Canadian Geotechnical Journal

Article

Rapid 3D lidar change detection for geohazard identification using GPU-based alignment and M3C2 algorithms

Luke Weidner, Alex Ferrier, Megan van Veen, and Matthew J. Lato

BGC Engineering Inc., Suite 500, 980 Howe Street, Vancouver, BC V6Z 0C8, Canada

Corresponding author: **Luke Weidner** (email: lweidner@bgcengineering.ca)

Abstract

Topographic change detection is increasingly being used to identify and monitor landslides and other geohazards in support of risk-informed decision-making. Expanding change detection from site specific to regional scales enables increasingly proactive asset management and contributes to improving the resilience of infrastructure to extreme events. It is widely known that change detection precision can be improved by applying three-dimensional algorithms, such as iterative closest point (ICP) and M3C2, directly to raw point clouds. However, this also increases the computational requirements compared to alternatives such as digital elevation model differencing. This study presents a novel graphics processing unit (GPU)-based implementation of the ICP-M3C2 workflow to address this limitation. In the proposed algorithm, point cloud data segments are automatically queued and served to the working GPU, which efficiently performs point cloud processing operations, while the central processing unit (CPU) performs data management operations in parallel. The developed method is estimated to be up to 54 times faster than CPU-based versions of the same algorithm. In this study, we present how the workflow has been applied to six regional-scale landslide identification and monitoring case studies in which landslides are causing the disruption of pipelines, highways, and rail corridors. Overall, in 2021 and 2022, over 17 500 linear km of change detection were processed using the demonstrated method.

Key words: landslides, GPU, change detection, alignment, DEM of Difference, geohazards

1. Introduction

Four-dimensional (4D) topographic change detection using three-dimensional (3D) data is increasingly being used to identify, monitor, and respond to natural hazards, such as rockfall, landslides, debris flows, floods, and earthquakes as well as anthropogenic changes due to construction and resource extraction. Repeated topographic data collection and comparison of these data facilitate a better understanding of when and where changes have occurred or could occur, which can be used to support risk-informed decision-making by engineers, scientists, asset owners, governments, and the public. Topographic change detection has played a key role in reconnaissance and response following recent geotechnical disasters, including the Oso landslide in Washington, USA (Wartman et al. 2016; Lato et al. 2019), and the November 2021 atmospheric river event in southern British Columbia (Hunter 2022; Vasquez 2022). Other recent examples include identifying landslides from change detection for regional-scale hazard mapping (Guzzetti et al. 2012; van Veen et al. 2017, 2022*a*), fault zone mapping (Nissen et al. 2014), coastal cliff erosion studies (Westoby et al. 2020), and monitoring landslide response to construction/mitigation activities (Lato et al. 2016; Donati et al. 2020).

Topographic datasets covering large spatial extents are typically collected using fixed-wing aircraft, helicopters, or uncrewed aerial vehicles (UAVs). These platforms are equipped with lidar sensors or cameras for photogrammetry model generation, each of which has its own unique set of advantages and disadvantages (Passalacqua et al. 2015). For projects with large spatial extents and vegetated ground, airborne lidar can be deployed to capture high-resolution and high-precision point data over hundreds or thousands of km^2 of land surface, while penetrating vegetation (Lato et al. 2015). For smaller sites, such as a mining project, UAVs enable datasets to be collected at high temporal frequencies (as often as daily or weekly).

3D data have become increasingly better in terms of accuracy and resolution (Pirotti et al. 2013; Harpold et al. 2015; Niculiță et al. 2020). This has enabled the detection of smaller changes, and in some cases, the ability to confidently identify precursors to large-scale movements (Kromer et al. 2015; Lato et al. 2019). However, this also means that 4D data processing is a time-consuming and computationally expensive task (Schovanec et al. 2021). Typical computational methods for performing change detection (Fig. 1) include the following:

Fig. 1. Illustration of four types of change calculation. Digital elevation model (DEM) of difference (DoD, panel A) is shown with change (d_1) being the result of a vertical difference between grid cells. Cloud to cloud (C2C, panel B) change and cloud to mesh (C2M, panel C) change are shown in the middle. Multiscale model to model cloud comparison (M3C2) (panel D) is shown with the change direction calculated dynamically for each point based on the local normal vector. The change value (d_4) is computed as the difference in mean projected locations of the two point clouds along the cylinder axis. Graphic modified from Winiwarter et al. (2021) under a Creative Commons BY 4.0 license.

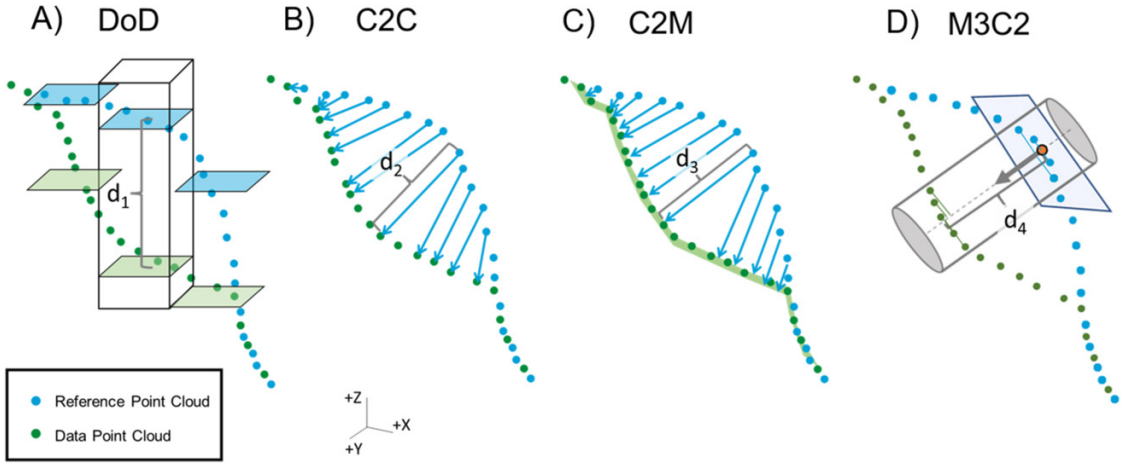

- Digital elevation model (DEM) differencing, or difference of DEMs (DoD), which measures changes in the vertical direction using gridded or interpolated versions of the datasets, and where the grid size is selected by the user.
- Point-to-point or cloud-to-cloud (C2C) change, which finds the distance between a point and its nearest neighbor in a reference point cloud (Girardeau-Montaut et al. 2005).
- Point-to-plane or cloud-to-mesh (C2M) change in which a point cloud is compared to a 3D meshed surface in a direction normal (i.e., perpendicular) to the mesh (e.g., Abellán et al. 2009; Fei et al. 2023).
- Multiscale model to model cloud comparison (M3C2) in which change is calculated in a direction normal to the local surface using an average of nearby points (Lague et al. 2013).

DoD is a widely used technique in research and engineering (Alotaibi et al. 2022; Conforti et al. 2021; Swirad and Young 2021; Tremblay-Auger et al. 2021). While the DoD method is computationally efficient and straightforward to calculate, it is subject to a variety of limitations. By gridding or interpolating data, DEMs introduce simplifications that limit the detail of the surface. DEMs may interpolate over areas where there are no data points present and therefore calculate a change value where no data exist in the raw point cloud (Lague et al. 2013). Further, change detection in the vertical direction may reasonably represent some ground movement processes (such as subsidence) but not others (such as some landslides, lateral bank migration or erosion, rockfall, and mining excavations) (Williams et al. 2021).

M3C2-like algorithms are often preferred for change detection in geomorphology for several reasons (Stumpf et al. 2015; Smets et al. 2017; DiFrancesco et al. 2020; Williams et al. 2021). First, change estimates that are orthogonal to the local surface are more meaningful in steep and complex topography than purely vertical change. For example, low-angle topography will result in a near-vertical change direction, whereas vertical change computed on steep rock slopes or incised riverbanks/canyons is not as informative as the horizonal retreat distance. According to Lague et al. (2013), a single change direction is not effective for datasets that contain multiple environments or changing slope angles/aspects. M3C2 and C2M automatically account for this and provide meaningful change in both instances. Second, spatial averaging of neighboring points results in the ability to detect smaller changes with higher statistical confidence and avoids unnecessary simplification in the gridding/meshing step (Abellán et al. 2009; Kromer et al. 2015; Wagner et al. 2017).

While change along the local normal direction may be preferable to DoD in some circumstances, neither DoD nor M3C2 measure or attempt to measure the true displacement vector of a given block or particle, which is often parallel to the surface, not orthogonal. Many methods for determining an appropriate change direction in complex topography have been developed recently, and this is an active area of research, but detailed interpretation of slope morphology is still required to infer true displacement when using these methods (Gojcic et al. 2021; de Gélis et al. 2022; Weidner et al. 2021; Williams et al. 2021;van Veen et al. 2022*b*; Zahs et al. 2022).

Change detection datasets must be finely co-registered to each other, as any registration errors will have a direct effect on the ability to distinguish real changes in the data. Here, several sources of uncertainty must be noted, including the georeferencing, positioning, and classification errors of both datasets (Passalacqua et al. 2015). For airborne lidar, these combined uncertainties typically result in a best-case limit of detectable change of several tens of centimeteres if no registration is applied (Passalacqua et al. 2015; Okyay et al. 2019), which can be problematic when the geomorphic

Fig. 2. Map illustrating the approximate locations of study sites and an overview of the extent of regional-scale lidar change detection (LCD) performed by the authors using the GPU method in 2021 and 2022. Pipeline and rail corridors associated with these change detection projects are included for reference. The width of the green highlighted areas is enlarged to improve visibility and is not reflective of the actual change detection area, which in most cases is much narrower, centered around the infrastructure.

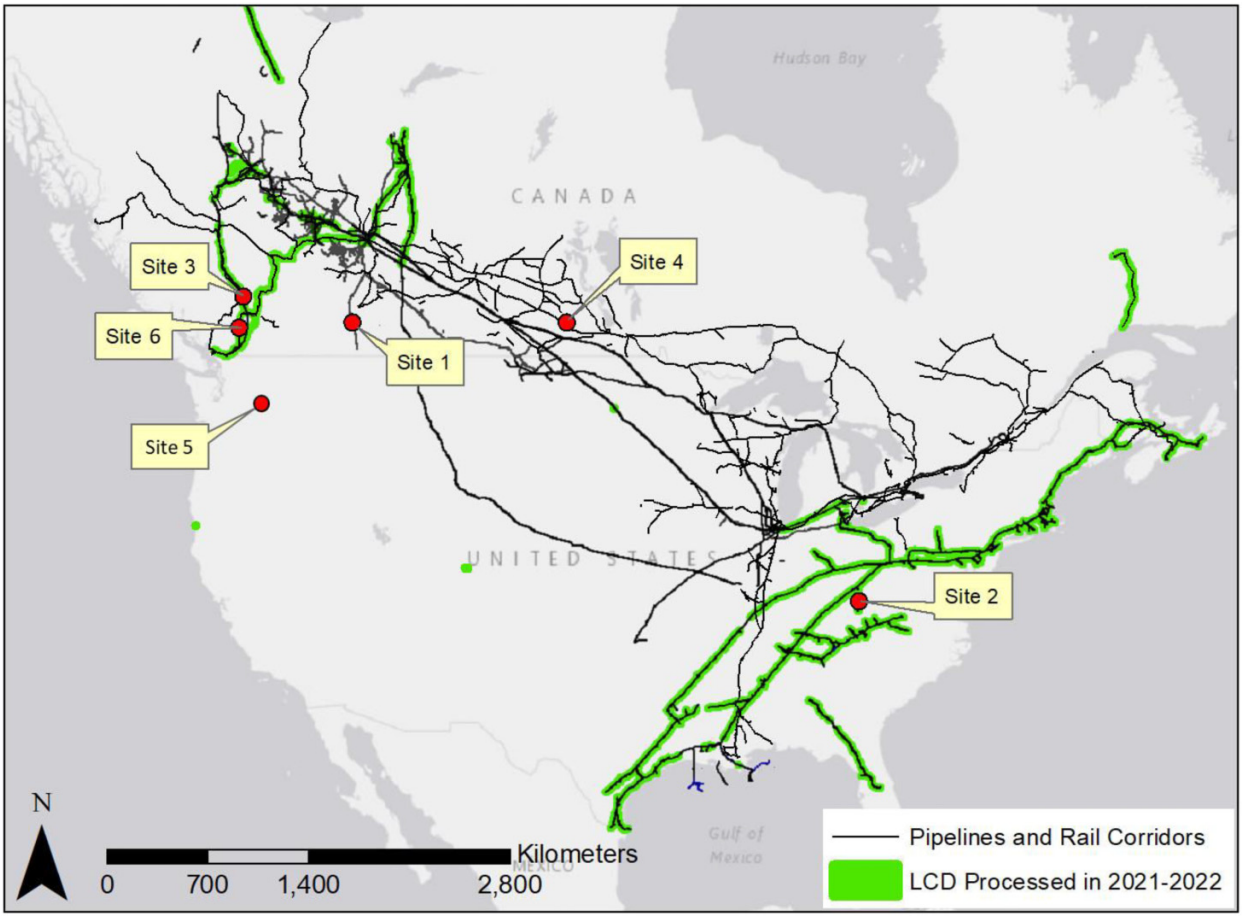

change of interest is of a similar magnitude (e.g., slow moving landslides, settlement, or erosion). Iterative closest point (ICP) algorithms, which estimate the rigid best-fit 3D transformation between two point clouds, are commonly used as an error reduction technique (Girardeau-Montaut et al. 2005; Guerin et al. 2017; Schovanec et al. 2021), effectively minimizing the network-scale georeferencing errors in the two point clouds in regions where no change is presumed to have occurred (e.g., flat ground, stable rock outcrops, roads, or buildings).

