DOI: https://doi.org/10.5592/CO/EUROENGEO.2024.238

PREDICTING ROCK TYPE AND QUALITY FROM MWD DATA IN EXPLORATORY DRILLHOLES - FOCUSING ON GEOLOGIC TRANSITION ZONES AND UNCERTAINTY ASSESSMENTS

TOM F. HANSEN [1,2], ZHONGQIANG LIU [1]

[1] *Norwegian Geotechnical Institute, Norway, tom.frode.hansen@ngi.no*

[2] *University of Oslo, Norway*

Abstract

Measure While Drilling (MWD) data, a high-resolution sensor dataset collected during rock tunnel excavation worldwide, is underutilised, mainly serving for geological visualisation. Recent studies have demonstrated that MWD-data can identify rock type and rock mass quality for data from short blasting holes. Currently, there are no reliable, efficient methods in tunnelling to accurately predict rock mass characteristics in advance, which limits planning of advance rock support, excavation design, and the logistical planning for reuse of excavated rock. Traditional exploratory methods, such as geophysics, core drilling, or subjectively interpreted hammer holes, suffer from being time-consuming, costly, and offering insufficient resolution. To overcome these limitations, we employed ensemble machine learning models on MWD sensor data from 24 m exploratory holes in infrastructure tunnels with diverse geology. This approach enables accurate predictions of rock type and rock mass quality, providing a planning horizon of several days. We provide confidence metrics for the predictions, enhancing decision support by identifying areas of uncertainty. Our approach, particularly focusing on the challenging geological transition zones, achieved balanced accuracies above 0.9 for rock type and 0.8 for rock quality (Q-class). This advancement significantly improves the planning and execution of rock tunnelling projects.

Keywords

Machine learning, MWD, rock mass characterisation, geological transition zones, uncertainty assessment.

1 Introduction

Measure While Drilling (MWD) data, a high-resolution sensor dataset collected during tunnel excavation, is currently underutilised, primarily serving for geological visualisation. Recent studies have demonstrated that MWD data can be translated into rock type and rock mass quality for short blasting holes (Hansen, Erharter, et al., 2024; Hansen, Liu, & Torresen, 2024). There are arguably no reliable, efficient methods in tunnelling to accurately predict rock mass characteristics several blasting rounds ahead, which does not limit the excavation process (Liu & Gan, 2023). Traditional exploratory methods, such as geophysics and core drilling, are time-consuming, costly, and typically offer insufficient resolution. Another common method involves drilling exploratory hammerholes and subjectively interpreting the penetration rate and colour of flushing water, which is subjective and provides no detailed information about rock type and rock mass quality. Existing methods exhibit high uncertainties, and the communication of these uncertainties is not well developed.

Accurately predicting rock type and rock mass quality over twenty meters in advance of tunnelling can optimise decisions about advance support, excavation design, and the reuse of excavated rocks, days before the rock mass is excavated. This capability is valuable for efficient tunnel construction and resource management. Previous attempts to predict from MWD-data have been hindered by the lack of a big dataset of high-quality drilling data, inadequate methods to capture the complexity of drilling data, and the computational demands of processing such data. To the author's knowledge, no studies have systematically evaluated the uncertainty of these predictions. For predicting rock type, several studies have employed machine learning algorithms with varying degrees of success (Kadkhodaie-Ilkhchi et al., 2010; Leung & Scheding, 2015; Silversides & Melkumyan, 2022). Despite their promising results, these studies are further limited by the narrow range of rock types considered, small datasets, and a focus on single holes, which restricts their applicability and generalizability. In the realm of predicting rock mass quality, recent research has explored the relationship between MWD data and rock mass metrics that describe overall stability or soundness, primarily through rock mass classification approaches (Fernández et al., 2023; Galende-Hernández et al., 2018; van Eldert et al., 2020, 2021). However, these studies also grappled with issues of small datasets, subjective human intervention, and limited representativeness for large-scale tunnelling projects

This study aims to leverage Measure While Drilling (MWD) data to provide accurate predictions of rock type and rock mass quality from long exploratory drillholes through geological transition zones, using data from infrastructure tunnels with diverse geology in Norway.
We introduce confidence metrics for the predictions, enhancing decision support by identifying areas of uncertainty. Our study focuses explicitly on challenging geological transition zones. From existing studies we know that such ensemble models and MWD-data work well for predicting rock type and rock mass quality from short blasting holes (Hansen, Erharter, et al., 2024; Hansen, Liu, & Torresen, 2024). In this study we extend the planning horizon significantly and provide uncertainty assessments through zones of rapidly changing geology. This way we provide a more comprehensive decision support system.

The structure of this paper is organised as follows: Section 2 details the methodology, including the dataset description, experimental setup, and evaluation metrics. Section 3 presents the results, covering prediction outcomes, communicating uncertainty, and evaluating the precision-recall trade-off. Finally, Section 4 provides conclusions and outlook.

2 Methods

2.1 Dataset

The dataset includes MWD data (features) and rock mass classification (labels) for 4,408 blasting rounds from 15 geologically diverse hard rock tunnels across four Norwegian infrastructure projects: UDK, UNB, RV4, and E39. The construction and pre-processing of the dataset are described in detail in a separate paper (Hansen, Liu, & Torressen, 2024). Precambrian Gneisses, Permian Basalt and Granite, Permian Rhomb porphyry, Cambro-Silurian shales, limestone, and claystone are key rock types. Figure describes a geological profile of the Permian rocks in the Drammens tunnel. Prediction results and confidence metrics will specifically be evaluated in this tunnel's geologic transition zones.

