

Wild ToFu: Improving Range and Quality of Indirect Time-of-Flight Depth with RGB Fusion in Challenging Environments

HyunJun Jung¹, Nikolas Brasch¹, Aleš Leonardis², Nassir Navab¹, Benjamin Busam¹

¹ Technical University of Munich, ² Huawei Noah’s Ark Lab

{hyunjun.jung, nikolas.brasch, nassir.navab, b.busam}@tum.de

Abstract

Indirect Time-of-Flight (I-ToF) imaging is a widespread way of depth estimation for mobile devices due to its small size and affordable price. Previous works have mainly focused on quality improvement for I-ToF imaging especially curing the effect of Multi Path Interference (MPI). These investigations are typically done in specifically constrained scenarios at close distance, indoors and under little ambient light. Surprisingly little work has investigated I-ToF quality improvement in real-life scenarios where strong ambient light and far distances pose difficulties due to an extreme amount of induced shot noise and signal sparsity, caused by the attenuation with limited sensor power and light scattering. In this work, we propose a new learning based end-to-end depth prediction network which takes noisy raw I-ToF signals as well as an RGB image and fuses their latent representation based on a multi step approach involving both implicit and explicit alignment to predict a high quality long range depth map aligned to the RGB viewpoint. We test our approach on challenging real-world scenes and show more than 40% RMSE improvement on the final depth map compared to the baseline approach [33].

1. Introduction

Depth sensing is a crucial element in modern 3D vision applications such as realistic geometry-aware augmented reality [18, 4] and an enabling technology for numerous robotics applications [20, 3]. A variety of cameras and technologies to estimate distances exist. Prominent passive depth prediction methods such as multi-view stereo suffer in low-textured areas and with limited ambient light while active methods such as structured light struggle in bright ambient light environments and may require multiple consecutive image acquisitions which hamper their real-time use [12].

Time-of-Flight (ToF) imaging is an active sensing method that can produce real-time depth maps. For indi-

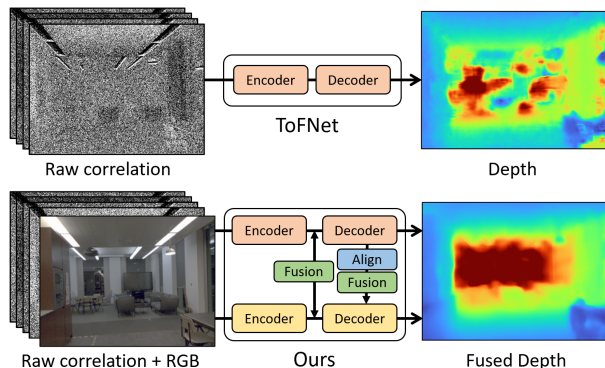


Figure 1. Unlike the baseline approach ToFNet [33] which only uses the raw correlation signals, we leverage an additional RGB input by fusing its latent representation in a multi-step alignment process and achieve a significant improvement of over 40% RMSE in challenging indoor scenarios.

rect ToF (I-ToF), the scene is illuminated with an amplitude-modulated light typically in the infrared spectrum. Due to the constant speed of light c , depth can be computed from the round-trip time of emitted photons. To measure the time delay, the sensor uses specific hardware to correlate a number of phase-shifted signals from the emitter diode with the received light reflected from scene geometry. Measuring 4 phase shifts per pixel generates the correlation images $C_i, i \in \{1, 2, 3, 4\}$ that are directly connected to the depth map D (see Sec.3). Complementary to stereo RGB, depth sensing from ToF works well in low-light scenarios and even for textureless objects due to its active illumination. Since the emitted light intensity decreases quadratically with the distance, ToF depth maps are typically less accurate at further distances. Ambient light sources can introduce high levels of noise on the correlation maps resulting in missing depth values or incorrect estimates. Strongly diffuse, absorptive (dark) and reflective materials as well as unwanted multiple reflections can additionally hamper the depth quality. The complementary advantages and the recent availability even on mobile phones [27] make Time-of-Flight an interesting modality that has recently been studied

mainly under ideal environmental conditions, in lab scenarios with low ambient light and on synthetically created images [15, 33, 1, 6].

Overview and Contribution. Previous correlation-based ToF pipelines target real applications in close range with controlled external illumination and utilize synthesized pixel-perfect ground truth signals for supervision (see Fig. 3). In this work, we investigate a way to improve the quality and range of single frequency indirect Time-of-Flight for challenging real indoor scenarios such as scenes with distances beyond the manufacturer’s maximum range with changing and strong ambient light by leveraging extra RGB input.

Using complementary multi-modal images can be problematic when they are not aligned. Both images can be aligned by calibration only if a high quality depth map is available. However, the quality of an I-ToF depth map is often not accurate enough for this task due to sparsity and depth errors caused by Multi Path Interference (MPI) or Motion Artifacts (MA). Matching algorithms also fail due to a large domain difference between I-ToF and RGB [16]. Qiu et al. [30] propose to use a flow based alignment between I-ToF depth and RGB which can be used to align the images to fix low frequency errors such as MPI. However, this approach is constrained to exploit the RGB information as the obtained depth map from the camera already lost information on the problematic low confidence regions during on-camera’s pre-processing, leaving no rooms to improve the depth map in those regions jointly with RGB from raw I-ToF signals. While aligning raw signals to the RGB frame using a flow algorithm can be problematic in the presence of extreme sparsity and shot noise typically present in real-life scenarios.

To exploit the richness of information in the raw signal jointly with an unaligned RGB image, we propose a multi step approach which gradually aligns both information while aggregating them at the same time. First, we align both camera inputs with the same intrinsic K , then feed them into independent encoders with a low resolution bottleneck (1/32 input resolution). Next, we fuse both feature maps from the bottleneck assuming both of them are weakly aligned to predict an improved depth at low resolution. Then we repeat the process estimating up-scaled depth maps until recovering the full resolution by 1) aligning I-ToF feature maps to RGB feature maps by warping with an intermediate depth and pose 2) applying a fusion module 3) hierarchical upsampling and depth estimation, which at the end predicts high quality depth on the RGB viewpoint. In the evaluation we show that our proposed method resolves the alignment issue while significantly improving over baseline approaches.