While the use of M3C2-like algorithms with ICP can lead to gains in the precision of point cloud change detection, they are computationally expensive and time consuming. Recent research has demonstrated that graphics processing units (GPUs) can be used to accelerate point cloud processing by a factor of two or more compared to central processing unit (CPU)-based processing. Problems such as shape detection (Chaiso and Ratanaworabhan 2020), ground point extraction (Baker and Sadowski 2013; Hu et al. 2013), and autonomous collision avoidance (Kaldestad et al. 2014) have all been explored using GPU acceleration. Anand et al. (2020) compared GPU- and CPU-based point cloud processing hardware and found a sixfold increase in speed by using a GPU for computation. This shows great opportunities for rapid point cloud processing, but to date, there have been very few published applications of GPUs for change computation from point clouds and none for geomorphic change with M3C2. Richter et al. (2013) developed a GPU-based change detection and visual rendering algorithm, but the implemented change method was C2C, not the more robust M3C2 method.

The objectives of this study are to (1) provide a benchmark of the precision of the M3C2-ICP workflow in comparison to DoD and (2) operationally apply a GPU-based M3C2-ICP workflow to manage landslide risks over large areas. We intend to demonstrate that automated processing enables analysis that was hitherto impractical or impossible to perform at a regional scale (e.g., a pipeline network, rail or highway corridor, hydro reservoir shoreline, or mine site).

2. Data and study sites

Data from six different geohazard sites have been used for this study. The six sites are located across North America in different physiographic and geological settings (Fig. 2). While

Table 1. Bare-earth point cloud resolutions for datasets used in this study.

Site No.	Year	Average point spacing (m)	No. of points	Overlap plan area (km^2)
1	2006	0.72	6 188 121	2.1
	2019	0.11	37 299 641	
2	2019	0.17	47 347 684	1.3
	2020	0.17	20 793 558	
3	2006	0.46	6 750 676	3.4
	2016	0.25	30 061 638	
4	2015	0.33	9 019 258	3.3
	2017	0.30	9 876 812	
5	2006	1.83	7 423 972	15.2
	2013	0.43	19 936 492	
6	2015	0.19	2 437 691 890	420
	2021	0.18	3 547 598 001	

the test locations used here are relatively small (a few square kilometeres), they are part of much larger change detection campaigns consisting of thousands of square kilometeres, as illustrated by the green highlighted areas in Fig. 2. The total length of infrastructure corridor change detection processed using our method in 2021 and 2022 is conservatively estimated at 17 500 km.

Site 1 is a pipeline slope crossing located in southern Alberta, Canada, where the buried infrastructure crosses a slope on approach to a large valley, with landslide morphology present. The slope is relatively shallow (10°), and evidence of both shallow and deep-seated landslides is present at this site. An active landslide is present, and the interpreted mechanism is sliding on a weak glaciolacustrine clay layer at depth.

Site 2 is a location in eastern Kentucky, USA, where a buried pipeline traverses the crest of a slope that comprises sedimentary bedrock (sandstone and shale). This slope contains an active landslide in bedrock within a gully-like feature. While this landslide has not been characterized in detail, other landslides in this physiographic region typically occur as relatively shallow, translational slides along a colluvium/bedrock interface in interbedded sedimentary units. Similar to Site 1, the slope at this location is shallow (10–12°). This site is part of an ongoing regional monitoring campaign across the eastern United States in which three airborne lidar epochs collected variously between 2006 and 2022 are compared over ~14 000 linear km of pipeline corridor.

Site 3 is a landslide in the Fraser River valley in British Columbia that is traversed by a highway and major railway line. The active landslide is within the toe of a much larger postglacial earth flow (which originated from volcanic sediments) and is sliding along a narrow shear zone at the base of the earth flow deposit. This site is described in detail by Lato et al. (2016).

Site 4 is located in the Assiniboine River Valley in Manitoba, where a railway traverses the valley slopes. The slopes in this area are subject to slow-moving landslides in clay-shale bedrock, often on shallow (<10°) slopes.

Site 5 is a river valley in western Washington, USA, at the site of the 2014 Oso, Washington landslide, a high-mobility landslide/debris avalanche in glacial and colluvial sediments (Wartman et al. 2016). The North Fork Stillaguamish Valley has been previously investigated in detail to understand its history of landslide activities (Iverson et al. 2015; LaHusen et al. 2016; Lato et al. 2019). In this region, the Washington Department of Natural Resources has made multiple overlapping vintages of aerial lidar data publicly available on the Washington Lidar Portal (lidarportal.dnr.wa.gov). For this study, we downloaded ~20 km^2 of data in the Oso, Washington area for the 2006 and 2013 data epochs.

Site 6 consists of a 236 km linear section of transportation corridor in southern British Columbia. The corridor is affected by a variety of geohazards, including landslides, debris flows, dry granular flows, rockfalls, and erosion, due to being situated in a steep river valley. This site was included primarily to test the capacity of our system to efficiently process large amounts of data in a single batch.

For each of the six sites, two datasets have been selected for demonstration in this study, as summarized in Table 1. We refer to the older dataset in a change detection analysis as the "reference", or baseline cloud, and the newer dataset as the "data", or comparison cloud. In the case of these sites, data collection was contracted out for the specific projects, or purchased from an "off-the-shelf" provider, in the case of the 2006 data at Site 3. Data for Site 5 were downloaded from the Washington Department of Natural Resources lidar data portal. The data resolution ranges from an average point spacing of 0.72 m in the worst case (at Site 1) to 0.11 m in the best case (also at Site 1). It should be noted that the 2019 data at Site 1 were collected using a lidar sensor mounted on a UAV, whereas the rest of the data were collected using a fixed-wing or helicopter-mounted sensor.

3. Methods

3.1. Proposed change detection workflow

The algorithms we have developed input a set of two point clouds from multiple epochs, and output calculated changes over time for each point in the newer input dataset. This

Fig. 3. Diagram showing the steps of the processing workflow.

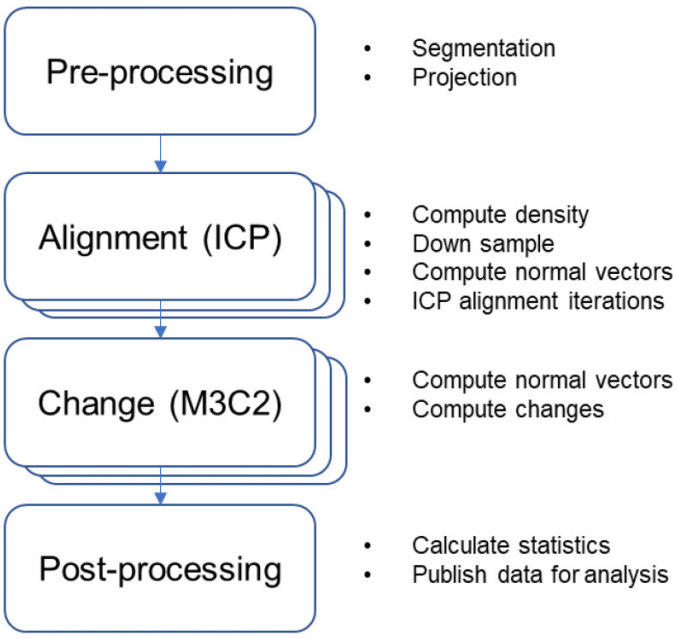

is accomplished through four main steps, as summarized in Fig. 3:

- Data pre-processing and segmentation
- Alignment using ICP (Section 3.1.1)
- Change calculation using M3C2 (Section 3.1.2)
- Post-processing and interpretation

Pre-processing typically includes converting raw data into the appropriate file formats, confirming data units and projection, and splitting the data into sections for batch processing. For large areas, data are split into segments of moderate size (less than 100 million points) to ensure that the clouds can be loaded into GPU memory. Splitting the data into segments before performing the ICP also helps to improve the local alignment within each section, rather than trying to align a few very large areas. Segments can be created manually by drawing polygons for each segment, which avoids having tile seams passing through important features, or automatically by creating square segments with a pre-defined size. The segment creation process is currently based on heuristics developed through experience, and as such there is a significant need for further work to define best practices.

Post-processing consists of rasterizing the resultant point clouds, merging the rasters into a single file, and uploading the layer to a custom web-viewer for interactive interpretation and interrogation.

3.1.1. ICP alignment

ICP alignment is used to ensure that the reference and data clouds are as closely aligned to each other as reasonably possible (Besl and McKay 1992). This maximizes the ability to detect real changes between two datasets. While the point clouds are assumed to have some level of initial correspondence due to both being georeferenced, the georeferencing quality may not be high enough to allow direct point cloud comparison at high-precision levels. This typically presents itself as a systematic "shift" or offset in one point cloud relative to the other. ICP alignment can be used to resolve these inconsistencies and increase the precision of change detection results. The steps of the ICP alignment are described in Algorithm 1.

Algorithm 1. Summary of the ICP implementation used in this study.

DO:

1. Compute point cloud densities and downsample the higher density cloud
2. Compute normal vector for each point in the reference cloud based on local plane fitting

While: Convergence criteria not met

3. Find point-to-plane nearest neighbor correspondences
4. Estimate and apply a rigid geometric transformation
5. Check if any convergence criterion has been met

END While
END DO

6. Return convergence results, move on to change calculation

This algorithm is based on the ICP implementation available in the Point Cloud Library (PCL; Rusu and Cousins 2011), with some modifications as described below.

In Step 1, the densities of the reference and data point clouds are estimated, and the higher density point cloud is downsampled using a voxel grid to be of a similar resolution to the lower density point cloud. Density is estimated as the average nearest-neighbor distances for 10 000 randomly sampled points. In our implementation, this downsampling process was observed to result in faster convergence and more reliable alignment results. Note that downsampling is only applied for the ICP process and is not carried forward to change calculation, where the full resolution point clouds are used.

In Step 2, normal vectors are computed for each point in the reference cloud by fitting a plane to a small spherical neighborhood of points. Normal vectors are needed for the point-to-plane distance metric used in Step 3. The neighborhood radius is computed automatically as three times the point spacing of the less dense cloud, or 1.5 m, whichever is larger. This value of the neighborhood radius was estimated empirically based on extensive testing on airborne laser scanning datasets. As such, the value would need to be adjusted based on user judgement if UAV or terrestrial-based datasets are used.

Steps 3 through 5 represent the core of the ICP algorithm in which the data point cloud is iteratively transformed until one of several convergence criteria is met (Fig. 4). First, a point-to-plane distance metric is used to find the nearest corresponding points between the two clouds (Step 3). After the initial correspondences are found, a percentile filter is applied such that the smallest 70% of the distance pairs are kept. This is done to ensure that outlier points and points cor-

Fig. 4. Illustration of iterative closest point (ICP) correspondence matching based on point-to-plane distance. Green arrows represent kept correspondences, while the orange arrows are a rejected correspondence as they have the largest distance.

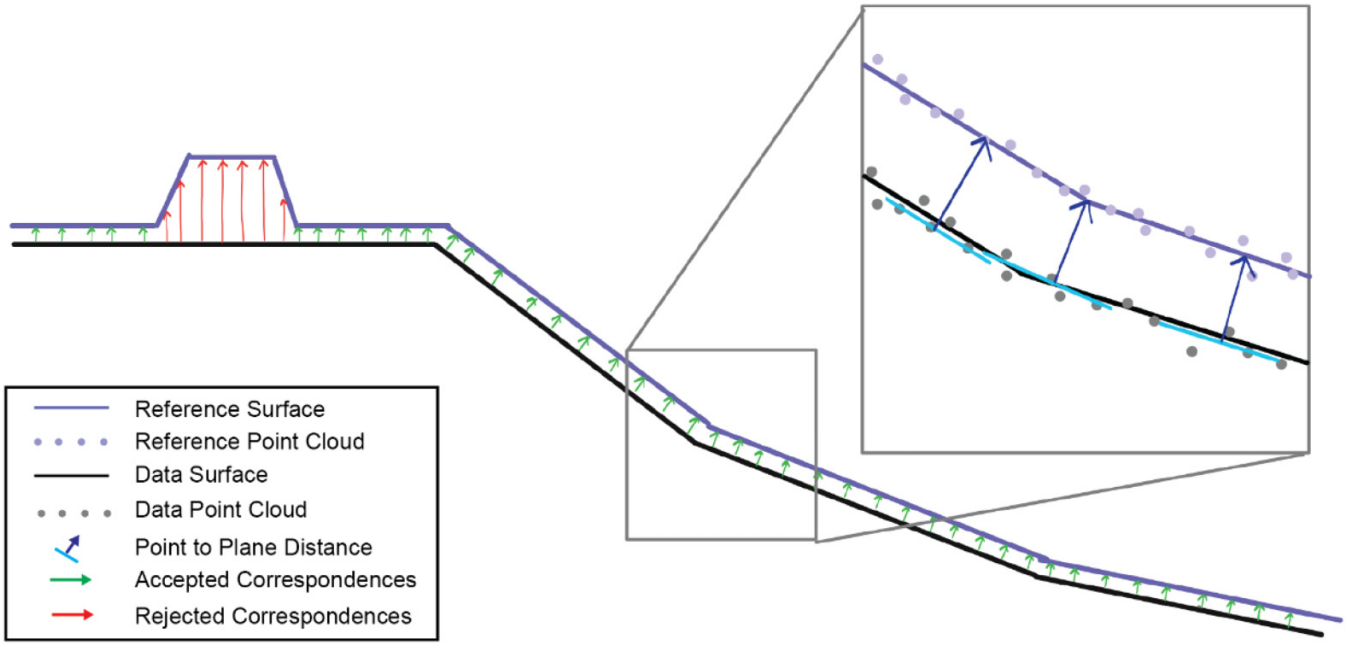

responding to large real changes on the slope are not used to estimate the geometric transformation.