Figure illustrates the data collection process conceptually for shorter drillholes. For every 15 m of tunnelling, the 24 m exploratory holes are drilled inside the profile and some holes with a five-degree inclination out from the tunnel profile. Pre-processed values from all drillholes within a 1 m tunnel section (approximately 5000 values) were used to compute mean, median, standard deviation, variance, skewness, and kurtosis for each MWD parameter, resulting in 48 MWD-feature values plus three geometrical parameters (overburden, tunnel width, and Jn-mult), totalling 51 feature values per sample. Mapped to the Q-systems'(reference) Q-class labels and rock type, resulting in 23,277 samples, enriching the dataset with feature variations to enhance model performance and generalizability. Q-values, grouped into stability classes from "Exceptionally good" (class A) to "Exceptionally poor" (class G), recommend specific measures for Q-values under 1.0, such as Reinforced Ribs of Shotcrete (RRS) and spiling bolts. An accurate binary prediction model identifying this threshold has significant practical and economic benefits.

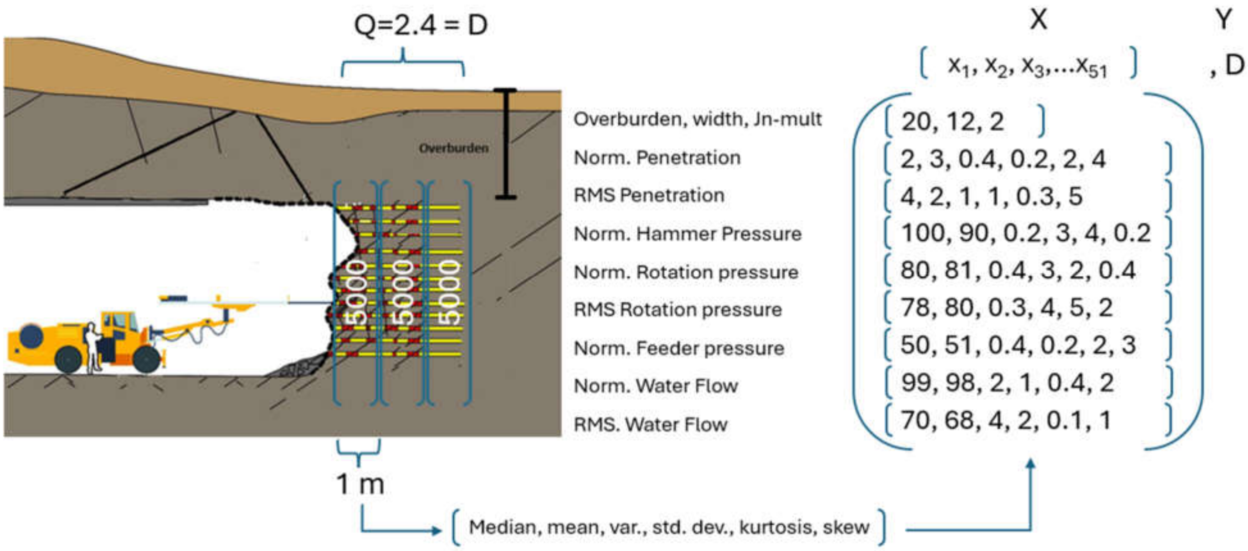

Figure 1. Illustrating the collection of MWD values for 1 m tunnel sections.

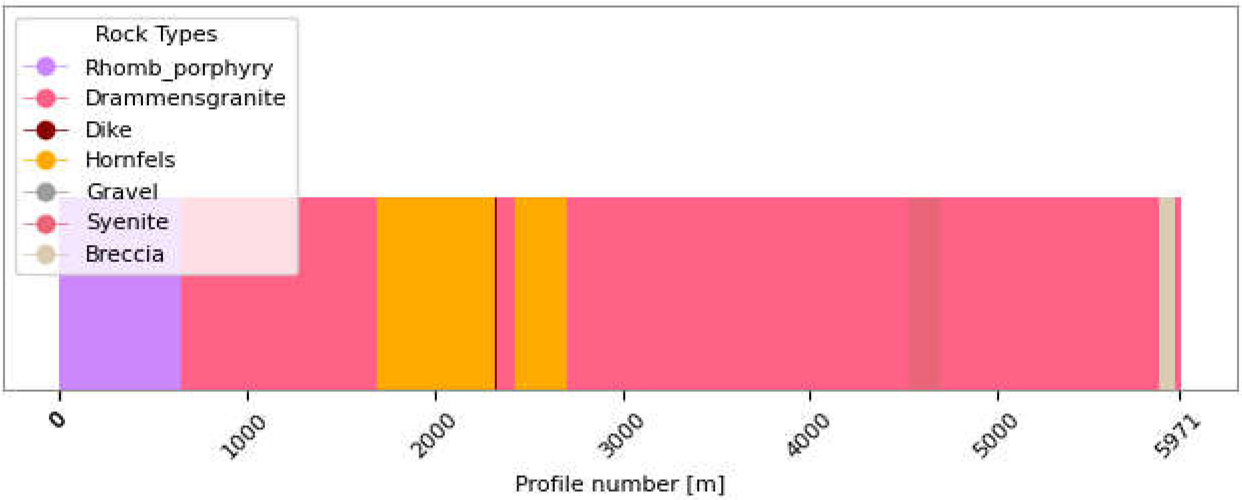

Figure 2. The geologic profile of the Permian rocks in the Drammens tunnel.

2.2 Experimental setup

Sample values of features and Q-class or rock type were trained and evaluated by several classifiers implemented in Scikit-learn (Random Forest, Extra Trees, KNN, logistic regression, MLP), LGBM, and XGBOOST (Pedregosa et al., 2011). Despite efforts to balance the dataset, medium Q-classes naturally had more samples. The SMOTE package (Chawla et al., 2002) was used to balance the training dataset before feeding it into the ML algorithms. Except for tree models, all tested models performed best when features were scaled to the [0,1] range using a MinMax scaler. To avoid data leakage, the pipeline functionality in Scikit-learn ensured scaling and balancing were based only on training data. The training and evaluation process followed the process outlined in Hansen et al. (Hansen, Liu, & Torresen, 2024) to ensure reproducibility. The dataset was split into 75% training and 25% test sets. Hyperparameters were tuned using Bayesian optimisation with the Optuna package (Akiba et al., 2019). Each new hyperparameter run used 5-fold cross-validation with random splits of the training set into training and validation sets. Final training used the entire set with optimised parameters before testing on the untouched test set and performing a 5-fold cross-validation on the full dataset. This process classified both Q-classes and rock types. The pipeline of processes used to train both target labels is illustrated in Figure .

