CIVIL ENGINEERING STUDIES
Illinois Center for Transportation Series No. 23-008
UILU-ENG-2023-2008
ISSN: 0197-9191

I-RIPRAP 3D Image Analysis Software: User Manual

Prepared By
Haohang Huang, PhD
Jiayi Luo
Kelin Ding
Erol Tutumluer, PhD
John M. Hart
Issam Qamhia, PhD
University of Illinois Urbana-Champaign

Research Report No. FHWA-ICT-23-007

A white paper on the deliverable software of

ICT PROJECT R27-214
3D Image Analysis Using Deep Learning for Size and Shape Characterization of Stockpile Riprap Aggregates—Phase 2

https://doi.org/10.36501/0197-9191/23-008

Illinois Center for Transportation

June 2023

ACKNOWLEDGMENT, DISCLAIMER, MANUFACTURERS' NAMES

The I-RIPRAP 3D image analysis software is a deliverable based on the findings of **ICT-R27-214: 3D Image Analysis Using Deep Learning for Size and Shape Characterization of Stockpile Riprap Aggregates—Phase 2**. ICT-R27-214 was conducted in cooperation with the Illinois Center for Transportation; the Illinois Department of Transportation; and the U.S. Department of Transportation, Federal Highway Administration.

Members of the Technical Review Panel (TRP) were the following:

- Andrew Stolba, TRP Chair, Illinois Department of Transportation
- Sheila Beshears, Riverstone Group, Inc.; Former TRP Chair with Illinois Department of Transportation
- Chad Nelson, Illinois Department of Transportation—District 3
- Del Reeves, Illinois Department of Transportation Bureau of Materials
- Kevin Tressel, Illinois Department of Transportation—District 2
- Megan Swanson, Illinois Department of Transportation
- Ryan Culton, formerly with Illinois Department of Transportation
- Sean Stutler, Illinois Department of Transportation
- Dennis Bachman, Federal Highway Administration
- Andrew Buck, Vulcan Materials Company
- Dan Barnstable, Vulcan Materials Company
- Jim Sergent, formerly with Lehigh Hanson

The authors would like to extend their appreciation to the current and past TRP Chairs and all TRP members for their time and significant contributions to this research project. The help of Andrew Stolba, Sheila Beshears, Chad Nelson, Andrew Buck, and Dan Barnstable in coordinating quarry field visits to sample and image individual riprap rocks and stockpiles is highly appreciated.

The authors would like to express their sincere gratitude to the University of Illinois Urbana-Champaign (UIUC) and Illinois Center for Transportation (ICT) students who helped on several tasks dealing with materials characterization, testing, image acquisition, and manual labeling of rocks in captured stockpile image datasets, namely: Zixu Zhao, Han Wang, Jiawei Fan, and Zhongyi Liu. Sincere acknowledgements are due to Prof. Maziar Moaveni (Savannah State University, formerly with University of Illinois) and Huseyin Boler for their insightful advice to the project.

The authors would like to acknowledge the technical support received from the Computer Vision and Robotics Laboratory in the Beckman Institute for Advanced Science and Technology at UIUC. Sincere acknowledgments are also due to ICT research engineer Greg Renshaw for his support throughout the project.

Trademark or manufacturers' names appear in this report only because they are considered essential to the object of this document and do not constitute an endorsement of product by the Federal Highway Administration, the Illinois Department of Transportation, or the Illinois Center for Transportation.

ABOUT THE I-RIPRAP 3D SOFTWARE

Riprap rock and aggregates are commonly used in various engineering applications such as structural, transportation, geotechnical, and hydraulic engineering. To ensure the quality of the aggregate materials selected for these applications, it is important to determine their morphological properties such as size and shape. There have been many imaging approaches developed to characterize the size and shape of individual aggregates, but obtaining 3D characterization of aggregates in stockpiles at production or construction sites can be a challenging task.

This research study introduces a new approach based on deep learning techniques that combines three developed research components: field 3D reconstruction procedures, 3D stockpiles instance segmentation, and 3D shape completion. The approach is designed to reconstruct aggregate stockpiles from multiple images, segment the stockpile into individual instances, and predict the unseen sides of each instance (particle) based on the partially visible shapes. The approach was validated using ground-truth measurements and demonstrated satisfactory algorithm performance in capturing and predicting the unseen sides of aggregates.

For better user experience, the integrated approach has been implemented into a software application named "I-RIPRAP 3D," with a user-friendly graphical user interface (GUI). This stockpile aggregate analysis approach is envisioned to provide efficient field evaluation of aggregate stockpiles by offering convenient and reliable solutions for on-site quality assurance and quality control tasks of riprap rock and aggregate stockpiles. The following document provides information for users of the I-RIPRAP 3D software to make the best use of the software's capabilities.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1: PREREQUISITES FOR I-RIPRAP 3D IMAGE ANALYSIS

SYSTEM REQUIREMENTS

Minimum System Requirements

Most state-of-the-art 3D imaging algorithms and applications have been built and distributed on a Linux Operating System (OS). The I-RIPRAP 3D software program was developed and tested in the Linux operating environment. The GPU requirement was designed for smoothly running 3D point cloud reconstruction. The minimum system requirements are given in Table 1.

Table 1. Minimum System Requirements

Operating System	Linux
CPU	4–8 core Intel or AMD processor, 2.0+ GHz
RAM	16–32 GB
GPU	NVIDIA or AMD GPU with 700+ CUDA cores/shader processor units (For example: GeForce GTX 1080 or Radeon RX 5700)

Advanced System Requirements

Using systems with more advanced capabilities (see Table 2) can benefit I-RIPRAP 3D in two aspects: (a) accelerate reconstruction by utilizing multi-threads and (b) obtain shorter run times using large-scale and high-quality reconstructed 3D point clouds.

Table 2. Advanced System Requirements

Operating System	Linux
CPU	6–24 core Intel or AMD processor, 3.0+ GHz
RAM	32–128 GB
GPU	1-2 NVIDIA or AMD GPU with 1920+ CUDA cores/shader processor units (For example: GeForce GTX 2080 Ti or Radeon VII)

Software Installation

UIUC’s team will pre-install the software on a high-end desktop computer with GPU.

CALIBRATION OBJECT PREPARATION

Calibration objects, which should be easily identifiable, are used during image acquisition and marker labeling. Based on the UIUC team's practice, it is recommended to use colorful wooden cubes as calibration objects because (a) color helps associate the cubes with marker numbers (e.g., a blue cube is marker #1), (b) wooden cubes have sufficient weight to remain in place in windy weather conditions, and (c) cubes in colors distinct from the background make the marker labeling process easier.

The set of wooden cubes can be purchased from different stores. For each block, use a black marker pen to mark the diagonal lines and the center point on the top surface (see Figure 1). One step in running the I-RIPRAP software entails a marker labeling procedure where the center point needs to be identified from image views. The diagonal lines can assist with locating it. Three calibration objects are needed during field image collection.

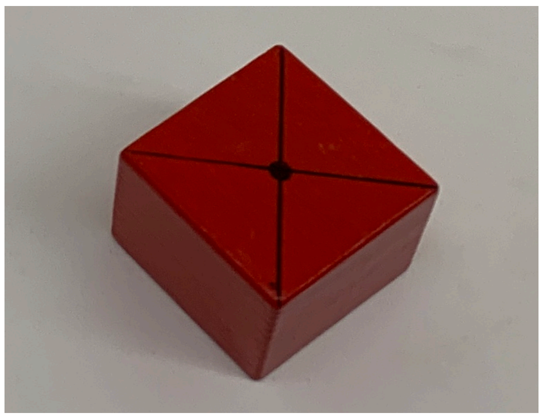

Figure 1. Calibration object.

CHAPTER 2: FIELD DATA ACQUISITION PROCEDURE

1. Identify the regions of interest (ROI) in the stockpile. The stockpiles at a quarry are typically large, so inspection and imaging should target smaller regions or parts of the stockpile to ensure sufficient reconstruction quality. Alternatively, smaller stockpiles can be formed on-site using a front loader or other means (see Figure 2).

2. Place the color-coded calibration objects around the stockpile and measure the center-to-center distance between two pairs (e.g., blue to red and blue to yellow). Remember to record these measurements for each inspected stockpile.

Figure 2. Mini stockpile with calibration objects.

3. Collect either approximately 40 to 50 multi-view images or approximately 60-second all-around videos of each stockpile, with most views covering the calibration objects.

    a. It is not strictly required that all three calibration objects be present in all images. In fact, it is neither possible nor recommended because of occlusion in certain views. Nevertheless, ensure the objects are mostly in sight without moving too close to the stockpile surface.

    b. It is recommended to vary the camera height, up and down, periodically while moving around the stockpile. This will help collect more visible sides of the aggregates from different side views.

Note: The UIUC team used smartphone cameras. Modern smartphone cameras with a minimum photo resolution of 8 MP are acceptable. For video, it is recommended to use video settings of 4K at 24 fps.

