Volume Change Analysis in Environmental Point Clouds

Gorazd Gorup, Ciril Bohak
University of Ljubljana, Faculty for Computer and Information Science
Večna pot 113
1000 Ljubljana, Slovenia
{gorazd.gorup, ciril.bohak}@fri.uni-lj.si

ABSTRACT

Observation of landscape change is essential for strategic planning, risk assessment, and disaster prevention in volatile environments, especially on steep rocky slopes and cliffs where rockfalls can endanger human lives. To support the analysis of large point cloud scans of these environments, we propose a novel method for estimating local terrain changes and computing corresponding material displacement by combining existing change detection algorithms, such as M3C2, with the identification of smaller areas of significant change and the construction of localized watertight meshes. Our approach processes complex terrain to extract and visualize critical information, providing more detailed volumetric analysis than standard software, which is often limited to point coloring. We evaluate our method on various simple synthetic scenarios, including hill, cube, and sphere tests, as well as a real UAV dataset of coastal cliffs captured over multiple years. The experimental results show that while the method can become unstable with complex geometry or noisy data, it produces plausible volume displacement values and demonstrates high potential for more efficient terrain deformation analysis. We conclude that our proposed pipeline provides a strong foundation for advanced geomorphological observation and the development of automated systems for land risk assessment.

Keywords

Point clouds, mesh reconstruction, M3C2, clustering, rockfalls, cliffs, geomorphological analysis.

1 INTRODUCTION

In geomorphological studies, land observation is essential for effective decision-making and risk assessment. On steep rocky slopes and cliffs, where rockfalls and unstable terrain can endanger human lives, experts must carefully analyze changes in slope structure and material displacement to identify potential hazards and respond appropriately. For effective analysis of rockfall frequency and trajectory, they need to obtain information about displacement location and the volume of change [HHH$^+$24]. In operational monitoring campaigns, this analysis is often repeated across many surveys acquired over months or years, making the scalability and interpretability of the workflow critical.

Laser-scanned data is useful in this regard because it provides detailed representations of real terrain in the form of point clouds. Various projects and initiatives have collected Terrestrial Laser Scanning (TLS) and/or Unmanned Aerial Vehicle (UAV) scans in critical areas over certain time intervals, sometimes spanning several years. These 4D data sets allow for detailed analysis of the terrain, but the volume of data can quickly become overwhelming. A single survey may contain millions of points, while long-term monitoring campaigns can produce dozens to hundreds of temporally aligned scans. Manual inspection of such data through point-wise visualization alone becomes impractical and error-prone, particularly when experts must distinguish true geomorphological change from sensor noise, vegetation artifacts, or registration uncertainty.

An effective terrain change analysis workflow should provide (i) robust detection of significant geometric change in noisy point cloud data, (ii) localization of changes into spatially coherent regions, (iii) quantitative estimation of displaced material volume, (iv) visualization that supports rapid expert interpretation, and (v) scalability to repeated surveys acquired over long monitoring periods.

The selected Slovenian coastal cliff dataset represents a particularly suitable case study for these challenges. The Strunjan cliffs are characterized by steep and highly irregular terrain, localized erosion, episodic rockfalls, and dense vegetation coverage, all of which complicate automated geometric analysis. Furthermore, the dataset spans multiple years and contains scans acquired under varying environmental and seasonal conditions, introducing additional variability in lighting, vegetation state, and point-cloud density. These properties make the dataset representative of realistic long-term monitoring scenarios faced by geomorphologists and environmental monitoring agencies.

At the same time, the intended application scope of the proposed method should be clearly distinguished from broader large-scale terrain deformation analysis. The primary focus of this work is the detection and volumetric estimation of localized cliff and rockfall-related surface changes rather than large-scale terrain deformation processes such as landslides, earthquakes, or seismic fault displacement. Such phenomena often involve substantially different spatial scales, deformation characteristics, and validation requirements, and may require alternative reconstruction and modeling assumptions. We therefore position the proposed workflow primarily as a tool for localized geomorphological monitoring of steep terrain and cliff environments.

While considerable work has been done on change detection in point cloud scans, to our knowledge, the problem of processing and presenting change information in a clear analytical way has yet to be explored in depth. Software such as CloudCompare[1] and MeshLab[2] provide some functionality to visualize changes, but are limited to point coloring and do not offer further analysis, such as quantifying how much material was displaced in specific areas. Experts could benefit from focused visualization of significant changes and the change in volume from one point cloud scan to the next.

Recent research has improved the robustness of point cloud change detection algorithms, especially through methods such as Multiscale Model-to-Model Cloud Comparison (M3C2). However, comparatively less attention has been given to how detected changes should be aggregated, quantified, and communicated to domain experts in an interpretable manner. In practice, experts are interested not only in where change occurred, but also in whether the change corresponds to a coherent rockfall event, the size of the displaced volume, and the confidence with which the event can be interpreted.

Our work focuses on processing point cloud data of complex terrain to extract and visualize important information. By combining existing change detection algorithms, such as M3C2, with the formation of smaller areas of significant change and the construction of localized watertight meshes, we propose a novel approach for fine local volume change estimation from point clouds. The primary contributions of this research include:

- A novel pipeline for local volume change estimation that integrates M3C2 distance computation with localized mesh reconstruction.
- The implementation of a workflow using PCA-driven Oriented Bounding Box (OBB) and mesh extrusion to create watertight volumes for accurate subtraction.
- A web-based visualization tool that overlays estimated volume changes and uncertainty markers directly onto 3D terrain models. An early version of the tool is available at `https://lgm.fri.uni-lj.si/gorazd/cliffs_viz/`.