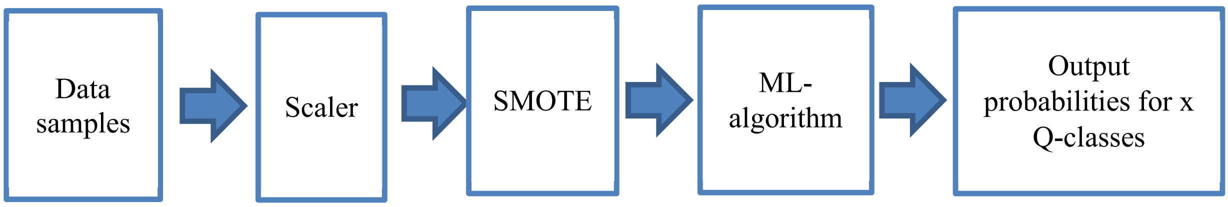

Figure 3. Defining the pipeline for the training process used to classify both target labels.

2.3 Metrics

To ensure robust analysis and mitigate misleading conclusions, we report several performance metrics (James et al., 2013). Each metric highlights different aspects of model performance, focusing on classifying rock masses into six categories.

Balanced Accuracy (Avg Recall Macro): Recall for a single class is defined as the ratio of true positives (correctly predicted instances of the class) to all actual instances of that class (sum of true positives and false negatives). Our primary metric, Balanced Accuracy, is the average recall across all classes, which is crucial for our unbalanced dataset. It ensures that larger classes do not disproportionately influence performance, vital for detecting weaker rock classes and preventing safety hazards.
Standard Accuracy: We include standard accuracy for comparison, measuring the proportion of correct predictions out of all predictions, providing an overall performance view.
Average Precision Macro: Avg. Macro Precision complements our recall focus by measuring the ratio of true positives to all instances identified as that class. This metric helps understand the false positive rate, which is critical in preventing the overestimation of rock strength.
F1 Score Macro: The F1 Macro Score combines recall and precision, offering a balanced view. It is essential for balancing recall and precision, acknowledging future insights that may highlight its importance.
Confusion Matrix: The confusion matrix provides an overview of precision and recall for each class, visualising class imbalances and prediction leakage. We present it in three versions: recall, precision, and non-normalized, showing correct and incorrect predictions.

2.4 Calibrating probabilities

To investigate model uncertainty in class predictions, we use the confidence information represented by the probabilities for each class. The class with the highest probability is predicted. A probability threshold can classify an instance as "unsure" if no class exceeds that threshold. In this study, we use a threshold of 0.8. However, existing probability scores often have flaws. Probabilities from several machine learning algorithms, such as tree-based models and KNN, tend to be biased, especially for imbalanced datasets (Niculescu-Mizil & Caruana, 2005; Zadrozny & Elkan, 2002). For tree-based models, this bias is due to the variance in the underlying base-trees, leading to less confident predictions. Conversely, Logistic Regression provides true probabilities as it directly optimises maximum likelihood based on cross-entropy loss. To address biased probabilities, one can calibrate them using isotonic and sigmoid calibration methods available in Scikit-learn. By comparing the true frequency of the positive label against its predicted probability, these methods transform probability values into interpretable statistical probabilities. For example, a probability of 0.85 for class B means that 85% of samples with similar feature values are correctly classified as class B.

We have constructed precision-recall graphs to visualise how we can adjust the probability threshold to vary the predictions of specific classes. To build a precision-recall graph for each class, we begin by splitting the dataset into training and testing sets, ensuring the testing set represents all classes. The classification model is trained using the training set, and the trained model then predicts probabilities for each class in the testing set. By varying the decision threshold from 0 to 1, we calculate precision and recall at each threshold level for each class. Precision is calculated as the ratio of true positives to the sum of true positives and false positives, while recall is the ratio of true positives to the sum of true positives and false negatives. These values are then plotted with recall on the x-axis and precision on the y-axis, using different line styles and colours to distinguish between classes. This results in a precision-recall curve for each class, visualising the trade-off between precision and recall across different thresholds.

3 Results

3.1 Predicting rock type and rock quality from long exploratory holes

The results for predicting Q-class and rock type from long holes closely aligned with existing prediction results from blastholes, albeit with approximately 4% lower scores for rock (Hansen, Erharter, et al., 2024; Hansen, Liu, & Torresen, 2024). This alignment was anticipated, given the comparability of data from long exploratory holes and shorter blast holes. We achieved the best performance by using a voting classifier in a pipeline that included Min-Max scaling and SMOTE for data balancing. For predicting rock type, the voting classifier combined KNN, Extra Trees, and LightGBM algorithms, while for predicting rock quality, it combined KNN, Extra Trees, and CatBoost algorithms. This ensemble approach excelled by enhancing recall and precision, improving balanced accuracy to 0.92, up from 0.87 with KNN alone, while maintaining the precision score of Extra Trees. Each algorithm was trained separately with its scaler in the voting classifier, and a majority vote determined the final classification. By leveraging the distinct advantages of KNN's non-parametric, lazy learning approach and the diverse strengths of both bagging (Extra Trees) and boosting (CatBoost for rock quality, LightGBM for rock type) methods, we were able to enhance the predictive performance for both rock type and rock quality.

In Figure , we present confusion matrices and metrics for a model trained to predict rock type and rock quality on the test set for all tunnels. The numbers are based on concatenating the results from the test sets for all splits in a 5-fold cross-validation run. The lower value for limestone seems to be caused by similarities in signature with the MWD signature from shale, "leaking" predictions to this rock type.

