Deep Mining 2024 – P Andrieux & D Cumming-Potvin (eds)

doi:10.36487/ACG_repo/2465_23

Automated ground and support deformation monitoring: a novel method with expanded features for geotechnical engineers

Jochen Franke [a,*], Carlos Gonzalez [a]

[a] *Caroni Geospatial, Australia*

Abstract

One of the regular tasks of geotechnical and mining engineers is the measurement and management of deformation in underground mine excavation networks. Although other methods have historically been used to do so, light detection and ranging (LiDAR) has emerged as the most suited technology for this application because it caters for quantitative and omission-free rather than merely qualitative tracking. Regular scanning has shown to be highly advantageous for mines with swelling or squeezing ground and those with rapid deformation, which is common for many deep mines.

Our paper provides an update on improved capabilities and expanded features of a novel solution of fully automated LiDAR point cloud data processing dedicated to drive and other void deformation tracking that was first introduced at the Australian Centre for Geomechanics Ground Support 2023 conference. This update covers new features to automatically extract and track rockbolt deformation. Our system not only enables geotechnical engineers to avoid having to learn unrelated skills, but also provides immediate output of sophisticated reporting deliverables not available elsewhere. This opens up the opportunity to monitor many more excavation volumes at a higher frequency and gain better insights than is possible with conventional manual processing tools and methodology. This in turn can make the difference between managing rather than missing the risk of rockfall and ultimately fatalities.

More specifically, we present how our processing methodology utilises latest generation cloud-based data storage and processing infrastructure to seamlessly integrate and automatically provide standard geotechnical reporting, accessible from anywhere, that relieves geotechnical engineers from having to compile this material manually. Such reporting includes sophisticated and comprehensive automated detection and tracking of ground support to aid cost efficient rehabilitation planning, which can lead to significant cost savings for deep and other sites experiencing dynamic conditions and squeezing ground. Case study material illustrates a range of these topics.

Keywords: *LiDAR, automated rockbolt tracking, deformation monitoring, convergence monitoring, instrumentation and monitoring*

1 Introduction

Light detection and ranging (LiDAR) technology was introduced to underground mining in the early 2000s and brought with it the ability for quantitative assessment across the entire excavation volume and its rock surface by collecting a three-dimensional (3D) image of the entirety of an excavation, such as an underground drive, decline, or a tunnel. This was a distinct improvement over traditional, pointwise deformation tracking of underground voids, because 3D LiDAR comparisons between epochs of complete 3D data coverage allows for change detection over time that doesn't feature otherwise typical omissions.

* Corresponding author. Email address: jochen.franke@caroni.com.au

Traditional individual point measurement methods can yield highly accurate individual point tracking, such as:

- bolt mounted systems including extensometers and point lasers
- handheld point lasers, such as the Leica Disto
- total station surveys of permanently mounted targets.

However, that becomes irrelevant when considering that they feature omissions that inevitably lead to mis-representation and almost always under-representation of actual worst deformations of the voids observed. As a result, and as illustrated in Figure 1, such pointwise tracking therefore almost always paints an overly optimistic view of deformation. This is dangerous because it is the worst deformations that are likely to cause adverse events like rockfalls, causing damage to equipment, injuries, or in the worst-case fatalities.

Fatalities through roc falls still occur in underground mining even in industrialised western countries with modern best practices, despite ongoing industry-wide efforts in using ground support to prevent them. Rigorous deformation tracking efforts provide quantified knowledge on ground support effectiveness over time. In situations where risks to personnel or equipment exist, simply installing ground support but then leaving its failure to chance because no rigorous monitoring effort is made, is not enough. At the time of writing, the most recent underground mining fatality caused by rockfall in Australia occurred in the state of Victoria in March 2024, despite its comparably stringent regulatory environment. As reported on by Worksafe Victoria (2024) two workers were trapped by a rockfall resulting in the death of one, a 37-year-old, while the other was rescued and airlifted to hospital with serious injuries.

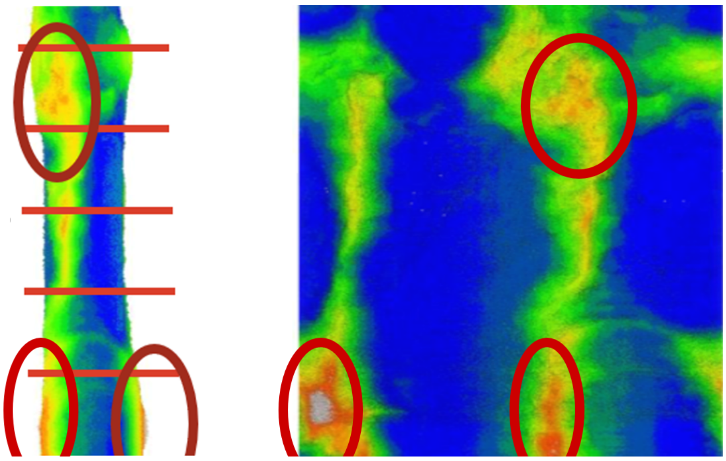

(a) (b)

Figure 1 An in–principle example for how pointwise measurement systems miss actual worst deformations in an underground drive. (a) shows colour coded deformations captured by 3D LiDAR, blue representing no deformation through to green, yellow, red and then white representing worsening convergence. (b) shows the same information in an unrolled or flattened representation for ease of interpretation. Red lines on the left represent pre-selected locations for tracking that any pointwise measurement system would naturally need to limit itself to. As can be seen in the areas circled in red, the pre-selected pointwise tracking locations miss actually occurring worst deformations, coloured red and white, whereas 3D LiDAR tracking does not

Unlike pointwise methods, regular 3D LiDAR scanning has been shown to be highly advantageous for mines in comprehensively tracking swelling or squeezing ground, or rapid, short-burst deformation through seismic exposure. In Australia and a number of other mining regions, the authors have observed that the adoption of 3D LiDAR equipment for underground deformation tracking has progressed considerably since Jones (2020) summarised mobile 3D LiDAR mapping techniques and their use in underground mining.

What has remained an obstacle though is that the workflow from data acquisition to final interpretation is not currently an automated algorithmic process, relying on manual methods used by practitioners. This has led to many mine sites having a rigorous and robust 3D scanning data regime, but little to show in terms of actionable reports and output data. A key reason for that is that working with point clouds is tedious to most geotechnical engineers. Point clouds have long been the realm of surveyors, but to fully adopt their use in geomechanics requires upskilling and allocated resources, which could already be stretched, thus there are often insufficient resources to be applied to bringing 3D LiDAR technology into effective use.

As already discussed in our previous related publication (Franke & Gonzalez 2023) the authors sought, and received, further validation on this issue through comprehensive industry feedback, whereby mine sites have limited or avoided 3D LiDAR data capture due to lack of time, processing personnel, or processing skills. The time, effort, the expertise and specialist software it takes to process point cloud data into meaningful reports are prohibitive for typical geotechnical engineering teams. These, and other findings, were originally published in 2022 (Franke & Gonzalez 2022) and led the authors to commit to the development of the automated processing, analysis and reporting software VoidMapper (Franke & Gonzalez 2024) that addresses these shortcomings and is described in more detail in this paper.

