POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Probabilistic Machine Learning for Mineral Prospectivity Modelling: Applications to Magmatic Nickel Sulphide Systems in Canada

LAHIRU MATHEESHA NAGASINGHA
Département des génies civil, géologique et des mines

Thèse présentée en vue de l'obtention du diplôme de *Philosophiæ Doctor*
Génie minéral

Mai 2026

© Lahiru Matheesha Nagasingha, 2026.

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Cette thèse intitulée :

Probabilistic Machine Learning for Mineral Prospectivity Modelling: Applications to Magmatic Nickel Sulphide Systems in Canada

présentée par **Lahiru Matheesha NAGASINGHA**
en vue de l'obtention du diplôme de *Philosophiæ Doctor*
a été dûment acceptée par le jury d'examen constitué de :

Gabriel FABIEN-OUELLET, président
Charles LAFRENIÈRE BÉRUBÉ, membre et directeur de recherche
Dany LAUZON, membre
Mohammad PARSASADR, membre externe

ACKNOWLEDGEMENTS

First and foremost, I would like to express my deepest gratitude to my supervisor, Prof. Charles L. Bérubé, for his guidance, support, and trust throughout this journey. This journey had its share of ups and downs, and through it all, Charles stood by me with patience and encouragement, always steering me in the right direction. Our weekly meetings and discussions were invaluable in shaping not only this research but also my growth as a researcher. I am especially grateful for the opportunity to be his first PhD student. Thank you, Charles, for believing in me from the start.

My sincere thanks go to my co-authors, Prof. Christopher J.M. Lawley and Prof. Reza Ghanati, for their valuable contributions to the articles presented in this thesis. Special thanks to Chris, who provided the geoscience data cube that allowed me to start this research and shared valuable insights and comments throughout.

I would like to thank the members of my examination committee, Prof. Gabriel Fabien-Ouellet, Prof. Dany Lauzon and Dr. Mohammad Parsasadr, for dedicating their time to reading and evaluating this thesis. I am also grateful to Gabriel for providing access to his computing server, which allowed me to train and run the models developed in this research.

This research is financially supported by Natural Resources Canada and the Natural Sciences and Engineering Research Council of Canada. I also gratefully acknowledge the Bourse des pionniers de la géophysique québécoise scholarship from the KEGS Foundation.

I am thankful to my colleagues at Polytechnique Montréal, including Jefferson, Jeremy, Camilla, Ali, Jean-Luc, Léopold, Sébastien, Clarissa, Remy and Rachel, for their incredible conversations and for dedicating their time to making the department a welcoming place. Special thanks to Jean-Luc and Sébastien, who participated in and won the VRIFY AI Student Challenge with me. That experience provided a valuable opportunity to apply the knowledge gained through this research in a real-world exploration challenge.

To my parents and my brother, thank you for your unwavering support and encouragement from across the world. Your belief in me has been a constant source of strength throughout this journey.

Finally, to my wife, Amali, thank you for your endless patience, love, and support. You stood beside me through every late night and every moment, even when things were not going well. None of this would have been possible without you.

RÉSUMÉ

La demande de minéraux critiques augmente rapidement. Cette hausse résulte de la transition vers des systèmes énergétiques à faible émission de carbone et de l'épuisement progressif des gisements proches de la surface terrestre. Elle accroît le besoin de méthodes quantitatives et interprétables capables de réduire l'espace de recherche des gisements, car ceux-ci se situent souvent dans des régions éloignées, à plus grande profondeur et dans des contextes géologiques complexes. Pour répondre à ce besoin, la modélisation de la prospectivité minérale (MPM) combine des ensembles de données géologiques et géophysiques avec des méthodes d'apprentissage automatique afin de cartographier le potentiel minéral à grande échelle. Les approches conventionnelles de MPM ont toutefois plusieurs limites. Notamment, les données d'entraînement sont fortement déséquilibrés en raison de la rareté intrinsèque des gisements, ce qui réduit le pouvoir prédictif des modèles. De plus, les prédictions sont souvent difficilement interprétables et incluent rarement une quantification de l'incertitude. Cette thèse répond à ces limites en développant des cadres et des outils probabilistes d'apprentissage automatique pour la MPM à l'échelle nationale des systèmes minéralisés de sulfures magmatiques de nickel (Ni) au Canada. Le Ni est reconnu comme l'un des principaux minéraux critiques. Guidée par le cadre conceptuel des systèmes minéralisés, l'approche proposée intègre des données géologiques et géophysiques comme variables indicatrices spatiales des composantes source, cheminement et piège de ces systèmes.

Cette thèse apporte trois contributions scientifiques originales et complémentaires, chacune présentée dans un article publié dans une revue à comité de lecture. La première présente un autoencodeur variationnel conditionnel (CVAE) permettant de générer des échantillons géoscientifiques synthétiques représentatifs des conditions de minéralisation. Cette approche surpasse plusieurs techniques d'équilibrage des données pour divers algorithmes de classification. Les cartes de prospectivité obtenues délimitent correctement les gisements connus et mettent en évidence de nouvelles zones prospectives dans des régions sous-explorées du Canada. La deuxième contribution exploite la nature stochastique du CVAE afin d'évaluer la confiance des prédictions. Elle quantifie l'incertitude aléatoire et épistémique à l'aide d'ensembles de modèles. Elle introduit également un cadre d'analyse des risques et des rendements fondé sur des mesures d'incertitude relative et sur le regroupement spatial afin de hiérarchiser les cibles d'exploration. Les zones à haut rendement concentrent 94,7 % des gisements connus dans environ 4 % de la zone d'étude. La troisième contribution met en oeuvre l'inférence gaussienne approximative tractable (TAGI), une méthode bayésienne analytique qui calcule les distributions a posteriori sans recourir à la rétropropagation par gradient. La TAGI per-

met d'obtenir un rappel plus élevé pour les gisements de Ni connus, une convergence plus rapide et des cartes d'incertitude spatialement plus cohérentes que l'inférence variationnelle et la méthode de Monte Carlo désactivation aléatoire de neurones. Une analyse d'attribution des variables identifie la lithologie et la géochronologie comme les principaux facteurs contrôlant à la fois les prédictions et la confiance du modèle. Enfin, cette contribution quantifie comment l'incertitude augmente lorsque les composantes du système minéralisé ne sont que partiellement satisfaites. Dans les trois contributions, les modèles sélectionnés présentent des aires sous la courbe caractéristique de fonctionnement du récepteur supérieures à 0,9. Ensemble, ces contributions montrent que l'apprentissage automatique probabiliste peut faire progresser la MPM à l'échelle nationale au-delà des cartes de prédiction déterministes. Cette thèse fournit ainsi de nouveaux outils d'aide à la décision qui intègrent la prospectivité et la confiance des prédictions et permet de hiérarchiser les cibles potentielles d'exploration.

ABSTRACT

The growing demand for critical minerals, driven by the transition to low-carbon energy systems and the progressive depletion of near-surface deposits, has increased the need for data-driven methods to identify new deposits in complex geological settings. To this end, mineral prospectivity modelling (MPM) typically integrates geoscience datasets with machine learning to reduce the exploration search space. However, conventional MPM approaches have several limitations. This is largely due to highly imbalanced training datasets, the lack of uncertainty quantification in model predictions, and limited model interpretability. This thesis addresses these limitations by developing probabilistic machine learning frameworks and tools for national-scale MPM of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems, where Ni is recognized as a critical mineral. Guided by the mineral systems framework, the proposed approach integrates geological and geophysical datasets as spatial proxies for the source, pathway, and trap components of magmatic Ni sulphide systems.

This thesis makes three original and complementary scientific contributions. The first introduces a conditional variational autoencoder (CVAE) to generate synthetic geoscience samples representative of mineralization conditions, outperforming other balancing techniques across multiple classification algorithms. The resulting prospectivity maps successfully delineate known deposits and highlight new prospective areas in underexplored regions of Canada. The second contribution leverages the stochastic nature of the CVAE to assess prediction confidence by quantifying aleatoric and epistemic uncertainty through ensemble modelling. It also introduces a novel risk-return framework based on relative uncertainty measures and spatial clustering to prioritize exploration targets, with high-return zones concentrating 94.7 % of known deposits within approximately 4 % of the study area. The third contribution implements tractable approximate Gaussian inference (TAGI), an analytical Bayesian method that computes posterior distributions without relying on gradient-based backpropagation. TAGI achieves higher recall for known Ni deposits, faster convergence, and more spatially coherent uncertainty maps compared to variational inference and Monte Carlo dropout. A feature-attribution analysis identifies lithology and geochronology as the principal factors controlling both predictions and model confidence. Notably, uncertainty increases where mineral system components are only partially satisfied. Across all three studies, the selected models achieve areas under the receiver operating characteristic curves greater than 0.9. Collectively, these contributions show that probabilistic machine learning can advance national-scale MPM beyond deterministic prediction maps, and this thesis provides new decision-support tools that integrate prospectivity with prediction confidence to prioritize exploration targets.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS iii

RÉSUMÉ iv

ABSTRACT vi

LIST OF TABLES xii

LIST OF FIGURES xiii

LIST OF SYMBOLS AND ACRONYMS xx

LIST OF APPENDICES xxiii

CHAPTER 1 INTRODUCTION 1
1.1 Global importance of minerals 1
1.2 Critical minerals in Canada 3
1.2.1 Nickel as a priority critical mineral 6
1.3 Addressing exploration challenges through data-driven approaches 7
1.3.1 The data integration challenge 7
1.3.2 Adaptation of machine learning for mineral prospectivity modelling . 9
1.4 Research objectives 12
1.4.1 General objective 12
1.4.2 Specific objectives 12
1.5 Dissertation structure 13

CHAPTER 2 LITERATURE REVIEW 14
2.1 Introduction to mineral prospectivity modelling 14
2.2 The evolution of mineral prospectivity modelling methodologies 17
2.2.1 Knowledge-driven models 18
2.2.2 Data-driven models 20
2.2.2.1 Statistical data-driven models 20
2.2.2.2 Machine learning models 21
2.2.2.3 Deep learning models 25
2.2.3 Hybrid models 29

2.2.4 Modelling with empirical functions 29
2.2.5 Logistic-based modelling 29
2.3 Key knowledge gaps and challenges in modern mineral prospectivity modelling 30
2.3.1 Data scarcity and class imbalance 30
2.3.2 Negative training sample selection 32
2.3.3 Input data quality and deep controls 34
2.3.4 Uncertainties in mineral prospectivity modelling 35
2.3.5 Model interpretability 37
2.4 Discussion and conclusions of the literature review 39

CHAPTER 3 PRESENTATION OF THE ARTICLES 41
3.1 Article 1: A balanced mineral prospectivity model of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems using conditional variational autoencoders 41
3.2 Article 2: A generative neural network approach to uncertainty and risk-return analysis in mineral prospectivity modelling 41
3.3 Article 3: Tractable approximate Gaussian inference for interpretable mineral prospectivity modelling 42

CHAPTER 4 ARTICLE 1: A BALANCED MINERAL PROSPECTIVITY MODEL OF CANADIAN MAGMATIC NI (±CU ±CO ±PGE) SULPHIDE MINERAL SYSTEMS USING CONDITIONAL VARIATIONAL AUTOENCODERS 44
4.1 Introduction 45
4.2 Canadian geological context and magmatic Ni (± Cu ± Co ± PGE) sulphide mineral systems 49
4.2.1 Geological provinces in Canada 49
4.2.2 Magmatic Ni (± Cu ± Co ± PGE) sulphide mineral systems 49
4.3 Data sets 51
4.3.1 Geological data sets 52
4.3.2 Geophysical data sets 52
4.3.3 Training and testing data sets 54
4.4 Methods 54
4.4.1 Data preprocessing 55
4.4.2 Data augmentation methods 56
4.4.2.1 Conditional variational autoencoder 56
4.4.2.2 SMOTE 59
4.4.2.3 Class weighting 60

