Self-Supervised Large Scale Point Cloud Completion for Archaeological Site Restoration

Supplementary Material

1. General Comparison with Other Methods

We list state-of-the-art methods related to our topic (e.g. supervised/unsupervised, point cloud based/view based, etc.) and state their capability of handling the discussed features under our addressed setting in Tab. A1.

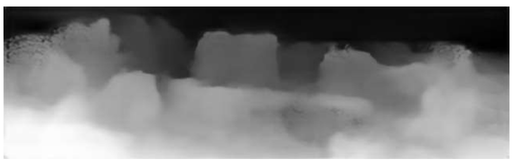

Figure A1. Completion results of S4C [4], a scene-level completion method, on the structure of Fig.1 main paper.

Scene-level completion methods target predicting the scene with RGB and semantic labels given a single-view depth map or RGB-D images [10]. Due to the limitied precision of the representation, the completion results of such methods are coarser and more inaccurate compared to ours, as shown in Fig. A1. In general, scene-level completion methods complete points occluded from one view, but cannot address points missing from all the views.

2. Additional Visualization & Analysis

In this section we show more visualizations of the completion to walls/structures from archaeological sites not shown in the main text due to limited space. Please see Fig. A2 for more completion results on different types of walls/structures.

Restoration results on *Kuelap-Raw* is shown in Fig. A3, which contains larger sloped walls.

We qualitatively analyze the importance of different loss terms in Fig. A4. As shown in Fig. A4b and A4e, L_{cons} serves as a strong regularization to enforce inpainted MCOP images to reproduce input patterns, thus improving inpainting fidelity and reconstruction quality. L_{adv} serves as the key component on fidelity of inpainted regions (see Fig. 10b,c,d and Fig. A4c). L_{sim} does not influence the fidelity of single patches, but helps maintain luminance and texture consistency at a larger receptive field (see Fig.9 and Fig. A4d), which notably improve quality, especially for large and complex structures.

3. Additional Notes on MCOP Images

We further visualize the details of our enhanced MCOP representation. The scanline trajectory when scanning the structure in the title figure of the main paper is shown in Fig. A5, which shows the transition from side to top for multiple consecutive slits.

We also show how the rotation channel enables efficient shape manipulation of the completion. By annotating a rough position of where the slits should turn from side to top (e.g. Fig. A6), we are able to control, for example, whether to complete a door or not in the output point cloud.

It is noted that the rotation channel annotation does not need to be precisely annotated by the user to receive perfect results. Due to efficient adversarial training on the rotation channel, the completer learns to predict a smooth and genuine looking rotation channel based on the coarse user annotation.

4. Dataset and training

We list a group of important statistics for the evaluation dataset in Table A2. As seen, the most important structures consist of nearly millions of points but are still missing many more points.

The corresponding histogram distribution for completion ratios of structures from *Huamanmarca Raw* is shown in Fig. A7. Unlike *Mawchu Raw* whose structures have completion of varying levels, structures in *Huamanmarca* only have uniformly 20% to 40% available points in the bottom part.

For sites which are reconstructed under illuminance imbalance, we try to recover (or deshadow) with a manual segmentation stage using PhotoShop and separate the relatively dark regions in the MCOP images. The pixels in the dark regions are then transformed using histogram matching with the relatively bright regions within the same structure (e.g. same MCOP image) to maintain better local consistency. A few deshadow results are shown in Fig. A8, where the shadow removed structures exhibits better illuminance consistency.

References

[1] Jingdao Chen, John Seon Keun Yi, Mark Kahoush, Erin S. Cho, and Yong K. Cho. Point cloud scene completion of obstructed building facades with generative adversarial inpainting. *Sensors*, 20(18), 2020. 2

[2] Xuelin Chen, Baoquan Chen, and Niloy J Mitra. Unpaired point cloud completion on real scans using adversarial training. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. 2

[3] Ruikai Cui, Shi Qiu, Saeed Anwar, Jiawei Liu, Chaoyue

Methods	Self Supervised	Spatially imbalanced distribution	Point cloud size	Color
Pcl2Pcl[2]	✗	✗	$10k$	✗
PointPnCNet[6]	✓	✗	$10k$	✗
P2C[3]	✓	✗	$10k$	✗
ACL-SPC[5]	✓	✗	$10k$	✗
Facade[1]	✗	✗	$> 100k$	✓
MCOP [8]	✗	✗	$> 100k$	✓
UAIR [7]	✓	✗	$10k$	✓
SSIM [9]	✓	✗	$10k$	✓
MCOP + UAIR/SSII	✓	✗	$> 100k$	✓
Ours	✓	✓	$> 100k$	✓

Table A1. Comparison of our method with other SOTA methods in the field of self-supervised point cloud completion. The methods are divided into purely point cloud based ones and view based ones. Our solution works best under the archaeological setting, for example.

