SSIM 2025 - JJ Potter & J Wesseloo (eds)
© 2025 Australian Centre for Geomechanics, Perth, ISSN 3083-399X
doi:10.36487/ACG_repo/2535_48

Automated televiewer data interpretation and validation: an applied case study at Kennecott copper

John Danielson [a,*], Jennifer Ellis [b], Connie Schultz [c], Eric Wilton [a], Kirk Schafer [b], Ian Stilwell [a]

[a] *BGC Engineering, Canada*

[b] *Rio Tinto, USA*

[c] *Rio Tinto, Australia*

Abstract

Manual interpretation of borehole televiewer data for geological and geotechnical site characterisation is essential, but hampered by considerable time and cost investment; variation in interpretation (inconsistency); and critical orientation biases – including a "visual salience bias" against low-angle features often relevant to slope stability. This paper introduces Structura, a tool employing deep learning to automate the interpretation of acoustic and optical televiewer logs, identifying structures, breakouts and drilling-induced tensile fractures, estimating roughness, and assessing borehole image quality. We present results from a collaborative pilot study at Rio Tinto Kennecott Copper, where Structura was compared against manual interpretation, refined based on operational needs, and validated using blind test data and physical core comparisons. Validation showed good agreement on structural orientations with potential bias reduction, a positive correlation for automated roughness requiring resolution-based adjustment, and reliable identification of breakout features for stress analysis. Structura demonstrates the potential of deep learning-based tools to significantly improve the efficiency, consistency, and objectivity of televiewer data analysis; complementing traditional oriented data collection workflows.

Keywords: *deep learning, data processing, automation, structural characterisation, televiewer data, machine learning, borehole image processing*

1 Introduction

The stability of hard rock slopes in open pit mines is fundamentally governed by the characteristics and spatial distribution of geological structures. Accurate mapping, classification, and characterisation of structures such as joints, faults, and intact fabric (e.g. bedding, foliation) are paramount for reliable geotechnical design and risk assessment (Hoek & Bray 1981; Read & Stacey 2009). Borehole televiewer surveys, both acoustic (ATV) and optical (OTV), provide high-resolution, oriented images of the borehole wall, offering invaluable data for identifying and characterising these structures – especially in areas where core quality is poor or recovery is limited. Recent image segmentation approaches (e.g. Li et al. 2024) support the feasibility of deep learning-based feature identification from borehole televiewer images.

At large mining operations like Rio Tinto Kennecott Copper (RTKC), televiewer surveys are routinely conducted to gather structural data critical for geological modelling and slope stability analysis. At RTKC, the collected structural data from televiewer surveys are particularly important because poor core quality prevents oriented core data collection. The traditional method of interpreting televiewer logs relies heavily on manual identification and picking of features by geoscientists. This manual process, while essential, faces several significant challenges:

[*] Corresponding author. Email address: jdanielson@bgcengineering.ca

- Time and cost: manually reviewing and interpreting potentially thousands of metres of televiewer data is labour intensive and slow. This process consumes considerable time from experienced geoscientists and engineers, whose expertise could arguably be better applied to higher-level tasks such as validating interpretations, integrating data into models, and performing engineering analysis.
- Inconsistency: interpretation can vary significantly between different geoscientists due to differences in experience, training, focus, and fatigue, leading to inconsistencies in the collected data.
- Bias: manual televiewer interpretation suffers from inherent biases that depend on the angle at which the structure intersects the borehole. While the tendency to under-sample structures parallel to the borehole is well-known, often manifesting as a “blind zone” on the stereonet, a less recognised secondary bias exists for structures perpendicular to the borehole. This bias, which we term “visual salience bias”, relates to how readily features are identified based on their apparent dip in the unwrapped televiewer image. Flat structures are more difficult to see. This can result in an under-representation of structures near-perpendicular to the borehole (low apparent dips) and a corresponding over-representation or concentration of structures with intermediate apparent dips (e.g. 30–75°).

As shown in Figure 1, plotting aggregated apparent orientation data from dozens of differently oriented boreholes reveal visual salience bias as a secondary blind zone as a data gap near the centre of the stereonet. Applying Terzaghi weighting, designed to correct the edge bias, unfortunately intensifies this central bias against flat structures. This phenomenon, also observed in individual borehole data from different sites (Figure 2), can result in 'donut-shaped' data distributions. The orientation bias severity appears to differ at different mine sites and may depend upon not only on the logger and logging guidance, but also upon the data resolution and the borehole diameter, which affect how ‘flat’ structures appear visually.

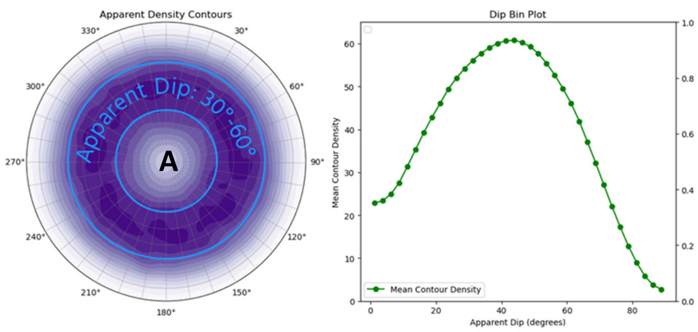

(a)

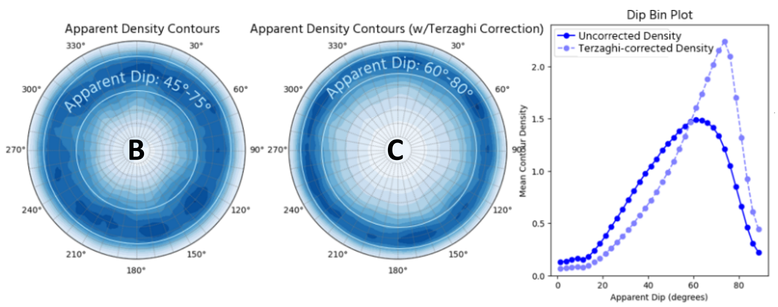

(b)

Figure 1 Aggregated apparent orientation data from (a) Mine Site 1 and (b) Rio Tinto Kennecott Copper (RTKC). Note the difference in positive bias between the two sites, concentrating at apparent dips of 30–60° in stereonet A and 45–75° in stereonet B. At RTKC, applying Terzaghi correction (stereonet C) increases the severity of the data bias, and expands the ‘secondary’ blind zone in the centre of the stereonet