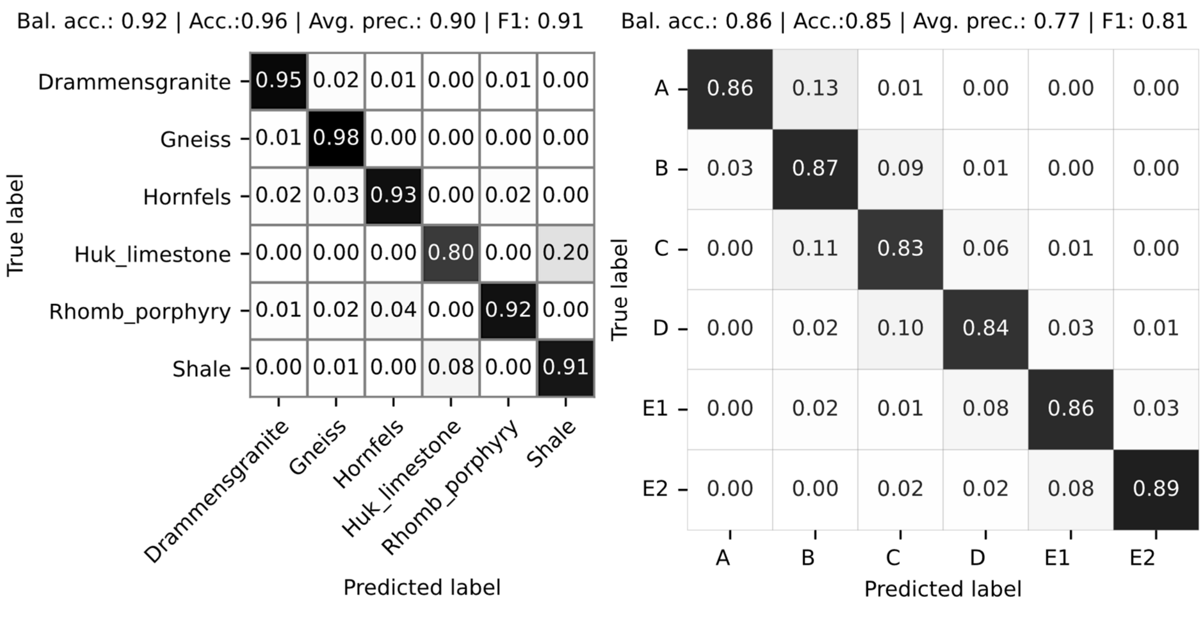

Figure 4. Confusion matrixes for predicting rock type and rock quality with a voting classifier.

3.2 Communicating uncertainty

Communicating uncertainty in predictions is crucial for a transparent AI-model. Transparency is considered one of the key factors in frameworks for ethical AI (Prem, 2023). When humans use ML models for decision-making, it is crucial to assess how confident the model is about its predictions. This assessment allows humans to determine how the prediction should be incorporated into their final decision-making process. By effectively conveying the uncertainty associated with ML predictions, we can ensure more informed, transparent, and responsible use of AI systems. This study emphasises assessing the uncertainty of predictions, expressed as probabilities, in geologic transition zones. Figure illustrates the calibration of probability values for predictions. In this figure, probabilities for the predicted class are plotted in a histogram of uncalibrated and calibrated probabilities using sigmoid and isotonic techniques. The calibrated values generally shift to the right, indicating increased confidence in the predicted class. Since the uncalibrated probabilities were biased towards extreme values for several samples, we chose sigmoid calibration for further analysis.

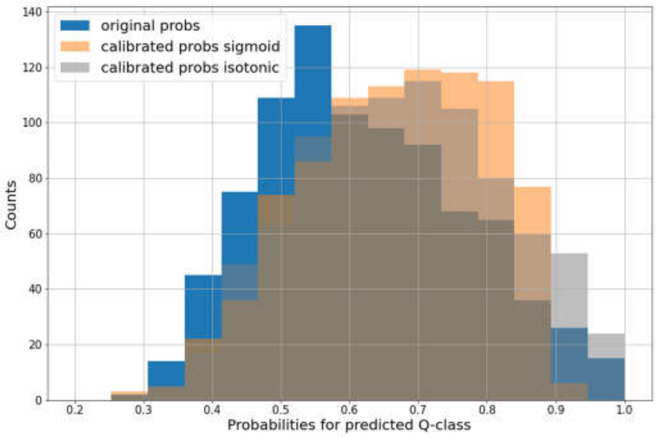

Figure 5. Comparison of ML probabilities with two types of calibrated probabilities.

To communicate the calibrated prediction-uncertainty in geologic transition zones, we selected two zones in the Drammen tunnel (see geologic profile in Figure). The plots show predicted rock types in five-meter sections along the exploratory holes drilled from the tunnel face. The y-scale indicates prediction probability. The tunnel face is on the left in the plot. Predictions below a 0.8 threshold are classified as "unsure". The pie charts below illustrate the distribution of probability values for each class for a specific prediction. In Figure , the correct prediction should be three Rhombporphyry followed by three Drammen Granite, indicating one incorrect prediction. For section 3, the model predicts Drammen Granite (60%) and Rhombporphyry (30%) but chooses incorrectly. This is consistent with the nearby rocks, suggesting that section 3 might be a mixture of Rhomb porphyry and Drammen granite.

In Figure , the true predictions should be three rounds of Drammen granite and three rounds of Hornfels, indicating one error. The incorrect prediction of Rhombporphyry in the transition zone with high probability is problematic. The reason for this is unknown but could be due to a mislabelled sample.

