

# RADU: Ray-Aligned Depth Update Convolutions for ToF Data Denoising

Michael Schelling, Pedro Hermosilla, Timo Ropinski  
 Ulm University, Germany

{michael-1.schelling, pedro-1.hermosilla-casajus, timo.ropinski}@uni-ulm.de

## Abstract

*Time-of-Flight (ToF) cameras are subject to high levels of noise and distortions due to Multi-Path-Interference (MPI). While recent research showed that 2D neural networks are able to outperform previous traditional State-of-the-Art (SOTA) methods on correcting ToF-Data, little research on learning-based approaches has been done to make direct use of the 3D information present in depth images. In this paper, we propose an iterative correcting approach operating in 3D space, that is designed to learn on 2.5D data by enabling 3D point convolutions to correct the points' positions along the view direction. As labeled real world data is scarce for this task, we further train our network with a self-training approach on unlabeled real world data to account for real world statistics. We demonstrate that our method is able to outperform SOTA methods on several datasets, including two real world datasets and a new large-scale synthetic data set introduced in this paper.*

## 1. Introduction

Time of Flight (ToF) cameras are devices that are able to capture depth information of a scene by measuring the time the light emitted by the device needs to travel back once intersecting with an object. In practice, timing the reception of a light impulse requires precise and costly hardware and, as a result, consumer-level ToF cameras perform indirect measurements. The most common types are so-called Amplitude-Modulated Continuous-Wave (AMCW) ToF cameras, as they are for example used by the Kinect. The working principle of these AMCW cameras is to emit a periodically modulated light signal and retrieve the phase shift of the received signal, through which the travel time of the light and, consequently, the distance of the object to the camera is given [12]. However, the continuous illumination of the scene inevitably leads to Multi-Path Interference (MPI), as not only the direct reflection is received but also indirectly illumination which significantly impairs the depth estimation. Furthermore, these ToF cameras suffer

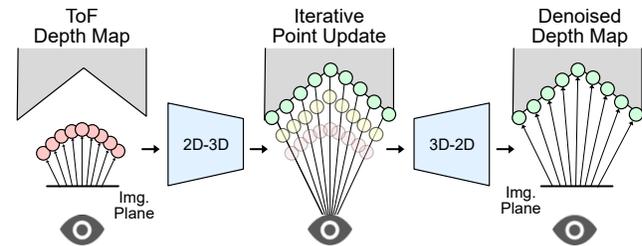


Figure 1. Given an initial ToF depth reconstruction our method projects the problem into a latent 3D space. The 3D point positions are updated iteratively along the camera rays using RADU point convolutions, in order to optimize the final 2D depth prediction.

from low Signal to Noise Ratios (SNR) on dark surfaces and the mixed pixel effect along sharp object edges [12].

Since deep learning approaches have shown great success in visual computing problems, many works have investigated the capabilities of 2D neural networks to correct ToF depth images [1, 3, 4, 10, 11, 29]. However, existing methods treat the task of ToF correction as a 2D image problem and do not take into account the explicit 3D information in their computations. In these works, the depth information is usually used as an input to standard Convolutional Neural Networks (CNN) for images [1, 3, 19, 27, 28], while the underlying 3D structure is not analyzed. In this work instead, we propose a new neural network architecture that projects the problem into the 3D domain and makes use of point convolutional neural networks [14] to analyze the erroneous reconstruction and adjust the point positions along the view direction, see Fig. 1. This iterative process makes small changes to the point positions to reduce the error level in between the convolutional layers and improve the final depth reconstruction. Further, we propose a novel fine-tuning procedure for Unsupervised Domain Adaptation (UDA) based on self-training methods, to transfer the knowledge acquired by our network from synthetic to real world ToF data. The effectiveness of this approach is evaluated on both synthetic and real datasets and proves to be able to outperform existing methods. Moreover, we introduce a large-scale high-resolution dataset consisting of challenging scenes, containing high MPI levels, materials which pro-

duce low SNR captures and objects with high frequency details. In summary, our contributions are:

- A novel architecture for correcting of ToF images in a latent 3D space.
- A two-staged training procedure with a cyclic self-training approach designed to bridge the gap between synthetic and real world ToF images.
- A large-scale high-resolution synthetic ToF dataset containing measurements for scenes with high MPI.

Our synthetic data set, code and trained networks are available at <https://github.com/schellmi42/RADU>.

## 2. Related Work

In this section we will briefly review existing work related to our approach in different fields.

**Learned ToF Correction.** Beginning with Marco *et al.* [19], several works on correcting ToF depth images using deep learning have been proposed. While the former uses a low-frequency ToF depth prediction as input, subsequent work by Agresti *et al.* [1–3] greatly improved the reconstruction by considering additional multi-frequency features derived from raw camera measurements to further reduce the error. The same year, Su *et al.* [29] proposed a generative approach that generates depths directly from raw measurements. Instead of predicting a corrected depth directly, Guo *et al.* [10] followed an inverse approach and used a 2D CNN to correct the raw camera measurements prior to the LF2 [33] depth reconstruction algorithm of the Kinect2. Further, several works on improving on various aspects and applications using machine learning followed, *e.g.* online calibration using RGB information [24], frame rate optimization [5], power efficiency [6], robotic arm setups [27] or translucent materials [28], to name a few. The aforementioned approaches all use standard 2D CNNs and thus consider the correction problem as an image task.

Recent work aims to estimate the real depth by reconstructing the transient image, *i.e.* the impulse response of the scene, through learning methods. Buratto *et al.* [4] used the assumption that the direct reflection reaches the camera sensor early and predict the intensity and arrival time of the first two peaks of the impulse response. The iToF2dToF method [11] first predicts ToF depths at various frequencies which are used to estimate the two leading coefficients of the Fourier-Transform of the impulse response. While transient reconstruction is a promising direction, these methods are, as of now, still very memory demanding, and can thus only query few pixels in parallel, and do not reach the correction capabilities of learned correction approaches.

