

Occlusion-Aware Depth Estimation with Adaptive Normal Constraints: Supplementary Materials

Xiaoxiao Long¹, Lingjie Liu², Christian Theobalt², and Wenping Wang¹

¹ The University of Hong Kong

² Max Planck Institute for Informatics

1 More Results

Depth prediction Figure 1 shows more visual results of our method, mvdepthnet [4] and neuralrgbd [2]. Compared with mvdepthnet [4] and neuralrgbd [2], our estimated depth map has less noise, sharper boundaries and spatially consistent depth values, as can also be seen in the surface normal visualization. Furthermore, the 3D point cloud exported from our estimated depth better preserves global planar features (e.g. the wall corner in the first scene) and local features (e.g. the head model and the keyboard in the second scene).

Surface normal accuracy In Figure 2, we show more visual results of surface normals calculated from the estimated depth. Our method outperforms GeoNet [3], Yin et al. [5] and Kusupati et al. [1] qualitatively. Compared with our model retrained using VNL, the normal values of our model trained with the *CNM* constraint are more consistent and smoother in planar regions. This indicates that *CNM*, which is a key novelty of our work, indeed contributes significantly to the overall improvement of performance, as compared with VNL.

Video-based 3D reconstruction In Figure 3, more video-based 3D reconstruction results are shown. For the white walls, the feature-less sofa and the table, our reconstructed result is much better than the results by the other two methods in terms of reconstruction quality of local and global structures.

2 Network structures

In this section, we illustrate the network structures used in our pipeline.

DepthNet. The structure of DepthNet is shown in Table 1, and we set $D = 64$ in our paper.

RefineNet. The structure of RefineNet is shown in Table 2, and we set $D = 64$ in our paper.

