



DDRNet: Depth Map Denoising and Refinement for Consumer Depth Cameras Using Cascaded CNNs

Shi Yan¹, Chenglei Wu², Lizhen Wang¹, Feng Xu¹, Liang An¹, Kaiwen Guo³,
and Yebin Liu¹(✉)

¹ Tsinghua University, Beijing, China
liuyebin@mail.tsinghua.edu.cn

² Facebook Reality Labs, Pittsburgh, USA

³ Google Inc, Mountain View, CA, USA

Abstract. Consumer depth sensors are more and more popular and come to our daily lives marked by its recent integration in the latest Iphone X. However, they still suffer from heavy noises which limit their applications. Although plenty of progresses have been made to reduce the noises and boost geometric details, due to the inherent illness and the real-time requirement, the problem is still far from been solved. We propose a cascaded Depth Denoising and Refinement Network (DDRNet) to tackle this problem by leveraging the multi-frame fused geometry and the accompanying high quality color image through a joint training strategy. The rendering equation is exploited in our network in an unsupervised manner. In detail, we impose an unsupervised loss based on the light transport to extract the high-frequency geometry. Experimental results indicate that our network achieves real-time single depth enhancement on various categories of scenes. Thanks to the well decoupling of the low and high frequency information in the cascaded network, we achieve superior performance over the state-of-the-art techniques.

Keywords: Depth enhancement · Consumer depth camera
Unsupervised learning · Convolutional neural networks
DynamicFusion

1 Introduction

Consumer depth cameras have enabled lots of new applications in computer vision and graphics, ranging from live 3D scanning to virtual and augmented reality. However, even with tremendous progresses in improving the quality and resolution, current consumer depth cameras still suffer from heavy sensor noises.

During the past decades, in view of the big quality gap between depth sensors and traditional image sensors, researchers have made great efforts to leverage RGB images or videos to bootstrap the depth quality. While RGB-guided filtering methods show the effectiveness [22, 34], a recent trend is on investigating

the light transport in the scene for depth refinement with RGB images, which is able to capture high frequency geometry and reduce the texture-copy artifacts [3, 12, 43, 46]. Progresses have also been made to push these methods to run in real time [30, 44]. In these traditional methods, before refinement, a smooth filtering is usually carried out on the raw depth to reduce the sensor noise. However, this simple spatial filtering may alter the low-dimensional geometry in a non-preferred way. This degeneration can never be recovered in the follow-up refinement step, as only high-frequency part of the depth is modified.

To attack these challenges, we propose a new cascaded CNN structure to perform depth image denoising and refinement in order to lift the depth quality in low frequency and high frequency simultaneously. Our network consists of two parts, with the first focusing on denoising while the second aiming at refinement. For the denoising net, we train a CNN with a structure similar to U-net [36]. Our first contribution is on how to generate training data. Inspired by the recent progress on depth fusions [11, 19, 26], we generate reference depth maps from the fused 3D model. With fusion, heavy noise present in single depth map can be reduced by integrating the truncated signed distant function (TSDF). From this perspective, our denoising net is learning a deep fusion step, which is able to achieve better depth accuracy than heuristic smoothing.

Our second contribution is the refinement net, structured in our cascade end-to-end framework, which takes the output from the denoising net and refine it to add high-frequency details. Recent progresses in deep learning have demonstrated the power of deep nets to model complex functions between visual components. One challenge to train a similar net to add high-frequency details is that there is no ground truth depth map with desired high-frequency details. To solve this, we propose a new learning-based method for depth refinement using CNNs in an unsupervised way. Different from traditional methods, which define the loss directly on the training data, we design a generative process for RGB images using the rendering equation [20] and define our loss on the intensity difference between the synthesized image and the input RGB image. Scene reflectance is also estimated through a deep net to reduce the texture-copy artifacts. As the rendering procedure is fully differentiable, the image loss can be effectively back propagated throughout the network. Therefore, through these two components in our DDR-Net, a noisy depth map is enhanced both in low frequency and high frequency.

We extensively evaluate our proposed cascaded CNNs, demonstrating that our method can produce depth map with higher quality in both low and high frequency, compared with the state-of-the-art methods. Moreover, the CNN-based network structure enables our algorithm to run in real-time. And with the progress of deep-net-specific hardware, our method is promising to be deployed on mobile phones. Applications of our enhanced depth stream in the Dynamic-Fusion systems [11, 26, 47] are demonstrated, which improve the reconstruction performance of the dynamic scenes.

2 Related Work

Depth Image Enhancement. As RGB images usually capture a higher resolution than depth sensors, many methods in the past have focused on leveraging the

RGB images to enhance the depth data. Some heuristic assumptions are usually made about the correlation between color and depth. For example, some work assume that the RGB edges are coinciding with depth edges or discontinuities. Diebel and Thrun [9] upsample the depth with a Markov-Random Field. Depth upsampling with color image as input can be formulated as an optimization problem which maximizes the correlation between RGB edges and depth discontinuities [31]. Another way to implement this heuristics is through filtering [23], e.g. with joint bilateral upsampling filter [22]. Yang et al. [45] propose a depth upsampling method by filtering a cost space joint-bilaterally with a stereo image to achieve the resolution upsampling. Similar joint reconstruction ideas with stereo images and depth data are investigate by further constraining the depth refinement with photometric consistency from stereo matching [50]. With the development of modern hardwares and also the improvements in filtering algorithms, variants of joint-bilateral or multilateral filtering for depth upsampling can run in real-time [6, 10, 34]. As all of these methods are based on the heuristic assumption between color and depth, even producing plausible results, refined depth maps are not metrically accurate, and texture-copy artifacts are inevitable as texture variations are frequently mistaken for geometric detail.