4.4.3 MPM classification methods 61
4.4.3.1 Extreme gradient boosting machines 61
4.4.3.2 Random forest 61
4.4.3.3 Multilayer perceptron 62
4.4.4 Evaluation metrics 62
4.4.5 Geoscience feature analysis 63
4.4.5.1 Sensitivity indices 63
4.4.5.2 Feature importance 63
4.5 Results 63
4.5.1 Success rate curves 64
4.5.2 Qualitative representation and geological comparison 65
4.5.3 Model analysis 67
4.6 Discussion 73
4.7 Conclusions 74

CHAPTER 5 ARTICLE 2: A GENERATIVE NEURAL NETWORK APPROACH TO UNCERTAINTY AND RISK-RETURN ANALYSIS IN MINERAL PROSPECTIVITY MODELLING 76
5.1 Introduction 77
5.2 The mineral system approach and related uncertainties 79
5.3 Magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems 82
5.4 Dataset 83
5.5 Methods 88
5.5.1 Data pre-processing 88
5.5.2 Data augmentation and CVAE 91
5.5.3 Extreme gradient boosting machines 93
5.5.4 Uncertainty quantification 94
5.5.4.1 Aleatoric uncertainty 94
5.5.4.2 Epistemic uncertainty 94
5.5.5 Risk-return analysis 96
5.5.6 Evaluation metrics 97
5.6 Results 97
5.6.1 Model comparison and Ni MPM 97
5.6.2 Uncertainty quantification 99
5.6.2.1 Aleatoric uncertainty 99
5.6.2.2 Epistemic uncertainty 101

5.6.3 Risk-return analysis . . . 103
5.7 Discussion . . . 104
5.7.1 Geological interpretation of the results . . . 104
5.7.2 Advancements and implications for exploration . . . 107
5.7.3 Limitations of the proposed framework . . . 107
5.8 Conclusions . . . 108

CHAPTER 6 ARTICLE 3: TRACTABLE APPROXIMATE GAUSSIAN INFERENCE FOR INTERPRETABLE MINERAL PROSPECTIVITY MODELLING . 110
6.1 Introduction . . . 111
6.2 Datasets . . . 113
6.3 Methodology . . . 117
6.3.1 Data preprocessing . . . 117
6.3.2 Data augmentation . . . 117
6.3.3 Uncertainty quantification methods . . . 119
6.3.3.1 Bayesian neural network with tractable approximate Gaussian inference . . . 119
6.3.3.2 Bayesian neural networks with variational inference . . . 122
6.3.3.3 Monte Carlo dropout . . . 124
6.3.4 Evaluation metrics . . . 124
6.3.4.1 Prediction performance . . . 125
6.3.4.2 Computational cost . . . 126
6.3.4.3 Robustness . . . 126
6.3.5 Model interpretation . . . 126
6.4 Results and discussion . . . 127
6.4.1 Model performance and selection . . . 127
6.4.2 TAGI prospectivity and uncertainty maps . . . 130
6.4.3 Geological interpretation through SHAP analysis . . . 134
6.4.4 Understanding TAGI's performance advantages . . . 138
6.4.5 Limitations of the proposed framework . . . 138
6.5 Conclusion . . . 139

CHAPTER 7 GENERAL DISCUSSION . . . 141
7.1 Cross-article methodological insights . . . 141
7.2 Geological implications . . . 143
7.3 Limitations . . . 143
7.4 Original contributions of the thesis . . . 144

CHAPTER 8 CONCLUSION 145
8.1 Summary 145
8.2 Recommendations for future work 146

REFERENCES 148

APPENDICES 181

LIST OF TABLES

Table 4.1 Area under the success rate curves for the testing data set with each data augmentation and classification method. 65

Table 5.1 Four key factors of geological uncertainty in MPM. 81

Table 5.2 Mappable criteria and datasets utilized for prospectivity modelling of the magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system (Lawley et al., 2021). 87

Table 6.1 Summary of the geoscience evidence layers used as input features, categorized by their role (drivers, sources, pathways and traps) in the magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system 115

Table 6.2 Performance comparison between the TAGI, BNN with VI, and MC dropout models evaluated on the unseen test set at classification threshold 0.5. The best result in each row is highlighted in bold. See Appendix B for performance at multiple thresholds. 128

Table 6.3 Summary of confusion matrix results for (a) TAGI, (b) BNN with VI, and (c) MC dropout models on the test set. TP, FN, TN, and FP denote true positives, false negatives, true negatives, and false positives, respectively. 130

Table A.1 Description of geoscience evidence layers used as input features for mineral prospectivity modelling of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems. 181

Table B.1 Performance comparison at multiple classification thresholds on the test dataset. Optimal thresholds are determined by maximizing the geometric mean score (GMS), which balances sensitivity and specificity. TP - true positives, FP - false positives, FN - false negatives, TN - true negatives. 184

LIST OF FIGURES

Figure 1.1 Average annual demand growth for Cu, Li, Ni, and Co minerals by scenario. kt = thousand tonnes; STEPS = Stated Policies Scenario; SDS = Sustainable Development Scenario (International Energy Agency, 2021). 2

Figure 1.2 Discovery cost and meters drilled per significant mineral deposit discovered (1979 – 2023). Data for 2013 to 2023 include an adjustment for unreported discoveries. Modified from Schodde (2025). 3

Figure 1.3 List of 34 Canadian critical minerals as of 2025 that shows primary and co-products (center) and common by-products (outer margin), separated by mineral system. Modified from Gadd et al. (2023); Government of Canada (2024). 4

Figure 1.4 Critical mineral development across Canada. Critical mineral mines, smelters, refineries, or advanced projects are located across all Canadian landmass, except for Prince Edward Island (Government of Canada, 2024). 5

Figure 1.5 Distribution of global Ni demand on major industrial applications in 2023 as stainless steel manufacturing remains the principal consumer, with batteries representing the second-largest and rapidly growing segment, followed by non-ferrous alloys, electroplating, and other minor sectors (Natural Resources Canada, 2023). 7

Figure 1.6 Exponential growth in MPM publications from the Web of Science Core Collection database from 2006 to 2024, showing budding (modest output), expansion (steady growth), and boom (rapid acceleration) phases. Modified from Zhang and Sun (2025). 10

Figure 1.7 AI-assisted prospectivity model for the Cadillac gold mining district (Val-d'Or, Québec), illustrating spatial variations in predicted gold prospectivity. Warmer colours, including pink-white, represent higher prospectivity, while cooler colours, including blue, indicate lower prospectivity values (VRIFY, 2025). 11

Figure 2.1 A simplified conceptual model of the mineral system, illustrating the translation from the critical process to its constituent processes and targeting elements, and finally to mappable criteria (Yousefi et al., 2019). 17

Figure 2.2 A chronological overview of MPM methodologies, showing the development from traditional methods to advanced DL models. Modified from Zuo (2020). 18
Figure 2.3 The modelling of epithermal Au prospectivity in the Aroroy district of the Philippines uses an inference network to combine input Boolean evidential maps. Modified from Carranza and Hale (2008). 20
Figure 2.4 Conceptual workflow illustrating the use of AI in mineral exploration. The process usually involves acquiring geoscience datasets, constructing AI models, applying these models to specific exploration tasks such as mapping, classification, or prospectivity analysis, and evaluating their performance (Yang et al., 2024a). 22
Figure 2.5 Computations performed at a single perceptron unit 23
Figure 2.6 Illustrations of (a) convolution and (b) max-pooling operations (Sun et al., 2020). 26
Figure 2.7 Schematic architecture of the cGAN used for data augmentation, showing label embedding, generator structure, and discriminator processing of real and synthetic samples (Wu et al., 2023). 28
Figure 2.8 Map of Canada's geological provinces illustrating the spatial distribution of various class label groups used for MVT prospectivity modelling. The map highlights the spatial patterns of traditionally selected negatives, magmatic deposits used for geological contrast, and dissimilarity-based negatives, compared to the known MVT deposits and occurrences (Parsa and Cumani, 2025). 33
Figure 2.9 Schematic illustrating the difference between two conceptual frameworks for mineral systems. The investigative lens (a) focuses on the causal links between multi-scale component processes (Hronsky and Kreuzer, 2019; Wyborn et al., 1994), whereas the predictive lens (b) highlights the necessary spatial and temporal conjunction of independent elements needed for ore genesis (McCuaig et al., 2010). 36
Figure 4.1 Map of the spatial distribution of geological provinces in Canada (Wheeler, 1996). Ni mineral deposits include active mining operations, past-producing mines, and advanced-stage exploration projects (Good et al., 2015; Lawley et al., 2022b), and Ni mineral occurrences include prospects, showings, and anomalous drill intercepts obtained through provincial and territorial databases. 50

Figure 4.2 Qualitative representation of the feature anomalies (a) depth to the lithosphere-asthenosphere boundary (LAB) and (b) depth to the Mohorovičić discontinuity (MOHO) demonstrates the cratonic edges and crustal thickness respectively; (c) magnetic first vertical derivative (1VD) map demonstrates the magnetic contrasts in shallow crust; (d) gravity horizontal gradient magnitude (HGM) highlights the edges of near-surface contacts or structures; (e) upward –continued (30 km) gravity HGM maps and (f) magnetic long –wavelength HGM emphasize the variations of deep density and magnetic susceptibility within the upper and lower crust respectively. 53

Figure 4.3 Spatial distribution of training and testing data set areas divided based on cross-validated technique. Based on the known Ni mineral deposits and occurrences, five cross-validation folds are obtained for the training data set and one cross-validated fold for the testing data set. 55

Figure 4.4 Flowchart of the proposed framework of this study. 56

Figure 4.5 The proposed CVAE framework. In the encoder, hidden layers $\mathbf{h}^{(l)}$ are interconnected by linear transformation parameters $\phi^{(l)}$, whereas in the decoder, they are connected by $\theta^{(l)}$. The dashed lines represent vector concatenation. 59

Figure 4.6 Success rate curves for the testing data set with (a) XGBM-MPM, (b) RF-MPM and (c) MLP-MPM models. The dashed lines are the success rate curves characteristic of a random classifier. 64

Figure 4.7 XGBM-CVAE prospectivity model results for Ni occurrences in Canada. 66

Figure 4.8 Ni mineral prospectivity around eastern Québec and western Labrador, Canada. The black rectangle on the top right side of the map denotes the exact zoomed area. 67

Figure 4.9 Detailed representation of the bedrock geology around eastern Québec and western Labrador, Canada. Modified from Lawley et al. (2024), which provides a comprehensive review of the area's geology. 68

Figure 4.10 Horizontal gradient magnitude of the gravity Bouguer anomaly across the Slave and Superior provinces in Québec, Canada. 69

Figure 4.11 Histogram plots of the distributions between generated data from the CVAE model and real geoscience data. One hot encoded data subset is deliberately excluded from this figure due to its high dimensionality. 70