To this end, our main contributions are:

1. We tackle I-ToF in a realistic and challenging scenario which involves strong ambient light and at ranges even beyond the manufacturer’s maximum depth.
2. We leverage unaligned RGB information jointly with I-ToF at low resolution while resolving their alignment issue with a special architecture capable of producing high resolution depth outputs.¹
3. We provide insight for fusion strategy between spatially separated sensors by extensive ablation studies.

2. Related Work

Previous works related to I-ToF depth correction can be separated into two branches. The first is directly connected to the I-ToF modality which mostly targets to improve specific noise patterns in the modality such as MPI from either raw signal or estimated depth from the camera. The other branch is not necessarily limited to I-ToF, but focuses on completing depth with RGB input in case of sparsity which can be applied to I-ToF depth as it often suffers from sparsity due to its limited power and scattering.

2.1. Indirect Time-Of-Flight

From raw correlations Early correction approaches using raw correlation images investigate mainly MA and MPI which are commonly addressed sources of I-ToF depth error. A few works [31, 23] focus on constraints in the correlation images to improve MA, and other works [10, 11, 21, 7, 14, 22] specifically target MPI artefacts. Some works focus on improving the overall range of the depth map by resolving the periodicity in the calculated phase (phase unwrapping, Sec.3) using intensity [29], multi-frequency modulation [8] or signal distribution [5].

While the aforementioned ToF approaches offer the possibility to easily exchange modules for denoising, phase unwrapping and MPI / MA correction, they work only independently and suffer from error propagation. Some scholars [33, 15] address this problem by applying learning based methods on the raw correlation data. Su et al. [33] understood the benefit of training with direct sensor input and use dual-frequency raw correlation images for an encoder-decoder network to predict depth in an end-to-end approach, while Guo et al. [15] use a window-based approach with a Kernel Prediction Network (KPN) [26] to pre-process the raw correlation images before extracting the depth. They both synthetically create a training dataset with a physics based renderer (Fig.2), which is capable of generating physical noise such as MPI and scattering as well as Gaussian and Poisson noise. However, the rendering resembles controlled lab scenarios which have limited ambient light and close distance objects which does not reflect challenging real world scenarios as exemplified in Fig.6. Although our

¹We also plan to release the code and weights upon paper acceptance.

work shares the similarity with these two works by taking raw correlation images as input and also uses a learning based approach, we rather focus on a more realistic and challenging setup, such as single frequency raw correlation images in uncontrolled and long distance indoor scenes to which previous works cannot generalize well.

From I-ToF depth Another series of works focuses on directly improving the depth map obtained by an I-ToF camera. At this stage the signal may have already undergone linear and nonlinear models from either hardware or software of the camera and information in areas with low confidence measures may be lost. Some scholars investigate I-ToF depth improvement as a post processing stage [32, 25, 6] focusing on removing MPI, MA, and shot noise while others fuse ToF with complementary information from an RGB camera [30] or use generative adversarial networks to bridge the gap between synthetic and real data [1]. These works are evaluated in relatively optimized scenarios where the main error source is MPI or MA. However, as shown in Fig.3 and Fig.6, I-ToF depth in real scenarios with strong ambient light suffers mainly from sparsity and an extreme amount of shot noise.

2.2. Depth Completion

The depth map obtained from the depth camera in general contains invalid pixels due to the hardware limitations. Although not ToF specific, there has been recent progress regarding the completion of invalid pixels especially for LiDAR depth jointly with RGB input, which shares similar principles with our work. Common approaches treat the sparse depth as a single channel image input for a network [35, 34, 24], while more geometric approaches being developed recently treat sparse depth in different form other than a single channel image. Hu et.al. [19] use an extra analytical feature map calculated by back-projecting coordinates with camera intrinsics and sparse depth in multiple layers and Chen et.al. [37] treat the sparse depth map as a point cloud to apply point-based networks. Applying those works directly on raw I-ToF correlation with RGB would be problematic. Geometric approaches are not feasible by design as the raw signal requires more steps to be used as distance measures as well as its strong noise. And different to LiDAR, the present misalignment between the raw signal and RGB can cause additional artefacts around object boundaries and depth discontinuities. In Sec.6 we show the impact of the misalignment on the final depth prediction.

3. Time-of-Flight Sensing

In this section we will briefly introduce the background of Time-of-Flight sensing and its possible sources of error.

3.1. Depth Acquisition from I-ToF

Correlation Images. When an Indirect Time-of-Flight (I-ToF) camera measures depth, it first emits an amplitude modulated infrared (IR) signal and then measures the returned signal which contains phase shift information. To recover the phase shift, cross-correlation is used. Cross-correlation between two signals $f(t)$ and $g(t)$ is defined as:

$$(f \otimes g)(\varphi) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} f(t) \cdot g(t + \varphi) dt \quad (1)$$

When we plug our returned signal $r(t)$ with amplitude a and phase shift φ into g and demodulation signal $d(t)$ with control phase τ into f of Eq 1,

$$r(t) = 1 + a \cdot \cos(\omega t - \varphi), \quad d(t) = \cos(\omega t - \tau) \quad (2)$$

we can obtain the function of both phase shift φ and control phase τ with

$$\begin{aligned} (r \otimes d)(\varphi) &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} [\cos(\omega t - \tau)] \cdot [1 + a \cdot \cos(\omega t - \varphi)] dt \\ &= \frac{a}{2} \cdot \cos(\varphi + \tau) \end{aligned} \quad (3)$$

With known control phases $\tau \in [0^\circ, 90^\circ, 180^\circ, 270^\circ]$, we retrieve functions of φ as

$$\begin{aligned} C_1 = c_0 &= \frac{a}{2} \cos(\varphi), & C_2 = c_{180} &= -\frac{a}{2} \cos(\varphi), \\ C_3 = c_{90} &= \frac{a}{2} \sin(\varphi), & C_4 = c_{270} &= -\frac{a}{2} \sin(\varphi). \end{aligned} \quad (4)$$

These constitute the four raw correlation values which are measured by an I-ToF camera for each pixels in the image and result in four correlation maps as shown in Fig.2(a).

