

Article

Detection of Removed Objects in 3D Meshes Using Up-to-Date Images for Mixed Reality Applications -- Appendix

Olivier Roupin ^{1,*} , Matthieu Fradet ¹ , Caroline Baillard ¹ and Guillaume Moreau ² 

¹ InterDigital; {olivier.roupin, matthieu.fradet, caroline.baillard}@interdigital.com

² IMT Atlantique; guillaume.moreau@imt-atlantique.fr

* Correspondence: olivier.roupin@interdigital.com

Version February 1, 2021 submitted to Electronics

This appendix showcases the datasets of scenes on which geometric change detection is performed. Each scene's past and present states are respectively described by an outdated 3D mesh and a sequence of current images. All the images of a sequence are taken within a narrow time frame to reduce structural or lighting changes. For clarity, the virtual cameras chosen to render the meshes align with the estimated camera poses of the image sequences, and a 2D change ground truth is provided for each such pose. The cameras' intrinsic and extrinsic matrices are taken from the original datasets [1,2] and can be approximate in the pose or not account for the distortion of the lenses. The last row presents the results of the 3D detection of removed objects, projected in 2D for comparison with the ground truth. The removed objects added to the original datasets were downloaded from [3].


More quantitative results are provided in the shape of ROC curves (*true positive rate* against *false positive rate*) produced by our removal detection algorithm. Every blue point on the curve is generated using a particular threshold (in pixel value) for what constitutes a change. Points near the origin are for higher thresholds (less pixels are classified as changes) and the best detection results are to the top and to the left. The red cross denotes the result obtained for the automatic threshold (based on the triangle method) which is used to produce the removal detection illustrations.

1. Scenes from Palazzolo Et Al. Dataset


[2]




(a) Up-to-date images.



(b) Ground truth : insertions (green) and removals (red).



(c) Outdated 3D mesh.



(d) Projected 3D removal detection (white: false negative, green(cyan): true positive(s), red(magenta): false positive(s)).

Figure 1. container-shelf.





Figure 2. ROC container-shelf.


The removed object is very close to the edge of the screen, making it harder to detect and leading to false positives on the trees in the background, whose geometries are approximated. The large inserted object does not affect the detection of removals.




(a) Up-to-date images.




(b) Ground truth : insertions (green) and removals (red).



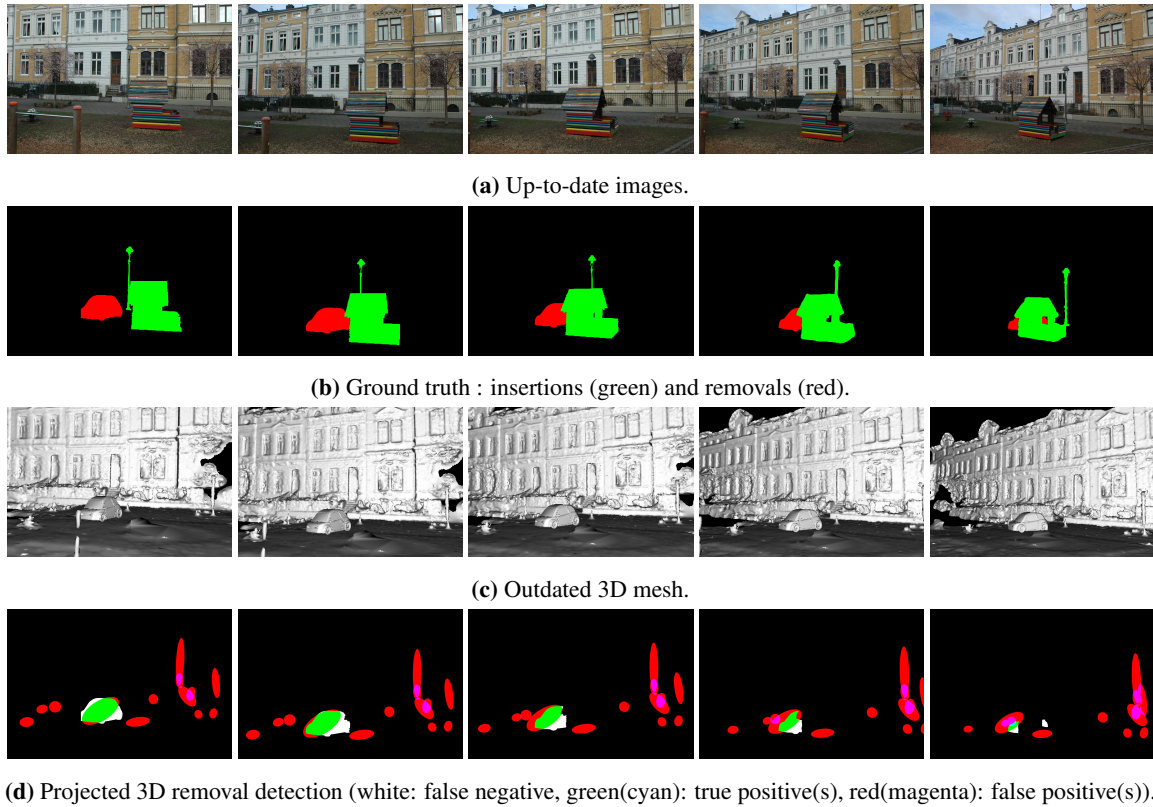

(c) Outdated 3D mesh.



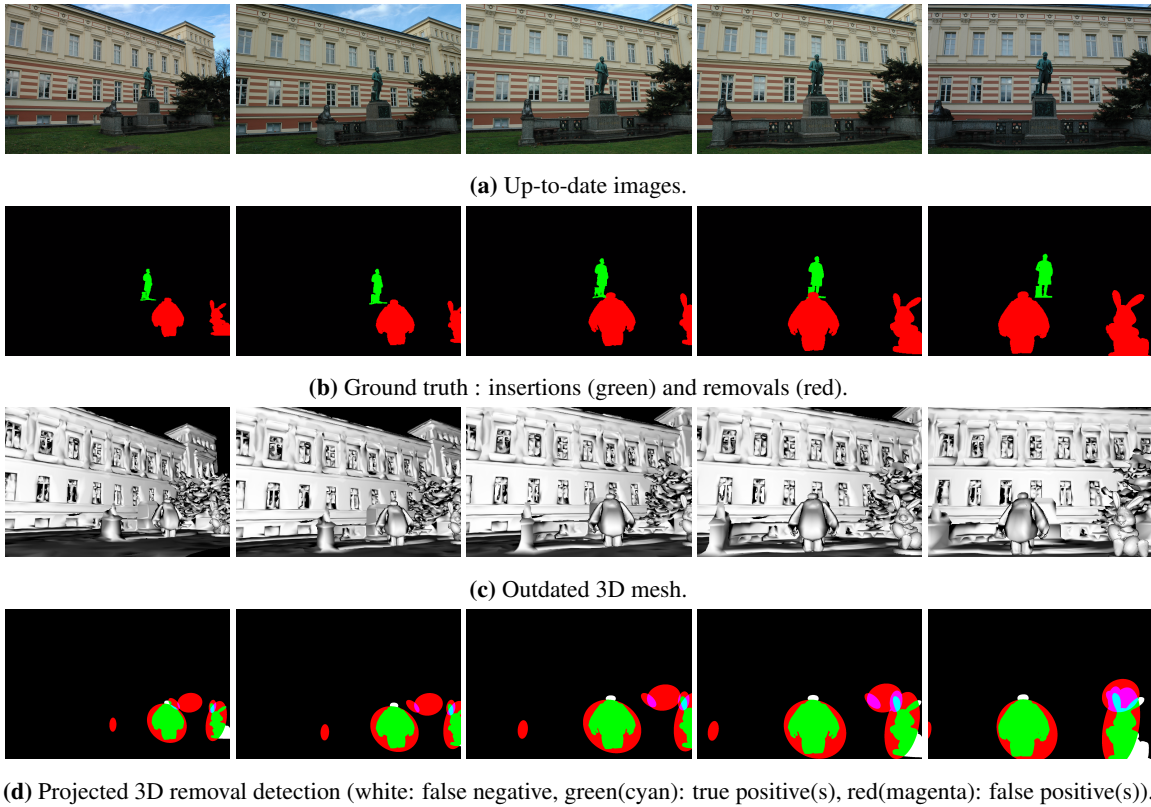

(d) Projected 3D removal detection (white: false negative, green(cyan): true positive(s), red(magenta): false positive(s)).

