geosciences

MDPI

Article

A Physics-Based Approach to Rock Bolt Detection and Spatial Monitoring

Munkhtsolmon Munkhchuluun and Davide Elmo *

NBK Institute of Mining Engineering, University of British Columbia, Vancouver, BC V6T 1Z4, Canada; tsomoo@student.ubc.ca

* Correspondence: davide.elmo@ubc.ca

Abstract

Rock bolts are the primary ground support mechanism in underground mining. Yet verification of their installation is rarely captured in a spatially precise, retrievable form, leaving operators without an auditable as-built record for regulatory review or post-incident reconstruction. This paper presents an automated rock bolt detection process that closes this documentation gap using dense point clouds from an underground hard rock mine acquired by terrestrial laser scan. The method computes per-point ambient occlusion (AO) on closure plane-sealed chambers using a PCV implementation of the ShadeVIS principle, forms candidates from a multi-scale protrusion field, and segments them by prominence watershed before classifying each candidate with PCA-based geometric descriptors, without machine learning or training data. Installation perpendicularity is applied as a per-detection confidence cue, and detections are reported in confidence tiers that concentrate human review on the ambiguous minority. Validated against a database of 1447 bolts across 20 walls in two areas of an underground mine, the system achieved an overall recall of 83.7%, with human review completing the inventory to 100%. The physics-based design transfers across bolt types and mine geometries through parameter re-tuning rather than retraining, addressing the core limitation of deep learning methods, which require site-specific labelled datasets.

Keywords: rock bolt detection; terrestrial laser scanning; underground mining; ambient occlusion; point cloud analysis; prominence watershed; PCA classification; geotechnical monitoring; convergence measurement

Check for updates

Academic Editor: John McBride

Received: 30 June 2026
Revised: 8 August 2026
Accepted: 17 August 2026
Published: 20 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Ground falls and ground control failures remain among the leading causes of serious injury and fatality in underground mining. A study of mining incidents in the period 2000–2023 attributed 2820 reported cases to ground control failures in underground mines [1]. The pattern is jurisdiction-independent: Ontario's provincial regulator identified ground instability as one of the largest causes of underground fatality [2]. These incidents do not occur solely because of insufficient ground support, but because engineers still lack reliable means to verify whether support has been installed correctly with respect to the observed structural conditions of the rock mass, or whether corrosion of support elements has compromised their capacity [3].

Rock bolts constitute the primary line of ground support in development drifts and stopes worldwide [4]. Despite their critical safety role, systematic documentation of rock bolt installation remains largely manual. Current practice largely relies on hand-drawn

sketches, visual inspection by geotechnical engineers, and paper-based logs. Gallwey et al. [5] describe these methods as not comprehensive given the number of bolts to be recorded, the time-consuming nature of manual surveying, and the human error inherent in repetitive inspection tasks. The same authors note that automatically detecting and recording the 3D coordinates of rock bolts would enable greater quality assurance and quality control, and that such records would be critical in a fall-of-ground situation where the exact bolting configuration installed prior to an incident must be determined [5].

The regulatory environment reinforces the requirement for better documentation. Mining regulations in Canada explicitly mandate systematic quality control of ground support installation. For example, Ontario Regulation 854 Section 73 requires employers to develop and maintain a Quality Control Program for underground rock bolt testing, including pull tests and verification within acceptable tolerance ranges [6]. In contrast, Ground Control Management Plans (GCMPs) require annual internal review, peer review, and biennial external audits, all of which depend on reliable, retrievable records of support installation. These requirements are difficult to satisfy in practice because installation quality is variable and the verification methods themselves are error-prone, as field evidence demonstrates.

A field review of resin-grouted rock bolts found that approximately 10% of bolts installed in a single quarter were incorrectly installed, with under-spinning as the dominant defect [7]. To complicate matters, when shotcrete is applied over a support, bolt locations are likely to be masked entirely, rendering after-the-fact verification of bolt placement and encapsulation extremely difficult [5]. Inter-observer studies of rock mass classification further show that several parameters vary significantly among trained practitioners [8], meaning that even where inspection records exist, their reliability cannot be assumed. In practice, bolt location, spacing, installation angle, and coverage are rarely captured in a spatially precise, retrievable format. When ground conditions change through blast damage, time-dependent deformation, or rehabilitation, engineers lack a reliable baseline against which to compare current conditions. Addressing this gap requires a survey technology capable of capturing the full geometric state of an underground excavation, including installed support elements, at sufficient resolution and repeatability for longitudinal comparison.

Terrestrial laser scanning (TLS) using phase-based or time-of-flight instruments is already widely deployed in underground mining operations for as-built surveys, void detection, and cavity monitoring. A single static scan produces a dense 3D point cloud capturing the full geometric surface of a drift section at millimetre-level resolution. The potential of TLS data for systematic ground support inventory, more specifically the automated detection of installed rock bolts, has not, however, been fully realized in operational practice. The use of TLS data would also serve a second objective by complementing existing convergence monitoring practices, which have historically relied on tape extensometers, which are time-consuming instruments that measure only at fixed, coarsely spaced locations. Simultaneous Localisation and Mapping (SLAM)-based mobile LiDAR mapping offers a scalable alternative to convergence monitoring but suffers from positional error that accumulates over long, geometrically uniform tunnels. The resulting positional error is understood to be a mathematical constraint rather than a sensor limitation, arising from the point-to-plane Iterative Closest Point (ICP) registration process that underpins most SLAM systems [9]. Current commercial systems achieve 6–20 mm accuracy with loop closure integration but can degrade to 1–10 m per kilometre under adverse geometry [9]. An automated rock bolt detection system addresses this constraint directly, as the dense, regularly spaced array of installed bolts provides exactly the in situ ground control features required for SLAM loop closure correction, at far greater density than conventional survey control points.

Automated rock bolt detection sits at the intersection of three research areas: geotechnical instrumentation, point cloud processing, and small object detection in computer vision. Image-based and photogrammetric approaches have demonstrated bolt detection from RGB imagery or reconstructions. Still, they depend on controlled lighting, unobstructed line of sight, and surface texture, and they detect visual appearance rather than the geometric signature of an installed bolt. Deep learning architectures designed primarily to process irregular 3D data, such as PointNet [10] and the dynamic graph convolutional neural network (DGCNN) [11], have produced the strongest reported detection performance to date. Recent applications to underground bolt detection have achieved precision in the range of 93–97% [5,12]. However, as detailed in Section 2, the same results reveal two structural constraints that impede routine deployment.