Figure 4.12 The bar chart illustrates the normalized sensitivities of the CVAE's decoded variables to different geological and geophysical data sets. The data are sorted in ascending order of sensitivity. The highest sensitivity corresponds to the augmented geoscience data set, which is most influenced by the mineralized and barren conditions. 71

Figure 4.13 Normalized feature importance in the preferred XGBM-CVAE prospectivity model. The importance value of each data set is averaged within its category for one-hot encoded data and directly obtained for non-encoded data. The bars show the average feature importance of each feature, providing insights into their relative significance in the XGBM model's predictions. 72

Figure 5.1 Conceptual model for magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems. The model illustrates mantle plumes as the primary drivers that trigger mantle melting and the generation of sulphur-poor (S-poor) magma that rises through pathways such as thinning crust and faults. Key attributes, including geological faults, magnetic and gravity anomalies, lithospheric thickness, and geological terrane boundaries, are used to map the movements of magma. Structural and lithological traps facilitate sulphide deposition, with Bouguer gravity providing information on these mechanisms. Modified from Lawley et al. (2021). 84

Figure 5.2 New bedrock geological map of Canada (Lawley et al., 2024). The bedrock geology shown on this map is used to define the source and trap components of the mineral systems approach. The known Ni mineral deposits and occurrences demonstrated here are used for model training. 85

Figure 5.3 Spatial distribution of the minimum geological age map of Canada. These minimum geological ages are used to identify rocks that are coeval with major mantle melting events, which act as a transient driver in the mineral system. 86

Figure 5.4 Examples of geoscience features used in this study: (a) faults dataset highlighting geological structures (Styron and Pagani, 2020), (b) shape index (SI) derived from the Gravity field and steady-state Ocean Circulation Explorer satellite gravity data representing variations in the Earth's gravitational field (Ebbing et al., 2018), (c) magnetic HGM map showing magnetic anomalies, and (d) gravity HGM map visualizing density variations (Lawley et al., 2022a). 89

Figure 5.5 Flowchart of key steps involved in the uncertainty and risk-return analysis of prospectivity modelling. 90
Figure 5.6 ROC curves for the testing dataset that indicate the classification performance of MPMs across three dataset configurations: (1) real geoscience data with CVAE-generated synthetic positive samples, (2) reconstructed data with synthetic positive samples, and (3) a hybrid dataset combining real, reconstructed, and synthetic positive samples. 98
Figure 5.7 Modelled mineral prospectivity across Canada based on the results from configuration 1. The prospectivity values are quantified on a scale from 0 to 1, where high-prospectivity zones, indicating areas with a higher likelihood of mineralization, are closer to 1. Low-prospectivity zones are closer to 0. 99
Figure 5.8 Box plots of the distribution of aleatoric uncertainty in the scaled input geoscience features. Plot (a) represents one-hot encoded features that include categorical geology data, while plot (b) comprises the rest of the non-encoded geoscience features. The red lines denote the median uncertainty, while the blue boxes illustrate the interquartile range, with whiskers representing the overall spread of uncertainty. 100
Figure 5.9 Spatial distribution maps of aleatoric uncertainty across Canada for two geoscience features: (a) geological faults and (b) the shape index of the gravitational field. Regions with higher uncertainty indicate areas where input features display greater variability. 101
Figure 5.10 Spatial distribution of relative uncertainty in mineral prospectivity across Canada. The map also shows five key geological province boundaries that contain over 90 % of Canada's known Ni deposits. 102
Figure 5.11 KDE plots of the relative uncertainty distributions in the testing dataset for positive and negative regions across major geological provinces in Canada. The density signifies the concentration of relative uncertainty values, whereas peaks denote the most frequent uncertainty levels. For clarity and because they contain more than 88 % of known deposits, only the Superior, Churchill, and Southern provinces are shown. . . . 103

Figure 5.12 Map (a) depicts the spatial distribution of risk-return categories across Canada. Subplot (b) provides a zoomed-in view of a region within the Churchill Province, whereas subplot (c) focuses on a region within the Superior Province. Known Ni deposits and occurrences are included in the subplots as reference points to demonstrate their alignment with the risk-return categories. 105
Figure 5.13 Bar chart comparing the area and known deposit percentages across risk-return categories. The percentages are calculated based on the proportion of the total area (blue bars) and the total number of known deposits (orange bars, out of a total of 1,452) in each category, respectively. 106
Figure 6.1 Simplified geological map showing the distribution of major rock categories across Canada. The locations of the known Ni deposits and occurrences used to train and test the models are indicated by white circles, and the boundaries of key geological provinces are highlighted to demonstrate that the majority of these deposits are situated within these specific terranes (Modified from Lawley et al. (2024)). 114
Figure 6.2 Schematic workflow of the proposed MPM framework comparing three Bayesian deep learning methods. 118
Figure 6.3 Figures showing TAGI models' (a) network architecture and (b) the compact representation of the analytical information flow (Modified from Goulet et al. (2021)). 121
Figure 6.4 Mean ROC curves and AUC values for the TAGI, BNN with VI, and MC dropout models on the unseen test set. The dashed diagonal line represents the performance of a random classifier with an AUC of 0.5. 129
Figure 6.5 MPM for magmatic Ni (±Cu ±Co ±PGE) sulphide deposits across Canada, generated by the TAGI model. High prospectivity scores, approaching one, are shown in yellow, while the lowest scores, approaching zero, are shown in blue. 131
Figure 6.6 P-A curve for the TAGI model, showing the relationship between the true positive rate and the proportion of total land area flagged as prospective at different classification thresholds. 132

Figure 6.7 Binary MPM for magmatic Ni (±Cu ±Co ±PGE) sulphide deposits in Canada, derived from the TAGI model using an optimal classification threshold of 0.5. Prospective target areas are shown in red, while non-prospective regions are shown in gray. Black dots represent known Ni deposit locations. Additionally, the testing dataset used to evaluate model performance is illustrated in purple regions. 133

Figure 6.8 Relative uncertainty map for the TAGI model across Canada, where purple areas indicate low relative uncertainty and high model confidence, while green areas indicate high relative uncertainty and low model confidence. The locations of known Ni mineral deposits and occurrences are shown as black dots. 134

Figure 6.9 SHAP summary plot showing the feature importance for the TAGI model's prospectivity predictions. Each feature's violin shows the distribution of its SHAP values across the test set. The position on the x-axis indicates the impact on the model's output, while the colour represents the feature's value from high (red) to low (blue). The features are ranked from most to least important. 136

Figure 6.10 SHAP violin plot showing the feature importance for the TAGI model's uncertainty across the test set. The features are ranked by importance, where a positive SHAP value on the x-axis corresponds to an increase in output variance. The colour bar represents the feature's value, with high values shown in red and low values in blue. 137

LIST OF SYMBOLS AND ACRONYMS

US	United States
Cu	Copper
Li	Lithium
Ni	Nickel
Co	Cobalt
SDS	Sustainable Development Scenario
STEPS	Stated Policies Scenario
Mn	Manganese
ESG	Environmental, social and governance
kt	Thousand tonnes
PGE	Platinum group element
ML	Machine learning
MPM	Mineral prospectivity modelling
NLP	Natural language processing
AI	Artificial intelligence
DL	Deep learning
GIS	Geographic Information System
WofE	Weights of evidence
ANN	Artificial neural network
MLP	Multi layer perceptron
SVM	Support vector machine
RF	Random forest
CNN	Convolutional neural network
RNN	Recurrent neural network
LSTM	Long short-term memory
GRU	Gated recurrent unit
GAN	Generative adversarial network
cGAN	Conditional generative adversarial network
SMOTE	Synthetic minority over-sampling technique
CVAE	Conditional variational autoencoder
PU	Positive-unlabelled
MVT	Mississippi Valley-type
SHAP	Shapley additive explanation

ICE	Individual conditional expectation
LIME	Local interpretable model-agnostic explanations
TAGI	Tractable approximate Gaussian inference
BNN	Bayesian neural network
VI	Variational inference
MC	Monte Carlo
XGBM	Extreme gradient boosting machine
PGE	Platinum group elements
DGGS	Discrete global grid system
LAB	Lithosphere-asthenosphere boundary
MOHO	Mohorovičić discontinuity
1VD	First vertical derivative
HGM	Horizontal gradient magnitude
IQR	Interquartile range
VAE	Variational autoencoders
D_{KL}	Kullback-Leibler divergence
ELBO	Evidence lower bound
SiLU	Sigmoid linear unit
SI	Sensitivity indices
FI	Feature importance
AUC	Area under the curve
RMSE	Root mean squared error
XGBoost	Extreme gradient boosting machine
CGMC	Canada Geological Map Compilation
NLL	Negative log-likelihood
ROC	Receiver operating characteristics
KDE	Kernel density estimation
ADASYN	Adaptive synthetic technique
ReLU	Rectified linear unit
GMA	Gaussian multiplication approximation
RTS	Rauch-Tung-Striebel
TP	True positive
TN	True negative
FP	False positive
FN	False negative
P-A	Prediction-area

GMS	Geometric mean score
MB	Megabytes
LLM	Large language model
PINN	Physics-informed neural network

LIST OF APPENDICES

Appendix A Description of geoscience evidence layers 181
Appendix B Threshold analysis for Bayesian model comparison 184

CHAPTER 1 INTRODUCTION

1.1 Global importance of minerals

Metals and minerals serve as essential elements in the development of human society, which form the basis for modern infrastructure and advanced technology. As a result, the mining industry functions as a key driver of global economic growth that shapes the development paths of resource-rich countries, such as China, Canada, Australia, and the United States (US). Indeed, metrics quantifying mineral wealth and annual production capacities are frequently utilized as critical indicators of national developmental status. This economic centrality is highlighted by the industry's significant scale. For example, in 2020, the global mining sector produced 21.8 billion tons of raw materials, representing an aggregate valuation of approximately 2.31 trillion US dollars (China Geological Survey, 2021; Liu et al., 2025a).

However, the context of mineral exploration is shifting as demand for a suite of commodities collectively termed critical minerals rapidly increases with the global energy transition. A mineral is generally considered critical if it is essential to economic or national security and required for the transition to a sustainable, low-carbon, and digital economy (Stinson and Cam, 2025). This transition, broadly characterized by the shift from fossil-fuel-based energy systems to low-carbon technologies like wind, solar, and lithium-ion batteries (Bazilian et al., 2019), is redefining exploration priorities worldwide.

This shift can be seen clearly in Figure 1.1, which highlights annual demand growth for selected critical minerals, including Copper (Cu), Lithium (Li), Nickel (Ni), and Cobalt (Co), from the period of 2010–2020 and projected trends through 2040 (International Energy Agency, 2021). Under the Sustainable Development Scenario (SDS), the demand for these essential battery metals is anticipated to increase at rates significantly higher than those observed in the previous decade. Most importantly, recent industrial data indicate that these aggressive projections are materializing. For example, the International Energy Agency (2025) reports that Li demand increased by nearly 30 % in 2024, while Ni and Co demand grew by 6-8 %, primarily due to the electric vehicle sector.