We structure this paper as follows: section 2 presents recent advancements and the state of the art in point cloud change detection and volume estimation. Section 3 describes our approach in detail, while section 4 presents our experimental setup and results. Finally, section 5 discusses the results and future work. We conclude with section 6.

2 RELATED WORK

Albawan *et al.* [AQL24] provide a comprehensive review of remote sensing methods for 3D landslide assessment, covering conventional field surveys, fusion-based multi-source approaches, and AI-driven methods, including deep learning and 3D geometric analysis. They note that hazard severity is primarily determined by the volume and rate of surface deformation, yet many existing approaches are still limited to 2D or 2.5D representations. This gap highlights the need for fully 3D, localized volumetric analysis. Žabota *et al.* [ŽBK23] place such monitoring within the broader context of sustainable geohazard management, emphasizing the importance of automated, scalable workflows for real-world deployment.

Although existing literature provides strong foundations for change detection and surface reconstruction, most methods address these tasks independently. Point-wise change detection methods often stop at identifying geometric differences, while volumetric approaches typically assume pre-segmented objects or require substantial manual intervention. In practice, domain experts require intuitive workflows that automatically progress from noisy point cloud observations to interpretable event-level quantities such as localized displaced volume.

2.1 Change Detection

The Multiscale Model-to-Model Cloud Comparison (M3C2) algorithm, introduced by Lague *et al.* [LBL13], is the leading method for 3D change detection in geoscientific point clouds. Unlike simpler

[1] `https://www.cloudcompare.org/`
[2] `https://www.meshlab.net/`

distance measures such as Cloud-to-Cloud (C2C) or Cloud-to-Model (C2M), M3C2 computes distances along locally oriented surface normals, making it well suited for geometrically complex terrain such as cliff faces. DiFrancesco *et al.* [DBH20] studied the sensitivity of M3C2 to the projection cylinder diameter, finding that a diameter of one to two times the point spacing provides the best balance between capturing true change and suppressing noise. Spatial averaging within the cylinder reduces residual noise but may mask small features if the diameter is too large.

Several works extend or challenge M3C2. Williams *et al.* [WAW$^+$21] proposed a multi-directional variant that estimates a dominant movement direction at each point rather than projecting along the surface normal, making it more appropriate for processes such as rock glacier creep or translational landslides, where displacement is not perpendicular to the surface. Yang *et al.* [YS23] introduced Patch-Based M3C2 (PB-M3C2), a patch-based approach that uses supervoxel over-segmentation and local plane fitting to estimate uncertainty more robustly. Jia *et al.* [JdM$^+$25] proposed a graph-based method that quantifies topographic change while remaining robust to registration errors. Most recently, Jia *et al.* [JdM$^+$26] described a Riemann manifold approach that captures both rigid and non-rigid surface deformations, representing a broader geometric characterization than M3C2 offers. A survey by Yang and Holst [YH25] compared the full landscape of correspondence-based methods, including C2M, M3C2, local Iterative Closest Point (ICP), Feature to Feature Supervoxel-based Spatial Smoothing (F2S3) [GZW20], and patch matching, highlighting how the choice of correspondence definition fundamentally determines what kinds of change are detectable. A related approach by Wang *et al.* [WBH$^+$26] fuses 3D geometry with co-registered RGB imagery to estimate dense displacement vector fields via hierarchical patch matching, conceptually similar to PB-M3C2 but enriched with photometric information. Winiwarter *et al.* [WAH21] introduced M3C2 Error Propagation (M3C2-EP), which incorporates sensor model information to compute Level of Detection (LoDetection).

Detecting per-point change is only the first step. Grouping changed points into meaningful objects enables event-level analysis. Anders *et al.* [AWM$^+$21] introduced the 4D Objects-by-Change (4D-OBC) method, which clusters spatio-temporal segments with similar change histories in Light Detection and Ranging (LiDAR) time series into discrete topographic change objects. These methods, along with M3C2 and M3C2-EP, are implemented in the open-source py4dgeo Python library [HKA$^+$26]. Weidner *et al.* [WW21] demonstrated a practical workflow combining M3C2-EP with DBSCAN clustering, providing insight into parameter selection for rock slope monitoring. A semi-automatic approach using k-NN-based segmentation of UAV photogrammetric point clouds has also been proposed as a lower-cost alternative to TLS-based workflows.

The literature demonstrates that modern change detection methods can robustly identify geometric differences even in noisy and irregular terrain. However, most approaches ultimately produce point-wise or patch-wise measurements that remain difficult to interpret directly in large-scale monitoring scenarios. Existing workflows generally rely on color-coded point cloud visualizations, which provide limited support for identifying coherent geomorphological events or quantifying their magnitude. Furthermore, many methods focus primarily on correspondence accuracy and uncertainty estimation, while leaving subsequent aggregation, volumetric interpretation, and visualization to manual post-processing. Our work addresses this gap by integrating change detection with automatic clustering and localized volumetric reconstruction.

2.2 Volume Estimation

Translating clustered change detections into volumetric estimates presents specific challenges. Bonneau *et al.* [BDH19] systematically compared surface reconstruction algorithms (e.g., Power Crust, Convex Hull, and Alpha-shape) for reconstructing 3D rockfall objects from point clouds, finding that 2.5D rasterization is inadequate for complex geometries and that watertight mesh construction is essential for reliable volume computation. The Alpha-shape algorithm requires careful parameter selection but, when combined with iterative mesh watertightness checks as demonstrated by Weidner *et al.* [WW21], produces accurate and geometrically faithful results. Walton *et al.* [WW23] conducted a dedicated accuracy study, showing that data resolution and the change filter threshold are the primary sources of volume estimation error, while point-level noise has comparatively little effect. VoxFall [FGTG25] takes a complementary voxel-based approach, wrapping extracted meshes into a voxel grid to compute volumetric differences. Mineo *et al.* [MCZP24] demonstrated TLS-based rockfall volume estimation using CloudCompare combined with traditional field surveys and 3D trajectory simulation.

