Time-Adaptive Change Analysis through Extension of the M3C2 Algorithm using Multi-Modal Laser Scanning Data in a Salt Marsh Environment

Xiaoyu Huang[1], Mathilde Letard[1], Dimitri Lague[2,3], Bernhard Höfle[4,5], Ronald Tabernig[4,5], Paul Leroy[2], Katharina Anders[1]

[1] Remote Sensing Applications, TUM School of Engineering and Design, Technical University of Munich, Ottobrunn, Germany - (xiaoyu.huang, mathilde.letard, k.anders)@tum.de
[2] Univ Rennes, Plateforme LiDAR, OSEREn, UAR 3343 CNRS, France - (dimitri.lague, paul.leroy)@univ-rennes.fr
[3] Univ Rennes, Géosciences Rennes, UMR 6118 CNRS, France
[4] 3DGeo Research Group, Institute of Geography, Heidelberg University, Heidelberg, Germany - (hoefle, ronald.tabernig)@uni-heidelberg.de
[5] Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University, Heidelberg, Germany

Keywords: Multitemporal, LiDAR, Time series, Change detection, Spatiotemporal, 4D point clouds

Abstract

Quantifying topographic dynamics from 3D time series is essential for a broad range of geoscientific applications. However, data acquired by laser scanning often vary between epochs, e.g. in point density and coverage. These heterogeneities present a challenge for change detection, particularly across multi-temporal and multi-modal data. We propose a new approach that enhances the analysis of spatiotemporally irregular point clouds, particularly when the sampling frequency exceeds the analysis time scale. In particular, our method extends the Multiscale Model to Model Cloud Comparison (M3C2) algorithm through adaptive temporal aggregation. Driven by a local point density requirement, our method employs both a spatial neighborhood (as in standard M3C2) and a temporal neighborhood for each core point within the point cloud time series. If the neighborhood of a core point is too sparse for surface estimation, an iterative temporal search incorporates data from adjacent epochs until either the density requirement or a maximum temporal window is reached. This adaptive process ensures sufficient local point density while preventing globally fixed temporal aggregation. We evaluate our method on a multi-modal laser scanning dataset from Mont-Saint-Michel Bay, France, comprising 38 epochs spanning a decade at daily to seasonal intervals. Results demonstrate that our method increases change detection completeness by more than 13% compared to standard M3C2 and increases accuracy by 31% compared to commonly used fixed-window averaging. Our approach thereby enhances 3D change detection in complex real-world 4D datasets, enabling higher accuracy and completeness in the analysis of surface dynamics at variable spatiotemporal scales.

1. Introduction

Monitoring Earth surface dynamics at high spatiotemporal resolution through 3D point cloud time series is essential for understanding geomorphic processes, particularly in highly dynamic environments like tidal salt marshes. Salt marshes are critical coastal habitats that provide valuable ecosystem services, including shoreline protection from storms, carbon sequestration, and water quality enhancement (Pinton et al., 2021; Stark et al., 2015).

The advent of 4D point clouds (3D + time) has significantly enhanced the ability to monitor surface dynamics in such environments, providing detailed information on the duration, location, and extent of geomorphic processes (Lindenbergh et al., 2025; Eitel et al., 2016). Technologies like Terrestrial Laser Scanning (TLS) and Unoccupied Aerial Vehicle (UAV)-based surveys are revolutionary for topographic monitoring, capturing high-resolution 3D data over inaccessible areas (Guisado-Pintado et al., 2019). However, the analysis of the resulting point cloud time series presents challenges, particularly in the methodology for change detection and quantification (Qin et al., 2016). The evolution of 3D change detection has moved from simple DEM of Difference (DoD) methods, which are limited to 2.5D analysis and cannot handle complex topography, to direct point cloud comparisons. Early approaches like Cloud-to-Cloud (C2C) (Girardeau-Montaut et al., 2005) or Cloud-to-Mesh (C2M) (Cignoni et al., 1998) comparisons are highly sensitive to variations in point density and noise. To address this, the Multiscale Model to Model Cloud Comparison (M3C2) (Lague et al., 2013) algorithm has become a well-established standard for bi-temporal analyses of surface change. By calculating distances along the local surface normal and averaging points within a cylindrical neighborhood, M3C2 provides robust change measurements on heterogeneous spatially sampled data and additionally provides a confidence interval (known as the Level of Detection, LOD). M3C2-EP further refines this by propagating sensor and alignment uncertainties to achieve a more physically-based and rigorous LOD (Winiwarter et al., 2021). Despite these advances, M3C2 and its variants remain fundamentally bi-temporal, quantifying the net change between two discrete snapshots in time, which aggregates all underlying processes and makes it difficult to infer how a surface evolved (Kharroubi et al., 2022). To overcome the limitations of bi-temporal comparisons and leverage full time series data for analyzing dynamic processes, methods such as M3C2 have been incorporated into multi-epoch analyses (Anders et al., 2021; Kuschnerus et al., 2021). However, these time-series-based approaches typically assume uniform, gap-free temporal sampling with consistent spatial density across epochs, which are conditions rarely met in environmental monitoring scenarios.