Figure 3. container-shelf2.**Figure 4.** ROC container-shelf2.

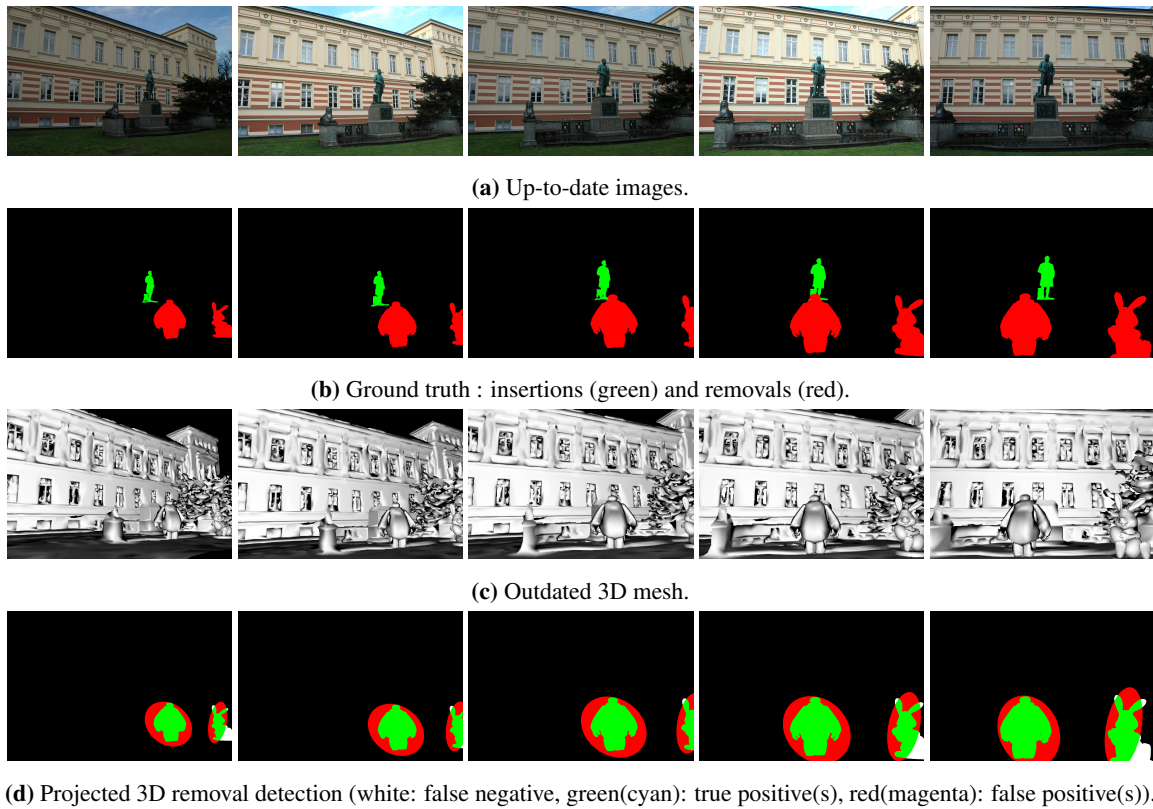

The same image sequence is used but the object is more centered in the frame, resulting in a much more accurate detection. The ROC curve is also less fragmented. The fragmentation is partly due to the grouping of changed pixels into regions, which can be discarded if their surface area falls under a certain threshold. The fragmentation occurs when there are multiple areas of detection (the trees and the object in the previous sequence).

**Figure 5.** playground-car.**Figure 6.** ROC playground-car.

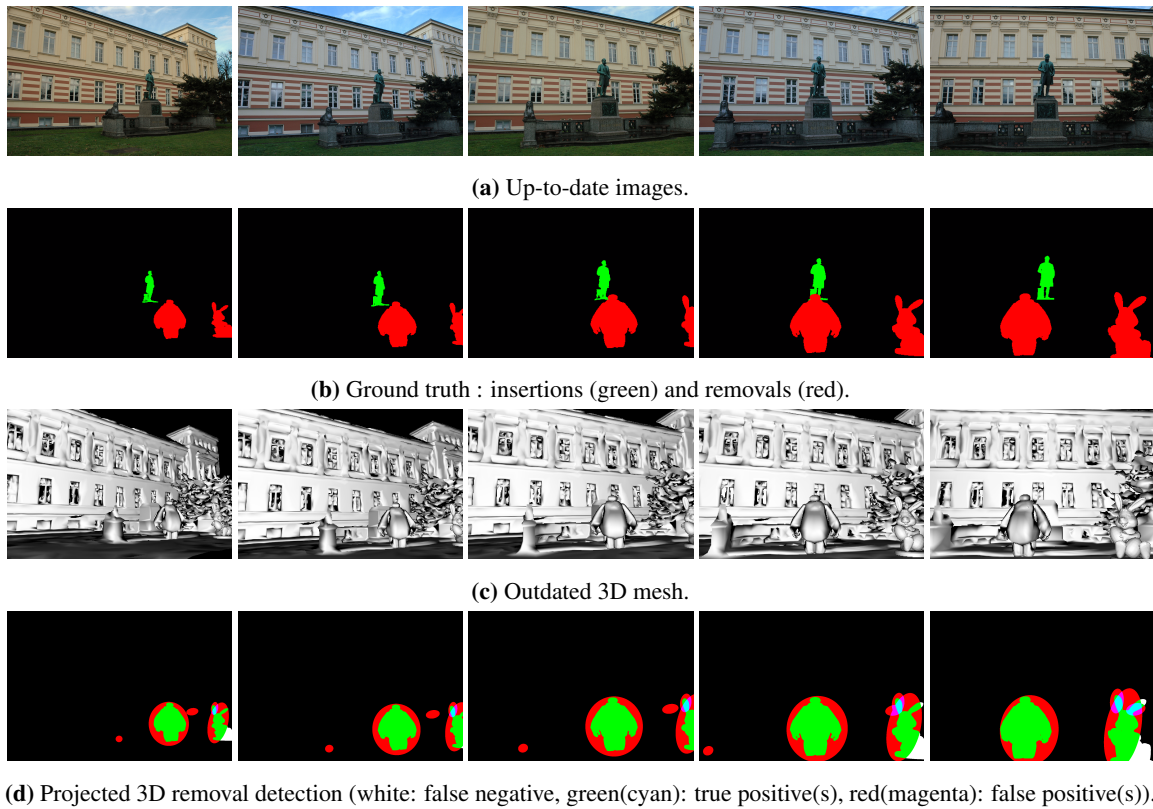

In the scene, an inserted object is in front of the removed object. As a result the object is only partially detected. The complexity of the scene also generates a lot of false positives because of the approximations in the mesh. In this particular scene, the automatic threshold is at a local minimum of performance. Also the results for the automatic threshold do not coincide with ones for any fixed threshold, as it can be different for each image of the sequence.