Existing volumetric reconstruction methods demonstrate that watertight mesh generation is essential for reliable rockfall volume estimation, especially in complex cliff geometries where 2.5D assumptions fail. However, many approaches either require manually isolated rockfall candidates or operate as standalone reconstruction procedures. As a result, the complete workflow from raw temporal point clouds to localized, interpretable volume estimates remains

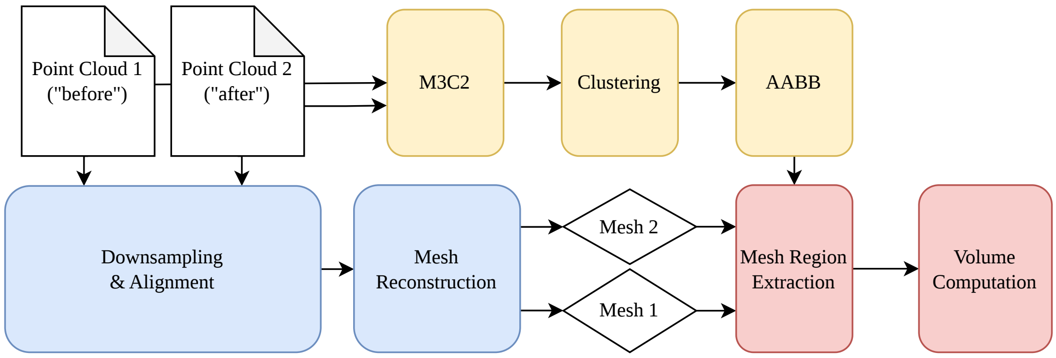

Figure 1: The pipeline of our method: A pair of point clouds is preprocessed to produce mesh reconstructions. Changes between the point clouds are computed using the M3C2 algorithm and mapped onto Point Cloud 1. The points are clustered, and Axis-Aligned Bounding Boxes are computed from the clusters. These bounding boxes are then used to extract parts of the mesh for final volume computation. The data preprocessing step is marked in blue, the change detection step is marked in yellow, and the volume estimation step is marked in red.

fragmented. Our method builds upon these reconstruction techniques by automatically deriving localized reconstruction regions directly from detected change clusters, enabling a unified pipeline for terrain change analysis and visualization.

3 METHODS

Since we want to find local changes in the point clouds, our approach consists of three steps. First, we preprocess the point clouds and reconstruct mesh surfaces. Next, we identify local regions of significant change between the point clouds, and in each region, we perform local volume estimation using the mesh surfaces. Throughout our method, we assume a pair of point clouds ordered by their time of creation from earliest to latest.

3.1 General Methodology

The proposed framework aims to detect localized geomorphological changes and estimate their associated volume displacement from temporally separated point cloud surveys. The pipeline combines point-based change detection with mesh-based volumetric analysis to provide both spatial and quantitative characterization of environmental change. An overview of the pipeline is presented below and visualized in Figure 1.

1. **Input preparation.** The method takes as input two temporally ordered point clouds, $C1$ and $C2$, together with their corresponding reconstructed mesh surfaces. Point clouds provide detailed geometric sampling of the environment, while meshes enable robust volumetric computations that are not directly possible on sparse point samples alone.

2. **Initial change detection.** We apply the M3C2 algorithm [LBL13] to estimate per-point distances and uncertainties between the two scans. Unlike direct Euclidean comparison methods, M3C2 measures displacement along local surface normals, making it more robust for irregular terrain and steep geomorphological structures such as cliffs and rocky slopes. This method is frequently used in terrain change analysis.

3. **Noise reduction and significance filtering.** Raw change estimates often contain noise caused by registration inaccuracies, varying point density, and surface roughness. To suppress insignificant changes, points below a threshold τ are removed. Remaining values are smoothed using neighborhood-based averaging to reduce local abrupt changes in positive and negative values, and prevent fragmentation of regions of coherent change.

4. **Localized change region extraction.** Smoothed change values are separated into positive and negative groups representing material accumulation and erosion. Each group is clustered independently using DBSCAN [EKSX96]. This step transforms sparse per-point displacement estimates into semantically meaningful regions of change, enabling more interpretable analysis of unstable terrain.

5. **Cluster characterization.** For each detected cluster, the spatial center and average uncertainty are computed. These descriptors provide a compact representation of the location and reliability of detected changes, supporting downstream interpretation and risk assessment.

6. **Volume estimation.** To quantify material displacement, each cluster is enclosed within an oriented bounding box derived using Principled Component Analysis (PCA) [Pea01]. The orientation adapts

to the local geometry, producing tighter spatial bounds than axis-aligned boxes. Corresponding mesh regions are clipped and converted into watertight meshes, enabling robust volume computation through surface integration. Thus we can derive differences between volumes of regions before and after deformation.

7. **Output generation.** The final outputs consist of cluster locations, associated OBBs, estimated volume changes, and M3C2 uncertainty measures. Together, these outputs provide both qualitative and quantitative insight into geomorphological change and support environmental monitoring and hazard assessment workflows.

3.2 Data Preprocessing

The number of points in the point cloud affects processing time in later stages, so for large point clouds, downsampling is necessary. In our experiments, we used CloudCompare's Octree downsampling with a tree cutoff at depth 10. The cutoff depth should ideally be selected according to the target geometric precision and the minimum feature size that must be preserved. At this stage, registration and point cloud alignment are also performed if necessary.

