SINTEF Community
Postal address:
Postboks 4760 Torgarden
7465 Trondheim, Norway
Switchboard: +47 40005100
community@sintef.no
Enterprise /VAT No:
NO 919303808 MVA

# Report

# Automated Detection of Geological Discontinuities and 3D Orientation Recovery from MWD Data

**VERSION**
2.0

**DATE**
2026-06-26

**AUTHORS**
Dirk van Oosterhout
Yared Bekele

**PROJECT NO.**
DINAMINE-101091541

**NO. OF PAGES/APPENDICES:**
22

## Abstract

This report presents an end-to-end unsupervised workflow to detect geological discontinuities and recover their three-dimensional orientations directly from Measurement-While-Drilling (MWD) data. The method integrates physics-informed pre-processing, transparent feature engineering, a consensus ensemble of anomaly detectors, and iterative multi-plane RANSAC fitting.

The approach is validated in two stages. First, detection is demonstrated on real boreholes from a Norwegian tunnel, where identified discontinuities coincide with independently catalogued optical-televiewer intervals. Second, a real-derived synthetic benchmark enables quantitative assessment of orientation recovery.

Results show that discontinuity planes are recovered with sub-degree accuracy when fracture aperture exceeds the sampling resolution, while detection completeness degrades sharply below this threshold. The method operates without labelled data and provides a scalable pathway for turning routine drilling telemetry into actionable geological information.

**PREPARED BY**
Dirk van Oosterhout, Yared Bekele

SIGNATURE

**APPROVED BY**
Mario Morales

SIGNATURE

**PROJECT MEMO NO.**
Project deliverable D3.1 (Annex)

**CLASSIFICATION**
Unrestricted

# Document history

| VERSION | DATE | Version description |
| --- | --- | --- |
| 1.0 | 2025-11-28 | Final version |
| 2.0 | 2026-06-26 | Updated version |

SINTEF

# Table of contents

# 1 Introduction

Measurement-While-Drilling (MWD) systems record drilling parameters at fine depth resolution during excavation, providing a continuous and spatially explicit representation of the rock mass along each borehole. These data constitute a dense record of both geological conditions and operational drilling behaviour, obtained without additional instrumentation or interruption to production.

Geological discontinuities, including fractures, joints and weakness zones, are among the most influential features controlling excavation performance. Their spatial distribution and orientation govern fragmentation, overbreak, support requirements and groundwater inflow. However, conventional mapping methods rely on post-excavation observation and are therefore retrospective, limited to exposed surfaces, and unable to resolve the three-dimensional geometry of discontinuities within a blast round.

This report presents an end-to-end methodology to detect discontinuities and recover their three-dimensional orientations directly from raw MWD data. The method operates without labelled training data and transforms drilling telemetry into a volumetric representation of the rock mass ahead of excavation.

# 2 Methodology

The methodology for detecting rock fractures from Measurement While Drilling (MWD) data is executed via an automated Python pipeline. This pipeline transforms raw, time-series drilling data into a set of processed features, which are then analysed using a multi-model unsupervised anomaly detection approach. The method follows a five-stage unsupervised workflow transforming raw drilling signals into three-dimensional discontinuity planes:

1. Raw data processing
2. Feature engineering
3. Ensemble anomaly detection
4. Anomaly aggregation
5. Discontinuity-plane fittin

Each stage consumes the output of the previous stage, forming a consistent chain from raw signals to geometric interpretation.

Two guiding principles underpin the design:

- The workflow is fully **unsupervised**, requiring no labelled fracture data
- Each stage is **explicit and configurable**, avoiding black-box processing


**Figure 1. Fracture plane detection workflow.**

## 2.1 Stage 1: Raw Data Processing

The first stage converts heterogeneous drilling records into a unified representation and removes non-geological artefacts that would otherwise dominate the analysis. Raw MWD data are ingested into a structured format that separates borehole geometry, depth-indexed drilling telemetry and processing metadata. This separation allows geometric interpretation to be decoupled from signal analysis and ensures consistency across different data sources.

Pre-processing is treated as a critical stage rather than a preliminary cleaning step, because the drilling signal contains systematic patterns that arise from the drilling process itself. These include the collaring phase, during which the drill establishes contact with the rock under non-standard conditions, the periodic disturbances introduced by rod extensions, and slow trends associated with drill control systems. If not removed, these effects can mask or mimic geological signatures.

The removal of the collaring interval is performed adaptively on a per-borehole basis. Instead of applying a fixed depth cut, the method monitors the evolution of control parameters, specifically feed and hammer pressures, to determine when the drilling system has reached stable operating conditions. The working level of each parameter is estimated from a reference depth window further along the borehole, and the collaring interval is defined as the portion of the signal before both parameters reach and maintain this level within a narrow tolerance band. This adaptive approach avoids both premature truncation of valid data and retention of start-up artefacts, which vary between boreholes.

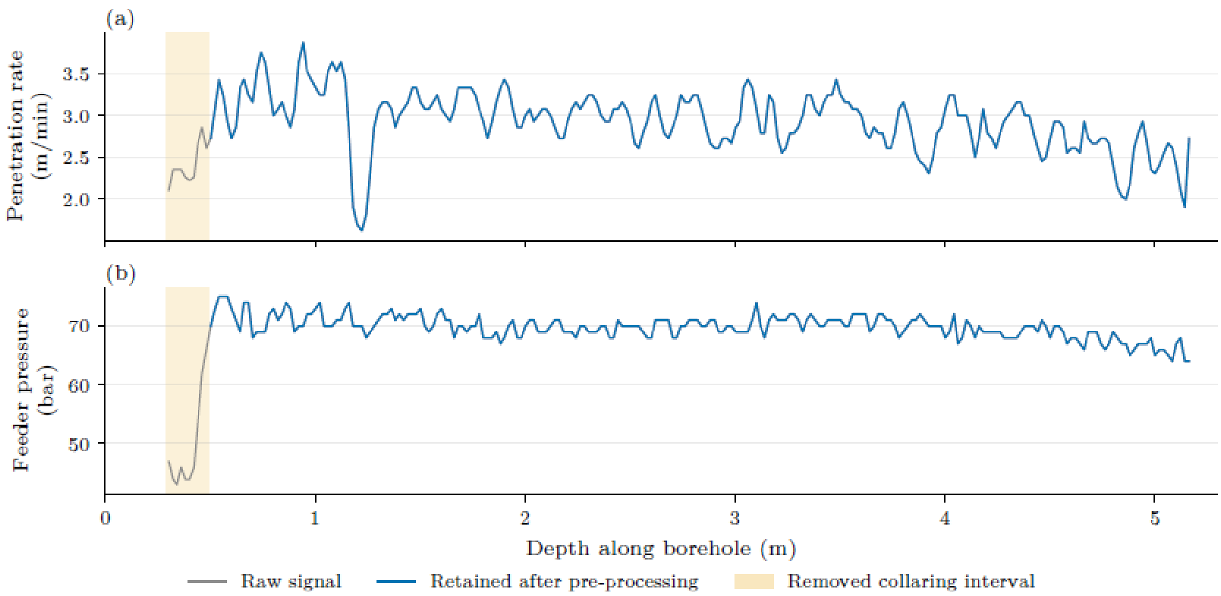


**Figure 2. Raw and pre-processed MWD signals for the real borehole BH7: (a) penetration rate and (b) feeder pressure. The shaded interval at the hole start is the collaring interval removed by pre-processing, within which the rig is still establishing the hole. No rod-junction signature appears in these records and no sample falls outside the physical-range limits, so no rod-junction windows and no implausible values are removed.**

Rod-change disturbances are identified within each rod segment by detecting deviations in feed pressure relative to the segment baseline. These disturbances are removed by excluding a short depth window around each detected coupling. The remaining samples are labelled with a rod index, allowing subsequent stages to account for segment-wise variations in the signal.

