

Time of the Flight of the Gaussians: Optimizing Depth Indirectly in Dynamic Radiance Fields

Runfeng Li Mikhail Okunev Zixuan Guo Anh Ha Duong
Christian Richardt[∞] Matthew O’Toole* James Tompkin

Brown University [∞] Meta Reality Labs * Carnegie Mellon University

Abstract

We present a method to reconstruct dynamic scenes from monocular continuous-wave time-of-flight (C-ToF) cameras using raw sensor samples that achieves similar or better accuracy than neural volumetric approaches and is 100× faster. Quickly achieving high-fidelity dynamic 3D reconstruction from a single viewpoint is a significant challenge in computer vision. In C-ToF radiance field reconstruction, the property of interest—depth—is not directly measured, causing an additional challenge. This problem has a large and underappreciated impact upon the optimization when using a fast primitive-based scene representation like 3D Gaussian splatting, which is commonly used with multi-view data to produce satisfactory results and is brittle in its optimization otherwise. We incorporate two heuristics into the optimization to improve the accuracy of scene geometry represented by Gaussians. Experimental results show that our approach produces accurate reconstructions under constrained C-ToF sensing conditions, including for fast motions like swinging baseball bats.

<https://visual.cs.brown.edu/gftorf>

1. Introduction

Active illumination sensing, like continuous-wave time-of-flight (C-ToF), can help reconstruct dynamic scenes with a single camera thanks to their depth estimates. C-ToF cameras derive depth with simple reconstruction models that assume that all surfaces are opaque and Lambertian, which is fast but can lead to depth errors. Recent optimization-based approaches [1, 26] attempt to resolve the scene with more sophisticated physics that model the emitted light and its reflection from an underlying transmissive 4D volume. While too slow to optimize for practical use, so-called neural ToF radiance fields [1] are still promising because, in principle, they are better at modeling superposition effects from multi-path light transport.

Our work considers two inter-related problems in this setting: how to make these methods faster, and how to make their optimizations more stable once we use faster—and brittle [20]—reconstruction methods. By the end of this paper, we will have accuracy comparable to or better than existing methods while increasing speed by 100×, so making single-camera 4D dynamic scene reconstruction more practical (Fig. 1).

First, we explain why such an optimization might be unstable. For a single camera, even with active illumination, the

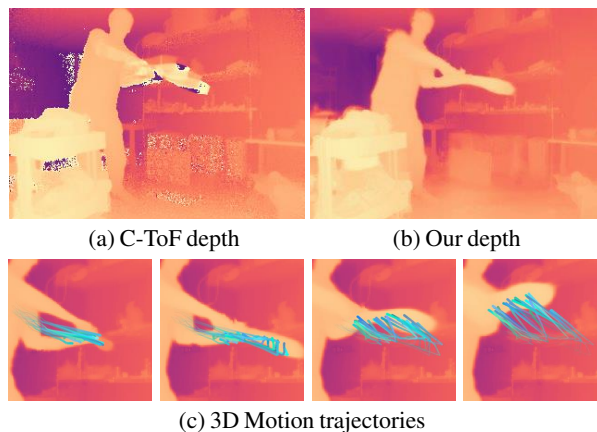


Figure 1. Measuring the depth of a fast-moving object is challenging for C-ToF cameras. By modeling raw C-ToF frames, our method can reconstruct the geometry and motion of a fast swinging baseball bat. It is 100× faster to optimize and render and achieves similar or better accuracy than prior neural volumetric approaches.

accurate reconstruction of depth is not sufficiently constrained by C-ToF sensor measurements under the transmissive image formation model. As depth is only ever indirectly optimized, the scene can have highly inaccurate depth but still produce high-quality reconstructions of sensor measurements (Fig. 2). Thus, as the problem is under-constrained, its optimization is sensitive to its initialization and hyperparameters. Past works have avoided this pernicious problem: TöRF [1] uses additional constraints to localize the depth by both moving the camera and by integrating RGB images. This reduces the problem, but is not suitable for static cameras. F-TöRF [26] provides a more challenging dataset with a static camera and fast-moving objects. But, without the additional constraints, its depth estimates can be worse than the simple C-ToF derived depth even for opaque surfaces. Some works add priors on depth, e.g., using learned single-image depth [17, 38] or priors over scenes [3, 35]. While pragmatic, these approaches are tangential to our goal of attempting to use what sensing we have to accurately measure a scene.

Second, we explain how to speed up the optimization. One way to speed up NeRFs is to use a fast Gaussian splatting (GS) based approach [13, 21, 36]. When there are sufficient cameras to constrain the optimization, GS methods can generally act as drop-in replacements for NeRFs. But, in our setting, we have only a single camera with a dynamic scene, and our desired property of depth is only indirectly optimized by the

All experiments were performed by university authors.



Figure 2. **Fitting C-ToF images \neq fitting depth.** *Top left:* Camera-derived depth from C-ToF. *Top right:* Rendering a GS scene reconstruction into C-ToF raw image samples, then deriving depth. As this is similar to the camera-derived depth to the left, the reconstruction objective was met. *Bottom left:* Rendered mean scene depth from Gaussians, which is highly inaccurate. *Bottom right:* Depth distortion error [10], which measures the Gaussian sparsity along each ray. Gaussians are not well localized.

reconstruction of the sensor measurements. This makes GS methods brittle, and it is difficult to produce accurate depth results for dynamic scenes with C-ToF imaging. While ToF NeRFs are slow, due to their MLP-based scene representation, they are more robust to unfavourable initializations and hyperparameters within this under-constrained setting.

Thus, the goal is to close this gap by making fast GS-based methods more robust in our challenging setting. We show that, with careful tuning of the optimization process, we can reliably achieve convergence in the GS model while maintaining the ability to represent complex geometry and motion. To summarize, our contributions include:

- A dynamic C-ToF GS method with 100× efficiency improvement in both optimization and rendering.
- An investigation into C-ToF radiance field optimization and the empirical biases critical for accurate reconstruction of C-ToF radiance fields.

Assumptions. We assume that the emitter is co-located with the camera, which can cause shadowing errors on close objects. We assume that motion can be modeled as a piecewise linear function between timesteps; higher-order motion models may improve results for rotational motions. We use optical flow estimates within our optimization, which may have errors that propagate to the final results.

2. C-ToF Imaging Principles

Following Okunev et al. [26], C-ToF cameras illuminate the scene with a continuously-modulating amplitude of light, usually as a sinusoid $\sin(2\pi ft)$.¹ As light returns to the camera, it is correlated with a reference sinusoid to produce

¹We can more accurately model light intensity with a non-negative signal $\frac{1}{2} \sin(2\pi ft) + \frac{1}{2}$; for clarity, we choose the simpler model.

an image with pixel intensities $A \sin(\psi + \phi) + B$. Here, A represents the amount of light received at each pixel—the amplitude—and B represents the bias and depends upon the ambient illumination. Phase ψ captures the time that light is in flight. There is also a programmable temporal shift of the reference signal, ϕ . A , B , and ψ are three unknowns, and so we need at least three intensity measurements captured at different offsets ϕ . Many cameras use four offsets for robustness, where $\phi \in \{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$, and the camera produces a quartet of raw frames $\mathbb{Q}_\phi = \{A \sin(\psi + \phi) + B\}$.

