Caving 2026 – A van As, D Cumming-Potvin & J Wesseloo (eds)

doi:10.36487/ACG_repo/2645_30

Artificial-intelligence-driven fragmentation assessment for safer block caving operations

Tungalag Ziinaa [a,*], Munkh-Erdene Enkhbayar [a], Myagmarsuren Sanaakhorol [a], Ankhbayar Boldbaatar [a], Khangerel Nasantogtogkh [a], Uuganbayar Ganbaatar [a], Bilguun Soyolbaatar [a], Bayasgalan Battulga [a]

[a] *Oyu Tolgoi LLC, Mongolia*

Abstract

In cave mining, drawpoint rock fragmentation influences operational safety, particularly with respect to inrush risk in block and panel caving operations. While fragmentation is linked to overall caving behaviour, the size distribution observed at drawpoints is often modified by material flow and progressive breakage within the caved column. Despite this, drawpoint fragmentation remains a key control on operational performance, influencing inrush hazards (both dry and wet), dilution entry, productivity and secondary blasting requirements. Adverse rock fragmentation – especially excessive fines and poorly distributed fragment sizes – increases the potential for inrush events, equipment damage and personnel exposure to hazardous conditions. Reliable assessment of drawpoint fragmentation is therefore essential for effective inrush risk mitigation and sustained production performance.

Despite its importance, fragmentation assessment in underground environments remains challenging. Visual inspection is widely used but is subjective and inconsistent, while existing commercial tools are unsuitable due to safety constraints and reliance on physical scale references. AI-based approaches offer potential but are limited by extensive labelling requirements. While foundation models can reduce this dependency, they demand substantial computational resources.

To address these challenges, we propose a computationally efficient and accurate fragmentation analysis system comprising three components: an offline image collection application, a rock detection model and a web application. The rock detection model is trained using Mask R-CNN and fine-tuned on 872 labelled images to accurately identify and categorise rock fragments. The model achieves 96% precision and processes images in 20–50 seconds, a 720-fold speed improvement compared to foundation models. Proof-of-concept trials across more than 200 drawpoints in Panel 0 at Oyu Tolgoi demonstrate the feasibility of this low computational effort, AI-driven workflow for assisting geotechnical engineers with rock size distribution analysis, drawpoint monitoring and operational decision-making in underground block caving operations.

Keywords: *rock fragmentation, AI, computer vision, rock detection, rock segmentation*

1 Introduction

Rock fragmentation at drawpoints plays an important role in the safety and performance of underground block caving operations. While fragmentation is linked to overall cave behaviour, the size distribution observed at drawpoints is often modified by material flow and breakage processes within the caved column. Drawpoint fragmentation directly influences material flow behaviour and inrush risk, particularly where excessive fine material accumulates within the muckpile. Reliable assessment of drawpoint fragmentation is therefore important for effective geotechnical monitoring and operational risk management.

Developing automated fragmentation assessment systems for underground environments presents challenges that differ significantly from those encountered in surface mining. Drawpoint images often

* Corresponding author. Email address: tungalagz@ot.mn

contain 500–1,000 densely packed rock fragments with irregular shapes and strong overlap, making accurate individual rock annotation extremely time consuming and difficult. In addition, existing commercial fragmentation analysis tools are primarily designed for surface blasting applications and rely on physical scale references, which are impractical and unsafe to deploy in underground drawpoints. These commercial systems are difficult to integrate with internal geotechnical monitoring systems, limiting their suitability for large-scale underground operations where seamless data exchange is essential.

This paper presents a computationally efficient AI-based framework for drawpoint fragmentation size distribution analysis developed specifically for block caving operations. The system is designed to operate reliably under challenging underground imaging conditions while supporting continuous geotechnical monitoring at scale. In addition to its lightweight and deployment-friendly architecture, the proposed framework is built for seamless integration with internal cave management systems, enabling automated data flow and operational insights.

2 Related work

Fragmentation size analysis has long been used in mining, initially through manual sieving and visual classification methods. With the development of digital imaging techniques, commercial tools such as WipFrag, Split-Desktop and FragScan were introduced to estimate fragment size distributions using edge detection and watershed-based algorithms (Palangio & Maerz 1995; Schleifer & Tessier 1996; Split Engineering 2025). While these systems are widely applied for blast fragmentation analysis in surface and open pit operations, they typically rely on physical scale references for calibration. This requirement limits their practical application in underground drawpoint environments, where placing scale objects can be unsafe or impractical.

More recently, deep-learning approaches have improved the ability to detect and segment rock fragments from images (Bamford et al. 2021). However, most of the existing studies have been conducted under surface mining conditions, where rock fragments are relatively well separated and imaging conditions are more controllable. Foundation models such as the Segment Anything Model (SAM) have significantly reduced the effort required for dataset annotation by enabling semi-automatic mask generation (Zhao et al. 2024) but often require substantial computational resources, which can limit their practicality for large-scale operational deployment in underground mining environments.

Despite these advances, there is a growing demand for computationally efficient methods that support continuous fragmentation monitoring in underground drawpoints for geotechnical assessment. This motivates the development of the framework presented in this study.

3 System design and architecture

The proposed fragmentation assessment system consists of three integrated components designed for underground mining conditions.

- Mobile application (Power Apps): enables structured image capture during drawpoint inspections and automatically associates each image with drawpoint metadata in offline underground environments. Once the geotechnical engineer reconnects to the network, the application automatically uploads the inspection images to a designated shared folder.
- Rock detection application: operating as a separate application, it automatically synchronises and retrieves data whenever new images are saved to the shared folder. Using an internet connection, it immediately processes these images using trained machine-learning models to detect individual rock fragments, calculate fragment areas and generate size distribution metrics.
- Web application: provides a user-facing interface for visualising results and monitoring fragmentation trends across drawpoints over time, allowing engineers to review and modify automated outputs based on engineering judgement.

