Machine-learning prediction of rock quality designation from AEM and geotechnical data in a Himalayan railway tunnel

Christensen C[1], Rodrigues P[1], Pacheco F[1], Pfaffhuber A[1], Skurdal G[1], Gangwar A[2]
[1] EMerald Geomodelling AS, Oslo, Norway
[2] Rail Vikas Nigam Limited, India

Abstract

Airborne electromagnetics (AEM), when combined with sparse borehole logs and supervised machine learning, can produce spatially continuous engineering-geology models along linear infrastructure corridors. Existing workflows for volumetric lithological classification are insufficient for existing engineering parameters that are continuous scalars rather than discrete classes.

We extend an existing volumetric classification workflow to numerical regression and apply it to Rock Quality Designation (RQD) along a 40 km segment of the Rishikesh–Karanprayag railway in the Indian Himalaya. An ensemble of artificial neural network regressors is trained on RQD logs from 65 boreholes using inverted SkyTEM 312 resistivity and spatial and terrain-derived support variables. Key adaptations from the classification workflow included a larger ensemble, longer early-stopping patience, and kernel-density sample reweighting to counter the strong zero-bias in RQD. Modelling RQD specifically also required post-processing clipping to physical bounds and window-averaging to reconcile the differing vertical resolution of AEM and borehole data.

Five-fold grouped cross-validation gives a mean absolute error of 11.1 percentage points; high-RQD intervals are underpredicted but flagged by elevated uncertainty. A subset experiment shows that 40–50% of boreholes recover the principal RQD structures. This study demonstrates AEM's utility as a route optimisation and borehole planning tool, especially in early-phase projects.

Introduction

In linear infrastructure projects, surface-based site investigations are generally slow and costly, especially in mountainous terrain. Airborne electromagnetic (AEM) surveys are increasingly used for engineering-geology characterisation. The resulting resistivity models, when combined with sparse geotechnical observations, can be turned into spatially continuous geological models using supervised machine-learning workflows. Development of such workflows began with interface detection (Lysdahl et al., 2022) and was later extended to volumetric classifiers for identifying different lithologies or mechanical behaviour groupings (Christensen et al., 2021). A weakness with the volumetric classifier workflow is that such models address only categorical outputs where uncertainty can only be expressed as a probability of a class. Many engineering parameters of interest to end-users are continuous scalar values, for which a numerical confidence interval is more desirable. To address that gap, we extend the volumetric workflow from classification to regression and apply it to Rock Quality Designation (RQD), a common measure of fracture intensity in tunnel design. Using data from a recent railway tunnel project in northern India, we demonstrate how the workflow was adapted for regression and the methodological considerations for RQD.

Site and data

The 40 km demonstration corridor lies along the first segment of the Rishikesh–Karanprayag rail line in Uttarakhand, India, a new broad-gauge route that crosses extremely steep Himalayan terrain. Seven planned tunnel sections account for the majority of the alignment. A SkyTEM 312 dual-moment time-domain electromagnetic system was flown along the alignment and surrounding terrain. Routine processing removed cultural-noise contamination from the soundings before we inverted the data to a smooth, laterally constrained 1D resistivity model. Sixty-five geotechnical boreholes were drilled along the 40 km of railway and reached depths of up to 390 m. RQD was logged in every borehole, though some gaps with no recovery exist in the records.

Methods

Like Lysdahl et al. (2022) and Christensen et al. (2021), we used an ensemble of predictors — in this case, artificial neural network (ANN) regressors implemented in Python using TensorFlow — where training targets are the borehole observations and electrical resistivity is the main support variable. A small set of additional supporting variables is included: horizontal coordinates (easting and northing), depth below surface, distance to the nearest borehole, and a measure of relative relief (height above channel base) derived from a Shuttle Radar Topography Mission (SRTM) terrain model. Each member of the ensemble used a multilayer perceptron trained independently on 80% of the boreholes, with 20% withheld as validation to monitor training and prevent overfitting. The mean and standard deviation across the ensemble were used, respectively, as the output RQD and associated uncertainty. We applied a post-processing clip to the interval [0, 100]% to enforce the physical bounds of RQD.

An exploratory data analysis shows two key characteristics of the RQD dataset that required extra consideration for the workflow design. First, more than half of all logged RQD values are zero, and fewer than one tenth exceed 60% (Fig. 1). To account for the strong bias towards zero values, we used kernel-density estimates to up-weight rarer high-RQD samples during training. Second, the moderate positive correlation between resistivity and RQD strengthens once the borehole logs are smoothed over larger window sizes, reflecting the difference in vertical resolution between AEM data (2 m at surface, ~30 m at depths of 300 m) and core logging (0.75 m intervals). Thus, rather than using raw RQD observations, we predicted mean RQD over a 10-m vertical window. This is closer to the resolution of the AEM system and dampens small-scale noise without being so large as to obscure engineering-relevant variations.