Given frames \mathbb{Q}_ϕ , a typical C-ToF camera recovers the phase ψ and computes the distance by multiplying time traveled with the speed of light c [9]:

$$d_{\text{ToF}} = \frac{c}{4\pi f} \psi \quad \text{where} \quad \psi = \arctan \left(\frac{Q_0 - Q_\pi}{Q_{\frac{\pi}{2}} - Q_{\frac{3\pi}{2}}} \right), \quad (1)$$

$$\begin{aligned} &\text{because} \quad \arctan \left(\frac{Q_0 - Q_\pi}{Q_{\frac{\pi}{2}} - Q_{\frac{3\pi}{2}}} \right) = \\ &\arctan \left(\frac{A \sin(\psi) + B - (-A \sin(\psi) + B)}{A \cos(\psi) + B - (-A \cos(\psi) + B)} \right) = \\ &= \arctan(\tan(\psi)) = \psi. \end{aligned}$$

We can represent the C-ToF signal as a complex phasor $a \cdot W(d_{\text{path}}) = a \cdot \exp(i \frac{2\pi d_{\text{path}} f}{c}) = (Q_0 - Q_\pi) + i(Q_{\frac{\pi}{2}} - Q_{\frac{3\pi}{2}})$, where $a = 2A$ and d_{path} is the light path length ($d_{\text{path}} = 2d_{\text{ToF}}$) [8]. Then, the phase $\psi = \angle W(d_{\text{path}})$ and the amplitude $A = \frac{1}{2} |a \cdot W(d_{\text{path}})| = \frac{1}{2} \sqrt{(Q_0 - Q_\pi)^2 + (Q_{\frac{\pi}{2}} - Q_{\frac{3\pi}{2}})^2}$. Furthermore, suppose that s is the emitted signal intensity, r is the reflectivity of a surface in a certain direction, and $1/d_{\text{ToF}}^2$ is the emitted light falloff, then $A = s \cdot r / d_{\text{ToF}}^2$.

C-ToF cameras have limitations. First, they can only measure depth unambiguously up to $d_u = \frac{c}{2f}$ since all depths over this range will map back to $[0, d_u]$ due to sinusoidal periodicity. Real cameras usually have d_u in the range of 5—10 m. Second, as we must capture four frames, achieving 30 Hz depth output requires raw frame capture at 120 Hz, and deriving depth in this way assumes that the quartet \mathbb{Q}_ϕ was captured simultaneously. This means that any motion within the quartet will cause depth errors as the light arriving at a pixel will come from different world points. Third, this model assumes that light reflects back to the sensor from a single surface within a vacuum, when in reality it travels in complex paths.

3. C-ToF GS Image Formation

Next, we explain how we extend Gaussian splatting methods for dynamic scene C-ToF radiance field reconstruction. We assume that the reader is familiar with both Kerbl et al. [13] and monocular 4D extensions for single-camera dynamic scenes, e.g., Yang et al. [36] or Liang et al. [21].

GS image formation model. Briefly, a scene is reconstructed by optimizing a large set of anisotropic Gaussians $\mathcal{G}_k \in \mathcal{G}$. Each is characterized by its center 3D position $\mathbf{x}_k \in \mathbb{R}^3$, covariance matrix Σ_k describing its 3D scale

and rotation, opacity $o_k \in [0, 1]$, and view-dependent color $\mathbf{c}_k \in [0, 1]^3$ parameterized by 16×3 spherical harmonic coefficients. Given a rasterizer to form a Gaussian as its 2D projection \mathcal{G}_k^{2D} on the sensor, the color at pixel \mathbf{x} in an output image is a weighted sum of contributing Gaussians:

$$\mathbf{c}(\mathbf{x}) = \mathbf{c}^{\text{bg}} T_N + \sum_{k=1}^N \mathbf{c}_k o_k \mathcal{G}_k^{2D}(\mathbf{x}) T_k, \quad (2)$$

where $T_k = \prod_{l=1}^{k-1} (1 - o_l \mathcal{G}_l^{2D}(\mathbf{x}))$ and the background signal intensity \mathbf{c}^{bg} is scaled by the final transmittance T_N .

C-ToF GS phasor formation model. Next, we adapt Eq. (2) to model C-ToF signals as phasors, following the image formation model from Attal et al. [1]. Phasor imaging [8] assumes that all \mathbb{Q}_ϕ are captured simultaneously. The phasor at pixel \mathbf{x} is given by:

$$\mathbf{p}(\mathbf{x}) = \mathbf{p}^{\text{bg}} T_N + \sum_{k=1}^N \frac{2sT_k}{d_k^2} W(d_k) o_k \mathcal{G}_k^{2D}(\mathbf{x}) T_k^2. \quad (3)$$

We note the differences from left to right. For the infrared C-ToF signal, rather than representing the scene radiance \mathbf{c}_k towards the camera directly as in Eq. (2), we must describe the emitted and returned light. We model the returned light as the product of the source intensity s and view-dependent surface reflectivity $r_k \in [0, 1]$ modeled with 16 spherical harmonic coefficients as a Gaussian property. This light undergoes an inverse square falloff $1/d_k^2$, where d_k is the distance from each Gaussian center to the camera’s optical center. Then, $W(d_k) = \exp(i\psi_k)$ models the light path importance, where the phase shift between the emitted and reflected light is $\psi_k = \frac{4\pi d_k f}{c}$, where f is the C-ToF camera modulation frequency and c is the speed of light. Finally, transmission $T_k = \prod_{l=1}^{k-1} (1 - o_l \mathcal{G}_l^{2D}(\mathbf{x}))$ is now squared to represent light traveling to and from the scene through the radiance field.

C-ToF GS raw image formation model. Rather than a single phasor assumed to be captured at a single time, Okunev et al. [26] expand this model to consider that the set of raw images \mathbb{Q}_ϕ are captured over time. To begin, we replace the phasor $W(d_k)$ in Eq. (3) with the modulated raw sample quads using $\phi(d_k) : \{\sin(\psi_k), \cos(\psi_k), -\sin(\psi_k), -\cos(\psi_k)\}$:

$$\mathbf{q}(\mathbf{x}) = \mathbf{q}^{\text{bg}} T_N + \sum_{k=1}^N \frac{sT_k}{d_k^2} \phi(d_k) o_k \mathcal{G}_k^{2D}(\mathbf{x}) T_k^2. \quad (4)$$

Given a dynamic scene, this lets us render each raw quad Q at an arbitrary time t . Then, given a real asynchronous \mathbb{Q}_ϕ , we must reconstruct the scene by finding correspondence across time between Q such that we can resolve depth even though the scene is moving. Okunev et al. [26] optimize for this correspondence while also solving for the scene’s geometry.

