Reconstructing the World's Museums: Supplementary Material

Jianxiong Xiao[1] and Yasutaka Furukawa[2]

[1] Massachusetts Institute of Technology
[2] Google Inc.

Abstract. This document provides further details of the view dependent model constructions proposed in the paper [1], and show more reconstruction results. We also provide more results and videos on our project website: http://mit.edu/jxiao/museum/.

1 View dependent model construction

CSG model manipulation also allows us to lower back-facing walls to improve visibility from a specific viewing angle (See Fig. 1). Let v be the viewing direction, for which the model is to be optimized, and denote h_f and h_b as the maximum height of the structure to be reconstructed for the front-facing and back-facing walls, respectively. For each primitive, we translate the primitive along each of the XYZ axes along v direction by $\Delta(= 0.5\text{m})$. Let us denote this translated model as T^v, then $(T^v - T^{\text{up}}) - \text{Rect}(h_f)$ generates front-facing facades with the height limited at h_f. Similarly, let $\overline{v}$ be the inverted reflection vector of v against the ground plane, then the back-facing facades are modeled by $(T^{\overline{v}} - T^{\text{up}}) - \text{Rect}(h_b)$. The finally view-dependent 3D model is obtained by taking their union:

$$((T^v - T^{\text{up}}) - \text{Rect}(h_f)) + ((T^{\overline{v}} - T^{\text{up}}) - \text{Rect}(h_b)). \quad (1)$$

h_f and h_b are set to 6 and 2 meters, respectively, in our experiments.

2 More Reconstruction Results

In this section, we show the texture-mapped 3D model and the extracted face-groups for our largest dataset, The Metropolitan Museum of Art. The model is shown in an opposite angle from the main paper with many close-ups, including explanations and assessment of failure cases. See Figures 2, 3, 4, 5 and 6 for the details.

References

1. Xiao, J., Furukawa, Y.: Reconstructing the world's museums. In: ECCV. (2012)

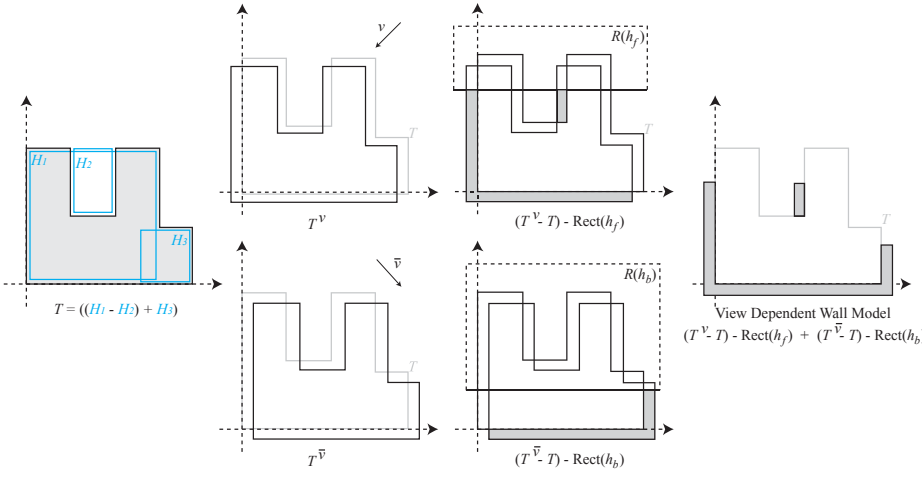

Fig. 1. View dependent model construction.

