PRACTICAL APPLICATIONS OF MACHINE LEARNING TO UNDERGROUND ROCK ENGINEERING

Josephine Sarah Morgenroth

A dissertation submitted to the Faculty of Graduate Studies
in fulfillment of the requirements of the Degree of

Doctor of Philosophy

Graduate Program in Civil Engineering
York University
Toronto, Ontario, Canada

July 2022

To Mama and Papa,

who each shaped me into the person I am.

Thank you for encouraging me to go on every adventure,

including this one.

ABSTRACT

Rock mechanics engineers have increasing access to large quantities of data from underground excavations as sensor technologies are developed, data storage becomes cheaper, and computational speed and power improve. Machine learning has emerged as a viable approach to process data for engineering decision making. This research investigates practical applications of machine learning algorithms (MLAs) to underground rock engineering problems using real datasets from a variety of rock mass deformation contexts. It was found that preserving the format of the original input data as much as possible reduces the introduction of bias during digitalization and results in more interpretable MLAs.

A Convolutional Neural Network (CNN) is developed using a dataset from Cigar Lake Mine, Saskatchewan, Canada, to predict the tunnel liner yield class. Several hyperparameters are optimized: the amount of training data, the convolution filter size, and the error weighting scheme. Two CNN architectures are proposed to characterize the rock mass deformation: (i) a Global Balanced model that has a prediction accuracy >65% for all yield classes, and (ii) a Targeted Class 2/3 model that emphasizes the worst case yield and has a recall of >99% for Class 2. The interpretability of the CNN is investigated through three Input Variable Selection (IVS) methods. The three methods are Channel Activation Strength, Input Omission, and Partial Correlation. The latter two are novel methods proposed for CNNs using a spatial and temporal geomechanical dataset. Collectively, the IVS analyses indicate that all the available digitized inputs are needed to produce good CNN performances.

A Long-Short Term Memory (LSTM) network is developed using a dataset for Garson Mine, near Sudbury, Ontario, Canada, to predict the stress state in a FLAC3D model. This is a novel method proposed to semi-automate recalibration of finite-difference models of high-stress environments. A workflow for optimizing the hyperparameters of the LSTM network is proposed. The performance of the LSTM network predicting the three principal stresses is improved as compared to predicting the six-component stress tensor, with corrected Akaike Information Criterion (AIC_c) values of -59.62 and -45.50, respectively.

General recommendations are made with respect to machine learning algorithm development for practical rock engineering problems, in terms of how to format and pre-process inputs, select architectures, tune hyperparameters, and determine engineering verification metrics. Recommendations are made to demonstrate how algorithms can be rendered interpretable with the application of tools that already exist in the field of machine learning.

ACKNOWLEDGEMENTS

Thank you first and foremost to my supervisors, Drs. Matthew Perras and Usman Khan. I am extremely grateful to have had the experience of being co-supervised by two experts in different fields, who also have the willingness to collaborate without ego that every PhD student dreams of. Matt and Usman were extremely supportive through both academic and personal struggles that arose in the unusual circumstances of completing this dissertation during a global pandemic. Their patience and (usually!) non-contradictory guidance were instrumental in developing machine learning techniques for rock engineering problems that not only satisfied my academic curiosity but also solved real-world problems. Thanks particularly to Matt for always being willing to "nerd out" about tunnels and rocks during our meetings, and to Usman for reassuring me when I was discouraged by MATLAB and questioning my coding abilities. Thanks also to Drs. Jit Sharma, Rashid Bashir, and Liam Butler for their questions and insights in the comprehensive and proposal examination phases that helped shape the direction of this research.

This research would not have been possible without the support of several funding agencies and institutions. I would like to thank: the Natural Sciences and Engineering Research Council of Canada for a three year Alexander Graham Bell Canada Graduate Scholarship – Doctoral as well as the Michael Smith Foreign Study Supplement; the Government of Ontario for the two term Ontario Graduate Scholarship; York University for the Enbridge Graduate Student Award and Graduate Development Fund awards; the Prospector's and Developers Association for the Joan Bath Bursary for Advancement in the Mineral Industry; and the Dr. Allan Carswell and Carswell Family Foundation for the Carswell Scholarship. I am also grateful for the numerous Teaching and Research Assistantships I held in the Department of Civil Engineering.

Thank you to the industry partners that provided datasets and informative conversations, particularly: Chris Twigs, Imre Bartha, and Kirk Lamont (Cameco), Kathy Kalenchuk (RockEng), and Lindsay Moreau-Verlaan (Vale). This research would not have been possible without the real data and practical insights given by these individuals. "Dankeschön" to Dr. Thomas Marcher, Paul Unterlaß, Tom Geisler, and Manuel Winkler for hosting me for four months at the Graz University of Technology in Austria.

My gratitude to my colleagues and friends at YorkU – Rodrigo Alcaino-Olivares, Pushpendra Sharma, Everett Snieder, Apostolos Vasileiou, Michael De Santi, Sebnem Boduroglu, Parin and Parnian Izadi, Marina Maciel Soares, and many others with whom I exchanged ideas, commiseration, and encouragement. It helped to know we were all in this together.

Finally, to my personal cheerleading squad – Katrina and Sebastian Morgenroth, Claire D'Amico, Ylena Quan, Natalie Narbutt, and Melissa Weaver, among many others far and wide. Your support over the last four years has been immeasurably important. "Thanks" will never be enough.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF EQUATIONS

LIST OF ABBREVIATIONS

%C	Percent Capture
1SHW	#1 Shear West (Garson Mine)
3D	Three-dimensional
Adam	Adaptive momentum estimation
AIC	Akaike Information Criterion
AIC_c	Corrected Akaike Information Criterion
ANN	Artificial Neural Network
BBN	Bayesian Belief Network
BEM	Boundary Element Modelling
CAS	Channel Activation Strength
CCM	Convergence-Confinement Method
CNN	Convolutional Neural Network
DEM	Discrete Element Modelling
DFN	Discrete Fracture Network
DISL	Damage Initiation and Spall Limit
DISP	Radial tunnel displacement
E	Easting
El	Elevation
EWD	East Wollaston Domain (Cigar Lake Mine)
FDM	Finite Difference Modelling
FEM	Finite Element Modelling
FREEZE	Ground freezing pattern
FVM	Finite Volume Modelling
GEO	Geotechnical zone
GMP	Ground Management Plan
GSI	Geological Strength Index
IO	Input Omission
IVS	Input Variable Selection
JBS	Jet Boring System (Cigar Lake Mine)
k-NN	*k*-Nearest Neighbours
LSTM	Long-Short Term Memory network
MAE	Mean Absolute Error
ML	Machine Learning
MLA	Machine Learning Algorithm
MLP	Multi-Layer Perceptron
MSE	Mean Squared Error
MSW	Municipal Solid Waste
N	Northing
NaN	Not-a-number

NSERC	Natural Sciences and Engineering Research Council of Canada
PC	Partial Correlation
PCA	Principal Component Analysis
R^2	Coefficient of determination
ReLu	Rectified Linear Unit
RF	Random Forest
RMR	Rock Mass Rating system
RMSE	Root Mean Square Error
RMSProp	Root Mean Square Propagation
RNN	Recurrent Neural Network
ROC	Receiver Operating Characteristic curve
RQD	Rock Quality Designation
SGD	Stochastic Gradient Descent
SSE	Sum of Squared Errors
SUPCL	Primary installed support class
SVM	Support Vector Machine
TBM	Tunnel Boring Machine
UCS	Unconfined Compressive Strength
WMTZ	Wollaston-Mudjatik Transition Zone (Cigar Lake Mine)
WWD	West Wollaston Domain (Cigar Lake Mine)
YIELD	Tunnel liner yield

CHAPTER 1. INTRODUCTION

1.1 Research Motivation

A significant cost of developing underground excavations is maintaining their stability by maximizing the utilization of the natural stress redistributions around the excavation and minimizing the need for human-made ground support. Conventional rock engineering favours two categories of analysis: support recommendations based empirical design charts, and numerical modelling based on a constitutive framework. In the former, case studies have been assembled to represent common rock mass deformation mechanics and are used to recommend rock support schemes, for example the Q tunneling index (Barton et al., 1974) and the Rock Mass Rating system (Bieniawski, 1993). In the latter, a numerical model is developed using a constitutive behaviour that most closely represents the observed rock mass behaviour. For example, the Hoek-Brown failure criteria modified by the Geological Strength Index (GSI) (Hoek & Brown, 2019) or the Damage Initiation and Spall Limit (DISL) (Perras & Diederichs, 2016) approaches used for brittle failure, or the phenomenological approach for squeezing ground (Barla & Borgna, 1999).

In general, empirical approaches are used in preliminary stages of design and then refined using numerical approaches in later stages of detailed design. In both cases, site specific data is used to validate the outputs of the approach being used. In the case of an observational design approach, as the project progresses more data becomes available and is used to calibrate the models and update design criteria (Peck, 1969). However, the calibration of physically based models is generally time-consuming, tedious, and in practice is often done using deterministic values or brute-force statistical methods. Uncertainty is introduced because it is uncommon that a complete dataset is available to evaluate rock mass behaviour, due to the time constraints on doing so (i.e., developing multiple models with all permutations of input data). Numerical model performance is evaluated by comparing observed ground conditions to modelled conditions, where inputs are manipulated until observed conditions are reproduced as closely as possible. Issues arise when the initially chosen framework does not capture the full range of the observed rock mass behaviour. In this case, the modelling technique chosen cannot account for all observed phenomena, and at times model refinement or redevelopment is time prohibitive.

There is often insufficient time, project schedule, and computational/expert resources to investigate complex rock mass phenomena in detail, particularly once a project is in progress. The large amount of multivariate data that characterizes complex and nuanced patterns of rock mass behaviour creates the ideal conditions to use machine learning algorithms. Machine learning allows for quick and efficient computation of large, multi-variate datasets to extract relationships between input variables and the chosen target. In the context of rock engineering, this allows practitioners to make use of all available data relevant to the predictand, and then analyze the trained algorithm to determine the most influential parameters to the present rock mass deformation mechanism. This understanding of the physical system can then be

used to inform design decisions, determine which data should be collected and in how much detail, and predict rock mass behaviour as additional data is collected to forecast excavation stability.

The purpose of this research is to advance the applications of machine learning algorithms, and Artificial Neural Networks (ANNs) in particular, to practical underground rock engineering design in a way that is useful and accessible for practicing rock engineers. To accomplish this, two datasets from mines representing endmember rock mass behaviour are used to develop algorithms to forecast targets relative to excavation behaviour.

The first dataset is from Cigar Lake Mine, located in northern Saskatchewan, Canada and owned by Cameco Inc. Cigar Lake Mine is situated in a challenging squeezing ground environment, and the ore extraction tunnels experience high tunnel convergences. For this dataset, conventional geotechnical mapping as well as radial tunnel survey measurements are used to forecast yield to tunnel liner elements such as shotcrete and rock bolts. The Cigar Lake Mine ground control engineers typically allow the tunnel convergence to occur up to a given threshold, and then budget to complete rehabilitation at necessary intervals. A Convolutional Neural Network (CNN) is developed for Cigar Lake Mine to forecast the tunnel liner yield classes to improve the mine's ability to schedule and budget for these rehabilitation activities.

The second dataset is from Garson Mine, located near Sudbury, Ontario, Canada and owned by Vale. Garson Mine experiences seismic events related to the high-stress environment and stope mining. This dataset consists of a database of microseismic events and a previously calibrated FLAC3D model of the mine. A consultant to the owner of the mine has developed the FLAC3D model to back-calculate the stresses in the mine using the major seismic events. This stress model is then used to determine stope removal sequencing to keep mine personnel and assets safe as mining progresses. Model recalibration is tedious, time-consuming, and computationally expensive, and so a Long-Short Term Memory (LSTM) network is developed to use the time-series microseismic data to forecast the stresses in the FLAC3D model, thereby allowing for real-time model recalibration as new microseismic events occur.

These case studies are used to demonstrate practical approaches to developing machine learning algorithms for application to real-world geomechanical data and rock mass phenomena. Based on these and a literature review, tools are recommended for the development of machine learning algorithms for practical rock engineering problems.

1.2 Background

This dissertation contains research on applying machine learning to time-dependent squeezing and high-stress seismogenic rock masses. Both rock mass deformation types are introduced briefly below, including the current state of the practice in terms of failure mechanism identification, characterization, and numerical modelling. These topics are provided as context for this dissertation and are not investigated in detail as part of the research scope. A brief introduction to machine learning in the context of rock engineering is also provided.

1.2.1 Time-Dependent Squeezing

Squeezing ground can be identified as large time-dependent convergence during tunnel excavation due to induced stresses and the material properties and takes place when the rock mass around the tunnel is pushed beyond the limiting shear stress and creep is initiated (Barla, 2002). The magnitude of tunnel convergence is related to geological conditions, in-situ stress relative to rock mass strength, and ground water and pore pressure. Excavation and ground support also play a crucial role in this time-dependent yielding process. For example, if support is installed immediately, it will build-up load, whereas if its installation is delayed after excavation, then stress redistribution occurs and the rock mass moves into the tunnel.

Empirical approaches exist to classify squeezing, for example Singh et al. (1992) based on Barton's *Q* (Barton et al., 1974), or Goel et al. (1995) based on the rock mass number Terzaghi's *N*. In general, squeezing is categorized by the tunnel convergence as a percentage of the tunnel diameter (Singh & Goel, 1999): mild squeezing is 1-3%, moderate squeezing is 3-5%, and high squeezing is >5%. Semi-empirical approaches have also been developed to further constrain these classifications (Aydan et al., 1993; Hoek & Marinos, 2000; Jethwa et al., 1984). Although these semi-empirical approaches are useful for estimating potential tunnelling problems due to squeezing conditions, they are not a substitute for more sophisticated methods of analysis (Barla, 2002).

Several computational models, such as the Convergence–Confinement Method (CCM) (Panet et al., 2001), axisymmetric models, and three-dimensional (3D) Finite Element and Finite Difference simulations, have been developed in the literature to represent tunnelling in rock with high squeezing potential (K. Zhao et al., 2015). Research has suggested that when the underlying assumptions of CCM or axisymmetric models are violated, true 3D modelling is required (Schürch & Agnagnostou, 2012). Ongoing research attempts to further improve 3D numerical modelling approaches to squeezing ground (e.g., Hasanpour et al., 2018; Kabwe et al., 2020), however these can be computationally expensive and difficult to parameterize correctly.

1.2.2 High-Stress Seismicity

Large seismic events are one of the greatest risks in deep hard rock excavations, particularly mines. These events have the potential to generate violent rocks bursts, which pose a risk to excavation stability, equipment, and personnel safety. Seismic hazard is defined as the likelihood of occurrence of a seismic event of a particular size (L. G. Brown et al., 2020), and each event gives information about the rock mass conditions at time of failure, such as increasing stress or yielding geological features (Hudyma et al., 2008).

Microseismic monitoring systems have gained popularity as a means of strategic risk mitigation (Yao & Moreau-Verlaan, 2010), with more than 70% of underground hard rock mines in Ontario, Canada using them (Hudyma & Brown, 2020). When the frequency of events is plotted versus their magnitude, the slope (or *b-value*) can be used to determine the mechanism of the seismicity occurring (e.g., fault slip, direct stress change due to mine blasting) (Hudyma et al., 2008). Frequency-magnitude analysis of seismic

events is one of the most widely used techniques for seismic hazard analysis. More sophisticated techniques, such as sequential spatial clustering and fractal dimension (Cortolezzis & Hudyma, 2018), have also emerged in recently to analyze microseismic data and extract information about the rock mass behaviour.

The data collected from microseismic monitoring systems can be combined with numerical modelling methods to improve the overall understanding of the rock mass behaviour. Numerical back analysis is a common approach to developing a model that replicates the observed ground behaviour, thus providing an estimation of the stress regime changes as mining progresses (Ju et al., 2019; Kalenchuk, 2017; Kalenchuk et al., 2014; Ma et al., 2018; Xue et al., 2021). Reducing numerical uncertainty of these complex numerical models is achieved via calibration, which must be undertaken with care and keeping in mind the practical limitations of the numerical models (Kalenchuk, 2019).

1.2.3 Machine Learning in Rock Engineering

Machine learning algorithms are tools that have gained popularity over the last several decades across medicine, technology, applied science, and engineering fields due to their flexibility in predicting a given output in the presence of appropriate input data. These data driven approaches are capable of processing large volumes of data with greater precision and accuracy than manual data analysis techniques (Liu & Yang, 2005; Marsland, 2014; Papadopoulos et al., 2000). Artificial Neural Networks (ANNs) are an algorithm type comprised of a series of highly interconnected nodes and a series of parallel nonlinear equations, and are particularly powerful for pattern recognition (M. Khan et al., 2017c; Lecun et al., 2015; Marsland, 2014). The two ANNs explored in detail in this dissertation are the CNN and the LSTM network.

CNNs are efficient at processing spatial and temporal dependencies in image or raster datasets. They were originally developed for handwritten digit classification (LeCun et al., 1989), and have evolved into a common tool in the field of computer vision (Zeiler & Fergus, 2014). CNNs incorporate spatial dependencies by extracting features from adjacent pixels, as well as their change over time, during the algorithm training process. LSTM networks are designed to process entire sequences of data using recurrent nodes, which allows them to “remember” and “forget” information during algorithm training (Mandic & Chambers, 2001). This allows the algorithm to retain useful parts of a time series input history, and discard information that is not useful for good prediction of the output.

To date, research at the intersection of machine learning and rock engineering has included predicting: rock mass properties (Sklavounos & Sakellariou, 1995; Song et al., 2015), constitutive behaviour (Kumar et al., 2013; Millar & Clarici, 2002), slope stability (Ferentinou & Fakir, 2018; Hibert et al., 2017; Mayr et al., 2018), tunnel performance (Delisio et al., 2013; Koopialipoor et al., 2019; Mahdevari & Torabi, 2012), and rock bursts (Pu et al., 2018; Ribeiro e Sousa et al., 2017). However, machine learning applications in rock engineering practice are in their infancy despite being widely used and generally accepted for decades in other risk assessment fields, such as flood forecasting (Abrahart et al., 2012).

A comprehensive review on machine learning in underground rock engineering was published (Morgenroth et al., 2019) and is presented in *CHAPTER 2. An Overview of Opportunities for Machine Learning Methods in Underground Rock Engineering Design*.

1.3 Thesis Objectives

Specific objectives for this research have been identified based on gaps in the current literature and in standard engineering practice. While there are increasing publications in the academic literature on the development of machine learning solutions to rock engineering problems, these have not yet become standard practice due to the perceived opacity of machine learning algorithms. This research aims to demonstrate the development of interpretable machine learning for both classification and regression problems. Objectives 1 and 2 are concerned with a classification algorithm developed for Cigar Lake Mine, where data sparsity is an obstacle. The input data is formatted similarly to familiar tunnel mapping, and an investigation and ranking of input variables is completed to uncover how the data is used in the algorithm. Objective 3 addresses a regression problem with an application to numerical model calibration for Garson Mine. Input data is formatted to preserve the spatial and temporal dependencies, and the algorithm outputs are imported into an industry-standard numerical modelling code. Objective 4 makes use of the lessons learned in fulfilling Objectives 1-3 to recommend machine learning tools and processes that can be applied to solving practical rock engineering problems.

The dissertation objectives are as follows:

1. *Develop a classification machine learning algorithm using standard geotechnical mapping data from a real project*

The first objective of this research is to develop a classification machine learning algorithm that takes standard geotechnical mapping data as inputs to forecast and classify yield to tunnel liner support elements. The aim of the algorithm is to forecast the worst liner yield, so that the ground control engineers can intervene in advance to prevent the need for complete tunnel reprofiling. This is currently a reactive and manual intervention process, whereby the intervention occurs after the yield is critical and the mine drifts cannot be accessed for ore extraction without rehabilitation. The motivation for this objective is to demonstrate that well-performing machine learning algorithm can be developed while preserving the original format of the tunnel mapping data familiar to practicing geotechnical engineers. A Convolutional Neural Network (CNN) approach is developed to process tunnel mapping images for tunnel liner yield forecasting, where the priority is to predict the worst-case tunnel liner yield. The approach must allow for the combination of numerical and categorical input variables.

The aim of this objective is to determine how much data should be digitized to obtain good predictive performance, to develop a hyperparameter optimization approach for the algorithm, and to ensure the engineering verification metrics for the algorithm align with the desired outcome for the ground control engineers.

2. *Develop and implement a methodology for assessing interpretability of machine learning algorithms applied to rock engineering*

The second objective is to investigate algorithm interpretability using Input Variable Selection (IVS) approaches. The developed IVS approaches must contribute to the two broad categories of algorithm interpretability: (1) algorithm transparency, and (2) post-hoc interpretations. The former category is concerned with understand how the model works, by considering the understanding of the entire model, the hyperparameters, and the training algorithm. The latter category reveals useful information about the model after it has been developed through visualization of learned features, as well as explanations by example.

The motivation for this objective is to debunk the perception that machine learning algorithms are “opaque” or “black boxes”, and that in fact much of how the algorithms use the provided input data parallels widely understood rock engineering principles. Three IVS methods will be developed and applied to demonstrate algorithm interpretability, including both model-based and model-free methods. The IVS methods must rank the available inputs in terms of importance for model performance and must inform the user about the sensitivity of the trained algorithm to the selected input variables. The IVS must be customized to account for the format of the tunnel images used to train the CNN algorithm, where each image channel is a discrete input.

3. *Develop a regression machine learning algorithm using standard geotechnical sensor data from a real project*

This research objective is to develop a regression machine learning algorithm that takes a large volume of sensor data as inputs to forecast the stress state in a seismogenic mining environment. The developed machine learning approach must highlight the formatting of time series data into sequences that can be combined with static categorical inputs, such as geology, for predictive modelling. The aim of this approach is to assist in the calibration of a finite difference model, which is typically recalibrated manually and only with the largest microseismic events, when new microseismic sensor data becomes available. A Long-Short Term Memory (LSTM) network is developed to process the microseismic event time-series and combine them with the static geomechanical and geological parameters to predict the stress state in a finite difference model. The approach must facilitate the calibration of the finite difference model while making use of as much useful microseismic data as possible.

The aim of this objective is to demonstrate the practical development of a regression algorithm, highlighting the pre-processing and formatting of input variables, the hyperparameters that should be investigated during algorithm development, and the integration of machine learning algorithms into standard rock engineering design processes.

4. *Formulate a guide for rock engineering practitioners in the selection, development, and engineering verification of machine learning algorithms for underground rock engineering problems*

The final objective is to combine all the lessons learned from the previous objectives to formulate a guide of tools and best practices for machine learning algorithms applied to underground rock engineering. The guide will also include recommendations from other authors publishing at the intersection of these fields. The framework will look at holistic machine learning algorithm development, from data pre-processing and balancing; to algorithm architecture selection; to engineering verification metric selection. Machine learning has not yet been widely adopted by the practicing rock engineering community, which would make this guide one of the first of its kind.

1.4 Research Significance

This research attempts to bridge the gap between two fields that are not yet intertwined, despite the opportunities to improve excavation safety and stability using techniques from machine learning. Perceived lack of algorithm interpretability is a major source of skepticism in the rock engineering community, and this dissertation aims to provide insight into how algorithms can be developed prudently and with appropriate engineering verification metrics.

Machine learning algorithms have the potential to ease the burden of data analysis, and also provide the advantage of reducing error and/or bias that can be introduced when data is manipulated manually. If the algorithm is developed with consideration, complex and nuanced relationships may be extracted that describe how a rock mass behaves and interacts with surrounding structures. Data driven methods do not require the dataset to conform to pre-determined empirical or constitutive frameworks, increasing the likelihood that the site-specific rock mass behaviour is captured accurately as compared to some conventional numerical modelling methods.

The contributions of this research are particularly relevant to the industries that rely on underground infrastructure, both in Canada and abroad, such as utility owners, mining companies, and transit projects. If rock mechanics practitioners can use machine learning to predict catastrophic rock mass behaviours earlier and more accurately, this could translate to cost savings in terms of reducing equipment loss, increasing efficiency of resource extraction or excavation rates, minimizing delays and shutdowns due to rock failure, and most importantly, increased safety for underground personnel and the public. All the data used in this research was supplied by partners in the mining industry In Ontario alone, mining is a $10.7 billion industry (2020), with 38 mining companies and 900 affiliated companies (Ontario Mining Association, 2022). Pushing the frontier of underground design using these emerging machine learning methods positions these mining companies to be more competitive in the international market, and presents the opportunity for safer, more economic underground excavations across the mining, infrastructure, and energy sectors.

1.5 Dissertation Outline

Several chapters of this dissertation are modified versions of journal or conference articles. Those chapters have their own introduction, methods, results, and discussions. Chapters that are not published elsewhere are included to enrich the context of the research presented in this dissertation.

Chapter 2 presents an in-depth literature review of machine learning applied in underground rock engineering, for audiences coming from both rock engineering and machine learning backgrounds. Thus, introductions of relevant rock engineering principles, as well as an introduction to the mechanics of common machine learning algorithms are provided. The opportunities for applying machine learning in rock engineering are highlighted, and relevant research published to date is cited to illustrate these opportunities.

Chapter 3 introduces the datasets and case studies used in the remainder of the dissertation, including sources, digitization, formatting, and project context. Chapter 4 describes the preliminary case studies conducted with the datasets presented in Chapter 3, and summarized the lessons learned for the development of subsequent, more sophisticated algorithms.

Chapter 5 contains the input data analysis and hyperparameter optimization performed to develop the Cigar Lake Mine CNN (related to Objective 1). The digitization of the input variables (geotechnical zones, primary installed support class, locations of ground freezing, and measured radial tunnel displacement) and their subsequent formatting into channels of a tunnel image are discussed. Convolution filter sizes, amount of training data, and error weighting schemes are evaluated based on a sensitivity analysis to determine the optimal hyperparameters for predicting the tunnel liner yield. Chapter 6 presents the IVS methods developed to enhance the interpretability of the Cigar Lake Mine CNN (related to Objective 2). Three IVS methods are developed: Channel Activation Strength (CAS), Input Omission (IO), and Partial Correlation (PC), where the latter two methods are novel approaches applied to CNNs. The input rankings resulting from the IVS methods are compared and contrasted in the context of the CNN algorithm's transparency and the post-hoc interpretations offered.

In Chapter 7, the development of the Garson Mine LSTM network is presented (related to Objective 3). The vast amount of available microseismic data is formatted to be compatible with the geomechanical properties obtained from a previously developed FLAC3D model. The iterative approach to the LSTM network development is thoroughly documented, including the choice of data pre-processing, hyperparameter optimization, and engineering verification metric selection.

Chapter 8 focuses on presenting the practical guide to best practices when developing a machine learning algorithm for a rock engineering problem (related to Objective 4). Recommendations are made in terms of holistic machine learning algorithm development, from data digitization, pre-processing, and formatting; to algorithm architecture selection and hyperparameter optimization; to engineering verification metric selection and output integration into standard design practice.

Chapter 9 summarizes major conclusions from each chapter and novel contributions of this dissertation. Suggestions for future work are also presented.

References for the entirety of this research, including chapters that were previously published, are included at the end of the dissertation. Appendix A includes the proofs of concept discussed in Chapter 4, Appendices B and C include the MATLAB code for the Cigar Lake Mine and Garson Mine datasets, respectively, and Appendix D includes a list of all the publications resulting from this dissertation.

CHAPTER 2. AN OVERVIEW OF OPPORTUNITIES FOR MACHINE LEARNING METHODS IN UNDERGROUND ROCK ENGINEERING DESIGN

2.1 Preface

This chapter focuses on a literature review on the state of the art of machine learning in underground rock engineering and related publications at the time of writing. The paper was published to a mixed audience of machine learning experts and rock mechanics engineers, and therefore explanations of standard practice for modelling in both machine learning and rock engineering are presented. Section 2.4 Current Practices in Rock Engineering Design focuses on standard numerical modelling methods and their applications to rock engineering design, while section 2.5 Review of Machine Learning Algorithms presents an overview of common machine learning algorithms and how the are parameterized. Section 2.6 Discussion of Machine Learning for Rock Engineering Design presents a discussion of where the fields of machine learning and rock engineering may intersect, as well as highlighting opportunities for the use of machine learning in rock engineering design. The content of this chapter was published in *Geosciences* in 2019 as follows:

Morgenroth, J., Khan, U. T., & Perras, M. A. (2019). An overview of opportunities for machine learning methods in underground rock engineering design. *Geosci* J, 9(12), 504–524. https://doi.org/10.3390/geosciences9120504

The contributions of the authors in the current chapter are as follows:

Josephine Morgenroth has conducted the literature review, and prepared and wrote the original manuscript of this publication. **Usman T. Khan** has supervised the research, provided the funding, contributed to the writing and editing the manuscript. **Matthew A. Perras** has supervised the research, provided the funding, contributed to writing and editing the manuscript.

This work would not have been possible without funding in part from the Natural Sciences and Engineering Research Council of Canada (NSERC) through the Discovery Grant program (funding reference numbers RGPIN-2018-05918 and RGPIN-2017-05661) and the National Research Council Canada's Industry Research Assistance Program – Artificial Intelligence Industry Partnership Fund. We would like to thank our industry partners, Yield Point Inc., RockEng, and Cameco for insightful conversations on the topic.

2.2 Abstract

Machine learning methods for data processing are gaining momentum in many geoscience industries. This includes the mining industry, where machine learning is primarily being applied to autonomously driven vehicles such as haul trucks, and ore body and resource delineation. However, the development of machine learning applications in rock engineering literature is relatively recent, despite being widely used and generally accepted for decades in other risk assessment-type design areas, such as flood forecasting. Operating mines and underground infrastructure projects collect more instrumentation data than ever before, however only a small fraction of the useful information is typically extracted for rock engineering design and there is often insufficient time to investigate complex rock mass phenomena in detail. This chapter presents a summary of current practice in rock engineering design, as well as a review of literature and methods at the intersection of machine learning and rock engineering. It identifies gaps, such as standards for architecture, input selection and performance metrics, and areas for future work. These gaps present an opportunity to define a framework for integrating machine learning into conventional rock engineering design methodologies to make them more rigorous and reliable in predicting probable underlying physical mechanics and phenomenon.

2.3 Introduction

The study of rock mechanics encompasses the theoretical and applied science of the mechanical behaviour of rock in response to its physical environment, and was formalized as a field of study in the 1960s (Hoek, 1966). Due to the high degree of variability in natural materials, a precise rock engineering design within a narrow tolerance is difficult to produce. Experience and expert knowledge are heavily relied upon in rock engineering practice and empirical design charts have become prolific for preliminary stage design. Numerical modelling methods, such as continuum and discrete methods, are also conventional tools in rock engineering design. These methods represent important tools for understanding rock mass behaviour and predicting its response to its environment and changes in in-situ stress conditions. In practice it is often difficult to integrate all the data collected into empirical and numerical models effectively due to time and budget constraints, as well as limitations in how constitutive behaviour is defined. Furthermore, it is sometimes not feasible to collect the quality and quantity of data needed, so extrapolation and interpolation techniques are often used.

Prior to the advent of “big data” and advances in machine learning, it was appropriate to base rock engineering design primarily on empirical data and expert knowledge because the geology, discontinuity, and in-situ stress data available for projects were generally sparse. Now, increasing amounts of some forms of data, such as displacements and pore water pressures, are being collected relatively inexpensively, while others are still infrequently collected because of the complexity and cost of the measurement methods, such as those used for stress measurement. Herein lies an opportunity to integrate machine learning into the existing rock engineering best practices to evaluate datasets more efficiently and to maximize the value (or information) extracted from the data. Machine learning algorithms are especially powerful because

simultaneous hypotheses can be tested much quicker than by a human being (Mitchell, 2015). This frees up the expert user or engineer to devote their judgment to selecting inputs and validating outputs in a more rigorous way, rather than manipulating data, which can be more time consuming and poses the risk of introducing bias.

Although machine learning has been used in other fields that rely heavily on numerical modelling, such as flood forecasting, for over 20 years (Abrahart et al., 2012), comparatively it is in its infancy in rock engineering applications. Some professionals in the mining industry believe the community is reluctant to accept data driven methods because they are not as interpretable as applying conventional design criteria determined by previous experience and empirical evidence (PDAC, 2019a). Work is being done to produce scientifically interpretable data driven models (Karpatne et al., 2017), however these frameworks have not migrated into engineering practice yet. Expert judgment is given much credence in the rock engineering industry, and empirical relationships form the foundation of fundamental rock engineering methods (W. J. Mcgaughey, 2019).

The progression of the design process is constrained by the analysis of the data that is collected. Machine learning algorithms have the potential to ease the burden of data analysis. If geotechnical professionals are able to predict geomechanical events or behaviours earlier and more accurately, this could translate to increased safety for underground personnel as well as cost savings in terms of reducing equipment loss, increasing efficiency of resource extraction or excavation rates, and minimizing delays and shutdowns.

This chapter aims to summarize the most common practices in rock engineering design and presents the opportunities for integrating machine learning into existing geomechanical design frameworks. It also presents a review of common machine learning algorithms and how they are currently being applied to rock engineering problems in the literature, as well as future opportunities for data driven approaches.

2.4 Current Practices in Rock Engineering Design

The inherent variability associated with natural materials makes standardization across geomechanical design processes difficult. Recent work to develop a geomechanical design framework has resulted in the publication of Eurocode 7 (LimiteStateInc., 2013). This code provides guidelines for tunnelling and underground excavations using a limit state approach, however in general customized designs are developed for most projects. Two of the most common categories of design approaches used in rock engineering design are: (i) empirical methodologies, and (ii) numerical modelling. These usually fall into a workflow together depending on the stage of the design, from prefeasibility to detailed design. Empirical design recommendations are often taken as a first estimate and are then verified and updated using the appropriate numerical tools, which are based on the geometry and complexity of the problem. For both these approaches, rock engineers use their judgement to combine numerical data (e.g. rock mass properties, stress conditions, pore water pressure) and categorical data (e.g. geological mapping,

discontinuity conditions) into an empirical or numerical framework. As context for the remainder of this chapter, both categories of conventional methodologies are briefly described below.

2.4.1 Empirical Design

In the early stages of conventional rock engineering design, it is common to turn to an established empirical rock support recommendation method to inform the design. The designed rock support may consist of a combination of rock or cable bolts, shotcrete, steel mesh, steel straps, concrete liner segments, etc. These methods are generally dependent on the size of the excavation and the quality of the rock mass and are used only as a first estimate followed by analytical or numerical methods to verify the design. It is typical to use multiple empirical methods and a range of input parameters representing the best and worst case to obtain a range of rock support recommendations. In this section, two common empirical rock support methodologies are described: the Rock Mass Rating (RMR) (Bieniawski, 1993) and the Q Tunnelling Index (Barton et al., 1974). These systems have been in wide use in the tunnelling and mining industries for 45 years, and have become engrained in standard rock engineering practice (Barton & Bieniawski, 2008). However, previous researchers (Barton & Bieniawski, 2008) advocate for exercising caution when using these empirical systems, and for practicing engineers to use them in the spirit for which they were developed – to assist with the design of excavations, but not as the sole tool for designing underground support. Both these classification systems result in a recommended category for rock support, which is meant to encompass the range of rock mass conditions found within that category.

2.4.1.1 Empirical Support Recommendation – Rock Mass Rating

Bieniawski was among the first to assert that no single index, such as Deere's Rock Quality Designation (Deere, 1963), was enough to capture that aggregate behaviour of the rock, and so developed the Rock Mass Rating (RMR) system (Bieniawski, 1993) to combine several measurable parameters. This system is the sum of five basic parameters to arrive at a score ranging from 0 to 100:

- Strength of intact rock
- Rock Quality Designation (RQD) (Deere, 1963)
- Spacing of joints
- Condition of joints
- Groundwater conditions

This information is obtained through field mapping and site investigations. The site engineer or geologist collects a score for each of the parameters using the scoring scheme set out by Bieniawski (Bieniawski, 1993). A statistical analysis is conducted to determine the appropriate value, or more likely values, for each parameter and then sums them to obtain an RMR. The RMR corresponds to a rock quality categorization that ranges from very good to very poor rock with five classes overall based on a linear relationship with the RMR value. Bieniawski created support guidelines based on the quality of the rock mass, i.e. the RMR value, and the stress condition the excavation is exposed to (Bieniawski, 1993). For example, low stress environments only require heavy support if the rock is of poor quality.

2.4.1.2 Empirical Support Recommendation – Q Tunnelling Index

The Q Tunnelling Index was developed primarily to predict the appropriate support for tunnels (Barton et al., 1974), and is defined by multiplication of three quotients: the rock block size, the roughness and frictional resistance, and a stress quotient. As with RMR, the parameters are collected during the site investigation stage of a project, and subsequently a statistical analysis is performed to determine the most representative value or values. The Q Tunnelling Index is a value that ranges from 0.001 to 1000, and the formula is given in Equation 2-1.

$$Q = \left(\frac{RQD}{J_n}\right) * \left(\frac{J_r}{Ja}\right) * \left(\frac{J_w}{SRF}\right) \quad \text{Equation 2-1}$$

Where:

RQD = Rock Quality Designation (Deere, 1963)

J_n = Number of joint/fracture sets

J_r = Roughness of most unfavourable joints

J_a = Alteration or infilling of joints

J_w = Water inflow

SRF = Stress reduction factor, quantifies stress conditions

The major difference between the Q Tunnelling Index support recommendation and others (such as RMR) is the inclusion of a parameter called the *equivalent dimension*, which uses an *excavation support ratio* to modify the span, or width, of the underground opening to capture a factor of safety correlating to the end use of the tunnel.

2.4.2 Numerical Methods

Numerical modelling is crucial for understanding fundamental intact rock and rock mass behaviour, assessing rock-structure interactions, and completing rock engineering designs (Jing & Hudson, 2002). These methods offer the tools to capture the mechanisms that are causing observed phenomena in a rock mass or intact rock sample, and subsequently incorporate them into design. Choosing inputs, geometry, boundary and stress conditions, among other defining aspects of a numerical model requires significant engineering judgement and the combination of qualitative and quantitative site observations.

A typical approach is to use back analysis of numerical models to calibrate their behaviour against observed site conditions. These calibrated models are then used to forward predict the excavation behaviour as the excavation is advanced, or as a new excavation is being designed in the same rock mass but at a different location (M. A. Perras et al., 2015a). Figure 2-1 illustrates this process, summarized as follows:

1. Site observations
2. Measurements and model calibration

3. Conceptual translation of the calibrated behaviour to the new site location
4. Compare with empirical approaches (Diederichs, 2007) to validate the new design.

For example, the observed overbreak, or damage beyond the designed diameter, of a tunnel can be back analyzed using a numerical modelling approach. Individual parameters, such as tunnel depth, are manipulated independently of the other input variables in order to calibrate them, by taking measurements along the tunnel axis at multiple locations, for example. This calibration process ensures that a robust set of input parameters are determined based on their individual impacts on the measured site conditions at a variety of locations along the tunnel alignment. The empirical knowledge from the tunnel site is translated conceptually to the new site in order to understand what similarities and difference may exist. With this conceptual understanding, a variety of models of the new site can be developed and the results compared to other empirical approaches, such as the damage depth prediction of Diederichs (Diederichs, 2007).

The most common numerical methods for rock engineering problems are (Jing & Hudson, 2002):

- Continuum methods – finite element modelling (FEM), finite difference modelling (FDM), boundary element modelling (BEM)
- Discrete methods – discrete element modelling (DEM), discrete fracture network (DFN) modelling
- Hybrid continuum/discrete methods

The choice between methods is made based primarily on the scale of the problem and the geometry of the discontinuity, or fracture, system in the rock mass (Jing & Stephansson, 2007). Continuum methods are most appropriate when detachment of discrete blocks is not a significant factor. Discrete methods are generally chosen when the number of fractures is too large to treat the rock mass as a continuum with fracture elements, or if discrete block detachment is anticipated. Hybrid models are selected to avoid the pitfalls of each of the former two approaches. Each of these methods, their applications, and limitations are briefly introduced.

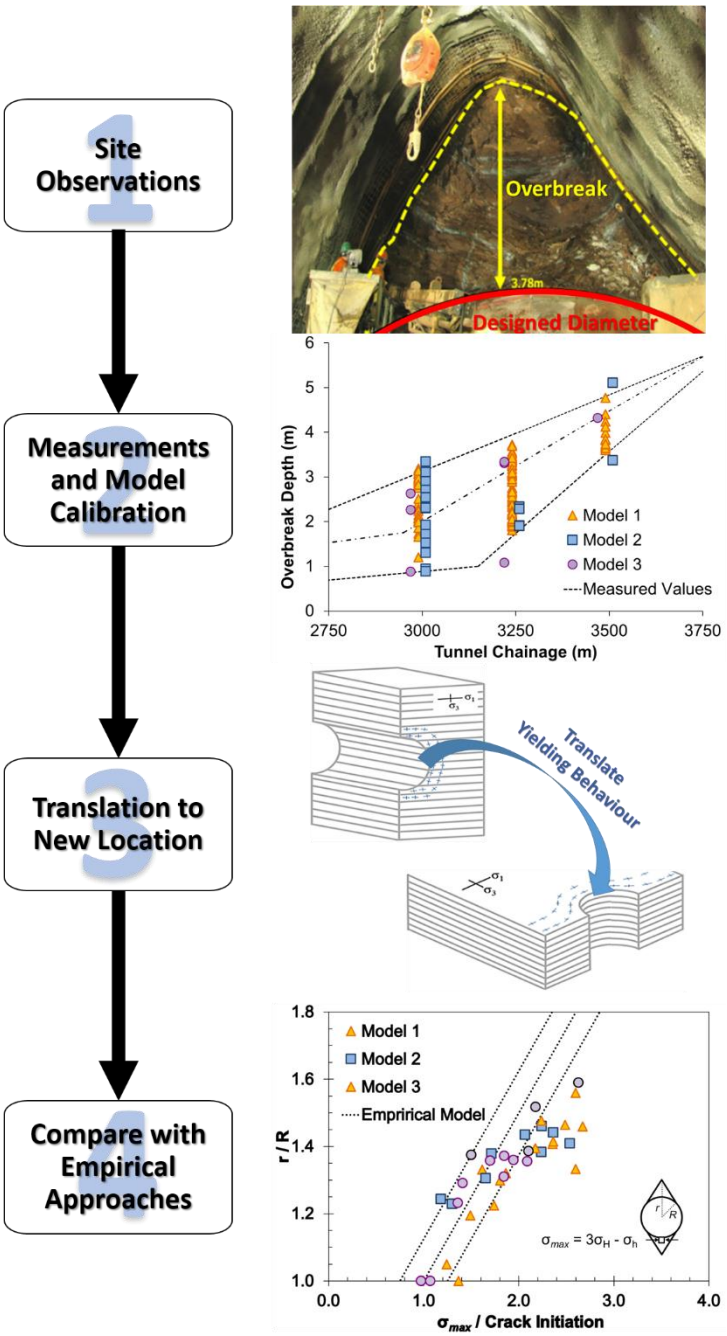

Figure 2-1. An illustration of a numerical modelling approach used to predict the overbreak depth around a vertical shaft (modified from (M. A. Perras et al., 2015a)). Step 1. Observed overbreak beyond the designed perimeter of the tunnel (M. A. Perras et al., 2015b). Step 2. Measured overbreak depths versus various numerical model results (Diederichs, 2007; ITASCA Consulting Group Inc., 2015; Perras, 2009; Perras & Diederichs, 2016). Step 3. Translation of the yielding behavior from the tunnel in the horizontal beds to a vertical shaft in the same rock mass. Step 4. Comparison of calibrated models and conceptual behaviour to an empirical approach for predicting normalized overbreak depth (r/R) using the maximum tangential stress (σ_{max}) calculated from the maximum and minimum horizontal stresses, σ_H and σ_h, respectively (NWMO, 2011; M. A. Perras et al., 2015a).

2.4.2.1 Continuum Methods

The basic concept for continuum methods is to discretize the material being modelled into a grid governed by partial differential equations. The partial differential equations at the grid points are in close enough spatial proximity that the errors introduced between them are insignificant and thus, acceptable.

The FDM is the most direct way to discretize a continuum, where points in space are replaced with discrete equations called finite difference equations. These equations are used to calculate displacement, strain, and stress in the material in response to conditional changes in the rock mass (Garza-Cruz et al., 2014). Solutions are formulated at grid points at the local scale, so no global matrix inversion is required, thus saving computational time and intensity. Complex constitutive behaviour can be captured without iterative solutions. However, FDM is inflexible with respect to fractures, complex boundary conditions, and material heterogeneity. FDM was historically unsuitable for rock mechanics problems due to these limitations, however advancements in irregular grid shapes gave rise to related Finite Volume Modelling (FVM) techniques (Jing & Hudson, 2002). FVM is more flexible in handling heterogeneity and boundary conditions and has been regarded as the bridge between FDM and FEM (ITASCA Consulting Group Inc., 2015). Commonly used commercial FDM software includes FLAC in two-dimensions and FLAC3D in three-dimensions (ITASCA Consulting Group Inc., 2015).

FEM is the most widely applied numerical method in science and engineering (Jing & Hudson, 2002). FEM also involves discretizing a continuum into a grid, however here the material is subdivided into parts called finite elements. The partial differential equation at each element is informed by the elements adjacent to it, allowing FEM to handle heterogeneity, plasticity and deformation, complex boundary conditions, in situ stresses and gravity (Hoek et al., 1990). However, detachment of the elements is not permitted since these models are based on continuum assumptions. The treatment of fractures has been the largest limitation of FEM in the past, and modern software packages include special algorithms to overcome this. Commercial FEM software available are RS2 (Rocscience, 2019), SIGMA/W (GEO-SLOPE, 2016), Plaxis 2D (Bentley, 2019), and ABAQUS (Dassault Systems, 2019), while open source options include Adonis (Geraili Mikola, 2019), OpenSees (UCRegents, 2006), and Code-Aster (EDF, 2019).

BEM differs from FDM and FEM in that is first seeks an approximate global solution. Initially, BEM was developed for underground stress and deformation analysis, soil-structure interactions, groundwater flow, and fracturing processes (Jing & Hudson, 2002). BEM approximates the solution of a partial differential equation inside an element by looking to the solution on the boundary and using that to inform the solution inside the element (Laforce, 2006). The main advantage of BEM over FDM or FEM is the simpler mesh generation and decreased computational expense. However, BEM is less efficient than FEM in handling material heterogeneity and plasticity. BEM has been used for: stress analysis of underground excavations, dynamic problems, back analysis of in situ and elastic properties, and borehole permeability tests (Jing & Hudson, 2002). Commercially available BEM software include Examine2D (Roscience Inc., 2019) and Map3D (Map3D, 2019), while open source BEM libraries are also available (Wieleba & Sikora, 2009).

2.4.2.2 Discrete Methods

Rock mechanics is one of the fields that originated DEM modelling, because highly fractured rock masses are not easily described mechanistically by a continuum approach. In discrete element approaches, the material is treated like an assemblage of rigid or deformable blocks or particles, the contacts between which are updated during the modelling process (Jing & Hudson, 2002). DEM solutions combine implicit and explicit formulations, based on FEM and FDM discretization, respectively. The main difference between DEM and continuum approaches are that the contacts between the elements are continuously changing, while they remain static for the latter. DEM methods are computationally demanding, however they offer an advantage when the rock mass experiences loss of continuity (from progressive failure, for example), as continuum constitutive models are inappropriate in that case (Lisjak et al., 2014). DEM methods have been popular for modelling a variety of rock engineering problems, including: underground works, rock dynamics, rock slopes, laboratory tests, hard rock reinforcement, borehole stability, acoustic emissions in rock, among others. Commercially available DEM codes include UDEC (ITASCA Consulting Group Inc., 1992) and PFC (ITASCA Consulting Group Inc., 2019b) for two-dimensional problems, and 3DEC (ITASCA Consulting Group Inc., 1994) and PFC3D (ITASCA Consulting Group Inc., 2019c), three-dimensional problems, respectively. Open source alternatives include Yade (Šmilauer, 2009) and LAMMPS (Sandia National Labs and Temple University, 2019).

The DFN model is a discrete method focused on fracture pattern simulation. It takes in statistical information about the fracture sets and can be used to generate a network for input into DEM codes for use in its behaviour, such as considering fluid flow through a series of interconnected fractures. It is a powerful method for studying fractured materials where an equivalent continuum cannot be established. DFN codes have been applied to the following problems: developments for multiphase fluid flow, hot dry rock reservoir simulations, permeability of fractured rock, and water effects on underground excavations and rock slopes (Jing & Hudson, 2002). Commercial DFN softwares include FracMan (Golder Associates Inc., 2019) and MoFrac (Miraco Mining Innovation, 2019), while open source options include ADFNE (Alghalandis Computing, 2019).

Discrete approaches are limited by the modeller's knowledge of the geometry of the fracture network, which can only be estimated based on geological mapping and interpretation of in situ information.

2.4.2.3 Hybrid Continuum/Discrete Methods

Hybrid numerical models have gained popularity to overcome the limitations of each of the numerical methods previously described. Hybrid methods are commonly used to address limitations in how fracture growth is addressed by the other methods. Some hybrid methods add fractures discretely and the mesh is adjusted dynamically as the crack propagates, while others have a very fine mesh and the fractures propagate along the boundaries. Care must be taken where two methods interact to ensure compatibility of the underlying assumptions.

The most common hybrid models are BEM/FEM (Varadarajan et al., 1985), DEM/FEM (Lorig et al., 1986), and DEM/BEM (Pan & Reed, 1991). There are many advantages to hybrid numerical models when applied correctly, for example the coupled DEM/FEM method explicitly satisfies equilibrium conditions of displacement at the interface between two domains. The FEM/DEM approach has been shown to realistically model the dynamic response of rocks, and simulations show good agreement with laboratory observations (Mahabadi et al., 2010). New developments in DEM/FEM approaches allow the modelling of rock masses with anisotropic strength and deformation characteristics, as well as the explicit modelling of fracture growth (Li, Kim, et al., 2019; Lisjak et al., 2014). Irazu is a commercially available DEM/FEM hybrid software package that explicitly models fracture processes in brittle materials, capturing complex non-linear behaviour (Geomechanica, 2019). In addition to Irazu, other commercially available software include the Hybrid Optimization Software Suite (Knight et al., 2014) and Elfen (Rockfield, 2019), while open source alternatives include Y-GEO (Grasselli's Geomechanics Group, 2019).

The choice of numerical method depends on the data available and the complexity of the problem being solved, and sometimes multiple methods may be explored before one is selected to use in subsequent design activities. Similar to empirical design approaches, significant expert judgement is required to determine the model inputs and interpret the outputs.

2.4.3 Discussion of Current Practices

Both empirical and numerical methods are strongly rooted in conventional rock mechanics design and lend insight into anticipated rock mass response to the construction of an excavation. The issues often arise not from the methods themselves, but rather how the data collected is used to make support decisions or how it is input into numerical models. The inherent variability of geomechanical datasets, as well as the variety of types (numerical measurements, categorical descriptions, photos) introduce error to the design process.

When making use of empirical support recommendations, it is common practice to determine the rock mass classification value based on a statistical analysis or using the best- and worst-case values. While this is appropriate for a prefeasibility estimate, these values are sometimes carried forward into detailed design. A data driven method would allow the design engineer to make use of all the data to inform the most appropriate parameters for further use. Some research has been conducted using Artificial Neural Networks (ANNs) to classify rock masses (Sklavounos & Sakellariou, 1995), determine strength properties (V. K. Singh et al., 2001), and model stress-strain behaviour (Millar & Clarici, 2002).

Numerical modelling of rock mass behaviour is difficult because rock is a discontinuous natural material with inherent variability and inhomogeneous properties. These need to be captured by fundamental equations and constitutive models which have geometrical and physical constraints. This can be overcome using data driven methods such as ANNs because the data trends are not conformed to these constraints, which sometimes understate the complexity of the problem (Jing & Hudson, 2002). Research has been performed using ANNs for predicting tunnel convergence (Mahdevari & Torabi, 2012) (inward radial displacement of the tunnel), rock bursts (Afraei et al., 2019) (an accumulation and sudden release of strain

energy), open pit stability (Ferentinou & Fakir, 2018) (stability of the mine slope geometry), among other applications.

2.5 Review of Machine Learning Algorithms

Machine learning is a branch of artificial intelligence that aims to program machines to perform their jobs more skillfully (Mitchell, 2015). This is done by using intelligent software that takes inputs to train a model to produce the desired result, thus replicating learning. Machines are better at performing repetitive tasks than humans, so harnessing this potential has been at the forefront of almost every industry since the beginning of the technological age. Sometimes the model is intuitively understandable, and other times the workings of the model cannot be easily explained. The choice of machine learning techniques is informed by what the output is, and what data is available (Mohammed et al., 2016a):

- Supervised learning: data is labelled, i.e. the training samples contains inputs with a corresponding output
- Unsupervised learning: data is unlabeled, i.e. the training samples do not have an associated output
- Semi-supervised learning: mixture of labelled and unlabeled data
- Reinforcement learning: no data; the algorithm maps situations to actions to maximize a reward (Marsland, 2014)

Before machine learning can be implemented in a project, the data must be acquired from various sources and cleaned for use. Data cleaning involves identifying incomplete, incorrect, inaccurate or irrelevant parts of the dataset and then replacing, modifying, or deleting raw data. Oftentimes the data acquisition and preparation are the most time consuming and onerous part of the process. Unlike other industries where machine learning already has a solid foothold or is widely used in practice, in rock engineering this is not yet the case (W. J. Mcgaughey, 2019).

Machine learning algorithms present an opportunity to reduce the error associated with data manipulation, and offer predictive capabilities to increase the efficiency of the design process. As discussed previously, early work has shown their suitability for determining rock mass parameters and constitutive behaviours, as well as predicting geomechanical phenomena and instabilities. This section presents a brief overview of the most common types of machine learning algorithms, with an emphasis on ANNs. These are broadly divided into categorical prediction models and numerical prediction models.

2.5.1 Categorical Prediction Models

Categorical machine learning algorithms are useful for classifying data or making a categorical prediction. For example, given a dataset comprised of site-specific inputs for RMR or Q classifications, the algorithm can predict what the RMR or Q value is for a point ahead of the excavation face (i.e. the current extent of the excavation).

2.5.1.1 Decision Trees

The Decision Tree approach is a supervised machine learning algorithm that places data into classes and presents the results in a flowchart. The data flows through a query structure from the "root," or selected attributes, through subsequent partitioning until it reaches a "leaf" where no sample remains, no attribute remains, or the remaining samples have the same attribute (Mohammed et al., 2016b). The goal of creating a Decision Tree is to have a generalizable model that can classify unlabeled samples. Attribute selection is a crucial step when applying decision tree algorithms, as they must be meaningful to split the dataset at hand in to "purer" subsets (Einstein et al., 1978). A simple schematic is shown in Figure 2-2.

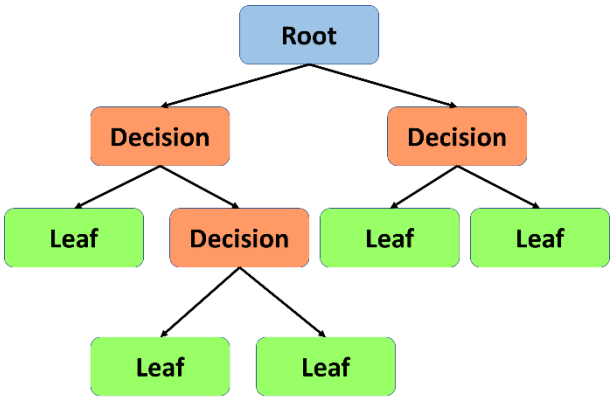

Figure 2-2. Decision Tree schematic showing root node, decision nodes, and leaf nodes.

In one case study, a Decision Tree method was employed to predict rock burst potential in a kimberlite pipe diamond mine (Pu et al., 2018). Here, the root is the linear elastic energy, the decision node is ratio between maximum tangential stress and unconfined compressive strength (UCS), and the decision node is the ratio between the UCS and the uniaxial tensile stress. This led to a classification, or leafs, of "no rockburst", "moderate rockburst", "strong rockburst", and "violent rockburst". The authors trained the algorithm with 132 training samples from real rockburst cases around the world, and subsequently the accuracy of the validation samples was shown to be 93%. Validation samples are those that have not been used to train the algorithm, and are therefor a metric for generalizability. The results of the study found that the mine under study was susceptible to moderate bursts, which matched the observed conditions.

2.5.1.2 Naïve Bayesian Classification

Bayesian networks are nodal networks that graphically represent probabilistic relationships between input variables, using particular simplifying assumptions. This technique is founded in Bayes' theorem, where the *posterior probability* of an event occurring is updated based on new *evidence* (M. Khan et al., 2017a):

$$Posterior\ Probability = \frac{Prior\ Probabiliy * Likelihood}{Evidence}$$ Equation 2-2

Naïve Bayesian classifiers are a type of Bayesian network that uses simple probabilistic classifiers that assumes independence between predictor variables (hence "naïve") (Araghinejad, 2014), as shown in Figure 2-3. To complete the classification, the numerator of Equation 2-2 is compared to each sample since

the evidence remains constant (M. Khan et al., 2017a). Though this type of algorithm can be advantageous because it is a simple representation of a problem and therefore easy to implement (W. J. Mcgaughey, 2019), a common criticism is that they assume independence between the input attributes (Ribeiro e Sousa et al., 2017).

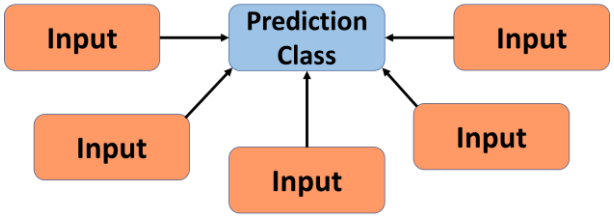

Figure 2-3. Naïve Bayesian classifier schematic showing predicted class and independent input variables.

A case study comparing different types of Bayesian network classifiers found that all the developed models showed a high accuracy rate when applied to predicting the magnitude of rock bursts for the dataset consisting of 60 cases (Ribeiro e Sousa et al., 2017). The naïve inputs included type and rock strength, geometry, stress state, and construction method, which were used to predict the magnitude of the rockburst. The naïve Bayesian model in particular classified 100% of overbreak cases, 83% of strong rock bursts, 25% of moderate rock bursts and 87.5% of slight rock bursts correctly, respectively.

2.5.1.3 k-Nearest Neighbours Classification

The *k*-nearest neighbours (*k*-NN) clustering algorithm is used for classification and is among the simplest of the machine learning algorithms. The *k*-NN approach can also be used to determine a numerical output using regression, as discussed in Section 0. The input consists of the *k* closest training examples and the output is a class membership. An object is classified by a majority vote of its closest neighbours, where the most common classifier is assigned (M. Khan et al., 2017b), as shown in Figure 2-4. The function is approximated locally, and computation is deferred until classification is complete. The user can assign weight to the contributions of the neighbours as a function of distance, for example. This algorithm is sensitive to the structure of the data, which may pose a limitation.

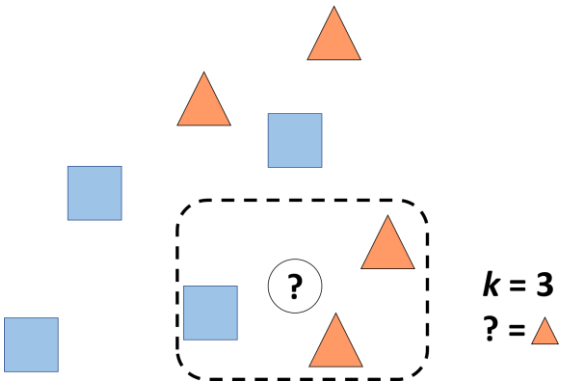

Figure 2-4. *k*NN classification schematic, showing *k* = 3 and therefore the three nearest neighbours are used to make the classification.

Work has been done comparing *k*-NN to four other supervised machine learning algorithms to classify geology using remotely sensed geophysical data (Cracknell & Reading, 2014). Inputs included a Digital Elevation Model, Total Magnetic Intensity, and four Gama-Ray Spectrometry channels, and the parameter *k* (number of nearest neighbours used for classification) was varied from 1 to 19. The authors conclude that as the spatial distribution of training data increases, the accuracy of the classifications also increase. They also conclude that explicit spatial information (coordinates) should be combined with geophysical data so that predictions are geologically plausible.

2.5.1.4 Support Vector Machine

Support Vector Machines (SVMs) are supervised clustering machine learning algorithms used for classification analysis (M. Khan et al., 2017d). The SVM maps the labelled training dataset as points in space divided into categories separated by a clear gap. New examples are mapped into the same space and are categorized depending on which side of the gap they fall into. The gap is defined by a hyperplane in two or three dimensions, depending on how many features are being used to classify the data. The main goal of the SVM is to find the hyperplane that maximizes the margin between the classes, as shown in Figure 2-5.

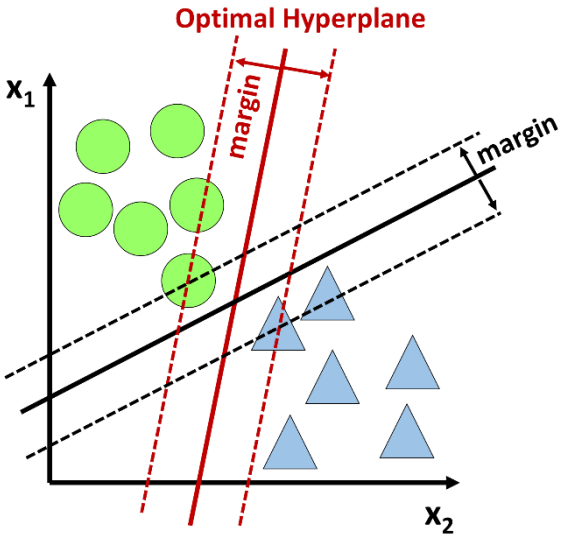

Figure 2-5. SVM schematic showing a two-dimensional classification problem, and the optimal hyperplane that should be used to classify future unlabeled data based on the smallest margin between the two closest data points.

A recent study applied SVM to predict tunnel squeezing based on four parameters: diameter, buried depth, support stiffness, and the Q Tunnelling Index (Sun et al., 2018). An 8-fold cross validation was used to create multiple models (called an ensemble) to get a measure of performance of the model. Cross validation is a common technique used in data driven methods to determine how well generalized the model is to an independent dataset, and how accurately the model will perform in practice. This process highlights whether the model is overfitting or if there is input selection bias. The resulting average performance of the algorithm was 88% and importantly this performance decreased to 74% when the support stiffness was not included in the SVM. The authors concluded that this method produced better performance in prediction accuracy as compared to existing empirical approaches, similar to other uses of the SVM cited there in. They were also able to estimate the severity of the potential squeezing based on the predicted squeezing class by introducing a multiclass SVM classifier trained using a database of 117 case histories. The multiclassifier was used to classify the database into the severity of the squeezing problem, similar to empirical classification schemes based on the strain or convergence.

2.5.1.5 Random Forests

Random Forest (RF) is a supervised ensemble classifying method that consists of many decision trees. The output is the majority vote of the classes output by the individual trees, as shown in Figure 2-6. Each decision tree is an individual learner, and the aggregate of each individual yields the prediction of the algorithm (PDAC, 2019a). RF is considered to be one of the most accurate classifying algorithms, running efficiently on large databases and effectively estimating missing data (PDAC, 2019b). However, RF has been noted to overfit for noisy datasets, and tends to be biased towards categories that are over represented in the training dataset.

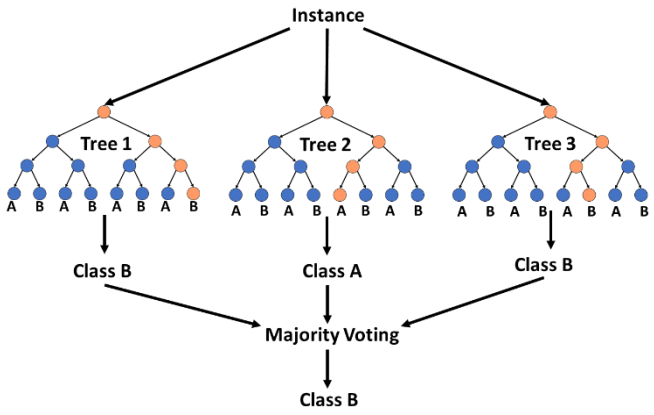

Figure 2-6. Random Forest schematic, showing an ensemble of tree decision tree classifications and the majority voting that determines the final class.

One case study applied a RF algorithm to predicting hanging wall stability, using a training dataset consisting of 115 cases (Qi et al., 2018). The inputs were subdivided into hanging wall geometry (stope dip, strike and height), geological properties (RQD, joint set number, joint set roughness, joint set alteration, dilution graph factors A, B, and C), and construction parameters (stope design method, undercut area, and stress category). Each of these inputs represents a “branch” of the RF structure. A 5-fold cross validation method was applied, and the grid search method was used to tune the hyper parameters. A common performance metric, the area under the receiver operating characteristic curve (ROC-AUC), was used to evaluate the accuracy of the classifications made by the RF algorithm and was shown to be 0.873 (out of a maximum of 1.0) for the testing dataset (the data subset withheld to ensure the generalized performance of the model). The authors state that this indicates their optimum RF model is excellent at predicting hanging wall stability, where “stability” is defined by the equivalent linearoverbreak\slough.

2.5.2 Numerical Prediction Models

Numerical prediction models are useful for obtaining an estimate where not enough information is available to do a conventional calculation (i.e. a discrete analytical solution), or where the relationships between the inputs and output are too complex for an analytical solution. For example, an algorithm trained on site-specific rock mass parameters and the in-situ stress field could predict the radial convergence of a tunnel before it is constructed. Or, a trained algorithm using an existing dataset comprised of typical rock mass parameters could determine the rock mass strength of an unmapped location. While other numerical predictors are briefly discussed, this chapter focuses on ANNs as a prediction tool, as they have been the focus of the research performed in machine learning methods for rock engineering design.

2.5.2.1 Support Vector Clustering

Support Vector Clustering (SVC) is an expansion of Support Vector Machines that is used when data is unlabeled or only some data is preprocessed (M. Khan et al., 2017d). SVC maps data points into a multi-dimensional feature space (where each feature is an input variable) using a kernel function (Ben-Hur et al.,

2001). The algorithm then searches for the smallest sphere that encloses the data in the feature space, and maps it back to the dataspace, where the sphere is transformed into contours that enclose the data that form part of the same group. Now unlabeled data has been classified.

Little work has been done applying SVC to geomechanical problems, however a case study from water quality literature is discussed here. A SVC algorithm was employed to model the electric conductivity and total dissolved solids in a river system, and was compared against a more conventional /genetic programming algorithm (Bozorg-Haddad et al., 2017). The authors concluded that the SVC method has better accuracy for modelling water quality parameters than the genetic programming algorithm (Jing & Stephansson, 2007).

2.5.2.2 k-Nearest Neighbours Regression

The k-nearest neighbours (k-NN) algorithm is used for regression, where the input consists of the k closest training examples and the output is the property value for the object (Khan et al., 2017). Similar to *k*-NN classification, the value is the average of its nearest neighbours, except the output is a numerical value rather than a classification.