3.1. Where the Problems Lie

C-ToF samples are not depth. Unlike other problem settings that directly optimize the property of interest, e.g., color images for novel view synthesis or depth maps for depth estimation [5, 13, 17], our challenge is that the optimized property—C-ToF measurements, either as phasors (Eq. (3))

or as raw quads (Eq. (4))—are not the same as the property of interest (depth). Suppose that we compute the depth of the scene as the mean Gaussian depth, as given by:

$$d(\mathbf{x}) = \sum_{k=1}^N d_k o_k \mathcal{G}_k^{2D}(\mathbf{x}) T_k. \quad (5)$$

Without additional constraints, fitting Eq. (3) or Eq. (4) volumetrically does not guarantee consistency with the depth rendered by Eq. (5). For example, when multiple density peaks exist along a ray but only a single surface exists, the rendered mean depth will diverge from the depth derived from the raw quads *even if* the raw quads are accurately reconstructed (Fig. 2). This discrepancy arises because the sine and cosine functions in $\phi(d_k)$ oscillate between convex and concave regions, preventing equality in the finite form of Jensen’s inequality. Only when all Gaussians are clustered at the same z -position along each ray does this depth gap reduce.

Previous approaches. Past works have integrated additional multi-view constraints from moving cameras or extra RGB input [1] to alleviate it. Of course, this works with a Gaussian renderer as well, but such an approach avoids the problem rather than finds a better solution to it. Further, experimentally, previous NeRF-based C-ToF radiance field methods [1, 26] exhibit this gap to a lesser extent, possibly due to the MLP’s inherent bias towards low-entropy (simple) solutions: a single high-density spike (a surface) has lower entropy along the ray than many locations with mid- or low-density values.

In contrast, GS has no such inherent bias; Gaussians can be positioned arbitrarily along the ray to fit the target, which harms depth accuracy. Supposing that we start from randomly initialized Gaussians, the flexibility in Gaussian reflectivities along each ray allows the optimizer to fit C-ToF data by forming high-entropy, multi-peak density distributions. Although this approach still causes a rendering of the scene to match the C-ToF data, the scene diverges from our goal of producing accurate mean depth maps.

Current approaches in RGB scene reconstruction—even in more-constrained multi-view settings—use a depth distortion loss to help push Gaussians towards a surface [10]:

$$\text{DD}(\mathbf{x}) = \sum_{k=1}^N \sum_{l=1}^N \omega_k \omega_l \|d_k - d_l\|^2, \quad (6)$$

where $\omega_k = o_k \mathcal{G}_k^{2D}(\mathbf{x}) T_k$. This implicitly introduces an opaque surface assumption, and often the scenes reconstructed by these methods do conform to this assumption. However, this implicit assumption reduces our ability to model complex light transport. If we only wanted to reconstruct opaque surfaces, we would just derive depth from C-ToF in a closed form. The loss also tends to produce oversmoothed depth reconstructions.

Heuristic 1: Occupancy bias. To address this problem, we note that, within the space of reconstruction parameters and mechanisms when trying to reconstruct scenes (reflectivity, position, opacity, densification), ‘where things are’ or ‘occupancy’ is the primary variance, and so reflectivity

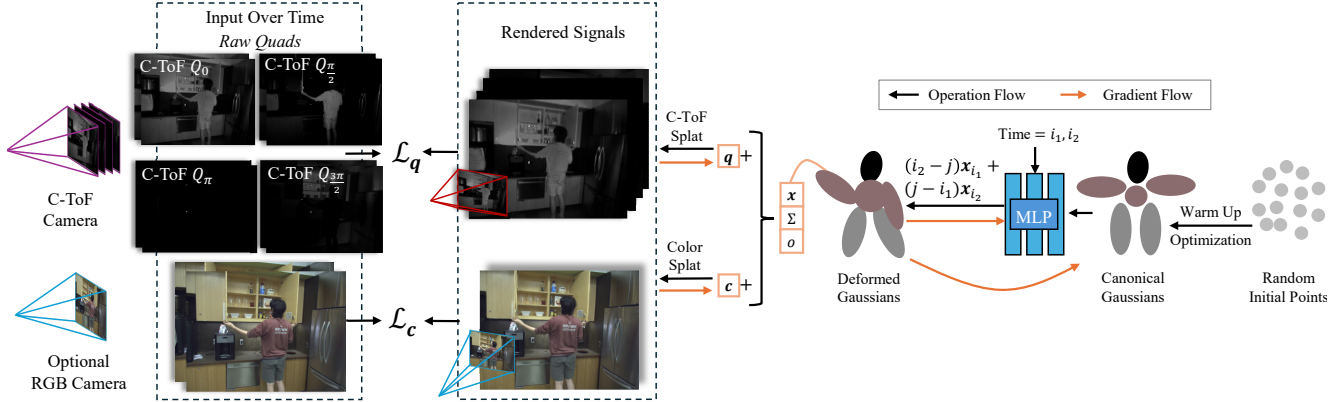


Figure 3. **Pipeline.** *Left:* We capture input raw quads (or phasors, not shown) from a continuous-wave time-of-flight camera with optional color camera. *Right to Left:* From randomly initialized Gaussians, the warm-up stage estimates canonical scene geometry for a static scene. Then, given time t , the MLP predicts offsets $(\delta \mathbf{x}_k)$ that reposition the canonical Gaussians. Then, we render the C-ToF and color images and compute losses.

changes should be rarer than occupancy changes (position, opacity, densification). Given the choice, we would rather move points around, make new density, or remove density than change the proportion of light that is returned. To achieve this, we can decrease the learning rate within the optimization for scene reflectivity $10\times$, which increases the effect of position and opacity variation instead. While simple, this adjustment better encourages Gaussians to conform to the true scene structure from the outset *without* forcing Gaussians to be close to each other as in Eq. (6).

Heuristic 2: Low-reflectivity bias. The initial reflectivity plays a crucial role in the optimization result, particularly for low-reflectivity regions of the scene. Assume that initial reflectivity values are high (e.g., 0.5). Then, with an occupancy bias, due to the inverse-square falloff term, we encourage Gaussians for low-reflectivity regions to be incorrectly placed far away from the camera. Before the reflectivity of these far Gaussians can be corrected, and as Gaussians have spatial extent once splatted, other closer splat-overlapping Gaussians interfere to prevent the reflectivity correction. This leads to multiple opacity peaks.

Now let us assume a lower initial reflectivity (0.01 to 0.1). This enables Gaussians for low-reflectivity regions to be placed more correctly because the lower initialization provides a closer starting fit to the true scene. With an occupancy bias, then Gaussians for high-reflectivity regions will be placed closer to the camera during initial optimization stages also due to the inverse-square falloff term. As they occlude the scene, they are less likely to interfere with other Gaussians, and so have more iterations to optimize their reflectivities and positions to reach the true depth.

Empirically, combining occupancy and low-reflectivity biases is effective for both low and high reflectivity scene regions in our test sequences, whereas only using an occupancy bias (no low-reflectivity bias) harms low-reflectivity regions. While simple, our approach is more effective than Eq. (6) without overly constraining the final result. Our supplemental document and video show evidence for this finding.

4. Pipeline

As input, our system takes a video sequence of raw quartet C-ToF images and an optional color sequence (Fig. 3). Our approach equivalently works with C-ToF phasors. We assume accurate camera poses, which align the scene scale with the C-ToF measurements. For our dynamic GS approach [36], we use an MLP that takes time t and Gaussian 3D positions \mathbf{x}_k as input and outputs Gaussian 3D position offsets $\delta \mathbf{x}_k$. Thus, Gaussians are implicitly deformed from a canonical space.