From the downsampled point clouds, we construct mesh surfaces, which are later used for volume computation. In our experiments, we used MeshLab's screened Poisson reconstruction [KH13] with depth 8, a minimum number of samples of 1.5, and an interpolation weight of 4, since that gave enough detail while still keeping the amount of geometry low enough for fast processing.

While the parameter values reported here were selected empirically for our experimental datasets, the parameters are not tied to a specific scene and can be adapted to new datasets based on geometric scale, point density, and expected change magnitude. In practice, most parameters scale with either the average point spacing or the minimum feature size that should remain detectable.

3.3 Change Detection

Our aim was to localize distinctive changes to smaller areas and compute volume changes in those areas. This approach allows for a more detailed analysis of volatile regions in the scanned environment. Change detection requires a pair of point clouds and a pair of corresponding mesh objects. We refer to the first scan in the pair as $C1$ and the second scan as $C2$, indicating their temporal sequence.

For rough change detection, we ran the M3C2 algorithm on the point clouds, storing change distances and uncertainties for each pair onto the points of point cloud $C1$. For each point, M3C2 constructs local surfaces of both point clouds by considering points within a cylinder of radius r, centered on the current point and oriented along the point's normal vector. It records the change distance as the distance between the surfaces and computes the uncertainty from the variance of points above and below the estimated surface. The distance values can be positive or negative depending on the type of change: accumulated or eroded material.

The M3C2 neighborhood radii r should be selected relative to the local point spacing and expected surface roughness. The smaller normal radius should exceed the average nearest-neighbor spacing to reduce noise, while the larger radius should approximate the scale of local geometric structures used for stable normal estimation. Similarly, the projection cylinder radius should be large enough to contain a sufficient number of points for reliable local surface estimation, but small enough to avoid blending distinct nearby structures. For denser point clouds, smaller radii are generally sufficient, whereas sparse or noisy datasets require larger neighborhoods for stable estimates. In our experiments, we set the cylinder radius r to 2.0 to capture enough neighboring points for each point. For computation of point normals, neighborhood radii of 0.1 and 1.0 were used. The maximum cylinder height was set to 10.0.

Points with insignificant M3C2 distances are filtered out using threshold τ, reducing noise and computation. Larger τ values suppress smaller changes, while lower values increase sensitivity but introduce noise. τ should reflect the minimum change magnitude considered meaningful in the target application. In practice, τ should typically exceed the expected registration and measurement noise level to suppress false detections. We empirically selected $\tau = 0.05$ (in our testing datasets, this corresponds to a cap of 5 cm) to balance detail and robustness.

Due to noise in data and unpredictable flow of the underlying shapes that the point clouds try to model, the M3C2 distances can vary between neighbouring points, sometimes also between positive and negative values. These variations can cause redundant fragmentation of clusters, making the analysis unnecessarily complex. We thus smooth out the values across points with Gaussian-like smoothing method, taking the neighbourhood of n points at each point and combining their weighted M3C2 distances. Given a standard deviation σ, M3C2 distance v_i and Euclidean distance d_i at point i, the equation

$$v_{\text{smooth}} = \frac{1}{2\pi\sigma^2} \sum_{i=1}^{n} v_i e^{-\frac{d_i^2}{2\sigma^2}} \tag{1}$$

gives a smoothed M3C2 distance value for an arbitrary point. In our experiments, we used $n = 30$ and $\sigma = 1.0$, which we determined based on the average point neighborhood density of our test clouds.

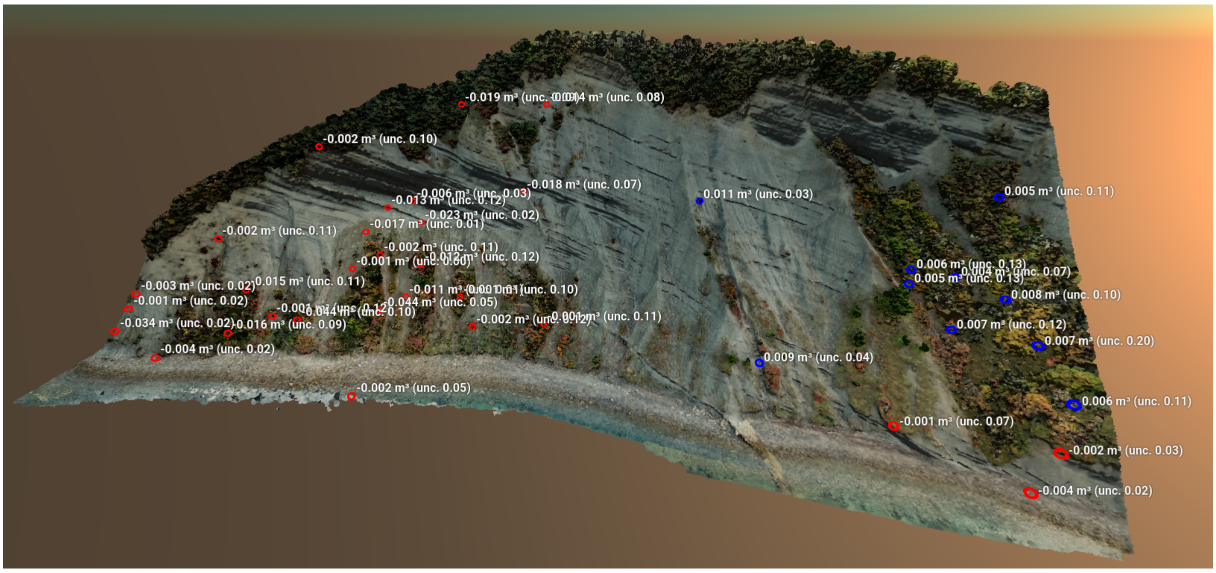

Figure 2: Visualization of interesting changes in the Slovenian cliff scan FG-20-11 from [KKT25]. Red circles mark the material falloff, and the blue circles mark the material accumulated. The numbers show estimated volume of change and the uncertainty.

