

Indoor Depth Completion with Boundary Consistency and Self-Attention

Yu-Kai Huang* Tsung-Han Wu* Yueh-Cheng Liu Winston H. Hsu

National Taiwan University

{r08922053, b05902013, r07922183, whsu}@ntu.edu.tw

Abstract

Depth estimation features are helpful for 3D recognition. Commodity-grade depth cameras are able to capture depth and color image in real-time. However, glossy, transparent or distant surface cannot be scanned properly by the sensor. As a result, enhancement and restoration from sensing depth is an important task. Depth completion aims at filling the holes that sensors fail to detect, which is still a complex task for machine to learn. Traditional hand-tuned methods have reached their limits, while neural network based methods tend to copy and interpolate the output from surrounding depth values. This leads to blurred boundaries, and structures of the depth map are lost. Consequently, our main work is to design an end-to-end network improving completion depth maps while maintaining edge clarity. We utilize self-attention mechanism, previously used in image inpainting fields, to extract more useful information in each layer of convolution so that the complete depth map is enhanced. In addition, we propose boundary consistency concept to enhance the depth map quality and structure. Experimental results validate the effectiveness of our self-attention and boundary consistency schema, which outperforms previous state-of-the-art depth completion work on Matterport3D dataset.

1. Introduction

In this task, we take single input RGB and raw sensing depth to complete depth value of missing parts. Depth sensing is widely used in real-world applications from autonomous driving, robotics, augmented reality to scene reconstruction. Most of the applications are for recognition purpose. However, commercially available RGB-D cameras, such as Microsoft Kinect and Intel RealSense, still fail on sensing depth map without corruption or noise along object edges. Sensors produce large amounts of missing pixels when surface is shiny, transparent, too close or too far. In indoor dataset Matterport3D [3], for example, over 15% of




Figure 1. Our work produces clearer structures and overcomes the hardness that previous works suffer from: interpolation and blurred boundaries. On the first row, our model provides clear complete depth but FCN interpolates depth value (marked in red) at the large missing area of raw depth map. On the second row, our model can assure boundary sharpness, while FCN generates blurred boundaries.

values in raw depth are missing. As a consequence, completing large missing parts of sensor depth is of crucial importance.

Traditional works reconstruct depth value by combining multiple views of sensor data. Simultaneous localization and mapping (SLAM) [1, 5] method incrementally builds a consistent depth map of the environment while simultaneously determines location within this map. Previous methods [7, 14] utilizes SLAM to estimate more accurate depth map. In our work, we do not consider multiple views of reconstruction. Rather, we put emphasis on single view for depth prediction, because single view prediction is more efficient when consider inference time and computing power.

For single view of depth completion, previous works that tried to solve the problem have their own issue. Traditional mathematical methods [4, 25] contain much hand-tuning hyper-parameters and reach their limits in filling high quality depth holes. Recent deep learning methods [6, 34] performed well yet only learned to interpolate or copy-and-paste depth value from neighboring pixels. Blurred boundaries and structure can be easily recognized by visualization of output depth. In Figure 1, the visualization results show the above issue of copy-and-paste, interpolation, blurred boundaries and structure. We conclude the issues of deep

*equal contribution

learning based methods in two orientations: estimating precise depth value and producing clear structure.

In this paper, we solve the issues mentioned above. Previous works pass the whole feature maps into the network and perform convolution. However, to complete the missing depths, some regions may be more important and some may be minor details depending on the semantics and geometric meanings of one scene. Inspired by Yu *et al.* [36], we leverage the self-attention network to encourage the model to pay attention on relevant parts, especially the semantics of the scene. Our model benefits from the attention mechanism and outputs depth maps with comprehensive consideration of the attentioned parts. Furthermore, with useful information from depth representation (surface normals and occlusion boundaries), proposed by [39], our self-attention model meets the needs of estimating more precise depth value rather than just interpolation.

In addition, we solve the issue of vague structures with boundary consistency concept. In order to generate clear depth structures, we equip another network to predict occlusion boundaries given the output depth from depth completion network. As a result, the depth completion network is indirectly refined to preserve clear boundaries in the output depth, making the complete depth image to be more structured and conformed to realistic situation. To the best of our knowledge, we are the first to apply self-attention mechanism and regard boundary as the main feature in indoor depth completion task.

