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

doi:10.36487/ACG_repo/2465_03

Automated extraction of discing regions from core photography using computer vision at Sunrise Dam

Sam Johnson [a,*], Yasin Dagasan [a], Callen Banff [b]

[a] *Datarock, Australia*

[b] *AngloGold Ashanti, Australia*

Abstract

As the global mining industry seeks to meet the ever-increasing demand for metals, mineral deposits are trending deeper, larger, and lower grade. Engineers play a crucial role in ensuring the extraction process is efficient. A key aspect of decision-making involves using subjectively logged geotechnical datasets, commonly collected by geologists, geotechnical engineers, and technicians. However, these datasets are often excluded from critical mining models due to a lack of quality, auditability, consistency and completeness.

Recent advances in computer vision, particularly deep learning, have provided algorithms capable of efficiently identifying features of interest in core photography. Coupled with domain expertise, these algorithms can provide a computer vision-assisted core logging process, significantly increasing the value of downhole datasets. We have developed a novel workflow using computer vision to automatically identify discing regions, highlighting areas of the rock mass exposed to high in situ stress. By detecting every fracture and measuring their angles, we can define discs, group adjacent discs to create a discing region, and then group nearby regions to extract a consistent dataset. This workflow allows geotechnical engineers to establish a standard discing definition, facilitating high confidence in the data.

The outputs are evaluated against traditionally logged geotechnical data to create a detailed comparison. This study demonstrates that detailed, consistent, and auditable geotechnical data can be extracted using computer vision and core imagery, significantly improving data collection workflows for the mining industry. This approach enables mine planners and geotechnical engineers to proactively manage potential hazards, integrate these risks into mine design and scheduling, and ultimately ensure safer and more efficient mining operations.

The developed workflow was implemented at the Sunrise Dam Gold Mine (SDGM), demonstrating the advantages of computer vision-assisted core logging over traditional methods. This implementation underscores the potential of our approach to enhance the efficiency and reliability of geotechnical data collection in the mining sector.

Keywords: *geotechnical data, drillcore, discing, geotechnical hazards, artificial intelligence, computer vision, machine learning*

1 Introduction

Discing, also known as core discing (or core disking) is a phenomenon where cylindrical rock core breaks into thin, disc-shaped fragments during drilling. When drilling into rock, especially at great depths, the rock surrounding the borehole is subjected to significant in situ stresses. These stresses are typically compressive and are due to the weight of the overlying rock and other geological forces (Jaeger & Cook 1963; Stacey 1982). As the drill bit advances into the rock, it creates a cylindrical hole. The material that is drilled out of this hole was previously bearing part of the in situ stresses. This process of stress alteration and relief causes the discing phenomenon (Zheng et al. 2020). The resulting disced cores often display smooth, saddle-shaped

[*] Corresponding author. Email address: samjohnson@datarock.com.au

fracture surfaces, indicative of the intense stress conditions they endure. In highly stressed environments, the spacing between adjacent discs tends to decrease, signalling higher stress concentrations (Stacey 1982).

Discing is a critical indicator of the stress state within a rock mass, signifying the presence of significant tensile stresses perpendicular to the borehole axis that exceed the rock's tensile strength, leading to fracturing (Jaeger & Cook 1963; Stacey 1982). Understanding the occurrence of discing is essential for grasping the mechanical behaviour of the rock and making informed decisions about mining methods and locating infrastructure. Furthermore, by analysing the patterns and characteristics of core discing, engineers can infer stress conditions and make more informed decisions regarding mine design and support systems (Jaeger & Cook 1963; Stacey 1982). As a result, having a high quality, deposit-wide discing dataset can facilitate a better understanding of the stress distribution within the rock mass which is important for several aspects of a mining operation.

Traditionally, discing data is collected manually by geologists and geotechnical engineers. This involves examining core samples for fractures, measuring the thickness and orientation of the discs, and recording these observations (Jaeger & Cook 1963; Stacey 1982; Li et al. 2019; Zheng et al. 2020). However, this process is time-consuming and prone to human error, resulting in inconsistencies and difficulties in data reproducibility (Ohta 2001; An et al. 2017; Belov & Ivanov 1992).

Discing regions are often described in terms of the number of discs per unit length or the thickness of the individual discs. The thickness of these discs can vary significantly depending on stress conditions, with thinner discs indicating higher stress levels (Jaeger & Cook 1963; Stacey 1982). The manual logging of discing regions is challenging due to the subjective nature of visual inspections and the variability in human observations (Ohta 2001; Corthésy & Leite 2008). This method is labour-intensive and often fails to capture the full extent of the discing phenomenon, especially in complex geological settings (An et al. 2017; Matsuki et al. 2004). Furthermore, the manual approach does not lend itself well to the high throughput demands of modern mining operations, where quick and accurate data collection is essential (Zheng et al. 2020).

Recent advancements in computer vision and deep learning offer promising solutions to these challenges. Automated systems can analyse core photographs to identify discing regions accurately and consistently. By detecting and measuring fractures and their orientations, these systems can define discs, group them into regions, and aggregate nearby regions to create comprehensive datasets. This automated approach ensures high confidence in the data, improving the reliability and efficiency of geotechnical assessments.

Computer vision-assisted logging not only speeds up the data collection process but also enhances the accuracy and auditability of the data. This is particularly beneficial for creating detailed comparisons between traditional logging methods and automated systems, highlighting the strengths and weaknesses of each approach. Studies such as those by Jutzeler et al. (2023) on volcanic stratigraphy reconstruction using machine learning, and Johnson et al. (2023) on extracting geotechnical data from drillcore imagery at the Carrapateena deposit, demonstrate the efficacy of these advanced techniques. Other works, such as those by Jutzeler et al. (2022) on machine-learning image analysis to reconstruct lithostratigraphy in mineralised terrains, and Jutzeler et al. (2021) on quantifying crystal size distribution in volcanic rocks, further validate the potential of machine learning in geological studies. Additionally, research by Dagasan et al. (2021) on inferring geological features masked by artefacts in core photography using neural networks underscores the importance of these technologies in enhancing geological interpretations. Ultimately, the integration of computer vision in core logging represents a significant advancement in geotechnical engineering, offering a robust tool for improving mining operations and safety.

In this paper, we explore the automated extraction of discing regions from core photography using computer vision. The proposed workflow leverages deep learning algorithms to detect and measure fractures, grouping them into discing regions. This method is evaluated against traditional manual logging, with the findings indicating significant improvements in data consistency, accuracy, and auditability. We applied this method at the SDGM in Western Australia, where it was identified that a reliable discing dataset was required to better characterise the rock mass. The integration of computer vision in geotechnical data collection at SDGM has the potential to enhance mining safety and efficiency, providing a standardised and reliable approach to understanding in situ stress conditions in rock masses.

2 Sunrise Dam Gold Mine

The Sunrise Dam Gold Mine (SDGM), located 55 km south of Laverton, Western Australia, is a significant producer of gold for AngloGold Ashanti. This large-scale operation leverages both open pit and underground mining methods to extract ore from the Laverton Greenstone Belt. Production commenced in 1997 from the open pit, which reached its final depth of 500 m by 2014. Underground mining, targeting ore deposits that extend to depths greater than 1,400 m below the surface, began in 2003 and has since become the primary source of ore. As outlined by AngloGold Ashanti (2023), SDGM has a gold mineral reserve estimated at 1.04 million ounces, and the mine produced over 252,000 ounces of gold in 2023. A dedicated workforce of 763 people keeps the mine operational.