In contrast to previous work, we propose to lift the problem into 3D and use 3D point CNNs. With this approach we follow the findings of recent research, that neural networks can benefit from latent 3D representations in various tasks,

such as object detection [31] using 3D voxel CNNs on lidar data, semantic segmentation [35] using 2.5D convolutions on RGB-D data or even image generation [22].

**Domain Adaptation for ToF Images.** Real ToF data with labeled ground truth data is only sparsely available, as its collection is rather complex. Thus various synthetic data sets have been introduced [3, 10, 19] to provide the amount of data needed to train deep neural networks. However the usage of data from a different source, in this case a synthetic simulation, comes at the cost of introducing a domain gap, *i.e.* real and synthetic images differ in their statistical properties. Consequently, one major challenge when learning the task of correcting ToF images is bridging the domain gap between synthetic data sets and real world data.

To improve performance on real data Marco *et al.* [19] pre-train an auto-encoder network on unlabeled real world ToF-data, and transfer the encoder layers to a encoder-decoder network for correction. Other works [1, 29] have suggested to use adversarial losses, where a second network acts as an adversarial agent which is trained to distinguish depths generated from real and synthetic data. During training the generator is optimized to deceive the discriminator.

Recently, self-training methods, which originated from a student-teacher approach by Hinton *et al.* [15], have shown great success in various variants of domain-adaptation [16–18, 20, 23, 34, 38, 39]. Building upon these works, we employ a cyclic self-training procedure to adapt our network to real world statistics by generating pseudo-labels for unlabeled real data, after pre-training the network on synthetic data.

**Point Cloud Denoising.** Several learned methods for denoising point clouds have been published [13, 25, 26, 36, 37]. In contrast to these approaches, which tackle general noise models, our approach aims to preserve the 2.5D structure of the data by performing 1D corrections. Further, MPI errors are not i.i.d., which, to our knowledge, is a typical assumption for unsupervised point cloud denoising methods [13].

## 3. Problem Statement

To estimate the distance  $d$  of an object, an AMCW ToF camera emits a light signal  $s_e$ , which is typically modulated by a sinusoidal periodic function, in the form of

$$s_e(t) = s_0 \cdot (1 + m \cdot \sin(2\pi ft)), \quad (1)$$

where  $s_0$  is the average intensity,  $m$  is the modulation coefficient,  $f$  is the modulation frequency, and  $t$  is the time. For compactness and *w.l.o.g.* we assume  $s_0 = 1, m = 1$ . In the optimal case of only direct reflection the received light  $s_r$  is a scaled and phase shifted version of the emitted signal

$$s_r(t) = r \cdot s_e(t - \Delta t) = r \cdot (1 + \sin(2\pi ft - \Delta\varphi)), \quad (2)$$

where  $r$  is the ratio of the light backscattered to the sensor from the surface, and  $\Delta\varphi$  is the phase delay after the signal has traveled the distance  $2d$ , *i.e.*  $\Delta\varphi = 2dc/f$ , where  $c$

is the speed of light. The received signal is averaged during the exposure time  $\delta t$  at the sensor with a phase shifted version of the emitted signal, resulting in the measurement

$$m_\theta = \frac{1}{\delta t} \int_{\delta t} s_r(t) \cdot s_e \left( t + \frac{\theta}{2\pi f} \right) dt. \quad (3)$$

Under the assumption  $\delta t \gg 1/f$  the measurement  $m_\theta$  can be approximated as [7]

$$m_\theta = I + A \cdot \cos(\Delta\varphi + \theta), \quad (4)$$

where  $I$  is called the intensity and  $A$  the amplitude. By measuring  $m_\theta$  for multiple phase offsets  $\theta$  the phase shift  $\Delta\varphi$  and the distance  $d$  can be reconstructed as [12]

$$\Delta\varphi = \arctan \left( \frac{\sum_\theta -\sin(\theta) \cdot m_\theta}{\sum_\theta \cos(\theta) \cdot m_\theta} \right), \quad (5)$$

$$d = \frac{c \cdot \Delta\varphi}{4\pi f}. \quad (6)$$

Due to the periodic nature of the signal, the reconstructed distance  $d$  using Eq. (6) is ambiguous for distances larger than  $d_{max} = c/(2f)$ . To resolve this so called phase-wrapping the common solution is to acquire measurements for different modulation frequencies  $f_k$ , which thus also have different maximal distances  $d_{max}$ .

In practice the received signal  $s_r$  is not only the direct reflection after the time  $\Delta t$ , as in Eq. (2), but a composition of light scattered along various paths  $P$  in the scene

$$s_r(t) = \int_P r(p) \cdot (1 + \sin(2\pi f t - \Delta\varphi(p))) dp. \quad (7)$$

While the intensity  $r(p)$  can be expected to be low after multiple reflections on an isolated path  $p$ , the accumulation over all possible paths  $P$  leads to the aforementioned notable MPI distortion in the distance recovery, Eq. (6). Additionally, ToF-sensors suffer from the common camera noise sources in the form of photon shot noise, thermal shot noise and read noise, which are typically modeled jointly as an additive Gaussian noise on the measurement  $m_\theta$ .

Although,  $d$  from Eq. (6) denotes the distance, it is commonly referred to as ToF depth, as we will do in this work.

## 4. Proposed Method

In this work we propose to incorporate 3D learning techniques to exploit the spatial structure of the scene geometry. In contrast to 2D convolutions, where spatial relations are encoded implicitly in the pixel grid location, a 3D convolution relies on explicitly given 3D coordinates [14]. In order to represent the 2.5D data in 3D space we use the inverse of the camera projection  $\mathcal{P}_{G \rightarrow C}$  to create an initial spatial position for each pixel, as illustrated in Fig. 2. This allows us

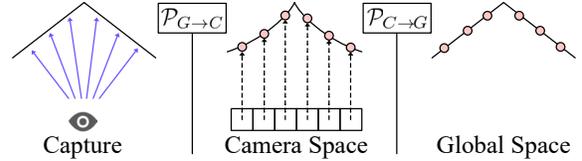


Figure 2. The non-parallelism of camera rays (left) leads to a distortion in the resulting depth image, as the camera measures the distance along the rays (middle). By applying the inverse camera projection  $\mathcal{P}_{C \rightarrow G}$  the depth map is projected into 3D space, and, in the ideal case, aligns with the scene geometry (right).

to correct the depth in global coordinates, allowing the 3D layers to consider the 3D neighborhoods of the individual pixels, and learn on the actual scene geometry, undistorted by the non-linear camera transformation.