Outlier removal is applied to eliminate physically implausible values or extreme deviations that are unlikely to have geological origin. Two complementary strategies are employed, one based on predefined physical limits for each drilling channel and another based on statistical thresholds relative to the borehole-wide distribution. In both cases, only clearly erroneous samples are removed, preserving the integrity of the signal for further analysis.

## 2.2 Stage 2: Feature Engineering

Following pre-processing, the cleaned drilling signals are transformed into a feature representation that emphasises the localised disturbances associated with discontinuities. Rather than relying on raw channel values alone, the method constructs a set of derived features that capture both the magnitude and the variability of the drilling response.

For a dataset comprising six drilling parameters, a total of fifty-six features are generated for each depth sample. These features are organised into five conceptual families that together provide a comprehensive description of the signal. The first family places all channels on a common scale through standardisation and min–max normalisation, ensuring comparability across parameters with different units and dynamic ranges. The second family captures local variability, which is the most diagnostic characteristic of discontinuities. Gradients with respect to depth quantify changes between consecutive samples, while rolling means and rolling standard deviations computed over trailing windows represent local trends and fluctuations. Two window sizes are used to ensure sensitivity to both narrow and broader disturbances.

The third family consists of pressure ratios formed between the scaled hammer, feed and rotation pressures. These ratios encode relative changes between different mechanical components of the drilling process, providing additional contrast for detecting anomalies that affect multiple channels differently.

The fourth family removes large-scale operational trends through rod-wise normalisation. Within each rod segment, smooth relationships between drilling parameters and depth or control variables are estimated using simple regression models. The original signals are then normalised by these trends, isolating local deviations that are more likely to have geological origin.

The fifth component is a composite fracturing index, designed to capture the increase in local variability that typically accompanies fractured or weakened rock. This index combines measures of local deviation in penetration rate and rotation pressure, normalised by dataset-wide statistics to ensure comparability between boreholes.

Before anomaly detection, the full feature matrix is standardised once more to remove differences in scale and to mitigate redundancy between correlated features.

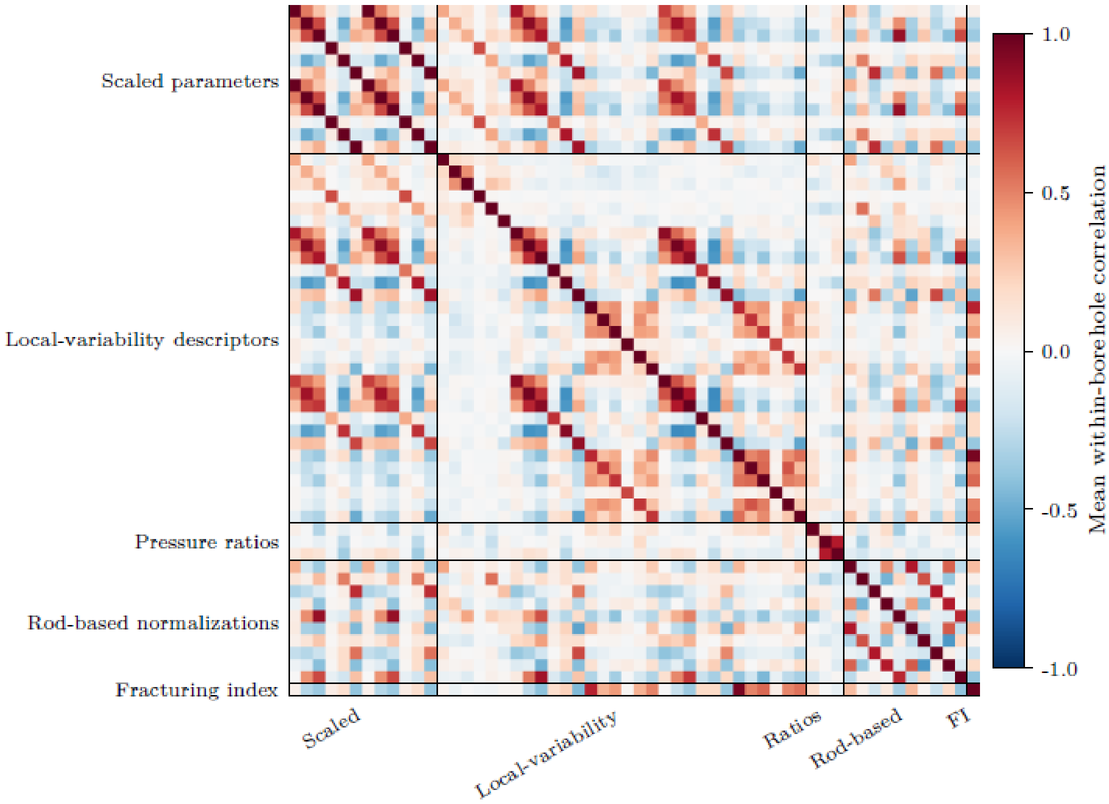


**Figure 3. Correlation structure of the fifty-six engineered features, computed within each of the six real boreholes of the detection demonstration and averaged across them so that the per-hole scaling of several families is not mistaken for a relationship between features.**

## 2.3 Stage 3: Ensemble Anomaly Detection

Discontinuity detection is formulated as an anomaly detection problem, reflecting the fact that fractures manifest as localised deviations from the dominant drilling response rather than as recurring classes with easily defined labels.

The methodology employs an ensemble of four unsupervised anomaly detectors, each capturing a different notion of abnormality. The Isolation Forest identifies samples that can be separated from the bulk of the data with few partitioning operations, reflecting their isolation in feature space. The Local Outlier Factor measures the density of the local neighbourhood around each sample, flagging points that lie in regions of comparatively low density. DBSCAN labels points that do not belong to any dense cluster as noise, while the one-class support vector machine constructs a boundary around the normal data and identifies samples that fall outside it.

Because each detector exhibits distinct strengths and limitations, their outputs are combined using a consensus mechanism. A depth sample is considered anomalous only when at least half of the detectors agree. In addition to binary flags, each detector produces a continuous score that is rescaled to a common range and averaged to form a consensus anomaly score. This score provides a graded measure of deviation while the consensus flag defines the subset of samples carried forward to the next stage.

The consensus threshold is deliberately set to be sensitive rather than conservative. The objective at this stage is not to produce a final classification but to ensure that potential discontinuity signatures are retained for aggregation and geometric interpretation.

## 2.4 Stage 4: Aggregation of Anomalies into Fracture Events

Individual anomalous samples typically occur in short sequences when the drill crosses a discontinuity. These sequences are aggregated into discrete fracture events to provide a more physically meaningful representation.

The aggregation proceeds by merging consecutive anomalous samples when the depth gaps between them are small relative to the sample spacing. Additional merging steps combine neighbouring groups that are separated by small gaps, ensuring that interruptions in the anomaly score do not split a single discontinuity into multiple events. Each resulting event retains its depth extent, allowing the thickness of the disturbed zone to be preserved.

To suppress residual noise, a minimum event size threshold is applied, filtering out isolated or very short anomaly sequences that are unlikely to correspond to meaningful geological features.

## 2.5 Stage 5: Discontinuity-Plane Fitting

The final stage reconstructs the three-dimensional geometry of discontinuities by combining information from multiple boreholes. Each fracture event is converted into a spatial point using the borehole geometry, advancing from the collar position along the borehole trajectory by the event depth.

