Predicting Complete 3D Models of Indoor Scenes supplementary material

Anonymous CVPR submission

Paper ID 1625

1. Sample Results

For more illustrations of the performance of our method, below we show 15 randomly selected examples of 3D models estimated automatically from the input RGBD images. Those examples are all from the 654 testing images in NYU v2 dataset. In each figure:

(1) **Input** represents the original input image;

(2) **Raw Depth** means the original raw depth provide by kinect;

(3) **Smoothed Depth** represents inpainted depth map based on raw depth;

(4) **Object Probability** shows the object probability map automatically computed based on the method in [1]. Lighter part in the figure represents a higher probability of being an object. This is used in the scene composition step to help encouraging rendered object and layout pixels to match the object probability;

(5) **3D Annotation** is the manual annotated 3D model for the given input image;

(6) **Rendered Segmentation** illustrates the 2D segmentation result based on the rendered object computed by our method;

(7) **Automatic 3D Model** provides results of automatically estimated 3D model based on our method. Two views are shown.

2. Video

For a more intuitive view of our proposed method and correlated results, please refer to the attached movie in the uploaded supplementary material package.

References

[1] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 1

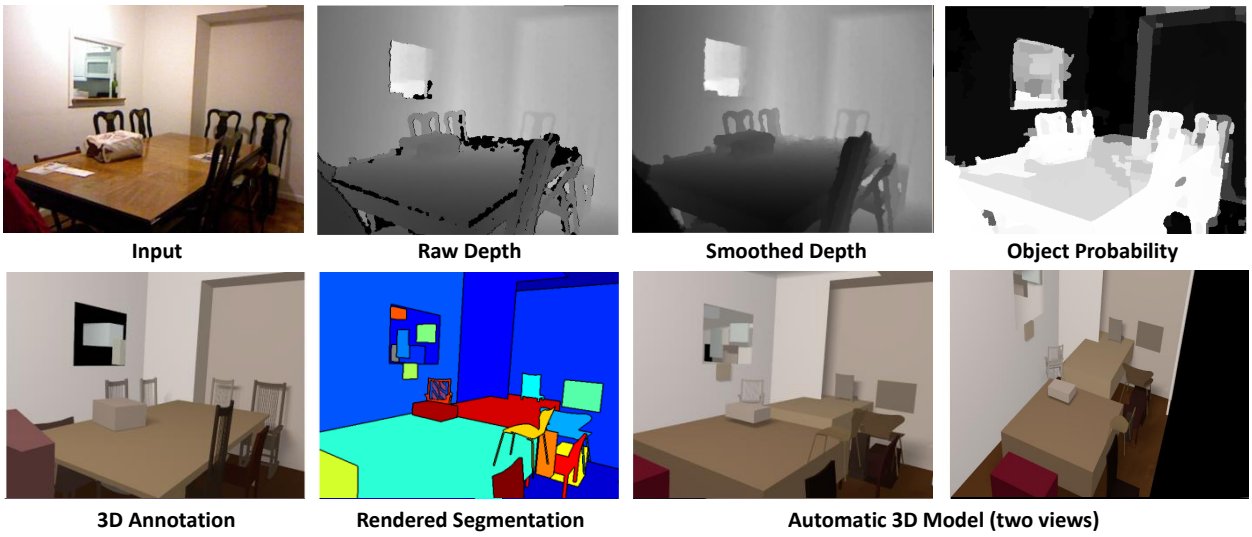

Figure 1: Occupancy Precision: 0.2669, Occupancy Recall: 0.5412, Freespace Precision: 0.9866, Freespace Recall: 0.9486, Depth Error Overall: 0.0227. The table is broken into two pieces, but otherwise this is a fairly accurate representation of the scene.