Little work has been done applying *k*-NN regression to geomechanical problems, however an analogous case study completed to forecast the municipal solid waste (MSW) generated by a city (Abbasi & El Hanandeh, 2016). Four machine learning algorithms were compared: SVM, ANN, adaptive neuro-fuzzy inference systems, and kNN regression. The authors found that the prediction ability of the kNN regression algorithm was in the middle in terms of matching the observed data and peaks in the trends, however it was the best at predicting the monthly average values of MSW generated.

2.5.2.3 Artificial Neural Networks

ANNs are inspired by biological neural networks, where a series of highly interconnected nodes and a series of parallel nonlinear equations are used simultaneously to process data and perform functions quickly (Khan et al., 2017). ANNs are known as universal predictors and can approximate any continuous function under certain conditions (e.g. availability of appropriate input parameters). ANNs are powerful for real-time or near real-time scenarios as they can function with high volumes of data as it is collected (Bell, 2015).

In general terms, ANNs are comprised of an input layer, hidden layers, and an output layer. Each node is linked to all the nodes in the layer preceding and following it, where each link has a function defined by a weight and bias (Figure 2-7). Each node also has an activation function, which compares the weighted sum of all the inputs to that node and compares it to a predetermined threshold. If the threshold is exceeded the node fires and the input is transmitted further in the network through the activation function. The most common activation functions are the step, sign, linear and sigmoid functions (Khan et al., 2017). Activation functions should be chosen with care as they inform the response of the network to the inputs, and therefore the resulting output.

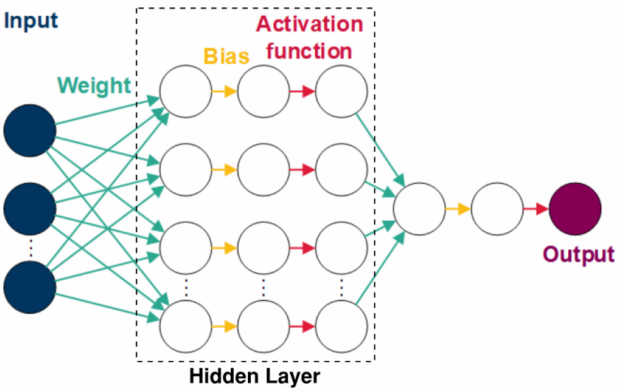

Figure 2-7. Simple schematic of an ANN

Multi-Layer Perceptron (MLP) are the simplest and most common form of ANN, and employ a learning technique called back propagation (Mahdevari & Torabi, 2012). Back propagation, short for backward propagation of errors, is used to adjust the weights and biases of the ANN by minimizing the error at the output. Back propagation is widely used in engineering and science because it is the most versatile and robust technique to find the global minima on an error surface (Trivedi et al., 2015). Back propagation consists of two steps: the propagation phase and the updating of the weight (Khan et al., 2017). In the propagation phase, the inputs are fed into the ANN and the values at the hidden and output nodes are calculated. In the second phase, the error is calculated at the output and then propagated backward to update the weights at the nodes using deterministic optimization to minimize the error sum (Bell, 2015).

The ANN is trained using a dataset to create a model that is able to make predictions using new information. In the case of an MLP, the network is inputs are fed forward and the error is back propagated to update the ANN iteratively until a solution is converged upon that matches the observed data within the error tolerance (M. Khan et al., 2017c). A concern with ANN development is “overfitting”, where the ANN predictions match all the observed data too closely and cannot handle new information well (Marsland, 2014). This can be avoided by using data partitioning techniques to avoid creating a biased model, which is a common approach for machine learning methods in general. The partitioned subsets are called the “training,” “validation,” and “testing” subsets. The testing set is withheld entirely during the ANN development. The training and validation sets can be used to create multiple models (called an ensemble) to get a measure of the performance of the model, called *k*-fold cross validation (Figure 2-8). This is good practice in ANN model development but currently rarely used in ANN applications in various engineering sub-disciplines (Abrahart et al., 2012).

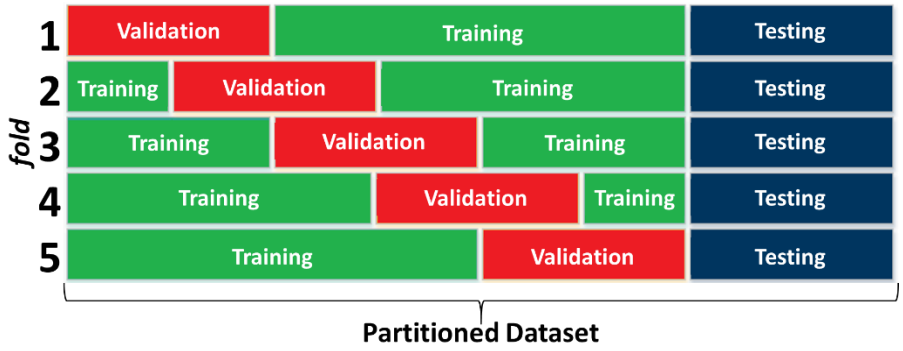

Figure 2-8. Example of data partitioning, showing 5-fold cross validation

ANN model uncertainties come from the choice of ANN architecture (i.e. number of hidden layers, number of neurons, choice of activation function, type of training algorithm, and data partitioning), as well as the performance metric chosen. Due to the data-driven nature of these models, propagating these uncertainties is easier (Khan & Valeo, 2016, 2017). One method of quantifying this uncertainty is by using fuzzy numbers to quantify the total uncertainty in the weights, biases, and output of the ANN (Khan et al., 2018; Khan & Valeo, 2017; Mosavi et al., 2018). This technique is useful for dealing with limited or imprecise datasets, and can be used to conduct risk analysis (Deng et al., 2011).

ANNs offer predictive and descriptive capabilities and have been applied to a range of rock property definition and rock engineering problems, including: intact rock strength, fracture aperture, rock mass properties, displacements of rock slopes, tunnel support, earthquake analysis, tunnel boring machine performance, among others (Jing & Hudson, 2002). Despite its wide applications, ANNs have not yet been proven to be a viable alternative to conventional methods to the rock engineering community.

2.6 Discussion of Machine Learning for Rock Engineering Design

As with all machine learning problems, the size and quality of the geomechanical datasets should inform the algorithm chosen for a given problem. The desired output is also a determining factor as to which algorithm is best suited to making the prediction. A summary of MLAs that have been researched for application to rock engineering problems are presented in Table 2-1.

Table 2-1 MLAs applied to various rock engineering problems

Rock Engineering Problem	MLAs	Opportunities
Rock mass properties	Categorical ANNs (Ching et al., 2019; Sklavounos & Sakellariou, 1995; Song et al., 2015)	Backwards predict rock mass properties based on observed site conditions Predict rock mass scale properties based on lab scale properties and rock mass behaviour
Laboratory testing and constitutive behaviour	Numerical ANNs (Kumar et al., 2013; Millar & Clarici, 2002) SVM (Kumar et al., 2013)	Use geology and peak Unconfined Compressive Strength to predict crack initiation and crack damage thresholds Use laboratory tests and field observations to predict constitutive behaviour Use rock mass scale classification (e.g. Q, RMR) to predict lab scale properties
Slope stability	Categorical ANNs (Fakir & Ferentinou, 2017; Ferentinou & Fakir, 2018) SVM (Kumar & Samui, 2014) RF (Hibert et al., 2017; Mayr et al., 2018) Clustering (Janeras et al., 2017)	Predict slope movements based on geometry, piezometers, inclinometer data, etc. Predict volume of structurally controlled failure based on mapped discontinuities
Point cloud analysis	RF (Weidner et al., 2019) *k*NN (Li et al., 2019)	Use successive tunnel scans to get volume differences and predict time-dependent deformations
Tunnel performance	Numerical ANNs (Bizjak & Petkovšek, 2004; Koopialipoor et al., 2019; Mahdevari & Torabi, 2012; Santos & Celestino, 2008) Categorical ANNS (Leu et al., 2001; Y. Xue & Li, 2018) SVM (Sun et al., 2018) RF (Einstein et al., 1978; Qi et al., 2018)	Predict stress/strain fields for input into numerical models Use preliminary/incomplete field mapping to prediction rock mass classification (e.g. Q, RMR) Predict tunnel support class based on rock mass classification (e.g. Q, RMR) Predict rock support performance based on geology, excavation method, environmental conditions, etc. Use microseismic monitoring arrays to predict rock mass deformation as excavation is developed
Rock bursts	Categorical ANNs (Ribeiro e Sousa et al., 2017) Naïve Bayesian classifiers (Ribeiro e Sousa et al., 2017) *k*NN (Ribeiro e Sousa et al., 2017) RF (Dong et al., 2013) SVM (Ribeiro e Sousa et al., 2017; Zhou et al., 2012) Decision Trees (Pu et al., 2018)	Predict magnitude and location of events using 3D excavation geometry, time series seismic events, mapped geology, etc. Use previous rock burst events, geology, etc. to predict magnitude of failed material and performance of rock support Use mapped rock classification (e.g. Q, RMR) to predict probability of rockburst
Blasting	Categorical ANNs (Liu & Liu, 2017; Vallejos & McKinnon, 2013) Numerical ANNS (Liu & Liu, 2017) SVM (Dong et al., 2011; Zhou et al., 2012) RF (Dong et al., 2011)	Use blast parameters and damage extent to predict optimum parameters for future blasts Predict blast parameters using mapped rock classification (e.g., Q, RMR)

In Table 2-1, column two summarizes the MLAs discussed in Section 2.5, while columns one and three tie these to various rock engineering problems discussed in Section 2.4. Note that the majority of the applications of MLAs are for categorical problems. To date, a variety of MLAs have been applied to classify

rock mass properties, rock bursts, and tunnel performance. The research that has been performed using MLAs to scale lab data and field observations to rock mass properties have been primarily categorical, as the rock mass classification schemes presented in Section 2.4.1 represent an industry accepted classification scheme. It is common for a project to have an incomplete dataset due to the cost associated with lab testing or field data collection. In practice, a basic statistical analysis may be performed to obtain conservative estimates of these properties, which are then scaled up to the rock mass scale properties using empirical or analytical methods. MLAs present an opportunity for determining the relationship between the lab scale and the rock mass scale properties, and potential to avoid the bias that can be injected by doing this manually. The algorithms for rock bursts and tunnel performance have been developed using inputs that represent physical properties, which that have been binned into categories that correspond to the severity of the rock bursts or tunnel deformation being predicted. Based on the classification schemes discussed in Section 2.4.1, the bins for these and other rock mass properties are already defined in engineering practice, and therefore practical to apply to a classification problem such as these. In practice, these properties are combined to calculate an overall quality score and then the support is designed based on this. MLAs present the opportunity to perform the rock support determination even if the dataset is incomplete, and allows the rock engineer to quantify the uncertainty associated with those predictions.

Little work in the literature has been done applying MLAs to the numerical modelling methods described in Section 2.4.2. Although sensitivity analyses are performed in practice, in general the numerical calibration process consists of manually adjusting the model parameters in a systematic manner until the numerical model outputs match the field observations. This requires careful adjustment of the input parameters followed by computation of the model to check the output against the observation or measured behaviour. For complex problems, running the model repeatedly is time consuming and therefore not done regularly in practice. MLAs present an opportunity to define the complex relationships between the input parameters, the numerical model behaviour and observed rock mass phenomena, and subsequently to conduct a more precise sensitivity analysis of the model inputs to ensure the rock mechanics relationships are being captured. An avenue for future research is surrogate modelling, where an MLA is used in conjunction with the numerical model outputs to systematically iterate through a distribution of inputs until the numerical model result matches the field observations. The MLA model may be used to calibrate the numerical model or simply for predicting an unknown state – hence the term “surrogate model”.

As shown in Table 2-1, several authors have found success applying ANNs to model complex rock mass behaviour over other algorithms, which is why they are emphasized herein. To date, the most common ANNs researchers are using are MLPs with one or two hidden layers (Trivedi et al., 2015). The main advantage of ANNs is that geometrical and physical constraints that govern rock mass constitutive behaviour and cause problems in numerical modelling approaches are not as problematic. This is because the data driven method does not rely on a function to capture all the anticipated rock mass behaviour, but

rather learns from the specific cases it is given to inform future predictions. An additional advantage is that ANNs can incorporate judgments based on empirical methods and can mimic the “perception” the human brain is capable of (Marsland, 2014). However, ANNs are limited to making predictions within the training parameters, meaning that they cannot make predictions outside the dataset it is given to train on. In other words, if there are not catastrophic events, such as a rock burst or falls of ground, included in the training dataset, then the ANN may not be able to predict these events given new data. Since ANNs are fundamentally a complex curve fitting algorithm, overfitting or underfitting may pose a concern to the general applicability of the final model. It is up to the developer of the ANN to ensure that the network has been validated and tested appropriately, and that suitable performance metrics (for example, the coefficient of determination (R^2), root mean square error (RMSE), precision/recall, receiver operating characteristic (ROC) curve) have been used to ensure the general applicability of the final network. Since ANNs are relatively new in the field of rock engineering design, there is a lack of verification and validation of ANN outcomes in the literature, and therefore cross-disciplinary literature review and independent validation using conventional design tools is necessary to prove their validity.

A unique challenge for rock engineering problems is the inherent spatial dependency of the variables, particularly when dealing with anisotropic and non-homogeneous rock mass properties. This may be considered explicitly as coordinate inputs, or implicitly by including inputs that are spatially variable. Some work has been done to determine the effects of these two input methods for ANNs, and it was noted that the model performed better when the spatial coordinates were included implicitly (Cracknell & Reading, 2014). An analogous problem has been addressed and is far more advanced in the field of image recognition, where the image is treated as a two-dimensional raster and inputs, or features, are mapped to a point in two-dimensional space on the image that is being processed. The techniques being used in that field may be applicable for encoding 2D and 3D spatial information in the geomechanical context. There is ample opportunity for further research in this area.

Time-dependent rock mass behaviour can be challenging to capture using numerical models (Paraskevopoulou & Diederichs, 2018). However, there may be an opportunity to couple an ANN with numerical models, particularly because ANNs can take time lagged variables as an input. Furthermore, convolutional layers may be added between the input node and the first hidden layer to allow a variety of input types (numerical, categorical, time-dependent) to be preprocessed and then combined to make a rock mass behaviour prediction by one network (Spence, 2018). A recurrent neural network (RNN) may also be developed, which allows the network to “remember” previous predictions made and use them to make further predictions (Mandic & Chambers, 2001). Further work in this avenue will have applications to solve both the time-dependent and spatial issues inherent to rock engineering modelling problems.

A continuum of data quantity and redundancy for geomechanical datasets is presented in Figure 2-9, based on this literature review and the authors’ collective experiences in both machine learning and rock mechanics. The data methods discussed in this chapter and those used in practice are plotted on this

continuum to show how much data and how many redundant data points are required to effectively use a particular data manipulation technique. Data redundancy indicates how many samples of a particular phenomenon or rock mass behaviour are required in order for the data method to be able to reliably calculate or predict it. Deterministic methods are located in the top left corner of Figure 2-9, with the lowest amount and redundancy of data required. This means that, if the correct data is available (through field or lab measurements), a deterministic method (e.g., a physics-based equation) can be easily implemented. However, common methods like regression have a higher data redundancy: several series of observations are required in order to “train” or “calibrate” a regression based model. The higher the redundancy in the dataset, the more confidence in the predictions of the regression model. The commonly used deterministic or probabilistic methods are usually supplemented by expert judgment or typical values when completing analytical or numerical rock engineering designs. Probabilistic, kNN, SVM, Decision Trees all have a relatively higher amount of data requirements than standard deterministic and are not generally suitable for smaller datasets, however reinforcement learning techniques (not covered herein) may become applicable as the algorithms become more advanced. Larger, more comprehensive, datasets, such as those collected in mining environments, are well suited to the machine learning algorithms discussed in this chapter, especially ANNs.

Operating mines often collect data continuously over time, for example: geological conditions, microseismic events, extensometer or crack meter data, in-situ stress measurements, groundwater and pore water pressure monitoring, among many, many others. This kind of multivariate dataset, where each parameter is intertwined and related to the overall stability in a complex way, is ideal for implementing an ANN. This is especially true if there are geomechanical events of interest that are ongoing, such as rock bursts or spalling, that have occurred over the life of the mine and were captured by some or all the instrumentation.

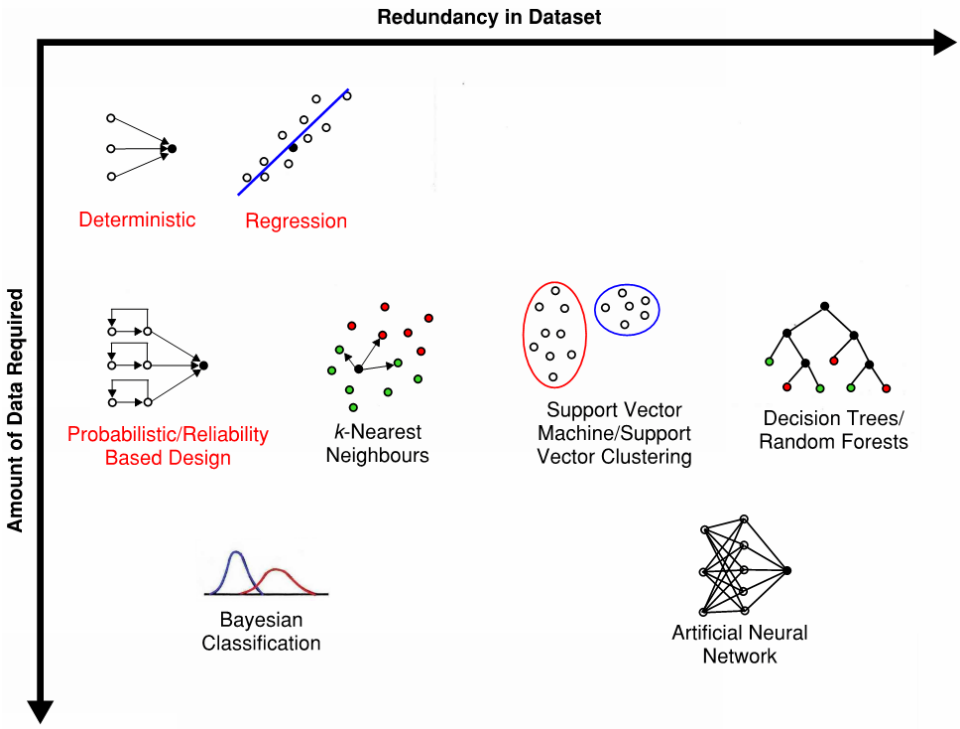

Figure 2-9. Comparison of data needs for machine learning algorithms included in this review (black) and some conventional rock engineering methods (red) in terms of amount of data and data redundancy required. Data redundancy indicates the representation of samples of the behaviour the data method should capture.

No matter the size or redundancy of the input dataset, it is crucial that the input selection and data partitioning scheme are appropriate for the problem to prevent overfitting and to ensure generalizability of the resulting model (Solomatine & Ostfeld, 2007a). Input variable selection is often done on an ad hoc basis (Abrahart et al., 2012), or using expert judgement and simple linear models (Martins & Miranda, 2013). Research in other geoscience fields suggests that input selection be done using a systematic approach based on a rigorous input ranking (Snieder et al., 2019). Formal input selection methods are used to determine which inputs from the larger input dataset are most useful in terms of relevance to the desired output prediction, while minimizing redundancies between input variables (Snieder et al., 2019). The framework for input selection has not been formalized and rarely receives the requisite attention (Maier et al., 2010a; Maier & Dandy, 2000; May et al., 2010).

While the importance of appropriate data partitioning is widely accepted (Shu & Burn, 2004), this process has not been formalized and is sometimes done arbitrarily (Maier et al., 2010b; Shahin et al., 2004). Poor data partitioning may result in a model with poor performance, and some work has been done to quantify the variability in the quality of training, validation, and testing subsets and their impacts on model performance (Daszykowski et al., 2002; May et al., 2010; Shahin et al., 2004). In particular, four main data division methods are prevalent in the literature:

1. Random data division
2. Data division ensuring statistical consistency within subsets
3. Data division using self-organizing maps
4. Data division using fuzzy logic methods

A thorough comparison on the impacts of these methods on model performance specifically for rock engineering datasets is needed to develop a framework for further application of machine learning methods in this field.

Model architecture and hyper parameter (i.e. number of hidden nodes, activation function, weights and biases, etc.) optimization are often done informally, despite the impact of these decisions on model behaviour (Maier et al., 2010a). Given that this approach is the accepted norm, Abrahart et al. (2012) asks if obtaining an optimal model structure is feasible or if the efforts required to acquire it are warranted (Abrahart et al., 2012). The authors go on to state that since many permutations of hyper parameters and architecture yield similar performances, and ad hoc approach may be appropriate for practical applications. However, if a more detailed and complex solution is required, numerous hyperparameter optimization methods exist, and have been applied in other geoscience fields, such as: *k*-fold cross validation to produce and ensemble of models (Xie & Peng, 2019), parameter regularization (Kumar & Samui, 2014), and metaheuristic optimization algorithms (Chou & Thedja, 2016).

2.7 Conclusions

This chapter briefly summarizes the current state of rock engineering design and presents the opportunities to integrate machine learning algorithms into the existing geomechanical design frameworks. A literature review has been conducted on the work that has been done specifically using ANNs on geomechanical datasets, and areas for future work have been identified. In practice, stability analyses are often performed relying on past experience and observational methods rather than on new data being collected on the geology and construction progress (Leu et al., 2001). ANNs are suitable for modelling complex rock mass behaviour and have been shown to be more efficient than regression functions (Leu et al., 2001). ANNs are especially powerful for repetitive construction processes as they can use real time data to update predictions ahead of the current stage of excavation. Future work on convolutional neural networks and recurrent neural networks will be invaluable to addressing spatial and temporal rock mass behaviour explicitly using ANNs.

Based on the body of literature available at this time, this review has found that there is a lack of standardization of the input selection process, data partitioning methods, model architecture and hyper parameter optimization, and performance measures in ANNs for rock engineering. These gaps present an opportunity to define a framework for integrating machine learning into rock engineering design to make the process more rigorous and reliable.

CHAPTER 3. DATASETS AND PROOFS OF CONCEPT

3.1 Chapter Introduction

This chapter presents a discussion of the two major datasets used to achieve the dissertation objectives, as well as the subsequent proofs of concept developed. A discussion of early work comparing ANNs to a Bayesian Belief Network (BBN) is also presented.

Two main datasets were used to complete the research in the present thesis, provided by research partner from mines in Canada with end member rock mass deformation behaviour. The Cigar Lake Mine data provided by Cameco Inc. represents a squeezing ground environment, while the Garson Mine data provided by Vale S. A. represents a high stress environment. These two datasets present different challenges with respect to developing machine learning algorithms, both in terms of developing the algorithm architecture to forecast the desired output and the data pre-processing and formatting requirements. This chapter introduces the datasets used in the remainder of the dissertation, including sources, digitization, formatting, and project context.

As with other types of model development, the development of a machine learning algorithm benefits from a gradual increase in complexity to ensure that the implications of the inputs selected, algorithm architecture, performance metrics, and engineering verification are understood and justified. This stepwise approach ensures that physical phenomena and sensitivity of model parameters are not overlooked. It is also important to understand the available data and its trends prior to increasing the complexity of the model. This chapter highlights some of the lessons learned during the data analysis and proof of concept stages of the research contained in this dissertation. These proofs of concept represent an experimental stage where alternative data driven tools, algorithm architectures, and output interpretations were explored prior to developing more sophisticated machine learning algorithms.

3.2 Comparison of Bayesian Belief Network and Artificial Neural Network

Early research was conducted to compare a probabilistic modelling tool, the BBN, with a simple machine learning algorithm, the ANN, as published in Morgenroth, Snieder, et al. (2019). The BBN and ANN tools were compared using a dataset from the Kemano hydroelectric power facility near Kitimat, British Columbia. The BBN developed was the subject of previously published research (Morgenroth, 2016). A data driven approach (the ANN) was developed to evaluate the prediction accuracy and level of effort required to develop it compared to the BBN, and to assess the comparative ease of data analysis for rock engineering purposes.

BBNs are compact, graphical representations of variables (called nodes) that are conditionally dependent on each other with direct links that reflect cause-effect relationships (Sousa & Einstein, 2012). The BBN nodes and conditional relationships are defined by the user, and inferences are made in the network using *a priori* knowledge to compute the query made to the BBN. BBNs have been applied in tunnel engineering

to forward predict and simulate tunnel construction progress as new data becomes available (Einstein et al., 1999; Sousa & Einstein, 2012; Špačková & Straub, 2013). The Kemano BBN was developed based on the geotechnical, topographical and geometrical aspects of the tunnel, to predict two Kemano-specific failure mechanisms: stress driven spalling failure and gravity/structure driven ravelling failure.

The Kemano ANN developed used the same inputs as the BBN: joint data (roughness, weathering, infilling, orientation), structure, groundwater, rock type, and in situ stress (major and minor principal stresses and orientations). The architecture was relatively simple consisting of one hidden layer with 8 neurons. Input data was partitioned into training (80%), validation (10%), and testing (10%) data. Bootstrap aggregating, called “bagging”, was utilized to generate an ensemble of ANN models and generate confidence estimates for prediction rates (Breiman, 1996). The Kemano ANN demonstrated overall classification accuracies of 88.2%, 81.7%, and 87.5% across the calibration, test, and complete datasets, respectively.

The BBN has the advantage of being capable of handling a mosaic of certainties in the data inputs and outputting a statistical distribution as an output instead of a deterministic value. The network can incorporate various degrees of certainty of inputs: single deterministic values, the likelihood of a particular state occurring, or a distribution that represents the possible values. The BBN is also transparent in that it allows the user to see the relationships between all the input nodes, the intermediary nodes, and the output nodes. All the conditional relationships can be controlled and coded according to empirical or analytical relationship that are widely accepted in literature, or according to the user’s expert judgment.

The BNN is a powerful tool if the rock deformation phenomena are fully understood, and the network is explicitly designed to capture that behaviour. However, developing and calibrating a BBN is time consuming and delicate work. It is difficult to develop a BBN that is generalized enough to apply to a multitude of sites. Due to the manual method of the network’s creation, it is difficult to avoid building in some form of bias into the BBN, leading to skewed outputs that may not capture the entire range of possible behaviours. Another limitation of BNNs and probabilistic methods in general is that the prior distributions of the input nodes must be specifically defined before their application. Often simplified distributions, such as a Gaussian curve, are selected for convenience. This has the potential for producing erroneous results.

ANNs also represent an advantage over conventional analytical approaches in terms of handling input uncertainty, in that inputs can be assigned a probabilistic distribution as opposed to a single, deterministic value. However, ANNs differ from BBNs because the user does not introduce bias into the model by encoding empirical or analytical relationships. Instead, the internal relationships between input parameters are determined by an objective training function. Machine learning methods can be used to identify useful inputs, validate results, and estimate uncertainty. The use of such methods facilitates model development compared to conventionally used models, as models can be trained and adapted with little expert knowledge. Instead, expert judgement is reserved for identifying candidate input parameters and interpreting model results.

Both BBNs and ANNs are powerful tools to streamline the data available for a tunneling project, though they differ is on the degree of oversight that is desired by the user. BBNs are entirely customizable and appear more transparent to a user that is unfamiliar with data driven methods. However, ANNs have the ability to learn complex patterns between inputs and the desired output which may not be easily discernible by the expert user, and where the anticipated geomechanical behaviour is too complex to be captured by a simplified network. This study comparing the two methods was found to be a validation of conducting further research on machine learning algorithms for geomechanical datasets.

The paper detailing the Kemano proof of concept is found in *APPENDIX A. Proofs of Concept* and is entitled *Comparison of Bayesian Belief Networks and Artificial Neural Networks for prediction of tunnel ground class*.

3.3 Cigar Lake Mine

3.3.1 Background

Cigar Lake Mine is owned and operated by Cameco, and is located in northern Saskatchewan, Canada (Figure 3-1). It is the world's second-largest uranium mine, with an ore grade approximately one hundred times the global average (Bishop et al., 2016). The ore body is unique due to its size, high grade, intensity of alteration, and a high degree of associated hydrothermal clay alteration (Bishop et al., 2016).

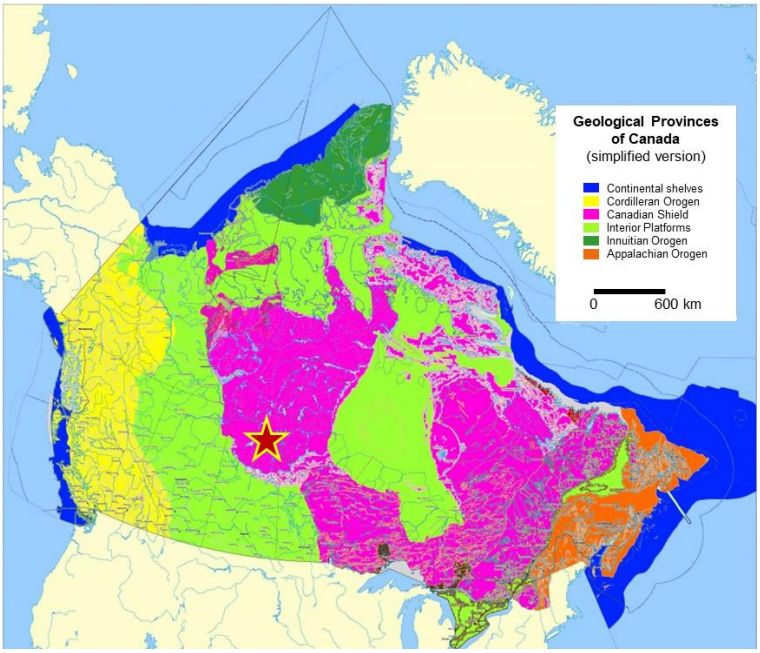

Figure 3-1. The location of Cigar Lake Mine in Saskatchewan, Canada (adapted from Geological Survey of Canada, 2009).

The geology of the Cigar Lake uranium deposit and environs is described in detail in *CHAPTER 4. A Convolutional Neural Network approach for predicting tunnel liner yield at Cigar Lake Mine* and is summarized here for context.

The uranium at Cigar Lake Mine is found in an unconformity type deposit. The deposit and host rock consist of three geological elements that also define the geotechnical domains: the deposit and associated hydrothermally altered clay cap, the overlying sandstone unit (Athabasca Group), and the underlying metamorphic basement rock (Wollaston Domain). The deposit and sandstone are highly fractured and water bearing, while the basement rock is impervious. The basement rock is composed mainly of pelitic metasedimentary gneisses and is considered the most favourable unit for uranium mineralization. In general, the Cigar Lake Mine operation and production tunnels are in three main rock mass types (Paudel et al., 2012): weak, highly weathered and saturated basement rock containing sand and clay; moderately weathered saturated basement rock; and strong unweathered basement rock.

The ore body is located above the 5.0 m lined diameter ore extraction tunnels, which are excavated using drill and blast. There are two main challenges facing the stability of the Cigar Lake Mine excavations: (i) controlling groundwater inflow, and (ii) supporting areas of weak rock (Bishop et al., 2016). The mine operators decided to freeze the rock mass surrounding the orebody to improve rock mass properties and to restrict groundwater inflow into excavated areas. However, the ground freezing operation results in complex time-dependent rock mass behaviour that is difficult to predict and presents challenges when designing support (Golder Associates, 2001; Roworth, 2013). The cavities created from the extraction of the ore are backfilled with concrete, and therefore provide additional ground support once they reach their full strength (Bishop et al., 2016). The general layout of the ore body and extraction method are shown in Figure 3-2.

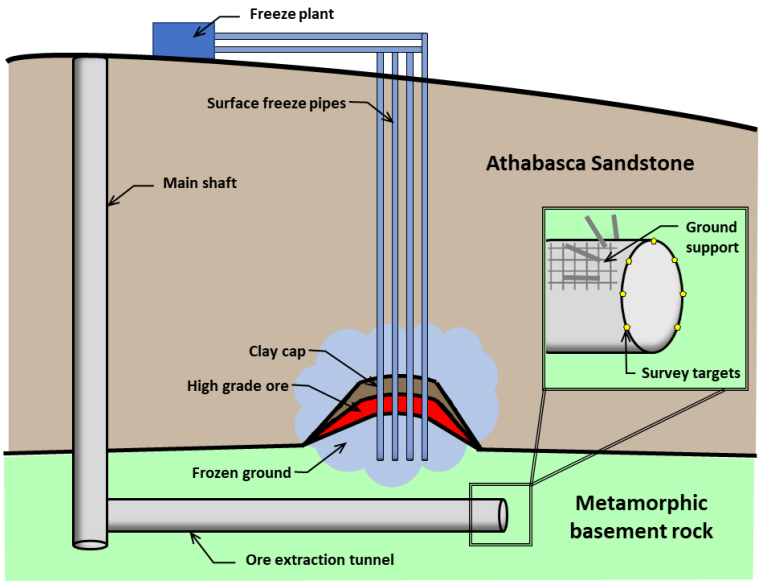

Figure 3-2. Schematic of the mining method at Cigar Lake Mine, showing a section view of the ore extraction tunnel and freeze holes from surface. The ore extraction tunnels are below the high grade ore body. Ground freezing is implemented to stabilize the excavations in the adverse geology and to manage ground water inflow from the Athabasca

Sandstone. Tunnel convergence is monitored using survey targets placed on the ground support around the circumference of the tunnel. Adapted from Morgenroth et al. (2022).

3.3.2 Data Preparation

The Cigar Lake Mine dataset is comprised of tunnel mapping, in the form of Ground Management Plans (GMPs), and survey data from the circumferential displacement measurements. The GMPs are produced at 6 to 8 weeks intervals depending on availability of geotechnical staff and ongoing mining operations. The GMPs contain data for an interval of time, including mapped geotechnical zones, observed damage or yield to tunnel support elements, surface and underground freeing operations, geotechnical monitoring stations, and locations of tunnel displacement monitoring. Figure 3-3 presents a typical GMP for Cigar Lake Mine.

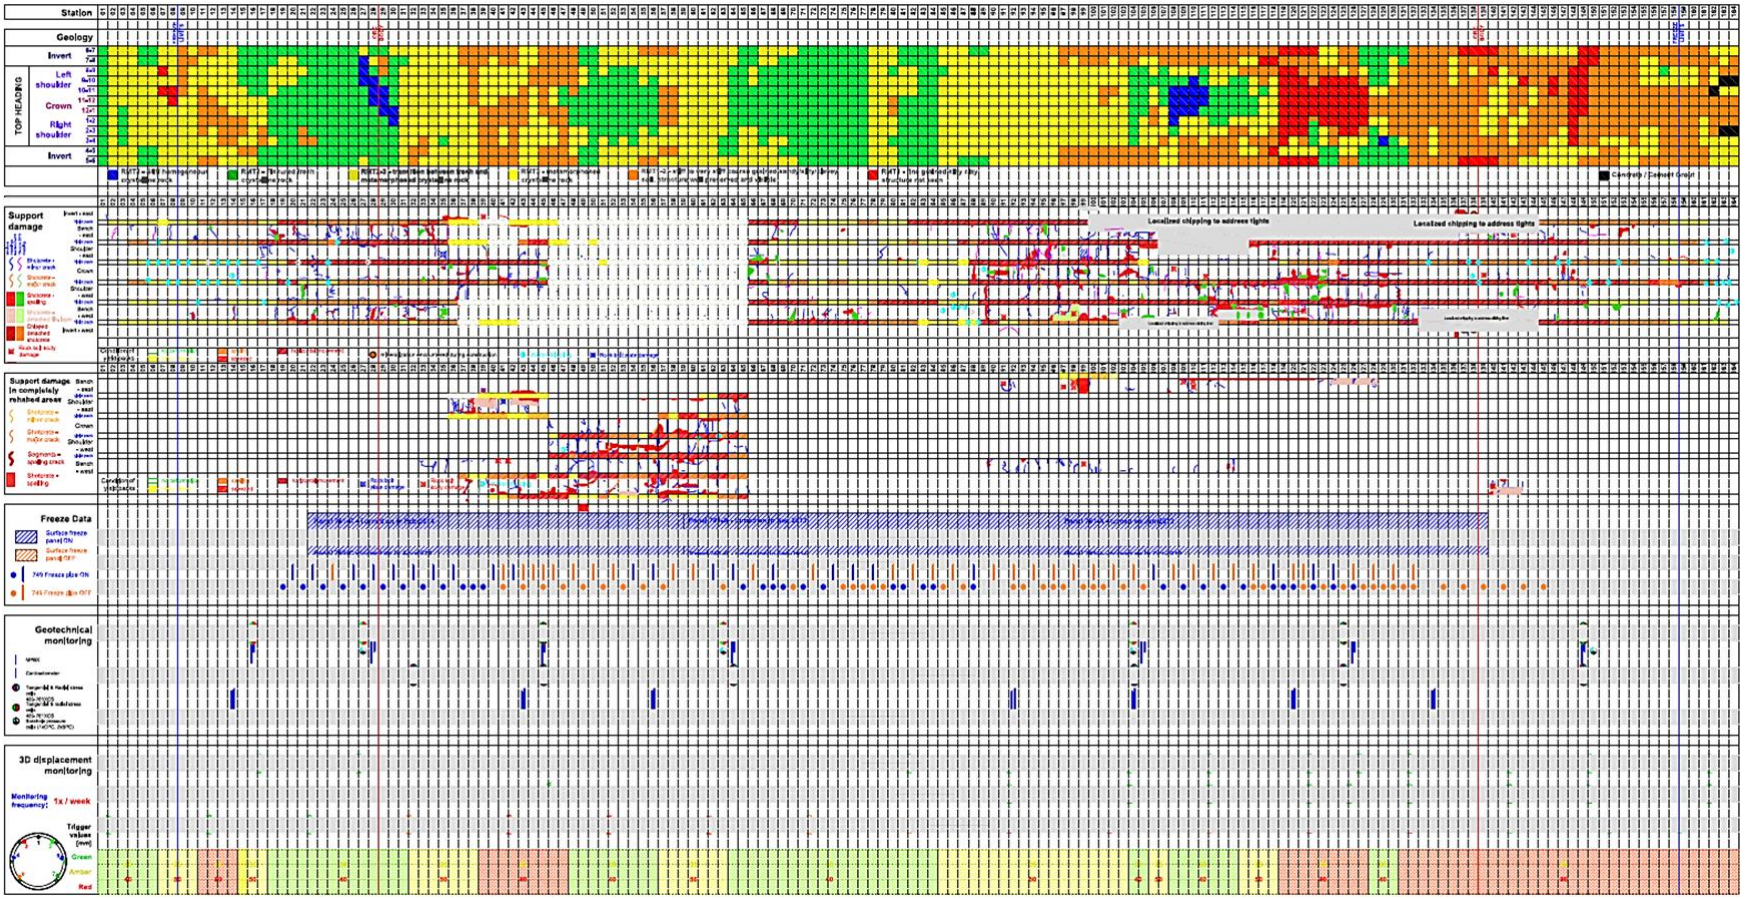

Figure 3-3. Sample Ground Management Plan (GMP) that was digitized to extract inputs for development of a Cigar Lake Mine machine learning algorithm (courtesy of Cameco).

The GMPs are essentially images, where the coordinates of the tunnel can be associated with pixel coordinates and various data can be translated to pixel intensity values. This allows for the translation of 3D spatial information into a 2D array of numerical or categorical values, which is a machine-readable form. This digitization exercise requires a coordinate system that will map locations on the tunnel to pixel coordinates, as well as a set of rules that will translate qualitative and categorical information to numerical values. Thus, the available datasets were digitized at a spatial resolution of 164 m (length of the tunnel) by 13 elements (representing elements around the circumference of the tunnel) as seen in Figure 3-4.

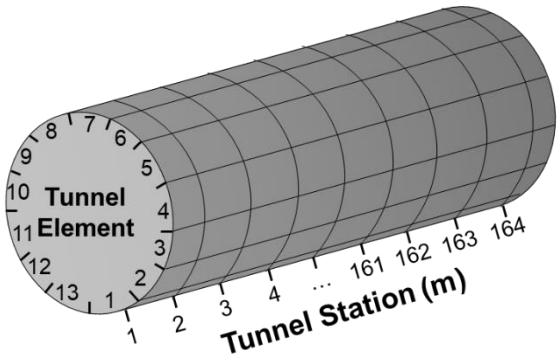

Figure 3-4: Nomenclature of mapped tunnel data

The GMPs correspond to 2015 Week 42, 2015 Week 50, 2016 Week 02, 2016 Week 10, and 2016 Week 24. This represents almost one year's worth of data. The input data that were digitized are: (i) the mapped geotechnical zones (GEO); (ii) as-built ground support class (SUPCL); (iii) ground freezing patterns (FREEZE); and (iv) radial tunnel displacement (DISP). Of the four inputs, GEO and SUPCL are spatially variable (i.e., inputs have a resolution of 164 m x 13 elements), while FREEZE and DISP are temporally variable (i.e., inputs have a resolution of 164 m x 13 elements x 5 time steps). A summary of each input digitized from the GMPs, and their corresponding definitions, are shown in Table 3-1. The categorical and binary inputs (GEO, SUPCL, FREEZE) were taken directly as presented on the GMP. These were the simplest to digitize. The DISP input was obtained from the displacement that is monitored using survey points around the circumference of the tunnel at 19 rings, with spacing between rings ranging from 4 to 12 m. Each ring has between 5 and 9 survey points that are captured when displacement measurements are taken. These measurements were interpolated around the tunnel circumference so that each tunnel element at each of the 19 rings had a DISP value. No interpolation between the measurement rings along the tunnel alignment was made, and instead the locations without measurements were assigned "not a number" during digitization. This prevents bias from being introduced, which would be possible if an analytical or numerical method was used to interpolate the magnitude of displacement between the measurement locations. This also results in a training dataset that is as "raw" as possible.

Table 3-1. Summary of inputs digitized from Cigar Lake Mine Ground Management Plans (GMPs).

GEO (164 x 13)	SUPCL (164 x 13)	FREEZE (164 x 13 x 5)	DISP (164 x 13 x 5)
1 – stiff homogeneous crystalline rock 2 – fissured fresh crystalline rock 3 – transition from fresh to metamorphosed crystalline rock 4 – metamorphosed crystalline rock 5 – stiff to very stiff coarse-grained soil, structure well preserved 6 – fine-grained soil, structure not seen	1 – class 1 support 2 – class 2 support 3 – class 3 support 4 – class 4 support 5 – class 5 support	0 – freezing off 1 – freezing on	Measured radial distance, in mm

The digitized output for this dataset was the yield to tunnel support elements (YIELD). The classes of YIELD are: Class 0 – no yield; Class 1 – minor yield; Class 2 – major yield; and Class 3 – total tunnel reprofiling required. These classes correspond with the amount of time, effort, and cost is associated with the required rehabilitation, and were defined in consultation with ground control engineers at Cameco. Based on these discussions, the YIELD classes were digitized from the GMPs as shown in Table 3-2.

Table 3-2. Summary of output digitized from Ground Management Plans (GMPs).

	YIELD (164 x 13 x 5)
Class	**Criteria**
Class 0	No yield recorded
Class 1	Minor isolated cracks in shotcrete, single water ingress spots, minor spalling of lining or yield elements
Class 2	High concentration of minor cracks, low number of major cracks, extensive spalling of the lining, spalling of yield elements, combination of cracks OR spalling AND water ingress, rock bolt yield
Class 3	Extensive spalling and major cracks, spalling AND cracks AND water ingress, cracks or spalling and rock bolt damage, squeezed yielding element, horizontal movement of yielding element

A detailed summary of all the digitized data and their distributions are presented in *CHAPTER 4. Section 4.6.1 Cigar Lake Mine Data*.

3.3.3 Lessons Learned from Proof of Concept

The Cigar Lake Mine algorithm development followed a stepwise process can be divided into three distinct phases: (i) preliminary data analysis and simple ANN development; (ii) CNN development and hyperparameter tuning; and (iii) Input Variable Selection (IVS) analysis. The onset of each of these phases is marked by a proof of concept, all of which are included in their entirety in *APPENDIX A. Proofs of Concept*. The relevant findings are summarized in this section.

Early study of the data patterns within the Cigar Lake Mine dataset, described in section *3.2 Comparison of Bayesian Belief Network and Artificial Neural Network*, attempted to characterize spatial correlations by analyzing the relationship between inputs in neighbouring pixels, not only focusing on a pixel-by-pixel study of the available tunnel mapping. This work was published in Morgenroth, Perras, et al., (2020). The mapping

data was treated as a raster, where correlations between a central pixel and its immediate neighbours were evaluated. For example, for pixel (i, j) the data from pixels (i±1, j±1) were also analyzed. The general data trends identified were as follows:

- Over time the severity of YIELD increases for all GEO classes, though the highest impact is to the areas with the poorest GEO.
- There is some variation of SUPCL assigned within each GEO class that impacts the severity of YIELD sustained at that location. This nuanced relationship is difficult to quantify from visual data inspection.
- YIELD in areas beneath where ground freezing is being implemented is higher than where the ground is not frozen. However, YIELD increases everywhere with time, indicating that other input variables are implicit drivers of yield.
- Poorer GEO classes experience higher values of DISP and therefore higher degrees of YIELD.

Prior to formatting the Cigar Lake Mine dataset as images, as described in section *3.3.2 Data Preparation*, a preliminary ANN was developed where the ANN was trained using 28 discrete inputs. The inputs are the GEO and SUPCL for pixel (i, j) and pixels (i±1, j±1), as well as the FREEZE and DISP across five timesteps. The ANN target was the YIELD at the fifth timestep. A relatively simple ANN architecture with one hidden layer containing 6 neurons was found to predict class 1 rock support damage reliably (84%), while over predicting class 0, under predicting class 2, and not being able to predict class 3. This inconsistent performance was theorized to indicate that the relative proportions of each of the YIELD classes is not adequately represented in the training dataset, and that the simplified method of including the pixels (i±1, j±1) was not sufficient to characterize the rock mass deformation phenomena. Thus, this proof-of-concept study determined that future work should focus on increasing or balancing the input dataset, investigating alternative ANN types and architectures, and validating the input selection.

Due to the limitations of the simple ANN developed for the Cigar Lake Mine dataset, an alternative algorithm architecture was selected for experimentation, to capitalize on the nuance in the spatial-temporal data and produce improved predictive performance. A CNN architecture was developed, where two main hyperparameters were investigated: the amount of temporal training data, and the convolution filter size. This proof-of-concept study is published in Morgenroth et al. (2020). The training data was varied so that the CNN was trained only on GMP 1 and tested on GMP 2, trained on GMPs 1-2 and tested on GMP 3, trained on GMPs 1-3 and tested on GMP 4, and finally trained on GMPs 1-4 and tested on GMP 5. The intention of this was to determine how many time steps are needed to get the most accurate prediction of YIELD. The convolution filter size was also varied from [2 m by 2 m] to [50 m by 50 m]. This meant that the YIELD was classified using information ranging from 2 to 50 m in distance from the pixel for which the classification was made, determining spatial sensitivity. This CNN proof-of-concept determined that the more temporal information the CNN is trained on, the more accurate the predictions become for all YIELD classes. This is because the CNN creates more robust relationships between the inputs and the output with

the addition of more temporal data. The optimal filter size was found to be [30 m by 30 m], which correspond to approximately 3 tunnel diameters (5 m tunnel diameter x 3 = 15 m in each direction).

Overall, this preliminary CNN was able to predict YIELD Classes 0-2 with >87% accuracy, however, was only able to predict Class 3 with 44% accuracy. Therefore, this study indicated that the CNN architecture was more suitable than the simpler ANN developed previously, however additional work was needed to achieve good prediction accuracy across all YIELD classes. A detailed study was conducted to fine tune several hyperparameters, including implementing an error weighting scheme, and further sensitivity analysis on the convolution filter size and the amount of training data. These detailed analyses are beyond proof of concept stage and are included in their entirety in *CHAPTER 4. A Convolutional Neural Network approach for predicting tunnel liner yield at Cigar Lake Mine*, which is also published in Morgenroth et al., 2021a.

The finalized Cigar Lake Mine CNN consists of two architectures: (i) a Global Balanced CNN, which uses an inverse frequency error weighting scheme to balance the dataset and achieve good prediction accuracy across all YIELD classes, and (ii) a Targeted Class 2/3 CNN, which uses a sigmoid error weighting scheme to prioritize Class 2 and 3 YIELD during training. In an attempt to increase the interpretability of these two CNNs, an IVS approach called Input Omission (IO) was adapted and applied as a proof of concept to increase algorithm transparency. This IO study was published in Morgenroth et al. (2021b) and is summarized here.

The IO method is a model-based IVS approach that estimates the usefulness of each input by iteratively examining model performance when an input is left out from the full set on which the model has been trained (Setiono & Liu, 1997). The significance, or lack thereof, of each input can then be compared based on the performance that is produced by leaving it out. Excluding redundant or irrelevant inputs from the CNN results in a model with a higher generalization capability (R. May et al., 2011). It is uncommon to apply IO to CNNs. This is because typical channels in input images are RGB or BW values, and for typical image processing it is not logical nor necessary to “turn off” individual channels. However, in the case of the Cigar Lake Mine CNN, each of the channels is a discrete and independent geotechnical input that has been formatted into an image channel. Therefore, a novel IO approach was adapted and tested as an IVS method for the Cigar Lake Mine CNN, as described in Morgenroth et al. (2021b).

The IO approach revealed that of the four categories of inputs used to develop the Cigar Lake Mine CNN, none of them could be omitted entirely to achieve better performance. The results also indicated that GEO and DISP contain the strongest signals for forecasting the severity of YIELD at Cigar Lake Mine. Practically speaking, these early IVS results indicated that the special emphasis should be placed on collecting the geotechnical zones in detail and the radial tunnel displacements with high temporal regularity. These results also suggested that the application of model-free IVS methods may offer further insights into the interpretation of the Cigar Lake Mine CNN, specifically with respect to post-hoc interpretation. This preliminary IO IVS study resulted in more in depth development of two further IVS methods.

The complete papers containing the Cigar Lake Mine proofs of concept described above are found in *APPENDIX A. Proofs of Concept* and are entitled:

- *An Artificial Neural Network approach for predicting rock support damage at Cigar Lake Mine: A Case Study* (Morgenroth et al., 2020)
- *Convolutional Neural Networks for predicting tunnel support and liner performance: Cigar Lake Mine case study.* (Morgenroth et al., 2020)
- *An Input Variable Selection approach for a Convolutional Neural Network that forecasts tunnel liner yield at the Cigar Lake Mine* (Morgenroth et al., 2021)

3.4 Garson Mine

3.4.1 Background

Garson Mine is a cooper-nickel mine owned and operated by Vale, and is located near Sudbury, Ontario, Canada (Figure 3-5). Garson Mine has been under development for over 100 years, with the first shaft sunk in 1907 (Sudbury.com, 2008). The mine is located on the southeast rim of the Sudbury Basin, and mining methods used are blasthole stoping and uppers retreat (Mining Data Solutions, 2022).

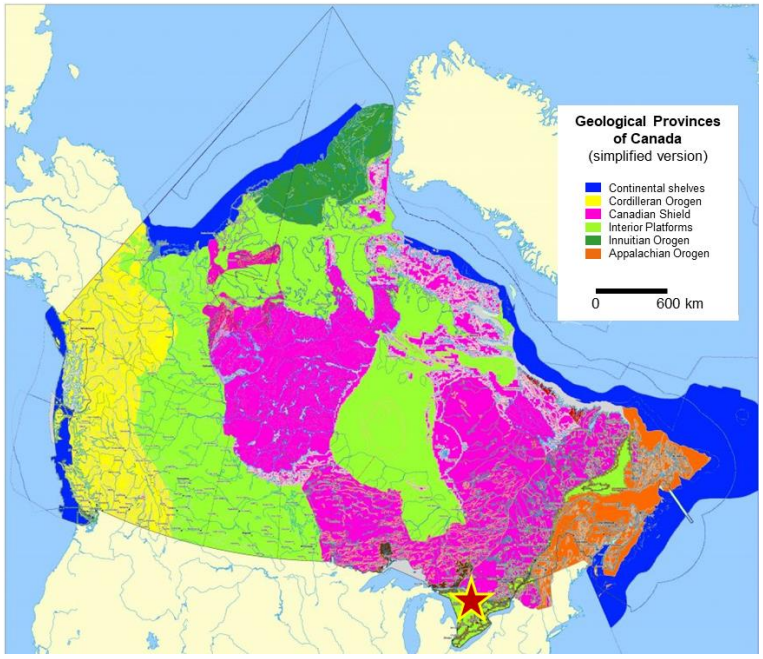

Figure 3-5. The location of Garson Mine near Sudbury, Ontario, Canada (adapted from Geological Survey of Canada, 2009).

The copper-nickel sulphide deposits at Garson Mine are hosted in parallel shears, which are offset by later stage dyke intrusions. The footwall is typically the lower zone norite of the Sudbury Igneous Complex and

Greenstone/Metabasalt metavolcanics, and the hangingwall consists of metasediments. The Garson Mine orebodies are more deformed than any other orebodies of the Sudbury Basin, and strike approximately east-west and dip south 75°. This research is focused on the #1 Shear West (1SHW) area of the mine, which is characterized by massive sulphide mineralization with sharp hangingwall and footwall contacts.

In response to large seismic events that occurred at Garson Mine between 2006 and 2008, strategic and tactical mitigation measures were implemented as a risk management approach to withstand future seismic impact (Yao & Moreau-Verlaan, 2010). The strategic measures relevant to this research are (i) the installation of a microseismic monitoring system to monitor ground movement and ground support effectiveness in high risk areas, and (ii) the ongoing numerical modelling that aimed to re-examine mining sequences. A microseismic monitoring system was installed to collect and evaluate continuous waveforms in real time, allowing for source-location and calculation of source parameters as events occur (Vale, 2015). In 2017 and 2018 the mine experienced an increase in frequency of large-magnitude seismic events in the 1SHW area. A FLAC3D model was developed by geomechanical consultants to back analyze the mechanisms that triggered these events, allowing for the evaluation of seismogenic risk associated with past, present, and future mine-induced stress evolution. The scope of the numerical modelling efforts included assessing the stability of the underground excavations and thus suggest operational changes to improve safety, reduce costs, and increase profitability (K. S. Kalenchuk, 2018). A sample simulation of this FLAC3D model is shown in Figure 3-6.

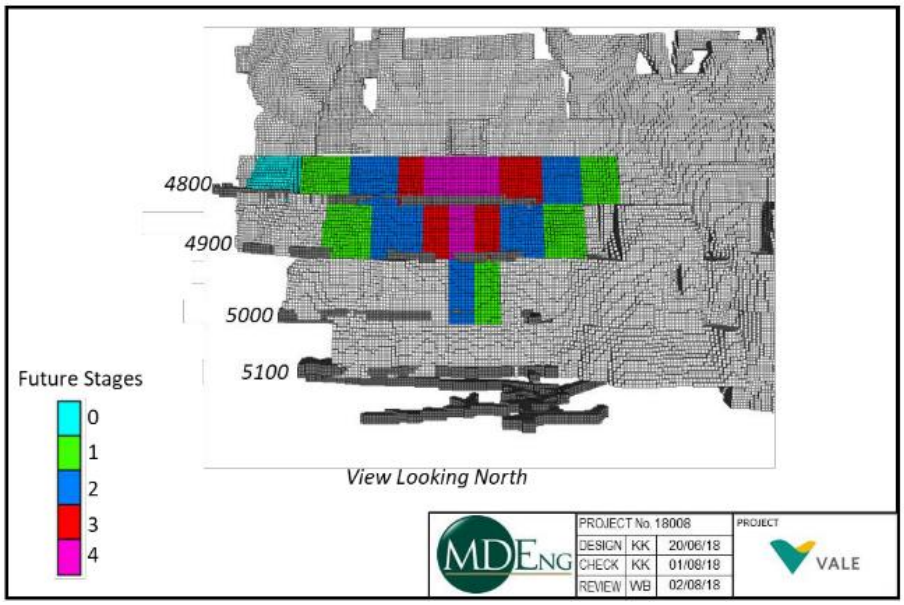

Figure 3-6. Garson Mine FLAC3D model, showing sequence of forward simulations with respect to stope mining. The intention of this simplified sequence was to demonstrate ongoing stress loading within the diminishing pillar over the remaining production in this stope block. (K. S. Kalenchuk, 2018).

3.4.2 Data Preparation

The Garson Mine dataset is comprised of the microseismic database dating from 2015-2018 and the previously calibrated FLAC3D model. The approach taken was to use microseismic events up to the date of the last manual FLAC3D model calibration to forecast the stresses at each zone in the FLAC3D model.

The microseismic parameters characterize the evolution of the stress regime leading up to larger seismic events and rock bursts, and inputs from the Garson Mine microseismic database were chosen following work by Zhang et al. (2021). However, a proof-of-concept study found that the microseismic data alone is not enough to accurately predict the stresses in the FLAC3D model, because all the seismic parameters in the database are highly correlated to each other (see *3.4.3 Lessons Learned from Proof of Concept*). For this reason, additional parameters from the FLAC3D model (e.g., material properties, geological zones, constitutive model) were added to the Garson Mine dataset to increase the variability and uniqueness within the dataset. The targets were the stresses from the FLAC3D model at each zone centroid. Two sets of targets were extracted for comparison: the principal stress tensor, and the six-component stress tensor. A list of all variable comprising the Garson Mine dataset is shown in Table 3-3.

The seismic database was filtered to the same Northing, Easting and Elevation as the FLAC3D model, to ensure the seismic events remaining in the database occurred near the 1SHW developments. Any seismic events with a source-location error greater than 20 ft (6 m) were removed from the database, to improve reliability of the data points.

Table 3-3. Summary of Garson Mine inputs (from microseismic database and FLAC3D model) and targets (from FLAC3D model).

Inputs	Targets
Time stamp	Major Principal Stress, σ_1
Microseismic event location (N, E, El.)	Intermediate Principal Stress, σ_2
Moment magnitude	Minor Principal Stress, σ_3
Seismic moment	
Energy	Normal Stress, σ_{xx}
E_s/E_p	Shear Stress, σ_{xy}
Apparent stress	Shear Stress, σ_{xz}
Source-location error	Normal Stress, σ_{yy}
Geological group	Shear Stress, σ_{yz}
Constitutive model	Normal Stress, σ_{zz}
Elastic modulus	
Poisson's ratio	
Hoek Brown m_b, s and a	
Accumulated plastic strain	

The Garson Mine dataset required formatting prior to processing by a machine learning algorithm. In this case a time-series (the microseismic events) will be used to perform a regression to values representing the stress state. Each microseismic event in the database was matched to the nearest zone centroid in the FLAC3D, then the seismic events that occurred near the same zone centroid were ordered into a sequence of seismic events (Figure 3-7).

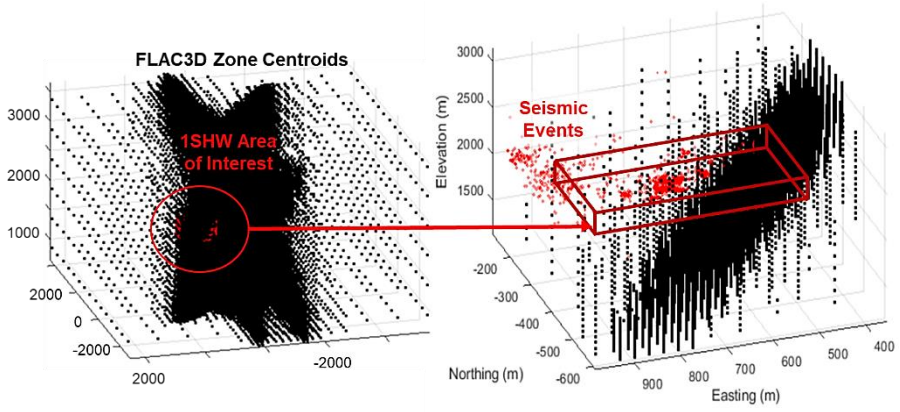

Figure 3-7. FLAC3D zone centroids, with 1SHW area indicated by the red cube. Seismic events (red points) were matched to the nearest zone centroid to train the Garson Mine LSTM network.

A detailed summary of the dataset and their distributions are presented in *CHAPTER 6. Section 6.5.1 Input Data*.

3.4.3 Lessons Learned from Proof of Concept

A simple preliminary LSTM network was developed for Garson mine due to the complex nature of the dataset, which is comprised of multivariate time-series microseismic data, as well as the desired outputs, namely the stress state in a complex FLAC3D model. This initial step was taken in advance of more complex sensitivity analyses and hyperparameter tuning. The preliminary LSTM network offered the opportunity to test the input data formatting described in section *3.4.2 Data Preparation*, as well as giving an initial indication about whether the microseismic database inputs were sufficient to provide good prediction of the stresses.

A simple sequence-to-one regression LSTM network was developed using the Garson Mine dataset, as described in Morgenroth et al. (2021). Two input format settings were investigated: the minimum number of seismic events that are included in a sequence for training, and the number of FLAC3D zone centroids that each discrete seismic event is assigned to. The former was an investigation of the minimum amount of temporal seismic data needed to make a reasonable forecast of the principal stresses. The latter was an investigation of one method to enhance the learning dataset, by assigning the same seismic event to multiple locations in the rock mass (represented by the FLAC3D zone centroids). This allows for more of the same seismic events to be included in multiple training sequences, creating richer sequences used for training the Garson Mine LSTM network.

A key finding of this proof of concept was that despite the large volume of microseismic data available, the input dataset is underdispersed. Underdispersion exists when outputs exhibit less variation than expected based on its distribution, and can occur when inputs are highly correlated with each other. This finding from the proof of concept was addressed in the subsequent Garson Mine LSTM network refinement. This is discussed in detail in Morgenroth et al. (2021) and is summarized here.

The problem of underdispersion found in the proof of concept stage was managed by increasing the number of inputs in the Garson Mine dataset. This provides the variability the LSTM network requires to learn the nuanced relationships between the inputs and targets. In the case of the Garson Mine LSTM network, the available inputs were increased by including additional parameters from the FLAC3D model (e.g., material properties, geological zones, etc.) in the dataset. The Garson Mine LSTM network was also developed to sample randomly with replacement during algorithm training, to allow all the training sequences to be used for both raining and calibration. Finally, four different loss functions were compared in order to select which function was best for capturing the non-normally distributed targets.

The paper detailing the Garson Mine proof of concept is found in *APPENDIX A. Proofs of Concept* and is entitled *Forecasting principal stresses using microseismic data and a Long-Short Term Memory network at Garson Mine*.

CHAPTER 4. A CONVOLUTIONAL NEURAL NETWORK APPROACH FOR PREDICTING TUNNEL LINER YIELD AT CIGAR LAKE MINE

4.1 Preface

This chapter focuses on the development of a Convolutional Neural Network (CNN) Cigar Lake Mine, Saskatchewan, Canada, to predict tunnel liner yield. Four inputs are used inn the CNN to make this prediction: geotechnical zone mapping, primary support class, ground freezing pattern, and measured tunnel displacement. A sensitivity analysis of CNN training parameters, called hyperparameters, is completed to optimize the final CNN performance. Hyperparameters analyzed include: the amount of training data, the convolutional filter size, and the error weighting scheme. Two final models are developed, one balanced model able to accurately predict tunnel liner yield across all classes of severity, and one targeted model that is calibrated to predict the higher classes of tunnel liner yield particularly well. This CNN represents a contribution towards Objective 1, which is to develop a classification machine learning algorithm using standard geotechnical mapping data from a real project.

The content of this chapter was published in *Rock Mechanics and Rock Engineering* in 2021 as follows:

Morgenroth, J., Perras, M. A., & Khan, U. T. (2021). A Convolutional Neural Network approach for predicting tunnel liner yield at Cigar Lake Mine. *Rock Mech Rock Eng*. https://doi.org/10.1007/s00603-021-02563-3

The contributions of the authors in the current chapter are as follows:

Josephine Morgenroth has conducted the literature review, performed the data collection and preparation, developed the algorithm using the required software to perform the analysis and modelling, validated and visualized the results, and prepared and wrote the original manuscript of this publication. **Matthew A. Perras** has supervised the research, provided the funding, and contributed to writing and editing the manuscript. **Usman T. Khan** has supervised the research, provided the funding, and contributed to the writing and editing the manuscript.

The authors would like to extend special thanks to Cameco, and particularly Chris Twiggs, Imre Bartha and Kirk Lamont for their constructive feedback and informative conversations. This work is funded in part by the Natural Sciences and Engineering Research Council of Canada through the Discovery Grant program and the joint Innovation York and National Research Council Canada's Industry Research Assistance Program – Artificial Intelligence Industry Partnership Fund, in partnership with Yield Point Inc. This work is also funded by the Ontario Graduate Scholarship program.

4.2 Abstract

As underground instrumentation improves and the storage of large volumes of data becomes more cost effective, the rock engineering community has access to larger datasets than ever before. Machine learning

algorithms (MLAs) present an opportunity to uncover nuanced rock mass deformation mechanics more efficiently than conventional data analysis tools, resulting in increased reliability of underground excavations. MLAs require appropriate pre-processing of inputs as well as ground truth validation of outputs. CNNs are an MLA that allow for the preservation of spatial and temporal dependencies within a dataset. CNNs were developed for image recognition and segmentation, such as video processing, and are efficient at analyzing sequential snapshots of an excavation as the environmental and in-situ factors change.

Herein a CNN is developed for Cigar Lake Mine, Saskatchewan, Canada, to predict tunnel liner yield. The mine experiences a complex time-dependent ground squeezing behaviour resulting from the poor geological conditions and the artificial ground freezing implemented to stabilize the ore cavities and to control ground water during the ore extraction process. A sensitivity analysis of the CNN training parameters, called hyperparameters, is completed to optimize the final CNN performance. Hyperparameters analyzed include: the amount of training data, the convolutional filter size, and the error weighting scheme. Two final models are developed, one balanced model able to accurately predict tunnel liner yield across all classes of severity, and one targeted model that is calibrated to predict the higher classes of tunnel liner yield particularly well. Model results demonstrate that the CNN is a promising tool for preserving the spatial and temporal dependencies between input variables, and for predicting tunnel liner yield. This is a novel approach for geomechanical datasets. In combination, the two final CNNs achieve a prediction precision of >87% across all classes and a recall of up to 99.9% for the higher yield classes. The activation strengths of the inputs were studied, and it was determined that the primary installed support class is the most dominant predictor of tunnel liner yield.

