

Artificial intelligence and big data analytics in mining geomechanics

J. McGaughey

Mira Geoscience Ltd, Canada

Mining geomechanics presents specific challenges to application of the closely-related methods of Artificial Intelligence (AI), big data, predictive analytics, and machine learning. This is because successful use of these techniques in geotechnical engineering requires four-dimensional (x, y, z, t) data integration as a pre-requisite, and 4D data integration is a fundamentally difficult problem.

This paper describes a process and software framework that solves the pre-requisite 4D data integration problem, setting the stage for routine application of AI or machine learning methods. The workflow and software system brings together structured and unstructured data and interpretation from drillhole data to all types of geological, geophysical, rock property, geotechnical, mine production, fixed-plant, mobile equipment, and mine geometry data, to provide a data fusion capability specifically aimed at applying machine learning to rock engineering problems.

The system does this by maintaining 3D earth model and 4D mine model geometrical data structures, upon which multiple data sets are projected, interpolated, upscaled, downscaled, or otherwise processed appropriately for each data type so that the variables of importance for each problem can be co-located in space and time, a requirement for the application of any analytics algorithm. Documents and files can be stored, managed, and linked to data and interpretation to provide relevant metadata and contextual links, providing the platform required for AI solutions. The system rationale and structure are described with reference to specific AI challenges in rock engineering.

INTRODUCTION

Most people are aware of the AI technology revolution. From self-driving cars to medical, financial, and marketing applications, we have been exposed to its predictive power. Why have these methods not yet had a significant impact on understanding or forecasting mining geomechanics outcomes? The rewards of AI should be immense as mines get deeper and forecasting of stress-related or other rock behaviour becomes a limiting factor on safety and production. The reason for lack of success is simple – there is a fundamental barrier that makes mining geomechanics different from traditional AI applications.

AI and its close relatives, predictive analytics, machine learning, and big data (all of which in practice are either broadly synonymous terms or subsets of each other), work well when you can measure many variables on a specific entity, such as a mining machine, a length of drill core, or even an industrial process, and simultaneously record a condition that you want to be able to forecast such as machine failure, or the mineral and geometallurgical properties of rock, or the output of a process. AI can uncover complex, predictive relationships among measured variables and the condition to be predicted. That is why it is already being used with success in some corners of the mining industry, such as understanding the relationship between fleet vehicle data and maintenance requirements or predicting geometallurgical parameters from core scans.

However, in mining geomechanics, its application is far from simple. The reason for this is that the condition being predicted, such as the location and timing of a geotechnical hazard (including rock fall, rock burst, or slope failure, seismic event probability forecasting, ore dilution forecasting, or drawpoint hang-up prediction), may be related to known factors (e.g., geology, rock mass properties, fault structures, mine geometry, stress, extraction, production, stope sequencing, deformation, seismicity, blasting, and support). But those factors are in many cases not easily estimated quantifiable variables at the location where the prediction is required. The condition to be forecast (e.g., the rock burst or the slope failure) exists when and where it does because of the properties of the complex, four-dimensional, spatial and temporal natural earth and engineered mine system. Not only are many of the factors affecting the prediction separated in space and time from the location and timing of the forecast event, but many of the factors can only ever partially be known, because they are inferred from models (geological models, geotechnical models, numerical stress models, etc.) that are themselves created from sparse measurement or drillhole data.

Nevertheless, in spite of these particular challenges of applying modern AI or machine learning methods to mining geomechanics, success can and has been achieved. The solution is to take the focus off the mechanics of AI itself and put the focus on how these problems are set up for the application of AI methods, which is where deep domain knowledge and a mining-specific, supporting computational framework are required.

HOW ARTIFICIAL INTELLIGENCE WORKS

There is much confusion in popular usage of the terms used to describe what amount to a collection of pattern recognition algorithms. In formal usage, AI is a broad term encompassing the general field of computer simulation of human intelligence. Machine learning is a narrower term, conventionally a subset of AI that uses computer algorithms to create a predictive mathematical model based on so-called historical training data that can be used to forecast the relative probability of future occurrences of given events.

Classes of machine-learning algorithms include decision trees, random forests, support vector machines, Bayesian inference, ensemble methods, and others. Deep learning is a subset of machine-learning algorithms that uses neural networks. The term predictive analytics is roughly synonymous with machine learning, but more often used in a business application context. The term big data is conventionally reserved for very large data sets, typically comprising both structured data (such as tables of numbers) and unstructured data (documents, photos). In popular use however, and for the purposes of this paper, I consider AI, machine learning, predictive analytics, and big data to all be effectively synonymous, and will use the term AI. For rock engineering applications, the choice of AI algorithm matters much less than correctly setting up the inputs to whatever algorithm is chosen.

“Artificial Intelligence is colossally hyped these days, but the dirty little secret is that it still has a long, long way to go... AI systems tend to be passive vessels dredging through data in search of correlations; humans are active engines for discerning how things work... Unlike human cognition, AI systems lack a theory of the world and how it works.” Marcus (2017).

The truth of the above quotation underlines what we can and what we cannot hope to achieve in applying these methods to mining geomechanics.

