Simulating virtual UAV surveys for benchmarking image-based 3D reconstruction methods: a comprehensive pipeline

Frederik Schulte[1], Beatrice Maier[2], Mia Ohlrogge[1], Lukas Winiwarter[1]

[1]Unit of Geometry and Surveying, Faculty of Engineering Sciences, University of Innsbruck, Innsbruck, Austria (frederik.schulte, mia.ohlrogge, lukas.winiwarter)@uibk.ac.at

[2]Unit of Engineering Mathematics, Faculty of Engineering Sciences, University of Innsbruck, Innsbruck, Austria beatrice.maier@uibk.ac.at

About the dataset

The data and code in this repository can be used to simulate various scenarios for evaluating the 3D reconstruction quality of different image-based methods. To simulate a wide range of applications, two scenes with distinct characteristics were digitized: A high-alpine rock glacier and a forested area. In this context, digitization implies that a real-world scene was surveyed, e.g., using light detection and ranging (LiDAR) and photogrammetric techniques, and subsequently imported as a 3D mesh into the software Blender, where virtual image flights with varying acquisition parameters can be simulated. Additionally, this repository contains Python scripts that enable fully automated planning of virtual flight missions, the processing of established photogrammetric methods (Structure from Motion (SfM) and Multi-View Stereo (MVS)) and novel radiance field methods (Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS)), as well as their conversion into point clouds, and the 3D evaluation of the results against the reference mesh. This enables analysis of 3D accuracy, completeness, and inlier ratio for each respective method.

This repository is part of a publication (Schulte et al., 2026). Detailed information on the processing and results is provided in this publication.

Content and workflow

This repository contains all the data and scripts required for the simulation. Specifically, the *blender* folder contains the Blender projects for the two scenes, including the mesh geometries and all camera locations and settings, so that each of the flights analyzed in the referenced paper can be recreated. This folder also contains the reference geometries as meshes and point clouds. These are required for evaluating accuracy, completeness, and the inlier ratio.

To execute the workflow, download and extract the complete repository, then follow the instructions provided in the *README.md* file. The folder structure of the *data.zip*

archive is shown below.
Only the scripts, the two Blender projects, the reference geometries (meshes and point clouds), the digital surface models (DSMs) used for terrain following during mission planning, and the bounding boxes of the reference areas are provided with the dataset. All other folders (*images*, *mission-parameter*, *nerfstudio*, *postshot* and *output*) are initially empty and are filled automatically when the pipeline is executed. The corresponding files and folders are marked as *(generated)* in the folder structure below.
The scripts always access the input files by their standard names (*dsm.txt*, *reference.ply*, *reference_pc.ply* and *bounding_box_ref.txt*), which refer to the rock glacier scene. To process the vegetation scene, the files with the suffix *_vege* have to be renamed to these standard names beforehand (e.g. *reference_pc_vege.ply* $\rightarrow$ *reference_pc.ply* and *dsm_vege_1m.txt* $\rightarrow$ *dsm.txt*).

```
data/
├── README.md
├── scripts/
│   ├── automation_simu.py
│   ├── import_blender.py
│   ├── mission_planning.py
│   ├── ms_batch.py
│   ├── parameter.ini
│   └── requirements.txt
├── blender/
│   ├── block_glacier_simu.blend
│   ├── vege_simu.blend
│   ├── dsm.txt
│   ├── dsm_vege_1m.txt
│   ├── reference.ply
│   ├── reference_vege.ply
│   ├── reference_pc.ply
│   └── reference_pc_vege.ply
├── metashape/
│   ├── bounding_box_ref.txt
│   ├── bounding_box_ref_vege.txt
│   └── <project_name>.psx (generated)
├── images/
│   └── <project_name>/ (generated)
├── mission-parameter/
│   └── <project_name>/ (generated)
│       ├── import_blender.txt
│       ├── import_metashape.txt
│       ├── import_realityscan.csv
│       ├── import_cameramodel_ms.xml
│       └── protocol_mission.txt
├── nerfstudio/
│   └── <project_name>/ (generated)
├── postshot/
│   └── <project_name>.psht (generated)
└── output/
    └── <project_name>/ (generated)
```

Software requirements

The pipeline combines several external software packages, which are listed in Table 1. The workflow is designed for Windows, but some parts use the Windows Subsystem for Linux (WSL). Agisoft Metashape Professional and Jawset Postshot require a license. The installation paths of the external tools have to be adapted in the scripts as described in the *README.md* file.