Local linearity within each four raw frames is critical for fitting as the quartet is asynchronously captured. Let us consider an image in a quartet to be at an ‘integer’ timestep i , and its matching-amplitude raw quad in four ticks time to be at $i + 1$. Then we only deform Gaussians to integer time steps but ensure local linearity at intermediate fractional time steps j by linearly interpolating Gaussian positions from the two nearest integer timesteps i_1 and i_2 :

$$\mathbf{x}_j = (i_2 - j)\mathbf{x}_{i_1} + (j - i_1)\mathbf{x}_{i_2}. \quad (7)$$

This is functionally equivalent to the phase-aware reprojection loss from Okunev et al. [26].

Optical flow weak supervision. To help correspond fast motions, we also weakly supervise the forward and backward 3D motion offsets of Gaussians against estimated optical flow from RAFT [31] with an L2 loss \mathcal{L}_f . Optical flow is projected from the estimated 3D position offsets, where the forward flow is computed as (similarly for backward flow):

$$\mathbf{f}(\mathbf{x}) = \Pi_i \left(\text{sg}(\mathbf{x}^d) + \sum_{k=1}^N \Delta \mathbf{x}_k o_k \mathcal{G}_k^{2D}(\mathbf{x}) T_k \right), \quad (8)$$

where $\Delta \mathbf{x}_k = \text{MLP}(\mathbf{x}, i + 1) - \text{MLP}(\mathbf{x}, i)$ is the Gaussian 3D motion, \mathbf{x}^d is the back-projected 3D point from the rendered mean depth, sg stops gradient computation, and Π_i is the camera’s perspective projection at time i .

Total loss. We combine three objectives:

$$\mathcal{L} = \alpha \mathcal{L}_q + \mathcal{L}_c + \beta \mathcal{L}_f, \quad (9)$$

where \mathcal{L}_q is the raw quad reconstruction loss and \mathcal{L}_c is the optional color reconstruction loss. Both use \mathcal{L}_2 and SSIM.

Initialization. We initialize Gaussians randomly within the camera frustum, bounded by near and far planes, and with a reflectivity of 0.1. Unlike other dynamic GS models, our Gaussian positions are in real-world units, potentially with large values. To ease optimization, we rescale input coordinates to the MLP (only) using the unambiguous depth range, ensuring most values lie within $[0, 1]$. Further, in an initial warm-up stage for 2 K iterations, we assume that the scene is static. This helps to place many of the Gaussians into useful positions.

Random background. For sensitive low-reflectivity scene areas, a low background contribution q^{bg} can lead to errors: Gaussians should be placed on a surface to contribute, but the background is already mostly sufficient to reproduce the C-ToF signal. This creates only low-contributing Gaussians. To fix this, we randomly change the background signal (uniformly between -1 and 1) at each optimization iteration.

MLP. The deformation MLP consists of 8 layers with 256 neurons each, using positional encoding frequencies of $L_x = 10$ and $L_t = 10$. MLP biases are initialized to zero, and weights are initialized with Xavier normal initialization, except for the final linear layer, which maps the final 256 features to the delta position. This layer’s weights are initialized from a normal distribution with a standard deviation of 10^{-5} to ensure stable optimization by starting Gaussian deformations from small δx_k values.

Hyperparameters. In the loss, we set $\alpha = 5$ (1 for synthetic scenes) and $\beta = 0.0008$. We use the Adam optimizer with $\beta = (0.9, 0.999)$ and $\epsilon = 10^{-15}$. We initialize reflectivity as 0.1 and use a reflectivity learning rate of $\gamma_r = 0.00016$. The MLP learning rate starts at $\gamma_{\text{MLP}} = 0.0008$ and decays exponentially to 1.6×10^{-6} over 60 K iterations.

Implementation details. Our code is based on Kerbl et al. [13], with custom CUDA implementations for forward and backward raw C-ToF image rasterization. Details on C-ToF Gaussian gradient calculations are provided in the supplemental material. Gaussians are rendered twice using a differentiable rasterizer: once for the C-ToF view to compute the quad loss, and once for the RGB view to compute the color loss.

5. Experiments

TöRF dataset. The dataset contains five real-world indoor dynamic monocular sequences (*Photocopier*, *Cupboard*, *DeskBox*, *PhoneBooth*, and *StudyBook*). All scenes have raw C-ToF and RGB images with 12 bits of information per pixel. C-ToF signals were captured with a Texas Instruments OPT8241 sensor (320×240 at 30 fps) with 5 m unambiguous depth range. We use the camera poses optimized by the TöRF model. The cameras are moving, which provides additional multi-view constraints, and scene motion is slow.

F-TöRF dataset. The dataset contains five sequences from the dataset: *Pillow*, *Baseball*, *JumpingJacks*, *Target*, and *Fan*. Again, these scenes were captured with a Texas Instruments

OPT8241 sensor (320×240 at 30 fps) with 5 m unambiguous depth range. The scenes are significantly more challenging: they contain no camera motion and they contain fast-moving objects such as falling pillows and swinging baseball bats.

F-TöRF synthetic dataset. We also use the seven synthetic scenes generated in Blender using a physically-based path tracer PBRT [27] adapted for C-ToF emission, at a resolution of 320×240 . The scenes show cubes undergoing axial, lateral, and rotational motions to induce varying disparities, with occlusion, both with and without texture, or with chairs to show thin features. The scenes are strictly monocular without camera motion. The fastest scenes (*3 Cubes Speed Test*, *3 Chairs Speed Test*, *Arcing Cube*, *Axial Speed Test*) have large disparity: maximally up to 30 pixels across raw frames, and often 9–18 pixels (*Sliding Cube*, *Occluded Cube*, *Orthogonal Speed Test*). Three challenging scenes test large axial motion (*Arcing Cube*, *Axial Speed Test*, *Orthogonal Speed Test*).

Metrics. For synthetic scenes, we report depth MSE from the ground truth. For real-world scenes, we only evaluate qualitatively. We report compute times for an NVIDIA 3090 GPU.

Baselines. We compare to the phasor-based C-ToF method TöRF (Attal et al. [1]) that assumes synchronous capture, and to F-TöRF (Okunev et al. [26]) that directly uses asynchronous raw C-ToF measurements. Both are neural volume based. We also compare to DeformableGS (Yang et al. [36]), a dynamic Gaussian splatting reconstruction method, where we add a loss between the C-ToF depth from Eq. (1) and the scene depth from Eq. (5). Lastly, we compare to a 2D baseline that uses softmax splatting (Niklaus and Liu [25]) and RAFT optical flow to warp raw quads to new timesteps.

5.1. Results

Computation time. As expected, our approach is significantly faster than TöRF [1] (60 hours) and F-TöRF [26] (72 hours) at 40–60 minutes per 4D scene reconstruction. Rendering is real time at 100+ Hz.

Quantitative results. Our method is competitive in terms of depth error (Tab. 1). For the reconstructed scene depth d , we show improved results on five out of seven synthetic scenes over all baselines and produce competitive results on the rest. Our predicted scene density is ‘denser’ than that of F-TöRF (*Arcing Cube*), as seen by the reduced gap between d and the depth d_{ToF} derived from the reconstructed sensor images using Eq. (1). Fast-moving thin structures (as in *3 Chairs Speed Test*) seem harder to reconstruct as they require a precise configuration of anisotropic Gaussians. Large-scale nonlinear motions (*Arcing Cube*) are also challenging to optimize under a canonical deformation model. *Orthogonal Speed Test* suffers from a local optimum on one of the cubes, which biases the error metric upward (see the supplemental videos).