Addressing the increasing demand poses significant supply challenges. Recent projections indicate that by 2035, critical imbalances will emerge between supply and demand. For example, Cu is expected to experience a potential 30 % supply shortfall, whereas Li confronts a 40 % deficit. These challenges are attributed to declining ore grades, rising capital expenditures, and a marked slowdown in the discovery of new resources (International Energy

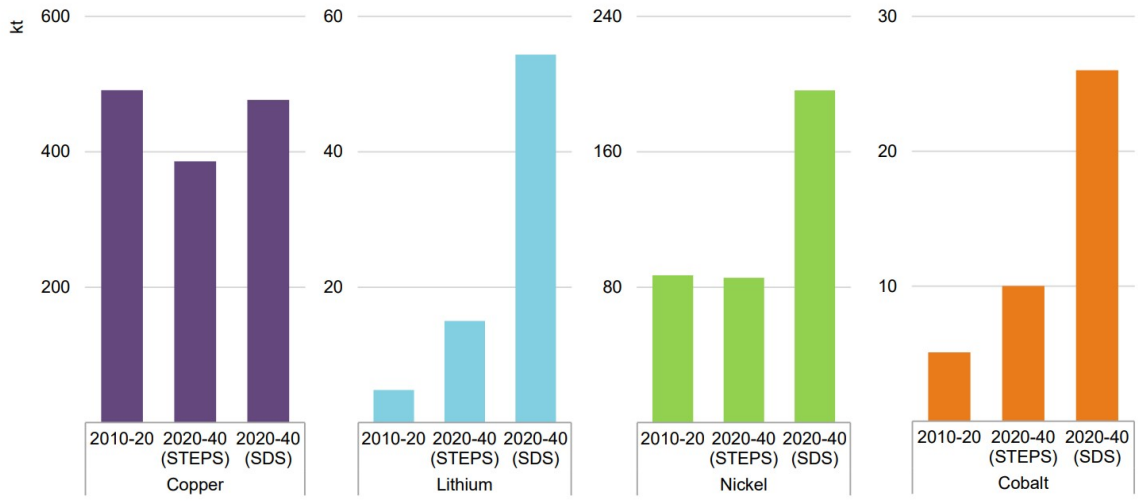


Figure 1.1 Average annual demand growth for Cu, Li, Ni, and Co minerals by scenario. kt = thousand tonnes; STEPS = Stated Policies Scenario; SDS = Sustainable Development Scenario (International Energy Agency, 2021).

Agency, 2025). Therefore, in order to address the imbalance between supply and demand, it is essential to discover new mineral deposits at an unprecedented rate. This situation presents both challenges and opportunities for resource-rich nations like Canada.

Compounding these supply challenges, the mineral exploration industry itself is also facing considerable difficulties as discovery rates for significant deposits have declined substantially from a peak of 202 deposits in 2009 to just 70-90 deposits annually in recent years (Figure 1.2). Similar to that, the cost per discovery has increased dramatically, rising fivefold from 53 million US dollars in 2003 to a peak of 248 million US dollars in 2015. Despite recent advancements, costs remain higher than pre-2005 levels, averaging approximately 170-190 million dollars per significant deposit. This decline stems from a fundamental shift in the nature of mineral exploration itself, in which over 90 % of known deposits discovered in the early 1900s were exposed at the near-surface, making them easily accessible to prospectors. In contrast, the current industry is required to utilize advanced exploration methods to identify concealed deposits at progressively greater depths (Schodde, 2025).

As illustrated in Figure 1.2, this progression toward deep-cover exploration has a dual trend of declining discovery rates and escalating unit costs. The meters drilled per discovery have more than doubled from 0.31 million meters in 2005 to 0.70 million meters in 2015, reflecting the twin challenges of increasing depth and diminishing availability of shallow targets. Additionally, the relationship between depth and discovery cost is non-linear, where increasing average exploration depth from 25 meters to 75 meters approximately doubles discovery costs,

with cost escalation accelerating significantly beyond 200 meters of cover. In well-explored terrains, commonly referred to as brownfield areas, exploration programs now drill two to three times deeper than in greenfield settings. Therefore, this indicates that near-surface targets are increasingly depleted and that exploration is progressively shifting toward deeper, covered and complex environments (Zadeh, 2025). This convergence of declining discovery rates, escalating costs, and increasing technical complexity necessitates the development of more effective, data-driven exploration methodologies that are capable of detecting the subtle, non-linear signatures of mineralization in these challenging geological settings.

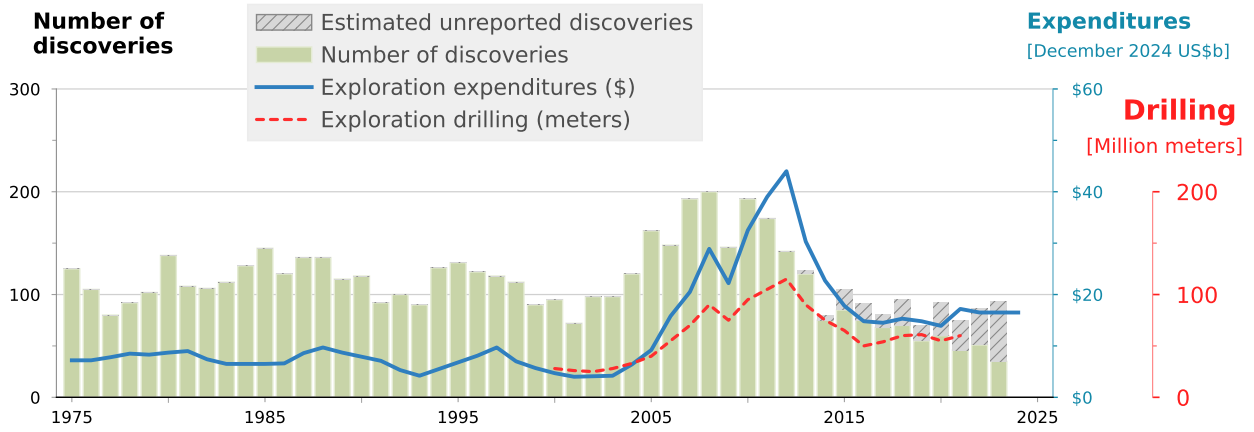


Figure 1.2 Discovery cost and meters drilled per significant mineral deposit discovered (1979 – 2023). Data for 2013 to 2023 include an adjustment for unreported discoveries. Modified from Schodde (2025).

1.2 Critical minerals in Canada

Canada's commitment to achieve carbon neutrality by 2050 requires a major shift away from hydrocarbon extraction, processing, and consumption (Government of Canada, 2021). For example, road transportation accounts for approximately 20 % of Canada's total greenhouse gas emissions and represents a key sector for reducing carbon footprint (Environment and Climate Change Canada, 2021; Gadd et al., 2023). A key component of this transition is the rapid adoption of electric vehicles, which is supported by the federal mandate that all new light-duty vehicle sales are expected to be zero-emission by 2035 (Transport Canada, 2021). However, the material intensity of this technological shift is profound. Electric vehicles require approximately six times as much mineral input as internal combustion automobiles, requiring notable volumes of Cu, Li, Ni, manganese (Mn), Co, and graphite (International Energy Agency, 2021). Achieving national net-zero objectives, therefore, demands an estimated

16-fold increase in electric vehicle sales, driving exceptional demand for critical minerals.

In recognizing both the strategic risks and economic opportunities inherent in the energy transition, the Government of Canada has defined a list of 34 critical minerals based on their importance to national economic security, the nation's progress toward a low-carbon transition, and its capacity to serve as a reliable supplier to allied nations (Government of Canada, 2024). Figure 1.3 depicts this compilation of 34 minerals, which are sourced from diverse mineral systems and can be extracted as primary commodities or obtained as co-products and by-products of base-metal operations. Six of these minerals, including Li, graphite, Ni, Co, Cu, and rare earth elements, have been identified as priority minerals due to their pronounced potential to support Canadian economic expansion (Government of Canada, 2024).

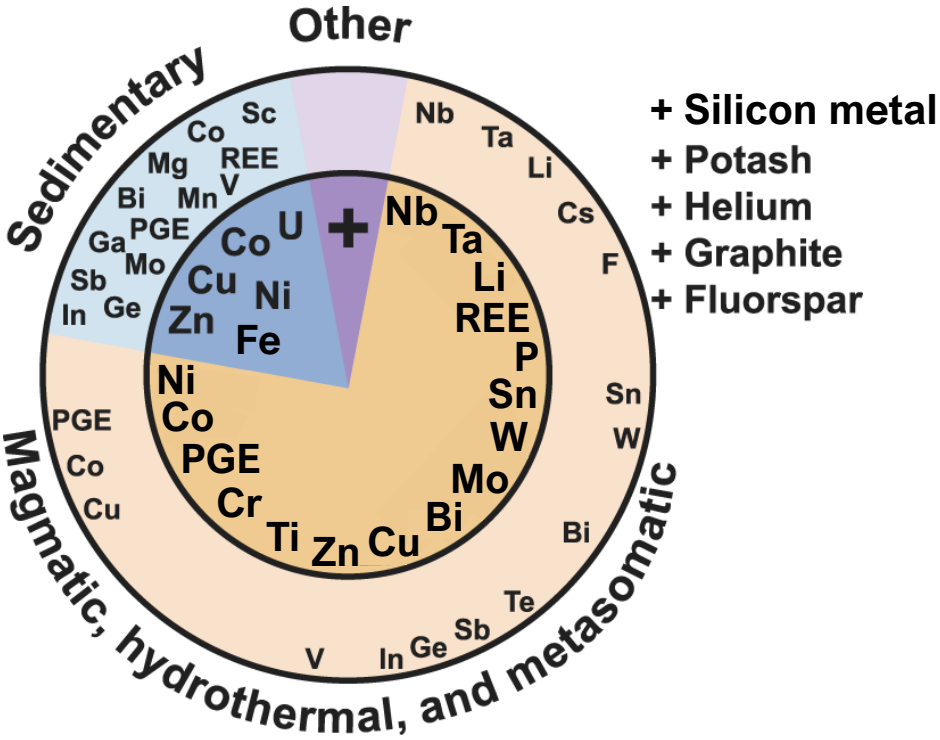


Figure 1.3 List of 34 Canadian critical minerals as of 2025 that shows primary and co-products (center) and common by-products (outer margin), separated by mineral system. Modified from Gadd et al. (2023); Government of Canada (2024).

Canada's position in the global critical minerals market is reinforced by several competitive strengths. The country hosts vast geological resources across ancient cratons and metallogenic regions spanning nearly all provinces and territories (Figure 1.4). These geological strengths are further supported by Canada's strong environmental, social, and governance

(ESG) performance. In contrast to many countries where critical mineral extraction is associated with environmental degradation and labour concerns, Canadian operations benefit from cleaner energy systems and stringent regulatory standards. For example, producing Ni in Canada generates nearly ten times fewer greenhouse gas emissions than production in Indonesia, which is currently the world's largest Ni producer (Government of Canada, 2024). Together, these advantages, including Canada's substantial mineral endowments, strong governance frameworks, and high ESG performance, have led to strong global rankings. As a result, an analysis conducted in 2024 by BloombergNEF ranked Canada as the world leader in the potential for a secure and sustainable battery supply chain, and second in overall competitiveness in critical minerals production (Government of Canada, 2024).

