MDAG.com Internet Case Study 68

The Rapid Evolution of Machine Learning Applied to Minesite-Drainage Chemistries and Flows

by K.A. Morin

www.mdag.com/case_studies.html

Abstract

Can minesite-drainage chemistries and flows be predicted reasonably well without any models or underlying assumptions? At first, this question sounds nonsensical. However, machine learning (ML), particularly artificial neural networks (ANN), can do this in some situations. Ripley (2004) succinctly stated,

> "To paraphrase provocatively, 'machine learning is statistics *minus* any checking of models and assumptions'."

Some partial success was achieved in the past two decades with ML of minesite drainage. However, in 2020 and 2021, the National Research Council of Canada published landmark studies showing that seasonal and long-term trends in both flow and aqueous acidity from full-scale waste rock could be reasonably predicted with an ANN. Despite the complex processes affecting full-scale flow and acidity, the predictions were based primarily on previous and current weeks of data on easily-measured mean temperature and total precipitation, plus "time tags" of when the data were collected.

Despite the worldwide hype, ML has now evolved enough that it is suitable for simulating and predicting minesite-drainage chemistries and flows.

Table of Contents

1. Introduction

Machine learning (ML) and the more generic category of artificial intelligence (AI) are often the subjects of exaggeration and hype (e.g., Witten et al., 2011; Chollet, 2018), which occur in cycles spanning many years. We are currently in one of the "upswings" of that hype.

There are many books and papers written about ML in its various forms; far too many to be listed and summarized here. Nevertheless, here is one example of a simple explanation:

> "Machine learning, in its simplest explanation, involves allowing a computer to vary its responses and introducing a feedback loop for good and bad responses. This means that machine learning algorithms are fundamentally different from the computer programs that have come before them." (Norman, 2018).

Since ML algorithms are fundamentally different from past computer programs, are they automatically more reliable? Emphatically, no:

> "The fundamental goal of machine learning is to generalize beyond the examples in the training set . . . The most common mistake among machine learning beginners is to test on the training data and have the illusion of success. If the chosen classifier is then tested on new data, it is often no better than random guessing." (Domingos, 2012).

How is ML fundamentally different?

> "To paraphrase provocatively, 'machine learning is statistics *minus* any checking of models and assumptions'." (Ripley, 2004)

No models or assumptions?!

As explained below in this MDAG Case Study, I think another explanation for ML, relative to past deterministic and stochastic modelling of minesite drainage, is: "let the ML algorithm substitute for the calibration factors of selected models and assumptions".

2. Traditional Modelling of Minesite-Drainage Chemistry and Flow

As a general statement, the modelling of minesite-drainage chemistry and flow often follows two basic steps:

(1) selecting a model that is thought to approximate a site and

(2) adjusting and calibrating that model until it matches inputted site data.

I have been involved with such modelling for 40 years, involving deterministic, stochastic, and statistical models (e.g., Morin, 1983; Morin and Cherry, 1988; Morin and Hutt, 1995; Morin et al., 2001; Betrie et al., 2012; Betrie et al., 2015a and 2015b; Morin, 2018; Liu et al., 2018a and 2018b and 2019; Ma et al., 2019). Thus, I believe my opinion here on Machine Learning counts for something. Many other people have also been involved in such modelling of minesite drainage over the decades, as referenced in many previous MDAG Case Studies (e.g., Morin, 2017a, 2017b, and 2017c).

Many models have been successfully calibrated to, and applied to, many minesites around the world. Is this evidence that the selected model was applicable to that specific site? Is it possible to calibrate more than one model to a particular site even when they literally are not appropriate?

My favourite example is the shrinking-core model for pyrite oxidation, acid generation, and acid rock drainage (ARD). This model has been used for, and successfully calibrated to, many studies of minesite drainage for decades (e.g., Davis and Ritchie, 1986; Wunderly et al., 1996; Hecht et al., 2002; Singh and Doulati Ardejani, 2003; Doulati Ardejani and Singh, 2004; Wang et al., 2019). The basic concept is that, as pyrite oxidizes, a progressively thickening outer "rind" of reaction products forms. As a result, the rate of pyrite oxidation decreases through time as the rind thickens and the unreacted inner core of pyrite shrinks. This is due to the increasing distance that oxygen has to diffuse through to reach the remaining pyrite and for the reaction products to diffuse outwards.

