749–1758.
- Liu, Q.; Liu, Y.; Wang, J.; Lv, X.; Wang, P.; Wang, W.; and Hou, J. 2024. MoDGS: Dynamic Gaussian Splatting from Casually-captured Monocular Videos. *arXiv preprint arXiv:2406.00434*.
- Liu, Y.-L.; Gao, C.; Meuleman, A.; Tseng, H.-Y.; Saraf, A.; Kim, C.; Chuang, Y.-Y.; Kopf, J.; and Huang, J.-B. 2023. Robust dynamic radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13–23.
- Luiten, J.; Kopanas, G.; Leibe, B.; and Ramanan, D. 2023. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4): 1–15.
- Ouyang, H.; Wang, Q.; Xiao, Y.; Bai, Q.; Zhang, J.; Zheng, K.; Zhou, X.; Chen, Q.; and Shen, Y. 2023. Codef: Content deformation fields for temporally consistent video processing. *arXiv preprint arXiv:2308.07926*.
- Park, J.-S.; and Oh, S.-J. 2012. A new concave hull algorithm and concaveness measure for n-dimensional datasets. *Journal of Information science and engineering*, 28(3): 587–600.
- Park, K.; Sinha, U.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Seitz, S. M.; and Martin-Brualla, R. 2021a. Nerfies: Deformable neural radiance fields. In *Proceedings of the*

- IEEE/CVF International Conference on Computer Vision*, 5865–5874.
- Park, K.; Sinha, U.; Hedman, P.; Barron, J. T.; Bouaziz, S.; Goldman, D. B.; Martin-Brualla, R.; and Seitz, S. M. 2021b. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*.
- Perazzi, F.; Pont-Tuset, J.; McWilliams, B.; Van Gool, L.; Gross, M.; and Sorkine-Hornung, A. 2016. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 724–732.
- Pont-Tuset, J.; Perazzi, F.; Caelles, S.; Arbeláez, P.; Sorkine-Hornung, A.; and Van Gool, L. 2017. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*.
- Pumarola, A.; Corona, E.; Pons-Moll, G.; and Moreno-Noguer, F. 2021. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10318–10327.
- Schonberger, J. L.; and Frahm, J.-M. 2016. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4104–4113.
- Schönberger, J. L.; Zheng, E.; Frahm, J.-M.; and Pollefeys, M. 2016. Pixelwise view selection for unstructured multi-view stereo. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, 501–518. Springer.
- Shao, R.; Zheng, Z.; Tu, H.; Liu, B.; Zhang, H.; and Liu, Y. 2023. Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16632–16642.
- Stearns, C.; Harley, A.; Uy, M.; Dubost, F.; Tombari, F.; Wetstein, G.; and Guibas, L. 2024. Dynamic gaussian marbles for novel view synthesis of casual monocular videos. In *SIG-GRAPH Asia 2024 Conference Papers*, 1–11.
- Sun, J.; Jiao, H.; Li, G.; Zhang, Z.; Zhao, L.; and Xing, W. 2024. 3dstream: On-the-fly training of 3d gaussians for efficient streaming of photo-realistic free-viewpoint videos. *arXiv preprint arXiv:2403.01444*.
- Tian, F.; Du, S.; and Duan, Y. 2023. Mononerf: Learning a generalizable dynamic radiance field from monocular videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17903–17913.
- Wang, Q.; Ye, V.; Gao, H.; Austin, J.; Li, Z.; and Kanazawa, A. 2024. Shape of motion: 4d reconstruction from a single video. *arXiv preprint arXiv:2407.13764*.
- Wang, S.; Leroy, V.; Cabon, Y.; Chidlovskii, B.; and Revaud, J. 2023. DUST3R: Geometric 3D Vision Made Easy. *arXiv preprint arXiv:2312.14132*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wang, Z.; Wu, S.; Xie, W.; Chen, M.; and Prisacariu, V. A. 2021. NeRF-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*.
- Wu, G.; Yi, T.; Fang, J.; Xie, L.; Zhang, X.; Wei, W.; Liu, W.; Tian, Q.; and Wang, X. 2023. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*.
- Xia, Y.; Tang, H.; Timofte, R.; and Van Gool, L. 2022. Sinerf: Sinusoidal neural radiance fields for joint pose estimation and scene reconstruction. *arXiv preprint arXiv:2210.04553*.
- Xu, H.; Zhang, J.; Cai, J.; Rezatofighi, H.; and Tao, D. 2022. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8121–8130.
- Xu, H.; Zhang, J.; Cai, J.; Rezatofighi, H.; Yu, F.; Tao, D.; and Geiger, A. 2023. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Yang, Z.; Gao, X.; Zhou, W.; Jiao, S.; Zhang, Y.; and Jin, X. 2023a. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101*.
- Yang, Z.; Yang, H.; Pan, Z.; and Zhang, L. 2024. Real-time Photorealistic Dynamic Scene Representation and Rendering with 4D Gaussian Splatting. In *International Conference on Learning Representations (ICLR)*.
- Yang, Z.; Yang, H.; Pan, Z.; Zhu, X.; and Zhang, L. 2023b. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. *arXiv preprint arXiv:2310.10642*.
- Ye, V.; Li, Z.; Tucker, R.; Kanazawa, A.; and Snavely, N. 2022. Deformable Sprites for Unsupervised Video Decomposition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yifan, W.; Serena, F.; Wu, S.; Öztireli, C.; and Sorkine-Hornung, O. 2019. Differentiable surface splatting for point-based geometry processing. *ACM Transactions on Graphics (TOG)*, 38(6): 1–14.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.