Figure 1 illustrates the high-level system design and unified workflow of the proposed fragmentation assessment system.

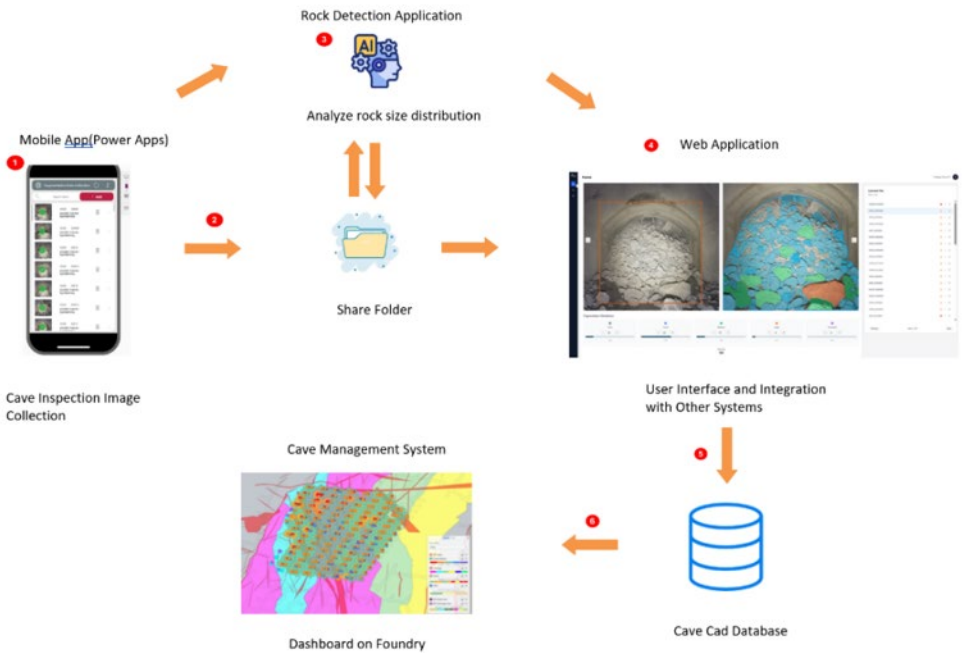

Figure 1 High-level system design

Together, these three components form an end-to-end workflow, as shown in Figure 1. While the initial cave inspection and image collection via Power Apps operate offline, the subsequent processes, including data synchronisation from the shared folder, individual rock segmentation and size categorisation, are executed once the device connects to an available underground network zone. The results of the rock detection model are displayed in the web application for user interaction.

3.1 Mobile application

The mobile application was developed using Microsoft PowerApps, a low-code platform within the Microsoft 365 ecosystem already used at Oyu Tolgoi. This approach eliminated the need for a separate authentication system, as engineers can log in using their existing corporate credentials, significantly reducing development time and maintenance overhead.

The application operates fully offline in underground environments, storing captured images locally and automatically synchronising them once the engineer returns to an area with network connectivity. In addition, the application automatically assigns image filenames based on the selected drawpoint, eliminating the previously manual and time-consuming process of renaming images to match drawpoint identifiers. Figure 2 shows the mobile application interface used for drawpoint image capture and data management.

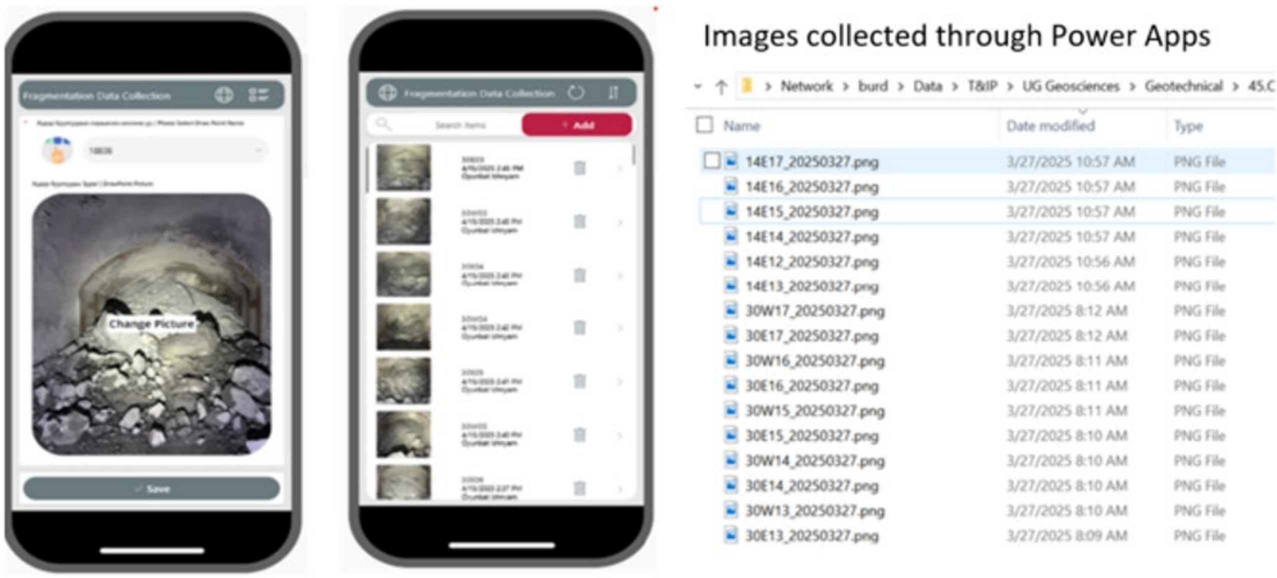

Figure 2 Mobile application (Power Apps)

3.2 Rock detection application

The rock detection application automatically synchronises images from the shared folder and processes them through the rock fragmentation algorithm. The generated fragmentation results are then integrated with the web application.