CHAPTER 3: IMAGE ANALYSIS WORKFLOW USING I-RIPRAP 3D

OPENING I-RIPRAP 3D SOFTWARE

1. Double click the I-RIPRAP 3D icon on the desktop (Figure 3) to open the software.

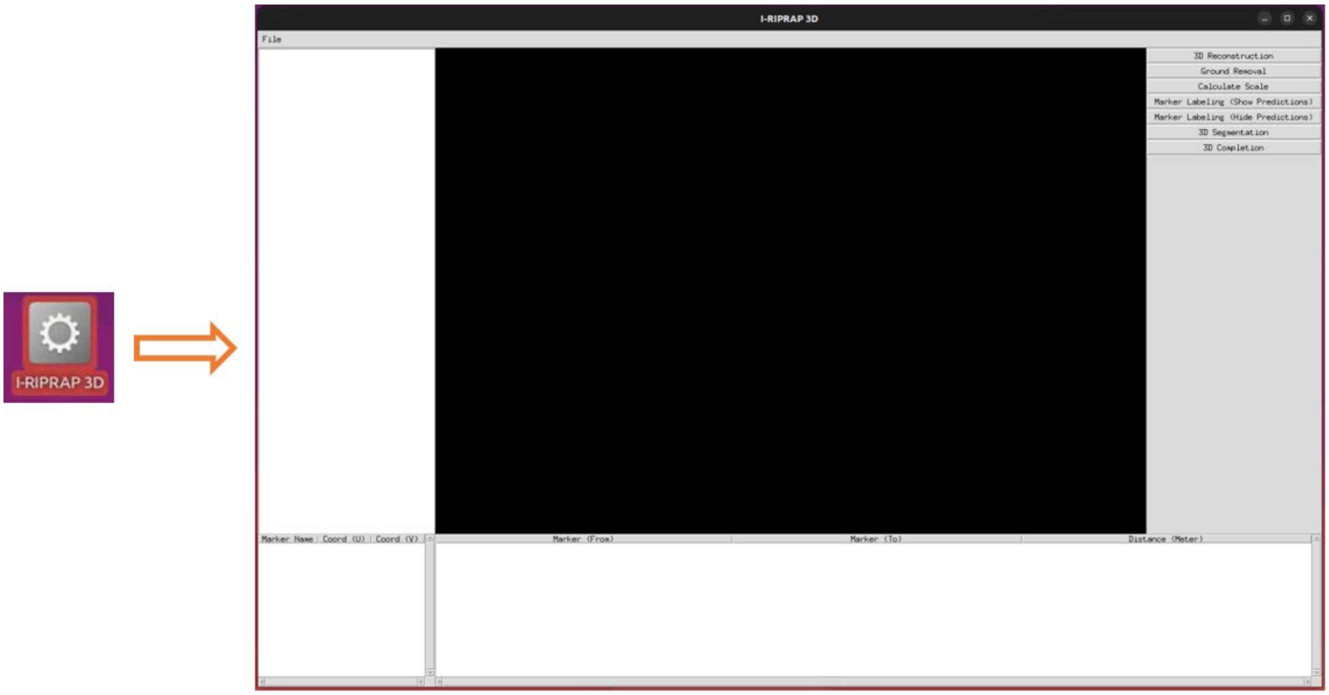

Figure 3. I-RIPRAP 3D software—welcome screen.

3D RECONSTRUCTION

Generating the 3D Point Cloud

Generating the 3D point cloud from multi-view images is a computer vision task in the areas of structure-from-motion (SfM) and multi-view stereo (MVS). To obtain a high-quality 3D reconstruction, an advanced general-purpose SfM and MVS pipeline, called COLMAP (Schonberger, 2016), is integrated in the software. Follow the steps listed below:

1. Click 3D Reconstruction to initiate the 3D reconstruction step. After clicking the button, the built-in COLMAP reconstruction window will pop up (Figure 4).

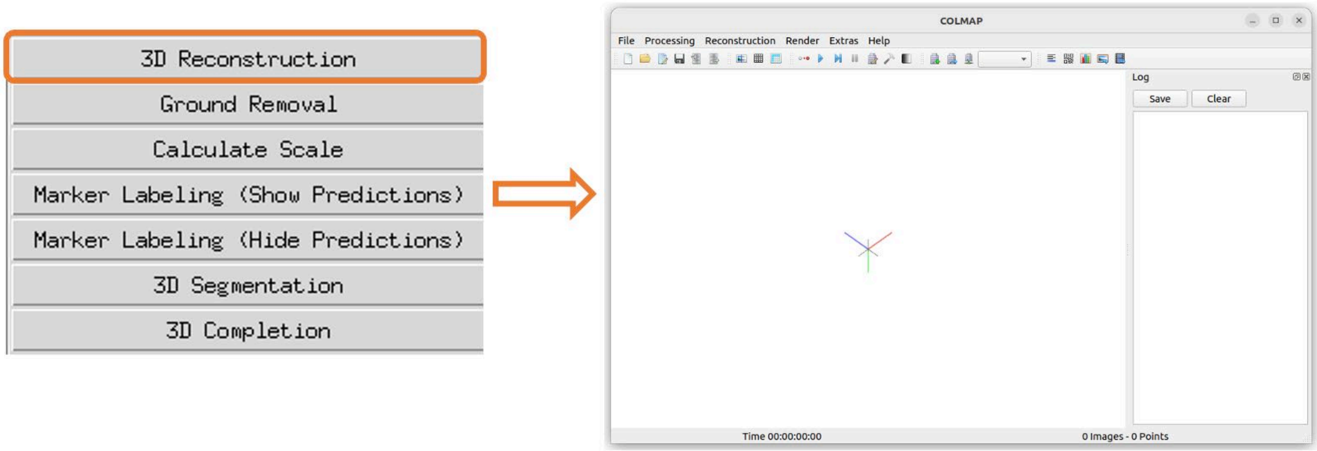

Figure 4. Open COLMAP for 3D reconstruction.

2. Click Reconstruction --> Automatic reconstruction, and set the parameters as shown in Figure 5.

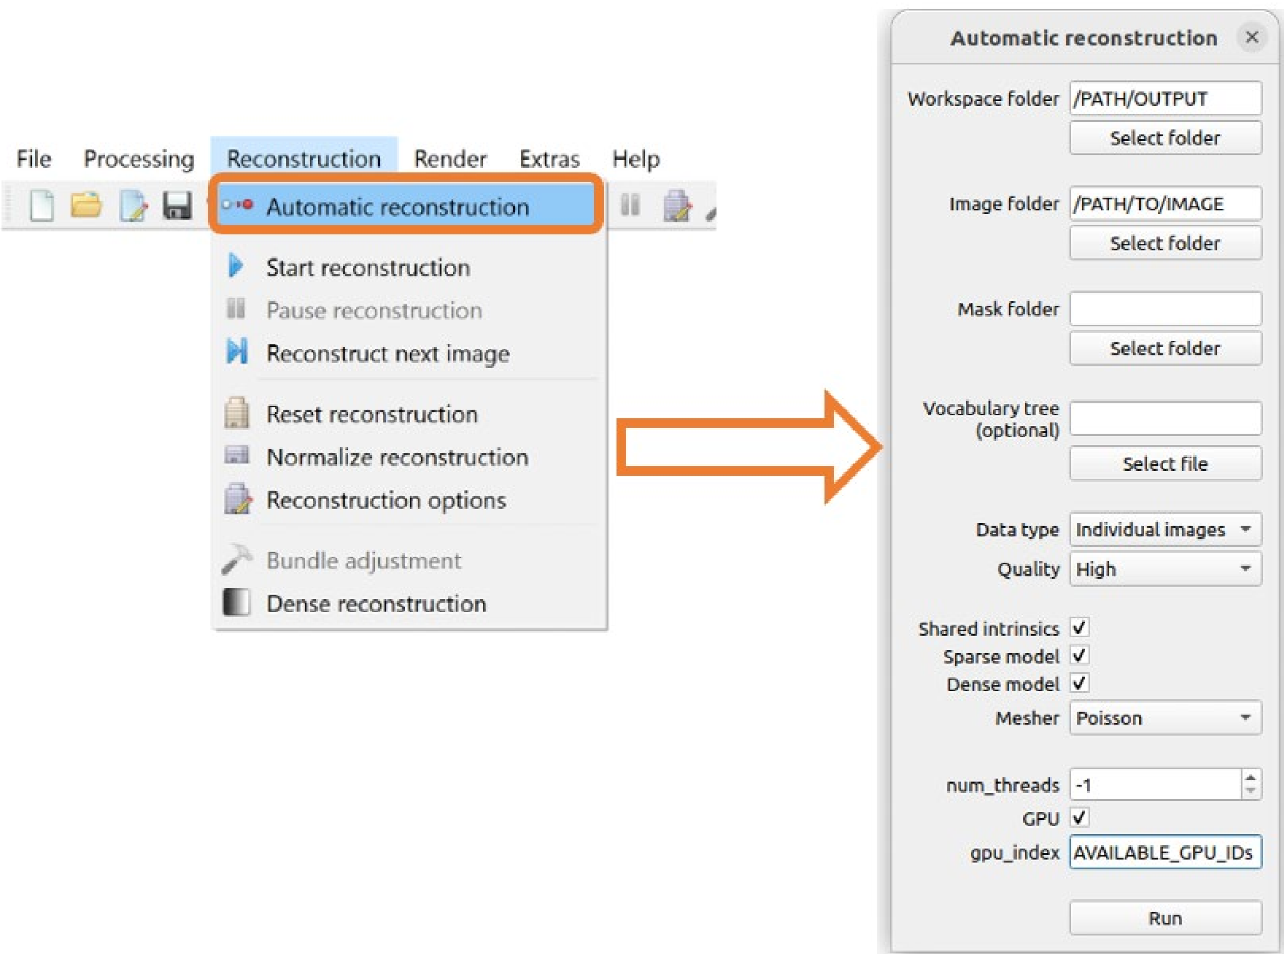

Figure 5. Initiating 3D reconstruction in COLMAP.

Notes:

- Workspace folder will contain all output files and should be empty before processing.
- Image folder should contain all images for reconstructing the 3D point cloud.
- The recommended Quality level is High.
- The gpu_index default value is −1.

3. Click Run, and the algorithm will conduct the 3D reconstruction. Log information will pop up in the Log sub-window (Figure 6).

