

FloatingFusion: Depth from ToF and Image-stabilized Stereo Cameras

Andreas Meuleman¹, Hakyong Kim¹,
James Tompkin², and Min H. Kim¹

¹ KAIST, South Korea

{ameuleman,hkkim,minhkim}@vclab.kaist.ac.kr

² Brown University, United States

Abstract. High-accuracy per-pixel depth is vital for computational photography, so smartphones now have multimodal camera systems with time-of-flight (ToF) depth sensors and multiple color cameras. However, producing accurate high-resolution depth is still challenging due to the low resolution and limited active illumination power of ToF sensors. Fusing RGB stereo and ToF information is a promising direction to overcome these issues, but a key problem remains: to provide high-quality 2D RGB images, the main color sensor’s lens is optically stabilized, resulting in an unknown pose for the floating lens that breaks the geometric relationships between the multimodal image sensors. Leveraging ToF depth estimates and a wide-angle RGB camera, we design an automatic calibration technique based on dense 2D/3D matching that can estimate camera extrinsic, intrinsic, and distortion parameters of a stabilized main RGB sensor from a single snapshot. This lets us fuse stereo and ToF cues via a correlation volume. For fusion, we apply deep learning via a real-world training dataset with depth supervision estimated by a neural reconstruction method. For evaluation, we acquire a test dataset using a commercial high-power depth camera and show that our approach achieves higher accuracy than existing baselines.

Keywords: Online camera calibration, 3D imaging, depth estimation, multi-modal sensor fusion, stereo imaging, time of flight.

1 Introduction

Advances in computational photography allow many applications such as 3D reconstruction [18], view synthesis [23,43], depth-aware image editing [49,53], and augmented reality [20,48]. Vital to these algorithms is *high-accuracy per-pixel depth*, e.g., to integrate virtual objects by backprojecting high-resolution camera color into 3D. To this end, smartphones now have camera systems with multiple sensors, lenses of different focal lengths, and active-illumination time-of-flight (ToF). For instance, correlation-based ToF provides depth by measuring the travel time of infrared active illumination with a gated infrared sensor.

We consider two challenges in providing high-accuracy per-pixel depth: (1) ToF sensor spatial resolution is orders of magnitude less than that of its compatriot color cameras. RGB spatial resolution has increased dramatically on

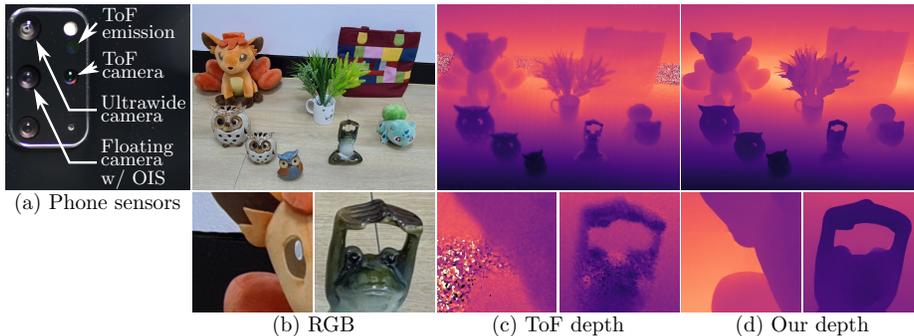


Fig. 1: (a) Multi-modal smartphone imaging. (b) Reference RGB image. (c) ToF depth reprojected to the reference. (d) Our depth from floating fusion.

smartphones—12–64 million pixels is common—whereas ToF is often 0.05–0.3 million pixels. One might correctly think that fusing depth information from ToF with depth information from color camera stereo disparity is a good strategy to increase our depth resolution. Fusion might also help us overcome the low signal-to-noise ratio in ToF signals that arises from the low-intensity illumination of a battery-powered device. For fusion, we need to accurately know the geometric poses of all sensors and lenses in the camera system.

This leads to our second challenge: (2) As RGB spatial resolution has increased, smartphones now use optical image stabilization [40,50]: a *floating* lens compensates for camera body motion to avoid motion blur during exposure. Two low-power actuators suspend the lens body vertically and horizontally to provide a few degrees of in-plane rotation or translation, similar to how a third actuator translates the lens along the optical axis for focus. The magnetic actuation varies with focus and even with the smartphone’s orientation due to gravity, and the pose of the stabilizer is not currently possible to measure or read out electronically. As such, we can only use a fusion strategy if we can automatically optically calibrate the geometry of the floating lens for each exposure taken.

This work proposes a *floating fusion* algorithm to provide high accuracy per pixel depth estimates from an optically-image-stabilized camera, a second RGB camera, and a ToF camera (Fig. 1). We design an online calibration approach for the floating lens that uses ToF measurements and dense optical flow matching between the RGB camera pair. This lets us form 2D/3D correspondences to recover intrinsic, extrinsic, and lens distortion parameters in an absolute manner (not ‘up to scale’), and for every snapshot. This makes it suitable for dynamic environments. Then, to fuse multi-modal sensor information, we build a correlation volume that integrates both ToF and stereo RGB cues, then predict disparity via a learned function. There are few large multi-modal datasets to train this function, and synthetic data creation is expensive and retains a domain gap to the real world. Instead, we capture real-world scenes with multiple views and optimize a neural radiance field [6] with ToF supervision. The resulting depth maps are lower noise and higher detail than those of a depth camera, and pro-

vide us with high-quality training data. For validation, we build a test dataset using a Kinect Azure and show that our method outperforms other traditional and data-driven approaches for snapshot RGB-D imaging.

2 Related Work

ToF and RGB Fusion. Existing data-driven approaches [1,2,36] heavily rely on synthetic data, creating a domain gap. This is exacerbated when using imperfect low-power sensors such as on mobile phones. In addition, current stereo-ToF fusion [14,10,15] typically estimates disparity from stereo and ToF separately before fusion. One approach is to estimate stereo and ToF confidence to merge the disparity maps [33,1,2,37]. In contrast, our ToF estimates are directly incorporated into our disparity pipeline before depth selection. Fusion without stereo [22] tackles more challenging scenarios than direct ToF depth estimation. However, Jung et al.’s downsampling process can blur over occlusion edges, producing incorrect depth at a low resolution that is difficult to fix after reprojection at finer resolutions.