The rock fragmentation algorithm consists of three integrated computer vision models, each responsible for a specific stage of the drawpoint analysis workflow. As illustrated in Figure 3, the pipeline converts raw drawpoint imagery into scale-corrected, individual rock fragmentation measurements.

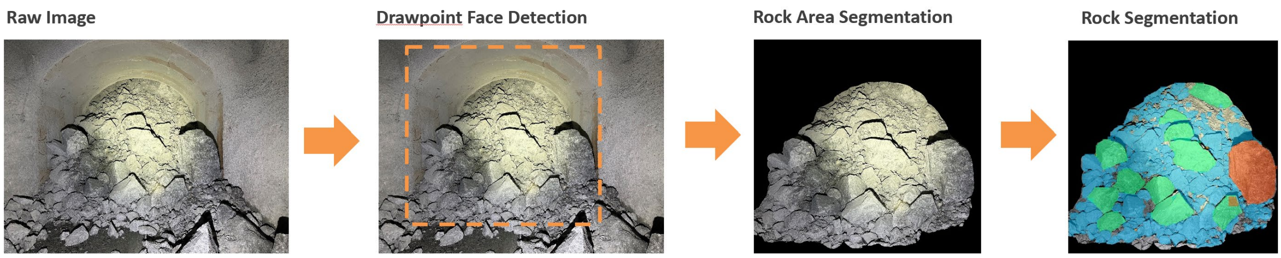

Figure 3 Rock fragmentation analysis pipeline for drawpoint images

The workflow proceeds sequentially:

- Drawpoint face detection identifies the drawpoint face to establish geometric calibration and determine the pixel-to-metre scale factor.
- Rock area segmentation: segments the active material region to restrict analysis to the operationally relevant area. While this system is developed as a customised solution for internal operational requirements, the underlying rock segmentation model is fully capable of detecting fragments independently without this preliminary bounding step. This module is implemented primarily to exclude tunnel walls and non-rock infrastructure within the mine environment. For broader applications or alternative deployment cases, this masking step can be skipped.
- Individual rock segmentation detects and segments individual rock fragments, calculates projected fragment areas and generates size distribution results.

Through this modular pipeline, raw underground images are transformed into quantitative fragmentation metrics suitable for operational analysis.

3.2.1 Drawpoint face detection model

Accurate detection of the drawpoint face is the foundation of the entire pipeline. The drawpoint face has a known fixed width of 3.9 m (defined by the legs of the steel arches), and once localised in an image, this dimension serves as the spatial anchor for converting all downstream pixel-level measurements into real-world units, regardless of camera distance, angle or field of view variation.

The model was developed using the SSDLite320 MobileNetV3 architecture, selected for its lightweight design and efficient CPU-based inference that make it suitable for deployment in underground environments. A dataset of 1,500 manually annotated images was prepared across two classes (drawpoint face and bund), captured under real underground operational conditions, including variable lighting, dust and partial occlusion. The model was trained for approximately 65 epochs, with an IoU (Intersection over Union) threshold of 0.5 used for evaluation. The model achieved AP@50 (Average Precision at an IoU threshold of 0.5) values close to 0.97–0.99 for both classes, while precision, recall and F1-score remained consistently above 0.96 across most epochs. Figure 4 shows the training and validation performance of the Drawpoint face detection model, including the training loss curve, validation AP@50 across both classes, validation precision/recall/F1 scores, and validation MAE (Mean Absolute Error).

Figure 4 Training and validation performance of the Drawpoint face detection model

These results demonstrate that the model generalises well and maintains reliable detection performance under varied underground conditions.

Figure 5 shows example detection outputs produced by the drawpoint face detection model, including predicted bounding boxes for the drawpoint face and bund classes along with their confidence scores.