Data acquired in dynamic natural scenes are commonly spatially and temporally irregular (Zahs et al., 2022). Spatial irregularity is inherent due to the unstructured nature of point clouds and can be amplified by varying sensor positions, (changing) occlusions, or variable meteorological conditions in near-continuous acquisitions (Chen and Yang, 2016). In long-term monitoring, this is

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-2-2026-623-2026

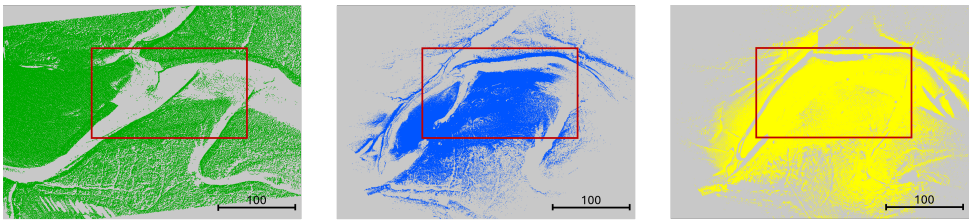

a) Epoch 1: Long-Range TLS
Timestamp: 2009-07-09

b) Epoch 2: TLS
Timestamp: 2010-10-07

c) Epoch 3: TLS
Timestamp: 2011-09-23

Figure 1. Variations in point cloud characteristics across different TLS sensors and scene conditions in different epochs for data acquired at Mont-Saint-Michel Bay, France. (a) Wide coverage with moderate, uniform point density. (b) A high-density, detailed scan with limited coverage. (c) Another high-density scan, concentrated in a different area, illustrates how data coverage varies between epochs. All subgraphs cover the same spatial extent, with a red bordered box highlighting the area of interest.

compounded by temporal irregularity, where regular sampling is rarely given due to varying survey schedules (for logistical reasons but also targeting different process time scales), temporary occlusions, adverse weather conditions, hardware limitations, or organizational and economic constraints. This spatiotemporal irregularity becomes critical when analyzing time series of scenes with a long monitoring history, as the data properties of epochs can vary strongly. When point clouds are locally sparse, the spatial resolution of detected changes is limited and the analysis coverage is reduced, because surface-based methods like M3C2 require a sufficient local point density to estimate a normal vector, to average points and calculate a robust mean distance, and to assess the uncertainty. Consequently, actual surface changes in sparsely sampled areas may be missed entirely or lead to an overestimation of the minimum detectable change and a failure to quantify certain parts of significant surface movements.

In situations where the temporal sampling frequency is higher than the required analysis scale, temporal redundancy can be leveraged to enhance the data basis for more comprehensive change detection and quantification (Kromer et al., 2015). Several methods have been developed to leverage the increased information content of 4D point cloud data in the temporal domain. One approach to mitigate data sparsity is temporal averaging when there is redundancy in temporal sampling compared to the analysis time scale. Kromer et al. (2015) have demonstrated how a moving window applied to filter distance values across both spatial and temporal neighbors reduces noise in topographic models. This fixed-window approach, where the size of the spatial and temporal neighborhood is globally defined, has become a standard preprocessing step in the analysis of 3D time series, commonly applied as a temporal smoothing step separate from the spatial interpolation involved in the M3C2 algorithm (Eltner et al., 2017; Anders et al., 2021). A recent approach by Tabernig et al. (2025) involves aggregating multiple consecutive point clouds to create a single, denser representation, which improves the geometric coverage of object-based change analysis of a vegetated landslide captured by permanent laser scanning. All these approaches require careful optimization of the temporal neighborhood size to avoid excessive smoothing of the (short-term) change while aggregating adequately for noise reduction or object representation. For dense 3D time series, another strategy is to use Kalman filtering to extend bi-temporal M3C2 analysis into the full spatiotemporal domain, treating the change time series at each point as a dynamic system (Winiwarter et al., 2022). This provides a smoothed best-estimate time series and can interpolate missing data, but standard Kalman filters assume linear process models and are not suited for accurately modeling abrupt, large-magnitude events.

To address these limitations, we propose a time-adaptive aggregation method that automatically integrates point cloud data considering multiple neighboring epochs when the points of a single epoch are too sparse for surface-based change estimation. This adaptive temporal aggregation continuously extends in time until it reaches a required local point number threshold, which is set based on the data density requirement of the chosen change analysis method (here, M3C2). This aggregation process is constrained by a maximum allowed temporal window that is automatically determined according to the analyzed time scale, so that signal mixing between compared epochs of the change analysis is avoided. The adaptive aggregation adjusts the temporal neighborhood for each local area based on its density, leaving areas with sufficiently dense sampling unaltered. This approach improves local point density for the M3C2 plane fit and distance quantification by leveraging (even irregular) temporal data availability.

We evaluate the effectiveness of our time-adaptive method by analyzing the quality of the aggregated point clouds in terms of their roughness (i.e., standard deviation of point distances to the local plane, thus representing plane fit quality) and the accuracy of the measured distances between epochs compared to original point cloud data left out in the experiments. The method is applied to a series of laser scanning point clouds collected from the Mont-Saint-Michel (MSM) Bay in France, comprising 38 epochs acquired using different TLS sensors between 2009 and 2017, with irregular temporal intervals ranging from several days to several months. Fig. 1 illustrates the inherent complexities within this dataset, highlighting significant differences in spatial coverage, point density, and sensor type across example epochs. This makes the application particularly suitable to investigate the effectiveness of our approach in enhancing topographic dynamics analysis for datasets with such temporal and spatial heterogeneity.

2. Study area and dataset

Our study focuses on a tidal channel within the mega-tidal environment of the Mont-Saint-Michel Bay on the north-western coast of France (Fig. 2a). The tidal regime is semi-diurnal, with a range that can reach up to 15 m (Tessier, 1993). The study site is an actively migrating meander bend located approximately 700 m southeast of the Mont-Saint-Michel Abbey, which drains a 2.2 km² area of salt marshes (Leroux, 2013). As depicted in Fig. 2b, this bend is characterized by a clear distinction between a depositional inner bar and an erosional outer bank. The inner bar’s pioneer vegetation, *Suaeda maritima*, is a key agent driving sediment accretion and the system’s overall morphodynamics (Leroux, 2013).