The major problem is that the shrinking-core model is wrong for most pyrite oxidation, despite many publications and studies of minesite drainage showing its successful application. This is obvious from the basic equation for pyrite equation, such as:

Pyrite + oxygen (from gas and/or water) + water (from liquid or vapour) =
Secondary minerals like ferric-iron oxyhydroxides and sulphates like gypsum + heat

When secondary minerals form, the solid-phase volume associated with each mole of Fe and S from pyrite increases substantially by taking O, H, and perhaps other elements from the liquid and gaseous phases. As a result, the reacted and oxidized rind "blows up like a balloon" and breaks away, leaving unreacted pyrite directly exposed to air and water. Other researchers understand this (e.g., Jerz and Rimstidt, 2003).

If the rind is not completely removed, then the extremely slow rate of diffusion surely would not apply to a fractured rind. Such abuse of diffusion in minesite-component studies was discussed in MDAG Case Study 67 (Morin, 2021).

Thus, the shrinking-core model is not appropriate for most pyrite oxidation. Yet, it reportedly simulates site data well after calibration according to many publications and researchers.

Therefore, let's take the stance that, upon calibration, we can potentially make more than one model fit a particular site even when they do not reliably apply to that site. In this case, the quotation above in Section 1 is worth repeating.

> "To paraphrase provocatively, 'machine learning is statistics *minus* any checking of models and assumptions'." (Ripley, 2004)

No models or assumptions!

To foreshadow the following sections, one popular type of ML is an Artificial Neural Network (ANN). In simple words, an ANN includes various numbers of "layers" (Figure 2-1 is a simple example of an ANN). The nodes in one layer are mathematically connected in multiple ways to nodes in adjacent layers using algorithm-calculated "weighting factors". Iteration through the ANN, over and over again, can further adjust each weighting factor between each node in each layer. The effort continues until the ANN is optimized and "trained", but not "overtrained", based on testing and then validation with withheld data.

Another way I think of an ANN is an extreme multivariate analysis that goes far beyond what most humans can envision and calculate by hand or spreadsheet. It is too complex to calculate the weighting factor for one input variable, because there could be hundreds of nodes and hundreds of individual weighting factors contributing to that one output variable.

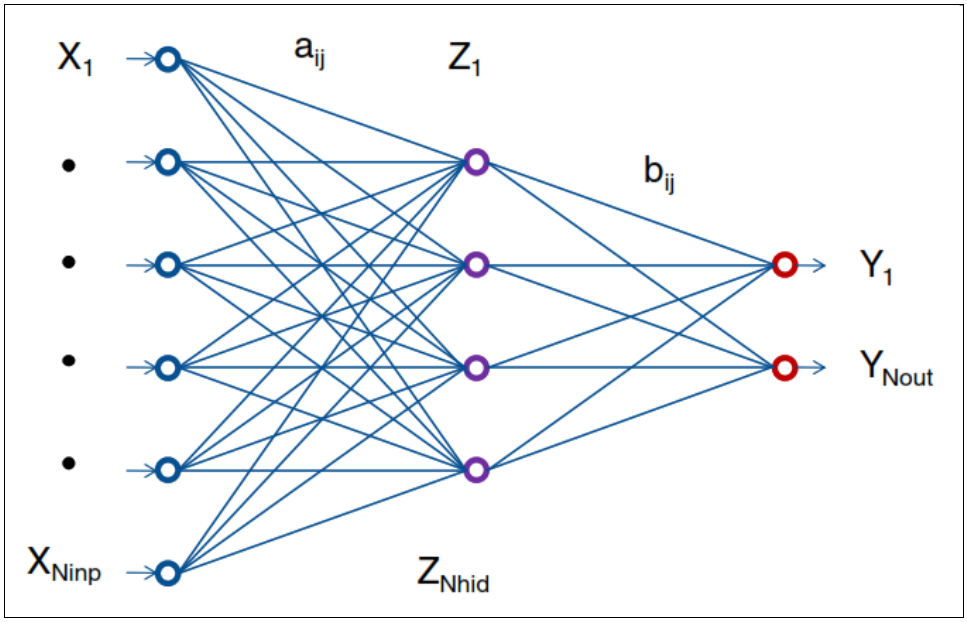

Figure 2-1. An example of a simple three-layer Artificial Neural Network (ANN) from Betrie et al. (2014); the "X" nodes are the input and the "Y" nodes are the output; "a" and "b" are the individual weighting factors between each layer.

For example, in Figure 2-1, there is no one single weighting factor to predict Output "Y1" from any of the "X" inputs when each line is traced from "Y1" to the left through two layers. In other words, there is no weighting factor that can be applied to Input X1 to get Output Y1, because (1) there are several sets of weighting factors ("a" and 'b") that connect X1 and Y1 and (2) the other Inputs also contribute to Y1. It is a complex approach that is not easily summarized or quantified, but that does not mean it is false and unreliable. This complex approach has only become available because of ongoing increases in computer speed and storage.

