

# Morpheus: Neural Dynamic 360° Surface Reconstruction from Monocular RGB-D Video

Hengyi Wang Jingwen Wang Lourdes Agapito  
 Department of Computer Science, University College London  
 {hengyi.wang.21, jingwen.wang.17, l.agapito}@ucl.ac.uk

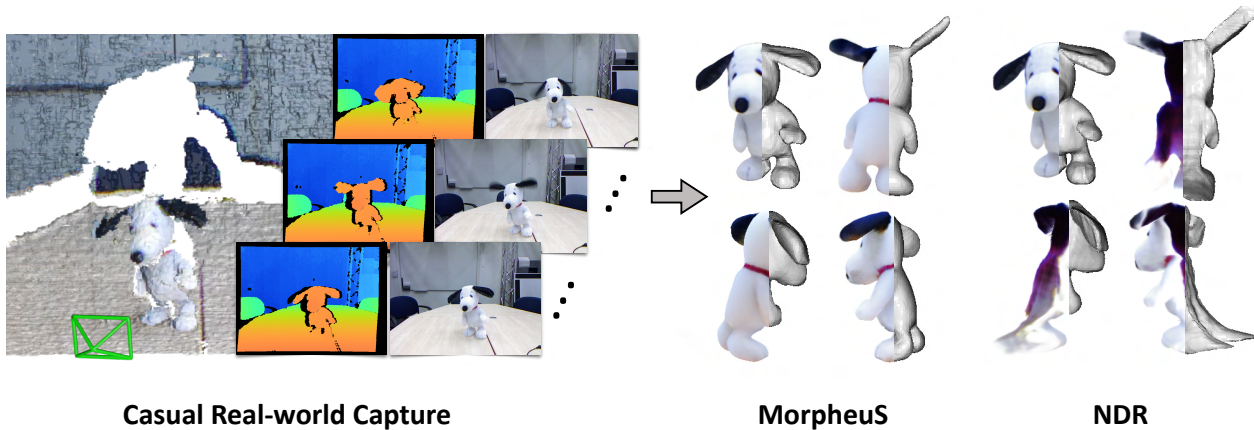


Figure 1. We propose Morpheus, a dynamic scene reconstruction method that leverages neural implicit representations and diffusion priors for achieving 360° reconstruction of a moving object from a monocular RGB-D video. Our approach can achieve both metrically accurate reconstruction of the observed regions and photo-realistic completion of unobserved regions of a dynamic scene.

## Abstract

Neural rendering has demonstrated remarkable success in dynamic scene reconstruction. Thanks to the expressiveness of neural representations, prior works can accurately capture the motion and achieve high-fidelity reconstruction of the target object. Despite this, real-world video scenarios often feature large unobserved regions where neural representations struggle to achieve realistic completion. To tackle this challenge, we introduce Morpheus, a framework for dynamic 360° surface reconstruction from a casually captured RGB-D video. Our approach models the target scene as a canonical field that encodes its geometry and appearance, in conjunction with a deformation field that warps points from the current frame to the canonical space. We leverage a view-dependent diffusion prior and distill knowledge from it to achieve realistic completion of unobserved regions. Experimental results on various real-world and synthetic datasets show that our method can achieve high-fidelity 360° surface reconstruction of a deformable object from a monocular RGB-D video. Project page: <https://hengyiwang.github.io/projects/morpheus>.

## 1. Introduction

The challenge of reconstructing the dense 3D geometry, appearance, and motion of dynamic scenes from videos has persisted for decades, offering a wide range of applications in virtual reality and augmented reality. Traditional dense non-rigid reconstruction methods often assume strong priors on the object’s shape and motion due to the inherent ambiguities of the problem [48, 89]. Model-based approaches [1, 11, 27, 28, 76, 87] require 3D parametric models [2, 6, 31, 40, 49, 100] of specific object categories such as humans [2, 40], faces [6, 31], or animals [100], resulting in promising but category-constrained reconstructions. In this paper, we focus on model-agnostic approaches that could generalize to generic non-rigid objects. This task is often approached by jointly fusing a canonical shape and a per-frame deformation field from an RGB-D video sequence [7, 8, 25, 47, 69, 70]. Due to the constraints of physical sensors and the reliance on traditional voxel-grid representations, they often struggle to produce high-fidelity surface details and can only reconstruct observed regions.

Recent attention has shifted towards leveraging neural representations for scene reconstruction. Neural Radiance Fields (NeRF) [44] represent scene density and color in the weights of a neural network. In combination with volume

rendering, NeRF achieves unprecedented performance on novel view synthesis. Many follow-up works adapt this idea into surface reconstruction [4, 83, 84, 93], SLAM [29, 72, 82, 98], dynamic scene reconstruction [9, 52, 53, 56] and more. Notably, NDR [9] adopts NeRF-style representations for dynamic RGB-D surface reconstruction. Thanks to the expressiveness of neural representations, they achieve accurate and smooth surface reconstruction. However, real-world captures often contain large unobserved regions. The inherent smoothness of neural representations, and the difficulty to incorporate domain-specific priors, cause current methods to fall short of achieving realistic completion.

In this paper, we present MorpheuS, a framework for neural dynamic 360° surface reconstruction from casually captured monocular RGB-D videos. We represent the target object in a hyper-dimensional canonical field and adopt a deformation field to deform the target object from observation space to hyper-dimensional canonical space, in similar spirit to HyperNeRF [53]. In terms of reconstruction of the unobserved regions, inspired by DreamFusion [55], we employ a diffusion prior, i.e. Zero-1-to-3 [38], and perform Score Distillation Sampling [55] (SDS) to distill knowledge from the diffusion prior to complete the unobserved geometry and appearance of the target object.

Our key contribution is to demonstrate the capability to learn metrically accurate geometry and deformations of dynamic objects from casually captured RGB-D videos while achieving realistic completion in unobserved regions with diffusion priors. We propose a temporal view-dependent SDS strategy to improve the realism of completion while learning an accurate deformation field. A canonical space regularization strategy is used to avoid the trivial solution of surface completion. To the best of our knowledge, MorpheuS is the first to achieve accurate, photo-realistic 360° surface reconstruction of an arbitrary dynamic object from casually captured monocular RGB-D video.