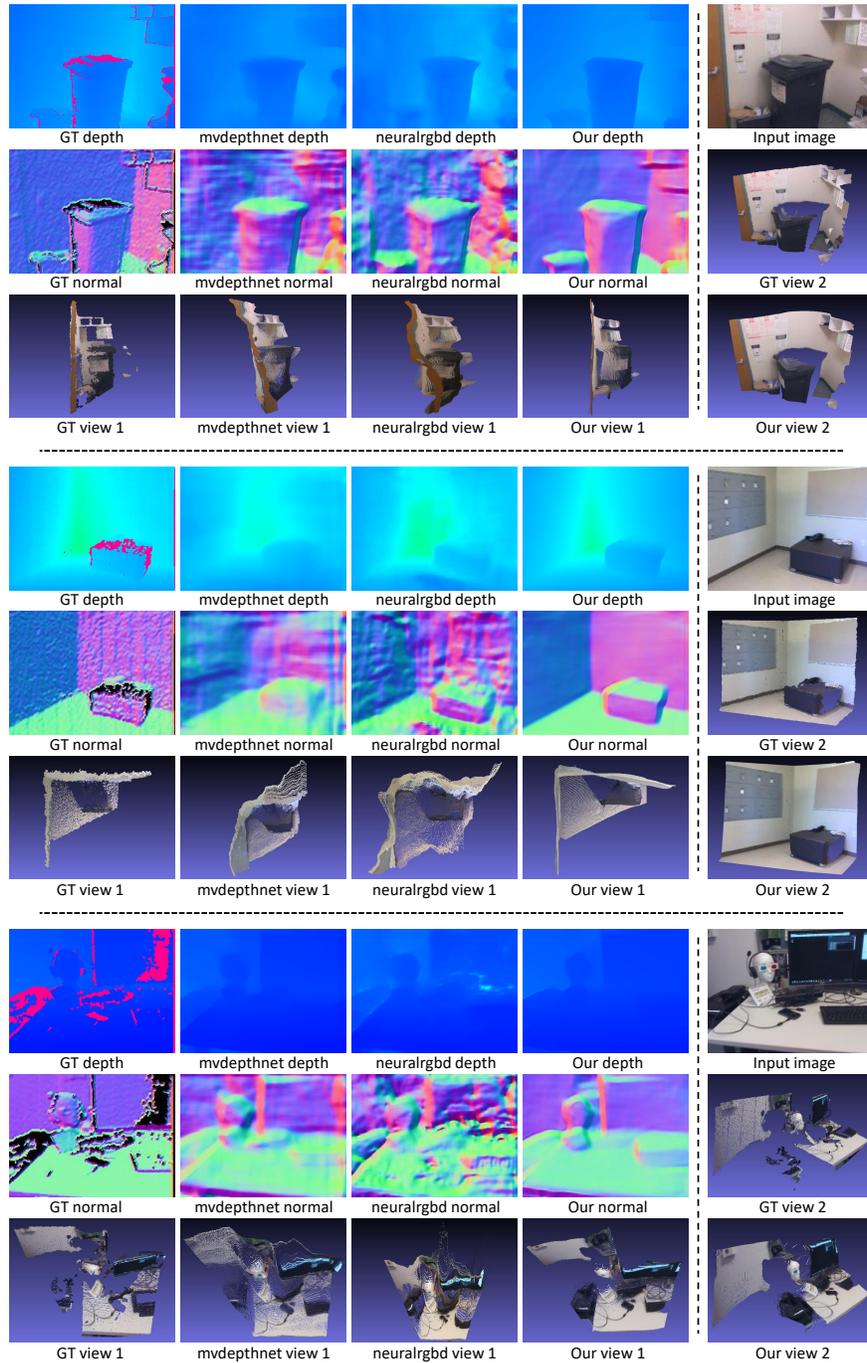


Fig. 1. Depth comparison with mvdepthnet [4] and neuralrgb [2].

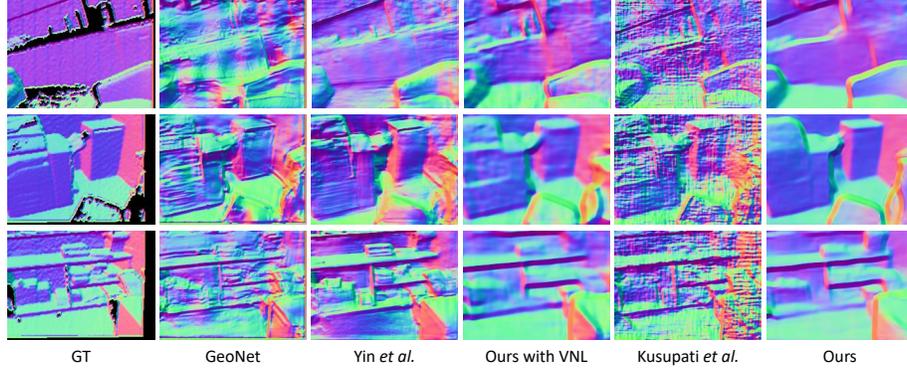


Fig. 2. Visual comparison of surface normal calculated from the estimated depth with GeoNet [3], Yin et al. [5] and Kusupati et al. [1].

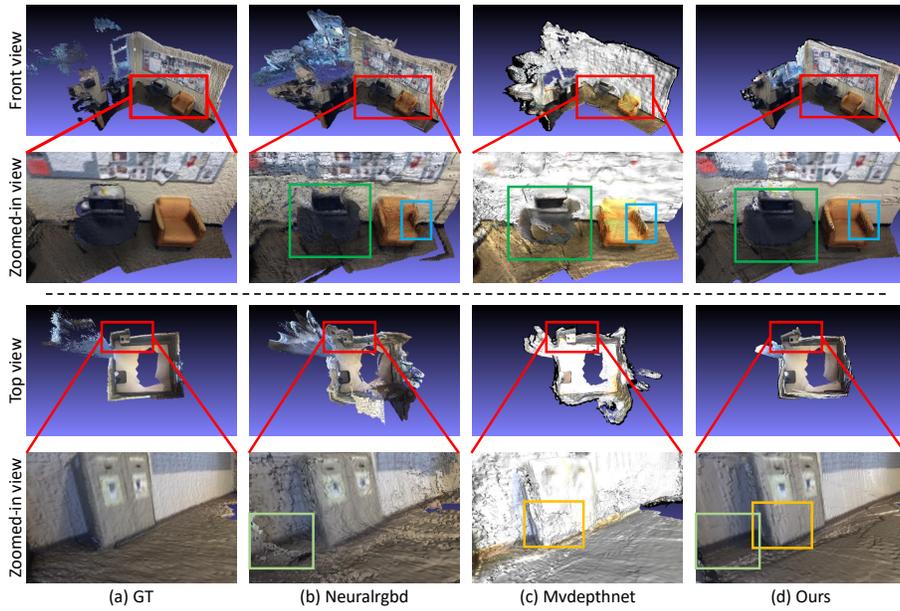


Fig. 3. Comparison with neuralrgb [2] and mvdepthnet [4] for 3D reconstruction on scenes from ScanNet. (a) With ground truth depth; (b) With estimated depth and confidence map from neuralrgb; (c) With estimated depth from mvdepthnet; (d) With our estimated depth and occlusion probability map. All reconstructions are done with TSDF fusion [6].