Phone and Multi-Sensor Calibration. DiVerdi and Barron [12] tackle per shot stereo calibration up to scale in the challenging mobile camera environment; however, absolute calibration is critical for stereo/ToF fusion. We leverage coarse ToF depth estimates for absolute stereo calibration. Gil et al. [16] estimate two-view stereo calibration by first estimating a monocular depth map in one image before optimizing the differentiable projective transformation (DPT) that maximizes the consistency between the stereo depth and the monocular depth. The method refines the DPT parameters, handling camera pose shift after factory calibration and improving stereo depth quality, but it still requires the initial transformation to be sufficiently accurate for reasonable stereo depth estimation. In addition, to allow for stable optimization, a lower degree of freedom model is selected, which can neglect camera distortion and lens shift. Works tackling calibration with phone and ToF sensors are not common. Gao et al. [15] use Kinect RGB-D inputs, match RGB to the other camera, use depth to lift points to 3D, then solves a PnP problem to find the transformation. Since it matches sparse keypoints, it is not guaranteed that depth is available where a keypoint is, leading to too few available keypoints. In addition, the method does not account for intrinsic or distortion refinement.

Data-Driven ToF Depth Estimation. Numerous works [44,32,17,3,46,39] attempt to tackle ToF depth estimation via learned approaches. While these approaches have demonstrated strong capabilities in handling challenging artifacts (noise, multi-path interference, or motion), our approach does not strictly require a dedicated method for ToF depth estimation as we directly merge ToF samples in our stereo fusion pipeline.

Conventional Datasets. Accurate real-world datasets with ground-truth depth maps are common for stereo depth estimation [41,34,45]. However, the variety of fusion systems makes it challenging to acquire large-high-quality, real-world

datasets. A majority of ToF-related works leverage rendered data [32,17], particularly for fusion datasets [1,2]. These datasets enable improvement over conventional approaches, but synthesizing RGB and ToF images accurately is challenging. A domain gap is introduced as the noise profile and imaging artifacts are different from the training data. Notable exceptions are Son et al. [44], and Gao and Fan et al. [39], where an accurate depth camera provides training data for a lower-quality ToF module. The acquisition is partially automated thanks to a robotic arm. However, this bulky setup limits the variety of the scenes: all scenes are captured on the same table, with similar backgrounds across the dataset. In addition, the use of a single depth camera at a different location from the ToF module introduces occlusion, with some areas in the ToF image having no supervision. In addition, this method only tackles ToF depth estimation, and the dataset does not feature RGB images.

Multiview Geometry Estimation. Several approaches are capable of accurate depth estimation from multiview images [42], even in dynamic environments [26,31,24]. Despite their accuracy, including ToF data to these approaches is not obvious. Scene representations optimized from a set of images [6,21,4] have recently shown good novel view synthesis and scene geometry reconstruction, including to refine depth estimates in the context of multiview stereo [51]. Since the optimization can accept supervision from varied sources, including ToF measurements is straightforward. For this reason, we select a state-of-the-art neural representation that has the advantage to handle heterogeneous resolutions [6] for our training data generation. TöRF [5] renders phasor images from a volume representation to optimize raw ToF image reconstruction. While efficiently improving NeRF’s results and tackling ToF phase wrapping, this approach is not necessary for our context as our device is not prone to phase wrapping due to its low illumination range (low power) and thanks to the use of several modulation frequencies. We also observe that, in the absence of explicit ToF confidence, erroneous ToF measurements tend to be more present in depth maps rendered from a TöRF. Finally, approaches based on ICP registration [18] cannot be applied directly to our data since depth maps from the low-power ToF module are too noisy to be registered through ICP.

3 Method

We use an off-the-shelf Samsung Galaxy S20+ smartphone. This has the main camera with a 12MP color sensor and a magnetic mount 79° lens for stabilization and focusing, a secondary 12MP color camera with a fixed ultrawide 120° lens, and a 0.3MP ToF system with an infrared fixed 78° lens camera and infrared emitter (Fig. 1a). As the ultrawide camera and the ToF module are rigidly fixed, we calibrate their intrinsics $K_{\text{UW}}, K_{\text{ToF}}$, extrinsics $[R|t]_{\text{UW}}, [R|t]_{\text{ToF}}$, and lens distortion parameters using an offline method based on checkerboard corner estimation. We use a checkerboard with similar absorption in the visible spectrum as in infrared. However, calibrating the floating main camera (subscript FM) is not possible offline, as its pose changes from snapshot to snapshot. OIS intro-

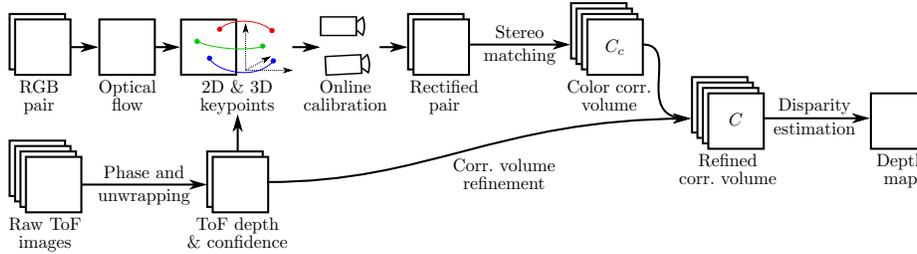


Fig. 2: Overview of our method. The first step is to estimate ToF depth via phase unwrapping (§ 3.1). Then, and after dense matching between the two RGB cameras, we use ToF depth to estimate the floating camera’s intrinsic and extrinsic parameters (§ 3.2). Once the stereo pair is calibrated and rectified, we extract features in each image to build a correlation volume C_c . This volume is refined using the ToF samples (§ 3.3) before the final disparity estimation.