The resulting set of points forms a sparse three-dimensional point cloud in which clusters of points lie along discontinuity planes. These planes are extracted using an iterative Random Sample Consensus (RANSAC) procedure. At each iteration, a candidate plane is fitted to a minimal subset of points, and the number of inliers within a specified distance threshold is counted. The best-supported plane is retained, its inliers are removed, and the process is repeated on the remaining points.

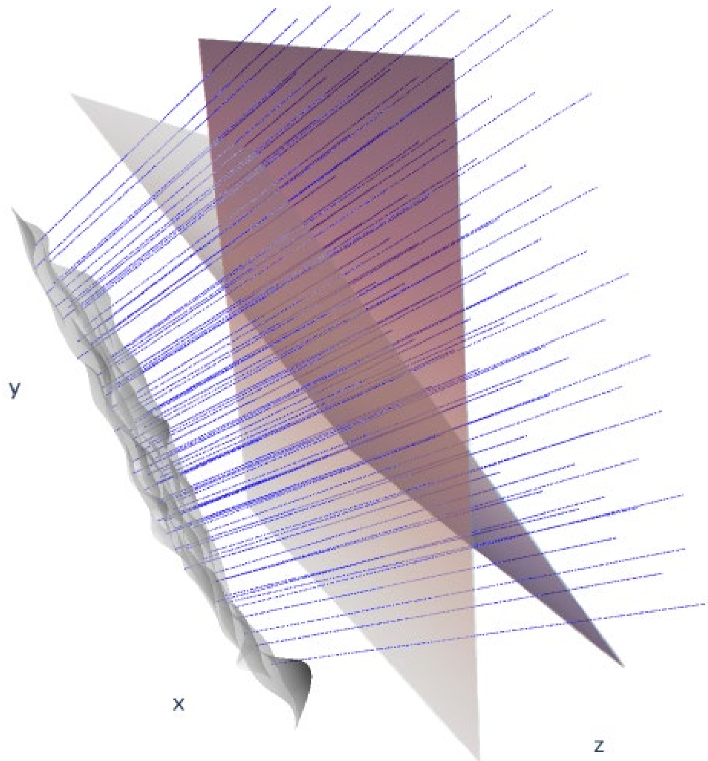


**Figure 4. Fracture plane detection using RANSAC, based on simulated data.**

To avoid over-fitting and to ensure that only meaningful structures are reported, a prioritisation step is applied. Planes with insufficient support are discarded, and near-parallel planes may be merged or filtered depending on their relative support. The final output consists of a set of discontinuity planes characterised by their orientation and the number of supporting observations.

## 3 Dataset and Benchmark

The methodology is demonstrated using real MWD data collected from tunnel excavation in Norway, complemented by a synthetic benchmark designed to provide ground truth for evaluation.

The real dataset consists of boreholes drilled in sedimentary formations and recorded at centimetre-scale resolution. Optical televiewer logs are used to identify fracture intervals independently from the MWD data, providing a reference for validating the detection stage.

Because real datasets lack dense three-dimensional ground truth for discontinuity planes, a synthetic benchmark is constructed. This benchmark inherits the statistical properties, background trends and noise characteristics of the real data while allowing discontinuity planes with known geometry to be inserted. Borehole geometry is derived from a realistic blast pattern, ensuring that the spatial configuration reflects operational conditions

# 4 Demonstration and Evaluation

The methodology is evaluated through a two-part demonstration that reflects both the capabilities and the intrinsic limitations of the available data. The real boreholes provide genuine drilling signals together with independently catalogued fracture intervals derived from optical televiewer logs, while the synthetic benchmark supplies the dense three-dimensional ground truth required to quantitatively assess orientation recovery. The evaluation therefore separates the validation of detection from the quantification of geometric reconstruction, ensuring that each claim is supported by appropriate evidence.

## 4.1 Evaluation Framework and Metrics

The evaluation uses different criteria for the two parts of the demonstration, reflecting the different types of available ground truth. In the real-data evaluation, the method produces depth-indexed anomaly scores and binary flags, which are aggregated into persistent responses and compared with fracture intervals catalogued independently from optical televiewer logs. Because the televiewer catalogue contains only confirmed open or soft-infilled fractures and does not guarantee completeness, the comparison is interpreted in terms of overlap and consistency rather than as a formal classification score.

In the synthetic evaluation, the full geometric output of the workflow is compared with known ground truth planes. A recovered plane is considered matched when its orientation deviates by less than five degrees from a planted plane and when its inlier set overlaps at least one true borehole intersection within a defined depth tolerance. For matched planes, angular error, recall, precision and F1 score are computed, together with geometric error measures based on point-to-plane distances. Each benchmark configuration is repeated across multiple independent runs, and all results are reproducible because each synthetic dataset records its generation parameters, including the real templates used for background signals and the signatures used for fracture injection.

## 4.2 Part One: Detection on Real Boreholes

The first part of the demonstration evaluates the one-dimensional detection capability of the workflow using a curated subset of boreholes from the Bjørnegårdtunnelen dataset. These boreholes are shown in their geological context in Figure 5 and include four positive cases with catalogued fracture intervals and two negative checks without such intervals.

The anomaly detection results are visualised in Figure 5 for the four positive boreholes. The figure present the consensus anomaly score along each borehole depth, together with the catalogued televiewer intervals and the persistent responses obtained after aggregation. The interpretation of these results depends not on the total number of flagged samples, which is relatively constant across boreholes by construction, but on the alignment between sustained anomaly responses and independently catalogued fracture intervals.