Table 1. The structure of DepthNet, which is built on [4]

Name	Layer components	Layer input	Output dimension
Input	concat (initial cost volume, reference image)		$W \times H \times 67$
conv1	<i>conv_2d</i> (7×7 , <i>ch_in</i> = 67, <i>ch_out</i> = 128, <i>stride</i> = 1), <i>BN</i> , <i>ReLU</i> <i>conv_2d</i> (7×7 , <i>ch_in</i> = 128, <i>ch_out</i> = 128, <i>stride</i> = 2), <i>BN</i> , <i>ReLU</i>	Input	$\frac{1}{2}W \times \frac{1}{2}H \times 128$
conv2	<i>conv_2d</i> (5×5 , <i>ch_in</i> = 128, <i>ch_out</i> = 256, <i>stride</i> = 1), <i>BN</i> , <i>ReLU</i> <i>conv_2d</i> (5×5 , <i>ch_in</i> = 256, <i>ch_out</i> = 256, <i>stride</i> = 2), <i>BN</i> , <i>ReLU</i>	conv1	$\frac{1}{4}W \times \frac{1}{4}H \times 256$
conv3	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 512, <i>stride</i> = 1), <i>BN</i> , <i>ReLU</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 512, <i>stride</i> = 2), <i>BN</i> , <i>ReLU</i>	conv2	$\frac{1}{8}W \times \frac{1}{8}H \times 512$
conv4	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 512, <i>stride</i> = 1), <i>BN</i> , <i>ReLU</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 512, <i>stride</i> = 2), <i>BN</i> , <i>ReLU</i>	conv3	$\frac{1}{16}W \times \frac{1}{16}H \times 512$
conv5	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 512, <i>stride</i> = 1), <i>BN</i> , <i>ReLU</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 512, <i>stride</i> = 2), <i>BN</i> , <i>ReLU</i>	conv4	$\frac{1}{32}W \times \frac{1}{32}H \times 512$
upconv5	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 512), <i>BN</i> , <i>ReLU</i>	conv5	$\frac{1}{16}W \times \frac{1}{16}H \times 512$
iconv5	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 1024, <i>ch_out</i> = 512), <i>BN</i> , <i>ReLU</i>	concat (upconv5,conv4)	$\frac{1}{16}W \times \frac{1}{16}H \times 512$
upconv4	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 512), <i>BN</i> , <i>ReLU</i>	iconv5	$\frac{1}{8}W \times \frac{1}{8}H \times 512$
iconv4	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 1024, <i>ch_out</i> = 512), <i>BN</i> , <i>ReLU</i>	concat (upconv4,conv3)	$\frac{1}{8}W \times \frac{1}{8}H \times 512$
depth4	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 1), <i>Sigmoid</i> <i>bilinear upsample</i>	iconv4	$\frac{1}{4}W \times \frac{1}{4}H \times 1$
upconv3	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 256), <i>BN</i> , <i>ReLU</i>	iconv4	$\frac{1}{4}W \times \frac{1}{4}H \times 256$
iconv3	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 513, <i>ch_out</i> = 256), <i>BN</i> , <i>ReLU</i>	concat (upconv3,conv2,depth4)	$\frac{1}{4}W \times \frac{1}{4}H \times 256$
depth3	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 1), <i>Sigmoid</i> <i>bilinear upsample</i>	iconv3	$\frac{1}{2}W \times \frac{1}{2}H \times 1$
upconv2	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 128), <i>BN</i> , <i>ReLU</i>	iconv3	$\frac{1}{2}W \times \frac{1}{2}H \times 128$
iconv2	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 257, <i>ch_out</i> = 128), <i>BN</i> , <i>ReLU</i>	concat (upconv2,conv1,depth3)	$\frac{1}{2}W \times \frac{1}{2}H \times 128$
depth2	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 128, <i>ch_out</i> = 1), <i>Sigmoid</i> <i>bilinear upsample</i>	iconv2	$W \times H \times 1$
upconv1	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 128, <i>ch_out</i> = 64), <i>BN</i> , <i>ReLU</i>	iconv2	$W \times H \times 64$
iconv1	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 65, <i>ch_out</i> = 64), <i>BN</i> , <i>ReLU</i>	concat (upconv1,depth2)	$W \times H \times 64$
depth1	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 64, <i>ch_out</i> = 1), <i>Sigmoid</i>	iconv1	$W \times H \times 1$
Output	depth1		$W \times H \times 1$

Table 2. The structure of RefineNet.