3. Summary of Machine Learning of Minesite Drainage, Prior to 2020

Machine learning (ML) has been applied to some aspects of mining, like ore-zone detection and delineation, for decades. However, ML applied to time series of minesite-drainage flow and chemistry, water contamination, and ARD, has been more limited, but not absent. The following prominent examples show how the application of ML to minesite drainage has evolved rapidly in the last 20 years.

Based on 40 water analyses from "simple field tests", Khandelwal and Singh (2005) used the ML technique of artificial neural networks (ANN; see Figure 2-1 above) for comparison to multivariate regression analysis. The ANN was trained with the 30 water analyses, and another 10 analyses were held back for testing, "validation", and prediction by the trained ANN. Input parameters were aqueous pH, temperature, and aqueous hardness; output parameters were aqueous sulphate, chloride, chemical oxygen demand, total dissolved solids, and total suspended solids.

Khandelwal and Singh (2005) reported the ANN predicted the 10 held-back analyses better than the multivariate regression analyses, although the comparisons in their Figures 12 to 16 are difficult to confirm. In any case, please note that the aqueous parameters of pH and hardness were used to predict other aqueous parameters. Aqueous pH is itself an important parameter that often has to be predicted and then affects predictions of other aqueous parameters. In other words, some of the water chemistry had to be known at a particular time step to estimate the remainder at that same point. Also, flow rates were not predicted, possibly because these analyses came from "simple field tests" and not full-scale minesite components.

Gholami et al. (2011), Rooki et al. (2011), and Aryafar et al. (2012) used support vector machine (SVM) and ANN for ML and compared the results to a general regression neural network. This work was based on 44 analyses of river water downstream of the Sarcheshmeh Copper Mine in Iran. The aqueous inputs from the river were pH, sulphate, and magnesium; predicted outputs were aqueous copper, iron, manganese, and zinc which had good correlations with the input parameters. SVM predicted the output concentrations better than the regression network, but ANN could be optimized to provide good predictions. As with Khandelwal and Singh (2005), some current water chemistry was needed as input in order to predict current aqueous concentrations of other elements. Also, flow rates were not predicted.

The previous examples from 2005 and 2012 used a few dozen water analyses, with three input parameters and four to five output parameters. Aqueous concentrations were taken from "simple field test" or diluted river water. In these cases, even visual examinations with graphical scatterplots can show correlations with relatively little effort.

In contrast, those visual and simple approaches are not always feasible with many hundreds to thousands of water analyses containing up to a dozen parameters collected near full-scale minesite components. These "big data" are where ML algorithms can perform well and can perform better than simpler visual and regression approaches, which leads to the next examples.

Betrie et al. (2013)[1] compared five linear ML techniques for predicting hundreds of aqueous copper concentrations at Island Copper Mine in British Columbia, Canada. These five techniques were: support vector machine with polynomial (SVM-Poly), support vector machine with radial base function (SVM-RBF) kernels, multilayer artificial neural networks (ANN), model tree (M5P), and K-nearest neighbors (K-NN). For each technique, 100 independent assessments of model errors were made for each ML technique. These showed the best predictions were made in the preceding order of the five techniques, with the first three providing similar error residuals.

Betrie et al. (2013) purposely avoided using strongly cross-correlated input parameters in the ML techniques. The input parameters were current-time aqueous pH, conductivity, acidity, and previous-time copper concentrations, and the output was current-time aqueous copper concentrations. Thus, knowledge of past and current aqueous concentrations of some parameters was needed to predict current aqueous copper. This was also pointed out in the previous examples as a weakness, because very recent examples are coming up (Section 4) that do not require current chemistry and flow in order to predict current chemistry and flow.

Betrie et al. (2013) made the following observations.

> "However, the predicted results [would] likely improve if more parameters (e.g., flow rate, internal gases concentrations, and temperature) are considered as more data become available in the future."
>
> "These results indicate that the process of AMD generation is highly nonlinear and could not be captured with techniques that build local linear models such as the M5P and K-NN techniques."
>
> "These [SVM and ANN] techniques take considerable time for training model since they have parameters to be optimized and the optimization is heuristically done."
>
> "Another limitation of ANN is that the function, which represents a given ANN model, is presented by interconnection weights and threshold values and is not easily understandable by decision-makers."

