INTEGRATING DRONE-BASED IMAGERY AND AI FOR GSI DETERMINATION AT HIDDEN VALLEY MINE

L. McKenzie[1], P. Paredes[2], Y. H. Rajcoomar[2] and R. Yang[2]

[1] Harmony Australasia, Brisbane, Australia [2] Stratum AI, Toronto, Canada

ABSTRACT

Determination of the Geological Strength Index (GSI) applicable to a rockmass is fundamental to slope stability assessment in open pit mining. Traditional visual assessment undertaken at the bench face can be constrained by access, safety, and time limitations. These constraints can lead to gaps in spatial coverage and hence poor estimation of the GSI in some areas. This paper presents a method that integrates drone-based imagery with machine learning–assisted interpretation to support consistent and spatially continuous GSI characterisation along pit walls. High-resolution, georeferenced bench-face imagery is acquired using frontal drone surveys, enabling detailed and continuous visual observation of exposed bench surfaces.

The workflow was applied at the Hidden Valley open pit gold mine in Papua New Guinea. The pit has complex geology and challenging structural characteristics. Drone-acquired imagery, supported by GSI assessment by a geotechnical practitioner, was used to train the machine learning model. The model subsequently assisted with and standardised GSI interpretation in the pit. The results showed improved spatial continuity, repeatability, and efficiency of GSI assessment compared to traditional approaches. They demonstrate the method can provide a practical framework for safer and more consistent geotechnical characterisation of rockmasses exposed in pit walls.

1 INTRODUCTION

Rock mass characterisation is a fundamental requirement in open pit geotechnical engineering for slope stability assessments, blast design, and operational decision-making. Hoek (1994) introduced the Geological Strength Index (GSI) to determine rock mass strength using a chart based on a visual assessment of rock mass structure and discontinuity surface conditions. GSI has subsequently been refined and extended (Hoek *et al*., 1998; Hoek and Marinos, 2000; Cai *et al*., 2004). Marinos *et al*. (2005) noted that, as the geological environment is inherently uncertain, the GSI is best expressed as a range rather than a single value.

As pits deepen and wall geometries become increasingly complex, the need for reliable and spatially continuous rock mass characterisation becomes progressively more important. However, the challenges of quantification are compounded by GSI assessment being subjective, relying on the experience and judgement of the geotechnical practitioner. And by restricted access to pit walls, safety constraints, and time pressures during active operations. Hence, traditional approaches can result in incomplete spatial coverage and inconsistent GSI interpretation across different benches.

Accessibility and time constraint issues can be addressed using drone-based photogrammetry for rapid acquisition of high-resolution photographs of the various rock masses exposed on a pit wall from safe distances. Combined with a machine learning–based tool, the photographs provide geotechnical practitioners with detailed knowledge of the deposit the data required to estimate the GSI. The tool does not replace geotechnical judgement but provides a valuable decision-support mechanism. On this basis, the GSI characterisation remains fully grounded in site-specific geological understanding and established engineering practice. To address this issue, a weakly supervised learning model has been developed that has provided visually validated predictions of GSI at Harmony's Hidden Valley Mine.

2 OPERATIONAL CONTEXT

2.1 MINE SETTING

The Hidden Valley Mine (HVM) is located at an elevation of 2,600 m within steep, densely forested mountainous terrain in Morobe Province, Papua New Guinea (PNG). It is on the northern flank of the Owen Stanley Range, approximately 15 km southwest of Wau. The site has annual rainfall of 2,800 mm and mean monthly temperatures ranging from 8°C to 22°C.

2.2 MINE GEOMETRY

The current pit is 1,400 m in length and 1,000 m wide, with a maximum wall height of 400 m. The final design has a maximum wall height of 465 m.

Bench heights are typically 18 m but are reduced to 9 m within the biotite-rich, highly foliated metaschist, which is susceptible to degradation when exposed to surface water. Inter-ramp angles vary from 36° to 51° depending on the lithology, geological structure and rock mass strength. Pre-split blasting is currently not implemented.

2.3 GEOLOGY AND GEOTECHNICAL

The mine lies within the seismically active New Guinea Mobile Belt, a tectonically complex region undergoing intense and ongoing deformation. The deposit is situated within the Wau-Bulolo Graben, a back-arc pull-apart rift basin, at the southern extension of the Mobile Belt in an area referred to as the Owen Stanley Foreland Thrust Belt. The basement geology of the graben is the Jurassic to Cretaceous age Owen Stanley Metamorphic (OSM) Complex and mid-Miocene Morobe Granodiorite. The base metal carbonate gold-silver Hidden Valley-Kaveroi deposit comprises a vein-stockwork hosted within the Morobe Granodiorite.