## 2. Related Work

**Dense Non-rigid Reconstruction.** Traditional dense dynamic scene reconstruction methods often require strong prior knowledge about the object shape and motion, fitting 3D prior parametric models [1, 11, 27, 28, 31, 40, 61, 76, 87, 100] to video sequences of a specific object category, and thus cannot work with arbitrary unseen objects. Early model-agnostic methods either require accurate dense correspondences [20] or a template shape of the object to be reconstructed first before tracking the dynamic motion [5, 73, 95, 99]. The seminal work DynamicFusion [47] was the first direct and template-free method to achieve real-time simultaneous tracking and reconstruction of generic non-rigid objects via joint optimization of a canonical shape and a per-frame deformation field. Follow-up methods improve the robustness of motion tracking and reconstruction

quality by incorporating photo-metric constraints [21], classical [25] or learned correspondences [7, 8], reformulating the optimization as direct dense SDF alignment [69, 70]. Although these methods have shown promising reconstruction results, they fail to complete the unobserved regions due to their inherent traditional scene representation.

**Neural Rendering for Dynamic Scenes.** The success of NeRF [44] in novel view synthesis has inspired many follow-up works [9, 16, 18, 19, 32, 34, 43, 45, 52, 53, 56, 71, 77, 79, 91, 92] in dynamic scene reconstruction. These approaches can be broadly classified into two paradigms: **a)** Directly modeling scene geometry and motion as a unified 4D space-time field in world (observation) space [10, 17, 18, 24, 32, 34, 64], where motion priors, such as scene flow or trajectory can be involved as an additional regularization [32, 34, 79]. **b)** Learning a canonical shape by warping the shape from observation space to canonical space via a separate deformation field [9, 15, 16, 19, 45, 52, 53, 56, 71, 77, 88, 91, 92]. Among these works, NDR [9] is particularly relevant to our work. They achieve high-fidelity dynamic surface reconstruction with neural representations. However, the continuity of neural representations does not allow realistic completion in unobserved regions.

**Neural Rendering with Diffusion Priors.** Recent years have witnessed remarkable progress in diffusion-based foundation models [60, 62], unlocking unprecedented capabilities for diverse applications such as image synthesis, editing, and shape generation [54]. DreamFusion [55] is a notable example that leverages pre-trained image diffusion priors to optimize a NeRF model through SDS, achieving high-quality text-to-3D generation. Building upon DreamFusion’s success, numerous follow-up works have emerged, particularly in text-to-3D [12, 35, 42, 81, 86, 97], image-to-3D [41, 57, 75, 90], scene editing [23, 50, 63], etc. Some other works [30, 38, 39, 66] try to improve the generation quality via fine-tuning pre-trained diffusion models. A noteworthy example is Zero-1-to-3 [38], which fine-tunes stable-diffusion [60] with relative pose conditioning.

**Concurrent works.** Recently, several works have been proposed for dynamic scene generation [26, 37, 51, 59, 68, 80, 94] and editing [50, 63]. MAV3D [68] adopts a text-to-video diffusion model [67] for text-to-3D generation. Control4D [63] utilizes a 2D diffusion prior to train a 4D GAN for 4D editing of human portraits. AvatarStudio [50] fine-tunes the diffusion prior with time stamps and viewpoints, enabling editing of human head avatars with a view-and-time aware SDS loss. Consistent4d [26], Efficient4D [51], and DpDy [80] use diffusion prior for RGB-based dynamic scene reconstruction. Different from these works, our focus is learning metrically accurate motion and geometry from monocular RGB-D video while achieving realistic completion in the unobserved region with a diffusion prior.

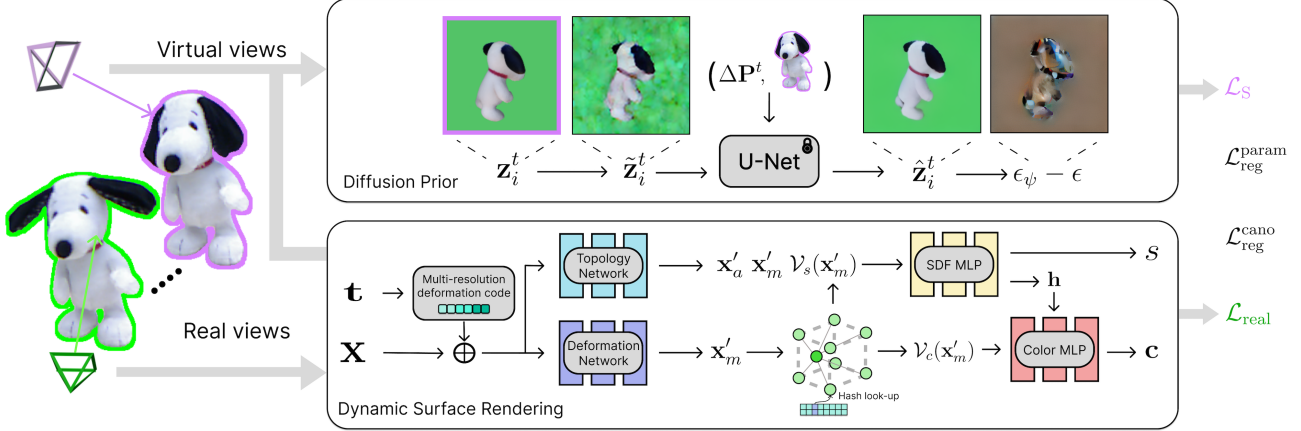


Figure 2. **Overview of MorpheuS.** 1) Dynamic surface rendering: we model the target dynamic scene via a deformation field that maps a point from observation space to a hyper-dimensional canonical space and a canonical field that decodes the point into SDF and color. 2) Diffusion prior: we leverage a diffusion prior and perform SDS to complete the unobserved region. Note here the de-noising process is in latent space. All visualization is generated via decoding the latent vector for illustration purposes. 3) Optimization: We optimize the scene representation using real view supervision  $\mathcal{L}_{\text{real}}$ , SDS loss  $\mathcal{L}_S$ , canonical regularization  $\mathcal{L}_{\text{reg}}^{\text{cano}}$ , and parameter regularization  $\mathcal{L}_{\text{reg}}^{\text{param}}$ .

### 3. Method

The goal of MorpheuS is to achieve 360° reconstruction of an arbitrary dynamic object from a casual monocular RGB-D video  $\{I_t\}_{t=1}^N, \{D_t\}_{t=1}^N$  with known camera intrinsics  $K \in \mathbb{R}^{3 \times 3}$  and camera poses  $P_t \in \mathbb{SE}(3)$ . Fig. 2 shows the overview of MorpheuS pipeline. The target scene is represented as a hyper-dimensional canonical field and a deformation field (Sec. 3.1). The unobserved region reconstruction is addressed through the distillation of knowledge from a diffusion prior (Sec. 3.2). We use both real-view observation and diffusion prior to produce supervision for our scene representation. Several regularization terms are used to improve the robustness of the optimization process (Sec. 3.3).

