Data-driven geotechnical hazard assessment: practice and pitfalls

WJ McGaughey *Mira Geoscience Ltd., Canada*

Abstract

Geomechanical risk in mining is universally understood to depend on many apparently disparate factors acting together such as stress, stiffness, mine geometry, rock mass character, rock type, structure, excavation rate and volume, blasting, and seismicity. We have worked on many case studies over the years in both underground and open pit mines with the objective of discovering and documenting the correlation of such factors with the experience of geomechanical failure. Whether that failure is slope failure, strainbursting, fault slip-induced rockbursting, roof fall, or any other of many possible failure types, statistical correlations among the different classes of data can be found, and predictive rules for understanding geohazard based on their quantitative combination can be established and deployed in day-to-day operations. This data-driven approach requires application of methods and avoidance of pitfalls that can be standardised into a universally applicable workflow. We discuss the workflow and the pitfalls in analysis to be avoided through case study examples.

Keywords: *geomechanical hazard assessment, data-driven analysis, data fusion, machine learning, artificial intelligence (AI), predictive analytics, rockburst, roof fall, slope failure*

1 Introduction

An obviously essential factor in successfully understanding, assessing, and acting upon mining geomechanical risk is the ability to forecast with some assurance the likelihood of experiencing geotechnical hazard incidents in location and time. It is uncontroversial to state that geotechnical hazard incidents, from rockbursts to roof falls to slope failures, are complex phenomena with multiple contributing factors. If this were not the case, the timing and location of their occurrence would be predictable and therefore avoidable.

Expert, knowledge-based combinations of contributing factors into an overall hazard score for particular classes of failure has often been applied by domain experts (e.g. Heal et al. 2006), particularly in mine design (e.g. Castro et al. 2011). However, expert knowledge also has real human constraints; consensus among experts, particularly when there is a need to apply quantitative weightings to multiple factors, is elusive; and more importantly, the multiple factors that contribute to geotechnical hazard incidents have both complex and site-specific interdependencies. Similar issues of complexity have led to the rise of artificial intelligence (AI) systems relying on machine learning (ML) algorithms or 'predictive analytics' in other fields to understand how to relate patterns in input data to the prediction of specific outcomes. The phrase 'data-driven geotechnical hazard assessment' in the title of this paper means using the history of geotechnical hazard incidents at a site to predict the likelihood in location and time of future incidents of the same class. We have undertaken such studies at many sites for many types of geotechnical hazard, and incorporated the methods into a working system for automated hazard assessment and reporting now implemented at several operations (McGaughey 2014; McGaughey et al. 2017). The subject of this paper is not the operational systems or methods of automated hazard assessment and reporting, but rather what we have learned in terms of the practices and pitfalls of applying data-driven methods to the problem of geotechnical hazard assessment.

We present a general workflow for analysing relationships among multiple input data streams and a site history of geotechnical hazard incidents, based on similar workflows in other application domains of ML. The goal of these algorithms, regardless of application domain, is to develop general rules based on historical data and then apply those rules to new data. This is an intensive general area of research in computational science and applications in other domains abound, from finance to marketing to the life sciences. The goal

of our workflow is to develop a quantitative model that combines weighted inputs of raw or transformed data available on the mine site into an optimum estimation of geohazard probability in location and time in the mine, and then use that model to provide routine automated computation and reporting to update the hazard assessment in response to new data.

There are many ML methods available to tackle this problem, but choice and application of particular ML algorithms is not the subject of this paper. For a useful, short introduction to the general application of ML, see Domingos 2015. The workflow described here concerns setting up the mining geohazard problem specifically for application of these techniques. Most of the examples we describe here used a standard 'weights of evidence' (WOE) or 'naive Bayes' approach, which has several advantages including conceptual simplicity and ease of implementation. It also has the benefit of a long history of acceptance in the mining industry, although primarily in mineral exploration applications (e.g. Agterberg et al. 1990; Agterberg & Cheng 2002; Bonham-Carter et al. 1989).

WOE describes the relationship between predictive variables such as stress, seismicity, rock type, or rock quality designation, and a binary target variable, such as hazard occurrence at a particular location and time in the mine (binary meaning the hazard happens at that location and time or it does not). "*Binary classification models are perhaps the most common use-case in predictive analytics. The reason is that many key client actions across a wide range of industries are binary in nature, such as defaulting on a loan, clicking on an ad, or terminating a subscription*" (Larsen 2015). These are examples to which we could add 'experiencing a geotechnical hazard incident.' Larsen (2015) goes on to list advantages of the WOE method generally, all of which are relevant to the mining geohazard case:

- Considers each variable's independent contribution to the outcome.

- Detects linear and non-linear relationships.

- Ranks variables in terms of 'univariate' predictive strength.

- Visualises the correlations between the predictive variables and the binary outcome.

- Seamlessly compares the strength of continuous and categorical variables without creating dummy variables.

- Seamlessly handles missing values without imputation.

- Assesses the predictive power of missing values.

Use of WOE also has drawbacks, such as an assumption of conditional independence of input predictive variables, which is never fully met in practice and must be kept in mind when interpreting results.

Most importantly, whatever the chosen method of ML or predictive analytics, it is imperative that we use systems that provide rules that are understandable by domain experts. Deep learning or other methods that rely on neural networks or other algorithms with 'hidden layers' between input and output may (or may not) provide more robust prediction, but in the mining case, prediction accuracy must be accompanied by an answer to the question 'why?' because site engineers have to act on the output. Depending on why a prediction of high probability of geohazard incident has occurred, actions taken may range from dismissing the prediction for site-specific reasons, investigating underground or in the pit, reinforcing support, adjusting operations, or abandoning an area. A 'black box' prediction will not provide the required engineering confidence in the output. Furthermore, the system must provide a foundation for a routine, automated computation and reporting system to update hazard assessment in response to new data as described in McGaughey et al. (2017).