Phase calculation and unwrapping. As correlation images are sine functions of phase shift φ , the arc-tangent formula can be used to extract φ . However, it can be only measured up to $0-2\pi$ range due to the nature of the periodic signal. The signal φ with its 2π -ambiguity is called the wrapped phase (see Fig.2(b)).

$$\varphi_{wrapped} = \arctan \left(\frac{C_3 - C_4}{C_1 - C_2} \right) \quad (5)$$

To obtain the full range of the depth map, it is required to find the correct multiplier for 2π . This disambiguation is called phase unwrapping.

$$\varphi = \varphi_{wrapped} + 2n\pi \quad (n \in \{0, 1, 2, \dots\}) \quad (6)$$

The phase unwrapping factor can be calculated either explicitly by non-learning based approaches such as by using amplitude [29], multi frequency modulation signal [8],

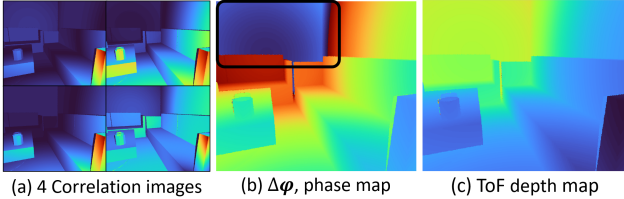


Figure 2. Synthetic I-ToF images from [15]. Each of the 4 correlation images (a) is function of the phase map. Note that the phase values drop suddenly in the marked area of the phase map (b) due to the phase wrapping. The full range depth map (c) can be obtained after phase unwrapping.

signal distribution [5], or implicitly by learning based approaches such as [33]. However, due to the high amount of noise under real lighting conditions the intensity and distribution methods do not work well. In Sec.6 we show that in the case of a single frequency I-ToF signal with high noise, the main source of depth error comes from the confusion of the unwrapping factor.

Depth conversion. After the phase φ is unwrapped, the distance D can be obtained with a conversion formula using the speed of light c and the modulation frequency ω as illustrated in Fig.2 (c) with

$$D = \frac{c}{2\omega} \cdot \frac{\varphi}{2\pi} \quad (7)$$

3.2. Time-of-Flight’s Source of Error

The acquisition process of I-ToF cameras introduces three main sources of error which we briefly describe.

Multi path interference and motion artefact. The superposition of return signals with different φ in one pixel is one of the main sources of I-ToF error. It happens when the returned signal contains multiple reflection paths (MPI) or sudden changes of the returned signal φ due to the camera motion (MA). In both cases, the final correlation signals become superimposed with signals of different φ and Eq.5 in turn fails to provide the accurate phase. Aforementioned artefacts can be characterised as low frequency noise around the corner (MPI) and over/under-shooting around the object boundary (MA). However, we will identify that these artefacts are not the main problem in challenging real scenes.

Sparsity and shot noise. The raw correlation images from the I-ToF camera rely on the returned IR signal from reflections. Any factors causing scattering of the IR signal can degrade the quality of the raw correlation images, such as the surface normal or the material properties of objects in the scene. Hansard et al. [16] shows that the amplitude of the return signal decreases and the depth error

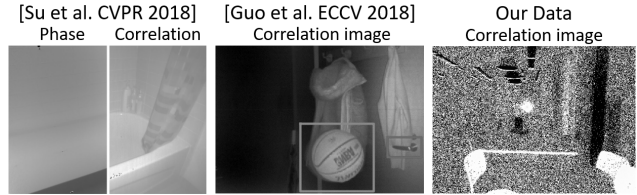


Figure 3. Comparison of input correlation noise on real data from Su et al. [33], Guo et al. [15], and us. The significantly higher noise level of our correlation images is immanent.

increases as the surface normal of the object deviates from the optical axis of the camera and when the object material is reflective due to a scattering effect. In real-life scenarios, when the aforementioned regions are combined with strong ambient light, the signals become highly noisy as shown in Fig.6. Non-static ambient illumination can additionally cause incorrect measurements of return signals in practice. To reduce the correlation deterioration by ambient light, previous investigations focused on ideal environmental conditions during data acquisition while we specifically target less constraint setups which manifests in significant differences in the input data as shown in Fig.3.

4. ToF Fusion

Our pipeline uses an encoder-decoder style fully convolutional network with raw I-ToF input similar to [33]. However, significant changes are made to deal with the extreme amount of noise in the raw signal and to leverage the unaligned RGB image.

4.1. Network Backbone

We use a two branch encoder-decoder network architecture with independent weights to predict the depth map from raw I-ToF input signal as well as RGB input. The overall architecture is illustrated in Fig.4. Different from [33] where the $1/4$ resolution feature map serves as bottleneck, we use the same encoder-decoder architecture with $1/32$ resolution bottleneck as in [13] for both RGB and I-ToF branch separately, but with additional feature sharing in the bottleneck and decoder layers with our custom fusion module (Sec.4.3). For the I-ToF branch, we modify the architecture by expanding the input channel dimension while removing all Batch-Norm layers so that the network can process 4 correlation images as input and keep their channel wise relation unchanged as in [33]. There are three reasons behind our choice of the architecture:

1. The $1/32$ resolution bottleneck helps to extract features in a more global context, which already improves the depth prediction in areas of unreliable I-ToF.
2. Misalignment between RGB and I-ToF is less problematic at $1/32$ resolution, making fusion possible between both feature maps without any alignment.