Software	Version	Purpose
Python	3.9	Mission planning, automation, evaluation
Blender	5.0	Rendering of the virtual image flights
Agisoft Metashape Professional	2.3.1	SfM and MVS (*ms_batch.py*)
Nerfstudio (nerfacto, via Docker)	1.1.5	NeRF training and point cloud export
Docker Desktop	4.40	Container for Nerfstudio
Jawset Postshot	1.0.110	3DGS training
WSL	2.6.3	Linux environment for 3DGS conversion
splat-transform[1]	1.10.2	Conversion of the 3DGS model
3DGS-to-PC[2]	-	3DGS to dense point cloud
CloudCompare	2.13.2	Filtering and C2C/C2M evaluation

Table 1: Software required to run the simulation and evaluation pipeline. The Python packages are listed in *requirements.txt* (C2C: cloud-to-cloud, C2M: cloud-to-mesh).

Data capturing and processing

This section describes the data acquisition and processing up to the generation of the meshes included in the Blender projects. All processing reports are provided in the *Reports.zip* archive.

1) Rock glacier

The first scene represents rough high-alpine terrain in the Ötztal valley. The Äußeres Hochebenkar rock glacier (WGS84: 46.834°N, 11.008°E) was selected for this purpose. It is located approximately four kilometers south of Obergurgl, Tyrol, and covers an area of about 42 ha. The north-west-facing rock glacier extends from about 2360 m (front) to approximately 2800 m above sea level. At an elevation of 2580 m, the rock glacier advances over a terrain break into steeper terrain (see Figure 1).
To capture the geometry as a simulation baseline, a measurement campaign was conducted on July 25, 2025. The lower part of the rock glacier was surveyed by uncrewed aerial vehicle (UAV) laser scanning (ULS) using a DJI Matrice 350 RTK (M350)[3] equipped with a DJI Zenmuse L2[4] scanner. As real-time kinematic (RTK) correction data are often unavailable in high-mountain regions, a local Global Navigation Satellite System (GNSS) base station (Trimble R12i[5]) was established, providing RTK correc-

[1] `https://github.com/playcanvas/splat-transform`, accessed 2026-10-19
[2] `https://github.com/Lewis-Stuart-11/3DGS-to-PC`, accessed 2026-10-19
[3] `https://enterprise.dji.com/matrice-350-rtk`, accessed 2026-10-19
[4] `https://enterprise.dji.com/zenmuse-l2`, accessed 2026-10-19
[5] `https://geospatial.trimble.com/en/products/hardware/trimble-r12i`, accessed 2026-10-19

tions to the UAV and raw data for post-processed kinematic (PPK) processing of the LiDAR trajectory. To determine the position of the base station, its raw data were post-processed against an EPOSA virtual reference station (VRS)[6], resulting in coordinates in ITRF2014 coordinate reference system. During the 55-minute data acquisition, approximately the lower half of the rock glacier was captured at high resoluion. To enable mesh texturing, over 1100 images were captured using the scanner's internal 20 MP camera. Table 2 shows the primary acquisition parameters.