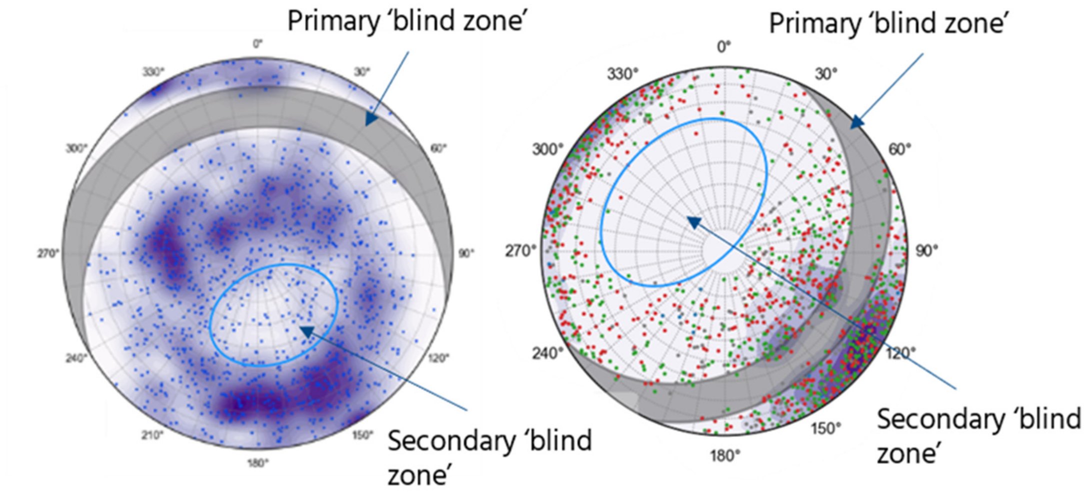

Figure 2 Manually logged acoustic televiewer data ('true' orientations) from individual boreholes from different mine sites

The implications are potentially significant: geotechnical boreholes are often oriented near-perpendicular to planned pit slopes specifically to intersect structures dipping at low angles relative to the borehole (i.e. sub-parallel to the pit slope being investigated). The bias against detecting these (apparent) flat-lying features means manual methods can systematically under-represent the structures most relevant to the slope stability assessment for which the drilling was conducted.

To address these limitations and improve the efficiency, consistency, and quality of structural data derived from televiewer logs, BGC Engineering developed Structura, an artificial intelligence (AI)-powered software tool (Danielson et al. 2023; Danielson et al. 2024). Structura employs deep learning algorithms to automate the detection, classification, and characterisation of geological features in televiewer data.

This paper introduces the Structura tool, details its capabilities, and presents the results of a collaborative pilot study conducted with RTKC. The study evaluated Structura's performance against traditional manual logging methods, explored its application on a large dataset, and led to refinements based on real-world mining needs. We discuss validation techniques, compare the outputs of manual and automated methods, and provide recommendations for integrating automated tools like Structura into the geotechnical workflow.

2 Structura: artificial intelligence powered televiewer interpretation

Structura is a software tool designed to automate the processing of ATV and OTV borehole logs using deep learning models. The core technology involves segmentation and classification algorithms trained on extensive datasets to rapidly and consistently identify, classify, and characterise key geological and geotechnical features.

2.1 Underlying technology and capabilities

Structura leverages deep learning techniques including convolutional neural networks for image segmentation and classification. In addition to the televiewer data, the workflow requires definition of user-specified parameters: declination, apparent north reference (typically 'high side' or 'magnetic north'), and hole diameter.

1. Preprocessing: televiewer data (typically in WellCAD format) are imported, normalised to standard ranges, and formatted appropriately for inference using the deep learning frameworks.
2. Segmentation: segmentation models are used to identify pixels belonging to geological structures, and breakouts and drilling-induced tensile fractures (DITFs). These models were trained on large and diverse datasets of manually labelled televiewer logs from various geological settings and

drilling projects. The output is a pixel-level 'mask' highlighting potential features. The current training dataset comprises many thousands of manually labelled structures from multiple deposit types; where a new geological setting differs substantially from those already represented, incorporating labelled selections from one to three representative holes can help maintain predictive accuracy.

3. Parameter estimation: for identified structural features, algorithms fit sinusoidal curves to the segmented pixels. Parameters like apparent dip, dip direction, aperture, and roughness are calculated from these fits and from the pixels belonging to each feature. Candidate picks with a model confidence below 0.2 or a roughness above a certain threshold are discarded to avoid over-picking. For breakouts/DITFs, parameters like length, average width, and orientation are determined from the segmented regions.
4. Classification: a separate classification model assigns a feature type to each identified structure feature based on its image characteristics and geometric properties. Classes are described in Table 1 (Section 2.2).
5. Quality assessment: another model analyses image characteristics like contrast, focus, and presence of artifacts to assign a quality score along the borehole depth (Figure 3). Quality can be affected by different causes, including drill-induced markings, poor rock mass qualities, survey resolution, centralisation, and coating (e.g. cuttings) on the borehole wall.
6. Rotation from apparent to true: using the user-specified parameters, data are rotated from apparent to true (or mine) north, summarised, and plotted.

This multi-stage process allows Structura to manage the complexity of televiewer images and extract a comprehensive set of parameters automatically.

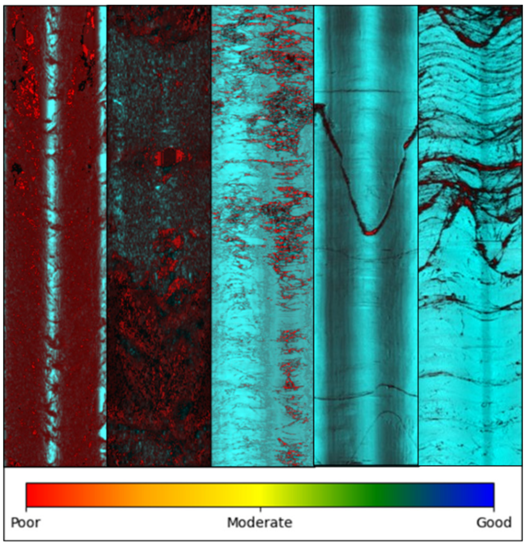

Figure 3 Survey image quality ranking examples

Structura automatically identifies and characterises:

- Geological structures: such as joints, faults, bedding, foliation, and veins. For each identified structure, Structura interprets the orientation (dip and dip direction), estimated aperture or thickness, a calculated roughness based on circumferential profile (Figure 4), and a classification of the feature type (e.g. joint, bedding). Examples of automatically identified structures on ATV and OTV logs are shown in Figure 5 and Figure 6, respectively.
- Borehole breakouts and DITFs: these features (Figure 7) provide information about the in situ stress field. Structura identifies these features and reports their length along the borehole, average width, and azimuth.

- Borehole image quality/survey quality: the tool assesses the quality of the televiewer log image itself, categorising intervals on a scale of 0 to 1, where 0 is poor and unusable and 1 is good quality. This allows users to understand the reliability of the input data and filter interpretations derived from low-quality intervals, an assessment often missing or inconsistently applied in manual logging workflows.

Multiple capabilities were added during the collaboration with RTKC (Section 3).