Qualitative results. On the easier TöRF dataset that has color and a moving camera (Fig. 6), we show that depth supervision is still critical: DeformableGS without the added

Table 1. **Depth error on the F-TöRF synthetic dataset.** Each number is a depth MSE $\times 100$. d is a standard volume-integrated depth from Eq. (5). d_{ToF} is depth derived from the reconstructed C-ToF raw images using Eq. (1). Bold marks best result for both d and d_{ToF} .

	<i>Sliding Cube</i>	<i>Occluded Cube</i>	<i>Axial ST</i>	<i>3 Cubes ST</i>	<i>3 Chairs ST</i>	<i>Arcing Cube</i>	<i>Ortho. ST</i>
C-ToF	0.096	0.130	2.068	3.131	0.787	36.768	41.931
2D Flow	0.018	0.088	0.662	0.916	0.229	1.678	19.805
F-TöRF d_{ToF}	0.023	0.114	0.251	0.501	0.324	0.470	27.488
Ours d_{ToF}	0.005	0.062	0.123	0.281	0.639	1.060	14.933
TöRF d	0.349	0.388	1.143	2.363	0.956	6.278	22.855
F-TöRF d	0.440	0.647	0.938	1.390	0.855	1.256	7.527
Ours d	0.037	0.369	0.525	0.641	1.023	1.023	22.772




Figure 4. **Our approach is competitive or better in terms of accuracy against the state of the art on synthetic scenes, while being two orders of magnitude faster.** We use the synthetic dataset from F-TöRF to demonstrate this. All images show rendered volumetric depth d . State-of-the-art NeRF model F-TöRF produces similar results to ours while being significantly slower. DeformableGS with C-ToF depth (Eq. (1)) fails to reconstruct dynamic objects well. Baseline 2D Flowed model is fast but cannot handle axial motion as well as our model, producing artifacts. *Inset on ground truth: corresponding RGB image. See our supplement for larger images.*

depth loss fails to reconstruct the scene well. With C-ToF depth, DeformableGS improves substantially, but is still inferior to our physically-based C-ToF supervision. Our method is comparable to TöRF, with improvements over some areas (e.g., dark hair region in *Cupboard*).

The challenging F-TöRF dataset (Figs. 4 and 5) reveals that our method is comparable with F-TöRF, and only these two methods can correct for fast motion artifacts. DeformableGS often completely fails to reconstruct a fast-moving object (*Arcing Cube*, *3 Chairs Speed Test*), and 2D Flowed cannot handle axial motion—something that our method can do. While faster, our method also often produces better static region reconstructions compared to F-TöRF (floor in *JumpingJacks*, *Target*, cart in *Baseball*). But, similar to the

quantitative analysis, our method somewhat struggles with fast-moving thin structures in *3 Chairs Speed Test* and can produce depth that is less sharp than F-TöRF. Finally, no approach can accurately reconstruct spinning fan blades in *Fan*.

6. Discussion

Limitations. *Depth wrapping.* Several of our reconstructions (*JumpingJacks*, *Target*, *Pillow*) contain depth wrapping in the background, which is a typical C-ToF artifact. Although some baselines happen to avoid this issue, this problem is fundamentally ill-posed for a static camera, and predictions heavily depend upon the optimization biases since the true solution is ambiguous.




Figure 5. **F-TöRF real scenes; rendered scene depth d .** As it models temporal dynamics and constrains geometry appropriately during training, our model produces comparable quality reconstructions with F-TöRF while being consistently better than other baselines. Our method tends to reconstruct static geometry better (floor in *JumpingJacks*). *Fan* scene presents a significant challenge as the motion is nonlinear—our canonical field struggles to track this geometry over time. F-TöRF can approximate this scene better due to the lack of such a global constraint. Depth wrapping effects in the background occur in some scenes; these represent an ambiguity that is out of our scope for this work.

Nonlinear motion. Our model can struggle when the motion is strongly nonlinear (as in *Fan* and to a smaller extent in *Arcing Cube*). This can, in principle, be resolved with a more flexible model that allows non-piecewise-linear trajectories, although it will likely make the problem even more ill-posed and require stronger regularization.

Sharp reconstruction. Our volume reconstruction can be visually less sharp than NeRF-based methods around the edges, despite having smaller depth error. This is likely a consequence of the spatial smoothness that Gaussians exhibit. More complex surface representations might help [7, 10].

Thin and fast-moving objects. Our deformation MLP might not maintain thin object structures during large deformation. Since the loss is computed in deformed space but Gaussians are densified in canonical space, this causes unstable densification and makes thin *and* fast-moving objects harder to reconstruct. Longer motion tracks or spacetime rigidity might help [30], and this is left for future work.

No universal guarantees. Our heuristics are based on empirical observations and may not be suitable for different scenes.

Limited indoor scenes. C-ToF imaging may struggle in outdoor scenes, and many motions exist in the real world that go beyond our set of planar, arcing, and axial motions.

7. Related Works

Accurately reconstructing geometry is important for many applications, such as in measurement or physics simulation [34]. When the capturing setup is equipped with a depth sensor, high-quality reconstruction becomes possible even with a small camera baseline. In TöRF [1], a NeRF model takes advantage of a C-ToF camera and uses phasors to optimize the geometry of a dynamic scene. In F-TöRF [26] the model works directly with raw C-ToF camera outputs and recovers motion in addition to geometry, while avoiding typical C-ToF camera motion artifacts. Our work adapts such techniques to fast Gaussian splatting. Neural structured light approaches can also recover geometry, normals, and direct and indirect illumination [24, 29]. For lidar, PlatoNeRF [14] recovers scene geometry from a single view using two-bounce signals captured by a single-photon lidar. Snapshot Lidar [6] and Lee et al. [16] improve C-ToF temporal resolution but they require hardware-level control.

Dynamic Gaussian splatting. Since 3DGS [13] was introduced, it has been explored and adapted widely for reconstruction due to its efficiency. Extending 3DGS, a plethora of works on dynamic Gaussian splatting emerged




Figure 6. **Our approach is competitive with other baselines on TöRF dataset.** Comparison of our method with TöRF and DeformableGS (with and without C-ToF depth prior). Images are novel views of rendered volumetric depth d and color, rendered along a spiral path around the training camera. Our approach demonstrates same or better depth reconstruction for both static and dynamic objects, particularly in dynamic regions with limited multi-view signals. Our method mitigates overfitting to noisy raw depth measurements in the *StudyBook* scene.

recently, differing in their motion representations, constraints and optimization approaches. In Dynamic 3D Gaussians [23], the model assumes a freeform rotation and translation of Gaussians between consecutive time steps and iteratively optimizes the scene with rigidity constraints, but it heavily relies on multi-view input. Some works [21, 33, 36] map each time moment to a canonical space to model the dynamic content. Another approach to model the motion is a global trajectory reconstruction using Fourier, polynomial or a learned basis approximation [12, 15, 22, 32]. In Dynamic Gaussian Marbles [30], isotropic Gaussians equipped with trajectories are learned through a divide-and-conquer optimization using a pretrained single-image depth prior. A local approach is taken in Spacetime Gaussians [18]: the temporal opacity of Gaussians is modeled through radial basis functions making them appear and disappear when needed. Finally, 4D Gaussian splatting [37] avoids using an explicit motion model at all by using 4D Gaussians, where the motion is a byproduct of time-conditioned projection to 3D space.