What we may achieve by applying AI in mining geomechanics:

1. find correlations among multiple data sets and conditions or events we would like to forecast.
2. create useful statistical models that quantitatively combine multiple input data sets into meaningful output forecasts of future geomechanical behaviour.
3. establish the relative importance of individual data types in understanding future behaviour.
4. confirm or refute assumptions concerning relationships between data, models, and experience and generally put our assumptions of site behaviour to the test of measured facts.

What we will not achieve (anytime soon) by applying AI in rock engineering:

establish new conceptual or physical models that describe rock engineering behaviour.

AI systems easily available to us today are indeed “passive vessels dredging through data in search of correlations. Yet that is of great value in itself in mining geomechanics. It provides us with a new, sophisticated capability to understand underlying patterns in very complex data and apply those patterns as a set of rules that can be used to predict future behaviour based on the patterns of past experience. AI works in any domain by measuring features of a great many examples of something and correlating those features with a condition to be predicted. For example, one could measure features (symptoms) of many individual patients in a medical application and label those patients according to the presence or absence of a specific medical condition. AI techniques could be deployed to comb through thousands of patient records, sort out the relative importance of multiple measured features (symptoms in the example), and create a mathematical model enabling the estimation of the probability of any new patient having the specific condition. AI does this by measuring the important features and combining them according to the learned relationship between the features and the probability of having the condition. The process of uncovering the relationship between measured features and the condition of interest is called training.

By analogy, the example above can be applied to many problems in rock engineering and, by further analogy to the medical diagnostic case, it can be of tremendous practical value to understand the likely existence of a specific condition of importance (e.g., high probability of failure) that can be addressed with practical remediation measures. That remains true whether or not the underlying root causes of the conditions to be remediated are fully understood. Nevertheless, in mining geomechanical applications of AI, unlike AI applications in many other domains, we prefer to use AI algorithms that are not black boxes, but rather reveal as much as possible about relationships among data, models, and outcomes.

Challenge in applying artificial intelligence to mining geomechanics

The central challenge in applying AI to mining geomechanics problems stems from a simple fact: the condition (including rock fall and slope failure) whose location and timing that we want to forecast results from a complex interplay of factors in a four-dimensional, dynamic system that can only partially be known. Capturing the important factors from this complex system for AI training, and subsequent application to new data for providing probabilistic forecasts of where and when conditions of interest may occur, is the key challenge. Meeting this challenge requires deep domain knowledge. It is here where mining geomechanics knowledge enters the AI workflow, and it is where the application of that knowledge to capturing the most meaningful system factors will mark the difference between success and failure.

To give some examples, consider rock bursts in underground mines or slope failures in open pit mines. Rock bursts may be possibly correlated to a host of factors such as depth, stress, stiffness, ground deformation, extraction ratio, production rate and sequencing, support, blasting, span and other mine geometry factors, rock type, rock quality, proximity to geological contacts, proximity to structures, proximity to structural intersections, and orientation of structures with respect to stress and mine geometry. Similarly, slope failures may be possibly correlated to a host of factors such as slope angle, face angle, inter-ramp angle, face height, berm width, rock quality generally, joint characterisation both generally and with respect to wall orientation, water, rock type, proximity to geological contacts, proximity to structures, proximity to structural intersections, orientation of structures with respect to pit geometry, and ground deformation. The factors in play are generally site-dependent; capturing the appropriate ones requires both general and site knowledge.

Co-location in space and time is the most important concept in properly capturing the rock engineering factors that may correlate to the conditions we want to forecast. The AI training algorithms require many examples of multiple measurements on the same thing. In the medical diagnostic analogy, that same

thing is the patient, and the algorithms require many patients on whom multiple factors are measured in addition to noting whether individual patients are afflicted with the condition of interest. In mining geomechanics, it is individual locations in space and time on the rock face that stands in for the patient of the medical analogy. At those individual locations on the rock face (along a drift, in a stope, on a pit wall), many factors can be measured, some of which (e.g., stress, deformation, seismicity) change over time. The data to be assembled for the AI training is of the form:

(x, y, z, t, observation 1, observation 2, ... observation m, condition = true or false).

In AI, this collection of measurements is called a feature vector. It contains the coordinates of the place (x, y, z, t) that specifies a unique location in space and time on the mine, a series of m observations (e.g. RMR, stress) that are observed or estimated at that location, and a condition or target variable that is most commonly a simple binary true or false indicating that the condition being investigated was present or absent at that place and time (for example a rock burst or slope failure). In practice, there are typically many thousands of individual feature vectors and a few tens of observations per feature vector. In fact, the number of feature vectors available to us in the rock engineering domain is virtually unlimited because we are sampling over the mine geometry and time, both of which we may discretise as finely or coarsely as we choose. The number m of observation variables per feature vector is also very much at our discretion, as it is not unusual in AI to include many secondary variables (such as mathematical derivatives to test for significance of both spatial and temporal rates of change) of the primary observed or inferred variables. This expansion of observations in the feature vector by mathematical manipulation such as taking derivatives can be carried out systematically. It is in establishment of the primary observations that the crux of the challenge lies.