The dataset used in this study is a time series of 38 epochs of

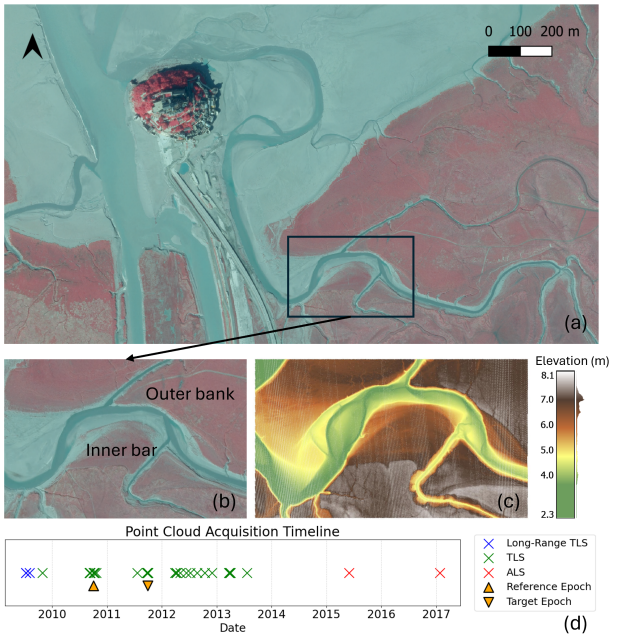

Figure 2. Study area in Mont-Saint-Michel Bay, France. (a) False-color composite (NIR-Red-Green) orthophoto, with the study area outlined by the black rectangle. (b) Detailed view of the study area. (c) Topography visualized through an ALS point cloud. (d) Timeline of the 38 point cloud acquisitions with TLS and ALS from 2009 to 2017. Data source for (a) and (b): (Ministère de la Transition Écologique et al., 2014).

3D point clouds acquired between 2009 and 2017. The dataset is characterized by a highly variable temporal sampling, with acquisition intervals ranging from daily during intensive monitoring campaigns to annually (Fig. 2d). It is a relatively heterogeneous collection of data from various sensors and platforms, including single-return ground-based TLS, long-range TLS (with range up to 2,000 m), and Airborne Laser Scanning (ALS). Detailed specifications of the sensors and acquisition parameters for the original surveys are documented in (Leroux, 2013). All epochs were co-registered with an estimated registration error of approximately 3 mm. The dense point clouds, with a point density ranging from 1 to 10 pts/cm², serve as the reference dataset for our evaluation. To validate the performance of our proposed method on sparser data, this reference dataset was subsequently downsampled using a voxel-based method with a voxel size of 0.2 m, creating the test data for our experiments.

3. Methods

3.1 M3C2 algorithm for 3D change quantification

Our method is an extension of the state-of-the-art M3C2 algorithm by Lague et al. (2013), a widely used method for robustly comparing multi-temporal point clouds to detect and quantify changes in 3D topography. For a set of two point cloud epochs, the algorithm first estimates a local surface normal ($\vec{N}$) for each core point (a defined sampling of coordinates) on the reference cloud (i.e., the temporally preceding epoch). This is achieved by fitting a plane to spatially neighboring points within a defined normal vector radius (R). This ensures that the direction of change is measured perpendicular to the local surface, based on the general assumption of local planarity (at a certain analysis scale). Since estimating a plane requires at least three points, the normal estimation is not possible in sparse areas of the point cloud where this condition is not met. Fitting a plane becomes typically more robust when the point density is even higher than the required minimum (i.e., denser local scene sampling), as measurement noise of individual points is averaged out. Consequently, where the normal vector cannot be estimated for the reference epoch, the algorithm returns a no data value (NaN), as no distance calculation can be performed. In the second step, the algorithm calculates the distance between the core point location of the reference (t_n) and target (t_m) clouds along this normal (Fig. 3b). It projects a cylinder of radius r that determines all original points from both clouds that fall within it. These points are then projected onto the axis of the normal vector estimated from the reference cloud, and their positions are averaged for each cloud. The final M3C2 distance is the distance between these two mean positions (i_n, i_m).
Finally, a local confidence interval associated to the quantified change is computed to distinguish statistically significant topographic changes from noise. However, the reliability of distance measurements and their confidence levels both depend on sufficient local point density across the two comparison epochs. Where this condition is not met, it can be potentially remedied through the temporal neighborhood, as proposed in this research.

3.2 Neighborhood and temporal window definition

Building upon the principles of the M3C2 algorithm, our proposed method requires a user-defined temporal scale as the basis for the bi-temporal change analysis (Fig. 3a). This requires the selection of a reference epoch and a target epoch to conduct change analysis between two states of the topography at certain points in time. For the experiments presented in this study, the TLS epoch from 2010-10-04 is selected as the reference epoch, and the TLS epoch from 2011-09-28 as the target epoch. We choose these epochs because they are both situated within periods of dense temporal sampling, providing a rich set of adjacent data for the analysis of our adaptive aggregation approach (Fig. 3c). Within the maximum temporal window, 9 additional epochs are available for the reference and 7 for the target epoch. The change detection is computed on a set of core points sampled from the reference epoch using a voxel grid filter with a size of 1.1 m.
The next step is to define the spatiotemporal neighborhood around each core point from which temporal information will be aggregated. As illustrated in Fig. 3d, this neighborhood is defined as a cylinder. The spatial extent is controlled by a radius parameter, set to 0.5 m in this study, and a cylinder depth, which defines the spatial search range along the cylinder axis relative to the core point, set to 2 m. The latter parameter addresses computational efficiency only and must be set at least higher than the maximum change magnitude expected in the analysis. The temporal extent is defined by an aggregation window radius, $W_t = W_r \times |\Delta T|$. Here, $|\Delta T|$ is the time difference (in days) between the reference and target epochs, and W_r is a user-defined, unitless ratio. Therefore, a point is a candidate for aggregation only if it falls within the 3D spatial cylinder and if its timestamp is within the symmetric temporal window spanned forward and backward in time around the core timestamp t_c of the epoch being processed, i.e., $t_c \pm W_t$.