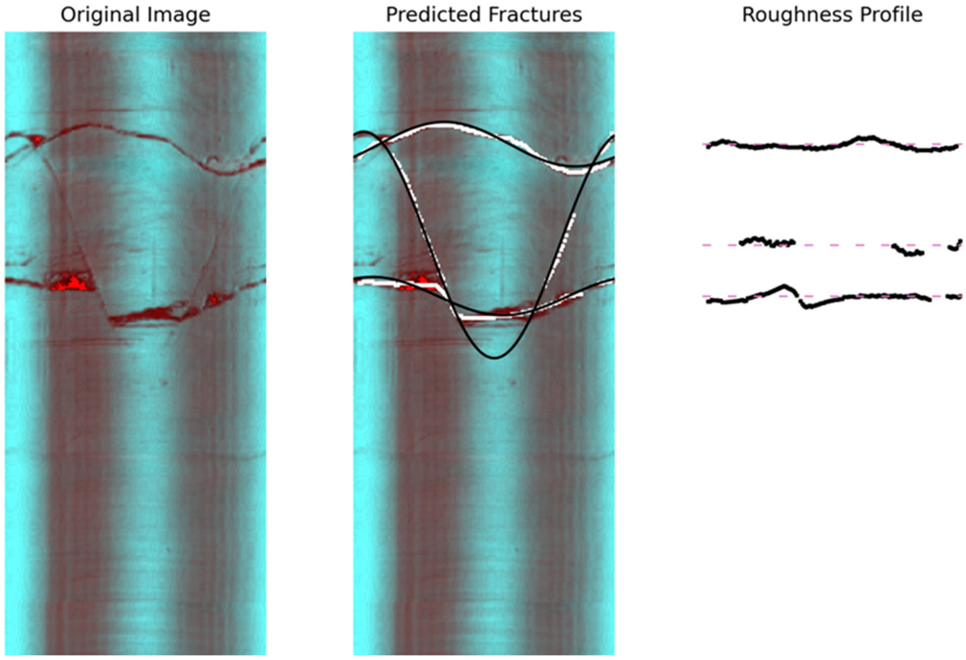

Figure 4 Example roughness profiles estimated using Structura. The roughness profiles shown are circumferential profiles projected onto the best fit plane of the structures

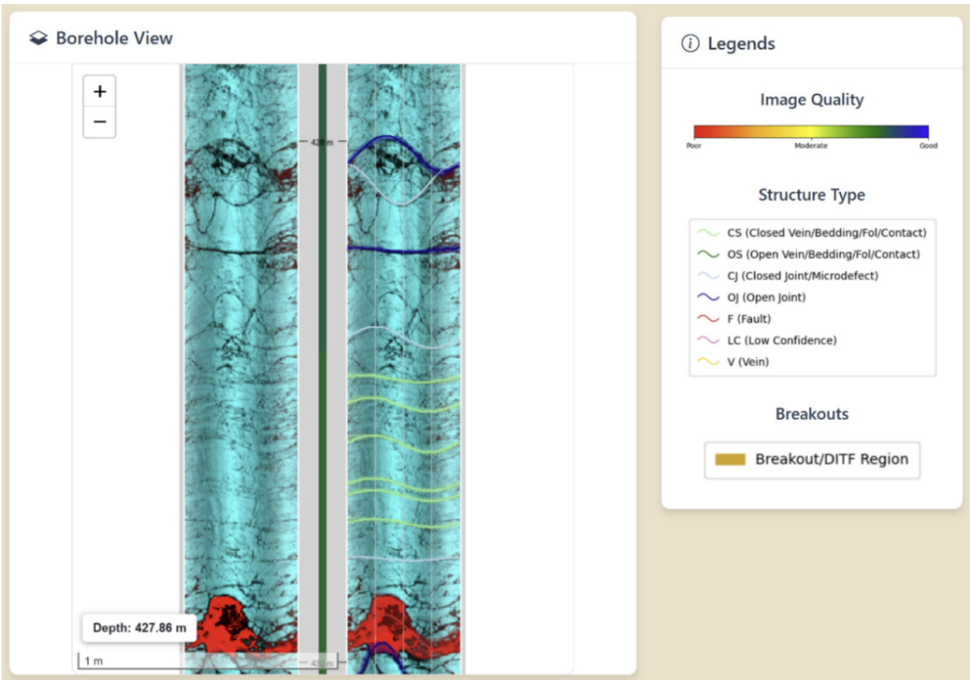

(a) (b)

Figure 5 Example (a) acoustic televiewer image with (b) overlain automatically identified structural picks

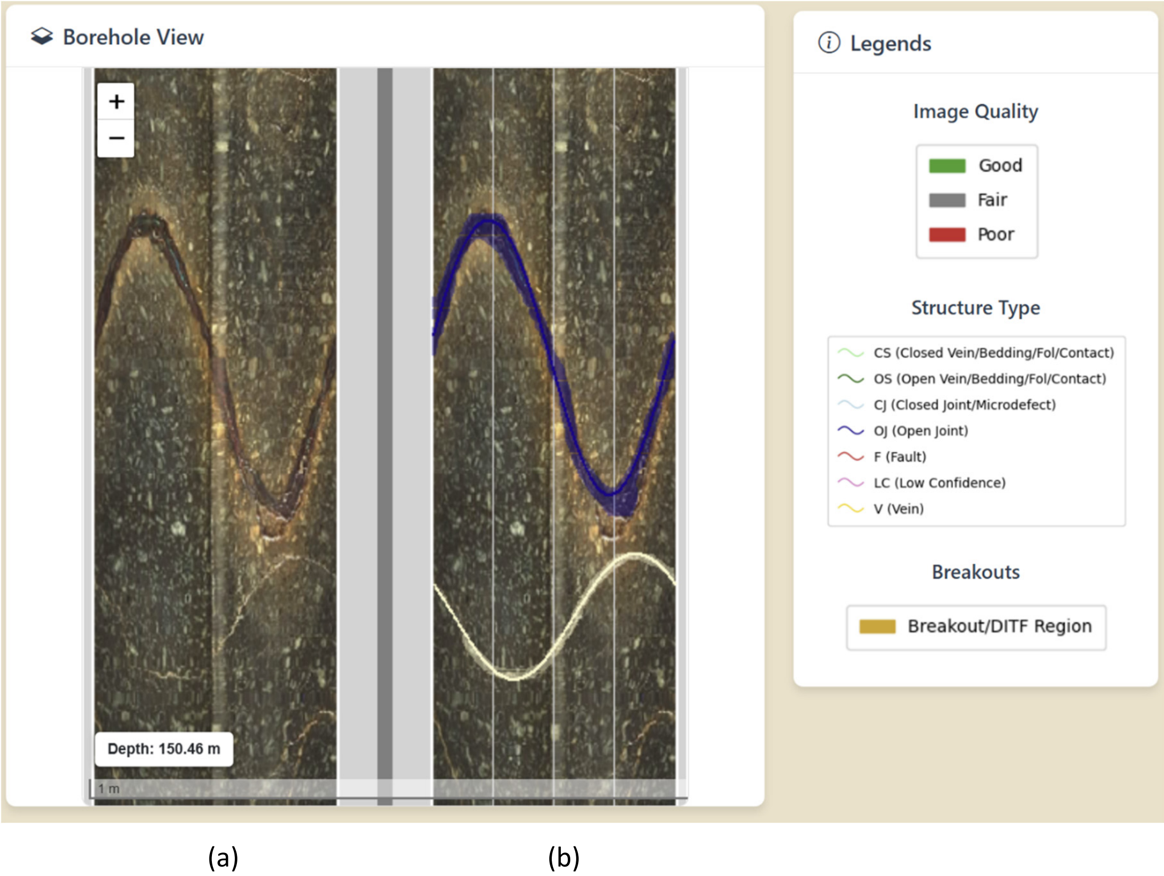

Figure 6 Example (a) optical televiewer image with (b) overlain automatically identified structural picks