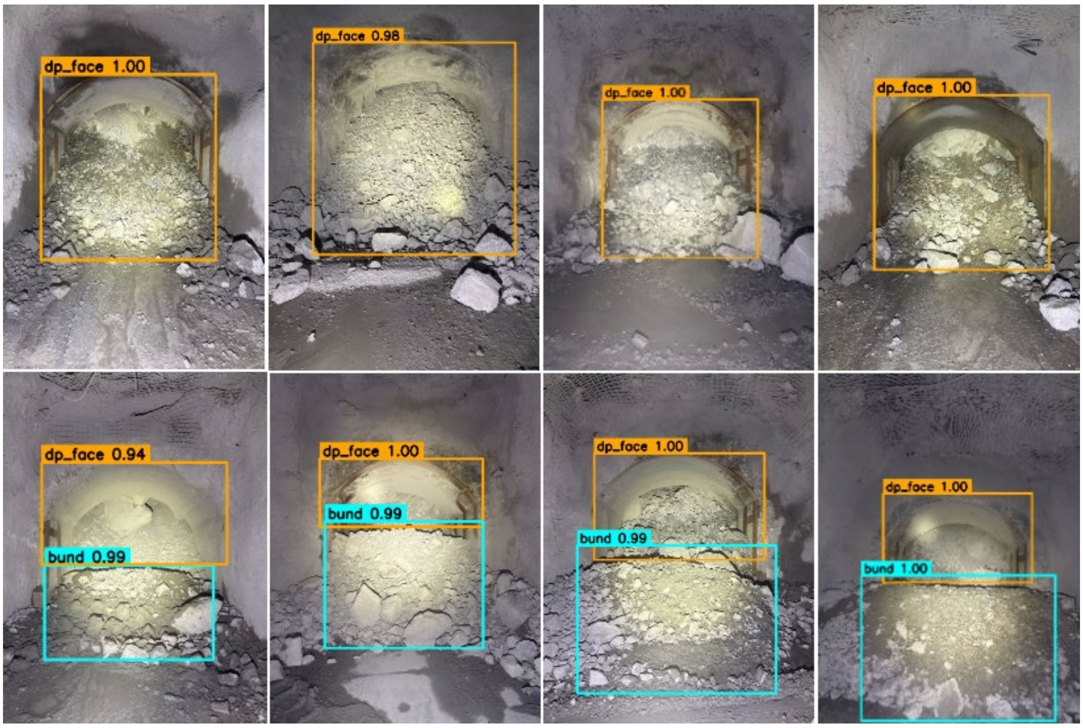

Figure 5 Example results of the Drawpoint face detection model

3.2.2 Rock Area Segmentation model

The model was implemented using the DeepLabV3+ MobileNetV3 architecture with a combined Cross Entropy and Dice Loss function, selected to balance pixel-level classification accuracy and boundary segmentation precision. A dataset of 600 manually annotated binary masks was used to train the model to distinguish fragmented rock regions from the surrounding background in drawpoint images. Training was performed for 100 epochs, during which the model demonstrated smooth and stable convergence. Figure 6 shows the training performance of the Rock Area Segmentation model.

Figure 6 Training performance of the Rock Area Segmentation model

Training loss decreased steadily from approximately 0.21 to below 0.03, while validation loss remained generally within the range of 0.05–0.08, indicating good generalisation performance without significant overfitting. Mean IoU increased rapidly during the early training phase (first 10–15 epochs) and stabilised around 0.93–0.94 in later epochs. These results indicate that the model achieved consistent segmentation accuracy and reliable delineation of the fragmented rock region across different drawpoint images, even under varying underground lighting and dust conditions. Figure 7 illustrates example segmentation outputs produced by the model, demonstrating accurate identification of the fragmented rock area.

Figure 7 Examples of Rock Area Segmentation results

3.2.3 Individual Rock Segmentation model

The Individual Rock Segmentation model performs instance-level segmentation of each rock fragment within the muckpile area identified by the previous model, enabling accurate computation of individual rock area, shape descriptors and size distribution.

3.2.3.1 Main challenges

The development of the Individual Rock Segmentation model involved overcoming two principal challenges.

First, preparing a high-quality training dataset for underground drawpoint conditions proved extremely difficult. Each image typically contains 500–1,000 fragmented rocks that are densely packed, heavily overlapped and highly irregular in shape. As shown in Figure 8, accurately manually tracing boundaries between fragments is labour intensive and technically demanding, particularly when rock surfaces visually resemble the surrounding cave walls and floor. This visual similarity often leads to inconsistent annotations, making dataset creation time consuming and complex. Figure 8 illustrates an example of manually labelled rock fragment images.

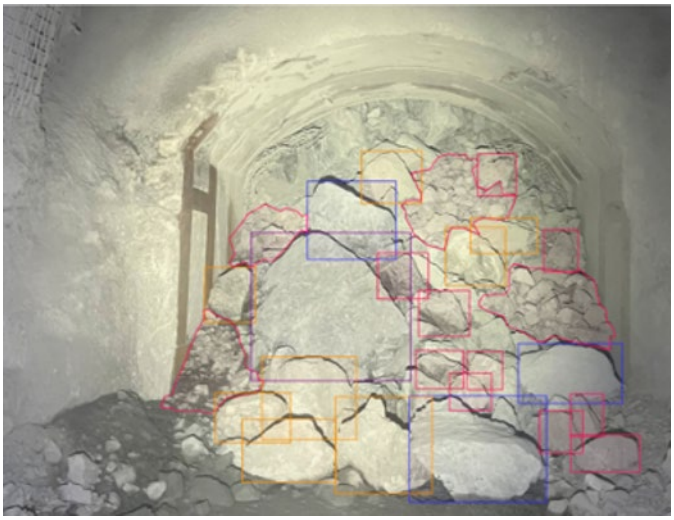

Figure 8 Example of manual rock fragment labelling

Second, although foundation models offer strong segmentation performance, they introduce substantial computational constraints. To reduce manual labelling requirements, we evaluated Meta's SAM2 base model, which demonstrated excellent rock boundary detection without the need for task-specific fine-tuning. Figure 9 illustrates the SAM2 labelling results together with the corresponding processing time. However, SAM2 is a prompt-based segmentation model, and its inference time is proportional to the number of rock fragments present in an image. Images with fewer and well-separated rocks could be processed in approximately 1–3 hours, whereas densely fragmented images required 4–12 hours on a server with 32 GB

RAM and 12 cores. Even on a Dell Precision 7960 workstation equipped with an RTX GPU, 128 GB RAM and 16 cores, heavily fragmented images still required 1–4 hours to process.