Locally known as Kaindi Metamorphics, the OSM at the mine consists of grey-black and green-brown, variably carbonaceous, schistose, quartz-rich psammites and pelites that have undergone regional greenschist metamorphism. The OSM intruded by the Morobe Granodiorite, a heterogenous coarse-textured granodiorite sill that crops out across most of the Mining Lease. Localised, high-grade contact metamorphism on the intrusive contact with the Morobe Granodiorite, has resulted in a unique biotite-rich schist unit. The Morobe Granodiorite is cut by thin aplite dykes and numerous porphyry dykes of the Edie Creek Suite which intrude both the Kaindi Metamorphics and the Morobe Granodiorite.

The mine has a well-developed geotechnical model; domains being predominantly based on lithology and weathering. The metasediments have a low strength when highly to moderately weathered, increasing to medium to high strength when slightly weathered and fresh. Bedding within the metamorphic sequence is typically weak, and the dominant discontinuities are joints, shears and faults. The granodiorite has a medium strength when highly to moderately weathered, increasing to high to very high strength when it is slightly weathered and fresh. Discontinuities within the granodiorite are predominantly joints, although shears and faults are common.

The rock mass is extensively faulted, comprising two distinct structural zones and four main fault sets. Several large, reactivated fault structures are characterised by zones of rock fragments in an unconsolidated clay matrix, while most faults are discrete and planar with minor clay and rock flour infill. Depth of weathering is highly variable (approximately 10 m to more than 50 m), strongly influenced by lithology, alteration intensity, and structural fabric, with alteration assemblages dominated by chlorite–epidote–carbonate ± pyrite propylitic alteration affecting both the granodiorite host and the overlying metasedimentary units. Pore water pressures are typically low within the granodiorite, although elevated pressures can develop behind major low-permeability fault zones.

2.4 GEOTECHNICAL CHALLENGES

Ongoing and timely assessment of GSI at the mine presents several operational and geotechnical challenges. The local geology is highly heterogeneous, with rock mass characteristics varying significantly over short distances. These rapid changes are influenced by the proximity of wall exposures to major fault zones, dyke intrusions, and zones affected by high-grade contact metamorphism at the metamorphic–granodiorite contact. As a result, the assessments are time-intensive and require detailed field observation.

Access to newly exposed wall surfaces is typically constrained by operational requirements and safety restrictions, limiting opportunities to undertake direct GSI assessments soon after scaling of benches. These access limitations introduce delays and reduce the ability to capture wall conditions before weathering, degradation, or additional deformation occurs.

The geotechnical environment is also complicated by high annual rainfall, which interacts with structure, and rock-mass-strength, controlled, failure mechanisms. Elevated pore pressures, degradation and loss of strength in fault and weathered materials, and frequent rainfall-triggered instability contribute to a complex and highly variable slope stability setting. These factors collectively result in an inherently elevated potential for wall instability and underscore the need for a systematic adaptable GSI assessment methodology.

3 METHODOLOGY

Figure 1 illustrates the end-to-end workflow from data collection through evaluation. Drone imagery flows through processing and labelling stages before entering the modelling pipeline, with validation results from the AI model feeding back into the labelling phase. This iterative refinement loop allows geotechnical practitioners to correct model predictions and continuously improves performance as new data becomes available.

Figure 1: Methodology Flow

3.1 DRONE-BASED DATA ACQUISITION

Data acquisition at the mine occurred during a four-day field campaign using a DJI Matrice 4D UAV in accordance with site-specific safety protocols. Bench faces were photographed with images captured at stand-off distances of 40 to 80m. Overlap exceeded 80%. Effective spatial resolution was approximately 1 to 4 cm per pixel. The drone was flown parallel to the walls to preserve the true bench geometry.

3.2 DATA CLEANING AND PROCESSING

The survey captured approximately 5,944 georeferenced images. A systematic filtering process was then applied to remove images unsuitable for GSI assessment. Images taken at extreme distances, or at steep downward pitch angles, were excluded. Also excluded were frames exhibiting poor lighting conditions such as overexposure, deep shadow, or atmospheric haze. These images were identified through EXIF metadata and luminance histogram analysis. Also excluded were frames that included vegetation, rubble, waste piles, water pipes, and other elements irrelevant to GSI estimation.