#### 3.1. Dynamic Neural Surface Rendering

**Deformation Field.** Inspired by the recent success of dynamic scene reconstruction [9], we model the deformation of the scene via a deformation network  $D(\cdot)$  and a topology network  $T(\cdot)$ . Given a point  $\mathbf{x}_t \in \mathbb{R}^3$  in observation space at frame  $t$ ,  $D(\cdot)$  and  $T(\cdot)$  map the point from the observation space to a point in the hyper-dimensional canonical space  $\mathbf{x}' \in \mathbb{R}^{3+m}$ :

$$\mathbf{x}' = \{\mathbf{x}'_m, \mathbf{x}'_a\} = \{\mathbf{x}_t + D(\phi(\mathbf{x}_t), \mathcal{V}_t(t)), T(\phi(\mathbf{x}_t), \mathcal{V}_t(t))\}, \quad (1)$$

where  $\mathbf{x}'_m \in \mathbb{R}^3$  is the 3D coordinate of the deformed point and  $\mathbf{x}'_a \in \mathbb{R}^m$  is the ambient coordinate predicted by the topology network  $T(\cdot)$ .  $\phi(\cdot)$  and  $\mathcal{V}_t(\cdot)$  are encoding functions of the location  $\mathbf{x}_t$  and time  $t$ . We use a frequency encoding  $\phi(\cdot)$  to model the position  $\mathbf{x}_t$  for simplicity. For  $\mathcal{V}_t(\cdot)$ , we use a multi-resolution 1-D feature grid instead of a per-frame deformation code [9, 53]. The feature vector for each frame  $t$  is obtained via linear interpolation.

**Hyper-dimensional Canonical Field.** The 3D geometry and appearance of the target object are represented as the zero-level set of the signed distance function (SDF) field  $s$  and color  $\mathbf{c}$ :

$$s, \mathbf{h} = f_\gamma(\mathcal{V}_s(\mathbf{x}'_m), \mathbf{x}'_a), \quad (2)$$

$$\mathbf{c} = f_\alpha(\mathcal{V}_c(\mathbf{x}'_m), \mathbf{h}), \quad (3)$$

where  $\mathcal{V}_s$  and  $\mathcal{V}_c$  represent the encoding of the input position  $\mathbf{x}'_m$ . Here we choose the Hash encoding [46] for the fast optimization speed. To incorporate the diffusion prior, a good initialization is important for the canonical field. Existing works in text-to-3D tasks often choose to initialize the radiance field with a fixed blob function [35, 55]. However, we find that the deformation field fails to converge with the fixed blob function as an initialization. Thus, we adopt the geometric initialization following SAL [3]. Since it is not straightforward to map the Hash features with ambient coordinates into a sphere, we append the 3D coordinate  $\mathbf{x}'_m$  and use zero initialization to effectively mask out the remaining coordinates in the input layer of the decoder  $f_\gamma(\cdot)$ .

**Volume Rendering.** We transform the SDF into volume density  $\sigma_\beta(s)$  following VolSDF [93], where  $\beta$  is a learnable parameter which controls the smoothness of the surface. The volume density is then used for color  $C$  and depth  $D$  rendering:

$$C = \sum_{i=1}^N T_i \alpha_i \mathbf{c}_i, \quad D = \sum_{i=1}^N T_i \alpha_i d_i, \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j). \quad (4)$$

Here  $T_i$  represents transmittance. The alpha value  $\alpha_i$  is calculated as  $1 - \exp(-\sigma_i \delta_i)$ , where  $\delta_i$  denotes the distance between neighboring sample points.

### 3.2. Diffusion Priors for 360° Reconstruction

We incorporate Zero-1-to-3 [38], a view-conditioned latent diffusion model, and distill knowledge from it to complete the unobserved region using SDS [55]:

$$\nabla_{\theta} \mathcal{L}_S = \mathbb{E}_{i, \epsilon} \left[ w(i) (\epsilon_{\psi}(\tilde{\mathbf{z}}_i; \mathbf{I}_r, \Delta \mathbf{P}, i) - \epsilon) \frac{\partial \mathbf{I}_v}{\partial \theta} \right], \quad (5)$$

where  $\mathbf{I}_v$  is the rendered image from the sampled virtual view, which has a relative pose  $\Delta \mathbf{P}$  with respect to the real image observation  $\mathbf{I}_r$  at the reference view. Here the relative pose  $\Delta \mathbf{P} = [\Delta r, \Delta \omega, \Delta \phi] = [r_v - r_r, \omega_v - \omega_r, \varphi_v - \varphi_r]$  is parameterized in polar-coordinate. Conditioned on the reference view  $\mathbf{I}_r$ , the relative pose  $\Delta \mathbf{P}$ , and the sampled time-step  $i \sim \mathcal{U}(0, 1)$ , the de-noising U-Net  $\epsilon_{\psi}(\cdot)$  applies a de-noising step on the noisy latent  $\tilde{\mathbf{z}}_i$  of the rendered view  $\mathbf{I}_v$ .  $\theta$  is the learn-able parameters of our dynamic scene model.  $w(i)$  is a time-dependent weighting term.

**Temporal View-dependent SDS.** As our input is a monocular video of a dynamic object, we condition the diffusion model on different frames so as to capture the motion of the target object. We select a keyframe for every  $k$  frame as the conditioning reference views of the diffusion prior. For any given rendered virtual view  $\mathbf{I}_v^t$  at frame  $t$ , we choose its nearest keyframe  $t_k$  as the reference view condition and modify the Eq. 5 as:

$$\nabla_{\theta} \mathcal{L}_S^t = \mathbb{E}_{i, \epsilon} \left[ \hat{w}(i) (\epsilon_{\psi}(\tilde{\mathbf{z}}_i^t; \mathbf{I}_r^{t_k}, \Delta \mathbf{P}^t, i) - \epsilon) \frac{\partial \mathbf{I}_v^t}{\partial \theta} \right]. \quad (6)$$

The latent vectors of each keyframe for conditioning are pre-computed to save the computational cost. To mitigate the negative effects of view-inconsistency stemming from the SDS loss, inspired by [74], we incorporate a view-dependant modulation term to the weighting term  $\hat{w}(i) = w(i)(\exp(|\Delta \mathbf{P}_{(1:2)}^t|) - 1)$ . This reduces the gradient for virtual views that have larger view differences with respect to the reference view:

$$|\Delta \mathbf{P}_{(1:2)}^t| = \arccos \left( \frac{\mathbf{o}_r \cdot \mathbf{o}_v}{\|\mathbf{o}_r\| \|\mathbf{o}_v\|} \right) / \pi, \quad (7)$$

where  $\mathbf{o}_r$  and  $\mathbf{o}_v$  are the Cartesian coordinates of the camera frame origins of the reference and virtual views.