2 Automated drive and void deformation tracking

Raw 3D LiDAR data collected in underground drives and other voids is, as such, not directly usable. It needs to be processed into a derivative form that lends itself for interpretation and analysis, and ultimately actionable reporting for decision-making to achieve safe and cost-effective outcomes for underground mining operations. The conventional state-of-the-art is for mine sites to employ their own survey or geotechnical personnel, or external service providers, to undertake manual 3D LiDAR data processing of all the steps necessary to produce such actionable deliverables. Aside from this manual process at best limiting sites to a restricted number of deformation tracking locations, due to the time and cost penalty involved, and at worst, leading to even those limited locations remaining unprocessed, the quality of manually produced deliverables can also be compromised.

The quality and accuracy of manual 3D LiDAR data processing results can be compromised through human error, which is more likely to come into play when the process in question, such as this, is tedious, repetitive, and often outside the expertise of geotechnical engineers tasked with producing the reports. In addition to that, by far the most common software tool used by geotechnical engineers or mine surveyors to produce manual underground deformation tracking results is the freeware CloudCompare (Girardeau-Montaut 2024).

While CloudCompare has many highly useful and a number of sophisticated software features to manually process 3D LiDAR point clouds for multiple purposes, its origins and software development are of academic rather than commercial nature, which frequently makes its use comparably less straight forward. Limitations, such as in using large file sizes common to 3D point clouds, can also apply, but arguably the biggest issue is its potential for intermittently producing false deformation tracking output when using it for underground drive and void deformation monitoring as is explained further in the following paragraph and Figure 2. CloudCompare was not specifically designed for underground voids, such as drives, which inherently are enclosed 3D spaces rather than un-enclosed geographic, or other open mapping spaces that are the mainstay focus for CloudCompare users.

Underground mining geotechnical users target the tracking of all actual deformation, frequently and simplistically referred to as convergence. The latter is inward in nature, when outward deformation also needs to be tracked, because it does occur in situ. Rather than calling the latter divergence, which has misleading connotations, the authors refer to it as expansion, which has a straightforward meaning. While possible to implement in CloudCompare, tracking of both convergence and expansion can unintentionally be reported by it in a false inverted manner, which is due to the nature of how CloudCompare orients itself for enclosed spaces. This orientation selection is done by the software and can’t be chosen by the user, which means false inversions are intermittent and not in the control of the user. Whether or not inversion occurs depends on the input data set in use.

Manual deformation tracking methods available in CloudCompare can also unintentionally lead to false large deformations at the end of scan cloud coverage. This is due to mismatching data coverage at these scan coverage ends which is natural for any repeat scanning of the same location area for tracking purposes. While such mismatching coverage can be minimised through robust scanning best practice, it is not feasible to completely avoid it.

These issues with the manual deformation tracking methods described in the preceding paragraphs are illustrated for a real-world example in Figure 2.

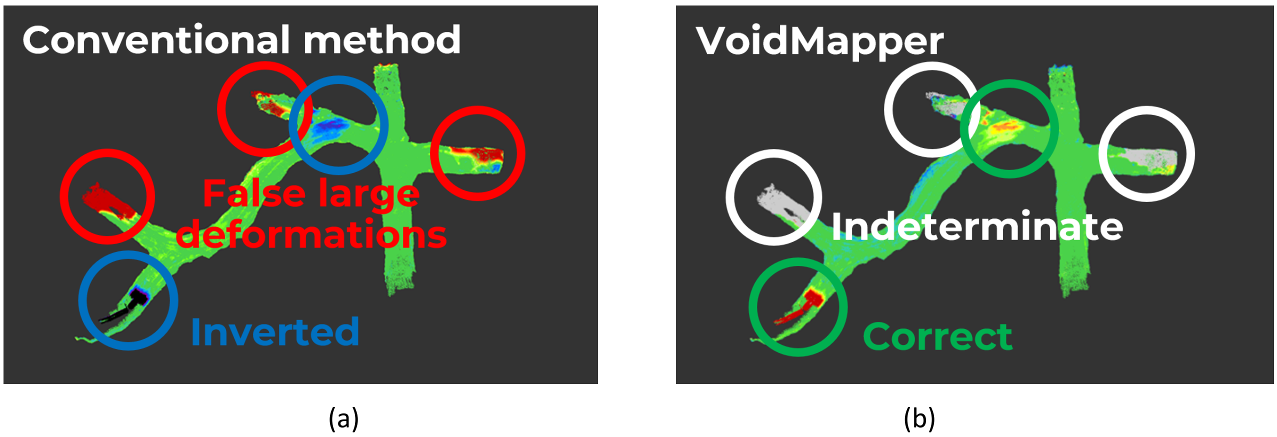

Figure 2 Real-world example of deformation tracking of an underground drive using scans collected at two different dates. It compares (a) conventional reporting, featuring false inversions, whereby convergence is reported as expansion and vice versa, and false large deformations due to mismatching end of scan coverage between the two scans and (b) correctly reported convergence and expansion and mismatching end of scan coverage reported as indeterminate, as produced by VoidMapper

The automated tracking system, VoidMapper, described in this paper, was designed for the specific underground mining void deformation tracking purpose required for underground geotechnical practitioners. Specialised and experienced software engineers have been engaged to overcome the issues caused by manual data processing, including false reporting. This process is ongoing as new features are added.

One of the consequences of manual data processing, and its finite resource of time to complete it, is that it is likely that not only are just a limited number of underground areas actively monitored, but deformation tracking is limited to comparing any newly collected scan to the previous scan. There is simply not enough time to manually produce anything else. This provides a snapshot of short-term deformations but does not provide any insight into whether or not deformations over a longer period of time may have reached a critical stage, which is naturally not obvious by analysing short-term tracking results. Geotechnical engineers have been trained, and over time become highly experienced in visually detecting obvious, and even sometimes less obvious deformations, in situ when inspecting underground drives. They do so frequently as part of their core role. However, quantifying such deformations and quantifying definitively whether or not these visually observed – not to mention visually undetected – deformations have reached a critical magnitude is another matter entirely. This is where an automatic detection system has its advantages.

When using a fully automated data processing system, no such time limitations apply because the algorithm only takes immaterially longer to produce additional tracking results showing long term deformations to supplement the short-term view. VoidMapper independently makes use of available scan data uploaded to it by detecting which scan dates have been collected for the particular drive area in question, and then reporting on deformation as follows:

- A short-term comparison (STC) which is produced once two or more scan date files for an area have been sent for automatic processing – this output file compares the previous with the current scan date.
- The maximum STC deformations containing the seven worst converging and the seven worst expanding regions automatically detected in the comparison, ranked by deformation volume.

- A long-term comparison (LTC) scan cloud file which is produced once three or more scan date files for an area have been sent for automatic processing – this output file compares the oldest available with the current scan date.
- The maximum LTC deformations containing the seven worst converging and the seven worst expanding regions automatically detected in the comparison, ranked by deformation volume.
- An automatically produced summary report PDF file containing all important analysis information for every scan file sent for automated processing.

VoidMapper not only overcomes the bottleneck of time and effort required to process raw 3D scan data into meaningful and actionable reporting, it also fast-tracks providing it because all steps for producing deliverables are fully automated and are completed by an algorithm rather than a human operator. Since it is a cloud-based system, there is also no limit to scalability regarding the number and size of scans being processed. Table 1 provides typical examples for VoidMapper automated processing times, when achieving the same outcome would take an in-house operator at best the most part of a day, or it could take days up to weeks waiting for results from a service provider.