On the whole, our main contributions of our work are:

- Enhance depth completion task using self-attention mechanism, which has never been used in RGB-D estimation before.
- Propose a novel idea, boundary consistency, to produce depth map of clear structure.
- Reach state-of-the-art performance on Matterport3D [3] on RGB-D depth completion task.

Extensive ablation study and visualization results validate our proposed idea and concept. The promising results on depth completion task makes contribution to recognition applications such as robotics and autonomous driving. Those applications suffer from various types of degradation of low-quality visual data, including large missing holes of sensor depth map. Our work enhances the robustness for recognition on sensor depth, and, consequently, contributes to real-world recognition from low-quality images.

2. Related Work

We introduce order as follows: depth estimation, super-resolution, reconstruction from sparse samples, image inpainting, and, finally, our task depth completion.

2.1. Depth Estimation

Depth estimation predicts depth value from monocular RGB image. This is a long-established problem in computer vision history. However, with little information of a single view image, this is still a hard task even for human beings to answer the exact depth distance [26]. Classic methods like Shape from Shading [27, 38] and depth from focus [9, 32] elaborated physical and mathematical property about light shading and focal setting at each pixel. Recent works extended classic methods by machine learning, like deep depth from focus [12] and deep estimation based on fourier domain analysis [19]. Fully connected convolution (FCN) networks are used to predict depth map [18] or refine coarse-scale depth value [6].

Previous works give us insight into the methods to solve the task, but different from what our topic focuses on. Specifically, depth estimation generate raw depth based on RGB image, while we utilize RGB to improve raw depth to become complete.

2.2. Depth Super-Resolution

The focus of depth super-resolution is the enhancement of spatial resolution. Generally, raw depth data are regularly sampled or quantized with low resolution. Yang *et al.* [35] iteratively refined depth map in terms of both its spatial resolution and depth precision with bilateral filtering. Mac *et al.* [23] utilized Markov random field to select candidate of depth patch. Some other works are shape-from-shading [10, 37], rigid body self-similarity [15], and deep learning based network [28].

Different from depth completion, these works take complete but low-resolution depth map as input. In our work, we aim at recovering the large missing parts of the depth map, which cannot be recovered by super-resolution methods.

2.3. Depth Reconstruction from Sparse Samples

Depth reconstruction from sparse samples is an interesting topic that tries to reconstruct full depth map from sparse one. Restrictions of camera costs and power consumption make it a necessary research topic. The difficulties of the topic are how to maintain decent performance, while saving costs and energy. Ma's method [24] can be used as a plug-in module in sparse SLAM and visual inertial odometry algorithms, creating more accurate, dense point cloud. Also, Ma [22] proposed a self-supervised framework without the need for dense labels, achieving great performance on KITTI [8] dataset. Some works combined semantic segmentation [16] to improve the prediction.

2.4. Image Inpainting

Image inpainting is related to depth completion in some ways. For example, free-form image inpainting [36] allows

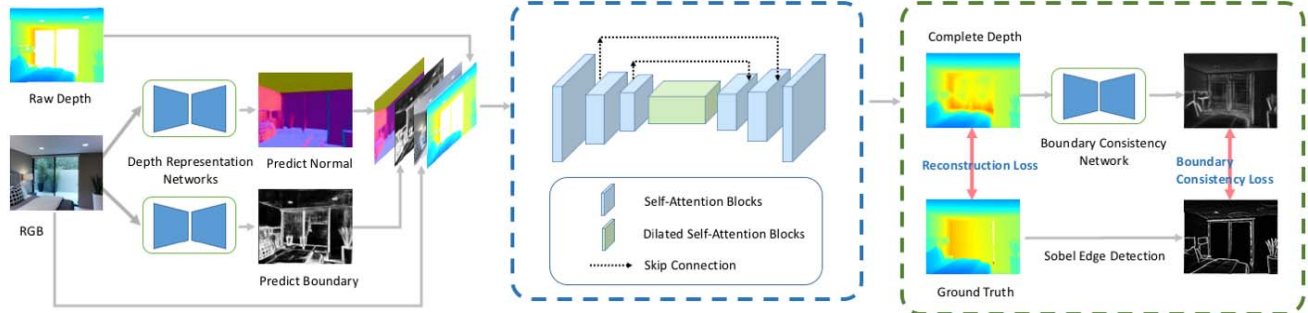


Figure 2. Overview of the network architecture. Main improvement of our method are twofold: self-attention mechanism 3.1 (blue dashed box) and boundary consistency 3.2 (green dashed box). Self-attention network pays attention on useful feature information. On the other hand, boundary consistency loss forces attention network to preserve boundary information, allowing network to reconstruct more structured depth map.