4.3 Introduction

Decision making in the field of rock engineering is challenging due to the inherent uncertainty associated with the behaviour of natural earth materials. Since the 1960s when rock mechanics was formalized as a field of study (Hoek, 1966), many empirical and numerical methods have emerged. These methods tend to base design decisions on previous knowledge of similar materials or locations and to account for the uncertainty by examining ranges of material properties (Jing & Hudson, 2002). Empirical tools proved to be a useful way of incorporating previous experience for otherwise unknown conditions, and numerical tools allowed for efficient computation of scenario analyses. While both these methods have many strengths, there is a common weakness: in practice, the data of the problem at hand is often conformed to a framework that does not exactly describe the local rock deformation phenomena. In empirical underground design methods (e.g., Barton et al., 1974; Bieniawski, 1993), the ground support recommendations are based on ground conditions from numerous case studies from around the world but may not match the rock deformation mechanics at the site for which the support is being designed. Numerical modelling, such as the Finite Element Method, Finite Difference Method, Discrete Fracture Network modelling, and hybrid methods, all rely on constitutive behaviours that may not capture the combination of factors that result in

the observed rock deformation mechanics (Morgenroth, Khan, et al., 2019). There is a risk of introducing bias to engineering decision making, because the actual rock mass deformation phenomena may be overly simplified or mischaracterized as a result of fitting the site-specific data available to an empirical or constitutive behaviour that does not capture all the observed rock mass deformation mechanics. This is commonly overcome at present by examining several scenarios, which increases the time needed to develop a suitable design.

These limitations may be overcome by applying a data driven approach, such as statistical or machine learning methods. Statistical methods, such as reliability-based design (Baecher & Christian, 2003; Bozorgzadeh et al., 2018; Langford, 2013), have been developed to quantify uncertainty in geomechanical datasets and build variable relationships that help engineers understand the associated rock and rock mass behaviour. Machine learning methods can identify patterns or relationships between datasets and does not require expert intervention to form inter-variable relationships (Khan & Valeo, 2017; Solomatine & Ostfeld, 2007). Instead, the expert's judgment is reserved to choose the relevant input data and to evaluate whether the outputs are mechanistically viable.

MLAs have emerged as a tool to be added to the rock engineering toolbox for investigating complex rock deformation mechanics (Morgenroth, Khan, et al., 2019), though they are not yet prevalent in rock engineering practice. MLAs are advantageous because of their ability to compute large volumes of numerical and categorical data (Marsland, 2014) – which today is collected at an increasing rate but is often not examined efficiently and thoroughly. MLAs provide a process to extract more value from the collected data, and to ease the burden of manual data manipulation and analysis. The availability of adequate and relevant data is a prerequisite for developing a well generalized and robust MLA. To date, research at the intersection of machine learning and rock engineering has included predicting: rock mass properties (Sklavounos & Sakellariou, 1995; Song et al., 2015), constitutive behaviour (Kumar et al., 2013; Millar & Clarici, 2002), slope stability (Ferentinou & Fakir, 2018; Janeras et al., 2017; Kumar & Samui, 2014), tunnel performance (Bizjak & Petkovšek, 2004; Leu et al., 2001; Qi et al., 2018; Sun et al., 2018), rock bursts (Pu et al., 2018; Ribeiro e Sousa et al., 2017), and blasting (Dong et al., 2011; Liu & Liu, 2017). These studies have been completed using a variety of MLAs, ANNs, a variation of which will be used in the present research.

Geomechanical datasets are generally challenging to work with due to the combination of data types and formats, including categorical (e.g., rock mass classification), numerical (e.g., instrumentation readings), spatial (e.g., geological mapping), and temporal (e.g., seismic monitoring) data. Combining these various data is challenging in most data driven methods due to the preprocessing required, which in addition to being time consuming and tedious, presents the risk of inadvertently injecting bias (Morgenroth, Khan, et al., 2019). This chapter presents the novel development and application of a CNN, which was originally designed for image classification and computer vision applications, and explores its suitability for processing rock engineering data with spatial and temporal dependencies.

Herein a CNN is developed and applied to a spatially and temporally dependent dataset pertaining to Cigar Lake Mine, located in northern Saskatchewan, Canada, where the ore production tunnels are experiencing squeezing ground conditions resulting in yield of the tunnel lining. The CNN was developed to predict the yield of the tunnel lining, where liner yield is defined as combined failure state of the rock support elements including shotcrete, mesh, and rock bolts. Four inputs are used inn the CNN to make this prediction: geotechnical zone mapping, primary support class, ground freezing pattern, and measured tunnel displacement. The CNN model was designed to predict high tunnel liner yield as accurately as possible, however, the overall performance of the model for lower classes of yield must not decrease as a result. In addition to prediction accuracy, the time-dependency of the input data is quantified by analyzing the impacts of the extent of historical data on model performance, i.e., how far back in time data is needed to make an accurate forecast. Additionally, the parameters used to train the CNN (called *hyperparameters*) were calibrated to determine the optimal values that result in the best model performance. Finally, the activation strengths of the CNN model inputs (geology, support class, ground freezing, tunnel displacement) were evaluated to determine the inputs that are the strongest predictor of tunnel liner yield at the Cigar Lake Mine.

4.4 Machine Learning and Underground Rock Engineering

Data driven machine learning methods have emerged as powerful tools because they are able to handle complex problems with large datasets much more efficiently and with a higher accuracy than manual data analysis techniques (Liu & Yang, 2005; Marsland, 2014; Papadopoulos et al., 2000). MLAs offer predictive abilities for nuanced data patterns within a framework that allow the developer to ensure that a generalized solution is found. These data driven methods pose an opportunity to the rock engineering community because complex rock mass phenomena are not constrained by generic constitutive behaviours, and so more complex behaviours may be captured by the algorithm (Morgenroth et al., 2019). However, their application in practical rock engineering is only in its initial stages.

Many types of MLAs exist – the choice of which one to use is a function of the question to be answered. For example, ANNs are well-suited to time series data and numerical predictions, Random Forests or k-Nearest Neighbours for classification, and Self-Organizing Maps for partially or completely unlabeled data (Morgenroth et al., 2019). In particular, a type of MLA, called ANNs, are gaining popularity in the research literature. Examples of specific rock mechanics research using ANNs include: classifying rock masses (Sklavounos & Sakellariou, 1995); determining rock strength properties (Singh et al., 2001); modeling stress-strain behaviour (Millar & Clarici, 2002); predicting tunnel convergence (Mahdevari & Torabi, 2012); predicting tunnel-induced ground settlement (Chen et al., 2019); back-analyzing geomechanical properties (Song et al., 2015); and predicting rock bursts (Afraei et al., 2019).

ANNs consist of a series of interconnected nodes in consecutive layers, represented by a set of parallel nonlinear equations, which are used to simultaneously process data and perform functions (Khan & Valeo, 2016). In their simplest form, they are analogous to linear or logistic regression. Thus, ANNs are objective

in the sense that the can recreate data trends by applying an optimizer for a given dataset, and no additional subjectivity is introduced beyond what already exists in the input data. The structure of these networks is inspired by biological neural systems, mimicking neural activity in human brains. Activation functions are used in ANNs to transform the input signals into an output or “activation strength” (Marsland, 2014). For example, a series of rock mass conditions may be the inputs used to predict rock mass deformation as shown in Figure 4-1. This architecture allows the algorithm to *learn* data patterns by back-propagating the error between a predicted output and the corresponding observed values. The calibration parameters, known as weights and biases, define the connections between the neurons, and are updated to minimize the error between the predicted and observed data (Marsland, 2014). These weights and biases are initialized randomly and converge on a solution through back-propagation during the learning process (also known as *training*).

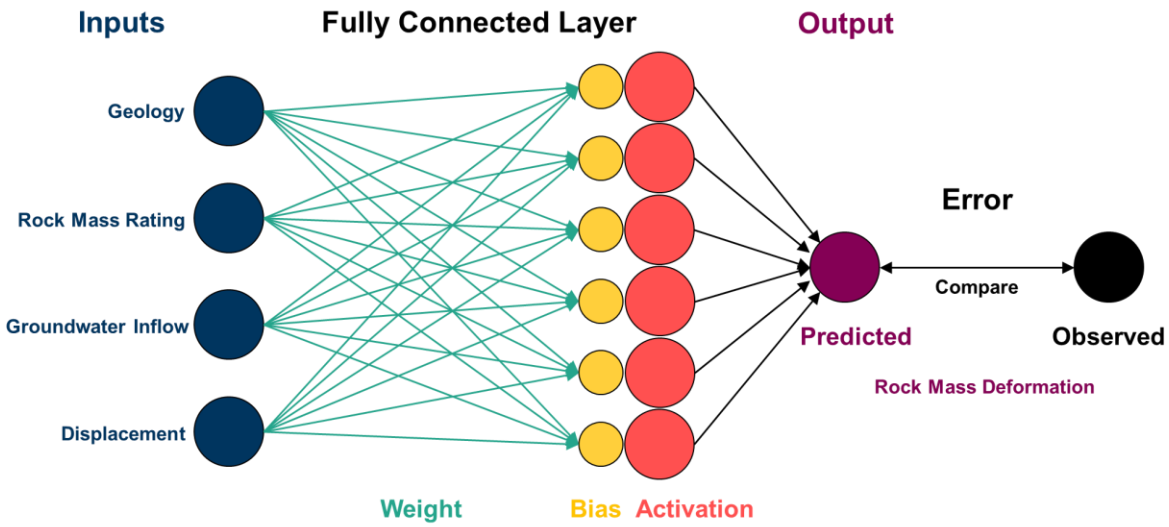

Figure 4-1. A schematic of an Artificial Neural Network (ANN) and its architectural elements, including the inputs and outputs, the calibration parameters (the weights and biases), and the activation functions.

A main concern when developing an ANN is *overfitting* the model to the dataset, i.e., producing an over parameterized model that matches the input data too closely without being generalized enough to tolerate new data (Snieder et al., 2019) This is undesirable because an overfitted model will reproduce the training data reliably, while not being able to generalize to an independent dataset (e.g., a new area of a mine). This pitfall is reduced by partitioning the dataset used for model development into training/validation and testing subsets (Khan et al., 2018). The training subset is used to *teach* the model the data patterns and to calibrate the model parameters, and the validation subset is used to determine when the model has been sufficiently trained and the stopping criteria has been met. Lastly, the testing subset is withheld entirely during this process, hence is an independent dataset, and is only used to determine the performance of the model once the final calibration parameters have been determined. This data partitioning exercise can be completed multiple times, using randomly reinitialized weights and biases each time, to create an *ensemble*

of models, each of which converge to a different solution. This ensemble-based approach makes it possible to quantify the uncertainty of the ANN.

ANNs are only able to make predictions that are included in the dataset used for training (Marsland, 2014). It is common that the phenomena the algorithm is trying to predict are under-represented in the training dataset, for example catastrophic events like rock bursts or large falls of ground. The imbalanced nature of typical rock engineering datasets means that trained ANNs will be most sensitive to the most represented phenomenon. This can be avoided by introducing *error weighting*, which instructs the algorithm which classification of the output should be given preference during the training phase (Seif, 2018). For example, an algorithm being trained to predict rock bursts would receive the highest weight on the error between the predicted large rock bursts and actual large rock bursts. This way the ANN prioritizes reducing this error for large rock bursts during the training process, whilst giving relatively lower priority to smaller or non-rock burst events. This leads to improved model performance in predicting large rock burst events.

This research focuses on CNNs, a type of ANN, which are efficient for processing spatial and temporal dependencies in image or raster datasets. CNNs were first developed for handwritten digit classification and have rapidly gained success in the field of computer vision since their introduction in the late 1980s (LeCun et al., 1989). They are a popular choice because they are computationally efficient at processing each pixel in an image, while considering both the surrounding pixels and their change over time. For example, in a video the frames before a selected frame can be used to train the network to increase its prediction accuracy of the image in the selected frame. This is relevant to rock engineering problems because the data usually contains a mapping component (e.g., mapped geology) onto which other data (e.g., measured displacement or groundwater inflow) can be transferred as a two-dimensional array and therefore plotted as an image or a map. These maps are akin to frames of a video, where some data (e.g., mapped geology) stay constant with time, while others (displacement and groundwater), change with each frame. When developing any kind of rock mechanics model, the spatial and temporal relationships between the various inputs must be preserved, and thus, CNNs are well suited for these types of datasets. Previous work utilizing the Cigar Lake Mine dataset presented a simple Multi-Layer Perceptron (MLP) ANN, and resulted in a poor classification accuracy for higher classes of tunnel liner yield (Morgenroth, Perras, Khan, et al., 2020).

CNNs allow for the tunnel mapping to be used as an image input, without any further processing required. CNNs process images by using a filter to scan, or *convolve*, over the pixels of the input image to create a *feature map* and which is then correlated to the output. Specifically, a square search area (i.e., the filter) is used to identify the presence of a feature of interest (e.g., a crack in the shotcrete liner) to produce a feature map. This feature map is then correlated to the predictor (e.g., the degree of yield the tunnel liner has sustained over time). The CNN will produce multiple feature maps that identify the combination of factors that produce the output, e.g., unfavourable geology plus other input data in a certain area results in a classification of high hazard within that area (a schematic of this process is shown in Figure 4-2). In addition

to processing images efficiently, CNNs have demonstrated success at more challenging computer vision tasks involving classifying objects in an image or video (Zeiler & Fergus, 2014), called *segmentation*. This type of CNN architecture allows for classification of multiple phenomena within an input as opposed to classifying the image as a whole (Morgenroth et al., 2020). For example, the CNN with segmentation will not classify an entire tunnel map with one level of deformation but rather identifies areas within the map with various amounts of deformation.

Along with the size of the filter used to convolve over the inputs, there are two other parameters to consider when developing a CNN: the *stride* and the *padding*. The *stride* is the amount by which the filter shifts when convolving over the input (i.e., how many cells it traverses at a time). When the stride is greater than 1, the size of the feature map becomes smaller than the size of the input image after convolution due to the overlap at the edges of the input image. The *padding* parameter adds a border of nonvalues around the input layer, so that the input and the feature map remain the same size, and the information contained in the original input is preserved. In the example shown in Figure 4-2, a 3 by 3 filter with a stride of 2 is used to convolve over the inputs. No padding is required in this example because the size of the input (12 by 12) agrees with the calculated size of the feature map (5 by 5) using those chosen filter parameters. Spatial and temporal relationships are crucial to the understanding of rock deformation mechanics, especially in a squeezing ground environment where progressive failure is occurring. Thus, it follows that the application of a machine learning approach that preserves these relationships should be chosen to address this type of problem. The complex nature of the rock mass deformations, as well as the time-dependency of the rock mass behaviour, resulted in the selection of a CNN for forecasting tunnel liner yield at the Cigar Lake Mine.

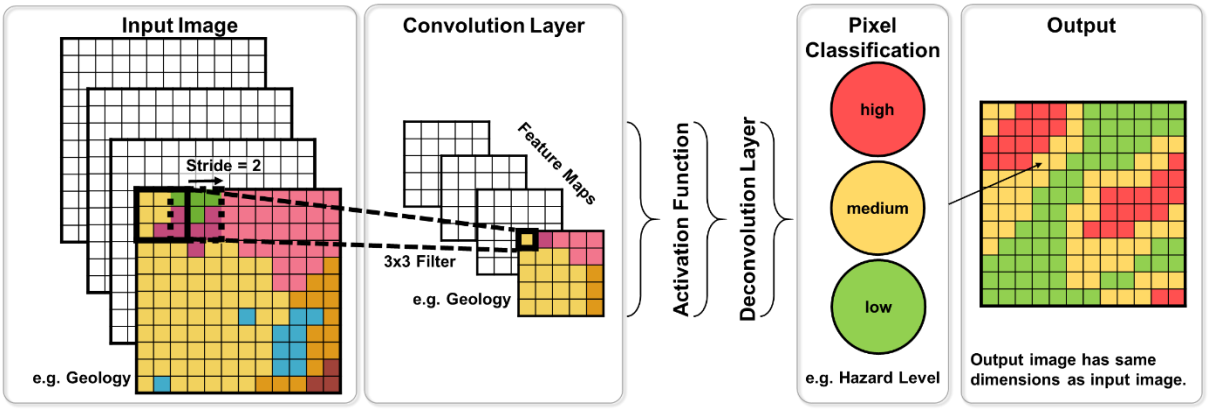

Figure 4-2. A schematic of a Convolutional Neural Network (CNN), where a 3 by 3 filter with a stride of 2 is used to convolve over the inputs to generate the feature maps. The example input shown is a 12 by 12 geologic map, with the colours representing the lithology. The 3 by 3 filter creates a 5 by 5 feature map that condenses that geology, in order for it to be combined with the feature maps produced for the other inputs. The CNN uses the feature maps to make a pixel-by-pixel classification based on whether the combined inputs of a particular cell has a high, medium, or low hazard level, which is defined when the CNN is developed.

MLAs are not infallible, and critical thinking must be applied when choosing them out of the toolbox of possible approaches or combination of approaches that may be applied to a geomechanical problem (Elmo & Stead, 2020). The hesitancy to adopt data driven methods in rock engineering practice is due in part to the perceived “black box” nature of the algorithms, and as with any modelling endeavour, there are challenges associated with developing MLAs. Standardized methods, such as those presented in (Mcgaughey, 2019), should be applied to the MLA workflow to avoid common pitfalls. It is crucial to understand the relationships between data streams, uncover how the data inputs are used to train the algorithms, examine the physical implications of the algorithm architecture, and to verify that the outputs are aligned with the knowledge and experience of the collective field of rock mechanics. MLAs are not an automated problem-solving machine – they require user insights just like any other model development. Developing MLA workflows and best practices is an intensive field of research in computational science, ranging in applications from health sciences, to computer vision, to transportation engineering. The rock engineering community has the advantage of reaping the benefits of fields that have already grappled with analogous problems, such as: handling noisy datasets with CNNs to extract insights into the physical system (e.g., labelling aerial imagery from maps (Mnih & Hinton, 2012), or using bootstrapping to label noise (Reed et al., 2015)), maximizing the benefit of small datasets (e.g., detecting diabetic retinopathy using small datasets (Samanta et al., 2020)), and deciding which of the available inputs to include in an MLA to minimize complexity and maximize performance (e.g., identifying salient inputs in an ANN using a variety of approaches (May et al., 2011)). With respect to the last point, Input Variable Selection (IVS) methods are one option for looking inside the algorithm and justifying the use of a subset of available inputs when developing MLAs for rock engineering problems (May et al., 2011). In this chapter, an analysis of the activation strengths of the inputs is presented as one possible IVS method. Methods such as this demonstrate how the MLA is a mathematical representation of the physical system being modelled, demystifying its internal workings.

4.5 Background – Cigar Lake Mine

Cigar Lake Mine is owned and operated by Cameco, and is located in northern Saskatchewan, Canada (Figure 4-3). It is the world’s second-largest uranium mine, with an ore grade approximately one hundred times the global average (Bishop et al., 2016). The ore body is unique due to its size, high grade, intensity of alteration, and a high degree of associated hydrothermal clay alteration (Bishop et al., 2016).

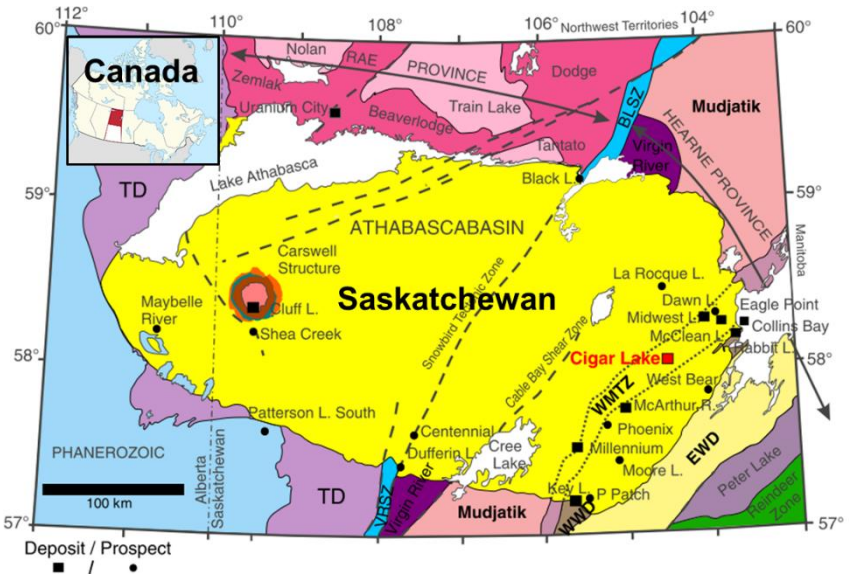

Figure 4-3. The location of Cigar Lake Mine in Saskatchewan, Canada showing regional geology, faults, shear zones, and other geological domains (adapted from Martz et al. (2017)), with the following abbreviations: TD Talston Domain, WMTZ Wollaston-Mudjatik Transition Zone, WWD West Wollaston Domain, EWD East Wollaston Domain, VRSZ Virgin River shear zone, BLSZ Black Lake shear zone. For detailed geological reference, the reader is referred to Martz et al. (2017) and Jefferson et al. (2007).

The geology of the Cigar Lake uranium deposit and environs has been described by (Bishop et al., 2016) and is summarized as follows. The uranium is found in an unconformity type deposit situated between the Athabasca Group (Athabasca Basin in Figure 4-3) and the underlying metasedimentary Proterozoic Wollaston Domain (WWD and EWD), in the Wollaston-Mudjatik Transition Zone (WMTZ in Figure 4-3). The deposit and host rock consist of three geological elements that also define the geotechnical domains: the deposit and associated hydrothermally altered clay cap, the overlying sandstone unit (Athabasca Group), and the underlying metamorphic basement rock (Wollaston Domain). There are three distinct styles of mineralization within the ore body: high grade mineralization at the unconformity, fracture-controlled mineralization in the sandstone, and fracture-controlled mineralization in the basement rock. The clay alteration of the ore body, a defining criterion of the geotechnical domains, is closely associated with major fault zones located within the deposit area. The deposit and sandstone are highly fractured and water bearing, while the basement rock is impervious. The basement rock is composed mainly of pelitic metasedimentary gneisses belonging to the Wollaston Domain. The Wollaston Domain is considered to be the most favourable unit for uranium mineralization. In general, the Cigar Lake Mine operation and production tunnels are located in three main rock mass types (Paudel et al., 2012): weak, highly weathered and saturated basement rock containing sand and clay; moderately weathered saturated basement rock; and strong unweathered basement rock.

The ore body is located above the 5.0 m lined diameter ore extraction tunnels, which are excavated using drill and blast. There are two main challenges facing the stability of the Cigar Lake Mine excavations:

controlling groundwater inflow, and supporting areas of weak rock (Bishop et al., 2016). The mine operators decided to freeze the rock mass surrounding the orebody to improve rock mass properties and to restrict groundwater inflow into excavated areas. A customized non-entry extraction method is used to extract the ore, where cutting cavities of frozen ore are created upwards from the ore extraction tunnels below with a high-pressure water jet, and then the ore is mixed with water to make a pumpable slurry (WSP (formerly Parson Brinkerhoff Quade & Douglas Inc.), 1999). However, the ground freezing operation results in complex time-dependent rock mass behaviour that is difficult to predict and presents challenges when designing support (Golder Associates, 2001; Roworth, 2013). A customized non-entry extraction method called the Jet Boring System (JBS) is used to extract the ore from the deposit. JBS consists of cutting cavities of frozen ore with a high-pressure water jet and then mixing the ore with water to make a pumpable slurry (WSP (formerly Parson Brinkerhoff Quade & Douglas Inc.), 1999). The cavities created from the extraction of the ore are backfilled with concrete to achieve early strength in frozen ground and therefore, provide additional ground support (Bishop et al., 2016).

Cigar Lake Mine was found to be an appropriate case study for a machine learning application because there are several phenomena that coalesce into a complex tunnel deformation, which is difficult to predict using existing techniques to anticipate squeezing ground (Barla, 2002; Barla et al., 2011; Barla & Borgna, 1999; K. Zhao et al., 2015). The combination of unfavourable geology (i.e., the weak and altered sandstone and basement rock) and the locally variable structural geology produces differential radial squeezing over the length of the ore extraction tunnels. The ground freezing regime, which is effective at stabilizing the rock mass for ore extraction, adds to the complex rock deformation mechanics by increasing the background stress during its implementation. These factors combine to result in tunnel liner and support yield that is difficult to forecast, which causes delays in production and adds complications for support rehabilitation scheduling and budgeting. Thus, accurately predicting tunnel liner and support yield is critical for the safe operation of the tunnel and for forecasting rehabilitation works needed to maintain operations. The CNN proposed herein is one approach to do this.

Previous experience developing a MLP ANN using this dataset (Morgenroth, Perras, Khan, et al., 2020) resulted in an average match of 40% between the predicted and observed tunnel liner yield, however poor performance was achieved for Class 2 and 3 yield (recalls of 18% and 0%, respectively). It was observed that the spatial dependencies in the dataset were cumbersome to include and were not preserved in the MLP ANN, and this led to the decision to use a CNN to analyze the Cigar Lake Mine dataset. To reduce the amount of data preprocessing and to limit the amount of bias introduced, image recognition and segmentation algorithms were explored, and the CNN was deemed to be a viable algorithm (Morgenroth et al., 2020). This resulted in a relatively simple digitizing and formatting process to transform the GMPs into images that could be processed by the CNN. This novel method for processing tunnel data and forecasting tunnel liner yield is advantageous because the complex behaviour is not constrained by a constitutive framework, and the data does not require interpretation or preprocessing that may inject bias.

4.6 Application of a CNN to the Cigar Lake Mine

4.6.1 Cigar Lake Mine Data

As part of their standard practice, Cigar Lake Mine produces as-built and production Ground Management Plans (GMPs) over the life of an ore extraction tunnel. The as-built GMPs summarize each tunnel's specific geotechnical domains as well as the installed support and ground freezing network layout, while the production GMPs track liner and support yield, ground freezing activity and tunnel rehabilitation efforts over time. In addition, tunnel displacement is monitored using survey points around the circumference of the tunnel. These are surveyed on an approximately weekly basis to determine the displacement of the extraction tunnel walls. The data from the GMPs as well as displacement measurements are used as the input dataset for the CNN developed herein. The GMPs and their corresponding dates are as follows:

- GMP 1 = Week 42 2015, 732 days elapsed since first displacement measurement
- GMP 2 = Week 50 2015, 56 days elapsed since GMP 1
- GMP 3 = Week 02 2016, 35 days elapsed since GMP 2
- GMP 4 = Week 10 2016, 56 days elapsed since GMP 3
- GMP 5 = Week 24 2016, 98 days elapsed since GMP 4

The dataset is comprised of a combination of spatial and temporal data that is both categorical (e.g., the support class, which is categorized as Class 1, 2, 3, etc.) and numerical (e.g., measured radial displacement within the tunnel). This study uses data from Cigar Lake Mine crosscut tunnel 765 to calibrate and test the CNN to predict tunnel liner yield forward in time. The dataset for tunnel 765 was digitized from the GMPs at a spatial resolution of 164 m, the length of the tunnel, by 13 elements, representing elements around the circumference of the tunnel, as illustrated in Figure 4-4. Three inputs were digitized from the GMPs for use in this case study: the mapped geology (GEO), as-built ground support class (SUPCL), and ground freezing patterns (FREEZE). The radial tunnel displacement (DISP) was obtained from the survey point measurements. The four inputs (GEO, SUPCL, FREEZE, DISP) were formatted into one image per GMP (Figure 4-5) and used to train the CNN to predict the tunnel liner yield (YIELD), which was also mapped on the GMPs. GEO and SUPCL are categorical inputs with an ordinal structure, and therefore have been encoded using integer values. FREEZE is a binary categorical input, while DISP is a numerical input. The YIELD categories represent the combination of damage to the liner elements (shotcrete, rock bolts, mesh) that would result in no rehabilitation time (Class 0), one day's rehabilitation work (Class 1), several days' rehabilitation work (Class 2), or total reprofiling of the tunnel (Class 3). A summary of each input, the output, and the corresponding digitized values, are presented in Table 4-1.

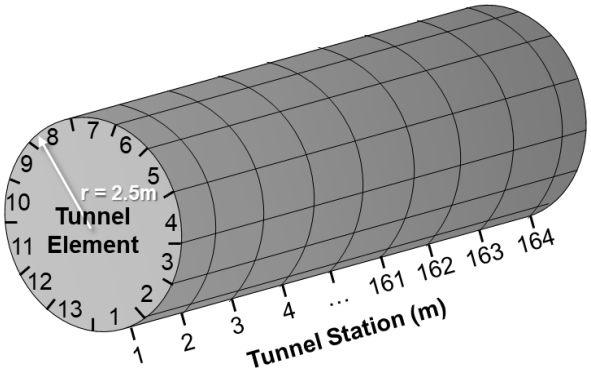

Figure 4-4. Nomenclature of digitized tunnel data, where r is the radius of the tunnel, showing the radial elements (1 to 13), along the 164 m length of the tunnel (Morgenroth et al., 2020).

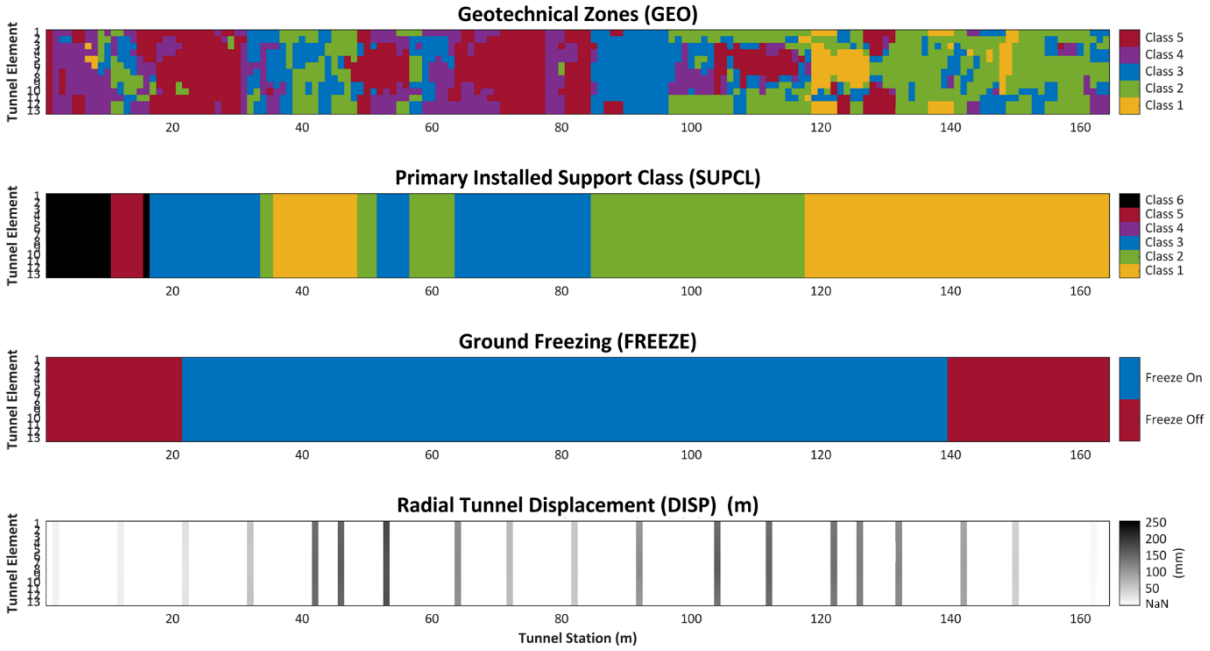

Figure 4-5. This image represents the digitized Ground Management Plan (GMP) 5. One digitized GMP image is formatted into four channels: geotechnical zones (GEO), primary installed support class (SUPCL), ground freezing (FREEZE), and radial tunnel displacement (DISP). This figure illustrates that the first three channels are integer-encoded categorical inputs. The fourth channel, DISP, is a numerical input and is only measured at particular locations along the tunnel. Where there is no DISP data, the pixel is assigned a value of not-a-number (NaN) and is not used by the Convolutional Neural Network during training.

Table 4-1. A summary of input variables (mapped geology, GEO; as-built ground support class, SUPCL; ground freezing patterns, FREEZE; and radial tunnel displacement, DISP) and output (tunnel liner yield, YIELD) used for the development of the CNN (Morgenroth et al., 2020).

Inputs							**Output**	
GEO (categorical)		**SUPCL (categorical)**		**FREEZE (categorical)**		**DISP (numerical)**	**YIELD (categorical)**	
Class	**Description**	**Class**	**Description**	**Class**	**Description**	**Description**	**Class**	**Description**
1	stiff crystalline rock	1	Class 1 support (lightest support)	0	freezing off	Measured radial displacement, ranging from 0 to 250 mm	0	no yield
2	fissured fresh crystalline rock	2	Class 2 support	1	freezing on		1	minor yield
3	transition from fresh to metamorphosed crystalline rock	3	Class 3 support				2	major yield
4	metamorphosed crystalline rock	4	Class 4 support				3	yield requiring total re-profiling
5	stiff to very stiff coarse-grained soil, structure well preserved	5	Class 5 support (heaviest support)					
6	fine-grained soil, structure not seen							

Three of the four inputs are objectively measured data: SUPCL represents the as-built support geometry, FREEZE represents the locations where thermistors indicate that the ground is frozen, and DISP is the measured radial tunnel displacement. GEO is slightly more subjective; however, the assigned geotechnical zones are standard across the mine site and are therefore based on site-specific experience. Of the four inputs, GEO and SUPCL are spatially variable (i.e., inputs have a resolution of 164 m by 13 elements), while FREEZE and DISP are temporally variable (i.e., inputs vary for each GMP). The time steps correspond to the five GMPs (Week 42 2015, Week 50 2015, Week 02 2016, Week 10 2016, and Week 24 2016). In tunnel 765, there are 19 rings where displacement measurements are taken with spacing between rings ranging from 4 to 12 m. The displacement survey targets are located around the circumference of the tunnel in Figure 4-4 at stations 0+002, 0+012, 0+022, 0+032, 0+042, 0+046, 0+053, 0+064, 0+072, 0+082, 0+092, 0+104, 0+112, 0+122, 0+126, 0+132, 0+142, 0+150, and 0+162. No interpolation between rings (i.e., along the length of the tunnel) was made in order to preserve the original form of the data available, and to prevent the introduction of bias. In cells where displacement measurements are not available, no value is given to the CNN. The displacement input is the maximum radial displacement at each survey point (7 survey points around the circumference of the tunnel) for the elapsed time between two consecutive GMPs.

The distribution of the parameters used to develop the Cigar Lake Mine CNN is shown in Figure 4-6 , including the four inputs (GEO, SUPCL, FREEZE, DISP) and the output (YIELD). Each column represents

a GMP, or snapshot in time of tunnel 765. The exceptions are GEO and SUPCL which do not vary with time. The third row, FREEZE, shows timing of the ground freezing front that was implemented in order to stabilize the ground and mitigate high groundwater inflows during the ore extraction process. The fourth row, DISP, shows a stacked bar graph where the magnitudes of displacement are differentiated by GEO class. As expected, the worse geologies experience higher displacement over time. It is also worth noting that Class 6 geology is not represented where the geology and displacement measurements intersect, i.e., the displacement measurement rings do not exist where Class 6 geology exists. The final row is YIELD, which shows that the tunnel liner support yield increases with time and progresses to high support classes as a consequence.

As with any modelling endeavour, the end goal is to reproduce as closely as possible the true rock mass behaviour observed. Because visualization of internal CNN processes can be difficult, an analysis of the input data from the GMPs and displacement measurements was completed prior to model development to identify both the dominant and subtle data trends. These data trends shown in Figure 4-6 are summarized as follows, as described by Morgenroth, Perras, Khan, et al. (2020). Over time the severity of the liner yield increases for all the geology classes, though the highest impact is to the areas with the poorest geology. There is some variation of support class assigned within each geology class that impacts the degree of liner yield sustained at that location. This nuanced relationship is difficult to quantify from visual data inspection. Liner yield where ground freezing is being implemented is higher than where the ground is not frozen, however, the liner yield increases everywhere with time indicating that other variables are implicit drivers of liner yield. Poorer quality geotechnical zones experience higher values of displacement and therefore, higher degrees of liner yield.

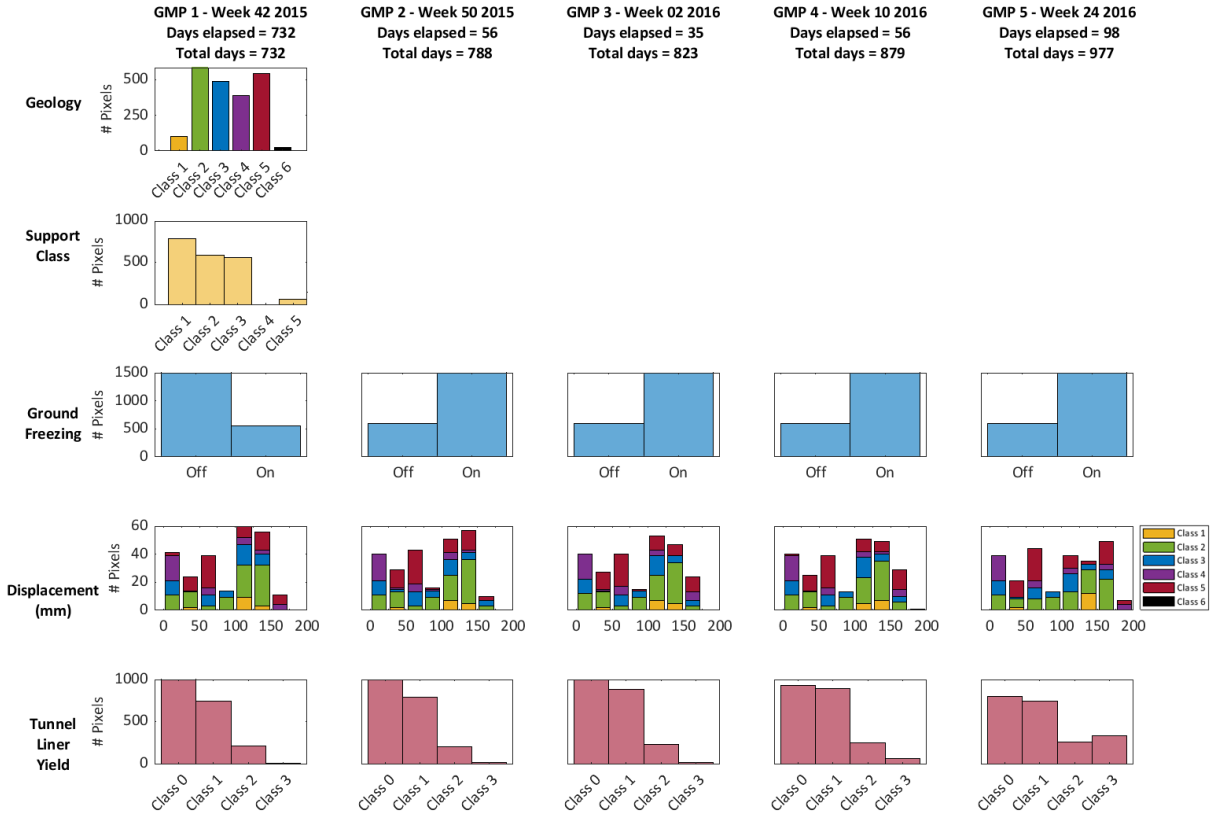

Figure 4-6. Histograms of Cigar Lake Mine dataset showing the distribution of each input for each of the five Ground Management Plans (GMPs), where each column of ground freezing, displacement, and liner yield histograms represents a GMP. The displacement histograms indicate the distribution of geology classes for each GMP. It is important to note that the rings where displacement is measured does not span areas where geology Class = 6, and as a result this class is not represented in the displacement bars.

A visualization of the tunnel liner yield with time is presented in Figure 4-7 , where each row represents a successive GMP for tunnel 765. Over time the areas of Class 3 yield (shown in red) increase, with the highest area of Class 3 yield occurring in GMP 5. It is important to note that the time elapsed between GMPs 4 and 5 is the largest, at almost double the time elapsed between any other consecutive pairs of GMPs. As shown in the figure, there is an imbalance of tunnel yield class in the GMPs: overall tunnel yield Class 3 is less represented across the GMPs, whereas Classes 0 to 3 are more uniformly represented. This demonstrates both the need for error weighting to account for this data imbalance problem, and the need for sensitivity analysis to determine how much of the historical (i.e., previous GMP) data is needed for accurate predictions. These analyses are detailed in Section 4.6.3.

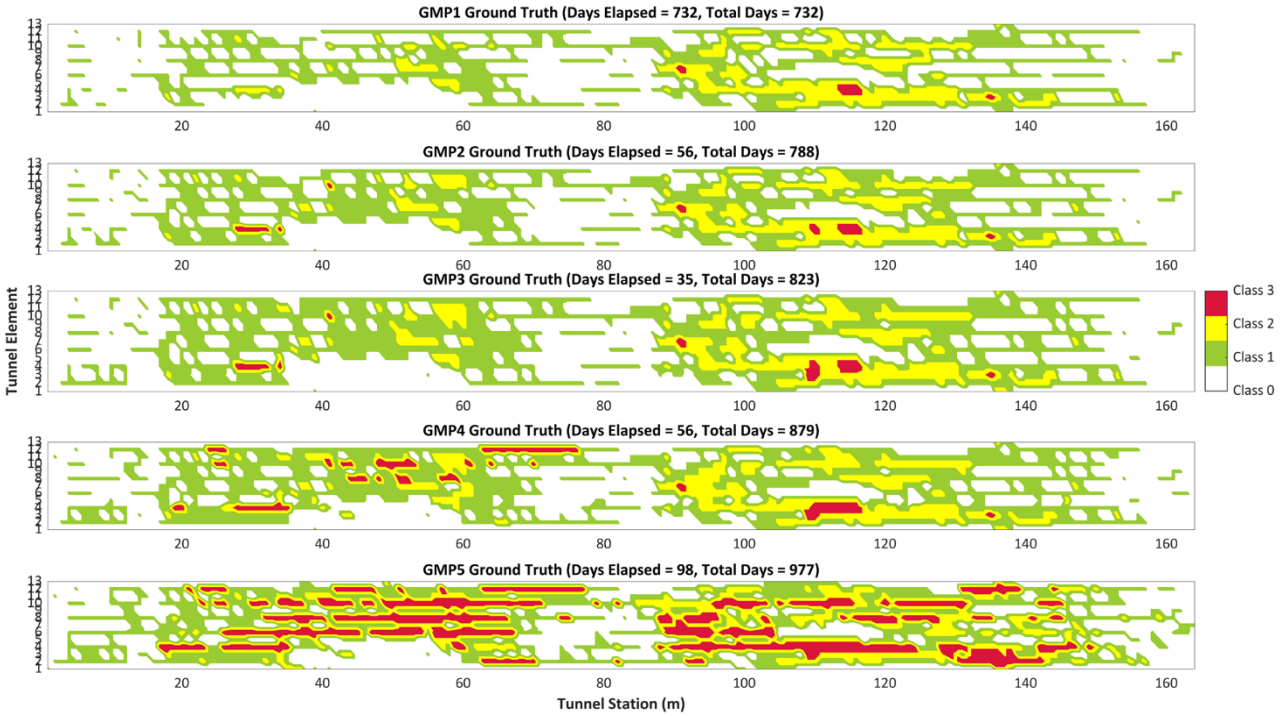

Figure 4-7. Tunnel maps of the tunnel liner yield (YIELD) for each successive GMP, showing the evolution of the YIELD with time.

The CNN was developed to predict the tunnel liner yield after the tunnel has been completed and the ore extraction process above the tunnel is underway, including the implementation of ground freezing. For this reason, data splitting was restricted to the temporal realm (i.e., subsequent GMPs) instead of spatially (i.e., along the length of the tunnel). The CNN needed to capture the data trends (i.e., the relative change in the inputs), ensuring that the time-dependent Cigar Lake Mine rock mass behaviour and subsequent tunnel liner yield are characterized properly for use in forecasting future yield. Forecasting the yield of the tunnel liner elements will allow Cameco to conduct more timely tunnel rehabilitation, minimize down time during repairs, and therefore increase confidence in costing and scheduling.

4.6.2 CNN Development

The CNN presented is designed to predict tunnel liner yield at the Cigar Lake Mine ore extraction tunnels. The CNN was developed in MATLAB R2019b (MathWorks Inc., 2019) using the Deep Learning and Computer Vision toolboxes. The goal of this model is to predict Class 2 and 3 tunnel liner yield as accurately as possible, as these are the levels of yield that require the most significant schedule and budget investment to mitigate. However, the overall performance of the model for other classes must not suffer as a result of the emphasis on Class 2 and 3.

The four inputs (GEO, SUPCL, FREEZE, DISP) were formatted into images for use in the CNN. Each image, made up of four channels that represent the four inputs, represents one GMP or point in time. As a result of this data formatting, there were 5 GMP images available for training and testing the CNN. Data

partitioning was applied by training on a subset of GMP images and withholding the rest. In order to test the CNN's ability to predict tunnel liner yield forward in time it was trained on the GMPs preceding the GMP for which the prediction was made, e.g., if the CNN is predicting yield for GMP 4, it would be trained on GMPs 1, 2 and 3. In order to test the sensitivity of the CNN's performance as a function of the amount of data it was trained on, an analysis was completed wherein the CNN was given a variety of preceding GMPs for training and the corresponding performance was compared. The ten permutations of training and testing datasets are shown in Table 4-2. These permutations simulate an operational scenario where the mine operators would use the available GMPs to date to forecast future yield.

Table 4-2. Summary of training and testing dataset permutations for the Cigar Lake Mine CNN.

Permutation	Training Dataset	Testing Dataset
1	GMP 1	GMP 2
2	GMP 2	GMP 3
3	GMPs 1 and 2	
4	GMP 3	GMP 4
5	GMPs 2 and 3	
6	GMPs 1, 2 and 3	
7	GMP 4	GMP 5
8	GMPs 3 and 4	
9	GMPs 2, 3 and 4	
10	GMPs 1, 2, 3 and 4	

The architecture of the CNN developed for the ore extraction tunnels at Cigar Lake Mine is shown in Figure 4-8, and consists of six layers. The first layer is the image input layer. Each image is sized 164 by 13 by 4 pixels corresponding to the tunnel map, which is 164 m long by 13 elements around the circumference for each of the 4 inputs for each GMP. Next is the 3D convolution layer, which uses a filter of varying sizes to convolve through each image with a stride of 1, and a padding value that is automatically generated to produce a feature map that is the same size as the input, which ensures all the input data is being used. This layer is followed by a rectified linear unit (ReLu) activation function, which performs a threshold operation on each pixel of the input and sets any value less than zero to zero. Next the transposed 3D convolution layer up-samples (or *deconvolves*) the feature maps using a cropping value (cropping is the opposite of padding) that is automatically generated to produce an output that is the same size as the input. This is given to the SoftMax activation function (a tan-sigmoid function), which normalizes the value of the input such that it sums to one, effectively producing a probability distribution that can be used for classification. Both ReLu and SoftMax activation functions are commonly used in CNN development (Bell, 2015). Finally, these probabilities are used by the pixel classification layer to assign a user-defined class for each pixel. It is here that the CNN makes the classification whether the predicted pixel is Class 0, Class 1, Class 2 or Class 3.

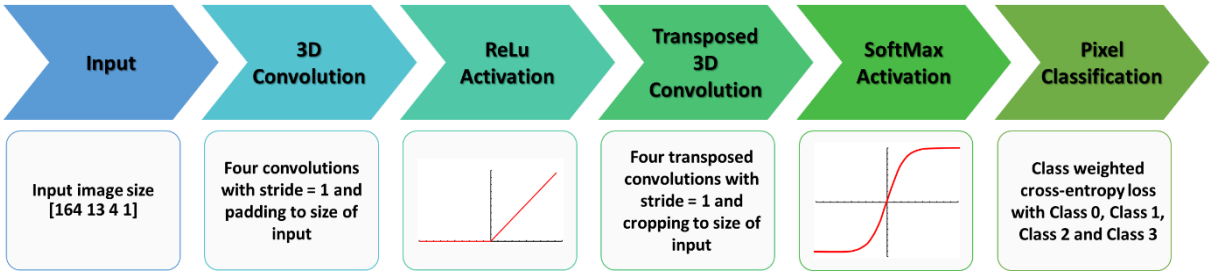

Figure 4-8. A graphic representation of the Cigar Lake Mine CNN, showing the six layers and their respective properties.

The size of the filter the convolution layer uses to convolve through the input images has a spatial and temporal significance. For a filter sized [*x y c*] the *x* and *y* represent the two-dimensional surface area of the filter, or the area around the pixel from which the CNN is drawing data to make a prediction on the rock mass behaviour at that pixel. The third dimension of the filter size, *c*, corresponds to the number of channels the filter should convolve through and therefore, the number of feature maps the convolution layer should create. In this study the input images have four channels (GEO, SUPCL, FREEZE, DISP) and the filter convolved through all of these to learn their internal relationships, therefore *c* was fixed at 4.

An important consideration for the Cigar Lake Mine CNN is the imbalance in the representation of the classes the CNN is predicting. The class representing the highest magnitude of tunnel liner yield, Class 3, is underrepresented by an order of magnitude as compared to the other classes. Specifically, the ratio of Class 0 to Class 3 representation is 1:90 in the training permutation where GMPs 1 through 4 are used to predict GMP 5. Other geomechanical datasets are liable to have similar problems, since there is generally the least amount of data for the most severe rock mass behaviour. For example, a dataset being used to predict rock bursts will have far more smaller magnitude rock bursts as compared to catastrophic rock bursts (Afraei et al., 2019). To address this, *error weights* may be assigned to the classes during the training phase in order to force the CNN to focus its prediction ability on whichever prediction the user deems to be most important. It does this by placing an emphasis on the errors that are backpropagated into the network of certain classes over others during training. In this study, four error weighting schemes were compared: uniform, linear, sigmoid, and inverse frequency (illustrated in Figure 4-9). The uniform weighting scheme gives equal preference to all the classes. The linear weighting scheme gives each successive class a higher weight on a linear scale. In this case, the sigmoid weighting scheme gives a much higher preference to Class 2 and 3, which correspond to the highest liner yield, where the weights are calculated using the sigmoid function shown in Equation 4-1.

$$Sigmoid\ Class\ Weight = \frac{1}{1 + e^{-(x-50)}} \qquad \text{Equation 4-1}$$

The inverse frequency scheme calculates the frequency of pixels of each class in the training images and assigns a preference inverse to this frequency.

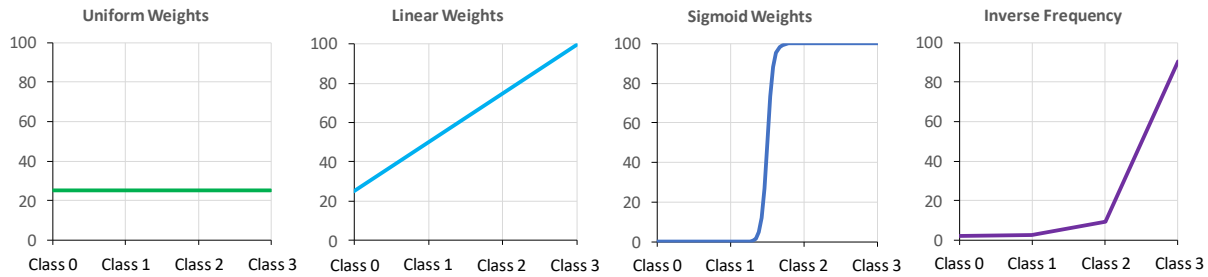

Figure 4-9. Four different error weight schemes applied to the classes in the training phase of the Cigar Lake Mine CNN.

Each combination of hyperparameters was run 20 times, using a different randomized set of initial weights and biases, to create an ensemble of models, allowing for a determination of confidence in the model predictions and a comparison between model ensembles. The accuracy of a model is determined by how often the predicted class matches the ground truth class. The performance of each ensemble is discussed in detail in Section 4.6.4.

4.6.3 CNN Sensitivity Analyses

To optimize the architecture for the Cigar Lake Mine CNN, a sensitivity analysis of three key hyperparameters is presented: the amount of training data, the filter size, and the error weighting scheme. To generate confidence intervals for each combination of model parameters, each permutation is run 20 times (the ensemble) for a total of 7200 CNN models. Recall that the best model in this context is defined as the model that has the highest recall of Class 2 and 3 tunnel liner yield, as these are the levels of yield that require the most significant schedule and budget investment to mitigate.

The parameters explored while designing the Cigar Lake Mine CNN were: the amount of data used for training, the size of the convolution filter, and the error weighting scheme. The filter size was varied from 10 by 10 by 4 through 50 by 50 by 4. Four different error weighting schemes are evaluated to determine which produced the best prediction accuracy on the test dataset – uniform, linear, sigmoid, and inverse frequency. Finally, the permutations of GMPs preceding the test dataset needed to make the most accurate predictions across all classes was analyzed (Table 4-2). Figure 4-10 shows a summary of the sensitivity analysis of the error weighting scheme, and Figure 4-11 shows a summary of the sensitivity analysis for filter size equal to 30 by 30 by 4. Figure 4-16 to Figure 4-19 in the Appendix show the detailed results of these analyses for all hyperparameters.

Figure 4-10 shows a representative comparison of the trends across error weighting schemes, where the CNN is trained on GMPs 1, 2, 3 and 4 and tested on GMP 5 for all filter sizes. Each box plot represents results from the ensemble with the same CNN parameters, where the box shows the 25th to 75th percentiles and the whiskers show the full range of performances for that suite of models. Each subplot shows the performance range for Class 0 (red), Class 1 (green), Class 2 (blue) and Class 3 (purple) predictions as well as the model's global performance (black). The global performance can be used as an indicator of the

overall model performance, and when compared across the four figures gives an indication of the impact of altering the CNN's weighting scheme. In general, the variance in model performances (indicated by the width of the box plots) decreases with increased training data (more GMPs used for training). It is also important to note that while increased data results in a narrower performance variance, convergence of the CNN to a global minimum is more difficult numerically, and therefore more time consuming. With respect to filter size, the performance for all combinations of error weighting schemes and data permutations plateau at a dimension of 30 by 30 by 4. Figure 4-11 shows all ten permutations of training and testing data and all error weighting schemes for a filter size of 30 by 30 by 4.

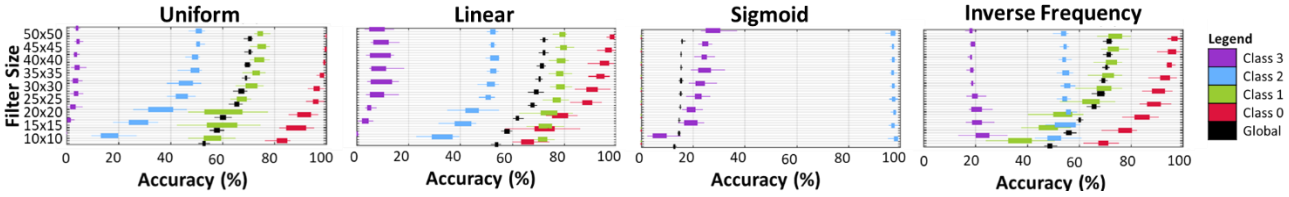

Figure 4-10. Summary of Cigar Lake Mine CNN hyperparameter sensitivity analyses, showing all error weighting schemes for models trained on GMPS 1, 2, 3, and 4 to predict GMP 5 for all filter sizes.

The uniform error weighting scheme assigns an equal weight to each YIELD class during the training process. Since all the classes are weighted the same the signals that are being learned by the CNN from each GMP is more clearly evident. For example, in Figure 4-11a GMP 1 contains almost no Class 3 damage for the CNN to learn, and so the performance is zero when using it to predict GMP 2. However, when using GMP 4 to predict GMP 5 (Figure 4-11d) the performance for Class 3 increases slightly because GMP 4 has more Class 3 samples for the CNN to learn. As shown in Figure 4-6, there is poor correlation between DISP and GEO Class 6, which is the worst geology, and which undergoes the highest deformation. This contributes to the overall poor performance of Class 2 and Class 3 YIELD seen in Figure 4-11 where uniform weighting is applied, i.e., no preferential weighting is applied. The uniform weighting scheme also demonstrates that GMP 5 (Figure 4-11d, g, i, j) is the most difficult to predict , because it contains the most extensive tunnel liner yield (Figure 4-7**)**, which is not present in prior GMPs. This further reinforces the need for a weighting scheme that prioritizes the higher YIELD classes, as this signal is underrepresented in the GMPs 1 through 4.

The linear error weighting scheme results show the change in model performance that can be achieved by applying an emphasis on higher YIELD classes. For these models, a simple linear scheme is applied, where each YIELD class is weighted 25% higher than the one preceding it. The variance of each box plot is lower than the uniform weighting scheme, indicating that the CNN is converging on a similar result with each run. In general, the YIELD Class 0 and 1 predictions drop slightly as compared to the uniform weighting, while there is an increase in Class 2 and 3 performance, so the global performance remains unchanged. While the linear error weighting scheme presents an improvement over the uniform scheme, it still does not place

sufficient emphasis on Class 2 and 3 YIELD to adequately capture these advanced stages of deformation (i.e., Class 2 and 3 still have lower performance than the other classes).

The sigmoid error weighting scheme is the most extreme scheme, where Class 0 and 1 are weighted close to 0% and Class 2 and 3 are weighted close to 100% (Equation 4-1). For all ten data splitting permutations, the performance for Class 2 predictions is close to 100%. Using the sigmoid scheme, Class 3 performance is increased as compared to the uniform and linear schemes. The improvement in Class 2 and 3 comes at the expense of Class 0 and 1 performance, which even the best of these models is only able to predict to a maximum of 20% accuracy. In general, the models trained on more data (i.e., each successive row) have a lower variance. Of note is that the difference between the performance for predicting GMP 4 (Figure 4-11c, f and h) and predicting GMP 5 (Figure 4-11d, g, i and j) is significantly lower as compared to the uniform and linear schemes. This is because the CNN prioritizes the performance of Class 2 and 3. Overall, the CNN with the sigmoid error weighting scheme can be applied to predict the worst of the tunnel liner YIELD.

The inverse frequency error weighting scheme results show the influence of weighting the CNN proportionally to the distribution of classes within the training samples. Class 3 YIELD is underrepresented in the input training data by an order of magnitude, and therefore shows poor prediction performance using a uniform weighting scheme. As with the previous schemes, the variance decreases as the amount of data increases (i.e., one GMP used in Figure 4-11d versus fours GMPs used in Figure 4-11j). The application of the inverse frequency scheme results in the highest global performance. While not as high as the sigmoid, Class 2 performance is higher than the uniform and linear schemes. In contrast to the sigmoid scheme, the inverse frequency scheme still allows the CNN to find a global optimum for all the Classes and not just for Class 2 and 3. For this reason, the inverse frequency weighting scheme results in the best global model performance of the four error weighting schemes investigated.

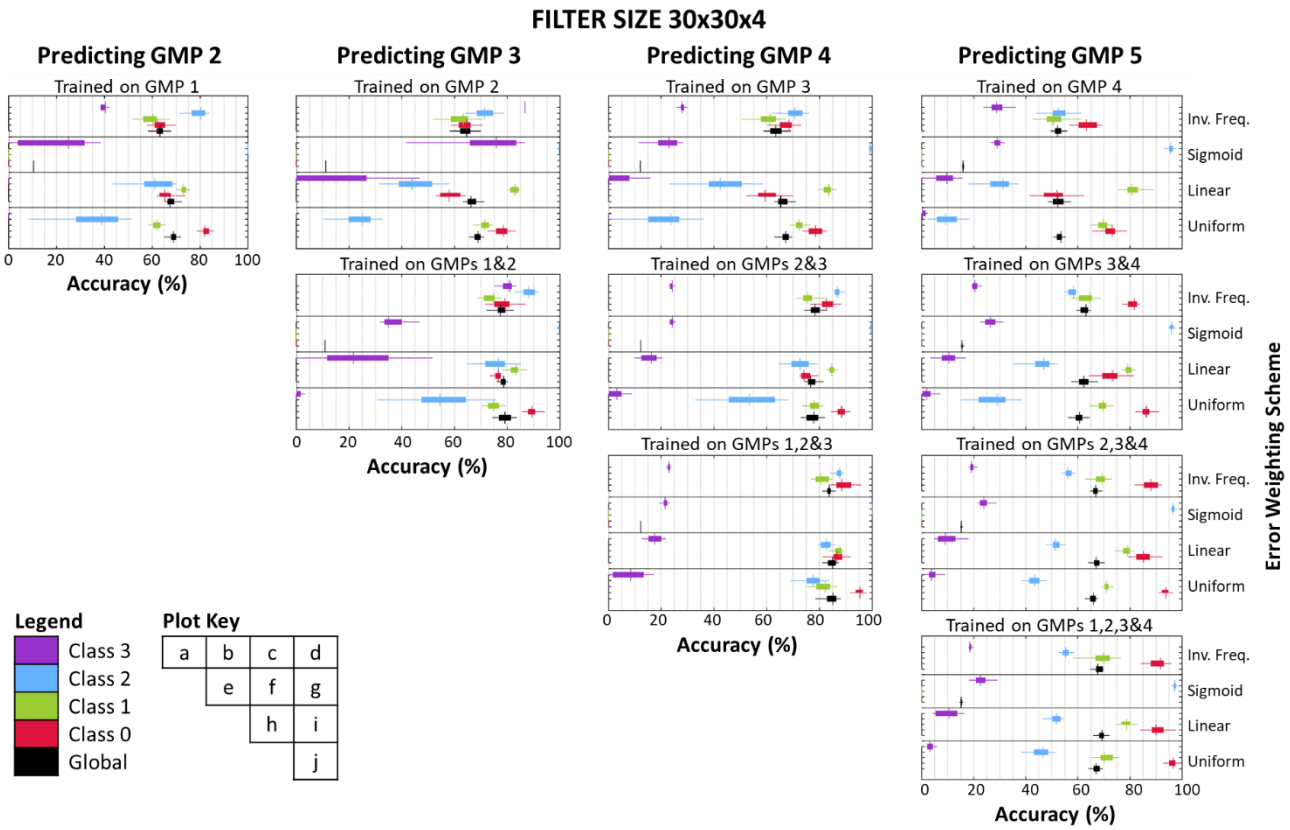

Figure 4-11. Cigar Lake Mine CNN sensitivity analysis results for filter size equal to 30 by 30 by 4, for different training data and error weighting schemes. Detailed results for all filter sizes are shown in the Appendix.

4.6.3.1 Input Activation Strengths

By extracting the activation strengths for each combination of model hyperparameters, it is possible to determine which of the four input parameters contain the strongest signal that is used by the CNN to make its classification, i.e., the most dominant input parameter(s). This analysis identifies which inputs are emphasized in the learning process. The results of this analysis are shown in Figure 4-12, where the most dominant input when predicting each GMP is shown as a percentage of the total models predicting the same GMP. This analysis is shown for the sigmoid and inverse frequency models, as these had the highest performance accuracy. SUPCL is the dominant input for this dataset, regardless of which GMP is being predicted or which error weighting scheme is applied. The DISP input shows slightly more dominance in the sigmoid scheme because these models emphasize YIELD Class 2 and 3 during training, which are more strongly correlated with DISP. In general, this analysis showed that GEO is never the strongest signal used by the CNN to predict YIELD. This may seem counter intuitive, since it is widely known in the rock mechanics community that geology plays a crucial role in rock mass deformation. However, in machine learning parlance this indicates that there is another input that is conveying the same information (either implicitly or explicitly) to the network or has a stronger correlation than GEO. In this case it is likely that the

SUPCL input is providing this signal, as the primary installed support is installed based on the observed geology during tunnel excavation and mine operation experience.

The activation strength analysis is a useful demonstration that some inputs that are known to be important from a physical perspective are not as important from a machine learning perspective due to redundancies or correlation in similar inputs. However, these results do not conclusively demonstrate that any of the inputs should be removed to reduce model complexity, as improved performance by removing them has not been demonstrated by this IVS method.

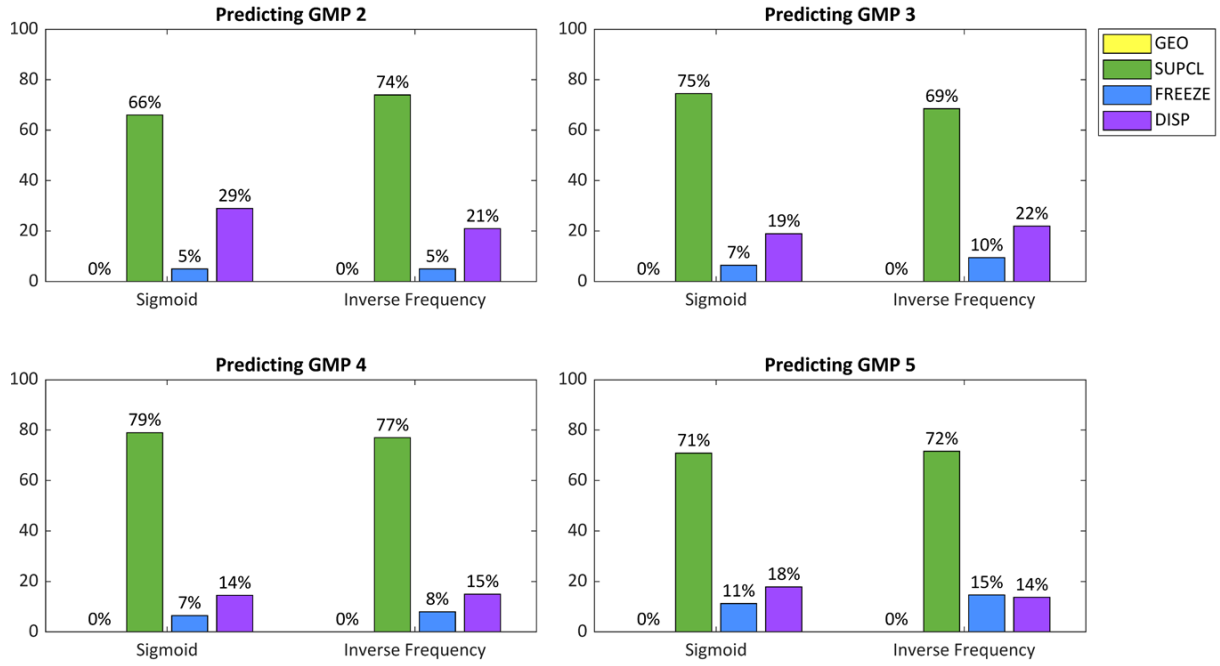

Figure 4-12. Input dominance for sigmoid and inverse frequency error weighting schemes, for each GMP that is being predicted. Dominance is shown as a percentage of the models predicting the same GMP.

4.6.4 CNN Performance

As a result of the sensitivity analyses, two CNN architectures were selected to predict the deformation phenomena at Cigar Lake Mine, depending on which Yield Class prediction is prioritized: the “Balanced Global” model and the “Targeted Class 2/3” model. The optimal model parameters are as shown in Table 4-3. Figure 4-13 and Figure 4-14 show visualizations of the CNN with both sigmoid and inverse frequency weighting schemes when using GMPs 1-3 to predict GMP 4, with smoothing applied for improved interpretability. The model accuracy is shown in Figure 4-13, where the white pixels have been identified correctly in all 20 models, the black pixels have been identified incorrectly in all 20 models, and the greyscale shows the percentage of the 20 models that were identified correctly. For each ensemble of 20 models, the pixel accuracy in the figure is defined as shown in Equation 4-2. The pixel score, indicating the number of correctly or conservatively identified pixels, is shown in Figure 4-14. White pixels represent where the models predict the YIELD class correctly or overpredicts it (i.e., the prediction is conservative), the black

pixels represent where the models under predict the YIELD class, and the greyscale shows the proportion of the 20 models that make a correct or conservative prediction for a given pixel. The score for each pixel in the figure is defined as shown in Equation 4-3.