**Implementation Details.** Following Dreamfusion [55], we use the shading model for rendering virtual views and randomly replace albedo with white color to produce texture-less rendering. A random background is used for the rendered image from a sampled virtual view. In order to ensure the convergence of the deformation field, we perform SDS every  $j$  iteration. The SDS loss is generated via a reparameterization trick [22, 74]:

$$\mathcal{L}_S^t = \|\text{stopgrad}(\mathbf{z}_i^t - \text{grad}) - \mathbf{z}_i^t\|^2, \quad (8)$$

where  $\text{grad} = \hat{w}(i) (\epsilon_{\psi}(\tilde{\mathbf{z}}_i^t; \mathbf{I}_r^{t_k}, \Delta \mathbf{P}^t, i) - \epsilon)$ .  $\mathcal{L}_S^t$  will be combined with a real view loss for optimization.

### 3.3. Optimization

The optimization of our scene representation can be formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{real}} + \mathcal{L}_S + \mathcal{L}_{\text{reg}}^{\text{cano}} + \mathcal{L}_{\text{reg}}^{\text{param}}. \quad (9)$$

**Real View Supervision.** Given a sampled input frame  $I_t$ , we randomly sample a batch of pixels and cast rays through those pixels to perform volume rendering. The real-view supervision is formulated as:

$$\mathcal{L}_{\text{real}} = \underbrace{\mathcal{L}_c + \mathcal{L}_d + \mathcal{L}_m}_{\text{rendering loss}} + \underbrace{\mathcal{L}_{\text{sdf}} + \mathcal{L}_{\text{surf}} + \mathcal{L}_{\text{smooth}}}_{\text{per-point loss}}. \quad (10)$$

We use  $\ell_2$  loss for color  $\mathcal{L}_c$  and depth  $\mathcal{L}_d$ , cross-entropy loss for mask  $\mathcal{L}_m$ . In terms of the per-point loss, following [4, 82, 83] we set up a truncation distance  $tr$  and apply pseudo-SDF loss  $\mathcal{L}_{\text{sdf}}$  to the ray-points within the truncation region. To speed up the convergence of our model, we employ  $\mathcal{L}_{\text{surf}}$  that directly supervises the SDF and color of the back-projected surface point using the input color and depth. Since Hash encoding does not have global continuity, resulting in noisy surface reconstruction, we further apply a normal smoothness term  $\mathcal{L}_{\text{smooth}}$  to near-surface points to encourage local smoothness of the SDF gradient in observation space:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{|S|} \sum_{\mathbf{x} \in S} \|\nabla s(\mathbf{x}) - \nabla s(\mathbf{x} + \delta \mathbf{x})\|^2, \quad (11)$$

where  $S$  is obtained by uniformly sampling  $n$  points from  $\hat{D} - tr/2$  to  $\hat{D} + tr/2$  along sampled rays. The perturbed point  $\mathbf{x} + \delta \mathbf{x}$  is sampled on a circle centered at  $\mathbf{x}$  with radius  $r$  orthogonal to the gradient vector  $\nabla s(\mathbf{x})$ . Following [33], we compute the SDF gradient via finite difference.

**Canonical Space Regularization.** We find that performing regularization using all points on sampled rays in observation space may lead to degenerate solutions. Instead, we propose to apply constraints directly to each individual slice of the hyper-dimensional canonical space by skipping the deformation network:

$$\mathbf{x}'_{\text{reg}} = \{\mathbf{x}_t, T(\phi(\mathbf{x}_t), \mathcal{V}_t(t))\}. \quad (12)$$

As we sample rays from 360° views, ensuring that the sampled points can span the entire space across different frames, using coordinates in Eq. 12 for regularization can be viewed as performing regularization in the entire hyper-dimensional canonical space. The regularization loss of the canonical space can be written as:

$$\mathcal{L}_{\text{reg}}^{\text{cano}} = \mathcal{L}_{\text{ori}} + \mathcal{L}_{\text{normal}} + \mathcal{L}_{\text{eik}}, \quad (13)$$

where  $\mathcal{L}_{\text{ori}}$  is the orientation loss used in DreamFusion [55].  $\mathcal{L}_{\text{eik}}$  is the eikonal loss that enforces the SDF gradient to

have a unit norm.  $\mathcal{L}_{\text{normal}}$  is the normal smoothness regularization on each sampled point.

**Parameter Space Regularization.** We also perform regularization on the parameters of our model:

$$\mathcal{L}_{\text{reg}}^{\text{param}} = \mathcal{L}_{\text{code}} + \mathcal{L}_{\beta}. \quad (14)$$

$\mathcal{L}_{\text{code}}$  is a regularization term on deformation code to encourage temporal smoothness of object motion:

$$\mathcal{L}_{\text{code}} = \frac{1}{|W_t|} \sum_{k \in W_t} \|\mathcal{V}_t(k) - \mathcal{V}_t(k-1) - \mathcal{V}_t(k+1)\|^2, \quad (15)$$

where  $W_t = \{\dots, t-1, t, t+1, \dots\}$  is a local window around  $t$ .  $\mathcal{L}_{\beta}$  is  $\ell_1$  loss used for minimizing  $\beta$  in  $\sigma_{\beta}(s)$ .

## 4. Experiments

### 4.1. Experimental Setup

**Datasets.** We evaluate our method on 9 real-world scenes that are from KillingFusion [69], DeepDeform [8], and iPhone dataset [19]. All videos in those datasets are captured with a consumer-grade RGB-D camera. Object masks are obtained with off-the-shelf object segmentation methods [13, 14, 36]. Additionally, we also evaluate our methods on 3 synthetic scenes from AMA [78] and BANMo [91] dataset which ground-truth meshes at all time-stamps are available. For AMA dataset [78], we render depth images from the calibrated 8 views aligned with the provided GT RGB images and masks and select the front-view observations for optimisation. For BANMo dataset [91], we follow a similar protocol except that all the RGBs, depths and masks are rendered. Following NDR [9], we compensate for the object’s rigid motion and relative transformation in the data pre-processing stage using Robust-ICP [96].

**Metrics.** We adopt accuracy (acc. [cm]) and completion (comp. [cm]) to evaluate the surface reconstruction quality and CLIP similarity [58] to evaluate the realism of completion. For 9 real-world scenes, we perform mesh culling for evaluation due to the lack of complete GT meshes. The CLIP similarity is computed between RGB frames and the novel view renderings generated by our optimized model. These renderings follow a 360° trajectory around the object, which is in motion over time. For 3 synthetic scenes, we evaluate the complete mesh and render RGB images from the other 7 views to obtain CLIP similarity.