users to erase values of a RGB image, and then network inpaints the missing parts of RGB values as real as possible. Compared to our task, the same part is the missing holes of depth sensor are similar to the erased parts of RGB image. The different part is depth completion network is trained to learn precise value of depth map supervised with correct answer, while image inpainting results do not require a correct answer but a real and reasonable one. Self-attention mechanism, which applies attention on each layer of convolution, is useful and powerful for image inpainting task. Yu *et al.* utilized self-attention mechanism on partial convolution [20] and gated convolution [36] and made astonishing performance on the inpainting image. We introduce the concept of self-attention on our task.

2.5. Depth Completion

Depth completion is to fill in missing holes in the relatively dense depth images. The definition of depth completion concept is proposed by Zhang *et al.* [39]. Earlier works [2, 30] evaluated performance on pixels captured by commodity RGB-D cameras. Under the circumstances, they can at best output raw depth (captured by the sensor), which contains large missing holes. Zhang extended the term "Depth Completion" to predicting the complete depth map generated from multi-view reconstruction.

In Zhang's work, he proved surface normals and occlusion boundaries from RGB are two geometric representations best for deep depth completion. Zhang also claimed that deep regression method simply learns to copy and interpolate depth value, so he optimized loss with sparse Cholesky factorization scheme [31] and reach state-of-the-art performance.

In our work, we re-verify the copy and interpolation issue of standard FCN. We overcome the obstacles that FCN encounters, and our model can generate depth map with clear boundaries and structures through an end-to-end neural network. Our approach outperforms the previous meth-

ods. Compared to Zhang's, we have faster inference time by replacing the Cholesky optimization, which is more desirable for real world applications.

3. Method

In this paper, we investigate how to complete indoor depth image given a single RGB-D image with effective deep learning method. Our work mainly focuses on the following two questions: "How to overcome the difficulty that network simply learns to interpolate depth value?" and "how to make clear structures of depth image?"

3.1. Self-Attention Mechanism

The first issue we address is that previous deep learning methods tend to copy and interpolate nearby depth value as outputs as shown in Figure 1. Since in most cases, predicting average depth values for unobserved area makes the loss drop rapidly, so the network can easily fall into local minima of copy and interpolation instead of predicting precise depth values. To solve this problem, we propose to use self-attention mechanism on each convolution layer. The network is allowed to focus on precise feature values at each convolution stage and forward useful information. With self-attention, the network can be more robust in predicting precise depth values and would not be easily trapped in local minima of copy and interpolation.

We use gated convolution, proposed by Yu *et al.* [36], as our backbone component. Yu *et al.* used gated convolution on free-form image inpainting task. We use gated convolution on our task, since, first, depth completion task aims to complete missing depth value, which is similar to free-form masks from users. Additionally, completing precise depth value is similar to inpainting with realistic RGB pixel values. As a result, we gain insight from the strong connection between the two different tasks.

For an input I of a convolution block, and convolution blocks for feature extraction $Conv_f$ and for gating $Conv_g$,

the mathematical formulation for self-attention mechanism can be presented as:

$$\text{Gating} = \sigma(SN(\text{Conv}_g(\mathbf{I}))) \quad (1)$$

$$\text{Feature} = \phi(SN(\text{Conv}_f(\mathbf{I}))) \quad (2)$$

$$\text{Output} = \text{Feature} \odot \text{Gating} \quad (3)$$

where σ is sigmoid function, SN is spectral normalization, ϕ is any activation function, and \odot is the pixel-wise multiplication. The self-attention mechanism is in the presence of an additional gating operation that Feature \odot Gating. It allows the model to learn effective dynamic feature selection, and also highlights the feature meanings for each channel map and each spatial location.

Since Conv_g learns to identify region that is useful and important, according to the above equation, the model preserves useful region of Feature in the Output. Therefore, the gated convolution can predict more accurate depth values by paying more attention on local and detailed information extracted by the self-attention mechanism.