Depth Fusion. With multiple frames as input, different methods have been proposed to fuse them to improve the depth quality or obtain a better quality scan. Cue et al. [8] has proposed a multi-frame superresolution technique to estimate higher resolution depth images from a stack of aligned low resolution images. Taking into account the sensors’ noise characteristics, the signed distance function is employed with an efficient data structure to scan scenes with an RGBD camera [16]. KinectFusion [27] is the first method to show real-time hand-held scanning of large scenes with a consumer depth sensor. Better data structures that exploit spatial sparsity in surface scans, e.g. hierarchical grids [7] or voxel hashing schemes [28], have been proposed to scan larger scenes in real time. These fusion methods are able to effectively reduce the noises in the scanning by integrating the TSDF. Recent progresses have extended the fusion to dynamic scenes [11, 26]. The scan from these depth fusion methods can achieve very clean 3D reconstruction, which improves the accuracy of the original depth map. Based on this observation, we employ depth fusion to generate a training data for our denoising net. By feeding lots of the fused depth as our training data to the the network, our denoising net effectively learns the fusion process. In this sense, our work is also related to Riegler et al. [35], where they designed an OctNet to perform the learning on signed distance function. Differently, our denoising net directly works on depth and by special design of our loss function, our net can effectively reduce the noise in the original depth map. Besides, high frequency geometric detail is not dealt with in OctNet, while by our refinement net we can achieve detailed depth maps.

Depth Refinement with Inverse Rendering. To model the relation between color and depth in a physically correct way, inverse rendering methods have been proposed to leverage RGB images to improve depth quality by investigating the

light transport process. Shape-from-shading (SfS) techniques have been investigated on how to extract the geometric detail from a single image [17, 49]. One challenge to directly apply SfS is that the light and reflectance are usually unknown when capturing the depth map. Recent progresses have shown that SfS can refine coarse image-based geometry models [4], even if they were captured under general uncontrolled lighting with multi-view cameras [42, 43] or an RGBD camera [12, 46]. In these work, illumination and albedo distributions, as well as refined geometry are estimated via inverse rendering optimization. Optimizing all these unknowns are very challenging by traditional optimization schemes. For instance, if the reflectance is not properly estimated, the texture-copy artifact can still exist. In our work, we employ a specifically structured network to tackle the challenge of reflectance and geometry separation problem. Our network structure can be seen as a regularizer which constrain the inverse rendering loss to back propagate only learnable gradient to train our refinement net. Also with a better reflectance estimation method than previous work, the reflectance influence can be further alleviated, resulting in a CNN network which extracts only geometry-related information to improve the depth quality.

Learning-Based and Statistical Methods. Data driven methods are another category to solve the depth upsampling/refinement problem. Data-driven priors are also helpful for solving the inverse rendering problem. Barron and Malik [2] jointly solve reflectance, shape and illumination, based on priors derived statistically from images. Similar concepts were also used for offline intrinsic image decomposition of RGB-D data [1]. Khan *et al.* [21] learn weighting parameters for complex SfS models to aid facial reconstruction. Wei and Hirzinger [40] use deep neural networks to learn aspects of the physical model for SfS. Note that even our method is also learning based, our refinement net does not take any training data. Instead, the refinement net relies on a pre-defined generative process and thus an inverse rendering loss for the training process. The closest idea to our paper is the encoder-decoder structure used for image-based face reconstruction [33, 38]. They take the traditional rendering pipeline as a generative process, defined as a fixed decode. Then, a reconstruction loss can be optimized to train the encoder, which directly regress from a input RGB image. However, these methods all require a predefined geometry and reflectance subspace, usually modeled by linear embedding, to help train a meaningful encode, while our method can work with general scenes captured by RGBD sensor.

3 Method

We propose a new framework for jointly training a denoising net and a refinement net from a consumer-level camera to improve depth map both in low frequency and high frequency. The proposed pipeline features our novelties both in training data creation and cascaded CNNs architecture design. Obtaining ground-truth high-quality depth data for training is very challenging. We thus have formulated the depth improvement problem into two regression tasks, while each one

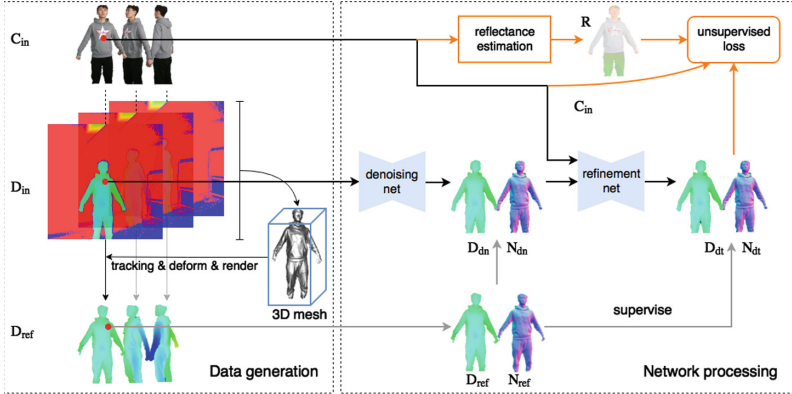


Fig. 1. The pipeline of our method. The black lines are the forward pass during test, the gray lines are supervise signal, and the orange lines are related to the unsupervised loss. Note that every loss function has a input mask W , which is omitted in this figure. D_{dn} and D_{dt} are denoised and refined output. N_{ref}, N_{dt} are reference normal and refined normal map, normals are only used for the training, not for the inference. (Color figure online)