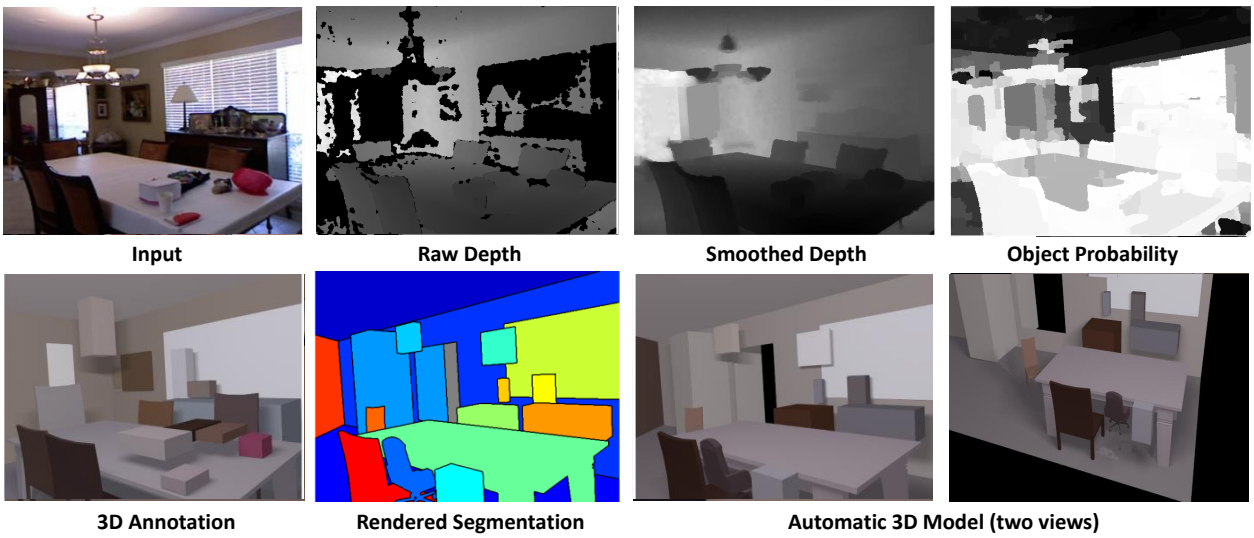

Figure 2: Occupancy Precision: 0.5221, Occupancy Recall: 0.2418, Freespace Precision: 0.8853, Freespace Recall: 0.9381, Depth Error Overall: 0.2569. Large pieces of furniture are well-modeled, though some small objects on them are missed.

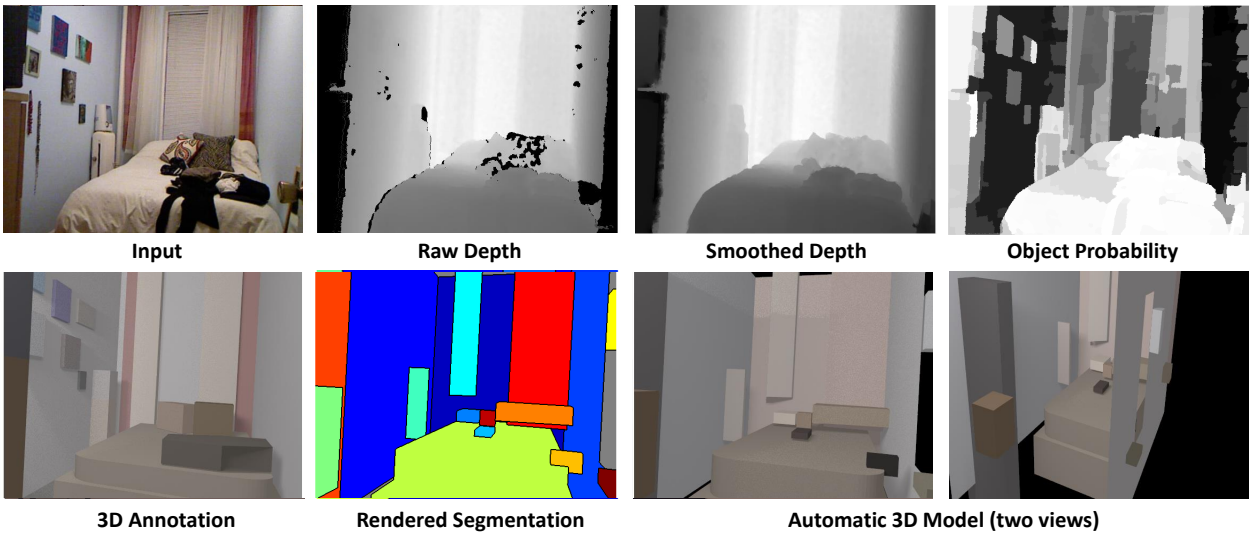

Figure 3: Occupancy Precision: 0.8943, Occupancy Recall: 0.6758, Freespace Precision: 0.9528, Freespace Recall: 0.9851, Depth Error Overall: 0.1444. Decent model of a cluttered scene.