3.2. Boundary consistency

Ordinary FCN methods are likely to generate depth image of copy and interpolation as explained above. This causes the output depth to have blurred boundaries and structures. For instance, line segments may be slightly distorted and object occlusion boundary is vague in most cases. Since depth completion is a complex regression problem, FCN models can hardly learn to generate depth map of sharp boundaries where neighboring pixel values differ greatly. In order to fix the problem, we introduce boundary consistency to encourage models to learn clear boundaries and structures.

Boundary consistency constrains the output depth map to have boundary information by constructing another network, boundary consistency network. The boundary consistency network, modified from U-Net [29], takes complete depth as input and predicts the occlusion boundary. It is trained end-to-end along with self-attention network. The boundary consistency network encourages the self-attention network to generate sharp edges so that boundary consistency network can estimate more precise boundary values. With moderate balance between the weights of losses, our self-attention network successfully completes depth map with sharper and clearer structures. The overview pipeline is shown in Figure 1, we add an additional encode-decode boundary consistency network to the original self-attention network. We name it "boundary consistency" because we use occlusion boundary as depth input feature and require our model to preserve the boundary information in the complete depth map.

The boundary estimation result is supervised by Sobel [17] boundary from ground truth depth image. There are

two reasons we use Sobel as our target loss. (1) The generated edges are sharp, clear and few variate compared to some other methods. (2) Sobel Algorithm can extract occlusion boundaries without being too sensitive to the noise.

With the help of occlusion boundary ground truth from Sobel, our generated depth map can learn sharper and clearer structures, preventing outputs like blurred boundaries and interpolation.

3.3. Depth Representation

In Zhang *et al.* [39] work, he mentioned two important depth representations: surface normals and occlusion boundaries. Surface normals relate strongly to local lighting variations, providing more information about local surface properties. Occlusion boundaries also represent local texture features observed from RGB images. Both of the two representations can be robustly detected by networks, called depth representation network in the following. Because it is proven by [39] that surface normals and occlusion boundaries are powerful, we simply combine the two representations with RGB and raw depth as the input of self-attention network.

3.4. Loss Functions

The following section would introduce each component of our loss in detail. The total loss can be written as:

$$L = L_{SA} - \lambda_S L_S + \lambda_{BC} L_{BC} + \lambda_N L_N + \lambda_B L_B \quad (4)$$

Given input RGB, depth representation networks predict normals N and boundaries B . The depth representation losses are $L_N = \frac{1}{|p|} \sum_p \|N(p) - N_0(p)\|^2$ and $L_B = \frac{1}{|p|} \sum_p \|B(p) - B_0(p)\|^2$, where N_0, B_0 are ground truths and p is a pixel in image channel.

Then, RGB, depth representations and raw depth are combined as input of self-attention network which predicts complete depth D . The ground truth depth D_0 comes from rendering multiple raw depth views of the same scene. The loss for self-attention network is $L_{SA} = \frac{1}{|p \in obs|} \sum_{p \in obs} \|D(p) - D_0(p)\|$. Note that we only consider observed pixels, which have valid depth values in ground truth D_0 , as mentioned in [39]. Missing parts in ground truth depth map are ignored. We denote $p \in obs$ to be observed pixels.

To enhance structural information, we add small fraction of structural related loss L_S , Structural Similarity Index (SSIM) [33]. SSIM measures quality assessment based on the degradation of structural information. In our task, the higher SSIM index is, the more structured the complete depth is. With SSIM, we expect our network to learn higher quality and undistorted depth map with better structure.

Finally, complete depth D is fed into boundary consistency network and outputs B . The ground truth depth D_0