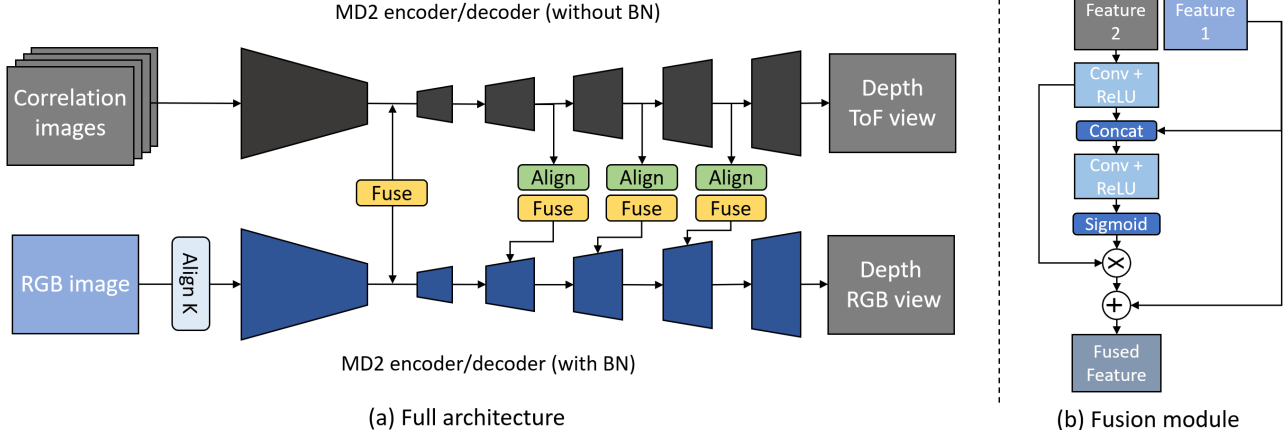


Figure 4. Details of our network architecture. While the lower branch in (a) follows [13], the upper branch in uses [13] without Batch-Norm and with a 4-channel input. The fusion in the bottleneck combines both feature maps bidirectionally without any alignment, while higher resolution fusions are performed unidirectionally from I-ToF to the RGB branch with intermediate depth and extrinsic calibration between the cameras. All fusion steps use the fusion module described in (b).

- Intermediate depth from the decoder can be used to align intermediate feature maps from one to the other.

4.2. Intrinsic Alignment

Although the 1/32 resolution bottleneck of ResNet-18 [17] helps to ease the misalignment between I-ToF and RGB camera, the difference in camera intrinsic K between the sensors can cause issues due to different resolution and field of view (FOV) during the fusion. To overcome this, we warp from the image I_1 with the larger field of view (RGB) onto the image I_2 with the smaller FOV (I-ToF) using the camera intrinsics K_1, K_2 of both cameras as depicted in Eq.8 such that both branches have the same resolution and least amount of invalid pixel (see Fig.5).

$$\begin{aligned} (x_2, y_2, 1)^T &= K_2 K_1^{-1} (x_1, y_1, 1)^T \\ I_{1 \rightarrow 2}(x_2, y_2) &= I_1(x_1, y_1) \end{aligned} \quad (8)$$

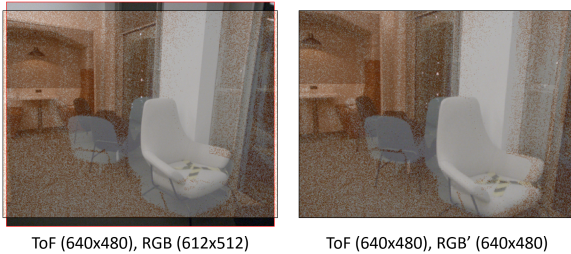


Figure 5. Two input images overlaid before (Left) and after (Right) intrinsic alignment. After the alignment, both images share equal resolution and the scene in the field of view only involves displacement, no scaling.

4.3. Fusion Module

Aligning images from different modalities is an important step before the fusion. Qiu et.al. [30] use a flow based network to align from I-ToF to RGB. However, this cannot be used in our setup due to the sparsity and strong noise which makes flow estimation difficult. Explicit alignment before fusion is not possible either with only RGB depth or I-ToF depth as I-ToF depth is unreliable due to the noise and phase unwrapping error, while RGB depth prediction produces poor depth when the scene has little texture or photometrically-changing geometry to infer depth from. To overcome this problem, we utilize a two step approach. First, we use encoders with low resolution bottleneck and then fuse both feature maps to predict improved intermediate depth maps at low resolution. Then we warp the I-ToF feature map onto the RGB reference frame with intermediate depth predictions and camera extrinsics and apply fusion inside the RGB decoder.

$$(x_2, y_2, 1)^T = K_2 h((R d_1 K_1^{-1} (x_1, y_1, 1)^T + t)) \quad (9)$$

Where camera 1 is I-ToF and 2 is RGB. R, t denote the extrinsics, d_1 intermediate depth prediction, and h indicates homogenization. We iterate this for the RGB decoder at resolution (1/16→1/8), (1/8→1/4), (1/4→1/2). The detailed structure of the entire network is described in our supplementary material. Passing the features from I-ToF to RGB has the advantage of relying on the finest depth possible in the RGB viewpoint at the end which is then ready for applications such as augmented reality.

The fusion module structure is shown in Fig.4 (b). As mentioned, both features are assumed to be aligned either implicitly (at 1/32 res.) or explicitly (with fixed extrinsics between the two sensors for intermediate depth). To convert

each feature map to the other, we use a Conv-layer followed by a ReLU-activation. To iterate on the confidence of the converted feature map, an additional Conv-layer with ReLU is used after concatenation with the additional feature map. A Sigmoid-activation is applied to normalize the output to the range $[0, 1]$. This attention weight is applied to the converted feature map before it is added to the other feature map.