To address these constraints, this paper presents an end-to-end automated rock bolt detection methodology developed as a research prototype at the University of British Columbia. The process integrates ambient occlusion (AO) computation via the ShadeVIS algorithm as a physics-grounded pre-filter, reducing the candidate point set by more than 99% without any learning, with candidate formation from a multi-scale protrusion field and prominence watershed segmentation to isolate individual bolt attachment points, and principal component analysis (PCA)-based geometric classification using physically grounded shape descriptors calibrated from verified bolt measurements. Detected bolt positions are further applied as in situ control features for SLAM loop closure, enabling survey-grade convergence monitoring from routine scanning operations without additional instrumentation. The system achieves 83.7% automated recall validated against manual ground-truth counts across multiple drift sections, with human review completing the inventory to 100% accuracy. The physics-based design requires no training data and generalizes across bolt types and mine geometries via parameter re-tuning, addressing the core limitation of existing deep learning methods, which require site-specific labelled datasets.

2. Research Framework

The introduction established that automated rock bolt detection lies at the intersection of geotechnical instrumentation, point cloud processing, and small object detection in computer vision. This section examines each class of prior work in detail and identifies the structural limitations that motivate the physics-based approach presented in this paper. Computer vision approaches operating on RGB imagery or photogrammetric reconstructions have demonstrated bolt detection in surface and partially lit environments. In underground mining operations, however, post-blast dust, variable LED lighting, shotcrete coatings, and surface contamination severely degrade image quality. More fundamentally, image-based methods detect appearance rather than the geometric signature of an installed bolt, which limits their applicability in safety-critical inventory contexts where verification of a physical protrusion is required, not merely a visual match.

Geometric deep learning architectures such as PointNet [10], DGCNN [11], and their successors operate directly on unstructured 3D point clouds and have achieved strong results on benchmark object detection tasks. Two recent applications to rock bolt detection are directly comparable to the present work. Gallwey et al. [5] combined a neural network and a random forest classifier over 65 hand-engineered point descriptors on TLS data from an underground tin mine, reporting 93% precision and 86% recall, while noting that a large descriptor set does not necessarily improve classification because many features behave similarly for bolts and non-bolts. DeepBolt, developed by Patra et al. [12], uses a DGCNN-based graph network to detect rock bolts in underground coal mine scans, achieving 96.4% precision and 97.0% recall after approximately eight hours of GPU training on annotated data, with geometry-sensitive post-filtering to address the approximately

100:1 non-bolt-to-bolt class imbalance. These results confirm that these methods can detect bolts at high accuracy. However, as noted in Section 1, there exist two important limitations: first, labelled training data must be collected for every site, but underground environments are sufficiently non-repeating across mines, drifts, and rock types that labelling would not transfer reliably between them. Second, these models provide no traceable rationale that decision-makers can use to defend a bolt inventory against regulatory audit or incident review. Arguably, this second limitation extends to all forms of artificial intelligence applications involving professional liability considerations, though this aspect is seldom addressed in the literature [13].

Classical geometric methods such as RANSAC-based fitting and Canupo-style multi-scale classifiers [14] are more interpretable but lack a physics-grounded pre-filter to reduce the detection search space in dense underground scans. The installed geometry of a bolt is itself diagnostic in this regard since rock bolts are, in principle, installed approximately perpendicular to the excavation surface to maximize anchorage and tensile capacity, and even minor angle deviations of only 10–15° can reduce support efficiency by an estimated 15–30% [15]. Orientation relative to the rock surface is therefore simultaneously a detection cue and an installation quality indicator, which is a signal that physics-based geometric methods can exploit directly, and the basis for the approach presented in this paper.

Hybrid approaches that combine LiDAR with co-registered RGB, thermal, or RGB-D sensors can, in principle, compensate for the geometric blind spots of either modality alone. In underground inspection contexts, however, the added hardware introduces calibration overhead, additional failure modes, and operational cost without addressing the core problem: the geometric signature of an installed bolt against an irregular host surface. Multi-sensor systems are valuable where appearance and geometry must be jointly interpreted; for bolt-as-built verification, geometry alone is sufficient and more reliable.

The prevailing industrial practice remains manual or semi-automated review of point clouds in 3D viewers, supplemented by hand-drawn sketch records and visual ground inspection. Gallwey et al. [5] characterize the documentation status quo as hand-drawn sketch-based and incomplete. Visual review is subjective, time-intensive, and scales poorly with mine extent. As a result, defensible audit-grade records of bolt installation rarely exist at the spatial resolution required by regulatory or incident response reviews, as discussed in the introduction. Table 1 summarizes the key characteristics and limitations of each method, as well as the approach presented in this paper.

Table 1. Key characteristics and limitations of bolt detection methods.

Method Class	Approach	Training Data	Geometry Only	Reported Accuracy	Key Limitation
Image photogrammetry	RGB + CV; 2D detection	Partial	No	Variable	Lighting; no geometric grounding
DL Gallwey 2021	NN + RF on 65 features	Required	Yes	93% P; 86% R	Site-specific labels; not auditable
DL—DeepBolt 2025	DGCNN graph network	Required	Yes	96.4% P; 97.0% R	8 h GPU training; black-box
Multi-sensor fusion	LiDAR + camera/ thermal	Partial	No	—	Hardware complexity; failure modes
Manual inspection	Visual review + sketches	None	Yes	Subjective	Labour-intensive; not auditable
Current methodology	AO + geometric pipeline	None	Yes	83.7% R; 100% w/human	Single-site validation; parameter sensitivity

Ambient occlusion has been used extensively in computer graphics rendering and architectural visualisation, and selectively in point cloud surface roughness analysis. To the

best of our knowledge, no prior work applies ambient occlusion as a candidate-set pre-filter for small embedded object detection in point clouds, exploiting the predictable geometric shadow that a protrusive element creates at its base on an irregular host surface. Equally, no prior bolt-detection pipeline combines a physics-based pre-filter with a training-free, fully interpretable geometric refinement chain. This is the gap the present work fills; the approach is described in Section 3.

3. Methodology

The methodology is applied to point clouds segmented into discrete processing units referred to throughout as chambers, each corresponding to a fixed axial window of a development drift section. The term drift section is used to describe the physical survey unit (a length of underground excavation bounded by survey control), while chamber refers to the subdivisions into which each drift section is partitioned for point cloud processing.