focuses on lifting the quality in different frequency domains. This also enables us to combine supervised and unsupervised learning to solve the issue of lacking ground truth training data. For denoising part, a function \mathcal{D} mapping a noisy depth map D_{in} to a smoothed one D_{dn} with high-quality low frequency is learned by a CNN with the supervision of near-groundtruth depth maps D_{ref} , created from a state of the art of dynamic fusion. For refinement part, an unsupervised shading-based criterion is developed based on inverse rendering to train and a function \mathcal{R} to map D_{dn} and the corresponding RGB image C_{in} to an improved depth map D_{dt} with rich geometric details. The albedo map for each frame is also estimated the CNN used in [25]. We concurrently train cascaded CNNs from supervised depth data and unsupervised shading cues to achieve state-of-the-art performance on the task of single image depth enhancement. The detailed pipeline can be visualized in Fig. 1.

3.1 Dataset

Previous methods usually take a shortcut to obtain the training data by synthesizing [37, 39]. However, what if noise characteristic varies from sensor to sensor, or even the noise source is untraceable? In this case, how to generate ground-truth (or near-ground-truth) depth map becomes a major problem.

Data Generation. In order to learn the real noise distribution of different consumer depth cameras, we need to collect a training dataset of raw depth data with corresponding target depth maps, which act as the supervised signal of our denoising net. To achieve this, we use the non-rigid dynamic fusion pipeline

proposed by [11], which is able to reconstruct complete and good quality geometries of dynamic scenes from single RGB-D camera. The captured scene could be static or dynamic and we do not impose any assumptions on the type of motions. Besides, the camera is allowed to move freely during the capture. The reconstructed geometry is well aligned with input color frames. To this end, we first capture a sequence of synchronized RGB-D frames $\{D_t, C_t\}$. Then we run the non-rigid fusion pipeline [11] to produce a complete and improved mesh, and deform it using the estimated motion to each corresponding frame. Finally the target reference depth map $\{D_{ref,t}\}$ is generated by rasterization at each corresponding view point. Besides, we also produce a foreground mask $\{W_t\}$ using morphological filtering, which indicates the region of interest in the depth.

Content and Novelty. Using the above method, we contribute a new dataset of human bodies, including color image, raw depths with real noises and the corresponding reference depths with sufficient quality. Our training dataset contains 36840 views of aligned RGB-D data along with high quality D_{ref} rendered from fused model, among which 11540 views are from structured light depth sensor and 25300 views are from time-of-flight depth sensor. Our validation dataset contains 4010 views. Training set contains human bodies with various clothes poses under different lighting conditions. Moreover, to verify how our method generalized to other scenes, objects such as furniture and toys are also included in the test set. Existing public datasets, *eg.* Face Warehouse, Biwi Kinect face and D3DFACS, lack geometry details, thus do not meet our requirement for surface refinement. ScanNet consists of a huge amount of 3D indoor scenes, but has no human body category. Our dataset fills the blank in human body surface reconstruction. Dataset and training code will be public available.

3.2 Depth Map Denoising

The denoising net \mathcal{D} is trained to remove the sensor noise in depth map D_{in} given the reference depth map D_{ref} . Our denoising net architecture is inspired by DispNet [24] with skip connections and multi-scale predictions, as shown in Fig. 2. The denoising net consists of three parts: encoder, nonlinearity and decoder. The encoder aims to successively extract low-resolution high-dimensional features from D_{in} . To add nonlinearity to the network without performance degradation, several residual blocks with pre-activation are stacked sequentially between encoder and decoder part. The decoder part upsamples encoded feature maps to the original size, together with skip connections from the encoder part. These skip connections is useful to preserve geometry details in D_{in} . The whole denoising net adopts the residual learning strategy to extract the latent clean image from noisy observation. Not only does this direct pass set a good initialization, it turns out that residual learning is able to speed up the training process of deep CNN as well. Instead of the “unpooling + convolution” operation, our upsampling uses transpose convolution with trainable kernels. Note that the combination of bilinear up-sampling and transpose convolution in our upsampling pass help to inhibit checkerboard artifacts [29,41]. Our denoising net is




Fig. 2. The structure of our denoising net consists of encoder, nonlinear and decoder. There are three upsampling levels and one direct skip to keep captured value.

fully convolutional with receptive field up to 256. As a result, it is able to handle almost all types of consumer sensor inputs with different size.

The first loss for our denoising net is defined on the depth map itself. For example, per-pixel L1 and L2 loss on depth are used for our reconstruction term:

$$\ell_{rec}(D_{dn}, D_{ref}) = \|D_{dn} - D_{ref}\|_1 + \|D_{dn} - D_{ref}\|_2, \quad (1)$$

where $D_{dn} = \mathcal{D}(D_{in})$ is the output denoised depth map, and D_{ref} is the reference depth map. It is known that L2 and L1 loss may produce blurry results, however they accurately capture the low frequencies [18] which meets our purpose.