**Implementation Details.** We run Morpheus on a desktop PC with an Intel Core i7-13700K CPU and NVIDIA RTX 4090 GPU with 24GB memory. The model takes around 2-3 hours to train and requires 10-22GB of memory depending on the size of the input frames. We employ a coarse-to-fine training strategy that progressively increases the number of frequency bands of the positional encoding

and hash encoding by a modulation term [9, 53, 85]. Additionally, we discover that when the learning rate is small, the learning of the canonical field (with hash encoding) is significantly faster than the deformation field (with positional encoding). This can let the model learn a good initialization of the canonical shape. Thus, we apply a learning rate warm-up strategy. In the warm-up phase, we freeze the deformation field when training the novel view with  $\mathcal{L}_S$  and set the timestep range of the diffusion prior to be [0.02, 0.5] to ensure the model learns a good canonical shape representation. In the second phase, we reduce the timestep range to [0.02, 0.2] to allow the model to learn accurate motion. The rendered resolution is doubled to improve the quality of the geometry and appearance completion.

### 4.2. Evaluation

Tab. 1-2 show the per-scene quantitative results on 9 real-world scenes and 4 synthetic scenes respectively. Both results show that our model demonstrates superior accuracy and completion in surface reconstruction, outperforming NDR which exhibits spurious surface completion due to the continuity MLPs (See Teddy in Fig. 3). Additionally, our method achieves a significantly better CLIP similarity in comparison to NDR, thanks to the effective knowledge distillation from the diffusion prior.

Fig. 3 shows a qualitative comparison of dynamic 360° reconstruction, highlighting the ability of diffusion priors to eliminate spurious surface prediction and achieve photo-realistic completion. Notably, our model exhibits robustness in completing unobserved regions even in the presence of significant motion changes, as illustrated in the last row of the Teddy sequence. Fig 4 shows a qualitative result of novel view synthesis. Training dynamic neural representations within a monocular setting typically leads to poor quality novel view synthesis in the case of large viewpoint differences. In contrast, our diffusion prior helps to alleviate this issue allowing 360° novel view synthesis. Please refer to supplementary materials for further quantitative analysis on novel view synthesis.

### 4.3. Analysis

We further analyze different components of our proposed method by conducting a set of ablation experiments on the 9 real-world scenes with quantitative results shown in Tab. 3 and qualitative demonstrations in Fig. 5-8.

**Temporal View-dependent SDS.** We start by analyzing the diffusion prior. As in Tab. 3, our full model (last row) exhibits a significant improvement in CLIP similarity compared to the variant without diffusion priors (first row). Notably, the results also highlight the diffusion priors as a regularization for surface reconstruction in the observed region, leading to the improvement in reconstruction accuracy (0.88 v.s. 0.99) by effectively eliminating spurious surfaces.



Figure 3. Real-world dataset reconstruction results (From left to right, top to bottom: FROG, TEDDY, HUMAN2, and MOCHI). NDR [9] achieves high-quality surface reconstruction in the observed region but fails to produce photo-realistic completion, resulting in spurious surfaces in unobserved regions. In contrast, our method can produce high-quality 360° surface reconstruction.

Method	Metric	KillingFusion [69]			DeepDeform [8]			iPhone [19]			Avg.
		Frog	Duck	Snoopy	Human1	Human2	Human3	Teddy	Mochi	Haru	
NDR	Acc. [cm] ↓	1.25	1.34	1.32	0.74	0.83	0.83	6.12	2.19	2.73	1.92
	Comp. [cm] ↓	0.99	1.15	1.05	0.63	0.68	0.65	1.91	1.06	1.73	1.09
	Clip sim. ↑	82.86	82.52	80.66	82.22	83.15	79.24	75.83	74.12	77.19	79.75
Ours	Acc. [cm] ↓	<b>0.77</b>	<b>0.88</b>	<b>1.05</b>	<b>0.57</b>	<b>0.61</b>	<b>0.78</b>	<b>0.99</b>	<b>0.77</b>	<b>1.51</b>	<b>0.88</b>
	Comp. [cm] ↓	<b>0.67</b>	<b>0.81</b>	<b>0.92</b>	<b>0.49</b>	<b>0.56</b>	<b>0.61</b>	<b>0.93</b>	<b>0.60</b>	<b>1.47</b>	<b>0.78</b>
	Clip sim. ↑	<b>90.24</b>	<b>93.15</b>	<b>90.82</b>	<b>86.41</b>	<b>88.75</b>	<b>82.67</b>	<b>80.50</b>	<b>81.57</b>	<b>86.87</b>	<b>86.77</b>

Table 1. Per-scene quantitative results on real world datasets, including KillingFusion [69], DeepDeform [8], and iPhone [19] dataset. The 3D metrics are obtained by comparing the back-projected depth and our culled mesh of each frame.

Method	Metric	AMA [78]		BANMo [91]		Avg.
		Samba	Swing	Eagle1	Eagle2	
NDR	Acc. [cm] ↓	2.32	2.70	6.51	15.42	6.73
	Comp. [cm] ↓	1.94	2.17	2.58	3.70	2.59
	Clip sim. ↑	86.72	88.03	86.23	81.64	85.65
Ours	Acc. [cm] ↓	<b>1.98</b>	<b>1.71</b>	<b>2.16</b>	<b>4.43</b>	<b>2.57</b>
	Comp. [cm] ↓	<b>1.88</b>	<b>1.83</b>	<b>2.07</b>	<b>3.67</b>	<b>2.36</b>
	Clip sim. ↑	<b>92.48</b>	<b>91.86</b>	<b>90.63</b>	<b>90.28</b>	<b>91.31</b>

Table 2. Per-scene quantitative results on synthetic datasets, including AMA [78] and BANMo [91] datasets. The 3D metrics are obtained by comparing the ground-truth mesh of each frame.

	Acc. [cm] ↓	Comp. [cm] ↓	Clip sim. ↑
w/o diffusion prior	0.99	<b>0.73</b>	82.68
w/o temporal condition	0.92	0.95	86.01
w/o angle weight	0.91	0.86	86.65
w/o depth	3.24	4.90	85.35
w/o canonical space	0.90	1.05	86.14
Full model	<b>0.88</b>	0.78	<b>86.77</b>

Table 3. Ablation studies on real-world datasets: a) w/o diffusion prior: optimize on real-view only; b) w/o depth: optimize with rgb only; c) w/o temporal condition: condition on first frame d) w/o angle weight: using original  $w(i)$ . e) w/o canonical space: model the scene as a unified 4D space-time field in world space.