**Figure 7.** statue-robot.**Figure 8.** ROC statue-robot.

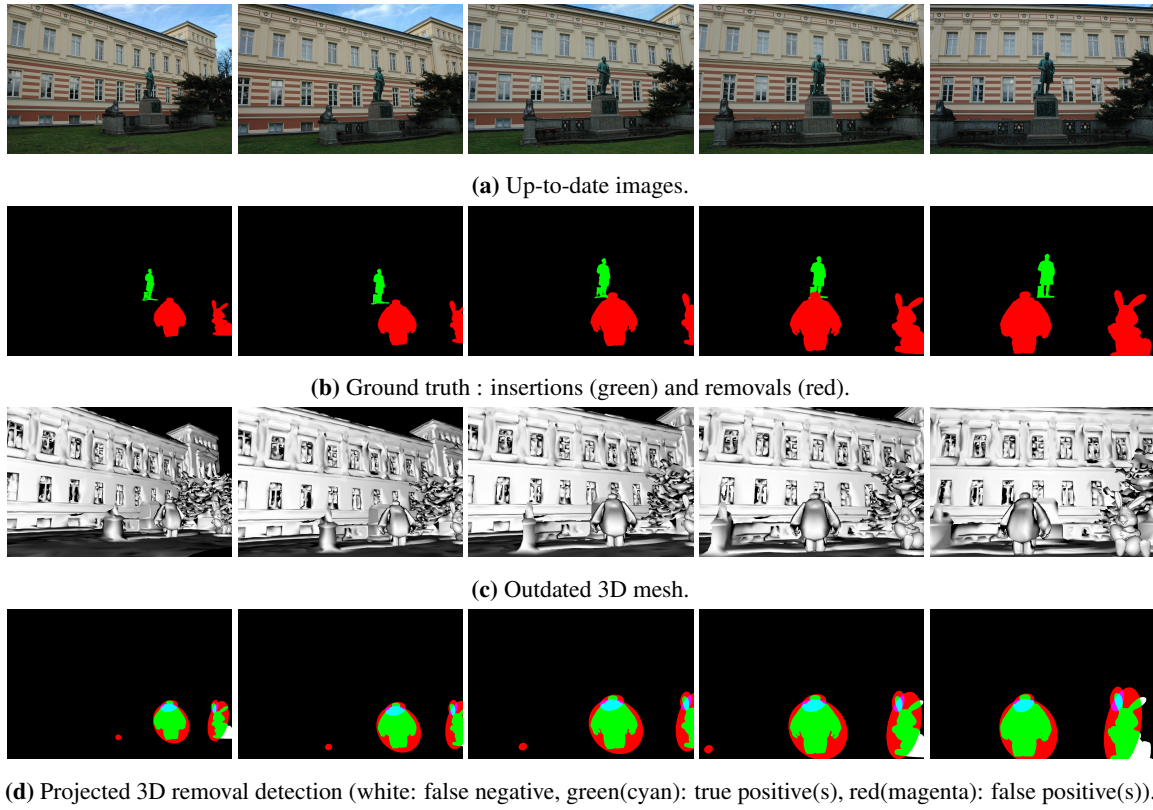

In this scene there are two disjointed removed objects. A few false positives occur on the trees on the right side of the frame and the statue on the left. The false positives can only occur on foreground objects, which cause occlusions.

**Figure 9.** statue-robot-bad-exp.**Figure 10.** ROC statue-robot-bad-exp.

This scene uses the same mesh and poses as the previous one, but the exposure in each image has been independently altered. This does not greatly affect the detection: the false positive rate is higher but it is the result of overestimating the size of the objects rather than detecting other objects altogether. The automatic threshold results are at a local maximum of performance: there is no other point to the top and left.

**Figure 11.** statue-robot-bad-temp.**Figure 12.** ROC statue-robot-bad-temp.

This scene uses the same mesh and poses as the previous one, but the color temperature in each image has been independently altered. This does not significantly impact the detection, and once again the automatic threshold results are at a local maximum.

**Figure 13.** statue-robot-temp.**Figure 14.** ROC statue-robot-temp.

This scene uses the same mesh and poses as the previous one, but the color temperature and exposure have been made more consistent across the images. This improves the results slightly.

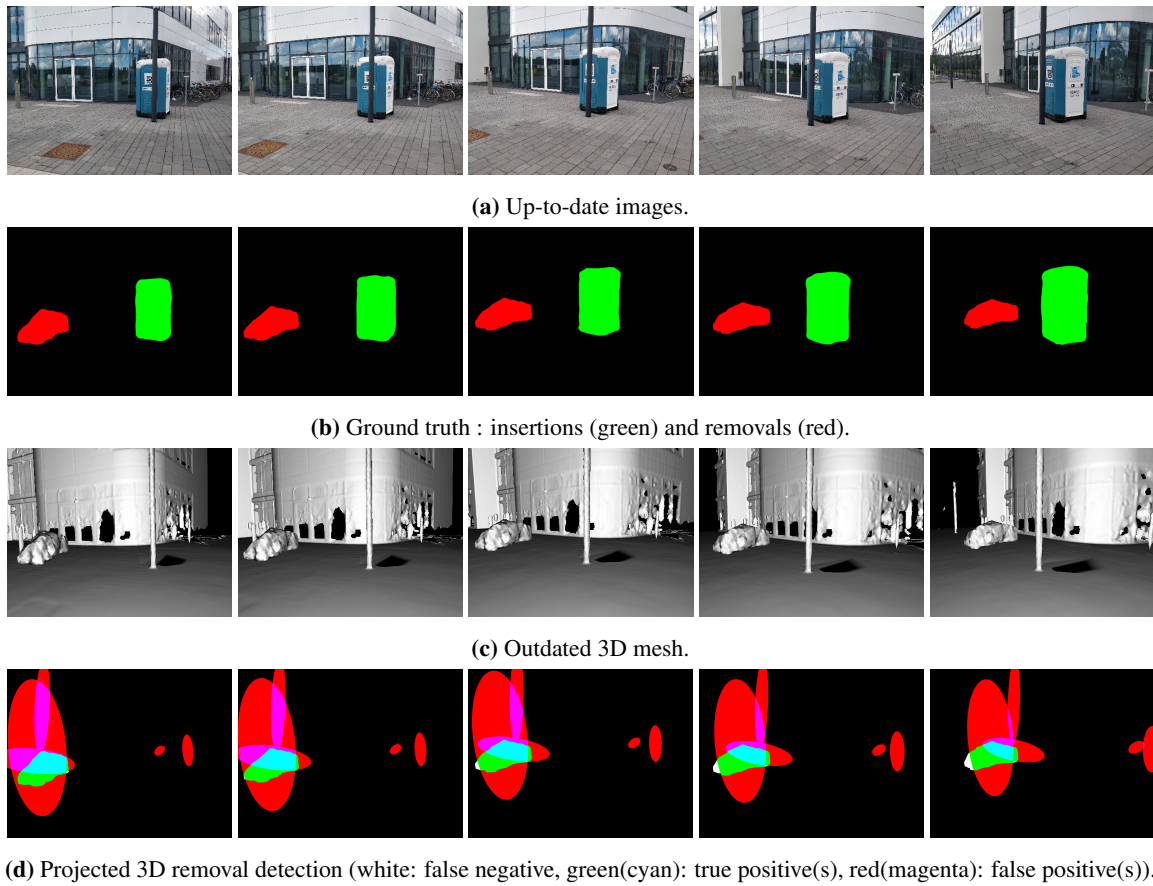


Figure 15. toilet-stone.




Figure 16. ROC toilet-stone.

43 The object is detected but it is located in front of an inaccurately modeled part of the building, which greatly
 44 increases the false positive rate. It should be noted that a lower threshold would significantly improve the result.