Prediction performance was evaluated in two ways. First, we validated performance with 5-fold grouped cross-validation in which the groups correspond to 1×1 km blocks. Second, we simulated performance in early-phase projects by using only 30%, 40%, and 50% of boreholes and compared to the full-data result.

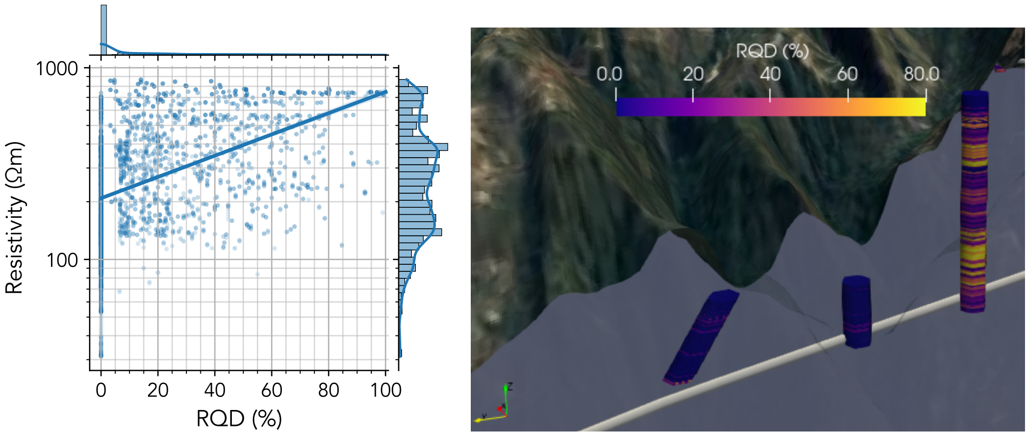

Figure 1 *(Left) Distribution of borehole RQD across the training dataset with a linear regression trendline; (Right) Three-dimensional view of three boreholes along Tunnel 1 coloured by observed RQD; vertical exaggeration 2×. Rightmost borehole is 390 m long.*

Results

The resultant RQD model is generated continuously across the whole resistivity model volume. Example sections are shown in Fig. 2 along the first 9 km of the railway alignment and as a slice orthogonal to the railway in Fig. 3. The prediction is largely steered by resistivity, depth, and observations at nearby boreholes, with deeper locations and higher resistivities generally having higher RQD predictions. The estimated uncertainties are highest in large gaps between boreholes (mostly highland areas), at deeper locations, at resistivities between 300 and 2000 Ωm, and where predicted RQD values are high. Five-fold grouped cross-validation produced a mean absolute error of 11.1 percentage points (p.p.) and a root-mean-square error of 17.0 p.p. Predictions in the low and intermediate ranges are generally accurate. The model systematically underpredicts the highest observed RQD values, but the uncertainty estimates are higher to compensate. We attribute this asymmetry primarily to two factors: the tendency of high RQD values to occur in spatially isolated, thin layers (Fig. 1); and lower EM signal sensitivity to deeper locations and high-resistivity materials where competent rock tends to occur.

Results also show that using a smaller subset of boreholes can lead to a model that is still useful and insightful (Fig. 4). The principal large-scale structures in the predicted RQD field — zones of generally low rock quality, interspersed with localised higher-RQD pockets — remain visible even when only 50% of the boreholes are used. These structures are not all visible when only 30% of boreholes are used, but many of the major transitions (like the low-to-high-RQD transitions at tunnel elevation around chainages 24.5 km and 30.5 km) are still visible. Hence, while not all details of the final model can be recovered, using fewer boreholes still lets one recover the major structures present in the resistivity model.

Discussion

Relative to the volumetric classification workflow on which this new regression workflow is built, the principal adjustments required to stabilise training were a larger ensemble and a longer early-stopping patience. KDE-based sample weighting and inclusion of physical limits on the parameter — either as a training loss term or as post-processing — were also necessary. The remainder of the architecture transferred directly. We expect the method could generalise to other engineering parameters, such as unconfined compressive strength, permeability, or porosity, so long as there is a correlation between the desired target, geophysical parameters, and other available support variables.