Table 1 Typical processing time examples for VoidMapper to automatically produce deliverables ready for analysis and action by geotechnical practitioners

Examples	Site A		Site B		
Dates	06/12/2022	28/12/2022	16/06/2020	04/12/2020	12/10/2021
File size [MB]	721	465	9	50	51
Point cloud clean up	5 min	6 min	2 min	1 min	2 min
LTC creation	–	–	–	–	2 min
LTC deformation ranking	–	–	–	–	+1 sec
STC creation	–	9 min	–	2 min	13 sec
STC deformation ranking	–	+1 sec	–	+0 sec	+1 sec
Word summary report	–	9 sec	–	17 sec	5 sec
PDF summary report	–	17 sec	–	15 sec	4 sec
Total	**–**	**15.5 min**	**–**	**3.5 min**	**4.5 min**

Once STC and LTC deformation results are produced, the next step towards actionable results and reporting is to analyse this data. Conventionally, this type of analysis involves geotechnical users viewing and interpreting deformation colour plots to visually quantify what they contain and to decide whether or not what is viewed poses a problem. While this manual digital process is superior to in situ visual inspections, because it ensures that no deformation is unrecorded, it has its limitations. Depending on the display parameters (such as colouring in use) and other factors (such as magnitude and distribution of deformation recorded, or the viewer tool utilised to visualise deformation) manual error can once again be a factor, this time in determining whether or not deformation colour plots contain issues that need to be addressed or not.

It is, therefore, preferable in this step to let an algorithm work out how big the problem is, and make sure that the worst recorded and observed deformation issues are objectively noted and reported on rather than missed because an operator subjectively judged a different area to be the main issue. VoidMapper fully automatically completes this task in seconds, as reported on in Table 1, and provides worst deformation zones as a separate 3D point cloud output that can be interrogated in its 3D viewer, as shown in Figure 3.

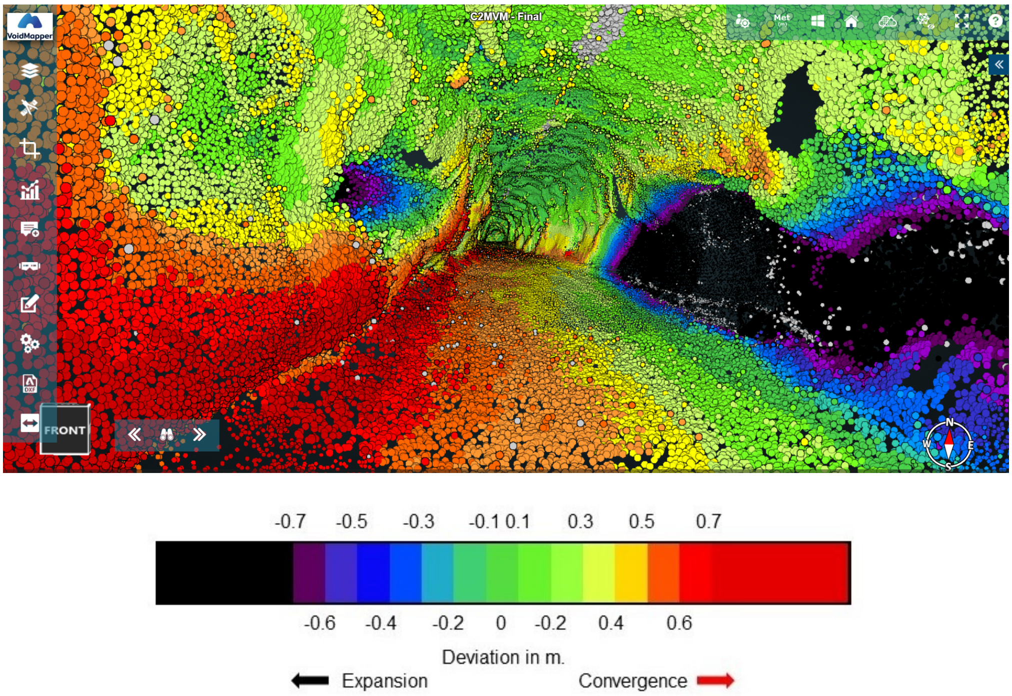

(a)

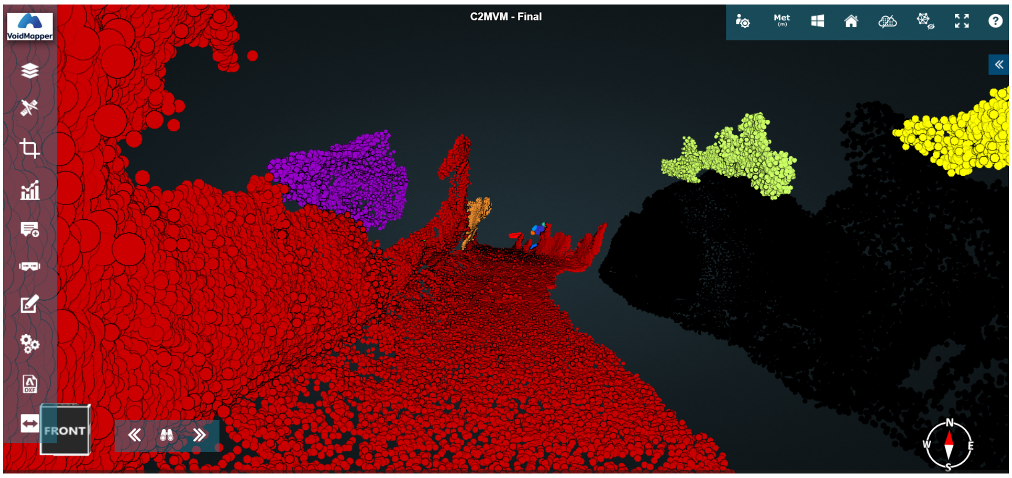

147.73	8.61	4.37	3.13	1.50	1.11	0.90	0	-0.06	-0.20	-0.26	-0.43	-1.94	-3.41	-63.53	Unit: m^3
Convergence								Expansion							

(b)

Figure 3 Example of (a) automatically determined worst convergence and expansion zones, viewed in 3D, extracted for an (b) automatically generated LTC deformation scan cloud. Worst deformation zones are automatically ranked and accordingly coloured by volume. In this example, with the worst convergence area featuring 147.73 m^3 of material pushed inward, and the worst expansion area featuring 63.53 m^3 having been dislodged

3D deformation monitoring of drive and underground void surfaces, as described in this section, generally speaking is, and can be used as, a proxy for tracking the integrity of installed ground support. There are, however, good reasons to explicitly track ground support deformation itself, which is explained further in subsequent sections. Even with a fully functional automated rockbolt tracking tool, as described therein, there is still a place for deformation tracking of the entirety of drive or void surfaces. Since LiDAR 3D point cloud spatial recording is based on a line-of-sight working principle, it does not provide any data on embedded rockbolt portions. Complementary drive surface deformation data can provide additional input into interpreting what might have occurred to them.

3 Automated rockbolt tracking

Accurate and comprehensive tracking of rockbolt integrity has long been targeted by underground mining practitioners because, if effective and actionable, it delivers a significant improvement for managing rockfall risk to personnel and equipment and managing the considerable cost of rehabilitation (Mercier-Langevin 2019).