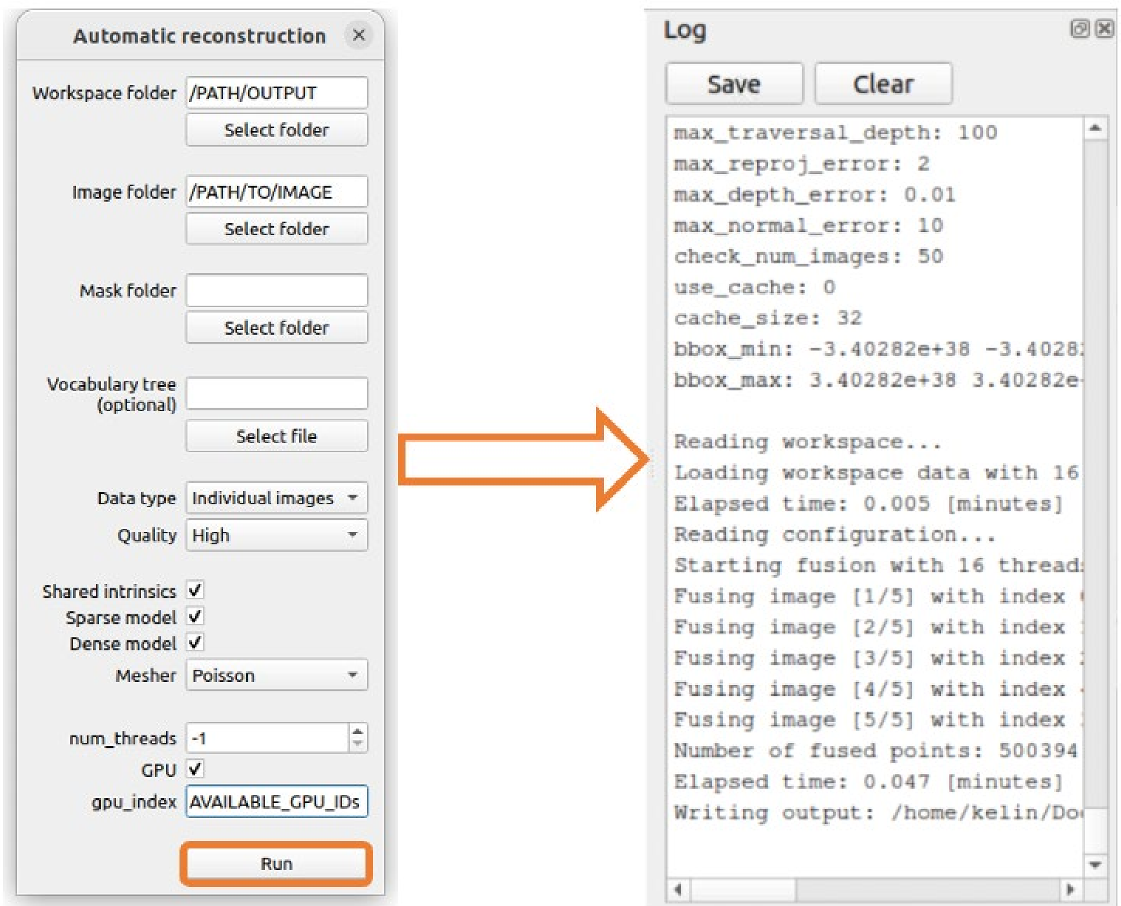

Figure 6. Running 3D reconstruction.

Note: This 3D reconstruction step will take approximately 10 minutes to 2 hours, depending on the system configuration and the number of images.

4. Close the COLMAP reconstruction window.

Cleaning the 3D Point Cloud

The reconstructed point cloud includes unwanted objects, for example, the ground. They need to be manually removed by professional 3D tools, like MeshLab (Cignoni, 2008).

1. Click Ground Removal to open MeshLab. After clicking the button, the MeshLab window will pop up (Figure 7).

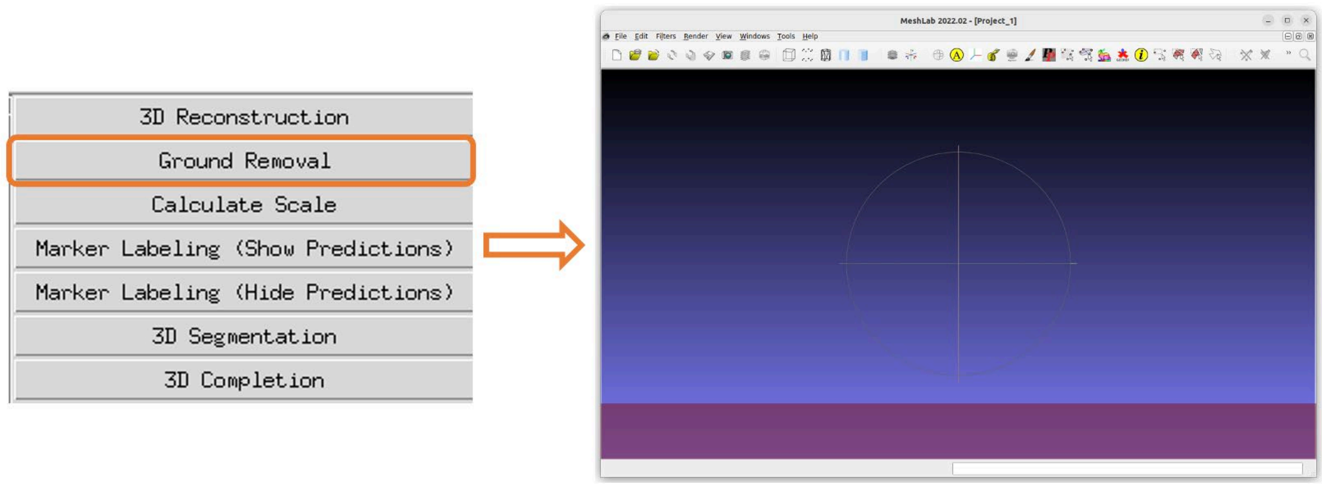

Figure 7. Opening MeshLab for ground removal.

2. Click Import Mesh (Figure 8) to load the reconstructed point cloud under path **/PATH/OUTPUT/dense/0/fused.ply**

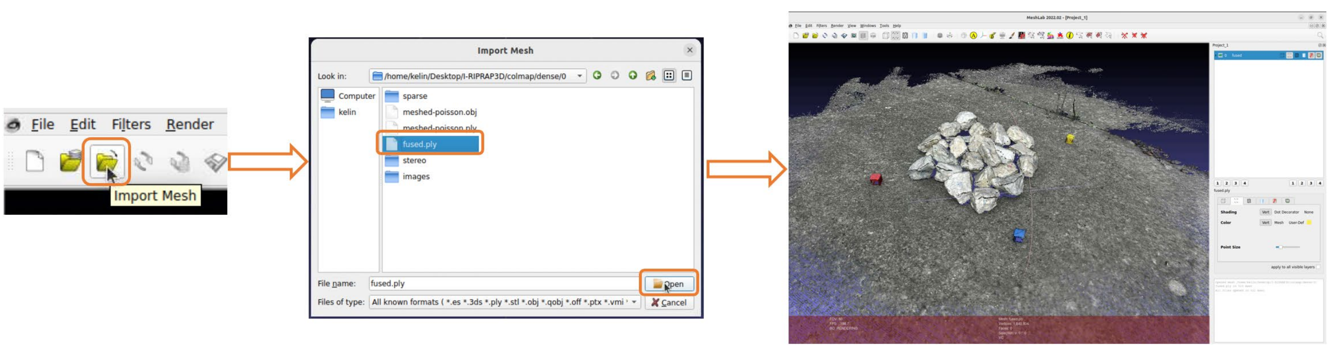

Figure 8. Importing reconstructed mesh in MeshLab.

3. Rotate the point cloud to make the view direction parallel with the plane of the ground (see Figure 9).

4. Click Select Vertices and then drag the selection area to select the unrelated vertices. (The selected vertices will turn red, as shown in Figure 9.)

Figure 9. Selecting unrelated vertices.

5. Rotate the viewport to make sure no related vertices are selected (Figure 10).

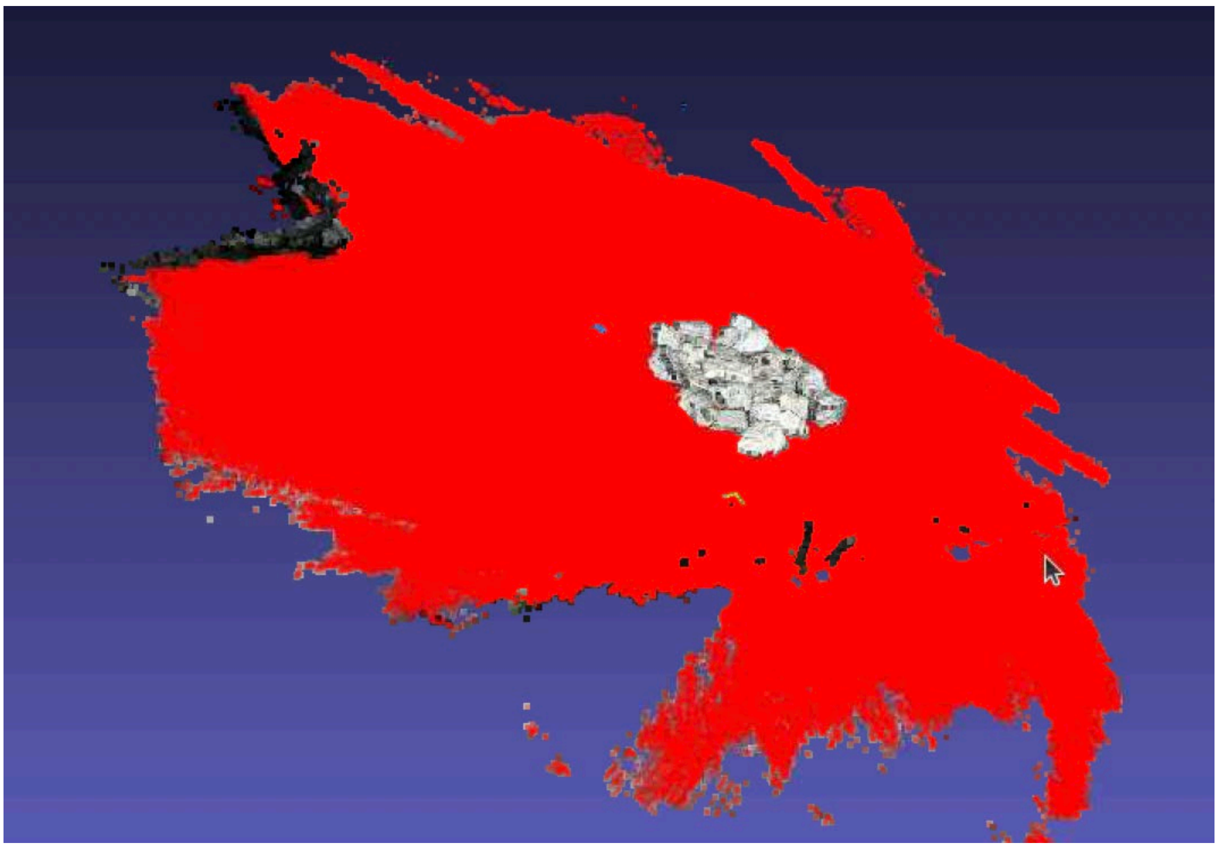

Figure 10. Checking selected vertices in a rotated view.