Over the SDGM mining area, more than 2,400 km of diamond drilling have been collected and geologically logged. In comparison, a total of 480 km of the diamond drilling has been geotechnically logged, which has been undertaken by a mix of consultants, graduates, geologists and geotechnical technicians or engineers over the life of operations depending on the data collection strategy at the time. Logging methodology also evolved over the years of operation as the focus on different data types was clarified as the orebody and rock mass knowledge developed. However, this meant it was difficult to quantify the accuracy of historical logs of some data types that may have not been collected as comprehensively as needed for new analysis. As some areas of the SDGM resource head deeper, stress-related hazards have the greater potential to affect mining. This has proactively been recognised to requiring more comprehensive logging of discing to identify areas of high stress to build into hazard mapping and risk amelioration strategies.

As historical data consistency and discing logs required review, building a robust stress data model collection was considered paramount for creating a comprehensive approach to risk management. However, the task of reviewing and accurately identifying areas of discing throughout hundreds of kilometres of core is both resource-intensive and slow.

Machine learning offers a cost-efficient and time-effective way to collect consistent, validated datasets over vast amounts of data with minimal human resources being taken away from operations or incurring large consulting fees. This methodology also provides meaningful data in a fraction of the time needed to review historical core photos. This has enabled the use of holes not previously geotechnically logged, such as exploration holes and diamond resource drilling, to identify areas of potential stress-related discing and better understand where the mine may experience issues during excavation.

3 Method

3.1 Overview

The method developed involves several steps to extract discing regions from core imagery. These steps can be seen in Figure 1 and will be explained in more detail in the following sections.

Core discing ranges from crushed through the thick discs (Lim & Martin 2010), and the visual differences within the core images across the range can be quite different so it is important to have a workflow that is robust to this variation. We have developed two key methods for identifying discs:

- Fracture angle analysis: This method assesses the angle of individual fractures to identify regions of medium to thick discs.
- Crushed/thin discing detection: This method identifies regions of discing where individual discs cannot be assessed and are instead detected as a whole region.

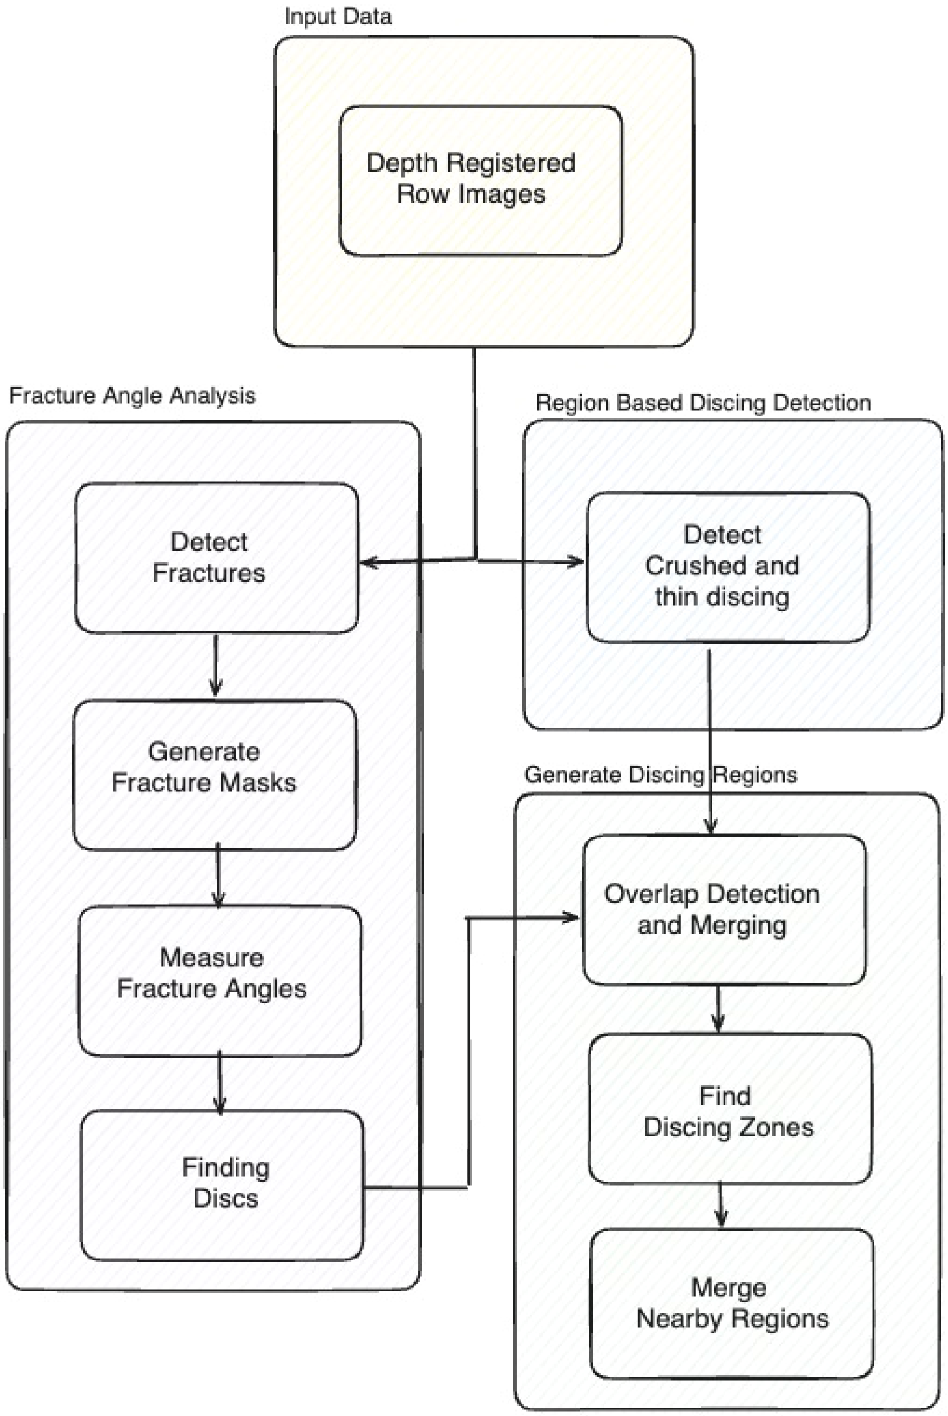

Figure 1 Overview of the workflow to extract discing regions from core imagery

3.2 Depth registration

Depth registration is the process of assigning accurate depths to core images, which is crucial for pinpointing locations of specific structures, such as discing. Accurate depth registration also enables reliable comparisons and calibration with other downhole datasets. This process involves several preprocessing steps to produce 'analytics-ready' imagery. These steps include de-warping images as needed, cropping core trays and rows, identifying coherent and incoherent rock within rows, quantifying compaction in incoherent rock, addressing core loss, correcting metadata entry errors, and applying optical character recognition to depth annotations. In this paper, we employed the Datarock Core product to prepare core row images for analysis. Examples of depth-registered rows and the assigned depths are illustrated in Figure 2.

Figure 2 Example of a depth-registered core rows generated by the Datarock Core product. Each row has assigned depths from and to, including detected depths based on handwritten metre marks (green markers) and inferred depths (yellow markers)

3.3 Angle-based disc detection

Fracture angle analysis is crucial for identifying discing, as discing usually occurs roughly perpendicular to the core axis. However, due to the handling and placement of the core on the tray, these angles cannot be perfectly perpendicular. Therefore, we have developed a method to assess whether the fracture angle relative to the core axis could potentially belong to a discing. The following sections will describe the method in more detail.