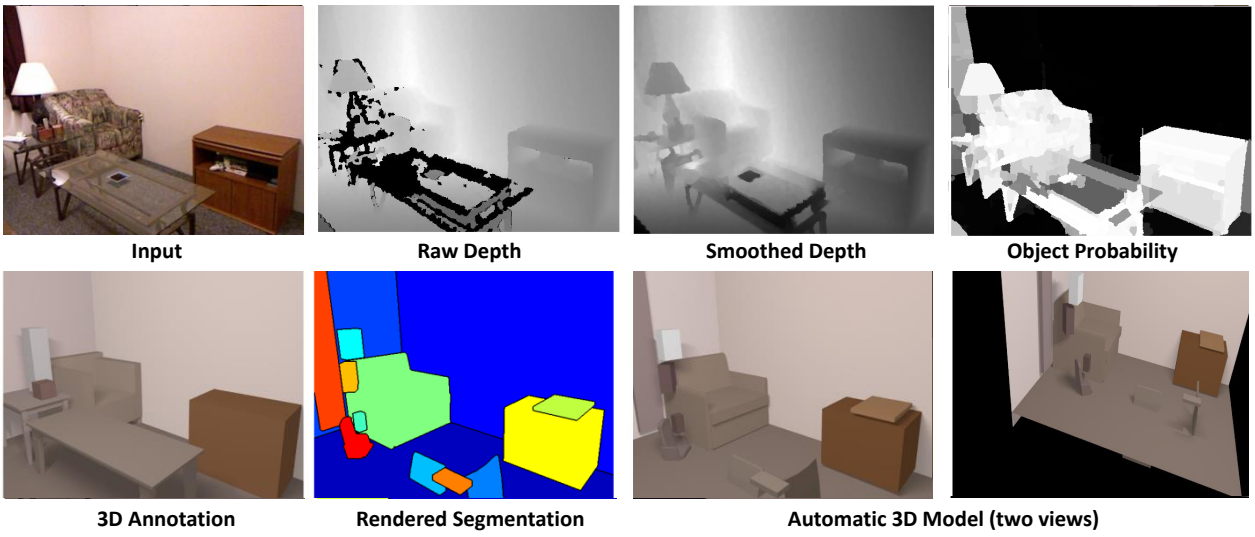

Figure 4: Occupancy Precision: 0.8206, Occupancy Recall: 0.7717, Freespace Precision: 0.9860, Freespace Recall: 0.9852, Depth Error Overall: 0.0193. Solid furniture is well-modeled, but glass tables are missed due to missing depth information (and transparency).

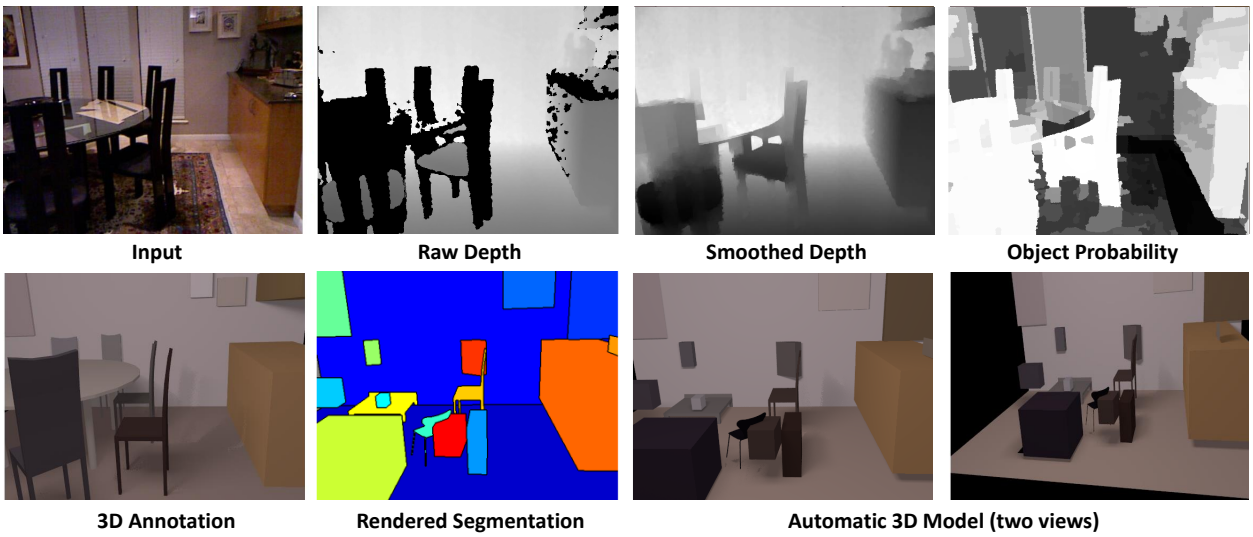

Figure 5: Occupancy Precision: 0.7337, Occupancy Recall: 0.7628, Freespace Precision: 0.9790, Freespace Recall: 0.9713, Depth Error Overall: 0.0160. Some mistakes are due to missing depth information. We use the smoothed depth image (from Silberman et al.), but depth is sometimes incorrectly completed.