However, with only the depth reconstruction constraint, the high-frequency noise in small local patch could still remain after passing denoising net. To prevent this, we design a *normaldot* term to remove the high-frequency noise further. Specifically, this term is designed to constrain the normal direction of the denoised depth map to be consistent with the reference normal direction. We define the dot production of reference normal N_{ref}^i and tangential direction as the second loss term for our denoising net. Since each neighbouring depth point j ($j \in \mathcal{N}(i)$) could potentially define a 3D tangential direction, we sum over all possible directions, and the final normaldot term is formulated as:

$$\ell_{dot}(D_{dn}, N_{ref}) = \sum_i \sum_{j \in \mathcal{N}(i)} [\langle P^i - P^j, N_{ref}^i \rangle]^2, \quad (2)$$

where P^i is the 3D coordinate of D_{dn}^i . This term explicitly drives the network to consider the dependence between neighboring pixels $\mathcal{N}(i)$, and to learn locally the joint distributions of the neighboring pixels. Therefore, the final loss function for training the denoising net is defined as:

$$\mathcal{L}_{dn}(D_{dn}, D_{ref}) = \lambda_{rec} \ell_{rec} + \lambda_{dot} \ell_{dot}, \quad (3)$$

where $\lambda_{rec}, \lambda_{dot}$ defines the strength of each loss term.

In order to get N_{ref} from the depth map D_{ref} , a *depth to normal* (d2n) layer is proposed, which calculate normal vector given depth map and intrinsic parameters. For each pixel, it takes the surrounding 4 pixels to estimate one normal vector. The d2n layer is fully differentiable and has been employed several times in our end-to-end framework as shown in Fig. 1.




Fig. 3. Refinement net structure. The convolved feature maps from D_{dn} are complemented with the corresponding feature maps from C_{in} possessing the same resolution.

3.3 Depth Map Refinement

Although denoising net is able to effectively remove the noises, the denoised depth map, even getting improved in low frequency, still lacks details compared with RGB images. To add high-frequency details to the denoised depth map, we adopt a relatively small fully convolutional network based on hypercolumn architecture [14, 33].

Denote the single channel intensity map of color image C_{in} as I .¹ The hypercolumn descriptor for a pixel is extracted by concatenating the features at its spatial location in several layers, from both D_{dn} and I of the corresponding color image with high-frequency details. We first combine the spectral features from D_{dn} and I , then fuse these features in the spatial domain by max-pooling and convolutional down-sampling, which end with multi-scale fused feature maps. The pooling and convolution operation after hypercolumn extraction constructs a new set of sub-bands by fusing the local features of other hypercolumns in the vicinity. This transfers fine structure from color map domain to depth map domain. Three post-fusion convolutional layers is introduced to learn a better channel coupling. \tanh function is used as the last activation to limit the output to the same range of the input. In brief, high frequency features in the color image are extracted, and used as guidance, to extrude local detailed geometry from the denoised surfaces by the proposed refinement net shown in Fig. 3. As high frequency details are mainly inferred from small local patches, a shallow network with relative small reception field has enough capacity. Without post-processing as in other two-stage pipelines [37], our refinement net generates high-frequency details on depth map in a single forward pass.

Many SfS-based refinement approaches [13, 44] demonstrate that color images can be used to estimate the incident illumination, which is parameterized by the rendering process of an image. For Lambertian surface and low-frequency illumination, we can express the reflected irradiance B as the function of the

¹ Intensity image I plays the same role as C_{in} . We study I for simplicity.



Fig. 4. Estimated albedo map and relighted result using estimated lighting coefficients and uniform albedo. The estimation is in line with the actual incident illumination.

surface normal N , the lighting condition \mathbf{l} and the albedo R as follows:

$$B(\mathbf{l}, N, R) = R \sum_{b=1}^9 l_b H_b(N), \quad (4)$$

where $H_b : \mathbb{R}^3 \mapsto \mathbb{R}$ is the basis function of spherical harmonics(SH) that takes unit surface normal N as input. $\mathbf{l} = [l_1, \dots, l_9]^T$ are the nine 2nd order SH coefficients which represent the low-frequency scene illumination.

Based on Eq. 4, a per-pixel shading loss is designed. It penalizes both intensity and gradient of the difference value between the rendered image and the corresponding intensity image:

$$\ell_{sh}(N_{dt}, N_{ref}, I) = \|B(\mathbf{l}^*, N_{dt}, R) - I\|_2 + \lambda_g \|\nabla B(\mathbf{l}^*, N_{dt}, R) - \nabla I\|_2, \quad (5)$$

where N_{dt} represents the normal map of the regressed depth from the refinement net, λ_g is the weight to balance shading loss term, R is the albedo map estimated using Nestmeyer’s “CNN + filter” method [25]. Then, the light coefficients \mathbf{l}^* can be computed by solving the least squares problem:

$$\mathbf{l}^* = \arg \min_{\mathbf{l}} \|B(\mathbf{l}, N_{ref}, R) - I\|_2^2. \quad (6)$$

Here N_{ref} is calculated by the aforementioned d2n layer in Sect. 3.2. To show the effectiveness of our estimated illumination, a per-pixel albedo image is calculated by $R_I = I / \sum_{b=1}^9 l_b H_b(N_{ref})$, as shown in Fig. 4. Note that pixels at grazing angles are excluded in the lighting estimation, as both shading and depth are unreliable in these regions. Additionally, to constrain the refined depth to be close to the reference depth map, a fidelity term is added:

$$\ell_{fid}(D_{dt}, D_{ref}) = \|D_{dt} - D_{ref}\|_2. \quad (7)$$

Furthermore, a smoothness term is added to regularize the refined depth. More specifically, we minimize the anisotropic total variation of the depth:

$$\ell_{smo}(D_{dt}) = \sum_{i,j} |D_{dt}^{i+1,j} - D_{dt}^{i,j}| + |D_{dt}^{i,j+1} - D_{dt}^{i,j}|. \quad (8)$$

With all the above terms, the final loss for our refinement net is expressed as:

$$\mathcal{L}_{dt}(D_{dt}, D_{ref}, I) = \lambda_{sh}\ell_{sh} + \lambda_{fid}\ell_{fid} + \lambda_{smo}\ell_{smo}, \quad (9)$$

where $\lambda_{sh}, \lambda_{fid}, \lambda_{smo}$ defines the strength of each loss term. The last two additional terms are necessary, because they constrain the output depth map to be smooth and also close to our reference depth, as the shading loss would not be able to constrain the low frequency component.