We further analyze a variant without the temporal condition (second row), where only the first frame is used as the condition for diffusion prior. Results show a degradation in both 3D metrics and CLIP similarity, demonstrating the importance of accurate temporal conditions for estimating the scene geometry when target object is in motion over time. However, as opposed to using a given frame as a direct data-term supervision, the use of SDS with slightly inconsistent view conditions could still achieve realistic scene completion without drastically affecting real-view reconstruction.

Additionally, we explore a variant w/o angle weight, results in Tab. 3 show angle weight can effectively improve the robustness of motion capture by weighting the  $\mathcal{L}_S$  based on the relative angle between real view and sampled view.

**Geometric Initialization.** As in Fig. 5, without the geometric initialization of the canonical SDF field, using SDS loss may not let the model converge to a coherent mesh. Note that here using fixed blob function [55] fails to converge.




Figure 4. Qualitative results of novel view synthesis. Please refer to suppl. material for more quantitative & qualitative comparisons.




Figure 5. Ablation study on geometric initialization. Geometric initialization with fixed blob function [55] fails here.




Figure 6. Ablation study on surface normal smoothness  $\mathcal{L}_{smooth}$ .

**Normal Smoothness.** Fig. 6 shows that applying normal smoothness loss  $\mathcal{L}_{smooth}$  to near-surface points in observation space leads to smoother reconstruction and better motion capture (e.g. hand).

**Effect of Depth Measurement.** As in Tab. 3, although our model achieves competitive CLIP similarity without depth measurement, it fails to capture accurate motion and geometry. We further illustrate the pitfalls of RGB-based methods in Fig. 7. Magic123 [57] is trained on a single frame with a monocular depth prior as a reference. Due to the scale ambiguity, it is very challenging to recover the accurate geometry (e.g., the human arm) for RGB-based methods.

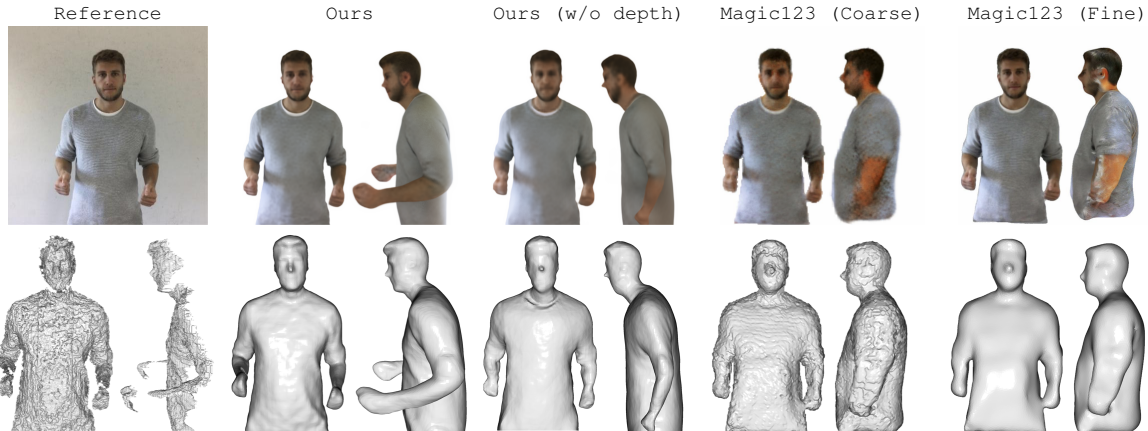


Figure 7. Ablation study: RGB+D vs. RGB-only. We train Magic123 [57] on a single frame for reference. Magic123 (Coarse) is the first stage of training using NeRF. Magic123 (fine) is the second stage of training using DMTet [65] that directly optimizes surface mesh. Due to the scale ambiguity, RGB-based methods fail to capture the accurate pose of the human arm.




Figure 8. Ablation study on canonical regularization: a) *obs. perturb.* w/ deformation field; apply perturbation in observation space. b) *cano. perturb.* w/ deformation field; apply perturbation in canonical space c) *w/o reg.* w/o per-point normal regularization d) *w/o deformation field*; apply perturbation in canonical space.

**Canonical Space Regularization.** We also evaluate the effect of different strategies when applying our regularization losses in Eq. 13. Apart from our proposed canonical space regularization (Eq. 12), we also include two variants *obs. perturb.* and *cano. perturb.* both of which involve the deformation network. The perturbation vector  $\Delta x$  is applied in the observation space and the canonical space respectively. Results in Fig. 8 reveal that both *obs. perturb.* and *cano. perturb.* fails to complete the tail of the cat, whereas our proposed canonical space regularization could enforce the 3D shape in hyper-canonical space to be consistent, leading to better completion of unobserved regions.

**Effect of Representations with a Canonical Shape.** To evaluate the impact of the canonical space, we compare our full model to a canonical space-free representation (w/o canonical space). In the latter, the world space is directly modeled as a 4D space-time grid, which is decomposed into a quad-cube representation with 4 3D hash grids. Without the implicit regularization provided by a canonical space, the quad-cube demonstrates worse temporal consistency in both geometry and appearance reconstruction as in Tab. 3.

**Limitations.** While we demonstrate photo-realistic 360° reconstruction, certain limitations are inherent in our current formulation. Specifically, real-world captures often




Figure 9. Limitations: Our method can fail in challenging situations such as incomplete views, motion blur, and challenging articulated poses of the target object.

present challenges such as incomplete views, motion blur, and intricate articulated poses of the target object (see Fig. 9). Our model, which relies on the diffusion prior conditioned on pure RGB frames, may fail in these scenarios and result in incoherent reconstruction or missing geometry. Moreover, our model cannot reconstruct complex motion in self-occluded regions due to the absence of motion priors. Introducing better diffusion priors (e.g., diffusion models that are conditioned on RGB-D images, time, etc) and motion priors can be a promising avenue for future research.

## 5. Conclusion

We presented MorpheuS, a novel framework designed for dynamic 360° surface reconstruction from a casual monocular RGB-D video. Through the effective integration of the diffusion prior with the dynamic scene reconstruction, MorpheuS goes beyond the conventional dynamic reconstruction methods by achieving both photo-realistic completion in unobserved regions and accurate motion and geometry reconstruction in the observed regions.

**Acknowledgements.** The research presented here has been supported by a sponsored research award from Cisco Research and the UCL Centre for Doctoral Training in Foundational AI under UKRI grant number EP/S021566/1. This project made use of time on Tier 2 HPC facility JADE2, funded by EPSRC (EP/T022205/1).