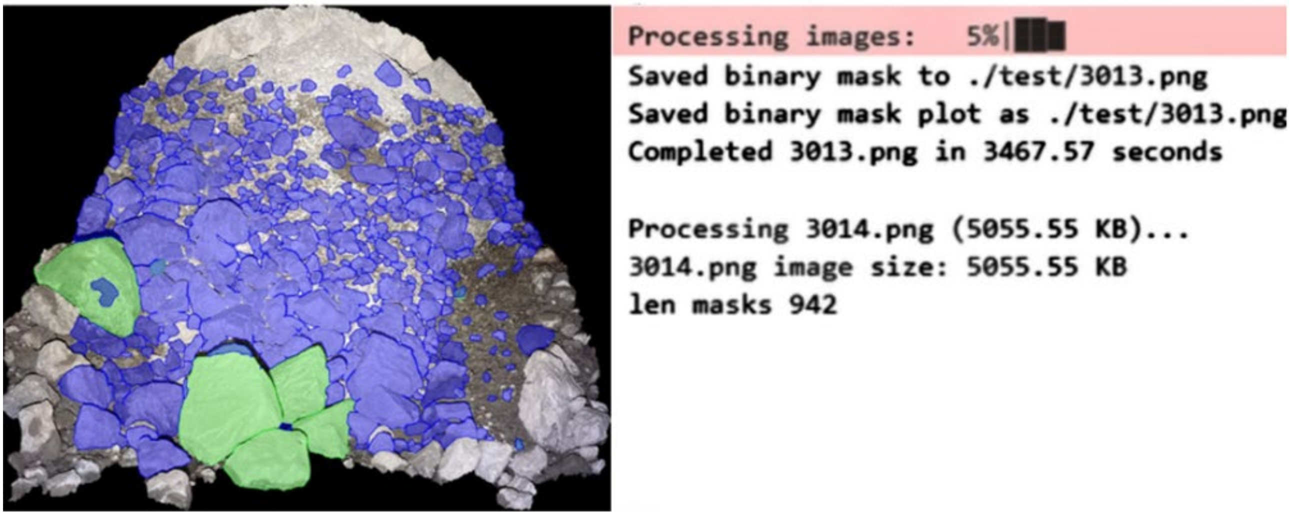

Figure 9 Example of SAM2 W rock segmentation output and processing time

Because SAM2 occasionally failed to capture very small fragments, we applied a secondary tiling-based inference step after the initial prompt-based masking. In this approach, each image was divided into overlapping high-resolution patches, allowing the model to focus on finer local details and improving small-rock boundary detection. However, this refinement step further increased computational effort, significantly extending overall processing time for images with high fragmentation. As a result, although the combined prompt-based SAM2 workflow and tiling strategy improved segmentation quality, the total computational load remains too high for large-scale dataset generation and impractical for operational deployment in typical mining environments.

3.2.3.2 Data preparation

To solve manual labelling issue, a custom AI-based rock labelling tool was developed using SAM2. Rather than labelling rocks manually, engineers use the tool to automatically generate segmentation masks for individual fragments, dramatically reducing annotation time while producing accurate polygon boundaries, even for irregular and partially occluded rocks. Figure 10 illustrates the semi-automated rock fragment labelling workflow using SAM2.

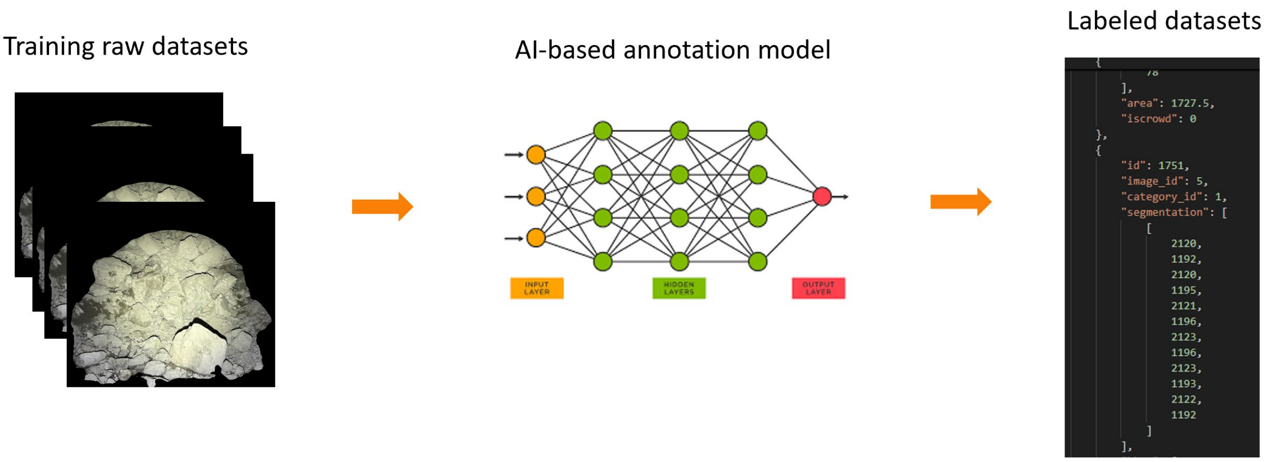

Figure 10 Semi-automated rock fragment labelling workflow using SAM2

A challenge is that SAM detects everything – walls, shadows, wet fines, background textures and other non-rock elements are all captured alongside the actual fragments. To address this, a second dedicated data-cleaning tool was developed to filter out these misdetected regions, enabling engineers to efficiently review and remove false annotations before finalising the dataset. Figure 11 illustrates the data-cleaning application interface.