3.4 End-to-End Training

We train our cascaded net jointly. To do so, we define total loss as:

$$\mathcal{L}_{total} = \mathcal{L}_{dn} + \lambda\mathcal{L}_{dt} \quad (10)$$

where λ is set to 1 during training. The denoising net is supervised by temporally fused reference depth map, and the refinement CNN is trained in an unsupervised manner. By incorporating supervision signals in both the middle and the output of the network, we achieve a steady convergence during the training phase. In the forward pass, each batch of input depth maps is propagated through the denoising net, and reconstruction L1/L2 term and *normaldot* term are added to \mathcal{L}_{total} . Then, the denoised depth maps, together with the corresponding color images, are fed to our refinement net. Shading, fidelity and smooth terms are added to \mathcal{L}_{total} . In the backward pass, the gradient of the loss \mathcal{L}_{total} are back-propagated through both network. All the hyper-parameters λ are fixed during training.

There are two types of consumer depth camera data in our training and validation set: structured light (K1) and time-of-flight (K2). We train the variants of our model on K1/K2 dataset respectively. To augment our training set, each RGB-D map are randomly cropped, flipped and re-scaled to the resolution of 256×256 . Considering that depth map is 2.5D in nature, the intrinsic matrix should be changed accordingly during data augmentation. This enables the network to learn more object-independent statistics and to work with sensors of different intrinsic parameters. For efficiency, we implement our d2n layer as a single CUDA layer. We choose *Adam* optimizer to compute gradients, with 0.9 and 0.999 exponential decay rate for the 1st and 2nd moment estimates. Base learning-rate is set to 0.001 and batch-size is 32. All convolution weights are initialized by Xavier algorithm, and weight decay is used for regularization.

4 Experiments

In this section, we evaluate the effectiveness of our cascade depth denoising and refinement framework, and analyze the contribution from each loss term. To the best of our knowledge, there is no public dataset for human body that contains raw and ground-truth depth maps with rich details from consumer depth cameras. We thus compare the performance of all available method on our own validation set, qualitatively and quantitatively.

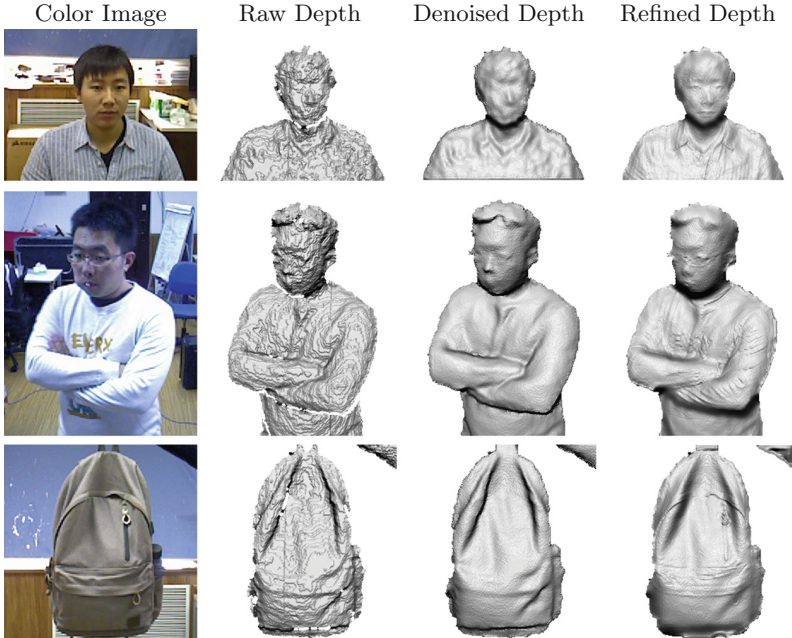


Fig. 5. Qualitative results on validation set. From left to right: RGB image, raw depth map, output of denoising net D_{dn} and output of refinement net D_{dt} . D_{dn} captures the low-dimensional geometry without noise, D_{dt} shows fine-grained details. Although trained on human body dataset, our model also produce high-quality depth map on general objects in arbitrary scenes, *eg.* the backpack sequence.

4.1 Evaluation

To verify the generalization ability of our trained network, we also evaluate on other objects other than human body, which can be seen in Figs. 5 and 8. One can see that although refined in an unsupervised manner, our results are comparable to the fused depth map [11] obtained using consumer depth camera only, and preserve thin structures such as fingers and folds in clothes better.

4.2 Ablation Study

The Role of Cascade CNN. To verify the necessity of our cascade CNNs, we replace our denoising net by a traditional preprocessing procedure, *eg.* bilateral filter, and still keep the refinement net to refine the filtered depth. We call this two-stage method as “Base+Ours refine”, and it is trained from scratch with shading, fidelity and smoothness loss. As we can see in the middle of Fig. 6, “Base+Ours refine” is not able to preserve distinctive structures of clothes in the presence of widespread structured noise. Unwanted high frequency noise leads to inaccurate estimation of illuminance, therefore shading loss term will

keep fluctuating during training. This training process will end up with non-optimal model parameters. However, in our cascade design, denoising net sets a good initialization for refinement net and achieves better result.