## References

- [1] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, 2018. 1, 2
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *ACM TOG*. 2005. 1
- [3] Matan Atzmon and Yaron Lipman. Sal: Sign agnostic learning of shapes from raw data. In *CVPR*, 2020. 3
- [4] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural rgb-d surface reconstruction. In *CVPR*, 2022. 2, 4
- [5] Adrien Bartoli, Yan Gérard, Francois Chadebecq, Toby Collins, and Daniel Pizarro. Shape-from-template. *IEEE TPAMI*, 2015. 2
- [6] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *CVPR*. 1999. 1
- [7] Aljaz Bozic, Pablo Palafox, Michael Zollhöfer, Angela Dai, Justus Thies, and Matthias Nießner. Neural non-rigid tracking. *NeurIPS*, 33, 2020. 1, 2
- [8] Aljaz Bozic, Michael Zollhofer, Christian Theobalt, and Matthias Nießner. Deepdeform: Learning non-rigid rgb-d reconstruction with semi-supervised data. In *CVPR*, 2020. 1, 2, 5, 7
- [9] Hongrui Cai, Wanquan Feng, Xuetao Feng, Yan Wang, and Juyong Zhang. Neural surface reconstruction of dynamic scenes with monocular rgb-d camera. In *NeurIPS*, 2022. 2, 3, 5, 6
- [10] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *CVPR*, 2023. 2
- [11] Chen Cao, Qiming Hou, and Kun Zhou. Displaced dynamic expression regression for real-time facial tracking and animation. *ACM TOG*, 2014. 1, 2
- [12] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023. 2
- [13] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Modular interactive video object segmentation: Interaction-to-mask, propagation and difference-aware fusion. In *CVPR*, 2021. 5
- [14] Ho Kei Cheng, Seoung Wug Oh, Brian Price, Joon-Young Lee, and Alexander Schwing. Putting the object back into video object segmentation. *arXiv preprint arXiv:2310.12982*, 2023. 5
- [15] Jaesung Choe, Christopher Choy, Jaesik Park, In So Kweon, and Anima Anandkumar. Spacetime surface regularization for neural dynamic scene reconstruction. In *ICCV*, 2023. 2
- [16] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *SIGGRAPH Asia*, 2022. 2
- [17] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *CVPR*, 2023. 2
- [18] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, 2021. 2
- [19] Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. In *NeurIPS*, 2022. 2, 5, 7
- [20] Ravi Garg, Anastasios Roussos, and Lourdes Agapito. Dense variational reconstruction of non-rigid surfaces from monocular video. In *CVPR*, 2013. 2
- [21] Kaiwen Guo, Feng Xu, Tao Yu, Xiaoyang Liu, Qionghai Dai, and Yebin Liu. Real-time geometry, albedo, and motion reconstruction using a single rgb-d camera. *ACM TOG*, 2017. 2
- [22] Yuan-Chen Guo, Ying-Tian Liu, Ruizhi Shao, Christian Laforte, Vikram Voleti, Guan Luo, Chia-Hao Chen, Zi-Xin Zou, Chen Wang, Yan-Pei Cao, and Song-Hai Zhang. threestudio: A unified framework for 3d content generation. <https://github.com/threestudio-project/threestudio>, 2023. 4
- [23] Ayaan Haque, Matthew Tancik, Alexei Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *ICCV*, 2023. 2
- [24] Mustafa Işık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Nießner. Humanrf: High-fidelity neural radiance fields for humans in motion. *ACM TOG*, 2023. 2
- [25] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. Volumedeform: Real-time volumetric non-rigid reconstruction. In *ECCV*. Springer, 2016. 1, 2
- [26] Yanqin Jiang, Li Zhang, Jin Gao, Weimin Hu, and Yao Yao. Consistent4d: Consistent 360° dynamic object generation from monocular video. *arXiv preprint arXiv:2311.02848*, 2023. 2
- [27] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 1, 2
- [28] Nikos Kolotouros, Georgios Pavlakos, Michael J. Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 1, 2
- [29] Xin Kong, Shikun Liu, Marwan Taher, and Andrew J Davison. vmap: Vectorised object mapping for neural field slam. In *CVPR*, 2023. 2
- [30] Xin Kong, Shikun Liu, Xiaoyang Lyu, Marwan Taher, Xiaojuan Qi, and Andrew J Davison. Eschnet: A generative model for scalable view synthesis. *arXiv preprint arXiv:2402.03908*, 2024. 2
- [31] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM TOG*, 2017. 1, 2
- [32] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 2

- [33] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *CVPR*, 2023. 4
- [34] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snavely. Dynibar: Neural dynamic image-based rendering. In *CVPR*, 2023. 2
- [35] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, 2023. 2, 3
- [36] Shanchuan Lin, Linjie Yang, Imran Saleemi, and Soumyadip Sengupta. Robust high-resolution video matting with temporal guidance. In *WACV*, 2022. 5
- [37] Huan Ling, Seung Wook Kim, Antonio Torralba, Sanja Fidler, and Karsten Kreis. Align your gaussians: Text-to-4d with dynamic 3d gaussians and composed diffusion models. *arXiv preprint arXiv:2312.13763*, 2023. 2
- [38] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *ICCV*, 2023. 2, 4
- [39] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023. 2
- [40] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 2015. 1, 2
- [41] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360° reconstruction of any object from a single image. In *CVPR*, 2023. 2
- [42] Gal Metzger, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *CVPR*, 2023. 2
- [43] Marko Mihajlovic, Sergey Prokudin, Marc Pollefeys, and Siyu Tang. Resfields: Residual neural fields for spatiotemporal signals. In *ICLR*, 2024. 2
- [44] 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 *ECCV*, 2020. 1, 2
- [45] Mirgahney Mohamed and Lourdes Agapito. Dynamicsurf: Dynamic neural rgb-d surface reconstruction with an optimizable feature grid. *arXiv preprint arXiv:2311.08159*, 2023. 2
- [46] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 2022. 3
- [47] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *CVPR*, 2015. 1, 2
- [48] Kemal Egemen Ozden, Kurt Cornelis, Luc Van Eycken, and Luc Van Gool. Reconstructing 3d trajectories of independently moving objects using generic constraints. *Computer Vision and Image Understanding*, 2004. 1
- [49] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. Npms: Neural parametric models for 3d deformable shapes. In *ICCV*, 2021. 1
- [50] Mohit Mendiratta Pan, Mohamed Elgharib, Kartik Teotia, Ayush Tewari, Vladislav Golyanik, Adam Kortylewski, Christian Theobalt, et al. Avatarstudio: Text-driven editing of 3d dynamic human head avatars. *arXiv preprint arXiv:2306.00547*, 2023. 2
- [51] Zijie Pan, Zeyu Yang, Xiatian Zhu, and Li Zhang. Fast dynamic 3d object generation from a single-view video. *arXiv preprint arXiv:2401.08742*, 2024. 2
- [52] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, 2021. 2
- [53] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: a higher-dimensional representation for topologically varying neural radiance fields. *ACM TOG*, 2021. 2, 3, 5
- [54] Ryan Po, Wang Yifan, Vladislav Golyanik, Kfir Aberman, Jonathan T Barron, Amit H Bermano, Eric Ryan Chan, Tali Dekel, Aleksander Holynski, Angjoo Kanazawa, et al. State of the art on diffusion models for visual computing. *arXiv preprint arXiv:2310.07204*, 2023. 2
- [55] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2022. 2, 3, 4, 7
- [56] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 2
- [57] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 2, 7, 8
- [58] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 5
- [59] Jiawei Ren, Liang Pan, Jiaxiang Tang, Chi Zhang, Ang Cao, Gang Zeng, and Ziwei Liu. Dreamgaussian4d: Generative 4d gaussian splatting. *arXiv preprint arXiv:2312.17142*, 2023. 2
- [60] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [61] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM TOG*, 2017. 2
- [62] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 2022. 2