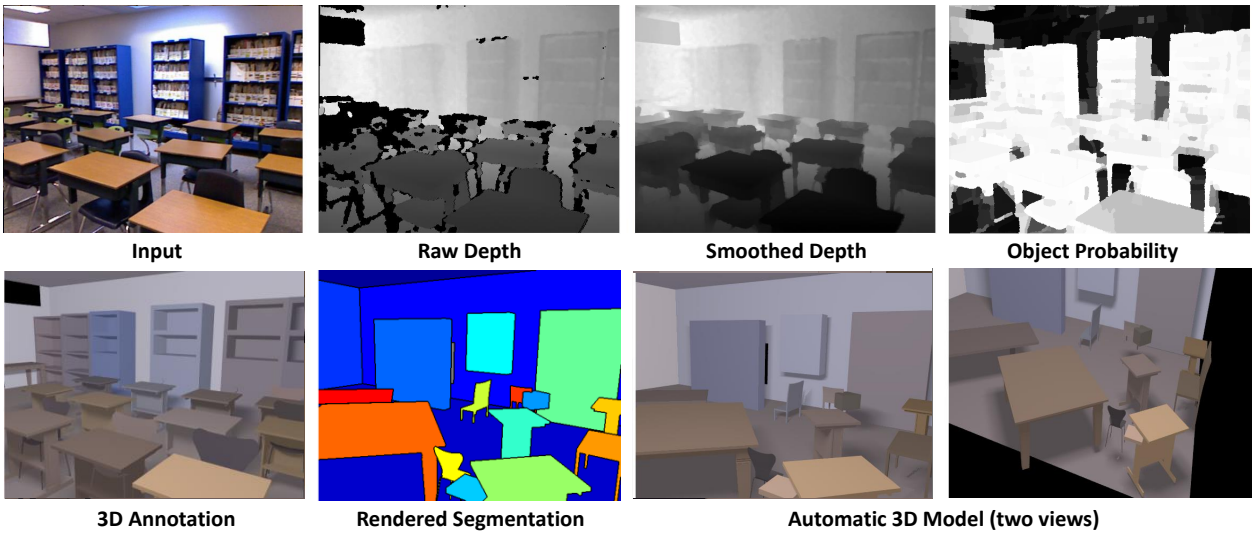

Figure 6: Occupancy Precision: 0.4954, Occupancy Recall: 0.2583, Freespace Precision: 0.9632, Freespace Recall: 0.9727, Depth Error Overall: 0.0.0558. Use of a prior that sets of tables and chairs are likely to have the same models and to be arranged in an orderly way could substantially improve results in some cases.

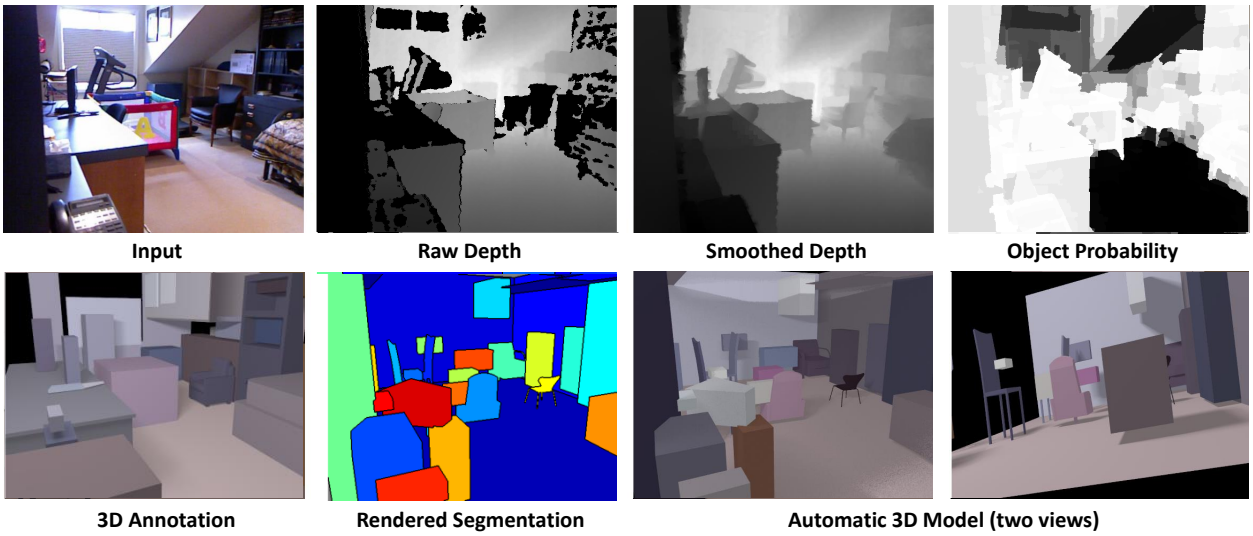

Figure 7: Occupancy Precision: 0.4662, Occupancy Recall: 0.1150, Freespace Precision: 0.9393, Freespace Recall: 0.9688, Depth Error Overall: 0.1015. Incorporation of support relationships would help ground objects and possibly improve segmentation and modeling.