We use these smoothed values to separate the points by positive and negative distances. This allows us to distinguish between different types of changes: accumulation and erosion. The two groups of points are then clustered separately using the DBSCAN algorithm. The resulting clusters represent regions of significant change of a specific type. DBSCAN parameters should be selected according to the spatial density of changed regions. The neighborhood radius ε should typically be slightly larger than the average point spacing within change regions, while the minimum sample count determines the smallest cluster considered significant. Increasing these parameters suppresses fragmented detections caused by noise but may merge nearby changes. We set the parameters $\varepsilon = 1.5$ and the minimum number of neighborhood samples to 10, which prevented over-fragmentation in our dataset.

We find the center of each cluster and compute its average uncertainty. We then perform volume estimation for each cluster.

3.4 Volume Estimation

We surround each cluster of points with an Axis-Aligned Bounding Box (AABB). We perform PCA [Pea01] on the cluster to extract directions of significant change. The extracted eigenvectors are combined into a transformation matrix, which we convert into a rotation matrix by normalizing the column vectors. We use this matrix to rotate the AABB and obtain the OBB.

In the following steps, a pair of mesh objects is used. These refer to the reconstructed surfaces of the point clouds mentioned in section 3.2. The mesh objects are aligned with the point clouds, so OBBs can be used for their processing without additional transformations. For operations on meshes, we use the Python library `trimesh`[3].

We perform mesh clipping for each mesh individually: we extract a patch inside the OBB to obtain a subset of the mesh. We then extrude the mesh surface downward along the Z-axis to create a watertight mesh. We calculate the volume of the watertight mesh using `trimesh` processing, which applies the surface integral method. After obtaining the volumes of both meshes, we subtract them to determine the volume difference between clusters of the two scans. The process is shown in fig. 3.

Volume change is mapped to clusters. The final algorithm outputs include cluster centers, their OBBs, volume changes, and M3C2 uncertainties.

3.5 Change Visualization

We developed a web-based visualization tool using Three.js[4] framework to observe computed regions of significant change. The tool displays the reconstructed mesh surface and overlays markers at the centers of clusters. Each marker shows the estimated volume change in that region and its overall uncertainty. We use colors to distinguish positive and negative changes. The visualization is visibile in fig. 2.

4 RESULTS

We conducted experiments to test our approach on artificially created point clouds and a set of scans of real

[3] `https://trimesh.org/index.html`
[4] `https://threejs.org/`

Scan Pair	Processing Time [s]	Amount of Clusters	Average Volume Change [m^3]	Average Uncertainty
FG-20-11 – FG-21-03	703	63	−3,955	0.07
FG-21-03 – FG-22-03	570	54	8,039	0.05
FG-22-03 – FG-22-12	533	59	−5,118	0.05
FG-22-12 – FG-23-04	744	76	52,740	0.08
FG-23-04 – FG-23-09	402	35	2,834	0.04

Table 1: The results of the proposed method on the Slovenian cliff dataset. We show results for each pair of point clouds. The pairs were constructed from consecutive scans, ordered by their capture date.

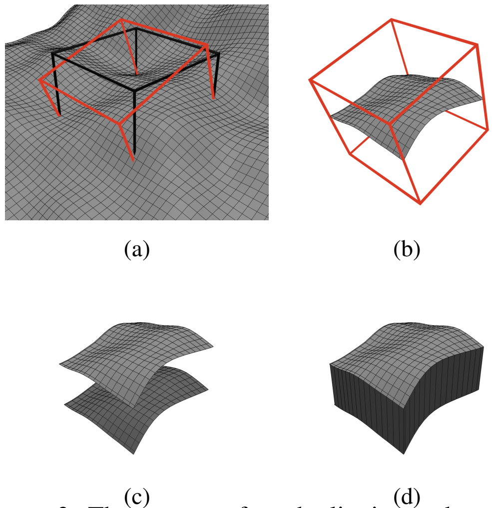

Figure 3: The process of mesh clipping and extrusion: (a) shows the entire mesh, the AABB in black, and the OBB in red. (b) shows the mesh clipped to the region inside the OBB. (c) and (d) show the extrusion process to create a watertight mesh by duplicating the surface and connecting its boundary vertices to those of the duplicate.

terrain. The synthetic tests allowed us to assess how our method performs in controlled environments, while the real dataset enabled us to observe how the approach handles complex and noisy data. However, our real datasets do not contain clear ground truth volumes, and to the best of our knowledge, no methods for precise local volume estimation on point clouds are publicly available.

4.1 Synthetic Tests

We constructed simplified mesh pairs with controlled changes between them. All tests were performed on a grid plane that was deformed in various ways (we refer to the first element in the pair as the $C1$ plane and the second as the $C2$ plane):

- In the Hill Test, the $C1$ mesh was left unchanged, while in the $C2$ mesh a small bump was pulled upward from the plane geometry.
- In the Cube Test, a cube-like shape was extruded from the $C2$ plane.
- In the Spheres Test, various spherical displacements were applied throughout the $C2$ plane.
- In the Cubes Test, cube-shaped displacements were applied to both the $C1$ and $C2$ planes. This test combines positive and negative changes.

To generate the point clouds, we placed points onto mesh faces using a Poisson distribution. The number of points in the resulting point clouds can be found in table 2.

The resulting computed volumes are shown in table 3. Most results returned one marker per test, while the last test returned six markers at the centers of the changes. In this case, the table presents the sum of all volumes detected at each marker. We also observe that the last test performs worse than the others (10% error), which may be attributed to conservative mesh clipping at different markers. In the other tests, the bounding box for mesh clipping includes most of the mesh.