3.3.1 Detecting fractures

The method for identifying discing regions is based on detecting and classifying fractures within depth-registered core images. We used RetinaNet (Lin et al. 2017) to train an object detection model that classifies fractures into six main categories. Examples of these fracture classes can be seen in Figure 3.

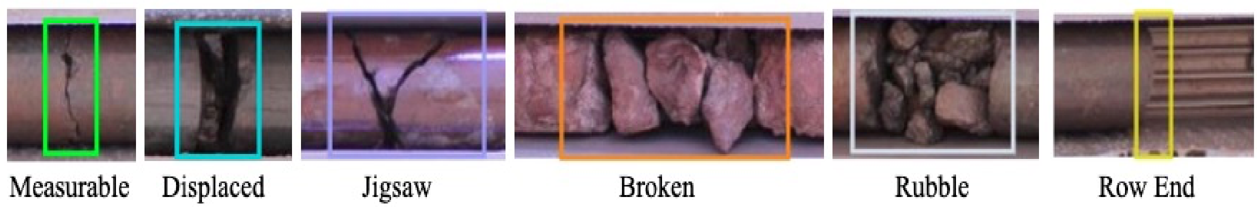

Figure 3 Example of the fracture detection classes used to identify the various types of fractures that are present within the core for identifying discing regions

We developed specific classes to categorise fractures within a geotechnical context, providing a consistent training dataset for our object detection algorithm. The training dataset was compiled by sampling and labelling individual row images to represent the global variability of the data. The dataset was then used to train the fracture detection model, which performed fracture detection on each depth-registered row image, as illustrated in Figure 4.

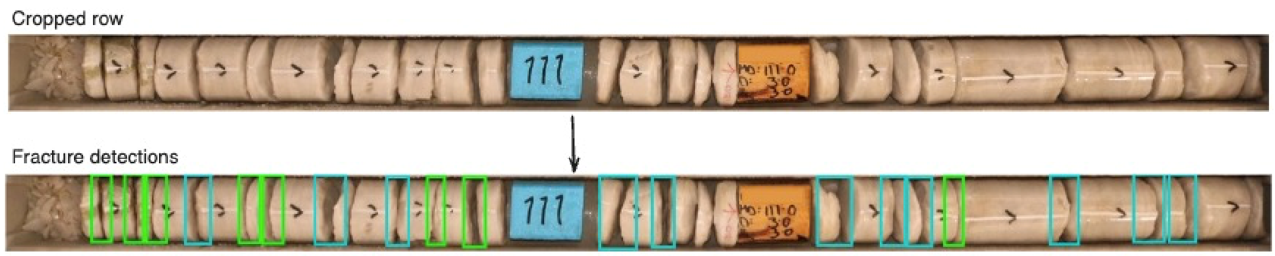

Figure 4 Example of fracture detections within a discing zone

3.3.2 Generating fracture masks

A key step in enabling detailed analysis of each individual fracture is the generation of a clean fracture mask. In the context of segmentation, a mask is a binary image where pixels corresponding to the fracture are marked, allowing for accurate geometric analysis. The following steps outline how the detected fracture is converted into this fracture mask, and are depicted in Figure 5:

1. Fracture crop: Once the bounding box of measurable or displaced fractures is identified, the bounding box is used to crop the image.
2. Raw fracture segmentation: The cropped image is then fed into a segmentation model, which predicts a segmentation mask for the fracture.
3. Thresholding: The output from the segmentation model undergoes thresholding to refine the mask, ensuring that only the pixels belonging to the fracture are selected.
4. Mask cleaning: The initial mask may undergo further processing to remove noise and ensure clean edges, resulting in a more accurate representation of the fracture.

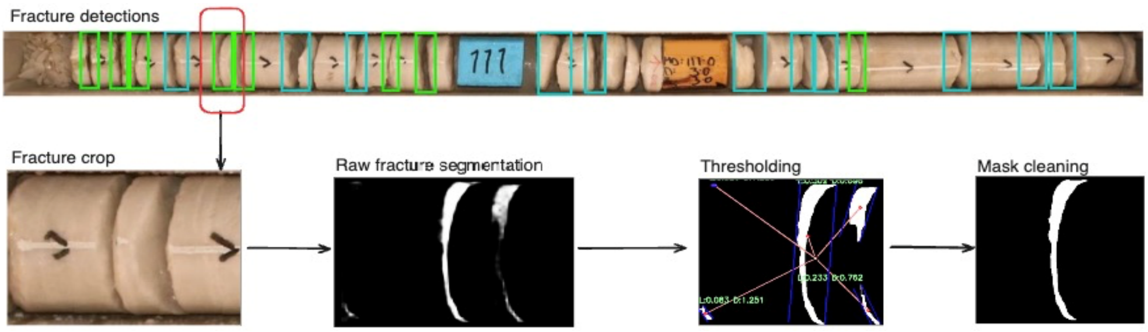

Figure 5 Example of the steps to generate a fracture mask. These steps are completed on every detected fracture

3.3.3 Measuring fracture angles

After generating the clean fracture masks, the next step involves calculating its angle with respect to the core axis. The two steps are described here, with an illustration of the line fitting in Figure 6.

1. Line Fitting: A line is fitted to the segmented fracture mask to represent the fracture's orientation.
2. Angle Measurement: The angle of the fitted line is measured with respect to the core axis.

Figure 6 Example of fitted lines used to calculate the angles of fractures. The fitted lines represent the orientation of the fractures, with angles measured relative to the core axis

3.3.4 Finding discs

Once the fracture locations are found in terms of depth, the distance between successive fracture centres is calculated, and any distance less than 4 cm is flagged as a potential disc. This serves as a filter based on disc size, ensuring that only relevant discing features are considered for further analysis. For each detected disc, the following properties are calculated:

- Right and left depths: The starting (left) and ending (right) depths of the disc.
- Right and left angles: The angles measured at the starting and ending depths.
- Disc Size: The difference between the right and left depths.
- Disc Depth: The midpoint between the right and left depths.

The data is further filtered based on the angles of the fractures to ensure consistency. Both the right and left angles are checked to ensure they do not deviate by more than 10° from 90° with respect to the core axis, indicating they are nearly perpendicular to the radial axis. Additionally, the absolute difference between the left and right angles is checked to ensure it does not exceed 5°. This dual-threshold approach ensures only discs with consistent angle measurements that exhibit discing characteristics are retained. Figure 7 illustrates some of the examples that pass the discing filtering.