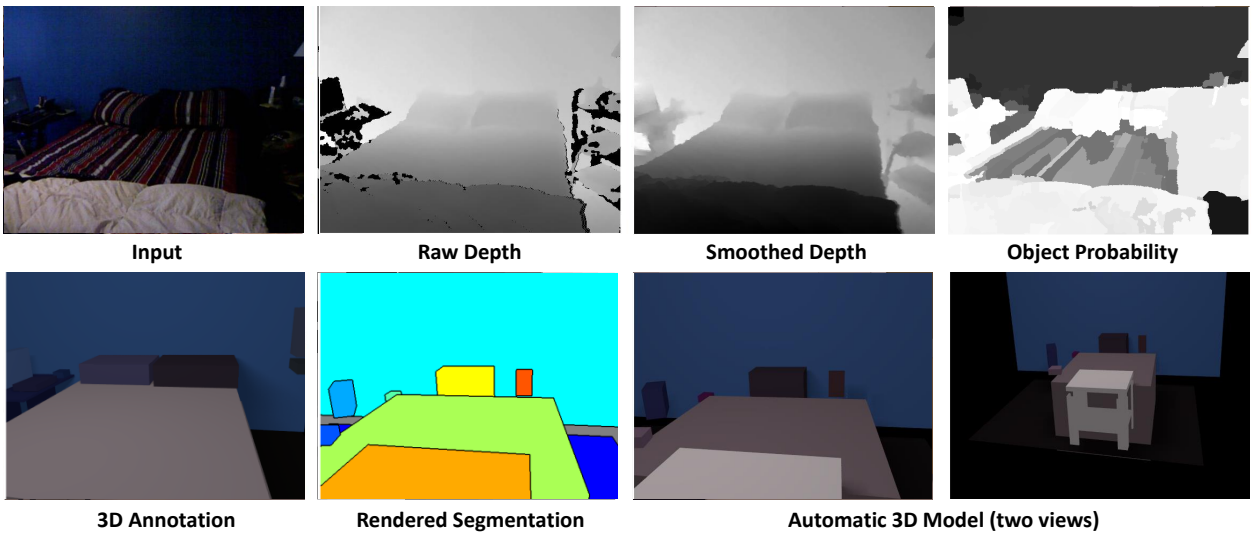

Figure 8: Occupancy Precision: 0.6472, Occupancy Recall: 0.6330, Freespace Precision: 0.9645, Freespace Recall: 0.8361, Depth Error Overall: 0.2578.

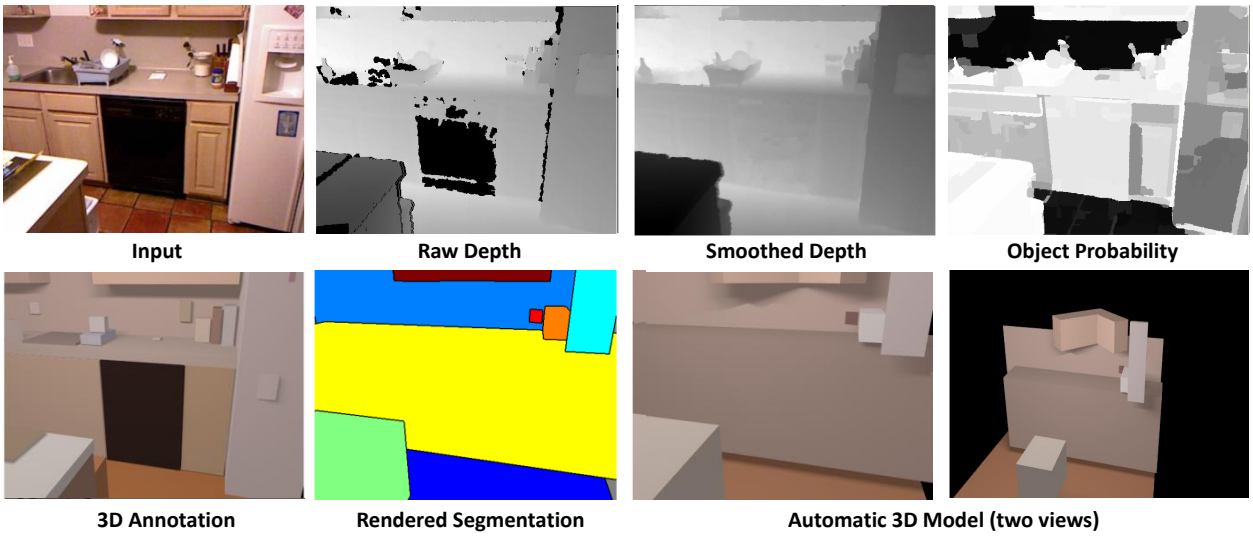

Figure 9: Occupancy Precision: 0.9180, Occupancy Recall: 0.4806, Freespace Precision: 0.7749, Freespace Recall: 0.9758, Depth Error Overall: 0.0158.