Terrestrial laser scans were acquired using a FARO Focus phase-based scanner positioned at regular intervals along the excavation centerlines. Each scan position produces a registered E57 point cloud file, and all scan positions were pre-registered into a unified point cloud prior to processing. Point densities vary with distance from the scanner position, which motivates the multi-scale clustering approach described below. The detection process is summarized in Figure 1 and Table 2. Raw point cloud data enter the pipeline as input and exit as a fully verified, spatially referenced bolt inventory.

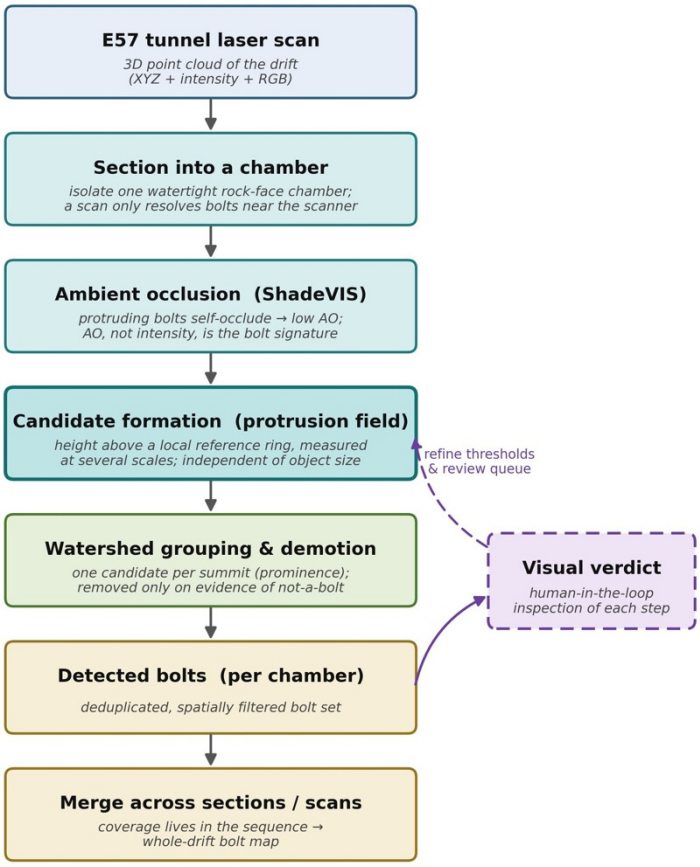


Figure 1. Rock bolt detection process using terrestrial laser scans, with particular reference to a case study (E57 tunnel) of an undisclosed mining operation.

Table 2. General stages for the detection process of rock bolts using terrestrial laser scans.

Stage	Process	Key Parameter (s)	Output
1	Side selection + chamber sectioning	Axial windows or polygon footprints	Per-chamber point cloud
2	Closure planes + ambient occlusion (PCV/ShadeVIS)	N = 256 rays; Leopardi sampling	Per-point illuminance I(p)
3	Candidate formation (multi-scale protrusion field)	terrain annulus 4–8 cm; ring ladder 6–30 cm	Candidate attachment point set
4	Prominence watershed segmentation	Minimum prominence 35 mm; admission 90 mm or 38 mm tip on 25 mm	Bolt candidates, one per summit
5	Footprint check	Fitted footprint radius 6–25 cm; 8 cm cap buffer	Candidates clear of closure membranes
6	Demotion on positive evidence + shape for disputes	See Table 3	Classified bolt candidates
7	Spatial merge	No-valley rule (no fixed separation constant)	Deduplicated bolt centre list
8	Human validation	Visual review in 3D viewer	Ground-truth verified bolt inventory (100%)

Table 3. Rules used to remove candidates, with each based on positive evidence that the object is not a bolt.

Rule	Trigger	Physical Basis
Upward-facing surface	Surface normal n_z > 0.55	Bolts are installed into walls and roof, not into the floor or muck
No rock overhead	Fewer than 6 surface cells in a rising cone 0.8–2.5 m above	Perimeter rock must exist above a wall-mounted bolt
Tilted plane	Plane fit $R^2 \geq 0.75$ and tilt ≥ 25 mm over 15 cm	The high end of a slope is terrain, not an object
Empty centre	Fewer than 3 measured points within 6 cm	A solid object has measured rock at its own centre
Close pair (disputed)	No descending height field between two candidates; weaker one requires axisymmetry ≥ 0.30	A bolt is a body of revolution; shape decides only contested cases

The first stage computes per-point ambient occlusion (AO) using the ShadeVIS algorithm, which casts rays from a sampled hemisphere above each surface point and records the proportion of unoccluded directions (Figure 2). Ray directions are generated using Leopardi equal-area sphere sampling to ensure uniform angular coverage, with N = 256 rays per point. Per-point illuminance is computed as:

$$I(p) = \left(\frac{1}{N}\right)\sum V(r_i)\ max(0, r_i \cdot n) \quad (1)$$

where $V(r_i)$ is the binary visibility function for ray i and n is the surface normal at p. The computation is deterministic (byte-identical across runs). The physical basis of this prefiltering step exploits a geometric property of installed rock bolts: attachment points are recessed fixtures that protrude into the excavation void, creating locally concave regions at their base that produce anomalously low illuminance values relative to the surrounding exposed rock surface.

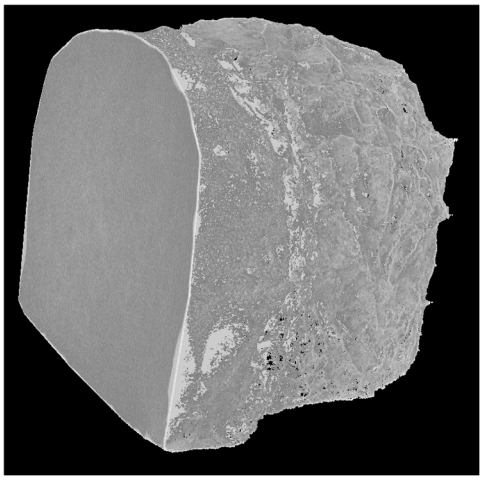


Figure 2. Example of ShadeVIS hemispheric ray sampling/closure plane for a mine tunnel 4.5 m wide, 4.5 m high.

A single fixed illuminance threshold does not transfer reliably between drift sections, since the absolute ambient occlusion scale shifts with chamber geometry and scan density. Candidate attachment points are therefore selected relative to each chamber rather than against a global value. In the configuration used for the validation reported in Section 4, candidates are formed from a multi-scale protrusion field computed on the same sealed chamber. The field is the height of each surface cell above the median of a surrounding annulus, which is measured at several ring radii so that the value is independent of the size of the object. Ambient occlusion and protrusion measure the same physical property, namely that a bolt stands proud of the rock surface; the protrusion field measures it directly rather than through the shadow it casts.