In Figure , the correct labels are three Hornfels and three Drammen granite, resulting in three incorrect and three uncertain predictions. The most significant error in this plot is the high probability prediction of Drammen granite in section 3, where the correct label was Hornfels. An incorrect prediction due to uncertainty is less concerning and reflects the changing geology. In all sections with incorrect predictions, the choice is between two logical rock types. Notably, there is no sudden high probability of Rhombporphyry within the Hornfels zone.

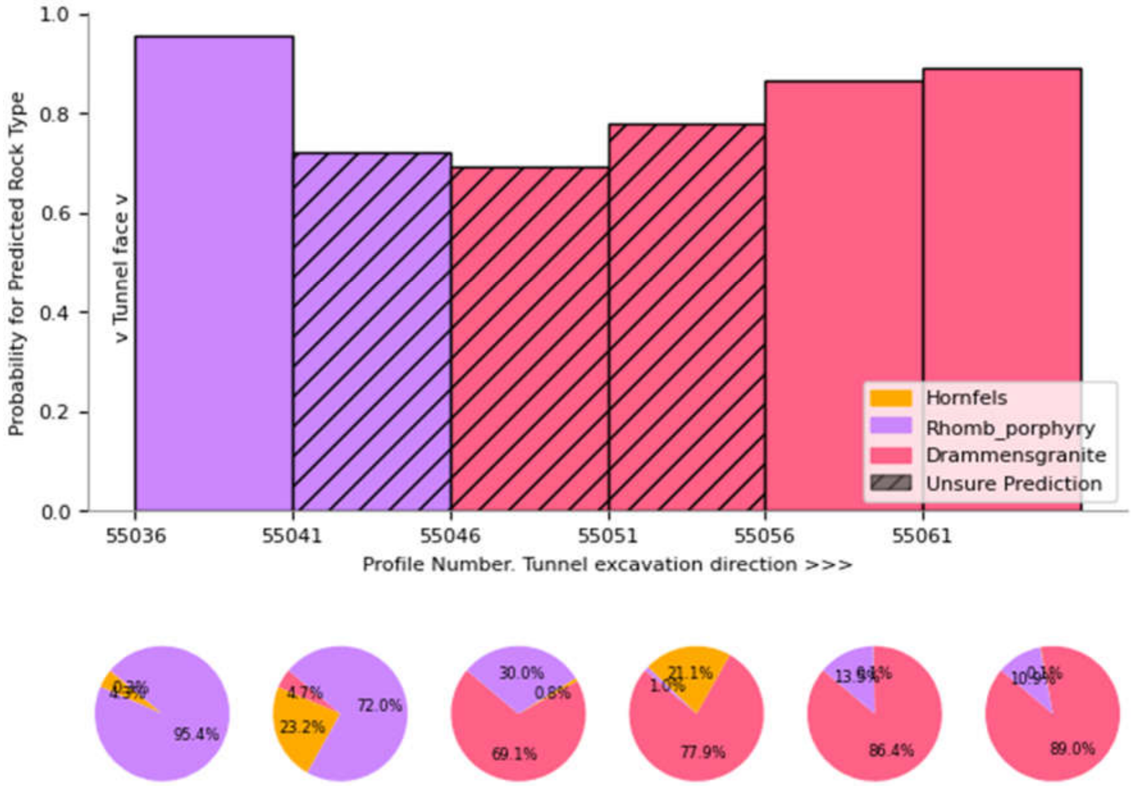

Figure 6. Prediction and uncertainty assessment in a geologic transition zone in Drammen tunnel, zone 1

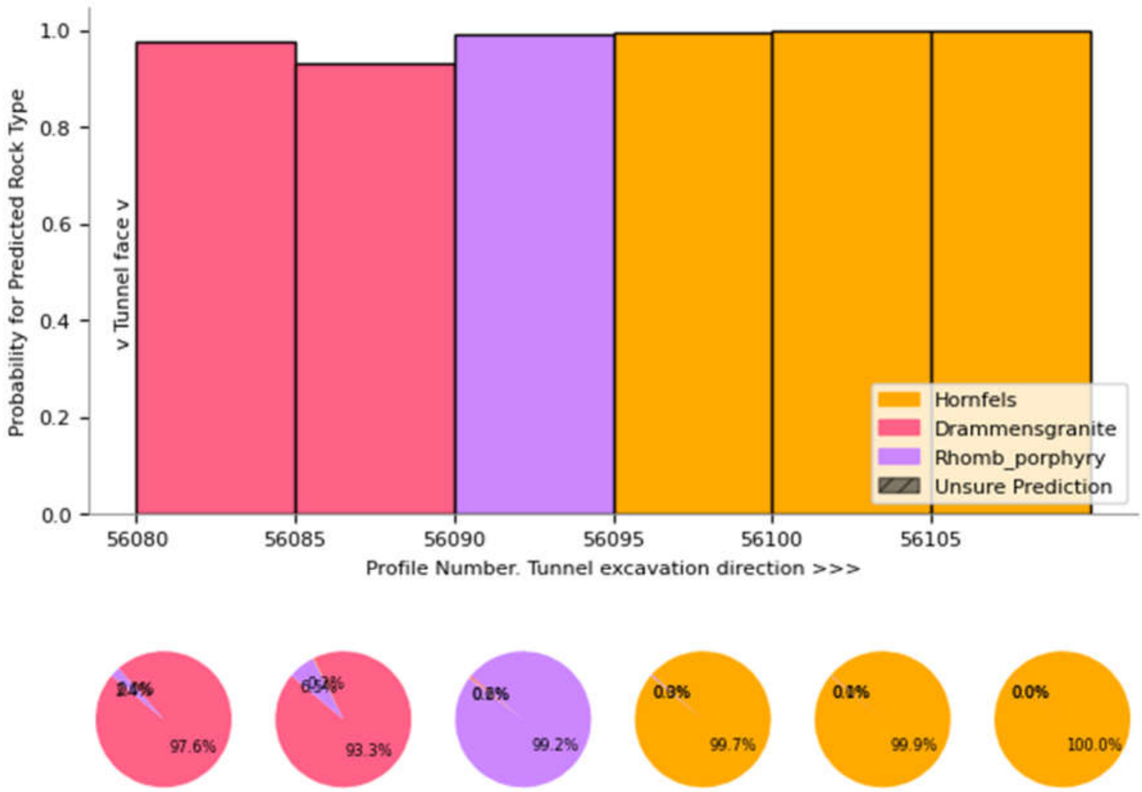

Figure 7. Prediction and uncertainty assessment in a geologic transition zone in Drammen tunnel, zone 2