2. Scenes from ScanNet Dataset

[1]

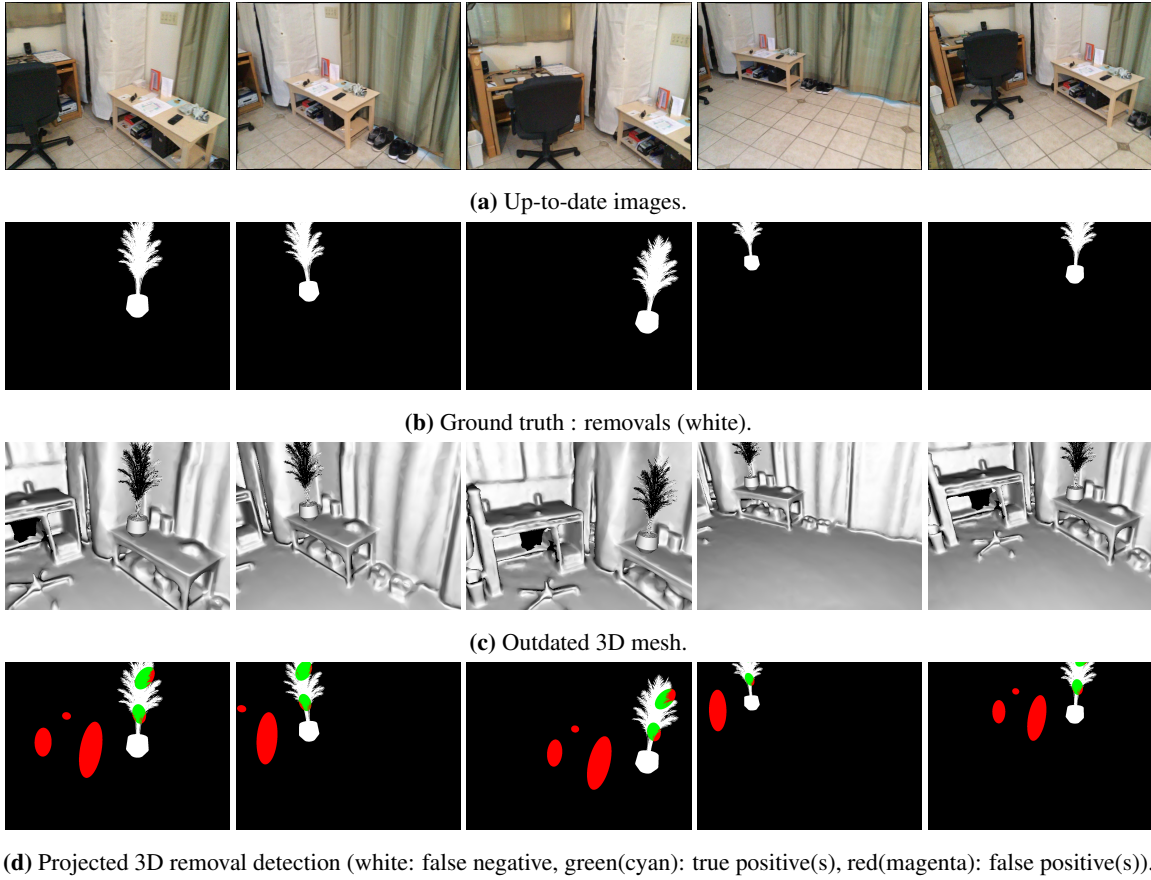


Figure 17. 0000_00+plant.