Test Name	Amount of Points	
	Before	After
Hill Test	3,161	6,995
Cube Test	6,443	6,735
Spheres Test	75,833	83,272
Cubes Test	70,955	70,955

Table 2: Details about test point clouds: We present the number of points in the first and second point clouds for each test pair.

4.2 Testing on Real Data

We tested our method on real-life UAV scans of coastal cliffs. This dataset presents several challenges that we aimed to address with our method: steep terrain, subtle changes in rock formations, and rockfalls. We used a subset of the Strunjan cliffs dataset, captured along the Slovenian coast over a 5-year period [KKT25]. Our subset covered the years 2020 to 2023 and included 6 scans. On average, the point clouds contain 120 million points, capturing both rocky slopes and vegetation. The scans were taken in different seasons, resulting in large variations over time.

We preprocessed the dataset by reconstructing mesh geometry using screened Poisson reconstruction [KH13],

Test Name	Ground Truth Volume [m^3]	Computed Volume [m^3]
Hill Test	0.62	0.62
Cube Test	0.82	0.82
Spheres Test	18.66	18.66
Cubes Test	40.00	35.89

Table 3: Tests for volume approximation on synthetic examples were constructed using Blender 3D software, with the `3D Print Toolbox` extension used to compute ground truth mesh volumes. The final test, which involved more complex changes, proved challenging for the pipeline. In some cases, the bounding box did not cover all of the relevant geometry.

which was robust enough for testing purposes. We found that the Ball Pivoting algorithm [BMR+99] produced better and more detailed geometry, but it also contained many holes that would hinder our volume computation, so we decided against using it.

The average processing time of our method on individual point clouds in the dataset was approximately 590 seconds. We also performed an analysis of the results, which are shown in table 1. The processing time is heavily affected by the number of detected clusters, as volume estimation is performed for each cluster. Average uncertainty remains relatively low for all scan pairs, indicating high confidence in change detection.

Evaluation on the real dataset was primarily qualitative in nature. To our knowledge, no publicly available benchmark dataset currently exists that contains validated ground-truth volumetric measurements for cliff erosion and localized rockfall events at this scale and level of detail. Consequently, we could not perform direct quantitative validation of the estimated volume changes against reference measurements. Instead, the evaluation relied on visual verification of the detected changes and assessment of whether the proposed workflow successfully identified morphologically plausible events.

More specifically, manual observation consisted of inspecting the detected clusters directly within the point clouds and reconstructed meshes, with particular focus on visually prominent rockfall regions, cliff retreat areas, and localized morphological changes visible across temporal scan pairs. We verified whether these visually identifiable changes were consistently detected by the method and whether the estimated change regions corresponded to plausible terrain modifications. This process was performed for all scan pairs included in the dataset.

On one hand, interesting areas of rockfall were detected, as seen in fig. 4. Based on the overall scale of the point cloud in real life (approximately 300 meters), the values of displaced volume seem plausible. However,

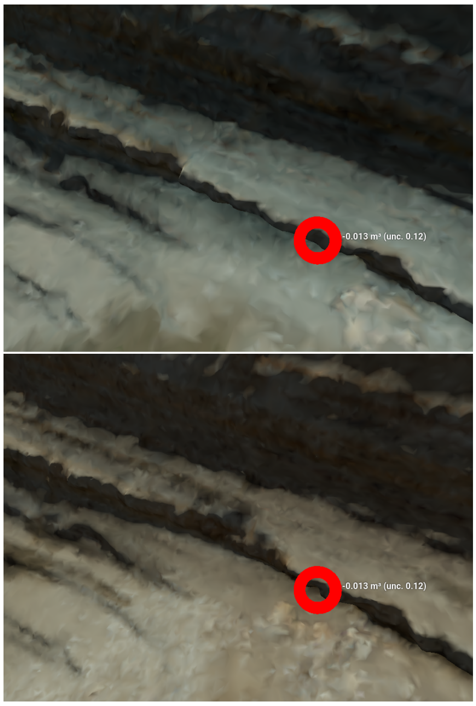

Figure 4: A comparison of two consecutive scans, FG-20-11 and FG-21-03, shows that the change is detected in a region of rockfall. The marker is slightly offset because the rockfall change is merged with smaller changes in the vicinity.

there are also areas where the estimated change does not appear to make sense. This may be due to noisy results from M3C2, which are further emphasized by unstable volume computation. The latter seems more prevalent in the scan pair "FG-22-12 – FG-23-04" due to vegetation. Many of the changes are marked on the vegetation portion of the terrain, which may not be suitable for rockfall analysis. However, existing popular methods for vegetation segmentation and removal, like SegmentAnyTree [WPX+24] or CloudCompare CANUPO plugin [BL12], perform poorly on steep terrain. Similarly problematic is primitive filtering based on geometric or visual features. Developing a robust vegetation filtering method is thus out of scope for this paper.

Overall, the real-data experiments demonstrate that the proposed workflow is capable of identifying visually plausible cliff changes and estimating corresponding displaced volumes in challenging large-scale UAV datasets. Nevertheless, the evaluation remains primarily qualitative, and further validation on datasets containing independent ground-truth measurements would be necessary to rigorously assess the quantitative accuracy and robustness of the proposed approach.

5 DISCUSSION AND FUTURE WORK

We observe that in simple scenarios with straightforward displacements, our method provides good results.

However, it begins to lose accuracy in more complex situations involving multiple displacements in different areas on the surface, of mixed types (accumulation and erosion), as seen in the Cubes Test in table 3. Manual inspection of cliff regions showed that observable differences in the terrain are pronounced and that the volume changes appear plausible, though further analysis of the extent of unnoticed changes may be necessary.