Two sources of variation are inherent to underground TLS acquisition: (i) bolt plate diameter varies across manufacturers and bolt types, and (ii) point density decreases with distance from the scanner. The spatial extent of a bolt cluster therefore varies across a single drift section. Candidates are grouped by a prominence watershed on the protrusion field, in which the field is flooded, and each summit becomes one candidate, with a minimum prominence of 35 mm required to separate two summits. This replaces the union-and-deduplicate step of the multi-scale clustering, which merged closely spaced bolts into a single over-extended cluster.

Bounding box extents along the principal axes provide length, width, and aspect ratio. In the configuration used for the validation reported in Section 4, a candidate is removed only where there is positive evidence that the object is not a bolt. Four cases apply (Table 3): the surface faces upward, so the candidate lies on the floor or on loose material; no rock stands above it, so it is not on the excavation perimeter; a tilted plane explains the local height, so the candidate is the high end of a slope; or too few points are measured at its centre for a solid object. Shape evidence settles one case only, namely two candidates close enough to be the same object, where the weaker of the two must be confirmed as a body of revolution.

Rock bolts are installed approximately perpendicular to the excavation surface to maximize anchorage, and the angle between a candidate's principal axis and the local rock surface normal is therefore simultaneously a detection cue and an installation quality indicator. This perpendicularity signal is used as a confidence weight rather than a hard rejection gate; when applied as a gate, it over-rejects genuine roof bolts and depresses recall. Each surviving candidate accordingly receives a confidence tier: high, where the axis is well aligned with the local surface normal or a bright, protruding plate signature is present; medium, where strong shape and likelihood support exist without confident

alignment; or low. A quality score combining intensity dispersion, ambient occlusion, and shape descriptors is assigned to each detection.

Two bolt candidates are merged only when the height field between them does not descend, which is a property of the two objects themselves and not of the installed pattern. No stage of the method uses the bolt pattern, either as a prior on where a bolt should be or as a filter on what to discard. The method is therefore not tuned to the support layout of the study site and could be applied at mines with different patterns without adjustment. Across chambers and overlapping scan positions, detections are merged in shared world coordinates, so that a bolt seen in two scans is counted once and full-drift coverage accrues from the scan sequence. Detected bolt axes are extended into the rock to represent the as-installed length for downstream spatial use.

Every stage of the pipeline emits a machine-parseable verdict, such as the detected count and the selected AO band, and a lockfile pins each stage's verdict alongside a cryptographic hash of the code that produced it. Any subsequent change is automatically classified as expected positional error, a regression, or no change. This makes each reported result reproducible and attributable to a specific code version, satisfying the traceability requirement that opaque learned detectors cannot fulfil for a safety-critical bolt inventory. A reviewer visually verifies high-confidence detections in a three-dimensional viewer, and medium- and low-confidence detections are adjudicated individually. This human validation step both completes the inventory to 100% accuracy and is the procedure by which the detection performance reported in Section 4 was established.

The detected bolt positions could serve as in situ control features for SLAM-based convergence monitoring, representing a natural extension of the pipeline beyond bolt inventory. SLAM-based underground LiDAR mapping accumulates positional error along long, geometrically uniform tunnels (typically 30–50 cm per 200 m of tunnel length) due to the ICP registration constraint described in Section 1. Correcting this error requires in situ reference points of known, stable position that are identifiable across multiple scan epochs. Installed rock bolts are well placed to fulfil this requirement at far greater density than conventional survey control, providing 60–100 bolt positions per 100 m of tunnel compared to 3–6 traditional ground control points. In principle, multi-epoch point cloud registration using detected bolt centres as correspondence anchors could enable SLAM loop closure and support displacement monitoring at each bolt location [16], enabling detection of deformation and proactive identification of accelerating convergence zones. Experimental validation of this application is beyond the scope of the present study and is identified as a priority direction for future work.

4. Case Study

Data were collected from an underground hard rock mine (location withheld per operator agreement). The study encompassed seven development drift sections with nominal cross-section dimensions of approximately 4.5 m × 4.5 m, with systematic rock bolt ground support installed at regular spacing consistent with the site Ground Control Management Plan. Scans were acquired from multiple static positions per drift section to ensure full perimeter coverage, with each position producing a registered E57 point cloud subsequently merged into a unified drift-level dataset prior to processing. Ground-truth bolt counts were established by an experienced geotechnical engineer performing manual visual inspection and physical counting of installed bolts in each drift section. In a representative chamber, the pipeline returned a reproducible set of 16 detected bolts with complete geometry (Figure 3). For the expanded validation reported below, two further walls surveyed in an earlier campaign were processed from a separate area of the same mine. These walls were held out; they were annotated and processed after the detection

parameters were fixed, without any adjustments. Figure 4 shows the section cuts and the ground-truth bolt database for the primary validation scan in plan view.

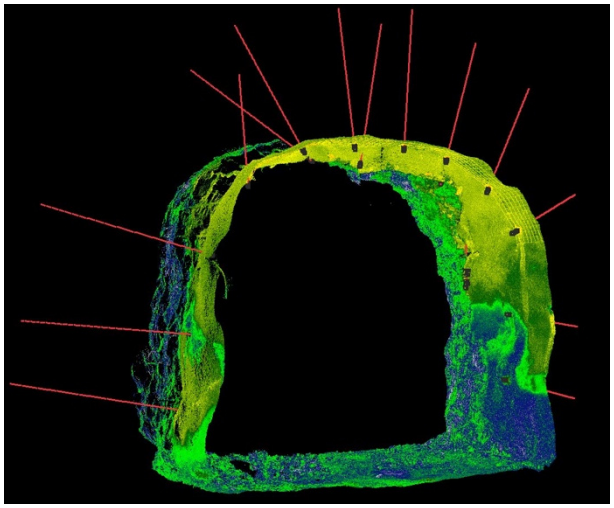


Figure 3. Detection result visualisation (point cloud with detected bolts overlaid) for a mine tunnel 4.5 m wide, 4.5 m high. In this specific example, the colours represent different shading when rotating the image in 3D.