6. Click Delete Selected Vertices to remove the selected unrelated vertices (Figure 11).

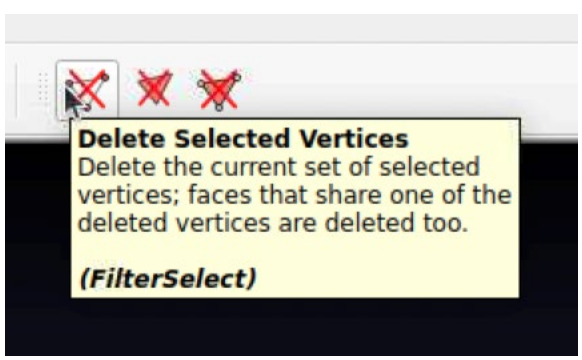

Figure 11. Deleting selected unrelated vertices.

7. Repeat Steps 4 to 6 above until all unrelated vertices are removed.

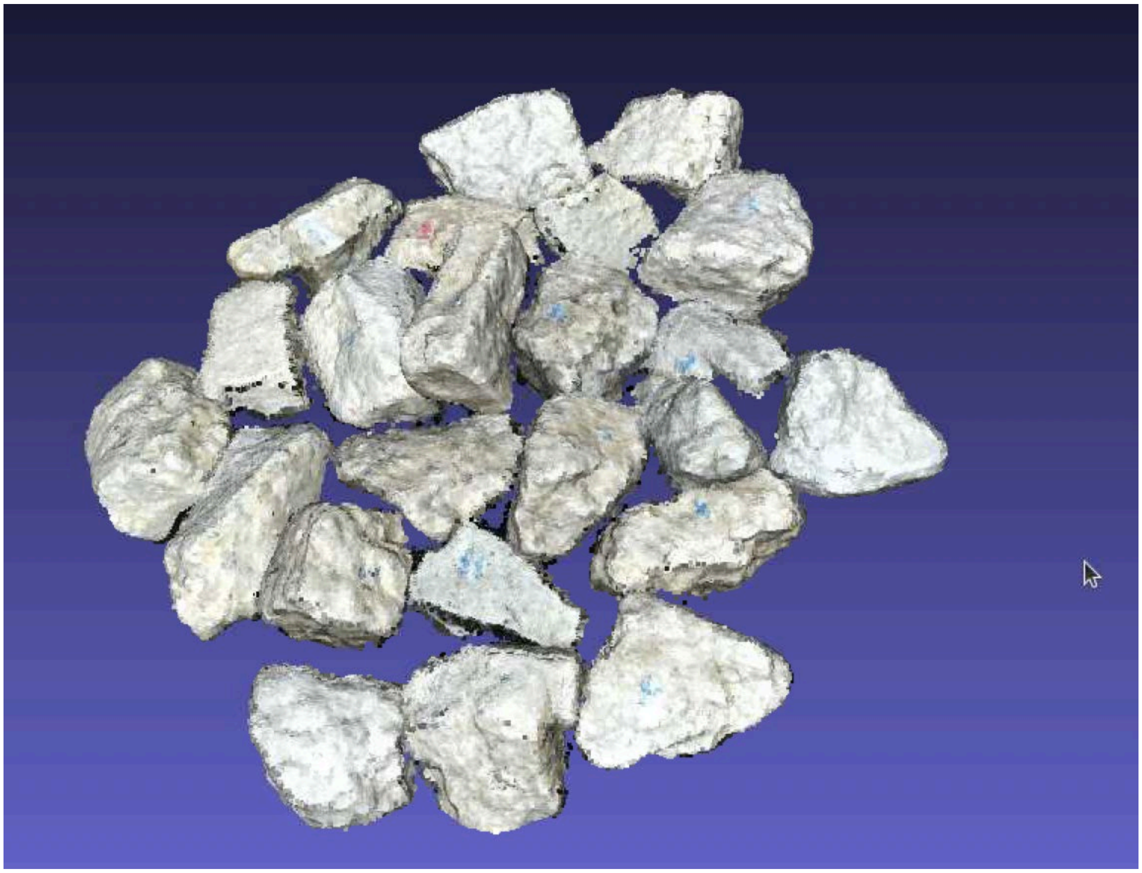

Figure 12. Cleaned stockpile point cloud.

8. Click File --> Export Mesh As... to save the cleaned point cloud as "**fused-clean.ply**" under the same directory as the original point cloud. Keep the default saving options and click OK (Figure 13).

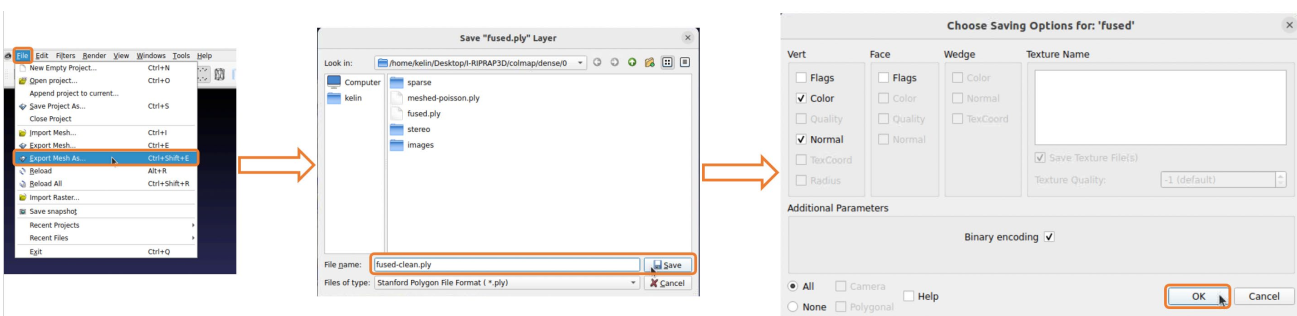

Figure 13. Saving the cleaned point cloud.

9. Close MeshLab window.

Scale 3D Point Cloud

The reconstructed 3D point clouds from structure-from-motion (SfM) algorithms have scale ambiguity. Therefore, manual labeling is necessary for calibrating the cleaned point cloud to the exact scale.

1. Click File to specify the images used for scale calculation (Figure 14).

Figure 14. Importing reconstructed project for scale calibration.

Note: The location of the project folder is under the Workspace folder specified in Step 2 of "Generating the 3D Point Cloud." Select the correct folder. The 3D reconstruction algorithm will generate the 3D point cloud results under dense directory starting from 0. Normally it is **/PATH/OUTPUT/dense/0**.

2. All images will be displayed in the window after Step 1 (Figure 15).

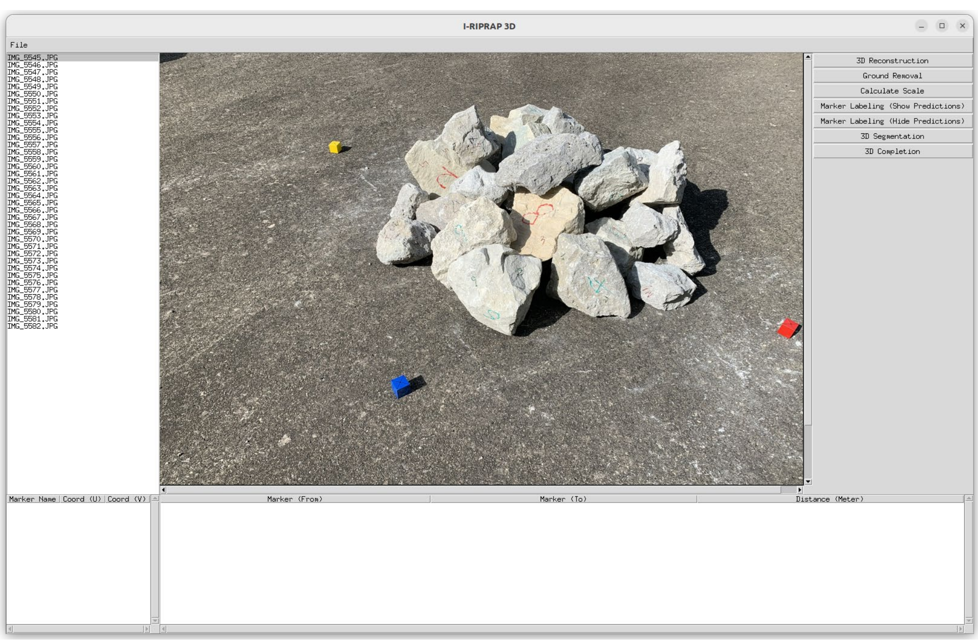

Figure 15. Images for scale calibration displayed in the I-RIPRAP 3D software.

Note: Users can use the mouse wheel to adjust the scale/size of the image and drag the image to adjust the location of the image.