Supervision of Refinement Net. For our refinement net, there are two choices for regularization depth map in fidelity loss formulation, using reference depth map D_{ref} or the denoised depth map D_{dn} . When using only output of denoising net D_{dn} in an unsupervised manner, scene illumination is also estimated using D_{dn} . We denote this unsupervised framework as “Ours unsupervised”. Output of these two choices are shown in Fig. 7. In the unsupervised case, refinement net could produce reasonable result, but D_{dt} may stray from input.



Fig. 6. Left: normal map of D_{in} . Middle: Base+Ours refine, bilateral filter can't remove wavelet noise, refinement result suffers from high-frequency. Right: Ours.



Fig. 7. Left: C_{in} and D_{in} . Middle: Ours unsupervised, output depth does not match input value in stripes area in the cloth. Right: Ours with more reliable result.

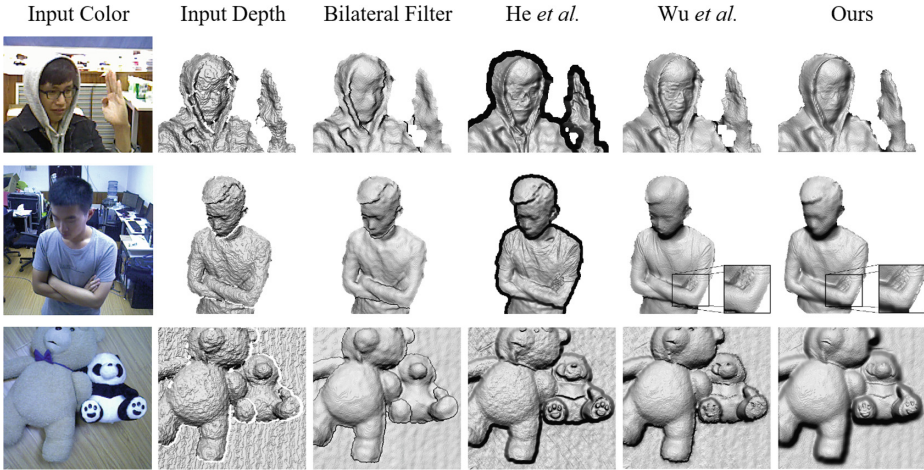


Fig. 8. Comparison of color-assisted depth map enhancement between bilateral filter, He *et al.* [15], Wu *et al.* [44] and our method. The closeup of the finger region demonstrates the effectiveness of unsupervised shading term in our refinement net loss.

4.3 Comparison with Other Methods

Compared with other non-data-driven methods, deep neural networks allow us to optimize non-linear loss and to add data-driven regularization, while keeping the inference time constant. Figure 8 shows examples of the qualitative comparison of different methods for depth map enhancement. Our method outperforms other methods by capturing cleaner structure of the geometry and high-fidelity details.

Quantitative Comparison. To evaluate quantitatively, we need a dataset with ground truth depth map. Multi-view stereo and laser scanner are able to capture static scene with high resolution and quality. We thus obtain ground truth depth value by multi-view stereo [32] (for K1) and Mantis Vision’s F6 laser scanner (for K2). Meanwhile, we collect the input of our method, the RGB-D image of the same scene by a consumer depth camera. The size of validation set is limited due to the high scan cost. Therefore, we also contribute a larger validation set labeled with the near-ground-truth depth obtained using mentioned method in Sect. 3.1. After reconstruction, the ground truth 3D model is rescaled and aligned with our reprojected enhanced depth map, using iterative closest point (ICP) [5]. Then the root mean squared error (RMSE) and the mean absolute error (MAE) between these two point clouds are calculated in Euclidean space. We also report the angular difference of normals, and the percentages of normal difference less than 3.0, 5.0, and 10.0°. Two sets of model are trained/evaluated on K1 and K2 data respectively. Quantitative comparison with other methods are summarized in Tables 1 and 2. Results shows that our method substantially outperforms other methods in terms of both metrics on the validation set.

Table 1. Quantitative comparison results on K1 validation set, error metrics in *mm*.

Method	Near-GT set		GT set	
	MAE	RMSE	MAE	RMSE
Bilateral [22]	15.9	4.8	15.1	3.7
He <i>et al.</i> [15]	46.5	14.7	41.1	15.2
Wu <i>et al.</i> [44]	14.5	4.3	15.7	4.4
Ours	10.9	4.1	11.0	3.6
Base+Ours refine	15.7	4.1	15.8	4.4
Ours unsupervised	16.1	5.2	14.9	5.5

Runtime Performance. At test time, our whole processing procedure includes data pre-processing and cascade CNN predicting. The preprocessing steps include: depth-to-color alignment, morphological transformation, and resampling if necessary. The forward pass takes 10.8 ms (256 × 256 input) or 20.4 ms (640 × 480 input) on TitanX, 182.56 ms or 265.8 ms per frame on Intel Core i7-6900K CPU. It is worth mentioning that without denoising CNN, a variant of our method, “Base+Ours refine” reaches a speed of 9.6ms per frame for 640 × 480 inputs.

Table 2. Average score of depth and normal error and on our K2 validation set.