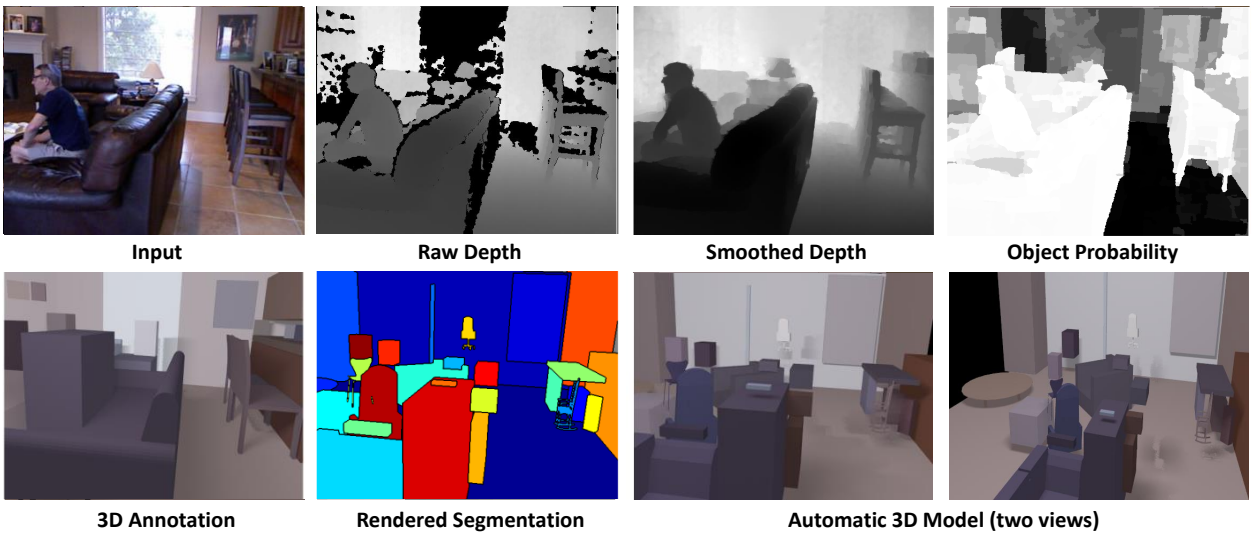

Figure 10: Occupancy Precision: 0.2985, Occupancy Recall: 0.1327, Freespace Precision: 0.9065, Freespace Recall: 0.9217, Depth Error Overall: 0.0853.

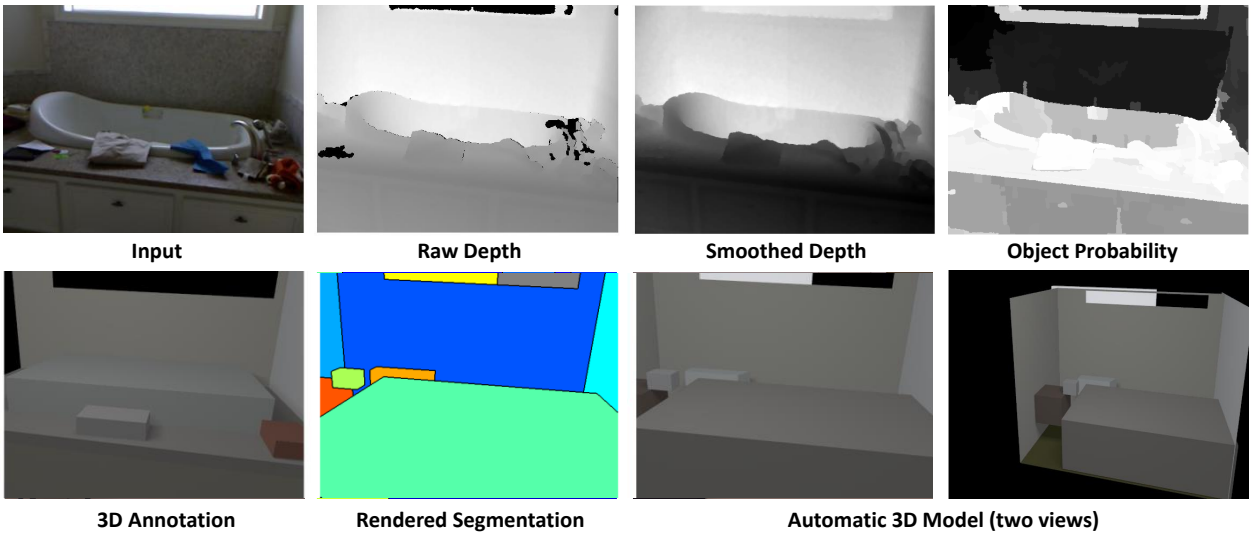

Figure 11: Occupancy Precision: 0.9714, Occupancy Recall: 0.6227, Freespace Precision: 0.9489, Freespace Recall: 0.9667, Depth Error Overall: 0.1394.