Change detection primarily depends on selecting appropriate factors for M3C2 computation, threshold values, and clustering parameters. Currently, these values must be set according to the dataset's size and level of detail, particularly the M3C2 cylinder width, which determines how much fine detail is included in the M3C2 distance computation. We note this because our method required special adjustment of the cylinder width to detect changes in steep elevation. Automatically adapting pipeline parameters to the dataset would be necessary to fully automate the process and avoid poor results due to superficial parameter tuning.

The unusually large average volume change of 52,740 reported in Table 1 reflects the highly skewed distribution of the detected events rather than a consistent trend across all observations. Specifically, a small number of extreme vegetation regions with very large volume changes substantially increased the arithmetic mean, while the majority of events exhibited comparatively moderate variations. This indicates that the vegetation and noise remain major challenges in the real-world scans.

The AABB estimation directly affects the calculation of volume change. In some cases, the AABB may capture only part of the desired region, potentially excluding significant portions of displaced volume. Thus, constructing OBB directly from eigenvectors may be more suitable in the future and reduce the overall computational burden. Because we compute the bounding box to fit the cluster of only one point cloud, we do not account for possible deformations in the other, so the AABB may not fit one mesh at all. We also perform mesh extrusion along a specific axis, while it would likely be better to extrude along the third PCA eigenvector. The nature of extrusion also raises the question of crevices and overhangs, where intersections with existing geometry may occur. Depending on meshing performance and cloud density, gaps and holes may occur. In the future, we may incorporate hole-filling methods into the pipeline, such as [BSKM20].

Another important limitation of the current work is the absence of quantitative ground-truth validation on real terrain data. To our knowledge, no publicly available benchmark dataset currently exists that provides validated volumetric measurements for localized cliff erosion and rockfall events in long-term monitoring scenarios. Consequently, the real-data evaluation remains primarily qualitative and relies on visual verification of prominent terrain changes. While the synthetic experiments provide partial controlled validation, further collaboration with geomorphology experts and acquisition of independently measured reference data would be necessary for rigorous large-scale evaluation.

Finally, the current visualization focuses primarily on presenting detected changes and estimated displaced volumes in an interactive and computationally efficient manner. More advanced visualization techniques for uncertainty representation, confidence communication, and displacement interpretation could substantially improve usability for domain experts. However, uncertainty visualization in this context is itself a challenging problem due to the highly localized and spatially heterogeneous nature of the uncertainty distributions. We therefore consider advanced uncertainty-aware visualization an important but separate research direction beyond the scope of the present work.

6 CONCLUSION

In this paper, we presented a method for detecting and quantifying volume changes within pairs of high-resolution point clouds, testing the workflow on both artificially generated scenarios and real-life UAV datasets. Our findings demonstrate that while the method remains somewhat volatile and highly dependent on the specific characteristics of the input data, it provides a solid foundation for advanced terrain analysis and geomorphological monitoring. By bridging the gap between raw point cloud comparison and localized volumetric measurement, the proposed approach allows for a more nuanced understanding of how material is displaced in volatile environments like coastal cliffs.

We believe that further improvements to the pipeline, particularly regarding the automation of parameter selection and the refinement of mesh extrusion techniques, could significantly stabilize the approach and accelerate the development of real-time systems for land risk assessments. Future work will focus on adapting the M3C2 cylinder width and clustering thresholds dynamically to suit varying levels of terrain detail, thereby reducing the need for manual tuning and minimizing the impact of measurement noise. Also, in this contribution we focus on rockfall and erosion, but in the future, analysis and adaptation to other natural processes like landslides and earthquake damage would have to be conducted. Ultimately, this research contributes to the broader goal of disaster prevention by providing experts with clearer analytical tools to identify, quantify, and respond to potential hazards in unstable landscapes.

7 ACKNOWLEDGMENTS

The research was conducted within the project Geospatial Information Technologies for Resilient and Sustainable Society (GeoAI) (GC-0006), funded by the Slovenian Research and Innovation Agency (ARIS).

We also thank the authors of [KKT25], the Faculty of Civil and Geodetic Engineering at the University of Ljubljana, and the Institute of Anthropological and Spatial Studies ZRC SAZU for sharing with us the Slovenian cliff dataset.

8 REFERENCES

[AQL24] Hessah Albanwan, Rongjun Qin, and Jung-Kuan Liu. Remote Sensing-Based 3D Assessment of Landslides: A Review of the Data, Methods, and Applications. *Remote Sensing*, 16(3), 2024.

[AWM+21] Katharina Anders, Lukas Winiwarter, Hubert Mara, Roderik Lindenbergh, Sander E. Vos, and Bernhard Höfle. Fully Automatic Spatiotemporal Segmentation of 3D LiDAR Time Series for the Extraction of Natural Surface Changes. *ISPRS Journal of Photogrammetry and Remote Sensing*, 173:297–308, 2021.

[BDH19] David Bonneau, Paul-Mark DiFrancesco, and D. Jean Hutchinson. Surface Reconstruction for Three-Dimensional Rockfall Volumetric Analysis. *ISPRS International Journal of Geo-Information*, 8(12), 2019.

[BL12] N. Brodu and D. Lague. 3d terrestrial lidar data classification of complex natural scenes using a multi-scale dimensionality criterion: Applications in geomorphology. *ISPRS Journal of Photogrammetry and Remote Sensing*, 68:121–134, 2012.

[BMR+99] Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. The Ball-Pivoting Algorithm for Surface Reconstruction. *IEEE transactions on visualization and computer graphics*, 5(4):349–359, 1999.