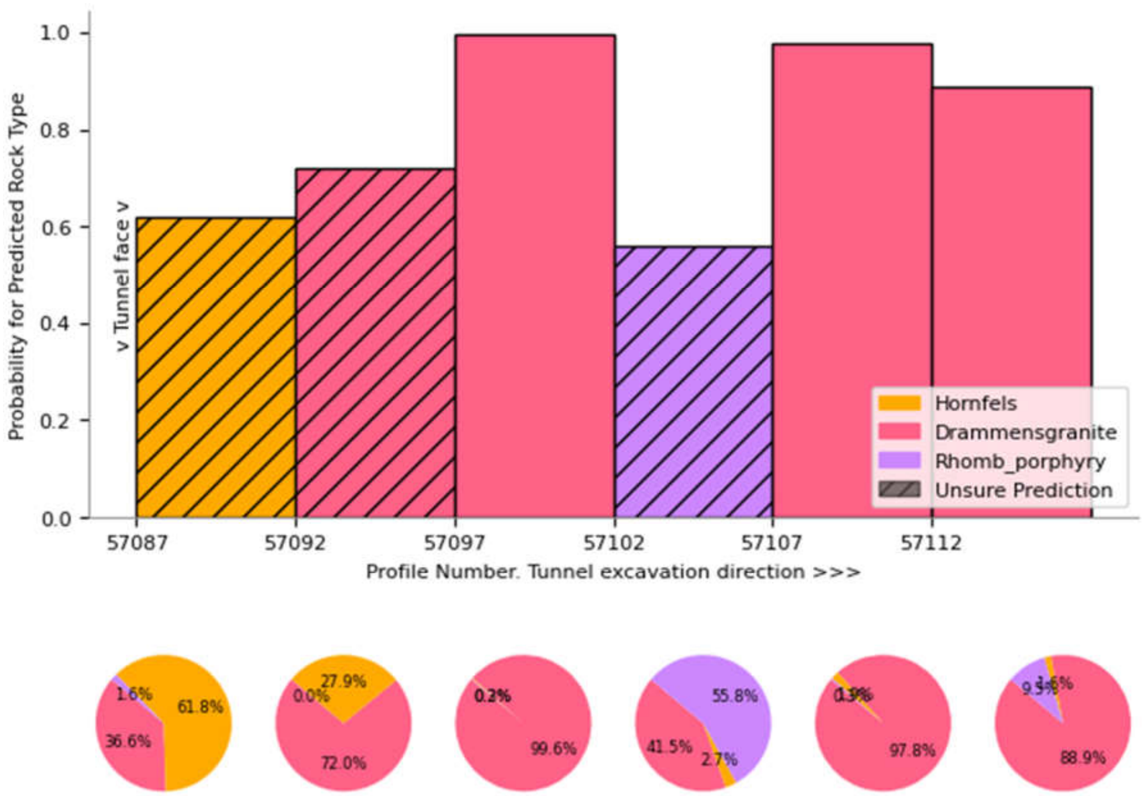

Figure 8. Prediction and uncertainty assessment in a geologic transition zone in Drammen tunnel, zone 3

Figure 12 illustrates prediction uncertainty conveyed through feature values and bar plots for rock mass classification. This visualisation aims to demonstrate alternative methods of communicating uncertainty in machine learning models, including providing feature values. Figure 12 presents plots of four random blasting rounds detailing feature values, true labels, predicted labels, and probabilities for each class within the Q-class labels A, B, C, D, E1, E2. In the figure, blasting rounds three and four are correctly predicted as class E2 and C with high confidence. Round one is incorrectly predicted as class E1 instead of C. This prediction should be considered uncertain, given the medium-sized probabilities for several classes. In round two, probabilities for classes B and C are observed. High probabilities for two neighbouring classes suggest a borderline sample between classes B and C. This is analogous to human physical mapping. For decision support, this should be conservatively interpreted as class B.
This leads to a suggested approach for handling prediction uncertainty in ML predictions:

- When only one class has a high probability, the prediction confidence is high.
- When two neighbouring classes have similar probabilities, the sample is likely a border case and should be conservatively interpreted as the class with the poorest ground.
- When several classes have significant probabilities (typically over 0.1), the classifier is uncertain, and the prediction should not be trusted.

The classifier can programmatically implement such probability-based confidence interpretation by adding simple target-based logic as the final prediction step. Thus, predictions strengthened with confidence information should significantly contribute to a data-driven decision support system for face engineers. In the next section, we will explore how to tune the precision-recall trade-off towards detecting certain classes, by varying the probability threshold for predictions.