3. Zoom in to the calibration objects in the images, double-click the centroids of the different color-coded calibration objects, and then add the corresponding labels (Figure 16).

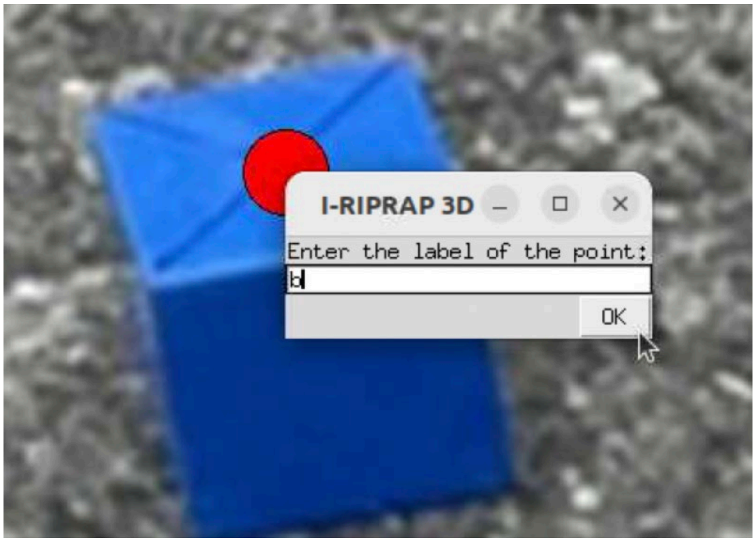

Figure 16. Labeling a calibration object in an image.

Notes:

- The labeled centroids will be displayed using a red circle.
- At least two randomly selected images should be labeled to enable the algorithm to calculate the scale. However, labeling four to five randomly selected images is recommended to provide a better ground truth for calculating the scale.
- After labeling, the image name will be highlighted in yellow in the image list bar (Figure 17).

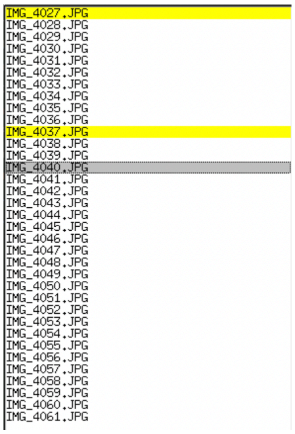

Figure 17. Image list bar in I-RIPRAP 3D software.

4. To remove an incorrectly labeled point, select its record in the bottom-left column and click DELETE on the keyboard (Figure 18).

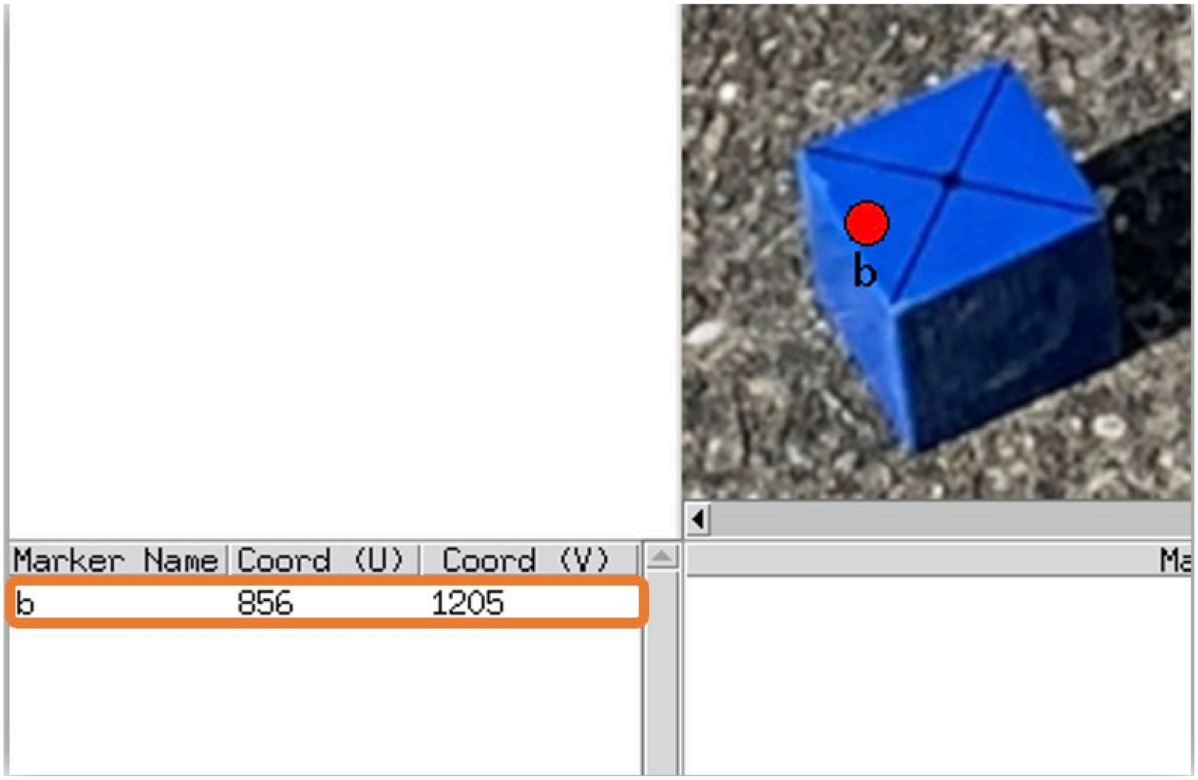

Figure 18. Calibration object coordinates record.

5. Assign the measured distances between each pair of color-coded calibration objects by double clicking the corresponding record at the bottom of the window (Figure 19).

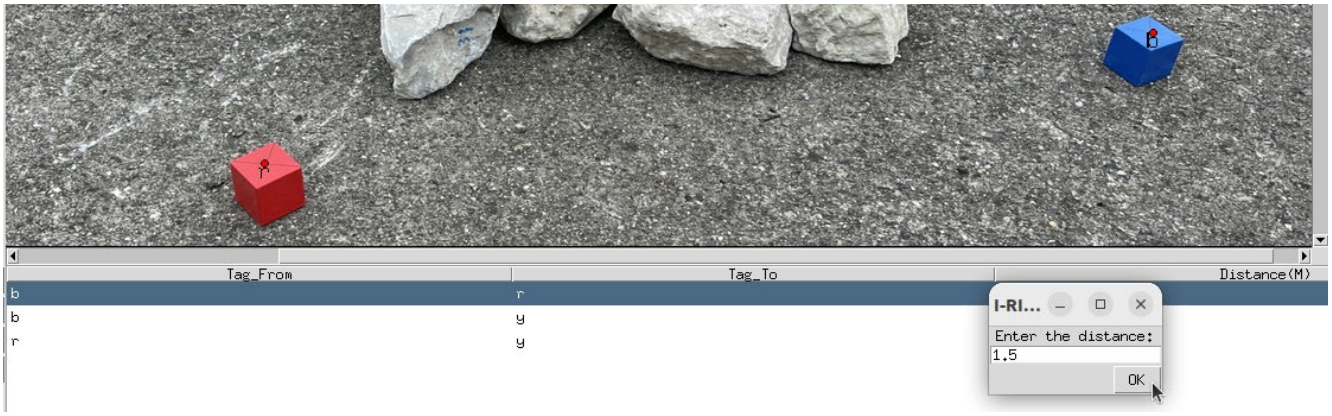

Figure 19. Assigning measured distances between calibration objects.

6. Click the Calculate Scale button to calculate the scale ratio from the reconstructed mesh space to the real-world space (Figure 20).

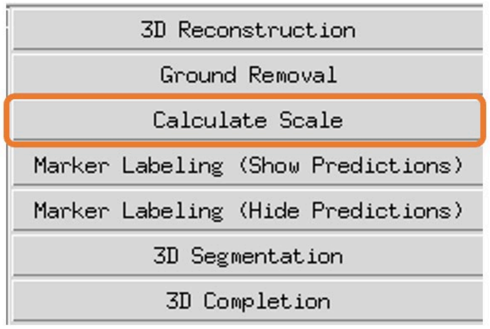

Figure 20. Calculate scale button.

7. After the scale is successfully calculated, the predicted locations of the color-coded calibration objects will be displayed as green triangles on each image (Figure 21).

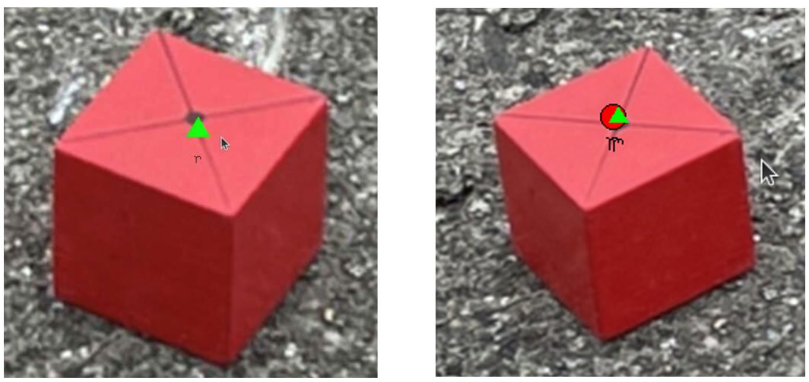

(a) (b)

Figure 21. Comparison between predicted locations on calibration objects in (a) an unlabeled image and (b) a manually labeled image.

8. Randomly check the predicted locations of the color-coded calibration objects on several images. If the predicted location is off compared to the real location, add new labels and repeat Step 6.

Note: The visibility of the predicted locations can be toggled by clicking the Marker Labeling (Show Predictions) button (Figure 22) and the Marker Labeling (Hide Predictions) button (Figure 23).