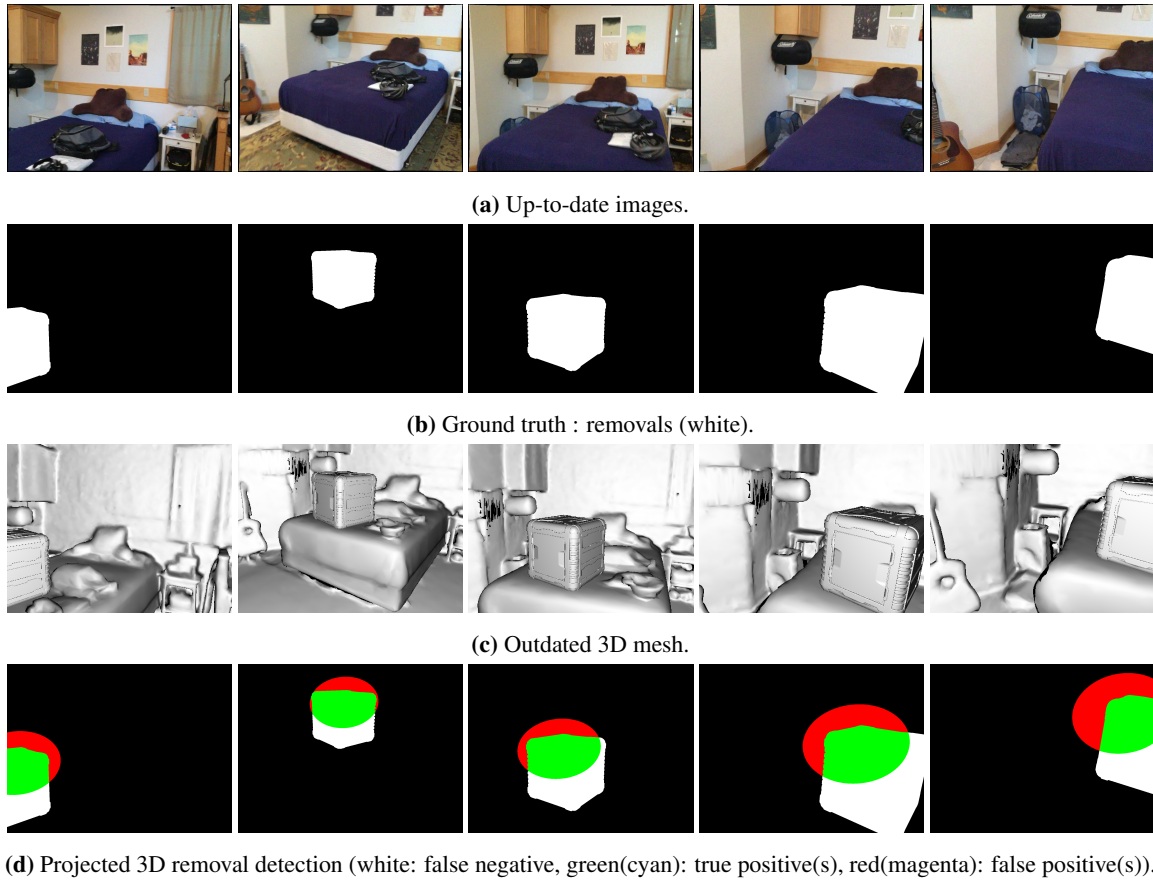

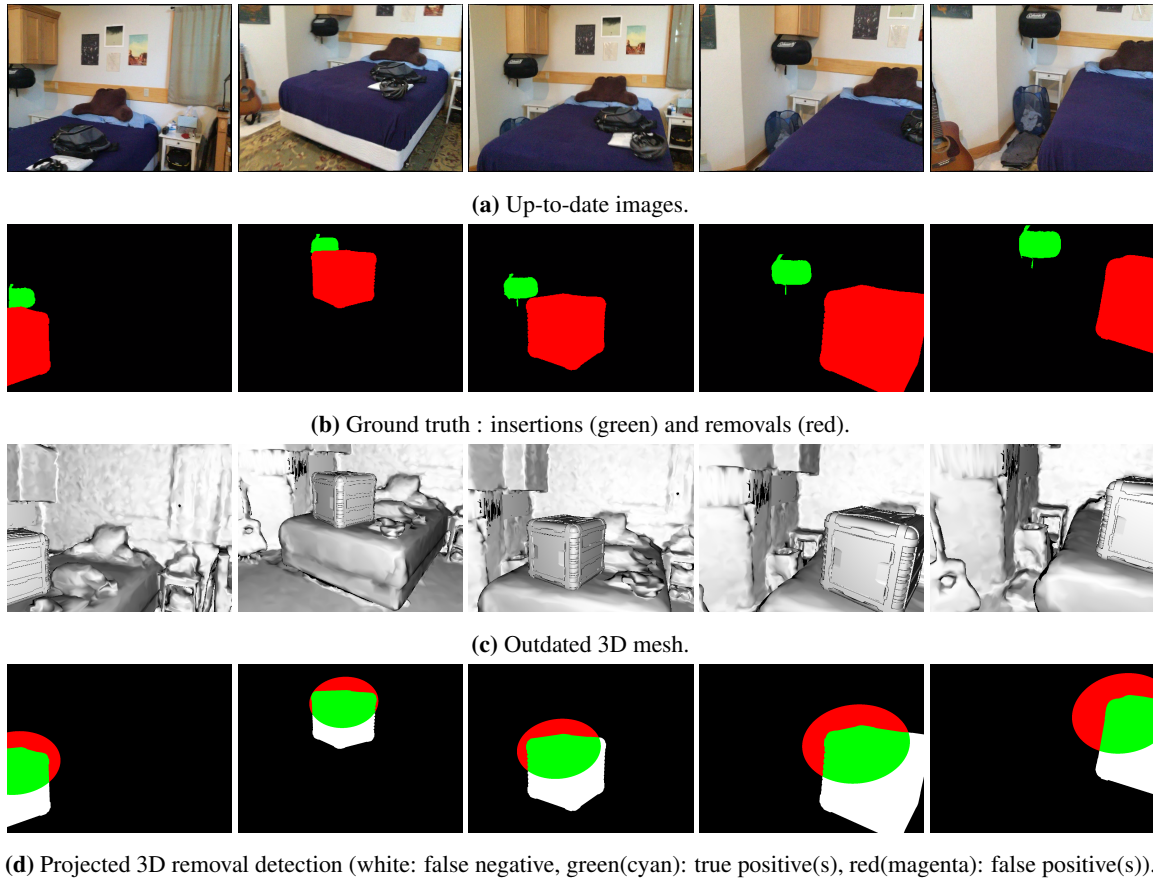



Figure 18. ROC 0000_00+plant.

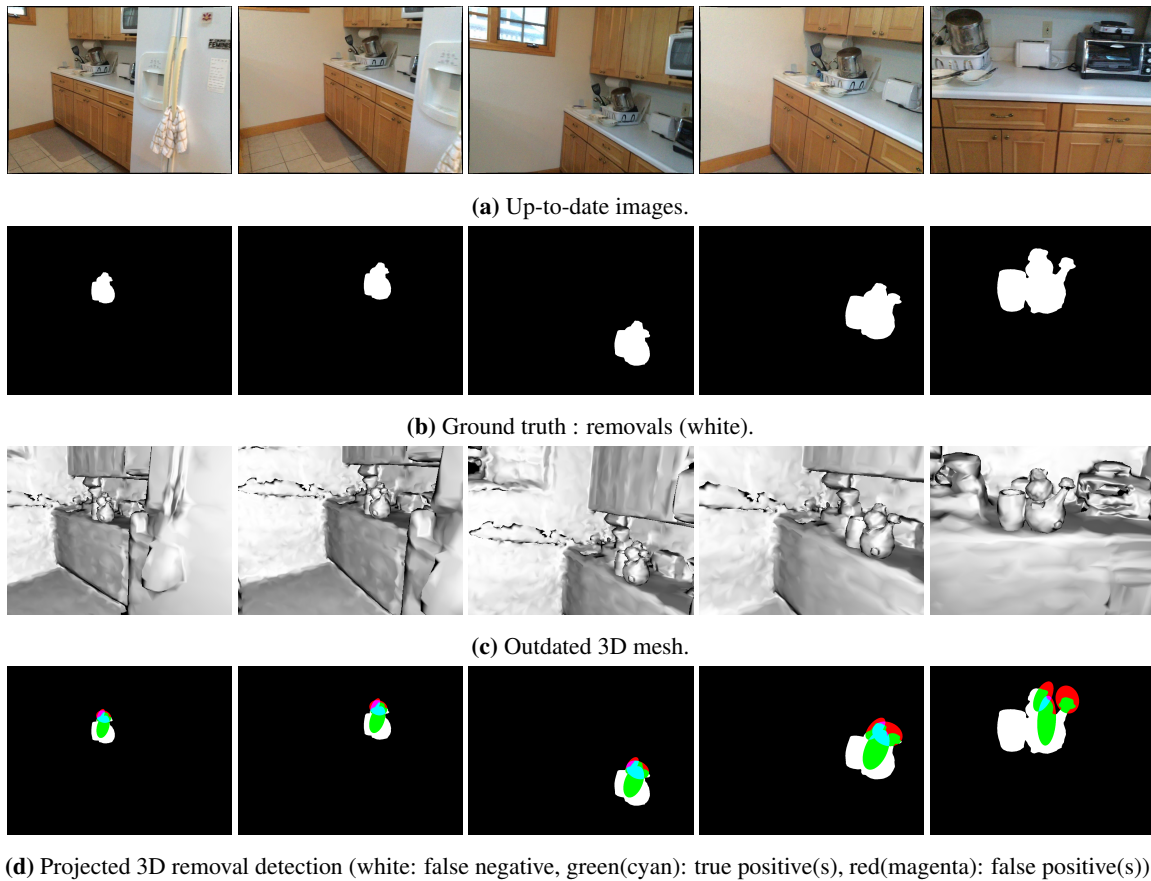

There is little parallax between the points of view as the camera is mostly rotating on it-self, and parts of the scene are inaccurately modeled. The true positive rate does not go over 0.45 signifying that some parts of the object cannot be detected.

**Figure 19.** 0000_01+box.**Figure 20.** ROC 0000_01+box.

50 The top of the object occludes the wall which is far away, while the bottom rests on top the bed. The parts
 51 of the object that are closer to the background are harder to detect.

**Figure 21.** 0000_01-insert+box.**Figure 22.** ROC 0000_01-insert+box.


The presence of an inserted object does not significantly affect the removal detection.

**Figure 23.** 0000_02+statue.**Figure 24.** ROC 0000_02+statue.


Even at the highest sensibility, some parts of the object cannot be recovered, they do not produce occlusions.




(a) Up-to-date images.




(b) Ground truth : removals (white).



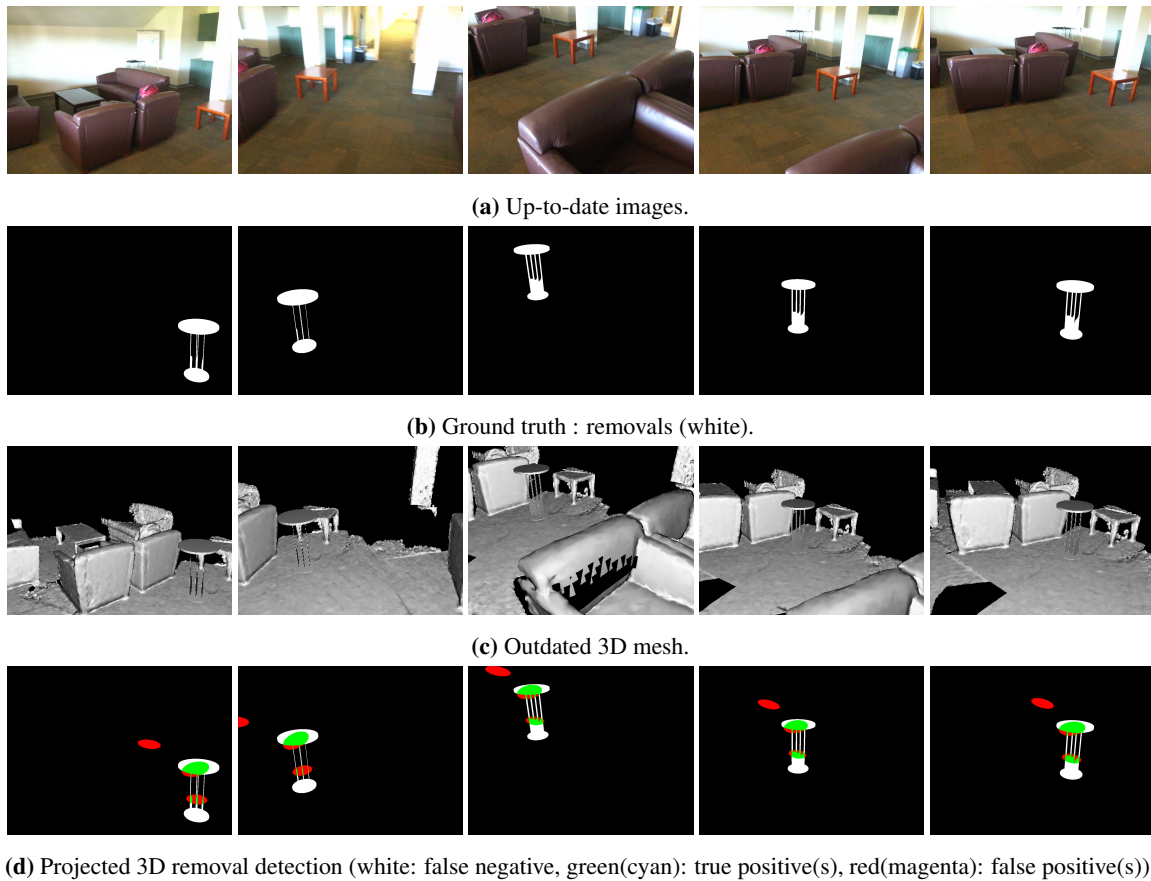

(c) Outdated 3D mesh.