Next, a least-squares geometric transform is calculated and applied (Arun et al. 1987). Finally, the convergence criteria are checked to see if the procedure should be stopped or continued iterating. We use the default convergence criteria settings in PCL.

3.1.2. M3C2 algorithm

M3C2 was introduced by Lague et al. (2013) to overcome the limitations of other point-cloud-based change detection algorithms. M3C2 operates in two main steps: (1) normal vector estimation for each point based on a set radius (R_N) of neighborhood points and (2) distance calculation along the normal vector direction. Points are collected from both data and reference clouds in a cylinder with a set length (L_C), an axis along the normal vector, and a set radius (R_C). Then, these points are projected along the cylinder axis, and the distance between the mean projected point locations for the two clouds is calculated. The main benefit of this approach is that by setting the radius parameters R_N and R_C appropriately, change can be calculated directly from the raw point clouds while also being robust to surface roughness arising from noisy data and complex natural environments.

The M3C2 algorithm requires three main parameters: normal radius (R_N), cylinder radius (R_C), and the cylinder length (L_C). By default, the cylinder length is set to 20 m in both forward and reverse normal directions from the core point. We use a normal radius of six times the average point spacing of the less dense point cloud (either the data or the reference clouds), or 3 m, whichever is larger. Similarly, the cylinder radius is set to the greater of three times the average point spacing of the less dense cloud or 1.5 m. Thus, in our implementation, $R_N = 2(R_C)$ such that the normal direction over flat ground is consistent and smooth.

The selection of the normal and cylinder radii is based on user judgement because while increasing the cylinder radius improves confidence in change results, it also increases the "smoothness" of change and could average out changes with a small spatial extent, such as rockfalls. A theoretical minimum radius that is equal to the point spacing is recommended to ensure that at least one point is captured in the search cylinder (Lague et al. 2013; DiFrancesco et al. 2020). This value is typical for terrestrial lidar and fine-scale rock slope analysis, but we opt to use values larger than this minimum based on the observation that typical landslides in aerial data are much larger than the point spacing, and higher confidence in change is desirable for detecting small amounts of movement over extensive areas. The cylinder length is selected to be larger than the maximum amount of expected change between two dates. A value of 20 m was found to be appropriate for most circumstances, with active mine sites a notable exception.

3.1.3. GPU implementation

We implemented the point cloud processing steps described above on GPU hardware to improve their speed. Point cloud processing requires the identification of various neighborhoods of points for plane estimation, locating ICP correspondences, and computing M3C2 distances within the search volume. K-dimensional (KD) tree and octree algorithms are commonly used to perform fast nearest neighbor searches on point clouds, and these are now commonly implemented on GPUs (Girardeau-Montaut et al. 2005; Zhou

et al. 2008; Richter et al. 2013). Tree algorithms enable efficient searches through arbitrary 3D space and are widely used for computer graphics applications, but there is some computational cost associated with the initial generation of the search tree structure. In addition to implementing M3C2 on a GPU, the two major differences of our approach from previous algorithms are summarized as follows:

1. Nearest neighbor searches are conducted using a two-dimensional grid generated for each search query.
2. The entire point cloud is loaded in GPU memory and native platform application programming interfaces (APIs), and GPU compute shaders are used to control all aspects of GPU core operations (as opposed to Nvidia CUDA, for example).

3.1.4. GPU-based neighborhood searches

Because geohazard mapping deals primarily with data collected aerially over large areas, the overall structure of the point cloud can be considered planar on average. As such, we opt to create uniform, horizontal, non-recursive grids to subdivide point clouds for neighborhood searches, instead of spending more computing resources creating a KD tree or octree structure for the entire cloud. The grids generated using this method are two dimensional (i.e., they extend infinitely in the Z direction) and are generated on-the-fly using a grid cell size that is dependent on the size of the neighborhood search needing to be performed. For spherical neighborhood searches, the grid side length is equal to the diameter of the sphere. For M3C2 cylinder searches, the side length is computed such that it contains the horizontal projection of the entire cylinder. Because the grid is generated for each search and is necessarily larger than the desired neighborhood, all points belonging to the neighborhood will be correctly identified in a single grid cell (i.e., no points will be missed). For example, if a GPU core needs to perform a spherical neighborhood search to estimate the normal vector, a grid is created centered on the query point, the returned points within the grid cell are sent to the GPU core, and then points outside the search sphere radius are rejected.

The use of an optimized grid search method entails that the entire point cloud segment is loaded into GPU memory. To enable more control of the data layout and data transfer to and from GPU memory, we used low-level graphics APIs, specifically Direct3D for Windows and Metal for iOS, to perform all calculations on GPU cores. As such, our implementation is not tied to specific hardware and can be deployed to multiple devices, including smartphones (Lato and Ferrier 2022).

3.2. Comparison tests

We performed a set of four experiments for each site to quantify the improvement in change precision obtained by using ICP-M3C2. Specifically, we compared M3C2 change calculation with DoD, and we compared applying or not applying ICP alignment. These options result in four possible configurations, as illustrated in Fig. 5.

Fig. 5. Schematic showing the four combinations of analysis steps performed in this study. Tests are assembled based on two different change detection methods (DoD and M3C2) and the presence or absence of the ICP alignment.

Change calculation method	Perform ICP? No	Perform ICP? Yes
	1. DoD*	2. DoD* w/ ICP
	3. M3C2**	4. M3C2** w/ ICP

Notes: * DoD implemented using Global Mapper
** M3C2 using novel GPU implementation

The DoD analysis consists of two main steps: conversion of the reference and data point clouds into a raster format and vertical subtraction of these two datasets in a regular grid. We performed these steps using Global Mapper (Blue Marble Geographics 2020). First, the raw point clouds are converted to triangulated irregular networks (TINs). There are many alternatives for creating a DEM from raw points, including various deterministic and probabilistic interpolation methods and gridding options, and there is not necessarily a single method which is expected to work best in all environments (Fisher and Tate 2006; Chen et al. 2017). We use TINs for DEM creation primarily because of their high efficiency, which is important when processing hundreds of square kilometeres of data, and because they are a commonly used approach in commercial software and are therefore useful for comparison. The TINs are then used to populate the raster DEM at a small cell size (less than 0.2 m) estimated by Global Mapper. A grid of resolution 0.5 m is then placed over the two DEMs, and the reference DEM elevation is subtracted from the data DEM elevation at each grid cell. A 0.5 m grid size represents a standard DEM size for high-resolution aerial lidar. A grid resolution of 0.1 m was also tested to check whether grid size would influence the comparison with point-cloud-based change, but the change statistics were substantially similar, so a 0.5 m grid was used to manage file size.

3.2.1. Speed comparison

To compare CPU- and GPU-based M3C2 computation, we ran non-ICP change detection using both our GPU implementation and the implementation provided in the open-source CloudCompare software (cloudcompare.org). The self-reported change computation times for both methods were recorded in milliseconds (both methods print the change detection time to the console). Our GPU method timing includes both the raw GPU computation time and the GPU "prologue" which consists primarily of copying data from CPU memory to GPU memory. The CloudCompare implementation was parameterized to make it comparable to our software: Normal

vectors and change were calculated using a fixed search radius of the same size as used in the GPU version, and the cylinder length was set to the same fixed length of 20 m in both tests. The option in CloudCompare to perform a progressive cylinder search, which extends the cylinder until a certain number of points are found, was disabled to be comparable to our implementation. These tests were performed on two different computer systems to evaluate the potential differences in the speed observed with different computer hardware. Computer 1 had a Nvidia Quadro RTX 5000 graphics cards (16 GB dedicated memory) and a 16 core, 32 thread Intel Xeon Gold 6226R CPU. Computer 2 had a Nvidia RTX A5000 graphics card (24 GB dedicated memory) and a 24 core, 48 thread AMD Ryzen Threadripper PRO 5965WX CPU.

Speed measurements were recorded for DoD and ICP steps on Computer 1 to provide additional context for the results of this study. DoD was computed as described above in Section 3.2 using Global Mapper. DoD times are meant to represent a typical example of what can be expected using common commercial software. We also timed individual runs of ICP for sites 1 through 5 as an estimate of the typical alignment runtime. However, we note that our ICP implementation runs with multiple tiles in parallel, and the time taken is highly dependent on the initial positions of the point clouds and how long they take to reach convergence.

A final speed test was performed on Computer 2 to demonstrate the combined system, which included data reading/writing, ICP alignment, and normal vector computation, M3C2 distance computation, and the rasterization of the outputs for convenient viewing in a web-based map. Site 6 was used for this comparison, as mentioned in Section 2, with a total combined number of points of 5.99 billion. The data were first divided into 562 segments and placed into separate folders for the baseline and comparison datasets. The ICP-M3C2 software was run, which loaded segment baseline/comparison pairs, estimated the point density, computed the ICP alignment and recorded the convergence result, performed M3C2 comparison, and wrote the output files to the specified folder.

3.3. Limit of detection

Assessing the confidence bounds of change is an important step in using change detection for earth science applications. For example, uncertainty assessment has a major impact on volume calculations (Wheaton et al. 2010) and can influence the detectability of slow-moving landslides (Lato et al. 2019). Assessment of change confidence and error is commonly reported as a limit (or level) of detection (LoD), analogous to a 95% "confidence" interval ($z = 1.96$). Many authors estimate LoD as twice the standard deviation of change computed in regions assumed to be unchanged (Abellán et al. 2009; Kromer et al. 2017; Williams et al. 2018). If a change value is within the 95% confidence interval (e.g., $\text{LoD}_{95\%} = \pm 0.05$ m), it is not considered for volume calculation or assessment of geomorphic processes. Several different methods have been proposed to estimate $\text{LoD}_{95\%}$, including different statistical tests, the inclusion of spatial variations in uncertainty, and error propagation of multiple datasets (Joerg et al. 2012; Lague et al. 2013; Winiwarter et al. 2021; Zahs et al. 2022).

In this study, we estimate a preliminary LoD by computing change statistics (mean, standard deviation, 2.5th and 97.5th percentiles) in a manually selected small region of each site which is morphologically similar (e.g., in terms of slope and curvature) to the changing area, but is away from areas of known or suspected change. The summary statistics are used to select a threshold for visualizing change results, where change values below the threshold are completely gray/transparent to the underlying hillshade. While the underlying statistics are functions of the raw data, the threshold chosen for visualization is not strictly data-driven and may in practice be chosen to be more conservative (i.e., a larger limit of detection and larger changes filtered) based on a worst-case scenario or the largest of multiple estimates. In this study, the thresholds were chosen to generally encompass the 2.5th and 97.5th percentiles of the data, while also highlighting visual differences in appearance between the different test results.

4. Results

4.1. Change statistics

Topographic change was computed for the datasets and analysis methods described in Sections 2 and 3, and the results are compared quantitatively using histograms and change statistics in Fig. 6 and Table 2.

Histograms in Fig. 6 illustrate the differences in change values obtained using DoD, M3C2, DoD with ICP alignment, and M3C2 with ICP alignment. The following two observations can be made for every site: (1) M3C2 has a smaller difference between the 2.5th and 97.5th percentiles than DoD (and corresponding smaller standard deviations) and (2) ICP generally, but not always, shifts the mean of the distribution closer to zero. Further, for the sites at which steeper slopes were used to calculate statistics (panels B and C in Fig. 6), a change in the shape of the distribution was observed depending on whether ICP was used or not. For example, the non-ICP M3C2 distribution at Site 2 (panel B, yellow curve) has two distinct peaks, but once ICP is applied, only a single high-density peak is observed. A similar pattern can be observed at Site 3 (panel C).

4.2. Visual interpretation of change maps

Change maps of results for five sites and four analysis methods are presented in Figs. 7 through 11. Areas of negative model differences (material loss, subsidence, erosion, or removal) are colored light to dark blue, while areas of positive model differences (material accumulation, bulging, or placement) are colored yellow to red. No color is shown if changes are less than the $\text{LoD}_{95\%}$.