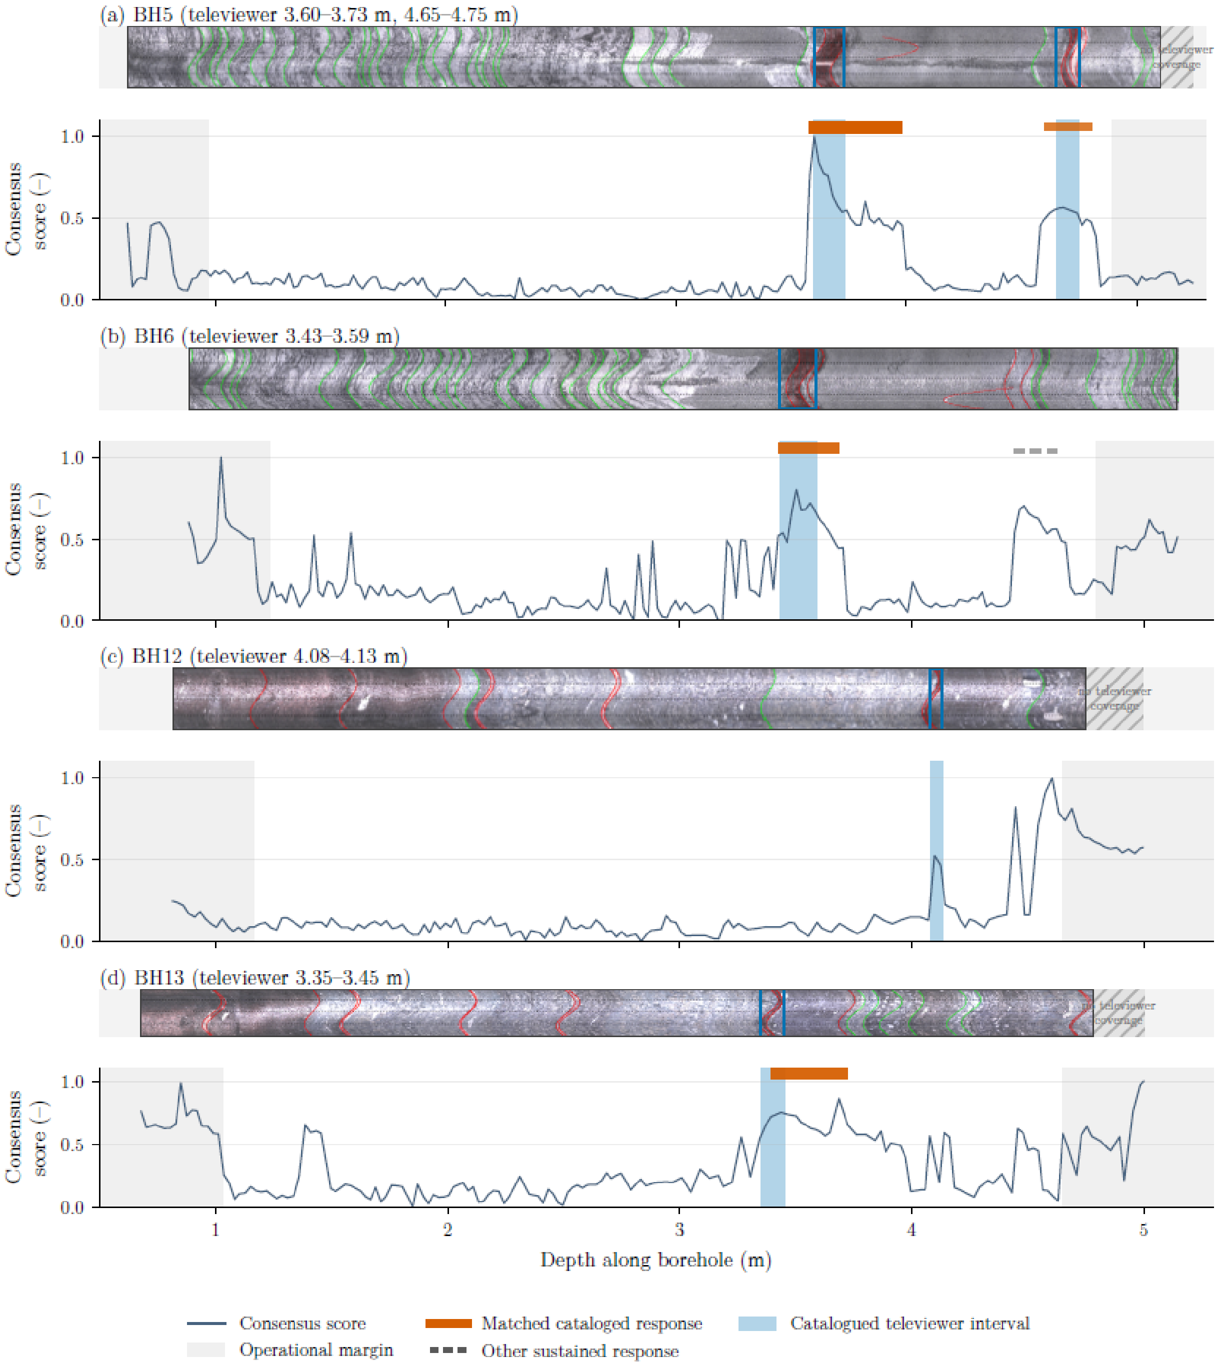


**Figure 5. Consensus anomaly scores for the four real boreholes with cataloged televiewer intervals (blue bands)**

In borehole BH5, shown in Figure 5a, two persistent anomaly responses coincide directly with the two catalogued televiewer intervals at depths of 3.60–3.73 m and 4.65–4.75 m. These responses correspond to the strongest peaks in the consensus score and demonstrate the ability of the unsupervised workflow to identify fracture-induced disturbances without labelled data. The detailed behaviour of the drilling channels and derived features in this borehole is examined in Figure 6, where penetration rate, rotation pressure, the composite fracturing index and the consensus score all show coincident disturbances at the catalogued depths. This illustrates that the detection signal arises from a multi-channel response rather than from a single parameter, and that the feature engineering stage successfully amplifies these coupled deviations.

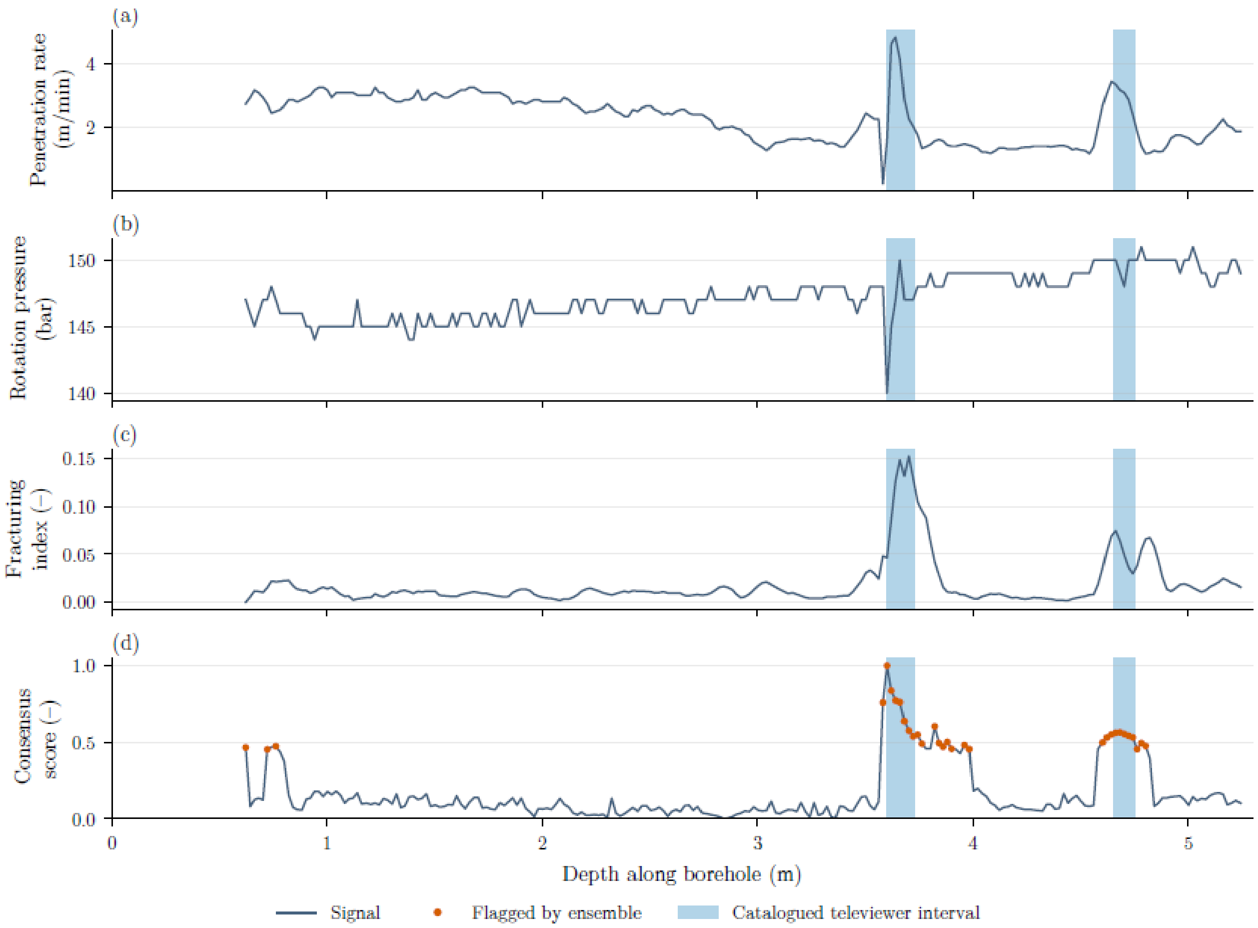


**Figure 6. Borehole BH5 in detail: (a) penetration rate, (b) rotation pressure, (c) the composite fracturing index, and (d) the consensus anomaly score with the flagged samples. The blue bands are the two cataloged televiewer intervals (3.60m to 3.73m and 4.65m to 4.75m).**

In borehole BH6, presented in Figure 5b, the catalogued interval at 3.43–3.59 m is matched by a persistent anomaly response extending from 3.42 to 3.68 m. A second sustained response appears at approximately 4.44–4.62 m but does not correspond to a catalogued interval and is therefore treated as an unmatched response. The internal structure of the feature representation for this borehole is shown in Figure 10, where the full fifty-six-feature matrix reveals a coherent depth interval in which multiple feature families deviate simultaneously. This demonstrates that the detection is driven by correlated behaviour across the feature space, rather than by isolated variations in individual channels, and explains why the ensemble detector identifies the interval even when individual signals appear weak.