3.3 Adaptive temporal aggregation

The adaptive temporal aggregation step addresses data sparsity by ensuring each core point has a sufficient number of neighbors

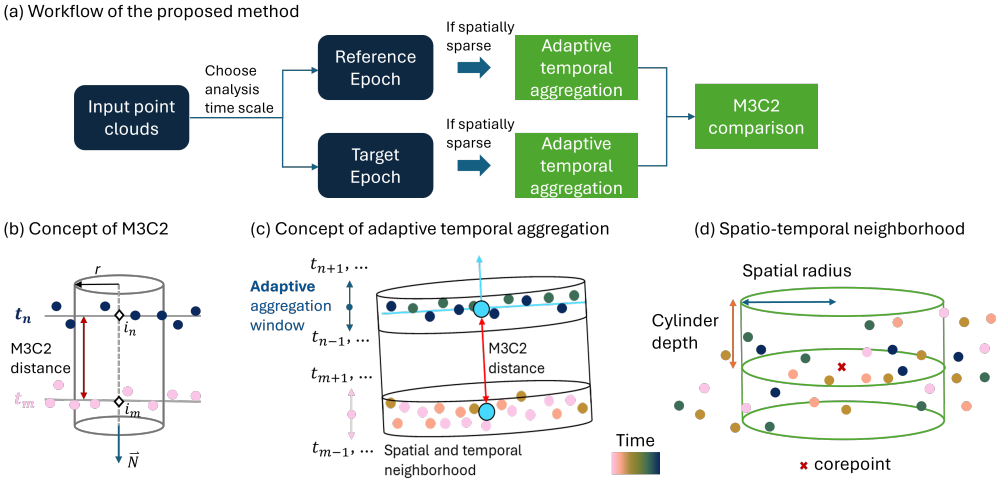

Figure 3. Overview of the proposed time-adaptive M3C2 extension. (a) Workflow showing the division of input point clouds into reference and target epochs according to a given analysis time scale, an adaptive aggregation step for locally sparse data, and the final comparison based on the original M3C2 algorithm. (b) Concept of the M3C2 distance calculation where a cylinder with a radius (r) is spanned along a pre-calculated normal vector ($\vec{N}$) of the reference epoch. All points from the reference (t_n) and the target cloud (t_m) inside the cylinder are projected onto this normal vector and their position averaged. The M3C2 distance is the distance between these two projected positions for the reference and target cloud. (c) The method concept of aggregating points in a spatiotemporal neighborhood into a representative core point within both the reference and target epochs before M3C2 distance calculation. (d) The spatiotemporal neighborhood is a local cylinder around the core point defined by a spatial radius and a cylinder depth. The color gradient in (b), (c) and (d) indicates points from different time epochs.

(N_{req}) within the cylinder. If the neighbor count within a core point's own epoch is below this threshold, a temporal search window expands symmetrically in time to incorporate the closest pair of past and future epochs (e.g., t_{n+1} and t_{n-1}) at each step (Fig. 3c). This symmetrical approach is crucial to prevent any temporal bias that could be introduced by actual surface change with one-sided aggregation. The iterative process terminates as soon as the neighbor count reaches at least N_{req} or the window expands to its predefined maximum limit (W_t). Since our aggregation step incorporates entire epochs at once rather than single points, the final neighbor count can be higher than the minimum requirement of N_{req}.

Crucially, this aggregation is applied on a core point basis, resulting in different levels of aggregation across different areas of the point cloud. Dense areas may conduct no temporal aggregation at all, while sparse areas are aggregated over a wider temporal window. The collective result of this process is the generation of a new, more complete epoch, with potential for more accurate change quantification through higher coverage. This new epoch then serves as the final input for the distance computation step, as conceptually illustrated in Fig. 3b.

Quantification of the 3D change between the reference and target epochs is performed by applying the M3C2 algorithm to the robust epochs generated through the adaptive temporal aggregation process. For this study, we used R = r = 0.5 m.

3.4 Experiments and evaluation

To quantitatively evaluate our proposed adaptive temporal aggregation method, we designed a set of experiments. The common workflow for all tests is a two-step process: (1) an aggregation method (our method or the fixed-window baseline) is applied to the sparse test data to generate an enhanced epoch, and (2) the M3C2 algorithm is then used to compute a distance result from the two enhanced epochs. The final accuracy of each method is determined by comparing its M3C2 results against the result derived from the dense reference data. Performance is measured using the metrics detailed below.

Our method is tested under two scenarios. The first scenario addresses the primary challenge of performing change analysis on sparse data. In this test, the accuracy of the M3C2 result is quantified by comparing it to the corresponding reference result. The second scenario tests the method's robustness to missing data by intentionally removing the entire point cloud of the reference epoch from the input dataset. Our method is then used to generate an estimation for this missing epoch by aggregating data from temporally adjacent epochs. The accuracy is subsequently evaluated by comparing the result against the reference data of the removed epoch.

3.4.1 Accuracy assessment: The accuracy of the method is evaluated by comparing its M3C2 distance (D_{test}) against the M3C2 result derived from reference data (D_r). This comparison is performed point-wise for the set of n common core points where both distance results are valid (i.e., not NaN). For each core point i, the distance difference is defined as $d_i = D_{test,i} - D_{r,i}$. From this set of differences $\{d_1, d_2, ..., d_n\}$, a set of statistical metrics is computed. The overall error magnitude is assessed using the Root Mean Square Error (RMSE) and the Mean Absolute Error (MAE). Systematic bias is evaluated using the mean difference, while the dispersion of errors is measured by the Standard Deviation (STD). To quantify the prevalence of significant differences, we also calculate the percentage of large errors, defined as the proportion of points where $|d_i|$ exceeds a 0.1 m threshold. In addition to these comparison metrics, we evaluate two metrics derived directly from the M3C2 output to assess the quality of the change detection itself. The LOD is the minimum change detectable at a 95% confidence level, a value derived from local roughness, point sampling, and registration error (Lague et al., 2013). The significance percentage is the proportion of points where the computed distance exceeds its corresponding LOD value.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-2-2026-623-2026 | © Author(s) 2026. CC BY 4.0 License.