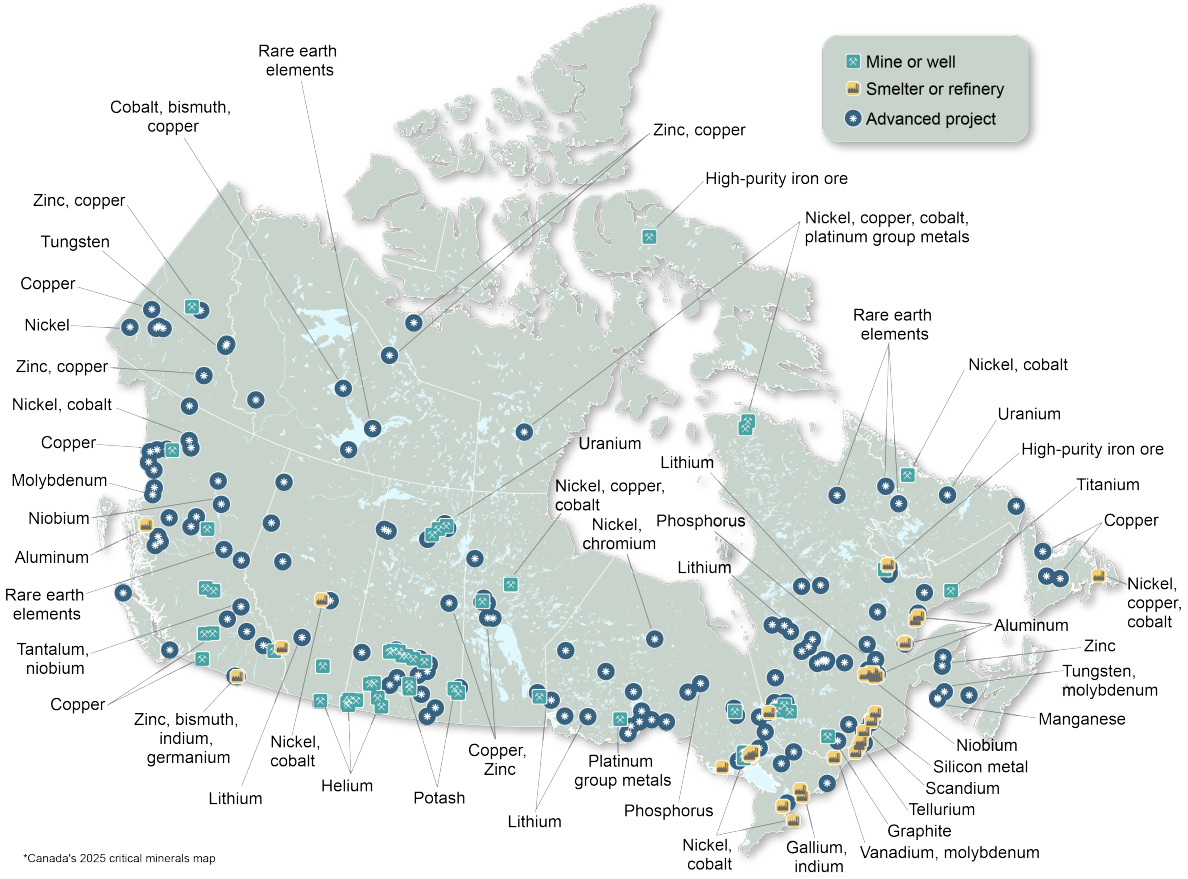


Figure 1.4 Critical mineral development across Canada. Critical mineral mines, smelters, refineries, or advanced projects are located across all Canadian landmass, except for Prince Edward Island (Government of Canada, 2024).

1.2.1 Nickel as a priority critical mineral

Among Canada's priority critical minerals, Ni is particularly important because of its central role in advanced battery technologies. It is a foundational component of Li-Ni-Mn-Co oxide and Ni-Co-aluminum oxide cathodes, where increased Ni content enhances energy density and extends driving range, which makes them key attributes for widespread electric vehicle adoption. As illustrated in Figure 1.1, global demand growth trends highlight this importance as under the Stated Policies Scenario (STEPS), which serves as a conservative benchmark based on existing policy frameworks and prevailing market trends, average annual demand growth for Ni is projected to remain relatively stable at approximately 85 thousand tonnes (kt) per year during both 2010 to 2020 and 2020 to 2040 periods. However, under the SDS, which models the transformation required to fully meet global climate and sustainable energy goals, it is expected that Ni demand growth will accelerate sharply to nearly 200 kt per year during 2020–2040, which is more than double the historical rate (International Energy Agency, 2021). The 2023 statistics show that stainless steel production has dominated global Ni consumption, accounting for approximately 64 % of demand, while the battery sector currently represents roughly 16 % and is expected to grow rapidly by 2040 (Natural Resources Canada, 2023).

With approximately 4–5 % of global Ni mine production, Canada remains one of the world's leading Ni producers and is well positioned to contribute to future supply needs (Natural Resources Canada, 2023). Mainly, the Sudbury Basin in Ontario, which is formed by a meteorite impact, is among the world's most significant Ni–Cu–platinum group element (PGE) systems, while Voisey's Bay in Labrador hosts another high-grade Ni-Cu-Co mineralization within layered mafic–ultramafic intrusions (Grieve et al., 1991; Latypov et al., 2019; Li and Naldrett, 1999). Additionally, Manitoba's Thompson Ni belt contains komatiite-hosted Ni sulphide deposits, and Ontario's Ring of Fire region comprises major, yet largely undeveloped, Ni-chromium resources within ultramafic complexes (Bleeker and Kamo, 2020; Gadd et al., 2023; Layton-Matthews et al., 2010).

Even with Canada's rich Ni resources, meeting future demand will require finding new deposits faster than ever before. Much of the untapped Ni potential in Canada, however, lies beneath thick cover, at greater depths, and in complex geological environments, which limit the effectiveness of standard surface exploration tools. Geophysical surveys can image subsurface anomalies, but these signals are often ambiguous and typically require costly drilling to confirm. These challenges underscore the need for new exploration methodologies that integrate diverse geoscience datasets to identify prospective areas in data-sparse, geologically complex regions.

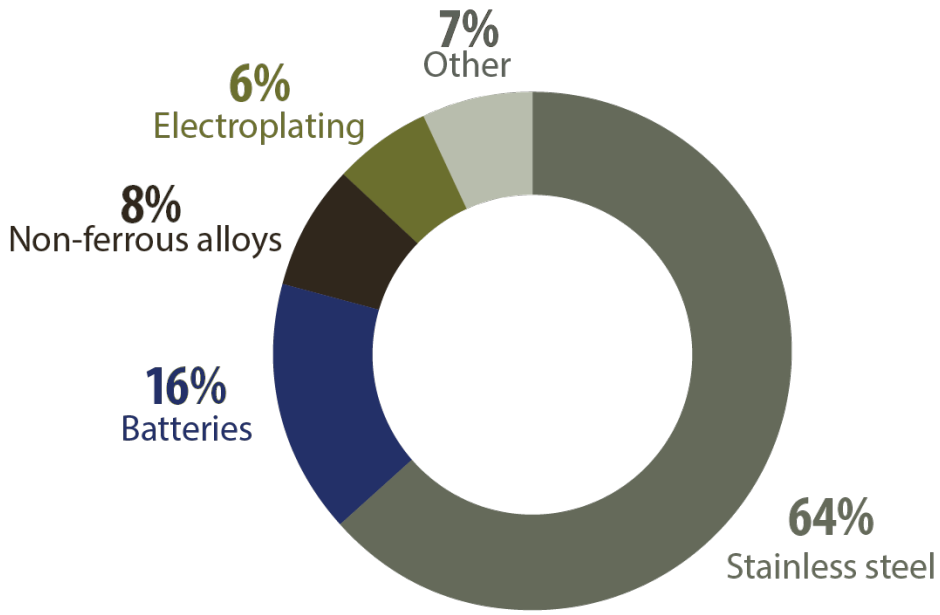


Figure 1.5 Distribution of global Ni demand on major industrial applications in 2023 as stainless steel manufacturing remains the principal consumer, with batteries representing the second-largest and rapidly growing segment, followed by non-ferrous alloys, electroplating, and other minor sectors (Natural Resources Canada, 2023).

1.3 Addressing exploration challenges through data-driven approaches

1.3.1 The data integration challenge

As mineral exploration programs have expanded over the years, they have produced large volumes of diverse, multi-scale geoscience data compilations that can be integrated to delineate prospective areas for further investigation. These datasets encompass multiple domains that includes airborne and ground-based geophysical surveys that assess variations in magnetic, gravity, seismic, electromagnetic, and radiometric properties; regional and detailed geochemical analyses of rock, soil, sediment, and water samples; geological mapping and structural interpretation based on field observations and remote sensing; as well as satellite-based multispectral and hyperspectral imagery that records surface mineralogy and alteration signatures (Sun et al., 2024; Zekang et al., 2025). Most importantly, each of these data types is acquired at different spatial resolutions, with varying degrees of coverage and quality, and represents fundamentally different physical or chemical properties of the subsurface (Lv and Wang, 2024; Porwal et al., 2003a). For example, aeromagnetic surveys are able to provide continuous coverage across entire regions, whereas geochemical sampling is typically con-

ducted at discrete, irregularly spaced locations constrained by the natural distribution of the target sampling medium (e.g., drainage networks for stream sediments, glacial deposits for till surveys, or bedrock exposure for rock samples). Additionally, some of these datasets may contain considerable missing values and noise.

However, traditional exploration methods face challenges when integrating heterogeneous datasets, especially in areas where multiple layers of post-mineralization cover obscure direct surface expressions of mineralization (Hajihosseinlou et al., 2024; Katiyar et al., 2024). This data integration challenge is compounded by the fact that geological, geophysical, and geochemical data often exhibit complex, non-linear relationships with mineralization that vary depending on deposit type, geological setting, burial depth, and degree of preservation. While knowledge-driven data integration approaches that integrate geological insights and experience have shown promising results, the outcomes can be subjective and vary significantly across interpreters' individual experience, theoretical perspectives, and regional familiarity (Zuo et al., 2023). This limitation has become problematic when dealing with dozens of data layers, each containing hundreds of thousands to millions of data points, where subtle but notable patterns may exist in high-dimensional feature spaces that are not readily seen through visual inspection or simple overlay analysis. These limitations can result in missed exploration opportunities, inefficient resource allocation, or suboptimal targeting decisions. Beyond data quality and integration challenges, exploration datasets are also affected by survival bias, where known deposits are geographically concentrated in well-explored regions with obvious surface expressions or strong geophysical signatures. Put simply, areas with weak or ambiguous mineralization signals may still host significant deposits beneath cover or in atypical geological settings, yet our training datasets contain few examples of deposits discovered in such challenging environments (Yousefi et al., 2024b).

Compounding the spatial bias in the training data, whereby known deposits are disproportionately concentrated in well-explored regions, is the inherent problem of class imbalance. Mineral deposit occurrences are rare events relative to the vast geological landscape being assessed, resulting in training datasets where known mineralized locations are outnumbered by barren grid cells by ratios that commonly exceed 1:100 or even 1:1000. Standard classification algorithms trained on such heavily skewed distributions tend to favour the dominant non-deposit class, which reduces their ability to correctly identify the mineralized locations that are of primary interest to exploration targeting. Addressing these challenges, including data heterogeneity, integration complexity, class imbalance, and training-set biases, has driven the adoption of machine learning (ML) approaches for mineral prospectivity prediction.

1.3.2 Adaptation of machine learning for mineral prospectivity modelling

ML techniques can overcome the limitations of the aforementioned approaches, because of their capability to automatically identify complex, non-linear patterns within high dimensional datasets; ability to infer spatial relationships directly from known deposit examples without the necessity for explicit mathematical formulation; and potential to analyze large datasets efficiently (Lary et al., 2016; Nykänen and Ojala, 2007; Sihombing et al., 2024). These advantages have positioned ML as a turning point in mineral prospectivity modelling (MPM), which helps to reduce the search area for exploration by identifying regions with higher potential for discovering new deposits (Zuo et al., 2023).