Change maps for Site 1 are presented in Fig. 7. The map highlights a series of recently active landslides along a road and larger landslide to the southeast of the road. While the headscarp of the larger landslide has dropped as much as 0.75 m in the analysis period, the toe of the landslide shows positive displacement magnitudes of only between 0.1 and 0.2 m, which is of a similar magnitude to the initial georeferenc-

Fig. 6. Comparison of histograms of change calculated for each of the five sites (panels A through E) and four analysis methods (different colored lines). Colored and dashed vertical lines indicate the 2.5th and 97.5th percentiles for each histogram. Statistics are calculated in an area of each point cloud where no geomorphic change was observed or interpreted.

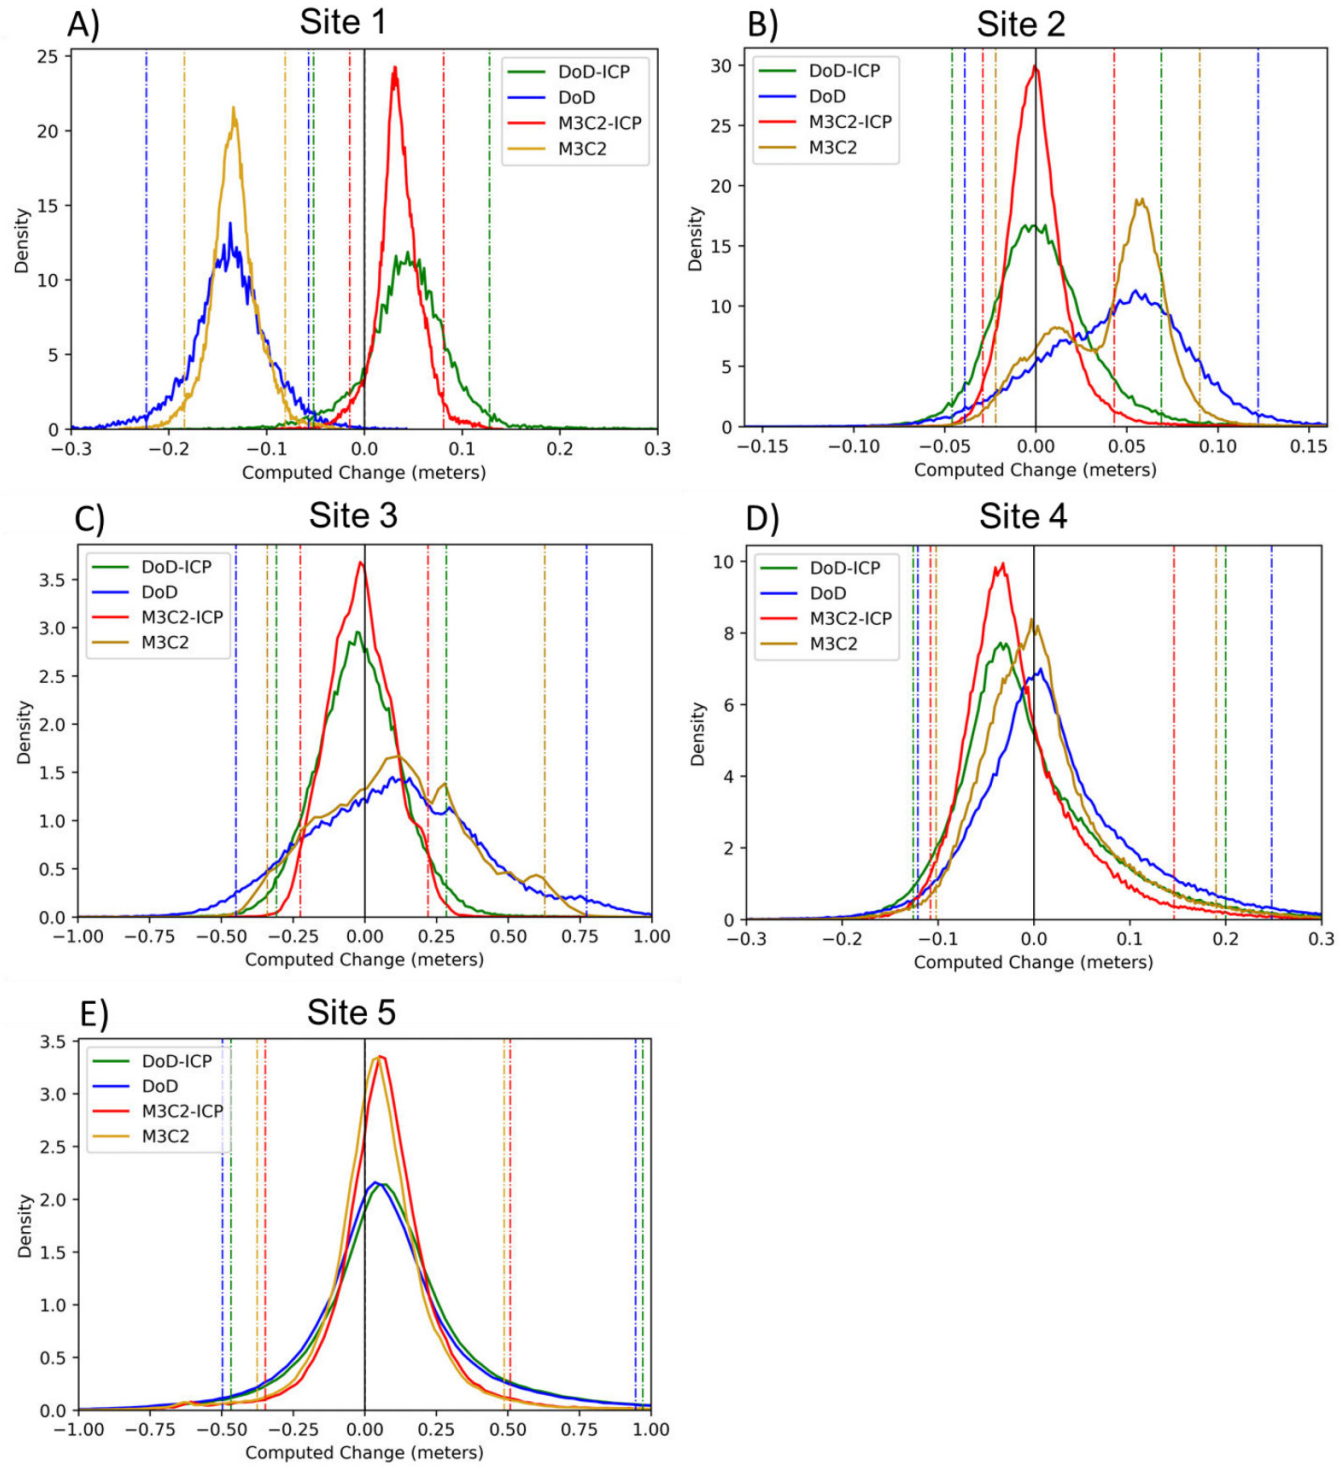

ing error between the two component datasets. Georeferencing error results in a larger $LoD_{95\%}$ for this comparison, and as a result, changes calculated with non-ICP aligned datasets are unable to resolve the toe of the landslide. The identification of distinct movement at both the scarp and toe of the landslide suggests the existence of a deeper rotational failure than would be clear if movement was only measurable at the scarp.

A similar situation occurs at Site 4 (Fig. 8), where the application of ICP changes both the magnitude and distribution of movement observed in this valley-scale landslide. In contrast to Site 1, however, the application of ICP shifts empha-

Table 2. Percentile change statistics for the four analysis types and five regions of interest.

Site	Test	Mean	SD	P 2.5%	P 97.5%	Range*	Change threshold for visualization
1	DoD-ICP	0.04	0.05	−0.05	0.13	0.18	±0.10
	DOD	−0.14	0.04	−0.22	−0.06	0.17	±0.30
	M3C2-ICP	0.03	**0.02**	−0.02	0.08	**0.10**	±0.10
	M3C2	−0.13	0.03	−0.18	−0.08	**0.10**	±0.30
2	DoD-ICP	0.00	0.03	−0.05	0.07	0.12	−0.05 to + 0.10
	DOD	0.04	0.04	−0.04	0.12	0.16	−0.05 to + 0.15
	M3C2-ICP	0.00	**0.02**	−0.03	0.04	**0.07**	−0.05 to + 0.10
	M3C2	0.04	0.03	−0.02	0.09	0.11	−0.05 to + 0.10
3	DoD-ICP	−0.02	0.16	−0.31	0.28	0.59	±0.30
	DOD	0.11	0.31	−0.45	0.77	1.22	−0.50 to + 0.80
	M3C2-ICP	−0.01	**0.12**	−0.23	0.22	**0.45**	±0.30
	M3C2	0.11	0.25	−0.34	0.63	0.97	−0.40 to + 0.70
4	DoD-ICP	0.00	0.08	−0.13	0.20	0.33	±0.15
	DOD	0.03	0.09	−0.12	0.25	0.37	±0.15
	M3C2-ICP	−0.01	**0.06**	−0.11	0.15	**0.25**	±0.15
	M3C2	0.01	0.07	−0.10	0.19	0.29	±0.15
5	DoD-ICP	0.12	0.36	−0.47	0.97	1.44	±0.90
	DOD	0.10	0.36	−0.50	0.95	1.44	±0.90
	M3C2-ICP	0.07	**0.23**	−0.35	0.51	**0.86**	±0.90
	M3C2	0.05	**0.23**	−0.38	0.49	**0.86**	±0.90

Note: All values are reported in meters; Statistics are calculated in an area of each point cloud where no geomorphic change was observed or interpreted. Bold values indicate the method with the lowest standard deviation and range for each site.
*Range is calculated as the difference between the 97.5th and 2.5th percentiles.

sis away from the toe of the landslide and toward the scarp, suggesting larger amounts of subsidence at the railroad grade and lower magnitudes of change at the toe. This is a situation where lower magnitudes of change may be an indication of predominantly translational movement which is parallel to the slope and therefore not able to be detected with an orthogonal change direction.

Differences between DoD and M3C2 change calculation methods are highlighted in Fig. 9, where an existing landslide shows recent movement. The DoD results are noticeably more speckled in appearance compared with M3C2 results with the same $LoD_{95\%}$. A second difference is that DoD is susceptible to artifacts in regions of steep topography: a sub-vertical scarp at the head of the landslide is reported as having moved upward (positive change) by more than 1.0 m in the DoD results, but the artifact is significantly reduced in the M3C2 results.

Figure 10 highlights a case where georeferencing error between the two original datasets was sometimes greater than 0.6 m, resulting in a relatively poor (higher) limit of detection. With the inclusion of ICP, change values above the main landslide become centered around zero and have a much smaller range compared to non-ICP DoD and M3C2. While the main landslide area appears visually similar in all four methods due to the large change magnitudes relative to the $LoD_{95\%}$, ICP alignment revealed 0.5–1.0 m of settlement in the railroad right of way that was within the limit of detection in non-ICP methods.

Figure 11 illustrates change detection results for Site 5 (Oso, Washington) between 2006 and 2013, prior to the catastrophic 2014 landslide. The error histograms for this site (Fig. 6, panel E) suggest that the baseline and comparison datasets were already closely aligned to each other relative to the overall error, and therefore the effect of ICP on results interpretation is minimal. The detail panels in Fig. 11 highlight differences between DoD and M3C2 results. The amount of speckled noise on steep slopes is markedly reduced in M3C2 compared to DoD, allowing for more confident and rapid identification of areas of erosion and deposition on the cliff face.

4.3. Speed comparison of CPU and GPU change calculation

CPU- and GPU-based M3C2 computation times for both test machines are compared in Table 3. The "speed-up factor" is the ratio of the compute times from CloudCompare to the GPU method. Note that the improvement in speed is dependent whether the "GPU prologue" is included in the total GPU compute time or not. Because the GPU prologue consists primarily of transferring point cloud data from the CPU to the GPU, it can be argued that a fair comparison of the GPU and CPU methods would not include the GPU prologue. However, a CPU-based method would not need to perform this data transfer. Consequently, we have reported speed-up factors for both including and excluding the GPU data transfer time. The speed-up factors varied for different sites and between the two computer systems. Speed-up factors for Computer 2 were relatively lower than for Computer 1, meaning that the relative processing capabilities of the CPU method were higher for Computer 2. This is consistent with expectations, since the CPU in Computer 2 had 16 more threads than the CPU in

Fig. 7. Change maps for Site 1, with each panel corresponding to a change calculation method. Change values above the $LoD_{95\%}$ are overlain on a hillshade rendering of the more recent DEM. $LoD_{95\%}$ values in the panel titles indicate the level of change below which change values are not shown. Red circled area highlights the toe of an active landslide which is only identifiable when ICP is applied. Dashed black box indicates region used for the calculation of the limit of detection.