Table 4-3. Summary of Cigar Lake Mine CNN architectures for the two models (the "Balanced Global model and the Targeted Class 2/3 model), showing the optimal error weighting and filter size hyperparameters, as well as the best GMPs to use for training to predict each GMP. GMP 2 is omitted because it can only be trained using GMP 1.

CNN Model Name (Model Goal)	Error Weighting Scheme	Filter Size	Training/Testing Permutation	
			To predict…	…train using…
"Balanced Global" (Best Overall Accuracy)	Inverse Frequency	Greater than 30 x 30 x 4	GMP 3	GMP 2*
			GMP 4	GMPs 1 to 3
			GMP 5	GMPs 1 to 4
"Targeted Class 2/3" (Best Class 2 and 3 Accuracy)	Sigmoid	Greater than 30 x 30 x 4	GMP 3	GMP 2*
			GMP 4	GMPs 1 to 3
			GMP 5	GMPs 1 to 4

*GMP 3 is best predicted without GMP 1 because there is a minimal difference between GMPs 1 and 2, and therefore no new information is provided for the network to learn.

$$Pixel\ Accuracy = \frac{\sum_{model=1}^{20} True\ Positives_{model}}{20} \quad \text{Equation 4-2}$$

$$Pixel\ Score = \sum_{model=1}^{20} True\ Positives_{model} + \sum_{model=1}^{20} Over\ Predictions_{model} \quad \text{Equation 4-3}$$

Combined, Figure 4-13 and Figure 4-14 illustrate that the sigmoid scheme is better at predicting Class 2 and 3 while over predicting the other classes, as expected due to the error weighting scheme. However, if a global prediction accuracy is desired then the inverse frequency scheme is preferable based on model performance. For example, in Figure 4-13 the sigmoid scheme misclassifies Class 0 and 1 between 0+060 and 0+085 (i.e., pixels are black), while the inverse frequency scheme generally classifies these pixels correctly (i.e., pixels are white). Upon examining this same chainage in Figure 4-14 for the sigmoid scheme, it is noted that the misclassified pixels are generally over predicted (i.e., these pixels are now white), indicating that the model is making a conservative forecast. However, the Class 3 pixels at tunnel element 13 between chainages 0+060 and 0+075 are not predicted correctly by either the sigmoid or the inverse frequency scheme, as indicated by the black pixels. Similarly, the Class 3 pixels at tunnel elements 6-9 between chainages 0+110 and 0+115 are not predicted correctly by either scheme.

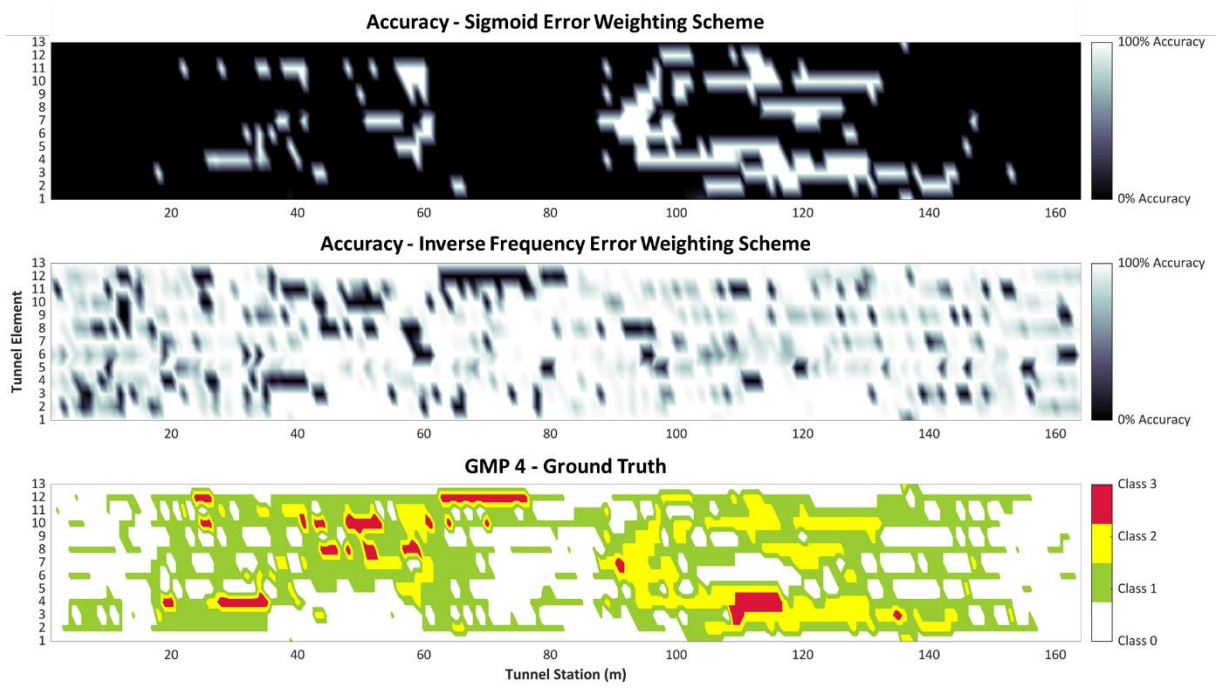

Figure 4-13. Pixel accuracy of Cigar Lake Mine CNN with sigmoid and inverse frequency schemes for predicting GMP 4 when the CNN is trained on GMPs 1 to 3, showing how often the model is correct across the 20 models.

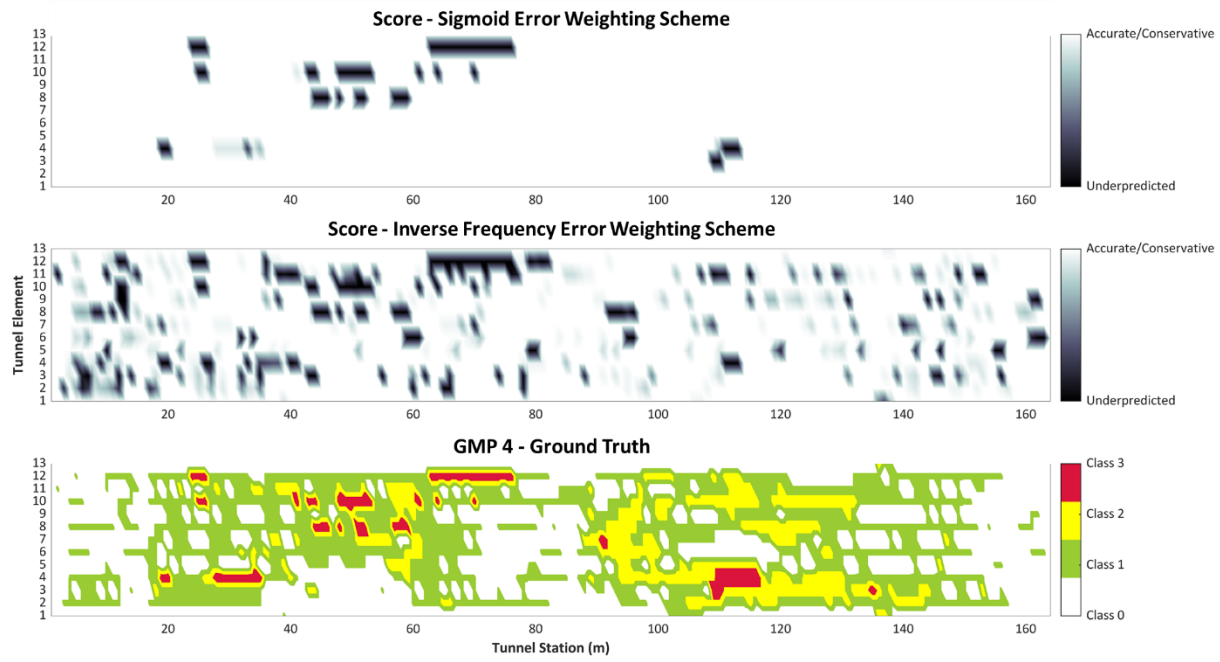

Figure 4-14. Pixel score of the Cigar Lake Mine CNN with sigmoid and inverse frequency schemes for predicting GMP 4 when the CNN is trained on GMPs 1 to 3, showing how often the model is correct or conservative across 20 models.

The performance of a classification model is commonly visualized using a confusion matrix (Figure 4-15). This type of visualization identifies where the system is confusing two classes, or mislabels one as another.

For example, the Balanced Global model (with the inverse frequency error weighting scheme) confuses actual Class 0 by predicting it as Class 1 a total of 25 times. The on diagonal values (3699, 3084, 877, 60) indicate how often the CNN makes correct classifications. The summary column on the right side of the confusion matrix shows how often the CNN correctly predicts a class as checked against the ground truth class. This is called the true positive rate or the recall. In this case, when Class 2 is predicted by the model, the ground truth (actual) is Class 2 (the prediction is correct) a total of 89.1% of the time. This metric gives an overall indication of how well the CNN is predicting the observed tunnel liner yield. For the Balanced Global model, Class 3 is predicted correctly with a recall of 24.6%. The summary row at the bottom of the confusion matrix indicates how often the CNN predicted a particular class correctly as a ratio of the total of the amount of times that class was predicted, whether the prediction was correct compared to the observed value or not. This is called the positive prediction rate or the precision. For example, this model has a Class 2 precision of 96.1% of the time out of the total number of times it predicted Class 2, i.e. it is correct 96.1% of the time, as compared to its own predictions.

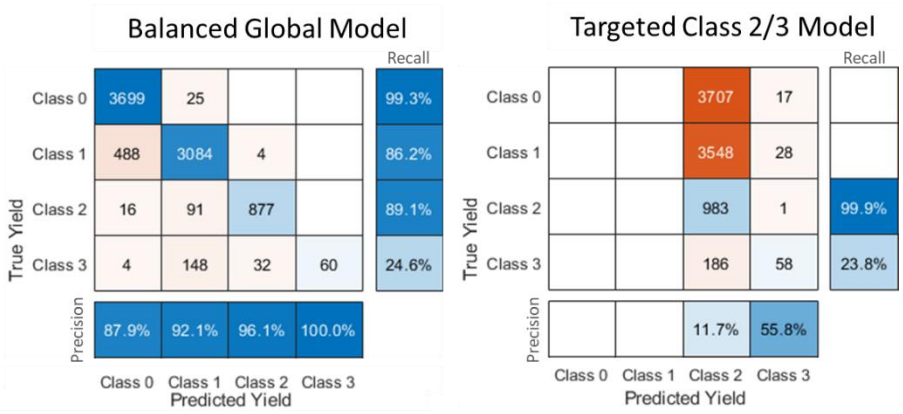

Figure 4-15. Confusion matrices for the Balanced Global model (left) and Targeted Class 2/3 model (right) for predicting GMP 4 when the CNN is trained on GMPs 1-3.

When comparing the confusion matrices for the Balanced Global and Targeted Class 2/3 (with the sigmoid error weighting scheme) models, it is evident that the Balanced Global model has higher classification precision across all the YIELD classes than the Targeted Class 2/3 model, as expected. It also has higher recall for Class 0 and 1, as it is not penalized for these during training. The Balanced Global model also has a slightly higher recall for Class 3 Yield as compared to the Targeted Class 2/3 model, which is attributed to the equal weighting of Class 2 and 3 received in the latter. In other words, convergence to a solution with a better total performance across Classes 2 and 3 was achieved by sacrificing a slight improvement to Class 3 recall from 24.6 to 23.8% for a 10% increase to Class 2 recall from 89.1 to 99.9%. The Targeted Class 2/3 model cannot predict Class 0 and 1 at all, which is to be expected since these classes were given a 0 weighting in the training process. The low precision of the Targeted Class 2/3 model reflects the inclusion of over-predicted pixels, which was deemed acceptable in order to attain a high recall for Class 2 pixels (i.e., minimal under-prediction). Class 2 pixels represent future Class 3s, and over-

predicting them may result in earlier rehabilitation intervention and therefore cost savings before total reprofiling is necessary (Class 3 occurs). Under-predicting Class 2 is not acceptable, as Class 3 represents a significant increase in cost for rehabilitation.

4.7 Discussion

It is imperative that the CNN be designed to appropriately address the problem it is intended to solve. A problem definition framework should be developed to ensure that the architecture of the CNN, in this case, is suited to making the predictions that were intended at the outset while maximizing the utility of the available inputs. In addition, the metric for sufficient performance should be defined. For the Cigar Lake Mine problem, the CNN was developed to predict the tunnel liner yield as a function of the geotechnical mapping and monitoring data available on the GMPs that were already produced. The measure of success was how well the model could predict all YIELD classes, with particular emphasis on Class 2 and Class 3 as these represent the highest capital investment for rehabilitation. High model performance for Class 2 in particular was targeted as these represent the future Class 3, and due to the difference in cost and scheduling associated with major repairs (Class 2) versus total reprofiling (Class 3).

The primary source of information for the Cigar Lake Mine are the GMPs that are produced at irregular intervals as tunnel maps, showing the geology, installed support, and progression of the ground freezing and tunnel liner yield. This study compared a combination of several hyperparameters for the CNN: ten permutations of data partitioning, convolution filter sizes ranging from 10 x 10 x 4 to 50 x 50 x 4, four error weighting schemes, and an ensemble of 20 models each, for a total of 7200 CNN models. In general, it was found that increased training data results in increased performance across all YIELD classes, as expected. The notable exception is when predicting GMP 5. This is attributed to the large gap in time between GMP 4 and 5, where a significant development in YIELD occurred throughout the tunnel. The CNN cannot predict a severity of YIELD that it has not seen before, and therefore some fine tuning of the hyperparameters is needed to improve Class 2 and 3 performance. For future work, GMPs could be produced at more regular time intervals in order to improve the CNN prediction accuracy and to capture the rock mass deformation mechanisms more continuously.

The error weighting scheme has a significant impact on the performance on the individual class performances. The uniform scheme was used as a baseline for comparison, as this represents the models where no emphasis on a particular class is introduced. The linear scheme represented a minor improvement in performance, and significant improvement in variance of each suite of models. The sigmoid scheme represents the best performance for Class 2, which is a good barometer for the support rehabilitation intervention that will be needed imminently. The inverse frequency scheme represents the best performance across all YIELD classes.

The Cigar Lake Mine deformation mechanisms and subsequent tunnel liner yield can be best described and predicted by two models – a Balanced Global model and a Targeted Class 2/3 model. The Balanced

Global model represents good prediction accuracy across all YIELD classes. This is achieved by applying an inverse frequency error weighting scheme during training to balance the samples of each class in the training dataset. This scheme results in the highest global prediction accuracy, and the error weights are adaptive to the distribution of samples within the training images. This means that the scheme updates itself for each new image that is added during training. The Targeted Class 2/3 model represents the best combined recall for Class 2 and 3, with a particularly high recall for Class 2. This is achieved using a sigmoid error weighting scheme that emphasizes those two classes during training over the others. The Targeted Class 2/3 model can be used when the mine operators are targeting high capital tunnel liner rehabilitation, and was developed to have as high a recall for Class 2 as possible. Over-predictions were deemed acceptable, while under-predictions were minimized since they are more detrimental to tunnel liner rehabilitation scheduling. From an operational standpoint, the Targeted Class 2/3 model can be used to plan and execute routine underground inspections. For example, the locations where Class 2 is predicted is added to a database where inspection by a ground control engineer is needed. If upon inspection a particular location's YIELD is less than Class 2 (i.e., over-predicted by the model), that location stays on the "watch list" but no rehabilitation action is taken. It would be far less desirable for the model to have under-predicted that location and it has already reached Class 3 YIELD, as that indicates that the optimal time for rehabilitation has already passed.

The dominance of each of the four model inputs (GEO, SUPCL, FREEZE, DISP) as predictors within the CNN were compared by analyzing their respective activation strengths as measured for each channel (i.e., the input) of the image. This analysis was focused on the models using the sigmoid and inverse frequency error weighting schemes, as these had shown the best performance during the sensitivity analysis of the hyperparameters. The analysis showed that support class was overwhelmingly the most dominant input parameter for the Cigar Lake Mine CNN, regardless of which GMP was being predicted and which error weighting scheme was employed. It is intuitive that this input has a strong correlation with the tunnel liner yield (the CNN output), as the suitability of the type and extent of primary rock support is highly relevant to the yield that it experiences as the rock mass deforms. This result does not demonstrate that any of the CNN inputs can be removed to reduce model complexity, as it only shows which inputs are the most dominant, and not which are not used by the CNN at all. Future work will explore how input dominance may be used as an IVS method.

In the context of tunnel design generally, and with respect to other rock mass deformation environments and mechanisms, there are some transferable learnings from the Cigar Lake Mine CNN and the mine's squeezing environment. The primary learning is that the spatial and temporal characteristics that are important to the mechanism at hand must be represented in the dataset used to develop the CNN. For example, the deformation mechanics at Cigar Lake Mine are known to depend on the nuanced relationship between the adverse geology and the ground freezing implemented during ore extraction. Thus, these data are explicitly given to the CNN as inputs. Similarly, in a gravity controlled failure environment, it may be

necessary to include explicit discontinuities as inputs into the CNN, rather than just the geotechnical zones, in order to be able to predict falling blocks at the excavation scale. In a bursting environment where stress accumulation is problematic, the CNN may need a stress related dataset (e.g., microseismic data) explicitly as this is the main driver of the mechanism that is observed locally at the excavation. This example raises a further degree of complexity in that stress is a 3D problem, and therefore a CNN that performs four-dimensional convolution (where the dimensions are x, y, z, and time) may need to be developed to capture the stress redistribution in three dimensions.

The amount of training data has also been found to be a sensitive CNN parameter, particularly in the temporal domain. For Cigar Lake Mine, it was generally true that as much displacement history as possible was needed for an accurate prediction, but this may not be the case for another rock mass deformation mechanism. For example, in a stress controlled environment only the seismic data collected just before the latest blast may be of consequence.

CNNs are powerful tools born in the field computer vision and segmentation, which present an exciting opportunity for processing geological and geomechanical data. The datasets that are traditionally collected in rock engineering practice do not require complex formatting for use, as their spatial nature is conducive for CNN processing. Geological mapping or geotechnical domains may be used as the baseline for the "images" that the CNN uses as an input, and instrumentation data can be overlaid and resampled to the resolution of the image pixels. This approach also allows for the combination of numerical and categorical information to be used within a single algorithm, which is novel at the intersection of rock engineering and machine learning (Morgenroth, Khan, et al., 2019). Also, this approach has the advantage of being intuitive to professionals that are not data scientists by training, because the data format is familiar and can be related to physical properties and mechanisms that are widely accepted. This study demonstrates that reasonable agreeance between CNN generated predictions and actual tunnel liner yield can be achieved, and that a CNN can be used to forecast future tunnel liner yield. This finding has powerful implications for mine operational planning, budgets and scheduling, and above all improved safety in underground excavations.

4.8 Conclusions

This chapter presents the novel application of a computer vision machine learning algorithm, the Convolutional Neural Network, to the rock engineering problem of predicting tunnel liner yield in the squeezing ground conditions of Cigar Lake Mine in Saskatchewan, Canada. Four conventional geological and geomechanical inputs were used in the CNN development: geology class, primary support class, ground freezing, and tunnel displacement. These data were obtained from the GMPs prepared by the mine and from tunnel liner displacement surveys. Previous work done to analyze these inputs revealed that there are nuanced and complex rock mass deformation mechanics at play, resulting from the ground freezing program interacting with the weak and faulted rock mass. These nuanced relationships are difficult to capture using a conventional numerical method, where the real constitutive behaviour may be overprinted

by one imposed by the numerical analysis, which provided the justification for exploring a machine learning algorithm.

The Cigar Lake Mine CNN was developed through a sensitivity analysis of three network hyperparameters: the amount of training data, the convolution filter size, and the error weighting scheme. A total of 7200 models were analyzed to determine the best combination of these parameters for this problem. It was found that an error weighting scheme was necessary to offset the imbalance of samples for each class in the dataset. A sigmoid error weighting scheme, which prioritizes Class 2 and 3 errors during training, was used to develop a Targeted Class 2/3 model that has higher recall for Class 2 and 3 tunnel liner yield. An inverse frequency error weighting scheme, where the proportion of samples in each class is used to determine the error weights, was used to develop a Balanced Global model that has reasonable prediction accuracy across all yield classes (average model accuracy >65% for all training data permutations).

Machine learning algorithms are another tool to add to the rock engineering toolbox, one that is filling the need for accessible tools to use when the rock mass deformation characteristics are complex and nuanced behaviour may be overlooked. The internal relationships between the parameters that are known to be important can be disentangled by a numerical system that is built on interconnectivity, a method which allows for more informed decisions when designing, excavating, and rehabilitating underground excavations. Pushing this frontier forward does not only present an exciting innovation of combining data science and rock engineering, but also an opportunity for more economic and safer underground excavations.

Appendix

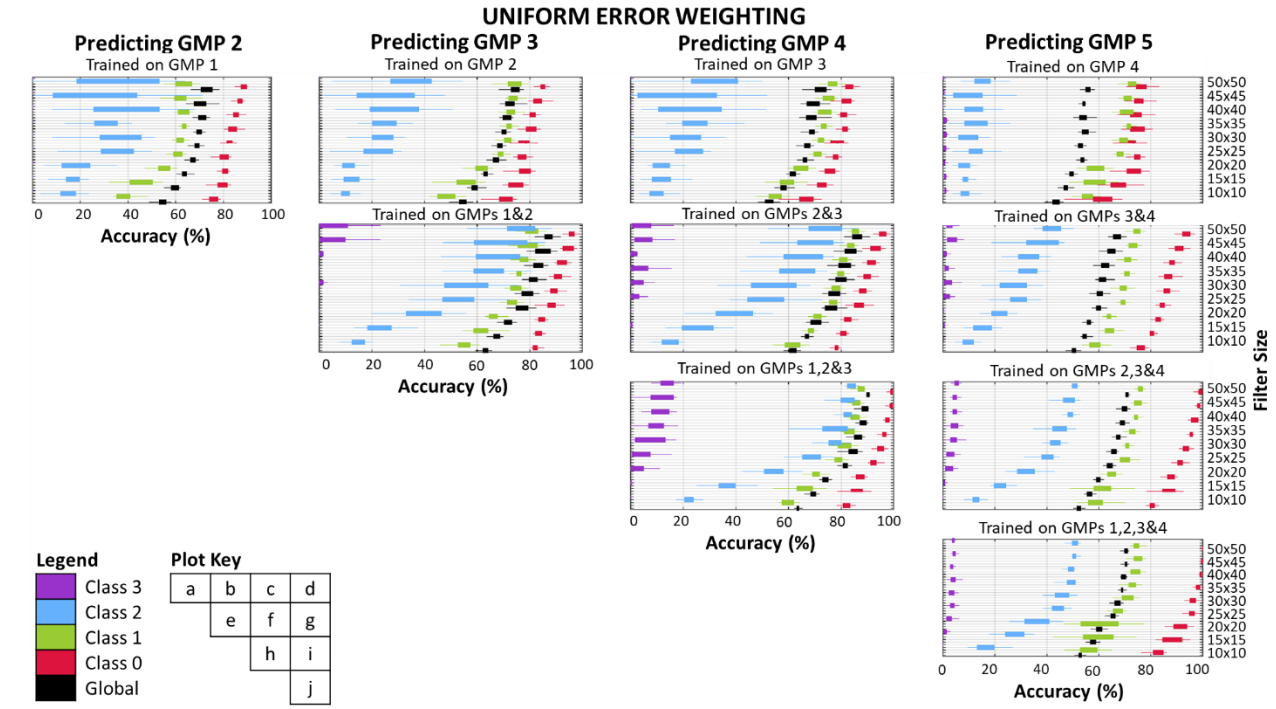

Figure 4-16. Cigar Lake Mine CNN results for the uniform error weighting scheme, for different training data and filter sizes.

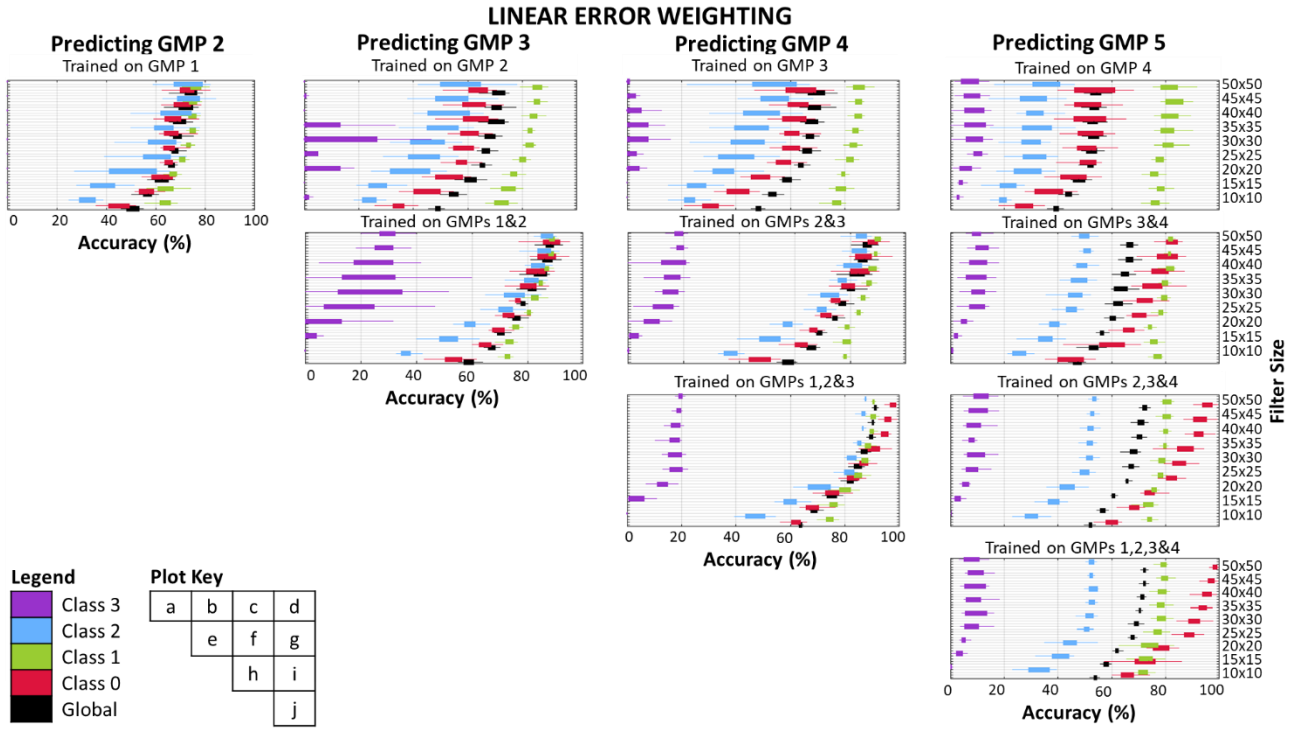

Figure 4-17. Cigar Lake Mine CNN results for linear error weighting scheme, for different training data and filter sizes.

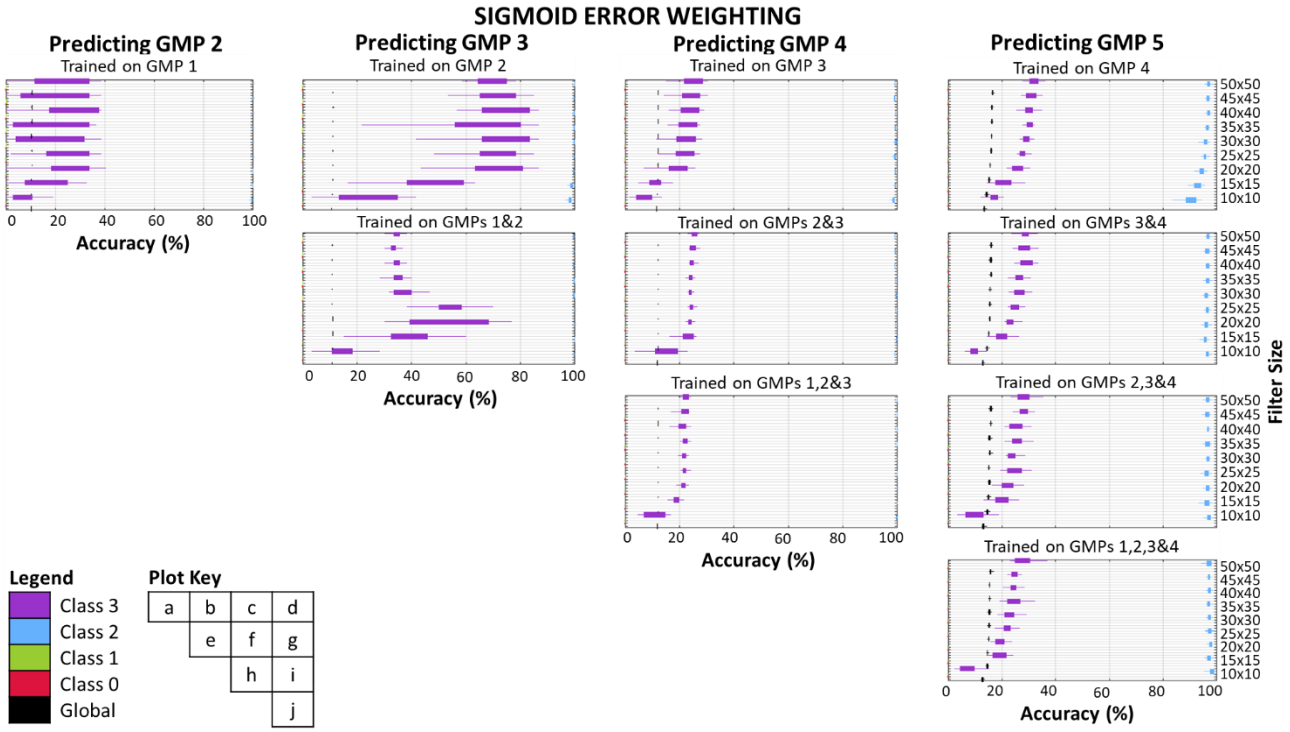

Figure 4-18. Cigar Lake Mine CNN results for sigmoid error weighting scheme, for different training data and filter sizes.

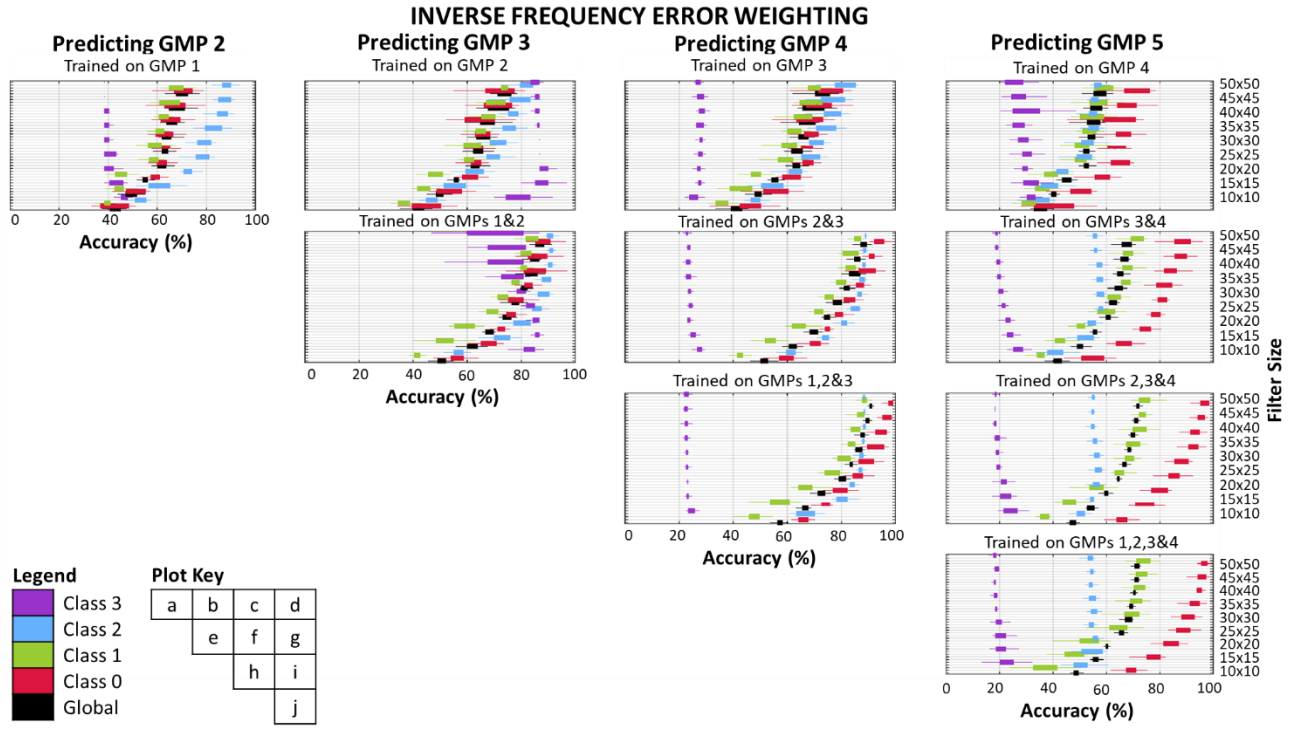

Figure 4-19. Cigar Lake Mine CNN results for inverse frequency error weighting scheme, for different training data and filter sizes.

CHAPTER 5. ON THE INTERPRETABILITY OF MACHINE LEARNING USING INPUT VARIABLE SELECTION: FORECASTING TUNNEL LINER YIELD

5.1 Preface

This chapter focuses on three Input Variable Selection (IVS) methods developed and applied to the Cigar Lake Mine CNN developed in *CHAPTER 4. A Convolutional Neural Network approach for predicting tunnel liner yield at Cigar Lake Mine*. IVS is an approach that examines how the CNN uses the given data, or inputs, to forecast rock mass behaviour. The three IVS methods investigated were Channel Activation Strength (CAS), Input Omission (IO), and Partial Correlation (PC). The IO and PC approaches proposed are novel for CNNs using a spatial and temporal geomechanical dataset. These IVS methods represent an approach to improve algorithm interpretability, both in terms of transparency and post-hoc interpretations, by providing insight into how the IVS methods are used and their impact on model performance.

The content of this chapter was submitted to *Rock Mechanics and Rock Engineering* in 2021 as follows:

Morgenroth, J., Perras, M. A., & Khan, U. T. (2022). On the Interpretability of Machine Learning Using Input Variable Selection: Forecasting Tunnel Liner Yield. *Rock Mech Rock Eng*. https://doi.org/10.1007/s00603-022-02987-5

The contributions of the authors in the current chapter are as follows:

Josephine Morgenroth has conducted the literature review, developed the IVS methods, applied the IVS methods using the required software, validated and visualized the results, and prepared and wrote the original manuscript of this publication. **Matthew A. Perras** has supervised the research, provided the funding, and contributed to writing and editing the manuscript. **Usman T. Khan** has supervised the research, provided the funding, and contributed to the writing and editing the manuscript.

The authors would like to extend special thanks to Cameco, and particularly Chris Twiggs, Imre Bartha and Kirk Lamont for their constructive feedback and informative conversations. This work is funded in part by the Natural Sciences and Engineering Research Council of Canada through the Discovery Grant program and the joint Innovation York and National Research Council Canada's Industry Research Assistance Program – Artificial Intelligence Industry Partnership Fund, in partnership with Yield Point Inc. This work is also funded by the NSERC Postgraduate Scholarships – Doctoral program.

5.2 Abstract

To validate the application of Machine Learning (ML) to rock engineering practice, it is crucial that algorithm developers use appropriate methods to quantify how closely the ML reproduces the observed rock mass deformation. Input Variable Selection (IVS) is one approach that examines how ML uses the given data, or inputs, to forecast rock mass behaviour. Three IVS methods were developed for two Convolutional Neural Network (CNN) architectures that predict tunnel liner yield at the Cigar Lake Mine, which exhibits time-

dependent squeezing deformation. One model architecture focused on accurately predicting the higher tunnel liner yield classes, while the second architecture prioritized prediction accuracy across all tunnel liner yield classes. The three IVS methods investigated herein were Channel Activation Strength (CAS), Input Omission (IO), and Partial Correlation (PC). The IO and PC approaches proposed are novel approaches proposed for CNNs using a spatial and temporal geomechanical dataset. Performance of all models was compared using the Corrected Akaike Information Criterion (AIC_c), where lower values indicate better performance.

Each IVS method was used to produce a unique ranking for each model architecture and training/testing data split: CAS produced an Activation Ranking, IO produced an Omission Ranking, and PC produced a Correlation Ranking. The Activation Rankings showed that the geotechnical zones input had the lowest activation strength in the CNN relative to the other inputs (ground freezing, primary installed support class, and radial tunnel displacement). Geology had the highest Omission Ranking, resulting from it having the most negative impact on performance as compared to the other inputs when it was omitted from the models entirely. The PC approach, using the Correlation Rankings, found that the highest model performances were reached when the most recent radial tunnel displacement was added into the pool of candidate inputs. The three IVS approaches and their respective rankings proved to be useful for analyzing the CNN inputs in terms of importance and confirming underlying assumption about the deformation mechanics at Cigar Lake Mine. Collectively, the IVS analyses indicated that all of the available digitized inputs for the Cigar Lake Mine CNNs are needed to produce good model performances. Each IVS method revealed different insights into this CNN development. Undertaking IVS for ML developed using geomechanical datasets allows for verification of the algorithms and thereby a better understanding of the nuance of the rock mass deformation. At Cigar Lake Mine, these findings may be used to assist in forecasting the schedule and budget for ground support rehabilitation.

5.3 Introduction

The rock engineering community has access to increasing volumes of data concerning underground infrastructure and rock mass behaviour, which are collected both manually by geomechanical professionals and automatically through instrumentation. These datasets present opportunities to explore rock mass deformation phenomena more efficiently and in more detail than ever before. However, using conventional empirical or numerical methods may be inefficient or may introduce bias to engineering decision making if they are not used correctly (Elmo et al., 2020; Elmo & Stead, 2021). Conventional methods can oversimplify or mischaracterize the rock mass behaviour due to fitting the site-specific data to an empirical or constitutive behaviour that does not capture all the complex deformation mechanics. In practice, this is commonly overcome by examining several scenarios (i.e., worst, median, best cases), which increases the time needed to come to a suitable engineering decision.

Data driven methods, and Machine Learning Algorithms (MLAs) in particular, have emerged as an alternative approach to rock engineering analysis in research literature to increase the efficiency of data

analysis and obtain new insights into data trends (Elmo et al., 2020; Lawal & Kwon, 2020; Mcgaughey, 2019; Morgenroth et al., 2019). Geomechanical datasets are generally complex and challenging to work with due to the combination of data types and formats, including categorical (e.g., rock mass classification), numerical (e.g., instrumentation readings), spatial (e.g., geological mapping), and temporal (e.g., seismic monitoring) data. MLAs offer a new tool in the rock engineering toolbox to combine these complex datasets to characterize rock mass deformations, and to quantify the uncertainty associated with each datatype within the machine learning model.

MLAs are not often applied in rock engineering practice, despite representing an opportunity for rock engineering professionals to explore the site-specific data that is collected without introducing the inherent bias of constitutive frameworks and conventional numerical methods (Morgenroth et al., 2019). MLAs can identify patterns or relationships between datasets and do not require expert intervention to form inter-variable relationships (Khan & Valeo, 2017). Instead, the expert's judgment is reserved to choose the relevant input data and to evaluate whether the outputs are mechanistically possible in the context of the physical system. To date, research at the intersection of machine learning and underground rock engineering have successfully applied a variety of MLAs, including ANNs, a variation of which will be used in the present research. ANNs have been researched for predicting: rock mass properties (Sklavounos & Sakellariou, 1995; Song et al., 2015), constitutive behaviour (Kumar et al., 2013; Millar & Clarici, 2002), tunnel performance (Bizjak & Petkovšek, 2004; Leu et al., 2001; Qi et al., 2018; Sun et al., 2018), rock bursts (Pu et al., 2018; Ribeiro e Sousa et al., 2017), and tunnel blasting damage zones (Liu & Liu, 2017).

The hesitancy to adopt these data driven methods in rock engineering practice is due in part to the perceived opaque nature of the algorithms (Mcgaughey, 2019). In order to increase confidence in the MLA developed, it is important to uncover how the data inputs are used to train the algorithms, examine the physical implications of the algorithm architecture, and to verify that the outputs are aligned with the knowledge and experience of the collective field of rock mechanics. As with any model development, there are best practices associated with developing MLAs. These best practices have been an intensive field of research in computational science, with applications ranging from health sciences, to computer vision, to transportation engineering. The rock engineering community has the advantage of borrowing MLA architectures and techniques from fields that have already grappled with problems associated with data uncertainty and variability. For example: handling large noisy geospatial datasets (e.g., labelling aerial imagery from maps (Mnih & Hinton, 2012), or using bootstrapping to label noise to improve image segmentation (Reed et al., 2015)), maximizing the benefit of small datasets (e.g., detecting diabetic retinopathy using small datasets (Samanta et al., 2020)), or deciding which of the available inputs to include in an MLA to minimize complexity and maximize performance (e.g., identifying salient inputs in an ANN (R. May et al., 2011)). The type of MLA proposed herein is a CNN, which was originally designed for image classification and computer vision applications. CNNs are efficient for geospatial datasets because the geological context forms an image-like backdrop for the datasets (Morgenroth et al., 2021).

Numerical modelling approaches are based on constitutive equations, where the underlying physical behaviour of the model is based on rock properties and other inputs derived from an interpretation of the most relevant parameters, from laboratory testing, field observations, experience. MLA modelling approaches do not benefit from this advantage, and instead the input variables are selected from the available data, and the MLA is developed subsequently. This presents challenges when developing MLAs where there is a large pool of candidate input variables, giving rise to a higher potential for correlation or redundancy between candidates. More input variables result in an MLA with higher complexity, and thus potentially poorer model performance due to the presence of candidate variables that have low or no predictive capacity (May et al., 2011).

Input Variable Selection (IVS) can be used to overcome these challenges. IVS methods typically identify the most useful inputs from a candidate pool of inputs, where usefulness is defined as having the maximum relevance to the output while minimizing the redundancy between the other inputs (May et al., 2011). The IVS process also allows for refinement of the hyper parameters of the data-driven model, resulting in less frequent convergence to local minima during algorithm training, and thereby reducing the variability of the model output (May et al., 2008). Completing this analysis is valuable in the context of rock engineering because it allows the user to rank the inputs with respect to their ability to accurately predict the output (i.e., which variables have the strongest link to the phenomena being predicted). This may inform decisions about parallel modelling efforts, such as choosing variables for conventional numerical models, and about which data should be collected how, and how frequently, such as determining the type and layout of instrumentation.

Recent research has shown the utility of CNNs for predicting tunnel liner yield for the Cigar Lake Mine in northern Saskatchewan, Canada (Morgenroth et al., 2021). However, further work was needed to determine which CNN inputs were most important to obtain higher prediction accuracy, and therefore and IVS analysis was needed to improve model performance and determine which data should be emphasized for forecasting. The CNN used in this research was developed to predict the yield of the tunnel lining based on four inputs: geotechnical zone mapping (spatially variable), primary support class (spatially variable), ground freezing pattern (temporally variable), and measured radial tunnel displacement (temporally variable). The CNN as designed to predict high tunnel liner yield as accurately as possible, however, the overall performance of the model for lower classes of yield must not decrease as a result. Three IVS methods are explored herein for the Cigar Lake Mine CNN: Channel Activation Strength (CAS), Input Omission (IO), and Partial Correlation (PC). The latter two of these methods are novel approaches that have been developed for application to CNNs using spatial and temporal datasets by the authors. A detailed explanation of these novel methods is provided to allow geomechanical engineers to apply IVS to data formats that are familiar in the industry, such as the tunnel mapping that is used herein. The IVS methods allow for ranking of the inputs in terms of their relative importance on determining the tunnel yield at Cigar Lake Mine. Three IVS methods and their respective rankings (Activation Ranking, Omission Ranking, and

Correlation Ranking) were compared and contrasted to extract a better understanding of the inputs that produce the complex rock mass deformation at the Cigar Lake Mine. This information can be used to inform how data is collected, and to better understand the factors influencing the rock mass deformation at the Cigar Lake Mine.

The remainder of this chapter is organized as follows: Section 2 contains the background and context of the Cigar Lake Mine case study; Section 3 explains the development and architecture of the Cigar Lake Mine CNN used for this IVS research; Section 4 explains in detail the mechanics of the three IVS methods that were developed and applied; Section 5 contains the IVS study results; and Sections 6 and 7 comprise the discussion and conclusions, respectively.

5.4 Background

5.4.1 Cigar Lake Mine

The Cigar Lake Mine is located in northern Saskatchewan, Canada. It is the world's second-largest uranium mine, with an ore grade approximately one hundred times the global average. The ore body is unique due to its size, high grade, intensity of alteration, and a high degree of associated hydrothermal clay alteration (Bishop et al., 2016).

The geology of the Cigar Lake uranium deposit and environs is described in (Bishop et al., 2016) and is summarized as follows. The uranium is found in an unconformity type deposit situated between the Athabasca Group and the underlying metasedimentary Proterozoic Wollaston Domain. The deposit and host rock consist of three geological units that double as the geotechnical domains: the deposit and associated hydrothermally altered clay cap, the overlying sandstone unit (Athabasca Group), and the underlying metamorphic basement rock (Wollaston Domain). The overlying sandstone unit is part of the Manitou Falls Formation within the Athabasca Group, which is on the eastern side of the sedimentary basin. There are three distinct styles of mineralization within the ore body: high grade mineralization at the unconformity, fracture-controlled mineralization in the sandstone, and fracture-controlled mineralization in the basement rock. The deposit and sandstone are highly fractured and water bearing, while the basement rock is impervious. The basement rock is composed mainly of pelitic metasedimentary gneisses belonging to the Wollaston Domain. The Wollaston Domain is considered to be the most favourable unit for uranium mineralization. In general, the Cigar Lake Mine operation and production tunnels are located in three main rock mass types (Paudel et al., 2012): weak, highly weathered and saturated basement rock containing sand and clay; moderately weathered saturated basement rock; and strong unweathered basement rock.

The high grade ore body is located above the 5.0 m diameter lined ore extraction tunnels (Figure 5-1), which are excavated using drill and blast methods. There are two main challenges facing the stability of the Cigar Lake Mine excavations: controlling groundwater inflow, and supporting areas of weak rock (Bishop et al., 2016). The mine operators freeze the rock mass surrounding the orebody downward from surface to improve rock mass properties and to restrict groundwater inflow into excavated areas (Figure 5-1). Tunnel

convergence is monitored using survey targets around the circumference of the tunnel. A customized non-entry extraction method is used to extract the ore, where cutting cavities of frozen ore are created upwards from the ore extraction tunnels below with a high-pressure water jet, and then the ore is mixed with water to make a pumpable slurry (WSP (formerly Parson Brinkerhoff Quade & Douglas Inc.), 1999). The cavities created from the extraction of the ore are backfilled with concrete with a custom mix design to achieve early strength in the frozen ground and therefore, provide additional ground support (Bishop et al., 2016). The ground freezing operation results in complex time-dependent rock mass behaviour that is difficult to predict and presents challenges when designing support (Golder Associates, 2001; Roworth, 2013).

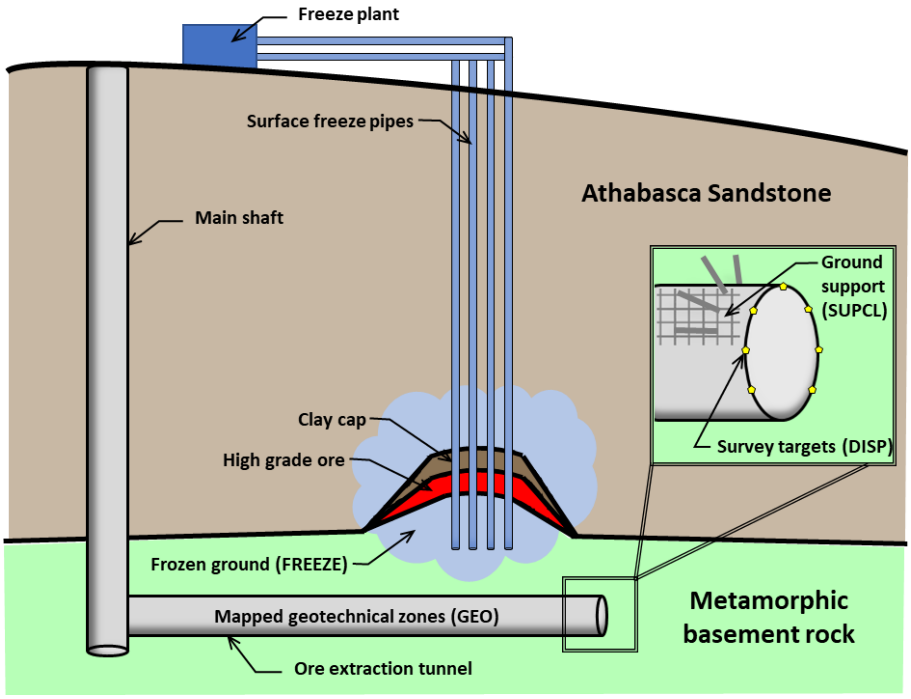

Figure 5-1. Schematic of the mining method at Cigar Lake Mine, showing a section view of the ore extraction tunnel and freeze holes from surface. A customized non-entry extraction method is used to extract the ore from the ore extraction tunnels which are below the high grade ore body. Ground freezing is implemented to stabilize the excavations in the adverse geology and to manage ground water inflow from the Athabasca Sandstone. Tunnel convergence is monitored using survey targets placed on the ground support around the circumference of the tunnel. The geotechnical zones (GEO), installed primary ground support class (SUPCL), ground freezing (FREEZE), and surveyed displacements (DISP) are used as inputs into in the proposed Convolutional Neural Network.

The Cigar Lake Mine was found to be an appropriate case study for an ML application because there are several phenomena that coalesce into a complex tunnel deformation, which is difficult to predict using existing techniques to characterize squeezing ground (Barla, 2002; Barla et al., 2011; Barla & Borgna, 1999; K. Zhao et al., 2015). The combination of unfavourable geology (i.e., the weak and altered sandstone and basement rock) and the locally variable structural geology produces differential radial squeezing over the length of the ore extraction tunnels. The ground freezing regime, which is effective at stabilizing the

rock mass for ore extraction, adds to the complex rock deformation mechanics by increasing the total stress during its implementation (Roworth, 2013). These factors combine to result in tunnel liner and support yield that is difficult to forecast, which causes delays in production and adds complications for support rehabilitation scheduling and budgeting. The priority for the CNN approach for Cameco, the mine operator, was to accurately predict tunnel liner and support yield, which is critical for the safe operation of the ore extraction tunnels and for forecasting rehabilitation works needed to maintain operations. The IVS approaches presented herein indicate which input variables have the highest impact on predicting the tunnel liner and support yield, thus giving the mine operators insight into which variables should be monitored and recorded more accurately, and which are not as strongly related to the yield. This information may help the mine prioritize its data collection and analysis efforts in order to forecast their annual rehabilitation budget and schedule more accurately.

5.4.2 Convolutional Neural Network

CNNs are a type of ANN that are efficient at processing spatial and temporal dependencies in image or raster datasets. Since their introduction in the late 1980s for handwritten digit classification (LeCun et al., 1989), CNNs have rapidly gained success in the field of computer vision as well as in other engineering applications (Zeiler & Fergus, 2014). They are computationally efficient at incorporating spatial dependencies in images by using data from surrounding pixels and their change over time in a pixel-by-pixel analysis. This is applicable to rock engineering problems because the data usually contains a mapping component, such as geology, onto which other data, such as measured displacement or groundwater inflow, can be transferred or overlain. Thus, some data stays constant with time while others change with each subsequent timestep. This approach maintains the spatial dependencies between the static and variable inputs, while also capturing the temporal change in the variable inputs. When developing any kind of rock mechanics model, the spatial and temporal relationships between the various inputs must be preserved, and thus, CNNs are well suited for these types of datasets.

CNNs processes conventional images in a similar manner to the manual processing of spatial geomechanical data and raster analysis in Geographic Information Systems. During convolution, CNNs scan the pixels of an input image using a square search area called a filter to identify the presence of a feature of interest, for example a crack in the shotcrete liner, to produce a feature map. This feature map is then correlated to the output, for example the degree of yield the tunnel liner has sustained over time. The CNN will produce feature maps that together identify the factors that produce the output. CNNs have also demonstrated success at more challenging computer vision tasks involving classifying objects in an image or video (Zeiler & Fergus, 2014), called *segmentation*. This type of CNN architecture allows for classification of multiple phenomena within an input image as opposed to classifying the image as a whole. For example, a CNN with segmentation can segment and identify areas within the input tunnel map with various amounts of ground support damage, producing an output image that is akin to a hazard map.

The Cigar Lake Mine CNN was developed in MATLAB R2019b (MathWorks Inc., 2019) using the Deep Learning and Computer Vision toolboxes. The input data was formatted as labelled images, which were passed to the built-in MATLAB function *trainNetwork* along with the desired algorithm architecture and hyperparameters to train the Cigar Lake Mine CNN. The MATLAB code for this research is publicly available (Morgenroth, 2021a).

The objective of the Cigar Lake Mine CNN is to forecast the ore extraction tunnel liner yield forward in time, in order to optimize the schedule and budget for liner rehabilitation. The dataset used to develop the Cigar Lake Mine CNN is comprised of a combination of spatial and temporal data that is both categorical (e.g., the geotechnical zones, which are categorized as Class 1, 2, 3, etc.) and numerical (e.g., measured radial displacement within the tunnel). This study used data from Cigar Lake Mine crosscut tunnel 765 to calibrate and test the CNN to predict tunnel liner yield forward in time. The dataset for tunnel 765 was digitized from five Ground Management Plans (GMPs) spanning from October 2015 to June 2016, at a spatial resolution of 164 m, the length of the tunnel, by 13 elements, representing elements around the circumference of the tunnel. Three inputs were digitized from the GMPs for use in this case study: the mapped geotechnical zones (GEO), as-built ground support class (SUPCL), and ground freezing patterns (FREEZE). The radial tunnel displacement (DISP) was obtained from survey point measurements made around the circumference of the tunnel. The inputs used to train the Cigar Lake Mine CNN are indicated in Figure 5-1, and their categories and values are presented in Figure 5-2. The GEO classes are defined as follows: Class 1 is stiff crystalline rock; Class 2 is fissured fresh crystalline rock; Class 3 is transition from fresh to metamorphosed crystalline rock; Class 4 is metamorphosed crystalline rock; Class 5 is stiff to very stiff coarse-grained soil, structure well preserved; and Class 6 is fine-grained soil, structure not seen. The SUPCL classes are defined as follows: Class 1 support (lightest support); Class 2 support; Class 3 support; Class 4 support; Class 5 support (heaviest support). The FREEZE input is binary: 0 is freezing is off at that tunnel station; and 1 is freezing is on at that tunnel station. Finally, DISP is the measured radial displacement, ranging from 0 to 250 mm.

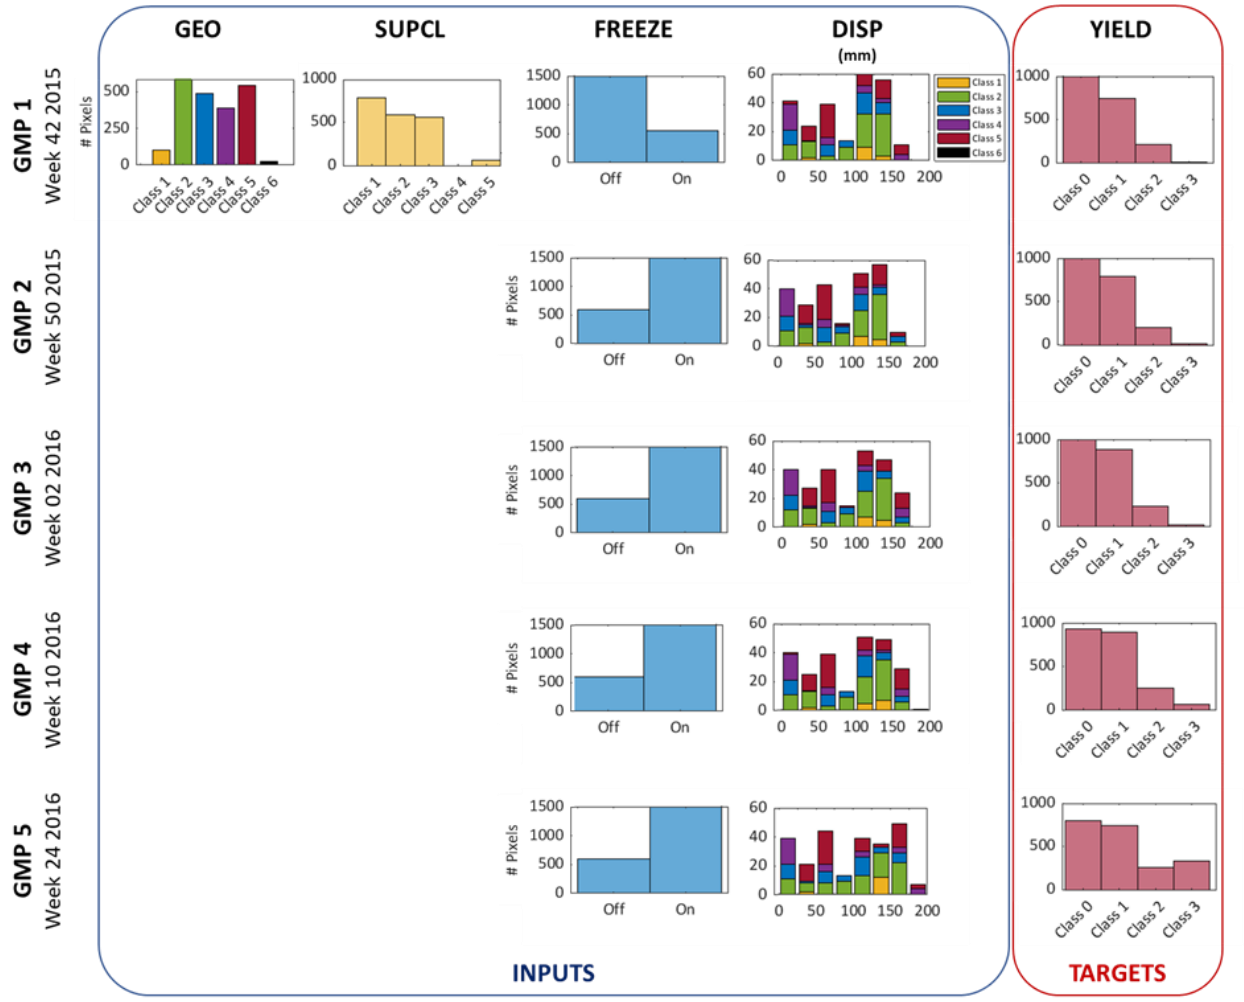

Figure 5-2. Histograms of Cigar Lake Mine dataset showing the distribution of each input for each of the five Ground Management Plans (GMPs), where each row of ground freezing, displacement, and liner yield histograms represents a GMP. The displacement histograms indicate the distribution of geology classes for each GMP. It is important to note that the rings where displacement is measured does not span areas where geology Class = 6, and as a result this class is not represented in the displacement bars (adapted from Morgenroth et al., 2021a).

The four inputs (GEO, SUPCL, FREEZE, DISP) were used to predict the tunnel liner yield (YIELD), which was also mapped on the GMPs in the form of failed rock bolts, spalled shotcrete, and compressed yield packs. YIELD classes are defined from discussions with Cameco as follows (Morgenroth et al., 2021):

- Class 0 – no yield of tunnel liner elements
- Class 1 – minor yield: minor cracks in shotcrete; isolated water ingress; minor spalling of lining or yield elements
- Class 2 – major yield: high concentration of shotcrete cracks; localized major shotcrete cracks; extensive spalling of lining; combination of cracks and/or spalling and/or water ingress; damaged rock bolts

- Class 3 – total reprofiling required: extensive liner spalling and major shotcrete cracks; combination of spalling, shotcrete cracks, and water ingress; compressed yield elements

The optimization and development of the CNN architecture for the Cigar Lake Mine is detailed in (Morgenroth et al., 2021), summarized here and illustrated in Figure 5-3. Each of the four inputs are formatted into an image, where each input represents a channel of the image, similar to how RGB values represent each of the three channels of a typical image, i.e., the image for GMP 1 consists of four channels which are GEO, SUPCL, FREEZE and DISP all digitized from GMP 1. The Cigar Lake Mine CNN is trained on a subset of GMP images and then tested on the subsequent GMP image, to simulate prediction tunnel liner yield forward in time. Specifically, the data was split temporally and not spatially, as the liner yield is being forecasted after the tunnel was complete during the ore extraction process. Ten permutations of training/testing data splitting were analyzed: train on GMP 1, test on GMP 2; train on GMP 2, test on GMP 3; train on GMP 3, test on GMP 4; train on GMP 4, test on GMP 5; train on GMPs 1 and 2, test on GMP 3; train on GMPs 2 and 3, etc. Detailed discussion of the Cigar Lake Mine input dataset and image formatting can be found in Morgenroth et al. (2020) and Morgenroth et al. (2021).

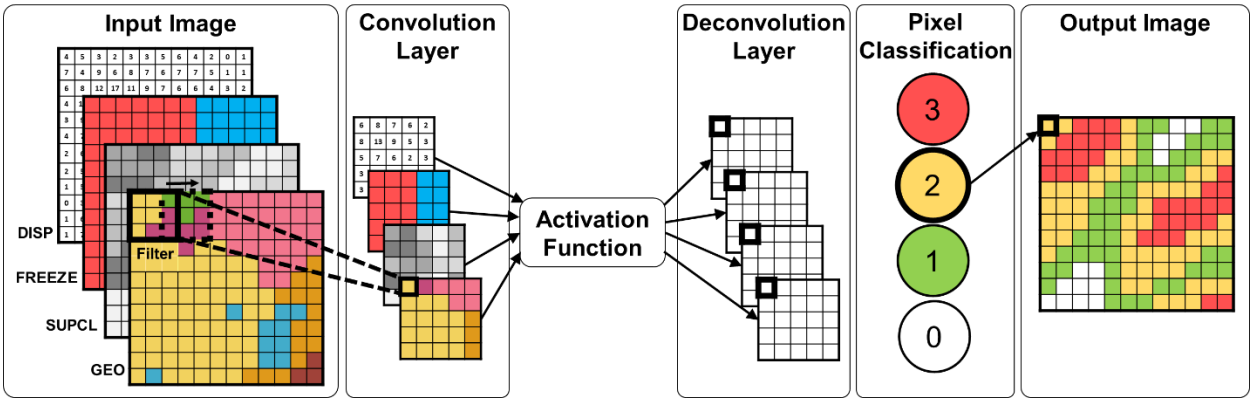

Figure 5-3. A schematic of the Cigar Lake Mine Convolutional Neural Network (CNN), where a filter is used to convolve over the inputs to generate the feature maps. The inputs are mapped geotechnical zones (GEO), primary installed support class (SUPCL), ground freezing (FREEZE), and radial tunnel displacement (DISP). The CNN uses the feature maps created in the convolution layer to make a pixel-by-pixel classification based on whether the combined inputs of a particular cell results in Class 0, 1, 2, or 3 tunnel liner yield (adapted from Morgenroth et al. (2021b)).

The Cigar Lake Mine CNN was developed through a sensitivity analysis of three network hyperparameters: the amount of training data, the convolution filter size, and the error weighting scheme (Morgenroth et al., 2021). A total of 7200 models were analyzed to determine the best combination of these parameters for this problem. It was found that more data in the temporal realm resulted in better performance, and the optimal filter size was determined to be 30 by 30 pixels. It was found that an error weighting scheme was necessary to offset the imbalance of samples for each class in the dataset. A sigmoid error weighting scheme, which prioritizes Class 2 and 3 errors during training, was used to develop a Targeted Class 2/3 model that has stronger prediction accuracy for Class 2 and 3 tunnel liner yield (model accuracy of 99.9% for Class 2). An inverse frequency error weighting scheme, where the proportion of samples in each class

is used to determine the error weights, was used to develop a Balanced Global model that has reasonable prediction accuracy across all yield classes (average model accuracy >65% for all training data permutations).

Herein, three IVS techniques are applied to both the Targeted Class 2/3 model and the Balanced Global model that were developed for the Cigar Lake Mine CNN, including one standard approach (CAS) and two novel approaches (IO and PC). In order to obtain a distribution of performances for each IVS method, an ensemble of 30 models was computed for each set of hyper parameters. Modelling an ensemble also allows for the quantification of uncertainty for each model. This approach resulted in 300 CAS models, 2,400 IO models, and 48,000 IO models analyzed as part of this study.

5.4.3 Performance Metrics

The performance metric selection for an MLA depends on the model architecture and the amount of training data. For this study it was important that the performance metric used to compare the models, and therefore the IVS methods, included the number of model parameters in its formulation. This is because the three IVS methods have different numbers of parameters associated with them, depending on the number of inputs used and the amount of training data. The chosen performance metric was the Corrected Akaike Information Criterion (AIC_c) (Hurvich & Tsai, 1989), which is based on the Sum of Squared Errors (SSE), number of samples, and number of model parameters, as described below.

The SSE is a measure of the variation of modeling errors, i.e., a measure of how the variation in the dependent variable in a model cannot be explained by the model. Generally, a lower SSE indicates that the regression model can better explain the data while a higher residual sum of squares indicates that the model poorly explains the data. The SSE calculation is shown in Equation 5-1, where y_i is the ground truth (observed data) and $\hat{y}_i$ is the predicted value.

$$SSE = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \qquad \text{Equation 5-1}$$

The Akaike Information Criterion (AIC) evaluates the model's fit on the training data, and adding a penalty term for the complexity of the model (Akaike, 1969). The desired result was to find the lowest possible AIC (the absolute value was not significant since the objective is to compare performance), which indicates the best balance of model fit with generalizability. Essentially, AIC is a measure of the model's entropy and is most frequently used in situations where one is not able to easily test the model's performance on a test set in standard ML practice (e.g., small dataset, time series data). The AIC calculation is shown in Equation 5-2, where n is the number of training samples and p is the number of parameters (weights and biases).

$$AIC = n * log\frac{SSE}{n} + 2 * p \qquad \text{Equation 5-2}$$

AIC assumes that the same data is used between models, the same outcome variable is predicted between models, and that the sample is of infinite size. This final assumption gave rise to a sample-size adjusted

formula known as the Corrected Akaike Information Criterion (AIC_c), which adds a correction term that gives a more accurate answer for smaller datasets, and which is used herein (Hurvich & Tsai, 1989). The AIC_c calculation is shown in Equation 5-3.

$$AIC_c = n * log\frac{SSE}{n} + \frac{n+p}{1-\frac{p+2}{n}} \quad \text{Equation 5-3}$$

AIC_c should be used when the ratio of training samples (n): number of parameters (p) is less than 40 (Burnham & Anderson, 2002), i.e., when there are less than 40 samples of data for every training parameter in the model. The AIC_c approaches and converges to the AIC value as the number of training samples approaches infinity.