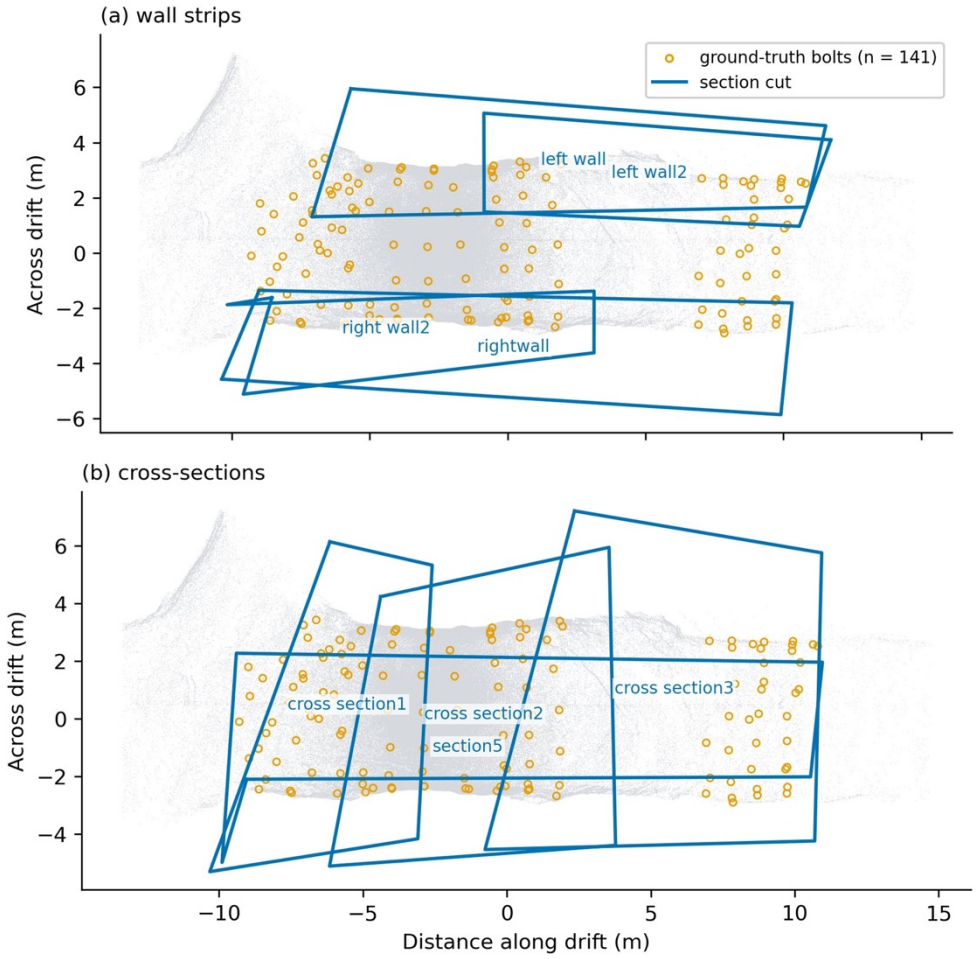


Figure 4. Section cuts and database ground-truth bolts on the primary validation scan. (**a**) in-plan view and (**b**) aligned to the drift axis.

Table 4 reports the distribution of geometric descriptors across these detections. The detected bolts are strongly rod-like, with a mean linearity of 0.89 and a mean sphericity of 0.018, and have dimensions consistent with the installed bolt type (mean length: 0.091 m; mean width: 0.022 m). All detections occur in highly occluded locations, with a mean ambient occlusion value of 0.038, thereby empirically confirming the Table 3 thresholds and the AO-based detection cue. Across the 16 confirmed detections, ambient occlusion is tightly bounded (mean 0.038, all below the 0.045 classifier threshold), and per-cluster intensity dispersion is high, with a mean intensity standard deviation of approximately 1.2×10^5. These are the two features the classifier weights most heavily, and the basis on which confirmed bolts separate from the rejected-candidate population. Note that Table 4 reports descriptor distributions for the 16 bolts detected in a single representative chamber, selected to illustrate the geometric consistency of detections; the full validation results across 20 walls are reported in Table 5.

Table 4. Geometric descriptors of detected bolts (n = 16 reference detections).

Descriptor	Mean	Range
Length (m)	0.091	0.058–0.173
Width (m)	0.022	0.016–0.029
Aspect ratio	4.3	2.80–7.42
Linearity	0.89	0.698–0.988
Sphericity	0.018	0.003–0.046
Mean AO	0.038	0.033–0.044
Points per cluster	27	8–49

Table 5. Validation results per drift section against available dataset (automated pipeline, before human review).

Drift Section	Bolts (Census)	Detections	Recall	Precision
DE11 30 June 2023	375	579	0.859	0.554
DE11 1 July 2023	4	59	0.250	0.017
DE11 9 July 2023	191	482	0.812	0.307
DE11 10 July 2023	177	304	0.797	0.474
DW11 24 June 2023	269	477	0.892	0.530
DW11 28 June 2023	121	264	0.785	0.398
DW11 13 July 2023	141	277	0.915	0.455
Held-out area (blind)	169	192	0.757	0.651
Overall	1447	2634	0.837	0.464

Detection accuracy was assessed against manual ground truth established by visual verification of detections in a three-dimensional viewer: detections were visually confirmed or reclassified as false positives where no corresponding bolt was present, and bolts present in the scan but missed by the pipeline were recorded as false negatives. For the expanded validation, ground truth was established as a census, in which an engineer marked every visible bolt in the full-density point cloud of each wall with no detector output displayed, giving 1447 bolts across 20 walls. Against the entire dataset, the proposed pipeline recovered 1211 bolts, an overall recall of 83.7% (Wilson 95% CI 81.7–85.5%), at an overall precision of 46.4% (Table 5). All figures are automated performance before human review; the review stage described in Section 3 resolves the remaining false positives and false negatives and completes the inventory.

Precision varies with the character of the wall rather than with any pipeline parameter setting. On clean drift walls (Figure 5), precision ranges from 0.64 to 0.74; on walls carrying services such as pipes, ventilation ducting, and cabling it falls to approximately 0.40, as

those fittings are bolt-scale objects that protrude from the rock face and are indistinguishable from bolts on geometric grounds alone. Precision is lowest on freshly blasted faces carrying few bolts, where blast-induced surface relief (the irregular micro-topography left by explosive fragmentation) dominates an otherwise bolt-free surface; on one such face, 72 detections were returned against a ground truth of 12 bolts. Figure 6 shows a face-on view of a representative validated wall at full scan density, illustrating the spatial distribution of true positives, false positives, and missed bolts. Recall is more stable, ranging between 0.83 and 0.96 across all wall types except the held-out area, which is discussed in Section 5.