(d) Projected 3D removal detection (white: false negative, green(cyan): true positive(s), red(magenta): false positive(s)).

Figure 25. 0001_00+dollhouse.**Figure 26.** ROC 0001_00+dollhouse.


54 The top of the object is the most subject to parallax and is the only portion detected. The false positive are
 55 caused by the photometric similarities between the objects of the scene and the background.

**Figure 27.** 0001_01+table.**Figure 28.** ROC 0001_01+table.

56 The bottom of the table is merged with the background and cannot be recovered. However the legs are
 57 detected despite their narrowness.




(a) Up-to-date images.




(b) Ground truth : removals (white).



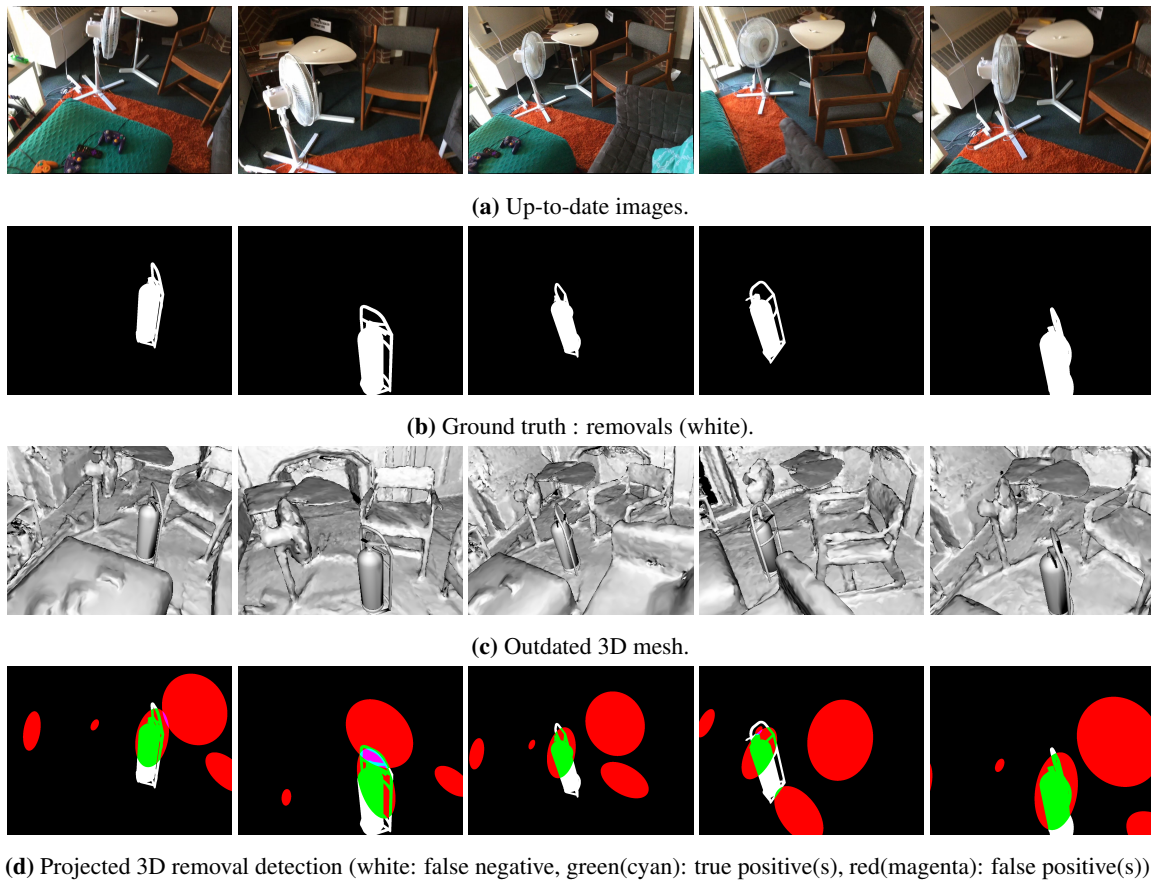

(c) Outdated 3D mesh.



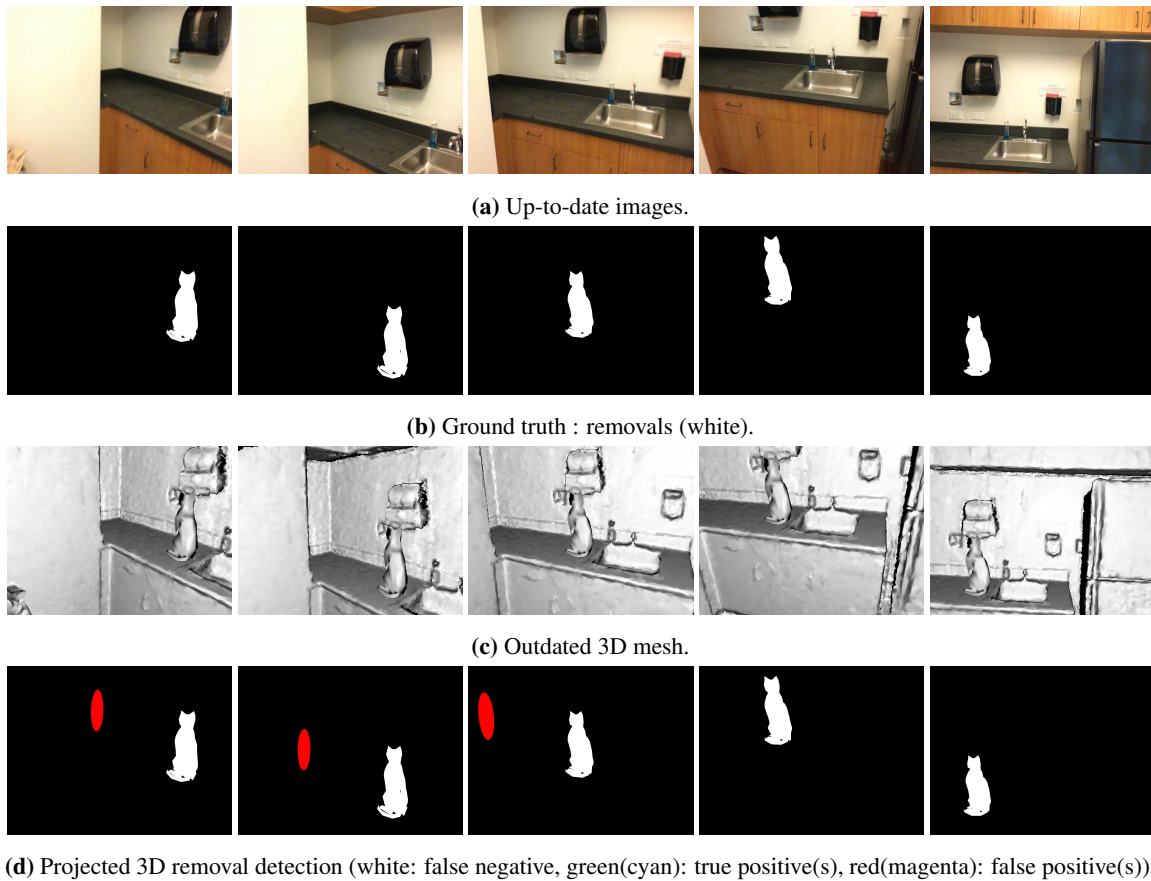

(d) Projected 3D removal detection (white: false negative, green(cyan): true positive(s), red(magenta): false positive(s)).