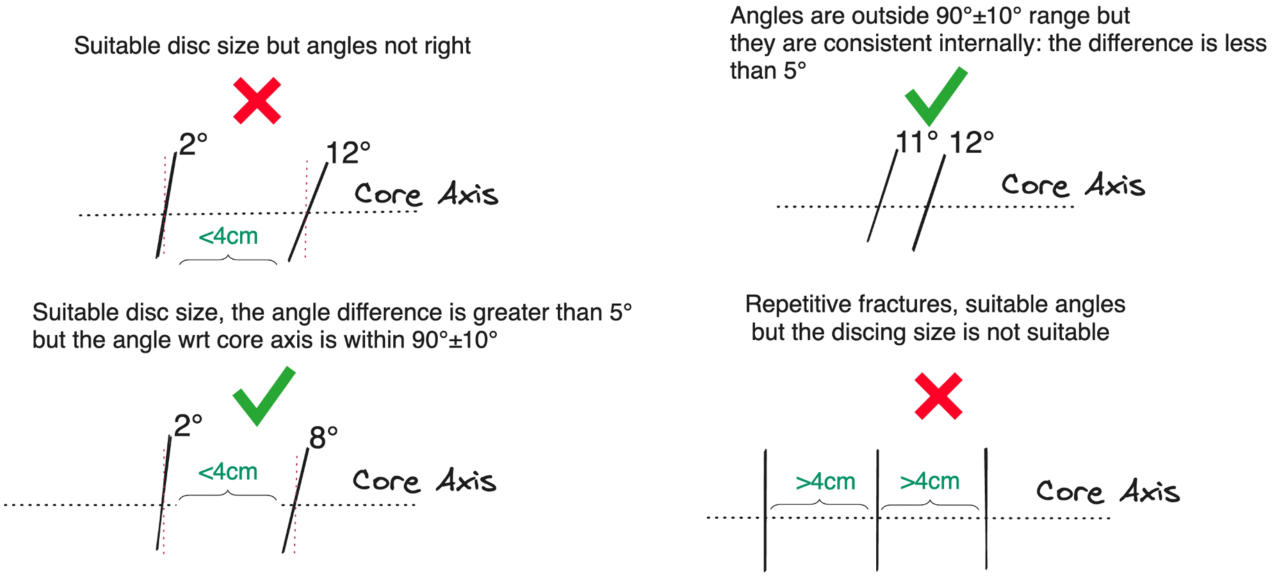

Figure 7 Examples of fracture angles that are either suitable or not suitable for potential discing

3.4 Region-based discing detection using an object detection algorithm

In areas with a high presence of broken rocks and irregular discs, a separate model is employed to identify discing regions. This model, an object detection model using RetinaNet, encloses the start and end locations of a discing region within a bounding box. This approach was necessary due to the wide-ranging visual appearance of discing regions. The object detection method is chosen for its ability to accurately detect and localise objects within an image, which makes it suitable for identifying complex discing regions that may not be easily captured through angle measurements alone. An example of a discing region that can be better captured by the object detection can be seen in Figure 8.

Figure 8 Complex discing regions shown on the left end of the image are better suited for detection by the region-based approach rather than the angle-based one

3.5 Overlap detection and interval merging

To combine the discing regions from angle-based and region-based methods, the algorithm checks for overlaps between discs and broken regions within each hole. The length of the overlap is computed, and when overlaps are detected, the intervals are extended. The overlap percentages are calculated to update the disc properties accordingly. This ensures that overlapping intervals are merged, and their properties are adjusted to reflect the combined region. Broken regions that do not overlap with any disc are processed separately to ensure accurate representation in the final analysis. These non-overlapping regions are handled by extending the intervals and updating their properties independently, maintaining the integrity of the overall dataset. The diagram in Figure 9 explains how the merging takes place.

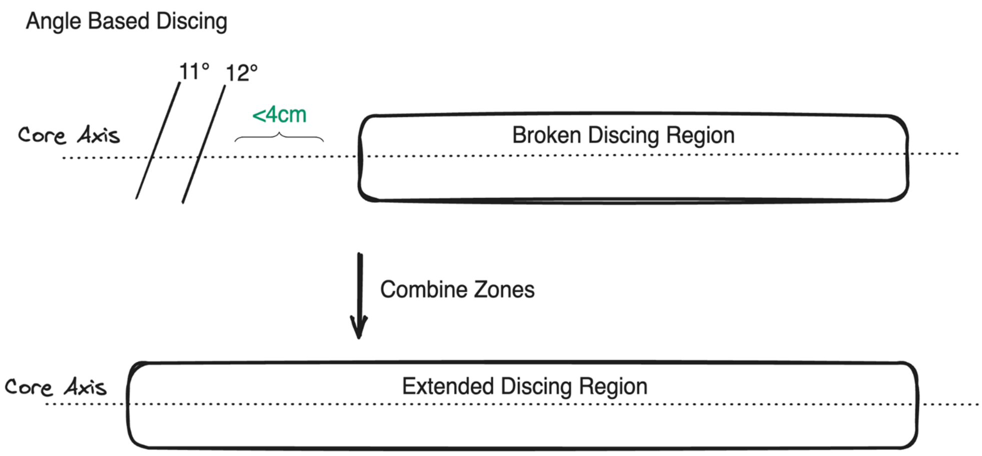

Figure 9 Diagram illustrating the merging of discing regions from angle-based and region-based methods

3.6 Finding discing zones

After merging angle-based and region-based discs, the discs are grouped based on depth differences within each hole. For instance, if two of the same angle-based discs are successive, they are designated as a single discing region. The resulting outcome is referred to as granular discing. Each discing region is summarised in terms of average disc size, number of discs, average angles and mean of absolute angle differences. These summaries provide a comprehensive overview of the discing patterns within the borehole. An example diagram illustrating this merging process is shown in Figure 10.

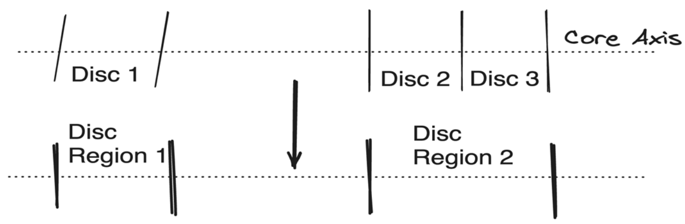

Figure 10 Example diagram showing the merging of angle-based and region-based discs into a single granular discing region-based on depth differences

3.7 Merging nearby disc regions

Discing logging is usually done by grouping (or lumping) discs together to define larger discing zones. One reason for this approach is that granular discing detection data can result in numerous small regions and grouping them facilitates easier analysis on a larger scale. This data is also provided along with the granular discing regions. The principle is that if granular discing zones are within 30 cm of each other, they are grouped together. An example of the grouping process is illustrated in the diagram in Figure 11.

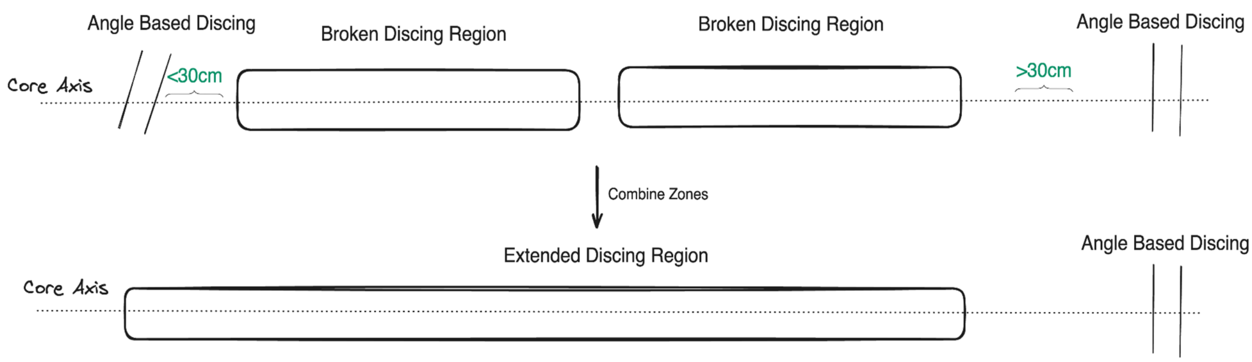

Figure 11 Example diagram illustrating the grouping of granular discing zones into larger discing regions based on proximity

4 Results

In this section we review the detected discing regions using two methods:

1. Using holes where manual site logging (including discing) has been completed, we compare detected discing against site logging. This dataset consists of nine holes, or 7,200 m of drillcore.
2. Using long (1,000 m+) holes across the site to provide vertical and lateral coverage across the deposit, we compare the detected discing against site modelled structures to determine if there is a correlation between the two. This dataset consists of 77 holes, or 85,000 m of drillcore.