In borehole BH13, shown in Figure 5d, the catalogued interval at 3.35–3.45 m is matched by a persistent response between 3.39 and 3.72 m. Together with BH5 and BH6, this confirms that the method consistently detects discontinuity signatures where the fracture aperture is sufficiently large relative to the sampling resolution.

The weakest positive case is borehole BH12, shown in Figure 5c. The catalogued interval at 4.08–4.13 m produces a clear local maximum in the consensus score, but this response does not reach the persistence required to form an aggregated event. The interval is the narrowest among the positive boreholes, spanning approximately 5 cm, which corresponds to only a few samples at the recorded spacing. This result anticipates the resolution boundary that is later quantified in the synthetic benchmark, demonstrating that very narrow discontinuities may produce detectable signal perturbations but fail to satisfy the criteria required for robust event formation.

The behaviour of the method in the absence of known fractures is illustrated by the two negative-check boreholes. Borehole BH9 produces no sustained anomaly responses outside the operational margins, while BH7 produces only a weak unmatched response. In both cases, clusters of flags near the collar and near the hole end are associated with operational effects rather than geological structure. These observations confirm that the anomaly detector responds broadly to localised disturbances, but that meaningful geological interpretation depends on persistence, depth location, and, in the full workflow, cross-hole consistency.

## 4.3 Part Two: Orientation Recovery on the Synthetic Benchmark

The second part of the demonstration evaluates the full workflow, including three-dimensional discontinuity reconstruction. Because real boreholes do not provide dense geometric ground truth, the evaluation is performed on synthetic datasets constructed from real data and arranged according to a realistic blast pattern shown in Figure 7.

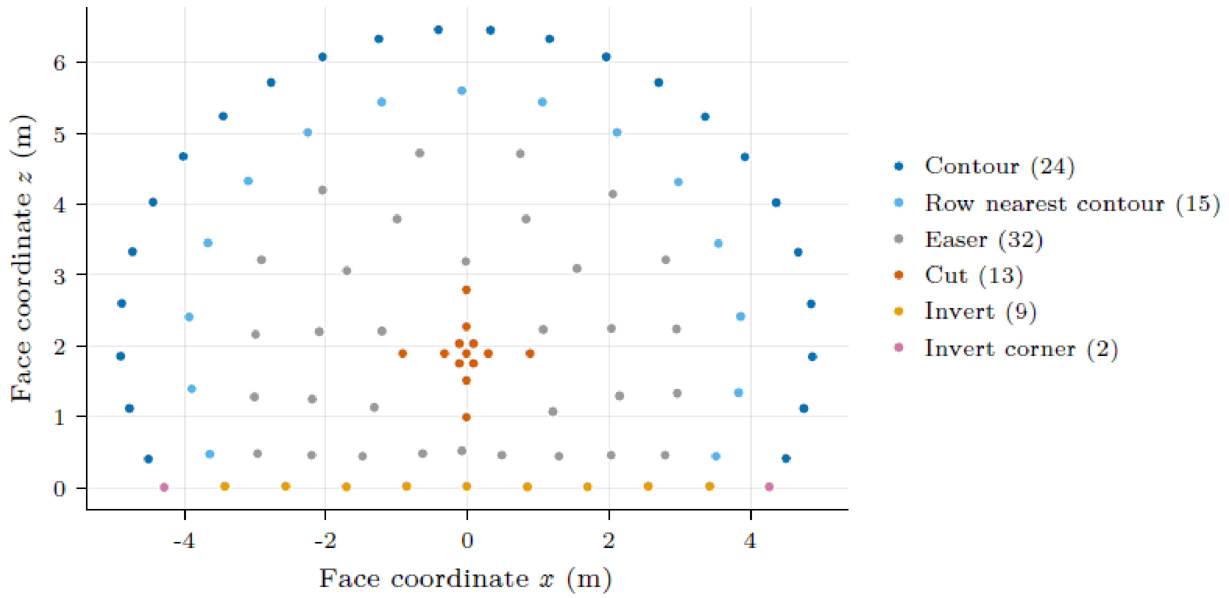


**Figure 7. The 59m2 blast pattern with 45mm hole.**

Synthetic drilling traces inherit their background trends and noise characteristics from real boreholes, while discontinuity planes of known orientation are inserted explicitly. An example of this construction is shown in Figure 8, where a synthetic borehole reproduces both the baseline behaviour and the localised disturbance introduced by a planted plane intersection. This ensures that the benchmark preserves the complexity of real signals while enabling exact measurement of recovery performance.