In the case of ToF-data we derive an initial 3D coordinate from the ToF depth given by Eq. (6). Of course this initial depth is inherently erroneous, and we assume that a 3D network would benefit from corrected 3D positions, which lie closer to the scene geometry. Thus, we propose to iteratively update the 3D positions of the points, enabling the network to optimize the latent point cloud in between the networks 3D layers. To keep the point position in alignment with the underlying structure given by the pixel position, we restrict the points movement to the respective camera ray.

However, compared to 2D Convolutions, 3D Convolutions are more demanding in compute power and memory consumption. To reduce this load we embed our proposed 3D layers in a 2D Network and introduce 2.5D pooling layers to reduce the spatial resolution in the 3D convolutions. This is necessary as for example a  $320 \times 240$  image would result in 76.8k points, for comparison common networks for 3D object recognition typically use 1k points per model.

In the following we will first define the 2.5D pooling and the RADU convolution layers, before describing our network architecture, the loss function and the self-cycled training procedure for domain adaptation.

### 4.1. 2.5D Pooling

Several methods have been proposed for pooling operations on 3D point clouds, e.g. Poisson disk sampling [14] and cell average sampling [30]. However these methods suffer from drawbacks when applied to 2.5D data. Cell average sampling creates a new point as the mean point in a pre-defined 3D grid, which in general does not lie on the camera ray of a pixel. Poisson disk sampling chooses a subset of the given points based on their distances in 3D space, making these points still align with the original camera-rays. However this method aims to optimize the distribution of points in 3D space, which can result in non-uniform distributions when projected back into 2D space.

Instead, we propose a 2.5D pooling operation which

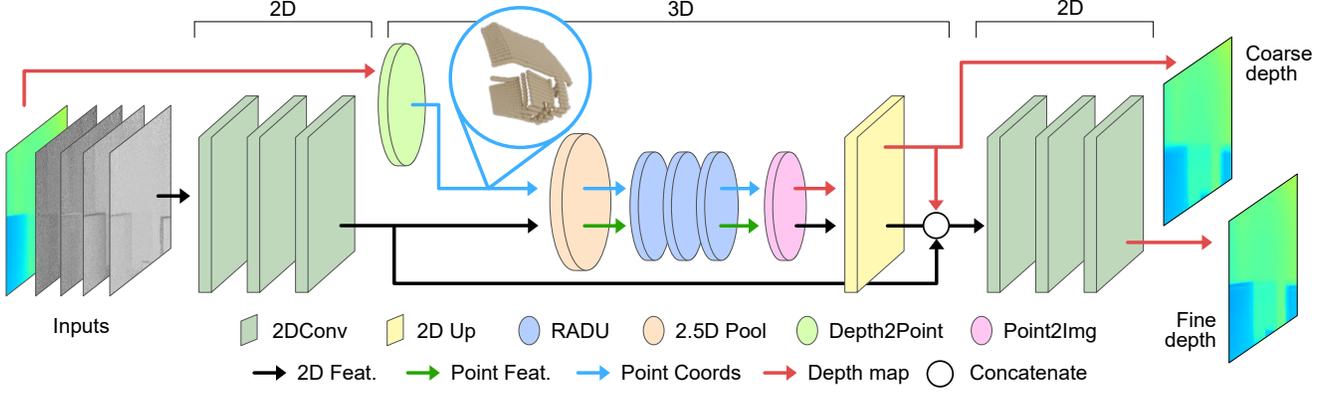


Figure 3. Proposed network architecture. We embed the 3D RADU convolutions in two 2D blocks and use an initial ToF depth given by Eq. (6) to compute initial spatial positions for the 2D features. After the 3D block the point coordinates updated by the RADU layers are projected back to a coarse depth image, which is used as an additional input feature to the following 2D block.

pools the points in 3D according to their structure induced by the pixel grid, by considering the points’ 2D neighborhoods given implicitly through the pixel grid and the camera ray direction associated with the corresponding pixel. For this we pool only the depth value of the 3D points of a  $k \times k$  2D patch using an order invariant operation, *e.g.* maximum or average pooling, and place a point with the pooled depth on the respective camera ray at the desired reduced resolution  $N/k \times N/k$ . The feature of the resulting point is computed by performing a second pooling operation on the features inside the points 2D neighborhood.

This way each resulting point represent a 2D neighborhood of fixed size and aligns with the camera rays. Note that this is different from (a) pooling the three point coordinates directly, where the result is not necessarily aligned with the camera rays, and (b) 2D pooling on the distance image, as the projections  $\mathcal{P}_{G \rightarrow C}$  and  $\mathcal{P}_{C \rightarrow G}$  are non-linear, see Fig. 2. Furthermore, this allows us to exploit existing 2D Pooling operations, dropping the need for costly 3D neighborhood and sampling computations for pooling.

## 4.2. Ray-Aligned Depth Update

In order to optimize the point position during the network execution, we propose to extend a given 3D convolution to predict an additional feature channel per layer. This value is used to shift the point along the respective camera ray, which we call Ray-Aligned Depth Update (RADU) convolution. By updating along the camera ray, the point position stays consistent to the 2.5D nature of the point cloud, see Fig. 1. In order to stabilize the depth updates we regularize the update  $p^u$  to the range  $(-\alpha, \alpha)$  using a scaled tanh function. In our experiments we set  $\alpha = 0.1$ , and provide an ablation study on different values for this hyperparameter in the supplementary. Formally, given a 3D convolution operator  $\text{conv}$ , a point cloud  $\{p_i^{in}\} \subset \mathbb{R}^3$  with associated camera rays  $\{r_i\} \subset \mathbb{R}^3$ , and input features  $\{f_i^{in}\} \subset \mathbb{R}^{C_{in}}$ , the RADU convolution on point  $p_j^{in}$  with neighborhood  $\mathcal{N}_j \subset \mathbb{N}$  is given as

$$(f_j^{out}, p_j^u) = \text{conv}(\{f_i^{in}\}_{i \in \mathcal{N}_j}, \{p_i^{in}\}_{i \in \mathcal{N}_j}), \quad (8)$$

$$p_j^{out} = p_j^{in} + \alpha \cdot \tanh(p_j^u) \cdot r_j, \quad (9)$$

where  $f_j^{out}$  is the output feature of the point  $p_j^{out}$ .