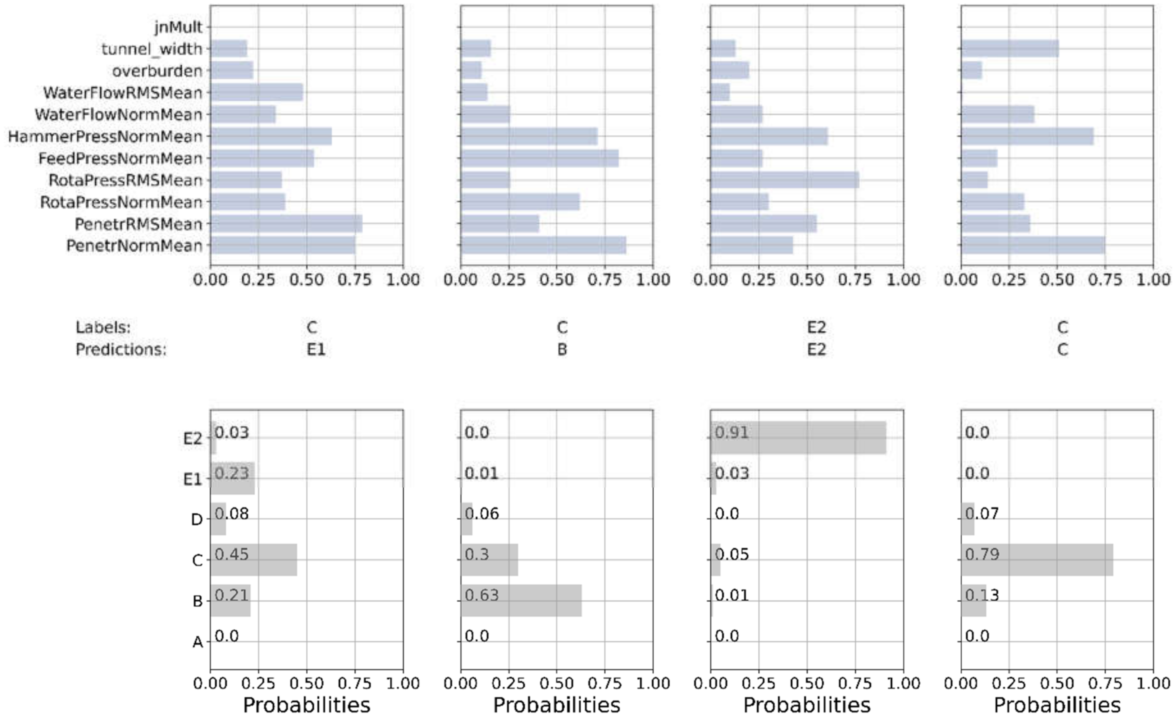

Figure 12. Calibrated probability plot of 4 Q-class predictions. Feature values are scaled to a 0-1 interval.

3.3 Tuning the prediction threshold

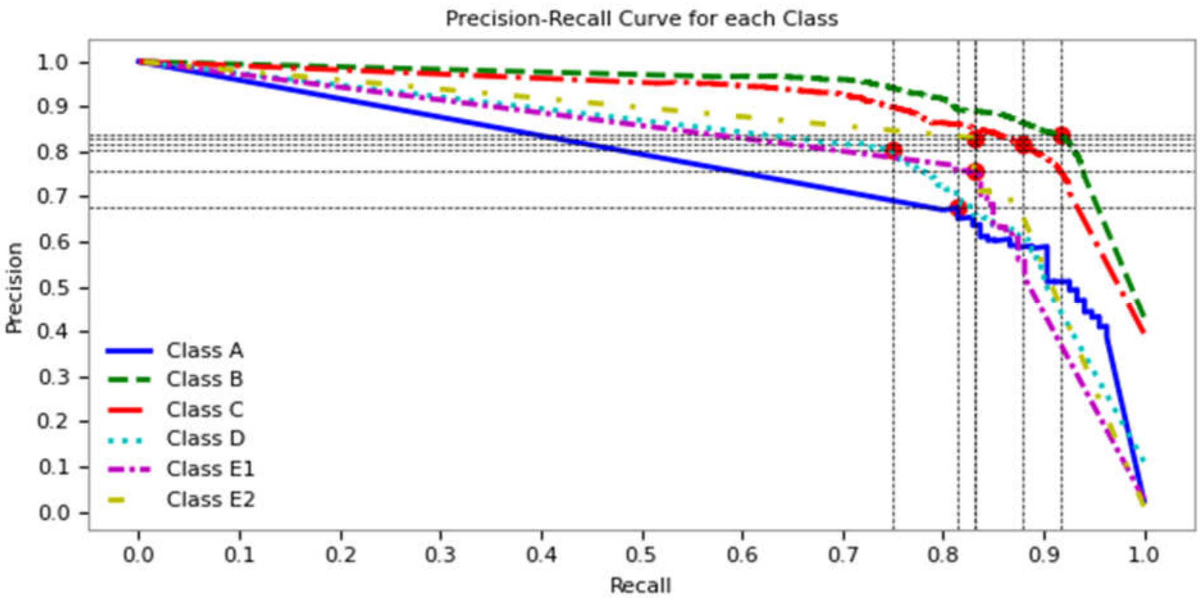

Figure 13. Precision - recall curves for rock quality classes

The plot in Figure 13 illustrates the precision-recall curves for six rock quality classes within our dataset, with each curve representing a different class. The red dots indicate the optimal trade-off points between precision and recall for each class, showing where we achieve the best balance of these metrics. For instance, examining Class E1 (represented by the magenta dashed line), we observe that adjusting the probability threshold for prediction can shift the balance between precision and recall. Specifically, lowering the threshold can increase recall, allowing us to detect more E1 samples. This adjustment would, however, come at the expense of precision, leading to a greater number of false positives. The implications of such an adjustment are significant in practical applications. In scenarios where detecting all instances of Class E1 is critical (e.g., identifying high-risk rock formations), it may be preferable to prioritise recall over precision. This approach ensures that most, if not all, E1 samples are identified, even if it means accepting a higher rate of incorrect predictions. Conversely, maintaining a higher precision may be more desirable in situations where the cost of false positives is high, even at the expense of missing some true E1 instances.

Balancing these metrics according to the specific requirements of the application can enhance the effectiveness and reliability of the predictive model, making it more suited to the practical needs of tunnelling and mining operations.