Co-location demands that we establish potentially useful quantities relating to each of the primary factors (e.g. rock quality, stress) that we may think have a relationship to the condition being analysed (rock fall, slope failure) at thousands of points (x, y, z, t) in the mine. In practice, this means creating a 4D model of the mine—a 3D model at several or many time steps—that contains all the primary observations believed to possibly have a relationship with the condition of interest. Creating that 4D mine model upon which AI algorithms can be trained to understand the patterns and relationships among data, interpretations, and the history of occurrence of specific events is the central challenge in applying AI methods to mining geomechanics. It is also in constructing the 4D model that rock engineering problems may indeed become big data. The number of data contained in the 4D model that is input to the AI algorithm is $(m \times n)$, where m is the number of observations per feature vector, n is the number of feature vectors (which is the number of digitised points on the mine model multiplied by the number of time steps, a quantity that can easily be in the millions).

I have described the practice and pitfalls associated with the application of AI algorithms to rock engineering problems elsewhere, McGaughey (2019). In the remainder of this paper I focus on the most pressing challenge in the overall workflow, which is construction of the 4D mine model from which the set of feature vectors used as input in AI are derived.

A framework for successful application of ai in rock engineering

A system, Geoscience INTEGRATOR, McGaughey et al (2017), has been created that provides simple computation of the variables required to address the application of AI to mining geomechanics problems, and provides a real, working data-structure definition to the notion of a 4D mine model. It accomplishes this by maintaining 3D earth model and 4D mine model geometrical data structures, upon which multiple data sets are projected, interpolated, upsampled, downsampled, or otherwise processed appropriately for each data type so that the variables of importance for each problem can be co-located in space and time. Documents and files can be stored, managed, and linked to data and models to provide relevant interpretational metadata and contextual links, providing the platform required for AI solutions.

The general system configuration is shown in Figure 1. A 4D data management system sits at the core of the system. The data management system manages all relevant data types, including geological models, mine infrastructure models, drillhole and sample data, production and blasting data, and instrument monitoring data of all types (e.g. convergence and extensometer station time series data, prism and radar data, seismic data). It is able to automatically ingest new data from instruments or external databases. Hazard occurrence or other relevant event conditions are input automatically or manually. Most importantly, the data management system maintains an explicit model of the mine, digitised in time and space, and provides the required mappings between input data streams, the 4D mine model, and output forecasts of rock engineering conditions or events.