5.4.4 Algorithm Interpretability

MLAs must be trained and validated in the specific context of the site for which they are developed and should not be transferred directly from another project context without revisiting its validation (Elmo et al., 2020), just as is done with numerical models in rock engineering practice currently. To this end, it is important to select engineering verification metrics for the algorithm outputs in advance of developing it to avoid confirmation bias, and also to select methods to interpret the inner workings of the MLA. Algorithm interpretability falls into two categories: model transparency (how does the model work?), and post-hoc explanations (what can be learned after its development?) (Lipton, 2016). An algorithm is interpretable if there is an understanding of how the model makes decisions based on its inputs and the learned components such as hyperparameters and architecture (Molnar, 2022).

Recent research literature has highlighted several options for algorithm interpretation, such as: input importance ranking, partial dependence among inputs, input interaction, and developing a secondary algorithm to recover features within inputs that are used by the algorithm (Fathipour-Azar, 2021; Isleyen et al., 2021; Pu et al., 2019). Researchers also place emphasis on practitioners focusing on developing the proper inputs for rock engineering applications, instead of focusing on the algorithms themselves (Mcgaughey, 2020). The IVS methods in this chapter were selected with the aim of making the Cigar Lake Mine CNN more interpretable, by uncovering both how the model works and what can be learned from it after its development.

5.5 Input Variable Selection Methods

The IVS process is not standard practice during MLA development, however it can serve to link these data-driven approaches to the physical system that is being modelled, thereby increasing user confidence that the MLA is producing results consistence with the fundamental principles of their field. For the Cigar Lake Mine CNN, the primary purpose of undertaking IVS is not to reduce the number of inputs, but rather to rank the available candidate variables to determine their relative importance in predicting tunnel liner yield.

IVS can be separated into two broad categories: model-based and model-free, as illustrated in Figure 5-4. Model-based methods are embedded into the training of the algorithm, while model-free methods are

distinct from the algorithm training and aims to measure the relevance and redundancy of individual input variables. For a more detailed taxonomy of IVS methods, the reader is referred to May et al. (2008). The IVS methods considered herein are channel activation strength (model-based), input omission (model-based), and partial correlation (model-free). These are explained in more detail in the subsequent sections.

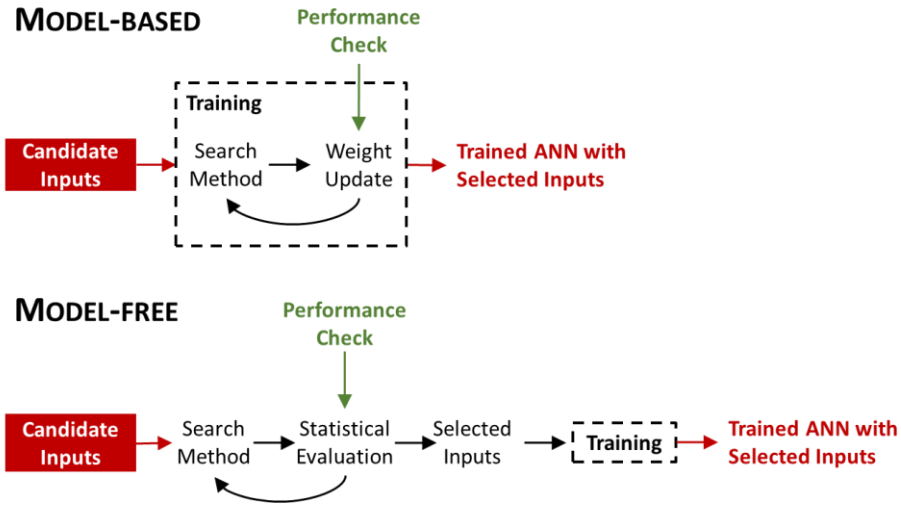

Figure 5-4. Conceptual schematics showing the implementation of model-free and model-based Input Variable Selection methods (adapted from May et al., 2008).

In the case of the Cigar Lake Mine dataset and CNN introduced above, the four input variables are not independent. In practice, SUPCL and FREEZE are chosen by the ground control engineers based on GEO, which result in DISP. Therefore, the use of IVS can uncover not only the significance of each individual input on predicting the target, but also how the interdependence among the inputs impacts overall model performance. The relationship between input variables is itself information the CNN can learn during training. For example, a tunnel location with a “worse” GEO class is correlated with a higher SUPCL, and over time may experience more tunnel liner yield. IVS can reveal whether the input correlations are additional useful information to the CNN, or redundancies that reduce model performance. To draw a parallel between the CNN and more conventional numerical modelling, for example, a finite difference model of the mine excavations would be initiated using the relevant GEO and FREEZE rock mass parameters and the SUPCL elements would be included explicitly. The model would be computed to determine DISP and after a predetermined amount of time and the resulting YIELD, which is compared to the observed tunnel liner behaviour. The input parameters are calibrated until the model output matches the true ground behaviour. In this example, the IVS analysis is comparable to completing a stress history analysis of the computation.

One possible approach to using non-independent input variables is to apply principal component analysis (PCA) to reduce the dimensionality of a dataset, with the goal of obtaining lower dimensionality data while preserving its variation. However, there are limitations to applying PCA to capture non-linear correlations,

such as those encountered in most geomechanical datasets. PCA is disadvantaged if the data has not been standardized, and if not performed properly there is a high likelihood of data loss (Shlens, 2014). For these reasons, caution should be used in applying PCA in rock engineering without first completing an input data analysis.

Each of the IVS methods explored as part of this research produced a different type of ranking that reveals information about the "usefulness" of each input to the CNN, and together provide a more comprehensive interpretation of the CNN. For clarity each ranking has been given a distinct name. CAS was used to produce an Activation Ranking, IO was used to produce an Omission Ranking, and PC was used to produce a Correlation Ranking. These are described in more detail in the following sections.

5.5.1 Channel Activation Strength (CAS)

The CAS method is a "model-based" approach that examines the activations of each channel for a particular layer within the CNN, revealing which features the network learns and which channels have the strongest impact on learning (Kudo et al., 1999). In other words, CAS measures the sensitivity of each input in an algorithm trained with all inputs. The activations of each input are their value after they are passed through the activation function, as illustrated in . The activations can be used to produce an Activation Ranking of all the candidate inputs. In the context of this research, the Activation Ranking reveals which input has the strongest impact on the predictand and therefore, which data is the most important indicator for a particular rock mass deformation mechanism within the constraints of the model. In the CAS approach, the candidate inputs are not being manipulated, left out, or inserted in any way. By examining the channel activations for each combination of model variables in the Cigar Lake Mine CNN, it is possible to determine which of the four inputs contain the strongest signal that is used by the CNN to make its classification for the Cigar Lake Mine tunnel. This type of analysis allows the CNN developer to see which inputs are emphasized in the learning process, and what the network learns at each layer.

5.5.2 Input Omission (IO)

The IO method is a "model-based" approach, which estimates the usefulness of each input by iteratively examining model performance when an input is left out from the full set on which the model has been trained (Setiono & Liu, 1997). In other words, IO measures the algorithm's sensitivity to removing an input entirely. This is illustrated in Figure 5-5, where each of the four input channels of one training GMP are iteratively left out. The significance, or lack thereof, of each input can then be compared based on the error that is produced by leaving it out. An Omission Ranking can be produced whereby the input with the highest negative impact it ranked first, and so on. It is generally understood that excluding redundant or irrelevant inputs from the CNN results in a model with a higher generalization capability (May et al., 2011). Relevant and irrelevant inputs are differentiated by their impact on performance, and inputs are selected for removal based on their saliency (Setiono & Liu, 1997). The saliency of an input is based on the derivative of the network error function, the weights of the network, or both. In order to obtain a confidence interval on the mean saliency value of a network, the network should be retrained with random weights at least 30 times

(Setiono & Liu, 1997). In this study, the saliency measure is the accuracy of the network when removing each of the inputs. The Omission Ranking and accuracy can be used in combination to determine if any inputs should be removed to improve overall model performance.

Herein, the IO method has been developed from its traditional application, where each input is a discrete value entered into the algorithm, to a format where each individual channel of an image is being omitted. This is a novel approach to the IO method because it is generally not desired to omit individual image channels in typical CNN image segmentation, as defined in section 5.4.2 Convolutional Neural Network.

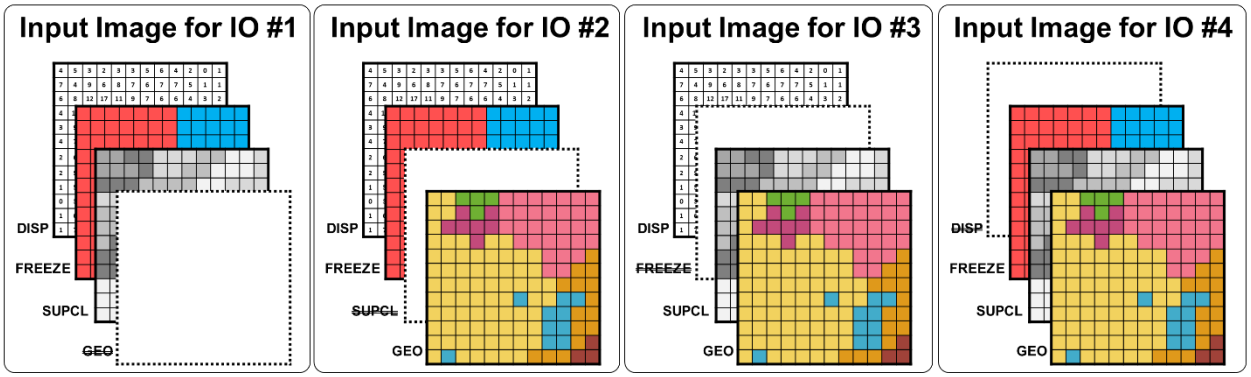

Figure 5-5. Conceptual schematic of the Input Omission (IO) method adopted for the Cigar Lake Mine Convolutional Neural Network (CNN), where one Ground Management Plan (GMP) image comprised of 4 channels is used as the training image for the subsequent GMP. The CNN is iteratively trained with one input channel omitted, and the resulting CNN performances are compared to discern whether any of the input candidates are not strictly required to obtain good model performance.

5.5.3 Partial Correlation (PC)

The PC method is an IVS method that it does not depend on a pre-existing model, hence “model-free”. This method iteratively selects inputs using a feed-forward algorithm in which the partial correlation between candidate inputs and the target are calculated. Subsequently a Correlation Ranking is produced by placing the candidate inputs in the order of highest correlation to lowest correlation, where the absolute value of the correlation is used for ranking to prevent strong negative correlations for being ranked the lowest. The inputs are then iteratively added to the pool of inputs used in training the model according to their respective Correlation Ranking (Figure 5-6). This selection process is repeated until the performance criterion does not improve, the desired number of inputs is reached, or there are no more candidate inputs (He et al., 2011). This method is not as sophisticated as methods that employ non-linear algorithms, however it is advantageous due to its low computational expense and reasonable accuracy (Snieder et al., 2019). This method is suitable for large candidate input sets, especially where little to no expert knowledge of the system exists. Currently, no formulation of the partial correlation exists for CNNs in the literature, and therefore it was developed for this study.

Similar to the IO method presented, the PC method developed in this study is novel in its application to CNNs. This is because each channel of all the available input GMPs is treated as a separate input for the initial partial correlation calculation, and the channels are added into the CNN regardless of which GMP image they originally belonged to. In other words, if the Correlation Ranking interweaves channels from different GMPs, they are added in that order without including the other channels from the same GMP at the same time. Like with the IO method, this is not generally desired in typical image segmentation algorithms.

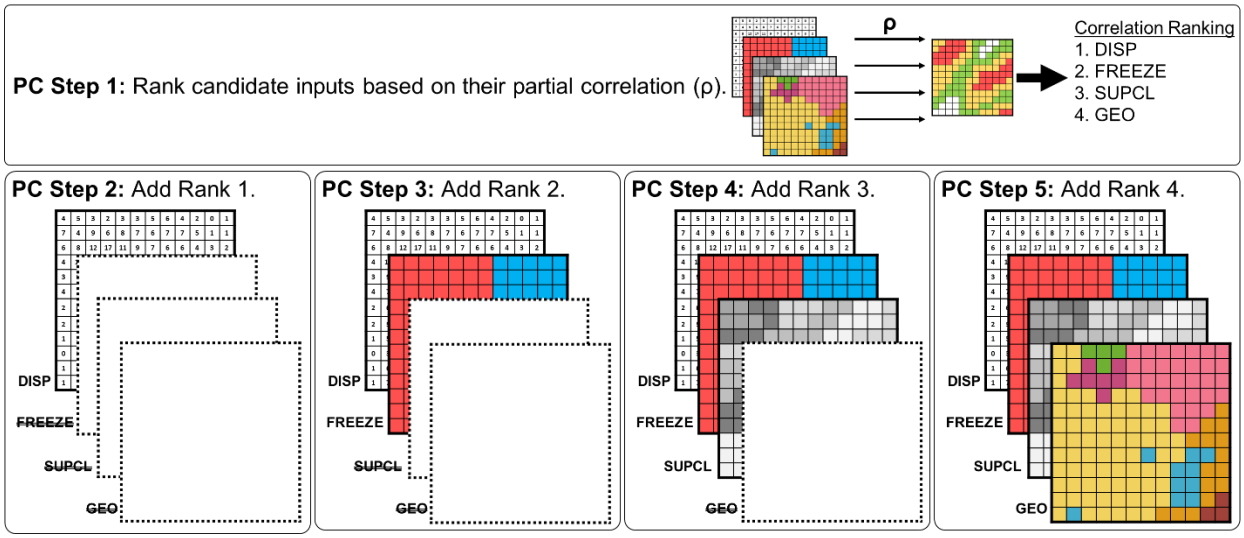

Figure 5-6. Conceptual schematic of the Partial Correlation (PC) method adopted for the Cigar Lake Mine Convolutional Neural Network (CNN), where one Ground Management Plan (GMP) image is used as the training image for the subsequent GMP. First, the candidate inputs are ranked based on their PC with the output (ρ) in the Correlation Ranking. Then the CNN is trained starting with the candidate input with the highest Correlation Ranking, iteratively adding the next input until all candidates have been added.

5.6 Results and Discussion

The three IVS methods described above were applied to two Cigar Lake Mine CNN architectures: the Global Balanced model, and the Targeted Class 2/3 model. All ten permutations of training/testing data splitting were analyzed for each IVS method. The results of each analysis are discussed in this section. All AICc values computed are presented in Table 5-4, Table 5-5, and Table 5-6 in the Appendix for the CAS, IO, and PC methods, respectively.

5.6.1 Channel Activation Strength (CAS)

The results of the CAS approach are shown in Figure 5-7. Each box in the figure represents an ensemble of models, and the whiskers represent the 25th and 75th percentiles for each ensemble. A total of 300 model runs were included in the CAS analysis. The left column of subplots (Figure 5-7a, c, e, g) shows AIC_c boxplots for each GMP that is predicted, i.e., Figure 5-7b shows Targeted Class 2/3 and Balanced Global model performance when GMP 3 is being predicted, where the boxplots represent the aggregate of the

models trained using only GMP 2 and GMPs 1-2. This performance represents the base line or control group of this IVS study, because for the CAS approach all the candidate inputs are included. The right column of subplots (Figure 5-7b, d, f, h) shows the activation strengths of the model inputs as a percentage of the total models predicting a particular GMP, aggregated across all combinations of training data, i.e., Figure 5-7d shows the Targeted Class 2/3 and Balanced Global input activation strengths for all models where GMP 3 is being predicted.

The left column of subplots showing the performance of the models in terms of AIC_c indicate that for all GMPs being predicted, the addition of more data results in a better performance in the form of a lower AIC_c. For example, Figure 5-7g shows that the AIC_c decreases when predicting GMP 5 using GMP 4 (AIC_c = -579 and -2001 for the Targeted Class 2/3 and Global Balanced models, respectively) versus predicting GMP 5 using GMPs 1 through 4 (AIC_c = -16278 and -26361 for the Targeted Class 2/3 and Global Balanced models, respectively). This confirms the intuitive notion that more information on the rock mass deformation with time will increase the model's ability to predict its behaviour forward in time. These plots also show that overall prediction ability increases with each subsequent GMP, i.e., where GMP 2 is being predicted in Figure 5-7a the AIC_c is higher (AIC_c = -403 and -3025 for the Targeted Class 2/3 and Global Balanced models, respectively) than Figure 5-7g where GMP 5 is being predicted (AIC_c = -579 to -16278 for the Targeted Class 2/3 and -2001 to -26361 for the Global Balanced models). Generally, the Targeted Class 2/3 models have poorer performance than the Global Balanced model for the same training/testing data split, e.g., Figure 5-7c where AIC_c = -442 and -4200 for the Targeted Class 2/3 and -3153 and -11727 for the Global Balanced models. This is expected as the Targeted Class 2/3 model was designed to have high accuracy in the higher tunnel yield classes while the Global Balanced model is trained to have good performance across all yield classes. The mean AIC_c values for all ensembles computed for the CAS approach are reported in Table 5-4 in the Appendix.

The right column of subplots in Figure 5-7 show the activations of each input and indicate that the primary support class (SUPCL) is the dominant input for this dataset, regardless of which GMP is being predicted or which of the two model architectures is being evaluated. The radial tunnel displacement (DISP) input shows approximately 5% higher strength on average in the Targeted Class 2/3 models because these models emphasize YIELD Class 2 and 3 during training, which are more strongly correlated with DISP (Morgenroth et al., 2021). In general, this analysis shows that geology class (GEO) is never the strongest signal used by the CNN to predict tunnel liner yield (YIELD). This may seem counter intuitive, since in the rock mechanics community it is known that geology plays a crucial role in rock mass deformation, but in ML parlance this indicates that there is another input that is conveying the same information to the network or has a stronger correlation than GEO, i.e., there may be redundancy amongst the inputs. In this case it is likely that the SUPCL input is also providing this signal, as the primary installed support is installed based on the observed geology during tunnel excavation.

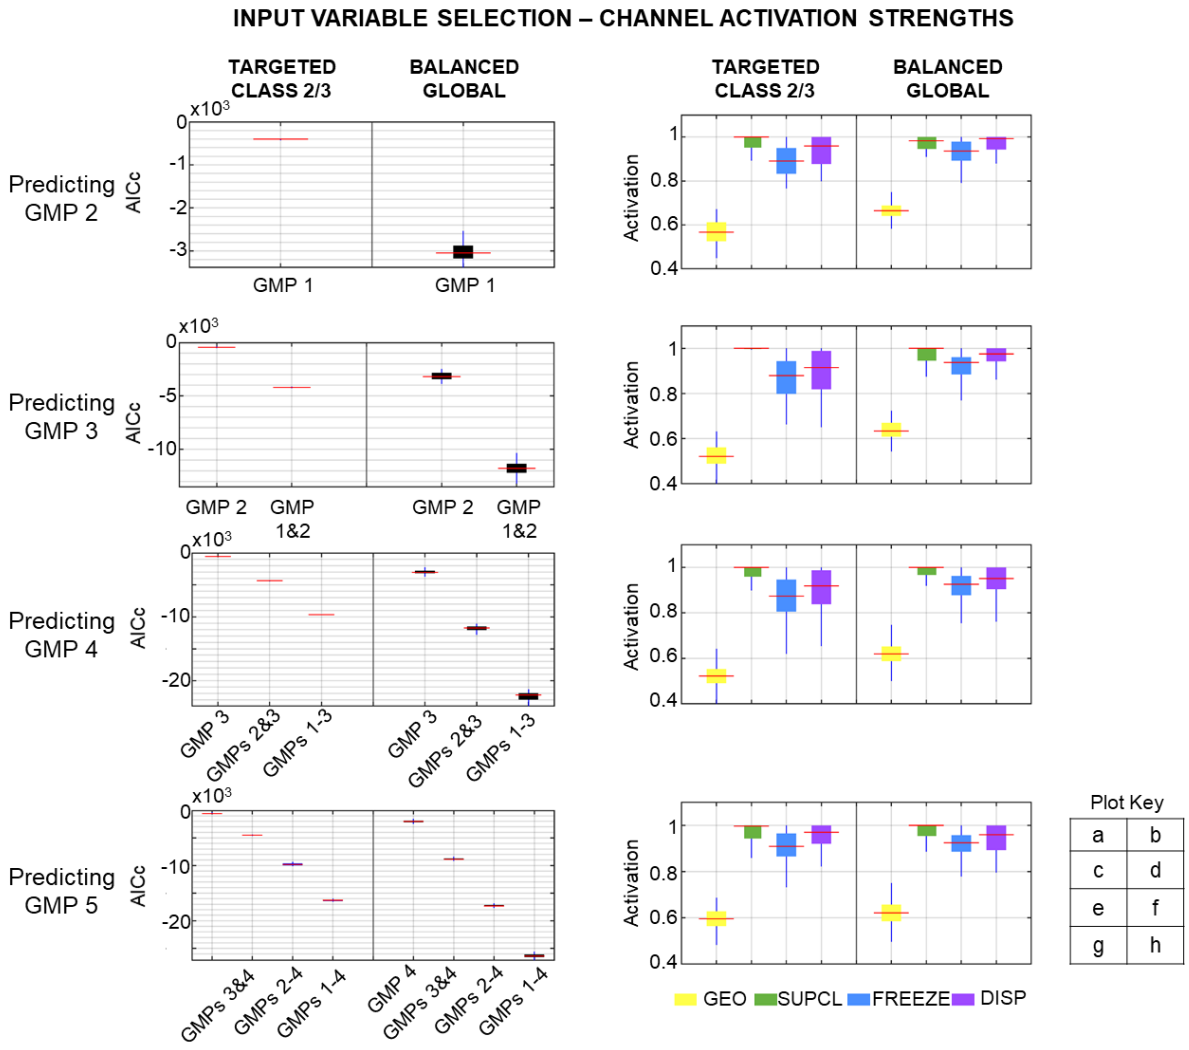

Figure 5-7. Results of the Channel Activation Strength (CAS) Input Variable Selection (IVS) approach for Cigar Lake Mine. The left column of plots shows the 30 model ensemble performance, AICc, for the Targeted Class 2/3 and Balanced Global models, for each permutation of training and testing Ground Management Plans (GMPs). The right column of plots shows the activations for each of the four inputs, for each GMP being predicted where the number of models in the ensemble is indicated at the bottom of the subplot.

Based on the CAS results, the Activation Ranking of the CNN inputs is SUPCL, DISP, FREEZE, GEO (Table 5-1). This ranking indicates which input signal is relied on most by the CNN to make a prediction of YIELD once it has been trained. It is important to note that the first three inputs are extremely close in terms of activation value (within 20%), while GEO is much less dominant with respect to the activation strength. The remaining IVS methods, namely IO and PC, can be used to determine whether this difference is attributed to input redundancy (i.e., redundant signals between inputs) or to GEO not being used at all during CNN training (i.e., it can be removed from the input pool altogether).

Table 5-1. Activation Rankings for both the Targeted Class 2/3 and Global Balanced models, produced by the Channel Activation Strengths (CAS) method.

	IVS Method & Data Split		INPUTS															
			GMP 1				GMP 2				GMP 3				GMP 4			
	Train	Test	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP
ACTIVATION RANKING	Targeted Class 2/3 Model																	
	GMP 1	GMP 2	4	3	1	2												
	GMP 2	GMP 3					4	3	1	2								
	GMP 3	GMP 4									4	3	1	2				
	GMP 4	GMP 5													4	3	2	1
	GMP 1&2	GMP 3	4	3	1	2	4	3	1	2								
	GMP 2&3	GMP 4					4	3	1	2	4	3	1	2				
	GMP 3&4	GMP 5									4	3	1	2	4	3	1	2
	GMP 1-3	GMP 4	4	2	1	3	4	2	1	3	4	2	1	3				
	GMP 2-4	GMP 5					4	3	1	2	4	3	1	2	4	3	1	2
	GMP 1-4	GMP 5	4	3	1	2	4	3	1	2	4	3	1	2	4	3	1	2
	Global Balanced Model																	
	GMP 1	GMP 2	4	3	2	1												
	GMP 2	GMP 3					4	3	1	1								
	GMP 3	GMP 4									4	3	1	1				
	GMP 4	GMP 5													4	2	3	1
	GMP 1&2	GMP 3	4	3	1	2	4	3	1	2								
	GMP 2&3	GMP 4					4	3	1	2	4	3	1	2				
	GMP 3&4	GMP 5									4	3	1	1	4	3	1	1
	GMP 1-3	GMP 4	4	2	1	3	4	2	1	3	4	2	1	3				
	GMP 2-4	GMP 5					4	3	1	2	4	3	1	2	4	3	1	2
	GMP 1-4	GMP 5	4	3	1	2	4	3	1	2	4	3	1	2	4	3	1	2

The first IVS method investigated was the CAS approach, which was the simplest to implement with 300 models run. This approach extracted the channel activations from the trained CNN to determine which inputs had the most influence on the output. This approach was not intended to determine which input should be removed, but rather produced an Activation Ranking to determine which inputs are used most, and are therefore the most dominant, in the Cigar Lake Mine CNN. The Activation Ranking was similar for both the Targeted Class 2/3 models and the Global Balanced models: SUPCL, DISP, FREEZE, GEO. The results indicated that GEO was by far the least important input in the CNN, followed by FREEZE, and then DISP and SUPCL which were separated by small margins. This indicates that the latter three inputs all had unique information with minimal redundancy between them, information which the CNN needed in order to have good performance. In comparison, GEO shared redundancy with some or all of the other three parameters, and therefore its activation strength was not as high. Since the primary installed support class was determined based on the geotechnical zone observed in the tunnel, it is intuitive that these two inputs had redundancy in terms of the signals they convey to the CNN. In practical terms, these findings indicate the ground freezing regime implemented and primary installed support class are strong determinants of tunnel liner yield, which is observed in the radial tunnel displacement over time.

5.6.2 Input Omission (IO)

The results of the IO approach are shown in Figure 5-8. Each column of subplots represents a different GMP being predicted, with an increasing amount of training data in each row, i.e., Figure 5-8b in the second column shows GMP 3 being predicted using GMP 2, while Figure 5-8e shows GMP 3 being predicted using GMPs 1 and 2. Within each subplot, the results for both the Targeted Class 2/3 model and the Global Balanced model are shown. Four boxplots are shown, where each represents an input that is being omitted

– radial tunnel displacement (DISP), ground freezing (FREEZE), geology class (GEO), and primary support class (SUPCL). For example, in Figure 5-8a, the fifth boxplot from the left shows the results for omitting DISP from the Global Balanced model when GMP 2 is being predicted using GMP 1. Each box in the figures represents an ensemble of 30 models, and the whiskers represent the 25^{th} and 75^{th} percentiles for each ensemble. A total of 2,400 model runs are included in the IO analysis.

When examining the results, it is important to remember that the training goals between the Targeted Class 2/3 models and the Global Balanced models are different – the former has emphasis on only the two higher tunnel yield classes during training due to the sigmoid error weighting scheme, while the latter is fitting a solution to all four yield classes using an inverse frequency error weighting scheme. The IO results indicate that the performance of Targeted Class 2/3 models is generally unaffected by the IO approach, as evidenced by the insignificant change to the AIC_c when subsequent inputs were removed. This confirms the strong correlation between all inputs and higher tunnel yield Class 2 and 3, which are given preferential treatment during the training of the Targeted models using the sigmoid error weighting scheme. The notable exceptions are Figure 5-8d, g, and j, where the variance when GEO is omitted is higher than when other inputs are omitted, although the mean performance when GEO is omitted is comparable to omitting the other inputs. The performance of the Global Balanced models is much more sensitive to the IO approach. For the Global Balanced models, when GEO is omitted, the models perform worst and have the largest variance as compared to omitting the other inputs, indicating that GEO contains information that is required for overall performance and precision. An example of this is Figure 5-8i, where AIC_c = 19.2 when GEO is omitted, and AIC_c = 19.0, 18.9, and 19.1 when FREEZE, SUPCL, and DISP are omitted, respectively. With a few exceptions (Figure 5-8e, h, j), removing DISP generally has the lowest impact on the performance, likely because this input is not as strongly correlated to all tunnel yield classes. The mean AIC_c values for all ensembles computed for the IO approach are reported in Table 5-5 in the Appendix.

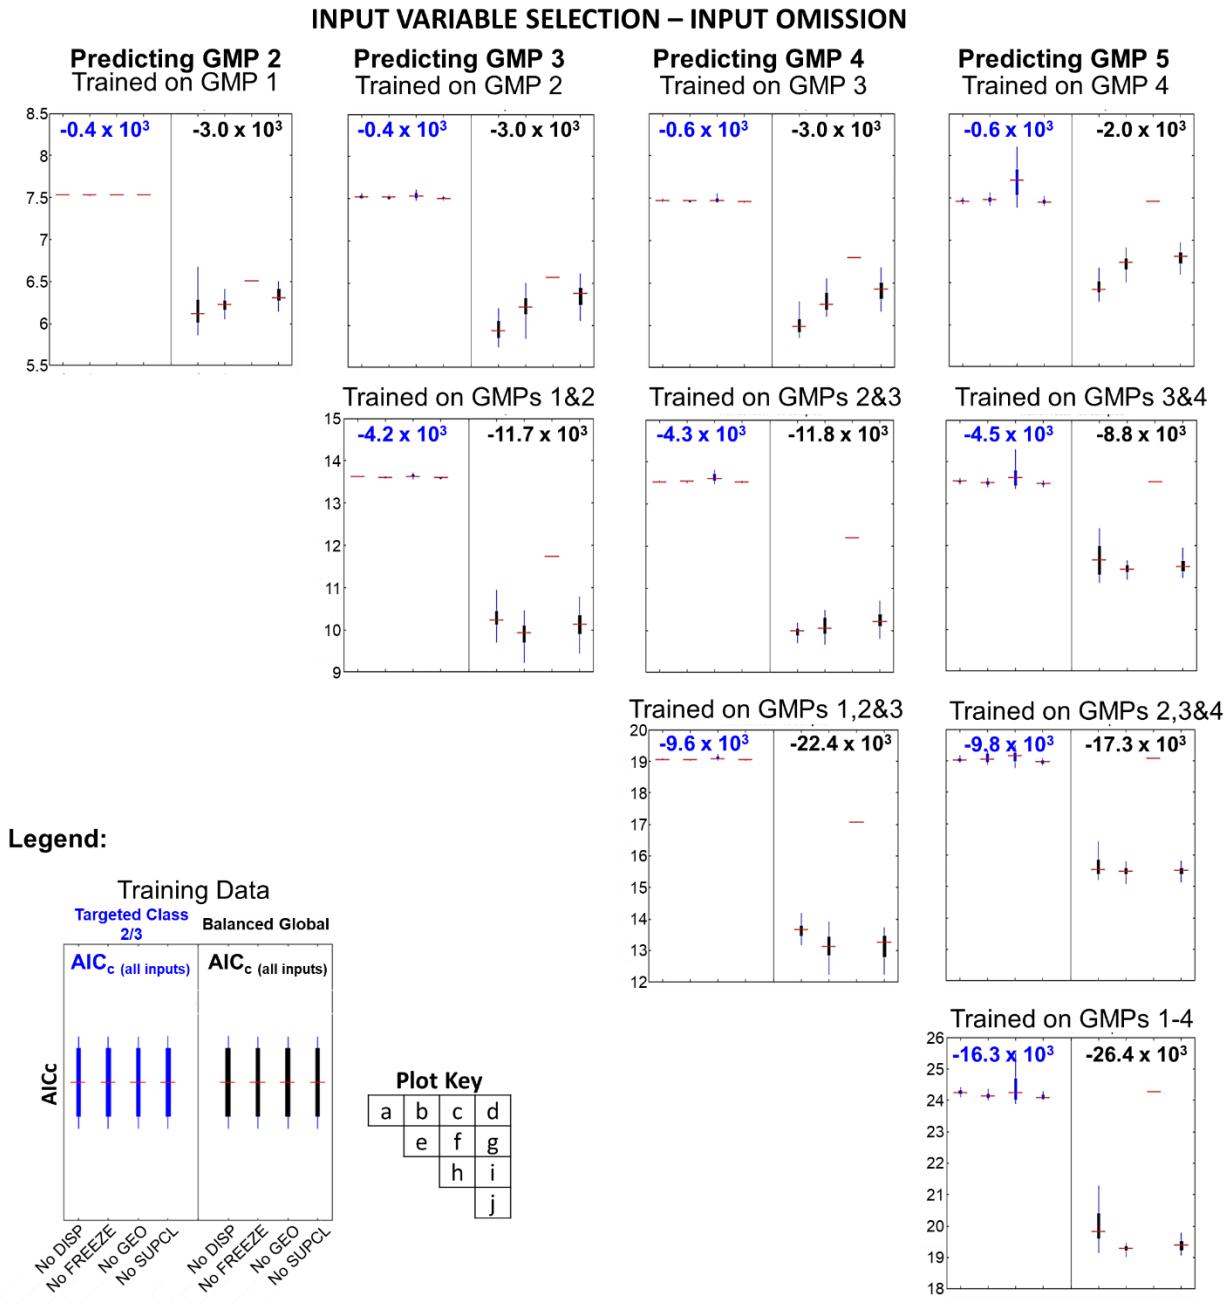

Figure 5-8. Results of the Input Omission (IO) Input Variable Selection (IVS) approach for Cigar Lake Mine. The model performance, AIC_c, is shown for the Targeted Class 2/3 and Balanced Global models for each permutation of training and testing Ground Management Plans (GMPs). Each boxplot represents an ensemble of 30 models, and the performance of the model with all inputs is shown a the top of the plots for reference.

Based on only the IO results, the Omission Ranking of the CNN inputs was GEO, DISP, FREEZE, SUPCL for the Targeted Class 2/3 model, and GEO, DISP, SUPCL, FREEZE for the Global Balanced Model (Table 5-2). This ranking indicates which inputs have the most negative impact on the overall CNN performance when they are removed entirely. Note that this ranking differs from the ranking produced by the CAS

approach indicating that none of the candidate inputs can be removed entirely, as the CNN learns the inter-dependencies between them during training in order to arrive at a YIELD prediction. This further implies that each of the four inputs is providing unique information to the CNN and therefore cannot be removed without decreasing overall performance. The difference in Omission Ranking between the two model architectures also indicates that SUPCL contains slightly less important data than FREEZE when choosing the Targeted Class 2/3 model to focus on areas of high tunnel liner yield.

Table 5-2. Omission Rankings for both the Targeted Class 2/3 and Global Balanced models, produced by the Input Omission (IO).

	IVS Method & Data Split		INPUTS															
			GMP 1				GMP 2				GMP 3				GMP 4			
	Train	Test	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP
OMISSION RANKING	Targeted Class 2/3 Model																	
	GMP 1	GMP 2	1	4	1	3												
	GMP 2	GMP 3					1	3	4	2								
	GMP 3	GMP 4									1	3	4	2				
	GMP 4	GMP 5													1	2	4	3
	GMP 1&2	GMP 3	1	3	4	2	1	3	4	2								
	GMP 2&3	GMP 4					1	2	4	3	1	2	4	3				
	GMP 3&4	GMP 5									1	3	4	2	1	3	4	2
	GMP 1-3	GMP 4	1	3	4	2	1	3	4	2	1	3	4	2				
	GMP 2-4	GMP 5					1	2	4	3	1	2	4	3	1	2	4	3
	GMP 1-4	GMP 5	1	3	4	2	1	3	4	2	1	3	4	2	1	3	4	2
	Global Balanced Model																	
	GMP 1	GMP 2	1	3	2	4												
	GMP 2	GMP 3					1	3	2	4								
	GMP 3	GMP 4									1	3	2	4				
	GMP 4	GMP 5													1	3	2	4
	GMP 1&2	GMP 3	1	4	3	2	1	4	3	2								
	GMP 2&3	GMP 4					1	3	2	4	1	3	2	4				
	GMP 3&4	GMP 5									1	4	3	2	1	4	3	2
	GMP 1-3	GMP 4	1	4	3	2	1	4	3	2	1	4	3	2				
	GMP 2-4	GMP 5					1	4	3	2	1	4	3	2	1	4	3	2
	GMP 1-4	GMP 5	1	4	3	2	1	4	3	2	1	4	3	2	1	4	3	2

The IO approach was the only IVS method investigated herein that iteratively omitted each individual input to determine its impact on performance, with a total of 2,400 models run. The goal was to determine whether the CNN was unaffected by the absence of one of the inputs, in which case that input could be removed from the input candidate pool. This IVS approach also produced an Omission Ranking, which revealed the relative importance of each input's presence in the training dataset for overall model performance. The Omission Ranking can assist in determining whether any of the data being collected is redundant or unnecessary, thus streamlining the data acquisition process and potentially resulting in cost savings. The Omission Ranking for the Targeted Class 2/3 model is GEO, DISP, FREEZE, SUPCL. The IO study had less impact on the performance of the Targeted Class 2/3 models as compared to the Global Balanced models, likely because they were less sensitive due to being trained to emphasize performance of Class 2 and 3 tunnel liner yield only and there being a higher correlation between those classes and the inputs (Morgenroth et al., 2021). For the Global Balanced models, the Omission Ranking was GEO, DISP, SUPCL, FREEZE. This ranking showed that omitting GEO had the highest negative impact on the model performance (i.e., the highest AIC_c), indicating that while it may have the lowest activation strength in the

CNN (as found by the CAS approach), it cannot be removed as it contains crucial information the CNN requires to make an accurate prediction of tunnel liner yield. This highlights the need for more than one IVS method to determine the saliency of all the candidate inputs. The second input in the Omission Rankings for both CNN architectures was DISP, which is the finest resolution time-dependent input. DISP is the only input that measures the tunnel liner deformation directly and has a relatively low level of subjectivity associated with it. The IO approach revealed that of the four inputs used to develop the Cigar Lake Mine CNN, none of them could be omitted entirely to achieve better performance. This gives rise to considerations for future work in refining the Cigar Lake Mine CNN – perhaps more inputs should be digitized from the GMPs, and/or the time intervals between radial tunnel displacement measurements should be increased, to further improve performance.

5.6.3 Partial Correlation (PC)

In contrast to the CAS and IO methods, the Correlation Ranking of the inputs is calculated prior to running the CNN and therefore the Correlation Ranking is the same for both the Targeted Class 2/3 and Global Balanced models. Absolute values were used to determine the Correlation Ranking to prevent strongly negative correlated inputs from being selected last. The detailed rankings for each of the ten training/testing data split permutations are presented in Table 5-3. In general, the first ranked input was from the most recent GMP available for training, e.g., if GMP 3 was being predicted, the first input selected was from GMP 2 rather than GMP 1. This finding emphasizes the importance of the temporal aspect of the mechanisms driving liner yield. Where there was only one training GMP, DISP always had the highest Correlation Ranking.

The results of the PC approach are shown in Figure 5-9 for the Targeted Class 2/3 models and Figure 5-10 for the Global Balanced models, with more detailed results figures in the Appendix (Figure 5-11 and Figure 5-12). Each subplot in Figure 5-9 and Figure 5-10 represents one permutation of training and testing data, with the performance in terms of AIC_c on the y-axis and the successively added candidate inputs (in descending order of absolute partial correlation) on the x-axis. Each point in the figures represents the mean of an ensemble of 30 models, and the whiskers represent the 25^{th} and 75^{th} percentiles for each ensemble. A total of 48,000 model runs are included in the PC analysis. For example, when reading Figure 5-9a, where the Targeted Class 2/3 model is trained on GMP1 to predict GMP 2, the inputs are added in the order DISP, FREEZE, SUPCL, GEO, and the lowest AIC_c is reached after SUPCL is added to the candidate input pool. When reading Figure 5-10e, the Global Balanced model is used to predict GMP 3, which is trained on GMPs 1 and 2, and where the inputs are added in the order GMP 2 FREEZE, GMP 2 GEO, GMP 1 FREEZE, GMP 1 SUPCL, GMP 2 DISP, GMP 2 SUPCL, GMP 1 DISP, GMP 1 GEO. In this example the lowest AIC_c occurs when the first five candidates have been added, shown in red. After GMP 2 DISP has been added, the performance decreases. The mean AIC_c values for all ensembles computed for the PC approach are reported in Table 5-6 in the Appendix.

The order of the candidate inputs on the x-axis of each subplot in Figure 5-9 and Figure 5-10 was determined by the Correlation Ranking. Where there was only one training GMP (Figure 5-9a, b, c, d), DISP and FREEZE were selected first as they were the highest correlated inputs with tunnel liner yield. Where there were two training GMPs (Figure 5-9e, f, g), FREEZE and GEO for the GMP immediately preceding the GMP that was being predicted were selected first. Where there were three training GMPs (Figure 5-9h, i), SUPCL for the preceding GMP and FREEZE from two GMPs prior were selected first. Finally, Figure 5-9j where four GMPs are used for training, the GEO input from the previous two GMPs are selected first. However, since this was a static input (i.e., GEO is identical across all GMPs), the performance did not change significantly when the second GEO was added. As was established by the above, there was no general pattern for the order in which the candidate inputs were added across all permutations of training and testing GMPs, beyond the general rule that the inputs from the GMPs immediately preceding the GMP being predicted are selected first.

As with the IO approach, the PC approach has minimal impact on the Targeted Class 2/3 models, similarly because the CNN is solving for a solution (or global minimum) that only emphasizes two of the four tunnel liner yield classes. As such, the algorithm is generally less complex and nuanced relationships between all four inputs and the output do not need to be extracted to achieve a good performance. This is illustrated by the low variance in AIC_c when successive inputs are added to the inputs for the same training/testing data split, e.g., Figure 5-9e where AIC_c remains approximately -4.2e3 for every new candidate input that is added to the input pool. Generally, the Targeted Class 2/3 models require less inputs than the Global Balanced models, needing only the four candidate inputs with the highest partial correlation to achieve the minimum AIC_c performance for a given training/testing data split.

In contrast, the PC approach produces very interesting results for the Global Balanced models. In particular, when one GMP is used for training (Figure 5-10a, b, c, d), only the training GMPs DISP input is needed to produce the minimum AIC_c (AIC_c = -3e3, -3e3, -3e3, -2e3 when predicting GMPs 2, 3, 4 and 5, respectively). However, these models have a performance that is an order of magnitude lower than the models that are more data rich, specifically the minimum AIC_c is approximately -8e3 with two training GMPs (Figure 5-10e, f, g), approximately -1.5e4 with three training GMPs (Figure 5-10h, i), and approximately -2.2e4 with four training GMPs (Figure 5-10j). In the more data rich models (Figure 5-10e to j), the optimal number of candidate inputs increases to the number of training GMPs minus one image, where each image is equal to four inputs plus one input. For example, in Figure 5-10h where three GMPs are used for training (twelve possible inputs in the candidate pool), the first nine candidates (eight inputs, equal to two images, plus one additional input) produce the minimum AIC_c of -2.2e4. Similarly, the minimum AIC_c of -2.6e4 in Figure 5-10j is produced when twelve inputs, equal to three images, plus one input, for a total of thirteen candidate inputs, are used for training. This pattern is represented by a zig-zag trend in each subplot, where each subsequent input worsens the AIC_c and increases the model variance, which indicates the model is receiving redundant or confusing signals from the combinations of inputs. Then the fifth, ninth, or thirteenth

input causes a large increase in model performance. This finding is related to the architecture of the CNN, where the convolution filter convolves through one image (or four inputs) at one time. It is also noteworthy that the input that produces the minimum AIC_c in all cases is the most recent DISP, which aligns with the rock mechanics context in the sense that the most recent deformation is the most important predictor of the tunnel liner yield, of the available candidate inputs.

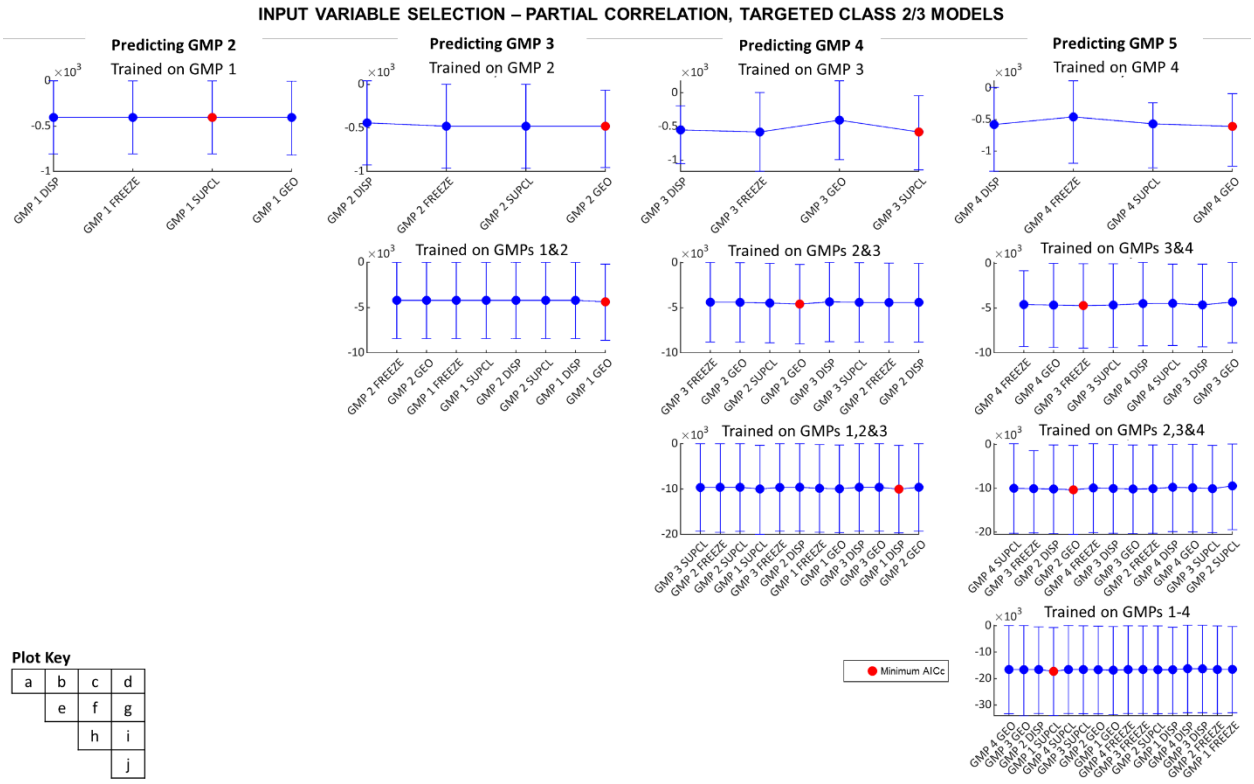

Figure 5-9. Results of Partial Correlation (PC) Input Variable Selection (IVS) approach for Cigar Lake Mine Targeted Class 2/3 models. Each plot shows the mean performance across an ensemble of 30 models of the Convolutional Neural Network, with the whiskers representing the 25th and 75th percentiles, as each successive input is added,

where the order is determined by the partial correlation of the candidate inputs and the target. The minimum AICc for each permutation is highlighted in red.

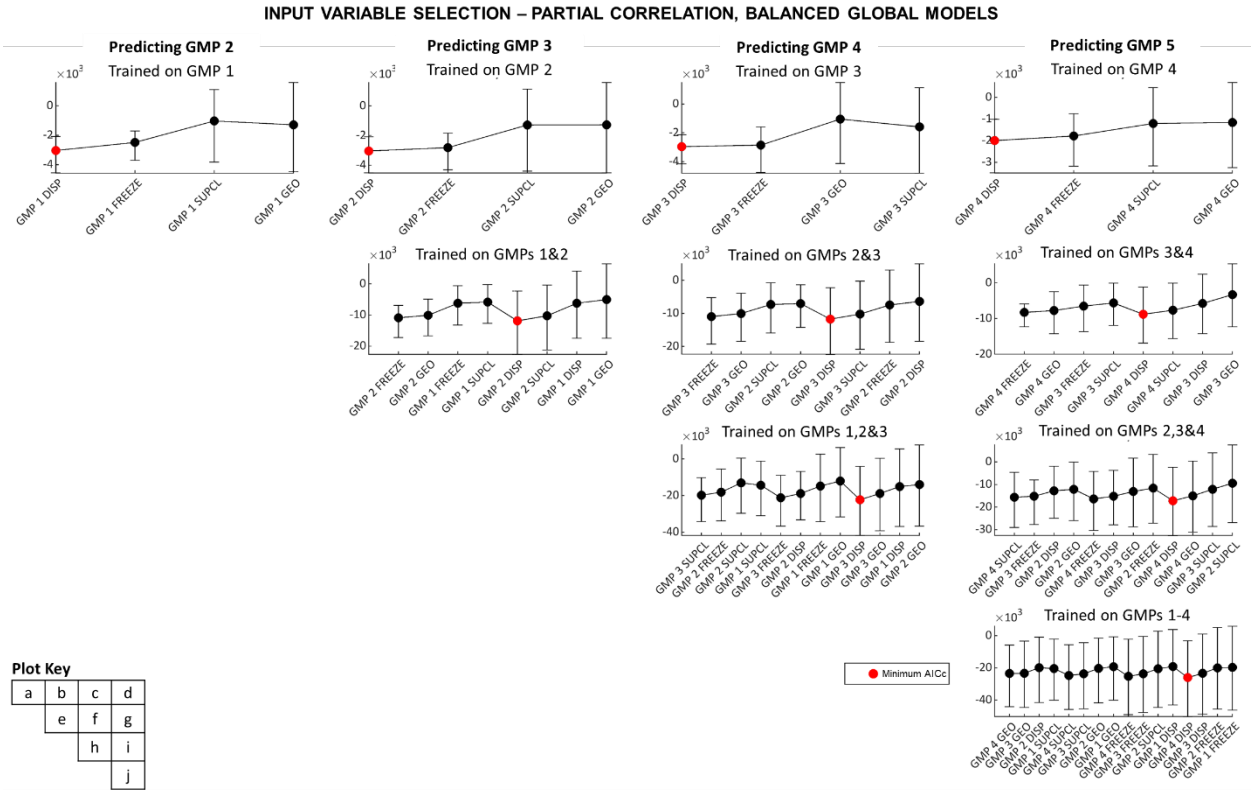

Figure 5-10. Results of Partial Correlation (PC) Input Variable Selection (IVS) approach for Cigar Lake Mine Global Balanced models. Each plot shows the mean performance across an ensemble of 30 models of the Convolutional Neural Network, with the whiskers representing the 25th and 75th percentiles, as each successive input is added, where the order is determined by the partial correlation of the candidate inputs and the target. The minimum AICc for each permutation is highlighted in red.

The PC approach was the most computationally expensive IVS approach investigated in this study, requiring 48,000 CNN models run. Despite the time and computational expense, this IVS approach revealed useful information with respect to optimizing the Cigar Lake Mine CNN and developing it further. The order in which the candidate inputs are selected for input into the CNN was based on the Correlation Ranking, which was determined by the partial correlation of each input with the tunnel liner yield (Table 5-3). The Correlation Ranking varied depending on the training/testing data split, however an input from the most recent training GMP always had the highest ranking. Where there was only one training GMP, the DISP input was selected first and also produced the best model results as compared to adding the subsequent inputs. This indicates that when there is limited data in the temporal realm, DISP is the best available predictand. However, it is worth noting that models with more training GMPs, i.e., more data in the temporal realm, resulted in an order of magnitude better performance. Where there is more than one GMP used for

training, the categorical inputs (GEO, SUPCL, FREEZE) were chosen first across all training and testing data permutations, as they had higher partial correlations as compared to the numerical input (DISP).

Table 5-3. Correlation Rankings produced by the Partial Correlation (PC) method. Note that ranking is identical for both the Targeted Class 2/3 and Global Balanced models as it is based on the partial correlation calculated and is independent of model performance.

			INPUTS															
IVS Method & Data Split			GMP 1				GMP 2				GMP 3				GMP 4			
	Train	Test	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP	GEO	FREEZE	SUPCL	DISP
CORRELATION RANKING	GMP 1	GMP 2	4	2	3	1												
	GMP 2	GMP 3					4	2	3	1								
	GMP 3	GMP 4									3	2	4	1				
	GMP 4	GMP 5													4	2	3	1
	GMP 1&2	GMP 3	8	3	4	7	2	1	6	5								
	GMP 2&3	GMP 4					4	7	3	8	2	1	6	5				
	GMP 3&4	GMP 5									8	3	4	7	2	1	6	5
	GMP 1-3	GMP 4	8	7	4	11	12	2	3	6	10	5	1	9				
	GMP 2-4	GMP 5					4	8	12	3	7	2	11	6	10	5	1	9
	GMP 1-4	GMP 5	8	16	4	12	7	15	11	3	2	10	6	14	1	9	5	13

Similar to the IO approach, PC had little impact on the Targeted Class 2/3 models, likely for comparable reasons stemming from how the model was trained to minimize error between less tunnel yield class, and therefore was less complex and sensitive. However, for the Global Balanced models, the PC approach revealed some interesting patterns with respect to how the CNN performance responded to using only a subset of the candidate inputs, specifically the most highly correlated ones. As a general rule, where there was only one GMP to train from, the minimum AIC_c was obtained when only the DISP from the previous GMP is used. If more than one GMP was available for training, the minimum AIC_c was obtained when the number of candidate inputs is equal to the number of training GMPs minus one, plus one channel, e.g., if there were three available training GMPs (3 GMPs x 4 channels = 12 available inputs), the minimum AIC_c was achieved with two GMPs (2 x 4 = 8 inputs) plus 1 channel, for a total of 9 inputs. In all training/testing data splits explored herein, this occurred when the most recent DISP was added to the candidate input pool. For example, when training on GMPs 1, 2 and 3 and predicting GMP 4, the optimal number of inputs is nine where the ninth input was DISP from GMP 3. It is difficult to discern exactly why this pattern is pervasive, although it likely has to do with the fact that the original input images have four channels each, and therefore the convolution filter has a depth of four. From a rock mass deformation mechanics perspective, it is also intuitive that the most recent DISP should be included in the model to obtain the minimum AIC_c. Perhaps the most important finding of the PC approach, however, is that fact that no matter how many training GMPs are used, it is not possible to drop out the older ones completely, as some of the inputs from those older GMPs always fall within the set of candidate inputs necessary to produce the minimum AIC_c. In practical terms, this indicates that when training CNNs for prediction of behaviour of underground excavations, as much historical data as possible should be used to train the algorithm, and only after an IVS investigation should the decision be made whether to drop any training data.

5.6.4 General Discussion

This study investigated three IVS methods to determine the relative significance of the four inputs in determining tunnel liner yield using the Cigar Lake Mine CNN, and specifically how the inputs differ between the Targeted Class 2/3 models and the Global Balanced models. Each of the three IVS methods give a different perspective on the inputs and how they are used to make the tunnel yield class prediction, thereby increasing the interpretability of the CNN. Interpretability is defined as an understanding of how the model makes decisions based on its inputs and the learned components such as hyperparameters and architecture (Molnar, 2022).

In comparing the three IVS methods, some patterns emerged that were consistent across all three methods and aligned with common knowledge in rock engineering. Of the available candidate inputs, none could be removed entirely. The Activation Rankings indicated that GEO had redundancy with the other inputs, which is logical considering SUPCL is determined based on the geotechnical zones. The Omission Rankings indicated that while there is redundancy, GEO has unique information that must be included for accurate model performance. This also appeals to the rock engineering experience that geology most often governs rock mass deformation. The Correlation Rankings showed that the most recent GMP had that highest correlation with the GMP being predicted, and that the most recent DISP input yielded that best overall model performance. This aligns with the reasoning that the most recent data on the rock mass deformation trajectory must be used to train the CNN to achieve good performance. All three IVS methods found that the more data that was used for training the CNN in the temporal realm, the more accurate the predictions became.

While the rankings produced by the model-based methods (CAS, IO) and the model-free method (PC) cannot be compared directly due to the inherent differences in how they are calculated, the results of this IVS study illustrate the need for multiple IVS methods before deciding to remove or add inputs to a CNN model. For example, an argument could have been made to remove GEO based on only the CAS results and Activation Rankings. However, IO and the Omission Rankings showed that GEO appears to be the common thread between the other three inputs, despite having a low activation strength on its own. PC and the Correlation Rankings revealed further complementary information about the importance of the temporal nature of the inputs and having recent GMPs to obtain higher prediction accuracy, as compared to IO which looked only at removing an input in its entirety.

An important point of discussion when considering all three IVS methods is that they collectively indicate the need for continuously adding the most recent data in agreeance with the idea that the updating of data-driven algorithms as new data becomes available is crucial (Elmo et al., 2020). In this research, the addition of new data is simulated by adding each subsequent GMP and predicting the next one. This essentially retrains the CNN in order to make the best possible use of any new signals that may be contained in the most recent data. The retraining and new learning is paramount in developing algorithms for rock engineering, as rock deformation mechanics cannot be adequately defined by a static snapshot in time. In

fact, this study finds that the more information the algorithm has on the trajectory of the historical deformation, the more accurate the predictions become. This is illustrated in the CAS method (Figure 5-7, bottom left subplot), the IO method (Figure 5-8d, g, i, j), and the PC method (Figure 5-10d, g, i, j), where consecutively increasing performances for predicting GMP 5 as another preceding GMP is added into the training dataset. This is a powerful conclusion, as conventional numerical modelling methods do not lend themselves to exploring every time step along the deformation path in detail, while CNNs are able to compute large volumes of data efficiently. The implications of this are that no data needs to be thrown out due to the time constraints associated with processing it, and that a more complete understanding of nuanced rock mass deformation mechanics can be gleaned from the available data. Thus, uncertainty between subsequent stages of engineering design may have less bias injected and reduced uncertainty associated with them.

An important aim of completing an IVS analysis of a given algorithm is to increase its interpretability, which broadly falls into two categories: model transparency and post-hoc explanations (Lipton, 2016). Previous work to complete hyperparameter tuning of the Cigar Lake Mine CNN falls into the former category, where “transparency” indicates an understanding of how the model works. In the previous research, parameters such as the convolution filter size, application of an error weighting scheme to rebalance the training data, and the amount of required training data were analyzed to this end (Morgenroth et al., 2021). In the current research, the CAS method also falls under algorithm transparency as it is concerned with extracting information on how the trained algorithm uses each available input (GEO, SUPCL, FREEZE, DISP) without altering the algorithm architecture. Post-hoc explanations do not generally indicate how the model works, but nonetheless offer useful interpretations for the end user of the algorithm. Common post-hoc interpretation methods include visual representations of learned features, and explanations by example. The IO and PC methods applied herein are post-hoc explanations, as the aim is to uncover how the inputs influence each other within the CNN by running additional models with modifications. The “explanation by example” in this case is the explanation of the influence removing certain inputs has, or rearranging the order in which they are added to the algorithm.

The Cigar Lake Mine CNN required customized IVS methods to be developed to interpret its results and performance due to the format of the input data. The original digitization scheme to convert the GMP data into images was chosen to mimic the format of the original tunnel mapping done by ground control engineers at Cigar Lake Mine (Morgenroth et al., 2021). The choice was made not to deconstruct the GMP format as not to introduce interpolation into the training dataset. Thus, the IO and PC IVS approaches were developed from their conventional application to ANNs for the Cigar Lake Mine CNN. This allows for a more intuitive understanding of how the CNN processes the tunnel mapping in its original form, without interpolated or augmented data. For example, omitting the DISP channel in the proposed IO method simulates a case where the tunnel convergence has not yet been surveyed and the mine’s ground control want to use the CNN to obtain an indication of liner performance. The findings from the IO and PC

approaches are transferable to other mine and tunnel projects, where the ground control or geotechnical engineers do not have all the data available to them that was used to train the CNN, or want to gain an understanding of whether all the data that is being collected is required to obtain a useful prediction from the CNN. In combination, CAS, IO, and PC enhance CNN interpretability by increasing model transparency and providing post-hoc explanations (Lipton, 2016).

5.7 Conclusions

Recent research at the intersection of rock mechanics and ML have provoked a sense of "accuracy and infallibility" when it comes to applying deep learning to rock engineering (Elmo et al., 2020), which may be attributed to the perception that MLAs are magic "black boxes" capable of predicting any phenomena given relevant inputs. However, it is important to continue to question and analyze the internal mechanics of these algorithms, including their architecture, hyperparameters and the underlying assumptions they are developed with. One way to probe the findings of ML is through applying an IVS approach to rank the candidate inputs, determine their usefulness to the algorithm, reduce redundancy within the model, and compare the aforementioned findings with the experience and knowledge of fundamental rock mechanics principles.

This research compared three IVS methods for two CNN architectures developed to predict tunnel liner yield at the Cigar Lake Mine, the Targeted Class 2/3 model and the Global Balanced model. For both models the three IVS approaches employed were Channel Activation Strength, Input Omission, and Partial Correlation. The latter of these two are novel approaches that have been developed for CNNs dealing with unconventional spatial and temporal image channels. The three IVS methods were used to produce Activation Rankings, Omission Rankings, and Correlation Rankings, respectively. The rankings presented in this chapter are parts of a whole interpretation of the Cigar Lake Mine CNN, rather than each being a stand-alone analysis.

All three IVS methods found that increasing data in the temporal realm increased the prediction accuracy of the models. The CAS approach and Activation Rankings found that the GEO input has the lowest activation strength across all models tested, while SUPCL has the highest. This suggests that GEO is used the least by the CNN as compare to the other three inputs, SUPCL, FREEZE, and DISP. CAS indicated that SUPCL has the high activation strength in the CNN. The IO approach was used to determine if any of the inputs could be omitted from the model entirely, however the results indicate that this is not the case for the Cigar Lake Mine CNN. The IO results and Omission Rankings for the Global Balanced models showed that removing GEO entirely caused the model performance to worsen as compared to the other inputs, indicating that GEO had some unique information that the other inputs were not able to convey to the CNN. Finally, the PC approach and Correlation Rankings were used to determine if the models could be optimized by using only the most highly correlated inputs to train the CNN. Where there was only one training GMP the DISP input was selected first, meanwhile where there was more than one GMP used for training the categorical inputs (GEO, SUPCL, FREEZE) were chosen first. The PC results for the Global

Balanced models showed that the minimum AIC_c was obtained when the number of candidate inputs was equal to the number of training GMPs minus one, plus one input channel. The minimum AIC_c was always obtained when the most recent DISP was added to the candidate pool. The PC results also indicated that no matter how many training GMPs were used, it was not possible to drop out the older ones completely.

An important contribution of this study from all three IVS approaches was that none of the inputs could be removed completely – i.e., all available data in the temporal realm had to be included to produce the best possible model performance, for both the Targeted Class 2/3 and the Global Balance model architectures. This finding was significant because it indicates that it is crucial to include all data in the initial development of a CNN for a rock engineering problem, and then only after completing an IVS study remove the inputs that are not contributing to good model performance. Since ML is efficient at processing large volumes of data, this is an advantage over conventional numerical modelling where exploring all temporal data is prohibitive.

This research demonstrates that no single IVS method should be used to determine input saliency. Instead, it should become standard practice that multiple IVS methods are applied to increase algorithm interpretability. This work finds that CNNs are a useful approach for characterizing rock mass deformation when investigating thoroughly using IVS approaches and may be used to complete temporal data analysis ahead of or in addition to more conventional modelling and design efforts.

Appendix

Table 5-4. Summary of mean AIC_c for each ensemble of models computed for the Channel Activation Strength (CAS) IVS method. Colour scale applied to each column shows higher AIC_c (yellow) to lower AIC_c (green), where a lower AICc indicates higher performance.

CHANNEL ACTIVATION STRENGTH (CAS)			
Data Split		AICc	
Train	Test	Targeted Class 2/3 Model	Global Balanced Model
GMP 1	GMP 2	-402.962	-3025.295
GMP 2	GMP 3	-442.238	-3153.048
GMP 3	GMP 4	-552.298	-2981.948
GMP 4	GMP 5	-579.346	-2001.418
GMP 1&2	GMP 3	-4200.894	-11727.397
GMP 2&3	GMP 4	-4352.077	-11787.607
GMP 3&4	GMP 5	-4487.933	-8837.556
GMP 1-3	GMP 4	-9633.779	-22408.528
GMP 2-4	GMP 5	-9768.008	-17323.756
GMP 1-4	GMP 5	-16278.276	-26360.995

Table 5-5. Summary of mean AIC_c for each ensemble of models computed for the Input Omission (IO) IVS method. Colour scale applied to each row for the two model types shows higher AIC_c (yellow) to lower AIC_c (green), where a lower AICc indicates higher performance.

Data Split		INPUT OMISSION (IO)							
		Omit GEO	Omit FREEZE	Omit SUPCL	Omit DISP	Omit GEO	Omit FREEZE	Omit SUPCL	Omit DISP
Train	Test	Targeted Class 2/3 Model				Global Balanced Model			
GMP 1	GMP 2	7.539	7.537	7.536	7.538	6.525	6.238	6.331	6.182
GMP 2	GMP 3	7.549	7.517	7.511	7.526	6.575	6.218	6.368	5.968
GMP 3	GMP 4	7.483	7.467	7.459	7.472	6.802	6.281	6.431	6.031
GMP 4	GMP 5	7.716	7.487	7.463	7.469	7.463	6.730	6.819	6.450
GMP 1&2	GMP 3	13.662	13.607	13.604	13.619	11.748	9.891	10.129	10.321
GMP 2&3	GMP 4	13.631	13.547	13.528	13.537	12.203	10.098	10.227	9.971
GMP 3&4	GMP 5	13.787	13.524	13.483	13.545	13.534	11.461	11.514	11.666
GMP 1-3	GMP 4	19.135	19.054	19.045	19.067	17.081	13.109	13.119	13.742
GMP 2-4	GMP 5	19.265	19.083	18.986	19.054	19.080	15.494	15.549	15.781
GMP 1-4	GMP 5	24.548	24.168	24.150	24.263	24.276	19.369	19.431	19.997

Table 5-6. Summary of mean AIC_c for each ensemble of models computed for the Partial Correlation (PC) IVS method. Note that the candidate inputs were added in the order indicated in . Colour scale applied to each row shows higher AIC_c (yellow) to lower AIC_c (green), where a lower AICc indicates higher performance.