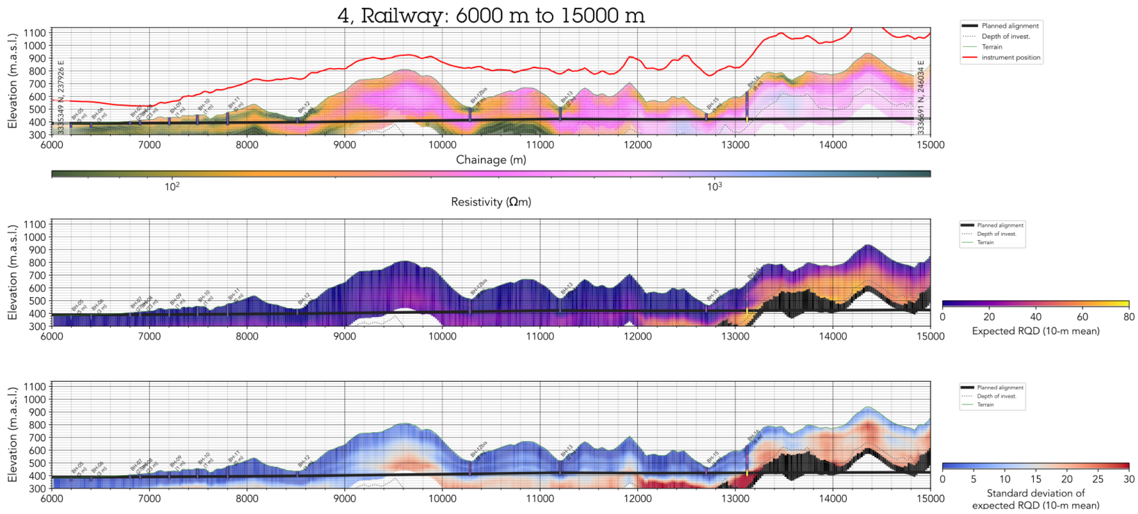

Figure 2 *Longitudinal section through the predicted RQD model along Tunnel 1 (chainage 6–15 km). From top to bottom: inverted resistivity, predicted RQD, and estimated standard deviation.*

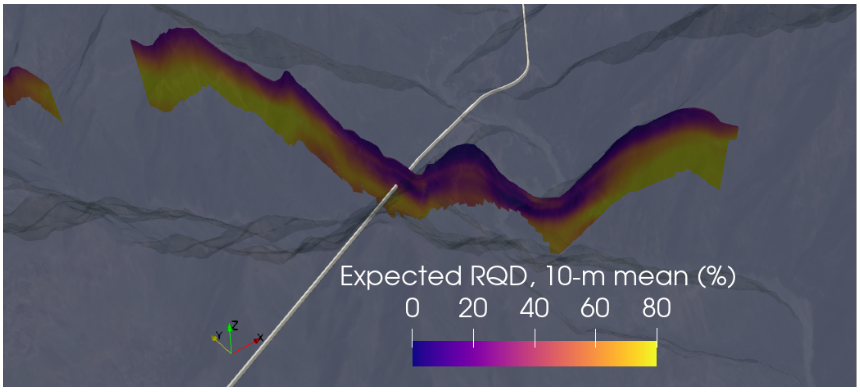

Figure 3 *Three-dimensional view of an orthogonal cut through the predicted RQD model surrounding the rail alignment (white tube).*

Though cross-validation showed good performance, two underlying limitations made precise prediction of RQD (especially high values) difficult. The first is petrophysical: the resistivity–RQD relationship is inherently imperfect, since lithology, weathering, and porosity each modulate resistivity independently of fracturing. Including supporting variables such as horizontal coordinates and terrain derivatives partially mitigates this, and we expect that incorporating additional data — for example, surface geological maps — would yield further improvement. The second relates to acquisition geometry. The nominal flight height for the SkyTEM 312 system is 35–60 m, but mean terrain clearance along this corridor was ~200 m and ranged from 50 to 400 m, which reduces signal sensitivity at tunnel elevation and contributes to the underprediction of high RQD values noted above.

Despite these limitations, the resulting predictions are useful and reasonably accurate. On-site engineers report that as-built observations along Tunnel 1 (Fig. 2) qualitatively align with the generally low rock-mass quality predicted by the model: low RQD values throughout but a significant increase towards the northeast. Cross-validation indicates that predicted RQD is typically within 10–20 p.p. of the observed value across the low and intermediate ranges. High RQD values are systematically underpredicted, but the corresponding uncertainty estimates, which are also elevated in those zones, alert users where to use greater caution.