First, mineralization results from complex interactions between source rocks, fluid pathways, chemical reactions, structural controls, and preservation mechanisms operating across multiple scales (Sun et al., 2024; Zekang et al., 2025), which will be further discussed in the literature review section. ML algorithms can better capture these highly nonlinear relationships without requiring explicit functional forms or independence assumptions, and can learn patterns from geoscience features. Second, ML models can provide reproducible integration frameworks where manual interpretation through multiple people could reach different conclusions from the same datasets. Most importantly, this is particularly valuable for large-scale prospectivity modelling studies that span over diverse geological terrains where maintaining consistent standards becomes challenging. Third, ML methods offer greater scalability, as once trained on one region, they can be adapted to new areas with similar characteristics, making them more efficient (Lauzon and Gloaguen, 2024; Mantilla-Dulcey et al., 2024).

Beyond integrating numerical geoscience datasets, ML techniques can also process the unstructured text data contained within geological map databases, which is a previously largely untapped source of exploration information (Parsa et al., 2025; Zhou et al., 2024). These geological databases contain vast volumes of detailed rock descriptions, lithostratigraphic relationships, and interpretations recorded as text attributes that are difficult to systematically incorporate into prospectivity analyses using conventional methods (Brodaric et al., 2004; Lawley et al., 2023; Laxton, 2017). Natural language processing (NLP) is a subfield of artificial intelligence (AI) that focuses on interpreting human language. It addresses this challenge by converting textual descriptions into numerical representations (word embeddings) that can be processed by ML algorithms (Devlin et al., 2019; Lawley et al., 2022c; Mikolov et al., 2013a,b; Pennington et al., 2014). Geoscience-specific language models, developed by training on geological publications and map databases, have demonstrated superior performance compared to generic language models for extracting geological knowledge from text (Fuentes et al., 2020; Padarian and Fuentes, 2019). With further developments in ML

models, they are now able to analyze mineral exploration reports, extracting geological entities and relationships for knowledge graph construction, and semantic search for deposit analogue (Lawley et al., 2023; Ma et al., 2022; Wang et al., 2022). Therefore, ML-based NLP unlocks previously inaccessible geological knowledge embedded in text data, expanding the scope and information content available for data-driven exploration targeting.

Publications applying ML to MPM have surged from dozens of studies in the early 2010s to hundreds by 2025 (Figure 1.6). This growth reflects not only increasing data availability and computational resources but also methodological shifts from Geographic Information System (GIS) based integration in the mid 2000s, through classical ML in the 2010s, to the current era dominated by deep learning (DL) architectures (Zekang et al., 2025).

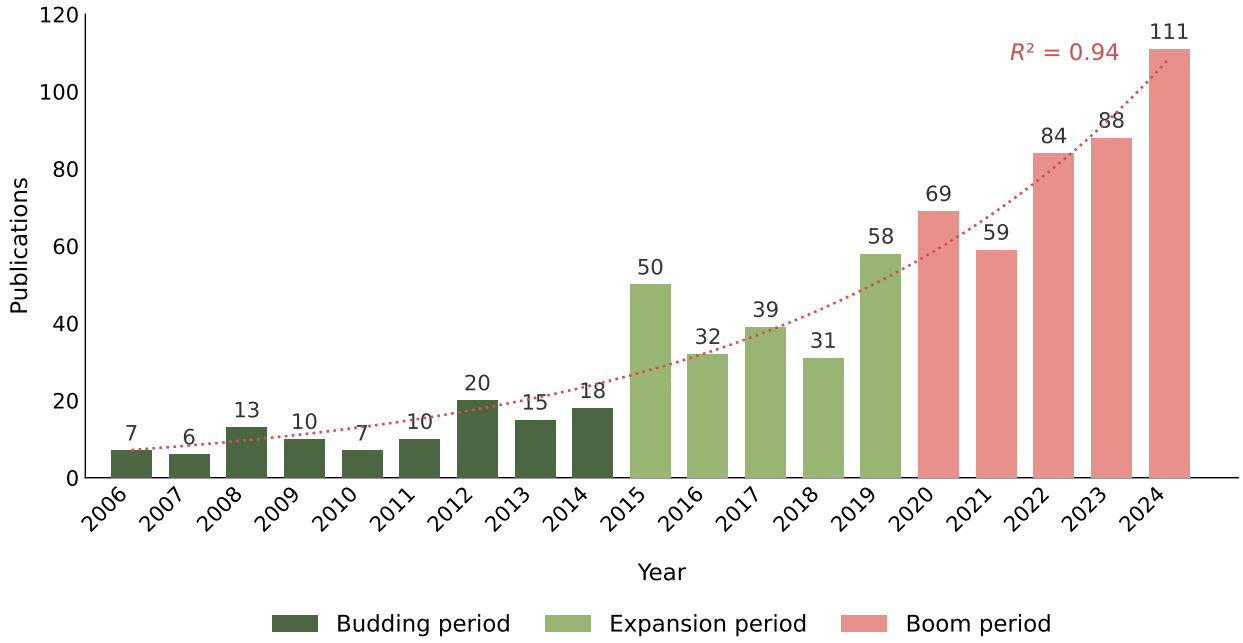


Figure 1.6 Exponential growth in MPM publications from the Web of Science Core Collection database from 2006 to 2024, showing budding (modest output), expansion (steady growth), and boom (rapid acceleration) phases. Modified from Zhang and Sun (2025).

Although industry adoption of MPM still remains limited, several recent industrial case studies demonstrate the practical value of data-driven targeting. One notable example is Cartier Resources, which implemented an AI-assisted prospectivity workflow in the Cadillac mining district of Val-d'Or, Québec, which is one of Canada's most historically productive gold-mining regions. More than 150,000 drill-hole assays, together with extensive geochemical, structural, and geophysical datasets, were standardized and integrated into a predictive modelling framework. The resulting prospectivity maps, as illustrated in Figure 1.7, high-

lighted previously underexplored zones, informed drill-budget prioritization, and supported transparent, reproducible decision-making through 3D visualization of targets. Notably, approximately one-quarter of Cartier's upcoming drilling program is dedicated to testing AI-generated targets, illustrating how ML can enhance target generation in data-rich exploration terrains (VRIFY, 2025). This data-driven learning complements geological expertise, in which geologists guide feature selection based on deposit models, while ML automates objective integration.

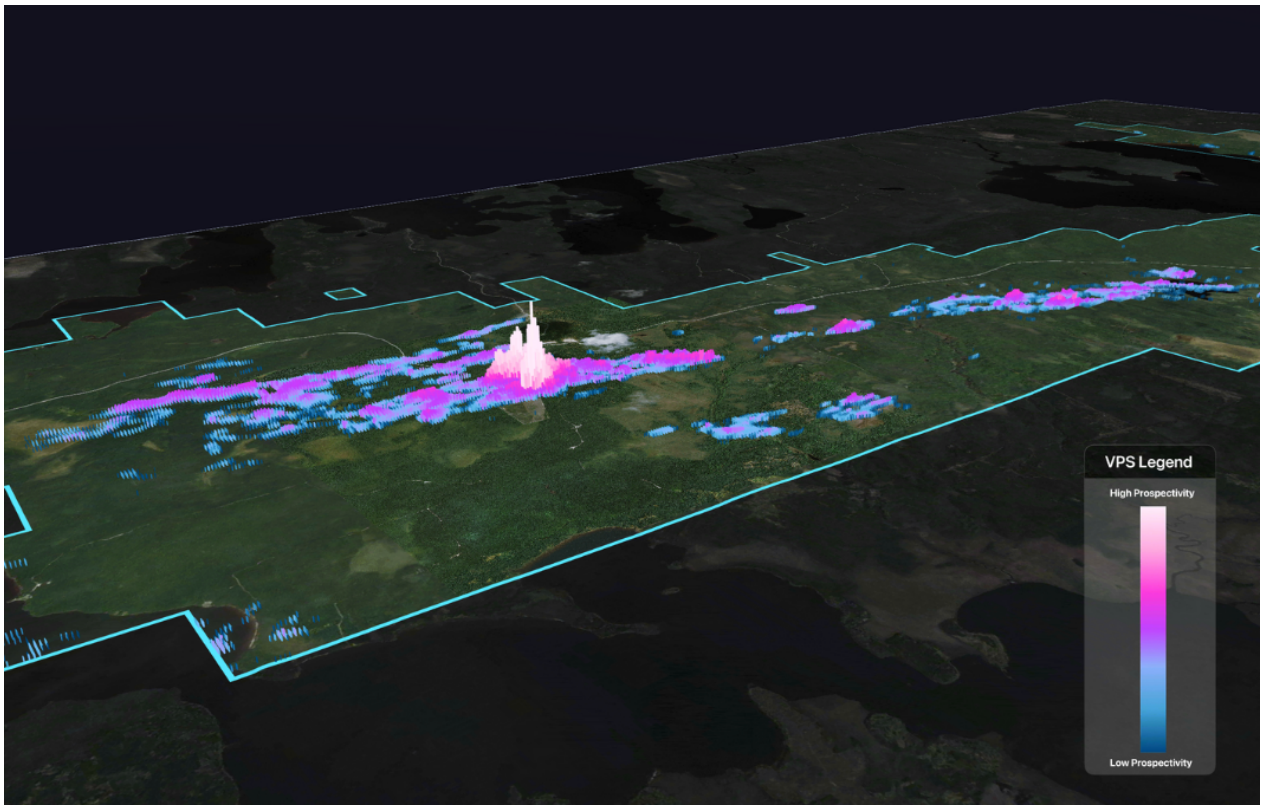


Figure 1.7 AI-assisted prospectivity model for the Cadillac gold mining district (Val-d'Or, Québec), illustrating spatial variations in predicted gold prospectivity. Warmer colours, including pink-white, represent higher prospectivity, while cooler colours, including blue, indicate lower prospectivity values (VRIFY, 2025).

While classical ML methods have demonstrated considerable value in prospectivity mapping, most applications remain deterministic, providing only point estimates of prospectivity without quantifying confidence in those predictions. This limitation is problematic in mineral exploration, where drilling decisions carry substantial financial risk: false positives waste millions of dollars on unproductive drill programs, while overlooking uncertain-but-prospective areas may result in missed discoveries. Moreover, deterministic predictions offer no mechanism to distinguish between high-confidence targets in well-constrained geological settings

and lower-confidence predictions in data-sparse or geologically complex regions.

Probabilistic ML techniques address these limitations by explicitly modelling both aleatoric uncertainty, which is the inherent randomness and noise in geoscience data, and epistemic uncertainty, which denotes the model's lack of knowledge due to limited training data or geological complexity (Caers, 2025). Crucially, these frameworks are capable of differentiating between predictions derived from data-rich geological environments and those subject to high uncertainty arising from sparse or geologically complex regions. Furthermore, probabilistic methods can address the data scarcity and class imbalance challenges that characterize MPM applications by leveraging generative modelling in order to augment limited training datasets while simultaneously quantifying the uncertainty introduced by synthetic data generation. Therefore, developing probabilistic frameworks for complex mineral systems at regional to national scales represents a critical frontier in advancing exploration targeting methodologies.

1.4 Research objectives

1.4.1 General objective

The general objective of this research is to develop and validate probabilistic ML frameworks and tools for national-scale MPM that address data imbalance, quantify prediction uncertainty, and provide interpretable decision-support tools for exploration targeting of magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems in Canada.