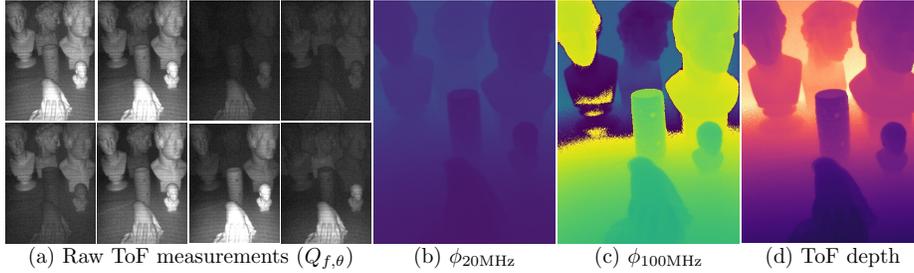


Fig. 3: ToF depth estimation from raw measurements. From raw ToF images for two different frequencies (a), we estimate a coarse but unwrapped phase map (b) and a finer but wrapped phase map (c). By unwrapping $\phi_{100\text{MHz}}$ using the lower frequency phase $\phi_{20\text{MHz}}$, we estimate a more accurate depth map.

duces lens shift in x, y for stabilization and in z for focus, with the z direction inducing additional lens distortion changes. The lens also tilts (pitch/yaw rotations) depending on the phone’s orientation because of gravity. As such, we must estimate per snapshot a new intrinsic matrix K_{FM} , new extrinsic matrix $[R|t]_{\text{FM}}$, and three radial and two tangential distortion coefficients $\{k_1, k_2, k_3, p_1, p_2\}_{\text{FM}}$ from the Brown-Conrady model [9,7] for the main floating camera. To tackle this challenge, we present a method to estimate these parameters at an absolute scale (not relative or ‘up to scale’).

3.1 ToF Depth Estimation

The ToF system modulates its infra-red light source by 20MHz and 100MHz square waves, and alternates between both frequencies sequentially. The sensor captures four shots per frequency, with the same modulation as the light

source shifted by 90 degrees. With two frequencies, we obtain eight raw ToF measurements per snapshot:

$$Q_{f,\theta}, \quad f \in \{20\text{MHz}, 100\text{MHz}\}, \quad \theta \in \{0, \pi/2, \pi, 3\pi/2\}, \quad (1)$$

where f is the modulation frequency and θ is the sensor phase shift. Figure 3(a) shows an example of captured ToF measurements. To estimate depth, the first step is to estimate phase (Figures 3(b) and 3(c)) for each modulation frequency:

$$\phi_f = \arctan2(Q_{f,\pi} - Q_{f,0}, Q_{f,3\pi/2} - Q_{f,\pi/2}). \quad (2)$$

From these phases, we can estimate the distance:

$$d_f = \frac{c}{4\pi f} \phi_f + k_f \frac{c}{2f}, \quad k_f \in \mathbb{N}, \quad (3)$$

with c the speed of light. This equation shows that, given a phase, the depth is known up to a $k_f \frac{c}{2f}$ shift. That is, we observe phase wrapping. The ToF depth variance is inversely proportional to the modulation frequency [25]. Therefore, the higher frequency tends to be more accurate; however, the wrapping range is shorter ($\approx 1.5\text{m}$ for 100MHz against $\approx 7.5\text{m}$ for 20MHz). Since the ToF illumination is low power, the signal becomes weak for further objects, making the phase estimation highly unreliable beyond the wrapping point of the lower frequency. Based on this, we assume that $\phi_{20\text{MHz}}$ does not show phase wrapping: $k_{20\text{MHz}} := 0$. Therefore, we unwrap $\phi_{100\text{MHz}}$ using $\phi_{20\text{MHz}}$ to benefit from the lower depth variance associated with a higher frequency without the phase wrapping ambiguity. In detail, we find $\hat{k}_{100\text{MHz}}$ that minimizes the depth difference between the two frequencies:

$$\hat{k}_{100\text{MHz}} = \underset{k}{\operatorname{argmin}} \left| d_{20\text{MHz}} - \frac{c}{4\pi \cdot 10^8} \phi_{100\text{MHz}} + k \frac{c}{2 \cdot 10^8} \right|. \quad (4)$$

From this, we can compute the distance $d_{100\text{MHz}}$, which will be used to obtain the depth d_{ToF} (Figure 3d). For more details on ToF, refer to [19,25].

With the estimated depth, we assign a confidence map ω based on the signal's amplitude, the concordance between $d_{20\text{MHz}}$ and $d_{100\text{MHz}}$ and local depth changes. First, we assign lower scores when the signal is weak:

$$\omega_A = \exp \left(-1 / \sum_f A_f / (2\sigma_A^2) \right), \quad (5)$$

With $A_f = \sqrt{(Q_{f,0} - Q_{f,\pi/2})^2 + (Q_{f,\pi} - Q_{f,3\pi/2})^2} / 2$ and $f \in \{20\text{MHz}, 100\text{MHz}\}$. We also take into account the difference between the estimated distance from the two frequencies:

$$\omega_d = \exp \left(-|d_{20\text{MHz}} - d_{100\text{MHz}}|^2 / (2\sigma_d^2) \right). \quad (6)$$

Since ToF is less reliable at depth discontinuities, we deem areas with large depth gradient to be less reliable:

$$\omega_{\nabla} = \exp\left(-\|\nabla d_{\text{ToF}}\|^2 / (2\sigma_{\nabla}^2)\right) \cdot \exp\left(-\|\nabla(1/d_{\text{ToF}})\|^2 / (2\sigma_{\nabla}^2)\right). \quad (7)$$

The complete confidence is: $\omega = \omega_A \omega_d \omega_{\nabla}$. For our experiments, we set $\sigma_A = 20$, $\sigma_d = 0.05$, and $\sigma_{\nabla} = 0.005$.

3.2 Online Calibration

To calibrate our floating main camera, we need to find sufficient correspondences between known 3D world points and projections of those points in 2D. We must find a way to correspond our 3D points from ToF with the main camera even though it does not share a spectral response. Our overall strategy is to use the known fixed relationship between the ToF and ultrawide cameras and additionally exploit 2D color correspondences from the optical flow between the ultrawide and floating cameras. In this way, we can map from main camera 2D coordinates to ultrawide camera 2D coordinates to ToF camera 3D coordinates. While ToF can be noisy, it still provides sufficient points to robustly calibrate all intrinsic, extrinsic, and lens distortion parameters of the main camera.