2 Workflow overview

Although each site is unique, over many case studies with different mining methods (both underground and open pit), different ore types, and different types of geotechnical hazard, we have developed a standard workflow for tackling the data-driven geotechncial hazard assessment problem. In this section, we provide a brief overview of the workflow, and in the following section, describe how we carry out each step as well as some of the issues to be aware of.

The major steps of the workflow are:

- Establish target hazard incident type.
- Establish hazard criteria and data sources.
- Establish geometrical mine model.
- Establish mapping from data sources to hazard criteria.
- Establish list of hazard incident occurrences for back-analysis.
- Create mine model snapshots.
- Data fusion.
- Training and validation.
- Implementation.

The practice and pitfalls of undertaking each step in the workflow are discussed in the following section.

3 Practice and pitfalls

In this section, we describe how we carry out each step of the workflow, using examples from case studies. From the initial site brainstorming to completion of the analysis, we follow a now standardised step-by-step process that we have continually refined over the past decade. The steps, in hindsight, now seem obvious but in fact, there was much trial and error in refining the workflow, which could not have been done without the experience of analysing many hazard types in underground and open pit coal, metal, and potash mines.

3.1 Establish target hazard

The first step is forming consensus among stakeholders as to the 'target' hazard incident type in question that we are attempting to model and understand. Sites may have different types of geotechnical hazard incidents and it is important to recognise each of them as separate problems to be solved, even if they are later combined into a holistic assessment of all hazard types. This workflow is to be applied to each hazard incident type separately. This step requires brainstorming-style discussion with site personnel. Examples from discussion with sites from our own work include recognition of the necessity of separating rockbursts into multiple types such as strainbursts (pillar bursts or development bursts) and fault slip induced rockbursts, which are distinguished interpretationally using site-specific criteria; strainburst-type rockbursts that occur during development versus strainburst-type rockbursts in older ground; slope failure in pits at different scales (e.g. face, inter-ramp, total slope); and water inflow hazard due to salt dissolution versus due to structural anomaly. We have worked on cases where multiple examples of what was believed to be a single type of hazard were revealed through the analysis to be separable types. The practice is to carefully think through what hazard type you are attempting to understand and ensure that all historical examples of it are truly of the same type, as success of the data-driven approach relies on consistent underlying statistics that will enable pattern recognition. Expect to revise your understanding of the hazard class as you carry the analysis through and possibly find unexpected or contradictory results. The potential pitfall is mixing multiple types of target hazard incident populations with different underlying statistical correlations to the input predictive variables, which will ultimately yield poor results.

3.2 Establish hazard criteria and data sources

The hazard criteria are a list of input variables that are recognised as possibly having predictive value, meaning a useful correlation to the geotechnical hazard in question. In practice, we approach the listing of hazard criteria through a combination of our own past experience with sites having some degree of similarity and brainstorming with stakeholders who know the site and its hazard experience well. In typical studies, we will list several tens of hazard criteria types. This is the key initial step toward quantification of the problem. Tables 1 and 2 show actual hazard criteria used in a couple of case studies, grouped by category.

We group hazard criteria into general categories appropriate to the target hazard, and organise them into a table in which each row corresponds to one candidate hazard criteria. The columns are hazard criteria name, quantification metric, data required to model the hazard criteria spatially and temporally, and how we will represent the hazard criteria in a 3D model of the mine at a given snapshot in time. Figure 1 is an image of a fragment of such a table from one of our first data-driven case studies, analysing the factors related to fault slip rockbursts in a deep underground mine in Canada.

Table 1   Sample candidate hazard criteria for fault slip rockbursting, adapted from a case study

Hazard criteria category	Example candidate hazard criteria
Mine development	Age of development, development rate, ground support category, age of support, span, orientation, proximity to intersections, depth.
Rock mass	Joint orientation, joint spacing, uniaxial compressive strength, fracture frequency, rock mass rating, rock quality designation.
Geology and structure	Rock type, proximity to contacts, proximity to waste gaps between ore zones, proximity to major structures, proximity to structural intersections, orientation of major structures, fault category, proximity to dykes.
Stress	Maximum principal stress, deviatoric stress, excavation ratio (as a local stress proxy), fault slip tendency.
Seismicity	Seismic event density, proximity to seismic cluster, local magnitude, E_s/E_p ratio, static stress drop, seismic moment, energy index.
Monitoring	Deformation from extensometers, convergence stations.

Table 2   Sample candidate hazard criteria for open pit slope failure, adapted from a case study

Hazard criteria class	Example candidate hazard criteria
Rock mass rating	Fracture frequency, rock quality designation, joint spacing, intact rock strength, rock mass ratings (Laubscher, Bieniawski rock mass rating, Q), alteration on joints, joint roughness, joint infill type, joint infill thickness, joint aperture, joint moisture, blasting.
Other geotechnical criteria	Core recovery, rock type, geotechnical zones; based on rock type, rock mass rating, structural orientation, and orientation of pit, failure mechanism; defined per geotechnical zone, but a zone may have more than one failure mechanism, rainfall, groundwater.
Pit geometry	Interaction of pit geometry with joint orientation, intersection of pit geometry with major structure orientation, overall slope angle, stack slope angle, bench angle, overall height of slope, height of stack, berm width, bench height.
Structure	Interaction between structures, continuous joints (no displacement), faults, folds, shears, dykes, veins, shear zones, water-bearing versions of above.
Mining operational	Quality of clean-up, limit blasting performance, achievement of limit (design), time since last blast.
Monitoring	Radar monitoring, laser scanning, survey prism monitoring, wireline extensometers, visual inspections.