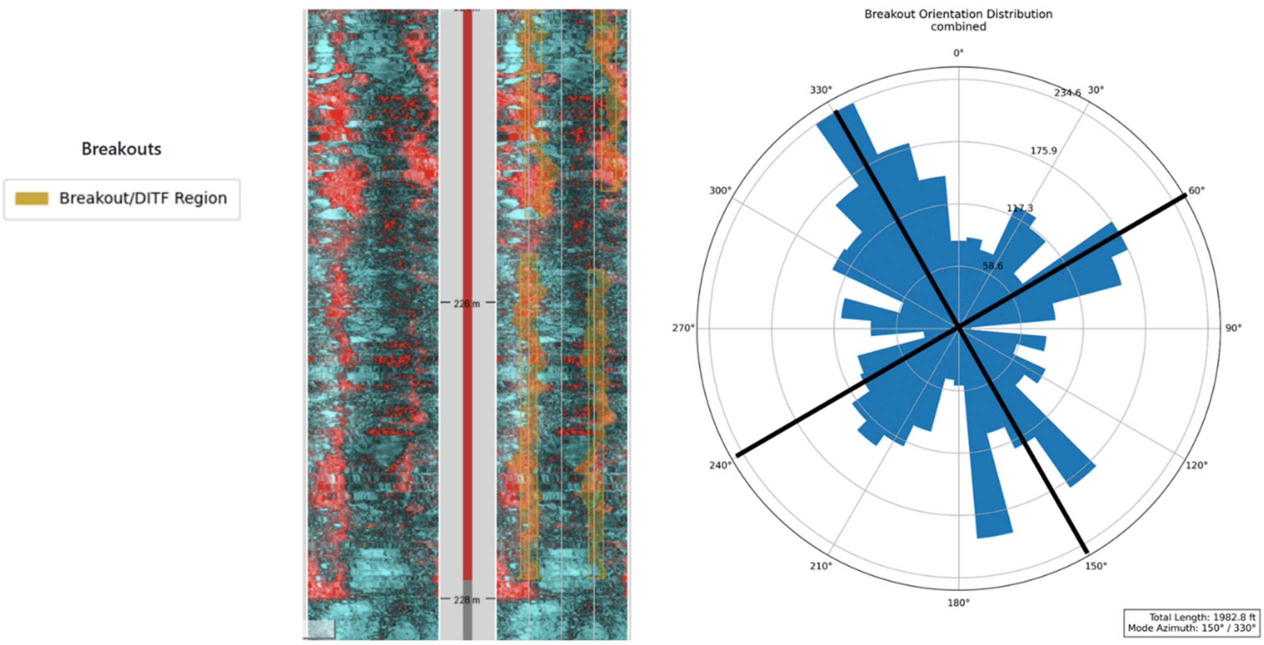

Figure 7 Borehole breakouts identified with Structura at a major mine site

2.2 Outputs and integration

Data are summarised into standard outputs per hole which include: a stereonet, a breakout/DITF rose plot, a strip log with overlain features, and tables summarising the identified structures, breakouts/DITFs, and image quality. QA can then be performed by bringing the data back into WellCAD or reviewing the generated strip log. The processing is typically performed via a cloud-based platform and usually requires less than five minutes per borehole.

Table 1 Default classes

Class name	Description	ATV example	OTV example
Closed bedding/foliation (**CB**)	Bedding/foliation with no or minimal separation. Can include intact veins or contacts if they form a repeating fabric.	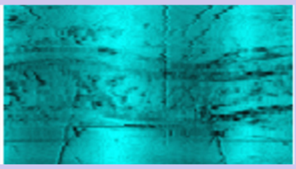	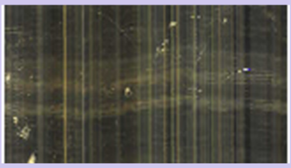
Open bedding/foliation (**OB**)	Bedding/foliation with separation. Can include open veins or contacts if they form a repeating fabric.	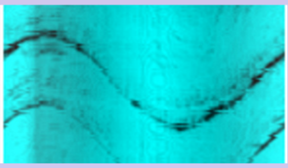	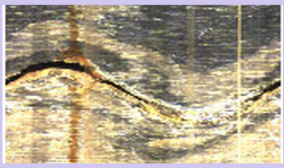
Fault (**F**)	Wider open discontinuities, typically with signs of displacement (brecciation) or gouge infill.	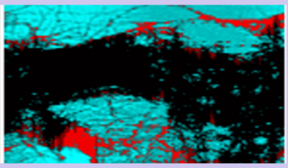	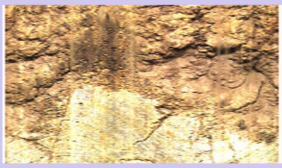
Closed joint (**CJ**)	Microdefects or tight joints.		
Open joint (**OJ**)	Joints with visible aperture.		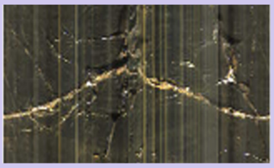
Vein (**V**)	Thin structures, often with differing colour and character compared to host rock. Only classified within OTV data.		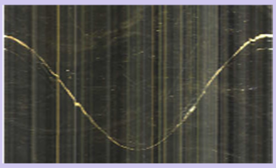

3 Rio Tinto Kennecott Copper pilot study

To evaluate and refine Structura's capabilities within a complex, real-world geological environment, a collaborative pilot study was undertaken with RTKC. As noted previously, interpretation of televiewer data is particularly crucial at Kennecott, where variable rock quality prevents the collection of oriented drillcore data. The study focused on the ATV data routinely collected at the site, which provided a robust dataset for testing and validation. The study was conducted in three main phases – an initial comparison, tool refinement, and validation.

3.1 Initial comparison and capability assessment

The initial phase involved processing 39 historical boreholes previously logged manually by experienced RTKC geoscientists. These same boreholes were independently processed using Structura. The primary goals were to compare the automated outputs directly against the manual interpretations and identify Structura's strengths and limitations in the context of RTKC's geological setting and data types.