From ToF to Ultrawide. The first step is to reproject the ToF depth estimates d_{ToF} to the ultrawide camera. We transform the depth map to a point cloud P :

$$P_{\text{ToF}} = K_{\text{ToF}}^{-1}[u, v, d_{\text{ToF}}]^{\top}, \quad (8)$$

where K_{ToF} is the known ToF camera matrix, and (u, v) are pixel coordinates with corresponding depth d_{ToF} . Then, the point cloud can be transformed to the ultrawide camera’s space:

$$P_{\text{UW}} = [R|t]_{\text{ToF} \rightarrow \text{UW}} [P_{\text{ToF}}^{\top} | 1]^{\top}, \quad (9)$$

where $[R|t]_{\text{ToF} \rightarrow \text{UW}}$ is the relative transformation from the ToF camera space to the ultrawide camera space. From point cloud P_{UW} , we obtain the pixel coordinate $[u_{\text{UW}}, v_{\text{UW}}]$ of the ToF point cloud reprojected to the ultrawide camera:

$$d_{\text{UW}}[u_{\text{UW}}, v_{\text{UW}}, 1] = K_{\text{UW}} P_{\text{UW}}. \quad (10)$$

The reprojected ToF points P_{UW} and their subpixel coordinates $[u_{\text{UW}}, v_{\text{UW}}]$ will be used to estimate calibration for our main camera in a later stage.

From Ultrawide to Floating Main. Next, we match the ultrawide camera to the floating main camera to be accurately calibrated. Since both cameras are located near to each other, they share a similar point of view, thus making sparse scale and rotation invariant feature matching unnecessary. As such, we use dense optical flow [47] to find correspondences. To use flow, we first undistort the ultrawide camera image given its calibration, then rectify it *approximately* to the floating main camera given an initial *approximate* offline

calibration. This calibration will be wrong, but as flow is designed for small unconstrained image-to-image correspondence, this rectification approach will still find useful correspondences. As output, we receive a 2D vector field $\mathcal{F}_{\text{UW} \rightarrow \text{FM}}$.

From ToF to Floating Main. We can now use the optical flow to form 2D/3D matches. For each ToF point reprojected to the ultrawide camera P_{UW} , we find the corresponding pixel in the floating camera:

$$[u'_{\text{FM}}, v'_{\text{FM}}] = [u_{\text{UW}}, v_{\text{UW}}] + \mathcal{F}_{\text{UW} \rightarrow \text{FM}}([u_{\text{UW}}, v_{\text{UW}}]). \quad (11)$$

We sample the flow using bilinear interpolation since $[u_{\text{UW}}, v_{\text{UW}}]$ are estimated with subpixel precision.

From this 2D/3D matching between the 2D points in the floating camera $[u'_{\text{FM}}, v'_{\text{FM}}]$ and the 3D points in the ultrawide camera space P_{UW} , we can estimate the floating camera’s calibration. We first solve the optimization through RANSAC for outlier removal, followed by Levenberg-Marquardt optimization, obtaining the transformation between the two RGB cameras $[R, t]_{\text{FM} \rightarrow \text{UW}}$, as well as the camera matrix K_{FM} and its distortion coefficients. Once the calibration is achieved, we rectify the two RGB images, enabling stereo/ToF fusion.

Discussion. Both Equations (10) and (11) are not occlusion-aware but, due to the small baseline, occluded world points are only a small portion of the total number of matched points. RANSAC helps us to avoid outlier correspondences from occlusion, incorrect flow estimates, and noisy ToF estimates for calibration.

Further, while we rely on a fixed RGB camera and ToF module, extending the approach to scenarios without fixed cameras is possible. Reliable feature matching between spectral domains has been demonstrated [13,55,8] as well as RGB/IR optical flow [38]. Using matching between the ToF module’s IR camera and the RGB cameras, calibration can likely be achieved even if no second RGB camera exists as fixed with respect to the ToF module.

3.3 Fusing ToF and Stereo

Given the now-calibrated color stereo pair, and the ToF depth samples, we will fuse these into an accurate high-resolution depth map for a color camera. The first step is to build a correlation volume \mathcal{C}_c from our RGB pair. A point in the volume \mathcal{C}_c at coordinate $[u, v, u']$ represents the correlation between a pixel $[u, v]$ in the reference image and a pixel $[u', v]$ in the target image at some disparity. Thus, the correlation volume’s shape is (width×height×width) since disparity is horizontal along the width direction. We compute correlation volume values by extracting 256-dim. image features from each view using RAFT-stereo [28]’s learned feature encoder, then by taking the dot product of the feature vectors from each RGB camera for each disparity amount. Note that the feature extraction process downsamples the images four times and a disparity map at original resolution is recovered through RAFT-stereo’s convex upsampling. For ease of evaluation, we use the fixed ultrawide camera as a reference, although the floating camera can be chosen without algorithmic change.

A 3D world point P_{ToF} corresponds to a 2D sensor coordinate $[u_{\text{UW}}, v_{\text{UW}}]$ in the ultrawide camera and in the floating main camera $[u_{\text{FM}}, v_{\text{FM}}]$ (see reprojection Equations (8),(9), and (10)). A ToF point also has a confidence ω estimated along with the ToF depth maps. We leverage those points by increasing the correlation at the location of the sample points in the volume. For a point, the corresponding location in the volume is $[u_{\text{UW}}, v_{\text{UW}}, u_{\text{FM}}]$. Coordinates given by a ToF point $[u_{\text{UW}}, v_{\text{UW}}, u_{\text{FM}}]$ have eight integer neighbors with defined values in the correlation volume. We inject the ToF point into RAFT-stereo’s correlation volume using linear-like weights:

$$\begin{aligned} \mathcal{C}([u_{\text{UW}}], [v_{\text{UW}}], [u_{\text{FM}}]) &= \mathcal{C}_c([u_{\text{UW}}], [v_{\text{UW}}], [u_{\text{FM}}]) + \\ &\quad \tau \cdot \omega (([u_{\text{UW}}] - u_{\text{UW}})([v_{\text{UW}}] - v_{\text{UW}})([u_{\text{FM}}] - u_{\text{FM}})). \end{aligned} \quad (12)$$

$[u]$ is the nearest integer $\leq u$, and $\lceil u \rceil$ is the nearest integer $\geq u$. We perform the same operation for all eight integer points with coordinates around $[u, v, u']$. If several ToF points affect the same point in the correlation, their contributions accumulate. τ is a scalar parameter that is optimized during the training process. We then use the updated correlation volume \mathcal{C} with the next steps of RAFT-stereo’s pipeline to estimate disparity. Thanks to this approach, we efficiently and robustly combine stereo correlation cues and ToF measurements before estimating the depth map.