3.4.2 Roughness assessment: In addition to accuracy, we assess the quality of the aggregated epoch by its local roughness, which quantifies the geometric integrity of the point cloud. For each core point, roughness is calculated within a local neighborhood (e.g., 0.5 m radius), provided it contains a minimum number of 3 points. A best-fit plane is determined for each core point using its neighboring points. These points are projected onto the plane's normal vector, and the roughness is defined as the standard deviation of these projected distances (which is the same procedure used by the original M3C2 to derive roughness for the LOD). This metric is used to evaluate how the temporal aggregation process, by incorporating points from different epochs, affects the local geometric structure. A significant increase in roughness after aggregation would indicate that the method is introducing noise or inconsistencies, which could impact the accuracy of the final change detection result.

4. Results

We first present a parameter sensitivity analysis to explore the impact of its core parameters, the required points per neighborhood, N_{req}, and the maximum temporal window ratio, W_r on the final results. Following this, a comprehensive comparative analysis is provided. Our method is evaluated against two baseline approaches, fixed-window aggregation and standard M3C2, across two defined scenarios: performance on spatially sparse data and in the presence of temporal gaps. The evaluation combines quantitative metrics with visual analysis to illustrate the performance differences.

4.1 Parameter sensitivity analysis

To investigate the influence of our method's key parameters, we conducted a sensitivity analysis by systematically varying the required points per neighborhood (N_{req}) and the maximum window ratio (W_r). As seen in Fig. 4a, the percentage of valid M3C2 results increases sharply even at a low N_{req} value of 1. This occurs because our aggregation step, once triggered, incorporates entire scans from adjacent epochs rather than single points. Concurrently, error metrics such as RMSE and MAE decrease as N_{req} approaches an optimal value around 5, after which they stabilize or slightly increase. This suggests that an excessively large N_{req} can incorporate temporally distant data that may slightly degrade accuracy. Fig. 4b shows the method is less sensitive to the maximum window ratio, W_r. The valid result percentage rises sharply as W_r increases from 0.0 to 0.2 and then flattens, indicating marginal improvements thereafter. The distinct plateau between W_r values of 0.2 and 0.5 is a characteristic of this dataset's specific temporal sampling, where no additional scans were available to be included in that range. This analysis confirms that W_r acts as an upper bound, while the density-driven N_{req} is the primary factor determining the algorithm's performance. Based on these results, we select $N_{req} = 3$ and $W_r = 0.2$ for our experiments to maximize result completeness while maintaining high accuracy.

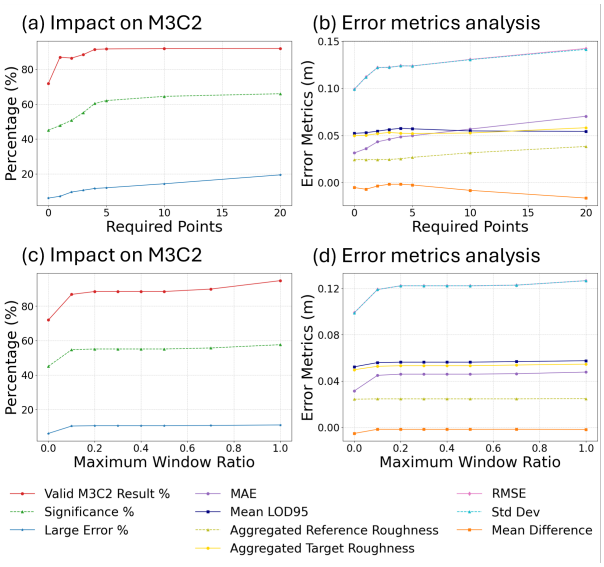

Figure 4. Parameter sensitivity analysis. (a) The impact of the required points per neighborhood (N_{req}) on the percentage of valid, significant, and large error results, and (b) its impact on error and roughness metrics. In (a) (b), the maximum window ratio (W_r) was fixed at 0.5. (c) The impact of W_r on result percentages, and (d) its impact on error and roughness metrics. In (c) (d), N_{req} was fixed at 3. The percentages of significance and large errors are subsets of the valid results.

4.2 Comparative analysis

The performance metrics for all methods and scenarios are summarized in Table 1. Overall, the results demonstrate a substantial increase in both the completeness and accuracy of the change detection using our method. A primary challenge of applying M3C2 to sparse data is the failure to compute results where local point density is insufficient. As shown in Fig. 5, the standard M3C2 baseline fails to yield a result for a large portion of the area (in gray), yielding 11,331 invalid points (NaNs). In contrast, both the fixed-window and our proposed adaptive method produce more complete change maps (Fig. 6a-c). This improvement in completeness is quantified in Table 1, where the percentage of valid M3C2 results increases. Beyond completeness, our method also demonstrates better geometric accuracy. According to Table 1, the proposed method achieves an RMSE of 0.122 m, a 31% reduction in error compared to the fixed-window approach (0.177 m). A similar trend is observed for the MAE (0.046 m vs. 0.110 m) and the percentage of large errors ($|d_i| > 0.1$m), which is nearly three times lower for our method (10.6% vs. 32.0%). The spatial distribution of these errors, shown in the bottom row of Fig. 6, visually confirms that our method results in significantly smaller errors compared to the fixed-window method. In all cases, the errors are mostly concentrated on the sparsely sampled channel banks. The higher error in the temporal gap scenario (b) compared to the sparse scenario (a) suggests that our method averages out genuine surface changes that occurred

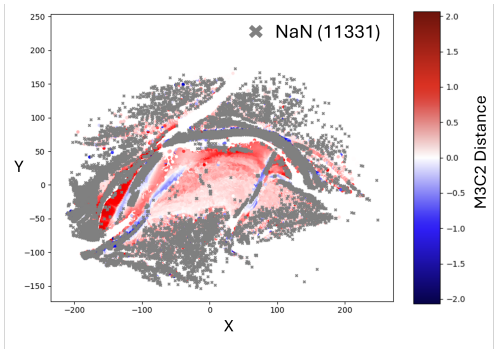

Figure 5. The M3C2 distance between the original reference and target epochs, visualized on the points of the reference epoch, with NaN values indicating where no result could be computed.