Gaussian splatting for other input domains. Gaussian splatting models are not restricted to only reconstruct from RGB color inputs. The method can be adapted to a non-RGB setting as well. Some works [2, 4, 11, 19, 28] focus on

extending 3DGS framework to HDR, X-ray, sonar, or thermal images. However, without a sufficient variation in camera poses, high quality reconstruction remains a challenge. To address that, some methods [5] integrate an additional depth input to improve the quality. However, these methods all rely on the quality of the depth priors. Within an under-constrained and more brittle optimization than NeRFs, we contribute a way to effectively use Gaussian splatting with C-ToF imaging.

8. Conclusion

We present a Gaussian splatting approach for fast and accurate reconstruction of dynamic monocular sequences with asynchronous C-ToF exposures. Adapting GS to this setting required us to consider how the underlying optimization affects the indirect estimation of depth, and devise two heuristics that better condition the optimization. These successfully avoid overfitting across a set of synthetic and real-world scenes. In sum, our method is $100\times$ faster than baselines, and produces comparable or better reconstructions of fast-moving objects even with a static monocular camera.

Acknowledgements. RL, MO, ZG, AHD, JT thank NSF CAREER 2144956, NASA RI-80NSSC23M0075, and Cognex. MOT thanks NSF CAREER 2238485.

References

- [1] Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’Toole. TöRF: Time-of-flight radiance fields for dynamic scene view synthesis. In *NeurIPS*, 2021. 1, 3, 5, 6, 7, 8, 14, 15, 16
- [2] Yuanhao Cai, Yixun Liang, Jiahao Wang, Angtian Wang, Yulun Zhang, Xiaokang Yang, Zongwei Zhou, and Alan Yuille. Radiative Gaussian splatting for efficient X-ray novel view synthesis. In *ECCV*, 2024. 8
- [3] David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelSplat: 3D Gaussian splats from image pairs for scalable generalizable 3D reconstruction. In *CVPR*, 2024. 1
- [4] Qian Chen, Shihao Shu, and Xiangzhi Bai. Thermal3D-GS: Physics-induced 3D Gaussians for thermal infrared novel-view synthesis. In *ECCV*, 2024. 8
- [5] Jaeyoung Chung, Jeongtaek Oh, and Kyoung Mu Lee. Depth-regularized optimization for 3D Gaussian splatting in few-shot images. In *CVPR*, pages 811–820, 2024. 3, 8
- [6] Sarah Friday, Yunzi Shi, Yaswanth Cherivirala, Vishwanath Saragadam, and Adithya Pediredla. Snapshot lidar: Fourier embedding of amplitude and phase for single-image depth reconstruction. In *CVPR*, 2024. 7
- [7] Antoine Guédon and Vincent Lepetit. SuGaR: Surface-aligned Gaussian splatting for efficient 3D mesh reconstruction and high-quality mesh rendering. In *CVPR*, 2024. 7
- [8] Mohit Gupta, Shree K Nayar, Matthias B Hullin, and Jaime Martin. Phasor imaging: A generalization of correlation-based time-of-flight imaging. *ACM Transactions on Graphics (ToG)*, 34(5):1–18, 2015. 2, 3
- [9] Miles Hansard, Seungkyu Lee, Ouk Choi, and Radu Horaud. *Time-of-Flight Cameras: Principles, Methods and Applications*. Springer, 2012. 2
- [10] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH*, 2024. 2, 3, 7, 11
- [11] Xin Jin, Pengyi Jiao, Zheng-Peng Duan, Xingchao Yang, Chun-Le Guo, Bo Ren, and Chongyi Li. Lighting every darkness with 3DGS: Fast training and real-time rendering for HDR view synthesis. In *NeurIPS*, 2024. 8
- [12] Kai Katsumata, Duc Minh Vo, and Hideki Nakayama. A compact dynamic 3D Gaussian representation for real-time dynamic view synthesis. In *ECCV*, 2024. 8
- [13] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 1, 2, 3, 5, 7, 11
- [14] Tzofi Klinghoffer, Xiaoyu Xiang, Siddharth Somasundaram, Yuchen Fan, Christian Richardt, Ramesh Raskar, and Rakesh Ranjan. PlatoNeRF: 3D reconstruction in Plato’s cave via single-view two-bounce lidar. In *CVPR*, 2024. 7
- [15] Agelos Kratimenos, Jiahui Lei, and Kostas Daniilidis. DynMF: Neural motion factorization for real-time dynamic view synthesis with 3D Gaussian splatting. In *ECCV*, 2024. 8
- [16] Jongho Lee, Ryan J Suess, and Mohit Gupta. Light-in-flight for a world-in-motion. In *European Conference on Computer Vision*, pages 204–220. Springer, 2024. 7
- [17] Jiahe Li, Jiawei Zhang, Xiao Bai, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. DNGaussian: Optimizing sparse-view 3D Gaussian radiance fields with global-local depth normalization. In *CVPR*, 2024. 1, 3
- [18] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime Gaussian feature splatting for real-time dynamic view synthesis. In *CVPR*, 2024. 8
- [19] Zhihao Li, Yufei Wang, Alex Kot, and Bihan Wen. From chaos to clarity: 3DGS in the dark. In *NeurIPS*, 2024. 8
- [20] Yiqing Liang, Mikhail Okunev, Mikaela Angelina Uy, Runfeng Li, Leonidas Guibas, James Tompkin, and Adam W Harley. Monocular dynamic Gaussian splatting is fast and brittle but smooth motion helps, 2024. arXiv:2412.04457. 1
- [21] Yiqing Liang, Numair Khan, Zhengqin Li, Thu Nguyen-Phuoc, Douglas Lanman, James Tompkin, and Lei Xiao. GauFRE: Gaussian deformation fields for real-time dynamic novel view synthesis. In *WACV*, 2025. 1, 2, 8
- [22] Youtian Lin, Zuozhuo Dai, Siyu Zhu, and Yao Yao. Gaussian-Flow: 4D reconstruction with dynamic 3D Gaussian particle. In *CVPR*, 2023. 8
- [23] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3D Gaussians: Tracking by persistent dynamic view synthesis. In *3DV*, 2024. 8
- [24] Parsa Mirdehghan, Maxx Wu, Wenzheng Chen, David B. Lindell, and Kiriakos N. Kutulakos. Turbosl: Dense accurate and fast 3d by neural inverse structured light. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 25067–25076, 2024. 7
- [25] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *CVPR*, 2020. 5
- [26] Mikhail Okunev, Marc Mapeke, Benjamin Attal, Christian Richardt, Matthew O’Toole, and James Tompkin. Flowed time of flight radiance fields. In *ECCV*, 2024. 1, 2, 3, 4, 5, 6, 7, 14, 15
- [27] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation*. Elsevier, 3rd edition, 2016. 5
- [28] Ziyuan Qu, Omkar Vengurlekar, Mohamad Qadri, Kevin Zhang, Michael Kaess, Christopher Metzler, Suren Jayasuriya, and Adithya Pediredla. Z-splat: Z-axis Gaussian splatting for camera-sonar fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 8
- [29] Aarrushi Shandilya, Benjamin Attal, Christian Richardt, James Tompkin, and Matthew O’Toole. Neural fields for structured lighting. In *ICCV*, 2023. 7
- [30] Colton Stearns, Adam W Harley, Mikaela Uy, Florian Dubost, Federico Tombari, Gordon Wetzstein, and Leonidas Guibas. Dynamic Gaussian marbles for novel view synthesis of casual monocular videos. In *SIGGRAPH Asia*, 2024. 7, 8
- [31] Zachary Teed and Jia Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 4
- [32] Qianqian Wang, Vickie Ye, Hang Gao, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4D reconstruction from a single video. arXiv:2407.13764, 2024. 8
- [33] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4D Gaussian splatting for real-time dynamic scene rendering. In *CVPR*, 2024. 8
- [34] Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. PhysGaussian: Physics-integrated 3D Gaussians for generative dynamics. In *CVPR*, 2024. 7