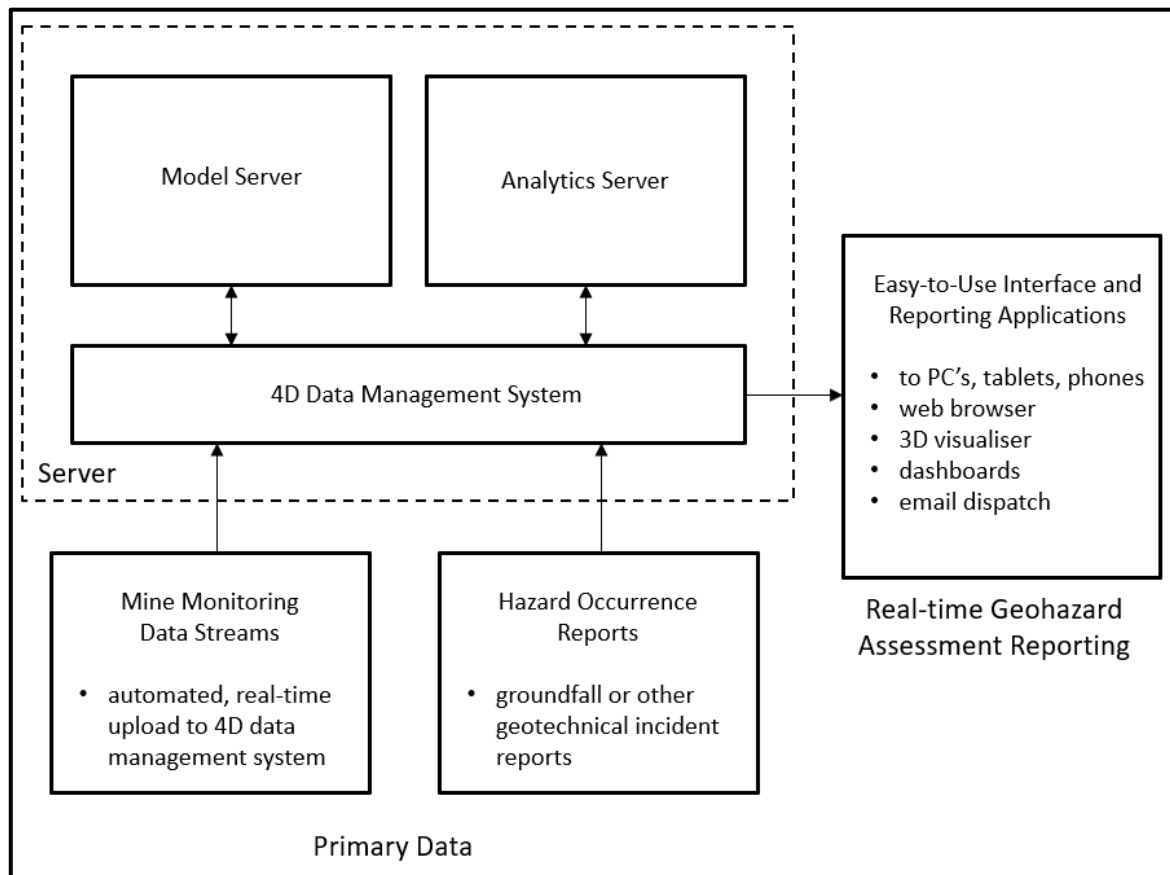


Figure 1. The Geoscience INTEGRATOR system configuration. A 4D data management system resides on a server, connected to a Model Server for automated computation of variables (feature vector observations described in the text) and an Analytics Server for applying AI rules and computing event probabilities

The data management system is directly connected to a Model Server, in this implementation a run-time version of the SKUA-GOCAD® modelling engine, and an Analytics Server which can apply AI rules to new data to deliver updated reports (typically hazard assessment reports). The Model Server is set up to compute required variables automatically, on user demand or on a set schedule (e.g., daily). It operates under the control of the data management system, which queues required computations, supplies the input data, triggers the Model Server to run one of many pre-defined scripts, and receives output as newly computed observations on its internal representation of the mine model at all relevant locations (x, y, z, t) .

Examples of computations that can currently be automatically run by the Model Server to update properties on the mine model (feature vector observations) include:

- interpolate rock quality variables in a block model based on a variety of simple interpolation and geostatistical estimation techniques.

- interpolate time-windowed seismic source properties.
- compute time-windowed seismic event density.
- compute maximum seismic PPV over given time windows.
- compute proximity to contacts and structures.
- compute proximity to intersections of any groups of faults, dykes, geological contacts.
- interpolate ground deformation.
- compute deviatoric stress.
- compute fault-slip tendency.
- compute extraction ratio based on mine infrastructure wireframes.
- compute wedge and planar joint failure parameters using kinematic bench analysis.

An example of the web browser user interface illustrating a sample list of computations set up on an automatic schedule for an actual case study is illustrated in Figure 2.

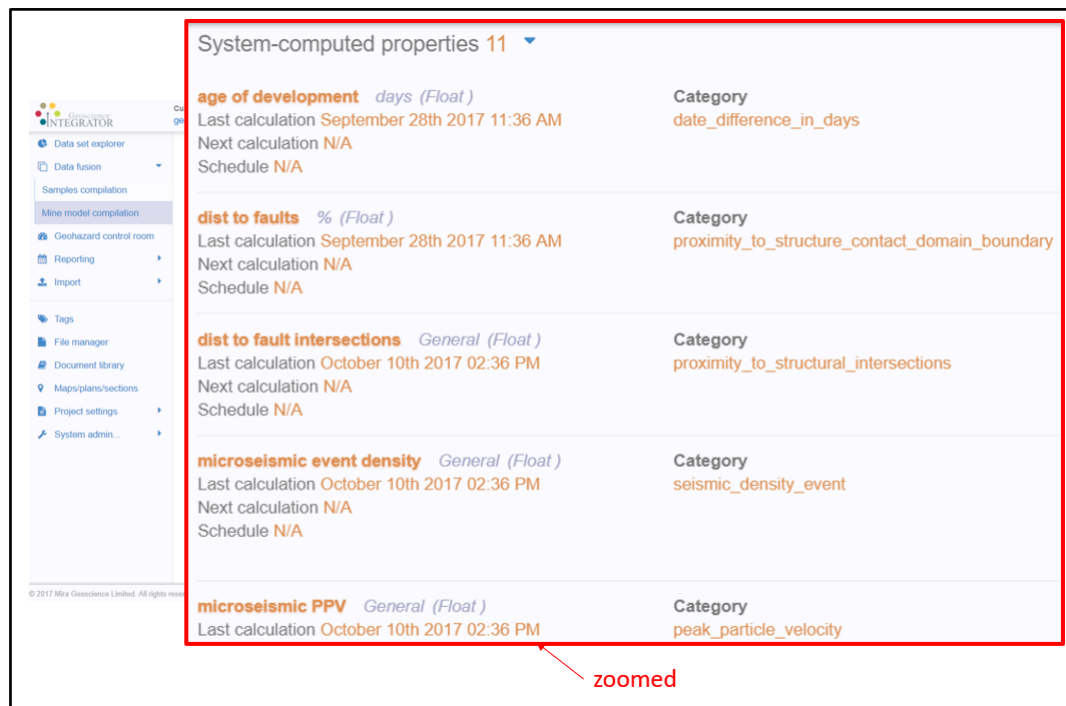
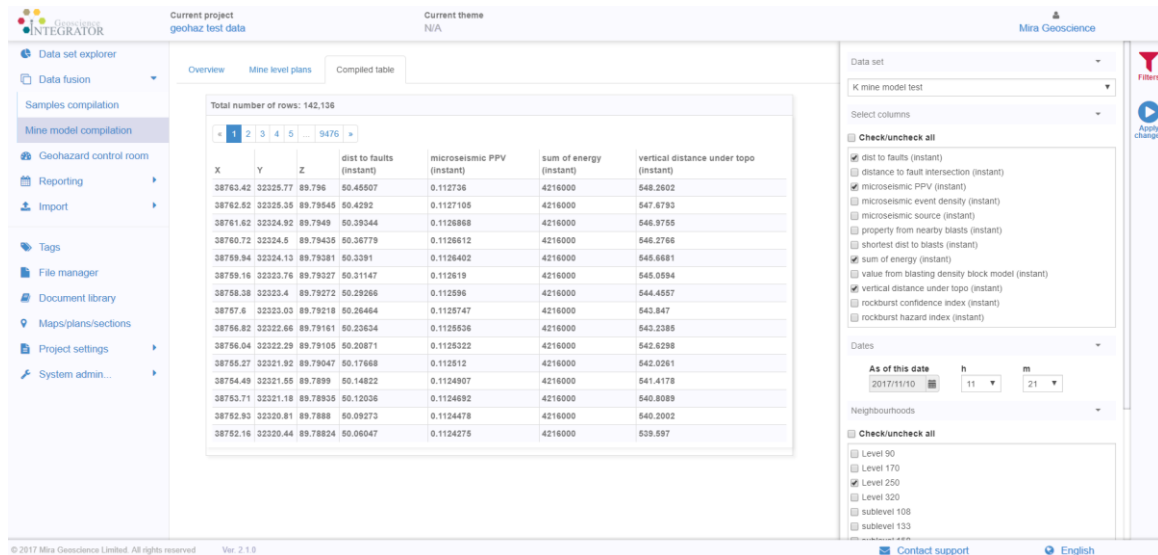