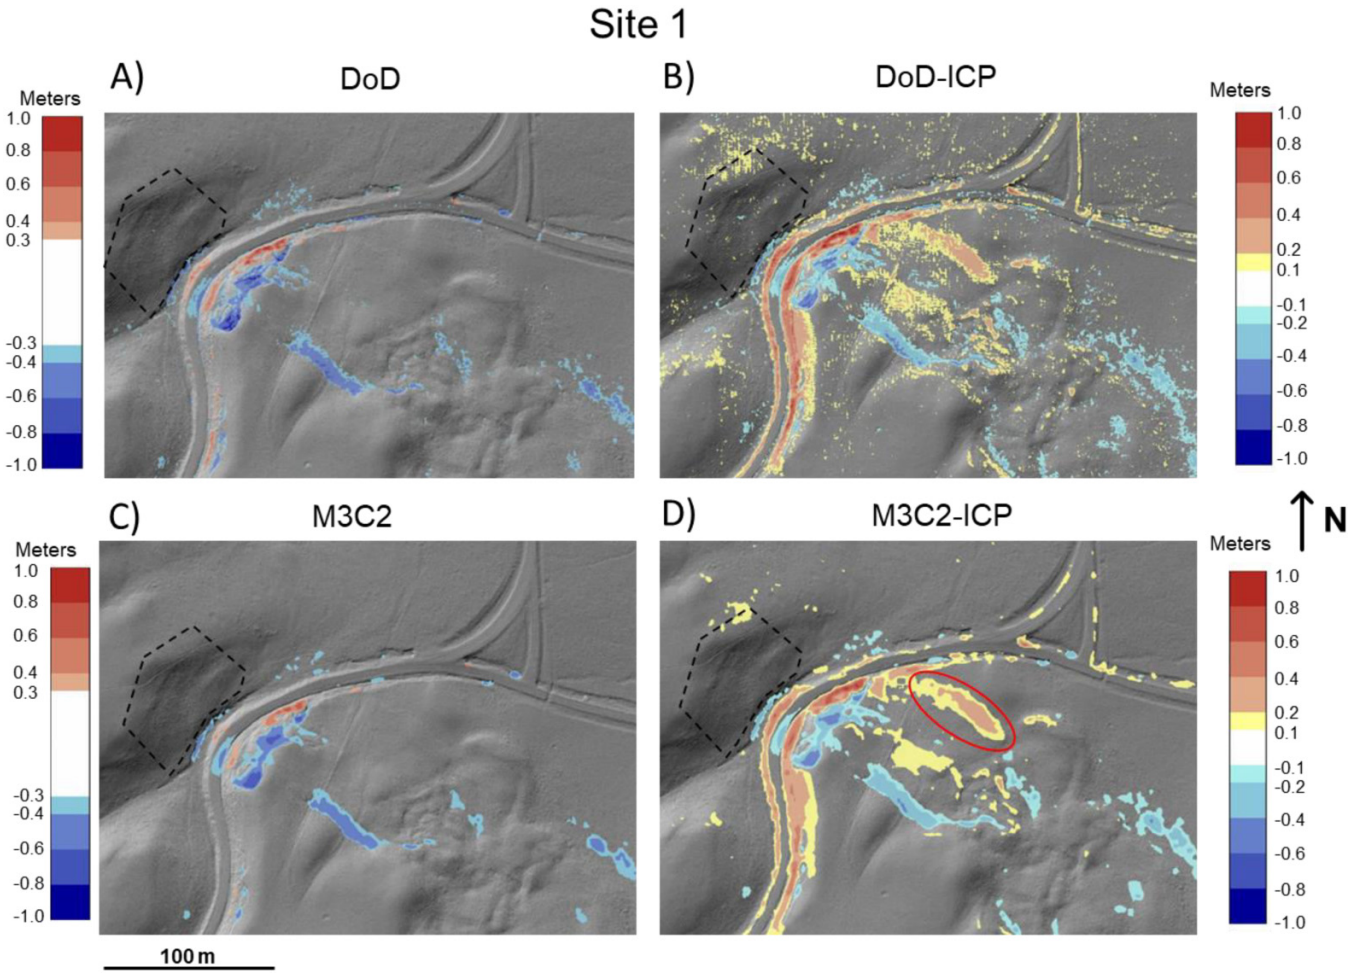

Computer 1, meaning that more parallel operations could be performed. Despite this, the GPU method compute time was still 9–46 times faster than the CPU.

Analysis of dataset characteristics with the speed results in Table 3 suggests a relationship between point density and computation time. Figure 12 illustrates the relationship observed between the point density of the comparison dataset and the M3C2 compute time, normalized by the aerial extend of the data (Tables 1 and 3). As indicated by the log-linear best-fit lines, the GPU method appears to be much less sensitive to point density compared to the CPU method. While this does indicate some influence of density on computation time, we note that several other factors likely play a role, including the density of the baseline cloud, the M3C2 search radii, and the steepness/classification quality of the terrain. The datasets used in this study were similar to each other but not identical in terms of number of points and search radii.

The large-scale speed test using Site 6 was also completed. The total time to complete all processing steps for 5.99 billion points, including ICP, change computation, reading, and writing, was ~3.2 h, or ~20 s/segment. Approximately 15 min of this time required user intervention for copying files and setting up the initial parameters, and the rest was running automatically in the background.

4.4. DoD and ICP speed

This section describes the results of the DoD and ICP speed tests, as presented in Table 4. For reference, we include both time taken to perform the gridding overlay and raster subtraction (differencing) and the initial creation of the TIN from the raw point cloud. In comparison to M3C2 (Table 3), the subtraction step of DoD was measured to be faster than the CPU change computation in most cases but slower than GPU change computation. We do not interpret this to suggest that GPU-M3C2 is faster than DoD in general, as we cannot confirm that Global Mapper is optimized similarly to CloudCompare and GPU-M3C2. Based on the number of computations required for each method, it is presumed that a GPU-based DoD computation would outperform M3C2, for example, but this is beyond the scope of the current study. However, these results suggest GPU acceleration may be useful for closing the

Fig. 8. Change maps for the Site 4, with each panel corresponding to a change calculation method. Change values above the $LoD_{95\%}$ are overlain on a hillshade rendering of the more recent digital elevation model (DEM). $LoD_{95\%}$ values in the panel titles indicate the level of change below which change values are not shown. Black dashed box indicates region used to calculate the limit of detection (region extends ~100 m to the north). DoD methods are not shown due to their visual similarity to the M3C2 results and to enlarge individual panels but are included in Fig. S1.

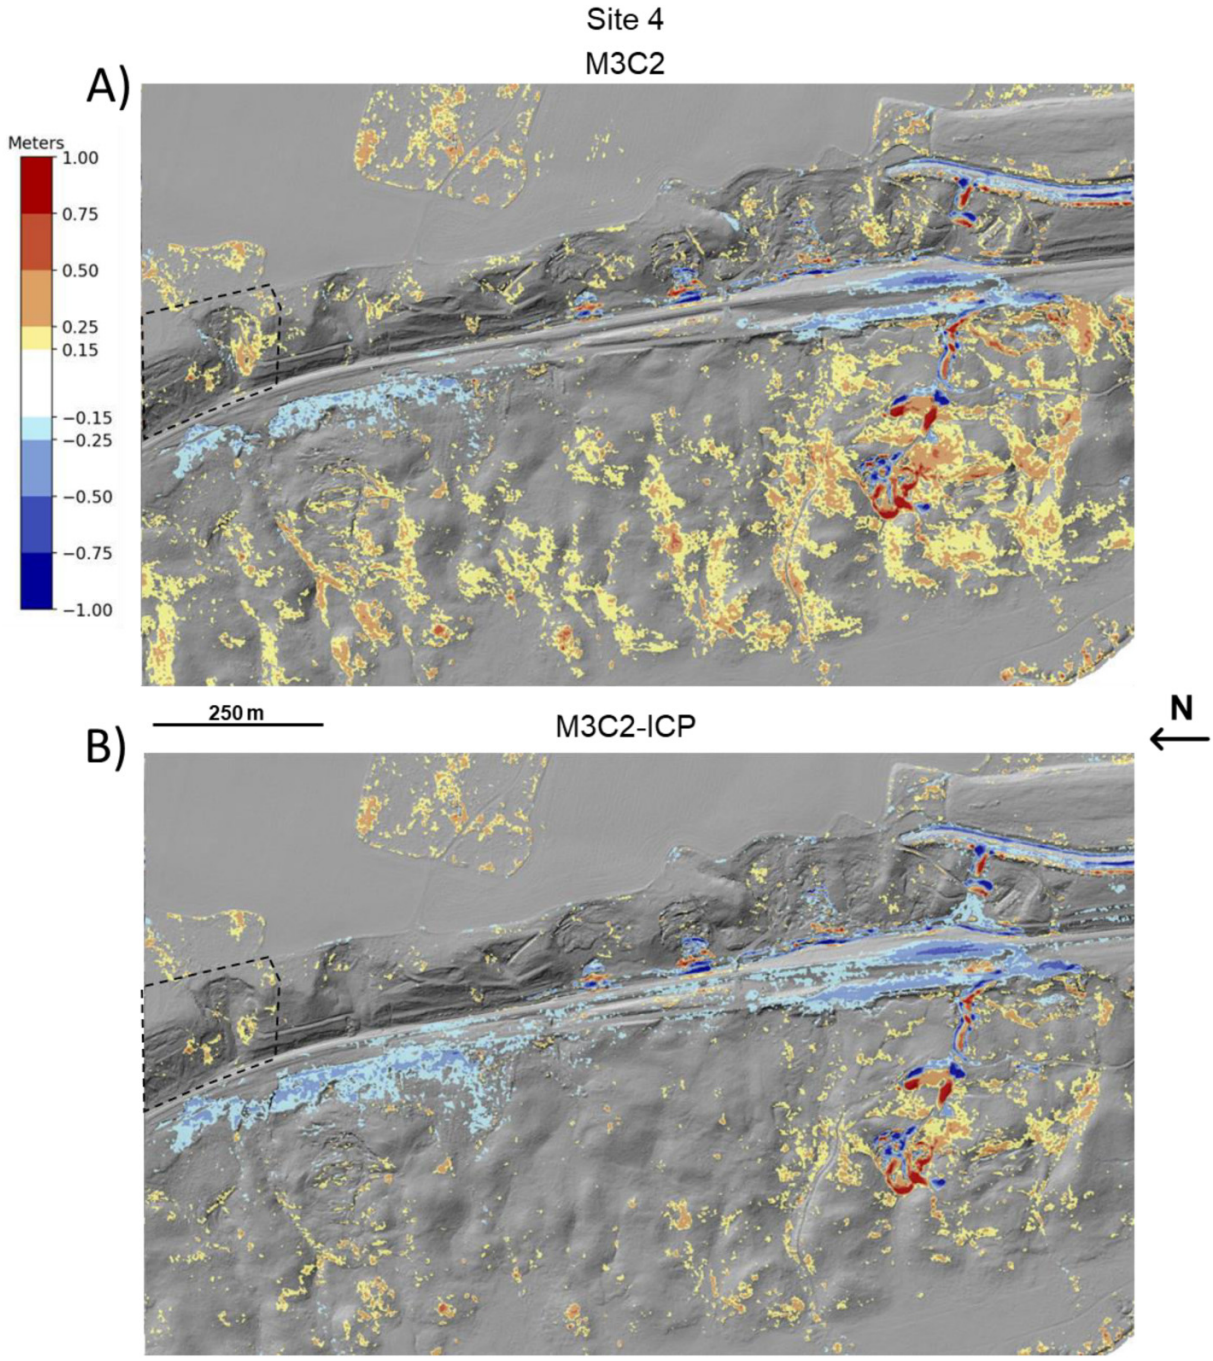

gap in speed between point-based and raster-based analysis methods.

5. Discussion

5.1. Comparison of DoD and M3C2

The results presented in this study highlight several limitations of the DoD method compared to direct point cloud comparison with spatial averaging. First, DoD computes change at grid nodes at a resolution that is necessarily less than the native point resolution, but M3C2 computes change based on an average of neighborhood raw points. In our tests, this resulted in an improvement (lowering) of the limit of detection, without reducing the resolution of results (see Fig. 6 and Table 2). In practice, simple DoD results could be "enhanced" by applying smoothing based on a neighborhood of grid cells (Abellán et al. 2009), but this neighborhood search would be

Fig. 9. Change maps for Site 2, with each panel corresponding to a change calculation method. Change values above the $LoD_{95\%}$ are overlain on a hillshade rendering of the more recent digital elevation model (DEM). $LoD_{95\%}$ values in the panel titles indicate the level of change below which change values are not shown. Black circled area highlights an artifact in difference of DEM (DoD) change results caused by steep topography. Black dashed box indicates regions used to calculate the limit of detection.