Attempts to do so have included various measurement techniques, including bolt mounted systems, such as a recent '…novel ultrasound based Rockbolt Sensor…to monitor load, deformation, and integrity of rockbolts' (Smith-Boughner et al. 2020). Even when accurate, bolt mounted systems carry the penalty of requiring the effort, time, material and associated cost to install and then maintain such systems on individual bolts, in a harsh environment. It is not practical and would be prohibitively costly to install such systems on more than a selected few bolts, not just in terms of material and labour cost, but also in terms of disruption to operations while installation takes place. Limiting the number of installed bolt-mounted sensors to a practical, but therefore low, quantity in turn carries the risk of missing the majority of compromised or failed bolts in all other unmonitored drive sections. As with pointwise, general drive surface deformation tracking, covered in earlier sections, this means that bolt-mounted tracking overall is likely also not sufficiently representative to underpin reliable and representative monitoring and planning.

Aside from bolt mounted sensors, the authors are aware of attempts by a Western Australian nickel operation to track considerable localised bolt deformations through conventional survey means using reflectorless total stations to collect bolt tip coordinates at regular time intervals. This approach had to be abandoned because of the inordinate amount of time required to do so, which left surveyors with no time to cover other day-to-day tasks.

One of the main challenges to the bolt tracking methods by others, described above, is the sheer number of bolts installed in any given drive, and therefore, the effort required to track if not all of them then at least a sufficient number. As is evident from the example shown in Figure 3, it is common for localised deformation to occur in large but sometimes also in relatively small patches, such as several of the orange or yellow sections shown in this deformation plot that may well contain failed bolts. Unless all or at least the majority of bolts are tracked, these localised patches of potentially failed bolts would likely be missed.

Mobile 3D LiDAR has the clear advantage of requiring only a short amount of time for data collection (Franke & Gonzalez 2023) in a drive or void without getting in the way of operations, which makes this method highly practical, while at the same time collecting a complete and very detailed digital representation of the entire drive or void surface, including steel mesh, and rockbolts. This means that, provided care was taken during data collection to avoid obstructions to the field of view of the mobile LiDAR, the resulting point clouds contain each and every rockbolt, as illustrated in Figure 4.

The issue with raw LiDAR point clouds is that file content, as such, does not readily distinguish between bolts and other features. It is up to the user to extract critical bolt characteristics for integrity tracking. One could argue that extraction of rockbolt positions from LiDAR point clouds could be completed by a human data processing operator. Assume a typical 300 m long mobile LiDAR scan was collected in a drive section, and the installed bolt pattern has the following quite common geometry:

- 11 bolts per ring
- 1.4 m spacing in the direction of advance.

This means that this point cloud contains over 2,300 bolts that a human operator would need to manually extract. Assuming this operator is highly skilled, experienced, and can manually extract one bolt per minute from this point cloud, it would take over 38 hours to extract all bolts for this 300 m scan. Such an effort for just one scan, when sites typically need to monitor many, sometimes tens of cumulative kilometres of drives per month, makes manual extraction of rockbolts prohibitively impractical.

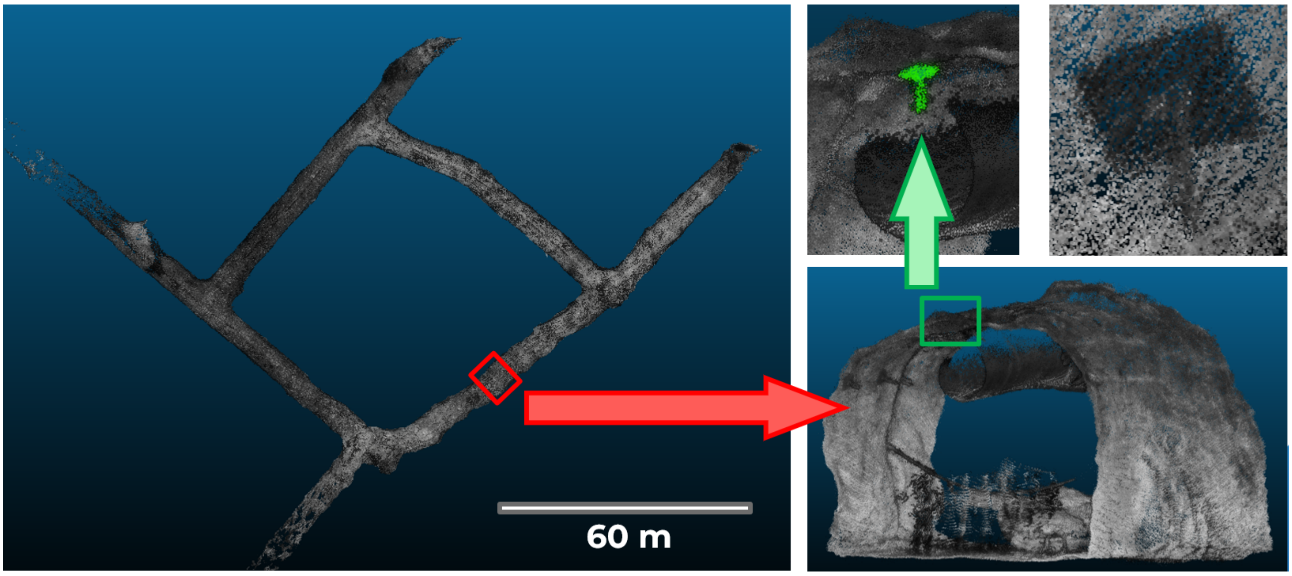

Figure 4 A typical mobile LiDAR scan of an underground drive showing the high resolution of points on rockbolts available for bolt characteristics extraction. It illustrates the advantage of evenly distributed high resolution across all surfaces and, therefore, all bolts that mobile LiDAR has over stationary LiDAR instruments, which feature significant losses in resolution in-between scan stations that make reliable rockbolt extraction for stationary LiDAR data impossible

A second issue for manually extracting rockbolt characteristics is that while rockbolt resolution and detail contained in mobile LiDAR point clouds is ample for the task, selecting any individual scan point to represent any particular bolt position is problematic because it may not be exactly in the centre of the bolt tip, which equates to a measurement error. It is worth noting that resolution in scan clouds produced by stationary LiDAR, is insufficient in-between stationary setup locations. This is due to the long and narrow enclosed space shape of a drive or tunnel that leads to dramatic resolution reduction along drive walls where bolts need to be captured. Mobile LiDAR instead features regular and dense point resolution all along drive walls. Aside from its much more cumbersome data collection process that is not practical in situations where scanning cannot be allowed to get in the way of other underground operations, and its more time-consuming post-processing methodology to produce georeferenced scan clouds, this definitively makes stationary LiDAR equipment unsuitable for automated rockbolt or steel mesh extraction.

As described in more detail in Figure 9 of Franke & Gonzalez (2023), consistency in what actual bolt position is extracted and then tracked over time is important to achieve accurate and robust tracking results. As mentioned previously, since LiDAR is a line-of-site technology, bolt sections embedded in rock cannot be tracked, but plates as well as bolt tip centres and bolt to plate centre contact points transferring load, or in other words bolt tails, can be, if the right methodology is applied.

So in summary, it stands to reason that rockbolt extraction for integrity tracking is only practical to achieve with an automated method, not only radically speeding up the process but achieving output bolt characteristics accuracy that is fit for purpose. The next section on 3D Point Classification provides details on how this is achieved.