Figure 2. Geoscience INTEGRATOR web interface screen shot illustrating the automated scheduling of several computations used in the application of AI-based geohazard assessment

The computations illustrated in Figure 2 serve to populate the 4D mine model data structure with calculated values for each observation type. The calculations are customised per site to account for the many specific parameters that typically must be set per computation (e.g. length of time windows), as well as the frequency of update per data type.

Figure 3 also shows a screen shot from the system's web browser interface. It is showing a view of its internal data fusion table, which is a tabular display of the values of system-computed observations on individual mine model points (x, y, z) for a given user-selected time t . The rows of this table correspond to individual feature vectors. The complete table is the input to the AI algorithms. The output of the AI algorithms is a probabilistic estimation of the given condition being analysed (e.g., rock fall or slope failure). The output estimation is in an additional, time-varying quantity on each mine model point (x, y, z, t), describing how the probability of manifesting the condition is varying across space and time. Figure 4 is a screen shot from the web browser interface showing a subset of rules, output from the AI algorithm, which are applied to the mine model points to determine, in the particular case study example shown, relative probability of rock burst occurrence across a mine. For the example shown in

Figure 4, the rock burst probability forecast is automatically updated weekly, but the schedule can be arbitrarily set to whatever is appropriate to the mine site. It is important to note that, without such an automated system, updating these computations is extremely laborious. Our experience over the years as consultants initially carrying out these computations manually, was that the computations were sufficiently burdensome that mines would carry out updates typically annually and at most quarterly, essential rendering it a tool for mid-to-long-term planning rather than a tactical operational guide to current areas in the mine that warrant concern.



X	Y	Z	dist to faults (instant)	microseismic PPV (instant)	sum of energy (instant)	vertical distance under topo (instant)
38763.42	32325.77	89.796	80.45507	0.112736	4216000	548.2602
38762.82	32325.26	89.79645	80.4292	0.1127105	4216000	547.8793
38761.62	32324.92	89.7949	80.39444	0.1126868	4216000	548.9758
38760.72	32324.5	89.79435	80.36779	0.1126612	4216000	548.2766
38759.94	32324.13	89.79381	80.3391	0.1126402	4216000	548.6681
38759.16	32323.76	89.79327	80.31147	0.112619	4216000	548.0594
38758.38	32323.4	89.79272	80.29266	0.112596	4216000	544.4557
38757.6	32323.03	89.79218	80.26464	0.1125747	4216000	543.847
38756.82	32322.66	89.79161	80.23654	0.1125536	4216000	543.2385
38756.04	32322.29	89.79105	80.20871	0.1125322	4216000	542.6298
38755.27	32321.92	89.79047	80.17668	0.112511	4216000	542.0261
38754.49	32321.55	89.7899	80.14822	0.1124907	4216000	541.4178
38753.71	32321.18	89.78935	80.12036	0.1124692	4216000	540.8089
38752.93	32320.81	89.7888	80.09273	0.1124478	4216000	540.2002
38752.16	32320.44	89.78824	80.06047	0.1124275	4216000	539.597

Figure 3. Geoscience INTEGRATOR web interface screen shot illustrating the table of feature vectors (also known as the data fusion table) for a case study. Each row of the table corresponds to one feature vector, with only four observations (columns in the table) selected for display. The table is shown for a selected time and area of the mine. The AI algorithms act on the complete table

Figure 5 shows a final, reportable operational output from the system. Once the AI rules (illustrated in Figure 4) are applied, the relative rock burst probability can be displayed as a property on the individual mine model points. The case study example shown is for one mine level only, with relative probabilities above a set threshold shown as large symbols as well as warmer colours to visually stand out. The mine level display can be captured in a PDF report and automatically dispatched on a schedule to a defined email group, or a trigger-alert can be set up if a given threshold is exceeded. All of the underlying variables, as well as the final output hazard assessment result, at each mine model point can be visualised for inspection and validation. All model components, variables, and hazard assessment results can also be easily visualised in 3D using the data management system's 3D visualiser client application (See Figure 6).