- [35] Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetzstein. GRM: Large Gaussian reconstruction model for efficient 3D reconstruction and generation. In *ECCV*, 2024. [1](#)
- [36] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3D Gaussians for high-fidelity monocular dynamic scene reconstruction. In *CVPR*, 2024. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#), [14](#), [15](#), [16](#)
- [37] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4D Gaussian splatting. In *ICLR*, 2024. [8](#)
- [38] Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. FSGS: Real-time few-shot view synthesis using Gaussian splatting. In *ECCV*, 2024. [1](#)

Time of the Flight of the Gaussians: Optimizing Depth Indirectly in Dynamic Radiance Fields

Supplementary Material

A. Supplemental Video + Full Figures

We have included a supplemental video showing all dataset sequences plus two ablation sequences. Further, we include full sequence qualitative figures for all three datasets at the back of this supplemental material; this reproduces some results from the main paper but groups and includes all results for completeness.

B. Ablations

We show quantitative (Tab. 2) ablations on F-TöRF synthetic scenes and also qualitative ablation results on F-TöRF real-world and synthetic scenes (Fig. 7). “No bias” refers to no consideration of either of our heuristics. “DD” means adding the depth distortion loss [10], and “H1” and “H2” means our control of the optimization using our occupancy bias (H1) and low initial reflectivity bias (H2). For result variants, ‘ d ’ refers to the mean depth from Eq. (5) in the main paper, and ‘ d_{ToF} ’ refers to the depth derived from the reconstructed quads using Eq. (1) in the main paper.

C. ToF Gradient Computation Details

We briefly explain the opacity gradient computation for our ToF image formation model, which is less trivial compared to other gradients computed via simple chain rules. Let L denote the loss function, and $\alpha_k = o_k \mathcal{G}_k^{2D}(\mathbf{x})$ the opacity term of the k -th Gaussian. Using the chain rule:

$$\frac{\partial L}{\partial \alpha_k} = \frac{\partial L}{\partial \mathbf{c}(\mathbf{x})} \frac{\partial \mathbf{c}(\mathbf{x})}{\partial \alpha_k}, \quad (\text{S10})$$

the term $\frac{\partial \mathbf{c}(\mathbf{x})}{\partial \alpha_k}$ is computed recursively in the original 3D Gaussian Splatting (3DGS) method [13]:

$$\frac{\partial \mathbf{c}(\mathbf{x})}{\partial \alpha_k} = T_k(\mathbf{c}_k - \text{acc}_k), \quad (\text{S11})$$

$$\text{acc}_k = \begin{cases} \alpha_{k+1} \mathbf{c}_{k+1} + (1 - \alpha_{k+1}) \text{acc}_{k+1}, & k < N, \\ 0, & k = N, \end{cases}$$

where T_k is the transmittance, \mathbf{c}_k is the Gaussian’s color, and acc_k aggregates contributions of later Gaussians.

In our model, this computation extends to C-ToF signals like phasor or quad pixels ($\mathbf{p}(\mathbf{x})$, $\mathbf{q}(\mathbf{x})$) as described in Eq. (4) in the main paper. We take $\mathbf{q}(\mathbf{x})$ as an example, the corresponding recursion becomes to:

$$\frac{\partial \mathbf{q}(\mathbf{x})}{\partial \alpha_k} = T_k^2(\mathbf{q}_k + 2(\alpha_k - 1)\text{acc}_k^q), \quad (\text{S12})$$

$$\text{acc}_k^q = \begin{cases} \alpha_{k+1} \mathbf{q}_{k+1} + (1 - \alpha_{k+1})^2 \text{acc}_{k+1}^q, & k < N, \\ 0, & k = N. \end{cases}$$

Here, \mathbf{q}_k represents the quad contribution of the k -th Gaussian, and acc_k^q accumulates later quad contributions.

Table 2. **Ablation depth error on the F-TöRF synthetic dataset.** Each number is a depth $\text{MSE} \times 100$. Bold marks the smallest result for each of d and d_{ToF} . On average, our approach is reliably more accurate without any catastrophic failures, though scene specific variations exist. *Sliding Cube* is too simple to induce large differences. *Ortho. ST* is the only sequence with large low reflectivity areas, and we see the advantage of our low initial reflectivity bias (H2) in the scene’s reconstructed d . Using both heuristics H1, H2 plus the depth distortion (DD) loss theoretically should reduce the gap between d and d_{ToF} , but in practice, because of the spatial extent of the Gaussians, this usually led to oversmooth depth, worse thin structures, and is sometimes unstable because of the existence of large Gaussians representing large homogeneous textureless areas (Fig. 7).

		Mean	<i>Sliding Cube</i>	<i>Occ. Cube</i>	<i>Axial ST</i>	<i>3 Cubes ST</i>	<i>3 Chairs ST</i>	<i>Arcing Cube</i>	<i>Ortho. ST</i>
d_{ToF}	No bias	144.134	0.013	13.227	0.286	1.152	0.885	5.884	987.492
	+DD	146.189	0.119	11.422	5.413	7.289	1.224	13.415	984.440
	+H1	3.641	0.017	0.024	0.206	0.916	0.605	1.113	22.610
	+H1 +H2 (ours)	3.360	0.008	0.073	0.154	0.294	0.650	1.526	20.818
	+H1 +H2 +DD	255.717	0.012	148.578	0.422	1.010	0.895	1537.952	101.148
d	No bias	365.686	0.049	15.604	55.834	24.892	1.482	4.456	2457.487
	+DD	274.316	0.172	11.675	5.922	7.805	1.421	13.067	1880.148
	+H1	44.520	0.058	0.310	0.726	1.610	1.011	1.025	306.898
	+H1 +H2 (ours)	5.824	0.033	0.348	0.730	1.059	1.072	1.479	36.048
	+H1 +H2 +DD	310.564	0.026	147.856	0.787	1.362	1.106	1537.840	484.997




Figure 7. **Ablations.** In the first column, the even rows show the reflectivity map, which is computed as input amplitude multiplied with the square of input depth (light falloff), and can be understood as the expected of Gaussian reflectivity at the corresponding surface. For visualization, the map is clipped to the range $[0, 1]$, where overexposed areas only indicate high target reflectivity. No bias led to arbitrary number of density peaks even for opaque surfaces, thus inaccurate depth from Eq. (5) in the main paper. Adding DD overly stack large Gaussians that led to overly smoothed mean depth. Occupancy bias improves the placement of Gaussians but still struggles in low-reflectivity areas (e.g., bottom left of the cart, and the pillow). Initializing Gaussians with low reflectivity mitigates this issue. Reintroducing the DD loss again after applying the two heuristics still led to oversmooth depth due to Gaussians’ spatial extent.