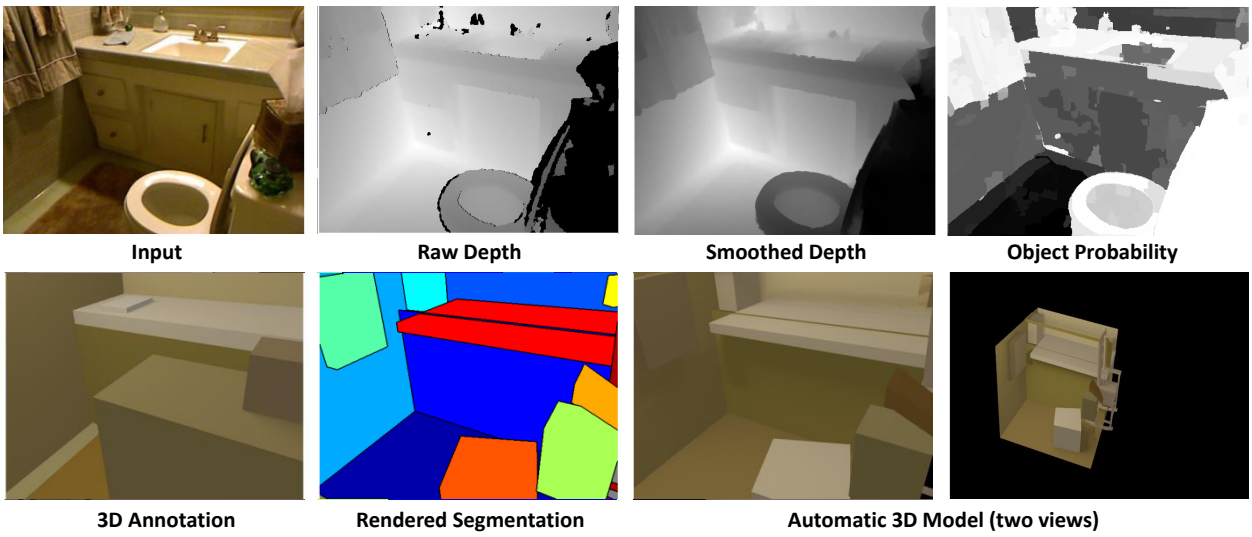

Figure 12: Occupancy Precision: 0.7183, Occupancy Recall: 0.0989, Freespace Precision: 0.9035, Freespace Recall: 0.9696, Depth Error Overall: 0.1513.

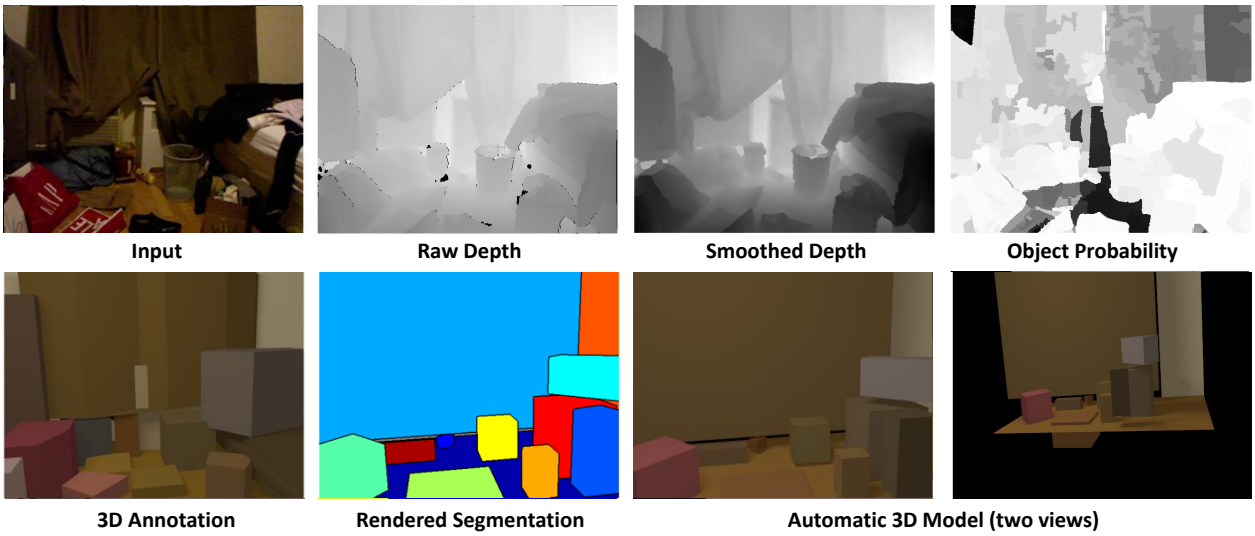

Figure 13: Occupancy Precision: 0.7150, Occupancy Recall: 0.2085, Freespace Precision: 0.9776, Freespace Recall: 0.8347, Depth Error Overall: 0.2530.