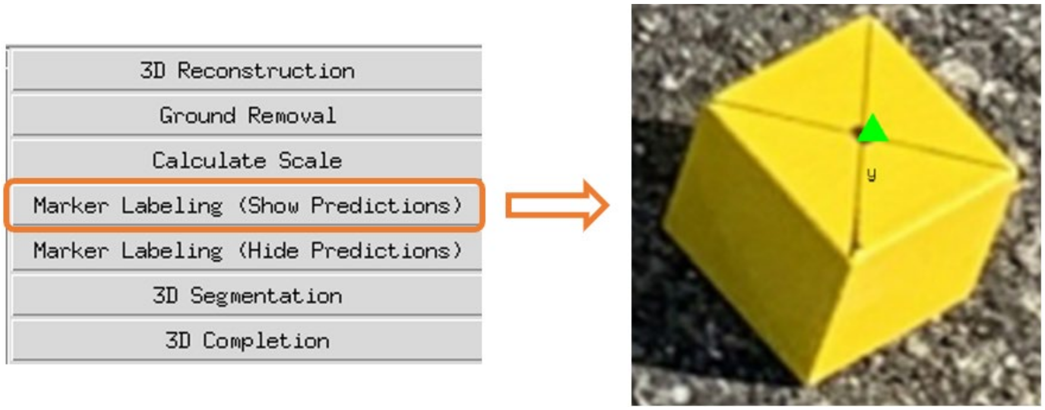

Figure 22. Showing predicted locations for calibration objects.

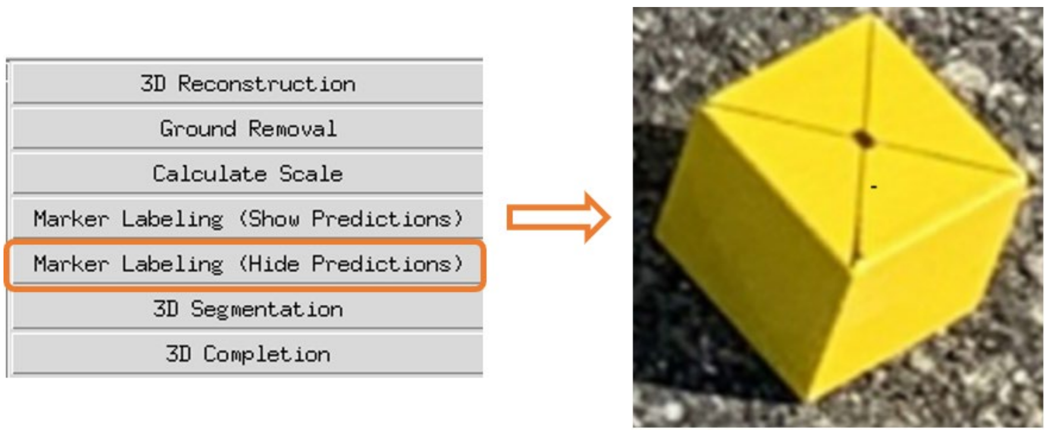

Figure 23. Hiding predicted locations for calibration objects.

3D SEGMENTATION

Segmenting a reconstructed 3D point cloud into instances is a cutting-edge technology for scene understanding in computer vision. I-RIPRAP 3D integrates PointGroup (Jiang, 2020), a deep-learning based high-performance 3D instance segmentation model, to help segment riprap instances.

1. Click the 3D Segmentation button (Figure 24), and the 3D segmentation result will pop out (Figure 25 and Figure 26).

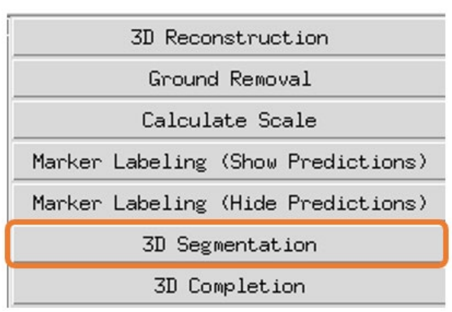

Figure 24. 3D segmentation button.

Note:

- The point clouds that pop up with window names “Fig1” and “Fig2” can be rotated by dragging the point cloud model. Zoom in and zoom out by rotating the mouse wheel.
- Rotate, zoom in, or zoom out on any point cloud in “Fig 1,” and the other two point clouds in “Fig 1” will rotate synchronously.

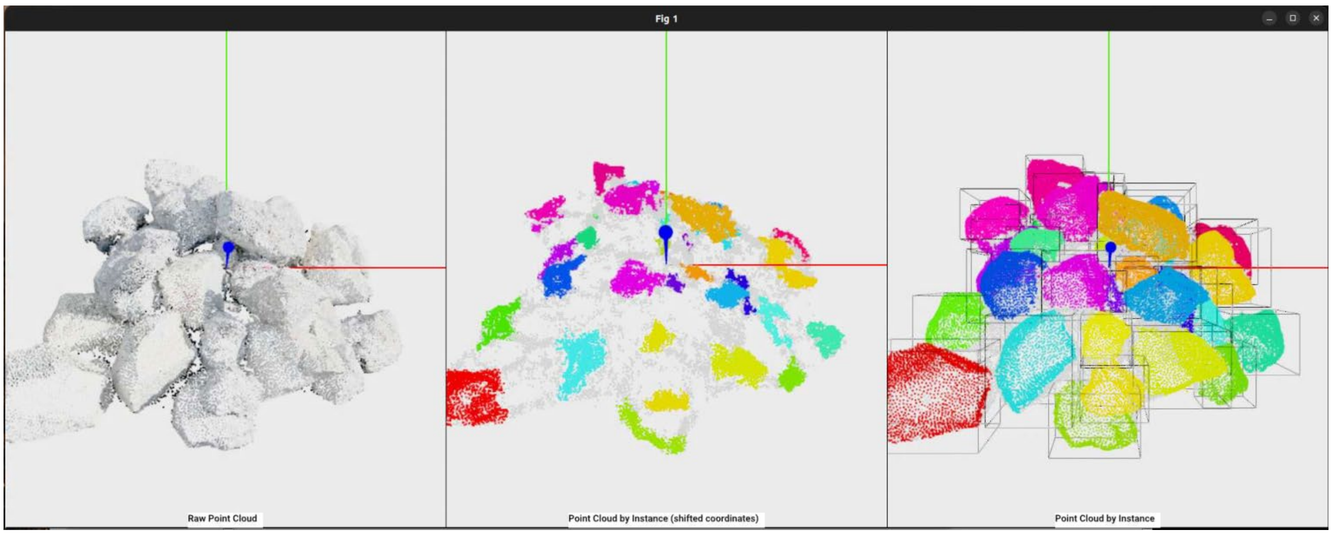

(a) (b) (c)

Figure 25. Visualizations of point cloud segmentation “Fig 1”: (a) raw input point cloud, (b) shifted coordinates point cloud by instance, and (c) point cloud by instance.

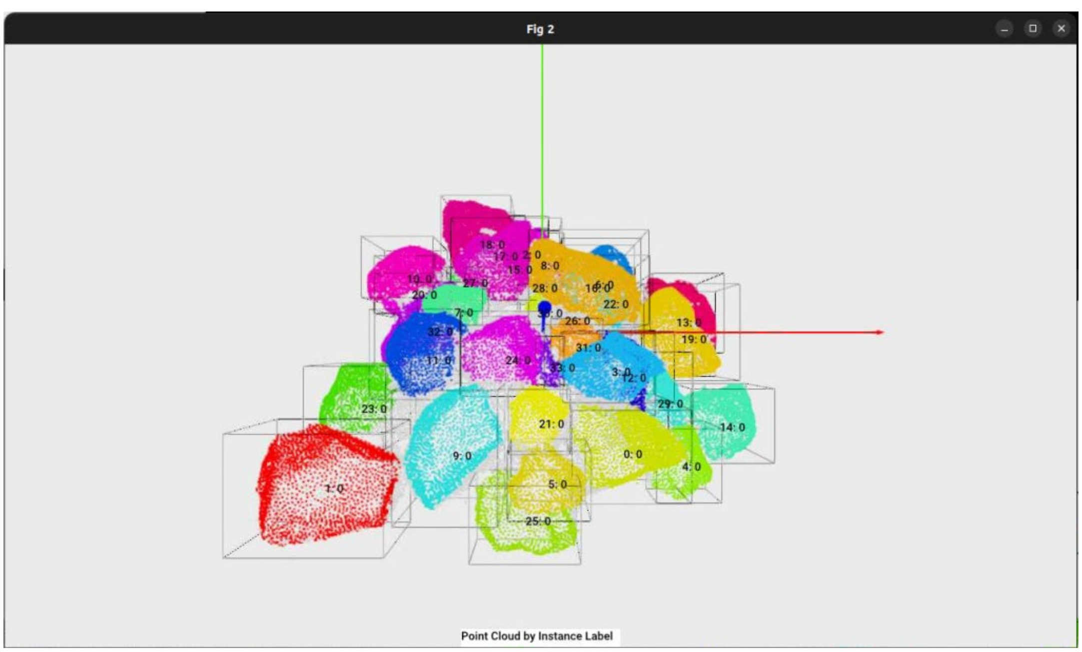

Figure 26. Point cloud by instance label “Fig 2.”

3D COMPLETION

Filling unseen parts of a segmented instance point cloud can be helpful for providing a complete description for each riprap instance. Another latest deep learning-based model, SnowflakeNet (Peng Xiang, 2021), is integrated for this task.