Section 4 below will confirm these observations. It is correct that (1) additional or alternative parameters are important, (2) minesite drainage is strongly nonlinear and thus complex, (3) ANN and SVM techniques require considerable time to train and optimize, and (4) the complexity of a neural network in an ANN is not easily explained or understood. The last one means that an ANN could be considered a "black box model" that is not easily critiqued and trusted (see the end of Section 2 above). However, it is this ANN complexity that is apparently needed to predict complex and nonlinear minesite-drainage chemistry and flow (Section 4).

Based on additional work after Betrie et al. (2013), Betrie et al. (2014)[1] showed that ANN performed better than SVM techniques. This latter study focussed on adding probability-based uncertainty (sort

[1] Disclosure: I contributed to a minor portion of this work.

of like error bars) to ML-based predictions. The input aqueous parameters were current-time aqueous pH, alkalinity, sulphate, acidity, and flow to predict current-time aqueous copper, cadmium, and zinc. Although flow did not show good correlations with aqueous chemistry, it was included because the ANN technique could still find some connection.

' ! " ! " (! ")*)*

! "#$% #$%&'
(( )
" ((* # ' + " + ')

,
) *

) #$#$ #$#% - . (/ - "
#$#$ % #$#% % 0) (/-
"%'
" #$#$ ' "#' " #$#%'

, ((
"%') 1 "#' 2
(

! % #

3
4
53 6)
. 7 "/ #$$&'
"#$$$ #$$%' 3)
)

& % 8 0 / 9 /)

0)
)
))
3) , .
) : "#$%;
' "#$%< #$#$ '

9 "#$$$ ')
=>$ %&$)
? ! - - 3
" '

week, and over six weeks. The focus was on Monitoring Station BD at Equity Silver with more than 2000 measurements of flow since 1998.

Because minesite-drainage chemistries and flows can display seasonality and long-term trends, Ma et al. (2020a) "refined" the ANN by adding two additional inputs from the monitoring data: month and year. These "time tags" turned out to be important because the ANN was then enabled to simulate and predict seasonal variations and long-term trends in flow rates.

Again, despite the known complexity and time variations of flow through and from a full-scale waste-rock pile, Ma et al. (2020a) limited input to precipitation, air temperature and time tags. No existing model known to me would predict basal flows from full-scale waste rock based only on these input parameters.

To start the training, only 70% of the ~2000 measured flows since 1998 were used to train the ANN; 15% were later used to test and improve the ANN; and the remaining 15% were used to determine the predictive performance of the ANN. Additionally, the last year of flow data was held back to further evaluate the predictive performance for flow by the ANN. Thus, there were two independent tests of predictive performance.

The best predictive performance of the ANN for flow rate came from the weekly averages for the preceding and current six weeks of (1) mean air temperature and (2) total precipitation. Daily mean air temperature and precipitation can be measured quickly and cheaply at minesites (and often are), whereas chemistry and flow are not. To be clear, the previous and current six, weekly averages of mean air temperature and total precipitation with time tags, combined with the complex weighting scheme of an ANN (see Figure 2-1 above), yielded reasonable predictions of (1) held-back monitoring data and (2) one year of flows not included in the training, testing, and validation (e.g., Figure 4-1).

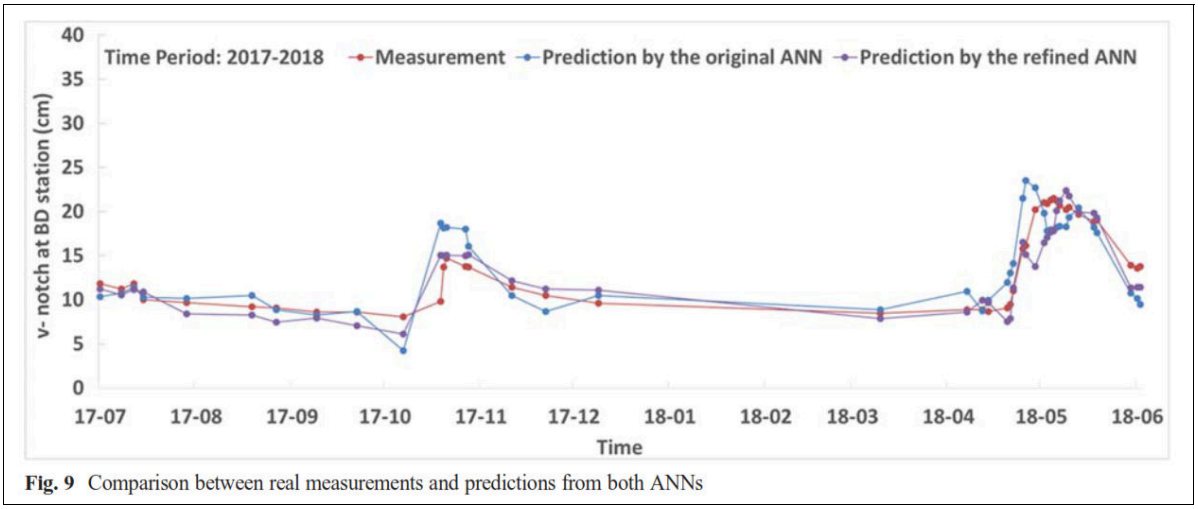

Figure 4-1. Comparisons of (1) measured data for flow rates not included in the training and validation, (2) the ANN without time tags, and (3) the refined ANN with time tags (from Ma et al., 2020a).