Criteria	Quantification Metric	Data Required	Representation (objects/properties/resolution)
Mine Development			
- ground support type	support code	support code legend definition and specification for actual development at Craig Zones 10 and 11	classification as discrete variable on development wireframe surfaces; resolution requirement on order of 10m
- ground support stiffness	stiffness code	support code legend definition and specification for actual development	classification as discrete variable on development wireframe surfaces; resolution requirement on order of 10m
- age of development	age in days	development history	stored as continuous property on development surfaces
- age of support	age in days	support placement history	stored as continuous property on development surfaces
- geometry	3D geometric model	3D dxf or 2D dxf with guidelines on how to extend to 3D	representation as 3D wireframe surfaces, resolution requirement on order of 1m
- orientation	strike of development axis	none; can be computed from modelled development	stored as continuous property on development surfaces
- intersections	proximity to drift intersections	none; can be computed from modelled development	create point set defining all intersections of the centerlines and compute proximity to this point set on voids surface
- development rate	rate	development rate historical records or possibly derived from "age of development" variable	stored as continuous property on development surfaces
- hydraulic radius	largest span length	approximated from modelled development	stored as continuous property on development surfaces
- depth	metres	depth reference	property on voids surface
Mining Method	mining method	mining method legend	classification as discrete variable on

Figure 1   Section of a table used in a case study to document hazard criteria, required source data, and model representation. Such tables are invaluable for organising and tracking project activity. The purpose of the figure is to exemplify useful process documentation rather than to prescribe criteria that should be used in any particular case

3.3   Establish geometrical mine model support

Geotechnical hazard incidents occur at the mine-rock interface, whether that is in development, a stope, or on the pit surface. We are therefore always conceptually modelling hazard on a surface and never in the 3D rock mass volume away from the mine workings. This is true even though many of the hazard criteria themselves are fundamentally of the 3D rock mass volume, e.g. seismicity or faults that do not intersect the mine workings, yet still pose an important geotechnical risk at the workings.

In our first attempts at this type of hazard analysis, we modelled hazard explicitly on wireframe surfaces; either the underground mine workings or pit shells as 3D wireframes. We modelled hazard as a property of the vertices of the wireframes (Figure 2). Wireframe surfaces are typically supplied to us as 3D DXF files output from mine planning software systems, and they are, as often as not, a mess, i.e. meaning highly irregular triangle density, many instances of collapsed or unstable triangles (two out of the three vertices co-located or nearly so), repeated vertices, missing segments, etc. These wireframes may look fine visually in a 3D viewer but are not appropriate for a mathematical hazard model. We spent a lot of time in early studies manually repairing wireframes. Eventually we gave up, replacing wireframes with a point model in which the mine is digitised as a series of points with a nominal spacing typically in the range 1–10 m. The points can exist on any rock interface where hazard is to be modelled, but typically it is drift centrelines if underground or bench and face centrelines if open pit. Figure 3 is a screenshot from the interface of our hazard reporting system, illustrating points digitised along development centrelines for one mine level, each point holding values for the many hazard criteria modelled at that point.
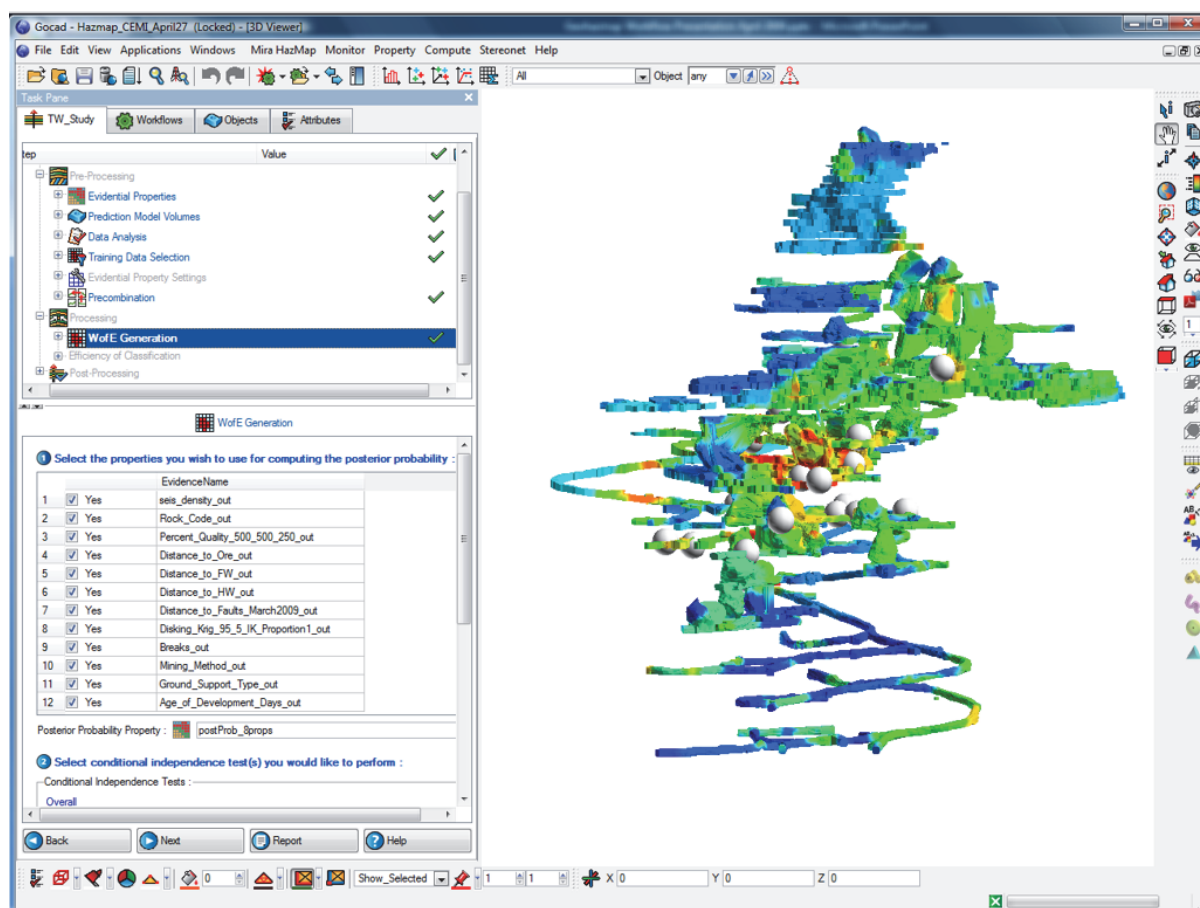