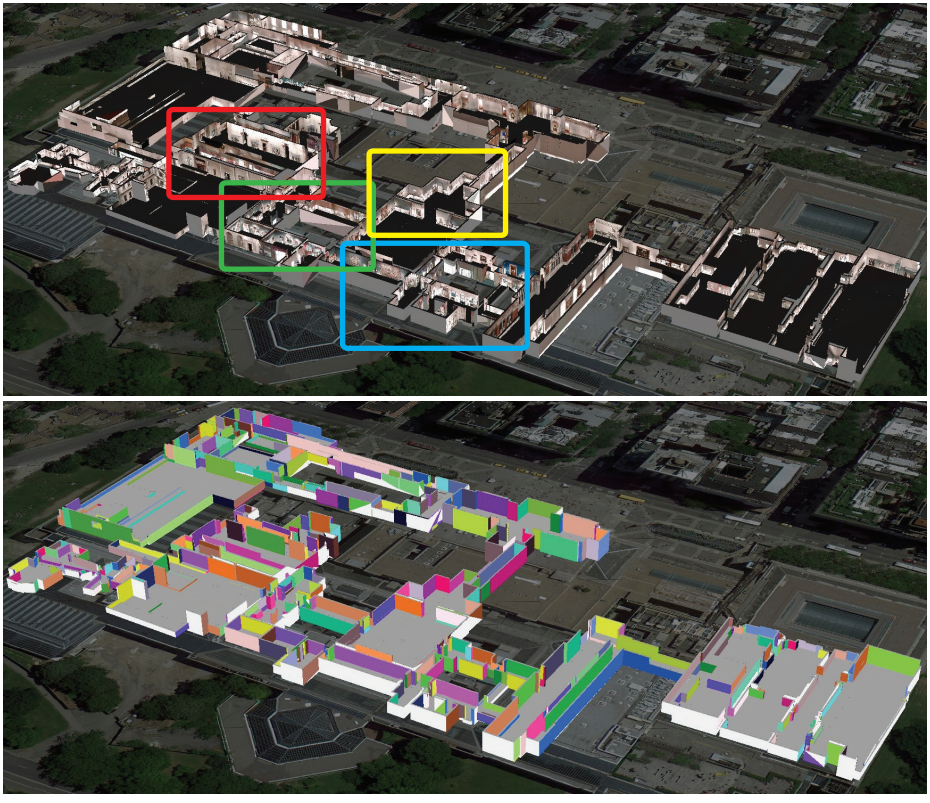

Fig. 2. The texture-mapped 3D model (top) and the extracted face-groups (bottom) for The Metropolitan Museum of Art. Further close-up views, highlighted in the four color rectangles, are given in Figures 3, 4, 5, and 6.

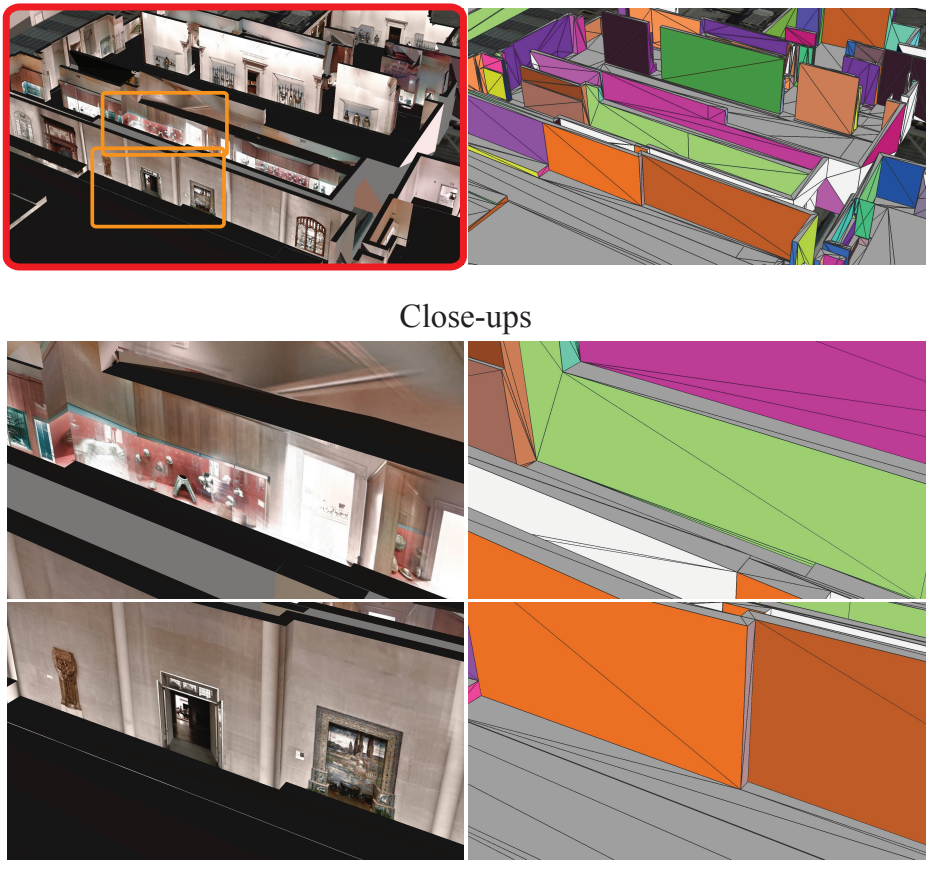

Fig. 3. The middle row shows ghosting artifacts due to the lack of precise 3D geometry, which however succeed in giving a *sense* of a place in a zoomed-out view for navigation purposes.