Each image was projected onto the pit surface using ray casting against a photogrammetric point cloud to reduce spatial redundancy. DBSCAN clustering (Ester *et al*., 1996) was applied to group images capturing overlapping regions of the same wall segment.

Individual bench faces were isolated from full-frame imagery to mimic how a geotechnical practitioner assesses rock mass quality. Doing so used an automated boundary detection algorithm that combines intensity, gradient, and texture signals to identify the horizontal patterns of bench crests and toes. This cropping step ensures the model learns exclusively from exposed rock surfaces where structure and surface condition can be observed, rather than from surrounding operational infrastructure. After processing, 2485 images were deemed suitable for subsequent processing.

3.3 PRACTITIONER ANNOTATION AND VALIDATION PLATFORM

Training the machine learning model to classify GSI from bench imagery required obtaining Structure, Surface, and GSI ratings from geotechnical practitioners familiar with the site characteristics. A web-based annotation and validation platform was developed to do so. The platform enables a dataset of the ratings for the cropped bench images to be created in a systematic and normalised manner. The platform served two purposes depending on the project stage. (1) During initial data collection, multiple geotechnical practitioners labelled bench images, assigning ratings for Structure, Surface Condition, and overall GSI range based on the Hoek and Marinos (2000) GSI chart (Figure 2). This approach allowed the inter-observer variance to be calculated and identified images where practitioner opinions diverged. (2) Once a trained model was deployed, the platform then shifted to a validation role, as subsequent drone surveys generated new imagery that the model predicted automatically. Geotechnical practitioners then reviewed the model predictions, corrected mistakes, and flagged disagreements. These corrections fed back into subsequent training iterations of the model, which enabled it to be continuously refinement as geological conditions changed across the pit.

Not every area of the walls was suitable for assessment. Some images included multiple GSI domains, major faults, waste dump material, or were of insufficient quality. These images were marked by annotators as being inconclusive and were excluded from the dataset. This filtering process ensured the model learnt from appropriately labelled examples where GSI could be confidently assigned. This process reduced the dataset from 2485 images to 2321. This dataset was randomly split into 80% (1856) for training and 20% (465) for validating the model.