In practice, this system can be easily set up at mine site on conventional hardware or as a cloud-hosted deployment (both have been done). Data sizes are manageable with large but not extraordinary demands required on storage capacity. Whether deployed on site or cloud-hosted, Geoscience INTEGRATOR can be connected to multiple data sources at the mine site in several ways. Users can manually update slowly changing data such as mine infrastructure geometry or block models through a manual drag and drop into specified folders on the file network system for automated import. These monitoring folders can also be used for machine-to-machine communication, typically as csv files automatically output from monitoring systems (such as microseismic or ground deformation). The system can also be customised to pull directly from third-party databases (such as production databases). Because all data relevant to the hazard assessment is contained within this single data warehouse, it provides a single point from which to query and access any relevant data. In fact, some mine sites use the system for this data warehouse purpose alone. Figure 7 provides a schematic representation of the data flow.

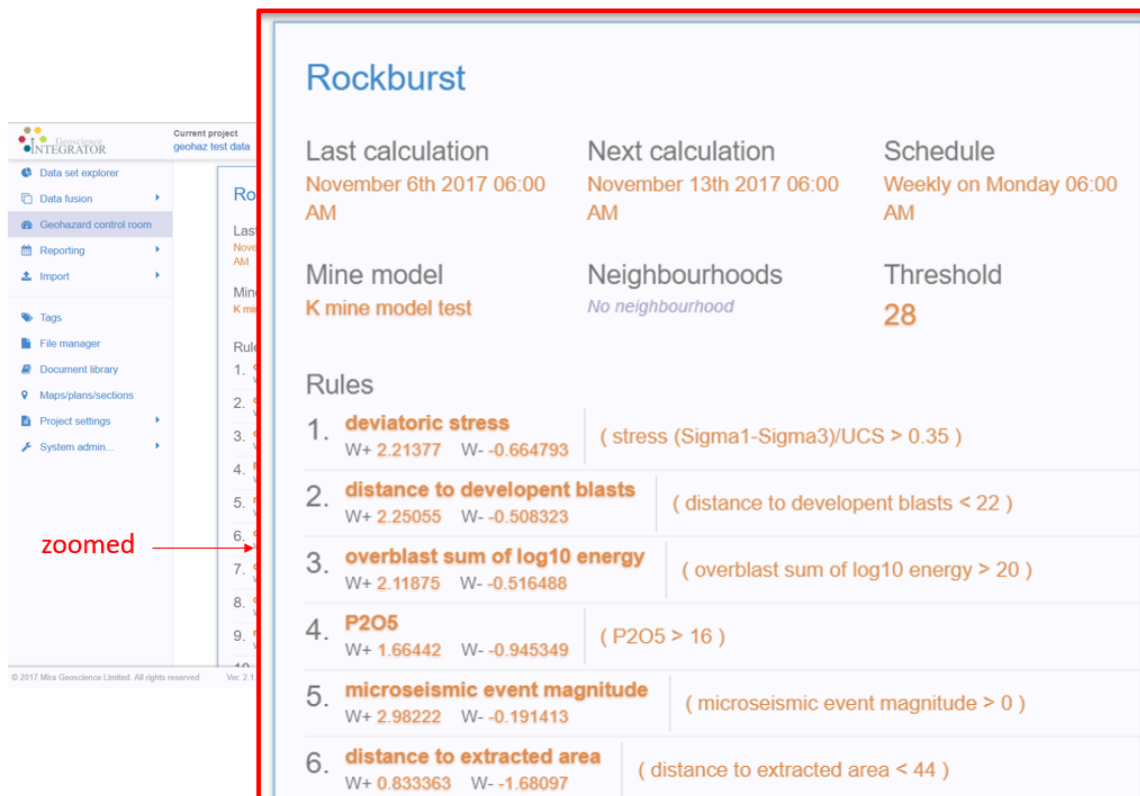


Figure 4. Screenshot from the Geoscience INTEGRATOR web browser interface showing a subset of rules output from an AI algorithm at a case study mine site where the objective was dynamically updating a model of rockburst hazard (in this case weekly), based on an automated update of several input data streams.