Figure 2   Rockburst hazard shown in a colour scale (warmer colours being more hazardous, in units of log{probability}) from one of our early underground case studies. Rockburst locations are shown as white spheres. We have since abandoned modelling hazard directly as a property of 3D wireframes because of the low quality wireframes typically available on mine sites, requiring great effort in manual clean-up

Zoomed for clarity

Figure 3   Screenshot of the web browser interface for our hazard management and reporting system, 'Geoscience INTEGRATOR', illustrating for one mine level the digitised centreline points which hold the values of both the individual hazard criteria and the computed hazard index based on them. Part of the interface has been zoomed into for clarity. The zoomed area shows the part of the interface where the user selects which value, and at what date, to display on the mine model points. In this case the 'rockburst hazard index' is shown, with values above a given threshold shown with larger radius for quick identification of hazardous areas

3.4   Establish mapping from data sources to hazard criteria

The transformations required to represent hazard criteria, such as shown in Tables 1 and 2, are the most complex step in the workflow. For each of the criteria we ask ourselves, 'How do we know the value of the criteria at any point in the mine model?' Many criteria have natural representations at the location of the digitised model points, such as rock type, rock quality, berm width, and drift orientation. Rock type and rock quality parameters are typically mapped directly from block model cells to the points. Drift orientation must be computed from the segment orientations of the centrelines. These are the easy ones.

More challenging are criteria we know to be important conceptually, but their mode of representation on the digitised mine model points is far from straightforward. Take, as an example, fault slip tendency (see Morris et al. 1996 for a useful description), which intuitively should find a correlation with fault slip induced rockbursting. Fault slip tendency is a property of fault surfaces, varying according to local ratio of shear to normal stresses as both the stress field and orientation of the fault varies across its surface geometry. As the natural representation of this criteria is a spatially varying property on a fault surface that may not intersect the mine workings at all, how do we represent its varying potential contribution to hazard on each of the mine model points? There is no definitive answer to this question. Our approach for any such variable has been to take a common sense view and try something we think is reasonable. In the specific case of fault slip tendency, we look at the statistical distribution of values over all faults in the model, and then compute proximity of the mine model points to sub-areas of the fault surfaces with fault slip tendency Z-scores above a chosen threshold. In practice, over several case studies in deep rockbursting mines, we have found this specific criteria, computed this way, to be very useful, although it is certainly possible that further research could reveal improved methods of its calculation. There are many other such properties (e.g. extraction ratio) for which the mapping from data to hazard criteria is intrinsically complex with no obvious right way to do it, but for which a common sense approach has consistently yielded good results.

In the ML community, the computation of such predictor variables (the hazard criteria in our case) is known as 'feature engineering'. A survey of the ML literature quickly reveals that feature engineering is an essential component of any ML workflow, and how it is done often marks the difference between success and failure:

> "At the end of the day, some machine learning projects succeed and some fail. What makes the difference? Easily the most important factor is the features [hazard criteria] used. If you have many independent features that each correlate well with the class [hazard incidents], learning is easy. On the other hand, if the class is a very complex function of the features, you may not be able to learn it. Often, the raw data is not in a form that is amenable to learning, but you can construct features from it that are. This is typically where most of the effort in a machine learning project goes. It is often also one of the most interesting parts, where intuition, creativity and "black art" are as important as the technical stuff. First-timers are often surprised by how little time in a machine learning project is spent actually doing machine learning. But it makes sense if you consider how time consuming it is to gather data, integrate it, clean it and pre-process it, and how much trial and error can go into feature design" (Domingos 2015).

It is in this mapping of data sources to hazard criteria (the feature engineering) where the crux of the data integration lies. The lack of reported back-analysis studies that are both quantitative and multi-disciplinary is testament to its difficulty. How can one get beyond seismic data in studies of seismically-induced hazards to quantitatively include the contributing influence of mine geometry, structure, stress, deformation, and rock mass properties? The answers are all in the mappings from the raw data sources to the representation of hazard criteria on the mine model points.

In practice, a powerful combination of spatial analysis, interpolation, statistical analysis, geological and mine modelling is required. For this, we use GOCAD® Mining Suite (Mira Geoscience Ltd 2017), as illustrated in Figure 4, both in the back-analysis and as run-time engine in our automated systems to compute features and update them automatically as new data become available in an operating mine.
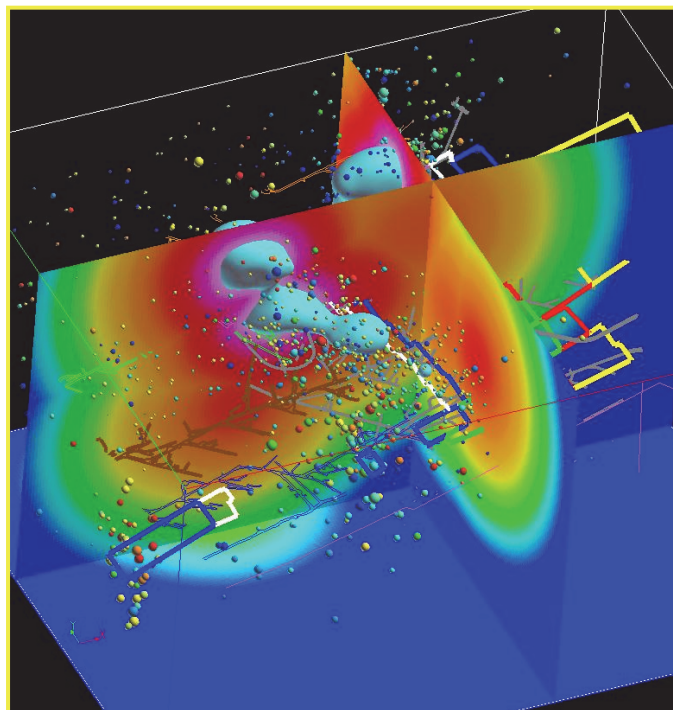