4 3D Point cloud classification

Since raw 3D LiDAR point clouds don't come with an inherent separator for any content feature, such as rockbolts, other ways to separate them from non-bolts need to be found. This, in essence, is a binary classification task. Automatically classifying rockbolts from a 3D point cloud is achieved through cloud descriptors. A large range of cloud descriptors have been developed for different purposes over the past decade or so, in particular for classifying natural or general mapping but also indoor building data collected with LiDAR scanners. A good overview of available cloud descriptors is provided by Han et al. (2018).

All cloud descriptors suitable for identifying rockbolts at a useful success rate exploit the three-dimensional spatial characteristics of point cloud data to classify elements within it. Modern mobile LiDAR scanners used to collect point cloud data in underground mine drives can detect the intensity value recorded for each point, which represents the reflectivity of a particular object at the wavelength of the laser in use. Most recent mobile LiDAR equipment also features optional panoramic cameras that can be used to colourise scan clouds with natural colours of the scene scanned. Intensity and camera colour values can be used to strengthen the ground support classification success rates of 3D classifiers.

3D point cloud classification methods can be grouped into the following types:

- Supervised/Semi-supervised Classifiers, such as CANUPO:
  - Built on manually pre-determined spatial and geometric attributes or relationships between points. They do not learn latent features. Downsides include sensitivity to training data and relatively slow processing.
- Convolutional Neural Network/Deep Learning:
  - Use deep neural networks to learn latent features from the 3D point cloud. They are built on convolutional mesh/voxel architectures rather than unordered point sets. Volumetric representation is constrained by its resolution and computation cost of 3D convolution, which can limit success rate.
- Point Cloud Neural Network/Deep Learning:
  - Use deep neural networks to learn latent features from the 3D point cloud. They directly consume unordered point sets as inputs. It is the most effective way overcoming constraints of other methods.

First attempts to use 3D point cloud data for a binary classification of rockbolts were based around the CANUPO classifier, as published in a University of New South Wales PhD thesis in 2022, which, by definition, was new research at the time (Singh SK 2022). Since any supervised or semi-supervised classifications do not learn latent features contained in a scan of an underground drive, using them limits how high the extraction success rate for rockbolts can go. It is impractical, and would be cost-prohibitive, requiring each classification to be trained with its own training data set that has to be prepared by a human operator. Ultimately, it needs to be possible to run a successful classification without human input into the process or the advantages of an automated method over manually extracting bolt positions directly from the point cloud are lost. Unless a cloud-based server system with substantial computing power is used, there are also serious limitations for classifiers such as CANUPO to how large input scan cloud file sizes can go because processing time can become prohibitive. For these reasons, the authors invested the time and resources to develop a deep learning algorithm for binary rockbolt classification of 3D point clouds instead.

VoidMapper uses a neural network, deep learning algorithm that also requires initial training to teach it what bolts and plates and what non-bolts look like, such that the classification can draw on its learning of latent features. While preparing the training data needs to be done by a human operator too, in contrast to supervised, or some semi-supervised classifiers, this process is only required once, upfront for any given mine site to appropriately teach the algorithm the relevant bolt and non-bolt characteristics. The wider a variety of bolt types and their specific characteristics, including changes in appearance due to deformation, can be included in the deep learning training stage, the higher the extraction success rate will be. When the deep learning algorithm has been fully trained for a particular mine site, any of their subsequently uploaded scan clouds can be classified using the existing upfront training data. This means that there is no more need for any human intervention and ongoing rockbolt and plate extraction occurs fully automatically. Figure 5 illustrates example output for such automated rockbolt and plate extraction.

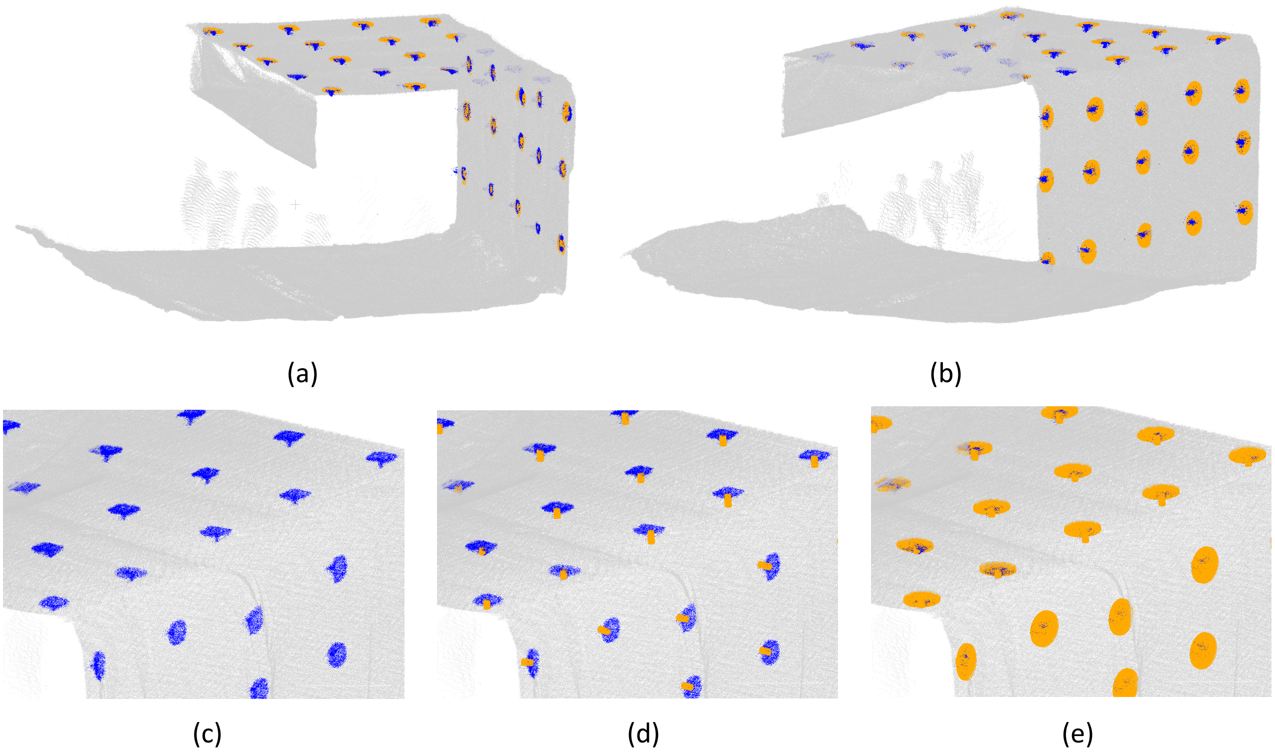

Figure 5 Example subsection of a fully automatically classified underground drive 3D scan cloud, shown in grey, using pre-existing deep learning training data. (a) and (b) show two different isometric views of the subsection, to allow for orientation, with only classified blue bolts and plates and extracted plates, in orange, visible; (a) Looks onto the scene from the right and upward and (b) looks onto the scene from the left and upward; (c) Shows a more detailed view of just the classified blue bolts and plates visible, (d) extracted orange bolt tails used for positional tracking overlaid and (e) both extracted orange bolt tails and plates for positional tracking overlaid

A wide range of test work has been undertaken to determine the success rate of the developed deep learning algorithm. This required extensive manual extraction of bolts and plates that could be used as ground truthing against the automatically classified data sets. One of the issues encountered in this process has been that it can be difficult for a human operator to manually extract rockbolts and their plates, even when using the best display settings of a stretched intensity histogram. This is because whether an operator can visually distinguish these features well enough from the rest of the drive or void walls to delineate them, depends on how well they stand out, aside from other factors like the software tools available or operator skill and experience. This varies quite widely between scan clouds, from cases where plates and bolts can be easily located and delineated, to cases where, despite the knowledge that they are there, they barely look any different to the surrounding rock. Reasons for this discrepancy can be traced to the type of scanner in use, the material, potential corrosion, and, therefore, appearance of plates and rockbolts, other site conditions, such as whether or not the surface was wet during scanning, dust that may cover installed plates and bolts, and more. Figure 6 illustrates a typical example of this phenomenon.