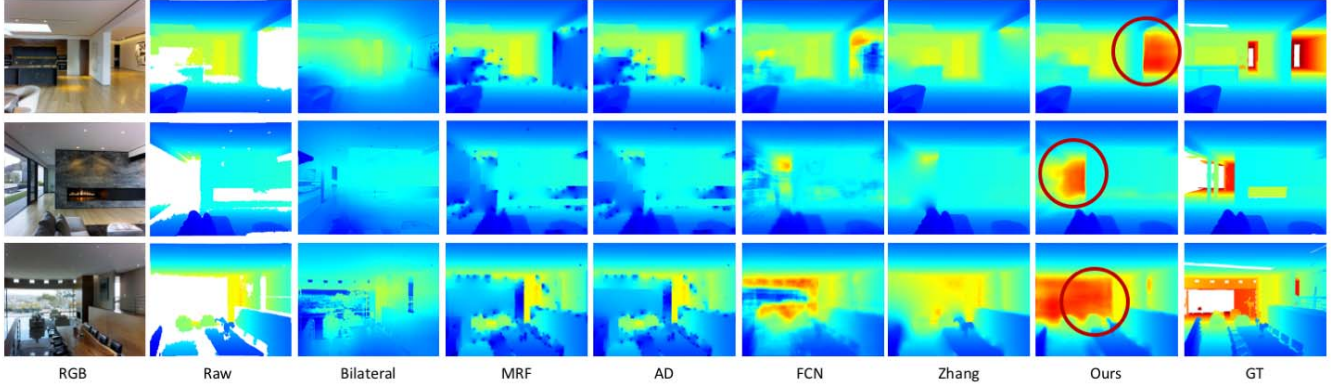


Figure 3. In the visualization result, we can see that our method learns better on the geometric meaning. For example, on the first row, our model notices the corridors on the right hand side next to the pillar and judge the depth to be deep (mark in red). Other methods tend to interpolate the large vacancy with nearby values, like Zhang fills in light blue depth value.

generates ground truth B_0 with Sobel [17]. The boundary consistency loss is $L_{BC} = \frac{1}{|p|} \sum_p \|B(p) - B_0(p)\|$.

4. Experimental Results

4.1. Dataset

We did not use the widely-known NYUv2 dataset [30], since the dataset does not provide complete ground-truth depth map for a single captured RGB-D image. Instead, we ran the experiments on Matterport3D dataset [3] to evaluate our proposed methods. Matterport3D is an indoor large-scale RGB-D dataset with 10.8k real panoramic views and 90 real indoor scenes. The dataset is scalable, diverse and representative of indoor depth completion task. We use the same training and testing lists as Zhang [39] did. We remove large bias image pairs (about 5%), and, finally, the dataset contains about 1M training data and 474 testing data. The ground truth of Matterport3D is generated from Zhang by multi-view reconstruction; that is, multiple camera views of the same scene can reconstruct a more complete indoor information, including depth map.

4.2. Evaluation Metrics

We use similar evaluation metrics in [39]. As mentioned in 3.4, we calculate loss based on observed data; that is, observed pixel values in ground truth depth. Given ground truth depth D_0 and complete depth D , the metrics include:

- Root Mean Square Error (RMSE):

$$\sqrt{\frac{1}{|obs|} \sum_{p \in obs} \|D(p) - D_0(p)\|^2} \quad (5)$$

- Mean Error (Mean):

$$\frac{1}{|obs|} \sum_{p \in obs} \|D(p) - D_0(p)\| \quad (6)$$

- SSIM [33]:

$$\frac{(2\mu_{D_0(p)}\mu_{D(p)} + c_1)(2\sigma_{D_0(p)D(p)} + c_2)}{(\mu_{D_0(p)}^2 + \mu_{D(p)}^2 + c_1)(\sigma_{D_0(p)}^2 + \sigma_{D(p)}^2 + c_2)} \quad (7)$$

where $c_1 = 0.0001, c_2 = 0.0009$.

- δ_t : percentage of pixels within the error range t , the error range is defined by equation 8.

$$\max\left(\frac{D(p)}{D_0(p)}, \frac{D_0(p)}{D(p)}\right) < t \quad (8)$$

where $t \in \{1.05, 1.10, 1.25, 1.25^2, 1.25^3\}, p \in obs$.

4.3. Performance

As shown in Table 1, our proposed method defeats all other works, including traditional methods and previous state-of-the-art [39], and shows significant performance gain on RMSE and Mean error. Not only SSIM metric but also most delta values show improvement with boundary consistency and self-attention. SSIM is an index standing for the depth quality and structure. Delta percentage is a measurement of relative depth difference, which means the pixels closer to depth sensor requires more precise prediction. Our model attains the highest performance on all of the above evaluation metrics.

Figure 3 shows the visualization results. Previous methods sometimes misunderstands the geometric meaning of one scene once raw depth is missing large coverage of depth value, which cause copy and interpolation. On the contrary, our self-attention network realizes the geometric features from RGB, normal and boundary and produces depth maps with more accurate geometric meaning than others. The red

¹We reproduce Zhang's work on github and find they use root median square error instead of root mean square error. Thus, we show the results reproduced ourselves.