Figure 29. 0002_00+chair.**Figure 30.** ROC 0002_00+chair.

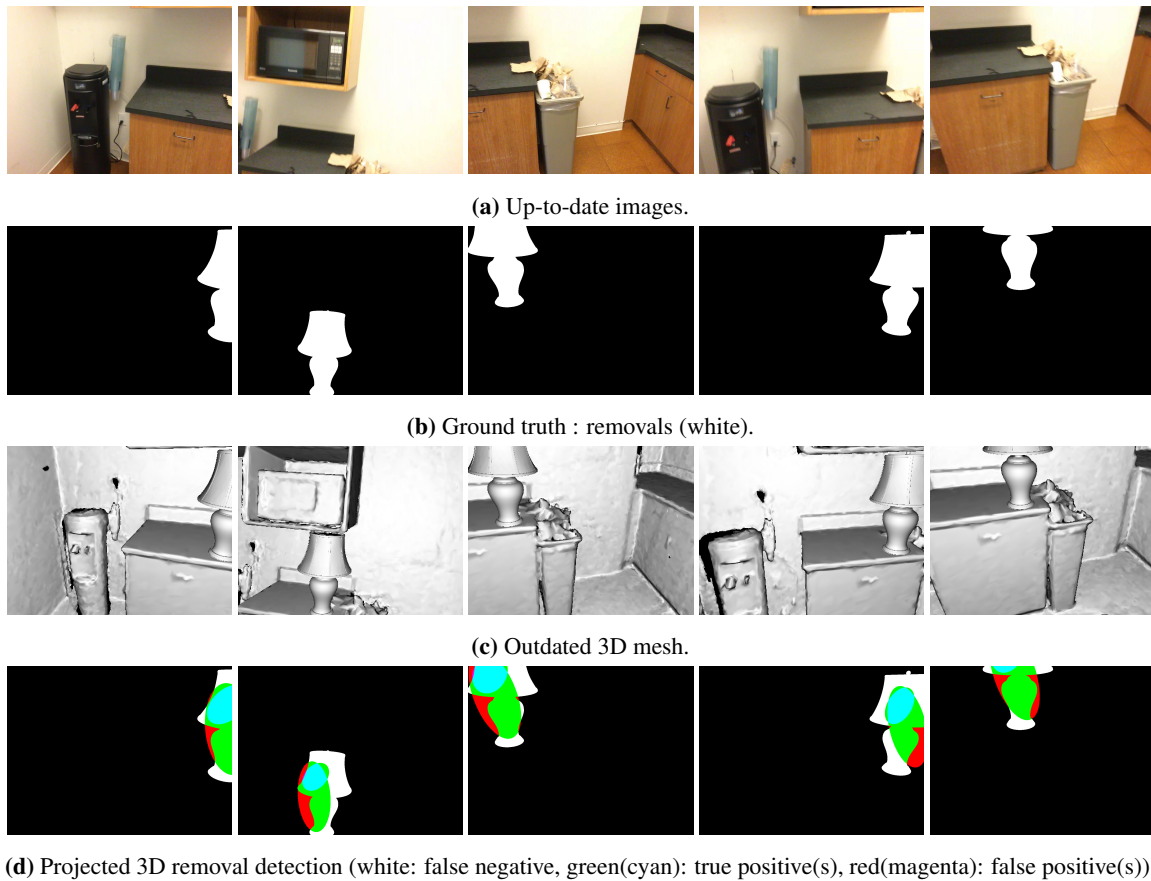

One arm of the foreground chair is wrongfully detected as removed since it does not perfectly align between the images and the mesh.

**Figure 31.** 0002_01+extinguisher.**Figure 32.** ROC 0002_01+extinguisher.

Once again, one arm of the foreground chair is wrongfully detected.

**Figure 33.** 0003_00+cat.**Figure 34.** ROC 0003_00+cat.

61 The object does not produce any occlusion because of the camera is mostly rotating on itself, creating little
 62 to no parallax.

**Figure 35.** 0003_01+desk lamp.**Figure 36.** ROC 0003_01+desk lamp.

The object is rarely seen in full, which makes evaluating its size more difficult.



(a) Up-to-date images.




(b) Ground truth : removals (white).



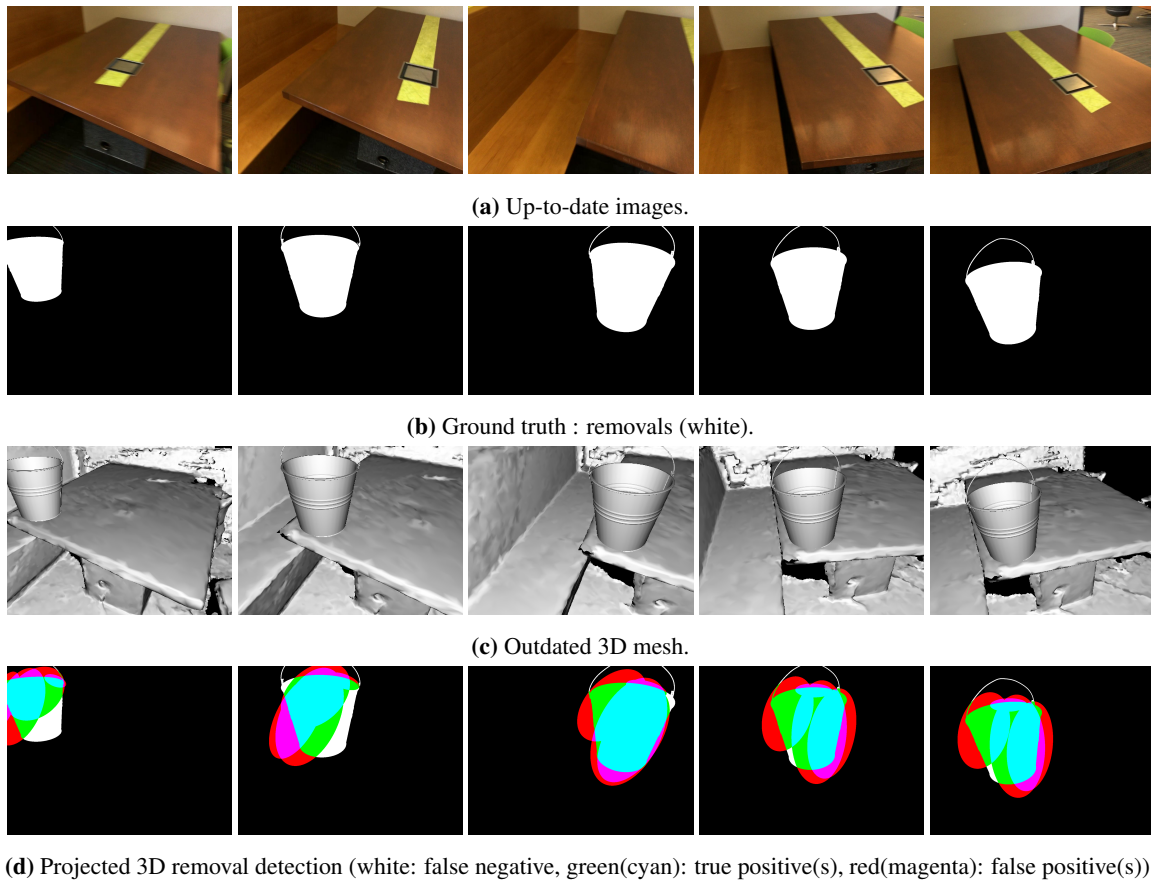

(c) Outdated 3D mesh.