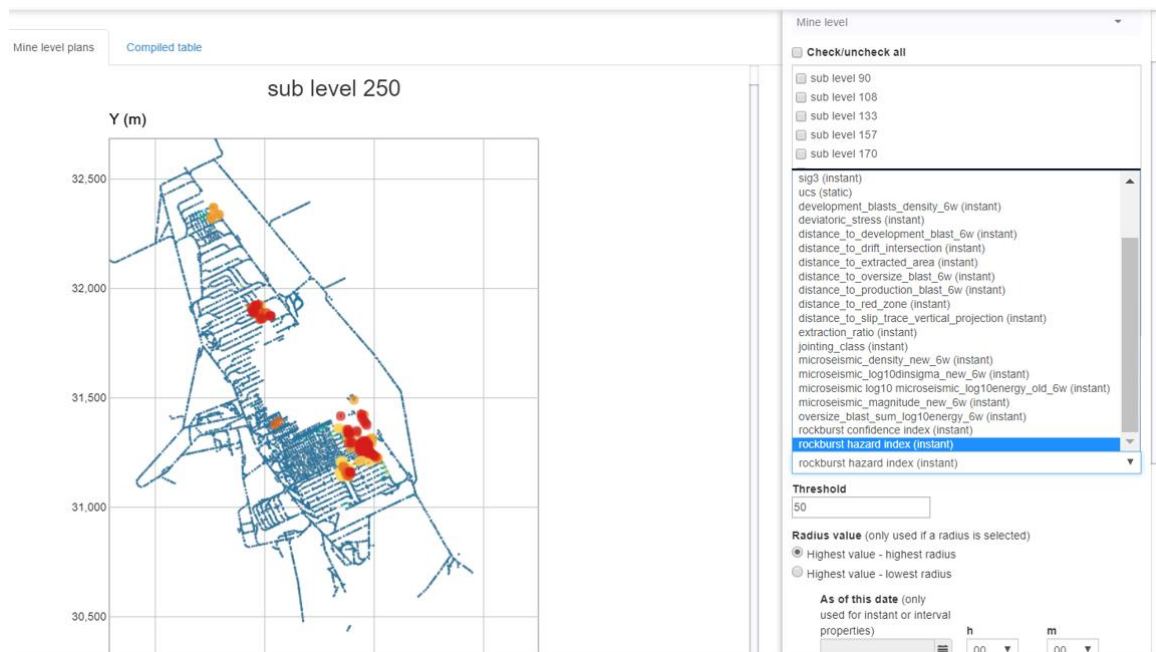


Figure 5. The mine model shown (by level) as the series of digitised points with computed observation values as well as the probabilistic AI-formulated hazard assessment. The display shows rock burst hazard index, as of a certain date and time, with values exceeding a given threshold shown with larger symbols to visually stand out

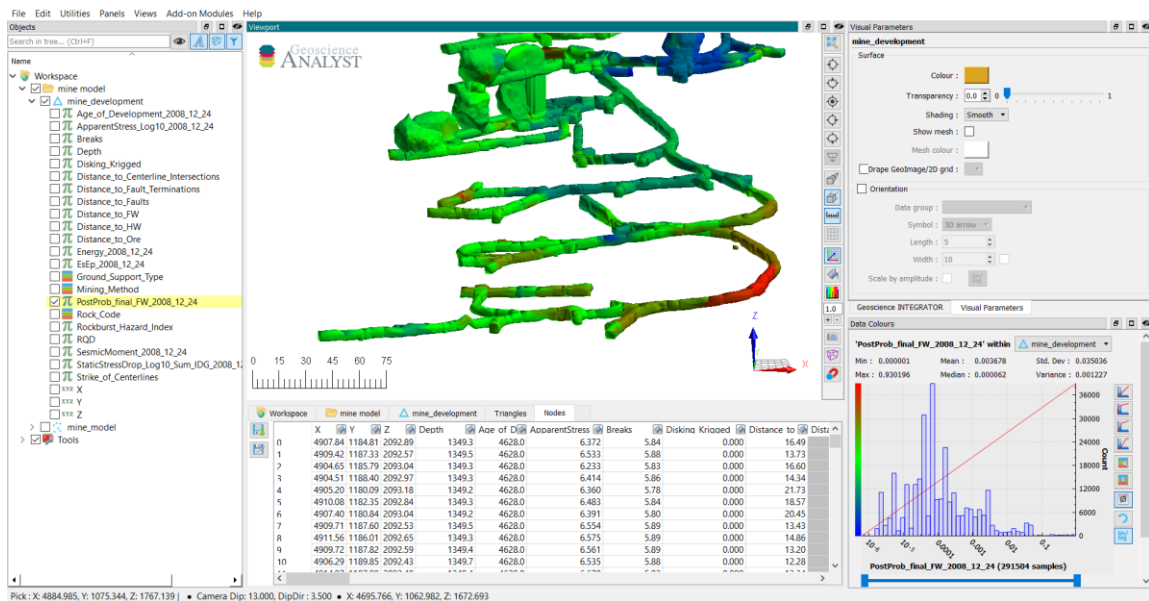


Figure 6. A 3D visualiser client called Geoscience ANALYST connects directly to the Geoscience INTEGRATOR server, enabling 4D query of the data management system to display hazard assessment results (as shown here with warmer colours indicating greater rock burst hazard probability) or any of the underlying data, model components, linked files, documents, and images