Method	Depth difference					Normal difference				
	Seq. 1	Seq. 2	Seq. 3	Seq. 4	Seq. 5	Mean ↓	Median ↓	3.0↑	5.0↑	10.0↑
Wu <i>et al.</i> [44]	27.60	22.19	21.34	22.41	25.67	11.20	5.02	29.81	50.24	76.62
Or-El <i>et al.</i> [30]	27.14	25.42	22.89	21.31	26.08	10.03	4.12	35.43	56.57	79.99
Ours D_{dn}	19.03	19.25	18.49	18.37	18.76	9.36	3.40	45.33	66.79	84.69
Ours D_{dt}	18.97	19.41	18.38	18.50	18.61	9.55	3.54	43.77	64.98	83.69

4.4 Limitation

Similar to other real-time methods, we consider simplified light transport model. This simplification is effective but will impose intensity image’s texture on depth map. With the learning framework, texture-copy artifacts can be alleviated due to the fact that network can balance fidelity and shading loss term during training. Another limitation comes with non-diffuse surface assumption, as we only consider second order spherical harmonics representation.

5 Applications

It is known that real-time single frame depth enhancement is applicable for low-latency system without temporal accumulation. We compare the result using depth refined by our method with result using raw depth, on Dynamic Fusion [11] and DoubleFusion [48]. The temporal window in fusion systems would smooth out noise, but it will also wipe out high-frequency details. The time in TSDF fusion blocks the whole system from tracking detailed motions. In contrast, our method runs on single frame and provide timely update of fast changing surface details (*eg.* deformation of clothes and body gestures), as shown in red circles in Fig. 9 and the supplementary video. Moreover, real-time single frame depth enhancement could help tracking and recognition tasks under interactive scenarios.



Fig. 9. Application on DynamicFusion (left) and DoubleFusion (right) using our enhanced depth stream. Left: color image, Middle: fused geometry using raw depth stream, Right: “instant” geometry using our refined depth stream.

6 Conclusion

We presented the first end-to-end trainable network for depth map denoising and refinement for consumer depth cameras. We proposed a near-groundtruth training data generation pipeline, based on the depth fusion techniques. Enabled by the separation of low/high frequency parts in network design, as well as the collected fusion data, our cascaded CNNs achieves state-of-the-art result in real-time. Compared with available methods, our method achieved higher quality reconstruction in terms of both low dimensional geometry and high frequency details, which leads to superior performance quantitatively and qualitatively. Finally, with the popularity of integrating depth sensors into cellphones, we believe that our deep-net-specific algorithm is able to run on these portable devices for various quantitative measurement and qualitative visualization applications.

Acknowledgement. This work is supported by the National key foundation for exploring scientific instrument of China No. 2013YQ140517, and the National NSF of China grant No. 61522111, No. 61531014, No. 61671268 and No. 61727808.

References

1. Barron, J.T., Malik, J.: Intrinsic scene properties from a single RGB-D image. In: Proceedings of CVPR, pp. 17–24. IEEE (2013)
2. Barron, J.T., Malik, J.: Shape, illumination, and reflectance from shading. Technical report, EECS, UC Berkeley, May 2013
3. Beeler, T., Bickel, B., Beardsley, P.A., Sumner, B., Gross, M.H.: High-quality single-shot capture of facial geometry. *ACM Trans. Graph.* **29**(4), 40:1–40:9 (2010)
4. Beeler, T., Bradley, D., Zimmer, H., Gross, M.: Improved reconstruction of deforming surfaces by cancelling ambient occlusion. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012. LNCS, vol. 7572, pp. 30–43. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33718-5_3
5. Besl, P.J., McKay, N.D.: Method for registration of 3-D shapes. In: Robotics-DL Tentative, pp. 586–606. International Society for Optics and Photonics (1992)
6. Chan, D., Buisman, H., Theobalt, C., Thrun, S.: A noise-aware filter for real-time depth upsampling. In: ECCV Workshop on Multi-camera & Multi-modal Sensor Fusion (2008)
7. Chen, J., Bautembach, D., Izadi, S.: Scalable real-time volumetric surface reconstruction. *ACM Trans. Graph.* **32**(4), 113:1–113:16 (2013)
8. Cui, Y., Schuon, S., Thrun, S., Stricker, D., Theobalt, C.: Algorithms for 3D shape scanning with a depth camera. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(5), 1039–1050 (2013)
9. Diebel, J., Thrun, S.: An application of Markov random fields to range sensing. In: Proceedings of the 18th International Conference on Neural Information Processing Systems, NIPS 2005, pp. 291–298. MIT Press, Cambridge (2005)
10. Dolson, J., Baek, J., Plagemann, C., Thrun, S.: Upsampling range data in dynamic environments. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 1141–1148, June 2010