Figure 4   A screenshot from GOCAD Mining Suite (Mira Geoscience Ltd 2017) showing feature engineering for the computation of hazard criteria. In this case, mine seismic events over a time window are shown as spheres, the cyan bubble is an isosurface enveloping a volume of high seismic density, and the coloured sections show proximity to the edge of that isosurface, which can be projected onto mine development

3.5   Establish list of hazard occurrences for back-analysis

Care must be taken in compiling the list of hazard occurrences to be used in the back-analysis. Because we are attempting to find a method of predicting the probability of hazards occurring based on a statistical treatment of past experience, we must be as certain as we can that the hazards used in the model are all of the same type; that is, have similar statistical relationships to the modelled hazard criteria, as discussed in Section 3.1.

We want the list of hazard incidents used in the back-analysis to be as extensive as possible, not only to improve the modelling statistics, but also because we have to hold some of the incidents back in testing to validate the model. On the other hand, as discussed in the next section, for each calendar date for which we have a hazard incident we have to create a separate mine model. The more calendar dates, the more mine models and therefore, more modelling time is required. The project time can increase significantly as the number of hazard incidents used in the back-analysis increases.

There are statistical tests that we normally carry out in a phase of exploratory data analysis at this point in the workflow. It is advisable to carry out unsupervised 'cluster analysis' on a population of historical hazard occurrences believed to be of the same class. Most cluster analysis algorithms (e.g. k-means clustering or self-organising maps) allow experimentation with selection of the number of natural sub-populations in an attempt to find the most natural true groupings. The example shown in Figure 5 illustrates isosurfaces drawn as envelopes around sub-populations of strainbursts showing statistical similarity in an underground mine. It becomes an interpretive decision as to whether sub-population groups should be split into separate problems (meaning separate runs through this workflow), depending on degree of similarity and necessity of maintaining a statistically significant population of target hazards for the analytics step of the workflow.

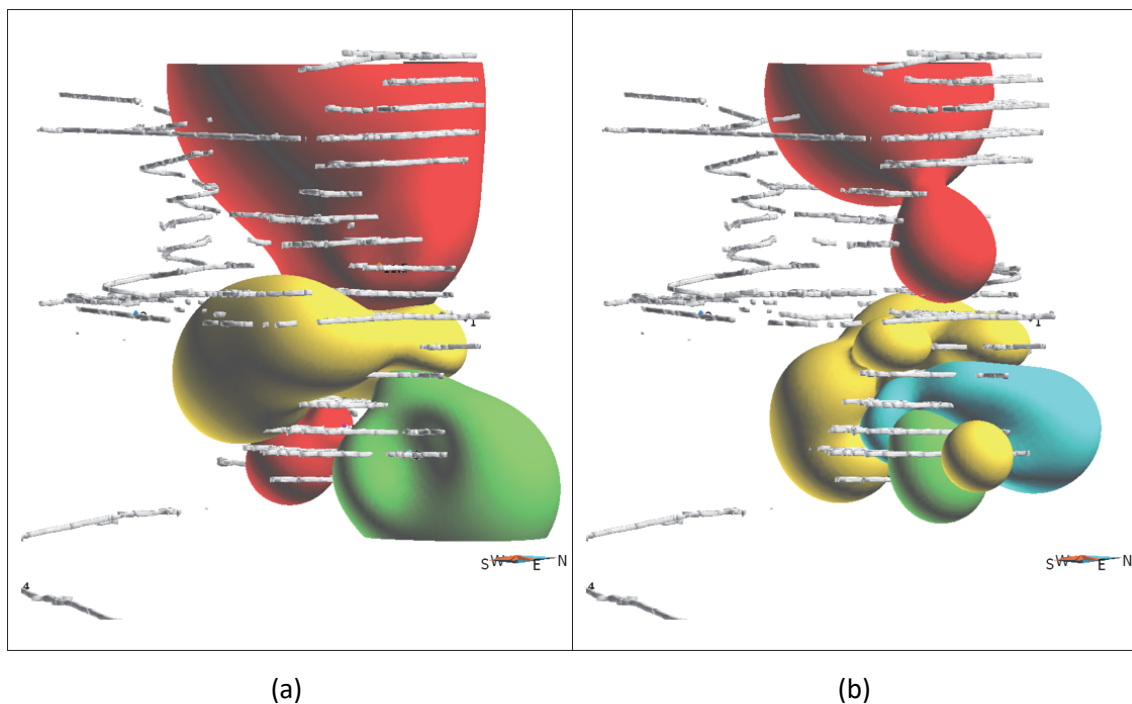(a)   (b)

Figure 5   Unsupervised classification of target hazards. In this case, a population of strainbursts was submitted to two different unsupervised classification methods. (a) HyperCube™ (BearingPoint France SAS 2015); (b) k-means clustering. Both methods appear to show at least three distinct populations

3.6 Create mine model snapshots

Geohazard assessment is fundamentally a 4D problem; hazard incidents occur at a location and a time. The individual hazard criteria represented on the digitised mine model points are either static or dynamic in time. For example, geology does not change with time as the mine evolves, but criteria related to development, stress, seismicity and other monitoring data do change in time.