1. Click the 3D Completion button (Figure 27), and the 3D completion result will be saved in the folder **/PATH/OUTPUT/dense/0/fused_clean_segmentation**

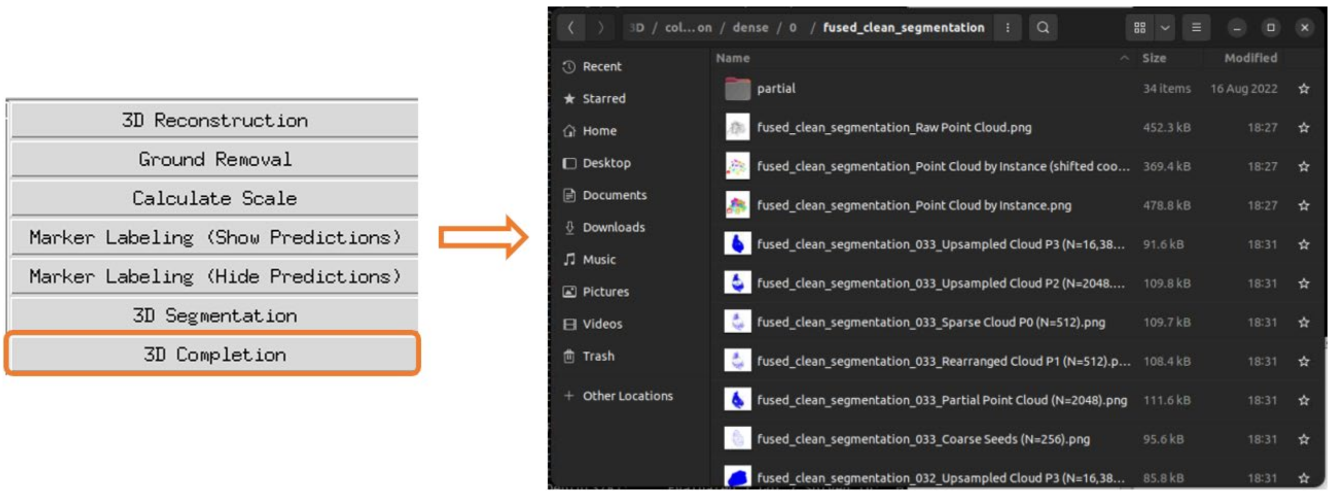

Figure 27. 3D completion results.

APPENDIX

EXAMPLE INPUT FILES

The input files for the I-RIPRAP 3D software should be multiple photos taken using a single camera from different viewpoints (Figure 28). Common image formats (JPG, PNG, BMP) are supported.

Input photos with larger resolutions offer better reconstruction quality but will also take more time to finish processing.

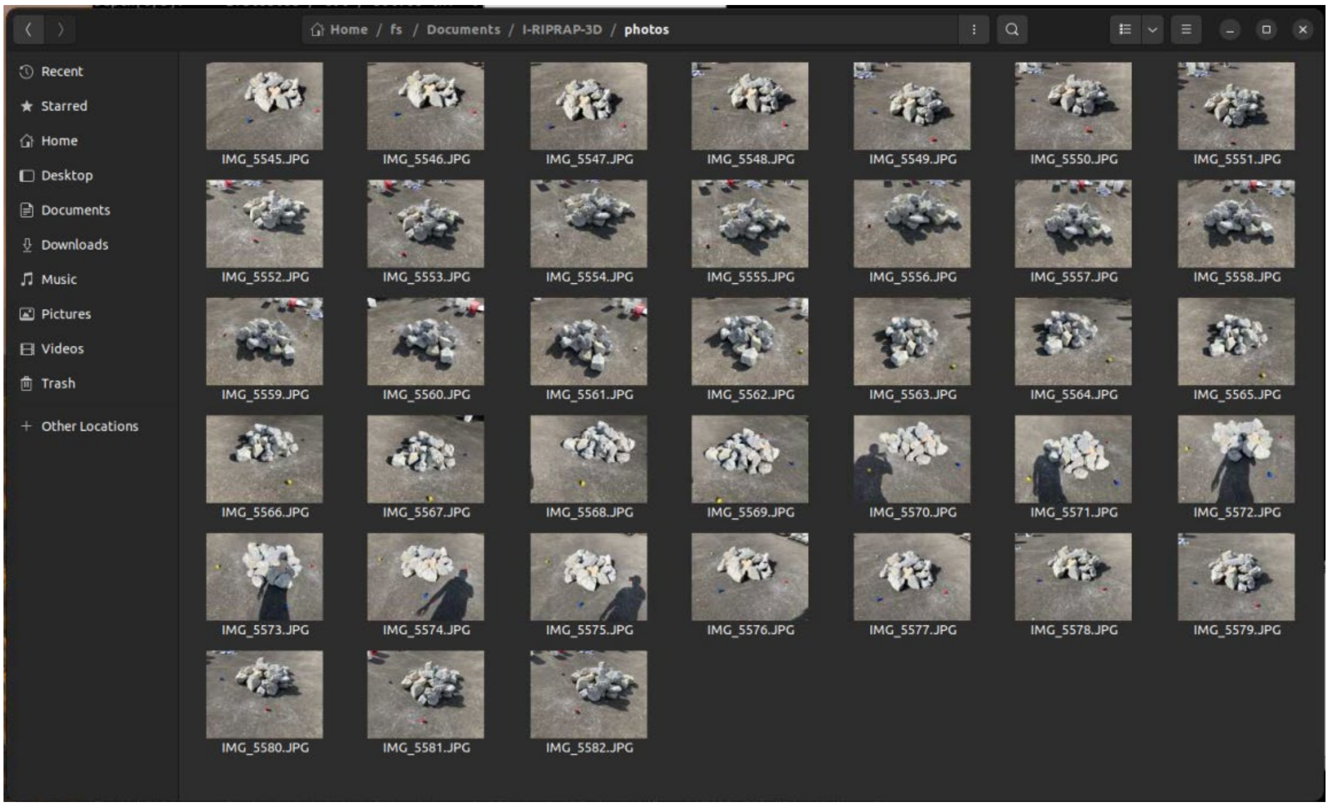

Figure 28. Example input images.

EXAMPLE OUTPUT FILES

Riprap Morphological Properties Summary Table

The fused_clean_segmentation.xlsx file summarizes the equivalent spherical diameter (ESD), shortest dimension (a), intermediate dimension (b), longest dimension (c), volume, surface area, 3D flat and elongated ratio (FER), and 3D sphericity of each individual riprap particle (Figure 29).

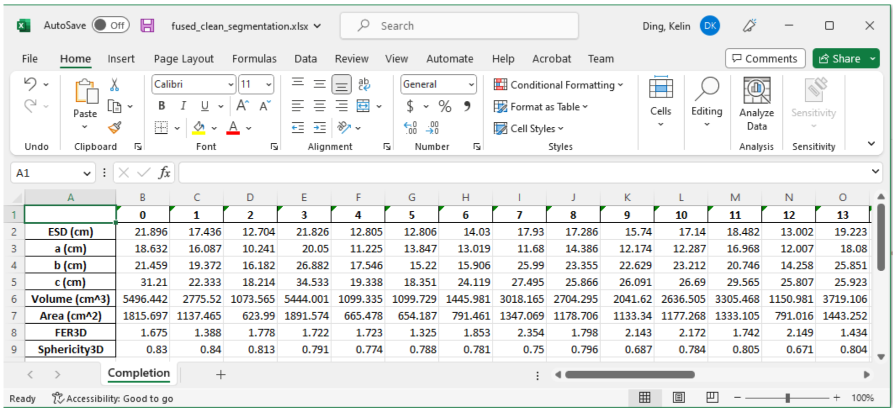

	0	1	2	3	4	5	6	7	8	9	10	11	12	13
ESD (cm)	21.896	17.436	12.704	21.826	12.805	12.806	14.03	17.93	17.286	15.74	17.14	18.482	13.002	19.223
a (cm)	18.632	16.087	10.241	20.05	11.225	13.847	13.019	11.68	14.386	12.174	12.287	16.968	12.007	18.08
b (cm)	21.459	19.372	16.182	26.882	17.546	15.22	15.906	25.99	23.355	22.629	23.212	20.746	14.258	25.851
c (cm)	31.21	22.333	18.214	34.533	19.338	18.351	24.119	27.495	25.866	26.091	26.69	29.565	25.807	25.923
Volume (cm^3)	5496.442	2775.52	1073.565	5444.001	1099.335	1099.729	1445.981	3018.165	2704.295	2041.62	2636.505	3305.468	1150.981	3719.106
Area (cm^2)	1815.697	1137.465	623.99	1891.574	665.478	654.187	791.461	1347.069	1178.706	1133.34	1177.268	1333.105	791.016	1443.252
FER3D	1.675	1.388	1.778	1.722	1.723	1.325	1.853	2.354	1.798	2.143	2.172	1.742	2.149	1.434
Sphericity3D	0.83	0.84	0.813	0.791	0.774	0.788	0.781	0.75	0.796	0.687	0.784	0.805	0.671	0.804

Figure 29. Example RIPRAP morphological properties summary table.

Top Views of Up-sampled Partial Point Clouds from Segmentation

Figure 30 presents intermediate and final reconstruction results from coarse seeds to upsampled clouds. “N” in the filename is the number of points in the point cloud. Point clouds with “P3” in the filename are the final results.