Ma et al. (2020a) concluded:

> "The proposed ANN approach is more advanced than many traditional mining-related models as it requires no simplification and assumption for hydrogeological mechanisms in waste rock piles before calculations. Even though the correlations are naturally sophisticated in field conditions, they can be captured by the advanced ANN approach through training by pairs of input and output data."

As explained in Section 2 above, any and all models have to be adjusted and calibrated to the site being studied. There is no type of generic model applicable unadjusted and uncalibrated to all minesites. ANN is the same, with this ANN for Monitoring Station BD at Equity Silver probably not directly applicable to other minesites.

4.2 Drainage Chemistry from Waste Rock

Ma et al. (2020a) recognized that aqueous chemistry of full-scale minesite drainage rarely correlates well with full-scale flow rates. This led to the second, geochemical study of Ma et al. (2021) of the same Equity Silver waste-rock pile.

The interacting physical, geochemical, and biological processes that create aqueous concentrations in minesite drainage are far more complex than those affecting flow rates. Thus, there was more uncertainty on whether an ANN could approximate any aqueous parameter. The primary target was acidity, which is a composite of several aqueous parameters and complexes, and which reflects costs of water treatment and environmental mitigation.

With a relatively large monitoring database, Ma et al. (2021) observed:

> "Compared to the traditional approach, the advantage of analyzing the [monitoring] database is that it generally requires much less simplification and fewer assumptions, because these monitoring data already inherently contain the information of all the underlying hydrogeochemical and biochemical mechanisms in the field."

In this geochemical study by Ma et al. (2021):

- A long-short-term memory unit was added to the ANN that remembers values only over certain time intervals.
- Liu et al. (2019) found that the snow season (November to April) and the rain season (April to October) had notably different patterns in drainage chemistry, making it difficult for traditional data analysis to define an annual correlation.
- The "time tags" were reduced to only the accumulating number of days to assist in simulating and predicting long-term temporal trends.
- The two main input parameters of mean air temperature and total precipitation used for flow were used again in this geochemical study, despite there being no good correlation between flow and chemistry and despite no traditional geochemical model focussing on total precipitation and air temperature (Section 2).

In this geochemical study, 80% of the monitoring data was used for training and 20% used for validation. Two years of subsequent data not included in the training and validation were used to assess predictive performance.

For those two last years not included in the training and validation, the previous and current 10 weeks of non-geochemical mean temperature and total precipitation simulated reasonably well the trends and values of aqueous acidity at the monitoring station (Figure 4-2). The close predictions in Figure 4-2 are rarely surpassed by traditional modelling (Section 2). In fact, the match in Figure 4-2 generally substantially surpass that from traditional modelling that does not include seasonality and long-term trends.