Through comparative analysis, using side-by-side comparison of individual logs and data aggregated in stereonets, several key observations emerged. As anticipated, manual interpretations exhibited variability between boreholes, attributable to the subjectivity of the task and differences among loggers. Structura's automated detection performance was found to be sensitive to input data resolution; thinner features were sometimes missed in RTKC's coarser resolution logs (which had vertical resolutions of 6 to 10 mm) relative to the finer (approximately 3 mm) resolution data the models were initially trained on. This initial phase identified necessary improvements to meet RTKC's operational requirements, including the addition of coarser resolution data to the training dataset and the development of automated structure classification and borehole breakout identification capabilities.

An internal RTKC analysis of the results indicated Structura reduced structure picking time from approximately 40 hours per hole to 5 minutes. Because manual depth matching, log preparation, and QA/QC were still necessary, overall time savings for a typical hole were estimated to be 67%. The analysis confirmed that Structura sometimes missed thinner features, correctly identifying approximately 80% of structures present. The analysis also indicated that some manual logs contained spurious/excess picks (over-identification), which can contribute to data bias.

3.2 Tool refinement

Based on the findings from the initial comparison and specific requirements from RTKC, Structura was customised and refined. Several key improvements were undertaken:

- Improved structure identification: the core segmentation model was refined, incorporating additional site-specific training data from RTKC to improve its performance and increase the quantity of picks on the coarser-resolution RTKC data.
- Structure classification: classification capabilities were added, enabling the tool to differentiate between various feature types, including the specific identification of intact (closed) and open bedding/foliation (Table 1). A post-processing step based on orientation clustering was implemented to assist in distinguishing bedding sets from other structures, reclassifying potential 'beds' outside the main sets as 'undifferentiated structures'. This addresses the challenge that beds, contacts, and veins can have a similar appearance in ATV data. Clustering is performed along a 30 m sliding window rather than the entire hole, so that real changes in bedding don't cause undesired reclassification.
- Borehole breakout identification: functionality was added to detect and characterise borehole breakouts and DITFs, reporting their orientation and length along the borehole (Figure 7). The tool recognises the signals indicative of these features, but does not currently distinguish between

stress-induced breakouts and DITFs. Compiled breakout orientations are summarised and plotted on rosettes after rotation to mine north.

- Roughness estimation: the algorithm for estimating discontinuity roughness (Figure 4) was improved and refined based on further analysis and validation.
- Borehole image quality assessment: the capability to automatically assess and report on the quality of the televiewer image data was integrated.

The following section describes how these refinements were validated against blind test data.

3.3 Validation and blind test results

To validate these improvements, Structura processed approximately 3,000 m (10,000 ft) of 'blind' test data provided by RTKC. This dataset consisted of 10 boreholes with no accompanying manual logs available to the BGC team during processing to maintain an unbiased assessment. Importantly, the blind test data were not used in the training of the refined models. The dataset included a mixture of borehole diameters, data resolutions (ranging from 2 mm to 10 mm vertically), and survey qualities.

Results from the blind test demonstrated the effectiveness of the refined tool. Processed structural data, when plotted on stereonets, generally revealed clear structural trends consistent with expected site geology, with identified bedding poles typically clustering appropriately. Structura successfully identified and characterised breakouts and DITFs. The test also confirmed that data resolution and image quality significantly influence the accuracy of automated interpretation, particularly for structure classification.

Beyond the blind testing, validation efforts included a site visit to RTKC. This allowed for direct comparison of Structura's predictions against physical core observations and manual logs where available. These validation steps were crucial for building confidence in the tool's outputs and understanding its practical application nuances. The collaborative process underscored the relative strengths of both manual and automated approaches and provided insights into developing efficient integrated workflows.

3.3.1 Roughness data

Discontinuity roughness is an important parameter influencing the shear strength along geological structures. Structura calculates a standardised two-dimensional roughness value, arithmetical mean roughness (Ra), for each identified discontinuity by analysing the deviation of the structure's centreline trace from a perfect sinusoid around the unwrapped televiewer borehole image (illustrated in Figure 4). This circumferential profile method provides a quantitative estimate based directly on the televiewer data.

To validate the accuracy and utility of this automated roughness calculation, a specific comparison exercise was undertaken as part of the RTKC pilot study. Fourteen discontinuities for which Structura had calculated Ra values from ATV logs were selected. Physical core samples containing these same discontinuities were retrieved and subjected to high-resolution photogrammetric scanning. From the resulting 3D point cloud data of the discontinuity surfaces, two standard roughness parameters were calculated:

1. Scan Ra: the arithmetic average roughness calculated along profiles approximating the circumference of the borehole wall trace on the fracture surface, analogous to the measurement method used by Structura.
2. Scan areal arithmetic mean height (Sa): the arithmetic average calculated over the entire scanned discontinuity surface area, providing a more comprehensive 3D equivalent to Ra.

The Structura-calculated Ra values were then compared directly against the Ra and Sa values measured from the 3D scans of the discontinuity surfaces (Figure 8).

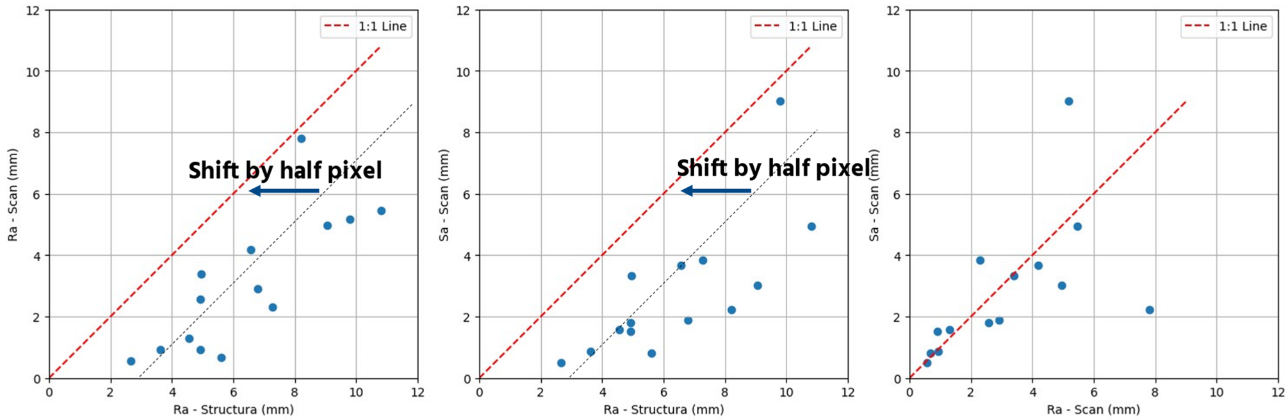

Figure 8 Comparison of automated Structura roughness (Ra) with roughness parameters derived from scans of the physical discontinuities

The comparison reveals a positive correlation between Structura Ra, Scan Ra, and Scan Sa measurements. This indicates the automated method can successfully capture relative differences in roughness between discontinuities. However, a systematic offset was observed, where the Structura Ra values were consistently higher than the corresponding scan measurements. The offset appears to be related to the resolution of the ATV data – pixel boundaries can look like small steps in the roughness profile. As indicated by the 'Half pixel shift' annotation in Figure 8, applying an adjustment equal to half the pixel resolution (e.g. subtracting 3 mm for a 6 mm/pixel log) to the Structural Ra values brings the central tendency of the data in line with the physical scan measurements.