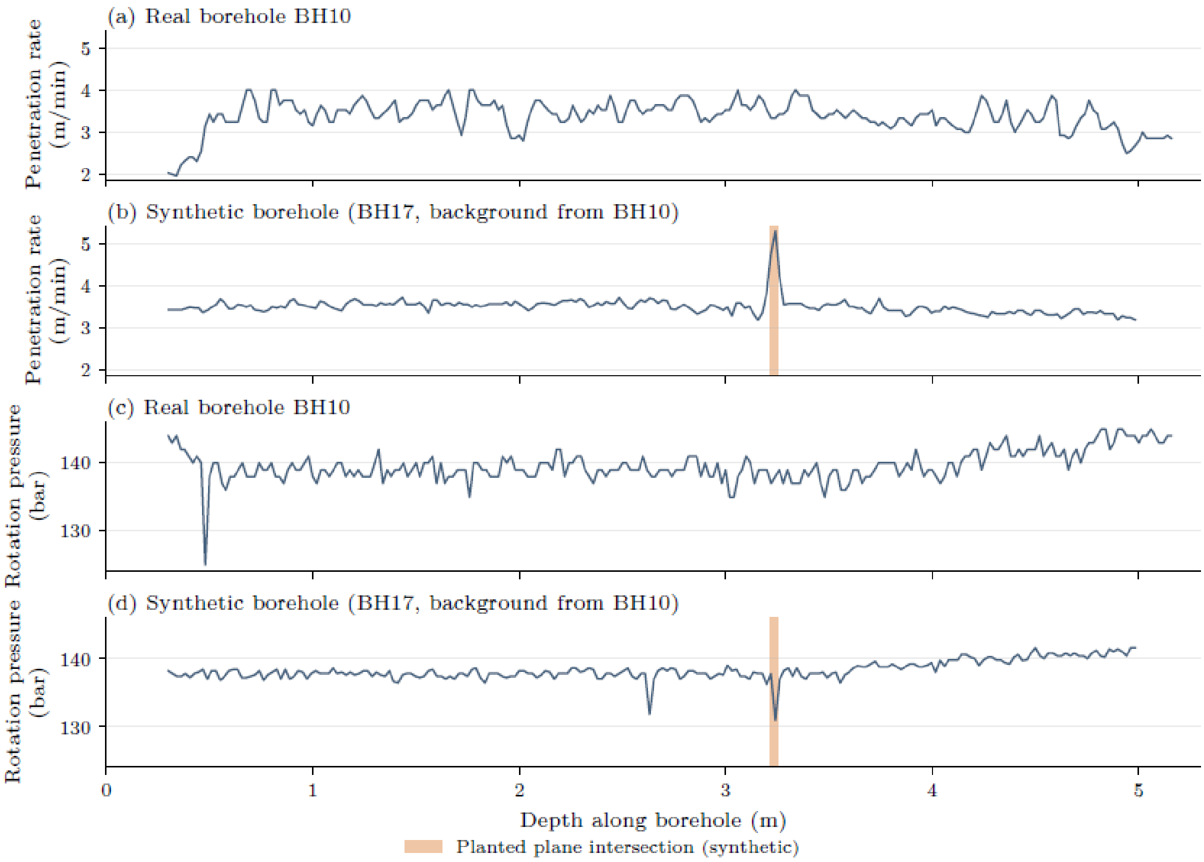


**Figure 8. Real and synthetic drilling traces, each panel spanning the full trace. (a, c) Penetration rate and rotation pressure of the real borehole BH10. (b, d) The same two channels for a synthetic borehole whose background is inherited from BH10, with the planted plane intersection shaded orange.**

In the baseline configuration, a single discontinuity plane is planted with an aperture exceeding the sampling spacing. Across multiple runs, the workflow recovers this plane consistently with high completeness and sub-degree orientation accuracy. A representative result is shown in Figure 9, where the recovered plane coincides with the planted plane in both three-dimensional view and orthographic projections. The inlier points trace the intersection of the plane with the borehole fan, while spurious detections are rejected as outliers. This demonstrates that the geometric stage successfully extracts coherent structures from a noisy and sparse point cloud.

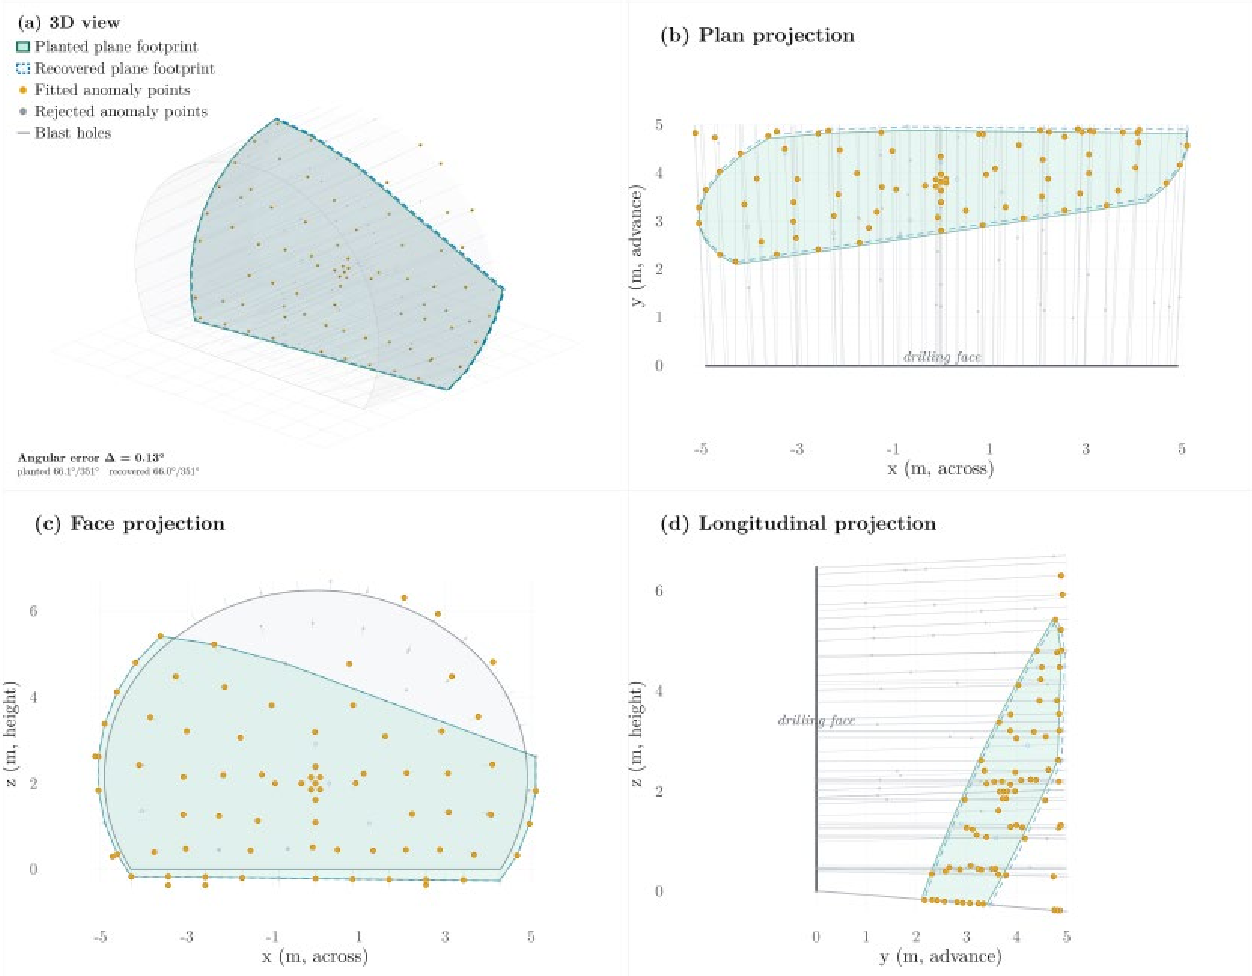


**Figure 9. Recovered versus planted discontinuity plane for one baseline run. (a) 3D view: the translucent gray volume is the rock mass to be excavated, the tunnel cross-section profile extruded along the advance direction, with the drilling face as its darker front wall at *y* = 0. The 95 borehole trajectories of the 59m2 pattern traverse the volume (contour holes overshoot it slightly: the look-out angle). Orange points are the detected fracture events assigned to the recovered plane, gray points are detected events rejected as outliers, the green footprint is the planted plane (strike 81◦, dip 66◦), and the dashed blue footprint is the recovered plane. (b) Plan, (c) face, and (d) longitudinal projections of the same volume; the footprints are orthographic projections of the plane–borehole intersections rather than single mid-section traces. The two planes coincide to within 0.13◦.**

## 4.4 Aperture-Resolution Relationship

The central experiment of the benchmark examines the relationship between fracture aperture and sampling resolution. The results of this comparison are summarised in Figure 10 and in the aggregate performance metrics shown in Figure 11.