4.4. Loss Formulation

We formulate our loss as a fully supervised multi-scale \mathcal{L}_1 loss in both the RGB and I-ToF branch with additional edge aware smoothness [13] in the RGB branch. For \mathcal{L}_1 , we adapted the multi-scale approach from [13], where each intermediate depth output with resolution i is upsampled in both branches to the full resolution d^i which is then compared to the ground truth \tilde{d} . The loss term reads as

$$\mathcal{L}_1 = \sum_i \left(\sum |d_{tof}^i - \tilde{d}_{tof}| + \sum |d_{rgb}^i - \tilde{d}_{rgb}| \right) \quad (10)$$

Edge aware smoothness is only applied on the RGB image I_{rgb} . Su et al. [33] apply this loss on the I-ToF branch by treating the I-ToF amplitude as a grey scale image. Due to sparsity and shot noise present in our data, we do not include such a term. The smoothness loss is calculated with

$$\mathcal{L}_s = \sum |\partial_x d_{rgb}| e^{-|\partial_x I_{rgb}|} + |\partial_y d_{rgb}| e^{-|\partial_y I_{rgb}|} \quad (11)$$

Our final loss is a weighted combination of \mathcal{L}_1 and \mathcal{L}_s :

$$\mathcal{L} = \mathcal{L}_1 + \lambda_s \mathcal{L}_s \quad (12)$$

5. Dataset and Training Details

In this section, we introduce the dataset we used for both training and evaluation as well as the detailed setup for the training

Training Dataset For our work, we leverage the indoor sequences from the multi-modal dataset [2] which constitutes ToF acquisitions under challenging light and scene conditions as illustrated in Fig.6. The dataset comprises synchronised image sequences capturing multiple modalities, at video-rate across various real-world indoor and outdoor scenes. A custom camera rig is constructed in order to capture synchronised data including i-ToF correlations from a Lucid HLS003S-001 camera (LUCID Vision Labs, Canada). The used RGB images are acquired with a resolution of 612×512 px while i-ToF captures 4 channel raw correlation images of 640×480 px resolution with 8 bit depth. A structured light sensor (RealSense D435i, Intel, USA) acquires active IR stereo depth maps in the RGB image viewpoint that serve as ground truth. Intrinsic and extrinsic camera parameters are calibrated via multi-view graph

bundle adjustment and provided together with the images. The indoor sequences are composed of 4 individual video sequences with around 800 frames each, recorded on one floor of large building with different trajectories and viewpoints. We use 3 full sequences for training which sums up to a total of 2230 frames and leverage the provided test split in the remaining sequence for validation.

Training Setup Before the training we preprocess both ground truth and input images. The ground truth is processed in two different ways to align to both RGB and I-ToF coordinate frames. We change the intrinsics of the ground truth in the same way as mentioned in Sec.4.2 for the RGB branch. For I-ToF supervision, we use forward warping to bring the ground truth depth onto the I-ToF viewpoint using the provided extrinsics. For data augmentation, we apply a sequence of brightness/contrast/saturation/hue transformations for RGB input. We use the pretrained ResNet-18 [17] encoder provided by PyTorch [28] for our RGB branch, while the I-ToF branch is trained from scratch. We use the ADAM optimizer for the entire training with an initial learning rate of 10^{-4} which decays 0.95 every 1000 iterations. For all setups in Tab.1, we train 80 epochs (approx. 12 hours using an Nvidia RTX 3090 GPU) with a loss weighting factor of $\lambda_s = 1.0$. We evaluate per epoch and used the result of best performing epoch for our tables.

6. Experiments

For evaluation, we first compare our method to two state of the art methods: ToFNet (raw I-ToF to depth [33]) and GuideNet (RGB guided depth densification [34]), and then justify pipeline components with an ablation study.

6.1. Quantitative and Qualitative Evaluation

To quantitatively evaluate our results, we calculate 4 depth metrics suggested in [9]: Root Mean Square Error (RMSE), absolute and square relative error (Rel(abs), Rel(sqr)), and depth accuracy with threshold (1.25^i). We first train ToFNet [33] and GuideNet [34] from scratch on the multi-modal dataset [2]. To optimize performance on GuideNet for this data, we make three adjustments to their pipeline: We remove Batch-Norm layers on the depth branch, increase the input dimension to 4, and also use intrinsic alignment (Sec.4.2) to resolve resolution differences. The training setups follow the original implementation. A quantitative evaluation is shown in Tab.1 and qualitative results are depicted in Fig.7.

In our experiments, ToFNet [33] achieves the worst overall results among the three methods. Fig.7 indicates that the largest errors come from the phase unwrapping (second column, marked red). As mentioned in Sec.3, the I-ToF signal has a periodicity and its phase (or depth) is wrapped

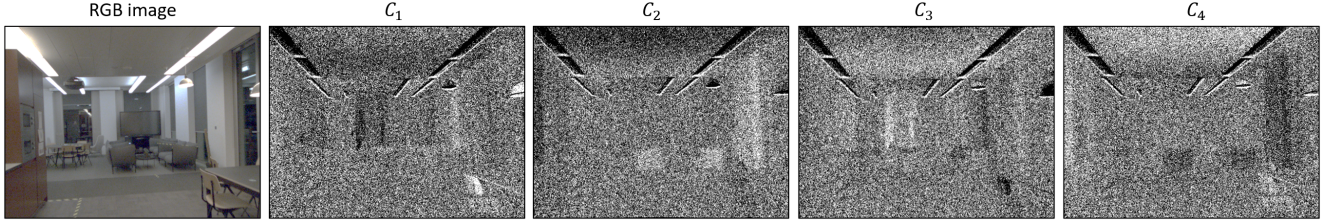


Figure 6. Example acquisition of a scene (RGB image) with corresponding raw I-ToF measurements as correlation images $C_i, i \in \{1, 2, 3, 4\}$. Due to ambient light and distant regions, the significant amount of shot noise is clearly visible.