Name	Layer components	Layer input	Output dimension
Input	concat (averaged cost volume, initial depths, the difference of initial depths)		$W \times H \times 67$
conv1	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 67, <i>ch_out</i> = 128, <i>stride</i> = 1), <i>BN</i> , <i>ReLU</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 128, <i>ch_out</i> = 128, <i>stride</i> = 2), <i>BN</i> , <i>ReLU</i>	Input	$\frac{1}{2}W \times \frac{1}{2}H \times 128$
conv2	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 128, <i>ch_out</i> = 256, <i>stride</i> = 1), <i>BN</i> , <i>ReLU</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 256, <i>stride</i> = 2), <i>BN</i> , <i>ReLU</i>	conv1	$\frac{1}{4}W \times \frac{1}{4}H \times 256$
conv3	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 512, <i>stride</i> = 1), <i>BN</i> , <i>ReLU</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 512, <i>stride</i> = 2), <i>BN</i> , <i>ReLU</i>	conv2	$\frac{1}{8}W \times \frac{1}{8}H \times 512$
Depth refinement branch			
upconv3_d	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 256), <i>BN</i> , <i>ReLU</i>	conv3	$\frac{1}{4}W \times \frac{1}{4}H \times 256$
iconv3_d	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 256), <i>BN</i> , <i>ReLU</i>	concat (upconv3_d, conv2)	$\frac{1}{4}W \times \frac{1}{4}H \times 256$
upconv2_d	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 128), <i>BN</i> , <i>ReLU</i>	iconv3_d	$\frac{1}{2}W \times \frac{1}{2}H \times 128$
iconv2_d	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 128), <i>BN</i> , <i>ReLU</i>	concat (upconv2_d, conv1)	$\frac{1}{2}W \times \frac{1}{2}H \times 128$
upconv1_d	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 128, <i>ch_out</i> = 64), <i>BN</i> , <i>ReLU</i>	iconv2_d	$W \times H \times 64$
iconv1_d	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 64, <i>ch_out</i> = 64), <i>BN</i> , <i>ReLU</i>	upconv1_d	$W \times H \times 64$
depth	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 64, <i>ch_out</i> = 1), <i>Sigmoid</i>	iconv1_d	$W \times H \times 1$
Occlusion probability prediction branch			
upconv3_p	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 256), <i>BN</i> , <i>ReLU</i>	conv3	$\frac{1}{4}W \times \frac{1}{4}H \times 256$
iconv3_p	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 512, <i>ch_out</i> = 256), <i>BN</i> , <i>ReLU</i>	concat (upconv3_p, conv2)	$\frac{1}{4}W \times \frac{1}{4}H \times 256$
upconv2_p	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 128), <i>BN</i> , <i>ReLU</i>	iconv3_p	$\frac{1}{2}W \times \frac{1}{2}H \times 128$
iconv2_p	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 256, <i>ch_out</i> = 128), <i>BN</i> , <i>ReLU</i>	concat (upconv2_p, conv1)	$\frac{1}{2}W \times \frac{1}{2}H \times 128$
upconv1_p	<i>bilinear upsample</i> <i>conv_2d</i> (3×3 , <i>ch_in</i> = 128, <i>ch_out</i> = 64), <i>BN</i> , <i>ReLU</i>	iconv2_p	$W \times H \times 64$
iconv1_p	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 64, <i>ch_out</i> = 64), <i>BN</i> , <i>ReLU</i>	upconv1_p	$W \times H \times 64$
probability	<i>conv_2d</i> (3×3 , <i>ch_in</i> = 64, <i>ch_out</i> = 1), <i>Sigmoid</i>	iconv1_p	$W \times H \times 1$
Output	depth and probability		$W \times H \times 1$

References

1. Kusupati, U., Cheng, S., Chen, R., Su, H.: Normal assisted stereo depth estimation. arXiv preprint arXiv:1911.10444 (2019)
2. Liu, C., Gu, J., Kim, K., Narasimhan, S.G., Kautz, J.: Neural rgb (r) d sensing: Depth and uncertainty from a video camera. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10986–10995 (2019)
3. Qi, X., Liao, R., Liu, Z., Urtasun, R., Jia, J.: Geonet: Geometric neural network for joint depth and surface normal estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 283–291 (2018)
4. Wang, K., Shen, S.: Mvdepthnet: real-time multiview depth estimation neural network. In: 2018 International Conference on 3D Vision (3DV). pp. 248–257. IEEE (2018)
5. Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction. In: The IEEE International Conference on Computer Vision (ICCV) (2019)
6. Zeng, A., Song, S., Nießner, M., Fisher, M., Xiao, J., Funkhouser, T.: 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In: CVPR (2017)