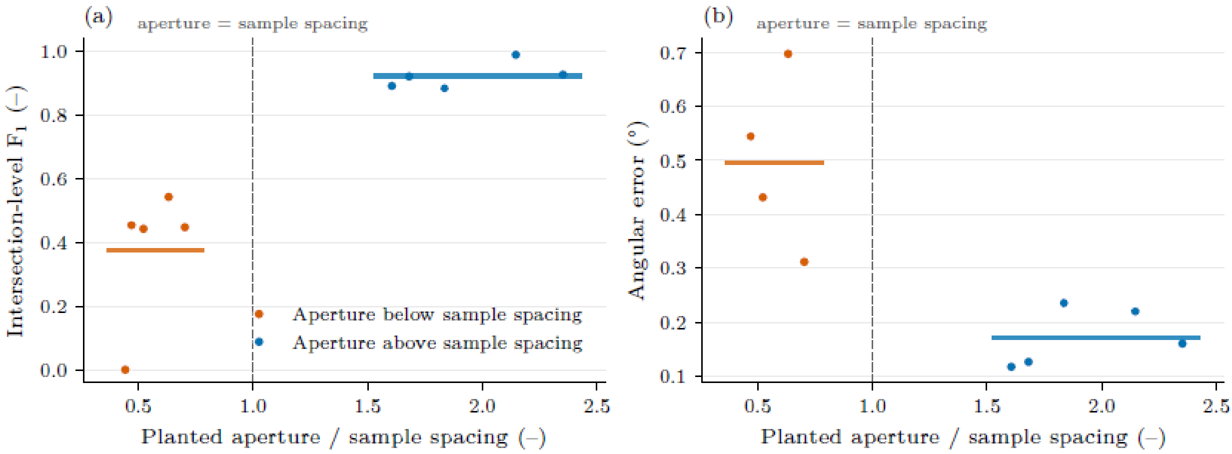


**Figure 10. The primary aperture comparison. (a) Intersection-level F1 and (b) angular error of the matched planes for five runs with the planted aperture below the 0.021m sample spacing (orange) and five with the aperture above it (blue), plotted against the ratio of the planted aperture to the sample spacing. Horizontal lines mark the means; the dashed line marks an aperture equal to the sample spacing.**

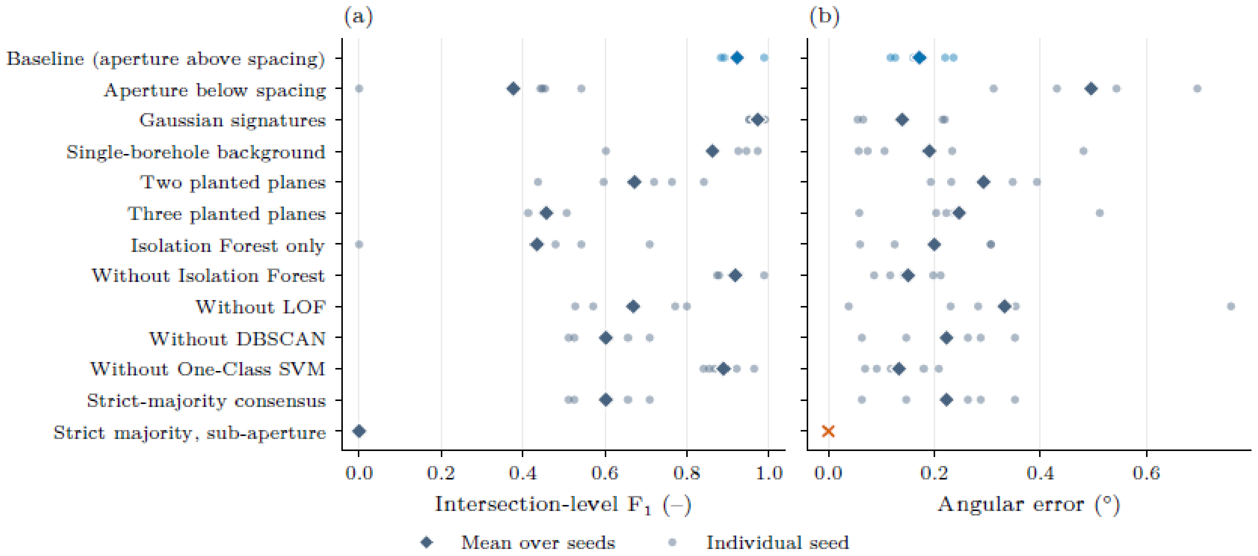


**Figure 11. Recovery across all benchmark configurations: (a) intersection-level F1 and (b) angular error of the matched planes. Small points are individual runs, diamonds their means; the baseline row is highlighted in blue. The sub-spacing runs at F1 = 0 are those whose planted plane was not recovered.**

When the fracture aperture exceeds the sampling spacing, detection and recovery are both reliable. The workflow consistently reconstructs the planted plane, and the recovered orientations are accurate to within a fraction of a degree. When the aperture is reduced below the sampling spacing, the behaviour changes fundamentally. The plane is still recovered in most runs, but its support is reduced, and in some cases the plane is not recovered at all. The F1 score decreases sharply, reflecting a loss of recall rather than a loss of precision.

The angular accuracy of recovered planes remains high even in the sub-resolution regime. This is because orientation estimation depends only on the subset of intersections that are successfully detected, provided that these are sufficiently distributed across the borehole fan. The primary failure mode is therefore a collapse in completeness rather than a degradation in geometric accuracy.

This behaviour directly explains the weak detection observed in BH12 in the real-data demonstration. The narrow televiewer interval produces a detectable signal spike but does not generate enough contiguous samples to form a persistent event. The synthetic benchmark generalises this observation, showing that the detection boundary corresponds to the limit at which fracture signatures span fewer than one or two samples in the MWD trace.

## 4.5 Sensitivity and Robustness Analyses

The robustness of the workflow is further evaluated through systematic variations of the benchmark configuration, the results of which are summarised in Figure 11.

Replacing real fracture signatures with idealised Gaussian profiles results in slightly higher detection performance, indicating that real signatures are inherently more complex and therefore more challenging to detect. This confirms that the use of real-derived signatures provides a conservative and realistic benchmarking framework.

The introduction of multiple discontinuity planes reduces completeness but preserves orientation accuracy, reflecting increased competition between intersecting structures within the same borehole fan. This demonstrates that the workflow remains geometrically stable even in more complex structural scenarios.

Comparison with a single-detector variant shows that the ensemble approach substantially improves consistency. While a single detector may recover a plane in some cases, its performance varies significantly across runs, and it may fail entirely under certain conditions. The ensemble reduces this variability by combining complementary notions of abnormality and stabilising the detection stage.

The contribution of individual detectors is further clarified through component-removal experiments, which show that no single detector is essential but that removing detectors reduces completeness and increases variability. This confirms that the strength of the method lies in the combined effect of multiple detection mechanisms rather than in any single component.

Finally, the sensitivity of the consensus threshold is evaluated. A stricter voting rule reduces the number of flagged samples and produces cleaner one-dimensional detection plots, but it significantly degrades three-dimensional recovery performance, particularly near the resolution limit. This demonstrates that the anomaly detection stage must be tuned for downstream geometric interpretation rather than for visual cleanliness in isolation.

## 4.6 Interpretation

The combined results demonstrate that the workflow successfully bridges the gap between one-dimensional signal disturbances and three-dimensional geological interpretation. The real-data demonstration shows that the method detects independently validated fracture intervals, while the synthetic benchmark confirms that these detections can be assembled into accurate discontinuity planes under realistic conditions.

The dominant limiting factor is the relationship between fracture aperture and sampling resolution. When discontinuities span multiple samples, they produce robust multi-channel signatures that survive feature extraction and detection. When they fall below this scale, detection becomes incomplete, although the geometry of the detected subset remains accurate.

The evaluation also highlights the importance of handling operational artefacts explicitly, as these can produce coherent patterns in both one-dimensional detection and three-dimensional reconstruction. Nonetheless, these artefacts are identifiable and can be mitigated through margin exclusion and geometric filtering.

Overall, the demonstration confirms that the methodology provides a practical and robust framework for deriving three-dimensional discontinuity information from routine MWD data, with performance bounded by physical data resolution rather than algorithmic limitations.

## 5 Summary and Conclusion

This annex has presented a complete, end-to-end methodology for detecting geological discontinuities and recovering their three-dimensional orientations directly from Measurement-While-Drilling (MWD) data. The workflow transforms raw drilling telemetry into spatially resolved structural information through a sequence of explicit processing stages, combining physics-informed pre-processing, multi-feature representation, ensemble anomaly detection, aggregation of localised disturbances, and iterative geometric reconstruction using Random Sample Consensus.