This validation exercise demonstrates that while Structura's automated roughness calculation captures the essential roughness characteristics, the raw output values are influenced by the televiewer data resolution. The identified systematic, resolution-dependent offset suggests that a simple correction factor can be applied to adjust the Structura Ra values, yielding estimates closer to physical reality. This corrected roughness data can then be more confidently used in subsequent set characterisation. However, this also highlights a limitation of the method: the minimum measurable Ra will be approximately equal to half of the data resolution; less rough discontinuity surfaces will be indistinguishable from one another. Assuming total asperity amplitude is approximately two times Ra, this would be approximately equivalent to a lower limit of a joint roughness coefficient (JRC) of 8 for 6 mm resolution data and a JRC of 4 for 3 mm resolution televiewer data. Further investigation across varied geological conditions and televiewer resolutions is warranted to confirm this observation.

3.3.2 Segmentation and orientation data

Evaluating the performance of Structura in accurately identifying geological structures (segmentation) and determining their orientation (dip and dip direction) formed a core part of the validation process. This involved multiple levels of comparison against the manual interpretations performed by experienced RTKC geoscientists.

Validation occurred at two primary scales – the individual structure scale and the aggregated borehole scale. To validate individual structures, direct, side-by-side comparisons were made by overlaying Structura's identified structure picks onto the televiewer logs alongside the corresponding manual picks (e.g. Figure 9), or through expert review of Structura's picks alone (similar to examples shown previously in Figures 4 and 5). This allowed for detailed visual inspection of the agreement on feature depth, type (where classification was available), and the quality of the sinusoidal fit representing the structure's orientation. These detailed reviews confirmed Structura's ability to identify significant geological features across varying data qualities.

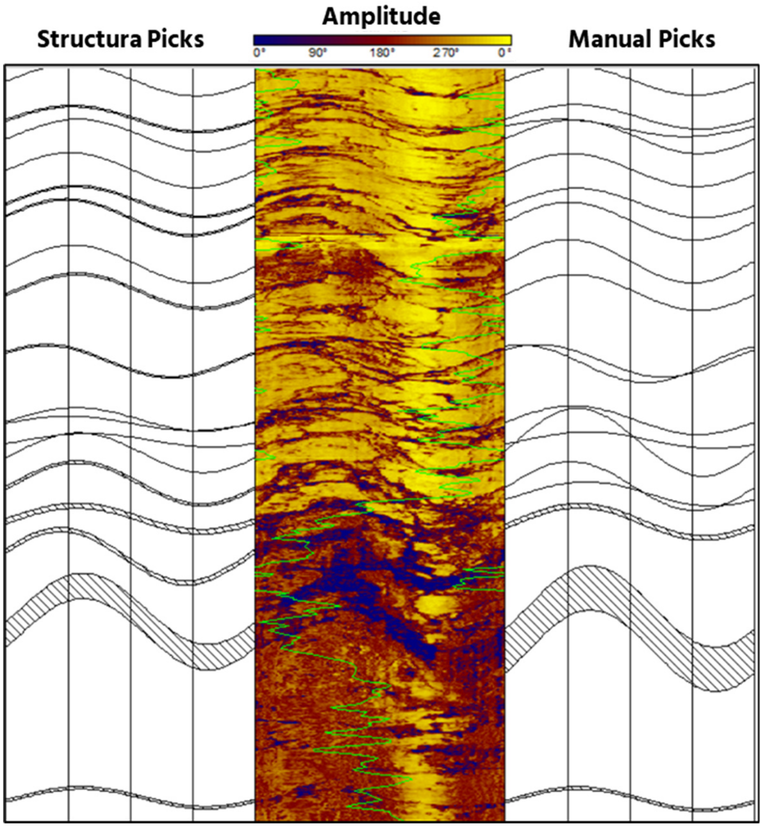

Figure 9 A comparison between structural interpretations between Structura and manual logging

To assess the tool's ability to capture the overall structural fabric within a borehole, orientation data from Structura was aggregated and compared to the aggregated manual data on equal area stereonets. This approach validates whether the automated method identifies the same dominant structural sets and patterns as the manual method. Figure 10 presents a typical example showing good agreement between the automated and manual methods. The overall distribution and density of poles are comparable, demonstrating that Structura effectively captured the primary structural signature of this borehole consistent with expert manual interpretation. In other borehole comparisons, the Structura-generated stereonets (e.g. Figure 11) occasionally exhibited more tightly defined clusters and a less 'donut-shaped' data distribution caused by the biases described in Section 1 compared with their manual counterparts. Based on a review of the underlying data, the apparent sub-vertical cluster in the manual stereonet is caused by spurious picks generated in that problematic angle range. The consistent application of Structura's algorithms can potentially reduce the variability and subjectivity associated with manual picking, particularly with interpreting logs across different quality intervals or by different geoscientists.