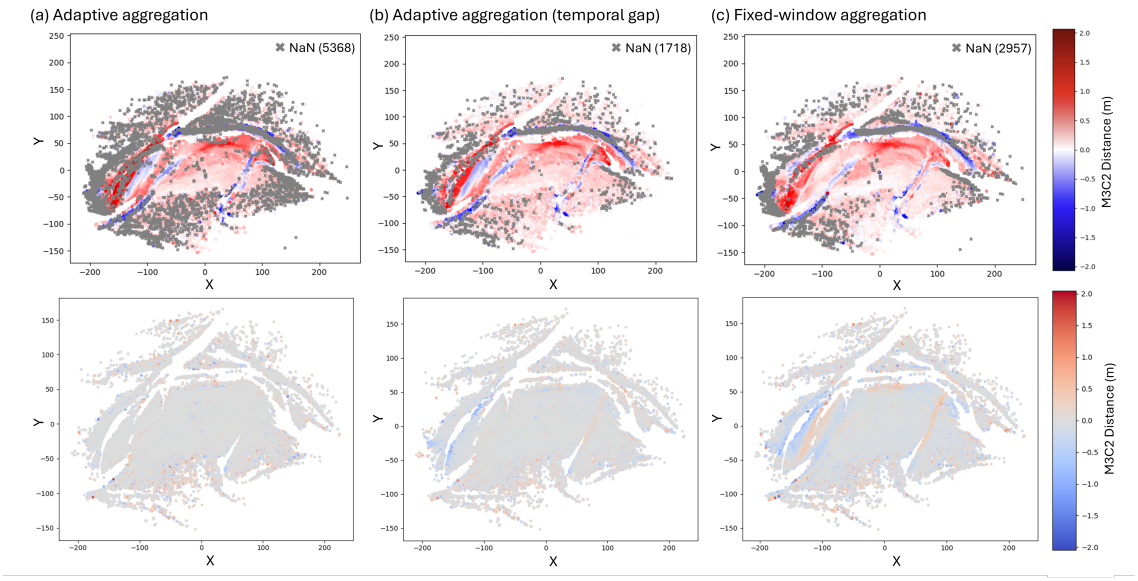

Figure 6. Visual comparison of M3C2 change detection results and corresponding error maps for the three tested experiments. The top row displays the M3C2 change maps, while the bottom row displays the error maps, calculated as the difference from the benchmark result (M3C2 on reference data). (a) Adaptive aggregation on spatially sparse data, (b) Adaptive aggregation on data with temporal gap (i.e. missing epochs), (c) Fixed-window aggregation on spatially sparse data.

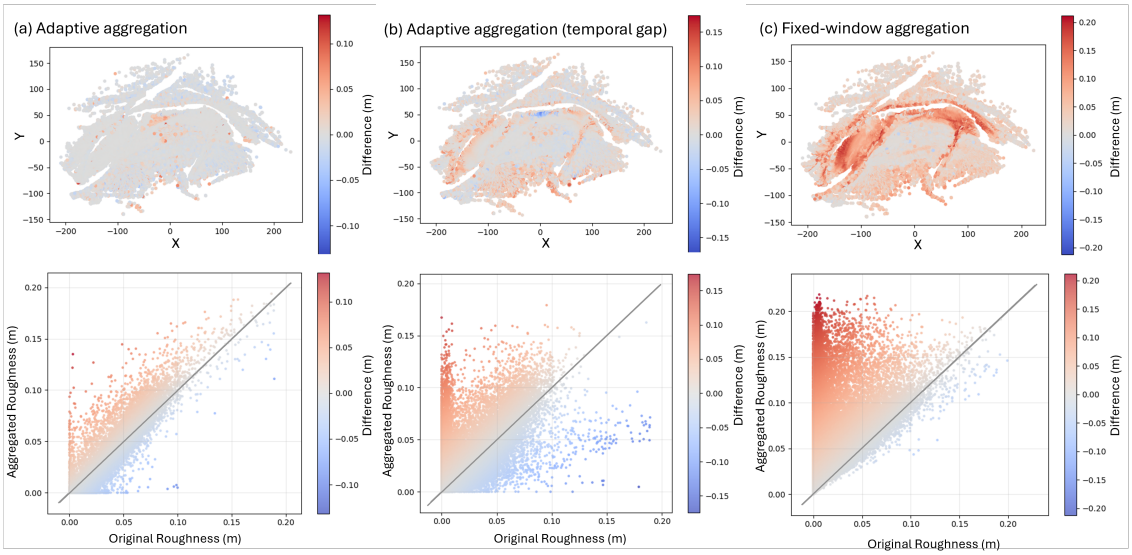

Figure 7. Aggregation impact on local roughness. The columns compare the proposed method in (a) the primary scenario, (b) the method in the temporal gap scenario, and (c) the fixed-window baseline. The top row displays the spatial distribution of roughness errors. The Difference (m) shows the roughness error, defined as the difference between each method’s roughness and the reference roughness computed from the original dense dataset. The bottom row presents scatter plots of the aggregated versus the original roughness.

Table 1. Comparison of the proposed method with the fixed-window baseline.