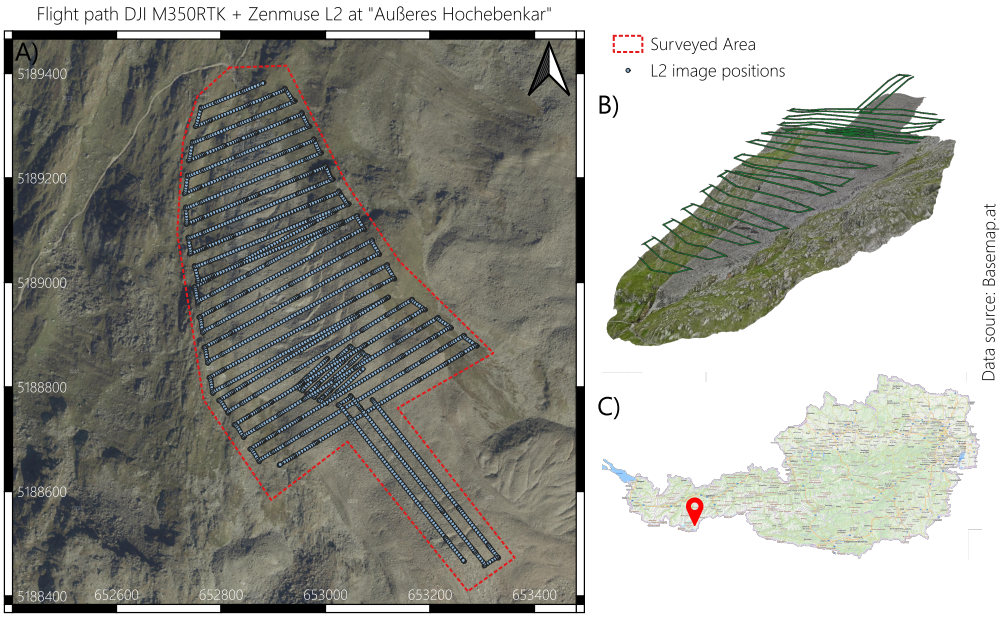


Figure 1: In A), the flight path of the M350 with the Zenmuse L2 used to generate the rock glacier reference mesh is shown. In B), the image positions are shown in 3D, and in C), the location of the survey area is illustrated.

[6] `https://www.eposa.at/`, accessed 2026-10-19

Parameter	Rock glacier mission	Vegetation mission
Flight altitude	75 m AGL	65 m AGL
Flight speed	6 m/s	3 m/s
Side overlap LiDAR	66 %	60 %
Side overlap photos	73 %	69 %
Returns	Up to 5 returns per pulse	Up to 5 returns per pulse
Frequency	240 kHz	240 kHz
Scanning mode	Repetitive	Repetitive
Ortho GSD	2 cm	1.75 cm
IMU calibration	Yes	Yes
Flight time	55 min	25 min
Ground beam diameter	$\sim 3 \times 9$ cm	$\sim 2.6 \times 7.8$ cm
Number of points	$\sim 800 \cdot 10^6$	$366 \cdot 10^6$
Average point density	1940 points/m^2	2900 (2×1450) points/m^2

Table 2: Comparison of the flight parameters of the rock glacier and vegetation missions (AGL: above ground level; GSD: ground sampling distance). For the vegetation mission, the point cloud density is multiplied by two because the area was flown twice in a cross pattern with identical settings, with the flight direction rotated by 90°. Detailed descriptions of the settings can be found in the DJI Zenmuse L2 User Manual[7]

Data processing was performed using DJI Terra[8] (V5.1.1). Strip adjustment was conducted using the acquired LiDAR, inertial measurement unit (IMU), and GNSS data. The results are documented in *Block_glacier_Terra_Report.pdf*. A fixed solution was obtained for 100% of the trajectory, and the average trajectory errors amount to 5–7 mm in position. The processed point cloud was smoothed and meshed in DJI Terra using the acquired images. Finally, the mesh with approx. 800 million points was transformed from ITRF2014/UTM32 to a local topocentric system (more information can be found in the main paper).

To ensure the scene is comparable to a real photogrammetric mission, an environment model was also created. This provides a particular advantage during the SfM process, as features in the background can also be detected and matched. For this purpose, the DSM (cell size 1 m) and the digital orthophoto (DOP; GSD = 0.2 m) of the surrounding area were downloaded from the tiris geodata portal[9]. The DSM was then imported into CloudCompare[10] (V2.13.2) as a point cloud. In the subsequent step,

[7] `https://dl.djicdn.com/downloads/Zenmuse_L2/Zenmuse_L2_User_Manual_v1.0_EN.pdf`, accessed 2026-10-19.