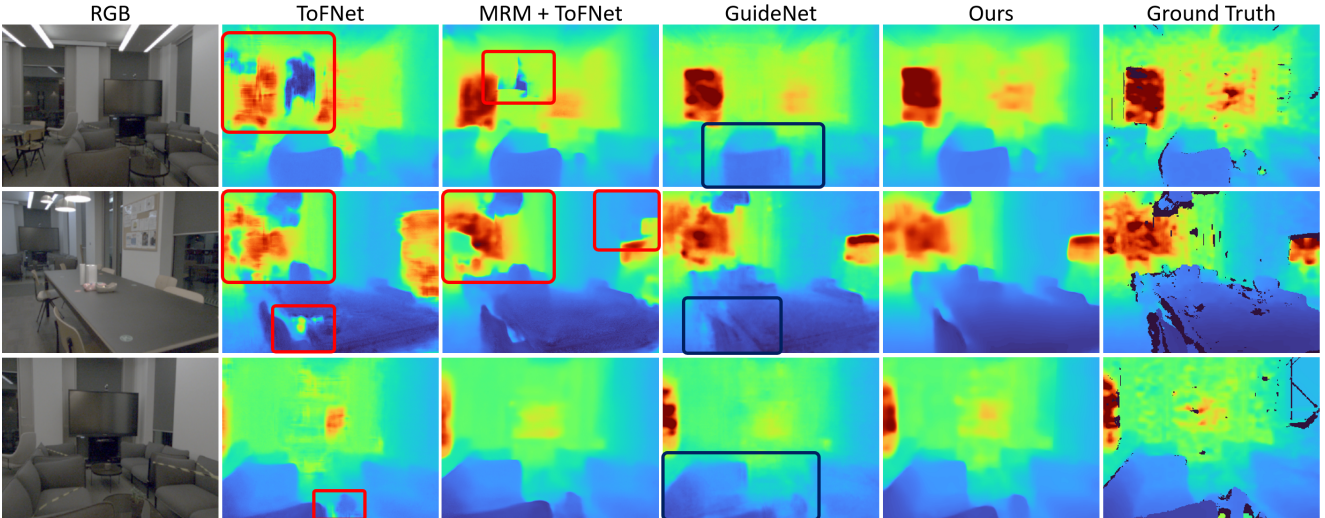


Figure 7. Qualitative comparison of ToFNet [33], GuideNet [34] and our pipeline. With or without pre-filtering (MRM [15]), ToFNet suffers from strong artefacts caused by a wrong unwrapping factor (marked red), which results in large regional errors and horizontal/vertical line artefacts. GuideNet [34] resolves this issue with RGB information, however unaligned RGB causes ghosting effects around the discontinuities (marked blue). Our method does not suffer from these artefact and achieves the best result. Note that corresponding correlation images of each RGB image are shown in the supplementary material.

by 2π . This requires extra information to obtain the full range (phase unwrapping). Su et.al applied ToFNet to a relatively clean signal (Fig.3) using dual frequency modulation as input such that the network learns the unwrapping factor easily both analytically and geometrically. This is not the case with the dataset [2] where I-ToF comes as a single frequency and is sparse and noisy. Improving the quality of the input signal by applying pre-filtering network [15] removes some of the artefacts, but not completely (the second col in Fig.7). GuideNet [34] achieved a huge improvement over ToFNet especially regarding phase unwrapping, showing the importance of leveraging dense RGB information in this scenario. The mid column of Fig.7 shows that the unwrapping related artefacts no longer exist in the prediction. However, the unaligned RGB input causes confusion around depth discontinuities which result in blurry predictions and ghosting effects (marked blue). Our proposed method resolves the misalignment issue in multiple steps and achieves the best result both quantitatively and qualitatively. The last column in Fig.7 shows that our proposed

method resolves both phase unwrapping and RGB alignment.

6.2. Ablation Study

As an ablation study, we train our network with different configurations to show their impact on the final results. Detailed qualitative results on the individual setups can be found in the supplementary material. First, we trained with two different architectures that provide different bottleneck resolution (MonoDepth2[13]: 1/32 resolution bottleneck, ToFNet[33]: 1/4 resolution bottleneck). We compare with and without fusion to see the impact of implicit alignment with regard to the bottleneck resolution and the global context for phase unwrapping. Results in Tab.2 indicate that without fusion, using 1/32 resolution bottleneck (ToFNet, No, I-ToF) already shows an improvement over 1/4 resolution bottleneck (MD2, No, I-ToF). The improvement is mostly on the phase unwrapping parts which implies the importance of the global context for this task in challenging scenes. Unwrapping related artefacts are reduced sig-

Backbone	Input	RMSE	Rel(abs)	Rel(sqr)	1.25 ¹	1.25 ²	1.25 ³
ToFNet [33]	Raw I-ToF	0.99	0.13	0.23	87.5	94.6	97.1
MRM [15] + ToFNet [33]	Raw I-ToF	0.84	0.098	0.17	91.8	96.5	97.9
GuideNet [34]	Raw I-ToF + RGB	0.61	0.10	0.091	91.2	97.7	99.3
Ours	Raw I-ToF + RGB	0.56	0.073	0.069	94.6	98.4	99.3

Table 1. Quantitative comparison between ToFNet [33], MRM [15]+ToFNet [33], GuideNet [34] and our pipeline on the dataset [2]. (RMSE, Rel: The lower the better, 1.25ⁿ: The higher the better)