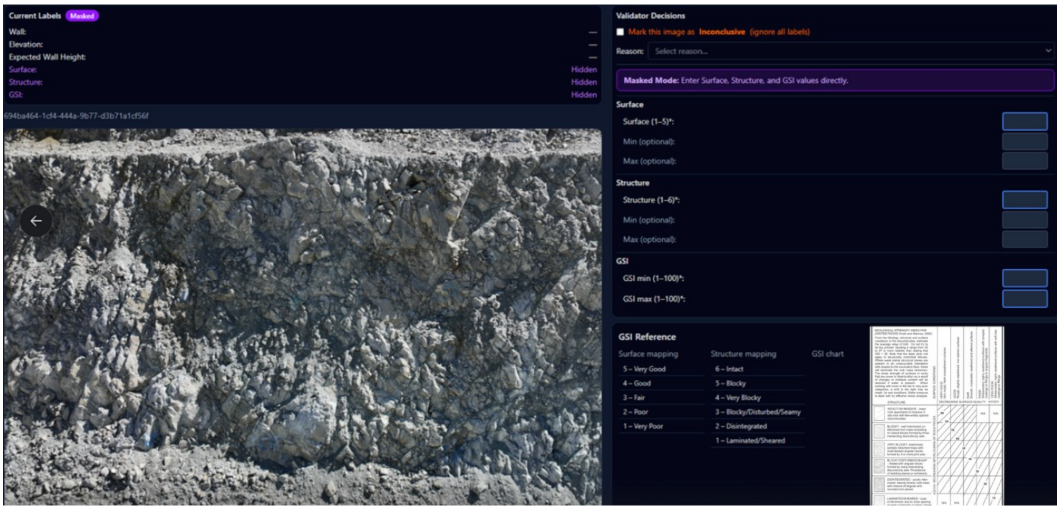

Figure 2: GSI Annotation Platform

3.4 MACHINE LEARNING FRAMEWORK - SEAM

When a geotechnical practitioner assesses a bench face, they scan the entire exposure, mentally weighing different regions such as fresh rock, weathered zones and visible joint sets before arriving at a single Structure or Surface rating. A useful predictive model should work in a similar way: not just producing a rating but also showing which parts of the bench the model relied on to reach that conclusion. In machine learning, this concept is known as attention (Vaswani *et al*., 2017): the model assigns a weight to each region of the image, the weight reflecting how important that region was to the final prediction. This process is analogous to a practitioner's eye being drawn to a prominent joint set or an area of extremely weathered rock. The process is challenging because training data consists only of whole-bench ratings; a conventional model trained on these image-rating pairs can predict bench-scale ratings but cannot reveal which regions contributed to the prediction. The alternative, labelling individual regions of each bench face, would be prohibitively time-consuming and does not address the issue of reconciling local assessments into an overall bench rating.

To address this issue, a weakly supervised learning approach (Zhou, 2018) was adopted, in which each bench image is divided into a grid of smaller areas. The model is trained with only the overall bench rating (e.g. "Structure: Very Blocky") and learns to identify which areas contributed most to that rating, thereby producing spatial attention maps that highlight influential regions and providing transparency without requiring additional labelling effort from practitioners.

To address the lack of region-level labels, the Stratum Embedding Attention Model (SEAM) was developed. This model is based on principles from Clustering-constrained Attention Multiple Instance Learning (CLAM) (Lu *et al*., 2020), originally designed for computational pathology where similar annotation constraints exist. The SEAM pipeline consists of three phases: (1) area extraction, where each bench image is divided into overlapping 64-pixel areas with a 32-pixel overlap; (2) feature encoding, where each area is processed by DINOv2 (Oquab *et al*., 2024), a self-supervised vision transformer selected for its strong generalisation to geological imagery without domain-specific fine-tuning; and (3) attention-based classification.

In the third phase, SEAM learns to weight each area by its relevance to the classification task rather than aggregating features through simple averaging. Areas showing clear fracture patterns or distinctive weathering receive higher attention scores. Featureless rock or imaging artefacts are down-weighted. The resulting weighted representation is passed to a classifier, which predicts the Structure or Surface rating. The training process further reinforces this procedure by encouraging separation between high- and low-attention areas: if the model assigns a "Very Blocky" rating, high-attention areas should exhibit blocky characteristics while low-attention areas should lack them, sharpening class understanding and improving generalisation to unseen bench images.

3.5 EXPLAINABILITY THROUGH ATTENTION MAPS

Attention scores are mapped back to their corresponding area coordinates to produce a spatial heatmap overlaid on the original bench image. Each area receives a score reflecting its contribution to the final prediction. The heatmap is rendered using a colour scale ranging from blue through green to yellow and red. Cooler tones indicate areas with lower attention and warmer tones highlight regions that most strongly influenced the model's classification (Figure 3). High-attention regions for Structure predictions typically coincide with visible joint sets or blocky fracturing, while Surface Condition attention concentrates on weathering, clay infill, or surface degradation. Geotechnical practitioners use the attention maps as a guide to support, not replace, practitioner judgement.