4 Conclusion and outlook

This study demonstrates the potential of using Measure While Drilling (MWD) data from long exploratory holes combined with ensemble machine learning models to predict rock type and rock mass quality in infrastructure tunnelling. The findings highlight the value of leveraging high-resolution sensor data to improve the planning and execution of rock tunnelling projects. Key findings:

- **Accurate Predictions**: Achieved balanced accuracies above 0.9 for rock type and 0.8 for rock quality (Q-class).
- **Extended Planning Horizon**: Enabled predictions several days in advance, providing critical information for logistical planning and support design.
- **Ensemble Models**: Ensemble machine learning models, including KNN, Extra Trees, and LightGBM/CatBoost, demonstrated the best performance.
- **Uncertainty Assessment**: Confidence metrics were introduced to enhance decision support, particularly in geological transition zones, ensuring more reliable and transparent predictions.
- **Tuning predictions**: Demonstrated how to tune predictions by varying the probability threshold illustrated by the consequences in the precision-recall trade-off and classifying to a class "unsure" below a certain threshold.
- **Practical Application**: Demonstrated practical benefits by predicting rock mass characteristics in advance, optimising excavation strategies, and improving resource management.

Future research should focus on refining these predictive models, particularly in the challenging geologic transition zones, to handle a broader range of geological conditions and extend their applicability to other tunnelling projects. Enhancing the dataset with more diverse geological samples and improving uncertainty assessment methods will be crucial. Integrating these predictive capabilities into real-time decision support systems for tunnelling projects can significantly enhance operational efficiency and safety.

Acknowledgements

The authors gratefully acknowledge Thorvald B. Wetlesen and Ivar Oppen from the tunnel software/hardware company Bever Control, which has helped in developing ideas, review, and facilitation of data from the clients Bane NOR, Statens Vegvesen, Nye Veier, and the contractor AF-Gruppen.

References

Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). *Optuna: A Next-generation Hyperparameter Optimization Framework*. arXiv. https://doi.org/10.48550/ARXIV.1907.10902

Chawla, N. V, Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, *16*, 321–357. https://doi.org/10.1613/jair.953

Fernández, A., Sanchidrián, J. A., Segarra, P., Gómez, S., Li, E., & Navarro, R. (2023). Rock mass structural recognition from drill monitoring technology in underground mining using discontinuity index and machine learning techniques. *International Journal of Mining Science and Technology*, *33*(5), 555–571. https://doi.org/10.1016/j.ijmst.2023.02.004

Galende-Hernández, M., Menéndez, M., Fuente, M. J., & Sainz-Palmero, G. I. (2018). Monitor-While-Drilling-based estimation of rock mass rating with computational intelligence: The case of tunnel excavation front. *Automation in Construction*, *93*(May), 325–338. https://doi.org/10.1016/j.autcon.2018.05.019

Hansen, T. F., Erharter, G. H., Liu, Z., & Torresen, J. (2024). A comparative study on machine learning approaches for rock mass classification using drilling data. *Http://Arxiv.Org/Abs/2403.10404*. http://arxiv.org/abs/2403.10404

Hansen, T. F., Liu, Z., & Torresen, J. (2024). Predicting rock type from MWD tunnel data using a reproducible ML-modelling process. *Tunnelling and Underground Space Technology*, *152*, 105843. https://doi.org/10.1016/j.tust.2024.105843

Hansen, T. F., Liu, Z., & Torressen, J. (2024). Building and analysing a labelled measure while drilling dataset from 15 hard rock tunnels in Norway. *Available at SSRN: Https://Ssrn.Com/Abstract=4729646 or Http://Dx.Doi.Org/10.2139/Ssrn.4729646*. https://doi.org/http://dx.doi.org/10.2139/ssrn.4729646

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An Introduction to Statistical Learning* (Vol. 103). Springer New York. https://doi.org/10.1007/978-1-4614-7138-7

Kadkhodaie-Ilkhchi, A., Monteiro, S. T., Ramos, F., & Hatherly, P. (2010). Rock recognition from MWD Data: A comparative study of boosting, neural networks, and fuzzy logic. *IEEE Geoscience and Remote Sensing Letters*, *7*(4), 680–684. https://doi.org/10.1109/LGRS.2010.2046312

Leung, R., & Scheding, S. (2015). Automated coal seam detection using a modulated specific energy measure in a monitor-while-drilling context. *International Journal of Rock Mechanics and Mining Sciences*, *75*, 196–209. https://doi.org/10.1016/j.ijrmms.2014.10.012

Liu, M., & Gan, Q. (2023). Applied research of comprehensive advance geological prediction in Daluoshan water diversion tunnel. *Scientific Reports*, *13*(1). https://doi.org/10.1038/s41598-023-36090-8

Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. *ICML 2005 - Proceedings of the 22nd International Conference on Machine Learning*, *1999*, 625–632. https://doi.org/10.1145/1102351.1102430

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-learn: Machine Learning in {P}ython. *Journal of Machine Learning Research*, *12*, 2825–2830.

Prem, E. (2023). From ethical AI frameworks to tools: a review of approaches. *AI and Ethics*, *3*(3), 699–716. https://doi.org/10.1007/s43681-023-00258-9

Silversides, K. L., & Melkumyan, A. (2022). Machine learning for classification of stratified geology from MWD data. *Ore Geology Reviews*, *142*, 104737. https://doi.org/10.1016/j.oregeorev.2022.104737

van Eldert, J., Funehag, J., Saiang, D., & Schunnesson, H. (2021). Rock support prediction based on measurement while drilling technology. *Bulletin of Engineering Geology and the Environment*, *80*(2), 1449–1465. https://doi.org/10.1007/s10064-020-01957-x

van Eldert, J., Schunnesson, H., Johansson, D., & Saiang, D. (2020). Application of Measurement While Drilling Technology to Predict Rock Mass Quality and Rock Support for Tunnelling. *Rock Mechanics and Rock Engineering*, *53*(3), 1349–1358. https://doi.org/10.1007/s00603-019-01979-2

Zadrozny, B., & Elkan, C. (2002). Transforming classifier scores into accurate multiclass probability estimates. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, *x*, 694–699. https://doi.org/10.1145/775047.775151