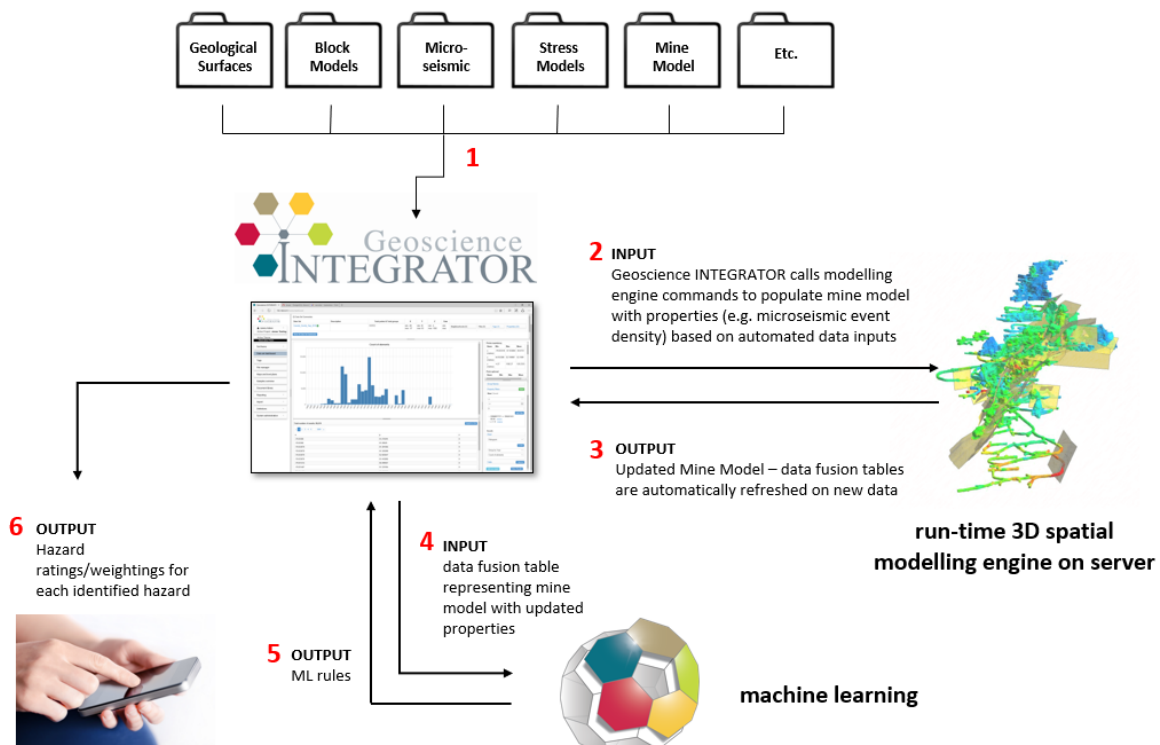


Figure 7. Schematic representation of data flow from independent data sources (top) to the data management system (centre), with connection to the server-based 3D spatial modelling engine for updating the 4D model in response to new data (right), input-output to the analytics (AI or machine learning) system (bottom), and finally hazard assessment reporting (left)

CONCLUSION

AI can be successfully applied to complex mining geomechanics problems. Doing so requires focusing on the primary challenge of setting up the problem rather than focusing on the AI algorithms themselves, most of which will provide value if the problem is properly set up. Developing the proper inputs for AI in rock engineering requires mapping the complex, 4D mine and earth model system to a proper data structure in which the many multi-disciplinary factors in play can be co-located in space and time. Doing so in a practical, operational sense requires implementation of a 4D data management system coupled with a powerful spatial modelling engine (the Model Server) and the AI algorithms (the Analytics Server). Inputs and outputs must be automated to support systematic update at a frequency that is operationally useful for tactical decision making by operators.

ACKNOWLEDGEMENTS

The author is grateful for the support of the Ultra Deep Mining Network (UDMN), administered by the Centre for Excellence in Mining Innovation (CEMI), in Sudbury, Canada, for early financial support in this research prior to its commercialisation; to Glencore for providing the case study R&D site, to JSC Apatit which provided permission to show some of the figures in this paper, and to many Mira Geoscience colleagues who have contributed to the work.

REFERENCES

- Marcus, G. (2017). Artificial Intelligence is Stuck. Here's How to Move it Forward. New York Times, July 29, 2017.
- 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.
- McGaughey, W.J. (2019). 'Data-driven geotechnical hazard assessment: practice and pitfalls', in J Wesseloo (ed.), Proceedings of the First International Conference on Mining Geomechanical Risk, Australian Centre for Geomechanics, Perth, pp. 219-232.



W. John McGaughey

Mira Geoscience Limited

Prior to founding Mira Geoscience in 1999, John spent 10 years at the Noranda Technology Centre in Montreal as a Senior Scientist in their rock mechanics group. He obtained his PhD in geological engineering at Queen's University. John has spent his career working in multi-disciplinary 3D and 4D modelling and interpretation for mineral exploration and mining. He is a pioneer in the application of quantitative, integrated interpretation of geophysics, geology, and engineering data for addressing difficult technical challenges in both exploration and geotechnical engineering.