3.4 Dataset Generation

Background. Learning-based stereo/ToF fusion methods require training data. For this, we optimize neural radiance fields [29,35,30] from RGB images. These allow querying density and appearance for every point in the scene, allowing us to render estimated depth maps. A typical volume rendering is as follows:

$$\hat{C}(r) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) c_i, \quad (13)$$

where $\hat{C}(r)$ is the rendered color for the ray r , σ_i and c_i are the density and color of the representation at i point. N points are sampled along the ray r , where δ_i is the distance between neighboring samples, and $T_i = \exp(-\sum_{i' < i} \sigma_{i'} \delta_{i'})$ is the approximated transparency between the ray origin and the sample.

This rendering is differentiable on the density and color of the sample points, allowing gradient-based optimization. To render depth maps from these representations, we swap the color term c_i in Equation (13) for the depth of the point w.r.t. the ray origin.

Design Choices and Our Approach. Despite improving on numerous aspects such as optimization time and novel view quality over neural representations, Plenoxels [4] produces fuzzy depth maps (see the supplemental material). We also observe that NerfingMVS’s [51] guided NeRF optimization based

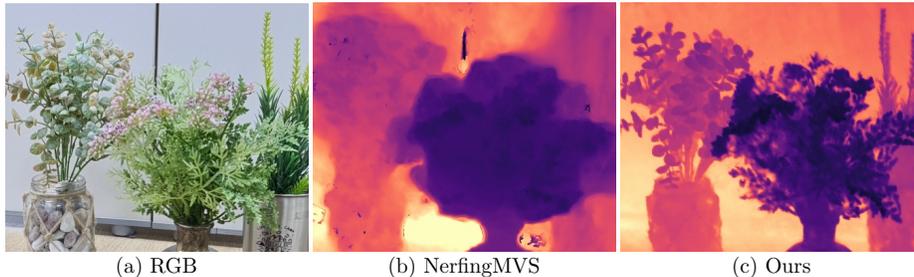


Fig. 4: NerfingMVS [51]’s results before filtering show more artifacts than ours.

on reducing the sampling range around the expected depth tends to create artifacts (Figure 4) that require an additional filtering step [48] for accurate depth maps. Since filtering at novel views is impossible, this approach is not suitable for creating training data for unseen views. Thus, we use MipNeRF [6] as a basis for our multiview depth estimation pipeline. This can naturally handle the resolution difference between the ToF and RGB cameras. For depth supervision [11], we use a straightforward approach similar to VideoNeRF [52]’s supervision on inverse depth maps computed from multi-view stereo: we add a loss using depth samples to the optimization:

$$\mathcal{L}_{depth} = \omega \|d_{rendered} - d_{ToF}\|_2^2, \quad (14)$$

with the ToF confidence ω (Section 3.1). In addition to this supervision, we implement extrinsic and intrinsic refinement [21] following a 6D continuous rotation representation [54]. Note that, since Mip-NeRF prevents too high frequencies in positional encoding and poses and that the intrinsics are well initialized, a coarse-to-fine positional encoding [27] is not required. Our training and validation datasets have eight scenes with around 100 snapshots per scene. The scenes feature varied depth range, background, objects, and materials.

4 Results

4.1 Evaluation Dataset

To evaluate our method, we build a real-world dataset with ground-truth depth obtained using a Kinect Azure. Since the depth camera is higher power, noise is reduced, and depth quality is much better than our phone’s ToF module. After securing the phone and the Azure Kinect on a joint mount, we calibrate the phone’s ToF module and the ultrawide camera w.r.t. the depth camera. Once the calibration is estimated, depth maps can be reprojected to the ultrawide and ToF cameras for comparison. We capture four scenes for a total of 200 snapshots.

In addition to this RGB-D dataset, we calibrate the floating camera using the conventional multi-shot offline pipeline with chessboards. This provides ground truth for our online calibration evaluation on four scenes.

Table 1: We compare our multiview fusion against a ToF-supervised scene representation [5] and multiview-stereo approaches [51,24,31]. Original Mip-NeRF [6] on which our implementation is based is given for reference.

	Bad ratio (%)		Depth error			
	>0.2	>0.05	MAE	Rel. RMSE	MAE	RMSE
TöRF [5]	1.82	26.48	0.055	0.075	0.041	0.062
NerfingMVS [51]	3.37	20.11	0.047	0.069	0.039	0.071
RCVD [24]	35.12	75.68	0.249	0.326	0.208	0.315
CVD [31]	1.68	11.67	0.033	0.056	0.028	0.056
Mip-NeRF [6]	17.47	61.71	0.156	0.199	0.115	0.170
Ours	0.81	7.07	0.028	0.047	0.022	0.044

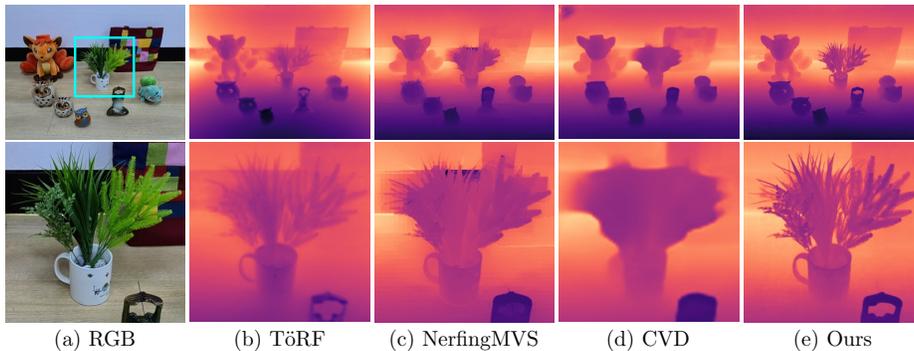


Fig. 5: Multiview depth estimation with ToF-supervised scene representation TöRF [5] and the multiview-stereo approaches NerfingMVS [51] and CVD [31].