4.1 Comparison against logging

Across the SDGM deposit, there were nine holes that included logged discing regions. These were used to compare against the detected discing created by the method described in this paper. As described in Section 2, historically there was not a significant need to understand the in situ stress regime, therefore it was expected that there would be regions of discing that had not been logged.

An example downhole plot in Figure 12 shows the areas where major discing takes place. As illustrated, there is strong agreement around 610–650 m, while there are other regions detected where logging was not recorded (as expected). Examples of these are described in this section.

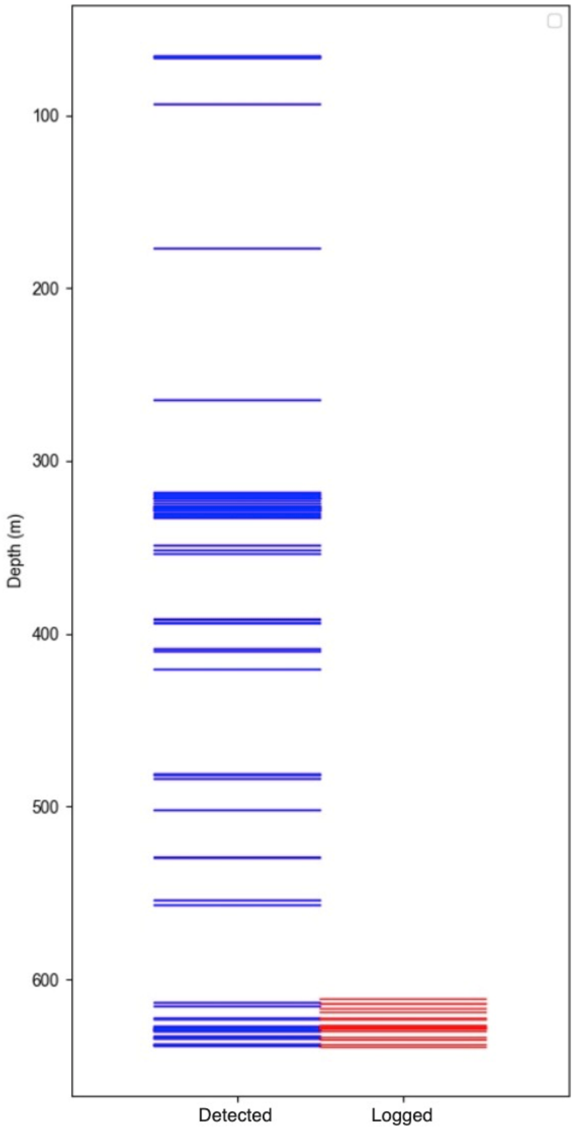

Figure 12 Downhole plot comparing detected and logged discing regions

4.1.1 Example of true positive where logged interval also exists (agreement between datasets)

A true positive is where a detected discing region corresponds to an actual discing region, regardless of whether it was logged onsite or not.

It can be observed that both datasets identify discing in the 610–640 m range. Within this range, both datasets have an interval from 626–632.2 m, meaning this is a good example where both the detections and the logging are correct.

On review of the logged intervals across this dataset, it was generally the case that where there was a logged region of discing, there was a corresponding detected region. An example of true positive discing region is shown in Figure 13.

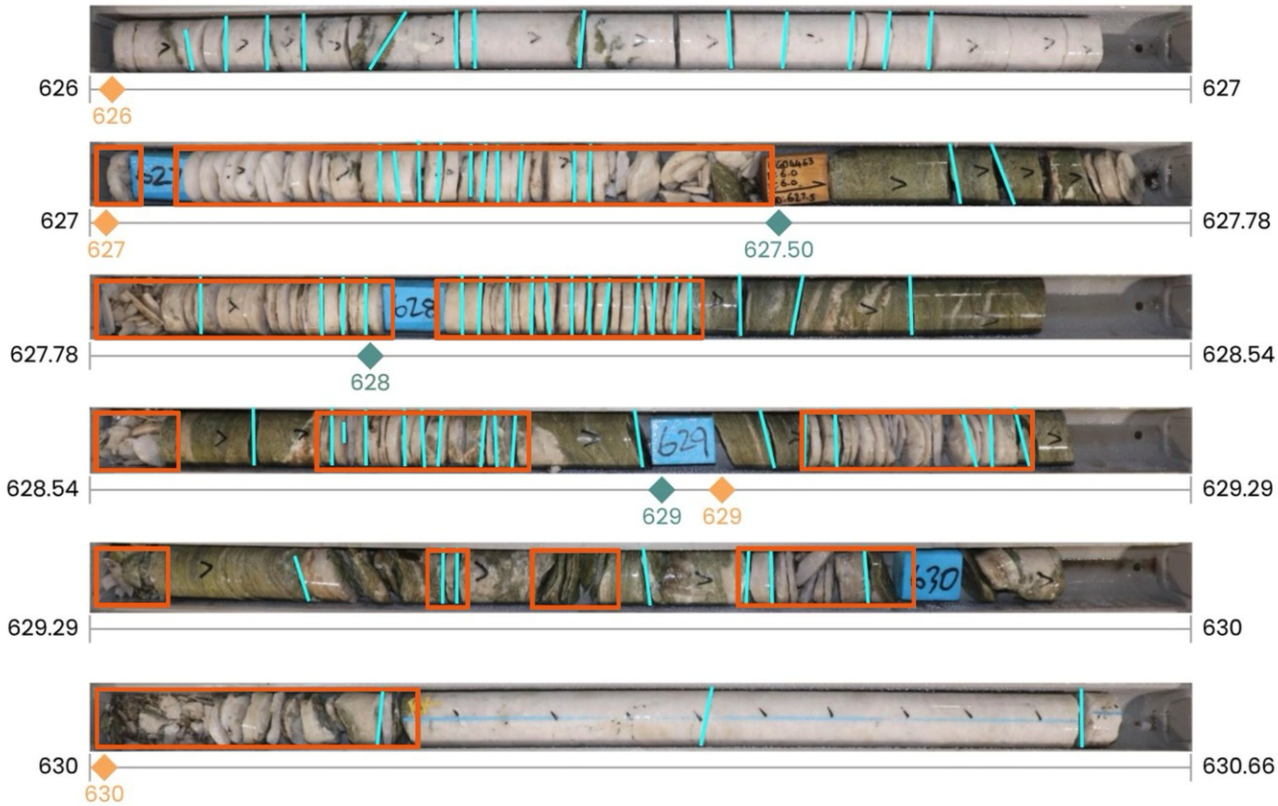

Figure 13 Example of a discing region that has strong correlation between logged and detected datasets. Angle lines are shown in light blue, region-based detections are shown in the orange boxes

4.1.2 Example of true positive detection where no logging exists

True positives also occurred in areas where there was no corresponding logged region.

An example of this can be observed in Figure 12 where there is a major discing region around 330–350 m which was not in the site logging dataset. This discing region was detected by the developed workflow, and a portion from 318–326 m of the core imagery from this region can be seen in Figure 14.