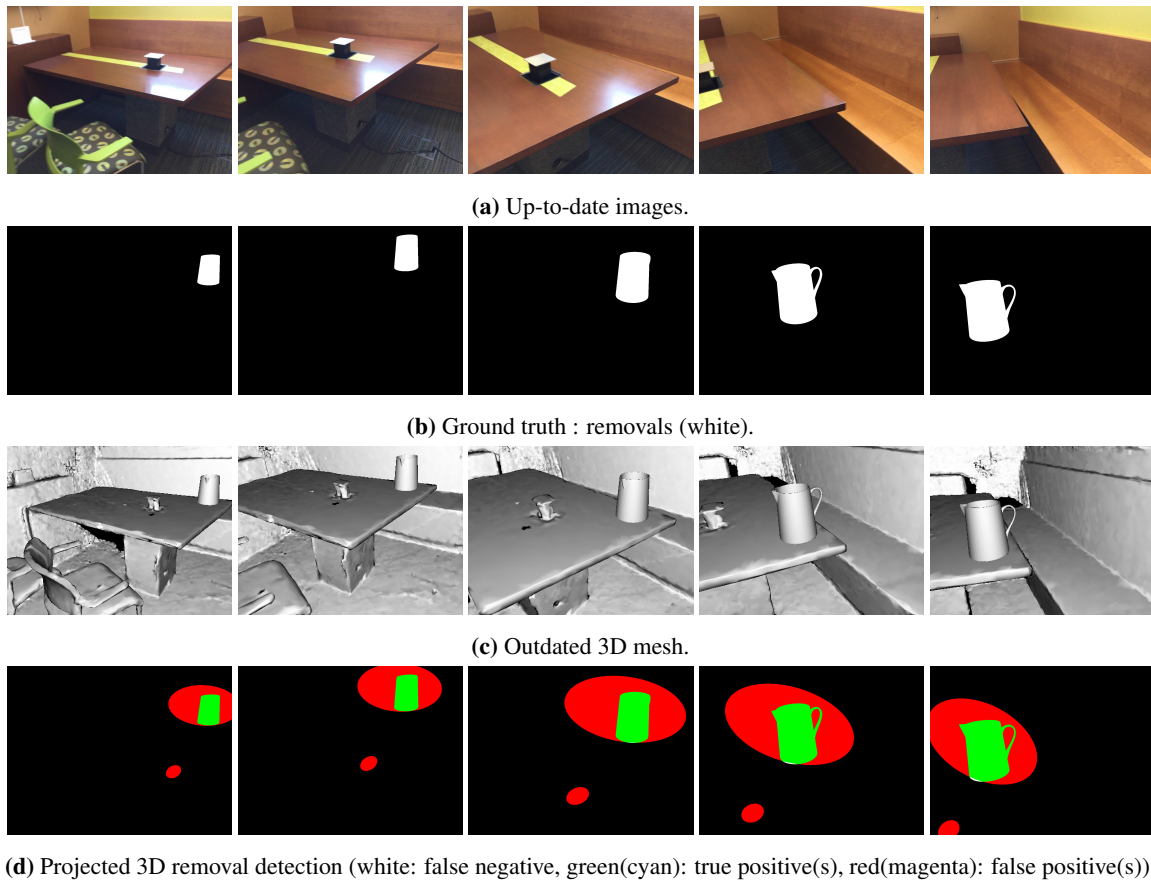

(d) Projected 3D removal detection (white: false negative, green(cyan): true positive(s), red(magenta): false positive(s)).

Figure 37. 0004_00+ghost.**Figure 38.** ROC 0004_00+ghost.

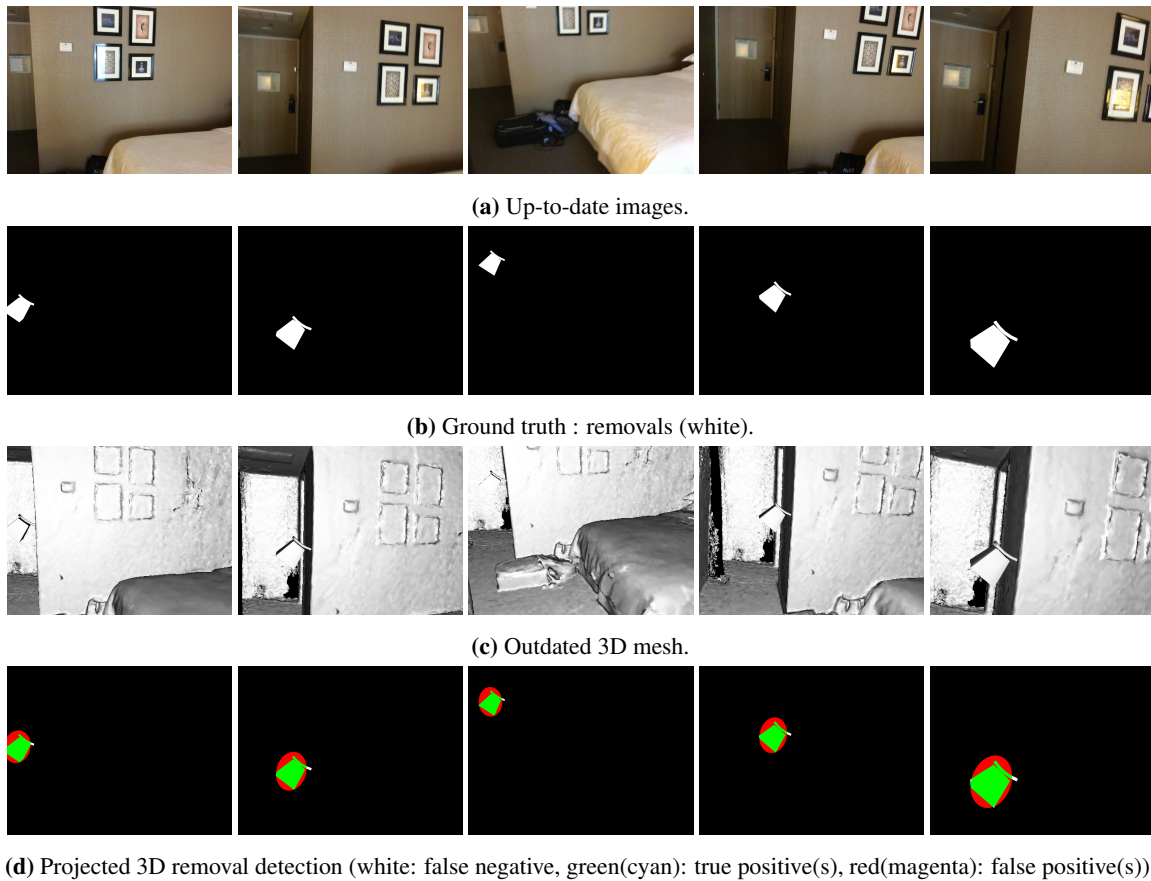

The downward angle of the scene makes the base of the object harder to detect.

**Figure 39.** 0005_00+bucket.**Figure 40.** ROC 0005_00+bucket.

The object is detected despite the incorrect camera pose estimation on the last 3 pictures.

**Figure 41.** 0005_01+pitcher.**Figure 42.** ROC 0005_01+pitcher.

66 The power socket on the surface of the table does not produce a false positive, despite its very approximate
 67 3D capture. However the size of the object is overestimated.

**Figure 43.** 0006_00+lamp.**Figure 44.** ROC 0006_00+lamp.

In this scene the geometry of the object is very simple, and the object itself is in front of a distant background. Most threshold value are able to retrieve the object in its entirety.

REFERENCES

- Dai, A.; Chang, A.X.; Savva, M.; Halber, M.; Funkhouser, T.; Nießner, M. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017.
- Palazzolo, E.; Stachniss, C. Fast Image-Based Geometric Change Detection Given a 3D Model. *Proceedings of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2018.
- Turbosquid. <https://web.archive.org/web/20201214222713/https://www.turbosquid.com/>. Accessed: 2020-12-14.

⁷⁶ **Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

⁷⁷ © 2021 by the authors. Submitted to *Electronics* for possible open access publication under the terms and conditions of the
⁷⁸ Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).