4.2 Depth Estimation for Training

Figure 5 shows that our method can preserve thin structures, and Table 1 confirms that our approach can efficiently merge ToF and stereo data from multiple views. While TöRF [5] is designed to handle ToF inputs, it performs worse than our method when ToF and RGB resolutions differ. We run CVD and RCVD at their default resolution since we observed degradation in accuracy when increasing image size. Note that progressive geometry integration methods such as [18] fail on our data: the raw ToF depth maps from our smartphone are not accurate enough for ICP registration.

4.3 Snapshot RGB-D Imaging

Calibration. Table 2 details the accuracy of our calibration method. Gao et al. [15] calibrate two GoPro cameras w.r.t. a Kinect RGB-D camera. In our comparison, we substitute the RGB-D camera by the ultrawide RGB plus the ToF module, and substitute the GoPro camera by the main camera. Gao et al. [15]’s

Table 2: Our online calibration versus Gao et al. [15] on our real dataset.

	Pose error (mm)		Rotation error (deg.)	
	MAE	RMSE	MAE	RMSE
Gao et al. [15]	5.664	5.997	1.292	1.431
[15] + DFM [13]	5.174	5.803	1.344	1.499
Ours	3.346	3.989	1.264	1.411

Table 3: Fusion evaluation. In the first rows, we evaluate other approaches for RGB/ToF fusions. Since their stereo matching method are not robust against noise and imaging artifacts, their fusion results are highly inaccurate. In the next rows, we replace their less robust stereo matching with a state-of-the-art method [28]. We use our calibration for all methods, except for “Ours (ignoring OIS)” to highlight the importance of our per-shot calibration. The results show that our fusion approach outperforms existing methods.

	Bad ratio (%)		Depth error			
	>0.2	>0.05	MAE	Rel. RMSE	Rel. MAE	RMSE
Marin et al. [33]	61.25	96.80	0.545	0.708	0.410	0.610
Agresti et al. [1,2]	91.29	98.29	0.962	1.172	0.722	1.012
Gao et al. [15]	14.95	40.82	2.262	6.493	1.394	5.066
[33] (RAFT-stereo [28])	2.01	13.87	0.043	0.092	0.034	0.085
[1,2] (RAFT-stereo [28])	1.48	13.96	0.037	0.065	0.031	0.065
[15] (RAFT-stereo [28])	1.67	7.71	0.031	1.021	0.026	0.854
Stereo only [28]	2.36	14.18	0.041	1.300	0.035	1.078
Ours (ignoring OIS)	9.06	29.54	0.082	0.163	0.073	0.164
Ours	1.40	7.17	0.028	0.050	0.024	0.051

calibration shows much lower accuracy than ours, even when the method is paired with a state-of-the-art feature matcher [13]. In addition, only our method is able to refine the camera matrix and distortion parameters.

Stereo/ToF Fusion. We evaluate our fusion approach against our real-world RGB-D dataset. For comparison, we implement [33,1,15] and we train Agresti et al.’s method [1] using their rendered SYNTH3 dataset. Since Gao et al. [15]’s calibration is too inaccurate, rectification fails severely on some snapshots. Total calibration failure occurred for 34 of the 200 snapshots (17%) in our test dataset. We show examples of poorly rectified stereo pairs in the supplemental material.

Thus, evaluate all methods using our calibration. We also evaluate if we can ignore OIS: we calibrate the main camera using a checkerboard while the phone is fixed, then we move the phone to capture our test scenes. We report the results in Table 3 under “Ours (ignoring OIS)”, showing a large decrease in depth accuracy. Thus, online calibration is both necessary and effective. Figure 6