This extension is independent of the type of 3D-convolution used, in our experiments we implement it based on the Monte-Carlo-Convolution [14], which is illustrated in Fig. 4 and described in more detail in the supplementary. This type of 3D convolution has been shown to work well with non-uniformly sampled point clouds, as in our case where the density varies with the distance of the points to the camera, which we also investigate in Sec. 6.4.

## 4.3. Network Architecture

Our network architecture is illustrated in Fig. 3. It consists of an initial stack of three 2D convolutions with kernel size  $3 \times 3$ , which is followed by a 2.5D pooling layer, using

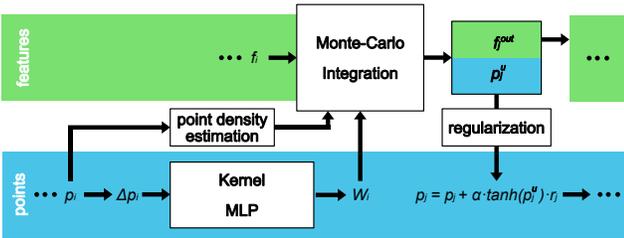


Figure 4. Realization of our RADU convolutional layer based on a 3D Monte-Carlo Convolution. In contrast to standard 3D point convolutions the spatial positions are not only used to predict the kernel weights but are also updated after each convolution.

average pooling on both depth and features, with a stride of  $8 \times 8$ . This 2D block is followed by a stack of three 3D RADU convolutions with increasing receptive fields of  $0.1m, 0.2m$ , and  $0.4m$ . After the 3D block we use bilinear upsampling to increase the spatial resolution of the intermediate features and the updated depth values. The coarse depth prediction of the 3D block is projected back into 2D using the inverse camera transform. Both the up-scaled depth and features are processed by a second block of three 2D convolutions with kernel size  $3 \times 3$ . We further introduce a skip connection between the 2D blocks. In between convolutions we use leaky ReLU as a non-linearity.

As input to our method we use the same multi-frequency ToF features used in the works of Agresti *et al.* [1, 3]. That is, given measurements at three different frequencies  $f_k$ ,  $k = 1, 2, 3$ , we use five input features,  $d_1, d_2 - d_1, d_3 - d_1, A_2/A_1 - 1$ , and  $A_3/A_1 - 1$ , where  $d_k$  is the ToF depth and  $A_k$  is the amplitude at frequency  $f_k$ . We follow the argumentation of Agresti *et al.* [3], that the differences  $d_1 - d_k, k = 2, 3$  encode the frequency dependent influence of MPI on the depth recovery, and the relative amplitudes  $A_1/A_k, k = 2, 3$  provide information about the strength of the MPI for a given pixel. The depth  $d_1$  further serves as initial ToF depth for the projection  $\mathcal{P}_{C \rightarrow G}$  into 3D space. Further, we augment the input data using rotation, mirroring and noise, a more detailed description of the data augmentation can be found in the supplementary.

#### 4.4. Coarse-Fine Loss

The 3D block of our network architecture produces an intermediate coarse depth estimate which is fed as an additional input to the subsequent 2D layers. To guide the network to predict an adequate representation of the 3D geometry we optimize both the final output  $\hat{d}_{out}$  and the coarse 3D representation of the 3D blocks  $\hat{d}_{3D}$ :

$$\mathcal{L} = \|d_{gt} - \hat{d}_{out}\|_1 + \|d_{gt} - \hat{d}_{3D}\|_1. \quad (10)$$

As the coarse depth  $\hat{d}_{3D}$  is not predicted by the final layer of the 3D network but is constructed iteratively as

$$(\hat{d}_{3D})_j = \mathcal{P}_{G \rightarrow C} \left( d_j^{init} + \sum_{l=1}^3 \alpha \cdot \tanh(p_{j,l}^u) \cdot r_j \right), \quad (11)$$

this allows each RADU layer to receive gradients directly from the loss function, preventing vanishing gradients, which is comparable to the influence on gradient flow of skip connections. However, as each layer is optimized to produce correctly denoised depths, this also increases the risk of overfitting. Ablations on the loss function and the choice of  $\alpha$  are provided in the supplementary.

#### 4.5. Unsupervised Domain Adaptation

To improve the performance of our method on real data we investigate a cyclic self-training procedure, derived from

---

#### Algorithm 1 Adapted Cycled Self-Training Procedure

---

**Require:**  $n_{cycle} \in \mathbb{N}, p \in [0, 1], D_{real}, D_{syn}$ , network  $N$

**for** *epoch* **do**

**if** *epoch* mod  $n_{cycle} \equiv 0$  **then**

$F_{in} \leftarrow D_{real}$  ▷ Get real world data

$\hat{d}_{out} = N(F_{in})$  ▷ predict, unaugmented

$S_2 \leftarrow \{d_{gt} : \hat{d}_{out}\}$  ▷ Save pseudo labels

**end if**

**for** *training step* **do**

**if**  $rand.unif([0, 1]) < p$  **then**

$(F_{in}, d_{gt}) \leftarrow D_{real}$  ▷ Real with pseudo label

**else**

$(F_{in}, d_{gt}) \leftarrow D_{syn}$  ▷ Synthetic with label

**end if**

$(F_{in}, d_{gt}) = augment(F_{in}, d_{gt})$

$(\hat{d}_{out}, \hat{d}_{3D}) = N(F_{in})$

*minimize*  $\mathcal{L}$ , Eq.(10), on  $(\hat{d}_{out}, \hat{d}_{3D})$

**end for**

**end for**

---

existing self-training methods for other tasks [16–18]. We evaluate a network, which is pre-trained on synthetic data, on unlabeled real data and use the predictions as pseudo labels in the following training phase. During training we choose randomly between synthetic data with labels and real data with pseudo-labels, to prevent the network from overfitting to the pseudo-labels. We repeat this process multiple times by updating the pseudo-labels every  $n_{cycle}$  epochs. To avoid providing the exact same input during pseudo label generation and training we create pseudo-labels on unaugmented data and use augmented input during training. The procedure is summarized in Algorithm 1.