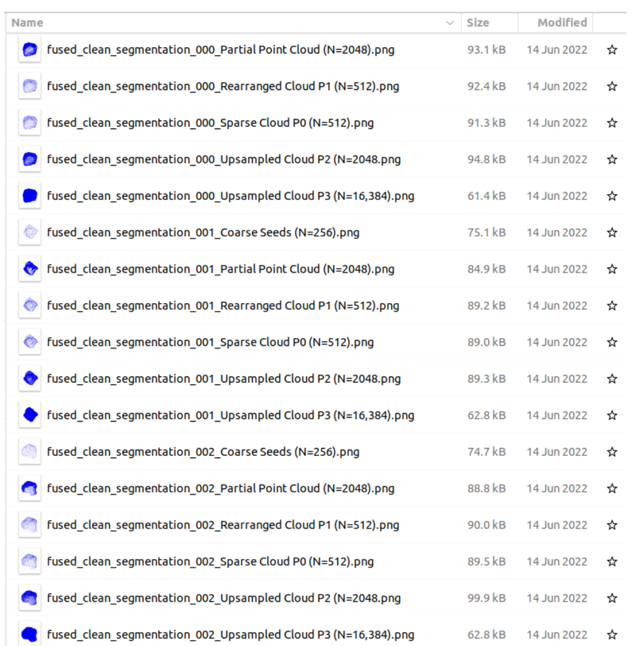

Name	Size	Modified
fused_clean_segmentation_000_Partial Point Cloud (N=2048).png	93.1 kB	14 Jun 2022
fused_clean_segmentation_000_Rearranged Cloud P1 (N=512).png	92.4 kB	14 Jun 2022
fused_clean_segmentation_000_Sparse Cloud P0 (N=512).png	91.3 kB	14 Jun 2022
fused_clean_segmentation_000_Upsampled Cloud P2 (N=2048.png	94.8 kB	14 Jun 2022
fused_clean_segmentation_000_Upsampled Cloud P3 (N=16,384).png	61.4 kB	14 Jun 2022
fused_clean_segmentation_001_Coarse Seeds (N=256).png	75.1 kB	14 Jun 2022
fused_clean_segmentation_001_Partial Point Cloud (N=2048).png	84.9 kB	14 Jun 2022
fused_clean_segmentation_001_Rearranged Cloud P1 (N=512).png	89.2 kB	14 Jun 2022
fused_clean_segmentation_001_Sparse Cloud P0 (N=512).png	89.0 kB	14 Jun 2022
fused_clean_segmentation_001_Upsampled Cloud P2 (N=2048.png	89.3 kB	14 Jun 2022
fused_clean_segmentation_001_Upsampled Cloud P3 (N=16,384).png	62.8 kB	14 Jun 2022
fused_clean_segmentation_002_Coarse Seeds (N=256).png	74.7 kB	14 Jun 2022
fused_clean_segmentation_002_Partial Point Cloud (N=2048).png	88.8 kB	14 Jun 2022
fused_clean_segmentation_002_Rearranged Cloud P1 (N=512).png	90.0 kB	14 Jun 2022
fused_clean_segmentation_002_Sparse Cloud P0 (N=512).png	89.5 kB	14 Jun 2022
fused_clean_segmentation_002_Upsampled Cloud P2 (N=2048.png	99.9 kB	14 Jun 2022
fused_clean_segmentation_002_Upsampled Cloud P3 (N=16,384).png	62.8 kB	14 Jun 2022

Figure 30. Top views of up-sampled partial point clouds from segmentation.

Segmented Partial Point Clouds

In the **/PATH/OUTPUT/dense/0/fused_clean_segmentation/partial** folder, segmented partial point clouds for each individual riprap particle are saved as .ply files (Figure 31). They can be opened by MeshLab.

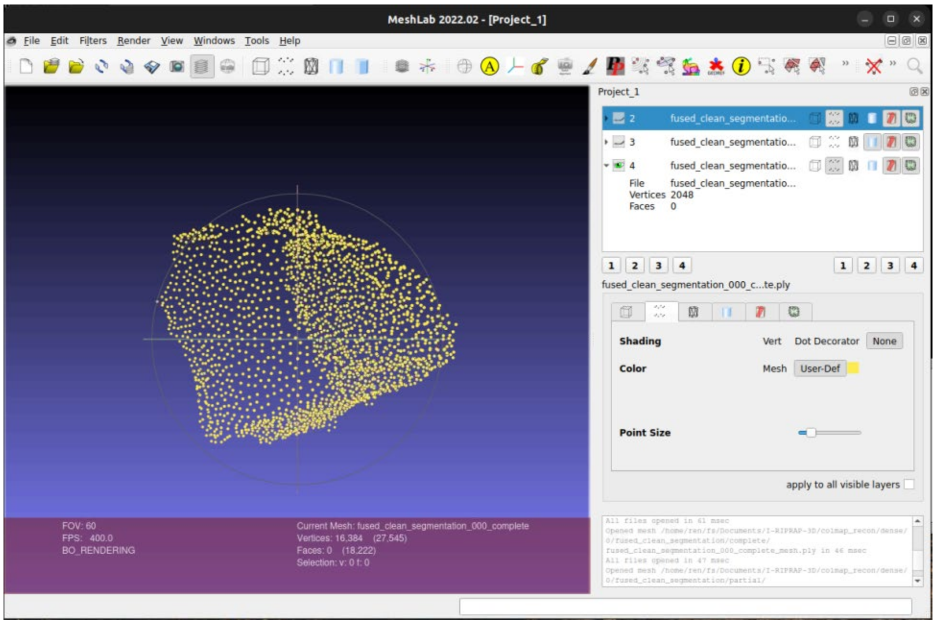

Figure 31. Segmented partial point cloud example.

Note: MeshLab can be opened by double clicking the MeshLab.AppImage icon on the desktop (Figure 32).

Figure 32. Opening MeshLab.

Completed Point Clouds and Meshes

In the **/PATH/OUTPUT/dense/0/fused_clean_segmentation/complete** folder, the completed point clouds and meshes for each individual riprap particle are saved as .ply files (Figure 33 and Figure 34). They can be opened by MeshLab.

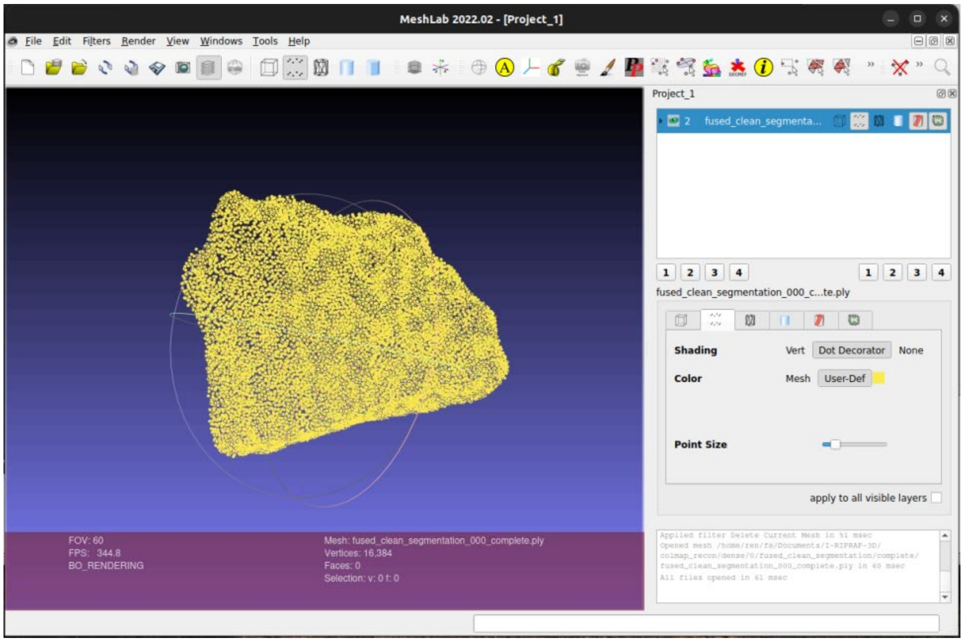

Figure 33. Completed point cloud example.

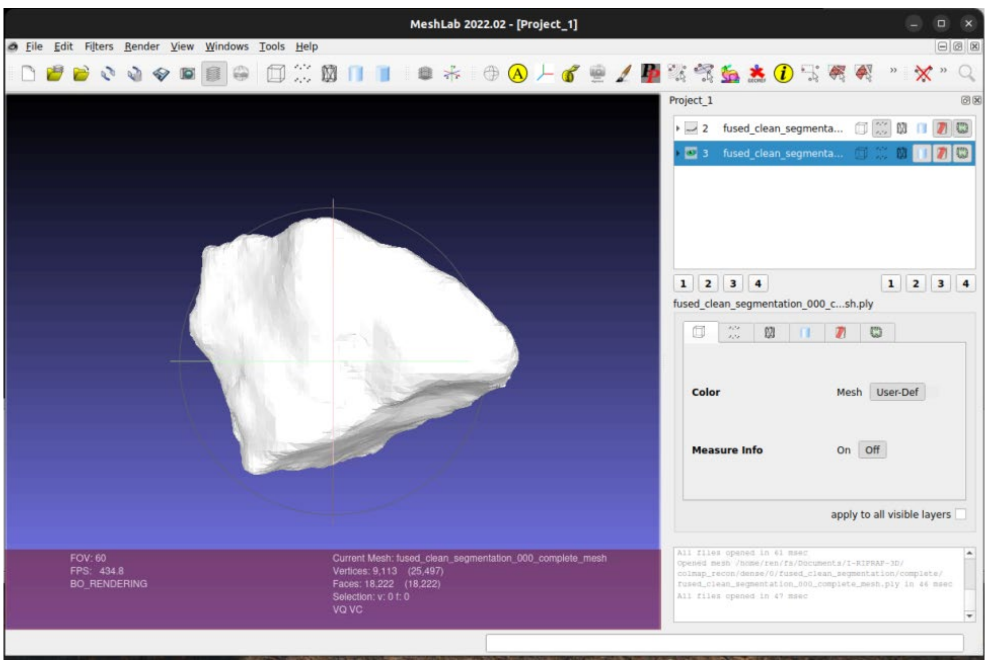

Figure 34. Completed mesh example.

REFERENCES

Cignoni, P. A. (2008). MeshLab: an Open-Source Mesh Processing Tool. *Eurographics Italian Chapter Conference.* The Eurographics Association.

Jiang, L. A.-W. (2020). PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Peng Xiang, X. W.-S.-P. (2021). SnowflakeNet: Point Cloud Completion by Snowflake Point Deconvolution With Skip-Transformer. (pp. 5499–5509). IEEE/CVF International Conference on Computer Vision (ICCV).

Schonberger, J. L.-M. (2016). Structure-from-Motion Revisited. *Conference on Computer Vision and Pattern Recognition (CVPR).*