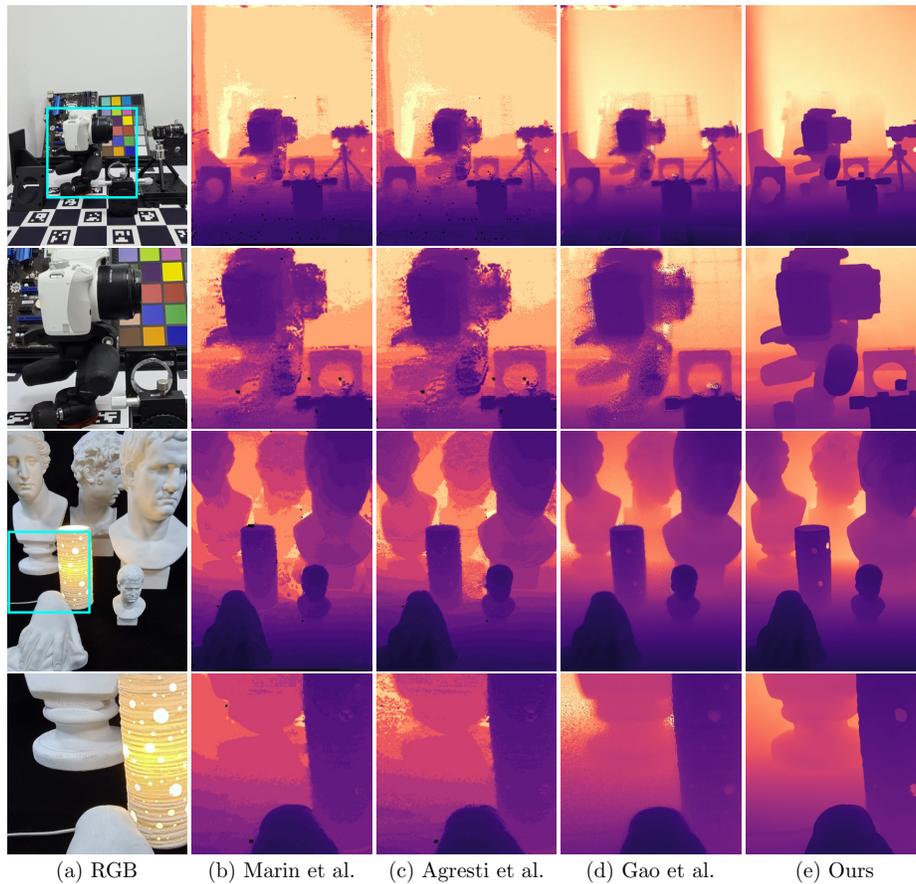


Fig. 6: ToF/stereo fusion results. We pair all other methods with a state-of-the-art stereo approach for fair comparison [28]. Marin et al. [33] and Agresti et al. [1] suffer from quantization as their sub-pixel resolution approach is not suitable for small phone camera baselines. Gao et al. [15] relies heavily on ToF measurements, degrading its performance when ToF is inaccurate (e.g., black parts of the camera). Refer to the supplemental material for additional results.

shows that our method allows for robust depth estimation with better edge and hole preservation. The low RMSE in Table 3 suggests that our method is robust against strong outliers. While other methods suffer from a less robust stereo matching—swapping theirs for RAFT-stereo [28] significantly improving their results—our approach maintains higher accuracy.

4.4 Dependency on the Device

Under the assumption of a narrow baseline on smartphones, the method should generalize as it allows accurate optical flow estimation between the two RGB cameras. In addition, we show the results of our fusion of datasets based on different hardware: a ZED stereo camera and a Microsoft Kinect v2 ToF depth camera for REAL3 [2], and two calibrated BASLER scA1000 RGB cameras and a MESA SR4000 ToF camera for LTTM5 [10] in the supplemental material.

4.5 Limitations

While our approach applies to indoor environments, the reliance on ToF and stereo prevents application in some scenarios. First, the ToF module cannot estimate depth accurately at large distances due to its low power. Second, ToF depth estimation is not reliable within strong IR ambient illumination (e.g., direct daylight). Since our calibration relies directly on ToF measurements, it becomes inaccurate if no ToF depth can be estimated. In addition, some materials—particularly translucent or specular materials—are challenging for both ToF and stereo depth estimation and cannot be tackled by our fusion approach.

5 Conclusion

Optical-image-stabilized lenses are now common but present problems for pose estimation when wanting to fuse information across multiple sensors in a camera system. This limits our ability to estimate high-quality depth maps from a single snapshot. Our method is designed for consumer devices, tackling calibration and robust sensor fusion for indoor environments. As our approach uses only a single snapshot and does not exploit camera motion for pose estimation, the acquisition is quick and could be used on dynamic scenes. Evaluated on real-world inputs, our method estimates more accurate depth maps than state-of-the-art ToF and stereo fusion methods.

Acknowledgement

Min H. Kim acknowledges funding from Samsung Electronics, in addition to the partial support of the MSIT/IITP of Korea (RS-2022-00155620, 2022-0-00058, and 2017-0-00072), the NIRCH of Korea (2021A02P02-001), Microsoft Research Asia, and the Samsung Research Funding Center (SRFC-IT2001-04) for developing 3D imaging algorithms. James Tompkin thanks US NSF CAREER-2144956.

References

1. Agresti, G., Minto, L., Marin, G., Zanuttigh, P.: Deep learning for confidence information in stereo and tof data fusion. In: ICCV Workshops (2017)
2. Agresti, G., Minto, L., Marin, G., Zanuttigh, P.: Stereo and tof data fusion by learning from synthetic data. *Information Fusion* (2019)
3. Agresti, G., Zanuttigh, P.: Deep learning for multi-path error removal in tof sensors. In: ECCV Workshops (2018)
4. Alex Yu and Sara Fridovich-Keil, Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks (2021)
5. Attal, B., Laidlaw, E., Gokaslan, A., Kim, C., Richardt, C., Tompkin, J., O’Toole, M.: Törf: Time-of-flight radiance fields for dynamic scene view synthesis. In: NeurIPS (2021)
6. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: ICCV (2021)
7. Brown, D.C.: Decentering distortion of lenses. *Photogrammetric Engineering* (1966)
8. Brown, M.A., Süssstrunk, S.: Multi-spectral sift for scene category recognition. In: CVPR (2011)
9. Conrady, A.E.: Decentred lens-systems. *Monthly Notices of the Royal Astronomical Society* (1919)
10. Dal Mutto, C., Zanuttigh, P., Cortelazzo, G.M.: Probabilistic tof and stereo data fusion based on mixed pixels measurement models. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2015)
11. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791* (2021)
12. DiVerdi, S., Barron, J.T.: Geometric calibration for mobile, stereo, autofocus cameras. In: WACV (2016)
13. Efe, U., Ince, K.G., Alatan, A.: Dfm: A performance baseline for deep feature matching. In: CVPR Workshops (2021)
14. Evangelidis, G.D., Hansard, M.E., Horaud, R.: Fusion of range and stereo data for high-resolution scene-modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2015)
15. Gao, Y., Esquivel, S., Koch, R., Keinert, J.: A novel self-calibration method for a stereo-tof system using a kinect V2 and two 4k gopro cameras. In: 3DV (2017)
16. Gil, Y., Elmalem, S., Haim, H., Marom, E., Giryes, R.: Online training of stereo self-calibration using monocular depth estimation. *IEEE Transactions on Computational Imaging* (2021)
17. Guo, Q., Frosio, I., Gallo, O., Zickler, T., Kautz, J.: Tackling 3d tof artifacts through learning and the flat dataset. In: ECCV (2018)
18. Ha, H., Lee, J.H., Meuleman, A., Kim, M.H.: Normalfusion: Real-time acquisition of surface normals for high-resolution rgb-d scanning. In: CVPR (2021)
19. Hansard, M., Lee, S., Choi, O., Horaud, R.: *Time of Flight Cameras: Principles, Methods, and Applications*. Springer Briefs in Computer Science, Springer (2012)
20. Holynski, A., Kopf, J.: Fast depth densification for occlusion-aware augmented reality. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* (2018)
21. Jeong, Y., Ahn, S., Choy, C., Anandkumar, A., Cho, M., Park, J.: Self-calibrating neural radiance fields. In: ICCV (2021)