We refrain from training a teacher network on labeled real data for pseudo-label generation to keep the assumption of U-DA and to make our approach applicable to settings where no labeled real data is available.

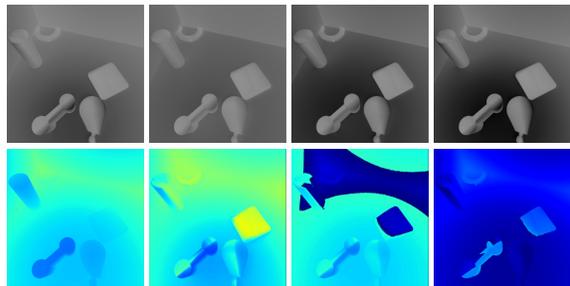


Figure 5. Example scene from our Cornell-Box dataset. The top row shows the four measurements  $m_\theta$  at 20MHz, bottom row shows, from left to right, ground truth depth and ToF depths using Eq. (6) at 20MHz, 50MHz, and 70MHz, with phase wrapping.

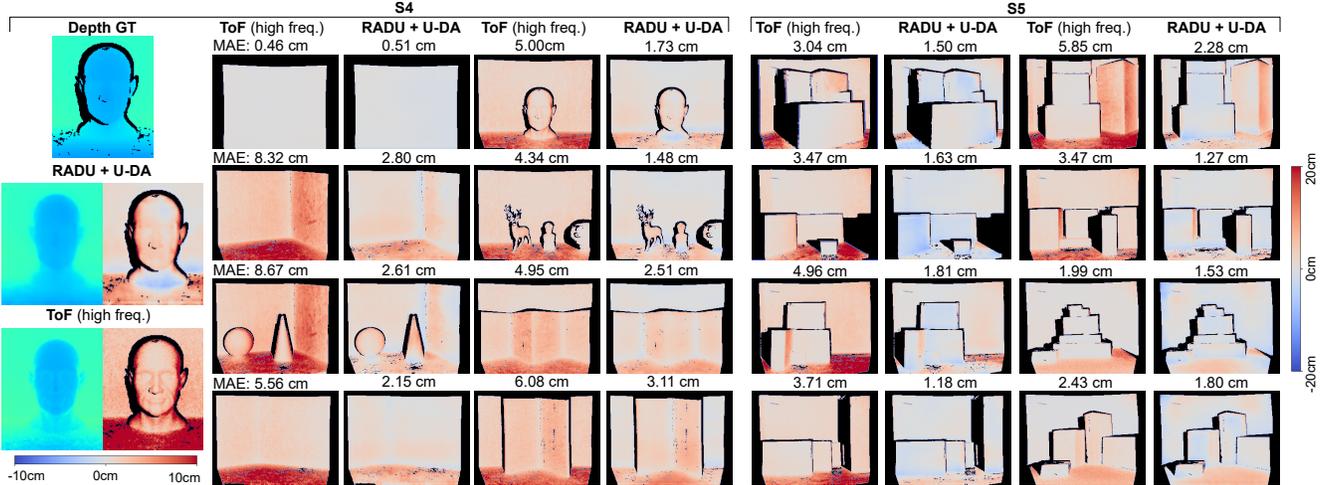


Figure 6. Depth error maps with MAE on the real datasets S4 (left) and S5 (right) of our method using both RADU convolutions and cyclic self-training for unsupervised domain adaptation. Areas where no ground truth depth is available are masked in black. The zoom in on the left displays both depth and error, with enhanced color coding on the error.

## 5. Cornell-Box ToF Dataset

To allow a more broad evaluation of our method, we generate a new large scale synthetic dataset. As recent advances in ToF-camera hardware [21] could allow for higher resolution ToF sensors in the future, we render our dataset at a comparably high resolution of  $600 \times 600$  pixels. Inspired by Miller *et al.* [21] we simulate the properties of a Raspberry Pi 3 camera equipped with an EAM for modulation.

To create a challenging setup for ToF correction we generate 142 scenes inspired by the Cornell box layout [9], to ensure high levels of MPI. We render each scene from 50 viewpoints with 3 different material properties, including dark materials for low SNR values, which results in 21.3k different renderings. The data is split into 116 training, 13 validation, and 13 test scenes. Each rendering is processed to simulate raw measurements for sinusoidal modulations at 20MHz, 50MHz, and 70MHz at phase offsets at  $0, \pi/2, \pi,$  and  $3\pi/2$ , resulting in 12 correlation measurements. To our knowledge our dataset exceeds existing ToF datasets in both

Dataset	Type	GT	$m_\theta$	Size	Resolution
S1 [3]	Syn.	Yes	No	54	$320 \times 240$
S2 [1]	Real	No	No	96	$320 \times 239$
S3 [1]	Real	Yes	No	8	$320 \times 239$
S4 [1]	Real	Yes	No	8	$320 \times 239$
S5 [3]	Real	Yes	No	8	$320 \times 239$
FLAT [10]	Syn.	Yes	Yes	1.2k	$424 \times 512$
Cornell-Box	Syn.	Yes	Yes	21.3k	$600 \times 600$

Table 1. Properties of the datasets considered in our experiments.

size and resolution, see Table 1. We refer the reader to the supplementary for further details about the dataset.