Data Split		PARTIAL CORRELATION (PC)															
		GMP 1				GMP 2				GMP 3				GMP 4			
Train	Test	Add GEO	Add FREEZE	Add SUPCL	Add DISP	Add GEO	Add FREEZE	Add SUPCL	Add DISP	Add GEO	Add FREEZE	Add SUPCL	Add DISP	Add GEO	Add FREEZE	Add SUPCL	Add DISP
Targeted Class 2/3 Model																	
GMP 1	GMP 2	-4.037E+2	-4.037E+2	-4.035E+2	-4.030E+2												
GMP 2	GMP 3					-4.810E+2	-4.800E+2	-4.794E+2	-4.422E+2								
GMP 3	GMP 4									-4.066E+2	-5.820E+2	-5.817E+2	-5.523E+2				
GMP 4	GMP 5													-6.079E+2	-5.668E+2	-4.565E+2	-5.793E+2
GMP 1&2	GMP 3	-4.374E+3	-4.199E+3	-4.199E+3	-4.198E+3	-4.197E+3	-4.196E+3	-4.201E+3	-4.201E+3								
GMP 2&3	GMP 4					-4.597E+3	-4.472E+3	-4.412E+3	-4.402E+3	-4.401E+3	-4.402E+3	-4.379E+3	-4.352E+3				
GMP 3&4	GMP 5									-4.314E+3	-4.656E+3	-4.715E+3	-4.638E+3	-4.646E+3	-4.466E+3	-4.602E+3	-4.488E+3
GMP 1-3	GMP 4	-9.943E+3	-9.982E+3	-9.843E+3	-1.001E+4	-9.618E+3	-9.621E+3	-9.632E+3	-9.633E+3	-9.628E+3	-9.642E+3	-9.644E+3	-9.634E+3				
GMP 2-4	GMP 5					-1.032E+4	-9.452E+3	-1.008E+4	-1.015E+4	-1.016E+4	-1.007E+4	-1.008E+4	-1.004E+4	-9.873E+3	-1.001E+4	-9.923E+3	-9.768E+3
GMP 1-4	GMP 5	-1.686E+4	-1.725E+4	-1.651E+4	-1.663E+4	-1.665E+4	-1.665E+4	-1.655E+4	-1.655E+4	-1.663E+4	-1.659E+4	-1.659E+4	-1.630E+4	-1.657E+4	-1.654E+4	-1.656E+4	-1.628E+4
Global Balanced Model																	
GMP 1	GMP 2	-1.289E+3	-1.030E+3	-2.491E+3	-3.037E+3												
GMP 2	GMP 3					-1.275E+3	-1.284E+3	-2.817E+3	-3.043E+3								
GMP 3	GMP 4									-1.048E+3	-1.591E+3	-2.852E+3	-2.955E+3				
GMP 4	GMP 5													-1.159E+3	-1.206E+3	-1.786E+3	-1.998E+3
GMP 1&2	GMP 3	-5.046E+3	-5.851E+3	-6.217E+3	-6.217E+3	-1.009E+4	-1.026E+4	-1.090E+4	-1.190E+4								
GMP 2&3	GMP 4					-7.034E+3	-7.320E+3	-7.464E+3	-6.406E+3	-1.005E+4	-1.023E+4	-1.096E+4	-1.174E+4				
GMP 3&4	GMP 5									-3.318E+3	-5.652E+3	-6.520E+3	-5.786E+3	-7.776E+3	-7.679E+3	-8.299E+3	-8.818E+3
GMP 1-3	GMP 4	-1.202E+4	-1.436E+4	-1.484E+4	-1.505E+4	-1.401E+4	-1.310E+4	-1.812E+4	-1.890E+4	-1.889E+4	-1.984E+4	-2.123E+4	-2.228E+4				
GMP 2-4	GMP 5					-1.210E+4	-9.419E+3	-1.157E+4	-1.282E+4	-1.314E+4	-1.210E+4	-1.520E+4	-1.523E+4	-1.504E+4	-1.567E+4	-1.636E+4	-1.721E+4
GMP 1-4	GMP 5	-1.928E+4	-2.040E+4	-1.974E+4	-1.921E+4	-2.029E+4	-2.052E+4	-1.999E+4	-1.989E+4	-2.337E+4	-2.373E+4	-2.374E+4	-2.342E+4	-2.348E+4	-2.476E+4	-2.528E+4	-2.604E+4

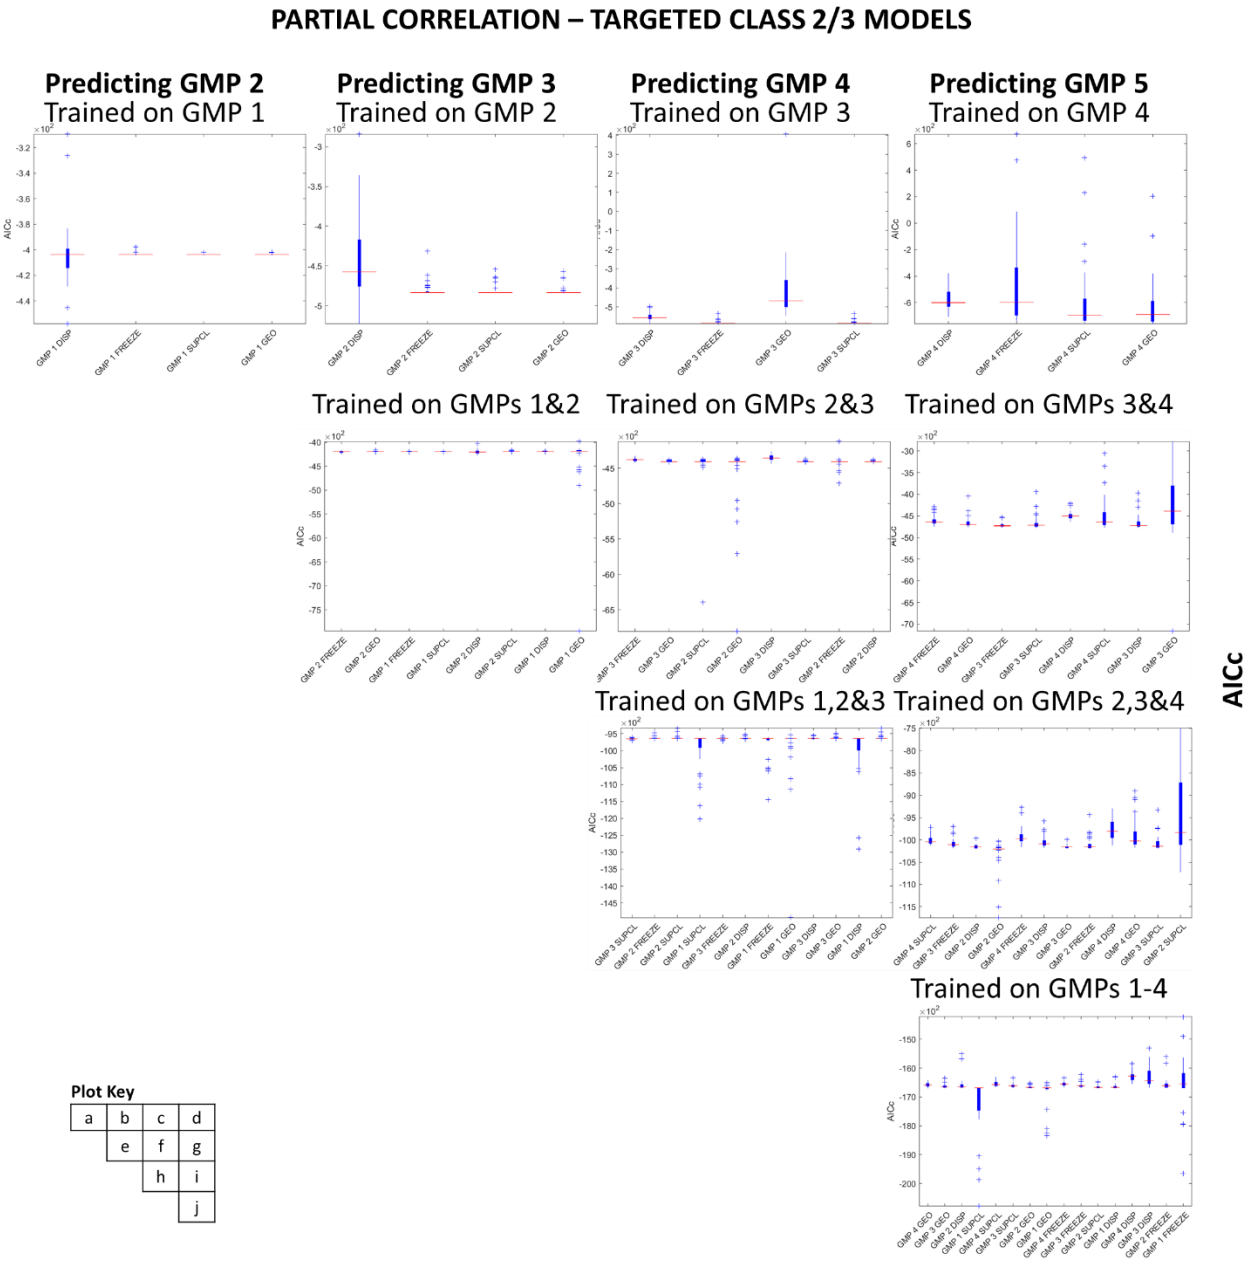

Figure 5-11. Results of Partial Correlation (PC) Input Variable Selection (IVS) approach for Cigar Lake Mine Targeted Class 2/3 models. Each plot shows the boxplot of performance across an ensemble of 30 models of the Convolutional

Neural Network as each successive input is added, where the order is determined by the partial correlation of the candidate inputs and the target.

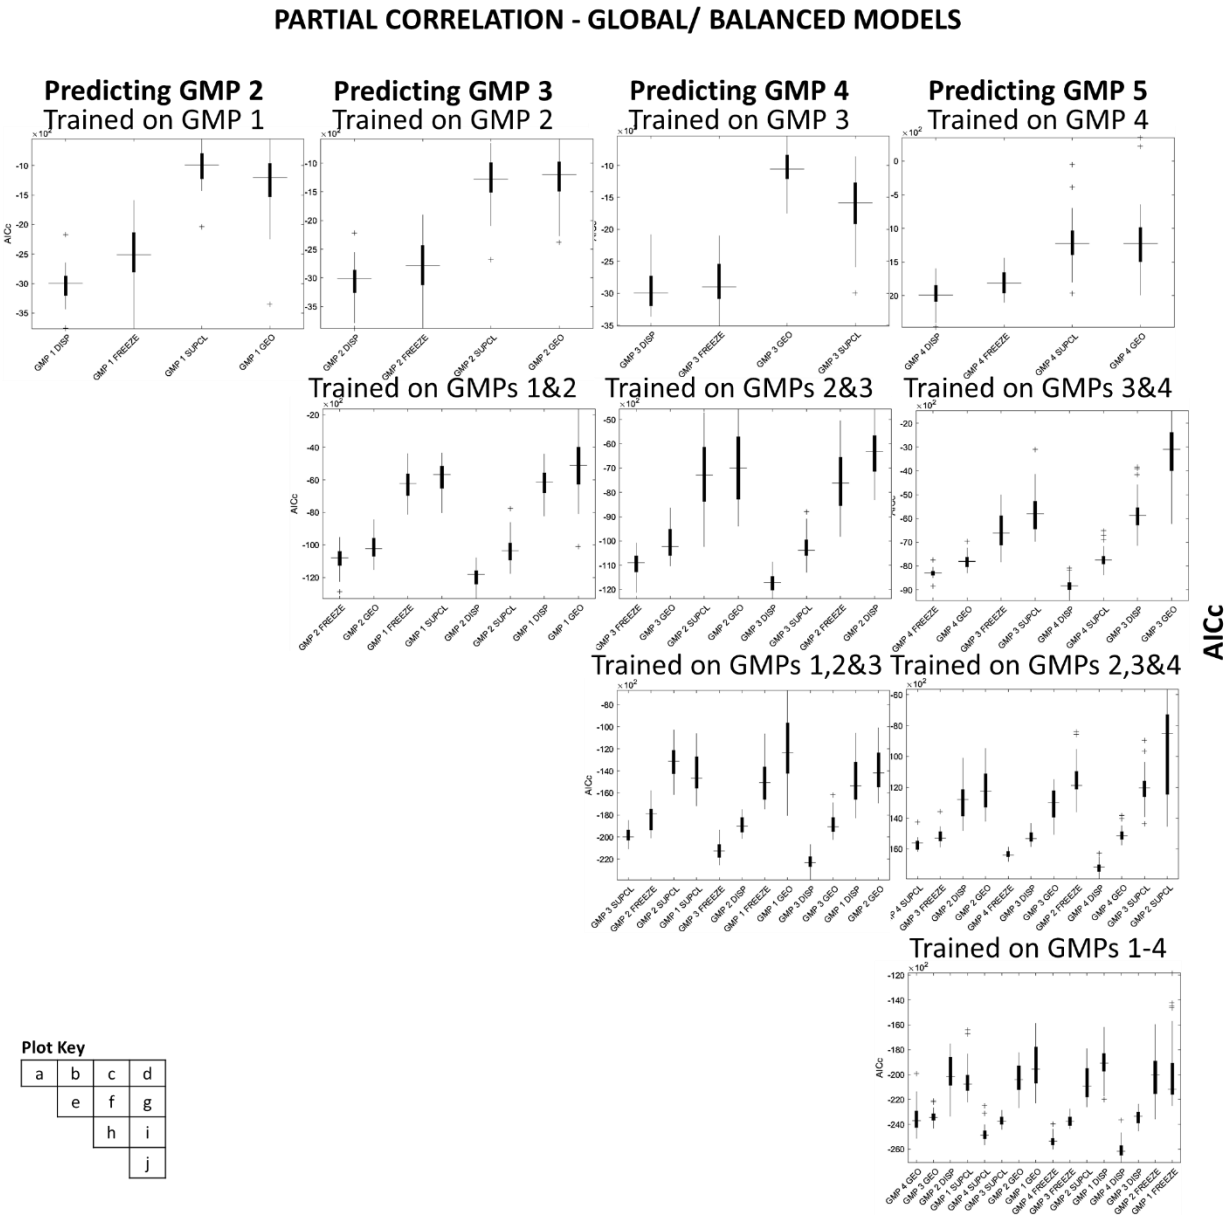

Figure 5-12. Results of Partial Correlation (PC) Input Variable Selection (IVS) approach for Cigar Lake Mine Global Balanced models. Each plot shows the boxplot of performance across an ensemble of 30 models of the Convolutional Neural Network as each successive input is added, where the order is determined by the partial correlation of the candidate inputs and the target.

CHAPTER 6. A NOVEL LONG-SHORT TERM MEMORY NETWORK APPROACH TO THE RECALIBRATION OF A FINITE DIFFERENCE MODEL FOR HIGH STRESS MINE EXCAVATIONS

6.1 Preface

This chapter focuses on the development of a Long-Short Term Memory (LSTM) network to assist in the calibration of a FLAC3D model, using the Garson Mine dataset as a case study. The training data consists of the microseismic database, geology, and geomechanical parameters from the FLAC3D model. Two LSTM networks are developed: (1) predicting principal stresses in the FLAC3D model, and (2) predicting the six-component stress tensors in the FLAC3D model. Hyperparameters optimized include input encoding and pre-processing methods, algorithm training solver, network layer architecture, and cost function. This research represents progress towards continuous, automated calibration of numerical models to allow for more accurate forecasts of changes in stress conditions.

The content of this chapter was submitted to a *Georisk* special issue entitled *Machine Learning and AI in Geotechnics* to be published in 2023 as follows:

Morgenroth, J., Kalenchuk, K., Moreau-Verlann, L., Perras, M. A., & Khan, U. T. (Under Review 2022). A novel Long-Short Term Memory network approach to the calibration of a finite difference model for high stress mine excavations. *Georisk – Machine Learning and AI in Geotechnics*. Submission ID NGRK-2022-0079.

The contributions of the authors in the current chapter are as follows:

Josephine Morgenroth has conducted the literature review, developed the LSTM network using the required software, validated and visualized the results, and prepared and wrote the original manuscript of this publication. **Matthew A. Perras** has supervised the research, provided the funding, and contributed to writing and editing the manuscript. **Usman T. Khan** has supervised the research, provided the funding, and contributed to the writing and editing the manuscript. **Kathy Kalenchuk** and **Lindsay Moreau-Verlaan** have reviewed the manuscript's technical content and provided feedback for improvement of the LSTM network for use in applied rock engineering.

The authors would like to extend special thanks to our industry partners, Kathy Kalenchuk and Lindsay Moreau-Verlaan, for their constructive feedback and informative conversations. This work is funded in part by the Natural Sciences and Engineering Research Council of Canada through the Discovery Grant program and the Postgraduate Scholarships – Doctoral program.

6.2 Abstract

Digitalization has increased access to large amounts of data for rock engineers. Machine learning presents an opportunity to aid data interpretation. The operators of Garson Mine use a microseismic database

calibrate a mine-scale finite difference model, which is used to assess seismic risk to inform mine operations. A Long-Short Term Memory (LSTM) network is proposed for numerical model recalibration. The model is trained using microseismic data, geology, and geomechanical parameters from the FLAC3D model. Two LSTM networks are developed for Garson Mine: (1) predicting far field principal stresses in the FLAC3D model, and (2) predicting the far field six-component stress tensors in the model. Various LSTM network hyperparameters were analyzed to determine the architecture for the targets: input encoding and pre-processing, training solver, network layer architecture, and cost function. Architectures were chosen based on the corrected Akaike Information Criterion (AICc), coefficient of determination (R2), and percent capture (%C). When predicting principal stresses, AICc = -59.62, R2 = 0.996, and %C = 97%, and when predicting the six-component stress tensor AICc = -45.50, R2 = 0.997, and %C = 80%. This research represents progress towards continuous, automated calibration of numerical models such that rapid, more accurate forecasts of changes in stress conditions will allow earlier reaction to challenging stress environments, increasing safety of excavations.

6.3 Introduction

Brittle rock mass deformations and their impact on underground excavations are difficult to monitor and predict using visual observations and conventional instrumentation. In the mining context, the ability to forecast brittle behaviour, including seismicity, is invaluable to planning mine operations, ore extraction sequencing, and budgeting. Seismic monitoring systems have gained popularity in the last two decades in mining operations to monitor seismogenic zones, or areas of active seismicity, to determine failure of the rock mass primarily through shearing and intact rock fracturing, and thus forecasting possible excavation instability (M. Hudyma et al., 2008).

Each seismic event captured by microseismic monitoring systems contains data about the rock mass failure that caused the event. Spatial-temporal analyses of these events can be used to forecast trends in the rock mass failure processes, and project changes in the mine-scale stress regime. Undertaking conventional analysis techniques, such as frequency-magnitude, magnitude-time history, S wave to P wave energy ratio, and apparent stress time history, among others (Hudyma et al., 2008), can be prohibitive due to the large volume of data that must be parsed to extract the important trends. The rock engineering industry is now shifting its focus on finding new ways to interpret the vast amounts of “big data” being collected from digitalization techniques, such as microseismic monitoring systems. One solution is the application of machine learning algorithms (MLAs) to process the data quickly and efficiently to extract relationships.

MLAs have emerged as a powerful technique for examining geomechanical data and extracting nuanced rock mass deformation phenomena (Elmo et al., 2020; Lawal & Kwon, 2020; Mcgaughey, 2019; Morgenroth et al., 2019; Morgenroth et al., 2021). In particular, MLAs have become the state of the art for microseismic signal processing and classifying rock burst events (Duan et al., 2021; Jiang et al., 2020; Pu et al., 2019). Some identified challenges in classification are the need for large amounts of labeled data, and the requirement to balance the “rock burst” and “non-rock burst” events in the training dataset (Pu et al., 2019).

MLA applications for stress magnitude prediction and microseismic data interpretation have gained momentum in academic literature only recently. Previous research developed regression MLAs has proposed a rock burst risk prediction MLA to estimate rock burst risk (Wojtecki et al., 2021); introduced an MLA to distinguish microseismic versus blasting events in seismic waveforms (Pu et al., 2020), and developing a capsule network embedded in a Convolutional Neural Network to predict in situ stress for a strain-softening model (W. Gao et al., 2020).

Previous authors caution that data-driven approaches should not be met by a cognitive resistance to the introduction of new approaches that is all too common in rock engineering practice (Elmo & Stead, 2021). However, the application of MLAs without critical thought and due diligence carries the risk of creating an illusion of technological advancement (Yang et al., 2021). Thus, this chapter also presents a detailed explanation of the process used to develop the MLA presented and a description of how common pitfalls were avoided or mitigated. This chapter proposes an MLA to assist in the recalibration of a complex numerical model for Garson Mine, near Sudbury, Ontario, Canada. Garson Mine installed a seismic monitoring system in 1996, which outputs large datasets that must be examined to remove irregular data points, such as those with large-source location errors, after which the data can be used for daily operations. At Garson Mine, these data are also used to calibrate sophisticated numerical models in FLAC3D (ITASCA Consulting Group Inc., 2019a) to determine stope sequencing and the resulting stress redistributions, and therefore, inform the schedule and budget of the mining operation. However, the manual recalibration of the Garson Mine FLAC3D model is time consuming and computational expensive, and therefore, is not recalibrated as frequently as needed by the Garson Mine ground control engineers. This chapter investigates the use of an MLA that uses microseismic events to recalibrate a finite difference model by forecasting in-situ stress. Recalibrating the stress model using an MLA saves computational time and expense, so this can be done more frequently, and also allowing more time for engineering decision making rather than manual recalibration.

Previous work by the authors successfully applied a type of Artificial Neural Network (ANN) called a Convolutional Neural Network to a mining case study with sparse data by strategically formatting the inputs, balancing the training dataset, and optimizing the hyperparameters to make useful predictions (Morgenroth et al., 2021). Lessons learned from this previous work were applied to the Garson Mine dataset, which required consideration with respect to data formatting to ensure the ANN was learning the relationships between the inputs and output efficiently. ANNs can identify relationships between input variables and the output using a framework that mimics the interconnected neurons in the brain, thus, being able to find relationships that can be used to reproduce the actual observed ground behaviour. However, common backpropagation ANNs cannot efficiently consider the relationships between the sequential events of the time-series data, which is important for real-time analysis (Z. Liu et al., 2021). Therefore, the type of ANN selected to model the Garson Mine dataset was a Long-Short Term Memory (LSTM) architecture. LSTM networks were developed specifically to process time-series data, where strong correlations are

"remembered" and carried forward to forecast into the future (Mandic & Chambers, 2001). Previous studies relevant to the current research have been done to classify seismic records (Vallejos & McKinnon, 2013), to predict TBM operating parameters (X. Gao et al., 2019), to predict lithology in TBM construction (Z. Liu et al., 2021), and to perform unsupervised learning on TBM operational data (Erharter & Marcher, 2020).

The Garson Mine LSTM network presented herein uses the microseismic database and geomechanical properties from a previously calibrated FLAC3D model to predict the stresses in the FLAC3D model, thereby assisting in the recalibration process, as shown in Figure 6-1. The LSTM network aims to assist with the FLAC3D model recalibration process by using all available seismic data to predict the stresses in the FLAC3D model. The microseismic database contains data from the microseismic array installed at the mine, including the location and timing of the microseismic events, as well as microseismic parameters such as moment magnitude, energy, and apparent stress. The geology and geomechanical parameters used to train the LSTM network were taken from work completed by a specialized geomechanics consultant, who also built and calibrated the FLAC3D model used in this research (Kalenchuk, 2018). The ability to forecast the stresses and automate the recalibration of the FLAC3D model will allow Garson Mine to run scenario analyses and update their excavation sequencing more frequently based on operations and ongoing seismicity. To the authors' knowledge, this is the first known application of using an LSTM network to recalibrate a FLAC3D model.

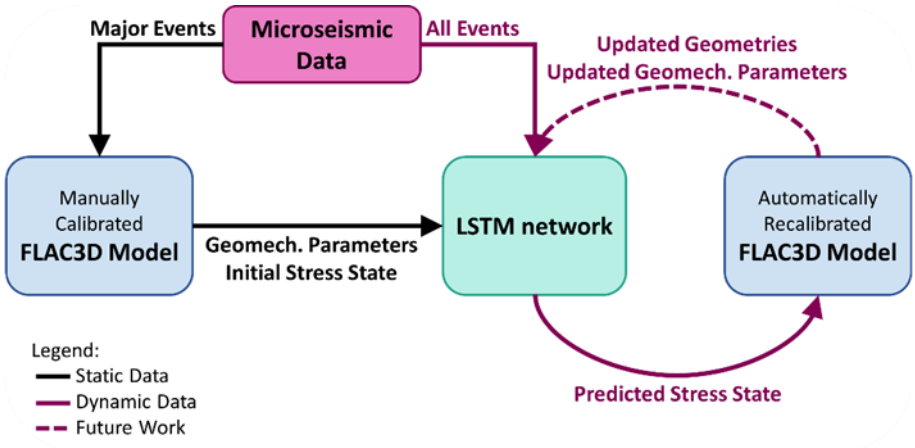

Figure 6-1. Workflow proposed in the present chapter, where the microseismic database and previously calibrated FLAC3D model are used to train an LSTM network to predict the changing stress state in the FLAC3D model. New microseismic events can then be passed to the trained LSTM network and predicted stresses are imported into the FLAC3D model, thereby recalibrating it. Future work is recommended to create dynamic interaction between the LSTM network and the FLAC3D model, automating the recalibration process.

6.4 Background

6.4.1 Long-Short Term Memory Networks

Over the last decade, various MLAs have been applied to characterizing stress regimes indirectly, namely by classifying rockburst hazards: random forests (L. J. Dong et al., 2013), support vector machine (J. Zhou

et al., 2012), ANNs (Ribeiro e Sousa et al., 2017), and decision trees (Pu et al., 2018). One study compared the performance of ten frequently used MLAs for rock burst class prediction (Pu et al., 2020), and another found that time-series prediction of microseismic parameters was possible using an LSTM network (Zhang et al., 2021). The current research uses an LSTM network.

LSTM networks are a type of ANN that contain recurrent feedback nodes. Unlike typical ANNs which contain basic feed-forward cells (Figure 6-2a), LSTM networks use the recurrent LSTM cells (Figure 6-2b) to process single data points as well as entire sequences of data. Input “remember” and “forget” gates operate based on the activation strength of the input being passed through them, shown as tanh or sigmoid functions in Figure 6-2b. Thus the recurrent LSTM nodes can forget part of previously stored memory and add partial new information that is useful (Mandic & Chambers, 2001). This is particularly powerful when using time-series data as inputs, since the LSTM network will remember the useful parts of the history of that input. LSTM networks have also been shown to process highly correlated datasets effectively (Z. Liu et al., 2021). In the context of geomechanical datasets, and microseismic data in particular, the LSTM network has the ability to preserve knowledge about minor seismic events, and therefore, the stress/strain history, that lead up to major seismic events.

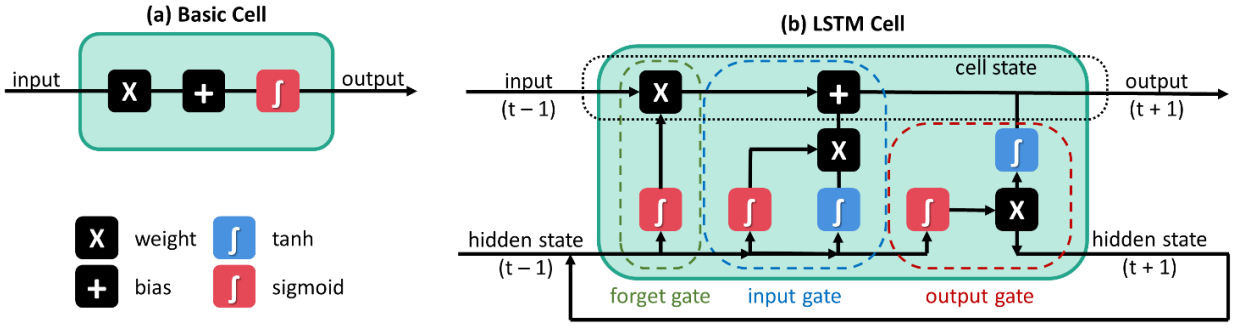

Figure 6-2. Schematic of (a) a basic feed-forward cell and (b) a Long-Short Term Memory (LSTM) cell. A basic cell applies a weight and a bias to the input and then fires it through an activation function. The recursive nature of the LSTM cell allows the algorithm to “remember” useful data for subsequent timesteps (t + 1) using information from the previous timestep (t – 1).

6.4.2 Case Study

6.4.2.1 Garson Mine

Garson Mine is a copper-nickel mine located near the town of Garson, Ontario, on the southeast rim of the Sudbury Basin (Figure 6-3). The copper-nickel sulphide deposits are hosted in parallel shears, which are offset by later stage dyke intrusions. The footwall typically consists of the lower zone Norite of the Sudbury Igneous Complex and metavolcanics (Greenstone/Metabasalt), while the hanging wall consists of metasediments. A bifurcated olivine diabase dyke crosscuts the host rock and the orebodies at Garson. The Garson Mine orebodies are more deformed than any other orebodies of the Sudbury Basin. The orebodies strike approximately east-west and dip south 75°. The geomechanical parameters applied to the

FLAC3D model and used as part of this research are provided in Table 6-6 and Table 6-7 (K. S. Kalenchuk, 2018). Note that it is not within the scope of this study to validate these rock mass parameters, as they have been based on previous work by geomechanical consultants spanning back to 2012.

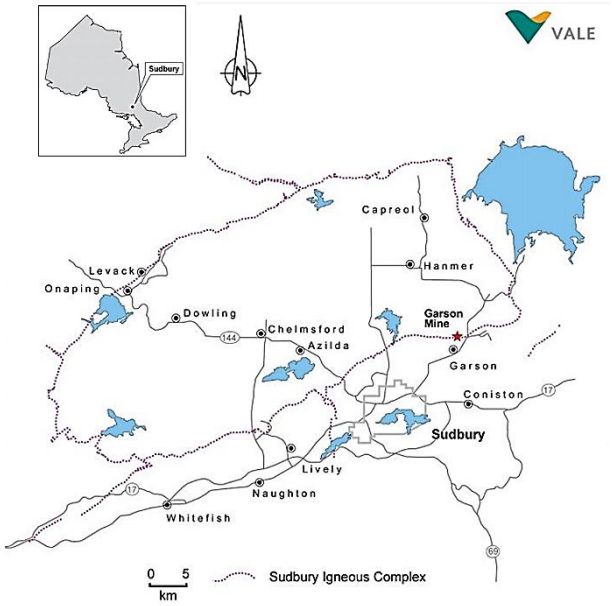

Figure 6-3. Location map of Garson Mine in relation to Sudbury, Ontario (courtesy of Vale).

In response to large seismic events that occurred at Garson Mine between 2006 and 2008, strategic and tactical mitigation measures were implemented to withstand future seismic impact (Yao & Moreau-Verlaan, 2010). The strategic measures relevant to this study are the ongoing numerical modelling that aimed to re-examine mining sequences, and the introduction of field instrumentation to monitor ground movement and ground support effectiveness in high risk areas. A microseismic monitoring system was installed in 1996 and updated in 2009 to collect and evaluate continuous waveforms in real time, allowing for source-location and calculation of source parameters as events occur (Vale, 2015). The current system consists of 22 uniaxial accelerometers and 5 triaxial accelerometers, which are installed between the 3400 Level and the 5100 Level. A strong ground motion seismic system was also installed to measure large scale seismic events. Events with local magnitudes less than +0.8 MN are classified as microseismic events, while events greater than +0.8 MN are classified as macroseismic events. The data that forms part of this study is primarily comprised of the former.

This research is focused on the #1 Shear West (1SHW) area of the mine (Figure 6-4), which is characterized by massive sulphide mineralization with sharp hangingwall and footwall contacts. In 2017 and 2018, the mine experienced an increase in frequency of large-magnitude seismic events in this area, including a +3.1 MN event on December 14, 2017 and a +2.3 MN event on July 29, 2018. A FLAC3D model was developed by geomechanical consultants to back analyze the mechanisms that triggered these events, allowing for

the evaluation of seismogenic risk associated with past, present, and future mine-induced stress evolution (K. S. Kalenchuk, 2018).

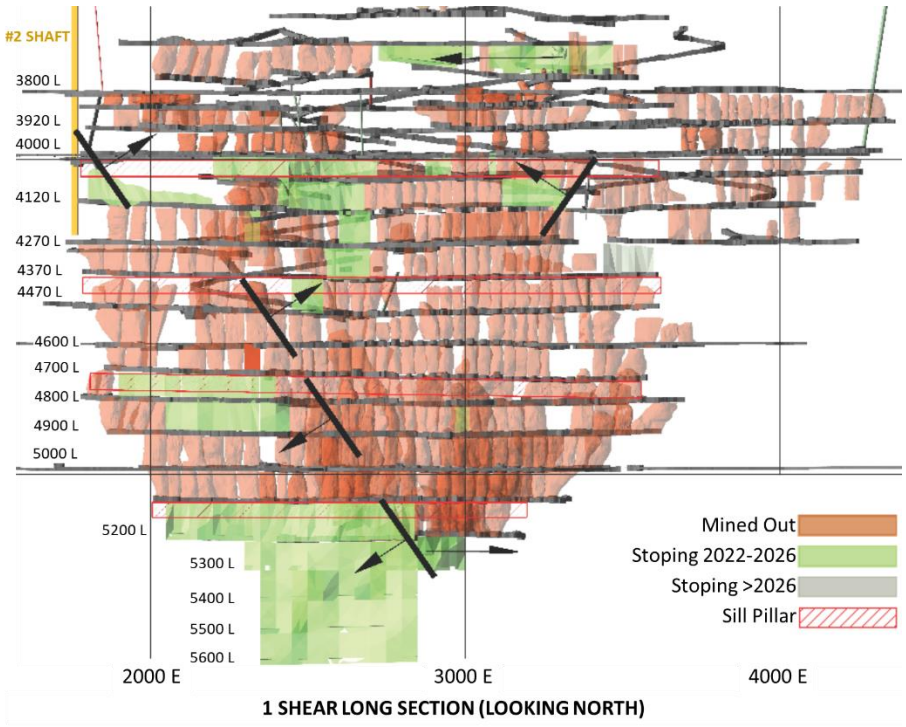

Figure 6-4. Garson Mine 1SHW area of interest, showing geometry of the mine overall with levels (L), location of faults (thick black lines with arrows), main access drifts and shafts (dark grey), and timing of stope removal (orange, green, grey) (adapted from Kalenchuk, 2018).

6.4.2.2 FLAC3D Model and Calibration

FDM approaches are often chosen to model time-dependent rock deformation behaviour due to the explicit time-marching component of its formulation scheme (ITASCA Consulting Group Inc., 2015). FDM is the most direct way to discretize a continuum, where points in space are replaced with discrete equations called finite difference equations, which are used to calculate displacement, strain, and stress in the material in response to conditional changes in the rock mass. Solutions are formulated at the grid points at the local scale, so no global matrix inversion is required for inelastic solutions, saving computational time and intensity. A popular code for stress analysis in rock using an FDM approach is FLAC (ITASCA Consulting Group Inc., 2015). Displacement or stress back-analysis, or synthetic rock mass models coupled with Discrete Fracture Networks (DFNs), can be used to calibrate FLAC model parameters (Farahmand et al., 2018; Ninić et al., 2017).

As part of the strategic and tactical measures to manage current and future seismic risk at Garson Mine, a FLAC3D model was developed and calibrated to back-analyze seismicity. The scope of the modelling efforts included assessing the stability of the underground excavations and suggest operational changes to improve safety, reduce costs, and increase profitability. The Garson Mine FLAC3D model was set-up

using strain-based stress initialization, where the in situ stresses were applied in the model using strain boundaries, thus allowing the stiffer units to load up more than the softer units (K. S. Kalenchuk, 2018). Two constitutive models were applied: Hoek-Brown (Hoek & Brown, 1997), and strain weakening Hoek-Brown (Hoek & Diederichs, 2006). The FLAC3D model focused on the back-analysis of the two large-magnitude events that occurred on December 14, 2017 and July 29, 2018.

The FLAC3D model was calibrated using semi-quantitative and qualitative approaches relying on microseismic data dating back to January 2006 (Kalenchuk, 2018), with the aim of reproducing the two large events mentioned previously. The calibration efforts used a subset of the available seismic database to improve data analysis efficiency by reducing the number of data points that need to be read into the model. The seismic data were filtered to include only events with a magnitude greater than -2 MN, a source-location error less than 30 m, and with an elevation in the area of interest between 1830 m and 3050 m, where surface elevation is 3660 m (Figure 6-5). The seismic data were temporally filtered to evaluate seismicity which occurs within each model stage, where each stage was set up to simulate historical mining according to the true extraction sequences at Garson Mine. Strain-based calibration was required to investigate the rate of post-peak decay to residual conditions, where the stress and strain conditions within the models were sampled at the location of each seismic event during the corresponding model stage. Then correlations between observations of actual rock mass behaviour and model results were evaluated to determine site-specific strength criteria and boundary conditions that should be applied in the model. The first phase of calibration compared the principal stress state sampled from the model to the estimated material strength. Conceptually, a well-calibrated model should produce statistically repeatable results for each geotechnical domain, however a perfect match in each model stage is exceptionally difficult to achieve. The second phase of model calibration involved a qualitative assessment of the first phase and adjusting material strengths to address discrepancies between the microseismic data and the queried model values. Each time the FLAC3D model was updated, this time-consuming recalibration process was repeated, and only a subset of the available database was used.

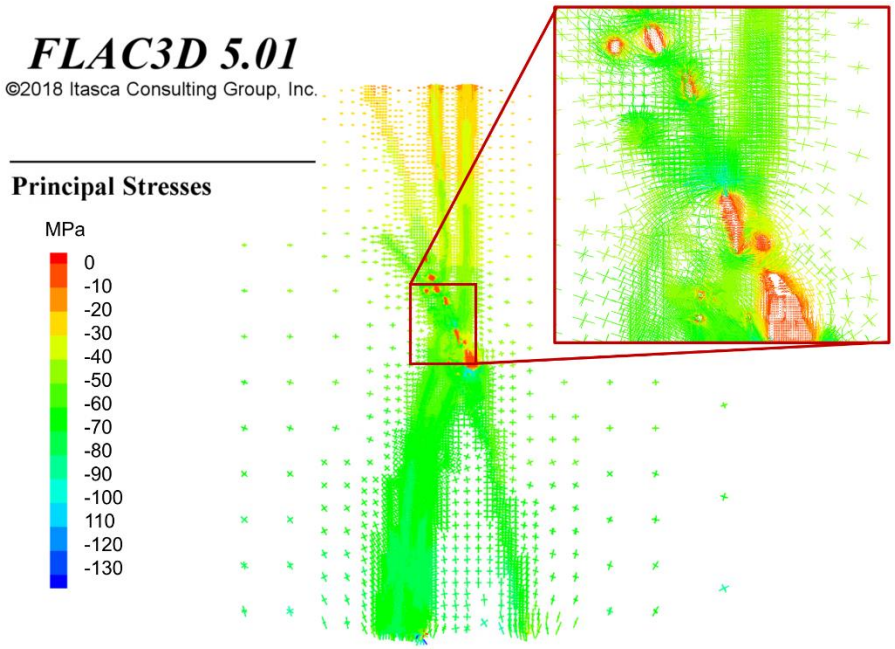

Figure 6-5. Manually calibrated FLAC3D model of 1SHW area of Garson Mine, showing volumetric stress contours (Kalenchuk, 2018).

It is worth note that even a well-calibrated numerical model has limitations in that no model is a perfect replica of the observed rock mass behaviour. This epistemic uncertainty is exceptionally difficult to eliminate from any modelling techniques, whether physics-based or data driven. In the present research the authors acknowledge the errors introduced by this uncertainty into the proposed LSTM network, however the stress values computed by the numerical model are the only viable ground truth available for training the algorithm.

6.5 Garson Mine LSTM Network Development

Numerical model recalibration is a time-consuming process, both in terms of computational time and work hours. Previous authors have proposed replacing the numerical model entirely with an MLA, whereby a deep learning model consists of three linked models that train the stress values for each point in space in parallel (Gao et al., 2020). While promising, the prospect of replacing widely accepted numerical models with MLAs is one that will require more validation both in research literature and in practical applications. However, recalibration of physical models using MLAs is a smaller step in this direction that still allows the user to apply their expertise about the physical system in the context of numerical modelling (Ninić et al., 2017). ANNs present an opportunity to expedite the recalibration process once the initial model set-up has been completed.

Critical thinking must be applied when parameterizing ANNs for rock engineering decision making (Elmo et al., 2020; Phoon et al., 2021). As with all modelling endeavours, development of an ANN should be an iterative process, starting with the simplest model that achieves satisfactory results and only adding complexity as justified by the real-world data and observed phenomena (Jakeman et al., 2006). This section

describes the setup of the Garson Mine LSTM network, and the steps taken for model selection, as supported by performance metrics throughout its development.

An LSTM network algorithm was chosen to maximize the usefulness of the data contained in the Garson Mine microseismic database. The Garson Mine LSTM network was developed to forecast the stresses at each zone centroid in the FLAC3D model. The input data was formatted into sequences of microseismic events that the LSTM network could “learn” and use to forecast the stresses (see section 6.5.2.3 Input Formatting for details). The forecasted stresses were compared to the stresses computed by the FLAC3D model (i.e., the targets). The square of the residual between the LSTM network outputs and the FLAC3D targets is minimized during algorithm training. MATLAB R2021 and the Deep Learning, Computer Vision, and System Identification toolboxes were used to develop the Garson Mine LSTM network (MathWorks Inc., 2021). One LSTM network was developed to forecast all the stress components, as opposed to individual LSTM networks for each stress component, because the latter would not preserve the relationship between stress components during algorithm training.

As with any modelling approach, there are a variety of best practices that should be applied to ensure that the chosen MLA is appropriate for the problem. In particular, applying techniques to confirm that a well generalized model has been produced is crucial. For example, it is standard practice in ANN development that data partitioning is applied to the available dataset such that the algorithm is trained on one subset and tested on another (Marsland, 2014). This ensures that the resulting ANN is well-generalized and not only able to reproduce the data it was trained on, but also can identify relevant patterns in new data. It is also typical to run an ensemble of models to quantify model uncertainty. Ensemble modelling yields a distribution of possible model forecasts from which model uncertainty can be derived. These techniques, among others, are describes throughout this section in the context of the Garson Mine LSTM network development. While MLAs are a new tool in the rock engineering toolbox, standards of practice exist and can be modified from other fields of science and engineering to ensure their reliability.

6.5.1 Input Data

The formatting of the geomechanical input data for use in ANNs is a crucial step in ensuring its success. Site-specific data is “ugly”, and as coined by (Phoon et al., 2021), can be described as *MUSIC-3X* (Multivariate, Uncertain and Unique, Sparse, Incomplete, and potentially Corrupted with “3X” three dimensional spatial variations). Therefore, careful consideration is required when determining how data is passed to an MLA, such that the useful information learned can be maximized during the learning process.

For the Garson Mine LSTM network, the inputs were obtained from the microseismic database (2015-2018) and from the previously calibrated FLAC3D model (Table 6-1). The microseismic parameters chosen as inputs characterize the evolution of the stress regime leading up to larger seismic events and rock bursts, following similar work by (Zhang et al., 2021). Previous work developing an LSTM network for Garson Mine found that the microseismic data alone is not enough to accurately predict the stresses in the FLAC3D model, because all the seismic parameters in the database are highly correlated to each other (Morgenroth,

Perras, Khan, et al., 2021). For this reason, additional parameters from the FLAC3D model (e.g., material properties, geological zones, constitutive model) were added to the candidate input pool to increase the variability and uniqueness between the inputs and the targets. Training, validation, and testing data was randomly split at 80%, 10%, and 10% for LSTM network development. Another option is to split the data spatially, as to investigate whether there are auto-correlation issues that may artificially inflate the performance of the final trained algorithm, which is addressed in Section 6.6 Results and Discussion.

Table 6-1. Summary of inputs and targets used to train the Garson Mine LSTM network (19 inputs and 3 or 6 targets).

	Inputs	Targets (primary)	Targets (secondary)
Microseismic data*	Time Stamp	Major Principal Stress, σ1	Normal Stress, σxx
	Microseismic event location (N, E, El.)	Int. Principal Stress, σ2	Shear Stress, σxy
	Moment magnitude	Minor Principal Stress, σ3	Shear Stress, σxz
	Seismic moment		Normal Stress, σyy
	Energy		Shear Stress, σyz
	Es/Ep		Normal Stress, σzz
	Apparent stress		
	Source-location error		
FLAC3D Parameters	Geological group		
	Constitutive model		
	Elastic modulus		
	Poisson's ratio		
	Hoek Brown mb, s and a		
	Accumulated plastic strain		

*Descriptions of these microseismic parameters and their derivation can be found in Hudyma et al. (2008) and Trifu & Young (1992).

As articulated by Yang et al. (2021), clear engineering geology information, whether in analogue or digital form should always accompany the use of real rock mass data for any empirical, numerical, or data-driven model. As MLAs become more prevalent in rock engineering research and practice, transparency of the input data used to develop algorithms is particularly crucial to ensure reproducibility. Thus, distributions of the inputs and targets used to develop the Garson Mine LSTM network are shown in Figure 6-6 and Figure 6-7, respectively. It is evident that the distributions are complex, arising from spatial and temporal dependencies. Some of the challenges include the mixture of categorical and numerical input data, the large difference in order of magnitude of inputs, and the non-Gaussian distributions of the targets. The Garson Mine LSTM network must learn the nuanced relationships within the inputs, between the inputs and targets, and amongst the targets.

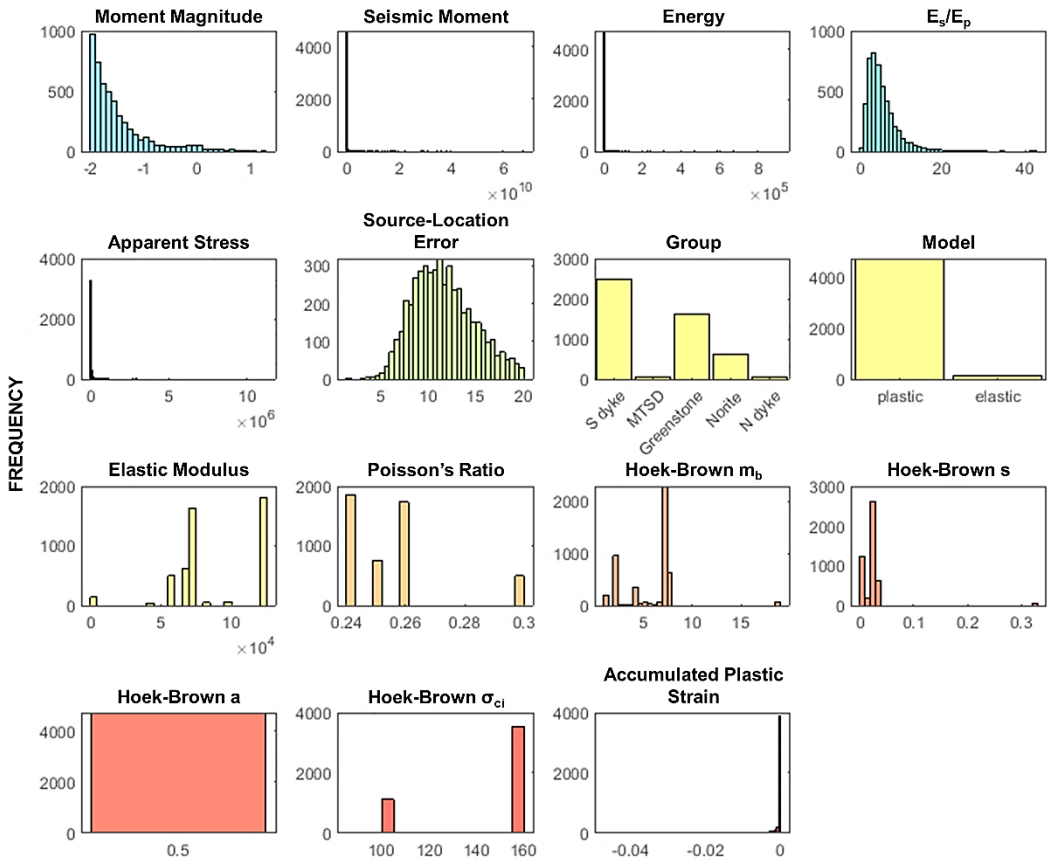

Figure 6-6. Histograms of all inputs used to train the Garson Mine LSTM network, apart from the time stamp and location (Northing, Easting, Elevation) of the microseismic events.

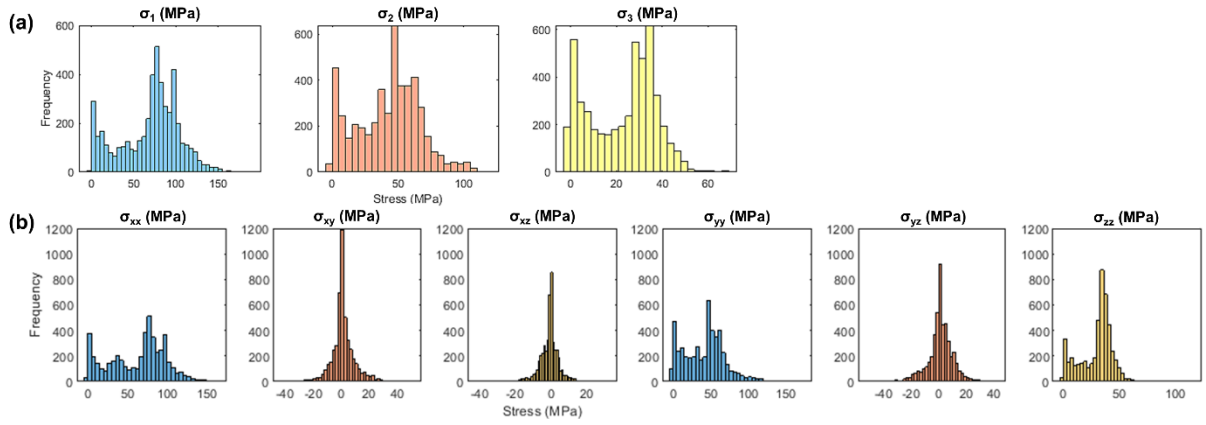

Figure 6-7. Histograms of the targets used to train the Garson Mine LSTM network, where the targets are (a) the three principal stresses or (b) the six-component stress tensor from the calibrated FLAC3D model.

6.5.2 LSTM Network Development

The Garson Mine LSTM network was developed using the authors previous experience and following examples in literature where MLAs were applied to analogous problems (Erharter & Marcher, 2020; Gao et

al., 2019; Liu et al., 2021; Vallejos & McKinnon, 2013). Following a literature review, several aspects were targeted for optimizing the performance of an LSTM network for predicting the stresses at Garson Mine: input encoding and pre-processing, training solver, network layer architecture, and cost function. Collectively, these will be referred to as *hyperparameters* throughout this chapter and are used to control the LSTM network learning process (Claesen & de Moor, 2015).

Figure 6-8 is an illustration of the process that was used to develop the Garson Mine LSTM network. The blue and purple nodes indicate the inputs and the outputs of the algorithm, respectively. The green nodes indicate where ensembles of 100 LSTM networks were run to evaluate the algorithm alternatives listed, and the alternative to proceed with was selected based on the AIC_c (see section 6.5.2.2 Model Selection and Performance Metrics). Grey nodes indicate where model performance was evaluated as a function of various outputs, and yellow nodes indicate the resulting algorithm revisions that could be considered.

As indicated by the green nodes, two main areas of LSTM network development were targeted as part of this process: input encoding and pre-processing, and algorithm architecture. Since the inputs were a combination of numerical data (e.g., microseismic parameters) and categorical data (e.g., geological zones), the categorical data had to be encoded for input into the LSTM network. Additionally, the wide range of magnitudes in the input data (e.g., energy with values $x10^4$ versus moment magnitude values as small as -2) needed to be normalized prior to training to avoid over emphasis of inputs with larger magnitudes. There are several options for network architecture optimization, however the focus here has been on the training solver type, how many layers (i.e., how “deep”) the network was, and whether the signal was normalized between successive layers. Finally, three cost functions were compared, where the cost function is a measure of how poorly the model is able to estimate the relationship between the inputs and outputs and is minimized during training.

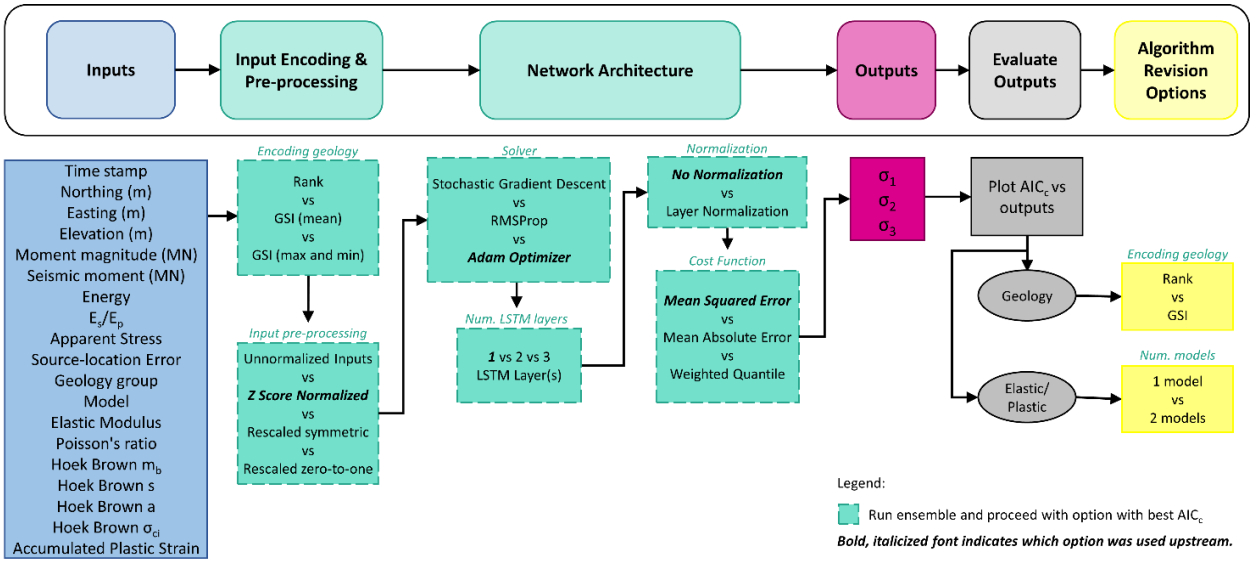

Figure 6-8. Process for developing the Garson Mine LSTM Network, where nodes indicate steps in the algorithm development, including inputs, architecture alternatives, outputs/targets, and evaluation of outputs for further algorithm refinement. Green nodes identify where each algorithm alternative was applied to an ensemble of LSTM models, and the alternative with the best performance (as determined by the AIC_c) was applied for subsequent algorithm development.

6.5.2.1 Ensemble Modelling

In order to obtain a measure of confidence in their predictive accuracy, it is common practice to run an ensemble of algorithms where the weights and biases are randomly reinitialized for each model run (Afraei et al., 2019; Benardos, 2008). Each member of an ensemble is trained independently, and the predictions can be combined to form a statistical distribution of possible outputs (Z. Zhao et al., 2007). An ensemble of models will generalize well to a dataset, as compared to a single, finite model that represents only one random initialization of weights and biases. This approach allows the developer of the algorithm to gain an understanding of how well the model is converging to the global optimum, as well as to determine the variance of the model output. In general, a minimum of 30 models should be trained as part of an ensemble (Setiono & Liu, 1997). For this research, an ensemble of 100 models was run for each combination of hyperparameters.

6.5.2.2 Model Selection and Performance Metrics

The hyperparameters selected for the Garson Mine LSTM network were based on the relative performance of the algorithm options evaluated. In general, model selection is based on the principle of parsimony – the model selected should be as simple as possible but as complex as necessary (Höge et al., 2018). Following this principle, the performance metric used to select hyperparameters and develop the Garson Mine LSTM network is the Corrected Akaike Information Criterion (AIC_c) (Hurvich & Tsai, 1989). The AIC_c penalizes a model for being too complex relative to its accuracy, preventing models that are accurate but overly complex from being selected as the preferred model architecture. The AIC_c is a sample-corrected AIC, which is

based on the Sum of Squared Errors (SSE), number of samples, and number of model parameters. The SSE is a measure of how the variation in the dependent variable in a model cannot be explained by the model, where a lower SSE indicates that the model can explain the data well while a higher SSE indicates that the model poorly explains the data. The SSE calculation is shown in Equation 6-1, where y_i is the ground truth (observed data) and $\hat{y}_i$ is the predicted value. The AIC is essentially a measure of entropy, evaluating the model's fit on the training data and adding a penalty term for the complexity of the model (Akaike, 1969). A lower AIC indicates the best balance of model fit with generalizability. The AIC calculation is shown in Equation 6-2, where n is the number of training samples and p is the number of parameters (weights and biases). AIC assumes that the same data is used between models, the same outcome variable is predicted between models, and that the sample is of infinite size. This final assumption gave rise to a sample-size adjusted AIC_c, shown in Equation 6-3. AIC_c should be used when the ratio of training samples (n) to the number of parameters (p) is less than 40 (Burnham & Anderson, 2002), i.e., when there are less than 40 samples of data for every training parameter in the model. The AIC_c approaches and converges to the AIC value as the number of training samples approaches infinity.

$$SSE = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \qquad \text{Equation 6-1}$$

$$AIC = n * log\frac{SSE}{n} + 2 * p \qquad \text{Equation 6-2}$$

$$AIC_c = n * log\frac{SSE}{n} + \frac{n+p}{1-\frac{p+2}{n}} \qquad \text{Equation 6-3}$$

In addition to AIC_c, the coefficient of determination (R^2) and percent capture (%C) are calculated for each model ensemble in this study. The R^2 value is a measure of the goodness of fit of any regression model. %C indicates the whether the true value is encompassed within the variance of the model ensemble predictions, i.e., whether the model is able to predict the target within the ensemble's variance (De Santi et al., 2021). %C should not be used exclusively to determine a model's performance, as models with extremely large variances (i.e., low precision) will also indicate a high %C.

6.5.2.3 Input Formatting

The seismic database was filtered to the same extent in terms of Northing, Easting and Elevation as the FLAC3D model, to ensure the seismic events used to train the LSTM network actually occurred in the vicinity of the 1SHW developments (Figure 6-9). The authors note that there may be error in the calculated stresses at the model boundary, resulting from the exclusion of microseismic events occurring just outside the model extents, however these effects are considered to be negligible as the event magnitudes are small (-2 MN). Any microseismic events with a source-location error greater than 20 ft (6 m) were removed from the database, to improve reliability of the data points. Each remaining seismic event in the database was matched to the nearest zone centroid in the FLAC3D model in order to assign the relevant FLAC3D zone locations to each seismic event. Finally, the seismic events that occurred near the same zone centroid were

ordered into a sequence of seismic events. The LSTM network training data were formatted so that each FLAC3D zone centroid was associated with a sequence of seismic events (from the microseismic database) that results in a stress state (from the calibrated FLAC3D model). Based on previous work by the authors, a minimum length of 10 microseismic events was selected to form the sequences used to train the Garson Mine LSTM network (Morgenroth et al., 2021). This previous study found that a smaller minimum event sequence length allowed for an increased number of discrete sequences to be created from the training dataset, thereby increasing the LSTM network's performance.

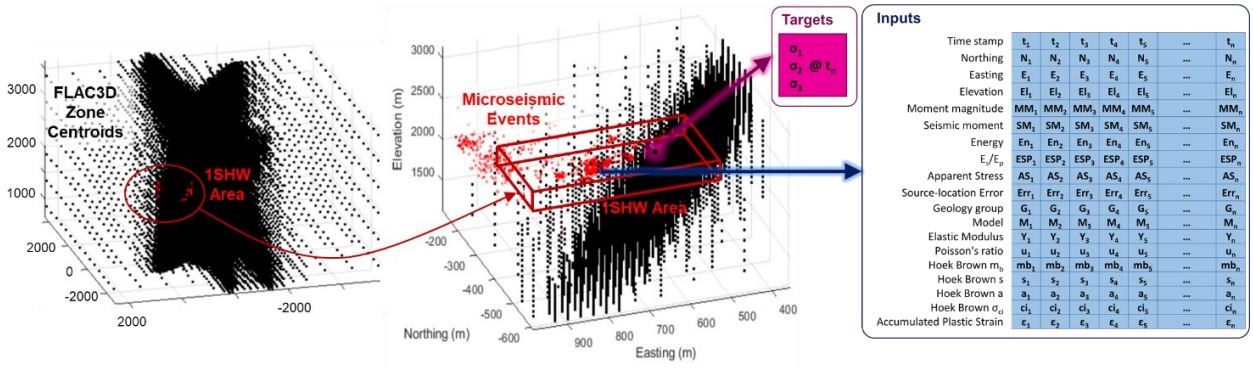

Figure 6-9. FLAC3D zone centroids from the previously calibrated model, with 1SHW area of interest indicated by the red cube. Microseismic parameters were formatted into a sequence and attributed to the nearest FLAC3D zone centroid in order to create a labelled sequence of events and corresponding stress state to form the training dataset for the Garson Mine LSTM network.

6.5.2.4 Input Encoding and Pre-processing

When combining numerical and categorical data into one algorithm, it is common practice to encode categorical data by replacing the classes with numerical substitutes (Marsland, 2014; Murphy, 2012). In the case of the Garson Mine dataset, there were two categorical inputs used for training – geology class and constitutive model. The latter was binary encoded, as there were only two constitutive models used in the FLAC3D model: elastic (backfill material) and plastic (rock). The geology classes were more complex to encode, as assigning a numerical value introduces an ordinal relationship between the geological classes. Therefore, three methods of encoding the geology were compared in the development of the Garson Mine LSTM network, as shown in Table 6-2. An ordinal encoding was applied based on the mean Geological Strength Index (GSI), and two numerical encoding methods were applied using the mean GSI and the maximum and minimum GSI values, respectively (Hoek & Brown, 2019). The purpose of testing these different encoding methods was to determine whether the performance of the Garson Mine LSTM network would benefit from learning with the explicit GSI values, or if an ordinal ranking was sufficient.

Table 6-2. Three methods of encoding Garson Mine geology groups (Kalenchuk, 2018).

Geology Class	Ordinal Encoding	Mean GSI Encoding	Max & Min GSI Encoding
Metasediments	0	28	35 & 20
South Dyke	1	65	75 & 55
Greenstone/Metabasalt	2	70	75 & 65
Norite	3	75	80 & 70
North Dyke	4	95	100 & 90

The relative magnitudes of the inputs used for LSTM network training can impact on the back-propagation process by biasing the training towards inputs with larger magnitudes (Morgenroth et al., 2021). To neutralize this, inputs may be normalized during training (Murphy, 2012). Previous research has found success normalizing inputs using the minimum and maximum value in order to improve forecasting accuracy and model convergence speed (Liu et al., 2021). For the development of the Garson Mine LSTM network, three input normalization schemes were tested against the base case of unnormalized inputs: z score (mean/standard deviation), rescaled to -1 to 1, and rescaled to 0 to 1. Rescaling does not impact the shape of the distributions of the inputs, while z score normalization includes the standard deviation and accounts for differences in relative distribution shape. Rescaling inputs from -1 to 1 impacts the training solver (discussed in 6.5.2.5.1 Cost Function and Training Solver) differently from rescaling inputs from 0 to 1, in that the negative values result in slightly different gradients calculated during stochastic gradient descent. Each stress target was normalized by the mean to ensure that the training of the LSTM network is not impacted by the relative magnitudes of the stress components.

6.5.2.5 Algorithm Architecture

The LSTM architecture applied in this study is a sequence-to-one regression network, where the sequence is the series of seismic events and related inputs that occurred at a particular zone centroid, and the one regression output is a vector comprised of the stresses at that same centroid. The hyperparameter options that were tested to develop the Garson Mine LSTM network are described in this section, including the training solvers, number of LSTM layers, layer normalization, and cost functions that were compared to produce the best performance.

6.5.2.5.1 Cost Function and Training Solver

The cost function is a hyperparameter used to calculate the error, or loss, between the output from the algorithm and the ground-truthed target. The loss is backpropagated into the LSTM network during training to update the weights and biases until the system converges on a solution, which is represented by the minimum of the cost function. A commonly used method of finding the minimum point of the cost function is gradient descent, discussed below.

Three cost functions are compared for the Garson Mine LSTM network: the mean squared error (MSE), Mean Absolute Error (MAE), and the weighted quantile functions. MSE is the most commonly used regression cost function and represents the sum of the squared distances between the predicted values

and the targets. MAE is the sum of the absolute differences between the predicted values and targets, i.e., it represents the average magnitude of the errors. MSE and MAE are more sensitive to outliers when it comes to non-normally distributed targets. For this reason, the weighted quantile cost function was also evaluated as a method of accounting for the unique distributions of the targets (Koenker & Hallock, 2001). The quantile cost function is useful for predicting an interval or set of values, instead of one regression value.

During LSTM network training the training algorithm, or *solver*, updates parameters (i.e., weights and biases in the network) to minimize the cost function. This is done by taking small steps at each iteration in the direction of the negative gradient of the loss and is referred to as gradient descent. In standard gradient descent, the gradient of the cost function is calculated using the entire training set at once. Stochastic Gradient Descent (SGD) uses a subset of the training data to calculate the gradient, and a different subset is used at each training iteration until the entire training set has been used (a full pass is called an *epoch*) (Murphy, 2012). Because of this, SGD is a more computationally efficient approach than standard gradient descent. In this research, three variations of SGD were compared for training the Garson Mine LSTM network: SGD with momentum, RMSProp (root mean square propagation), and Adam (adaptive moment estimation).

Standard SGD may oscillate on its path toward the optimum during training. Adding a momentum term is one way to reduce oscillation, where the momentum term determines the contribution of the previous gradient to the current iteration. SGD with momentum uses a constant learning rate for all parameters, however RMSProp uses learning rates that differ by parameter. Specifically, RMSProp uses a moving average of the squares of the parameter gradients to normalize the updates of each parameter individually. This approach decreases the learning rates of parameters with large gradients and increases the learning rates of parameters with small gradients. Similar to RMSProp, Adam uses a parameter-wise update but with an added momentum term, using an element-wise moving average of both the parameter gradients and their squared values (Kingma & Ba, 2015). The Adam approach allows the parameter updates to pick up momentum in a certain direction, making it computationally efficient and well suited to problems that are large in terms of parameters or data.

6.5.2.5.2 LSTM Network Layers

In keeping with the principle of parsimony that governed the model selection criteria, relatively simple network architectures were evaluated for the Garson Mine data. Networks with layer normalization were tested, as it has been found to increase LSTM network precision by previous authors (Liu et al., 2021).The network layers that comprise the Garson Mine LSTM network architectures that were tested are listed in Table 6-3. Architectures with 1, 2 and 3 LSTM layers, and both with and without layer normalization layers were tested, as shown in Figure 6-10.

Table 6-3. Layers used to develop the Garson Mine LSTM network.

Layer Type	Settings	Description
Sequence input layer	Size = 19	Inputs sequence data into network. Size corresponds to number of inputs. Input data normalization is set here.
LSTM layer	Size = 100 Output Mode = Last	Learns long-term dependencies between timesteps in sequence data. Output last timestep of the sequence.
ReLu layer	n/a	Performs a threshold operation to each element of the input, where any value less than zero is set to zero. Generally follows a normalization layer.
Dropout layer	Probability = 0.5	Randomly sets input elements to zero between iterations to prevent overfitting.
Layer normalization layer	n/a	Normalizes a mini-batch of data across all inputs for each observation independently. Generally follows an LSTM layer.
Fully connected layer	Size = 100 or 3 (or 6 when predicting stress tensor)	Multiplies the input by a weight matrix and adds a bias vector. Connects to all the neurons in the previous layer, combining all features learned by the previous layers to identify the larger patterns. If preceding regression layer, size corresponds to number of targets.
Regression layer	Size = 3 (or 6 when predicting stress tensor)	Computes the loss of the regression network based on the chosen cost function. Size corresponds to number of targets.