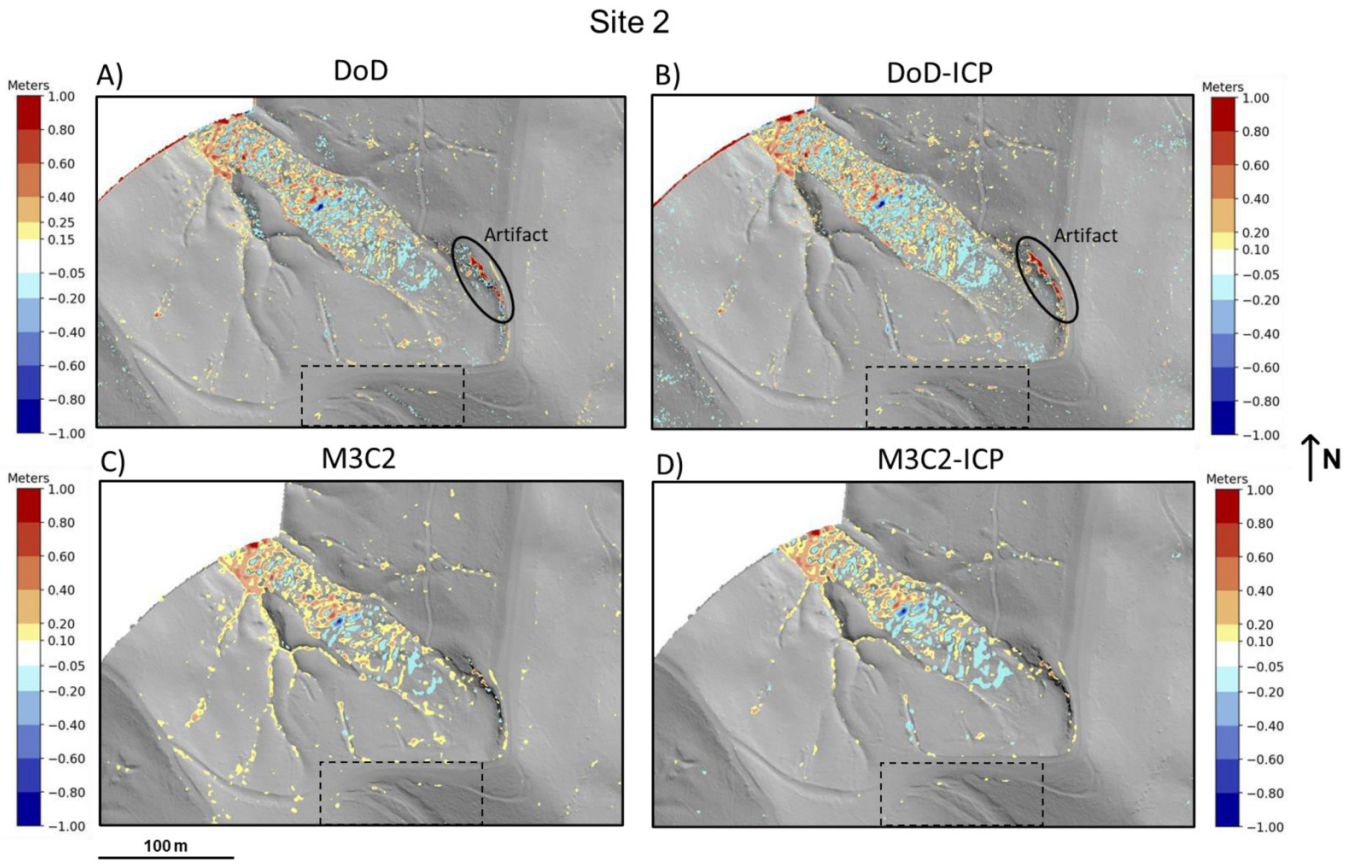

based on 2D proximity and would not preserve change in steep slopes.

The practical differences between DoD and M3C2 results depend on the precision and overall quality of the input datasets. It is well known that in complex topography, DoD is unlikely to accurately represent 3D geohazard processes and that change in the orthogonal direction is more relevant (Lague et al. 2013; DiFrancesco et al. 2020; Williams et al. 2021). In our experience, the improvements (reduction) in the limit of detection afforded by M3C2 also enables detection of subtler landslide movements (<15 cm) with higher confidence due to the reduction in ambient noise. However, if the quality and precision of both baseline and comparison datasets are high, and only large changes are of interest, it is possible that the improvement in limit of detection will have a relatively small impact on interpretation of geohazard processes and risks. Similarly for change direction, if the topography of interest consists of only relatively smooth, low-angle surfaces, the change direction computed using normal-vector-based M3C2 would be relatively close to vertical. Generally, we are interested in environments prone to geohazards, which are more likely to have complex topography and datasets of variable quality.

5.2. Use of ICP for error reduction in airborne lidar

This study also highlights the added value of ICP in reducing errors in airborne lidar change detection. However, several assumptions are made when applying ICP over such large scales, and these assumptions must be checked regularly. First, ICP assumes that the two point clouds being compared are already approximately aligned to each other, an assumption that often cannot be made for terrestrial lidar, for example (Schovanec et al. 2021). Without the pre-existing approximate alignment, ICP would be much more likely to converge to a "local minimum" solution which is not near to the actual optimal solution. In our experience, however, georeferencing errors are typically small enough to avoid this except in rare cases.

ICP also assumes that areas that have changed in the analysis period are small relative to the total size of the region of interest. Large-scale changes, such as underground mining subsidence, fault ruptures, open pit mining excavations, changes in water level, or agricultural changes, all have the potential to violate this assumption. In our experience, the 70% rejection criterion can overcome some of these situations, but others continue to pose challenges.

Fig. 10. Change maps for Site 3, with each panel corresponding to a change calculation method. Change values above the $LoD_{95\%}$ are overlain on a hillshade rendering of the more recent digital elevation model (DEM). $LoD_{95\%}$ values in the panel titles indicate the level of change below which change values are not shown. Dashed black box in the M3C2 panel indicates the region of points used to calculate statistics in Fig. 6. Black arrow indicates the location of settlement and erosion identified in panels B and D.

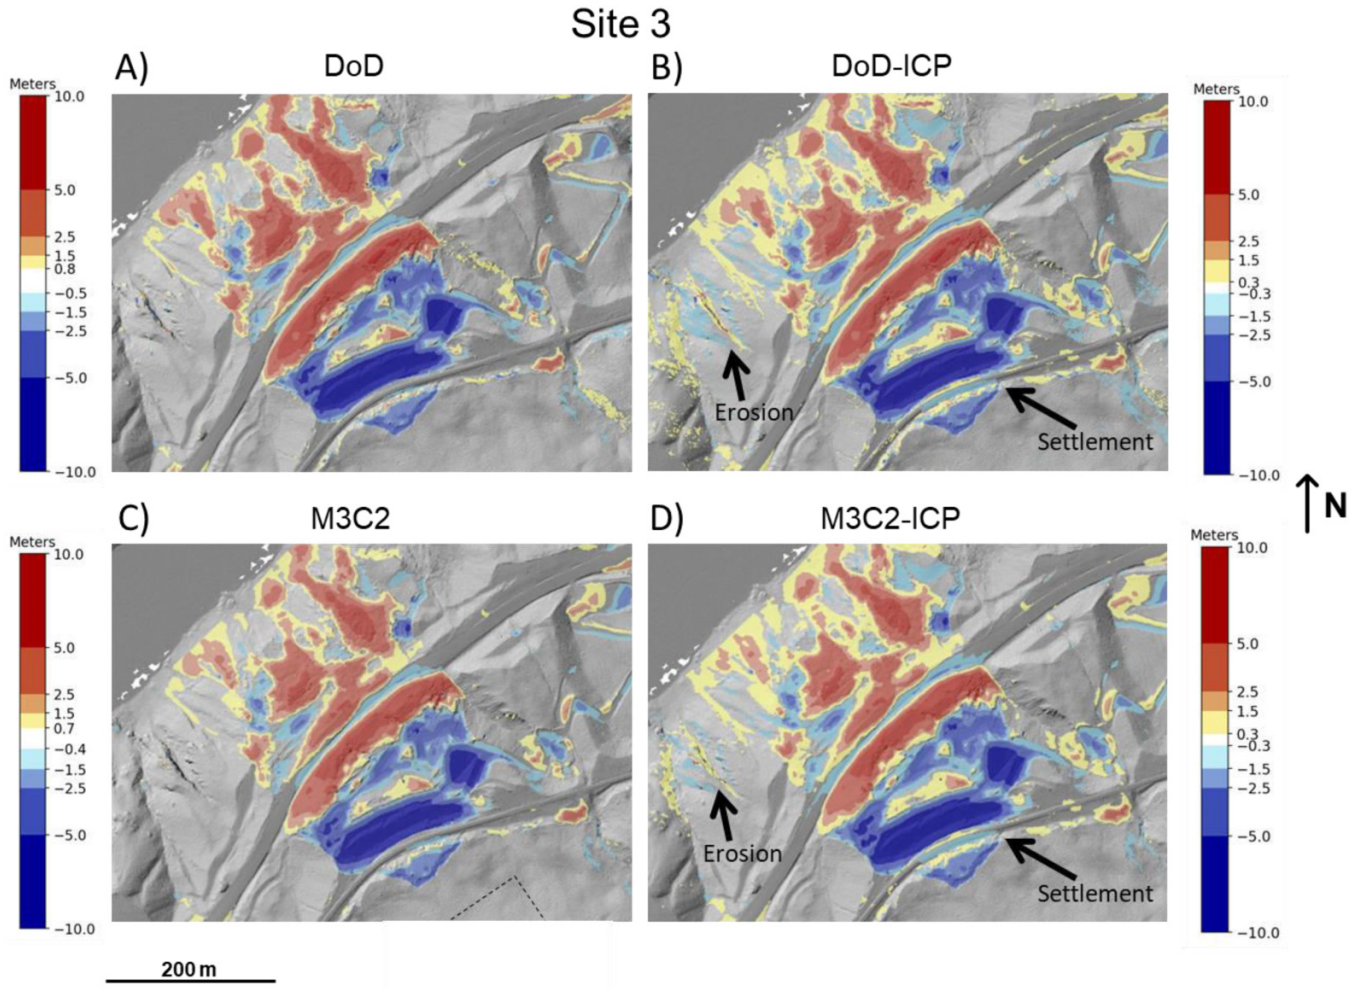

For example, the benefit of applying ICP versus no ICP is somewhat ambiguous in the change detection results presented in Fig. 8. As observed in Fig. 6D, non-ICP histograms are centered on zero, but once ICP is applied, the histograms shift in a negative direction. This negative shift could be caused by millimeter- to centimeter-scale subsidence occurring along the entire slope length, but it could also be caused by a change in the elevation from crop harvesting and tilling in adjacent farmlands. Overall, the benefits of wide application of ICP are clear, but the potential for occasional complications should be acknowledged.

5.3. Adoption of 3D point cloud processing in engineering

The data presented in Fig. 12 imply that using the existing CPU-based point cloud processing may not scale effectively as data become increasingly dense and, more importantly, are collected over larger areas. The GPU-based methodology presented here is a promising alternative in this case. Regardless of the specific hardware or software, however, developing more automated and efficient data processing workflows in general would lower the barrier to entry to state-of-the-art methods and would be more cost-effective for engineers to use in practice. It is important to note that this does not eliminate the need for experience and domain knowledge in interpreting geomorphic change. Future engineers will need to be trained specifically in the interpretation of point-cloud-based change detection as it is a major departure from raster-based methods.

6. Conclusion

Geohazards present major risks to critical linear infrastructure in North America, such as pipelines, roads, and rail networks. Recent high-profile disasters highlight the need for engineers to take a proactive approach to manage these risks over increasingly large areas, and high-precision lidar change detection is an essential tool in this space. The 2021 British

Fig. 11. Change maps for Site 5, with the upper panel (A) corresponding to the difference of DEM-iterative closest point (DoD-ICP) method and the lower panel (B) corresponding to the multiscale model to model cloud comparison (M3C2)-ICP method. Detail panels for each method highlight differences between DoD and M3C2 for a specific area. Non-ICP methods are not shown due to the small influence of ICP at this site. Dashed box indicates the region used to calculate the limit of detection.