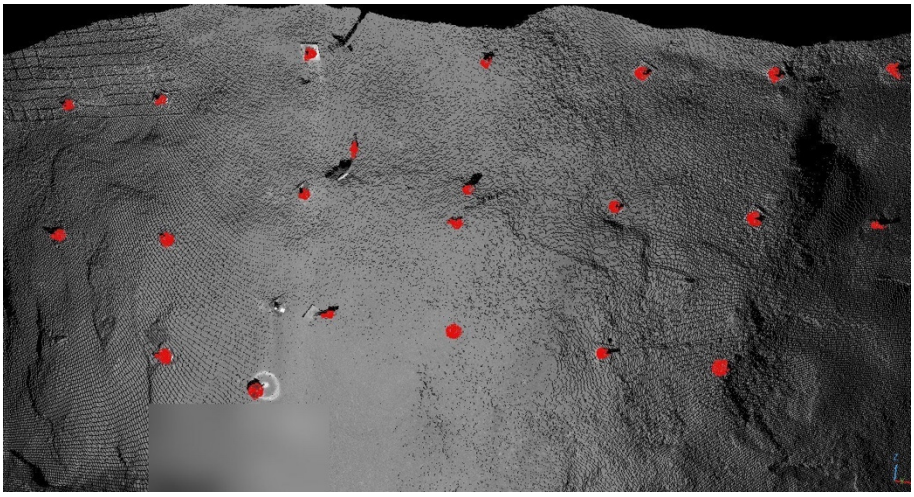


Figure 5. Example of detected bolts overlaid on a coloured point cloud (22 detected bolts, shown in red, including 2 missed bolts and 3 false positives).

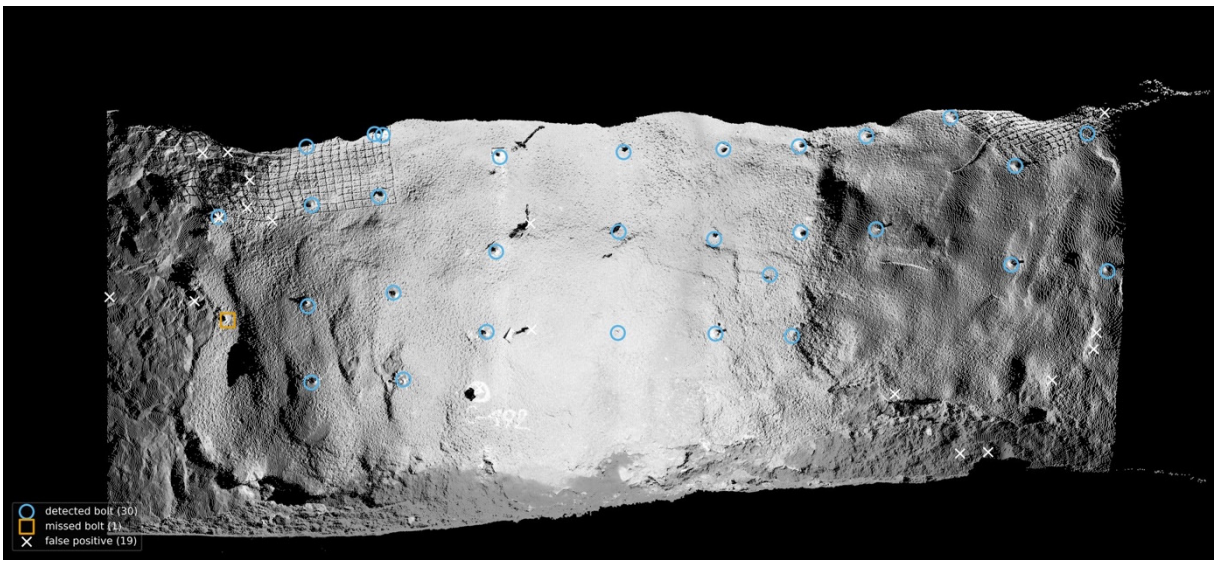


Figure 6. Face-on view of a validated wall at full scan density. Circles mark detected bolts, squares mark bolts the pipeline missed, and crosses mark false positives.

The detected and rejected candidate bolts were compared across the complete database using four geometric descriptors: (i) axisymmetry about the surface normal, (ii) peak protrusion height, (iii) point count at the candidate centre, and (iv) a shape descriptor. All four descriptors separate the two populations (Mann–Whitney U, $p \leq 5 \times 10^{-28}$). Axisymmetry about the surface normal (a measure of how symmetrically the candidate protrudes relative to the local rock surface) is the strongest discriminator, with Cliff's δ

of +0.45 and +0.57 on the two largest sections. Median axisymmetry is 0.63 for confirmed bolt detections and 0.35 for false positives. Peak protrusion height is the second strongest discriminator (Cliff's δ +0.33 and +0.42), followed by point count at the candidate centre (+0.38 and +0.41). Sorting the human review queue by axisymmetry therefore concentrates reviewer effort where discrimination is most uncertain.

A sensitivity analysis was carried out by varying each detection threshold in turn by 25 to 50 per cent around its adopted value on the largest validation section. Bolt recall remains between 0.83 and 0.93 and precision between 0.41 and 0.58 across this range. The watershed prominence is the single parameter that moves recall and precision against each other; performance degrades gradually across the tested range rather than collapsing at any single threshold value. The false-positive population is a structural property of the rock surface geometry rather than a consequence of parameter choice, and it cannot be eliminated by threshold adjustment alone.

Processing time was measured on a laptop-class CPU (Intel i5-1145G7, 8 threads, 16 GB RAM, no GPU). Ambient occlusion computation on a chamber of approximately 800,000 points requires 97 s; the complete preprocessing chain requires approximately ten minutes per chamber from cold; and detection itself requires 25 s per chamber. No training phase is required at any site, so these figures represent the total processing time for a new drift section without exception.

The effect of the perpendicularity criterion on pipeline performance merits specific comment. On the reference chamber, applying perpendicularity as a hard rejection gate reduced the detected set from 16 to 9, incurring a recall loss due to genuinely installed roof bolts whose recovered axes deviate from the local surface normal. The reproducibility harness flagged this change automatically and attributed it to the gate. Demoting perpendicularity to a confidence weight restored the full set of detections while preserving the orientation information as a per-detection score. This result directly motivates the design choice described in Section 2.

The candidate formation stage remains unsupervised and adapts to each chamber without manual intervention since the protrusion field is measured relative to the local rock surface rather than against an absolute scale. The detection thresholds listed in Table 2 are physical lengths taken from the geometry of the installed bolt type. Therefore, deployment to a site with a different bolt product requires a one-time recalibration against a small verified sample before automated processing proceeds.