The evaluation framework was deliberately structured in two complementary parts to match the inherent constraints of the available data. The real borehole dataset provides genuine, field-measured drilling signals together with independently catalogued fracture intervals derived from optical televiewer logs. This enables direct validation of the detection stage without recourse to artificial assumptions. At the same time, real boreholes do not supply dense, surveyed three-dimensional ground truth for discontinuity geometry. This limitation is addressed through a real-derived synthetic benchmark, in which signals inherit the statistical structure and residual variability of the real data while discontinuity planes are inserted with known geometry. The benchmark therefore preserves the realism of the field signals while making the full geometric recovery problem measurable.

The first part of the demonstration establishes that the unsupervised detection stages of the workflow respond to genuine geological disturbances. Persistent anomaly responses coincide with independently catalogued fracture intervals in the majority of positive boreholes, specifically BH5, BH6 and BH13, while boreholes without catalogued open fractures remain comparatively quiet outside operational margins. The detailed inspection of individual boreholes shows that the detection signal arises from coordinated perturbations across multiple drilling parameters and derived features, rather than from isolated excursions in a single channel. This confirms the underlying assumption that discontinuities are expressed as localised multi-channel disturbances, and demonstrates the value of the feature-engineered representation and consensus detection approach.

At the same time, the real-data results reveal the limits of purely one-dimensional detection. The case of BH12 shows that narrow discontinuities may produce a detectable anomaly response without forming a persistent event when the number of affected samples is too small. In addition, residual anomaly responses appear systematically near collaring and breakthrough zones, reflecting the influence of operational processes. These observations motivate the subsequent aggregation criteria and the geometric filtering stage, and they establish that the anomaly detector must be interpreted as an intermediate stage rather than as a final classification in itself.

The second part of the demonstration shows that when these detections are carried through the full workflow and combined across multiple boreholes, the resulting discontinuity planes can be reconstructed with high geometric accuracy. Under baseline conditions in the synthetic benchmark, in which the fracture aperture exceeds the sampling spacing, the workflow recovers the planted planes consistently across multiple independent runs. The recovered planes achieve near-perfect precision, meaning that detected inliers correspond almost exclusively to true plane intersections, and sub-degree angular accuracy, demonstrating that the geometric stage effectively extracts coherent structures from a sparse and noisy point cloud. The remaining discrepancy between unity precision and lower recall is attributable to missed intersections rather than contamination of the recovered planes, indicating that the limiting factor lies in the detection and aggregation stages rather than in the geometric reconstruction.

The central finding of the benchmark is that the performance of the method is controlled by a physically interpretable boundary defined by the ratio between fracture aperture and sampling resolution. When the aperture exceeds the sampling spacing, discontinuity signatures span multiple depth samples and produce stable, multi-channel disturbances that survive feature extraction, anomaly detection and aggregation. When the aperture falls below the sampling spacing, the situation changes fundamentally. Individual intersections are often expressed as single-sample excursions that are indistinguishable from noise under the consensus detection and event persistence criteria. As a result, the number of detected intersections collapses, and the completeness of the recovered planes degrades sharply. In the most extreme cases, the plane may not be recovered at all because insufficient support remains after aggregation.

Crucially, this transition affects completeness rather than geometric accuracy. Even in the sub-resolution regime, when a subset of intersections survives, the fitted plane retains sub-degree orientation accuracy because RANSAC requires only a sufficiently distributed subset of points across the borehole fan to constrain the plane orientation. The failure mode is therefore not a degradation of the geometric model but a loss of support for that model. This distinction is essential for practical interpretation: the presence of a recovered plane can be trusted geometrically, but its absence does not necessarily imply the absence of a discontinuity when the aperture is near or below the sampling limit.

The evaluation further demonstrates that the ensemble structure of the anomaly detection stage is not an implementation detail but a key contributor to the robustness of the method. Single-detector configurations are able to recover planes in favourable cases but exhibit significant variability between runs and may fail entirely under certain conditions. By combining four complementary detectors through a consensus mechanism, the workflow reduces sensitivity to the failure modes of any individual method and produces stable detection behaviour without per-site tuning. The detector-removal experiments show that while no single detector is indispensable, each contributes to coverage and stability, especially near the detection boundary.

The sensitivity analysis of the consensus threshold reinforces the same point from a different perspective. A stricter voting rule produces visually cleaner one-dimensional anomaly traces but significantly reduces the number of flagged samples available to the geometric stage. This leads to a marked degradation in three-dimensional recovery performance, particularly in the vicinity of the resolution boundary and in the sub-aperture regime, where the stricter rule prevents any plane from being recovered. This demonstrates that the appropriate operating point of the anomaly detection stage is defined not by the clarity of its standalone output, but by the ability of the downstream aggregation and plane-fitting stages to recover coherent structures from its input.

The use of a real-derived synthetic benchmark is itself an important element of the contribution. By constructing synthetic boreholes that inherit baseline trends, noise characteristics and fracture signatures from real data, the benchmark ensures that the evaluation reflects realistic signal conditions while enabling reproducible ground truth. The comparison between real extracted fracture signatures and idealised Gaussian injections shows that simplified benchmarks may overestimate performance. The real-derived approach therefore provides a more conservative and reliable basis for assessing the full workflow.

In operational terms, the results demonstrate that routine MWD data can be used as a volumetric sensing tool for the rock mass ahead of excavation. Unlike conventional mapping, which observes the excavated surface retrospectively, the present methodology infers the presence, location and orientation of discontinuities within the volume intersected by the borehole fan before blasting. The recovered planes provide directly usable information for engineering decisions, including support design, blast optimisation and hazard identification. At the same time, the method identifies its own limits through the aperture-resolution relationship, making clear under which conditions discontinuities can be expected to be detected reliably.

The principal limitation of the method is therefore not algorithmic but physical. The sampling interval of the drilling system determines the minimum aperture that can be resolved, and discontinuities narrower than this interval will remain only partially detectable. Additional limitations arise from the absence of dense field ground truth for plane geometry, the simplifying assumptions of planarity in the synthetic benchmark, and the presence of operational artefacts that can form coherent but non-geological structures across the borehole fan. These limitations define the scope of applicability rather than invalidating the approach.

In conclusion, the methodology presented in this annex establishes a practical, fully unsupervised framework for converting drilling telemetry into three-dimensional geological information. It demonstrates that discontinuity detection and orientation recovery can be achieved directly from MWD data, without labelled datasets and without additional instrumentation. The workflow performs robustly under realistic conditions, achieves sub-degree orientation accuracy where sufficient signal support exists, and degrades in a physically interpretable manner at the resolution boundary. As such, it represents a significant step toward integrating real-time geological awareness into drill-and-blast excavation workflows.

## 6 Future work

Future work will focus on extending the methodology toward more realistic geological and operational conditions, as well as enhancing its practical deployment potential. On the validation side, the most immediate priority is the application of the workflow to datasets with fully surveyed borehole geometry, enabling direct comparison between recovered discontinuity planes and independently mapped structures. The synthetic benchmark can also be extended beyond ideal planar geometries to include partially persistent, curved or discontinuous features, as well as to explicitly incorporate borehole deviation and survey uncertainty. From a methodological perspective, further developments will explore the estimation of additional structural attributes, such as discontinuity aperture and persistence, and the incorporation of uncertainty quantification in the recovered orientations. Increasing the depth sampling resolution, which is a controllable acquisition parameter, remains a key avenue for improving detectability at the fine scale identified by the aperture–resolution boundary. Finally, the underlying workflow can be generalised to related applications, including the detection of rock-mass transitions, water-bearing zones and ore boundaries, providing a broader framework for transforming drilling telemetry into actionable subsurface information.

[This page intentionally left blank]