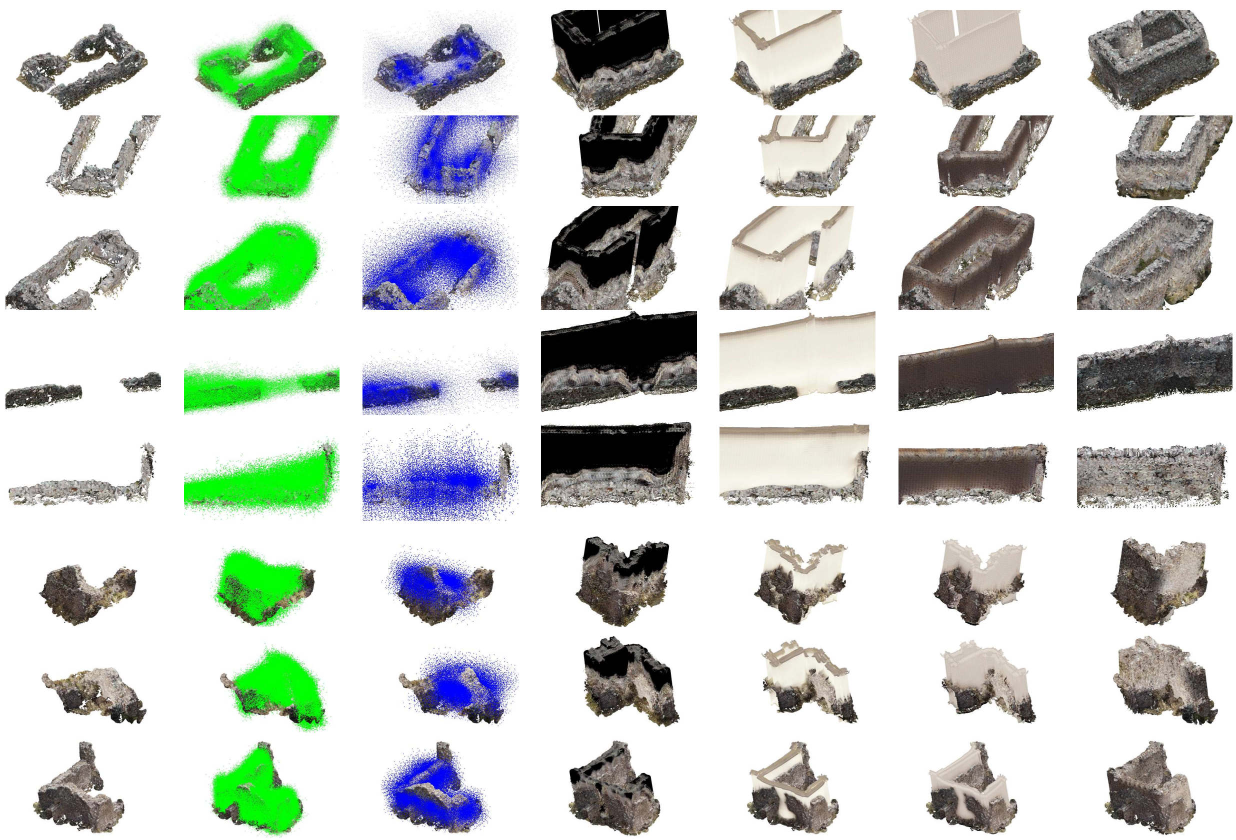

(a) Inputs (b) P2C (c) ACL-SPC (d) UAIR (e) SSII (f) SSII(Patch) (g) Ours

Figure A2. Additional Visualization of Completion Results of Different Methods.

Source Completed Source Completed

Figure A3. Additional restoration results on site Kuelap.