## 6. Experiments

First we test our method on the established data sets of Agresti *et al.* [1,3], who provide a synthetic and several real data sets for the SoftKinect camera, containing both MPI and shot noise, to evaluate our method including U-DA. We refer to these data sets using the same notation as previous authors [1,4], see Tab. 1. As the real datasets are rather small, we further evaluate our method on larger synthetic datasets, our dataset described in Sec. 5, where we additionally consider phase wrapping in the input features, and the FLAT dataset [10], which simulates a Kinect2 sensor, to also investigate the influence of non-sinusoidal modulations of  $s_e(t)$ , which violate the assumptions from Sec. 3.

Further we perform an ablation experiment, comparing RADU layers to other convolution layers. For additional information about hyperparameter settings and other aspects of the experiments we refer to the supplementary.

### 6.1. Experiment 1: Real World Data

The dataset S1-S5 contain intensities, amplitudes and phase unwrapped ToF depths. Similar to Agresti *et al.* [1] we train our network on the synthetic dataset S1 and use the real dataset S3 for validation. In the second stage we use the unlabeled real dataset S2 in our cycled self-training procedure for U-DA. We evaluate our method with and without U-DA on the two real data sets S4, with lower MPI levels but more detailed objects, and S5, a 'box' dataset with higher MPI levels but fewer details. We compare to the Coarse-Fine-Network (CFN), with [2] and without

Method	S4		S5	
	MAE [cm]	Relative Error	MAE [cm]	Relative Error
ToF (low freq.)	7.28	-	5.06	-
ToF (high freq.)	5.43	-	3.62	-
SRA [8] <sup>†</sup>	5.11	94.1%*	3.37	93.1%*
DeepToF [19] <sup>†</sup>	5.13	70.5%*	6.68	132%*
+calibration [3] <sup>†</sup>	5.46	75.0%*	3.36	66.4%*
TIR [4] <sup>†</sup>	2.60	47.9%*	1.88	52.0%*
CFN [3] <sup>†</sup>	3.19	58.7%*	2.22	60.5%*
+ U-DA [2]	2.31	42.5%*	1.64	45.3%*
RADU	<b>1.83</b>	<b>33.7%*</b>	2.59	71.5%*
RADU + U-DA	2.11	38.8%*	<b>1.63</b>	<b>45.0%*</b>
RADU + S-DA	1.89	34.8%*	1.53	42.3%*

Table 2. Results of various methods on the real world datasets S4 and S5. Each row reports the MAE and the relative error with respect to the phase unwrapped high frequency ToF depth. (\*: relative to low frequency, †: numbers taken from Buratto *et al.* [4])

U-DA [3], DeepToF [19], Transient Image Reconstruction (TIR) [4], and the non-learned Sparse Reconstruction Analysis (SRA) [8]. The Mean Absolute Error (MAE) and the remaining relative error are reported in Tab. 2.

The combination of our RADU network with cycled self-training for U-DA outperforms existing approaches on both datasets and successfully removes noise and MPI from the unwrapped ToF depth images, as can be seen in Fig. 6. To evaluate the stability of the cyclic self-training we repeat the U-DA of our method 10 times and measured a standard deviation of the MAE at 0.072cm on S4 and 0.021cm on S5, which indicates that the proposed U-DA approach is stable. During our experiments, we noticed that the RADU’s performance correlates on the sets {S1, S4} and {S2, S3, S5}, which is indicated by the drop of performance on the dataset S4 after U-DA. For comparison, we also report the results of our method after a Supervised Domain Adaptation (S-DA) using the small labeled real dataset S3 for fine-tuning.

## 6.2. Experiment 2: Cornell-Box Dataset

We train instances of the previously mentioned CFN, DeepToF and our RADU network on our Cornell-Box dataset described in Sec. 5. Since our dataset contains raw measurements  $m_\theta$ , we further compare to the End2End network [29]. For a fairer comparison, we perform a hyperparameter tuning for each method, we refer to the supplementary for details. As our dataset is purely synthetic we do not use domain adaptation strategies.

We evaluate on the test images where our method achieves the lowest MAE of the mentioned methods as can be seen in Tab. 3. The remaining relative error is comparable to the previous experiment, showing that our network,

Method	Cornell-Box		FLAT	
	MAE [cm]	Relative Error	MAE [cm]	Relative Error
ToF (low freq.)	29.0	-	59.34	-
ToF (high freq.)	11.14	-	-	-
DeepToF [19]	10.17	35.1%*	23.0	38.8%*
CFN [1, 3]	3.99	35.8% <sup>†</sup>	6.29	10.6%*
End2End [29]	5.99	53.8% <sup>†</sup>	6.20	10.5%*
RADU	<b>3.64</b>	<b>32.7%<sup>†</sup></b>	<b>3.31</b>	<b>5.58%*</b>

Table 3. Results of various methods on unseen data from our synthetic Cornell-Box dataset and the FLAT dataset. Each row reports the depth MAE and the relative error with respect to a phase unwrapped ToF depth. (\*: relative to low frequency, †: relative to high frequency)

as well as CFN, are able to handle phase wrapping in the input features. The iterative correction in the latent 3D space in between the three RADU convolutions is shown in Fig. 8.

However, while our method yields better results than the others, objects with low SNR and scenes with high MPI can still lead to failures in the reconstruction, as shown in Fig. 7.

## 6.3. Experiment 3: FLAT Dataset

In a third experiment we train our method on the FLAT dataset which contains nine raw correlations for three fre-

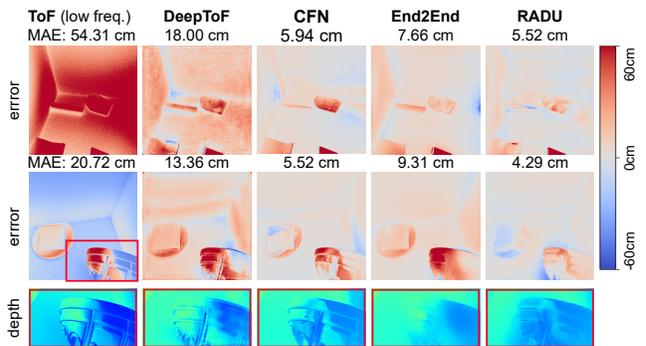


Figure 7. Visual results for two challenging scenes from our Cornell-Box dataset. The top scene contains an object with low SNR at the bottom, where all methods fail to retrieve the correct depth. The bottom scene exhibits high MPI, and shows a failure case of our method, where the object boundaries are blurred.