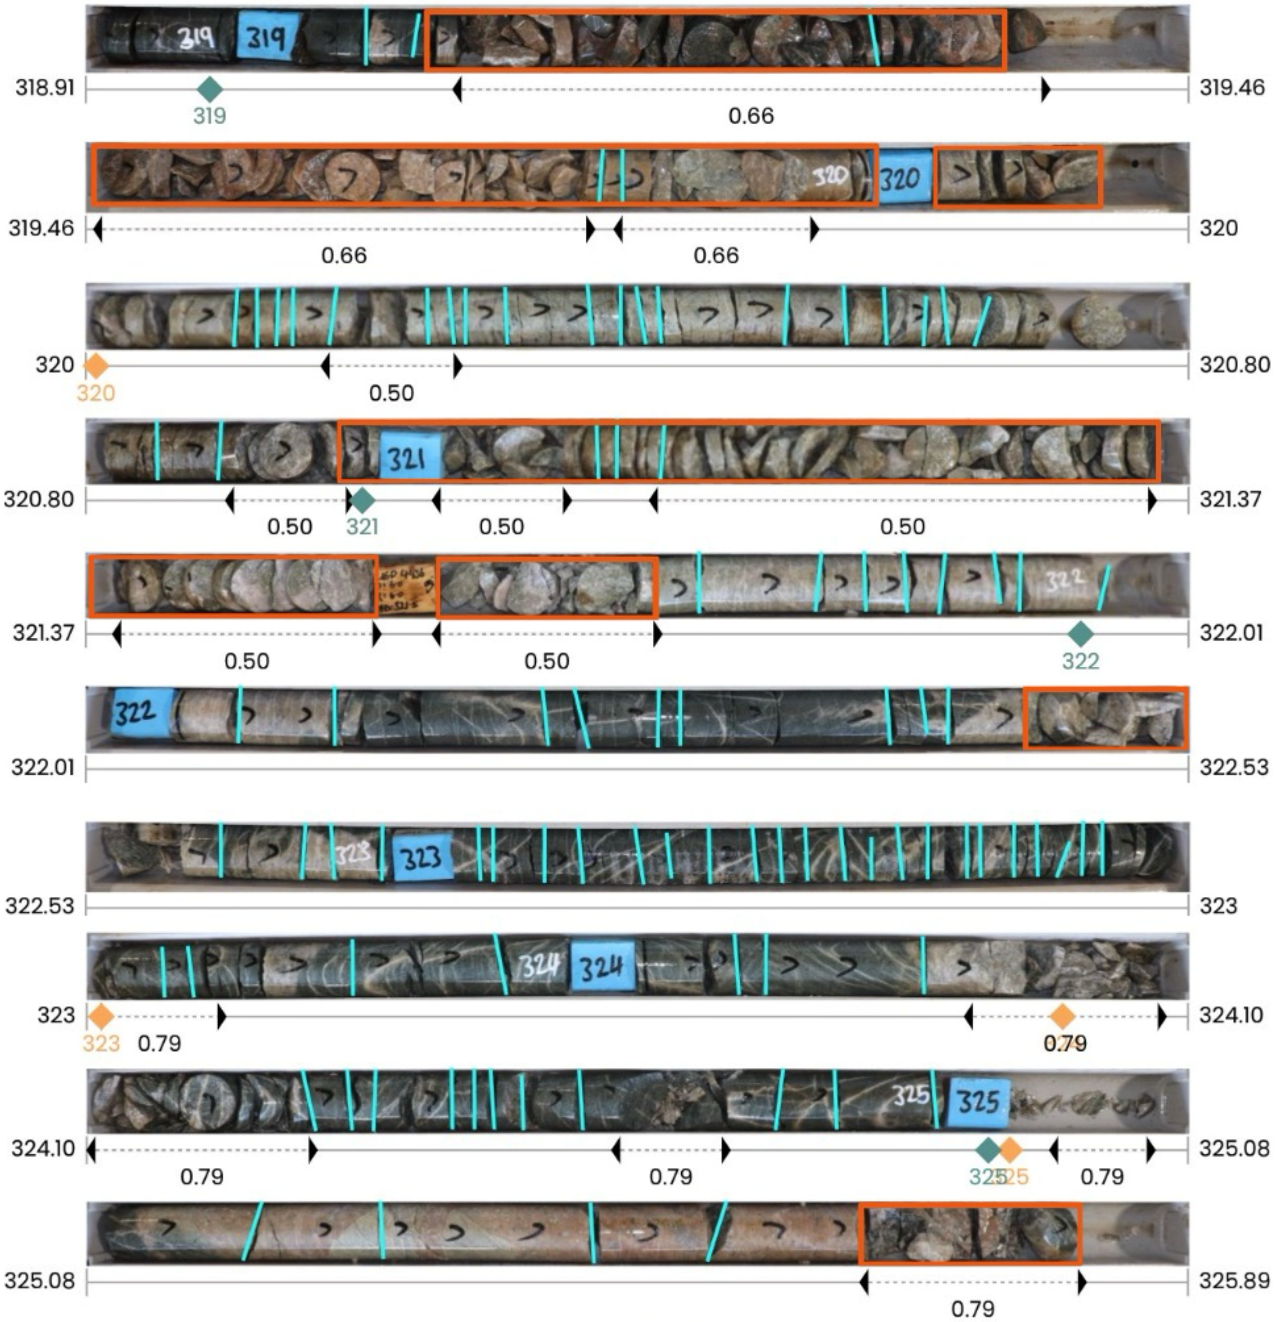

Figure 14 Examples from the major discing region detected by the developed algorithm. Angle lines are shown in light blue, region-based detections are shown in the orange boxes

4.1.3 Example of false positive detection where no logging exists

A false positive is where a detected discing region does not correspond to an actual discing region.

An example of this can be observed in Figure 15 where there is a region-based detection at 93.3 m. There is broken rock with rounded surfaces within the zone, resulting in the model incorrectly identifying the region as discing. Minimising false positives will be a key area of further development to enhance its accuracy and robustness, with two key areas of focus being:

1. Incorporating additional manually labelled training data to cover more variability in the core imagery.
2. Refining the model architecture to ensure it is best fit for this use case.

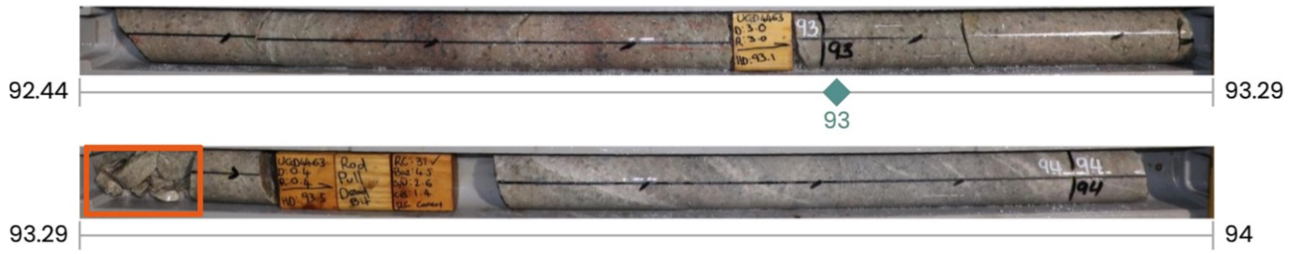

Figure 15 Example of a false positive discing detection in the orange box

4.1.4 Comparison to site modelled structures

To gain a better understanding of the performance of the discing detections, detected regions were modelled against the site modelled structures which may have higher stress concentrations such as the shear or faulted zones.