5. Discussion

This section interprets the results reported in Section 4 in the context of existing approaches to automated rock bolt detection. The first part examines the technical and practical innovations of the physics-based pipeline relative to machine learning methods, including its generalisability, interpretability, and secondary application to convergence monitoring. The second part addresses the limitations of the current work, identifying the boundaries of the validation scope and the conditions under which the pipeline requires recalibration or human oversight before operational deployment.

5.1. Innovation and Practical Contribution

The absence of a widely adopted commercial bolt detection solution, despite a growing body of research, is attributable to the difficulty of generalising applications across different mine sites. Three factors explain why the physics-based approach described here generalizes more readily than existing deep learning methods:

1. The proposed method requires no training data. Deep learning methods require labelled datasets of underground bolts to train segmentation or detection models;

such datasets are scarce, site-specific, and expensive to produce under the constraints of underground safety and operations. The physics-based methodology requires only the physical dimensions of the bolt type installed at a given site to calibrate the PCA criteria. This information is available from standard bolt specifications.

2. Failure modes are interpretable. When the proposed approach fails to detect a bolt, the reason is traceable to a specific stage: insufficient point density at range, cluster merging with adjacent geometry, or plate deformation beyond the dimensional criteria. This interpretability allows a site engineer to diagnose and address failures without needing to understand the internal mechanics of a neural network.
3. Bolt geometry is physically invariant. Rock bolt dimensions (plate diameter, diameter, and exposed length) are constrained by product and installation standards that hold regardless of mine site, rock type, or geographic location. The PCA classification criteria encode these physical invariants directly. In contrast, deep learning models tend to overfit to site-specific appearance features such as reflectivity, surface texture, and corrosion pattern that do not transfer across sites.

The approach was calibrated to a specific mine and bolt product range. The detection thresholds reflect the bolt geometry and scan density of the primary study sections, and application to a site with different bolt specifications could require recalibration, though the underlying physical reasoning carries over directly. The held-out area quantifies this. Its bolts protrude approximately half as far as those in the study sections, with median values of 62 mm and 122 mm, respectively. Without adjustment, the pipeline recovered 72 per cent of the bolts on the drift wall and all of them on the end wall. Scaling the thresholds to the measured bolt class raised recall on the drift wall to 94 per cent but reduced precision from 0.74 to 0.62; applying the same scaling to the end wall, which did not require it, reduced precision from 0.42 to 0.12. Recalibration is therefore a deliberate trade-off to be made per wall type, guided by the measured protrusion of confirmed bolts, and not an automatic step. The support layout is a separate matter: since the bolt pattern is not used anywhere in the method, a site with different spacing requires no change on that account.

The secondary contribution proposed in this paper (using detected bolt positions as in situ control features for SLAM localisation error correction) has broader implications for underground geotechnical monitoring practice if validated experimentally. Survey-grade convergence monitoring has historically required either installation of physical targets such as prisms and reflectors, or dedicated survey crews with total station equipment, both of which are episodic and resource-intensive. The dense, regularly spaced array of installed bolts offers a candidate solution to this problem that would require no additional instrumentation beyond the existing scanning program. Demonstrating this application experimentally across multiple scan epochs, with comparison against prism or total station baselines, is a priority follow-on activity to the present work.

The 3D bolt inventory described in this paper constitutes a spatially anchored geotechnical knowledge layer. When tied to a mine centreline coordinate system and timestamped at each scan, the resulting dataset enables longitudinal tracking of ground support conditions—detecting degradation, missing bolts, spacing violations, and rehabilitation needs without reliance on manual inspection cycles. This connection between automated geotechnical sensing and persistent spatial data infrastructure represents a direction for future work with particular implications for institutional knowledge retention in operations with high personnel turnover.

5.2. Limitations

The validation reported in Section 4 was conducted on TLS data from a single underground hard rock mine, and the detected bolt population, host geology, drift geometry,

support pattern, and scanner configuration are accordingly drawn from a single operational context. This is a recognized limitation, though it is worth noting that single-site or single-dataset origins are not uncommon in geotechnical practice [8,13]. The expectation that a first-validation automated detection system should demonstrate cross-site generalisability before publication would be inconsistent with the standard applied to foundational methods in the discipline [17].

The present work establishes that the physics-based design is inherently transferable via parameter re-tuning rather than retraining. Deployment to a new site therefore requires a one-time expert recalibration step against a small verified sample before automated processing proceeds. The underlying physics of the ambient occlusion pre-filter is geometry-invariant and does not depend on rock type or bolt material, which is why re-tuning rather than retraining is sufficient. Shifts of more than approximately 50% in the AO percentile threshold produce material changes in the detected set, confirming that the recalibration step is not optional. Notwithstanding, cross-site generalisability remains to be established through validation on additional mine sites with differing rock conditions, bolt types, and scan acquisition protocols. A multi-site validation program is a priority follow-on activity to this work.

The dataset produced after the AO stages, before the statistical scoring filter is applied, may contain a substantial proportion of non-bolt features (rock surface fractures, crevices, and irregular geometry that share the locally occluded signature of installed bolts). The recall of 83.7% reported in Section 4 reflects the performance after this stage has filtered the candidate set. The approach relies on the scoring stage for discrimination, and any modification or removal of that stage could produce many false positives. However, such a deliberately permissive pre-filter is designed to preserve recall across heterogeneous rock surfaces, allowing human decisions to remain a crucial part of the validation step.

The current approach operates on completed scan datasets in batch mode. Real-time operation during scan acquisition would require GPU-accelerated ambient occlusion computation and incremental clustering, both of which are achievable in principle but not implemented in the present work. The current implementation is well-suited to post-acquisition analysis at an end-of-shift or end-of-day cadence; in-cycle scanning workflows that demand detection within the acquisition window remain outside the scope of this contribution.

Finally, the case study mine uses standard rebar-and-plate rock bolts with approximately 20–30 mm diameter and face plates of approximately 150 mm width. The approach detects these bolts via the geometric signature of the plate-and-rebar assembly. Performance on alternative bolt types, such as friction-set bolts (Split Set or Swellex), cable bolts, or bolts with non-standard plate geometries, has not been completed at this stage. Bolts that have been over-coated with shotcrete are not detectable by any geometry-based method, including the one presented in this paper; the system is intended for verification before shotcrete application or in operations that scan before each shotcrete cycle.

6. Conclusions

This paper presented a physics-based, training-free pipeline for automated rock bolt detection from routine terrestrial laser scan data, addressing the as-built ground support documentation gap that current manual practice cannot reliably close. The system achieved an overall recall of 83.7% against a census ground truth of 1447 bolts across 20 walls, with human review completing the inventory.