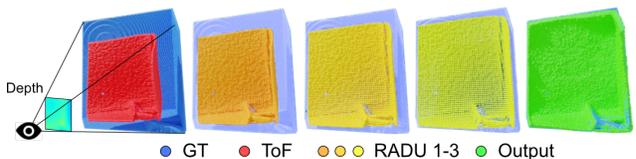


Figure 8. Visualization of the latent point clouds on a corner scene. The depth reconstruction improves after every RADU layer.

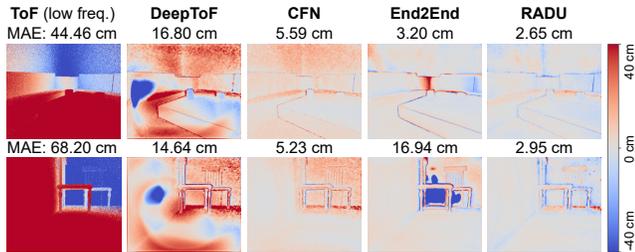


Figure 9. Error maps on the FLAT dataset. The top row shows an example where both RADU and End2End produce an almost error free prediction. The bottom row shows a failure case of End2End.

quencies, simulating the Kinect2. Unlike in the previous cases the low frequency signal is non-sinusoidal [10] which introduces additional distortions in the depth estimation using Eq. (6). A further challenge is the domain gap between the training data, to a large part consisting of images of isolated floating objects and thus low MPI levels, and the test data, containing complete scenes. We compare our method to the same methods as in the previous experiment, again performing a hyperparameter search for each method.

We evaluate on the 120 test images, where our method achieves the lowest MAE compared to the other learned methods, see Tab. 3. In Fig. 9 we show error images for the different methods. We recognize that End2End can produce competitive results in some cases but has a tendency to create artifacts, which we assume to stem from the domain gap between training and test data.

#### 6.4. Ablation: Latent 3D Representation

To validate the benefits of our 3D RADU convolutions we evaluate the performance when replacing the RADU convolutions with 2D, 2.5D or 3D point convolutions. In detail we compare to 2.5D convolutions [35], which performs three 2D convolutions on separate depth ranges, KPConv [30] which uses spherical kernel points to represent the 3D kernel function, MConv [14], which uses kernel MLPs and a density estimation, and PointConv [32] which

Method	Training (S1) MAE [cm]	Validation (S3) MAE [cm]
2D Conv	9.68	2.72
2.5D Conv [35]	8.79	2.61
3D KPConv [30]	10.86	4.00
3D PointConv [32]	10.49	3.42
3D MConv [14]	8.38	2.51
3D MConv + RADU	7.87	2.28

Table 4. Results of our network architecture with different layer types in the latent block. We report MAE on training (synthetic) and validation (real) data after hyperparameter optimization.

uses a kernel MLP with a learned density estimation. We conduct a smaller experiment without domain adaption, using S1 for training and S3 for validation. For a fair comparison we conduct a hyperparameter search for each method, and report results of the validation MAE in Tab. 4.

The results show that latent 3D information can help to improve the performance, but the choice of the convolution type is critical. Both the 2.5D convolution and MConv yield better results than the 2D convolution, on both training and validation data, outperformed only by our proposed RADU extension of the MConv layer.

## 7. Limitations

While we covered several error sources present in ToF data in our experiments, we did not investigate motion artifacts that occur in dynamic scenes [12], which would require additional considerations in the network architecture as shown in previous works [10, 24]. However, we believe that the latent corrected 3D point cloud representation of our network can potentially be used to improve image alignment. Further, the results on our Cornell-Box dataset indicate that high MPI, low SNR and fine details can have a drastic impact in the depth estimation, not only for our method but for all evaluated methods. Finally learned methods, including ours, are trained for specific sensor data and are thus not able to generalize to for example different modulation frequencies. We believe learned methods which are able to treat sensor properties as input parameters, as in traditional approaches [8], is a promising line of research.

## 8. Conclusion

In this work we presented an extension of unstructured 3D convolutions for ToF correction, which exploits structured information from depth images to iteratively correct point clouds in 3D space. The experiments indicate that latent 3D representations improve the correction capabilities of neural networks for various error sources present in ToF data. Further, we demonstrated that cyclic self-training using pseudo-labels can effectively be used for unsupervised domain adaptation on ToF data and, applied to our network, outperforms existing methods on two real data sets.

While we demonstrated our RADU layers in the context of ToF correction, they can in principle be applied to any task where 2.5D information is present, in order to benefit from both a latent 3D representation and an iterative correction of the former.

## 9. Acknowledgements

This project is financed by the Baden-Württemberg Stiftung gGmbH. We thank Markus Miller and Rainer Michalzik for their insights on the AMCW ToF working principle and the noise simulation.