22. Jung, H., Brasch, N., Leonardis, A., Navab, N., Busam, B.: Wild tofu: Improving range and quality of indirect time-of-flight depth with rgb fusion in challenging environments. In: 3DV (2021)
23. Kopf, J., Matzen, K., Alsisan, S., Quigley, O., Ge, F., Chong, Y., Patterson, J., Frahm, J.M., Wu, S., Yu, M., Zhang, P., He, Z., Vajda, P., Saraf, A., Cohen, M.: One shot 3d photography. In: ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH) (2020)
24. Kopf, J., Rong, X., Huang, J.B.: Robust consistent video depth estimation. In: CVPR (2021)
25. Li, L.: Time-of-flight camera – an introduction (2014), <https://www.ti.com/lit/wp/sloa190b/sloa190b.pdf>
26. Li, Z., Dekel, T., Cole, F., Tucker, R., Snavely, N., Liu, C., Freeman, W.T.: Learning the depths of moving people by watching frozen people. In: CVPR (2019)
27. Lin, C.H., Ma, W.C., Torralba, A., Lucey, S.: Barf: Bundle-adjusting neural radiance fields. In: ICCV (2021)
28. Lipson, L., Teed, Z., Deng, J.: RAFT-Stereo: Multilevel Recurrent Field Transforms for Stereo Matching. In: 3DV (2021)
29. Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehrmann, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. ACM Transactions on Graphics (2019)
30. Lombardi, S., Simon, T., Schwartz, G., Zollhoefer, M., Sheikh, Y., Saragih, J.: Mixture of volumetric primitives for efficient neural rendering. ACM Transactions on Graphics (2021)
31. Luo, X., Huang, J., Szeliski, R., Matzen, K., Kopf, J.: Consistent video depth estimation. ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH) (2020)
32. Marco, J., Hernandez, Q., Muñoz, A., Dong, Y., Jarabo, A., Kim, M.H., Tong, X., Gutierrez, D.: Deeptof: Off-the-shelf real-time correction of multipath interference in time-of-flight imaging. ACM Transactions on Graphics (2017)
33. Marin, G., Zanuttigh, P., Mattoccia, S.: Reliable fusion of tof and stereo depth driven by confidence measures. In: ECCV (2016)
34. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: CVPR (2015)
35. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
36. Pham, F.: Fusione di dati stereo e time-of-flight mediante tecniche di deep learning (2019), <https://github.com/frankplus/tof-stereo-fusion>
37. Poggi, M., Agresti, G., Tosi, F., Zanuttigh, P., Mattoccia, S.: Confidence estimation for tof and stereo sensors and its application to depth data fusion. IEEE Sensors Journal (2020)
38. Qiu, D., Pang, J., Sun, W., Yang, C.: Deep end-to-end alignment and refinement for time-of-flight rgb-d modules. In: ICCV (2019)
39. Rongrong Gao and Na Fan, Li, C., Liu, W., Chen, Q.: Joint depth and normal estimation from real-world time-of-flight raw data. In: IROS (2021)
40. Sachs, D., Nasiri, S., Goehl, D.: Image stabilization technology overview. InvenSense Whitepaper (2006)
41. Scharstein, D., Hirschmüller, H., Kitajima, Y., Krathwohl, G., Nešić, N., Wang, X., Westling, P.: High-resolution stereo datasets with subpixel-accurate ground truth. In: German Conference on Pattern Recognition (GCPR) (2014)
42. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: CVPR (2016)

43. Shih, M.L., Su, S.Y., Kopf, J., Huang, J.B.: 3d photography using context-aware layered depth inpainting. In: CVPR (2020)
44. Son, K., Liu, M.Y., Taguchi, Y.: Learning to remove multipath distortions in time-of-flight range images for a robotic arm setup. In: ICRA (2016)
45. Song, S., Lichtenberg, S.P., Xiao, J.: Sun rgb-d: A rgb-d scene understanding benchmark suite. In: CVPR (2015)
46. Su, S., Heide, F., Wetzstein, G., Heidrich, W.: Deep end-to-end time-of-flight imaging. In: CVPR (2018)
47. Teed, Z., Deng, J.: Raft: Recurrent all pairs field transforms for optical flow. In: ECCV (2020)
48. Valentin, J., Kowdle, A., Barron, J.T., Wadhwa, N., Dzitsiuk, M., Schoenberg, M., Verma, V., Csaszar, A., Turner, E., Dryanovski, L., Afonso, J., Pascoal, J., Tsotsos, K., Leung, M., Schmidt, M., Guleryuz, O., Khamis, S., Tankovitch, V., Fanello, S., Izadi, S., Rhemann, C.: Depth from motion for smartphone ar. SIGGRAPH Asia (2018)
49. Wadhwa, N., Garg, R., Jacobs, D.E., Feldman, B.E., Kanazawa, N., Carroll, R., Movshovitz-Attias, Y., Barron, J.T., Pritch, Y., Levoy, M.: Synthetic depth-of-field with a single-camera mobile phone. SIGGRAPH (2018)
50. Wang, J., Qiu, K.F., Chao, P.: Control design and digital implementation of a fast 2-degree-of-freedom translational optical image stabilizer for image sensors in mobile camera phones. *Sensors* **17** (2017)
51. Wei, Y., Liu, S., Rao, Y., Zhao, W., Lu, J., Zhou, J.: Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In: ICCV (2021)
52. Xian, W., Huang, J.B., Kopf, J., Kim, C.: Space-time neural irradiance fields for free-viewpoint video. In: CVPR (2021)
53. Zhang, X., Matzen, K., Nguyen, V., Yao, D., Zhang, Y., Ng, R.: Synthetic defocus and look-ahead autofocus for casual videography. SIGGRAPH (2019)
54. Zhou, Y., Barnes, C., Jingwan, L., Jimei, Y., Hao, L.: On the continuity of rotation representations in neural networks. In: CVPR (2019)
55. Zhu, R., Yu, D., Ji, S., Lu, M.: Matching rgb and infrared remote sensing images with densely-connected convolutional neural networks. *Remote Sensing* (2019)