It is worth noting that, as is evident in this example, we have consistently observed that the deep learning algorithm performs better, sometimes considerably, in locating and correctly classifying bolts and plates than the human operator can, when instructed to only manually extract them provided it can be done with high confidence. This can be explained by the fact that the deep learning algorithm learns latent 3D and intensity related features in the point cloud for both the bolt and plate and the non-bolt class, which are both provided in the training files, whereas a human operator can only go by whatever contrast in intensity values suits the human eye in the given point cloud. The operator is also hindered by practicalities of visually finding bolts and plates in an enclosed space, with many other features potentially confusing the visual search.

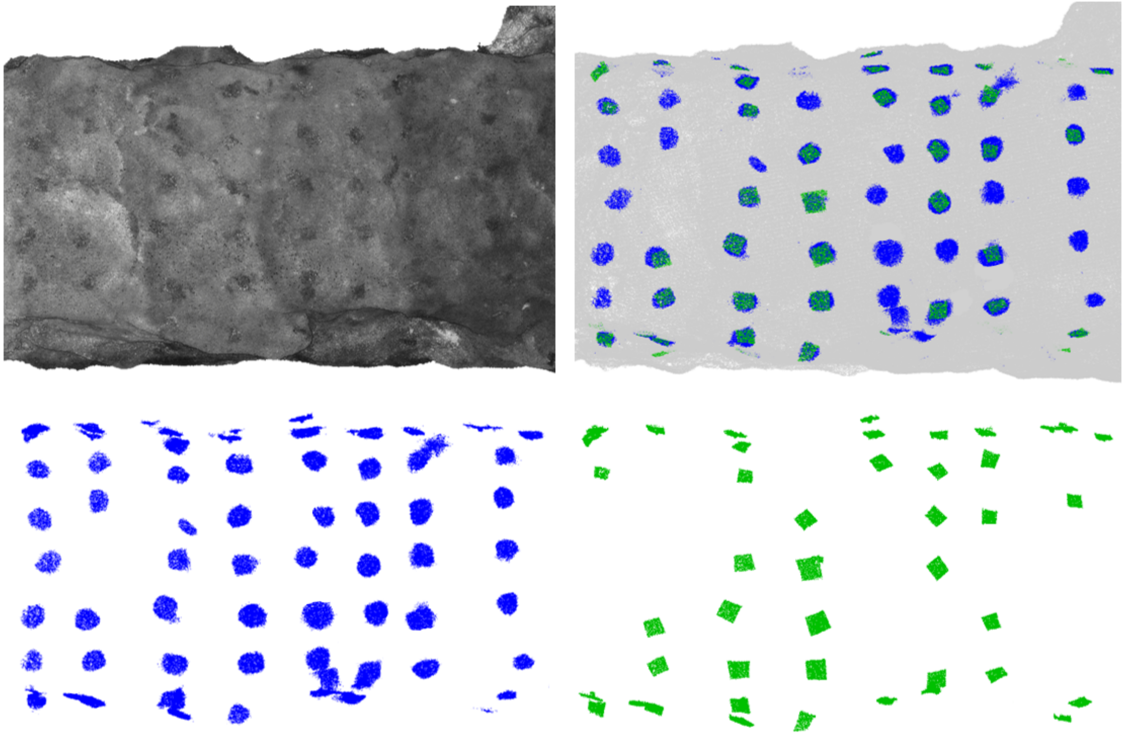

Figure 6 Example for an automatically classified point cloud section of an underground drive displayed in optimised intensity values, viewed from the outside top down, showing rockbolts and their plates installed in the backs as darker grey patches (top left). Green patches are bolts and plates that were manually extracted by a skilled operator for ground truthing purposes. Blue patches show bolts and plates found by the deep learning algorithm, based on previous training on other data sets from this mine. As illustrated by missing green patches where there are blue patches in the overlay image on the top right, it is obvious that the algorithm was able to find notably more bolts and plates than the manual operator could

The 3D point cloud specific, neural network classifier adopted as the core basis for our custom developed code to extract rockbolts and plates, comes with standard quality assurance (QA) tools that allow for an assessment of success by several measures. These include metrics to track bolt and plate true positives, false positives, true negatives, and false negatives. The native QA method of tracing each and every individual classified 3D point to one of these QA classes, is suited for many different classification applications, but it did not turn out to be suitable for tracking our specific case of extracting bolts and plates. The reason is that bolt and plates constitute only a minute portion of the entire scan cloud dominated by the non-bolt class and that even only a few individual point discrepancies between classified and ground-truth bolt and plate patches could result in a distorted and unrepresentative success rate percentage. It is irrelevant to the success of the process if there is a discrepancy by a number of individual 3D scan points in the coverage between classified and ground-truth bolt and plate patches, provided the shape of the classified bolt and the shape of the plate can be accurately extracted from the classified bolt and plate patch. This is typically the case. If it can, then each such individual extraction can be added to the total count of successful extractions for QA purposes. Figure 7 illustrates this case through bolt patches of differing quality classified by the deep learning algorithm, which were suitable to be used for accurate bolt and plate positional extraction, regardless.