Backbone	Fusion	Branch	RMSE	Rel(abs)	Rel(sqr)	1.25 ¹	1.25 ²	1.25 ³
ToFNet [33]	No	I-ToF	0.99	0.13	0.23	87.5	94.6	97.1
	Bottleneck (implicit align)	I-ToF	0.75	0.098	0.14	92.1	97.0	98.0
		RGB	0.80	0.11	0.15	88.6	96.0	98.1
MD2 [13]	No	I-ToF	0.77	0.09	0.14	91.9	96.7	98.4
		RGB	0.75	0.16	0.17	78.5	95.0	98.6
	Bottleneck (implicit align)	I-ToF	0.58	0.079	0.095	94.0	98.1	99.1
		RGB	0.57	0.076	0.071	94.3	98.4	99.3
	Bottleneck (attention)	I-ToF	0.68	0.088	0.12	92.6	97.6	99.0
		RGB	0.68	0.11	0.11	88.2	97.3	99.0
	All resolutions (implicit + explicit align)	I-ToF	0.57	0.076	0.093	94.4	98.2	99.2
		RGB	0.56	0.073	0.069	94.6	98.4	99.3

Table 2. Ablation study of different pipeline configurations. We change the backbone and compare various fusion strategies.

nificantly, while the artefacts on the area where I-ToF is not reliable still remain. These can be removed with the additional RGB input. With fusion in the bottleneck, both architectures show an improvement, however, unwrapping related artefacts still exist after fusion with a 1/4 resolution bottleneck. Most of these artefacts are removed after fusion with the 1/32 resolution bottleneck. We argue that a smaller perceptive field combined with larger displacement in pixels at a higher resolution makes it harder to fuse and combine information in the former case.

For bottleneck fusion, we compare an attention based fusion which does not rely on an implicit alignment at low resolution. We first add a few layers on the bottleneck of RGB and I-ToF features and flatten them to obtain query, key, and value similar to [36]. Then we aggregate the values with a weight obtained from the dot product between cross modal query and key. Result in Tab.2 show that attention based fusion (MD2, Bottleneck(attention)) has some improvement over no fusion (MD2, No) while it does not perform as good as our implicit alignment strategy (MD2, Bottleneck(implicit align)). We believe that a huge domain difference between two features cannot be resolved with a few layers and a simple attention mechanism. In the supplementary material, we show details of the obtained attention map in three scenarios: I-ToF vs RGB, I-ToF vs I-ToF (temporal), RGB vs RGB (temporal) to support this argument.

At last, we train our network with and without explicit alignment to show the advantage of sharing higher resolution but aligned feature maps. As shown in Tab.2, quantitatively fusing higher resolution features (MD2, All resolutions) does not show a significant difference in compari-

son to bottleneck fusion (MD2, Bottleneck (implicit align)), however, it improves details such as texture copy artefacts which exist in the RGB depth branch, which we show this improvements later in the supplementary material.

7. Conclusion

We present a network architecture capable of dealing with realistic yet challenging scenarios for single frequency I-ToF where signals are acquired under strong ambient light with distances beyond the manufactures maximum depth range. Our pipeline uses both I-ToF and RGB information in a specially designed network which can incrementally resolve misalignments between the images while fusing their information from a low resolution. Our results show more than 40% RMSE improvement over the state of the art raw I-ToF depth pipeline and around 10% improvement over RGB guided depth completion. We strongly believe that this work can pave the way towards wider use of I-ToF in unconstrained environments and our fusion strategy of implicit alignment with follow-up explicit refinement can be a general contribution for spatially separated sensors fusion for depth estimation beyond ToF and RGB. Further work could exploit the inter-modal disparity of RGB and raw I-ToF (or depth respectively) to constrain the search space for the attention based fusion by epipolar geometry.