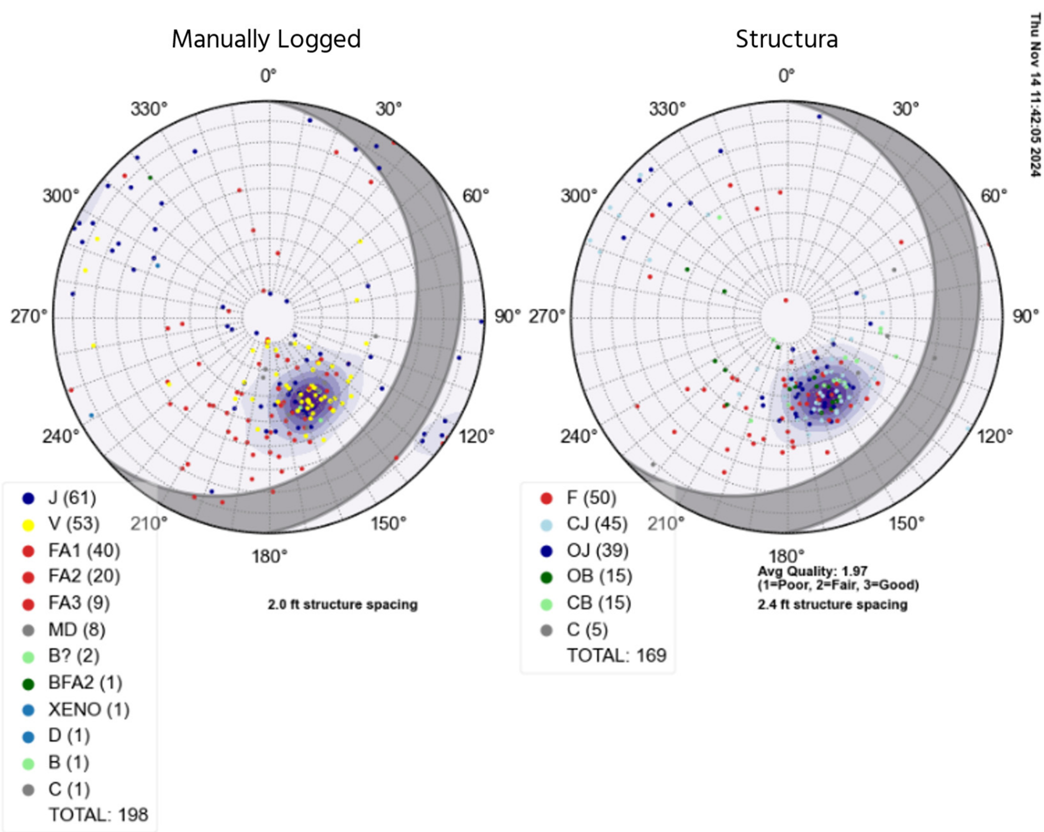

Figure 10 Comparison of stereonets for a drillhole, showing good agreement

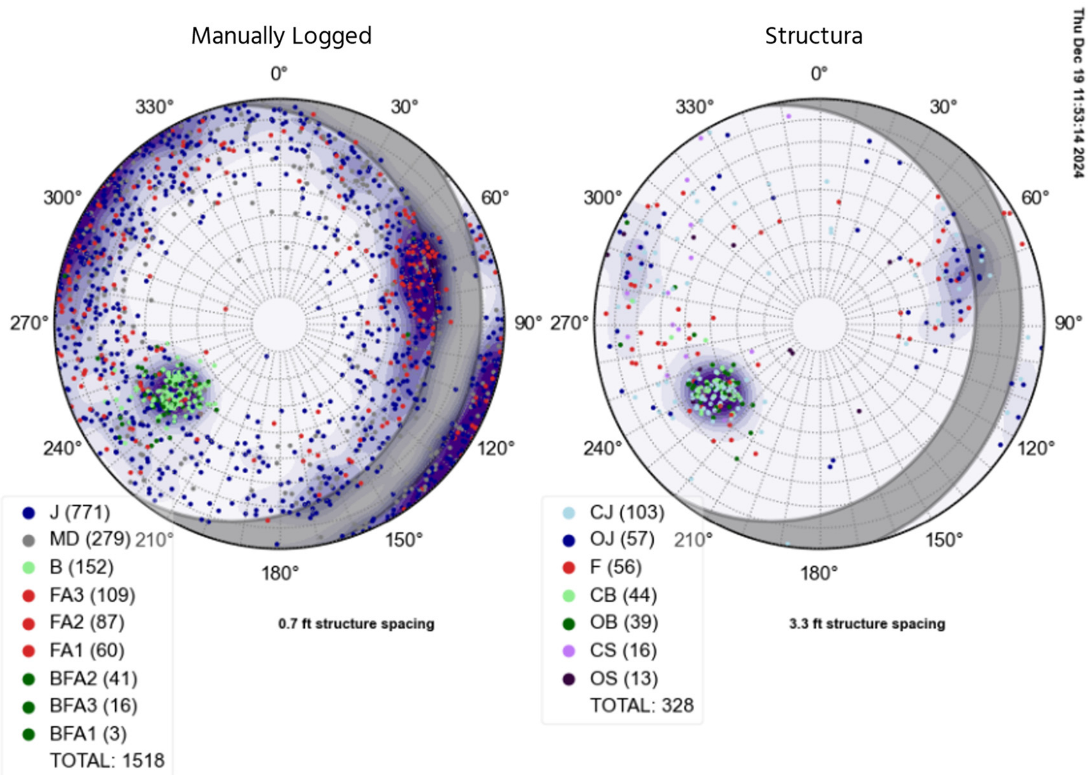

Figure 11 Comparison of stereonets for a drillhole, where the Structura data show reduced visual salience bias

3.3.3 Breakout and drilling-induced tensile fracture data

Validation of the automated identification of borehole breakouts and potential DITFs was performed through qualitative review, which involved scanning through the results and reviewing the identified features overlain on the ATV logs (see Figure 12). These checks confirmed that the algorithm was identifying characteristic borehole wall failure features – such as breakouts and thinner, axially aligned fractures (possible DITFs). The current algorithm identifies these features based on their morphology in the televiewer image but does not yet differentiate between stress-induced breakouts and DITFs.

Qualitative review also involved compiling the distribution of breakout/DITF orientation data. Compiled data from the 10 blind test holes show breakouts/DITFs oriented in a relatively clear quadrimodal distribution, however, the pattern is affected by scatter and would indicate principal horizontal stress orientations approximately 30° oblique to the regional stress orientation (Figure 13). Some discrepancy is expected; the 10 blind test holes are located around differently oriented walls of the pit and include data near the pit that are likely affected by its presence. Aggregated breakout data from a larger dataset of holes, only including data located 150 m (500 ft) or more from the pit, shows a more definitive trend closely aligned with the regional stress data (Figure 13) highlighting the local influence of pit geometry and unloading.

While this study placed less emphasis on direct, feature-by-feature comparison for breakouts than for structural features, the reviews and analysis of aggregated data confirmed that Structura reliably identifies geomechanically significant borehole failure features. The automated detection provides a consistent method for mapping potential stress indicators, offering valuable input for site-scale stress field characterisation, complementing the structural interpretations. Further work could involve direct comparison with calliper logs, or more detailed manual breakout logs if available and refining the algorithm to potentially distinguish between borehole breakouts and DITFs.

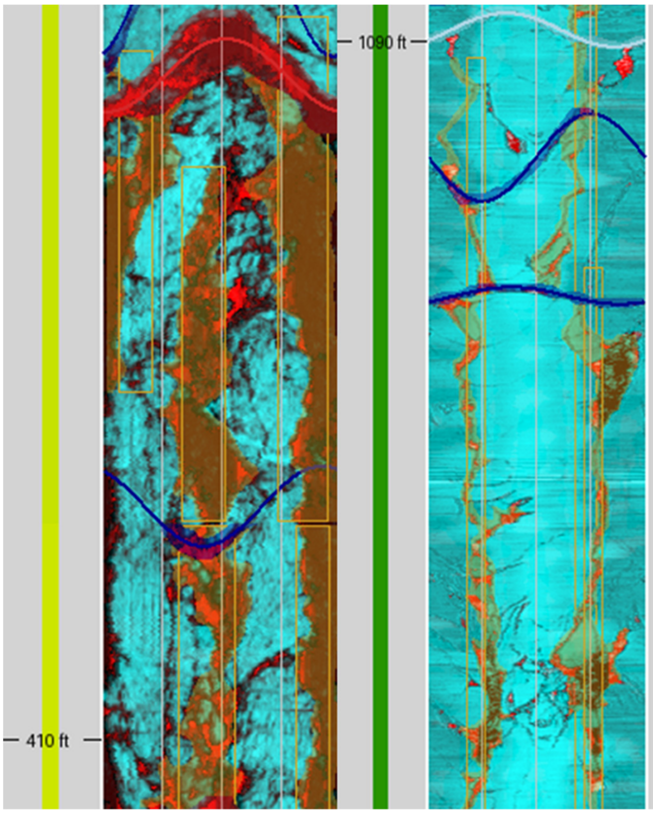

Figure 12 Identified breakouts and thinner features which may represent drilling-induced tensile fractures (yellow highlights). Yellow rectangles highlight the central tendency of the identified features