In practice, we create a version of the mine model for each date on which a geotechnical hazard incident occurred to capture the state of the mine at a snapshot in time associated with a hazard incident. Each snapshot is created by duplicating the mine model points from the previous snapshot, adding or removing points where necessary to account for loss of mine development (where it has been removed by the mining process) or new development, and copying the static criteria such as rock type or proximity to structure to the points representing the mine model at the next snapshot in time. Dynamic criteria, such as those related to stress, seismicity, or other monitoring data, must be recomputed. For seismic data we typically create numerous hazard criteria on the mine model points by analysing windows in time prior to the incident. For example, if one of our criteria is 'microseismic event density in the previous 10 days', we have to interrogate the microseismic database for the events that correspond to the 10-day windows prior to each of our mine model snapshots and repeat the event density computation for each of the snapshots.

This process is time consuming in proportion to how well the data are organised. Structured data managements systems such as Geoscience INTEGRATOR (McGaughey et al. 2017), built expressly for this purpose, make the problem of creating multiple mine model snapshots tractable.

3.7 Data fusion

Although an important level of integration or 'data fusion' is achieved simply by compiling the mine model snapshots, each of which contain numerous criteria derived from disparate data sets on each of the mine model points, we still must combine the snapshots with the hazard incident history in a structured way, suitable for application of ML or any other type of statistical analysis. There are many ML algorithms and codes available and virtually all of them expect data to have the same input structure. That structure is a flat table, in practice typically a csv file. The rows correspond to observations and the columns correspond to 'features', the individual hazard criteria in our case. In supervised learning algorithms where the goal is to 'train' the system to learn the relationships among features that can be used to generalise predictions, an additional 'training' or 'target' column records the prediction class, which in our case is a simple binary flag indicating whether or not a hazard incident occurred for that observation. We refer to this flat table as the 'data fusion table'.

The structure of the data fusion table is illustrated in Figure 6.

We construct the data fusion table by first constructing a version of it for each of the snapshots in time, one table per calendar date when an incident occurred, documenting the state of the mine on that date. In practice we use the Geoscience INTEGRATOR data management system to carry this out automatically. The mine model point(s), corresponding to rows of the table, where the hazard occurred are flagged in the final training column of the table with the value X and rows corresponding to points where the hazard did not occur are given the value O. The tables for each snapshot are then appended, resulting in a single csv file that corresponds to a flat table in which each row represents a mine model point at a certain date (x, y, z, t) with its associated hazard criteria values as the columns, plus a final training column with the value X or O depending on whether a hazard incident occurred at that location and time (x, y, z, t).

Figure 6   Structure of the data fusion table. Each row ('feature vectors' in machine learning parlance) corresponds to a single (x, y, z, t) observation, in practice a single digitised mine model point at a snapshot in time when a hazard incident occurred. All mine model snapshots are concatenated into this single table. Each table column contains the values of one hazard criteria on the mine model points over all snapshots in time (x, y, z, t). The final 'training' column, labelled 'hazard', records whether a hazard incident was experienced (X) or not (O) at the observation point (x, y, z, t) corresponding to that row. The data fusion table, in an ASCII (csv) representation, is nearly universally applicable as input to machine learning or other statistical analysis algorithms. Its construction is a key achievement of our work and the principal subject of this paper

3.8   Training and validation

Training is the process through which ML methods analyse relationships in the data fusion table and generalise rules based on the hazard criteria from the specific examples of hazard incidents. It is not the purpose of this paper to discuss the analysis of the data fusion table, i.e. the running of the ML algorithms, in detail because the most important and difficult parts of the entire exercise are all the steps leading up to the creation of the data fusion table.

Nevertheless, there are a couple of general points worth discussion. The size of the data fusion table is approximately (number of points in the mine model) × (number of calendar dates from which hazard incidents have been represented) rows by (number of modelled hazard criteria) columns. Typically, we have a few thousand mine model points, perhaps a couple of dozen calendar dates of historical hazard incidents, and a few tens of hazard criteria, i.e. hundreds of thousands of rows by twenty or thirty columns. Over the hundreds of thousands of rows, perhaps only one row per 1,000 or one per 10,000 will have an X and the remainder will be Os, corresponding the reality that hazard incidents worth modelling are relatively rare events in time and space. This needle-in-a-haystack problem is known as 'biased' or 'unbalanced' training data in ML jargon. It is not uncommon. For example, banks wanting to predict which customers are more likely to fraudulently misuse a credit card (like hazard, a binary target class—it happens or it does not happen) must also work with unbalanced data. There are several techniques for dealing with this, for example undersampling the Os or oversampling the Xs, both of which have established methods described in the literature (e.g. Chawla 2010). Another important issue to consider in the analysis is correlation among predictor variables. Correlations among all variables are examined prior to analysis. If any variables are highly correlated, only one is kept. Post-analysis output 'rules' are also examined for correlation, to avoid

over-weighting individual components in the final hazard function. There are a host of other detailed issues that must be considered when doing this type of analysis, all of which are general issues of ML, and not specific to geotechnical hazard.