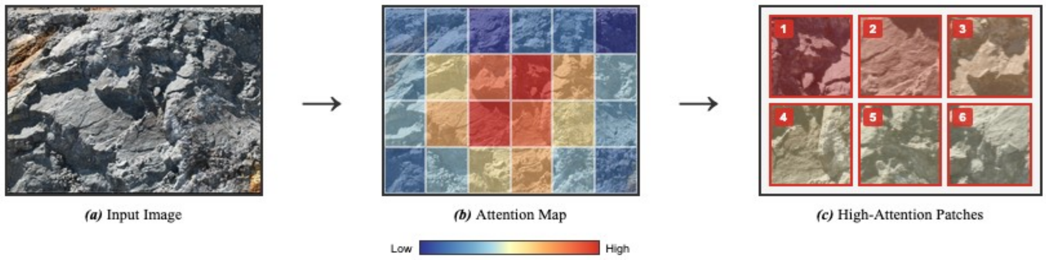

Figure 3: Interpretation of Attention Maps

4 RESULTS

4.1 LABELER VARIABILITY AND MODEL EXPECTATIONS

The nature of GSI assessment is inherently subjective. While the Hoek and Marinos (2000) chart provides a structured framework for assessing Structure and Surface Condition, the GSI rating depends on the interpretation of the geotechnical practitioner. Factors beyond the visible rock face influence this judgement, including familiarity with the local lithology, and understanding of the mine's structural geology. Two practitioners may assess the same bench exposure and assess different GSI ranges. This occurs, because each person weighs observable features differently based on their interpretation. This inter-observer variability has been demonstrated empirically by Yang and Elmo (2022) who found that different practitioners assessing the same core samples produced JCond89 ratings ranging from 15 to 25 and RQD values varying by up to 78 points. No study to date has systematically quantified inter-rater variability for GSI based on observation of rock masses exposed in pit walls. Given that variability is observed both between practitioners and between the methods used to derive the GSI, it can be concluded that GSI variability is an inherent characteristic of the metric rather than an artefact of poor practice.

To quantify this variability, 100 bench images were provided to two geotechnical practitioners, each with several years of experience. Both reviewers independently assigned Structure and Surface Condition ratings through the validation platform described previously. On average, the two practitioners' ratings differed by 0.54 class units for Structure and 0.62 for Surface Condition. Agreement occurred on only 20% of images for both axes. This variability resulted in the GSI differing by 8.4 points on average with a standard deviation of approximately 8 GSI points. This result indicates that any individual GSI estimate carries an inherent observer-dependent uncertainty of ±8 points. These results confirm that GSI classification is inherently subjective, and any automated system is being evaluated against a ground truth that carries significant inter-observer variability. This variability carries important implications for model evaluation. Since the training data reflects the subjective judgements of practitioners, the model inherits a degree of that same variability in its predictions. Rather than converging on a single objective value, the model's outputs falls within the range. Therefore, a model that predicts GSI within the variability observed among practitioners is reproducing the uncertainty inherent in the task itself.

4.2 MODEL PERFORMANCE

The performance of the model was evaluated against the manually annotated validation set. The Hoek and Marinos (2000) GSI chart divides rock mass quality into discrete cells, each spanning a range of approximately 15-20 GSI units. Our model predicts Structure and Surface Condition as discrete classes, which map directly to these chart cells, therefore a Structure prediction of “Very Blocky” and a Surface Condition of “Good” results in a GSI prediction of 45 to 65. A weighted rating is also predicted using the probability distribution across all classes. This approach captures the model's uncertainty: when the model is confident, the rating falls near the centre of the predicted cell. But when probabilities are distributed across adjacent classes, the rating shifts accordingly. The standard GSI range is reported based on the predicted class and the weighted GSI range derived from the full probability distribution.

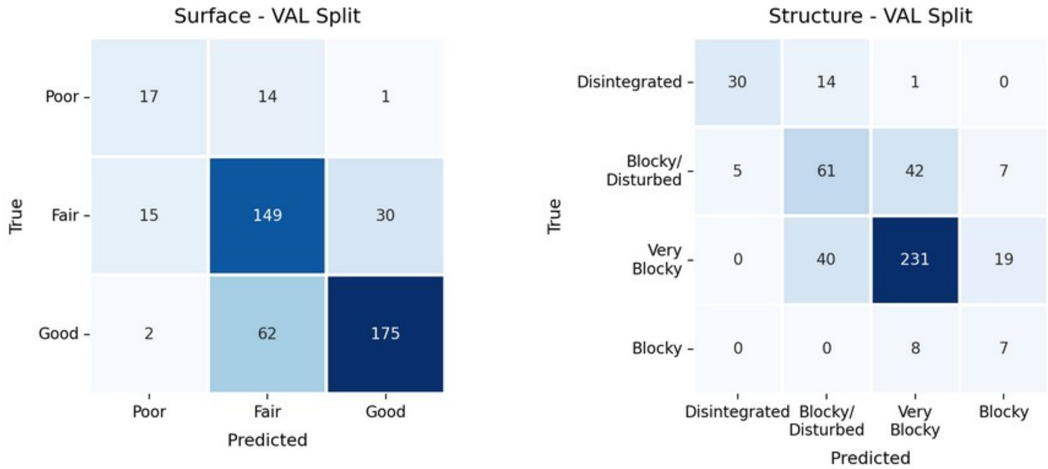

Figure 4 - Validation Confusion Matrix

Model performance was evaluated on 465 unseen bench images. The Structure model achieved an exact-class accuracy of 70.8%; the Surface model achieved 73.3%. Both models performed strongest on their respective majority classes ("Very Blocky" for Structure and "Good" for Surface). Less frequent classes showed lower accuracy due to limited training examples (Figure 4). Misclassifications occurred almost exclusively between adjacent classes, which is consistent with the ordinal nature of the rating scale where neighbouring classes share visual characteristics that challenge both human annotators and automated systems. Allowing a one-class tolerance, consistent with the inter-labeler variability (Section 4.1) where experienced practitioners routinely disagreed by one class on the same bench face, accuracy rises to 98.3% for Structure and 99.4% for Surface.