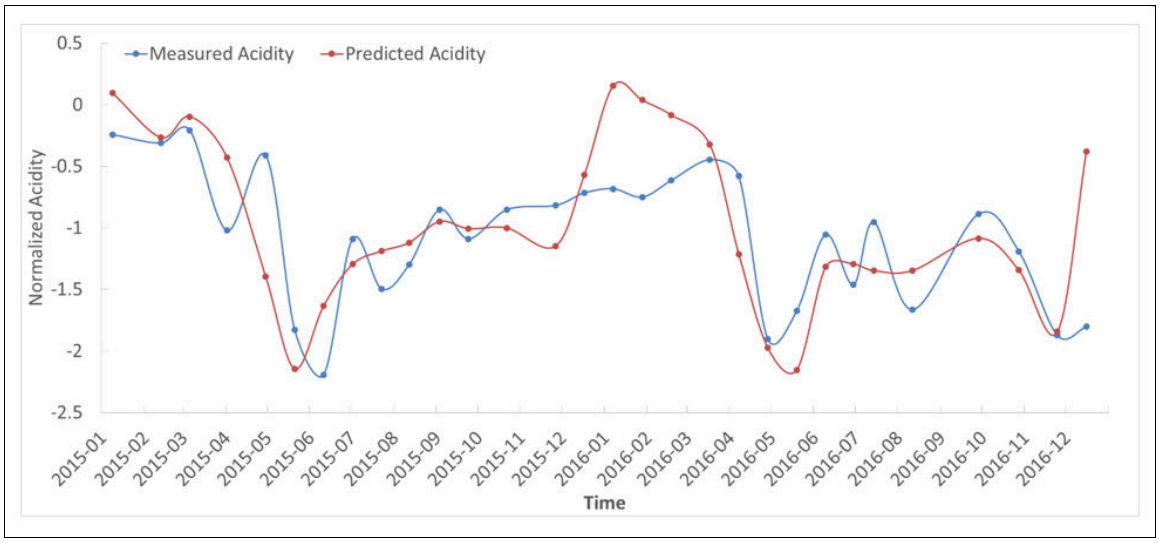

Figure 4-2. Comparisons of (1) measured acidity for two years not included in the training and validation, (2) the ANN without time tags, and (3) the refined ANN with time tags (from Ma et al., 2021).

Ma et al. (2021) concluded,

> "The advantage of this machine learning based approach generally requires no pre-defined simplifications and assumptions for hydrogeochemical and biochemical mechanisms inside full-scale waste rock piles."

5. Conclusion

This MDAG Case Study summarized "traditional" approaches for modelling of minesite-drainage chemistries and flows based on deterministic, stochastic, or statistical equations. The equations typically represent a user-chosen subset of the applicable, myriad, interacting physical, geochemical, and biological processes at full-scale minesites.

These traditional models would be acceptable if the underlying model, equations, and assumptions can be confirmed as reliably applying to a particular site. However, it appears likely that more than one model can often be applied to a site through calibration. In fact, the shrinking-core model for pyrite oxidation does not apply to most situations, despite decades of minesite-drainage models reporting success in using it.

This leads to the question of whether minesite-drainage chemistries and flows can be predicted reasonably well without any models or underlying assumptions? At first, this question sounds nonsensical. However, machine learning (ML), particularly artificial neural networks (ANN), can do this in some situations. Ripley (2004) succinctly stated,

> "To paraphrase provocatively, 'machine learning is statistics *minus* any checking of models and assumptions'."

Between 2000 and 2019, progress was made on the ultimate goal of predicting minesite-drainage chemistries using ML techniques. However, these studies showed that some current aqueous concentrations were needed to predict other current concentrations, and flow was not predicted.

In 2020 and 2021, the National Research Council of Canada (NRC) published landmark studies showing that seasonal and long-term trends in flows and aqueous acidity from full-scale waste rock could be reasonably predicted with an ANN. Despite the complex processes affecting full-scale flow and acidity, the predictions were based primarily on previous and current weeks of data on mean temperature and total precipitation, plus "time tags" for when data were collected. Daily temperature and precipitation are easily and cheaply measured at minesites, and often are, whereas flows and aqueous chemistries are not, providing additional value to the NRC results.

This initial groundbreaking work by NRC Canada will likely be built on and expanded in the future.

6. References

Aryafar, A., R. Gholami, R. Rooki, F. Doulati Ardejani. 2012. Heavy metal pollution assessment using support vector machine in the Shur River, Sarcheshmeh copper mine, Iran. Environmental Earth Sciences, 67, p. 1191-1199.

Betrie, G.D., R. Sadiq, K.A. Morin, and S. Tesfamariam. 2015. Ecological risk assessment of acid rock drainage under uncertainty: The fugacity approach. Environmental Technology and Innovation, 4, p. 240-247.

Betrie, G.D., R. Sadiq, C. Nichol, K.A. Morin, and S. Tesfamariam. 2015. Environmental risk assessment of acid rock drainage under uncertainty: The probability bounds and PHREEQC approach. Journal of Hazardous Materials, 301, p. 187-196.

Betrie, G.D., R. Sadiq, K.A. Morin, and S. Tesfamariam. 2014. Uncertainty quantification and integration of machine learning techniques for predicting acid rock drainage chemistry: A probability bounds approach. Science of the Total Environment, 490, p. 182-190. DOI: 10.1016/j.scitotenv.2014.04.125

Betrie, G.D., S. Tesfamariam, K.A. Morin, and R. Sadiq. 2013. Predicting copper concentrations in acid mine drainage: A comparative analysis of five machine learning techniques. Environmental Monitoring and Assessment, 185, p. 4171-4182. DOI: 10.1007/s10661-012-2859-7 [Note: this paper was first published online in 2012 and is thus sometimes referred to as Betrie et al., 2012]

Betrie, G.D., S. Rehan, K.A. Morin, S, Tesfamariam. 2012. Uncertainty analysis of an aquivalence-based fate and transport model using a hybrid fuzzy-probabilistic approach. IN: Proceedings of the 12th International Environmental Specialty Conference, Canadian Society for Civil Engineering, Edmonton, Alberta, June 6-9, 2012.