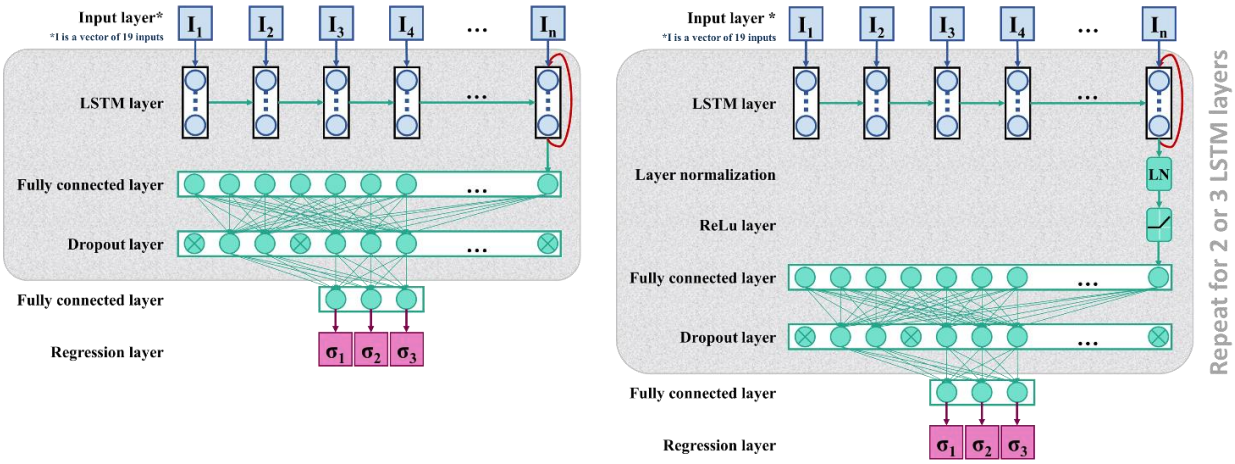

Figure 6-10. Schematics of architectures investigated for developing the Garson Mine LSTM networks for predicting the principal stresses and six component stress tensor, with and without layer normalization. Area shown in gray was repeated consecutively to test 2 and 3 LSTM layers.

6.6 Results and Discussion

The study to develop this Garson Mine LSTM network comprised of 5800 models run in total. Following the process detailed in Section 6.5.2 LSTM Network Development and Figure 6-8, the optimal architecture for predicting the

principal stresses and six-component stress tensors, respectively, are shown in Table 6-4. All ensemble performances are shown in Table 6-5.

Table 6-4. Final architecture of the Garson Mine LSTM network for two different targets.

	Targets	
	3 Principal Stresses	**6 Component Stress Tensor**
Geology encoding	max & min GSI	mean GSI
Input pre-processing	z score normalization	z score normalization
Cost function	Mean Squared Error	Mean Squared Error
Training solver	Stochastic Gradient Descent	Stochastic Gradient Descent
Number of LSTM layers	1	1
Normalization	none	none
AIC_c	-59.62	-45.50
R^2	0.996	0.997
%C	97 %	80 %

Table 6-5. Performance of all ensembles evaluated to develop the Garson Mine LSTM networks.

			3 Principal Stresses			6 Component Stress Tensor		
Scenario			**AIC_c**	**R^2**	**% Capture**	**AIC_c**	**R^2**	**% Capture**
Encoding Geology		Rank	-12.68	0.421	100	-3.63	0.896	100
		GSI (mean)	-11.08	0.233	100	-8.14	0.824	100
		GSI (max and min)	-12.68	0.446	100	0.16	0.389	100
Input Pre-Processing		None	-6.62	0.069	10	10.11	0.047	68
		Z score	-34.66	0.897	100	-2.37	0.600	98
		Rescale symmetric	-33.11	0.930	80	5.57	0.044	62
		Rescale zero-one	-7.08	0.091	63	8.32	0.050	60
Training Solver		SGD	-59.62	0.996	97	-45.50	0.999	80
		RMSProp	-35.78	0.950	100	-22.20	0.944	78
		Adam	-34.66	0.897	100	-22.84	0.980	78
Cost Function	MSE	1 layer	-59.62	0.996	97	-45.50	0.999	80
		2 layers	-27.44	0.991	77	-18.70	0.997	68
		3 layers	-11.52	0.953	67	1.50	0.967	23
		1 layer + layer norm.	-53.24	0.993	90	-42.22	0.998	85
		2 layers + layer norm.	-17.34	0.974	80	-6.51	0.994	65
		3 layers + layer norm.	4.81	0.540	70	3.11	0.960	27
	MAE	1 layer	-53.59	0.992	100	-42.26	0.997	85
		2 layers	-28.06	0.990	93	-15.62	0.994	68
		3 layers	-12.34	0.953	57	1.47	0.968	43
		1 layer + layer norm.	-48.85	0.987	100	-42.26	0.994	88
		2 layers + layer norm.	-15.54	0.955	100	-3.81	0.987	62
		3 layers + layer norm.	3.77	0.574	80	3.35	0.958	30
	Quantile	1 layer	-53.64	0.993	100	8.69	0.540	75
		2 layers	-26.38	0.987	97	30.82	0.599	62
		3 layers	-15.94	0.962	83	30.94	0.540	50
		1 layer + layer norm.	-48.16	0.986	97	8.76	0.566	73
		2 layers + layer norm.	-15.38	0.952	93	30.84	0.669	60
		3 layers + layer norm.	6.50	0.470	87	30.96	0.610	53

The performances of the test dataset of the two final LSTM network architectures for predicting the principal stresses and stress tenor, respectively, are shown in Figure 6-11. These results are presented for the test dataset only, following a random data split of 80% training, 10% validation, and 10% testing. The y = x line represents where the predicted stress is equal to the target stress from the FLAC3D model. The y error

bars represent the range of the 100 models in the ensemble. The three performance metrics are displayed for each model ensemble: the AIC_c, R^2, and %C. Recall that AIC_c is a sample corrected metric that accounts for model complexity, while R^2 is a measure of goodness of fit. %C illustrates whether the model ensemble can capture the target stresses, as demonstrated by whether the error bars cross the y=x line.

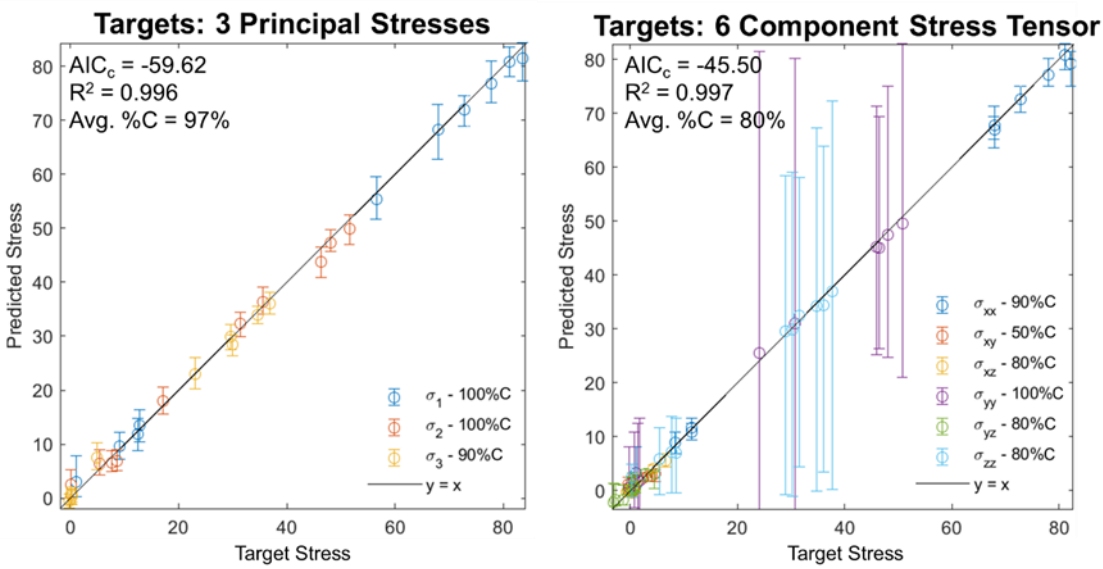

Figure 6-11. Performance of Garson Mine LSTM network on testing data, predicting 3 principal stresses (left) and 6 component stress tensors (right). Y-axis error bars indicate the variance of the ensemble predictions, where an ensemble is comprised of 100 LSTM networks. The AIC_c, coefficient of determination, and average %C are presented on the top left of each plot. The %C of each individual stress component is displayed in the legend.

The Garson Mine LSTM network was developed to assist in recalibrating the FLAC3D model used to assess the operations, including stope blasting and removal, and the impact of operations on ground reaction at Garson Mine. Several hyperparameters were targeted for optimizing the performance the Garson Mine LSTM network: input encoding and pre-processing, training solver, network layer architecture, and cost function. These hyperparameters were evaluated for predicting the three principal stresses and the six-component stress tensor for a manually recalibrated FLAC3D model of the Garson Mine 1SHW area. The performances of all the ensembles are presented in Table 6-5. Some general trends for each hyperparameter selection can be observed from this study.

For visualization purposes, the principal stresses predicted by the Garson Mine LSTM network were imported into the original FLAC3D model geometry and contoured (Figure 6-12). These figures show the predicted stresses in comparison to the stresses computed by FLAC3D (graphs below models), as well as the distribution of the predicted stresses from the ensemble of 100 LSTM network for five chosen locations. The stresses computed by FLAC3D fall within the range of the stresses predicted by the 100-model ensemble in all cases, indicating that the Garson Mine LSTM network is able to converge on a similar stress

state as the FLAC3D model. This indicates that the LSTM network provides a similar quality of FLAC3D recalibration as the manual recalibration approach, however, in a shorter period of time.

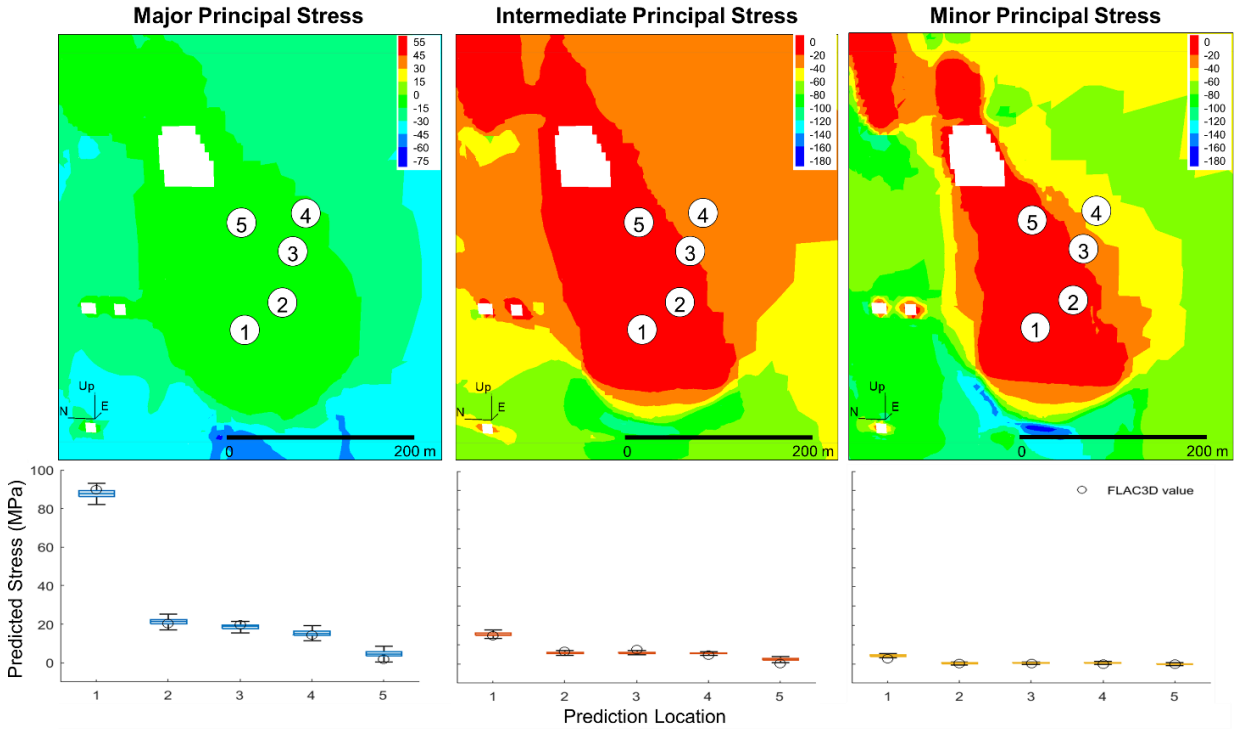

Figure 6-12. Principal stresses predicted by the Garson Mine LSTM network imported into FLAC3D model. Five locations are chosen for which the distribution of predictions of the ensemble of 100 LSTM networks are plotted (boxplot) versus the FLAC3D computed value (black circles).

Geology was encoded from the FLAC3D model for use in the LSTM network using three different approaches: ordinal encoding, representing the geological zones by their mean GSI, and representing the geological zones by the max and min GSI. When predicting the three principal stresses, the ordinal encoding and max/min GSI encoding had comparable performances, while the mean GSI encoding produced much poorer results. When predicting the six-component stress tensors, the three encoding methods performed similarly with the ordinal encoding performing slightly better. This investigation reveals that multiple encoding approaches should be tested when developing an MLA for rock engineering applications. It is also recommended that the raw parameters of classification systems, such as the components of the Q Tunnel Index (Barton et al., 1974) or Rock Mass Rating (Bieniawski, 1993), and not the classifications themselves should be used in MLAs to overcome the inherent limitations of those systems (Yang et al., 2021).

Input pre-processing was investigated to neutralize the relative magnitude effects of the various input parameters on LSTM network training. Ensembles with unnormalized inputs and rescaling from 0 to 1 performed an order of magnitude worse in terms of AIC_c than those with z score normalization or rescaling

from -1 to 1. For example, when the major principal stresses were predicted with unnormalized inputs and with rescaling from 0 to 1 AIC_c = -7, while z score normalization resulted in an AIC_c = -35 and rescaling -1 to 1 yielded an AIC_c = -33. Z score normalization produced a better AIC_c but lower R^2 as compared to rescaling from 0 to 1. The %C was greater than 98% for both predicting the principal stresses and six-component stress tensor as compared to other input normalization schemes (<80%). The z score normalization approach was selected for both predicting the principal stresses and stress tensor, because z score takes input variance into account.

Three cost functions were compared for calculating the loss during LSTM network training to assess their relative impacts on the final performance. The MSE, MAE, and quantile cost functions were compared. It was found that the performance was similar across LSTM networks with the same layer architecture, where MSE showed a slightly improved performance as compared to the MAE or quantile functions. Three training solvers were investigated to determine if there was an optimal solver for performing gradient descent during LSTM network training. SGD performed better than the RMSProp or Adam solvers in terms of AIC_c for both predicting the principal stresses and the six-component stress tensor, while the coefficient of determination and %C across the three solvers were similar. This is likely due to the increased complexity in the LSTM network when the momentum terms are added for the latter two solvers, resulting in the selecting of SGD.

The addition of learning layers in the LSTM network architecture was investigated to determine if increased depth of the algorithm would increase performance, despite the increased complexity that accompanies the addition of layers. LSTM network architectures with 1, 2 and 3 LSTM layers were compared, with and without layer normalization between them. Across all the combinations of hyperparameters for both predicting the principal stresses and the six-component stress tensors, the increased complexity from adding additional layers resulted in a decreased LSTM network performance in terms of AIC_c and %C. The coefficient of determination was largely unaffected by the change in number of layers. The additional of layer normalization had a slightly negative impact for all scenarios, indicating that it was not required for the LSTM network to learn the data patterns for the Garson Mine dataset.

When comparing the Garson Mine LSTM networks for predicting two different sets of targets, namely the principal stresses and the six-component stress tensors, the former generally had higher performance in terms of AIC_c, coefficient of determination, and %C. This is to be expected, since increasing the number of targets introduces more complexity during training when the training solver attempts to find a global minimum and balances the performance of all the targets during gradient descent. The accuracy when predicting the three principal stresses is similar across all three principal stresses, where the variance in the ensemble predictions is comparable across major, intermediate, and minor principal stress (σ_1, σ_2, and σ_3) predictions. The mean prediction accuracy when predicting the six-component stress tensor is similar across all stress components, however the variance for the y and z stress components (σ_{yy} and σ_{zz}) is far greater than any of the other components. This can be attributed to the training solver stepping towards

good performance across all six targets. Future work will consider error weighting schemes to distribute the ensemble variance evenly across all targets.

As mentioned previously, random data splitting was applied when training the Garson Mine LSTM network in order to prevent overfitting to the training dataset. This is standard practice in MLA development to ensure that the model is trained on as much of the dataset as possible, however, in spatially dependent datasets this method can inadvertently inflate the model performance due to autocorrelation. To investigate this, the final architecture for predicting the principal stresses was retrained using block splitting, which splits the training data sequentially instead of randomly during training. Since the Garson Mine training data is comprised of microseismic sequences that are built based on their proximity to a FLAC3D zone centroid, block splitting achieves a spatial splitting of the training data. An ensemble of 100 models was run using this data splitting method, and the results were found to be comparable to random splitting (AICc +5.12, R2 -0.003%, and %C +10%), indicating that spatial autocorrelation is not an issue for this particular problem.

For this research, one LSTM network was developed for the entire area of interest to capture the stress regime evolution with time, whether the targets be the principal stresses or stress tensors. Due to the minimum sequence length of 10 microseismic events used to format the training data, all the sequences used to train the Garson Mine LSTM network were located in plastic regions of the FLAC3D model, and primarily within three of the geology units (South Dyke, Greenstone, and Norite). Future work will investigate whether shortening the sequence length requirement, and therefore increasing the number of sequences available for training, increases the LSTM network performance. Sampling with replacement, bagging, and other resampling methods (Snieder et al., 2021) are additional approaches to augment the training data set (Breiman, 1996). Increasing the number of training sequences will also allow for a comparison of performance of the Garson Mine LSTM network for different geologies, indicating whether separate LSTM networks should be developed based on lithology. Those findings can be used to determine the quantity of algorithms needed to recalibrate the FLAC3D most accurately and efficiently.

The development of this LSTM network to recalibrate the FLAC3D model raises interesting considerations with respect to the conventional process of recalibrating a numerical model. Many model inputs are qualitative properties that are assigned numerical values, for example, GSI, and consequently the Hoek-Brown parameters m_b and s, and joint characterizations. In practice, rock mechanics engineers evaluate a distribution of these qualitative assessments and then choose a representative (but deterministic) value in the numerical model, an approach which is expedient but does not always respect the limitations of the underlying empirical systems (Yang et al., 2021). The results of the current research underscore this concept, particularly when considering that using the max/min values of GSI produced improved predictive performance over using the mean GSI in the LSTM network. This finding implies that using raw rock mass classifications instead of a mean or standard deviation in machine learning is preferable. Using machine learning to recalibrate numerical models also presents the opportunity to overcome the limitations of

empirically based data by allowing engineers to run ensembles of algorithms to recalibrate the numerical model, instead of choosing only one deterministic value.

The eventual goal of this research project is to establish a dynamic interaction between the Garson Mine FLAC3D model and the LSTM network to automate the recalibration process. With the trained LSTM network, new partial inputs can be provided to the LSTM network to obtain updated stress states in the FLAC3D model, either as new microseismic data becomes available or as excavation advances and stope sequencing are planned. This is possible because a trained algorithm is able to handle gaps in the input dataset as “unknown” values, and still provide an updated prediction based on its previous learning on complete inputs. In other words, stress model recalibration efficiency can be improved by creating a microseismic data feed through the proposed LSTM network for continuous, automated calibration of the FLAC3 model. This will allow the LSTM network to update the stress states in the Garson Mine model during successive stope mining, informing operational decisions forward in time, allowing the operations team at Garson Mine to produce more current re-interpretations of mining induced stress conditions as stoping advances, as drilling and blasting is performed, and as new rock mass and geological data are collected. Similarly, the operations team will be able to forecast how stope removal in one location in the mine will affect the stress potential in another, by feeding the LSTM network updated model geometries and FLAC3D parameters where new excavations are created or backfilled. This will inform decisions on managing the stress conditions through design and hierarchical controls. Earlier and more accurate forecasts of changes in stress conditions will allow earlier intervention and reaction to challenging stress environments, leading to increased safety of excavations and mine personnel.

6.7 Conclusions

This study presents an LSTM network using microseismic data and a FLAC3D model to predict the stresses at the zone centroids of the FLAC3D model for the 1SHW area at Garson Mine. For the Garson Mine LSTM network, microseismic data and material parameters from the FLAC3D model were used to train the algorithm. The inputs were formatted into sequences of microseismic events that could be assigned to a resulting stress state extracted from a FLAC3D model manually calibrated previously. Two LSTM network architectures were developed for two sets of targets: the principal stresses and the six-component stress tensors. Various hyperparameters were analyzed to determine the optimal architecture for the two sets of targets: input encoding and pre-processing, training solver, network layer architecture, and cost function. In general, the two architectures were similar with the exception of how the geology was encoded for use in algorithm training. In terms of AIC_c, where a lower value indicates better performance, the performance of the LSTM network predicting the three principal stresses was improved as compared to predicting the six-component stress tensor, with AIC_c values of -59.62 and -45.50, respectively. The R^2 values were 0.996 and 0.997, respectively, and the %C values were 97% and 80%, respectively. When visualizing the Garson Mine LSTM network results within the original FLAC3D model geometry, it was found that the principal

stresses computed by FLAC3D fell within the range of stresses predicted by the 100-model LSTM network ensemble.

Future work on the Garson Mine LSTM network should include improvement to the LSTM network itself, and how it is integrated with the FLAC3D model. An error weighting scheme could be applied to the six-component stress tensor to balance the respective component performances during algorithm training. The number of sequences available for training could be increased by both reducing the minimum number of microseismic events in a sequence, and by applying a bootstrapping algorithm to sample with replacement during training. The ultimate goal would be to create an automated recalibration process whereby stope mining in the FLAC3D model is used to re-train the LSTM network, and then the updated stress predictions are used to automatically recalibrate the FLAC3D model.

The development of an LSTM network to assist in numerical model recalibration presents an opportunity that is thus far largely untapped by the rock engineering community. If the same care is given to developing the ANN as is given to the manual recalibration of numerical models, this previously tedious and time-consuming process can be made more efficient and may allow in more reliable modelling of underground excavations.

Appendix

Table 6-6. Intact rock mass parameters for Garson Mine FLAC3D model (Kalenchuk, 2018)

Domain	UCS (MPa)			E_i (GPa)			v	m_i
	Avg	Min	Max	Avg	Min	Max		
Metasediments	130	85	175	68	46	90	0.24	20
Norite	120	75	165	83	71	95	0.25	22
Olivine diabase dyke (North and South)	200	160	240	117	83	150	0.26	27
Greenstone	200	140	240	167	95	190	0.23-0.28	25

Table 6-7. Rock mass parameters for Garson Mine FLAC3D model (Kalenchuk, 2018)

Domain	GSI		E_{rm} (GPa)			m_b			s			a
	Min	Max	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	
Metasediments	20	35	3.1	4.7 (*40)	7.7	1.1	1.5	2	0.0001	0.0003	0.0007	0.5
Norite	70	80	60.8	67.8	73.1	7.5	9	108	0.0357	0.0622	0.1084	0.5
Olivine diabase dyke (North)	90	100	112.2	116.3	116.3	18.9	22.6	27	0.3292	0.5738	1	0.5
Olivine diabase dyke (South)	55	75	73.9	95.5	95.5	5.4	7.7	11.1	0.0067	0.0205	0.0622	0.5
Greenstone	65	75	122.4	136.3	136.3	7.2	8.6	10.2	0.0205	0.0357	0.0622	0.5

* The E_{rm} value for the metasediments unit was considered to be uncharacteristically low, and resulted in numerical instabilities during stress initialization, so an approximation of 40 GPa was used.

CHAPTER 7. PRACTICAL RECOMMENDATIONS FOR MACHINE LEARNING IN UNDERGROUND ROCK ENGINEERING

7.1 Preface

This chapter presents recommendations of tools from the field of machine learning and how to apply them in the context of a practical rock engineering problem. The tools and approaches are illustrated with case studies.

The content of this chapter was submitted as a conference paper to *the 71st Austrian Geomechanics Colloquium* in 2022, and subsequently published in *Geomechanics and Tunnelling* as follows:

Morgenroth, J., Unterlaß, P. J., Sapronova, A., Khan, U.T., Perras, M. A., Erharter, G. H. & Marcher, T. (2022). Practical recommendations for machine learning in underground rock engineering – On algorithm development, data balancing, and input variable selection. *Geomechanics and Tunnelling*. https://doi.org/10.1002/geot.202200047

The contributions of the authors in the current chapter are as follows:

Josephine Morgenroth has conducted the ideation of the topics contained in the manuscript, assembled the illustrative case studies, coordinated with the co-authors on their contributions, and prepared and wrote the original manuscript of this publication. **Paul J. Unterlaß** and **Alla Sapronova** have contributed to writing sections pertaining to algorithm architecture selection and engineering verification, as well as the case studies to illustrate those topics. **Georg H. Erharter** has contributed to ideation of the topics and editing the manuscript. **Matthew A. Perras, Usman T. Khan, and Thomas Marcher** have supervised the research, provided the funding, and contributed to editing the manuscript.

This work is funded in part by the Natural Sciences and Engineering Research Council of Canada through the Postgraduate Scholarships – Doctoral program and the Michael Smith Foreign Study Supplement.

7.2 Abstract

Research has demonstrated that machine learning algorithms (MLAs) are a powerful addition to the rock engineering toolbox, and yet they remain a largely untapped resource in engineering practice. The reluctance to adopt MLAs as part of standard practice is often attributed to the “opaque” nature of the algorithms, the complexity in developing them, and the difficulty in determining how the datasets are used by the algorithms. This chapter presents tools and processes for the development of MLAs, input data selection, and balancing for practical underground rock engineering. MLAs for classification and regression - two main applications of machine learning - are presented in terms of developing MLA to extract information from the dataset to obtain the desired output. Engineering verification metrics are selected based on their suitability for specific output. Methods for input data balancing and selection are discussed with a focus on selecting input data that is appropriate for the problem without introducing bias or excess

complexity. Each tool and process for algorithm development, data preparation, and input selection is illustrated with a case study.

This chapter demonstrates that a geotechnical practitioner does not need to become an expert data scientist to apply MLAs to practical rock engineering problems. Once an understanding of the functions of MLAs is reached, the building blocks and open-source code are available to be adapted to suit the rock mass behaviour of interest.

7.3 Introduction

A common criticism of data driven methods, including machine learning, is that they are opaque and that the processes therein are non-intuitive or not representative of the physical system being modelled. This chapter challenges veracity of these perceptions by presenting tools that render machine learning more accessible and interpretable to practicing rock mechanics engineers.

The topics illustrated in this chapter are summarized in Figure 7-1. This chapter is divided into two main sections – algorithm development and input data. The first section introduces the two main applications of MLAs, namely classification and regression. Best practices for selecting an algorithm architecture as well as the corresponding verification metrics are presented. The second section provides insights into how rock engineering data may be balanced for use in an MLA. Input Variable Selection (IVS) methods are presented that can be used to rank input variables. Case studies are provided throughout the chapter to illustrate the development and utility of an MLA in the given case study context.

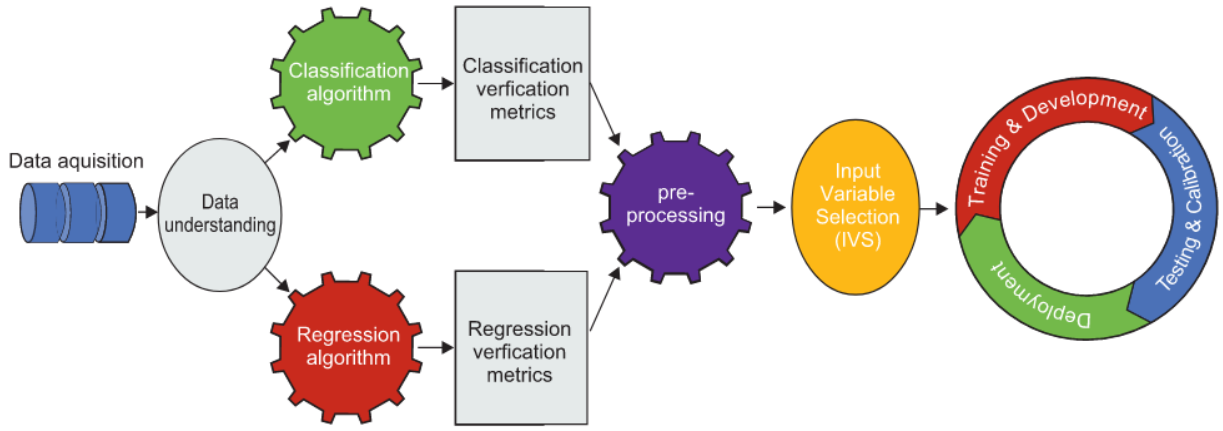

Figure 7-1. Illustration of the workflow presented in this chapter.

7.4 Algorithm Development

Depending on whether the desired output of a MLA is a classification or a numerical value, different algorithm architectures may be developed. This section discusses MLAs for classification and regression problems in the context of their selection, development, and engineering verification.

7.4.1 On Classification Algorithm Architecture Selection and Verification

The case study presented is a classifier that classifies the rock mass class ahead of the tunnel face (Unterlaß et al., 2022). Selection of a classification algorithm that outperforms others can only be achieved for the given dataset and cost function (Wolpert & Macready, 1996). Typically, an algorithm ranking is constructed where top-ranked algorithms are selected for evaluation (Shawkat & Smith, 2006). Data quality will significantly affect the accuracy of the classification algorithms. Due to its sparsity and imbalance, geotechnical data must be pre-processed in the form of standardization to keep the data quality at the highest possible level.

One universal approach can be suggested: start the classification with "white box" methods (e.g., methods from a decision trees family). These methods are significantly easier to explain and interpret, though they provide less predictive capacity and have limitations with respect to capturing the inherent complexity of the dataset. In addition, the white-box methods can be used to explore feature importance and check for the coherence of a dataset. To overcome data quality problems, one can shift from operating multiclass classifiers to building a cascade of binary classifiers (Figure 7-2).

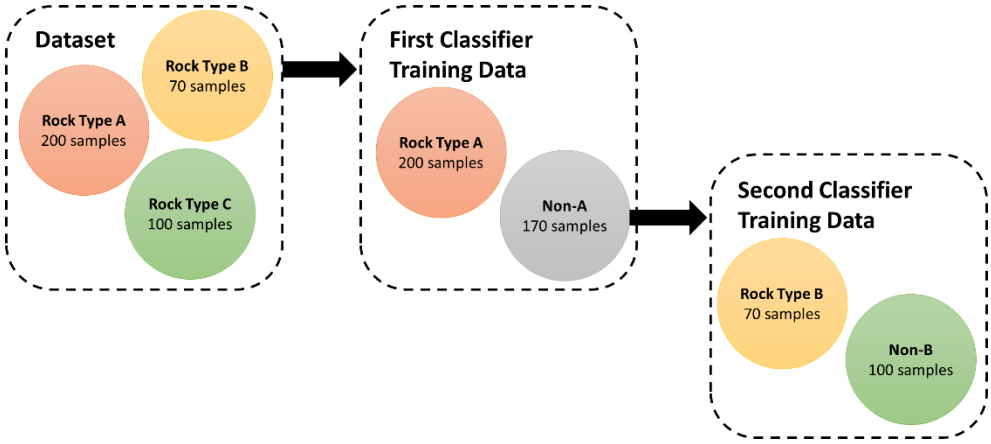

Figure 7-2. Improving the data quantity and imbalance by building a cascade of binary classifiers for a multiclass dataset.

The proposed cascade strategy allows for the ranking of classification algorithms since the verification results from the binary classifiers are relatively simple. When ranking classifiers, it is important to keep in mind the accuracy metric can become misleading with classes imbalance in the training dataset (Xie et al., 2011) Therefore it is important to include e.g., precision and sensitivity metrics derived from the confusion matrix. The example in Figure 7-3 illustrates how the performance metrics of a Random Forest classifier change with different ratios between classes. Two classifiers were trained similarly on the dataset containing the same number of samples for rock mass types A, B, and C. However, the binary classes were composed differently: the first classifier predicts the rock mass type B (70 samples) vs. non-B (classes A and C, a total of 300 samples), and the second classifier predicts the rock mass type A (200 samples) vs. non-A (classes B and C, total 170 samples). The ratio between classes is 7:30 for the first classifier and

17:20 for the second. As shown in Figure 7-3, Classifier 2 exhibits increased overall performance over Classifier 1, despite Classifier 1 having better accuracy.

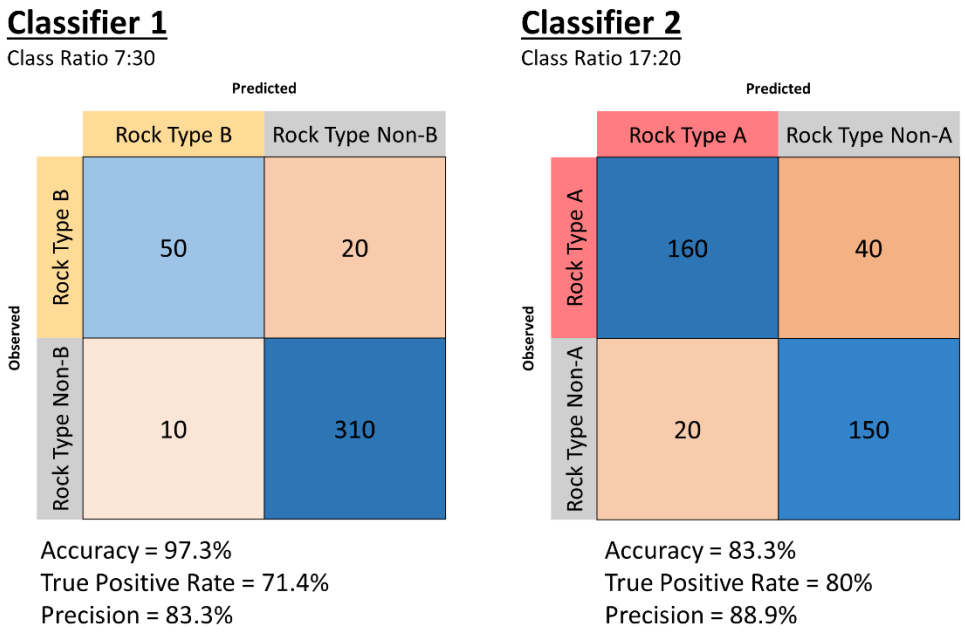

Figure 7-3. Example of changing performance metrics of a Random Forest classifier with different ratios between classes.

7.4.2 On Regression Algorithm Architecture Selection and Verification

In this section, recommendations on the selection and verification of algorithms for regression problems are given based on a case study where an artificial neural network (ANN) is used for Q-value estimation of tunnel boring machine (TBM) data. Regression algorithms approximate a mapping function from input variables (i.e., TBM data) to a continuous output variable (i.e., Q-values) (Figure 7-4 a.).

The goal of this case study is to train an ANN to determine Q-values on a TBM dataset from Norway's Ulriken Tunnel, excavated with an open gripper TBM. TBM operational data results from the sensor readings during operation of a TBM on a regular time interval. The resulting dataset is of sequential nature (i.e., time series), meaning that each datapoint represents an observation at a certain point in time and is dependent on the other observations in the dataset (Figure 7-4 a.). Based on the sequential nature of the data the implementation of a recurrent neural network (RNN) architecture was chosen to be the best fitting. RNN architectures are specifically designed to handle sequential data, able to take both current and previously received input data into account, thus being able to memorize earlier received input data (Figure 7-4 b.). A comparison of ANNs for a classification task based on TBM data is given in Erharter et al. (2019). Nevertheless, training of RNNs introduces problems like vanishing / exploding gradients (Pascanu et al., 2013), which can be overcome with strategies like Long Short-Term Memory (LSTM) architectures (Hochreiter & Schmidhuber, 1997).

The deployed RNN shown in Figure 7-4 a. consists of three LSTM layers with 32 units each (for further information on LSTMs see (Olah, 2015). The output layer consists of a fully connected linear layer with a ReLU activation function. To prevent the model from overfitting a dropout regularization method, where input and recurrent connections to the LSTM units are probabilistically excluded from activation during training, set to 0.05 has been implemented. For model optimization the Adam training algorithm was used.

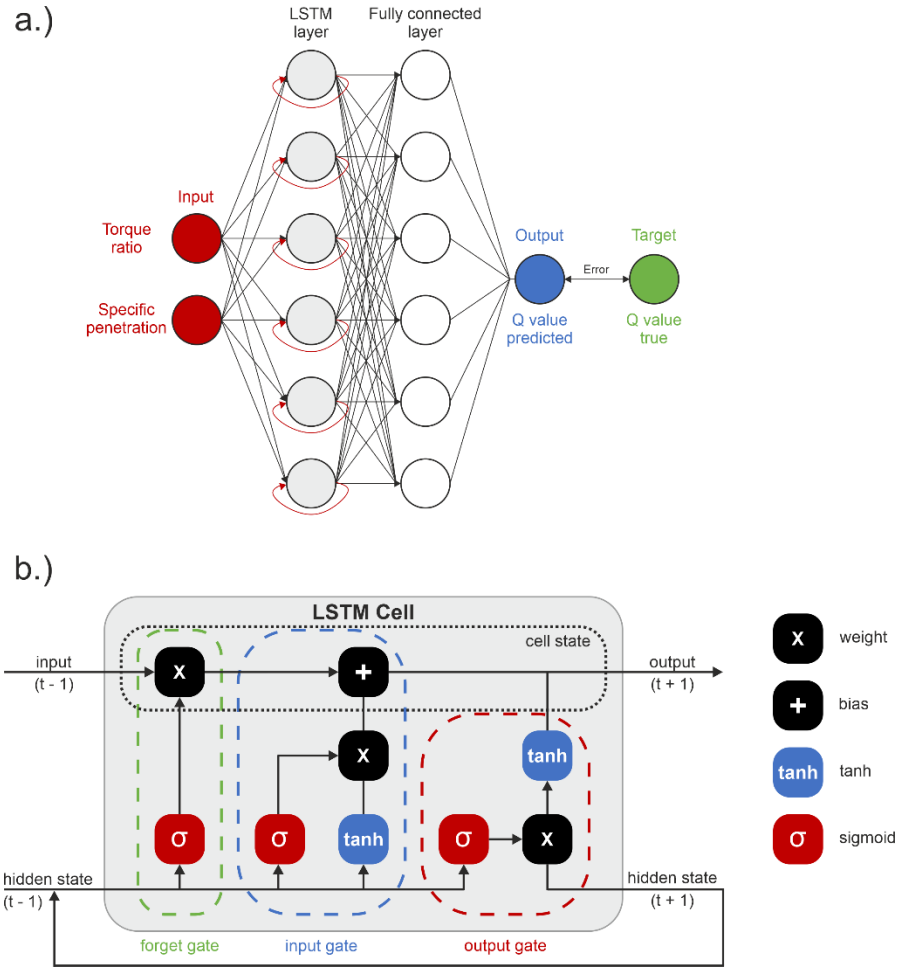

Figure 7-4. a) Graphical representation of the deployed LSTM model architecture. b) Schematic sketch of a LSTM cell with its four interacting layers, comprising of the cell state and its protecting and controlling gates (i.e., forget-, input and output gate).

Specific penetration and torque ratio are used as input features to the model, one set of input data consists of a matrix of these two over a certain length of the overall dataset. One input matrix for a length of 50m and a data-point spacing of 3 cm has the shape 1666x2. Output features consist of continuous Q-values. The model has been trained on data obtained from tunnel metre 1000 to 2000 and 3000 to 6900 and tested on data between tunnel metre 2000 and 3000. Results of the predicted Q-values for the geotechnically most relevant fault are shown in Figure 7-5.

Because a regression model predicts a quantity instead of a class or category as with classification models, the performance of the model must be reported as an error instead of an accuracy. To do so and to diagnose the variation in the errors, mean absolute error (MAE) and mean squared error (RMSE) are 0.12 and 0.18, respectively. Both metrics measure the average magnitude of the distance between prediction and target value and are negatively-oriented scores – lower values mean smaller differences. RMSE gives a relatively high weight to large errors, when using them together RMSE will always be larger than or equal to MAE, thus the difference between them reflects the variance of individual errors in the test set.

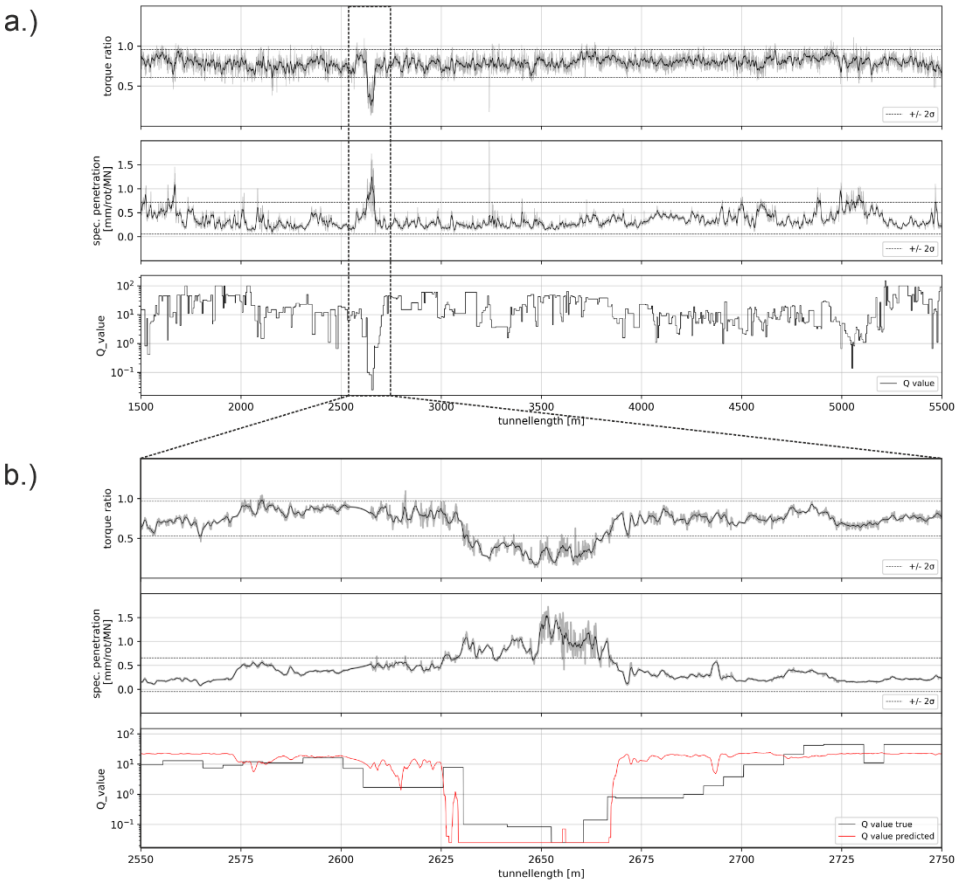

Figure 7-5 a) Example of sequential TBM data between tunnel metre 1500 and 5500; note the reaction of the features (torque ratio - first row and specific penetration - second row) as response to the fault (i.e. lowest Q-values in the third row). b) Tunnel metre 2550 to 2750: Input features (first and second row); comparison of Q-values, “ground truth” - human classification in black vs. Predicted Q-values in red (third row).

7.5 Input Data

The intuitive selection and manipulation of input parameters is a familiar concept in many forms of geomechanical modelling. For example, the geomechanical properties of a numerical model would be fine-tuned as a result of observed ground conditions and failure mechanisms (M. A. Perras et al., 2015c), or pre-processed to remove outliers and unreliable data, or to smooth data gaps. The field of machine learning

can offer repeatable methods for input data selection and pre-processing. This section highlights some such strategies for developing MLAs in rock engineering.

7.5.1 On Input Data Pre-processing and Balancing

Geomechanical datasets has been described as MUSIC-3X: Multivariate, Uncertain and Unique, Sparse, Incomplete, and potentially Corrupted with "3X" spatial data variations (Phoon & Chiang, 2021). As such, it is often necessary to balance input datasets prior or during MLA training. Balancing has the aim of ensuring that the MLA learns the nuanced relationships in the input dataset without being hindered by the data format or sparsity of the available inputs.

Categorical data, such as rock mass classifications, often need to be combined with numerical data within the MLA, and thus must be encoded into a numerical equivalent. One option is to use a numerical proxy, for example mean GSI for geotechnical zones. If this approach is not appropriate, ordinal or nominal encoding may be applied. Ordinal encoding assigns a relative rank to the classes, for example a weaker geotechnical zone would receive a lower rank than a stronger one. If no such ordinal ranking exists in the classes, then nominal ranking can be assigned via one-hot encoding. These encoding approaches are demonstrated in Table 7-1. The encoding approach applied should be carefully considered, as not to introduce noise during algorithm development.

Table 7-1. Example of categorical data encoding of geology for use in algorithm training (modified from Morgenroth et al., 2021)

Geology Class	Numerical Proxy (e.g., Using Mean GSI)	Ordinal Encoding	One-hot Encoding
Metasediments	28	0	1000
Dyke	65	1	0100
Metabasalt	70	2	0010
Norite	75	3	0001

Data pre-processing can have a significant impact on model performance. It is common, even necessary by some authors (Maier & Dandy, 2000), to apply an input standardization in order to ensure all inputs get equal attention during training (Z. Liu et al., 2021; Murphy, 2012). Variables should also be scaled in such a way as to be compatible with the limits of the activation functions in the output layer. The impacts of various standardization schemes are illustrated using a case study, where an LSTM network applied to a dataset from Garson Mine. The LSTM network is trained on microseismic event data and geomechanical parameters to predict the stress state in a FLAC3D model, where the performance was compared based on the corrected Akaike Information Criterion (AIC_c) (Akaike, 1969). Three input normalization schemes were tested against the base case of unnormalized inputs: z score (mean/standard deviation), rescaled to -1 to 1, and rescaled to 0 to 1. Rescaling does not impact the shape of the distributions of the inputs, while z score normalization includes the standard deviation and accounts for differences in relative distribution shape. This study of the Garson Mine LSTM found that models with unnormalized inputs and rescaling from 0 to 1 performed an order of magnitude worse in terms of AICc than those with z score normalization or

rescaling from -1 to 1. The z score normalization approach was selected to normalize the inputs because it produced the best LSTM network performance.

Imbalance in the input dataset can result in a model that is not well-generalized. Balancing may be built into the MLA itself by using random resampling with replacement during training (called bagging), or by applying error weighting schemes. The bagging algorithm generates subsets of data by randomly sampling from the original dataset, which are then used to train and evaluate each MLA (Breiman, 1996).

Error weighting is a technique used in classification to give particular targets preference during training by penalizing others (Seif, 2018). This is illustrated with a case study where a Convolutional Neural Network (CNN) was developed to classify tunnel liner yield at Cigar Lake Mine (Morgenroth et al., 2021). Four error weighting schemes were compared: uniform, linear, sigmoid, and inverse frequency (Figure 7-6). In this case, the inverse frequency scheme resulted in the best performance across all target classes.

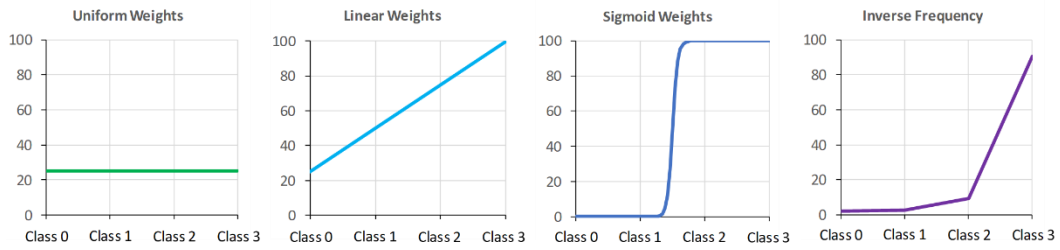

Figure 7-6. Four different error weight schemes applied to the classes in the training phase of the Cigar Lake Mine CNN (Morgenroth et al, 2021).

7.5.2 On Input Variable Selection

Input Variable Selection (IVS) methods typically identify the most useful inputs from a candidate pool of inputs used to develop an MLA. “Usefulness” of an input is defined as having the maximum relevance to the output, while minimizing the redundancy between the other inputs (May et al., 2011). Completing an IVS analysis of an MLA can result in less frequent convergence to local minima during algorithm training, thereby reducing the variability of the model output (May et al., 2008). In the context of rock engineering, IVS allows the user to rank the inputs with respect to their ability to accurately predict the output, i.e., which variables have the strongest link to the phenomena being predicted. An experienced geotechnical practitioner may choose to select input data based on expert judgement and can then confirm the selection using an IVS method. A novice practitioner may apply IVS from the outset of MLA development, to avoid overlooking critical input data. Both approaches are valid and can be adapted to the level of expert judgment that the MLA developer wants to inject.

Herein, two categories of IVS methods are discussed: *model-free* and *model-based*. Model-free methods aim to quantify the relevance and redundancy of individual input variables prior to MLA training, while model-based methods are embedded into the training of the MLA and measure the impact of a particular input on its performance.

Model-free methods distinctly separate the IVS task from MLA training, and instead make use of statistical techniques to measure the relevance of individual or combinations of inputs. An analogous example is determining the partial factors when designing using Eurocode 7 (Limit State Inc., 2013). There are several possible approaches to obtaining statistical ranks for the input parameters, for example: partial correlation, Pearson correlation, Spearman correlation, partial mutual information, among others. Model-based methods search through the available input variables during algorithm training and selects those that result in an optimal generalized performance of the trained MLA. This process differs from conventional geomechanical modelling, where, for example, all parameters for a given constitutive behaviour or rock mass classification scheme must be selected to compute the model. Model-based IVS methods offer the opportunity to decouple all the candidate inputs to investigate whether fewer inputs produce better MLA performance.

A model-free partial correlation (PC) IVS method is presented in Figure 7-7, where the available inputs are ranked based on partial correlation with the output, and then are iteratively added to the MLA in a forward selection process (He et al., 2011). The addition of inputs ceases when MLA performance has been satisfactorily fortified relative to the baseline. This method has the advantage of increasing the complexity (through the addition of inputs) of the MLA only if the performance improves as a result. A model-based IVS method called Input Omission (IO) is presented in Figure 7-7, sometimes also called *ablation*. IO estimates the usefulness of each input by iteratively examining model performance when an input is left out from the full set on which the model has been trained (Setiono & Liu, 1997). This approach allows the model to be examined for redundancy after it has been developed.

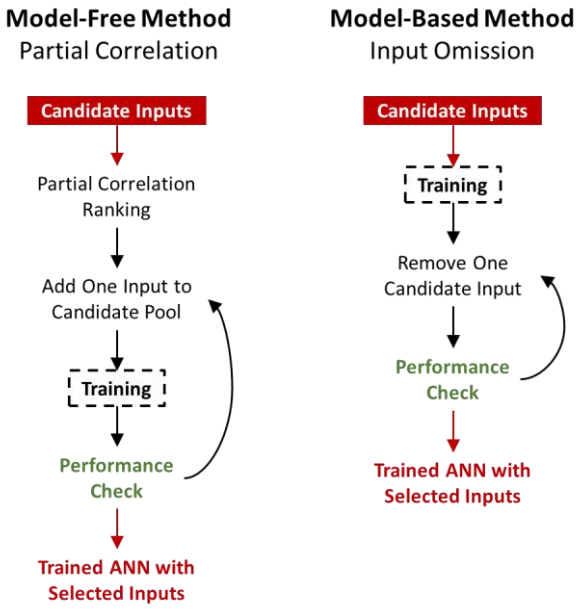

Figure 7-7. Workflows of model-free Input Variable Selection (IVS) method *Partial Correlation* (left), and model-based IVS method *Input Omission* (right).

These two categories of IVS methods are illustrated using the Cigar Lake Mine CNN case study discussed above. IVS approaches were applied to determine if there is redundancy in the inputs and to rank them in terms of their importance for predicting tunnel liner yield (Morgenroth et al., 2022). The PC method determined that the best performance was reached when the most recent DISP input was added to the candidate input pool. The IO method concluded that none of the inputs used to train the CNN could be omitted, and that the most sensitive inputs are GEO and DISP, indicated that these inputs should be measured with more spatial and temporal frequency to improve performance of the CNN.

IVS methods are powerful tools to interrogate MLAs developed for rock engineering problems, particularly to increase confidence in how the algorithm represents the physical system. IVS allows the user to determine how the inputs are being used to predict the output, and may also reveal the importance of the inputs relative to each other.

7.6 Conclusions

This chapter aims to demonstrate practical applications of machine learning to increase the confidence the rock engineering community has in applying these tools. The specific recommendations highlighted herein are summarized as follows:

1. The correct model type (e.g., classification vs regression) and sub-type (specific MLA architecture) must be chosen.
    a) Classification – start with a simple “white-box” method and increase complexity as required

   b) Regression – ensure the selected cost function (e.g., RMSE, MAE) is appropriate for the desired target
2. Input data must be balanced (e.g., oversampling, bagging, error weighting) so that the MLA can learn the patterns represented in the input dataset without inadvertently emphasizing overrepresented data.
3. The trained MLA should be analyzed using IVS techniques to remove redundant data, or to rank inputs so those ranking highest can be captured at higher spatial or temporal resolution.

Machine learning is a powerful tool that frees up rock engineers to interrogate their assumptions and ensure the correct model is built (validation), and that the model is also built correctly (verification). The reluctance to apply machine learning methods in applied rock engineering is justified insofar as it highlights the need for critical thought when evaluating the algorithm outputs in the practical context. As with any modelling technique, the model's success is not only predicated on its ability to replicate observed ground behaviour, but also its transparency in how it is using the available data to do so. Techniques to interrogate machine learning exist in data science and adjacent fields to render algorithms interpretable, all we need to do is seek them out and adapt them to advance digital innovation in rock engineering.

CHAPTER 8. CONCLUSIONS

"The situation has provided a cue; this cue has given the expert access to information stored in memory, and the information provides the answer. Intuition is nothing more and nothing less than recognition."

– Thinking, Fast and Slow (Kahneman, 2011)

8.1 Summary of Research

The overall objective of this research was to develop and apply practical machine learning solutions to real geomechanical datasets to predict underground excavation and rock mass behaviour. Four specific objectives were reached:

1. Develop a classification machine learning algorithm using standard geotechnical mapping data from a real project
2. Develop and implement a methodology for assessing interpretability of machine learning algorithms applied to rock engineering
3. Develop a regression machine learning algorithm using standard geotechnical sensor data from a real project
4. Formulate a guide for rock engineering practitioners in the selection, development, and engineering verification of machine learning algorithms for underground rock engineering problems

The major conclusions of each objective are provided below.

1. *Develop a classification machine learning algorithm using standard geotechnical mapping data from a real project*

A major obstacle when using geotechnical data for machine learning is its inherent complexity, sparsity, and the often incomplete nature of the datasets (Elmo & Stead, 2020; Jing & Hudson, 2002; Morgenroth et al., 2019; Phoon et al., 2019). In *CHAPTER 4. A Convolutional Neural Network approach for predicting tunnel liner yield at Cigar Lake Mine*, the original format of geotechnical mapping data is preserved by formatting the raw data as images with the machine learning inputs comprising the channels of the image. A CNN was developed to predict tunnel liner yield using these input images, using the case study of Cigar Lake Mine, Saskatchewan, Canada. CNNs are designed to process spatial and temporal dependencies in image or raster datasets and are computationally efficient at processing each pixel in an image while considering both the surrounding pixels and their change over time. When developing any kind of rock mechanics model, the spatial and temporal relationships between the various inputs must be preserved, and thus, CNNs are well suited for these types of datasets. The data usually contains a mapping component

(e.g., mapped geology) onto which other data (e.g., measured displacement or groundwater inflow) can be transferred as a two-dimensional array and therefore plotted as an image or a map.

A CNN is developed to predict tunnel liner yield in the squeezing ground conditions of Cigar Lake Mine using four geological and geomechanical inputs: mapped geotechnical zones (GEO), primary installed support class (SUPCL), ground freezing (FREEZE), and radial tunnel displacement (DISP). These data were obtained from the GMPs prepared by the mine and from tunnel liner displacement surveys. The Cigar Lake Mine CNN was developed through a sensitivity analysis of three network hyperparameters: the amount of training data, the convolution filter size, and the error weighting scheme. A total of 7200 models were analyzed to determine the best combination of these parameters for this problem.

It was found that the best CNN classification performance was obtained when all available data in the temporal realm was used for training, and the convolution filter was 30 x 30 pixels. An error weighting scheme was necessary to offset the imbalance of samples for each class in the dataset. An inverse frequency error weighting scheme, where the proportion of samples in each class is used to determine the error weights, was used to develop a Balanced Global model that has reasonable prediction accuracy across all yield classes (average model accuracy >65% for all training data permutations). A sigmoid error weighting scheme, which prioritizes Class 2 and 3 errors during training, was used to develop a Targeted Class 2/3 model that has higher recall for Class 2 tunnel liner yield (>99%).

2. *Develop and implement a methodology for assessing interpretability of machine learning algorithms applied to rock engineering*

This dissertation contains an exploration of algorithm interpretability using Input Variable Selection (IVS) approaches. The developed IVS approaches address two categories of algorithm interpretability: (1) algorithm transparency, and (2) post-hoc interpretations. IVS allows the user to analyze the internal mechanics of algorithms by ranking the candidate inputs, determining their usefulness to predict the output, reducing redundancy within the model, and comparing the findings with the experience and knowledge of fundamental rock mechanics principles.

In *CHAPTER 5. On the Interpretability of Machine Learning Using Input Variable Selection: Forecasting Tunnel Liner Yield*, three IVS methods were compared for two CNN architectures developed to predict tunnel liner yield at the Cigar Lake Mine – the Targeted Class 2/3 model and the Global Balanced model. For both models the three IVS approaches employed were Channel Activation Strength (CAS), Input Omission (IO), and Partial Correlation (PC).

All three IVS methods found that additional training GMPs increased the prediction accuracy of the models. The CAS results suggested that GEO is used the least by the CNN as compared to the other three inputs, SUPCL, FREEZE, and DISP, and that SUPCL has the high activation strength in the CNN. The IO approach determined that none of the inputs could be omitted from the model entirely. The IO results showed that removing GEO entirely caused the model performance to decrease most significantly as compared to

omitting the other inputs, indicating that GEO had some unique information that the other inputs were not able to convey to the CNN. Finally, the PC approach was developed to determine if the models could be optimized by using only the most highly correlated inputs to train the CNN, according to their partial correlation with the output. PC demonstrated that where there was only one training GMP the DISP input was selected first, meanwhile where there was more than one GMP used for training the categorical inputs (GEO, SUPCL, FREEZE) were chosen first. The PC results showed that the minimum AIC_c was obtained when the most recent DISP was added to the candidate pool. The PC results also indicated that no matter how many training GMPs were used, it was not possible to drop out the older ones completely.

The IVS approaches completed as part of this objective indicated that it is good practice to include all data in initial algorithm development for a rock engineering problem, and then only after completing an IVS study remove the inputs that are not contributing to good model performance. This research demonstrated that no single IVS method should be used to determine input saliency. Instead, multiple IVS methods must be applied to increase algorithm interpretability.

3. *Develop a regression machine learning algorithm using standard geotechnical sensor data from a real project*

Datasets gathered using digital rock mass instrumentation do not suffer from sparsity but rather pose challenges with respect to large volumes of data that must be processed before meaning can be extracted. This completed objective involved the development of a LSTM network using microseismic data at Garson Mine, where millions of microseismic events were available for algorithm training. The process of formatting the input data and optimizing the LSTM network architecture and hyperparameters are presented in *CHAPTER 6.*

The LSTM network was developed using microseismic data and a FLAC3D model as inputs, where the targets were the stresses at the zone centroids of the FLAC3D model for the 1SHW area at Garson Mine. The inputs were formatted into sequences of microseismic events that could be assigned to a resulting stress state extracted from a FLAC3D model manually calibrated previously. Two LSTM network architectures were developed for two sets of targets: the principal stresses and the six-component stress tensors. Various hyperparameters were analyzed to determine the optimal architecture for the two sets of targets: input encoding and pre-processing, training solver, network layer architecture, and cost function.

Both optimized LSTM networks (targets = principal stresses and targets = six-component stress tensors) had inputs normalized by their z score, MSE as the cost function, an SGD training solver, and one LSTM layer. The main difference was that where targets = principal stresses, the geology was encoded using the maximum and minimum GSI, and where targets = six-component stress tensors the geology was encoded using the mean GSI. In terms of AIC_c, where a lower value indicates better performance, the performance of the LSTM network predicting the three principal stresses was improved as compared to predicting the six-component stress tensor, with AIC_c values of -59.62 and -45.50, respectively. The R^2 values were 0.996

and 0.997, respectively, and the %C values were 97% and 80%, respectively. The principal stresses computed by FLAC3D fell within the range of stresses predicted by the 100-model LSTM network ensemble.

The completion of this objective demonstrates semi-automatic numerical model recalibration using a machine learning algorithm and represents progress towards establishing a dynamic interaction between a FLAC3D model and the LSTM network to automate the recalibration process.

4. *Formulate a guide for rock engineering practitioners in the selection, development, and engineering verification of machine learning algorithms for underground rock engineering problems*

A guide of tools and processes for developing machine learning for practical rock engineering problems has been prepared. The guide focuses on architecture selection for classification and regression problems, including the selection of appropriate engineering verification metrics. Recommendations with respect to input data pre-processing, balancing, and selection are also made. The tools and processes covered in this guide were chosen based on experience gathered completing Objectives 1-3.

In general, the research contained in this dissertation demonstrates the practical applicability of machine learning algorithms to rock engineering problems. A review of the state of the art of machine learning in underground rock mechanics research literature is presented. Two different algorithms are developed, one for classification and one for regression, for two end member rock mass deformation mechanics, time-dependent squeezing and high stress rock bursting. A guide of common tools and processes in machine learning that can be applied to develop geomechanical is presented. This dissertation has demonstrated that good performance and predictive ability can be achieved when developing machine learning for geomechanical datasets, provided that the algorithms are developed within the context of the physical system for which they are being developed.

8.2 Novel Contributions

This dissertation explored the development of machine learning algorithms for practical rock engineering problems, including a squeezing ground environment and a high-stress environment. The novel contributions include:

1. a review of the state of the art of machine learning applications applied to underground rock engineering, which did not exist in the research literature previously
2. a novel application of a Convolutional Neural Network to a squeezing ground environment to predict tunnel liner yield, where the original format of the mapping data is preserved as images to reduce the introduction of bias during the data digitalization process
3. a method to increase algorithm interpretability in rock engineering using Input Variable Selection methods, where two novel methods have been developed for Convolutional Neural Networks

4. a semi-automatic method of recalibrating a finite difference model using a Long-Short Term Memory network and a database of microseismic events, whereby the stresses in the numerical model are predicted as new microseismic events occur
5. a guide of best practices for developing machine learning solutions for underground rock engineering problems

8.3 Limitations of Research

A common response to this research, both when presented at conferences and when submitted for peer-reviewed publications, is that algorithms are “black boxes” or that “machine learning will take our jobs,” where the latter is often said in half-jest. These comments encapsulate the limitations of this dissertation, and indeed other research at the intersection of machine learning and rock engineering.

Machine learning is not a substitute for a geotechnical engineer having an understanding of the rock mass characteristics and deformation mechanisms. For the algorithms developed in this dissertation, the prevalent failure mechanisms in the respective mining environments were well understood by the mine operators, and this knowledge was transferred during research meetings. When this is not the case, effort must be expended to characterize the rock mass behaviour and ensure the appropriate data is available before developing a machine learning algorithm. If not enough data is available to develop an algorithm, geotechnical engineering judgment should be used to develop a site investigation and/or laboratory testing program to collect additional data.

Even the most optimized machine learning algorithm cannot make up for poor input data quality. In the research presented in this dissertation, the datasets were collected in line with standard practice and checked for consistency with the observed rock mass deformation phenomena by the industry partners that provided the data. In other words, the algorithms were able to provide insights for the mine operators because the available data contained the relevant correlations that could be extracted by machine learning. If this is not the case, significant data cleaning and pre-processing is necessary, which is not in the scope of this dissertation. It is also possible that additional field and lab data must be collected.

As has been eloquently described by other authors, for example Elmo & Stead (2021), Marcher et al. (2020), and Yang et al. (2021), there is a temptation to take the output of a machine learning algorithm as veritable fact without interrogating its blind spots and potential short comings. Machine learning is one of many tools available in the rock engineering toolbox and is not infallible. Just as numerical models are examined to verify their outputs, algorithms must undergo analogous investigation after they are developed. An example of this is the IVS analysis presented in *CHAPTER 5.* . Further work is needed to standardize the engineering verification and interpretability of machine learning for rock engineering practice, whereby the confidence of their use in designing underground excavations can be increased.

8.4 Future Work

The results presented in this dissertation may be used as a basis to continue to enhance the practical scope of the machine learning algorithms developed for rock engineering. Future research could focus on:

1. *Further work to refine figure on continuum of algorithm types for rock engineering*

As part of the literature review presented in CHAPTER 2. An Overview of Opportunities for Machine Learning Methods in Underground Rock Engineering Design, Figure 2-9 was developed to illustrate to continuum of machine learning algorithms as a function of data redundancy and quantity (reproduced below for clarity). In order to be more practically useful, future work on this figure could involve a meta analysis of research completed to date at the intersection of machine learning and rock engineering to quantify the amount of data and data redundancy, and also to overlay the stage of engineering design the algorithms were successful for (i.e, pre-feasibility, feasibility, detailed design).