Metric	Reference	Adaptive Aggregation	Adaptive Aggregation (Temporal gap)	Fixed-Window
RMSE	-	0.122 m	0.158 m	0.177 m
STD	-	0.122 m	0.157 m	0.172 m
MAE	-	0.046 m	0.081 m	0.110 m
Mean difference	-	-0.002 m	-0.023 m	-0.042 m
Percentage of large errors	-	10.6 %	21.6 %	32.0 %
Percentage of valid M3C2 results	75.7 %	88.5 %	96.3 %	93.7 %
Mean LOD	0.041 m	0.056 m	0.075 m	0.056 m
Percentage of significant changes	48.1 %	55.1 %	56.9 %	65.5 %

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-2-2026-623-2026

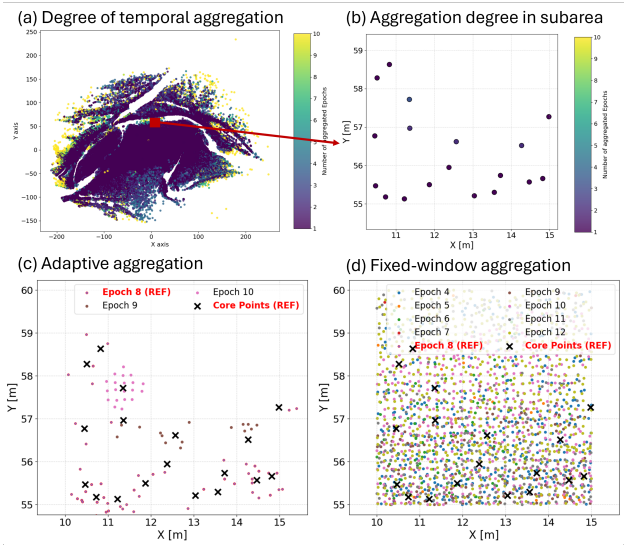

Figure 8. Illustration of the temporal aggregation method. (a) Spatial distribution of aggregation degree (number of epochs). (b) Zoom-in subarea of aggregation degree for core points from (a). (c) Neighborhood enhancement by our method. (d) Comparison using fixed-window aggregating 9 epochs over two months.

during the temporal gap.

To illustrate the mechanism behind our method, Fig. 8a confirms that the aggregation is not spatially identical across the scene. In dense areas, only a single epoch is used, preserving the original data. Along sparse channel edges, the algorithm adaptively incorporates data from multiple adjacent epochs to ensure a sufficient point density. To investigate this process at the core point level, we selected a 5 m $\times$ 5 m subarea for a detailed view. Fig. 8b shows the aggregation degree for these core points. Our method (Fig. 8c) selectively aggregates a minimal yet sufficient number of points from the closest temporal neighbors to enhance the core points (black crosses). This approach preserves geometric fidelity by avoiding the introduction of noise and the averaging of real changes across temporally distant data. In contrast, the fixed-window aggregation (Fig. 8d) merges all points from a wide temporal window of nine epochs (4 to 12), spanning a total of two months, which yields a denser but less accurate estimation.

Furthermore, we evaluated the impact of aggregation on data quality by analyzing the change in local point cloud roughness (Fig. 7). For the majority of points, the aggregated roughness remains very close to the original roughness, confirming that our method successfully increases coverage without significantly degrading the local geometric quality. Consequently, this ensures that the subsequent uncertainty estimation and the resulting LOD remain reliable. Finally, in the temporal gap scenario, the robustness of our method to missing epochs is confirmed. Table 1 shows an RMSE of 0.158 m and an MAE of 0.081 m, demonstrating a reasonable level of accuracy even when the entire reference epoch is missing.

5. Discussion

The results demonstrate that our adaptive temporal aggregation provides an effective method for improved change detection on spatially and temporally heterogeneous point clouds, particularly for 3D time series with highly variable acquisition intervals and non-linear dynamics. While M3C2 is a state-of-the-art algorithm for accurate 3D change detection, its applicability is fundamentally dependent on a certain data coverage, a condition rarely met for all epochs in long-term environmental monitoring. Specifically, core M3C2 operations, such as local normal estimation, position averaging for distance determination, and uncertainty assessment, require a minimum number of points in local neighborhoods. In heterogeneous spatiotemporal data, this fundamental requirement is not met across all epochs and areas, leading to less complete and thus less accurate change detection and quantification. Our method directly addresses this limitation through a temporal neighborhood search and its density-driven mechanism. It intelligently aggregates the minimum necessary information from adjacent epochs only where needed, ensuring that a robust normal can be estimated and a distance can be computed subsequently. This targeted, locally adaptive aggregation enhances the applicability of M3C2 while mitigating the effects introduced by fixed-window aggregation. The method is currently implemented as a Python workflow utilizing the M3C2 implementation from py4dgeo (py4dgeo Development Core Team, 2022). A core assumption of the method is that the local topography remains sufficiently stable within the temporal aggregation window, relative to the change scales to be detected. In scenarios where rapid or high-magnitude changes occur between adjacent epochs, our method would lead to averaging them out. Our error and roughness maps (Fig. 6 and Fig. 7) reveal that the largest errors and highest post-aggregation roughness values are concentrated in the dataset’s sparsest peripheries, which forces the algorithm to apply the highest degree of temporal aggregation (Fig. 8a). This suggests that while our method handles sparsity effectively, its performance can degrade in extremely poor data regions, potentially introducing geometric inconsistencies. Future work will focus on automated parameter tuning and developing enhanced aggregation criteria, such as temporal weighting, to remedy the effect of incorporating actual changes into the aggregation. To ensure broader generalization, we will evaluate the method across diverse environmental scenarios using Virtual Laser Scanning (VLS) for simulated dynamic scenes (Weiser and Höfle, 2026). Finally, we plan to integrate this method as a spatiotemporal neighborhood search in py4dgeo, including an assessment of the computational cost for large-scale datasets to optimize its applicability in real-time 4D monitoring.

6. Conclusion

In this paper, we propose a novel adaptive temporal aggregation method to overcome the fundamental limitations of 3D change detection on sparse and irregularly sampled point cloud time series. Our approach addresses incomplete data coverage, a primary limitation when applying high-precision algorithms like M3C2 to sparse datasets. It ensures sufficient local points for M3C2 to perform robust surface estimation, compute average distances, and assess uncertainty, which in our results increased M3C2 result coverage by 13% compared to the standard M3C2. Furthermore, our method leverages temporal redundancy by aggregating adaptively, which minimizes unnecessary data aggregation and reduces distance errors by over 30% compared to a non-adaptive fixed-window approach. While the method assumes a degree of local stability and may mask extremely rapid inter-epoch changes, its primary contribution is to provide a practical pathway for reliably quantifying 3D change in dynamic environments. This work effectively extends the utility of high-precision comparison methods to the heterogeneous datasets common in long-term topographic monitoring.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-2-2026-623-2026 | © Author(s) 2026. CC BY 4.0 License.