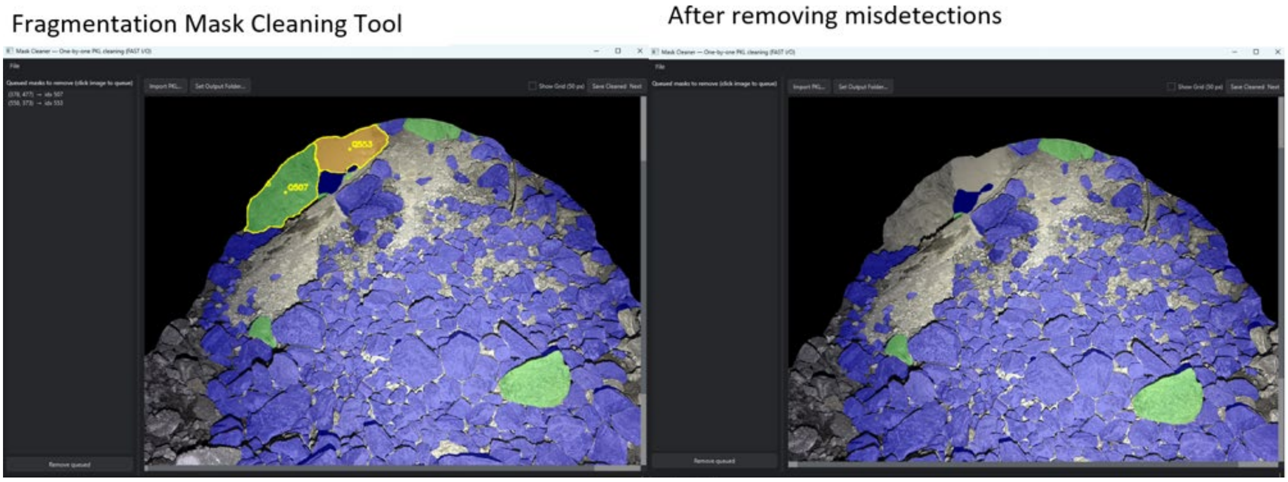

Figure 11 Example of cleaned background images

Together, these two tools (the SAM-based labeller and the data-cleaning interface) transformed what would have been an infeasible manual process into a practical, repeatable workflow. From 20,000 images, 1,000 high-quality samples were filtered for review. After automated labelling, data cleaning and quality control, 872 final images were selected for model training – a curated dataset that balances annotation quality with coverage of the diverse fragmentation conditions encountered across Panel 0.

The main contribution of this work is enabling cost-effective, computationally efficient preparation of high-quality training datasets, making advanced instance-segmentation models usable in real underground mining conditions.

3.2.3.3 Training an individual rock segmentation model

Several segmentation architectures were evaluated to assess robustness under dense fragment overlap and irregular boundary conditions. Mask R-CNN with a ResNet-50 FPN backbone demonstrated the most stable instance-level boundary delineation and was therefore selected for the proof-of-concept implementation.

The curated dataset of 872 images was formatted in COCO (Common Objects in Context) style with a single rock class and split into 654 training images (75%), 87 validation images (10%) and 131 test images (15%). Images were resized to 1,024 pixels to preserve fine fragment detail. The model's training performance is shown in Figure 12.

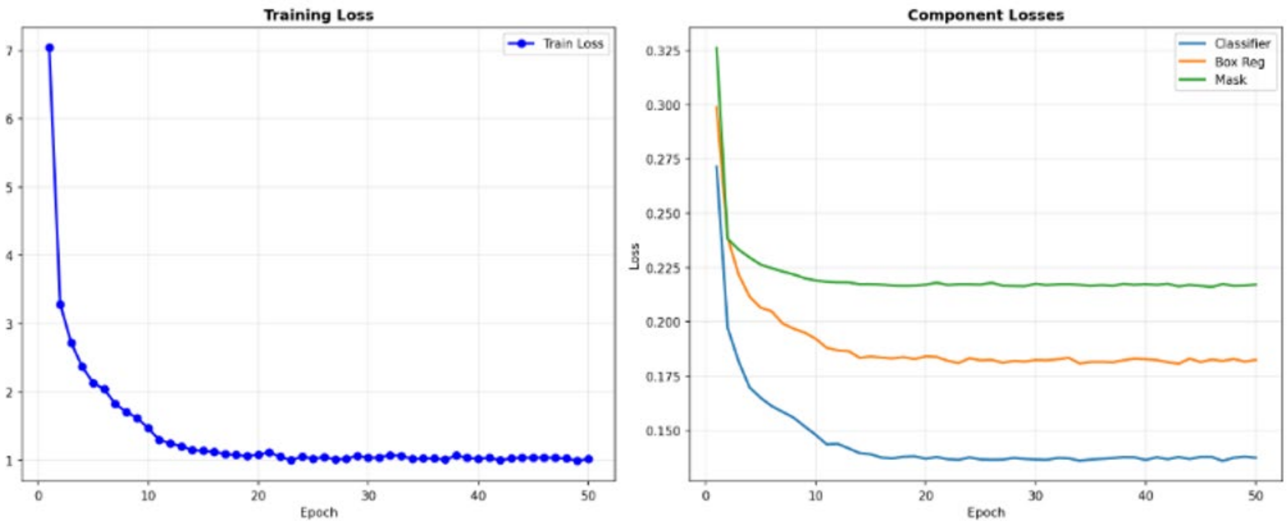

Figure 12 Training performance of the Individual Rock Segmentation model

Training was conducted for 50 epochs using SGD (stochastic gradient descent) optimisation with gradient clipping to ensure stability in dense segmentation scenarios. The model converged by approximately epoch 15 and remained stable through epoch 50, achieving 96% precision on the test set. These results confirm reliable instance-level detection performance across varied underground conditions. The individual rock segmentation model's results are shown in Figure 13.