## 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. 1, 2, 5, 6, 7
- [2] Gianluca Agresti, Henrik Schaefer, Piergiorgio Sartor, Yalcin Incesu, and Pietro Zanuttigh. Unsupervised domain adaptation of deep networks for ToF depth refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2, 6, 7
- [3] Gianluca Agresti and Pietro Zanuttigh. Deep learning for multi-path error removal in ToF sensors. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, September 2018. 1, 2, 5, 6, 7
- [4] Enrico Buratto, Adriano Simonetto, Gianluca Agresti, Henrik Schäfer, and Pietro Zanuttigh. Deep learning for transient image reconstruction from ToF data. *Sensors*, 21(6):1962, 2021. 1, 2, 6, 7
- [5] Yan Chen, Keyuan Qian, Xuanye Cheng, and Jingkun Zhou. A learning method to optimize depth accuracy and frame rate for time of flight camera. In *IOP Conference Series: Materials Science and Engineering*, volume 563, page 042067. IOP Publishing, 2019. 2
- [6] Yan Chen, Jimmy Ren, Xuanye Cheng, Keyuan Qian, Luyang Wang, and Jinwei Gu. Very power efficient neural time-of-flight. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2257–2266, 2020. 2
- [7] Mario Frank, Matthias Plaue, Holger Rapp, Ullrich Köthe, Bernd Jähne, and Fred A Hamprecht. Theoretical and experimental error analysis of continuous-wave time-of-flight range cameras. *Optical Engineering*, 48(1):013602, 2009. 3
- [8] Daniel Freedman, Yoni Smolin, Eyal Krupka, Ido Leichter, and Mirko Schmidt. SRA: Fast removal of general multipath for ToF sensors. In *European Conference on Computer Vision*, pages 234–249. Springer, 2014. 7, 8
- [9] Cindy M Goral, Kenneth E Torrance, Donald P Greenberg, and Bennett Battaile. Modeling the interaction of light between diffuse surfaces. *ACM SIGGRAPH computer graphics*, 18(3):213–222, 1984. 6
- [10] 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. 1, 2, 6, 8
- [11] Felipe Gutierrez-Barragan, Huaijin Chen, Mohit Gupta, Andreas Velten, and Jinwei Gu. iToF2dToF: A robust and flexible representation for data-driven time-of-flight imaging. *arXiv preprint arXiv:2103.07087*, 2021. 1, 2
- [12] Miles Hansard, Seungkyu Lee, Ouk Choi, and Radu Patrice Horaud. *Time-of-flight cameras: principles, methods and applications*. Springer Science & Business Media, 2012. 1, 3, 8
- [13] Pedro Hermosilla, Tobias Ritschel, and Timo Ropinski. Total denoising: Unsupervised learning of 3D point cloud cleaning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 52–60, 2019. 2
- [14] Pedro Hermosilla, Tobias Ritschel, Pere-Pau Vázquez, Àlvar Vinacua, and Timo Ropinski. Monte Carlo convolution for learning on non-uniformly sampled point clouds. *ACM Transactions on Graphics (TOG)*, 37(6):1–12, 2018. 1, 3, 4, 8
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [16] Ananya Kumar, Tengyu Ma, and Percy Liang. Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning*, pages 5468–5479. PMLR, 2020. 2, 5
- [17] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896, 2013. 2, 5
- [18] Hong Liu, Jianmin Wang, and Mingsheng Long. Cycle self-training for domain adaptation. *arXiv preprint arXiv:2103.03571*, 2021. 2, 5
- [19] 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. 1, 2, 7
- [20] Alexander Mey and Marco Loog. A soft-labeled self-training approach. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 2604–2609. IEEE, 2016. 2
- [21] Markus Müller, Hongwang Xia, Mina Beshara, Susanne Menzel, Karl Joachim Ebeling, and Rainer Michalzik. Large-area transmission modulators for 3d time-of-flight imaging. In *Unconventional Optical Imaging II*, volume 11351, page 113511F. International Society for Optics and Photonics, 2020. 6
- [22] Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7588–7597, 2019. 2
- [23] Viraj Prabhu, Shivam Khare, Deeksha Kartik, and Judy Hoffman. Sentry: Selective entropy optimization via committee consistency for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8558–8567, 2021. 2
- [24] Di Qiu, Jiahao Pang, Wenxiu Sun, and Chengxi Yang. Deep end-to-end alignment and refinement for time-of-flight RGB-D module. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9994–10003, 2019. 2, 8
- [25] Marie-Julie Rakotosaona, Vittorio La Barbera, Paul Guerrero, Niloy J Mitra, and Maks Ovsjanikov. POINTCLEAN-NET: Learning to denoise and remove outliers from dense point clouds. In *Computer Graphics Forum*, volume 39, pages 185–203. Wiley Online Library, 2020. 2
- [26] Riccardo Roveri, A. Cengiz Öztireli, Ioana Pandele, and Markus Gross. Pointprone: Consolidation of point clouds with convolutional neural networks. In *Computer Graphics Forum*, volume 37, pages 87–99. Wiley Online Library, 2018. 2

- [27] 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. 1, 2
- [28] Seongjong Song and Hyunjung Shim. Depth reconstruction of translucent objects from a single time-of-flight camera using deep residual networks. In *Asian Conference on Computer Vision*, pages 641–657. Springer, 2018. 1, 2
- [29] 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, 7
- [30] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. KPConv: Flexible and deformable convolution for point clouds. *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 3, 8
- [31] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudolidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8445–8453, 2019. 2
- [32] Wenxuan Wu, Zhongang Qi, and Li Fuxin. PointConv: Deep convolutional networks on 3D point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2019. 8
- [33] Lingzhu Xiang, Florian Echtler, Christian Kerl, Thiemo Wiedemeyer, Lars, hanyazou, Ryan Gordon, Francisco Facioni, laborer2008, Rich Wareham, Matthias Goldhoorn, alberth, gaborpapp, Steffen Fuchs, jmtatsch, Joshua Blake, Federico, Henning Jungkurth, Yuan Mingze, vinouz, Dave Coleman, Brendan Burns, Rahul Rawat, Serguei Mokhov, Paul Reynolds, P.E. Viau, Matthieu Fraissinet-Tachet, Ludique, James Billingham, and Alistair. libfreenect2: Release 0.2, Apr. 2016. 2
- [34] Sang Michael Xie, Ananya Kumar, Robbie Jones, Fereshte Khani, Tengyu Ma, and Percy Liang. In-n-out: Pre-training and self-training using auxiliary information for out-of-distribution robustness. *arXiv preprint arXiv:2012.04550*, 2020. 2
- [35] Yajie Xing, Jingbo Wang, Xiaokang Chen, and Gang Zeng. 2.5D convolution for RGB-D semantic segmentation. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 1410–1414. IEEE, 2019. 2, 8
- [36] Wang Yifan, Shihao Wu, Hui Huang, Daniel Cohen-Or, and Olga Sorkine-Hornung. Patch-based progressive 3D point set upsampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5958–5967, 2019. 2
- [37] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. EC-net: an edge-aware point set consolidation network. In *Proceedings of the European conference on computer vision (ECCV)*, pages 386–402, 2018. 2
- [38] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018. 2
- [39] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5982–5991, 2019. 2