Chollet, F. 2018. Deep Learning with Python. Manning Publications Co., New York. ISBN 9781617294433.

Davis, G.B., and A.I.M. Ritchie. 1986. A model of oxidation in pyritic mine wastes, 1, Equations and approximate solutions. Applied Mathematical Modelling, 10, p. 314-322.

Doulati Ardejani, F., and R.N. Singh. 2004. Two-Dimensional simulation of pyrite oxidation and pollutant transport in backfilled open cut coal mines. IN: International Mine Water Association Symposium, Proceedings of the Symposium: Mine Water 2004- Process, Policy and Progress, University of Newcastle, Newcastle upon Tyne, United Kingdom, p. 1-20.

Domingos, P. 2012. A Few Useful Things to Know about Machine Learning. Communications of the ACM, 55, October 2012. Accessed at: https://sites.astro.caltech.edu/~george/ay122/cacm12.pdf DOI: 10.1145/2347736.2347755

Gholami, R., A. Kamkar-Rouhani, F. Doulati Ardejani, and Sh. Maleki. 2011. Prediction of toxic metals concentration using artificial intelligence techniques. Applied Water Science, 1, p. 125-134. DOI: 10.1007/s13201-011-0016-z

Hecht, H., M. Kölling, and N. Geisler. 2002. DiffMod7 – Modelling Oxygen Diffusion and Pyrite Decomposition in the Unsaturated Zone Based on Ground Air Oxygen Distribution. IN: Geochemical Processes: Conceptual Models for Reactive Transport in Soil and Groundwater, Research Report, p. 55-78. Wiley-VCH. ISBN 978-3-527-27764-3.

Jerz, J.K., and J.D. Rimstidt. 2003. Efflorescent iron sulfate minerals: Paragenesis, relative stability, and environmental impact. America Mineralogist, 88, p. 1919-1932.

Khandelwal, M., and T.N. Singh. 2005. Prediction of mine water quality by physical parameters. Journal of Scientific

and Industrial Research, 64, p. 564-570.

Liu, S., L. Ma, E. Dy, Z. Xie, K. Tufa, E. Fisher, J. Zhou, K. Morin, M. Aziz, C. Meints, M. O'Kane, L. Tallon. 2019. The Characteristic Properties of Waste Rock Piles in Terms of Metal Leaching. Journal of Contaminant Hydrology 226, published online August 2019.

Liu, Z-S, C. Huang, L. Ma, E. Dy, Z. Xie, K. Tufa, E. Fisher, J. Zhou, K. Morin, M. Aziz, C. Meints, M. O'Kane, L. Tallon. 2018a. Rate-control quotient of mineral dissolution from waste rock dumps. IN: 41st Annual BC Mine Reclamation Symposium, Williams Lake, British Columbia, Canada, September 17-20.

Liu, Z-S, C. Huang, L. Ma, E. Dy, Z. Xie, M. Aziz, C. Meints, K. Morin, M. O'Kane, L. Tallon. 2018b. Experimental models of metal leaching for scaling-up to the field. IN: CIM 2018 Convention, Presentation 2188, Vancouver, British Columbia, Canada, May 6-9, Canadian Institute of Mining, Metallurgy and Petroleum.

Ma, L., C. Huang, Z-S Liu, K.A. Morin, M. Aziz, and C. Meints. 2021. The correlation between drainage chemistry and weather for full-scale waste rock piles based on artificial neural network. Journal of Contaminant Hydrology, May 2021, 103793. DOI: 10.1016/j.jconhyd.2021.103793.

Ma, L., C. Huang, Z-S. Liu, K.A. Morin, M. Aziz, and C. Meints. 2020a. Artificial neural network for prediction of full-scale seepage flow rate at the Equity Silver Mine. Water, Air, and Soil Pollution 231:179. Published online 14 April 2020. DOI: 10.1007/s11270-020-04541-x.

Ma, L., Z-S Liu, E. Dy, K. Tufa, J-Q. Zhou, E. Fisher, M. Goulet, C. Huang, and K. Morin. 2020b. Additional Mechanisms and Measurements to Improve ARD/ML Prediction. IN: 2020 BC MEND ML/ARD Virtual Workshop, Challenges and Best Practices, December 1-3.