References

- [1] Gianluca Agresti, Henrik Schaefer, Piergiorgio Sartor, and Pietro Zanuttigh. Unsupervised domain adaptation for tof data denoising with adversarial learning. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5584–5593, 2019. 2, 3
- [2] Anonymous Authors. Anonymous title. In *Under Submission*, 2021. 6, 7, 8
- [3] Benjamin Busam, Marco Esposito, Simon Che’Rose, Nassir Navab, and Benjamin Frisch. A stereo vision approach for cooperative robotic movement therapy. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 127–135, 2015. 1
- [4] Benjamin Busam, Matthieu Hog, Steven McDonagh, and Gregory Slabaugh. Sterefo: Efficient image refocusing with stereo vision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 1
- [5] Ouk Choi, Hwasup Lim, Byongmin Kang, Yong Sun Kim, Keechang Lee, James D. K. Kim, and Chang-Yeong Kim. Range unfolding for time-of-flight depth cameras. In *2010 IEEE International Conference on Image Processing*, pages 4189–4192, 2010. 2, 4
- [6] Guanting Dong, Yueyi Zhang, and Zhiwei Xiong. Spatial hierarchy aware residual pyramid network for time-of-flight depth denoising. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 35–50, Cham, 2020. Springer International Publishing. 2, 3
- [7] Adrian A Dorrington, John P Godbaz, Michael J Cree, Andrew D Payne, and Lee V Streeter. Separating true range measurements from multi-path and scattering interference in commercial range cameras. In *Three-Dimensional Imaging, Interaction, and Measurement*, volume 7864, page 786404. International Society for Optics and Photonics, 2011. 2
- [8] David Droschel, Dirk Holz, and Sven Behnke. Multi-frequency phase unwrapping for time-of-flight cameras. pages 1463 – 1469, 11 2010. 2, 3
- [9] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *arXiv preprint arXiv:1406.2283*, 2014. 6
- [10] Stefan Fuchs. Multipath interference compensation in time-of-flight camera images. In *2010 20th International Conference on Pattern Recognition*, pages 3583–3586. IEEE, 2010. 2
- [11] Stefan Fuchs, Michael Suppa, and Olaf Hellwich. Compensation for multipath in tof camera measurements supported by photometric calibration and environment integration. In *International Conference on Computer Vision Systems*, pages 31–41. Springer, 2013. 2
- [12] Jason Geng. Structured-light 3d surface imaging: a tutorial. *Advances in Optics and Photonics*, 3(2):128–160, 2011. 1
- [13] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth prediction. October 2019. 4, 5, 6, 7, 8
- [14] John P Godbaz, Michael J Cree, and Adrian A Dorrington. Closed-form inverses for the mixed pixel/multipath interference problem in amcw lidar. In *Computational Imaging X*, volume 8296, page 829618. International Society for Optics and Photonics, 2012. 2
- [15] Qi Guo, Iuri Frosio, Orazio Gallo, Todd Zickler, and Jan Kautz. Tackling 3d tof artifacts through learning and the flat dataset. In *The European Conference on Computer Vision (ECCV)*, September 2018. 2, 4, 7, 8
- [16] Miles Hansard, Seungkyu Lee, Ouk Choi, and Radu Patrice Horaud. *Time-of-flight cameras: principles, methods and applications*. Springer Science & Business Media, 2012. 2, 4
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 5, 6
- [18] Aleksander Holynski and Johannes Kopf. Fast depth densification for occlusion-aware augmented reality. *ACM Transactions on Graphics (ToG)*, 37(6):1–11, 2018. 1
- [19] Mu Hu, Shuling Wang, Bin Li, Shiyu Ning, Li Fan, and Xiaojin Gong. Penet: Towards precise and efficient image guided depth completion. *arXiv preprint arXiv:2103.00783*, 2021. 3
- [20] Minh Hwang, Daniel Seita, Brijen Thananjeyan, Jeffrey Ichnowski, Samuel Paradis, Danyal Fer, Thomas Low, and Ken Goldberg. Applying depth-sensing to automated surgical manipulation with a da vinci robot. In *2020 International Symposium on Medical Robotics (ISMR)*, pages 22–29. IEEE, 2020. 1
- [21] David Jiménez, Daniel Pizarro, Manuel Mazo, and Sira Palazuelos. Modeling and correction of multipath interference in time of flight cameras. *Image and Vision Computing*, 32(1):1–13, 2014. 2
- [22] Ahmed Kirmani, Arrigo Benedetti, and Philip A Chou. Spumic: Simultaneous phase unwrapping and multipath interference cancellation in time-of-flight cameras using spectral methods. In *2013 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2013. 2
- [23] Seungkyu Lee, Byongmin Kang, James DK Kim, and Chang Yeong Kim. Motion blur-free time-of-flight range sensor. In *Sensors, Cameras, and Systems for Industrial and Scientific Applications XIII*, volume 8298, page 82980U. International Society for Optics and Photonics, 2012. 2
- [24] Lina Liu, Xibin Song, Xiaoyang Lyu, Junwei Diao, Mengmeng Wang, Yong Liu, and Liangjun Zhang. Fcfr-net: Feature fusion based coarse-to-fine residual learning for monocular depth completion. *arXiv preprint arXiv:2012.08270*, 2020. 3
- [25] Julio Marco, Quercus Hernandez, Adolfo Munoz, Yue Dong, Adrian Jarabo, Min H Kim, Xin Tong, and Diego Gutierrez. Deeptof: off-the-shelf real-time correction of multipath interference in time-of-flight imaging. *ACM Transactions on Graphics (ToG)*, 36(6):1–12, 2017. 3
- [26] B. Mildenhall, J. T. Barron, J. Chen, D. Sharlet, R. Ng, and R. Carroll. Burst denoising with kernel prediction networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2502–2510, 2018. 2
- [27] James Noraky and Vivienne Sze. Low power depth estimation of rigid objects for time-of-flight imaging. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(6):1524–1534, 2019. 1
- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming

- Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 6
- [29] Jann Poppinga and Andreas Birk. A novel approach to efficient error correction for the swissranger time-of-flight 3d camera. In Luca Iocchi, Hitoshi Matsubara, Alfredo Weitzenfeld, and Changjiu Zhou, editors, *RoboCup 2008: Robot Soccer World Cup XII*, pages 247–258, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. 2, 3
- [30] Di Qiu, Jiahao Pang, Wenxiu Sun, and Chengxi Yang. Deep end-to-end alignment and refinement for time-of-flight rgb-d modules. In *International Conference in Computer Vision (ICCV)*, 2019. 2, 3, 5
- [31] Mirko Schmidt, Klaus Zimmermann, and Bernd Jähne. High frame rate for 3d time-of-flight cameras by dynamic sensor calibration. In *2011 IEEE International Conference on Computational Photography (ICCP)*, pages 1–8. IEEE, 2011. 2
- [32] Kilho Son, Ming-Yu Liu, and Yuichi Taguchi. Learning to remove multipath distortions in time-of-flight range images for a robotic arm setup. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3390–3397. IEEE, 2016. 3
- [33] Shuochen Su, Felix Heide, Gordon Wetzstein, and Wolfgang Heidrich. Deep end-to-end time-of-flight imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6383–6392, 2018. 1, 2, 4, 6, 7, 8
- [34] Jie Tang, Fei-Peng Tian, Wei Feng, Jian Li, and Ping Tan. Learning guided convolutional network for depth completion. *IEEE Transactions on Image Processing*, 30:1116–1129, 2020. 3, 6, 7, 8
- [35] Wouter Van Gansbeke, Davy Neven, Bert De Brabandere, and Luc Van Gool. Sparse and noisy lidar completion with rgb guidance and uncertainty. In *2019 16th international conference on machine vision applications (MVA)*, pages 1–6. IEEE, 2019. 3
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 8
- [37] Shanshan Zhao, Mingming Gong, Huan Fu, and Dacheng Tao. Adaptive context-aware multi-modal network for depth completion. *IEEE Transactions on Image Processing*, 2021. 3