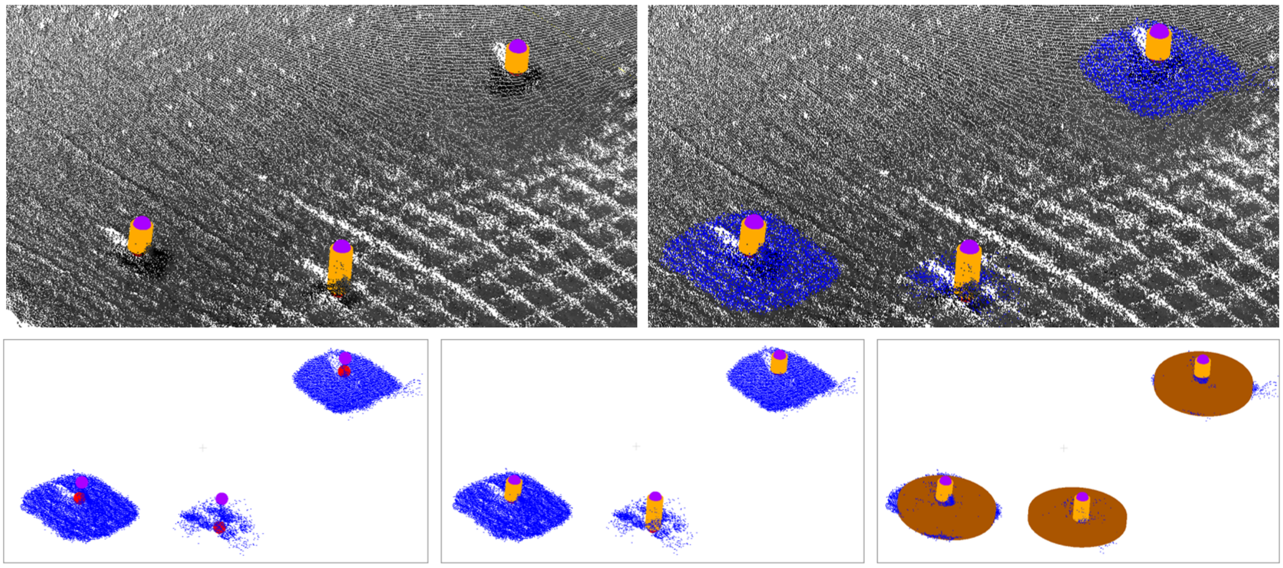

Figure 7 Example for accurate positional extraction of three bolts and plates from the classified scan cloud visible in the background in dark grey (top row). Note that there is barely any plate discernible by the naked eye in the grey scan and that the middle plate and bolt shown appear to have been installed underneath the steel mesh, which is visible in the scan cloud as a regular square pattern. This is likely the cause for the middle bolt and plate patch, classified and shown in blue, to be sparse and irregular compared to the other two patches. Yet the data was sufficient for the deep learning algorithm to extract accurate bolt tips (purple spheres), bolt to plate contact points (red spheres), and, therefore, orange bolt tips and brown plates

A more detailed view to gauge how the bolt extraction process works is provided in Figure 8. It also shows how the full bolt length can be tracked if bolt vendor specifications are provided as input. In this case embedded bolt length is added to the extracted and measured bolt tip length which allows for visualisation of compliance of bolt installation with design for any scans of newly or at least recently installed bolts for example.

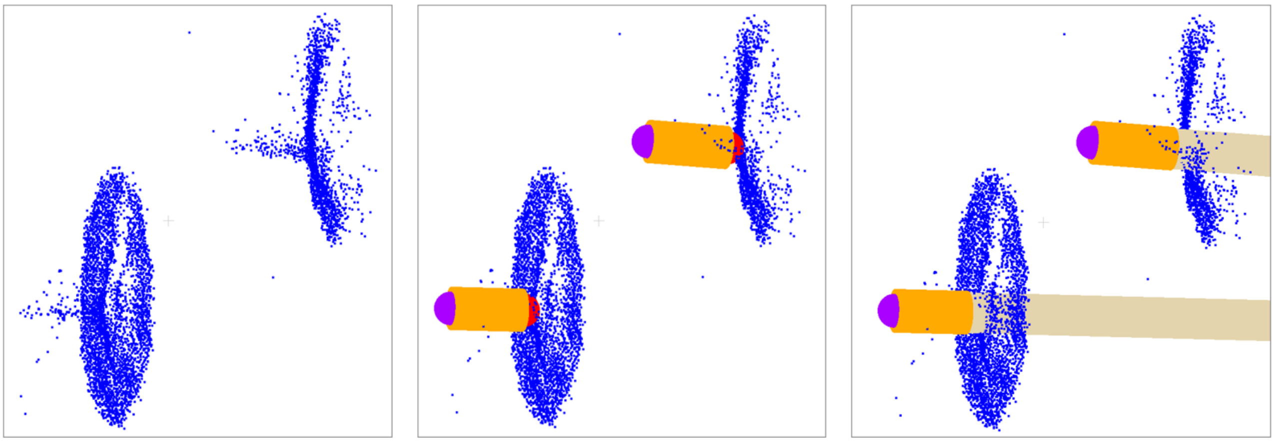

Figure 8 Side view example of sparse 3D rockbolt data of two neighbouring plates and bolts, classified by the deep learning algorithm in blue, used to automatically extract accurate bolt tips (purple spheres), bolt to plate contact points (red spheres) and, therefore, orange bolt tails. Embedded bolt sections are shown in beige

Based on the shortcomings of the native QA method of the deep learning core code, we developed our own custom QA method, based on a voxel rather than an individual scan point count, to make ascertaining classification success more representative and accurate. The procedure is as follows:

- Perform voxelisation on the point cloud, then count the number of points with ground truth and predicted labels respectively in each voxel.
- Perform maximum peak finding on the voxels with ground-truth labels. This returns the local maxima positions. Apply a threshold on the number of points in these voxels and count the number of voxels satisfying the condition (N).
- Switch to the voxels with predicted labels. Apply the same threshold on the number of points in the voxels at the previously found local maxima positions and count the number of voxels satisfying the condition (X).
- The voxel-based recall or, in other words, true positive rate is calculated as X/N.

Ideally, each voxel that passes the number-of-points-threshold should correspond to a real bolt. This is not always the case, as it also depends on the configuration parameters of voxel size and number-of-points-threshold, as well as other factors. We calibrated configuration parameters based on a test point cloud featuring visually very distinguishable and countable bolts and plates returning an output of 133 voxels, while the manual visual count returned 132 real bolts. Since these two results were very close, this configuration was adopted for the current round of voxel-based QA testing presented in this paper.

As described in the previous paragraphs, the key metric for determining the degree of success of an automated classification is expressed as recall or, in other words, the percentage of bolts and plates detected by the algorithm when compared to manually extracted ground-truth bolts and plates. When conducting such a QA process, it is naturally only possible to report the success rate using the manually extracted ground-truth data as the reference. If manual extraction missed bolts and plates that were actually contained in the point cloud, but the operator could not confidently isolate and delineate them, then the comparison recall rate can only represent a portion of the total of all bolts contained in the point cloud. Importantly, and as pointed out in Figure 6, the deep learning algorithm is demonstrably better at finding bolts and plates patches than the human operator. This means that although not reported this way, the recall rate can de facto be in excess of 100%, considering that this QA factor does not account for points found by the algorithm that are in addition to those found by the human operator.

5 Classification based rockbolt integrity tracking

Previous sections of this paper describe how our method and deep learning algorithm automatically extract rockbolt tail and plate positional characteristics from a mobile LiDAR scan cloud. These bolt tail and plate positions, by themselves, are already useful for some geotechnical considerations, such as assessing installation to design compliance, but their real value comes into play when tracking them over time. In VoidMapper, this is done by automatically allocating each extracted bolt and plate a unique ID, which is reallocated again for any subsequent scans of the same drive or void area. This allows for bolt and plate deformation tracking and trending. Figure 9 shows an example for a deformed section of drive, with deformation magnitude and direction between bolt tips shown as red vectors.

Automatically tracked bolt deformation data are the input starting point for tracking bolt integrity, as introduced in the authors' previous paper on the subject (Franke & Gonzalez 2023). Quantified and dimensioned bolt tip deformation data can be automatically correlated to vendor specifications for installed bolt elongation to failure, provided this data is used as a priori input before processing. This output quantifies a measure of failed bolts and is automatically compiled in report documents for every scan survey uploaded to the system. Since it is done for the majority if not nearly all bolts contained in a scan, and since tracking output can be viewed in VoidMapper's 3D viewer to assess any localised bolt failure patches, such integrity tracking provides a much more robust decision-making tool than what is currently best industry practice. Such bolt integrity tracking is the basis for more efficient ground support rehabilitation planning that may well have a significant positive impact on a mine's bottom line, considering ground support costs involved (Mercier-Langevin 2019).