Acknowledgements

This research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), project 535733258 (Extract4D), and partly supported by the Bavarian State Ministry of Science and the Arts in the framework of the bidt Graduate Center for Postdocs (AI4ENV project).

References

Anders, K., Winiwarter, L., Mara, H., Lindenbergh, R., Vos, S. E., Höfle, B., 2021. 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.

Chen, C., Yang, B., 2016. Dynamic occlusion detection and inpainting of in situ captured terrestrial laser scanning point clouds sequence. *ISPRS Journal of Photogrammetry and Remote Sensing*, 119, 90–107.

Cignoni, P., Rocchini, C., Scopigno, R., 1998. Metro: measuring error on simplified surfaces. *Computer Graphics Forum*, 17(2), 167–174.

Eitel, J. U., Höfle, B., Vierling, L. A., Abellán, A., Asner, G. P., Deems, J. S., Glennie, C. L., Joerg, P. C., LeWinter, A. L., Magney, T. S. et al., 2016. Beyond 3-D: The new spectrum of lidar applications for earth and ecological sciences. *Remote Sensing of Environment*, 186, 372–392.

Eltner, A., Kaiser, A., Abellan, A., Schindewolf, M., 2017. Time lapse structure-from-motion photogrammetry for continuous geomorphic monitoring. *Earth Surface Processes and Landforms*, 42(14), 2240–2253.

Girardeau-Montaut, D., Roux, M., Marc, R., Thibault, G., 2005. Change detection on points cloud data acquired with a ground laser scanner. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 36(3), W19.

Guisado-Pintado, E., Jackson, D. W., Rogers, D., 2019. 3D mapping efficacy of a drone and terrestrial laser scanner over a temperate beach-dune zone. *Geomorphology*, 328, 157–172.

Kharroubi, A., Poux, F., Ballouch, Z., Hajji, R., Billen, R., 2022. Three dimensional change detection using point clouds: A review. *Geomatics*, 2(4), 457–485.

Kromer, R. A., Abellán, A., Hutchinson, D. J., Lato, M., Edwards, T., Jaboyedoff, M., 2015. A 4D filtering and calibration technique for small-scale point cloud change detection with a terrestrial laser scanner. *Remote Sensing*, 7(10), 13029–13052.

Kuschnerus, M., Lindenbergh, R., Vos, S., 2021. Coastal change patterns from time series clustering of permanent laser scan data. *Earth Surface Dynamics*, 9(1), 89–103.

Lague, D., Brodu, N., Leroux, J., 2013. Accurate 3D comparison of complex topography with terrestrial laser scanner: Application to the Rangitikei canyon (NZ). *ISPRS Journal of Photogrammetry and Remote Sensing*, 82, 10–26.

Leroux, J., 2013. Chenaux tidaux et dynamique des prés-salés en régime méga-tidal: approche multi-temporelle du siècle à l’événement de marée. PhD thesis, Université Rennes 1; Université Européenne de Bretagne.

Lindenbergh, R., Anders, K., Campos, M., Czerwonka-Schröder, D., Höfle, B., Kuschnerus, M., Puttonen, E., Prinz, R., Rutzinger, M., Voordendag, A. et al., 2025. Permanent terrestrial laser scanning for near-continuous environmental observations: Systems, methods, challenges and applications. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, 100094.

Ministère de la Transition Écologique, Cerema, IGN, 2014. Ortho Littorale v2 (coastal orthophotography, tile no. 53). GéoLittoral. Open Licence / Etalab v2.0. Available at: `https://www.geolittoral.developpement-durable.gouv.fr/ortho-littorale-a1101.html`.

Pinton, D., Canestrelli, A., Wilkinson, B., Ifju, P., Ortega, A., 2021. Estimating ground elevation and vegetation characteristics in coastal salt marshes using UAV-based LiDAR and digital aerial photogrammetry. *Remote Sensing*, 13(22), 4506.

py4dgeo Development Core Team, 2022. py4dgeo: library for change analysis in 4D point clouds. https://github.com/3dgeo-heidelberg/py4dgeo.

Qin, R., Tian, J., Reinartz, P., 2016. 3D change detection – approaches and applications. *ISPRS Journal of Photogrammetry and Remote Sensing*, 122, 41–56.

Stark, J., Van Oyen, T., Meire, P., Temmerman, S., 2015. Observations of tidal and storm surge attenuation in a large tidal marsh. *Limnology and Oceanography*, 60(4), 1371–1381.

Tabernig, R., Albert, W., Weiser, H., Fritzmann, P., Anders, K., Rutzinger, M., Höfle, B., 2025. Temporal aggregation of point clouds improves permanent laser scanning of landslides in forested areas. *Science of Remote Sensing*, 100254.

Tessier, B., 1993. Upper intertidal rhythmites in the Mont-Saint-Michel Bay (NW France): perspectives for paleoreconstruction. *Marine Geology*, 110(3-4), 355–367.

Weiser, H., Höfle, B., 2026. Advancing vegetation monitoring with virtual laser scanning of dynamic scenes (VLS-4D): Opportunities, implementations and future perspectives. *Methods in Ecology and Evolution*, 17(1), 33–51.

Winiwarter, L., Anders, K., Höfle, B., 2021. 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.

Winiwarter, L., Anders, K., Schröder, D., Höfle, B., 2022. Full 4D change analysis of topographic point cloud time series using Kalman filtering. *Earth Surface Dynamics Discussions*, 2022, 1–25.

Zahs, V., Winiwarter, L., Anders, K., Williams, J. G., Rutzinger, M., Höfle, B., 2022. Correspondence-driven plane-based M3C2 for lower uncertainty in 3D topographic change quantification. *ISPRS Journal of Photogrammetry and Remote Sensing*, 183, 541–559.