Figure 7 illustrates the general approach of the ML algorithms we employ. The Xs and Os from the data fusion table are investigated in the multi-variate hazard criteria space (two of many hazard criteria dimensions being shown in the figure for simplicity) for areas of 'over-concentration' of Xs. Or to put it another way, what combinations of hazard criteria values, such as local fault slip tendency and excavation ratio, have historically been associated with the occurrence of hazards such as rockbursts? Searching for such tell-tale associations in high-dimensional spaces (each hazard criteria being a dimension) cannot be readily done with traditional statistical, visual, or observational approaches. However, ML algorithms can effectively comb through the multidimensional space using brute force or many other approaches to discover relationships otherwise easily missed. Once ML has revealed the associations among predictor variables that correlate to the history of hazard incidents (the 'rules'), we establish the final hazard function at each mine model point as a weighted sum of scores (positive scores for each rule that a mine model point matches, and negative where it does not). The scores for each rule are computed using conventional Bayesian statistics based on binary membership classification, with the sum presented as a hazard incident probability estimate as illustrated in Figure 2. It is important to note that, unlike with neural network or similar systems, users can always see why an overall hazard score is high, i.e. what hazard criteria value and hazard incident rules contributed to it.
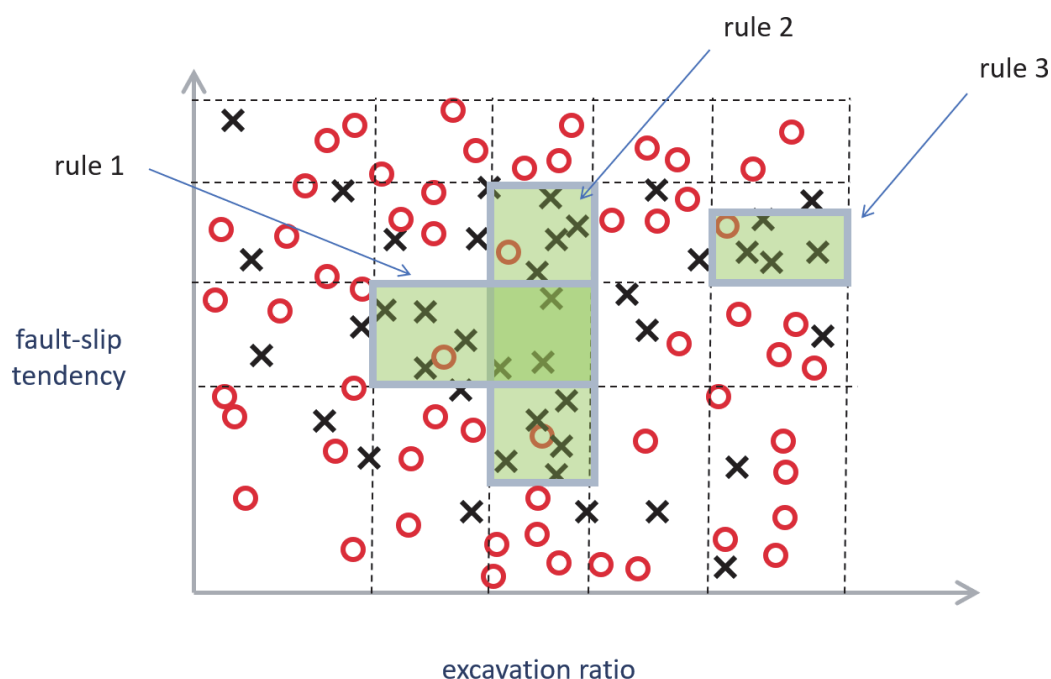

Figure 7   Machine learning discovery of 'rules' as bounded ranges of multiple criteria (only two in this example) in which there is a relatively high concentration of Xs versus Os, i.e. in which there is a relatively high probability of geohazard

Validation is essential when using these techniques. It is accomplished by withholding a subset of the hazard incidents from the analysis and testing the effectiveness of the rules generated with the remaining hazard incidents on the subset withheld. If all the hazard incidents available are used to develop the rules, there is a likelihood of overfitting the data. A common strategy is to withhold a randomly-selected fraction (typically a quarter or a third) of the training data, and repeat the analysis several times withholding different randomly-selected subsets, in each case measuring the effectiveness of predicting the withheld data. There are several standard tests available to provide metrics of success or failure in generalising predictive rules. In practice, we have found that in all cases these methods yield demonstrably superior results to knowledge-driven approaches.

3.9   Implementation

In practice, we implement the results of data-driven geotechnical hazard assessment through a data management system called Geoscience INTEGRATOR. It holds a representation of the discretised points of the mine model, updating each of the time-dependent hazard criteria as new data becomes available. It does this by using a run-time version of the SKUA-GOCAD geological and spatial modelling system that is contained within it and capable of carrying out complex spatial calculations under the automated control of the system. For example, as new seismic data becomes available, Geoscience INTEGRATOR can automatically send an updated time window of data to the modelling engine to compute updated values for hazard criteria such as microseismic event density. After the modelling engine updates the values on the mine model as required, the rules developed in the data-driven training are applied and new values of hazard are computed and reported. The system (Figure 8) for operational deployment was discussed at greater length in McGaughey et al. (2017).
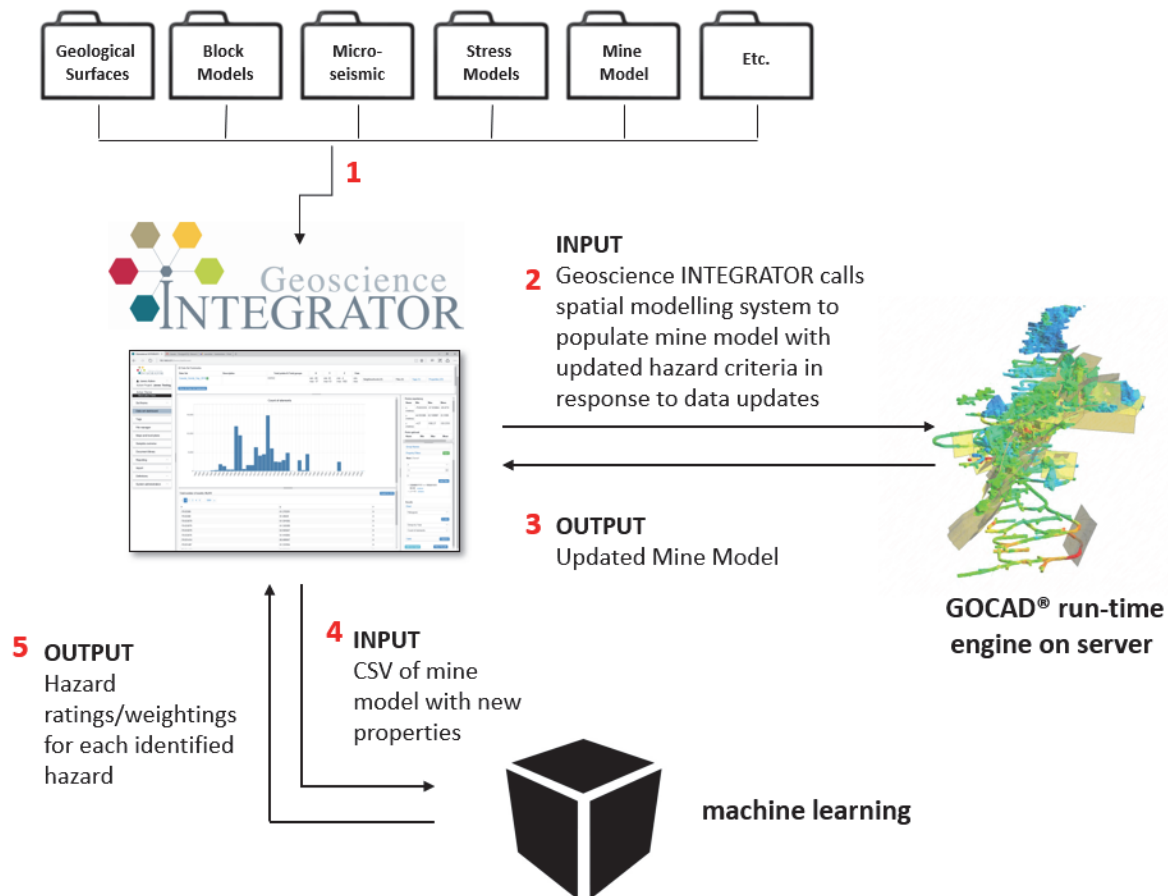