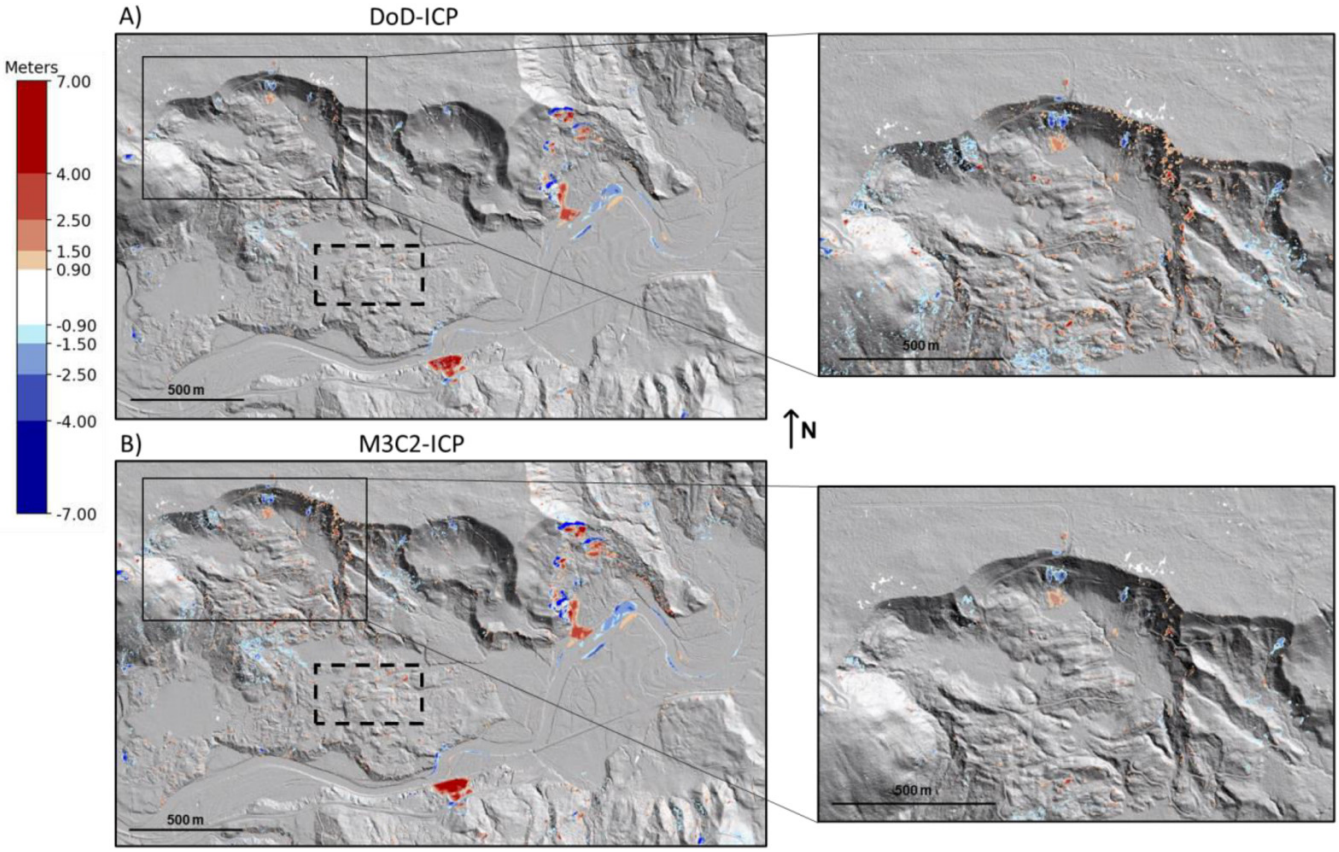

Table 3. Comparison of runtimes in seconds reported for calculation of change using M3C2. The CPU version is run using CloudCompare, while the GPU version is run using our implementation.

Site number	GPU prologue (s)	GPU compute (s)	CPU (s)	Speed-up factor (with prologue)	Speed-up factor (without prologue)
			Computer 1		
1	1.1	1.4	73.1	30	54
2	1.8	1.3	40.0	13	30
3	1.1	1.1	33.0	15	30
4	0.5	0.4	8.9	10	22
5	0.8	0.7	31.5	21	46
			Computer 2		
1	0.7	0.9	42.3	26	46
2	1.0	0.8	14.3	8	17
3	0.8	0.7	11.9	8	16
4	0.4	0.4	3.3	5	9
5	0.4	0.5	9.8	11	22

Columbia floods are a key example where rapid turn-around, regional change detection provided immense value for disaster response, recovery, and resilient redesign (Hunter 2022). However, the geotechnical industry is lagging the state of the art in change detection processing, resulting in the delivery of lower precision, noisy data products which hamper interpretation.

While M3C2-based change detection is gaining popularity, our experience indicates that traditional non-ICP, raster-based change detection is still widely used in practice (e.g.,

Fig. 12. Comparison of point density and multiscale model to model cloud comparison (M3C2) compute time for central processing unit (CPU) and graphics processing unit (GPU) methods for both computers. Compute times are normalized by the plan area of each dataset to account for differences in the total compute time by dataset size.

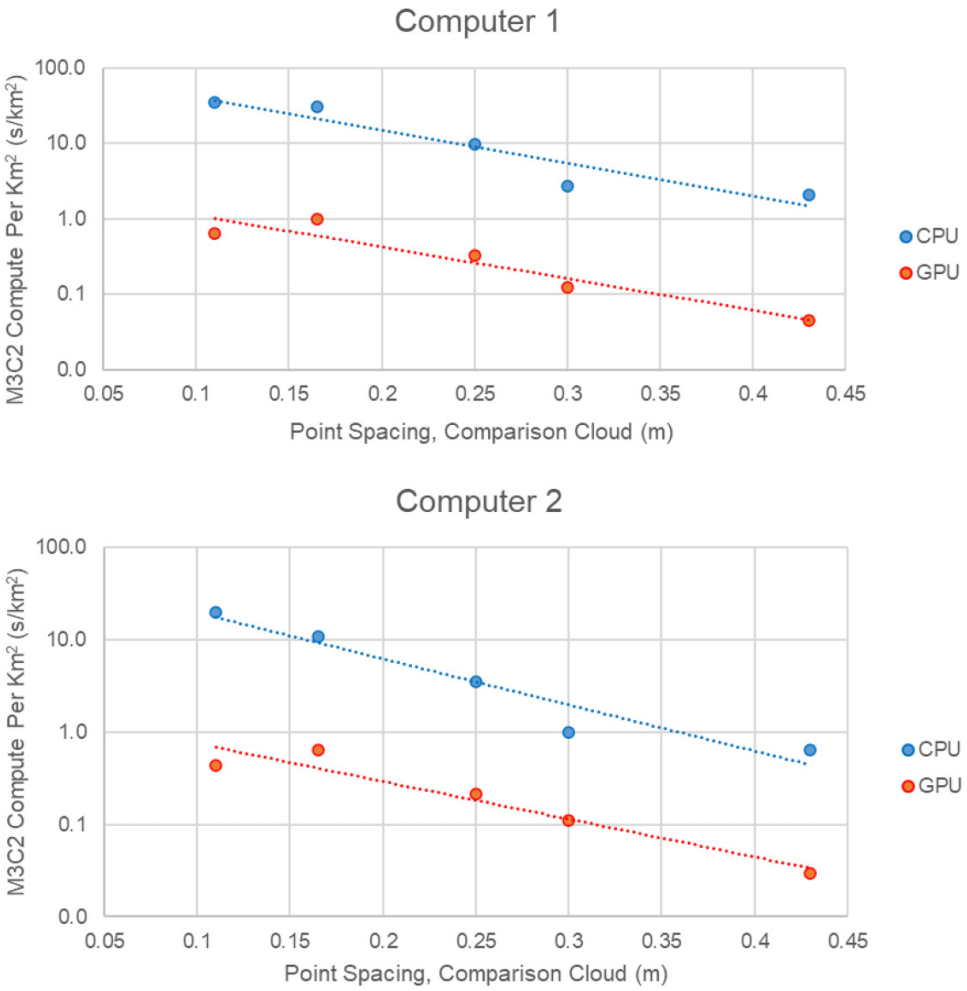

Table 4. Speed results for the difference of DEM (DoD) in seconds using Computer 1.

Site number	ICP (s)	TIN creation, cloud 1 (s)	TIN creation, cloud 2 (s)	Differencing (s)
1	117	17	254	5
2	547	219	134	8
3	365	20	140	6
4	167	34	34	9
5	128	19	58	26

Piovan et al. 2023). Our results demonstrate that the tools now exist for the broader adoption of direct point-cloud-based change detection processing and that this can be performed at the regional scale. The major conclusions of this study are summarized as follows:

- GPU-based algorithms show promise in accelerating change calculation, with our results showing a speed-up factor of up to 54× compared to a CPU-based algorithm. Continued optimization of GPU algorithms in the future will likely improve these speeds.
- M3C2 improves (reduces) the limit of detection, meaning that smaller changes can be reliably detected, and it reduces artifacts compared to DoD, even in generally low-gradient topography.
- ICP reduces systematic georeferencing errors when applied at regional scales while also preserving many types of large geomorphic and anthropogenic changes.

The results of this study have direct implications for quantitative risk management applications. First, applying more robust computational methods means that one can have

higher confidence in smaller amounts of geomorphic change, which can lead to more informed decision-making. Second, large-scale change detection enables decision makers to make more proactive choices, discovering problems at an earlier stage in their evolution.

Article information

History dates

Received: 14 February 2023
Accepted: 14 September 2023
Accepted manuscript online: 4 October 2023
Version of record online: 8 March 2024

Copyright

Data availability

Point cloud data for Site 5 (Oso, Washington area) are publicly available to download from the Washington Department of Natural Resources lidar portal (lidarportal.dnr.wa.gov). Other point cloud data analyzed during this study are not publicly available due to confidentiality agreements with commercial clients. However, derivative data, such as statistics and speed comparison data, are available from the corresponding author upon reasonable request.

Author information

Author ORCIDs

Luke Weidner https://orcid.org/0000-0002-6874-3308

Author notes

Matthew Lato served as Editorial Board Member at the time of manuscript review and acceptance and did not handle peer review and editorial decisions regarding this manuscript.

Author contributions

Data curation: MvV
Investigation: LW
Methodology: LW, AF, MvV, MJL
Software: AF
Supervision: MJL
Validation: LW
Visualization: LW
Writing – original draft: LW
Writing – review & editing: AF, MvV, MJL

Competing interests

This research is sponsored by BGC Engineering Inc. and may lead to the development of products of which said sponsor has financial interests.

Supplementary material

Supplementary data are available within the article at https://doi.org/10.1139/cgj-2023-0073.

References

Abellán, A., Jaboyedoff, M., Oppikofer, T., and Vilaplana, J.M. 2009. Detection of millimetric deformation using a terrestrial laser scanner: experiment and application to a rockfall event. Natural Hazards and Earth System Sciences, **9**: 365–372. doi:10.5194/nhess-9-365-2009.

Alotaibi, A.A., Maerz, N.H., Boyko, K.J., Youssef, A.M., and Pradhan, B. 2022. Temporal LiDAR scanning in quantifying cumulative rockfall volume and hazard assessment: a case study at southwestern Saudi Arabia. The Egyptian Journal of Remote Sensing and Space Science, **25**: 435–443. doi:10.1016/j.ejrs.2022.03.010.

Anand, B., Patil, A.G., Senapati, M., Barsaiyan, V., and Rajalakshmi, P. 2020. Comparative run time analysis of LiDAR point cloud processing with GPU and CPU. 2020 IEEE International Conference on Computing, Power and Communication Technologies (GUCON). Presented at the 2020 IEEE International Conference on Computing, Power and Communication Technologies (GUCON), IEEE, Greater Noida, India, pp. 650–654. doi:10.1109/GUCON48875.2020.9231067.

Arun, K.S., Huang, T.S., and Blostein, S.D. 1987. Least-squares fitting of two 3-D point sets. IEEE Transactions on Pattern Analysis and Machine Intelligence, **5**: 698–700. doi:10.1109/TPAMI.1987.4767965.

Baker, S.P., and Sadowski, R.W. 2013. GPU assisted processing of point cloud data sets for ground segmentation in autonomous vehicles. 2013 IEEE Conference on Technologies for Practical Robot Applications (TePRA). Presented at the 2013 IEEE Conference on Technologies for Practical Robot Applications (TePRA), IEEE, Woburn, MA, USA, pp. 1–6. doi:10.1109/TePRA.2013.6556352.

Besl, P.J., and McKay, N.D. 1992. A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence, **14**: 239–256. doi:10.1109/34.121791.

Blue Marble Geographics 2020. Global Mapper Version 22. https://www.bluemarblegeo.com/

Chaiso, K., and Ratanaworabhan, P. 2020. Accelerating cylinder detection in point clouds using GPU. 2020 17th International Joint Conference on Computer Science and Software Engineering (JCSSE). Presented at the 2020 17th International Joint Conference on Computer Science and Software Engineering (JCSSE), IEEE, Bangkok, Thailand, pp. 134–138. doi:10.1109/JCSSE49651.2020.9268309.

Chen, Z., Gao, B., and Devereux, B. 2017. State-of-the-art: DTM generation using airborne LIDAR data. Sensors, **17**: 150. doi:10.3390/s17010150. PMID: 28098810.

Conforti, M., Mercuri, M., and Borrelli, L. 2021. Morphological changes detection of a large earthflow using archived images, LiDAR-Derived DTM, and UAV-based remote sensing. Remote Sensing, **13**: 120. doi:10.3390/rs13010120.

de Gélis, I., Bessin, Z., Letortu, P., Jaud, M., Delacourt, C., Costa, S., et al. 2022. Cliff change detection using Siamese Kpconv deep network on 3D point clouds. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences. Presented at the XXIV ISPRS Congress "Imaging today, foreseeing tomorrow," Commission III - 2022 edition, 6-11 June 2022, Nice, France, Copernicus GmbH, pp. 649–656, ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.V-3-2022. doi:10.5194/isprs-annals-V-3-2022-649-2022.

DiFrancesco, P.-M., Bonneau, D., and Hutchinson, D.J. 2020. The implications of M3C2 projection diameter on 3D semi-automated rockfall extraction from sequential terrestrial laser scanning point clouds. Remote Sensing, **12**: 1885. doi:10.3390/rs12111885.