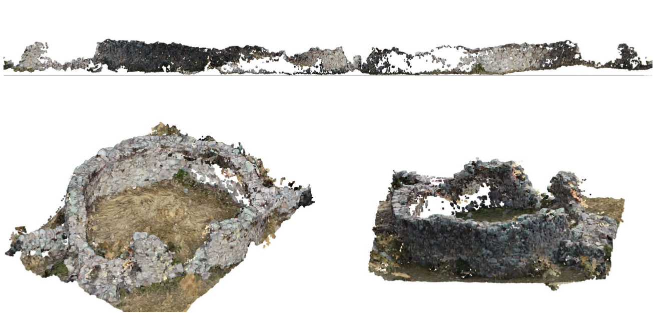

(a) Inputs

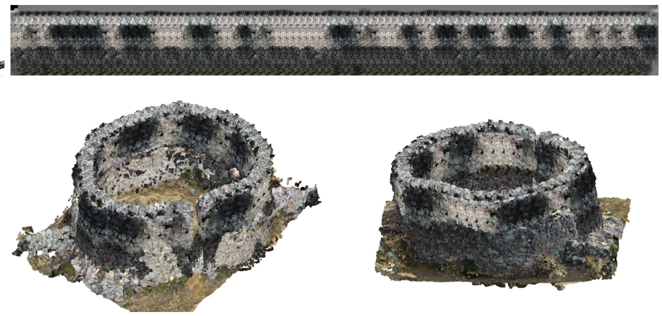

(b) Ours without L_{cons}

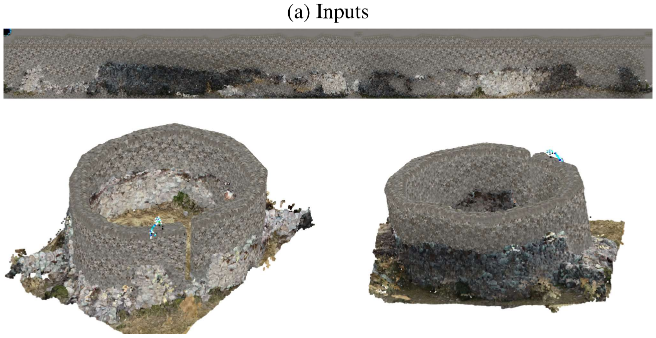

(c) Ours without L_{adv}

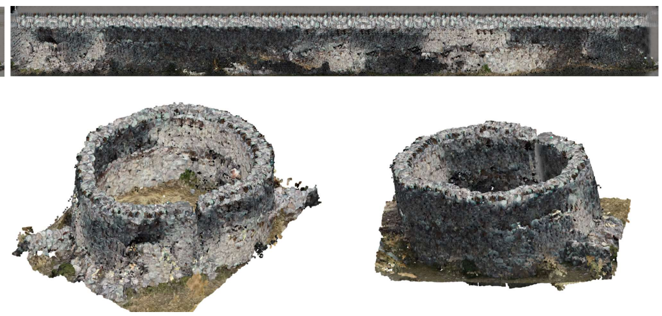

(d) Ours without L_{sim}

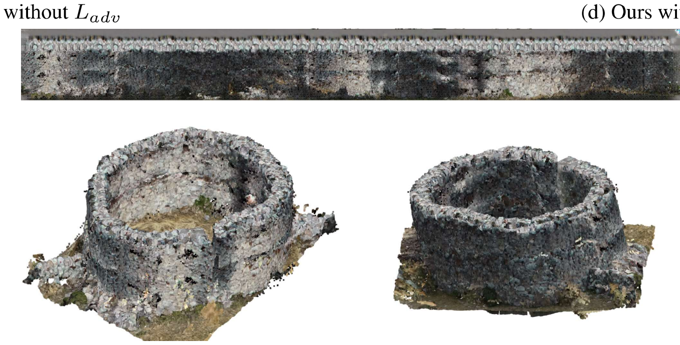

(e) Ours

Figure A4. Inpainted MCOP images and restorations on the same structure with/without L_{cons}, L_{adv}, and L_{sim}.

Xing, Jing Zhang, and Nick Barnes. P2c: Self-supervised point cloud completion from single partial clouds, 2023. 2

[4] Adrian Hayler, Felix Wimbauer, Dominik Muhle, Christian Rupprecht, and Daniel Cremers. S4c: Self-supervised semantic scene completion with neural fields, 2023. 1

[5] Sangmin Hong, Mohsen Yavartanoo, Reyhaneh Neshatavar, and Kyoung Mu Lee. Acl-spc: Adaptive closed-loop system for self-supervised point cloud completion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2

[6] Himangi Mittal, Brian Okorn, Arpit Jangid, and David Held. Self-supervised point cloud completion via inpainting. *arXiv preprint arXiv:2111.10701*, 2021. 2