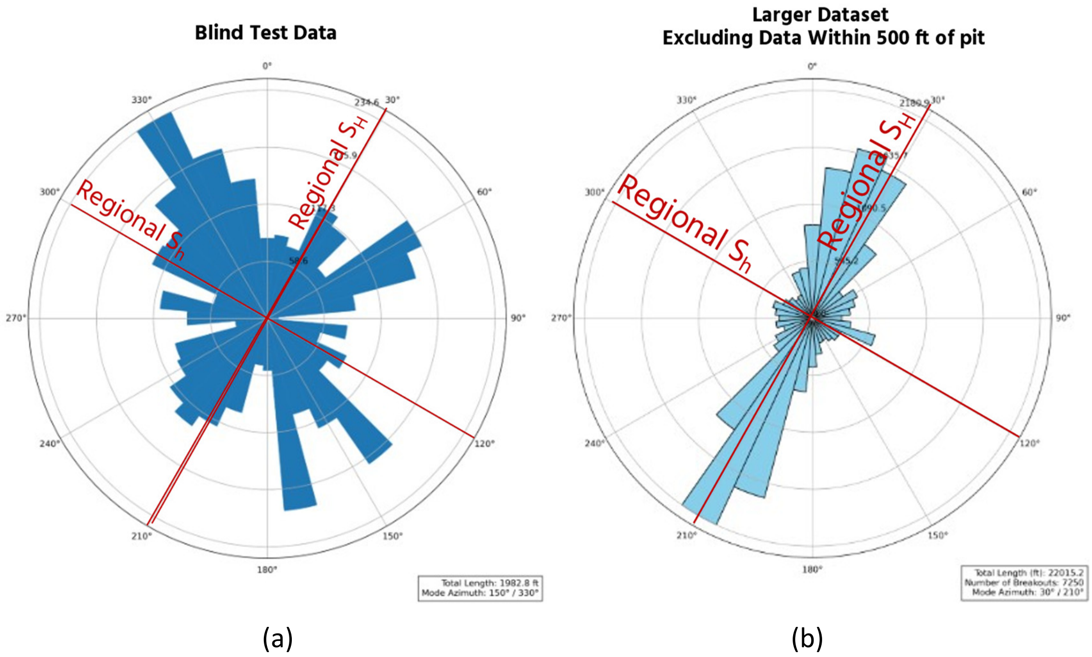

Figure 13 (a) Breakout/DITF orientation distributions for the blind test dataset; (b) Larger aggregated dataset excluding data within 500 ft of the pit

4 Conclusion

Manual interpretation of borehole televiewer data, while fundamental to geological and geotechnical site characterisation in mining, suffers from significant limitations; including high time and cost investments, inherent interpreter inconsistency, and potentially critical orientation biases affecting both high-angle features and low-angle features ("visual salience bias"). This paper presented Structura, an AI-powered software tool developed to automate the interpretation process and detailed its application and validation through a collaborative pilot study with RTKC.

The collaborative pilot study at RTKC demonstrated Structura's potential for enhancing efficiency, consistency, and bias reduction in a complex operational mining environment; however, important limitations remain. The performance of automated tools like Structura is inherently dependent on the quality and resolution of the input televiewer data. Fine-scale features may be missed in lower-resolution logs, and poor image quality can impede accurate identification and classification. The automated roughness estimation requires resolution-based adjustment, and further work is needed to reliably distinguish between breakouts and DITFs or between visually similar planar features (e.g. beds, foliation, lithological contacts) without additional geological context.

The pilot study highlighted significant practical benefits for Kennecott. Key outcomes included improved consistency in structure identification and an estimated 67% reduction in overall ATV data turn-around time, factoring in the essential manual steps like data preparation and QA/QC. The successful refinement of Structura to manage coarser resolution data, identify breakouts, and assess image quality showed its potential to enhance the efficiency of geological interpretation within Kennecott's operational context, provided it is used as a tool integrated within and validated by established expert workflows.

AI-driven automated televiewer interpretation tools like Structura represent a significant advancement in geological data interpretation. They offer the potential to substantially increase the speed, consistency, and objectivity of structural and geomechanical characterisation from borehole imagery. While not replacing the need for geological expertise, these tools can manage the bulk of rote processing, allowing geoscientists and engineers to leverage their expertise for higher level analysis, validation of critical features, and integration of the data into geological models and engineering design. They also provide the opportunity to quickly reprocess historical datasets to replace interpretations of questionable quality or improve consistency.

The use of deep learning technology also introduces the risk of systematic issues, where algorithms might inadvertently introduce errors or biases based on the data they were trained on. This could lead to misinterpretations and skewed results. Careful oversight and validation of results are critical to ensure accuracy and reliability.

Future work should focus on continued model refinement, further validation across diverse geological settings and data types, and developing best practices for integrating these automated outputs into standard data collection workflows – potentially through hybrid approaches combining automated processing with targeted expert review.

Acknowledgement

The authors wish to thank Jens Gibbs, Julianne Laumann, Robin Butz and Jessica Williams of RTKC. Their expert geological and operational insights, provision of data, coordination of site activities, and constructive feedback throughout the pilot study were essential to the refinement and validation of the Structura tool presented in this paper. We also gratefully acknowledge the broader Rio Tinto team, particularly Domenica Cambio, Mario Zelic, and Megan Gaida for their valuable input and support for this collaborative effort.

References

Danielson, J, Clayton, MA, Kelly, L & Kinakin, D 2023, 'Getting more out of drillhole televiewer data: geotechnical toolbox edition', in PM Dight (ed.), *SSIM 2023: Third International Slope Stability in Mining Conference*, Australian Centre for Geomechanics, Perth, pp. 421–434, https://doi.org/10.36487/ACG_repo/2335_26

Danielson, J, Kinakin, D, Stilwell, I & Ladyman, E 2024, 'AI powered discontinuity characterization in televiewer data', *Proceedings of Slope Stability 2024*, Brazilian Geotechnical Society, Belo Horizonte.

Hoek, E & Bray, JW 1981, *Rock Slope Engineering*, 3rd edn, Institution of Mining and Metallurgy, London.

Li, J, Wu, J, Wu, S, Wang, T & Xi, Y 2024, 'Deep learning-based method for rock discontinuity recognition in complex-stratum borehole images', *Journal of Tsinghua University (Science & Technology)*, vol. 64, no. 7, pp. 751–760, https://doi.org/10.16511/j.cnki.qhdxxb.2024.26.020

Read, J & Stacey, P 2009, *Guidelines for Open Pit Slope Design*, CSIRO Publishing, Collingwood.