Model	RMSE↓	Mean↓	SSIM↑	1.05 ↑	1.10 ↑	1.25 ↑	1.25 ² ↑	1.25 ³ ↑
Bilateral	1.978	0.774	0.507	0.385	0.497	0.613	0.689	0.730
MRF [11]	1.675	0.618	0.692	0.506	0.556	0.651	0.780	0.856
AD [21]	1.653	0.610	0.696	0.503	0.560	0.663	0.792	0.861
FCN	1.262	0.517	0.605	0.397	0.527	0.681	0.808	0.868
Zhang [39] ¹	1.316	0.461	0.762	0.657	0.708	0.781	0.851	0.888
Ours	1.092	0.342	0.799	0.661	0.750	0.850	0.911	0.936

Table 1. We achieve state-of-the-art performance on Matterport3D dataset. We compare our method with some inpainting methods, including traditional methods, such as anisotropic diffusion, joint bilateral filter, FCN-based methods (ResNet18) and Zhang [39] et al. The result shows our significant improvement on all of the evaluation metrics.

circle in Figure 3 shows the power of self-attention mechanism on our task. Our network notices detailed geometric properties, like walls and corridors. Besides, with boundary consistency Figure 4, our network learns architectural details, like the wall’s edge, and produces depth with desirable quality.

4.4. Ablation Studies

For better understanding of our works, we investigate the effects of each component of our method. The following ablation study concentrates on the following three things: self-attention, SSIM loss and boundary consistency.

4.4.1 Self-Attention

The first two rows of Table 2 exam whether self-attention mechanism has positive effect on the result of complete depth. Indeed, self-attention makes huge gains of performance over traditional FCN models. In this setting, we take ResNet18 [13], which has similar parameters, as the classic FCN method.

As we can see in Figure 5, on the right hand side of corridors next to the pillar, model with self-attention(SA) notices those depth value should be distant, while without SA misunderstands the geometric meaning. We believe that the improvement comes from the attention on convolution features to helps the model focus on important regions and features. In this case, self-attention mechanism helps model to learn more geometric meanings.

4.4.2 SSIM Loss

By adding small weight of SSIM loss to optimize, the self-attention network learns to balance structural information without degradation much on RMSE and delta percentage. The straightforward results that SSIM score improves 8.6% after adding SSIM loss represent network successfully produce higher quality of depth map values.

As we can see in Figure 5, on the second row of columns w/SA and SA+SSIM, the background surface is smooth and accurate. Also, the edge detected from Sobel shows the enhancement of depth image quality and reduction in noise,

which can be seen from the layered lines on the floor and ceiling.

4.4.3 Boundary Consistency

The last two rows of Table 2 show boundary consistency loss helps a lot in generating better depth map. There is extraordinary improvement in SSIM and delta percentage. 65.7% of $\delta 1.05$ and 75.8% of $\delta 1.10$ require highly sophisticated prediction of depth map for most of depth values. Besides, 3% of SSIM gains implies boundary consistency helps enhance the structure and sharpness of boundary and validates the effectiveness of boundary consistency.

As we can see in Figure 5, the structure of the bottom left sofa is clearer both in depth map and Sobel edge detection. Besides, the floor and ceiling are smoother with boundary consistency shown in the Sobel edge detection.




Figure 4. With the help of boundary consistency and SSIM loss, our method produces depth map of clearer structures. The upper row is original output depth, and the lower row is the partial amplification of the upper one. In the left two columns, our wall structure is clear. In the right two columns, our depth quality is higher, since boundary detail is not blurred.

4.5. Comparison

In this section, we compare our main advantages over previous state-of-the-art work [39]. Mentioned in Section 4.3, self-attention mechanism can be easier to learn geometric information than other methods. We find Zhang’s method depended heavily on raw depth. Once there are large holes in raw depth, like views of a wide and deep

Model	RMSE↓	Mean↓	SSIM↑	1.05 ↑	1.10 ↑	1.25 ↑	1.25 ² ↑	1.25 ³ ↑
W/O SA	1.262	0.517	0.605	0.397	0.527	0.681	0.808	0.868
SA	1.095	0.400	0.706	0.497	0.629	0.785	0.881	0.923
SA+SSIM	1.096	0.397	0.767	0.488	0.626	0.787	0.884	0.926
SA+SSIM+BC	1.092	0.342	0.799	0.661	0.750	0.850	0.911	0.936

Table 2. Effect of each component of our model. We show that our proposed self-attention mechanism (SA), SSIM loss and boundary consistency (BC) enhances completion performance.