1.4.2 Specific objectives

To achieve the general objective, the following three specific objectives need to be fulfilled:

1. Evaluate whether generative models such as conditional variational autoencoder (CVAE) can address class imbalance in MPM by generating synthetic geoscience data representative of mineralized areas, and benchmark their performance against alternative data balancing techniques across multiple classification algorithms.

2. Develop a novel MPM risk-return analysis framework that integrates aleatoric and epistemic uncertainty quantification using generative ensemble modelling and a spatial clustering technique to categorize exploration targets based on both prospectivity and confidence.

3. Compare the MPM performance of analytical Bayesian inference method against gradient-based Bayesian approximations, and evaluate the interpretability of both prospectivity

predictions and their associated uncertainties through feature attribution analysis.

1.5 Dissertation structure

The work presented in the following chapters addresses the specific research objectives while collectively achieving the general objective of developing probabilistic frameworks for MPM. Chapter 2 provides a literature review that introduces the fundamental concepts of MPM and traces its evolution through different methodologies. This chapter also examines key challenges in MPM applications, including data imbalance, uncertainty quantification, and model interpretability, thereby establishing the theoretical foundation for this research and identifying the knowledge gaps that motivate the subsequent studies. Next, Chapter 3 provides a roadmap for the three published scientific articles presented in Chapters 4, 5, and 6, explaining how each study addresses a specific research objective outlined in Section 1.4.2 and contributes to the overarching goal. The contributions from these articles form a cohesive framework that supports the achievement of the general objective. The articles are organized in the following order:

1. « A balanced mineral prospectivity model of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems using conditional variational autoencoders »
   Published in *Ore Geology Reviews* (2024).

2. « A generative neural network approach to uncertainty and risk-return analysis in mineral prospectivity modelling »
   Published in *Ore Geology Reviews* (2025).

3. « Tractable approximate Gaussian inference for interpretable mineral prospectivity modelling »
   Published in *Applied Computing and Geosciences* (2026).

Following the three articles, Chapter 7 presents a general discussion that connects the findings across the studies and examines their implications for mineral exploration. Chapter 8 concludes with the main findings, contributions, and recommendations for future work.

CHAPTER 2 LITERATURE REVIEW

This chapter reviews the key concepts and methodologies that form the foundation of the research presented in this thesis. We begin by introducing MPM, including the geoscience datasets it relies on and the mineral system framework that guides the construction of predictive models. We then trace the methodological evolution of MPM, from early knowledge-driven and statistical approaches through to modern ML and DL architectures. Next, we examine the critical challenges that remain unresolved in current MPM workflows, including data scarcity, uncertainty quantification, and model interpretability for both prediction and confidence. Finally, we synthesize the insights from this review and connect them to the contributions developed in the subsequent chapters of this thesis.

2.1 Introduction to mineral prospectivity modelling

MPM is a pivotal tool in modern mineral exploration, which is designed to delineate areas with elevated potential for hosting a specific mineral deposit type (Agterberg et al., 1993; Bonham-Carter, 1994; Cheng and Agterberg, 1999; de Quadros et al., 2006; Partington and Sale, 2004). The formation of mineral deposits is a long and complex process influenced by numerous geological factors operating across multiple spatial and temporal scales. Understanding and mapping these factors requires a comprehensive analysis of diverse geological information, including structural background, lithological characteristics, mineral assemblages, alteration patterns in host rocks, geophysical and geochemical anomalies, and exploration drilling data, all of which form the foundation of prospecting and prediction research.

Early MPM work primarily focused on delineating prospective areas or target zones through qualitative metallogenic prediction, which relied heavily on metallogenic theory, an understanding of metallogenic systems, expert geological experience, and detailed field geological surveys. However, as the purpose of geological work extends beyond simply identifying mineralization to quantifying the amount and economic value of resources, the field has necessarily evolved toward quantitative prediction methods. These quantitative approaches estimate the quantity, grade, and economic potential of undiscovered mineral resources by demonstrating the economic feasibility related to mineral exploration and extraction, making quantitative prospecting prediction an essential component of modern exploration programs (Sun et al., 2024).

MPM integrates and analyzes diverse, multidimensional geoscience datasets, including geo-

logical, geophysical, geochemical, and remote-sensing data, in order to identify characteristic spatial patterns and anomalies that may predict the occurrence of mineralization. With the increasing availability of large, complex geoscience databases, MPM has now evolved into a data-intensive and multidisciplinary approach. Therefore, it is necessary to understand how such datasets can be effectively mined and integrated, which is essential for converting raw data into meaningful information, generating knowledge, and deepening insights into the geological processes that control mineralization (Zuo, 2020). The effectiveness of MPM is usually dependent on the quality, diversity, and integration of input geoscience datasets that capture the geological processes and signatures associated with mineralization.

Firstly, geological datasets serve as the foundation for MPM predictions and include information on mineralogy, ore-deposit geology, and ore-forming structures. These data help determine geological structures and rock types that extract key geological information related to mineralization. Geological data typically encompass lithological maps, structural interpretations, fault and fold geometries, stratigraphic sequences, and information about known mineral occurrences and deposits. Secondly, geophysical data are obtained by measuring variations in the Earth's physical fields, such as gravity, magnetic field intensity, electromagnetic induction responses, and seismic wave propagation. These data can respectively reveal differences in the density, magnetic susceptibility, electrical conductivity, and elastic properties of subsurface materials and provide information related to the interaction of geomaterials with natural or artificial physical fields within a selected area. Geophysical surveys are valuable because they can be used to detect mineralization and ore-controlling structures beneath surface cover. Therefore, they help to locate mineral deposits in concealed terrains where traditional geological mapping is limited. Thirdly, geochemical data are obtained by analyzing the chemical compositions of soil, rock, sediment, and water samples. Abnormal concentrations of specific pathfinder elements and diagnostic element ratios may indicate the presence of mineral deposits or their geochemical footprint. Anomaly zones are delineated through the integration and interpretation of various types of geochemical data within the study area, including stream and lake sediment surveys, soil geochemistry, lithogeochemical sampling, and indicator mineral analysis. Geochemical data are particularly effective at mapping the spatial extent of mineralized systems and their dispersion halos in surficial environments. Finally, remote sensing data are used with satellite and aerial imaging technologies to obtain information about surface features across large areas efficiently and cost-effectively. These data help identify surface geology, lithological boundaries, structural lineaments, and hydrothermal alteration minerals associated with mineralization. Advanced remote sensing techniques, including multispectral and hyperspectral imaging, can detect subtle spectral signatures of alteration assemblages such as iron oxides, clay minerals, and hydroxyl-bearing

minerals, thereby delineating areas with key indicators for mineralization (Imasu et al., 2024; Srivastava and Jain, 2025). Therefore, MPM involves the integration of anomalies and patterns identified across all these data types to generate prospectivity maps within the study area.

The MPM workflow is generally guided by a conceptual understanding of the targeted ore-forming system. In practice, this conceptual framework is known as the mineral system approach, which has become the dominant methodology for exploration targeting (Hagemann et al., 2016; Knox-Robinson and Wyborn, 1997; McCuaig et al., 2010; Wyborn et al., 1994). This mineral system approach considers all geological processes operating over a wide range of scales that are required to form and preserve a mineral deposit. Therefore, this makes the concept ideal for the MPM as it provides a clear structure for translating a single, abstract critical process (e.g., the formation of an ore deposit) into its constituent processes (e.g., source, pathway, trap) and then into targeting elements (Figure 2.1). Most importantly, it connects these geological concepts to the final, mappable criteria that we can actually use (Ford et al., 2019; Hronsky, 2011; Kreuzer et al., 2007, 2015). However, this translation requires understanding both syn-mineral subsystems and post-mineral subsystems (Yousefi and Kreuzer, 2024). Firstly, syn-mineral subsystems capture processes active during ore formation, such as reactions between ore-bearing fluids and surrounding country rocks triggered by magmatic or hydrothermal events, which can result in metal deposition within favourable trap sites along structurally active, permeable fault-fracture systems. Secondly, post-mineral subsystems capture processes that occur after mineralization, such as the interaction of tectonic forces and weathering that leads to exhumation and the dispersion of mineralized particles into surficial environments. Depending on the types of mineralization processes and scales of an exploration program, different survey methods (e.g., stream sediment or lithogeochemical surveys) can effectively map the spatial extent of geochemical anomalies and serve as typical evidence layers for MPM.

Understanding these subsystems within the broader mineral system framework governs the selection of the spatial geoscience datasets to be used in the targeting model (e.g., Carranza, 2008, 2017). The specific application of this framework to magmatic nickel sulphide mineral systems in Canada, including the identification of relevant subsystems and their corresponding mappable criteria, is detailed in the subsequent chapter 4 and 5 of this thesis. Following this, the workflow then proceeds to the next crucial step, which converts these mappable criteria into evidence layers (or predictor layers) within a GIS framework (Porwal and Carranza, 2015). These evidence layers represent the final, processed maps derived from raw geoscience features. Additionally, it is important to note that these conceptual models often need to be adapted to different geological settings, as the characteristics of local geology can

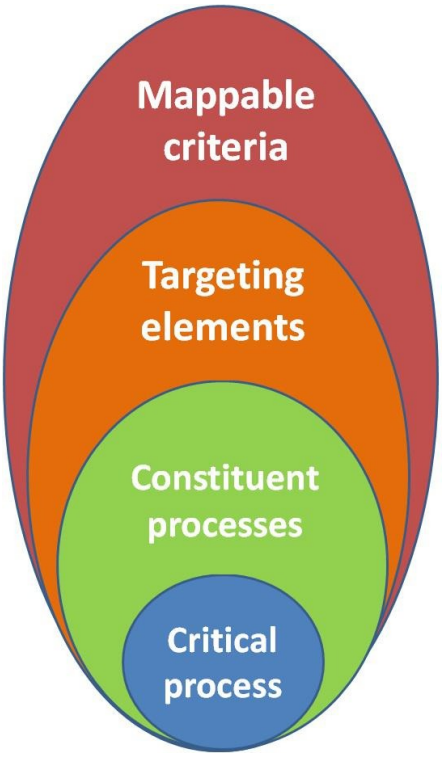


Figure 2.1 A simplified conceptual model of the mineral system, illustrating the translation from the critical process to its constituent processes and targeting elements, and finally to mappable criteria (Yousefi et al., 2019).

be diverse (Yousefi et al., 2021). Next, the evidence layers are combined using a specific integration function or algorithm to create the final prospectivity map, which is then validated against known deposit locations and expert geological knowledge. Several categories of these integration methods have been developed, each with its own advantages and disadvantages. The main difference among these categories is how they quantify the relationships between the spatial evidence layers and the mineral deposits: that is, how they weight them according to their predictive capacity.

2.2 The evolution of mineral prospectivity modelling methodologies

The integration functions used to combine evidence layers have progressed significantly over the last several decades, as illustrated in Figure 2.2. This timeline shows a clear evolution from traditional methods used in the 1980s to the rise of ML models in the 2000s and, most recently, to the DL algorithms during the last decade.