Ma, L., C. Huang, Z-S. Liu, K.A. Morin, M. Aziz, and C. Meints. 2019. Prediction of Acid Rock Drainage and Metal Leaching in Waste Rock Piles. Part 1: Water Film Model for Geochemical Reactions and Application to a Full-scale Case Study. Journal of Contaminant Hydrology, 220, p. 98-107

Ma, L., C. Huang, Z-S. Liu, K. Morin, M. Aziz and C. Meints. Submitted. Prediction of Acid Rock Drainage in Waste Rock Piles. Part 2: Water Flow Patterns and Leaching Process.

Morin, K.A. 2021. Diffusion-Dominated Flux of Poregases and Aqueous Ions in Minesite Components - Commonly Modelled but Rarely Applicable Because Realistic Fluxes are Much Faster. MDAG Internet Case Study #67, www.mdag.com/case_studies.html

Morin, K.A. 2018. The Third Way: Worthy of Research and Evaluation. Presented at the International Canada-Germany-UK ARD-ML Workshop, National Research Council Canada, Vancouver, Canada, December 2018. MDAG Internet Case Study #58, www.mdag.com/case_studies/cs58.html

Morin, K.A. 2017a. A Case Study Applying First-Passage Percolation Theory and Reservoir Modelling to Full-Scale Waste-Rock Drainage after a One-Day Dynamic Storm. MDAG Internet Case Study #52, www.mdag.com/case_studies/cs52.html

Morin, K.A. 2017. A Case Study of High-Frequency Time Series for Individual Chemical Elements in Waste-Rock Drainage. MDAG Internet Case Study #46, www.mdag.com/case_studies/cs46.html

Morin, K.A. 2017c. A Case Study of Rapid Water Flow through Full-Scale Waste-Rock Piles. MDAG Internet Case Study #45, www.mdag.com/case_studies/cs45.html

Morin, K.A. 1983. Prediction of Subsurface Contaminant Transport in Acidic Seepage from Uranium Tailings Impoundments. Ph.D. Thesis, Department of Earth Sciences, University of Waterloo, Ontario. Supervisor: Dr. John A. Cherry.

Morin, K.A., and J.A. Cherry. 1988. Migration of acidic groundwater seepage from uranium-tailings impoundments, 3. Simulation of the conceptual model with application to Seepage Area A. Journal of Contaminant Hydrology, 2, p.323-342.

Morin, K.A., and N.M. Hutt. 1995. MINEWALL 2.0: A technique for predicting water chemistry in open-pit and underground mines. IN: Proceedings of the Conference on Mining and the Environment, Sudbury, Ontario, May 28 - June 1, Volume 3, p.1007-1016.

Morin, K.A., N.M. Hutt, and S. Hutt. 2001. A compilation of empirical drainage-chemistry models (EDCMs). IN: Proceedings of Securing the Future, International Conference on Mining and the Environment, Skellefteå, Sweden, June 25-July 1, Volume 2, p. 556-565. The Swedish Mining Association.

Norman, A.T. 2018. Machine Learning in Action - A Primer for the Layman. Independently published. ISBN 9781717819994.

Ripley, B.D. 2004. Data Mining: Large databases and methods. IN: useR! 2004 - The R User Conference. Austrian Association for Statistical Computing, Technische Universität Wien in Vienna, May 20-22. Accessed at: https://www.r-project.org/conferences/useR-2004/Keynotes/Ripley.pdf

Rooki, R., F. Doulati Ardejani, A. Aryafar, and A. Bani Asadi. 2011. Prediction of heavy metals in acid mine drainage using artificial neural network from the Shur River of the Sarcheshmeh porphyry copper mine, Southeast Iran. Environmental Earth Sciences, 64, p. 1303-1316.

Singh, R.N., and F. Doulati Ardejani. 2003. Evaluation of factors affecting pyrite oxidation and subsequent pollutant generation in backfilled open cut coal mines. IN: 8th International Congress on Mine Water and the Environment, Johannesburg, South Africa, October, p. 173- 186.

Wang, H., P.A. Dowd, and C. Xu. 2019. A reaction rate model for pyrite oxidation considering the influence of water content and temperature. Minerals Engineering, 134, p. 345-355.

Witten, I.H., E. Frank, and M.A. Hall. 2011. Data Mining - Practical Machine Learning Tools and Techniques, Third Edition. Morgan Kaufmann Publishers. ISBN ISBN 978-0-12-374856-0.

Wunderly, M.D., D.W. Blowes, E.O. Frind, and C.J. Ptacek. 1996. Sulfide mineral oxidation and subsequent reactive transport of oxidation products in mine tailings impoundments: A numerical model. Water Resources Research, 32, p. 3173-3187.