Using a total of 77 long holes (typically greater than 1,000 m in length) to provide vertical and lateral coverage across the deposit, approximately 85,000 m of drillcore was processed using the described workflow, generating a total of 774 individual discing regions. It should be noted that this dataset was selected to provide coverage, however it is only 6% of the available core imagery, and it is the authors' intention to process the full dataset once further improvements of the method have been completed (as the following discusses).

Comparison of the workflows identified that discing was highly correlated with the interpreted large-scale structures as illustrated in Figure 16. However, additional zones were also highlighted away from modelled areas, indicating that additional zones of high stress or large-scale structures may be present and will need further interpretation.

Of the discing regions furthest away from modelled structures, several are false positives associated with closely spaced broken rock, generally at the relatively shallow areas of the mine. However additional true positives have also been identified that can be used to further inform the structural model in the region. As additional drill holes are added to the dataset, it is expected the structural trends will be more easily seen and false negatives will become visibly identifiable as outliers.

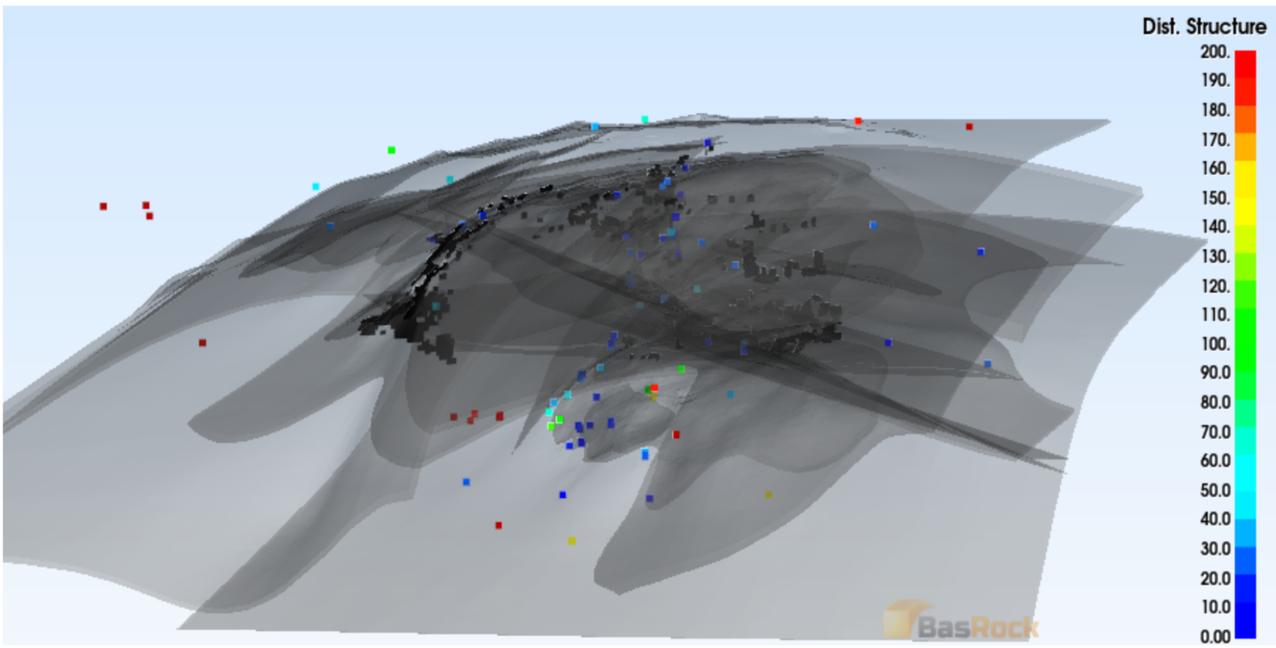

Figure 16 Correlation between discing regions and shear/fault structures. The colours indicate the distance between the discing locations and the shear/fault structures

Another interesting observation was made from some of the discing regions which were perpendicular to the dominant direction of faulting and could potentially be un-modelled structures, as can be seen in Figure 17.

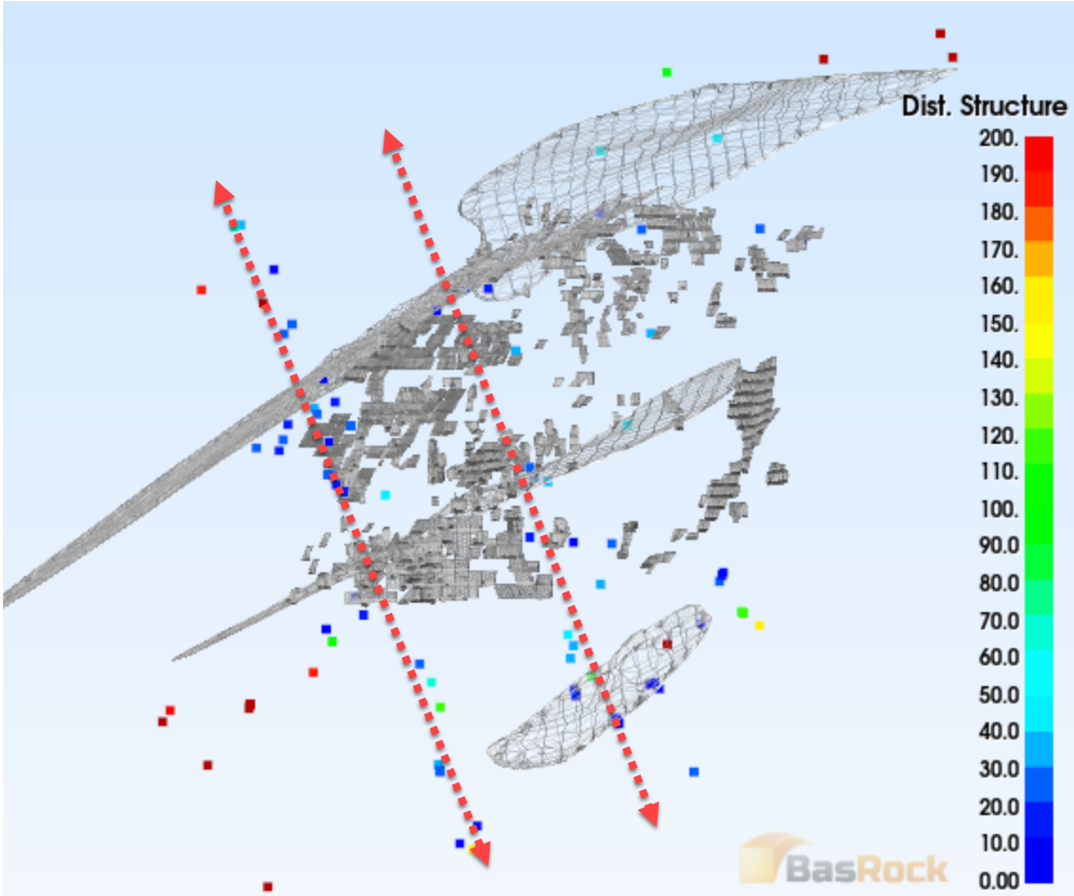

Figure 17 Discing regions at roughly 400 m spacing that are perpendicular to the dominant direction of faulting, which could indicate un-modelled structures. Indicative large structures shown to indicate the primary modelling trend of structures

5 Discussion

The results show that the detected discing regions extracted by this method correlate well with the actual discing regions. While diamond drilling itself is quite expensive, drillcore photographs are among the least expensive yet highly informative datasets that all exploration projects, development projects, and operational mines routinely collect. The implications of having a robust method for extracting discing regions from these images are significant, especially in the context of deposits trending deeper, becoming larger, and presenting lower grades.