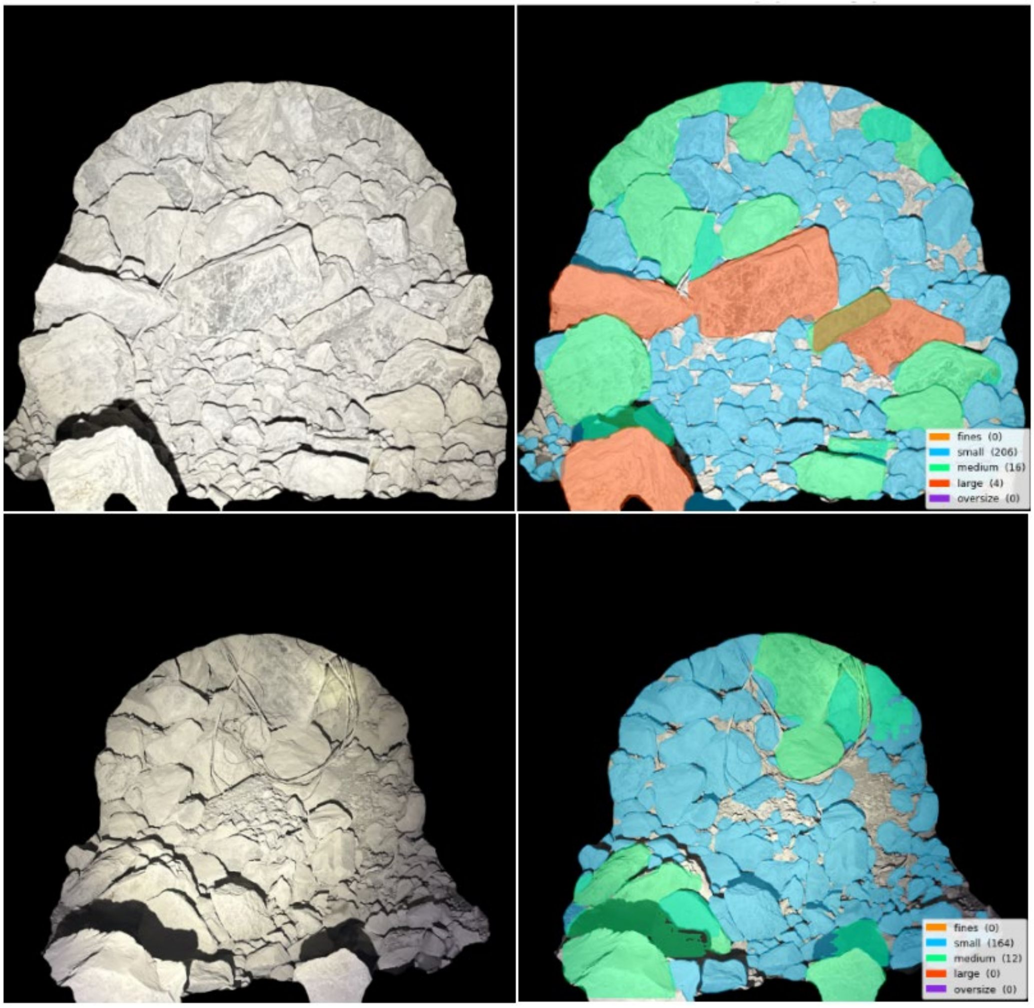

Figure 13 Example results of the individual rock segmentation model

After training the individual rock segmentation model, the predicted masks were used to compute the area of each detected rock fragment within the drawpoint image. The calculated areas were then converted to equivalent fragment sizes using the established scale reference and classified into five fragmentation categories: oversized fragments (>2.0 m), large fragments (1.0–2.0 m), medium fragments (0.5–1.0 m), small fragments (0.05–0.5 m) and fine material (<0.05 m). The areas of oversized, large, medium and small fragments were obtained directly from the segmentation outputs. Because fine material cannot be reliably segmented due to its extremely small size and dense distribution, the fine fraction was estimated indirectly. Specifically, the total visible fragmentation area was first calculated, and the fine area was then derived by subtracting the combined area of all detected fragments larger than 0.05 m. This approach enables consistent estimation of fines while leveraging the strengths of the instance segmentation model for larger fragments.

3.3 Web application

Once the rock detection application processes images and performs fragmentation size distribution analysis, the results are automatically transmitted to a web-based application for visualisation and monitoring of fragmentation trends across multiple drawpoints. The developed web application displays segmented rock size distribution results together with additional operational monitoring inputs, enabling engineers to review fragmentation conditions and track changes over time. Figure 14 illustrates the web application interface.