[8] `https://enterprise.dji.com/dji-terra`, accessed 2026-10-19

[9] `https://www.tirol.gv.at/sicherheit/geoinformation/geodaten-tiris/`, accessed 2026-10-19

[10] `https://www.cloudcompare.org/`, accessed 2026-10-19

the point cloud was transformed into the same topocentric system as the UAV-based point cloud, colorized using the RGB values of the DOP, and finally converted into a colorized mesh using 2.5D Delaunay triangulation.
To achieve an optimal compromise between the level of detail and performance in the later simulation, the final scene was composed of both datasets. For the lower rock glacier area (420 m × 620 m), the ULS mesh was used to maintain a high resolution (16 million faces) in the scene. The larger environment model (2300 m × 1800 m) was inserted around this core model, completing the scene with a sufficient level of detail.

2) Vegetation

The second scene represents a forested area and serves as an example for assessing the reconstruction accuracy of vegetation geometries. The study site is located in Scharnitz, Austria (WGS84: 47.382°N, 11.276°E). The forest is relatively sparse, so that substantial parts of the ground surface remain visible in the simulated images. To capture this area, a combined terrestrial laser scanning (TLS) and ULS approach was employed on May 13, 2026. A particular challenge was the co-registration of the two point clouds. Although the TLS instrument (RIEGL VZ-600i[11]) is equipped with an RTK-capable GNSS receiver, poor satellite reception beneath the forest canopy prevented a consistently accurate georeferencing solution throughout the survey.
To enable the accurate integration of the ULS and TLS point clouds, an initial flight was conducted using an M350 equipped with a DJI Zenmuse P1[12] camera and a 35 mm lens. The primary objective of this flight was to survey coded 8-bit targets distributed throughout the study area. To achieve highly accurate georeferencing, a GNSS base station (Trimble R12i) was established, as described for the first scene. The camera centers were subsequently corrected in a PPK workflow based on an EPOSA VRS solution and the software RedToolBox[13]. Since the UAV operated above the tree canopy, PPK-based georeferencing was not problematic, in contrast to the TLS survey. To enable precise estimation of the target centers, the mission was flown at an altitude of 60 m AGL using terrain-following mode with a forward overlap of 90% and a side overlap of 80%, resulting in a GSD of 7.5 mm. To ensure the highest possible image quality, the flight was conducted with an f-number of at least f/5.6, a shutter speed of 1/1000 s, and automatic ISO, kept as low as possible.
Subsequently, image orientation and bundle adjustment of the 440 acquired images were performed using Agisoft Metashape (V2.3.1). Finally, the coordinates of six targets, which had been positioned in areas with unobstructed sky visibility within the study area, were automatically determined. The high redundancy of the observations (11–49 automated image measurements per target) resulted in a low median reprojection

[11] `https://www.riegl.com/en-austria/products/detail/riegl-vz-600i`, accessed 2026-10-19
[12] `https://enterprise.dji.com/zenmuse-p1`, accessed 2026-10-19
[13] `https://www.redcatch.at/redtoolbox/`, accessed 2026-10-19