11. Guo, K., Xu, F., Yu, T., Liu, X., Dai, Q., Liu, Y.: Real-time geometry, albedo, and motion reconstruction using a single RGB-D camera. *ACM Trans. Graph.* **36**(3), 32:1–32:13 (2017)
12. Han, Y., Lee, J.Y., Kweon, I.S.: High quality shape from a single RGB-D image under uncalibrated natural illumination. In: *Proceedings of ICCV* (2013)
13. Han, Y., Lee, J.Y., Kweon, I.S.: High quality shape from a single RGB-D image under uncalibrated natural illumination. In: *IEEE International Conference on Computer Vision*, pp. 1617–1624 (2013)
14. Hariharan, B., Arbelaez, P., Girshick, R., Malik, J.: Hypercolumns for object segmentation and fine-grained localization, pp. 447–456 (2014)
15. He, K., Sun, J., Tang, X.: Guided image filtering. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(6), 1397–1409 (2013)
16. Henry, P., Krainin, M., Herbst, E., Ren, X., Fox, D.: RGB-D mapping: using kinect-style depth cameras for dense 3D modeling of indoor environments. *Int. J. Robot. Res.* **31**(5), 647–663 (2012)
17. Horn, B.K.: Obtaining shape from shading information. In: *The Psychology of Computer Vision*, pp. 115–155 (1975)
18. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks (2016)
19. Izadi, S., et al.: KinectFusion: real-time 3D reconstruction and interaction using a moving depth camera. In: *Proceedings of UIST*, pp. 559–568. ACM (2011)
20. Kajiya, J.T.: The rendering equation. In: *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1986*, pp. 143–150. ACM, New York (1986)
21. Khan, N., Tran, L., Tappen, M.: Training many-parameter shape-from-shading models using a surface database. In: *Proceedings of ICCV Workshop* (2009)
22. Kopf, J., Cohen, M.F., Lischinski, D., Uyttendaele, M.: Joint bilateral upsampling. *ACM Trans. Graph.* **26**(3), 96 (2007)
23. Lindner, M., Kolb, A., Hartmann, K.: Data-fusion of PMD-based distance-information and high-resolution RGB-images. In: *2007 International Symposium on Signals, Circuits and Systems*, vol. 1, pp. 1–4, July 2007
24. Mayer, N., et al.: A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In: *Computer Vision and Pattern Recognition*, pp. 4040–4048 (2016)
25. Nestmeyer, T., Gehler, P.V.: Reflectance adaptive filtering improves intrinsic image estimation. In: *CVPR*, pp. 1771–1780 (2017)
26. Newcombe, R.A., Fox, D., Seitz, S.M.: Dynamicfusion: reconstruction and tracking of non-rigid scenes in real-time. In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 343–352, June 2015
27. Newcombe, R.A., Izadi, S., et al.: KinectFusion: real-time dense surface mapping and tracking. In: *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 127–136 (2011)
28. Nießner, M., Zollhöfer, M., Izadi, S., Stamminger, M.: Real-time 3D reconstruction at scale using voxel hashing. *ACM Trans. Graph. (TOG)* **32**(6), 169 (2013)
29. Odena, A., Dumoulin, V., Olah, C.: Deconvolution and checkerboard artifacts. *Distill* (2016)
30. Or El, R., Rosman, G., Wetzler, A., Kimmel, R., Bruckstein, A.M.: RGBD-fusion: real-time high precision depth recovery. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015

31. Park, J., Kim, H., Tai, Y.W., Brown, M.S., Kweon, I.: High quality depth map upsampling for 3D-TOF cameras. In: 2011 International Conference on Computer Vision, pp. 1623–1630, November 2011
32. RealityCapture (2017). <https://www.capturingreality.com/>
33. Richardson, E., Sela, M., Or-El, R., Kimmel, R.: Learning detailed face reconstruction from a single image. In: CVPR (2017)
34. Richardt, C., Stoll, C., Dodgson, N.A., Seidel, H.P., Theobalt, C.: Coherent spatiotemporal filtering, upsampling and rendering of RGBZ videos. *Comput. Graph. Forum* **31**(2pt1), 247–256 (2012)
35. Riegler, G., Ulusoy, A.O., Bischof, H., Geiger, A.: OctNetFusion: learning depth fusion from data. In: 2017 International Conference on 3D Vision (3DV), pp. 57–66. IEEE (2017)
36. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
37. Sela, M., Richardson, E., Kimmel, R.: Unrestricted facial geometry reconstruction using image-to-image translation (2017)
38. Tewari, A., et al.: MoFA: model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In: The IEEE International Conference on Computer Vision (ICCV), vol. 2, p. 5 (2017)
39. Varol, G., et al.: Learning from synthetic humans. In: CVPR (2017)
40. Wei, G., Hirzinger, G.: Learning shape from shading by a multilayer network. *IEEE Trans. Neural Netw.* **7**(4), 985–995 (1996)
41. Wojna, Z., et al.: The devil is in the decoder (2017)
42. Wu, C., Stoll, C., Valgaerts, L., Theobalt, C.: On-set performance capture of multiple actors with a stereo camera. *ACM Trans. Graph. (TOG)* **32**(6), 161 (2013)
43. Wu, C., Varanasi, K., Liu, Y., Seidel, H., Theobalt, C.: Shading-based dynamic shape refinement from multi-view video under general illumination, pp. 1108–1115 (2011)
44. Wu, C., Zollhöfer, M., Nießner, M., Stamminger, M., Izadi, S., Theobalt, C.: Real-time shading-based refinement for consumer depth cameras. *ACM Trans. Graph. (TOG)* **33**(6), 200 (2014)
45. Yang, Q., Yang, R., Davis, J., Nister, D.: Spatial-depth super resolution for range images. In: 2007 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8, June 2007
46. Yu, L., Yeung, S., Tai, Y., Lin, S.: Shading-based shape refinement of RGB-D images, pp. 1415–1422 (2013)
47. Yu, T., et al.: BodyFusion: real-time capture of human motion and surface geometry using a single depth camera. In: The IEEE International Conference on Computer Vision (ICCV). IEEE, October 2017
48. Yu, T., et al.: DoubleFusion: real-time capture of human performance with inner body shape from a depth sensor. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
49. Zhang, Z., Tsai, P.S., Cryer, J.E., Shah, M.: Shape from shading: a survey. *IEEE PAMI* **21**(8), 690–706 (1999)
50. Zhu, J., Wang, L., Yang, R., Davis, J.: Fusion of time-of-flight depth and stereo for high accuracy depth maps. In: 2008 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8, June 2008