Figure 8   The flow of data between components of the Geoscience INTEGRATOR system for operational implementation of data-driven geotechnical hazard assessment

4   Conclusion

After several years of research and development, we have successfully brought together several methods and technologies for operational deployment of automated, data-driven geotechnical hazard assessment:

1. Methods of quantitatively integrating disparate geological, geotechnical, and mine-operational data streams.

2. Methods of representation of hazard criteria and hazard incidents to data-driven approaches such as ML to document the underlying correlations and optimally estimate hazard probability.

3. The technologies of data management and real-time connection to mine data sources to support automated computation and reporting to operators as new data becomes available. Standardising the workflow, documenting, and continually paying attention to the potential pitfalls in this type of analysis has demonstrated enormous value in its now routine application.

Acknowledgement

References

Agterberg, FP, Bonham-Carter, GF & Wright, DF 1990, 'Statistical pattern integration for mineral exploration', in G Gaal & DF Merriam (eds), *Computer Applications in Resource Estimation Prediction and Assessment for Metals and Petroleum*, Pergamon Press, Oxford.

Agterberg, FP & Cheng, Q 2002, 'Conditional independence test for weights-of-evidence modelling', *Natural Resources Research*, vol. 11, no. 4, pp. 249–255.

BearingPoint France SAS 2015, *HyperCube*, computer software, BearingPoint France SAS, Paris, https://www.hcube.io/en

Bonham-Carter, GF, Agterberg, FP & Wright, DF 1989, 'Weights of evidence modelling: a new approach to mapping mineral potential', in FP Agterberg & GF Bonham-Carter (eds), *Statistical Applications in the Earth Sciences*, Geological Survey of Canada, Ottawa, pp. 171–183.

Castro, L, Cottrell, B, Barker, P & Sintim, K 2011, 'Application of geohazmap to the pit slope design for the Detour Lake project', in E Eberhardt & D Stead (eds), *Proceedings of the 2011 International Symposium on Rock Slope Stability in Open Pit Mining and Civil Engineering*, Canadian Rock Mechanics Association.

Chawla, NV 2010, 'Data mining for imbalanced datasets: an overview', in O Maimon & L Rokach (eds), *Data Mining and Knowledge Discovery Handbook*, Springer, Berlin, pp. 875–886.

Domingos, P 2015, *A Few Useful Things to Know About Machine Learning*, viewed 1 October 2018, https://homes.cs.washington.edu/~pedrod/papers/cacm12.pdf

Heal, D, Potvin, Y & Hudyma M 2006, 'Evaluating rockburst damage potential in underground mining', *Proceedings of the 41st US Symposium on Rock Mechanics*, American Rock Mechanics Association, Golden.

Larsen, K 2015, *Data Exploration with Weight of Evidence and Information Value in R*, viewed 1 October 2018, MultiThreaded, https://multithreaded.stitchfix.com/blog/2015/08/13/weight-of-evidence

McGaughey, WJ 2014, '4D data management and modelling in the assessment of deep underground mining hazard', in M Hudyma & Y Potvin (eds), *Proceedings of the Seventh International Conference on Deep and High Stress Mining*, Australian Centre for Geomechanics, Perth, pp. 93–106.

McGaughey, WJ, Laflèche, V, Howlett, C, Sydor, JL, Campos, D, Purchase, J & Huynh, S 2017, 'Automated, real-time geohazard assessment in deep underground mines', in J Wesseloo (ed*.), Proceedings of the Eighth International Conference on Deep and High Stress Mining*, Australian Centre for Geomechanics, Perth, pp. 521–528.

Mira Geoscience Ltd 2017, *GOCAD Mining Suite*, version 17, computer software, Mira Geoscience Ltd, Montreal, http://www.mirageoscience.com/our-products/software-product/gocad-mining-suite

Morris, A, Ferrill, D & Henderson, D 1996, *Slip-tendency Analysis and Fault Reactivation*, viewed 1 October 2018, https://www.researchgate.net/publication/252699966_Slip-tendency_analysis_and_fault_reactivation