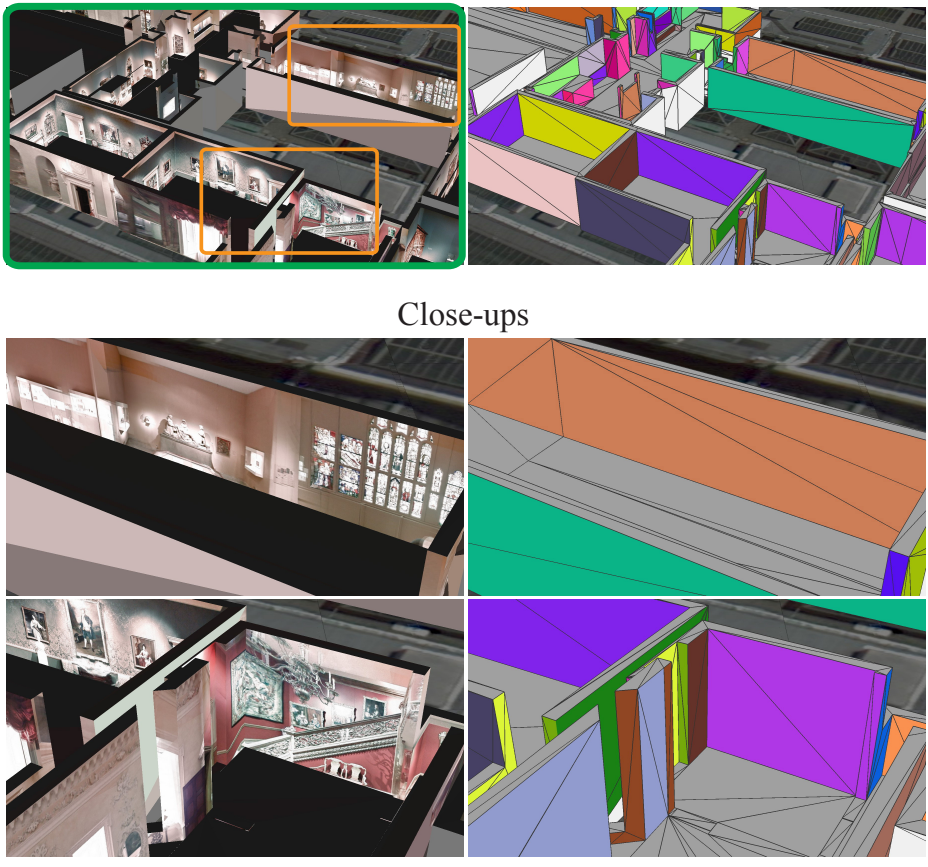

Fig. 4. In the bottom row, complex scene structures such as a chandelier and stairs are visible in the stitched texture. While their geometry is not represented in the CSG model, the texture is free from major artifacts except for some distortions.