[7] Arthur Pajot, Emmanuel de Bezenac, and Patrick Gallinari. Unsupervised adversarial image reconstruction. In *International Conference on Learning Representations*, 2019. 2

[8] Paul Rademacher and Gary Bishop. Multiple-center-of-projection images. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, page 199–206, New York, NY, USA, 1998. Association for Computing Machinery. 2, 4

[9] Sriram Yenamandra, Ansh Khurana, Rohit Jena, and Suyash P. Awate. Learning image inpainting from incomplete images using self-supervision. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 10390–10397, 2021. 2

[10] Xumin Yu, Yongming Rao, Ziyi Wang, Jiwen Lu, and Jie Zhou. Adapointr: Diverse point cloud completion with adaptive geometry-aware transformers, 2023. 1

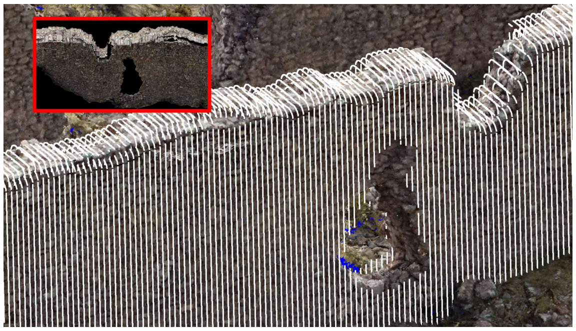

Figure A5. Zoom in view for the structure from the teaser figure in the main paper with the scan line of the camera overlaid. Unlike typical vertical scan trace in [8], we move the camera smoothly from the side to the top of the target structure, which effectively maintains the geometric locality for points distributed around the transition part.

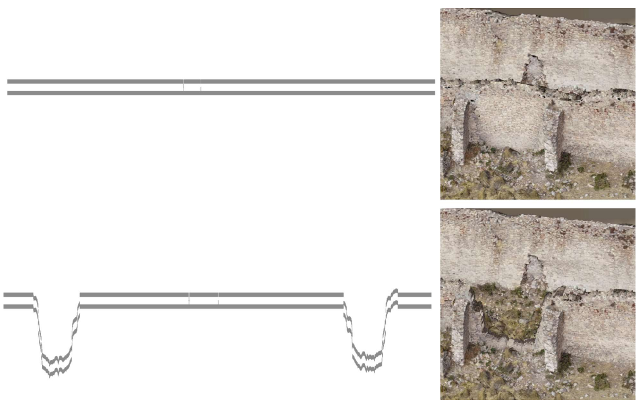

Figure A6. Illustration of how different rotation channels could encode different shape priors for the final completion shape. A uniform rotation map would encourage completion of same height everywhere, while leaving out the necessary boundaries would enable carving out important openings like doors.

Stats	**Mawchu**	**Huaman**
# of structures	483	141
Min size	0.74k	485k
Max size	15 m	11m
Mean size	712 k	1.6m
Min MCOP size	0.49 k	103k
Max MCOP size	1.9 m	1.5m
Mean input MCOP size	70 k	575k
Min output size	48 k	348k
Max output size	3.1 m	4.46m
Mean output size	0.31 m	1.87m
Min missing ratio	16.8%	57.2%
Max missing ratio	83.2%	76.5%
Mean missing ratio	77.0%	68.7%

Table A2. Dataset statistics used for qualitative/quantitative evaluation.

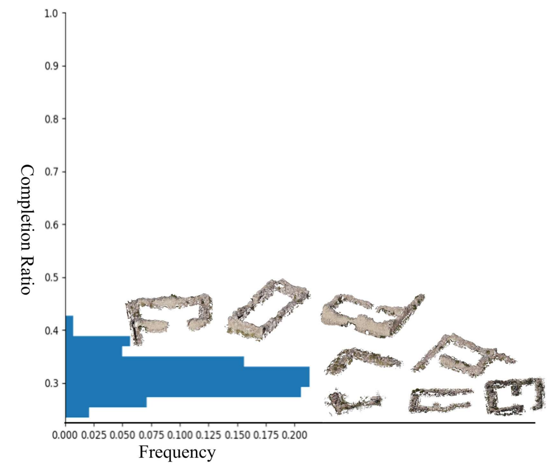

Figure A7. Histogram of random point cloud samples from the Huamanmarca dataset, with high completion ratios at the top to low-completion ratios at the bottom.

(a) Original MCOP images

(b) Shadow removed MCOP image

Figure A8. Results of manual shadow removal using histogram matching in the preprocessing stage.