Beyond individual Structure and Surface classification, the combined GSI predictions were evaluated against practitioner annotations. The midpoint of the GSI range was taken to simplify calculations. When comparing predicted GSI to practitioner ratings, the model achieved an average absolute error of 3.9 GSI units, and 4.5 GSI units when using the probability-weighted GSI. The model matched the exact cell on the Hoek-Marinos chart in 55.5% of validation images, correctly predicting both Structure and Surface ratings. The agreement rate rises to 93% when applying an inter-observer tolerance of ±8 GSI points, derived from the standard deviation of GSI errors between the two practitioners in Section 4.1.

4.3 ATTENTION MAP REVIEW

The attention maps generated during inference were examined to verify that the model focused on geologically meaningful features. As described in Section 3.5, the attention mechanism assigns a score to each area thereby reflecting its contribution to the final prediction.

A Surface Condition attention map for a bench rated as "Good" is shown in Figure 5, where red areas indicate high attention, yellow areas medium attention, and areas with no overlay received low attention. The model isolates three distinct components: the upper bench portion, the brownish weathered granodiorite on the left, and darker fresh granodiorite along the bottom. Notably, attention is distributed across all exposed fresh rock surfaces rather than focusing solely on weathered material, while rubble and blast debris at the base received low attention.

A Structure attention map for a bench rated as "Very Blocky" is presented in Figure 6. The model assigns high attention to exposed fresh rock whereas water pipes and the toe of the bench are ignored, thereby focusing on visible joints in the centre and bottom left of the image. The model was never shown labelled discontinuity data nor programmed to explicitly seek fractures. Through training, it learnt that the most reliable way to predict Structure rating is by focusing on discontinuities visible in fresh rock exposures, from which fracture frequency and degree of interlocking can be inferred.

The attention maps in Figures 5 and 6 provide a window into the model's reasoning, challenging the common misconception that AI operates as an opaque black box unsuitable for safety-critical applications. Confidence in the prediction increases when attention aligns with features a geotechnical practitioner examines. The discrepancy signals an opportunity for review when attention falls on unexpected regions. Practitioners can flag such cases through the validation platform, steering the model toward geologically appropriate features in subsequent training iterations.

Figure 5: Surface Attention Map. Rating - Good

Figure 6: Structure Attention Map. Rating - Very Blocky

5 DISCUSSION

5.1 BENEFITS

The proposed workflow has been demonstrated to be capable of providing spatially continuous GSI characterisation for the Hidden Valley open pit. The workflow enabled faster turnaround of GSI assessments compared to conventional field methods. Consistency was enhanced by reducing inter-observer variability using AI-assisted interpretation while maintaining validation by a geotechnical practitioner. Doing so improves repeatability across benches.

5.2 LIMITATIONS

A natural constraint of training the model at a single mine site is that certain GSI classes may be poorly represented or entirely absent. At Hidden Valley, for example, the geological setting results in some Structure and Surface classes, such as “Laminated/Sheared” or “Very Good”, occurring infrequently or not at all. This class imbalance is inherent to any mine and means the model has limited exposure to the full range of rock mass conditions it may encounter elsewhere.

Additionally, the model does not account for the spatial arrangement of areas within an image; each area is assessed independently based on its visual content alone, without knowledge of where it sits relative to other areas. This limitation likely implies that the model cannot directly reason about joint spacing or the spatial relationship between fracture sets. This approach was deliberately selected for its simplicity as encoding positional information introduces training complexity for deep learning models to ensure they do not memorise location relationships in training images rather than learning to recognise geological features. This is particularly problematic in an active pit where bench faces change continuously with each blast, so positions in future imagery will not match those seen during training. Future work may explore spatially aware approaches once larger and more diverse datasets are available.