The present case in Uttarakhand serves primarily as a technology demonstration, since some tunnels are already excavated and the railway alignment is fixed. The potential benefit for future, early-phase projects is much greater. The off-alignment predictions (Fig. 3) demonstrate how AEM-derived rock-quality models can support route adjustments in early-phase projects, similar to applications in road-alignment optimisation elsewhere (Christensen et al., 2020). The borehole-subset experiment indicates that drilling campaigns can be reduced or, more usefully, redistributed, allowing one to prioritise drilling within critical zones rather than at uniform spacing along the alignment.

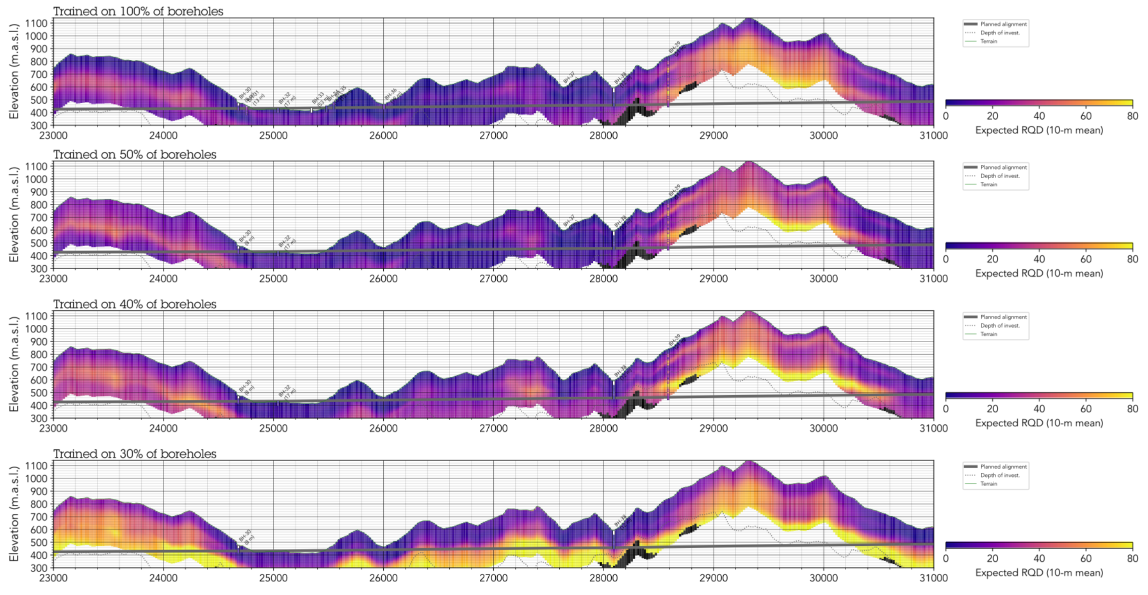

Figure 4 *Predicted RQD at 23–31 km using 100%, 50%, 40%, and 30% of the boreholes.*

Conclusions

We extended an existing volumetric machine-learning workflow from classification to regression and applied it to the prediction of Rock Quality Designation from AEM resistivity along a 40 km Himalayan rail-tunnel corridor. The cross-validated regression model achieves a mean absolute error of 11.1 percentage points and recovers the principal large-scale structures of the RQD field; predictions along Tunnel 1 are qualitatively consistent with as-built rock conditions. A subset experiment showed that approximately 40-50% of the available boreholes are sufficient to recover those structures, which suggests that future ground-investigation campaigns on similar corridors could be substantially reduced or redistributed when supported by an AEM survey of comparable quality.

References

Christensen, C.W., Harrison, E.J., Pfaffhuber, A.A., Lund, A.K., 2021. A machine learning–based approach to regional-scale mapping of sensitive glaciomarine clay combining airborne electromagnetics and geotechnical data. Near Surface Geophysics 19, 523–539. https://doi.org/10.1002/nsg.12166

Christensen, C.W., Skurdal, G.H., Pfaffhuber, A.A., Rønning, S., Lindgard, A., Sellgren, K.C., 2020. Airborne geoscanning and efficient geotechnical ground investigation workflows: a road-building case study from Central Norway. 18th Nordic Geotechnical Meeting, 25–27 May 2020, Helsinki, Finland.

Lysdahl, A.K., Christensen, C.W., Pfaffhuber, A.A., Vöge, M., Andresen, L., Skurdal, G.H., Panzner, M., 2022. Integrated bedrock model combining airborne geophysics and sparse drillings based on an artificial neural network. Engineering Geology 297, 106484. https://doi.org/10.1016/j.enggeo.2021.106484