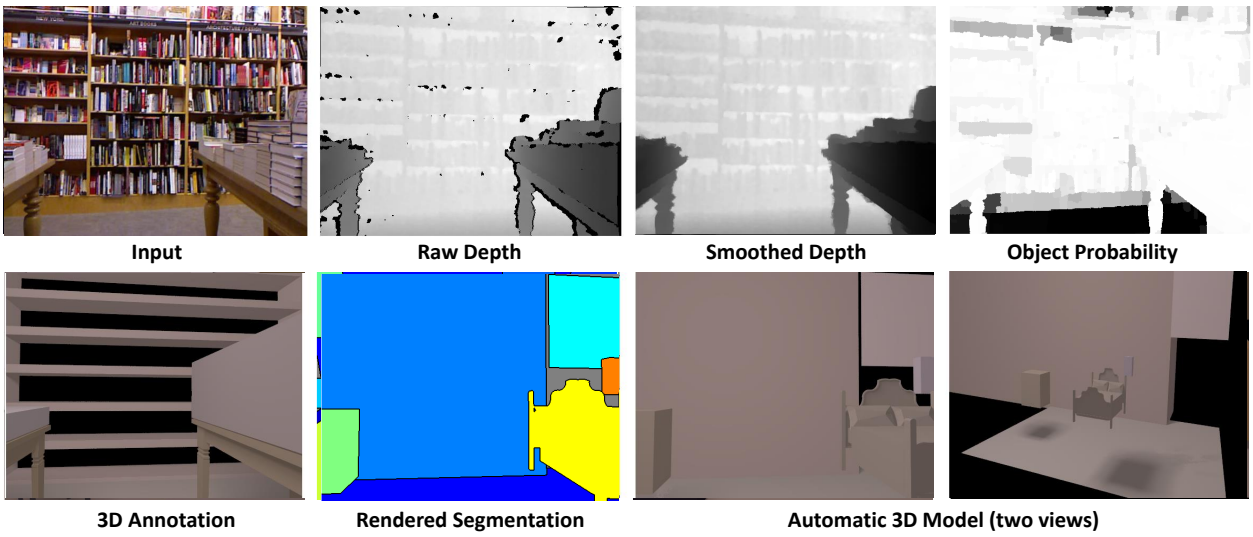

Figure 14: Occupancy Precision: 0.6706, Occupancy Recall: 0.4353, Freespace Precision: 0.9979, Freespace Recall: 0.9973, Depth Error Overall: 0.0448.

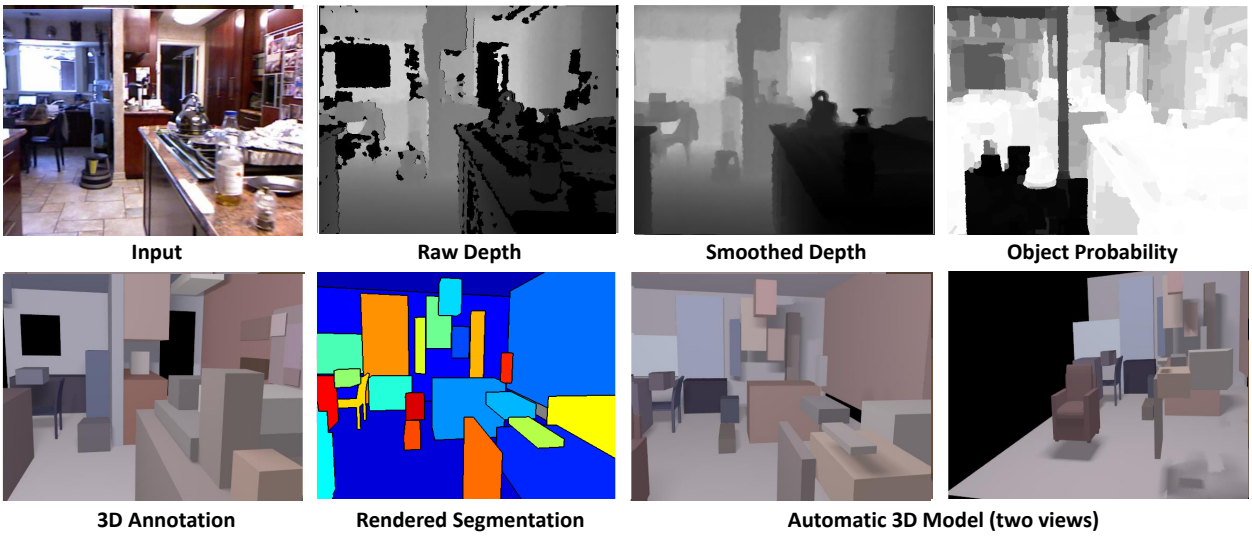

Figure 15: Occupancy Precision: 0.4330, Occupancy Recall: 0.1800, Freespace Precision: 0.9099, Freespace Recall: 0.8911, Depth Error Overall: 0.4816.