[BSKM20] Ciril Bohak, Matej Slemenik, Jaka Kordež, and Matija Marolt. Aerial lidar data augmentation for direct point-cloud visualisation. *Sensors*, 20(7):1–17, 2020.

[DBH20] Paul-Mark DiFrancesco, David Bonneau, and D. Jean Hutchinson. The Implications of M3C2 Projection Diameter on 3D Semi-Automated Rockfall Extraction from Sequential Terrestrial Laser Scanning Point Clouds. *Remote Sensing*, 12(11), 2020.

[EKSX96] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, page 226–231. AAAI Press, 1996.

[FGTG25] Ioannis Farmakis, Davide Ettore Guccione, Klaus Thoeni, and Anna Giacomini. VoxFall: Non-parametric Volumetric Change Detection for Rockfalls. *Engineering Geology*, 352:108045, 2025.

[GZW20] Zan Gojcic, Caifa Zhou, and Andreas Wieser. F2S3: Robustified Determination of 3D Displacement Vector Fields Using Deep Learning. *Journal of Applied Geodesy*, 14(2):177–189, 2020.

[HHH+24] Jian Huang, Xiang Huang, Tristram C. Hales, Nengpan Ju, and Zicheng He. Volume estimation for high-locality fragmented rockfall using uav-based photogrammetry. *Natural Hazards*, 121(6):7347–7364, December 2024.

[HKA+26] Bernhard Höfle, Dominic Kempf, Katharina Anders, Ronald Tabernig, Thomas Isensee, Xiaoyu Huang, and Hannah Weiser. py4dgeo: Library for Change Analysis in 4D Point Clouds., February 2026.

[JdM+25] Shoujun Jia, Lotte de Vugt, Andreas Mayr, Chun Liu, and Martin Rutzinger. Location and Orientation United Graph Comparison for Topographic Point Cloud Change Estimation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 219:52–70, 2025.

[JdM+26] Shoujun Jia, Lotte de Vugt, Andreas Mayr, Katharina Anders, Chun Liu, and Martin Rutzinger. Change Tensor: Estimating Complex Topographic Changes From Point Clouds Using Riemann Manifold Surfaces. *ISPRS Journal of Photogrammetry and Remote Sensing*, 232:766–786, 2026.

[KH13] Michael Kazhdan and Hugues Hoppe. Screened Poisson Surface Reconstruction. *ACM Transactions on Graphics (ToG)*, 32(3):1–13, 2013.

[KKT25] Klemen Kregar and Klemen Kozmus Trajkovski. Combining UAV Photogrammetry and TLS for Change Detection on Slovenian Coastal Cliffs. *Drones*, 9(4), 2025.

[LBL13] Dimitri Lague, Nicolas Brodu, and Jérôme Leroux. Accurate 3D Comparison of

Complex Topography with Terrestrial Laser Scanner: Application to the Rangitikei Canyon (N-Z). *ISPRS Journal of Photogrammetry and Remote Sensing*, 82:10–26, 2013.

[MCZP24] S. Mineo, D. Caliò, G. Zocco, and G. Pappalardo. Implementing Close-range Remote Surveys for the Digitally Supported Rock Mass Stability Analysis. *Engineering Geology*, 328:107382, 2024.

[Pea01] Karl Pearson. LIII. On Lines and Planes of Closest Fit to Systems of Points in Space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.

[WAH21] Lukas Winiwarter, Katharina Anders, and Bernhard Höfle. M3C2-EP: Pushing the Limits of 3D Topographic Point Cloud Change Detection by Error Propagation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 178:240–258, 2021.

[WAW+21] Jack G. Williams, Katharina Anders, Lukas Winiwarter, Vivien Zahs, and Bernhard Höfle. Multi-directional Change Detection Between Point Clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 172:95–113, 2021.

[WBH+26] Zhaoyi Wang, Jemil Avers Butt, Shengyu Huang, Tomislav Medić, and Andreas Wieser. Dense 3D Displacement Estimation for Landslide Monitoring via Fusion of TLS Point Clouds and Embedded RGB Images. *International Journal of Applied Earth Observation and Geoinformation*, 146:105093, 2026.

[WPX+24] Maciej Wielgosz, Stefano Puliti, Binbin Xiang, Konrad Schindler, and Rasmus Astrup. Segmentanytree: A sensor and platform agnostic deep learning model for tree segmentation using laser scanning data. *Remote Sensing of Environment*, 313:114367, 2024.

[WW21] Luke Weidner and Gabriel Walton. Monitoring the Effects of Slope Hazard Mitigation and Weather on Rockfall along a Colorado Highway Using Terrestrial Laser Scanning. *Remote Sensing*, 13(22), 2021.

[WW23] Gabriel Walton and Luke Weidner. Accuracy of Rockfall Volume Reconstruction from Point Cloud Data—Evaluating the Influences of Data Quality and Filtering. *Remote Sensing*, 15(1), 2023.

[YH25] Y. Yang and C. Holst. How to Find Geometric Changes in Laser Scanning Point Clouds? A Perspective on Correspondence Definitions. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-G-2025:1003–1010, 2025.

[YS23] Yihui Yang and Volker Schwieger. Patch-based M3C2: Towards Lower-uncertainty and Higher-resolution Deformation Analysis of 3D Point Clouds. *International Journal of Applied Earth Observation and Geoinformation*, 125:103535, 2023.

[ŽBK23] Barbara Žabota, Frédéric Berger, and Milan Kobal. The Potential of UAV-Acquired Photogrammetric and LiDAR-Point Clouds for Obtaining Rock Dimensions as Input Parameters for Modeling Rockfall Runout Zones. *Drones*, 7(2), 2023.