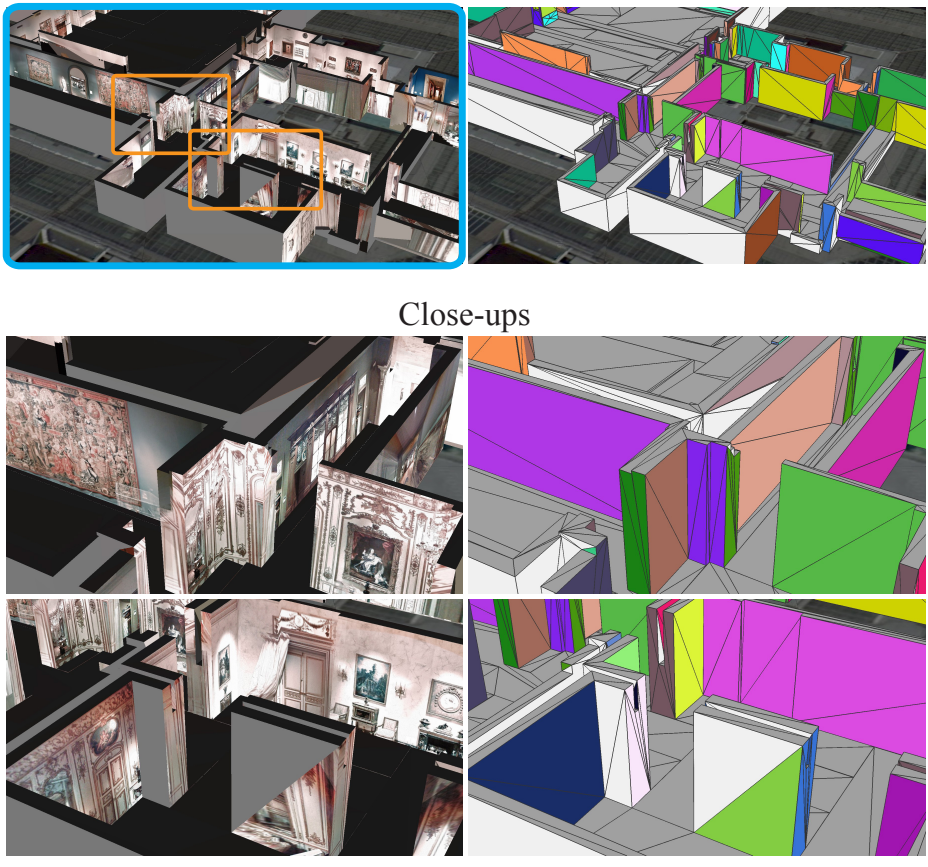

Fig. 5. In the bottom row, textures are missing in several faces probably due to the input pose errors and the lack of input images suitable for texture mapping.

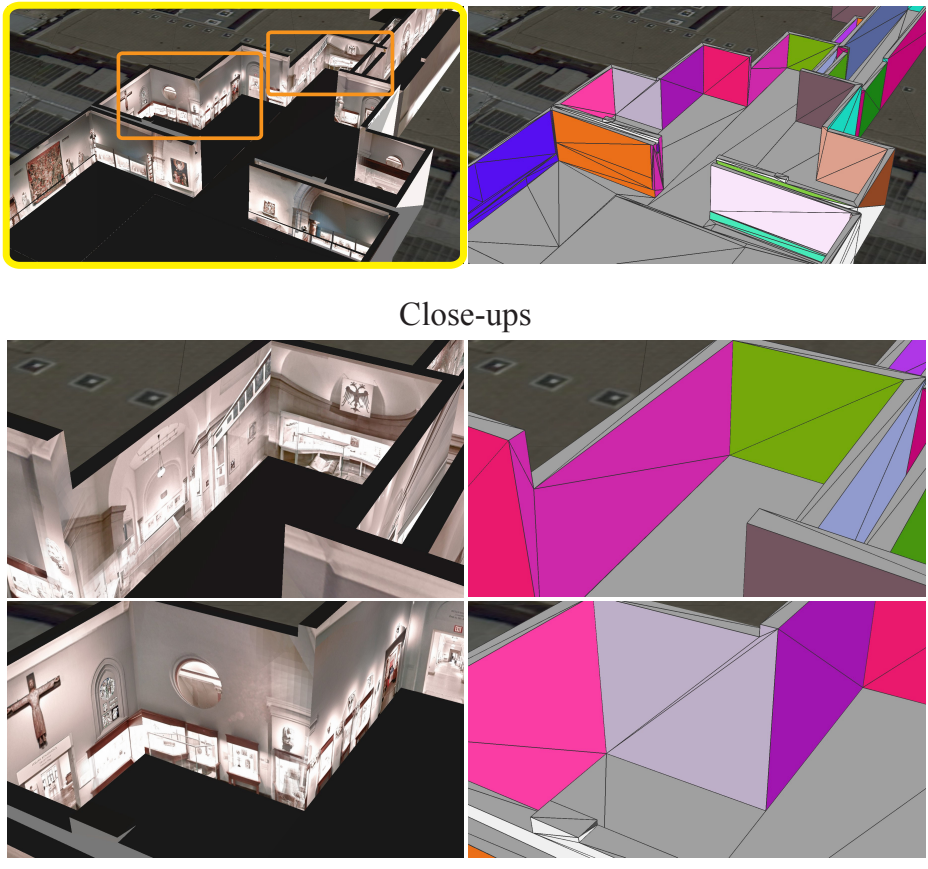

Fig. 6. In the middle row, the texture of an entire room is mapped onto a single plane due to the lack of precise geometry. However, the stitched texture is free from major artifacts, except for an inconsistent perspective.