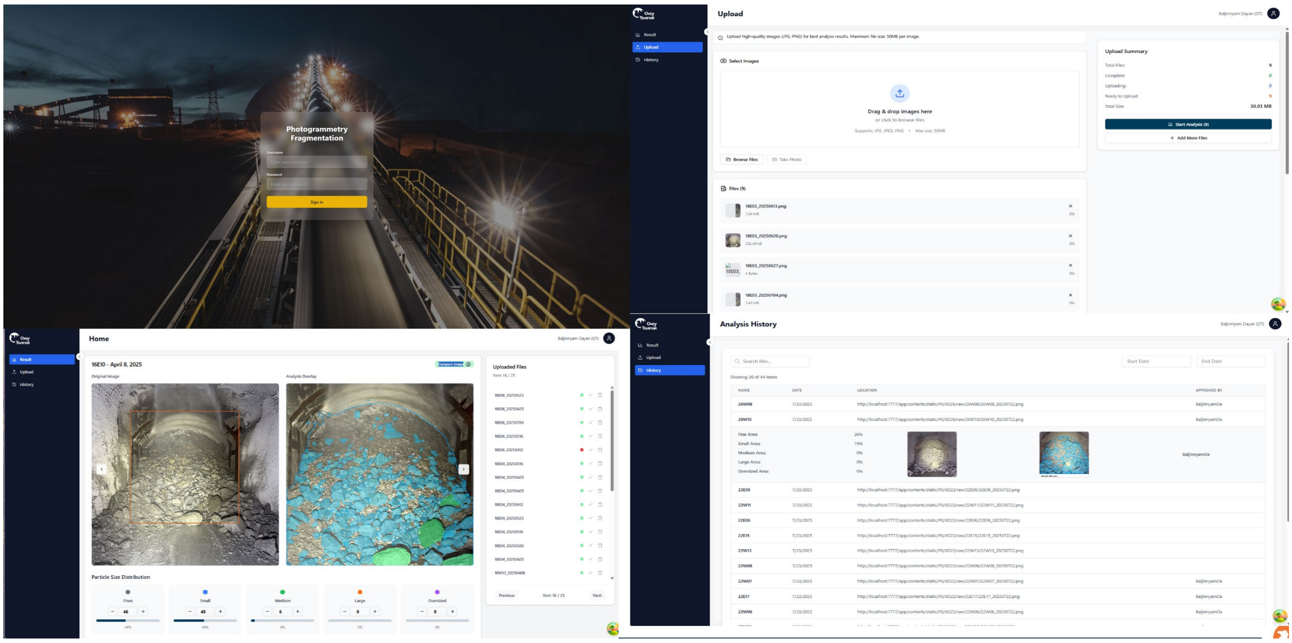

Figure 14 Web application interface

4 Limitations

A primary limitation of any image-based approach is its reliance on visible surface characteristics, which reflect only the muckpile face at the drawpoint rather than the internal volumetric composition of the caved column. Furthermore, characterising fine-grained material remains challenging because fine materials share highly similar visual characteristics at this scale. Consequently, differentiating specific micro-fractions – such as clay, silt, sand and gravel – falls beyond the intended scope of this framework, which focuses primarily on rock size distribution analysis.

From a hydrogeological standpoint, identifying specific clay and silt ratios is critical for predicting dry or wet inrush susceptibility. While this computer vision system provides the rapid, automated size distribution data necessary to flag anomalous fragmentation changes, it cannot serve as a standalone hazard predictor. To support safer operations, these surface indicators must be integrated with pre-existing spatial hazard maps and lithological data from the mine's geological block models to evaluate drawpoint inrush susceptibility safely.

5 Conclusion

This study developed and demonstrated an AI-driven multi-model pipeline tailored for automated rock size distribution analysis at Oyu Tolgoi's Panel 0. A key contribution of the study was addressing the challenge of preparing high-quality training datasets under dense fragment overlap conditions through the development of an AI-assisted annotation workflow. This workflow enabled the efficient generation of training data and the development of a computationally efficient individual rock segmentation model that achieved 96% precision on the test dataset, supporting scalable rock size distribution analysis in underground operational environments.

By processing full-resolution drawpoint images in approximately 20–50 seconds on standard CPU infrastructure, the framework addresses operational efficiency constraints associated with routine drawpoint monitoring. Integrated into a web application and mobile Power Apps platform, the system provides geotechnical engineers with a scalable and cost-effective tool for routine daily workflows without requiring specialised underground computing hardware.

While the inherent limitations of surface imagery mean that the framework reflects visible muckpile conditions rather than the complete internal composition of the caved material, its value lies in providing rapid, repeatable and objective rock size distribution information at scale. The framework is not intended to serve as a standalone inrush hazard prediction tool; rather, it provides an automated source of rock size distribution information that can support broader geotechnical assessment and monitoring workflows.

When integrated with existing geotechnical monitoring systems, spatial hazard mapping and geological datasets, the framework provides a practical source of continuous rock size distribution information to support underground operational decision-making and risk management. Future work will focus on broader operational validation, improving model robustness under lower-quality imaging conditions, refinement of rock size distribution characterisation methodologies, and deeper integration with underground operational and geotechnical monitoring workflows.

Use of artificial intelligence disclosure statement

In the preparation of this manuscript, the authors used Microsoft 365 Copilot (Enterprise) to improve the clarity and language of the text. The authors reviewed and edited the content where required and take full responsibility for said content.

References

Bamford, T, Esmaeili, K & Hamdi, A 2021, 'A deep learning approach for rock fragmentation analysis', *International Journal of Rock Mechanics and Mining Sciences*, vol. 145, no. 104670, https://doi.org/10.1016/j.ijrmms.2021.104839

Palangio, TC & Maerz, NH 1995, 'Case studies using the *WipFrag image analysis system*', *Proceedings of the 21st Annual Conference on Explosives and Blasting Technique*, International Society of Explosives Engineers, Cleveland, https://www.researchgate.net/publication/228739792_Case_studies_using_the_WipFrag_image_analysis_system

Schleifer, N & Tessier, B 1996, 'FRAGSCAN: a tool to measure fragmentation of blasted rock', *Proceedings of the 5th International Symposium on Rock Fragmentation by Blasting (Fragblast 5)*, Balkema, Rotterdam, pp. 81–88, https://doi.org/10.1201.

Split Engineering 2025, *Split-Desktop Fragmentation Analysis*, Hexagon Mining, viewed 8 March 2026, https://hexagon.com/products/hexagon-split

Zhao, J, Li, D & Yu, Y 2024, 'Identification of rock fragments after blasting by using deep learning-based Segment Anything Model', *Minerals*, vol. 14, no. 7, 654, https://doi.org/10.3390/min14070654