error of 0.3 pixels (2 mm), ensuring that the targets were accurately surveyed within the georeferenced coordinate system. These targets were subsequently used as ground control points (GCPs) for the georeferencing of both the ULS and TLS point clouds.
The ULS point cloud was acquired using a DJI Zenmuse L2 sensor. A cross-flight pattern with high overlap was selected (see Table 2), ensuring good penetration of the tree canopy. These data were also processed using the aforementioned VRS-supported PPK workflow in DJI Terra (V5.2.5). This resulted in a 100% fixed trajectory solution. The photogrammetrically surveyed targets were imported as check points in the strip adjustment and exhibited an average 3D deviation of 1.1 cm. To eliminate these remaining discrepancies, the processed point cloud was transformed into the coordinate reference system (CRS) defined by the GCPs using a 6-parameter transformation. To generate a surrounding environment model, an additional mission covering the area surrounding the study site was flown. For this mission, a flight altitude of 120 m, a flight speed of 6 m/s, and a LiDAR side overlap of 40% were selected. The mesh was generated in DJI Terra and textured using the 110 images acquired by the internal camera of the Zenmuse L2. A comparatively low resolution of 540,000 faces was chosen in order to improve rendering performance.
In the ULS data, the forest floor and the tree stems are only captured by the comparatively few returns penetrating the canopy. To capture these areas with considerably higher point density, an additional 49 TLS scan positions were acquired using the RIEGL VZ-600i. To transform the TLS point cloud into the photogrammetric CRS, a bireflex target was mounted on each GCP using a custom 3D-printed adapter with a known offset, allowing the targets to be scanned precisely by the TLS system (see Figure 2B).



Figure 2: In A), the flight path of the M350 with the Zenmuse L2 used to generate the reference point cloud is shown. In B), the targets used to link the TLS and photogrammetric surveys can be seen, and in C), the location of the survey area is shown.

The scans were processed using the manufacturer's software RiSCAN PRO[14] (V2.24.2). All scan echoes were retained; however, points with reflectance values outside the range of −25 to 5 dB or a pulse shape deviation greater than 10 were filtered out. Subsequently, a plane-based Multi Station Adjustment (MSA2) was performed, in which the GCPs were introduced as fixed control points. According to the results (*Vege_RiScan_MSA2_Report.pdf*), the estimated 3D position accuracy in the global coordinate system is better than 2 cm for 36 of the 49 scan positions (median: 1.4 cm) and the residuals at the GCPs amount to 7 mm on average (maximum: 2.0 cm). Since the resulting point cloud consisted of more than 1.1 billion points, it was cropped to the area of interest and subsampled using an octree with a cell size of 5 mm, resulting in a final dataset containing 603 million points.

To create the final model, both the TLS and ULS point clouds were imported into Agisoft Metashape. A high-resolution mesh consisting of 555 million faces was generated from the laser scanning data and textured using the images of the Zenmuse P1 sur-

[14] `https://www.riegl.com/en-austria/products/detail/riscan-pro`, accessed 2026-10-19

vey, whose orientations from the previous bundle adjustment were retained. To enable efficient use of the mesh within the simulation pipeline, the model was cropped to an extent of 118 m × 116 m and the number of faces was reduced to 50 million. Finally, both the detailed model and the surrounding environment model were transformed into a topocentric coordinate system.

Reference model

To evaluate the accuracy of the methods, the reference mesh was exported for both scenes. In addition, point clouds were sampled from these meshes to facilitate the analysis of completeness and inlier ratio. Table 3 shows the extent of the reference area, the number of faces in the mesh and the number of points in the point cloud. These files are all located in the *blender* folder of this repository. Since the simulated images are rendered directly from the scene meshes, the reference geometry is error-free by definition within the simulation. The accuracy of the field surveys described above therefore only affects the realism of the scenes, but not the evaluation of the reconstruction methods.

Scene	**Extent [m]**	**Number of faces [M]**	**Number of points [M]**
Rock glacier	120 × 80	1.2	48
Vegetation	50 × 50	17	50

Table 3: Information about the reference data for the two scenes. The extent of the reference area, the number of faces in the mesh, and the number of points in the point cloud are listed.

License

The DSM and DOP used for the environment model of the rock glacier scene were provided by Land Tirol (tiris) as open government data (source: Land Tirol – data.tirol.gv.at).

References

Schulte, F., Maier, B., Ohlrogge, M., Akwensi, P., and Winiwarter, L. (2026). Best-practice UAV data acquisition for NeRF and 3DGS: Surpassing conventional MVS in complex scenarios. Unpublished Script.