Donati, D., Stead, D., Lato, M., and Gaib, S. 2020. Spatio-temporal characterization of slope damage: insights from the Ten Mile Slide, British Columbia, Canada. Landslides, **17**: 1037–1049. doi:10.1007/s10346-020-01352-3.

Fei, L., Jaboyedoff, M., Guerin, A., Noël, F., Bertolo, D., Derron, M.-H., et al. 2023. Assessing the rock failure return period on an unstable Alpine rock wall based on volume-frequency relationships: the Brenva Spur (3916 m a.s.l., Aosta Valley, Italy). Engineering Geology, **323**: 107239. doi:10.1016/j.enggeo.2023.107239.

Fisher, P.F., and Tate, N.J. 2006. Causes and consequences of error in digital elevation models. Progress in Physical Geography: Earth and Environment, **30**: 467–489. doi:10.1191/0309133306pp492ra.

Girardeau-Montaut, D., Roux, M., Marc, R., and 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**: W19.

Gojcic, Z., Schmid, L., and Wieser, A. 2021. Dense 3D displacement vector fields for point cloud-based landslide monitoring. Landslides, **18**: 3821–3832. doi:10.1007/s10346-021-01761-y.

Guerin, A., Abellán, A., Matasci, B., Jaboyedoff, M., Derron, M.-H., and Ravanel, L. 2017. Brief communication: 3-D reconstruction of a collapsed rock pillar from Web-retrieved images and terrestrial lidar data—the 2005 event of the west face of the Drus (Mont Blanc massif). Natural Hazards and Earth System Sciences, **17**: 1207–1220. doi:10.5194/nhess-17-1207-2017.

Guzzetti, F., Mondini, A.C., Cardinali, M., Fiorucci, F., Santangelo, M., and Chang, K.-T. 2012. Landslide inventory maps: new tools for an old problem. Earth-Science Reviews, **112**: 42–66. doi:10.1016/j.earscirev.2012.02.001.

Harpold, A.A., Marshall, J.A., Lyon, S.W., Barnhart, T.B., Fisher, B.A., Donovan, M., et al. 2015. Laser vision: lidar as a transformative tool to advance critical zone science. Hydrology and Earth System Sciences, **19**: 2881–2897. doi:10.5194/hess-19-2881-2015.

Hu, X., Li, X., and Zhang, Y. 2013. Fast filtering of LiDAR point cloud in urban areas based on scan line segmentation and GPU acceleration. IEEE Geoscience and Remote Sensing Letters, **10**: 308–312. doi:10.1109/LGRS.2012.2205130.

Hunter, J. 2022. The long road ahead: remaking the Coquihalla Highway after a one-in-1,000-year storm. The Globe and Mail.

Iverson, R.M., George, D.L., Allstadt, K., Reid, M.E., Collins, B.D., Vallance, J.W., et al. 2015. Landslide mobility and hazards: implications of the 2014 Oso disaster. Earth and Planetary Science Letters, **412**: 197–208. doi:10.1016/j.epsl.2014.12.020.

Joerg, P.C., Morsdorf, F., and Zemp, M. 2012. Uncertainty assessment of multi-temporal airborne laser scanning data: a case study on an Alpine glacier. Remote Sensing of Environment, **127**: 118–129. doi:10.1016/j.rse.2012.08.012.

Kaldestad, K.B., Haddadin, S., Belder, R., Hovland, G., and Anisi, D.A. 2014. Collision avoidance with potential fields based on parallel processing of 3D-point cloud data on the GPU. 2014 IEEE International Conference on Robotics and Automation (ICRA). Presented at the 2014 IEEE International Conference on Robotics and Automation (ICRA), IEEE, Hong Kong, China, pp. 3250–3257. doi:10.1109/ICRA.2014.6907326.

Kromer, R., Abellán, A., Hutchinson, D., Lato, M., Edwards, T., and Jaboyedoff, M. 2015. A 4D filtering and calibration technique for small-scale point cloud change detection with a terrestrial laser scanner. Remote Sensing, **7**: 13029–13052. doi:10.3390/rs71013029.

Kromer, R., Lato, M., Hutchinson, D.J., Gauthier, D., and Edwards, T. 2017. Managing rockfall risk through baseline monitoring of precursors using a terrestrial laser scanner. Canadian Geotechnical Journal, **54**: 953–967. doi:10.1139/cgj-2016-0178.

Lague, D., Brodu, N., and Leroux, J. 2013. 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. doi:10.1016/j.isprsjprs.2013.04.009.

LaHusen, S.R., Duvall, A.R., Booth, A.M., and Montgomery, D.R. 2016. Surface roughness dating of long-runout landslides near Oso, Washington (USA), reveals persistent postglacial hillslope instability. Geology, **44**: 111–114. doi:10.1130/G37267.1.

Lato, M., and Ferrier, A. 2022. Systems and methods for evaluating changes in terrain topography over time. US 11,288,826 B1. U.S. Patent and Trademark Office.

Lato, M., Porter, M., Hensold, G., and McDougall, S. 2016. Understanding landslide movement and kinematics with airborne lidar. Proceedings of the Canadian Geotechnical Conference, GeoVancouver. Presented at the GeoVancouver, Canadian Geotechnical Society, Vancouver, British Columbia.

Lato, M.J., Anderson, S., and Porter, M.J. 2019. Reducing landslide risk using airborne lidar scanning data. Journal of Geotechnical and Geoenvironmental Engineering, **145**: 06019004. doi:10.1061/(ASCE)GT.1943-5606.0002073.

Lato, M.J., Hutchinson, D.J., Gauthier, D., Edwards, T., and Ondercin, M. 2015. Comparison of airborne laser scanning, terrestrial laser scanning, and terrestrial photogrammetry for mapping differential slope change in mountainous terrain. Canadian Geotechnical Journal, **52**: 129–140. doi:10.1139/cgj-2014-0051.

Niculiță, M., Mărgărint, M.C., and Tarolli, P. 2020. Using UAV and LiDAR data for gully geomorphic changes monitoring. Developments in Earth Surface Processes, Remote Sensing of Geomorphology *Edited by* P. Tarolli and S.M. Mudd. Elsevier, pp. 271–315. doi:10.1016/B978-0-444-64177-9.00010-2.

Nissen, E., Maruyama, T., Ramon Arrowsmith, J., Elliott, J.R., Krishnan, A.K., Oskin, M.E., and Saripalli, S. 2014. Coseismic fault zone deformation revealed with differential lidar: examples from Japanese Mw ∼7 intraplate earthquakes. Earth and Planetary Science Letters, **405**: 244–256. doi:10.1016/j.epsl.2014.08.031.

Okyay, U., Telling, J., Glennie, C.L., and Dietrich, W.E. 2019. Airborne lidar change detection: an overview of Earth sciences applications. Earth-Science Reviews, **198**: 102929. doi:10.1016/j.earscirev.2019.102929.

Passalacqua, P., Belmont, P., Staley, D.M., Simley, J.D., Arrowsmith, J.R., Bode, C.A., et al. 2015. Analyzing high resolution topography for advancing the understanding of mass and energy transfer through landscapes: a review. Earth-Science Reviews, **148**: 174–193. doi:10.1016/j.earscirev.2015.05.012.

Piovan, S.E., Hodgson, M.E., Mozzi, P., Porter, D.E., and Hall, B. 2023. LiDAR-change-based mapping of sediment movement from an extreme rainfall event. GIScience & Remote Sensing, **60**: 2227394. doi:10.1080/15481603.2023.2227394.

Pirotti, F., Guarnieri, A., and Vettore, A. 2013. State of the art of ground and aerial laser scanning technologies for high-resolution topography of the earth surface. European Journal of Remote Sensing, **46**: 66–78. doi:10.5721/EuJRS20134605.

Richter, R., Kyprianidis, J.E., and Döllner, J. 2013. Out-of-core GPU-based change detection in massive 3D point clouds. Transactions in GIS, **17**: 724–741. doi:10.1111/j.1467-9671.2012.01362.x.

Rusu, R.B., and Cousins, S. 2011. 3D is here: Point Cloud Library (PCL). IEEE International Conference on Robotics and Automation (ICRA). IEEE, Shanghai, China.

Schovanec, H., Walton, G., Kromer, R., and Malsam, A. 2021. Development of improved semi-automated processing algorithms for the creation of rockfall databases. Remote Sensing, **13**: 1479. doi:10.3390/rs13081479.

Smets, B., d'Oreye, N., Kervyn, M., and Kervyn, F. 2017. Gas piston activity of the Nyiragongo lava lake: first insights from a Stereographic Time-Lapse Camera system. Journal of African Earth Sciences, **134**: 874–887. doi:10.1016/j.jafrearsci.2016.04.010.

Stumpf, A., Malet, J.-P., Allemand, P., Pierrot-Deseilligny, M., and Skupinski, G. 2015. Ground-based multi-view photogrammetry for the monitoring of landslide deformation and erosion. Geomorphology, **231**: 130–145. doi:10.1016/j.geomorph.2014.10.039.

Swirad, Z.M., and Young, A.P. 2021. Automating coastal cliff erosion measurements from large-area LiDAR datasets in California, USA. Geomorphology, **389**: 107799. doi:10.1016/j.geomorph.2021.107799.

Tremblay-Auger, F., Locat, A., Leroueil, S., Locat, P., Demers, D., Therrien, J., and Mompin, R. 2021. The 2016 landslide at Saint-Luc-de-Vincennes, Quebec: geotechnical and morphological analysis of a combined flowslide and spread. Canadian Geotechnical Journal, **58**: 295–304. doi:10.1139/cgj-2019-0671.

van Veen, M., Funk, A., Fish, C., and Porter, M. 2022a. A landslide velocity database for the Site C Reservoir in Northeastern British Columbia. GeoCalgary 2022: Reflection on Resources. Presented at the GeoCalgary 2022, Calgary, CA.

van Veen, M., Lato, M., Hove, J., Hunchuck, G., Babcock, J., and Bracic, J. 2017. The use of airborne LiDAR in understanding ground hazards for large pipeline networks. Proceedings of GeoOttawa 2017, the 70th Canadian Geotechnical Conference. Presented at the GeoOttawa 2017, Canadian Geotechnical Society, Ottawwa, ON.

van Veen, M., Porter, M., Lato, M., Mitchell, A., Fish, C., and Van Gassen, W. 2022b. Using tree stems in multi-temporal terrestrial lidar scanning data to monitor landslides on vegetated slopes. Landslides, **19**: 829–840. doi:10.1007/s10346-021-01815-1.

Vasquez, T. 2022. How an atmopsheric river flooded British Columbia. Weatherwise, **75**: 19–23. doi:10.1080/00431672.2022.2021767.

Wagner, W., Lague, D., Mohrig, D., Passalacqua, P., Shaw, J., and Moffett, K. 2017. Elevation change and stability on a prograding delta. Geophysical Research Letters, **44**: 1786–1794. doi:10.1002/2016GL072070.

Wartman, J., Montgomery, D.R., Anderson, S.A., Keaton, J.R., Benoît, J., dela Chapelle, J., and Gilbert, R. 2016. The 22 March 2014 Oso landslide, Washington, USA. Geomorphology, **253**: 275–288. doi:10.1016/j.geomorph.2015.10.022.

Weidner, L., van Veen, M., Lato, M., and Walton, G. 2021. An algorithm for measuring landslide deformation in terrestrial lidar point clouds using trees. Landslides, **18**: 3547. doi:10.1007/s10346-021-01723-4.

Westoby, M., Lim, M., Hogg, M., Dunlop, L., Pound, M., Strzelecki, M., and Woodward, J. 2020. Decoding complex erosion responses for the mitigation of coastal rockfall hazards using repeat terrestrial LiDAR. Remote Sensing, **12**: 2620. doi:10.3390/rs12162620.

Wheaton, J.M., Brasington, J., Darby, S.E., and Sear, D.A. 2010. Accounting for uncertainty in DEMs from repeat topographic surveys: improved sediment budgets. Earth Surface Processes and Landforms, **35**: 136–156. doi:10.1002/esp.1886.

Williams, J.G., Anders, K., Winiwarter, L., Zahs, V., and Höfle, B. 2021. Multi-directional change detection between point clouds. ISPRS Journal of Photogrammetry and Remote Sensing, **172**: 95–113. doi:10.1016/j.isprsjprs.2020.12.002.

Williams, J.G., Rosser, N.J., Hardy, R.J., Brain, M.J., and Afana, A.A. 2018. Optimising 4-D surface change detection: an approach for capturing rockfall magnitude–frequency. Earth Surface Dynamics, **6**: 101–119. doi:10.5194/esurf-6-101-2018.

Winiwarter, L., Anders, K., and 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. doi:10.1016/j.isprsjprs.2021.06.011.

Zahs, V., Winiwarter, L., Anders, K., Williams, J.G., Rutzinger, M., and 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. doi:10.1016/j.isprsjprs.2021.11.018.

Zhou, K., Hou, Q., Wang, R., and Guo, B. 2008. Real-time KD-tree construction on graphics hardware. ACM Transactions on Graphics, **27**: 126. doi:10.1145/1409060.1409079.