The GSI values determined using the discussed procedure apply to rocks exposed in the bench faces. These rocks have been damaged by blasting. The values are therefore conservative. The GSI values applicable to rocks beyond the damaged zone will typically be greater than those determined.

5.3 OPPORTUNTIES FOR FUTURE WORK

Future development is to extend the coverage to include additional benches to produce a pit-scale GSI model. This model could be further enhanced by integrating geotechnical drillhole data to provide vertical continuity in the rock mass characterisation. Doing so may highlight differences in the GSI values applicable to damaged and undamaged rock. The methodology also presents opportunities to move beyond two-dimensional image-based interpretation toward three-dimensional geotechnical characterisation. High-resolution photogrammetric point clouds and textured surface meshes generated from drone surveys can be used to identify and characterise geological structures, including major structures and joint sets, in three dimensions. Integrating AI-assisted interpretation with these 3D datasets would support more advanced structural mapping and improve understanding of rock mass fabric at bench and larger scales.

6 CONCLUSIONS

This study demonstrates that the integration of drone-based imagery with AI-assisted interpretation provides a robust and operationally viable framework for GSI characterisation in open-pit mining. High-resolution, georeferenced bench-face images captured through frontal drone surveys enables safe, continuous, and detailed observation of rock mass, overcoming the access and safety limitations inherent to traditional face mapping.

To ensure consistent reconstruction quality and reliable geotechnical interpretability, data acquisition should be conducted using automated drone flight plans, with controlled stand-off distances and image overlap, targeting a spatial resolution on the order of 1–2 cm per pixel. Automated flight execution reduces operator-induced variability and improves repeatability of the acquired datasets, particularly when surveys are conducted over multiple campaigns or by different operators.

A key finding of this project is the inherent subjectivity of GSI assessment, reflected in measurable inter-observer variability between geotechnical practitioners. Rather than treating this variability as noise to be eliminated, it was explicitly captured, quantified, and incorporated into the model. The observed inter-observer variance defines a realistic uncertainty envelope for GSI interpretation and establishes an appropriate performance benchmark for AI-assisted predictions.

Overall, the findings demonstrate that reliable and scalable GSI characterisation is achieved not through artificial intelligence alone, but through the combination of high-quality drone-acquired imagery data, transparent machine learning methods, and continuous practitioner validation. The proposed methodology provides a practical pathway toward safer, faster, and more consistent geotechnical characterisation, supporting improved ground control, blast design, and geotechnical risk management in active open-pit operations.

REFERENCES

Cai, M., Kaiser, P. K., Uno, H., Tasaka, Y., and Minami, M. (2004). Estimation of rock mass deformation modulus and strength of jointed hard rock masses using the GSI system. *International Journal of Rock Mechanics and Mining Sciences*, 41(1), 3-19.

Ester, M., Kriegel, H-P., Sander, J. & Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. *Proc. 2nd Int. Conf. on Knowledge Discovery and Data Mining (KDD-96)*, Portland, Oregon.

Hoek, E. (1994). Strength of rock and rock masses, *ISRM News Journal*, 2(2), pp. 4–16.

Hoek, E., Marinos, P. & Benissi, M. (1998). Applicability of the geological strength index (GSI) classification for very weak and sheared rock masses: the case of the Athens schist formation. Bulletin of Engineering Geology and the Environment, 57(2), 151–160.

Lu, M.Y., Williamson, D.F.K., Chen, T.Y., Chen, R.J., Barbieri, M. & Mahmood, F. (2020). Data efficient and weakly supervised computational pathology on whole slide images. arXiv preprint arXiv:2004.09666.

Marinos, P. & Hoek, E. (2000). GSI: a geologically friendly tool for rock mass strength estimation. *Proc. GeoEng 2000*, Melbourne, Australia, pp. 1422–1446.

Marinos, V., Marinos, P. & Hoek, E. (2005). The geological strength index: applications and limitations. *Bulletin of Engineering Geology and the Environment*, 64(1), 55–65.

Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.-Y., Li, S.-W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A. & Bojanowski, P. (2024). DINOv2: Learning robust visual features without supervision. Transactions on Machine Learning Research.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008.

Yang, B. & Elmo, D. 2022. Why engineers should not attempt to quantify GSI. *Geosciences*, 12(11), 417.

Zhou, Z.H. 2018. A brief introduction to weakly supervised learning. *National science review*, 5(1), pp.44-53.