Figure 5. In the visualization result, we can see our proposed SA+SSIM+BC (self-attention, SSIM loss and boundary consistency) performs the best. The upper row is the Sobel edge detection results from the depth images of bottom row. From left to right images, we can see the depth background surface and bottom left sofa gradually improve both in structure and boundary. Model can learn better geometric meanings and semantics with the help of SA, SSIM and BC.

space, Zhang’s method fell short of understanding the geometric meaning and easily fit onto noises in raw depth. The results seems to be simple interpolation but not realization of the whole view via clues from RGB and depth representation features.

Running real-time of sensing depth and training time can be an issue for most depth completion models. Our work is end-to-end trainable and training time is faster than optimization-based algorithms. Zhang proposed model that was not end-to-end trainable and used global optimization. Our proposed method overcomes those difficulties and outperforms previous works.

4.6. Attention Map

To show the self-attention mechanism can capture the structure information, we randomly sample the attention output from our network. The result is shown in Figure 6. The attention map is able to distinguish the occlusion boundaries and the geometric shape of the objects in the scene.

5. Conclusion

In this work, we propose to use self-attention mechanism and boundary consistency to improve performance in depth completion task. Self-attention mechanism allows network to learn more geometric meanings and complete depth with

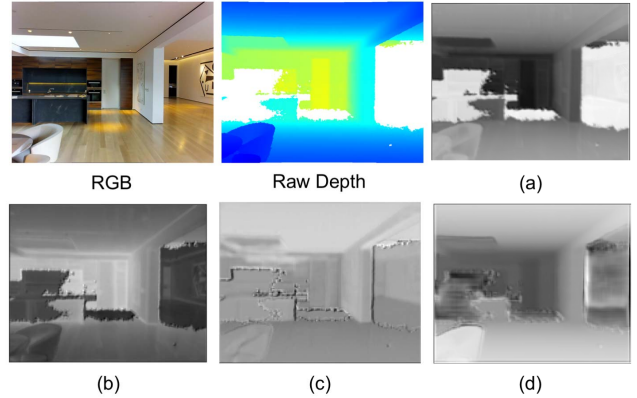


Figure 6. Visualization of some self-attention maps. Some filters pays attention on the missing holes (a), and some on the background and structures (b). More importantly, some filters pay attention on the semantics and object boundaries (c). For decoding phase, (d) shows the geometric meaning that the missing depth value behind the white walls at the right of the image should be relatively distant.

more precise values. Boundary consistency improves both the depth boundary and image quality, producing clearer and sharper structures. Extensive experiments demonstrate that our proposed method reach state-of-the-art on Matterport3D dataset. Moreover, ablation study validates the effectiveness of each component we propose.

Acknowledgement

This work was supported in part by the Ministry of Science and Technology, Taiwan, under Grant MOST 108-2634-F-002-004, FIH Mobile Limited, and Qualcomm Technologies, Inc., under Grant NAT-410477. We are grateful to the National Center for High-performance Computing.