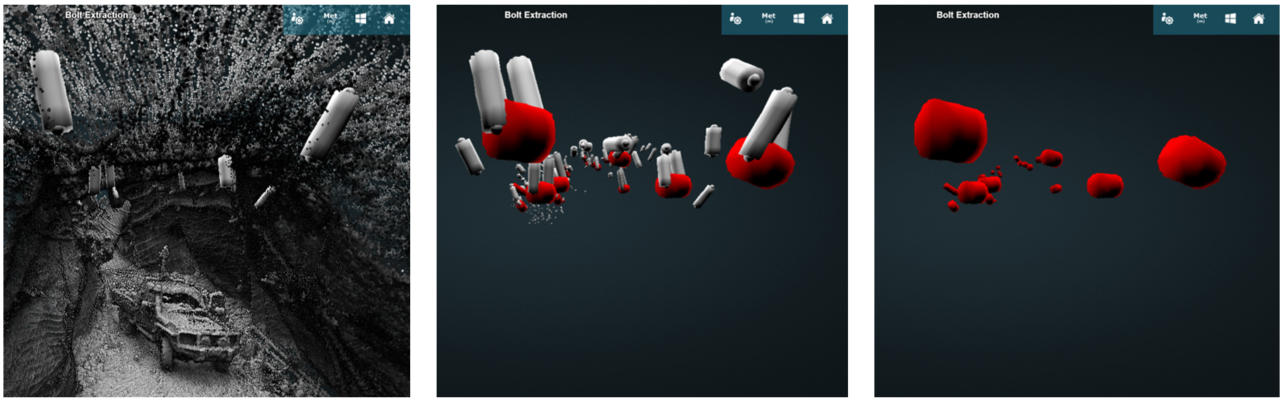

Figure 9 Example for tracking the deformation of bolt tips in two scans of a drive, collected at different dates. Deformation is shown as red vectors

6 Automated rockbolt classification case studies

Utilising the voxel QA metric for recall, as described in previous chapters, we have processed and analysed the success of our deep learning classification method in extracting rockbolts from scan clouds when compared to ground truth data for many real-world scan clouds of different mine sites featuring different ground conditions, deformation states, and types of rockbolts installed. We provide two case studies with specific details in the following.

It is worth noting that split sets, which are a commonly installed type of rockbolt, can actually be classified by the deep learning algorithm quite successfully, despite its lack of a tail protruding past the plate which gives the deep learning algorithm less of a distinct 3D feature to work with when trying to separate bolt or, in this case, mostly plate patches from non-bolts. Case B quantifies the success rate for a split set example.

6.1 Case A – classic rockbolts

Scanner:

- FARO Orbis, SLAM not Flash sensor.

Installed rockbolt type:

- One type throughout entire drive featuring classic steel bolt with plate fastened by a nut as shown in Figure 5.

Voxel-based recall representing success rate in classifying rockbolts and plates:

- 98.24%.

6.2 Case B – cable bolts and split sets

Scanner:

- Emesent Hovermap.

Installed rockbolt type:

- Two types throughout drive, with twin tailed cable bolts in certain sections of the backs and split sets in side walls and elsewhere.

Voxel-based recall representing success rate in classifying rockbolts and plates:

- Cable bolts: 99.36%.
- Split set bolts: 91.12%.

7 Conclusion

This paper has built on our previous introduction (Franke & Gonzalez 2022 & 2023) of a novel method, to provide detail on our approach of extracting accurate rockbolt tail and plate positional data. Our description covers many of VoidMapper's features and how they can contribute to a more robust method for tracking drive or other void surface deformation than continuing with unquantified or manual processing methods, currently in industry use.

We introduce in more detail our new 3D point cloud specific, deep learning classification algorithm dedicated to rockbolt and plate extraction and deformation and, with that, integrity tracking. The latter is key to ensuring safety to personnel and equipment, as well as for efficiently planning for rehabilitation.

In the case studies shown, using different mobile scanners our deep learning algorithm yields high classification success rates of 98.24% for classic steel bolt and nut type bolts, 99.36% for twin tailed cable bolts, and 91.12% for split sets. As for all other output, richly featured rockbolt, as well as drive or void surface tracking, is fully automatically generated without requirement for any manual intervention. Presentation of results includes 3D viewer data for detailed visual interrogation, as well as automatically generated reports of key observations.

References

Franke, J & Gonzalez, C 2022, 'Automated omission-free geotechnical deformation monitoring – a new method deployable by non-specialists', *Proceedings of Ausrock Conference 2022*, Australian Institute of Mining and Metallurgy, Carlton, pp. 15–17.

Franke, J & Gonzalez, C 2023, 'Automated ground support deformation monitoring: a novel method with new opportunities for geotechnical engineers', in J Wesseloo (ed.), *Ground Support 2023: Proceedings of the 10th International Conference on Ground Support in Mining*, Australian Centre for Geomechanics, Perth, pp. 75–90, https://doi.org/10.36487/ACG_repo/2325_04

Franke, J & Gonzalez, C 2024, *VoidMapper*, version 4, computer software, Caroni Geospatial, Perth, https://voidmapper.com.au

Girardeau-Montaut, D 2024, *CloudCompare*, version 2.13.1, computer software, CloudCompare, Paris, https://cloudcompare-org.danielgm.net/

Han, X, Sun, S, Song, X & Xiao, G 2018, '3D point cloud descriptors in hand-crafted and deep learning age: state-of-the-art', *arXiv Computing Research Repository*, Cornell University, Ithaca, https://arxiv.org/abs/1802.02297

Jones, EW 2020, 'Mobile LiDAR for underground geomechanics: learnings from the teens and directions for the twenties', in J Wesseloo (ed.), *UMT 2020: Proceedings of the Second International Conference on Underground Mining Technology*, Australian Centre for Geomechanics, Perth, pp. 3–26, https://doi.org/10.36487/ACG_repo/2035_0.01

Mercier-Langevin, F 2019, 'Ground support: a mine manager's perspective', in J Hadjigeorgiou & M Hudyma (eds), *Ground Support 2019: Proceedings of the Ninth International Symposium on Ground Support in Mining and Underground Construction*, Australian Centre for Geomechanics, Perth, pp. 29–40, https://doi.org/10.36487/ACG_rep/1925_0.03_Mercier-Langevin

Singh, SK 2022, *Optimising Mobile Laser Scanning for Underground Mines*, PhD thesis, University of New South Wales, Sydney.

Smith-Boughner, L, Anderson, Z, Sun, Z, Levesque, D, Kruger, S, Quenneville, Y, ... Lacroix, R 2020, 'Ground support loading monitored with Rock Bolt Sensor (RBS™) and relationship to nearby production and seismicity', in J Wesseloo (ed.), *Underground Mining Technology 2020*, Australian Centre for Geomechanics, Perth, pp. 273–284, https://doi.org/10.36487/ACG_repo/2035_12

Worksafe Victoria 2024, *WorkSafe Investigates Fatal Mine Collapse*, Melbourne, viewed 26 May 2024, https://www.worksafe.vic.gov.au/news/2024-03/worksafe-investigates-fatal-mine-collapse