Three conclusions stand out. First, a physics-grounded pre-filter can detect bolts without any learning. Per-point ambient occlusion computed on closure plane-sealed chambers, combined with PCA geometric classification, isolates recessed bolt-attachment zones and identifies installed bolts reliably. The conjunction of the occlusion signal and

shape descriptors is more discriminative than either signal alone, and the approach requires no labelled training data, addressing the core limitation of existing deep learning methods.

Second, per-chamber adaptive band selection makes the method transferable across sites. Selecting the ambient occlusion candidate band per chamber using a percolation model, rather than applying a fixed global threshold, allows the same pipeline to carry across chambers and sites through parameter re-tuning rather than retraining. This is precisely the property that deep learning detectors lack, and it is what makes the physics-based approach a practical candidate for deployment across the varied geometries and bolt types encountered in operational underground mining.

Third, auditability is a first-class outcome of the pipeline design. Verdict-pinned stages make every reported result reproducible and attributable to a specific code version. This property was demonstrated concretely by the automatic detection and attribution of the perpendicularity recall regression reported in Section 4, where demoting the perpendicularity criterion from a hard gate to a confidence weight restored full recall without any loss of orientation information. This level of traceability is what qualifies the pipeline output as a defensible as-built record suitable for regulatory audit and incident review—a standard that opaque learned detectors cannot currently meet.

Future work will extend validation across additional drift sections, mine sites, and bolt types to establish the cross-site generalisability that the current single-site results cannot yet confirm. Downstream applications enabled by the resulting spatial bolt inventory, including longitudinal condition monitoring, rehabilitation planning, and integration with mine-wide geotechnical data infrastructure, represent a significant practical direction for ground control management in underground hard rock mining.

Author Contributions: Conceptualisation, M.M.; methodology, M.M.; software, M.M.; formal analysis, M.M.; data curation, M.M.; writing—original draft preparation, M.M. and D.E.; writing—review and editing, M.M. and D.E.; visualisation, M.M.; supervision, D.E. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The code is available from the corresponding author upon reasonable request. However, the datasets presented in this article are not readily available due NDAs signed by the authors.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

DGCNN	Dynamic graph convolutional neural network
GCMP	Ground Control Management Plan
ICP	Iterative Closest Point
LiDAR	Light detection and ranging
SLAM	Simultaneous Localisation and Mapping
TLS	Terrestrial laser scanning

References

1. Mensah, E.K.; Yeboah, A.; Saah, B.P.; Nii-Okai, E.; Miezah, A.A.; Oppong Wiafe, G.F. Trend Analysis of Accident Frequency and Severity in Surface and Underground Mines Using Publicly Available Safety Databases. *J. Front. Multidiscip. Res.* **2025**, *6*, 520–529. [CrossRef]
2. Ontario Ministry of Labour, Immigration, Training and Skills Development. Compliance Initiative Results: Ground Control. 2025. Available online: https://www.ontario.ca/page/compliance-initiative-results-ground-control (accessed on 28 June 2026).
3. Aziz, N.; Craig, P.; Nemcik, J.; Hai, F. Rock bolt corrosion—An experimental study. *Min. Technol.* **2014**, *123*, 69–77. [CrossRef]

4. Li, C.C. Principles of rockbolting design. *J. Rock Mech. Geotech. Eng.* **2017**, *9*, 396–414. [CrossRef]
5. Gallwey, J.; Eyre, M.; Coggan, J. A machine learning approach for the detection of supporting rock bolts from laser scan data in an underground mine. *Tunn. Undergr. Space Technol.* **2021**, *107*, 103656. [CrossRef]
6. Ontario Regulation 854–Mines and Mining Plants, s.73: Quality Control Program, R.R.O. 1990. Available online: https://www.ontario.ca/laws/regulation/900854 (accessed on 28 June 2026).
7. Maepa, T.G. Installation of resin-grouted rockbolts in hard rock mining: Challenges and solutions for improved safety. *J. S. Afr. Inst. Min. Metall.* **2017**, *117*, 329–336. [CrossRef]
8. Yang, B. Examining the Reliability of Integrating Machine Learning with Rock Mass Characterization and Classification Data (T). Ph.D. Thesis, University of British Columbia, Vancouver, BC, Canada, 2024. Available online: https://open.library.ubc.ca/collections/ubctheses/24/items/1.0447167 (accessed on 28 June 2026).
9. Tuna, T.; Nubert, J.; Pfreundschuh, P.; Cadena, C.; Khattak, S. Informed, constrained, aligned: A field analysis on degeneracy-aware point cloud registration in the wild. *arXiv* **2024**, arXiv:2408.11809.
10. Qi, C.R.; Su, H.; Mo, K.; Guibas, L.J. PointNet: Deep learning on point sets for 3D classification and segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*; IEEE: New York, NY, USA, 2017; pp. 652–660.
11. Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S.E.; Bronstein, M.M.; Solomon, J.M. Dynamic graph CNN for learning on point clouds. *ACM Trans. Graph.* **2019**, *38*, 1–12. [CrossRef]
12. Patra, D.; Ranasinghe, P.; Banerjee, B.; Raval, S. A deep learning approach to identify rock bolts in complex 3D point clouds. *Remote Sens.* **2025**, *17*, 2701. [CrossRef]
13. Elmo, D.; Adams, S.K. Rock Engineering Knowledge and Radical Uncertainty: From Empirical Methods to Professional Practice. *Geosciences* **2026**, *16*, 73. [CrossRef]
14. Brodu, N.; Lague, D. 3D point cloud classification of complex natural scenes using a multi-scale dimensionality criterion: Applications in geomorphology. *EGU Gen. Assem.* **2012**, *14*, 4368.
15. Isfahani, N.S.; Azhari, A.; Motra, H.B.; Hashemalhoseini, H.; Hosseinabadi, M.H.; Baghbanan, A.; Bazargan, M. Optimizing Rock Bolt Support for Large Underground Structures Using 3D DFN-DEM Method. *Geosciences* **2025**, *15*, 293. [CrossRef]
16. Wang, S.; Zheng, D.; Li, Y. LiDAR-SLAM loop closure detection based on multi-scale point cloud feature transformer. *Meas. Sci. Technol.* **2025**, *35*, 036305. [CrossRef]
17. Elmo, D.; Zoorabadi, M. Examining the case for accuracy and precision when determining the Hoek-Brown parameters. *Rock Mech. Rock Eng.* **2026**, *in press*. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.