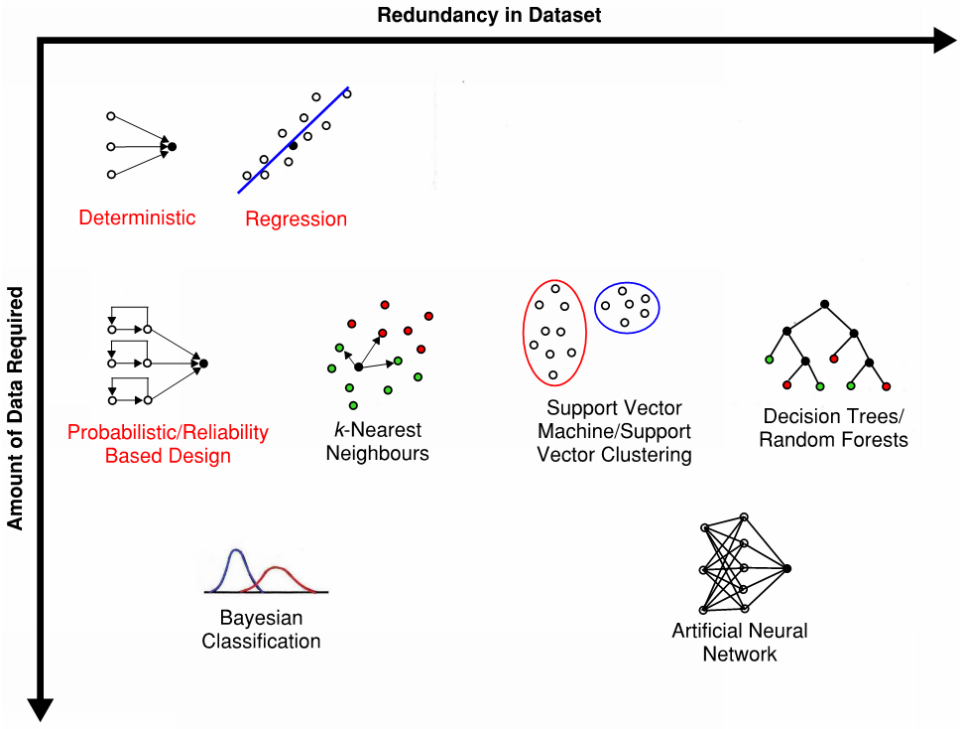

Figure 2-9. Comparison of data needs for machine learning algorithms included in this review (black) and some conventional rock engineering methods (red) in terms of amount of data and data redundancy required. Data redundancy indicates the representation of samples of the behaviour the data method should capture.

2. *Further optimization of the Cigar Lake Mine CNN*

Given the small number of inputs used to develop the Cigar Lake CNN, future work could consider the additional digitalization of inputs from the GMPs. The inputs digitalized in this dissertation were chosen based on discussion with the mine operators at Cigar Lake Mine, however there is potential to further increase the tunnel liner yield prediction accuracy by increasing the input candidate pool.

The CNN in this dissertation was developed by conducting a sensitivity analysis of a select subset of hyperparameters. This CNN could be further refined using an exploration of an automatic hyperparameter tuning approach, such as the tree-structured Parzen estimator (Zhou et al., 2021) or Optuna (Akiba et al., 2019). Such approaches systematically tune hyperparameters using a dynamic search space, as opposed to grid searches such as the one completed in this dissertation. These automatic tuning approaches may prove more computationally efficient and could increase confidence that a global solution has been converged upon during CNN training.

3. *Further development of the Garson Mine LSTM network*

Future work using the Garson Mine dataset and the LSTM network developed in this dissertation could be approach in various ways. The candidate input dataset could be adapted to include the raw geotechnical mapping data as opposed to the geomechanical parameters from the FLAC3D model. Other authors have suggested decomposing classification methods and using their components in modelling, for example developing algorithms to use rock mass classification inputs to predict failure phenomena or excavation performance directly instead of arriving at a single classification value (Yang et al., 2021). The Garson Mine dataset could be further enriched by increasing the number of sequences available for training by both reducing the minimum number of microseismic events in a sequence, and by applying a bootstrapping algorithm to sample with replacement during training.

Improvements to the Garson Mine LSTM network could involve refinement of the model efficiency, particularly when predicting the six-component stress tensor. An error weighting scheme could be applied to ensure good performance across all six targets during algorithm training. IVS methods may also be applied to determine if all the microseismic parameters are needed for good model performance, or if these may be reduced to improve performance and reduce computational time. Research could also be expanded to other algorithm types, such as the recently developed hybrid CNN-LSTM model (Shi et al., 2015), or an unsupervised method, such as Self-Organizing Maps.

4. *Integrating the Garson Mine LSTM network into a larger hazard prediction framework*

The Garson Mine LSTM in this dissertation could be used to develop a larger hazard prediction framework, whereby the LSTM network is used to automatically recalibrate the FLAC3D model, and then the outputs of the recalibrated FLAC3D model are used in a second LSTM network that forward predict microseismic event parameters (e.g., moment magnitude and location). These outputs could be used to create a hazard map of future seismogenic zones, which in turn could be used to plan future mine developments and stope sequencing.

5. *Investigate combination of statistical analysis and machine learning.*

Several publications over the past decade have questioned the validity of reducing qualitative rock mass classifications to probabilistic distributions, especially without respecting their inherent limitations (Barton &

Bieniawski, 2008; Yang et al., 2021). To this end, future research should investigate the sensitivity of machine learning algorithms to sampling from true input data populations as opposed to relying on a probability distribution function that represents those qualitative parameters. Additionally, these characterizations are subjective and are known to contain a high degree of uncertainty. Fuzzy logic or Bayesian approaches could be considered to incorporate the inherent uncertainty in these inputs.

6. *Formalize guidelines for quantifying uncertainty when using machine learning in rock engineering*

It is well understood that engineering epistemology can pose limits on the usefulness of various types of models in rock engineering, and this is also true of data driven models such as machine learning algorithms (Elmo et al., 2022). Future work should explore the impacts the uncertainty in the input candidates has on algorithm performance, such as by applying probabilistic methods or constructing fuzzy numbers from the raw inputs.

Ensemble modelling methods could also be explored in more detail. In this dissertation, ensemble sizes were chosen based on model convergence and computational speed, using the rule of thumb that ensembles should have at least 30 members (Setiono & Liu, 1997). Ensemble learning models could be explored in future work, whereby predictions are updated by voting on the next iteration which is more computationally efficient since models become more accurate during training. Ensemble learning models can resolve three types of issues: representation challenges with high bias, computational challenges with high computational variance, and statistical challenges with high variance. (Zhou et al., 2021).

REFERENCES

Abbasi, M., & El Hanandeh, A. (2016). Forecasting municipal solid waste generation using artificial intelligence modelling approaches. *Waste Management*, *56*, 13–22. https://doi.org/10.1016/j.wasman.2016.05.018

Abrahart, R. J., Anctil, F., Coulibaly, P., Dawson, C. W., Mount, N. J., See, L. M., Shamseldin, A. Y., Solomatine, D. P., Toth, E., & Wilby, R. L. (2012). Two decades of anarchy? Emerging themes and outstanding challenges for neural network river forecasting. *Prog Phys Geog*. https://doi.org/10.1177/0309133312444943

Afraei, S., Shahriar, K., & Madani, S. H. (2019). Developing intelligent classification models for rock burst prediction after recognizing significant predictor variables, Section 1: Literature review and data preprocessing procedure. *Tunn and Undergr Sp Tech*, *83*, 324–353. https://doi.org/10.1016/j.tust.2018.09.022

Akaike, H. (1969). Fitting autoreggressive models for prediction. *Ann I Stat Math*, *21*, 243–247.

Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A Next-generation Hyperparameter Optimization Framework. *KDD '19: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2623–2631. http://arxiv.org/abs/1907.10902

Alghalandis Computing. (2019). *ADFNE* (1.1). https://alghalandis.net/products/adfne/

Araghinejad, S. (2014). Chapter 5 : Artificial Neural Networks. In *Data-Driven Modeling : Using MATLAB ® in Water Resources and Environmental Engineering* (Vol. 67). https://doi.org/10.1007/978-94-007-7506-0

Aydan, Ö., Akagi, T., & Kawamoto, T. (1993). The squeezing potential of rock around tunnels: theory and prediction. *Rock Mech Rock Eng*, *2*, 137–163.

Baecher, G. B., & Christian, J. T. (2003). *Reliability and statistics in geotechnical engineering*. John Wiley & Sons, Ltd.

Barla, G. (2002). Tunnelling under squeezing rock conditions. In D. Kolymbas (Ed.), *Tunnelling Mechanics - Advances in Geotechnical Engineering and Tunnelling* (Vol. 5, pp. 169–268). Springer Science & Business Media.

Barla, G., Bonini, M., & Semeraro, M. (2011). Analysis of the behaviour of a yield-control support system in squeezing rock. *Tunn and Undergr Sp Tech*, *26*(1), 146–154. https://doi.org/10.1016/j.tust.2010.08.001

Barla, G., & Borgna, S. (1999). Tunnelling in Squeezing Rock Conditions. *Proceedings of ROCKSITE-99*, 97–108.

Barton, N. R., & Bieniawski, Z. T. (2008). RMR and Q - Setting records straight. *Tunnels & Tunnelling International*, *February*.

Barton, N. R., Lien, R., & Lunde, J. (1974). Engineering classification of rock masses for the design of tunnel support. *Rock Mechanics*, *6*(4), 189–236.

Bell, J. (2015). Chapter 5: Artificial Neural Networks. In *Machine Learning : Hands-on for Developers and Technical Professionals* (2nd Editio). John Wiley & Sons (US). https://doi.org/10.1002/9781119183464

Benardos, A. (2008). Artificial intelligence in underground development: A study of TBM performance. *WIT Transactions on the Built Environment*, *102*, 21–32. https://doi.org/10.2495/US080031

Ben-Hur, A., Horn, D., Siegelmann, H., & Vapnik, V. (2001). Support Vector Clustering. *Journal of Machine Learning Research*, *2*, 125–137. https://doi.org/10.1109/IJCNN.2014.6889518

Bentley. (2019). *Plaxis 2D* (2019.00).

Bieniawski, Z. T. (1993). Chapter 22: Classification of rock masses for engineering: the RMR system and future trends. In H. Konietzky (Ed.), *Comprehensive Rock Engineering* (pp. 553–573).

Bishop, S., Mainville, A., & Yesnik, L. (2016). *Cigar Lake Operation, National Instrument 43-101 Technical Report*.

Bizjak, K. F., & Petkovšek, B. (2004). Displacement analysis of tunnel support in soft rock around a shallow highway tunnel at Golovec. *Eng Geol*, *75*(1), 89–106. https://doi.org/10.1016/j.enggeo.2004.05.003

Bozorg-Haddad, O., Soleimani, S., & Loáiciga, H. A. (2017). Modeling Water-Quality Parameters Using Genetic Algorithm-Least Squares Support Vector Regression and Genetic Programming. *Journal of Environmental Engineering (United States)*, *143*(7), 1–10. https://doi.org/10.1061/(ASCE)EE.1943-7870.0001217

Bozorgzadeh, N., Escobar, M. D., & Harrison, J. P. (2018). Comprehensive statistical analysis of intact rock strength for reliability-based design. *Int J of Rock Mech & Min Sci*, *106*(May 2017), 374–387. https://doi.org/10.1016/j.ijrmms.2018.03.005

Breiman, L. (1996). Bagging predictors. *Machine Learning*, *24*(2), 123–140.

Brown, L. G., Hudyma, M. R., Carusine, O., & Reimer, E. (2020, June). Evaluating Short Term Seismic Hazard Forecasting in Hard Rock Mines. *ISRM International Symposium - EUROCK 2020*.

Burnham, K. P., & Anderson, D. R. (2002). *Model Selection and Multi-Model Inference: A Practical Information-Theoretic Approach* (2nd ed.). Springer-Verlag New York. https://doi.org/10.1007/b97636

Chen, R., Zhang, P., Wu, H., Wang, Z., & Zhong, Z. (2019). Prediction of shield tunneling-induced ground settlement using machine learning techniques. *Frontiers of Structural and Civil Engineering*, *13*(6), 1363–1378. https://doi.org/10.1007/s11709-019-0561-3

Ching, J., Phoon, K., Li, K., & Weng, M. (2019). *Multivariate probability distribution for some intact rock properties*. *18*(October 2018), 1–18.

Chou, J. S., & Thedja, J. P. P. (2016). Metaheuristic optimization within machine learning-based classification system for early warnings related to geotechnical problems. *Automation in Construction*, *68*, 65–80. https://doi.org/10.1016/j.autcon.2016.03.015

Claesen, M., & de Moor, B. (2015). Hyperparameter Search in Machine Learning. *The XI Metaheuristics International Conference*, 10–14. http://arxiv.org/abs/1502.02127

Cortolezzis, D., & Hudyma, M. (2018). Application of sequential spatial clustering and fractal dimension to caving seismic event parameters of time, distance, and intensity. *Caving 2018: Proceedings of the Fourth International Symposium on Block and Sublevel Caving*, 799–814. https://doi.org/10.36487/acg_rep/1815_63_cortolezzis

Cracknell, M. J., & Reading, A. M. (2014). Geological mapping using remote sensing data: A comparison of five machine learning algorithms, their response to variations in the spatial distribution of training data and the use of explicit spatial information. *Computers and Geosciences*, *63*, 22–33. https://doi.org/10.1016/j.cageo.2013.10.008

Dassault Systems. (2019). *ABAQUS* (No. 2019).

Daszykowski, M., Walczak, B., & Massart, D. L. (2002). Representative subset selection. *Analytica Chimica Acta*, *468*(1), 91–103. https://doi.org/10.1016/S0003-2670(02)00651-7

de Santi, M., Khan, U. T., Arnold, M., Fesselet, J. F., & Ali, S. I. (2021). Forecasting point-of-consumption chlorine residual in refugee settlements using ensembles of artificial neural networks. *Clean Water*, *4*(1), 1–16. https://doi.org/10.1038/s41545-021-00125-2

Deere, D. U. (1963). Technical description of rock cores for engineering purposes. *Mech.Eng. Geol.*, *1*(18–22).

Delisio, A., Zhao, J., & Einstein, H. H. (2013). Analysis and prediction of TBM performance in blocky rock conditions at the Lötschberg Base Tunnel. *Tunnelling and Underground Space Technology*. https://doi.org/10.1016/j.tust.2012.06.015

Deng, Y., Sadiq, R., Jiang, W., & Tesfamariam, S. (2011). Risk analysis in a linguistic environment: A fuzzy evidential reasoning-based approach. *Expert Systems with Applications*, *38*(12), 15438–15446. https://doi.org/10.1016/j.eswa.2011.06.018

Diederichs, M. S. (2007). The 2003 Canadian geotechnical colloquium: mechanistic interpretation and practical application of damage and spalling prediction criteria for deep tunnelling. *Canadian Geotechnical Journal*, *44*, 1082–1116. https://doi.org/10.1139/T07-033

Dong, L. J., Li, X. B., & Peng, K. (2013). Prediction of rockburst classification using Random Forest. *Transactions of Nonferrous Metals Society of China (English Edition)*, *23*(2), 472–477. https://doi.org/10.1016/S1003-6326(13)62487-5

Dong, L., Li, X., Xu, M., & Li, Q. (2011). Comparisons of random forest and Support Vector Machine for predicting blasting vibration characteristic parameters. *Procedia Engineering*, *26*, 1772–1781. https://doi.org/10.1016/j.proeng.2011.11.2366

Duan, Y., Shen, Y., Canbulat, I., Luo, X., & Si, G. (2021). Classification of clustered microseismic events in a coal mine using machine learning. *J Rock Mech Geotech Eng*, *13*(6), 1256–1273. https://doi.org/10.1016/j.jrmge.2021.09.002

EDF. (2019). *Code_Aster* (No. 14).

Einstein, H. H., Indermitte, C., Sinfield, J., Descoeudres, F. P., & Dudt, J.-P. (1999). Decision Aids for Tunneling. *Transportation Research Record*, *1656*, 6–13.

Einstein, H. H., Labreche, D. A., Markow, M. J., & Baecher, G. B. (1978). Decision analysis applied to rock tunnel exploration. *Engineering Geology*, *12*, 143–161. https://doi.org/10.1016/0013-7952(78)90008-X

Elmo, D., Mitelman, A., & Yang, B. (2022). Examining Rock Engineering Knowledge through a Philosophical Lens. *Geosciences*, *12*(4), 174. https://doi.org/10.3390/geosciences12040174

Elmo, D., & Stead, D. (2020). Disrupting rock engineering concepts: is there such a thing as a rock mass digital twin and are machines capable of learning rock mechanics? *Proceedings of the 2020 International Symposium on Slope Stability in Open Pit Mining and Civil Engineering*, 565–576. https://doi.org/doi.org/10.36487/ACG_repo/2025_34

Elmo, D., & Stead, D. (2021). The Role of Behavioural Factors and Cognitive Biases in Rock Engineering. *Rock Mech Rock Eng*, *54*, 2109–2128. https://doi.org/doi.org/10.1007/s00603-021-02385-3

Elmo, D., Stead, D., Yang, B., Tsai, R., & Fogel, Y. (2020). Can new technologies shake the empirical foundations of rock engineering ? In J. Wesseloo (Ed.), *Proceedings of the Second International*

Conference on Underground Mining Technology (pp. 107–116). Australian Centre for Geomechanics. https://doi.org/doi.org/10.36487/ACG_repo/2035_01

Erharter, G. H., & Marcher, T. (2020). MSAC: Towards data driven system behavior classification for TBM tunneling. *Tunn and Undergr Sp Tech*, *103*(May). https://doi.org/10.1016/j.tust.2020.103466

Erharter, G. H., Marcher, T., & Reinhold, C. (2019). Comparison of artificial neural networks for TBM data classification. *Proceedings of the 14th International Congress on Rock Mechanics and Rock Engineering*.

Fakir, M., & Ferentinou, M. (2017). A holistic open pit mine slope stability index using artificial neural networks. In *ISRM AfriRock - Rock Mechanics for Africa* (Vols. 2017-Octob).

Farahmand, K., Vazaios, I., Diederichs, M. S., & Vlachopoulos, N. (2018). Investigating the scale-dependency of the geometrical and mechanical properties of a moderately jointed rock using a synthetic rock mass (SRM) approach. *Computers and Geotechnics*, *95*, 162–179. https://doi.org/10.1016/j.compgeo.2017.10.002

Fathipour-Azar, H. (2021). Data-driven estimation of joint roughness coefficient. *Journal of Rock Mechanics and Geotechnical Engineering*, *13*(6), 1428–1437. https://doi.org/10.1016/j.jrmge.2021.09.003

Ferentinou, M., & Fakir, M. (2018). Integrating Rock Engineering Systems device and Artificial Neural Networks to predict stability conditions in an open pit. *Eng Geo*, *246*, 293–309. https://doi.org/10.1016/j.enggeo.2018.10.010

Gao, W., Lu, X., Peng, Y., & Wu, L. (2020). A Deep Learning Approach Replacing the Finite Difference Method for in Situ Stress Prediction. *IEEE Access*, *8*, 44063–44074. https://doi.org/10.1109/ACCESS.2020.2977880

Gao, X., Shi, M., Song, X., Zhang, C., & Zhang, H. (2019). Recurrent neural networks for real-time prediction of TBM operating parameters. *Automat Constr*, *98*, 225–235. https://doi.org/10.1016/j.autcon.2018.11.013

Garza-Cruz, T., Pierce, M., & Kaiser, P. K. (2014). Use of 3DEC to study spalling and deformation associated with tunnneling at depth. *Deep Mining 2014. 7th International Conference on Deep and High Stress Mining*, 421–436.

Geological Survey of Canada. (2009). *Atlas of Canada* (6th Edition). National Resources Canada. https://open.canada.ca/data/en/dataset/dbe45821-8893-11e0-a8da-6cf049291510

Geomechanica. (2019). *Irazu* (No. 2019).

GEO-SLOPE. (2016). *GeoStudio* (8.16.0).

Geraili Mikola, R. (2019). *Adonis*.

Goel, R. K., Jethwa, J. L., & Paithakan, A. G. (1995). Tunnelling through the young Himalayas – a case history of the Maneri-Uttarkashi power tunnel. *Eng Geol*, *39*, 31–44.

Golder Associates. (2001). *Numerical Analysis to Estimate Stresses in Crosscut Linings at Cigar Lake Mine*.

Golder Associates Inc. (2019). *FracMan Geotechnical Edition* (2.6).

Grasselli's Geomechanics Group. (2019). *Y-GEO*.

Hasanpour, R., Rostami, J., Thewes, M., & Schmitt, J. (2018). Parametric study of the impacts of various geological and machine parameters on thrust force requirements for operating a single shield TBM in squeezing ground. *Tunn Undergr Sp Tech*, *73*, 252–260. https://doi.org/10.1016/j.tust.2017.12.027

He, J., Valeo, C., Chu, A., & Neumann, N. (2011). Prediction of event-based stormwater runoff quantity and quality by ANNs developed using PMI-based input selection. *Journal of Hydrology*, *400*(1–2), 10–23. https://doi.org/10.1016/j.jhydrol.2011.01.024

Hibert, C., Provost, F., Malet, J. P., Maggi, A., Stumpf, A., & Ferrazzini, V. (2017). Automatic identification of rockfalls and volcano-tectonic earthquakes at the Piton de la Fournaise volcano using a Random Forest algorithm. *Journal of Volcanology and Geothermal Research*, *340*, 130–142. https://doi.org/10.1016/j.jvolgeores.2017.04.015

Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, *9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Hoek, E. (1966). Rock Mechanics - an introduction for the practical engineer Parts I, II and III. *Mining Magazine*, 1–67.

Hoek, E., & Brown, E. T. (1997). Practical estimates of rock mass strength. *Int J of Rock Mech & Min Sci1*, *34*, 1165–1186.

Hoek, E., & Brown, E. T. (2019). The Hoek–Brown failure criterion and GSI – 2018 edition. *Journal of Rock Mechanics and Geotechnical Engineering*, *11*(3), 445–463. https://doi.org/10.1016/j.jrmge.2018.08.001

Hoek, E., & Diederichs, M. S. (2006). Empirical estimation of rock mass modulus. *Int J of Rock Mech & Min Sci*, *43*(2), 203–215. https://doi.org/10.1016/j.ijrmms.2005.06.005

Hoek, E., Grabinsky, M. W., & Diederichs, M. S. (1990). Numerical Modelling for Underground Excavation Design. *Trans Inst Min Metall., Section A: Mining Industry*, 22–30.

Hoek, E., & Marinos, P. (2000). Predicting tunnel squeezing problems in weak heterogeneous rock masses. *Tunnels and Tunnelling International*, 33–51.

Höge, M., Wöhling, T., & Nowak, W. (2018). A Primer for Model Selection: The Decisive Role of Model Complexity. *Water Resources Research*, *54*(3), 1688–1715. https://doi.org/10.1002/2017WR021902

Hudyma, M., Potvin, Y., & Allison, D. (2008). Seismic Monitoring of the Northparkes Lift 2 Block Cave - Part 1 Undercutting. *Journal of the South African Institute of Mining and Metallurgy*, *108*(7), 405–419. https://www.researchgate.net/publication/291758699%0ASeismic

Hudyma, M. R., & Brown, L. G. (2020). Using Seismic Data to Identify Temporal Increases in Mining-Induced Stress. *ISRM International Symposium - EUROCK 2020*.

Hurvich, C., & Tsai, C. L. (1989). Regression and time series model selection in small samples. *Biometrika*, *76*(2), 297–307. https://doi.org/10.1093/biomet/76.2.297

Isleyen, E., Duzgun, S., & McKell Carter, R. (2021). Interpretable deep learning for roof fall hazard detection in underground mines. *Journal of Rock Mechanics and Geotechnical Engineering*, *13*(6), 1246–1255. https://doi.org/10.1016/j.jrmge.2021.09.005

ITASCA Consulting Group Inc. (1992). *UDEC Manual*.

ITASCA Consulting Group Inc. (1994). *3DEC Manual*.

ITASCA Consulting Group Inc. (2015). *FLAC 8 Basics* (First Edit).

ITASCA Consulting Group Inc. (2019a). *FLAC3D — Fast Lagrangian Analysis of Continua in Three-Dimensions*. *Ver. 5.0*.

ITASCA Consulting Group Inc. (2019b). *PFC* (6.0).

ITASCA Consulting Group Inc. (2019c). *PFC3D* (No. 6).

Jakeman, A. J., Letcher, R. A., & Norton, J. P. (2006). Ten iterative steps in development and evaluation of environmental models. *Environmental Modelling and Software*, *21*(5), 602–614. https://doi.org/10.1016/j.envsoft.2006.01.004

Janeras, M., Jara, J. A., Royán, M. J., Vilaplana, J. M., Aguasca, A., Fàbregas, X., Gili, J. A., & Buxó, P. (2017). Multi-technique approach to rockfall monitoring in the Montserrat massif (Catalonia, NE Spain). *Eng Geol*, *219*, 4–20. https://doi.org/10.1016/j.enggeo.2016.12.010

Jefferson, C. W., Thomas, D. J., Gandhi, S. S., Ramaekers, P., Delaney, G., Brisbin, D., Cutts, C., Portella, P., & Olson, R. A. (2007). Unconformity-associated uranium deposits of the Athabasca basin, Saskatchewan and Alberta. *Bulletin of the Geological Survey of Canada*, *588*, 23–67. https://doi.org/10.4095/223744

Jethwa, J. L., Singh, B., & Singh, B. (1984). Estimation of ultimate rock pressure for tunnel linings under squeezing rock conditions – a new approach. In E. T. Brown & J. A. Hudson (Eds.), *Design and Performance of Underground Excavations, ISRM Symposium* (pp. 231–238).

Jiang, R., Dai, F., Liu, Y., & Wei, M. (2020). An automatic classification method for microseismic events and blasts during rock excavation of underground caverns. *Tunn Undergr Sp Tech*, *101*. https://doi.org/10.1016/j.tust.2020.103425

Jing, L., & Hudson, J. A. (2002). Numerical methods in rock mechanics. *Int J of Rock Mech & Min Sci*, *39*, 409–427.

Jing, L., & Stephansson, O. (2007a). Chapter 4 Fluid Flow and Coupled Hydro-Mechanical Behavior of Rock Fractures. In *Developments in Geotechnical Engineering* (Vol. 85, pp. 111–144). https://doi.org/10.1016/S0165-1250(07)85004-8

Jing, L., & Stephansson, O. (2007b). Constitutive Models of Rock Fractures and Rock Masses - The Basics. In *Developments in Geotechnical Engineering* (Vol. 85, pp. 47–109). https://doi.org/10.1016/S0165-1250(07)85003-6

Ju, Y., Wang, Y., Su, C., Zhang, D., & Ren, Z. (2019). Numerical analysis of the dynamic evolution of mining-induced stresses and fractures in multilayered rock strata using continuum-based discrete element methods. *Int J of Rock Mech & Min Sci*, *113*, 191–210. https://doi.org/10.1016/j.ijrmms.2018.11.014

Kabwe, E., Karakus, M., & Chanda, E. K. (2020). Isotropic damage constitutive model for time-dependent behaviour of tunnels in squeezing ground. *Comput Geotech*, *127*. https://doi.org/10.1016/j.compgeo.2020.103738

Kahneman, D. (2011). *Thinking, Fast and Slow*. Penguin Books.

Kalenchuk, K. (2017, June). The Contribution of Time-Dependent Rib Deterioration to Pillar Bursting Seismicity at Westwood Mine. *ISRM Progressive Rock Failure Conference*.

Kalenchuk, K. S. (2018). *MDEng Technical Memo #18008-105: Back Analysis of Large Seismic Events in 1SHW at Garson Mine, and Assessment of Future Risk to Mining and Development in 1SHW 4800-5200 Levels*.

Kalenchuk, K. S. (2019). 2019 Canadian Geotechnical Colloquium : Mitigating a fatal flaw in modern geomechanics : understanding uncertainty , applying model calibration , and defying the hubris in numerical modelling. *GEO St. John's 2019 – The 72nd Canadian Geotechnical Conference*.

Kalenchuk, K. S., Crockford, A. M., Hume, C. D., Milner, N., & Watson, J. (2014). Application of numerical modelling to predicted seismic probability and mitigate associated risks during the craig pillar

extraction at morrison mine, KGHM international. *48th US Rock Mechanics / Geomechanics Symposium 2014*.

Karpatne, A., Atluri, G., Faghmous, J., Steinback, M., Banerjee, A., Ganguly, A., Shekhar, S., Samatova, N., & Kumar, V. (2017). Theroy-Guided Data Science: A New Paradigm for Scientific Discovery from Data. *IEEE Transactions on Knowledge and Data Engineering*, *29*(10).

Khan, M., Bashier, E., & Bashier, M. (2017a). Chapter 4 Naïve Bayesian Classification. In *Machine Learning : Algorithms and Applications* (Issue c).

Khan, M., Bashier, E., & Bashier, M. (2017b). Chapter 5 The k -Nearest Neighbors Classifiers. In *Machine Learning : Algorithms and Applications* (Issue c).

Khan, M., Bashier, E., & Bashier, M. (2017c). Chapter 6 Neural Networks. In *Machine Learning : Algorithms and Applications*.

Khan, M., Bashier, E., & Bashier, M. (2017d). Chapter 8 Support Vector Machine. In *Machine Learning : Algorithms and Applications* (Issue c).

Khan, U. T., He, J., & Valeo, C. (2018). River flood prediction using fuzzy neural Networks: An investigation on automated network architecture. *Water Sci Technol*, *2017*(1), 238–247. https://doi.org/10.2166/wst.2018.107

Khan, U. T., & Valeo, C. (2016). Dissolved oxygen prediction using a possibility theory based fuzzy neural network. *Hydrol. Earth Syst. Sci.*, *20*, 2267–2293. https://doi.org/10.5194/hess-20-2267-2016

Khan, U. T., & Valeo, C. (2017). Optimising fuzzy neural network architecture for dissolved oxygen prediction and risk analysis. *Water (Switzerland)*, *9*(6), 381. https://doi.org/10.3390/w9060381

Kingma, D. P., & Ba, J. L. (2015). Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 1–15.

Knight, E. E., Rougier, E., Lei, Z., & Munjiza, A. (2014). *Hybrid Optimization Software Suite*.

Koenker, R., & Hallock, K. F. (2001). Quantile regression. *J Econ Perspect*, *15*(4), 143–156. https://doi.org/10.1257/jep.15.4.143

Koopialipoor, M., Fahimifar, A., Ghaleini, E. N., Momenzadeh, M., & Armaghani, D. J. (2019). Development of a new hybrid ANN for solving a geotechnical problem related to tunnel boring machine performance. *Engineering with Computers*. https://doi.org/10.1007/s00366-019-00701-8

Kudo, M., Toyama, J., & Shimbo, M. (1999). Multidimensional curve classification using passing-through regions. *Pattern Recognition Letters*, *20*(11–13), 1103–1111. https://doi.org/10.1016/S0167-8655(99)00077-X

Kumar, M., & Samui, P. (2014). Analysis of epimetamorphic rock slopes using soft computing. *Journal of Shanghai Jiaotong University (Science)*, *19*(3), 274–278. https://doi.org/10.1007/s12204-014-1499-1

Kumar, M., Samui, P., & Naithani, A. K. (2013). Determination of uniaxial compressive strength and modulus of elasticity of travertine using machine learning techniques. *International Journal of Advances in Soft Computing and Its Applications*, *5*(3), 2074–8523.

Laforce, T. (2006). *PE281 Boundary Element Method* (Issue June, pp. 1–12). Stanford University, CA.

Langford, J. C. (2013). *PhD Thesis - Application of reliability methods to the design of underground structures*.

Lawal, A. I., & Kwon, S. (2020). Application of artificial intelligence to rock mechanics: An overview. *J Rock Mech Geotech*, *13*(1), 248–266. https://doi.org/10.1016/j.jrmge.2020.05.010

Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. In *Nature* (Vol. 521, Issue 7553, pp. 436–444). Nature Publishing Group. https://doi.org/10.1038/nature14539

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., & Jackel, L. D. (1989). Backpropagation applied to handwritten zip code recognition. *Neural Comput*, *1*(4), 541–551. https://doi.org/doi.org/10.1162/neco.1989.1.4.541

Leu, S. S., Chen, C. N., & Chang, S. L. (2001). Data mining for tunnel support stability: Neural network approach. *Automat Constr*, *10*(4), 429–441. https://doi.org/10.1016/S0926-5805(00)00078-9

Li, X., Chen, Z., Chen, J., & Zhu, H. (2019). Automatic characterization of rock mass discontinuities using 3D point clouds. *Engineering Geology*, *259*(May), 105131. https://doi.org/10.1016/j.enggeo.2019.05.008

Li, X., Kim, E., & Walton, G. (2019). A study of rock pillar behaviors in laboratory and in-situ scales using combined finite-discrete element method models. *International Journal of Rock Mechanics and Mining Sciences*, *118*(March), 21–32. https://doi.org/10.1016/j.ijrmms.2019.03.030

Limit State Inc. (2013). *Eurocode 7 Design Approach 1: Basic Procedure*. http://www.limitstate.com/geo/eurocode7

LimiteStateInc. (2013). *Eucocode 7 Design Approach 1: Basic Procedure*. http://www.limitstate.com/geo/eurocode7

Lipton, Z. C. (2016, June 10). The Mythos of Model Interpretability. *2016 ICML Workshop on Human Interpretability in Machine Learning*. http://arxiv.org/abs/1606.03490

Lisjak, A., Tatone, B. S. A., Grasselli, G., & Vietor, T. (2014). Numerical modelling of the anisotropic mechanical behaviour of opalinus clay at the laboratory-scale using FEM/DEM. *Rock Mechanics and Rock Engineering*, *47*(1), 187–206. https://doi.org/10.1007/s00603-012-0354-7

Liu, F., & Yang, M. (2005). Verification and Validation of Artificial Neural Networks. In Z. S. & J. R. (Eds.), *AI 2005: Advances in Artificial Intelligence. Lecture Notes in Computer Science, vol 3809.* (pp. 1041–1046). Springer. https://doi.org/doi.org/10.1007/11589990_137

Liu, K., & Liu, B. (2017). Optimization of smooth blasting parameters for mountain tunnel construction with specified control indices based on a GA and ISVR coupling algorithm. *Tunn and Underground Sp Tech*, *70*, 363–374. https://doi.org/10.1016/j.tust.2017.09.007

Liu, Z., Li, L., Fang, X., Qi, W., Shen, J., Zhou, H., & Zhang, Y. (2021). Hard-rock tunnel lithology prediction with TBM construction big data using a global-attention-mechanism-based LSTM network. *Automat Constr*, *125*(August 2020), 103647. https://doi.org/10.1016/j.autcon.2021.103647

Lorig, L. J., Brady, B. H. G., & Cundall, P. A. (1986). Hybrid distinct element-boundary element analysis of jointed rock. *International Journal of Rock Mechanics and Mining Sciences And*, *23*(4), 303–312. https://doi.org/10.1016/0148-9062(86)90642-X

Ma, T., Tang, C., Tang, S., Kuang, L., Yu, Q., Kong, D., & Zhu, X. (2018). Rockburst mechanism and prediction based on microseismic monitoring. *Int J of Rock Mech & Min Sci*, *110*, 177–188. https://doi.org/10.1016/j.ijrmms.2018.07.016

Mahabadi, O. K., Cottrell, B. E., & Grasselli, G. (2010). An example of realistic modelling of rock dynamics problems: FEM/DEM simulation of dynamic brazilian test on Barre Granite. *Rock Mechanics and Rock Engineering*, *43*(6), 707–716. https://doi.org/10.1007/s00603-010-0092-7

Mahdevari, S., & Torabi, S. R. (2012). Prediction of tunnel convergence using Artificial Neural Networks. *Tunn and Undergr Sp Tech*, *28*(1), 218–228. https://doi.org/10.1016/j.tust.2011.11.002

Maier, H. R., & Dandy, G. C. (2000). Neural networks for the prediction and forecasting of water resources variables: A review of modelling issues and applications. *Environmental Modelling and Software*. https://doi.org/10.1016/S1364-8152(99)00007-9

Maier, H. R., Jain, A., Dandy, G. C., & Sudheer, K. P. (2010a). Methods used for the development of neural networks for the prediction of water resource variables in river systems: Current status and future directions. *Environmental Modelling and Software*, *25*(8), 891–909. https://doi.org/10.1016/j.envsoft.2010.02.003

Maier, H. R., Jain, A., Dandy, G. C., & Sudheer, K. P. (2010b). Methods used for the development of neural networks for the prediction of water resource variables in river systems: Current status and future directions. In *Environmental Modelling and Software*. https://doi.org/10.1016/j.envsoft.2010.02.003

Mandic, D. P., & Chambers, J. A. (2001). *Recurrent neural networks for prediction: learning algorithms, architectures, and stability*. John Wiley & Sons, Inc.

Map3D. (2019). *Map3D* (No. 2019).

Marcher, T., Erharter, G. H., & Winkler, M. (2020). Machine Learning in tunnelling – Capabilities and challenges. *Geomechanik Und Tunnelbau*, *13*(2), 191–198. https://doi.org/10.1002/geot.202000001

Marsland, S. (2014). Machine learning: An algorithmic perspective. In *Machine Learning: An Algorithmic Perspective, Second Edition*. Chapman & Hall. https://doi.org/10.1201/b17476

Martins, F. F., & Miranda, T. F. S. (2013). Prediction of hard rock TBM penetration rate based on Data Mining techniques. *Proceedings of the 18th International Conference on Soil Mechanics and Geotechnical Engineering*, 1751–1754.

Martz, P., Cathelineau, M., Mercadier, J., Boiron, M. C., Jaguin, J., Tarantola, A., Demacon, M., Gerbeaud, O., Quirt, D., Doney, A., & Ledru, P. (2017). C-O-H-N fluids circulations and graphite precipitation in reactivated Hudsonian shear zones during basement uplift of the Wollaston-Mudjatik Transition Zone: Example of the Cigar Lake U deposit. *Lithos*, *294–295*(October), 222–245. https://doi.org/10.1016/j.lithos.2017.10.001

MathWorks Inc. (2019). *MATLAB R2019b 9.7.0.1261785*.

MathWorks Inc. (2021). *MATLAB R2021 9.10.0.1684407*.

May, R., Dandy, G., & Maier, H. (2011). Review of Input Variable Selection Methods for Artificial Neural Networks. In K. Suzuki (Ed.), *Artificial Neural Networks - Methodological Advances and Biomedical Applications* (pp. 19–44). Intech. https://doi.org/10.5772/16004

May, R. J., Dandy, G. C., Maier, H. R., & Nixon, J. B. (2008). Application of partial mutual information variable selection to ANN forecasting of water quality in water distribution systems. *Environ. Model. Softw.*, *23*(10–11), 1289–1299. https://doi.org/doi.org/10.1016/j.envsoft.2008.03.008

May, R. J., Maier, H. R., & Dandy, G. C. (2010). Data splitting for artificial neural networks using SOM-based stratified sampling. *Neural Networks*, *23*(2), 283–294. https://doi.org/10.1016/j.neunet.2009.11.009

Mayr, A., Rutzinger, M., & Geitner, C. (2018). Multitemporal analysis of objects in 3D point clouds for landslide monitoring. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences - ISPRS Archives*, *42*(2), 691–697. https://doi.org/10.5194/isprs-archives-XLII-2-691-2018

Mcgaughey, J. (2020). Artificial intelligence and big data analytics in mining geomechanics. *The Journal of the Southern African Institute of Mining and Metallurgy*, *120*. https://doi.org/10.17159/2411

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* (pp. 219–232). Australian Centre for Geomechanics. https://doi.org/doi.org/10.36487/ACG_rep/1905_11_McGaughey

Millar, D., & Clarici, E. (2002). Investigation of back-propagation artificial neural networks in modelling the stress-strain behaviour of sandstone rock. *Genet Sel Evol*, *47*, 3326–3331. https://doi.org/10.1109/icnn.1994.374770

Mining Data Solutions. (2022). *Garson Mine*. MDO Data Online Inc. https://miningdataonline.com/property/41/Garson-Mine.aspx#Owners

Miraco Mining Innovation. (2019). *MoFrac*.

Mitchell, T. M. (2015). Chapter 1 : What is Machine Learning? In *Machine Learning : Hands-on for Developers and Technical Professionals* (Issue c).

Mnih, V., & Hinton, G. (2012). Learning to label aerial images from noisy data. *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, *1*, 567–574.

Mohammed, M., Khan, M., Bashier, E., & Bashier, M. (2016a). Chapter 1 Introduction to Machine Learning. In *Machine Learning : Algorithms and Applications* (Issue c, pp. 1–34). https://doi.org/10.1201/9781315371658-2

Mohammed, M., Khan, M., Bashier, E., & Bashier, M. (2016b). Chapter 2 Decision Trees. In *Machine Learning : Algorithms and Applications* (Issue c, pp. 1–204). https://doi.org/10.1201/9781315371658

Molnar, C. (2022). *Interpretable Machine Learning: A Guide for Making Black Boxes Explainable* (2nd ed.). christophm.github.io/interpretable-ml-book/

Morgenroth, J. (2016). *Elastic stress modelling and prediction of ground class using a Bayesian Belief Network at the Kemano tunnels*.

Morgenroth, J. (2021a). *Cigar Lake Mine Convolutional Neural Network*.

Morgenroth, J. (2021b). *Cigar Lake Mine Convolutional Neural Network*. https://doi.org/10.5281/ZENODO.5755063

Morgenroth, J. (2022). *Garson Mine Long Short-Term Memory Network*.

Morgenroth, J., Khan, U. T., & Perras, M. A. (2019). An overview of opportunities for machine learning methods in underground rock engineering design. *Geosci J*, *9*(12), 504–524.

Morgenroth, J., Perras, M. A., & Khan, U. T. (2020). Convolutional Neural Networks for predicting tunnel support and liner performance: Cigar Lake Mine case study. *Proceedings of the 54th US Rock Mechanics/Geomechanics Symposium*.

Morgenroth, J., Perras, M. A., & Khan, U. T. (2021a). A Convolutional Neural Network approach for predicting tunnel liner yield at Cigar Lake Mine. *Rock Mech Rock Eng*. https://doi.org/10.1007/s00603-021-02563-3

Morgenroth, J., Perras, M. A., & Khan, U. T. (2021b). An Input Variable Selection approach for a Convolutional Neural Network that forecasts tunnel liner yield at the Cigar Lake Mine. *Rocscience International Conference 2021 - The Evolution of Geotech: 25 Years of Innovation*.

Morgenroth, J., Perras, M. A., & Khan, U. T. (2022). Tunnel liner yield forecasting: what can we discover using Machine Learning and Input Variable Selection? *Rock Mech Rock Eng*.

Morgenroth, J., Perras, M. A., Khan, U. T., Kalenchuk, K. S., & Moreau-Verlaan, L. (2021). Forecasting principal stresses using microseismic data and a Long-Short Term Memory network at Garson Mine. *GEO Niagara 2021 - Creating a Sustainable and Smart Future*.

Morgenroth, J., Perras, M. A., Khan, U. T., & Vasileiou, A. (2020). An Artificial Neural Network approach for predicting rock support damage at Cigar Lake Mine : A Case Study. *ISRM International Symposium Eurock 2020 – Hard Rock Engineering*.

Morgenroth, J., Snieder, E., Perras, M., & Khan, U. (2019). Comparison of Bayesian Belief Networks and Artificial Neural Networks for prediction of tunnel ground class. *Proceedings of the ISRM 14th International Congress of Rock Mechanics*.

Mosavi, A., Ozturk, P., & Chau, K. W. (2018). Flood prediction using machine learning models: Literature review. *Water*, *10*(11), 1–40. https://doi.org/10.3390/w10111536

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. The MIT Press.

Ninić, J., Freitag, S., & Meschke, G. (2017). A hybrid finite element and surrogate modelling approach for simulation and monitoring supported TBM steering. *Tunnelling and Underground Space Technology*, *63*, 12–28. https://doi.org/10.1016/j.tust.2016.12.004

NWMO. (2011). *Geosynthesis, Technical Report DGR-TR-2011-11*.

Olah, C. (2015). *Understanding LSTM Networks*. http://colah.github.io/posts/2015-08-Understanding-LSTMs/

Ontario Mining Association. (2022, March 18). *GDP contribution from mining in Ontario expected to grow 25% by 2025*. GlobeNewswire. https://www.globenewswire.com/news-release/2022/03/18/2405974/0/en/GDP-CONTRIBUTION-FROM-MINING-IN-ONTARIO-

EXPECTED-TO-GROW-25-BY-2025.html#:~:text=In%202020%2C%20Ontario's%20mining%20sector,50%25%20from%202010%20to%202019.

Pan, X. D., & Reed, M. B. (1991). A coupled distinct element-finite element method for large deformation analysis of rock masses. *International Journal of Rock Mechanics and Mining Sciences And*, *28*(1), 93–99. https://doi.org/10.1016/0148-9062(91)93238-2

Panet, M., Givet, A. G., Duc, J. L., Piraud, J., & Wong, H. T. (2001). *The convergence–confinement method*.

Papadopoulos, G., Edwards, P. J., & Murray, A. F. (2000). Confidence estimation methods for neural networks : a practical comparison. *ESANN'2000 Proceedings - European Symposium on Artificial Neural Networks*, 75–80.

Paraskevopoulou, C., & Diederichs, M. (2018). Analysis of time-dependent deformation in tunnels using the Convergence-Confinement Method. *Tunnelling and Underground Space Technology*. https://doi.org/10.1016/j.tust.2017.07.001

Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. *Proceedings of the 30th International Conference on Machine Learning*.

Paudel, B., Jafarpour, M., & Brummer, R. (2012). *Cigar Lake Mine MDS Tunnel Liner Loading Analysis at Cameco*. ITASCA International Inc.

PDAC. (2019a). *Concepts and application of machine learning to mining geoscience: A practical course*.

PDAC. (2019b). *Concepts and application of machine learning to mining geoscience: A practical course*.

Peck, R. B. (1969). 9th Rankine lecture: Advantages and limitations of the observational method in applied soil mechanics. *Geotechnique*, *19*(2), 171–187.

Perras, M. (2009). *MASc Thesis: Tunnelling in horizontally laminated ground: the influence of lamination thickness on anisotropic behaviour and practical observations from the Niagara Tunnel Project*.

Perras, M. A., & Diederichs, M. S. (2016). Predicting excavation damage zone depths in brittle rocks. *Journal of Rock Mechanics and Geotechnical Engineering*. https://doi.org/10.1016/j.jrmge.2015.11.004

Perras, M. A., Wannenmacher, H., & Diederichs, M. S. (2015a). Underground Excavation Behaviour of the Queenston Formation: Tunnel Back Analysis for Application to Shaft Damage Dimension Prediction. *Rock Mechanics and Rock Engineering*. https://doi.org/10.1007/s00603-014-0656-z

Perras, M. A., Wannenmacher, H., & Diederichs, M. S. (2015b). Underground Excavation Behaviour of the Queenston Formation: Tunnel Back Analysis for Application to Shaft Damage Dimension Prediction.

Rock Mechanics and Rock Engineering, *48*(4), 1647–1671. https://doi.org/10.1007/s00603-014-0656-z

Perras, M. A., Wannenmacher, H., & Diederichs, M. S. (2015c). Underground Excavation Behaviour of the Queenston Formation: Tunnel Back Analysis for Application to Shaft Damage Dimension Prediction. *Rock Mechanics and Rock Engineering*, *48*(4), 1647–1671. https://doi.org/10.1007/s00603-014-0656-z

Perras, M., & Diederichs, M. S. (2016). Predicting excavation damage zone depths in brittle rocks. *Journal of Rock Mechanics and Geotechnical Engineering*, *8*(1), 60–74.

Phoon, K. K., & Chiang, J. (2021). Keynote Paper: Advances in data-driven subsurface mapping. *Rocscience International Conference 2021*.

Phoon, K. K., Ching, J., & Shuku, T. (2021). Challenges in data-driven site characterization. In *Georisk: Assessment and Management of Risk for Engineered Systems and Geohazards* (pp. 1–13). Taylor & Francis. https://doi.org/doi.org/10.1080/17499518.2021.1896005

Phoon, K. K., Ching, J., & Wang, Y. (2019). Managing risk in geotechnical engineering-from data to digitalization. *Proceedings of the 7th International Symposium on Geotechnical Safety and Risk*. https://doi.org/dx.doi.org/10.3850/978-981-11-2725-0-SL-cd

Pu, Y., Apel, D. B., & Hall, R. (2020). Using machine learning approach for microseismic events recognition in underground excavations: Comparison of ten frequently-used models. *Eng Geo*, *268*, 1–13. https://doi.org/10.1016/j.enggeo.2020.105519

Pu, Y., Apel, D. B., & Lingga, B. (2018). Rockburst prediction in kimberlite using decision tree with incomplete data. *Journal of Sustainable Mining*, *17*(3), 158–165. https://doi.org/10.1016/j.jsm.2018.07.004

Pu, Y., Apel, D. B., Liu, V., & Mitri, H. (2019). Machine learning methods for rockburst prediction-state-of-the-art review. *Int J Min Sci Technol*, *29*(4), 565–570. https://doi.org/10.1016/j.ijmst.2019.06.009

Qi, C., Fourie, A., Du, X., & Tang, X. (2018). Prediction of open stope hangingwall stability using random forests. *Natural Hazards*, *92*(2), 1179–1197. https://doi.org/10.1007/s11069-018-3246-7

Reed, S. E., Lee, H., Anguelov, D., Szegedy, C., Erhan, D., & Rabinovich, A. (2015). Training deep neural networks on noisy labels with bootstrapping. In Y. Bengio & Y. LeCun (Eds.), *3rd International Conference on Learning Representations, ICLR 2015 - Workshop Track Proceedings May 7–9, 2015* (pp. 1–11). ICLR.

Ribeiro e Sousa, L., Miranda, T., Leal e Sousa, R., & Tinoco, J. (2017). The Use of Data Mining Techniques in Rockburst Risk Assessment. *Engineering*, *3*(4), 552–558. https://doi.org/10.1016/J.ENG.2017.04.002

Rockfield. (2019). *Elfen*.

Rocscience. (2019). *RS2* (9.018).

Roscience Inc. (2019). *Examine2D* (8.0).

Roworth, M. (2013). *Understanding the Effect of Freezing on Rock Mass Behaviour as Appllied to the Cigar Lake Mining Method (MASc Thesis)*. The University of British Columbia.

Samanta, A., Saha, A., Satapathy, S. C., Fernandes, S. L., & Zhang, Y. D. (2020). Automated detection of diabetic retinopathy using convolutional neural networks on a small dataset. *Pattern Recognition Letters*, *135*, 293–298. https://doi.org/10.1016/j.patrec.2020.04.026

Sandia National Labs and Temple University. (2019). *LAMMPS*.

Santos, O. J., & Celestino, T. B. (2008). Artificial neural networks analysis of São Paulo subway tunnel settlement data. *Tunnelling and Underground Space Technology*. https://doi.org/10.1016/j.tust.2007.07.002

Schürch, R., & Agnagnostou, G. (2012). The applicability of the ground response curve to tunnelling problems that violate rotational symmetry. *Rock Mech Rock Eng*, *45*, 1–10.

Seif, G. (2018). *Handling Imbalanced Datasets in Deep Learning*. Towards Data Science. https://towardsdatascience.com/handling-imbalanced-datasets-in-deep-learning-f48407a0e758

Setiono, R., & Liu, H. (1997). Neural-network feature selector. *IEEE Transactions on Neural Networks*, *8*(3), 654–662. https://doi.org/10.1109/72.572104

Shahin, M. A., Maier, H. R., & Jaksa, M. B. (2004). Data division for developing neural networks applied to geotechnical engineering. *Journal of Computing in Civil Engineering*, *18*(2), 105–114. https://doi.org/10.1061/(ASCE)0887-3801(2004)18:2(105)

Shawkat, A., & Smith, K. A. (2006). On learning algorithm selection for classification. *Appl Soft Comput*, *6*(2), 119–138.

Shi, X., Chen, Z., Wang, H., Yeung, D., Wong, W., & Woo, W. (2015). Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. *Adv Neur In*, *8*.

Shlens, J. (2014). *A Tutorial on Principal Component Analysis*. http://arxiv.org/abs/1404.1100

Shu, C., & Burn, D. H. (2004). Artificial neural network ensembles and their application in pooled flood frequency analysis. *Water Resources Research*, *40*(9), 1–10. https://doi.org/10.1029/2003WR002816

Singh, B., & Goel, R. K. (1999). *Rock mass classification: a practical approach in Civil Engineering*.

Singh, B., Jethwa, J. L., Dube, A. K., & Singh, B. (1992). Correlation between observed support pressure and rock mass quality. *Tun Undergr Sp Tech*, *7*, 59–74.

Singh, V. K., Singh, D., & Singh, T. N. (2001). Prediction of strength properties of some schistose rocks from petrographic properties using artificial neural networks. *Int J of Rock Mech & Min Sci*. https://doi.org/10.1016/S1365-1609(00)00078-2

Sklavounos, P., & Sakellariou, M. (1995). Intelligent classification of rock masses. *Trans on Info Comm Tech*, *8*, 387–393. https://doi.org/10.2495/AI950411

Šmilauer, V. (2009). *Yade User's Manual*.

Snieder, E., Abogadil, K., & Khan, U. T. (2021). Resampling and ensemble techniques for improving ANN-based high-flow forecast accuracy. *Hydrology and Earth System Sciences*, *25*(5), 2543–2566. https://doi.org/10.5194/hess-25-2543-2021

Snieder, E., Shakir, R., & Khan, U. T. (2019). A Comprehensive Comparison of Four Input Variable Selection Methods for Artificial Neural Network Flow Forecasting Models. *J Hydrol*, *583*, 124299. https://doi.org/10.1016/j.jhydrol.2019.124299

Solomatine, D. P., & Ostfeld, A. (2007a). Data-driven modelling: some past experiences and new approaches. *Journal of Hydroinformatics*, *10*(1), 3–22. https://doi.org/10.2166/hydro.2008.015

Solomatine, D. P., & Ostfeld, A. (2007b). Data-driven modelling: some past experiences and new approaches. *J Hydroinform*, *10*(1), 3–22. https://doi.org/doi.org/10.2166/hydro.2008.015

Song, Z., Jiang, A., & Jiang, Z. (2015). Back Analysis of Geomechanical Parameters Using Hybrid Algorithm Based on Difference Evolution and Extreme Learning Machine. *Math Probl Eng*, *2015*. https://doi.org/10.1155/2015/821534

Song, Z. P., Jiang, A. N., & Jiang, Z. Bin. (2015a). Back Analysis of Geomechanical Parameters Using Hybrid Algorithm Based on Difference Evolution and Extreme Learning Machine. *Mathematical Problems in Engineering*, *2015*. https://doi.org/10.1155/2015/821534

Song, Z. P., Jiang, A. N., & Jiang, Z. B. (2015b). Back Analysis of Geomechanical Parameters Using Hybrid Algorithm Based on Difference Evolution and Extreme Learning Machine. *Math Probl Eng*, *2015*, 1–11. https://doi.org/10.1155/2015/821534

Sousa, R. L., & Einstein, H. H. (2012). Risk analysis during tunnel construction using Bayesian Networks: Porto Metro case study. *Tunn Undgr Space Tech*, *27*(1), 86–100. https://doi.org/10.1016/j.tust.2011.07.003

Špačková, O., & Straub, D. (2013). Dynamic Bayesian Network for Probabilistic Modeling of Tunnel Excavation Processes. *Comput-Aided Civ Inf*, *28*(1), 1–21. https://doi.org/10.1111/j.1467-8667.2012.00759.x

Spence, E. (2018). *CO Summer School Central : Introduction to Neural networks with Python I* (Issue June).

Staff writer. (2008, September 16). *A glimpse into Garson Mine's 100 years of evolution*. Sudbury.Com. https://www.sudbury.com/local-news/a-glimpse-into-garson-mines-100-years-of-evolution-221460

Sun, Y., Feng, X., & Yang, L. (2018). Predicting Tunnel Squeezing Using Multiclass Support Vector Machines. *Advances in Civil Engineering*, *2018*, 1–12. https://doi.org/10.1155/2018/4543984

Trifu, C.-I., & Young, R. P. (1992). *Analysis of Strong Ground Motion: An Overview and Application to Rockburst Studies. Internal Research Report MRD-SM0034*.

Trivedi, R., T.N.Singh, Mudgal, K., & Gupta, N. (2015). Application of Artificial Neural Network for Blast Performance Evaluation. *International Journal of Research in Engineering and Technology*, *03*(05), 564–574. https://doi.org/10.15623/ijret.2014.0305104

UCRegents. (2006). *OpenSees*.

Unterlaß, P. J., Sapronova, A., Hecht-Méndez, J., Dickmann, T., & Marcher, T. (2022). Cascade learning system to improve the interpretation of geological conditions ahead of the tunnel face. *EAGE Digital 2022*.

Vale. (2015). *Garson Complex Mine Design Package Update (2015)*.

Vallejos, J. A., & McKinnon, S. D. (2013). Logistic regression and neural network classification of seismic records. *Int J of Rock Mech & Min Sci*, *62*, 86–95. https://doi.org/10.1016/j.ijrmms.2013.04.005

Varadarajan, A., Sharma, K. G., & Singh, R. K. (1985). Some aspect of coupled FEBEM analysis of underground openings. *International Journal for Numerical and Analytical Methods in Geomechanics*, *9*(6), 557–571.

Weidner, L., Walton, G., & Kromer, R. (2019). Classification methods for point clouds in rock slope monitoring : A novel machine learning approach and comparative analysis. *Eng Geo*, *263*. https://doi.org/10.1016/j.enggeo.2019.105326

Wieleba, P., & Sikora, J. (2009). Open source BEM library. *Advances in Engineering Software*, *40*(8), 564–569. https://doi.org/10.1016/j.advengsoft.2008.10.007

Wojtecki, L., Iwaszenko, S., Apel, D. B., & Cichy, T. (2021). An Attempt to Use Machine Learning Algorithms to Estimate the Rockburst Hazard in Underground Excavations of Hard Coal Mine. *Energies*, *14*(21). https://doi.org/https://doi.org/10.3390/en14216928

Wolpert, D. H., & Macready, W. G. (1996). No free lunch theorems for search. In *Technical Report SFI-TR-95-02-010*. The Santa Fe Institute.

WSP (formerly Parson Brinkerhoff Quade & Douglas Inc.). (1999). *Cigar Lake Mine Global Mine Model Study: Final Report*.

Xie, Q., & Peng, K. (2019). Space-Time Distribution Laws of Tunnel Excavation Damaged Zones (EDZs) in Deep Mines and EDZ Prediction Modeling by Random Forest Regression. *Advances in Civil Engineering*, *2019*. https://doi.org/10.1155/2019/6505984

Xue, R., Liang, Z., & Xu, N. (2021). Rockburst prediction and analysis of activity characteristics within surrounding rock based on microseismic monitoring and numerical simulation. *International Journal of Rock Mechanics and Mining Sciences*, *142*(January), 104750. https://doi.org/10.1016/j.ijrmms.2021.104750

Xue, Y., & Li, Y. (2018). A Fast Detection Method via Region-Based Fully Convolutional Neural Networks for Shield Tunnel Lining Defects. *Computer-Aided Civil and Infrastructure Engineering*. https://doi.org/10.1111/mice.12367

Yang, B., Mitelman, A., Elmo, D., & Stead, D. (2021). Why the future of rock mass classification systems requires revisiting their empirical past. *Quarterly Journal of Engineering Geology and Hydrogeology*. https://doi.org/10.1144/qjegh2021-039/5327071/qjegh2021-039.pdf

Yao, M., & Moreau-Verlaan, L. (2010). Strategies for mining in highly burst-prone ground conditions at Vale Garson Mine. In M. V. S. J. & Y. Potvin (Ed.), *Proceedings of the Fifth International Seminar on Deep and High Stress Mining* (pp. 549–560). Australian Centre for Geomechanics.

Zeiler, M. D., & Fergus, R. (2014). Visualizing and Understanding Convolutional Network. In D. Fleet, T. Pajdla, B. Schiele, & T. Tuytelaars (Eds.), *Computer Vision – ECCV 2014* (Vol. 8689, pp. 818–833). Springer International Publishing Switzerland.

Zhang, H., Zeng, J., Ma, J., Fang, Y., Ma, C., Yao, Z., & Chen, Z. (2021). Time Series Prediction of Microseismic Multi-parameter Related to Rockburst Based on Deep Learning. *Rock Mechanics and Rock Engineering*. https://doi.org/10.1007/s00603-021-02614-9

Zhao, K., Bonini, M., Debernardi, D., Janutolo, M., Barla, G., & Chen, G. (2015). Computational modelling of the mechanised excavation of deep tunnels in weak rock. *Computers and Geotechnics*, *66*, 158–171. https://doi.org/10.1016/j.compgeo.2015.01.020

Zhao, Z., Gong, Q., Zhang, Y., & Zhao, J. (2007). Prediction model of tunnel boring machine performance by ensemble neural networks. *Geomech Geoeng*, *2*(2), 123–128. https://doi.org/10.1080/17486020701377140

Zhou, J., Li, X., & Shi, X. (2012). Long-term prediction model of rockburst in underground openings using heuristic algorithms and support vector machines. *Safety Science*, *50*(4), 629–644. https://doi.org/10.1016/j.ssci.2011.08.065

Zhou, M., Chen, J., Huang, H., Zhang, D., Zhao, S., & Shadabfar, M. (2021). Multi-source data driven method for assessing the rock mass quality of a NATM tunnel face via hybrid ensemble learning models. *Int J Rock Mech Min*, *147*. https://doi.org/10.1016/j.ijrmms.2021.104914

APPENDIX A. PROOFS OF CONCEPT

This appendix contains the publications that collectively represent the proofs of concept completed as part of this PhD dissertation. The publications are as follows:

1. Morgenroth, J., Snieder, E., Perras, M., & Khan, U. (2019). Comparison of Bayesian Belief Networks and Artificial Neural Networks for prediction of tunnel ground class. *Proceedings of the ISRM 14th International Congress of Rock Mechanics*. September 13 – 19, 2019. Foz do Iguassu, Brazil.
2. Morgenroth, J., Perras, M. A., Khan, U. T., & Vasileiou, A. (2020). An Artificial Neural Network approach for predicting rock support damage at Cigar Lake Mine: A Case Study. *ISRM International Symposium EUROCK 2020 – Hard Rock Engineering*. June 14 – 19, 2020. Virtual.
3. Morgenroth, J., Perras, M. A., & Khan, U. T. (2020). Convolutional Neural Networks for predicting tunnel support and liner performance: Cigar Lake Mine case study. *Proceedings of the 54th US Rock Mechanics/Geomechanics Symposium*. June 28 – July 1, 2020. Virtual.
4. Morgenroth, J., Perras, M. A., & Khan, U. T. (2021). An Input Variable Selection approach for a Convolutional Neural Network that forecasts tunnel liner yield at the Cigar Lake Mine. *Rocscience International Conference 2021 - The Evolution of Geotech: 25 Years of Innovation*.
5. Morgenroth, J., Perras, M. A., Khan, U. T., Kalenchuk, K. S., & Moreau-Verlaan, L. (2021). Forecasting principal stresses using microseismic data and a Long-Short Term Memory network at Garson Mine. *GEO Niagara 2021 - Creating a Sustainable and Smart Future*. September 26 – 30, 2021. Niagara Falls, Canada.

APPENDIX B. CIGAR LAKE MINE CONVOLUTIONAL NEURAL NETWORK MATLAB CODE

All MATLAB code for developing the Cigar Lake Mine Convolutional Neural Network is hosted at a permanent DOI (Morgenroth, 2021), and may be cited as follows:

Morgenroth, J. (2021). *Cigar Lake Mine Convolutional Neural Network*. https://doi.org/10.5281/zenodo.5755063

APPENDIX C. GARSON MINE LONG-SHORT TERM MEMORY NETWORK MATLAB CODE

All MATLAB code for developing the Garson Mine Long-Short Term Memory Network is hosted at a permanent DOI (Morgenroth, 2022), and may be cited as follows:

Morgenroth, J. (2022). *Garson Mine Long Short-Term Memory Network*. https://doi.org/10.5281/zenodo.6606521

APPENDIX D. LIST OF PUBLICATIONS

All contributions published as part of this dissertation are summarized below.

Journal Articles Published in a Refereed Journal

Morgenroth, J., Perras, M. A., & Khan, U. T. (2022). *On the Interpretability of Machine Learning Using Input Variable Selection: Forecasting Tunnel Liner Yield*. Rock Mech Rock Eng. https://doi.org/10.1007/s00603-022-02987-5

Morgenroth, J., Perras M. A., & Khan, U. T. (2021). *A Convolutional Neural Network approach for predicting tunnel liner yield at Cigar Lake Mine*. Rock Mech Rock Eng. https://doi.org/10.1007/s00603-021-02563-3.

Morgenroth, J., Khan, U.T., & Perras, M. A. (2019). *An Overview of opportunities for machine learning methods in underground rock engineering design*. Geosciences 2019, 9(12), 504.

Journal Articles Submitted to a Refereed Journal

Morgenroth, J., Kalenchuk, K., Moreau-Verlann, L., Perras, M. A., & Khan, U. T. (Under Review 2022). *A novel Long-Short Term Memory network approach to the recalibration of a finite difference model for high stress mine excavations*. Georisk – Machine Learning and AI in Geotechnics. Submission ID: NGRK-2022-0079.

Conference Articles Published in Refereed Conference Proceedings

Morgenroth, J., Unterlaß, P. J., Sapronova, A., Khan, U.T., Perras, M. A., Erharter, G. H. & Marcher, T. (2022). *Practical recommendations for machine learning in underground rock engineering.* The 71st Austrian Geomechanics Colloquium. Salzburg, Austria. October 13-15, 2022.

Morgenroth, J., Perras, M. A., Khan, U.T., Kalenchuk, K., & Moreau-Verlaan, L. (2021*). Forecasting principal stresses using microseismic data and a Long-Short Term Memory network at Garson Mine*. Canadian Geotechnical Society. Proceedings of GEO Niagara 2021 – Celebrating a Sustainable and Smart Future. *(Recipient of the Best Student Paper award)*

Morgenroth, J., Perras M. A., & Khan, U. T. (2021). *An Input Variable Selection method for a Convolutional Neural Network predicting tunnel liner yield at Cigar Lake Mine*. Proceedings of the Rocscience International Conference.

Morgenroth, J., Perras M. A. & Khan, U. T. (2020). *A Convolutional Neural Network for predicting tunnel support and liner performance: A Case Study*. American Rock Mechanics Association 54th Symposium. Virtual. *(ranked in the top 40 papers)*

Morgenroth, J., Khan, U. T., Perras, M. A. & Vasileiou, A. (2020). *An Artificial Neural Network approach for predicting rock support damage at Cigar Lake Mine: A Case Study*. EUROCK 2020. June 13-19, 2020. Trondheim, Norway.

Morgenroth, J., Snieder, E., Perras, M. A. & Khan, U. T. (2019). *Comparison of Bayesian Belief Networks and Artificial Neural Networks for prediction of tunnel ground class.* Proceedings of the International Society of Rock Mechanics 14th International Congress of Rock Mechanics. Iguassu, Brazil. September 13-19, 2019.

Conference Articles Published in Other Conference Proceedings

Morgenroth, J., Khan, U. T., and Perras, M. A. (2022). *Machine Learning and Underground Geomechanics – data needs, algorithm development, uncertainty, and engineering verification*, EGU General Assembly 2022, Vienna, Austria, 23–27 May 2022, EGU22-799, https://doi.org/10.5194/egusphere-egu22-799, 2022.

Code Published at a Permanent DOI

Morgenroth, J. (2021). Cigar Lake Mine Convolutional Neural Network. https://doi.org/10.5281/ZENODO.5755063

Morgenroth, J. (2022). Garson Mine Long Short-Term Memory Network. https://doi.org/10.5281/zenodo.6606521