Figure 8. **Results on all F-TöRF synthetic scenes.** DeformableGS is given C-ToF-derived depth as an additional input. *Inset on ground truth: corresponding RGB image. Full RGB images can be found in Fig. 11.*



Figure 9. Results on all F-TöRF real-world scenes. DeformableGS is given C-ToF-derived depth as an additional input.

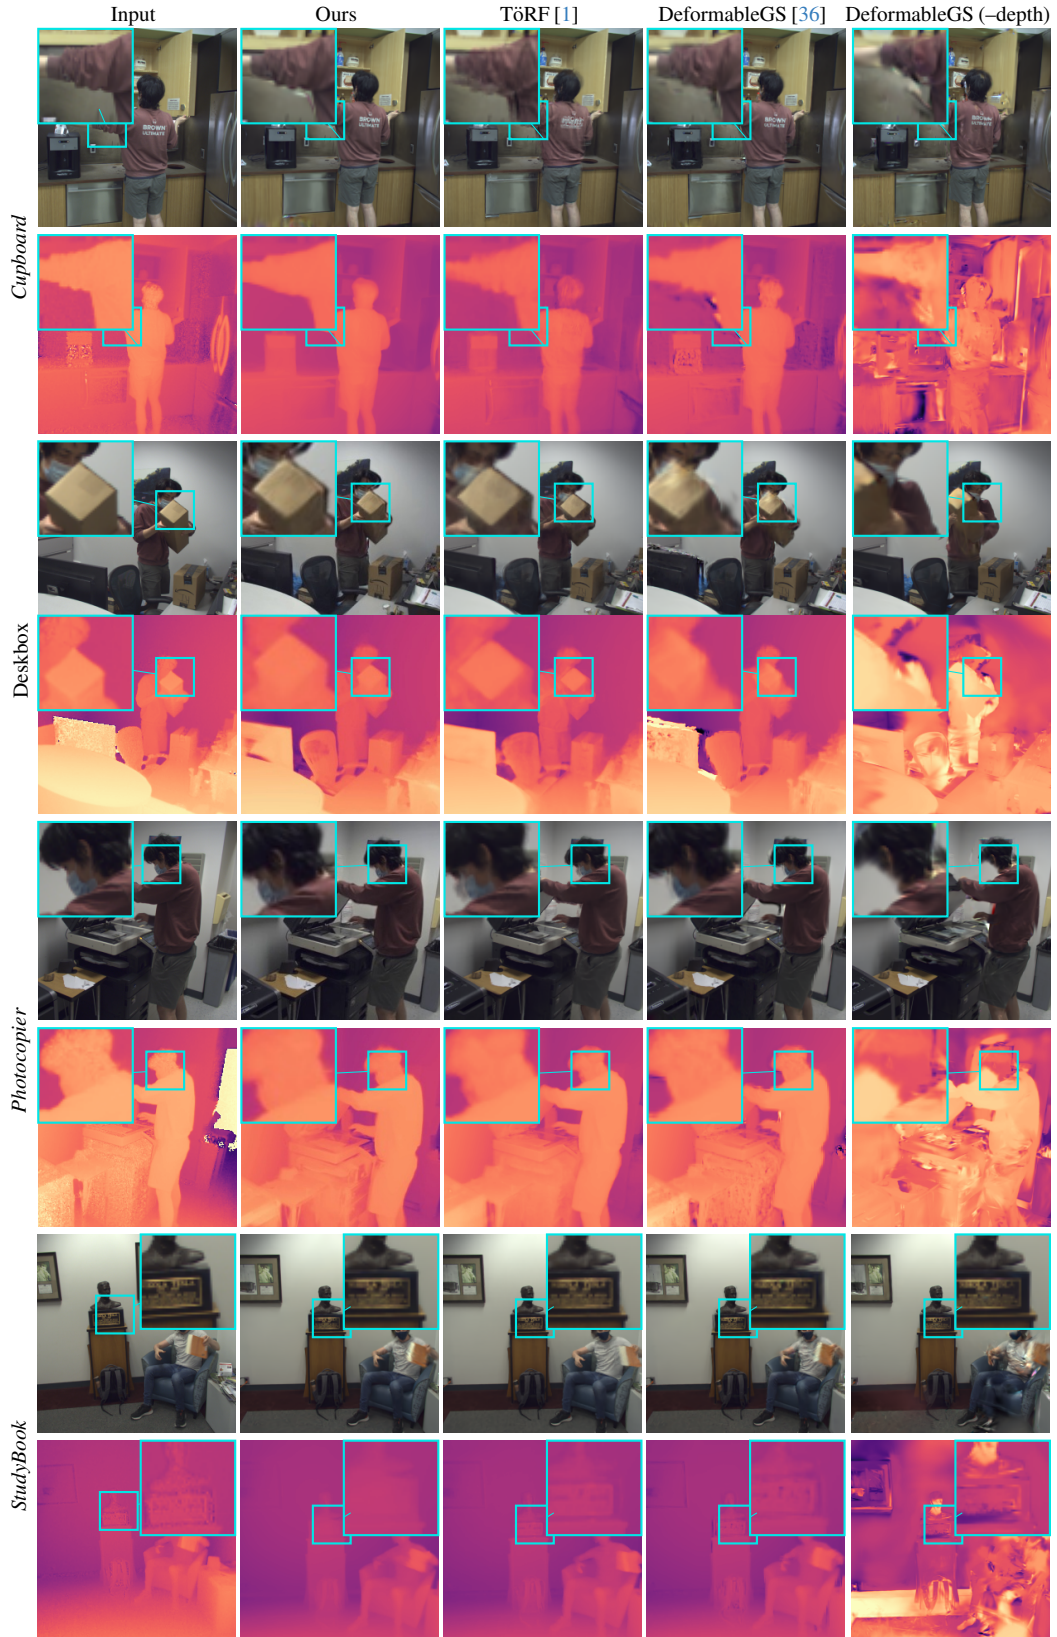




Figure 10. More results on TöRF real-world scenes. DeformableGS is given C-ToF-derived depth as an additional input.




(a) *Sliding Cube*




(b) *Occluded Cube*



(c) *Ortho. ST*



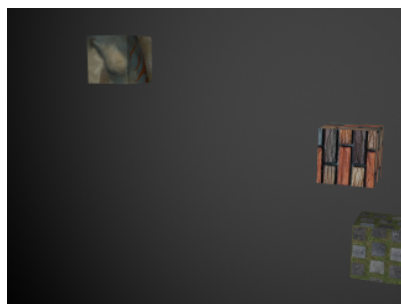
(d) *Arcing Cube*



(e) *Axial ST*



(f) *3 Chairs ST*



(g) *3 Cubes ST*

Figure 11. Corresponding RGB color images for the synthetic scenes for Fig. (4) in the main paper and Fig. 8. Note that *Sliding Cube*, *Occluded Cube*, and *Arcing Cube* have no texture.