Identifying geotechnical structures is paramount for a mining operation. Proactively managing potential hazards enables mine planning engineers and geotechnical teams to integrate these risks into the mine's design and scheduling. This leads to more robust schedules for mine planners, as understanding rock mass conditions allows for adjustments in production rates and strategic planning of stoping areas, thus managing ore delivery to the crusher more effectively.

Geotechnical engineers can utilise this information to identify areas requiring additional rigour in design, adjust ground support plans, and reduce uncertainty in their designs. This proactive approach ensures safer and more efficient mining operations.

Using the machine learning approach outlined here provides a more accurate and consistent methodology applicable across the entire deposit. This approach allows site or corporate engineering teams to create a standard discing definition, facilitating high confidence in the data. It eliminates the need for significant time and resources to manually review historical information. For instance, processing and logging 85 km of core data would require an experienced geologist or engineer several months to complete, whereas machine learning can compile this data within hours, free from the subjectivity and errors associated with logging fatigue.

While the detailed approach works very well, there are several areas that will benefit from additional development to maximise its performance and utility:

- Improvements on fracture detection and segmentation: Although the computer vision models are highly effective in identifying core discing, it's essential to acknowledge that no model is perfect. Incorporating additional training data and refining the model architecture could further enhance its accuracy and robustness. Improving these models will likely have the most significant impact on reducing false positive detections.
- Better refinement of discing regions: Currently, disc sizes, discing angles, and disc angle deviations are specified to define the discing regions. Further investigation of these parameters could lead to a more generalised and accurate model.
- Incorporation of roughness and curvature: Currently, the angle-based method uses fracture angles alone. In the future, it may be beneficial to expand it to a fracture analysis method that also incorporates joint roughness (by assessing the profile of the fracture mask against an inferred fracture ellipse) and fracture curvature. These additions could better describe and distinguish the discing regions.
- Further classification of region-based detections: Currently, only one category of broken regions is detected. Introducing various classes into the region-based detection model could identify additional information, such as disc thickness (crushed, very thin, thin, etc.), leading to a better understanding of discing intensity.

6 Conclusions

We have presented a computer vision-based workflow capable of detecting, classifying, and extracting discing regions from historical core photography. This method has been successfully applied to imagery from the SDGM, producing data that is highly comparable to the actual discing that exists in situ, and has been used to extract a deposit-wide dataset.

Additionally, the method successfully detected discing regions that strongly correlate with interpreted large-scale structures within the deposit, and it has the potential to identify un-modelled structures. The nature of this approach instils high confidence in the data for engineering teams, significantly reducing the time and resources needed to review historical information manually and minimising the potential for rework due to inaccurate logging.

It is anticipated that once the discussed improvements are explored and, where appropriate, implemented, the Sunrise Dam team will incorporate the data into various aspects of the mining operation, enhancing operational efficiency and decision-making.

Acknowledgement

We would like to thank AngloGold Ashanti for permission to publish this paper using their core imagery. We would also like to acknowledge Datarock for providing analysis-ready processed imagery and the compute resources required to implement the described methods.

References

An, L, Jin, C, Liu, D, Ding, C, & Dai, X 2017, 'Mechanism of core discing in the relaxation zone around an underground opening under high in situ stresses', *Bulletin of Engineering Geology and the Environment*, vol. 77, no. 3.

AngloGold Ashanti 2023, *Operational Profile 2023: Sunrise Dam*, viewed 24 June 2023, https://reports.anglogoldashanti.com/23/wp-content/uploads/2024/04/AGA-OP23-SUNRISE-DAM.pdf

Belov, NI & Ivanov, VI 1992, ‘Numerical modelling of the stressed state of the foot of a borehole with a study of drill core disking’, *Journal of Mining Science*, vol. 28, no. 6, pp. 500–503.

Corthésy, R & Leite, MH 2008, ‘A strain-softening numerical model of core discing and damage’, *International Journal of Rock Mechanics & Mining Sciences*, vol. 45, no. 3, pp. 329–350.

Dagasan, Y, Nguyen, H, & Grujic, M 2021, ‘Inferring geological features masked by artefacts in core photography using neural networks’, *AEGC 2021 - Australasian Exploration Geoscience Conference*, Brisbane.

Jaeger, JC & Cook, NGW 1963, ‘Pinching-off and disking of rocks’, *Journal of Geophysical Research*, vol. 68, no. 6, pp. 1759–1765.

Johnson, S, Dagasan, Y & Stryk, A 2023, ‘Extracting consistent geotechnical data from drill core imagery using Computer Vision at the Carrapateena deposit’, *Australasian Exploration Geoscience Conference*.

Jutzeler, M, Dagasan, Y, Carey, R & Ila'ava, M 2023, ‘Volcanic stratigraphy reconstruction based on machine-learning crystal size distribution’, *AGU23*, American Geophysical Union, New Orleans.

Jutzeler, M, Dagasan, Y & Carey, R 2022, ‘Machine-learning image analysis on phenocrysts to reconstruct lithostratigraphy in mineralised terrains: An example with dacites in the Mt Read Volcanics’, *AusIMM Annual Meeting,* Tullah.

Jutzeler, M, Dagasan, Y, Ila'ava, M & Carey, R 2021, ‘Quantification of crystal size distribution in volcanic rocks with machine-learning image analysis’, *AGU Fall Meeting Abstracts*, American Geophysical Union, New Orleans.

Li, J Wang, E, He, X & Zhang, L 2019, ‘The Stress conditions of rock core disking based on an energy analysis’, *Rock Mechanics and Rock Engineering*, vol. 52, pp. 465–470.

Lim, SS & Martin, CD 2010, ‘Core disking and its relationship with stress magnitude for Lac du Bonnet granite’, *International Journal of Rock Mechanics & Mining Sciences*, vol. 47, pp.254–264.

Lin, TY, Goyal, P, Girshick, R, He, K & Dollár, P 2017, ‘Focal loss for dense object detection’, *Proceedings of the IEEE International Conference on Computer Vision*, Institute of Electrical and Electronics Engineers, New Jersey, pp. 2980–2988.

Matsuki, K, Kaga, N, Yokoyama, T & Tsuda, N 2004, ‘Determination of three dimensional in situ stress from core discing based on analysis of principal tensile stress’, *International Journal of Rock Mechanics & Mining Sciences*, vol. 41, no.7, pp. 1167–1190.

Ohta, T 2001, ‘Core disking and 'rockburst' in soft tuffaceous rock masses at Iwate Tunnel’ *Quarterly Report of RTRI*, vol. 42, no.3, pp. 130–135.

Stacey, TR 1982, ‘Contribution to the mechanism of core discing’, *Journal of the South African Institute of Mining and Metallurgy*, vol. 82 no. 9, pp. 269–276.

Zheng, M-z, Li, S-j, Yao, Z, Zhang, A-d, Xu, D-p & Zhou, J-f 2020, ‘Core disking characteristics and mitigation approach by a novel developed drill bit in deep rocks’, *Journal of Central South University*, vol. 27,pp. 2822–2833.