References

- [1] T. Bailey and H. Durrant-Whyte. Simultaneous localization and mapping (slam): Part ii. *IEEE robotics & automation magazine*, 13(3):108–117, 2006.
- [2] M. Camplani and L. Salgado. Efficient spatio-temporal hole filling strategy for kinect depth maps. In *Three-dimensional image processing (3DIP) and applications II*, volume 8290, page 82900E. International Society for Optics and Photonics, 2012.
- [3] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017.
- [4] D. Doria and R. J. Radke. Filling large holes in lidar data by inpainting depth gradients. In *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 65–72. IEEE, 2012.
- [5] H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006.
- [6] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014.
- [7] J. Engel, T. Schöps, and D. Cremers. Lsd-slam: Large-scale direct monocular slam. In *European conference on computer vision*, pages 834–849. Springer, 2014.
- [8] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [9] P. Grossmann. Depth from focus. *Pattern recognition letters*, 5(1):63–69, 1987.
- [10] Y. Han, J.-Y. Lee, and I. So Kweon. High quality shape from a single rgb-d image under uncalibrated natural illumination. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1617–1624, 2013.
- [11] A. Harrison and P. Newman. Image and sparse laser fusion for dense scene reconstruction. In *Field and Service Robotics*, pages 219–228. Springer, 2010.
- [12] C. Hazirbas, S. G. Soyer, M. C. Staab, L. Leal-Taixé, and D. Cremers. Deep depth from focus. In *Asian Conference on Computer Vision*, pages 525–541. Springer, 2018.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox. Rgb-d mapping: Using kinect-style depth cameras for dense 3d modeling of indoor environments. *The International Journal of Robotics Research*, 31(5):647–663, 2012.
- [15] M. Hornacek, C. Rhemann, M. Gelautz, and C. Rother. Depth super resolution by rigid body self-similarity in 3d. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1123–1130, 2013.
- [16] M. Jaritz, R. De Charette, E. Wirbel, X. Perrotton, and F. Nashashibi. Sparse and dense data with cnns: Depth completion and semantic segmentation. In *2018 International Conference on 3D Vision (3DV)*, pages 52–60. IEEE, 2018.
- [17] J. Kittler. On the accuracy of the sobel edge detector. *Image and Vision Computing*, 1(1):37–42, 1983.
- [18] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab. Deeper depth prediction with fully convolutional residual networks. In *2016 Fourth international conference on 3D vision (3DV)*, pages 239–248. IEEE, 2016.
- [19] J.-H. Lee, M. Heo, K.-R. Kim, and C.-S. Kim. Single-image depth estimation based on fourier domain analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 330–339, 2018.
- [20] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 85–100, 2018.
- [21] J. Liu and X. Gong. Guided depth enhancement via anisotropic diffusion. In *Pacific-Rim Conference on Multimedia*, pages 408–417. Springer, 2013.
- [22] F. Ma, G. V. Cavalheiro, and S. Karaman. Self-supervised sparse-to-dense: self-supervised depth completion from lidar and monocular camera. *arXiv preprint arXiv:1807.00275*, 2018.
- [23] O. Mac Aodha, N. D. Campbell, A. Nair, and G. J. Brostow. Patch based synthesis for single depth image super-resolution. In *European conference on computer vision*, pages 71–84. Springer, 2012.
- [24] F. Mal and S. Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE, 2018.
- [25] K. Matsuo and Y. Aoki. Depth image enhancement using local tangent plane approximations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3574–3583, 2015.
- [26] E. Mingolla and J. T. Todd. Perception of solid shape from shading. *Biological cybernetics*, 53(3):137–151, 1986.
- [27] E. Mingolla and J. T. Todd. Perception of solid shape from shading. In *Shape from shading*, pages 409–441. MIT Press, 1989.
- [28] G. Riegler, M. Rüther, and H. Bischof. Atgv-net: Accurate depth super-resolution. In *European conference on computer vision*, pages 268–284. Springer, 2016.
- [29] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [30] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision*, pages 746–760. Springer, 2012.

- [31] G. Strang and K. Aarikka. *Introduction to applied mathematics*, volume 16. Wellesley-Cambridge Press Wellesley, MA, 1986.
- [32] S. Suwajanakorn, C. Hernandez, and S. M. Seitz. Depth from focus with your mobile phone. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3497–3506, 2015.
- [33] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [34] J. Xie, R. Girshick, and A. Farhadi. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In *European Conference on Computer Vision*, pages 842–857. Springer, 2016.
- [35] Q. Yang, R. Yang, J. Davis, and D. Nistér. Spatial-depth super resolution for range images. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007.
- [36] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang. Free-form image inpainting with gated convolution. *arXiv preprint arXiv:1806.03589*, 2018.
- [37] L.-F. Yu, S.-K. Yeung, Y.-W. Tai, and S. Lin. Shading-based shape refinement of rgb-d images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1415–1422, 2013.
- [38] R. Zhang, P.-S. Tsai, J. E. Cryer, and M. Shah. Shape-from-shading: a survey. *IEEE transactions on pattern analysis and machine intelligence*, 21(8):690–706, 1999.
- [39] Y. Zhang and T. Funkhouser. Deep depth completion of a single rgb-d image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 175–185, 2018.