- [63] Ruizhi Shao, Jingxiang Sun, Cheng Peng, Zerong Zheng, Boyao Zhou, Hongwen Zhang, and Yebin Liu. Control4d: Dynamic portrait editing by learning 4d gan from 2d diffusion-based editor. *arXiv preprint arXiv:2305.20082*, 2023. 2
- [64] Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *CVPR*, 2023. 2
- [65] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *NeurIPS*, 2021. 8
- [66] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023. 2
- [67] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 2
- [68] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, et al. Text-to-4d dynamic scene generation. *arXiv preprint arXiv:2301.11280*, 2023. 2
- [69] Miroslava Slavcheva, Maximilian Baust, Daniel Cremers, and Slobodan Ilic. Killingfusion: Non-rigid 3d reconstruction without correspondences. In *CVPR*, 2017. 1, 2, 5, 7
- [70] Miroslava Slavcheva, Maximilian Baust, and Slobodan Ilic. Sobolevfusion: 3d reconstruction of scenes undergoing free non-rigid motion. In *CVPR*, 2018. 1, 2
- [71] Chonghyuk Song, Gengshan Yang, Kangle Deng, Jun-Yan Zhu, and Deva Ramanan. Total-recon: Deformable scene reconstruction for embodied view synthesis. In *ICCV*, 2023. 2
- [72] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *ICCV*, 2021. 2
- [73] Supasorn Suwajanakorn, Ira Kemelmacher-Shlizerman, and Steven M Seitz. Total moving face reconstruction. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*. Springer, 2014. 2
- [74] Jiaxiang Tang. Stable-dreamfusion: Text-to-3d with stable-diffusion, 2022. <https://github.com/ashawkey/stable-dreamfusion>. 4
- [75] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 2
- [76] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Niessner. Face2face: Real-time face capture and reenactment of rgb videos. In *CVPR*, 2016. 1, 2
- [77] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *CVPR*, 2021. 2
- [78] Daniel Vlasic, Ilya Baran, Wojciech Matusik, and Jovan Popović. Articulated mesh animation from multi-view silhouettes. *ACM TOG*, 2008. 5, 7
- [79] Chaoyang Wang, Ben Eckart, Simon Lucey, and Orazio Gallo. Neural trajectory fields for dynamic novel view synthesis. *arXiv preprint arXiv:2105.05994*, 2021. 2
- [80] Chaoyang Wang, Peiye Zhuang, Aliaksandr Siarohin, Junli Cao, Guocheng Qian, Hsin-Ying Lee, and Sergey Tulyakov. Diffusion priors for dynamic view synthesis from monocular videos. *arXiv preprint arXiv:2401.05583*, 2024. 2
- [81] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *CVPR*, 2023. 2
- [82] Hengyi Wang, Jingwen Wang, and Lourdes Agapito. Coslam: Joint coordinate and sparse parametric encodings for neural real-time slam. In *CVPR*, 2023. 2, 4
- [83] Jingwen Wang, Tymoteusz Bleja, and Lourdes Agapito. Go-surf: Neural feature grid optimization for fast, high-fidelity rgb-d surface reconstruction. In *3DV*, 2022. 2, 4
- [84] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 2
- [85] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *ICCV*, 2023. 5
- [86] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv preprint arXiv:2305.16213*, 2023. 2
- [87] Thibaut Weise, Sofien Bouaziz, Hao Li, and Mark Pauly. Realtime performance-based facial animation. *ACM TOG*, 2011. 1, 2
- [88] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *CVPR*, 2022. 2
- [89] Jing Xiao, Jin-xiang Chai, and Takeo Kanade. A closed-form solution to non-rigid shape and motion recovery. In *ECCV*, 2004. 1
- [90] DeJia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In *CVPR*, 2023. 2
- [91] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *CVPR*, 2022. 2, 5, 7
- [92] Gengshan Yang, Chaoyang Wang, N Dinesh Reddy, and Deva Ramanan. Reconstructing animatable categories from videos. In *CVPR*, 2023. 2

- [93] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *NeurIPS*, 2021. [2](#), [3](#)
- [94] Yuyang Yin, Dejia Xu, Zhangyang Wang, Yao Zhao, and Yunchao Wei. 4dgen: Grounded 4d content generation with spatial-temporal consistency. *arXiv preprint arXiv:2312.17225*, 2023. [2](#)
- [95] Rui Yu, Chris Russell, Neill DF Campbell, and Lourdes Agapito. Direct, dense, and deformable: Template-based non-rigid 3d reconstruction from rgb video. In *ICCV*, 2015. [2](#)
- [96] Juyong Zhang, Yuxin Yao, and Bailin Deng. Fast and robust iterative closest point. *IEEE TPAMI*, 2021. [5](#)
- [97] Joseph Zhu and Peiye Zhuang. Hifa: High-fidelity text-to-3d with advanced diffusion guidance. *arXiv preprint arXiv:2305.18766*, 2023. [2](#)
- [98] 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 *CVPR*, 2022. [2](#)
- [99] Michael Zollhöfer, Matthias Nießner, Shahram Izadi, Christoph Rehmann, Christopher Zach, Matthew Fisher, Chenglei Wu, Andrew Fitzgibbon, Charles Loop, Christian Theobalt, et al. Real-time non-rigid reconstruction using an rgb-d camera. *ACM TOG*, 2014. [2](#)
- [100] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *CVPR*, 2017. [1](#), [2](#)