While this timeline provides the historical context, we can also classify these integration

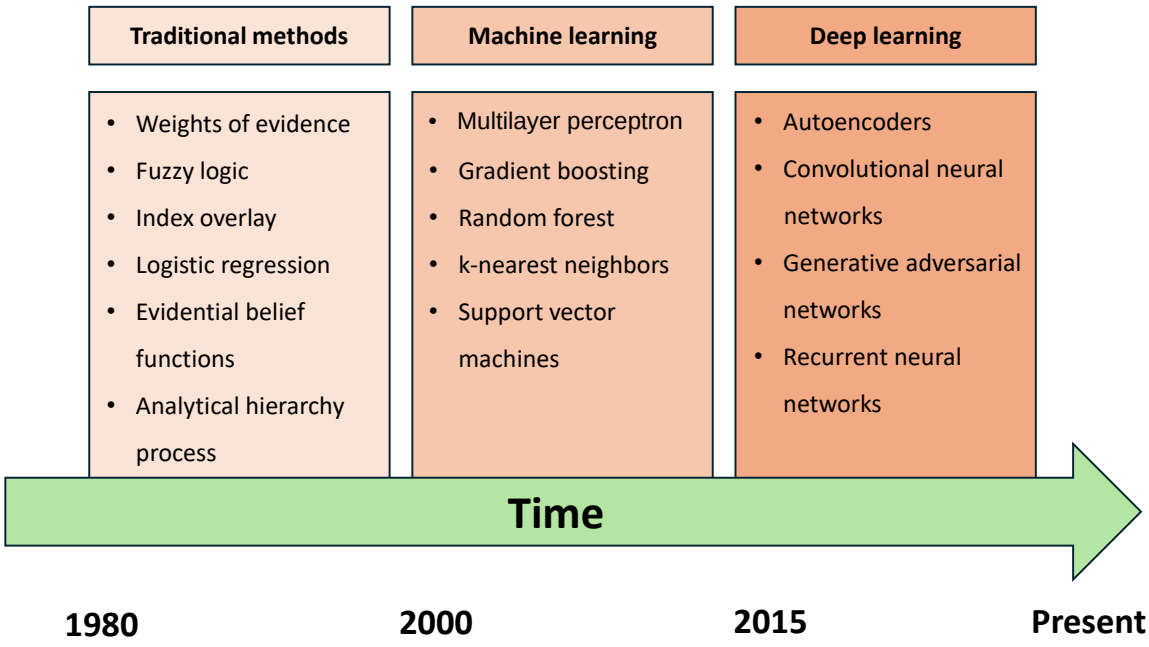


Figure 2.2 A chronological overview of MPM methodologies, showing the development from traditional methods to advanced DL models. Modified from Zuo (2020).

functions by their underlying principles. In particular, they differ in how they quantify the spatial relationships between exploration variables and known mineral deposits, and weight each evidence layer according to its predictive capacity. Several categories of methods have been developed, each with distinct advantages and limitations. Accordingly, these methods can be broadly grouped into five main categories: (1) knowledge-driven models, (2) data-driven models, (3) hybrid models, (4) modelling with empirical functions, and (5) logistic-based approaches (Yousefi et al., 2021).

2.2.1 Knowledge-driven models

Knowledge-driven MPM is appropriate in uncharted or under-explored (greenfield) regions where known occurrences of the desired mineral deposits are absent or scarce (Porwal et al., 2003b). When creating evidentiary maps (i.e., estimating evidential map weights and evidential class scores) and integrating these evidentiary maps, it is taken into account whether the location is promising for mineral deposits of the targeted type. Thus, the phrase knowledge-driven describes the qualitative evaluation or weighting of evidence in relation to a claim. The estimates of weights for each evidence map and the estimates of scores for each category reflect an expert's evaluation of the spatial link between the sought-after mineral deposits

and each set of indicative geologic features. Therefore, it is also known as expert-driven MPM (Carranza and Hale, 2008). In fact, the expert knowledge used in knowledge-driven MPM is usually acquired through extensive field experience in mineral exploration or through the application of spatial analytical techniques to study the spatial distributions of sought-after mineral deposits and their spatial associations with specific geological features.

The most common knowledge-driven methods can be differentiated by how they represent this expert judgment, moving from simple binary models to more complex, ambiguous ones. The simplest method is Boolean logic, which requires the expert to set a hard threshold for each evidence layer, resulting in a binary map of TRUE (1) for favourable areas and FALSE (0) for unfavourable ones. These binary maps are then combined using logical operators like AND or OR (Bonham-Carter, 1994). An example of this method is shown in Figure 2.3, which depicts a simple inference network for combining evidential maps. In this network, the OR operator is used to combine similar features (e.g., different fault trends) into the evidence of structural permeability, while the AND operator is used to combine essential, co-existing factors, such as evidence of hydrothermal fluid circulation and a geochemical anomaly, to delineate the final prospectivity zone (Carranza and Hale, 2008). A more flexible version is the index overlay (or weighted overlay) method. In this approach, each evidence layer is still often classified, but the expert assigns a numerical weight (w_i) to each layer to represent its relative importance, and the layers are then combined using a weighted average (e.g., Carranza and Hale, 2008; de Araújo and Macedo, 2002; Sadeghi and Khalajmasoumi, 2015). However, the most sophisticated and widely used method is fuzzy logic (Zadeh, 1965), which is well-suited to geology-related applications as it explicitly handles ambiguity. Instead of a hard 0 or 1 class label, an expert assigns a fuzzy membership score (a continuous value from 0.0 to 1.0) that represents the degree to which a location is favourable (e.g., very close to a fault). These fuzzy maps are then integrated using flexible fuzzy operators, such as the fuzzy AND, OR, or GAMMA, and have been applied to a wide range of studies over the years (Feltrin, 2008; Rogge et al., 2006; Tangestani and Moore, 2003; Zhang and Zhou, 2015).

The primary strength of this knowledge-driven approach is its ability to operate in greenfield or data-poor regions where a sufficient training dataset of known deposits are unavailable. Their main weakness, however, is that the final prospectivity map is heavily dependent on the subjective judgments of the expert who defines the weights and fuzzy membership curves; two different experts may produce different results. Furthermore, many of these methods, such as index overlay, combine data in a simple, linear, additive way, which does not naturally capture the complex, non-linear interactions of geological processes. These significant limitations led researchers to develop and focus more on data-driven techniques.

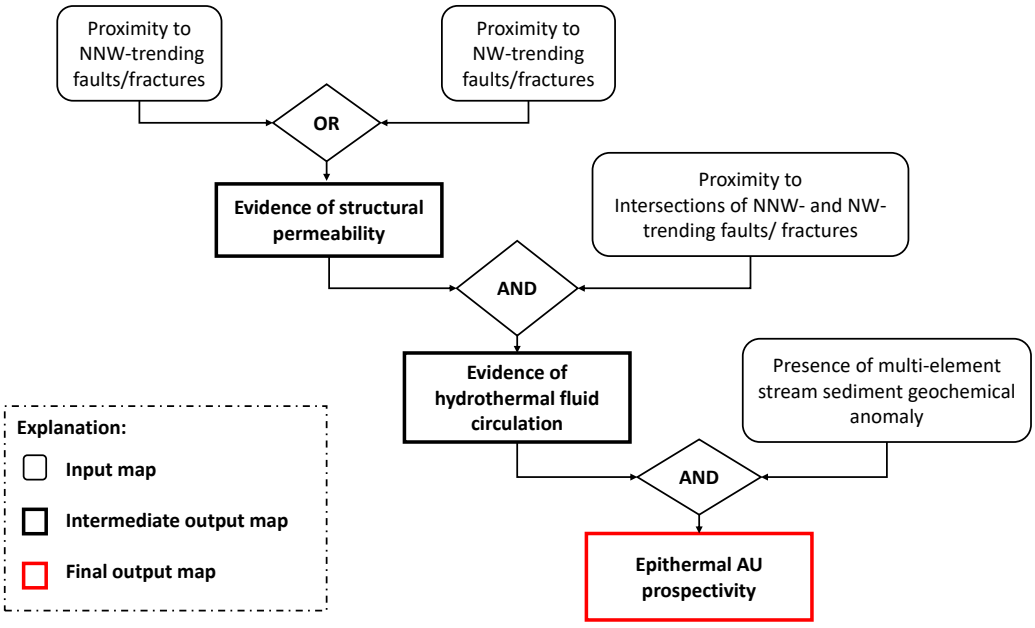


Figure 2.3 The modelling of epithermal Au prospectivity in the Aroroy district of the Philippines uses an inference network to combine input Boolean evidential maps. Modified from Carranza and Hale (2008).

2.2.2 Data-driven models

In contrast to knowledge-driven models, data-driven models are appropriate for brownfield or well-explored, data-rich areas. These models require a training dataset that includes known mineral occurrences such as deposits, prospects, and showings of the deposit type we are interested in exploring. These models are built on the hypothesis that undiscovered deposits will share similar geological features with the known ones used in the training set. Furthermore, the integration function is not defined by an expert but is instead quantified statistically by the algorithm, which learns the complex spatial and genetic relationships between the evidence layers and the known deposits. Significant advancements in MPM have been made during the last few decades thanks to a broad range of data-driven modelling techniques (Porwal et al., 2006a; Leite and de Souza Filho, 2009; Li et al., 2015; Yousefi and Carranza, 2016).

2.2.2.1 Statistical data-driven models

The foundational data-driven methods, which fall under the traditional era in Figure 2.2, are mostly probabilistic statistical models. The most common among them is the weights-of-evidence (WofE) method, which is based on Bayesian probability theory and was adapted for

MPM in the late 1980s (Agterberg, 1989; Bonham-Carter et al., 1989). It statistically assesses the spatial correlation between each evidence layer and the known deposits by computing positive and negative weights. These weights are then summed to create the final prospectivity map (Bonham-Carter, 1994). The strengths of WofE are its interpretability and its capacity to handle missing data and non-linear interactions between variables. While WofE is interpretable and widely applied across MPM applications (Ford et al., 2016; Fu et al., 2021; Zeghouane et al., 2016), its primary weakness is the assumption of conditional independence across all evidence layers, which can be problematic in complex geological systems.

To overcome this, logistic regression models are used, as they do not assume conditional independence (Chung and Agterberg, 1980; Harris et al., 2003; Reddy and Bonham-Carter, 1991). However, logistic regression models struggle to capture the highly non-linear relationships between geoscience features. Therefore, due to these limitations in WofE and logistic regression, researchers have focused on the more powerful non-linear capabilities of ML.

2.2.2.2 Machine learning models

Before diving into the specific ML algorithms, it is worth taking a look at how these methods fit within the broader role of AI in mineral exploration. As shown in Figure 2.4, the model construction follows a range of expert systems, statistical learning, and ML techniques. These models are then applied to key exploration tasks, such as geological mapping, feature extraction, deposit classification, and MPM. Their performance is evaluated using various metrics. Within the broader AI categories, ML methods have become popular because of their ability to learn patterns directly from data rather than relying solely on expert-defined rules.

However, the success rate of these ML models often depends more on the quality of the input representations than on the choice of algorithm. This necessitates feature engineering, which is a method that creates a geoscience feature set that accurately represents the target task of mineral exploration, leveraging both ML and expert knowledge (Ustundag et al., 2022). This is an important step in MPM because the feature engineering serves as the bridge between raw geological data and predictive algorithms. For example, Montsion et al. (2024) demonstrates the efficacy of combining ML-based feature engineering with expert examination to construct feature sets representing various mineral system fertility indicators. This process aligns closely with the human-in-the-loop paradigm described by Caers (2025), where the objective is not to replace the geologist, but to use data science in order to enhance domain expertise. Therefore, feature engineering ensures geological plausibility, reduces data dimensionality while retaining critical signals, and is then applied to construct the MPMs.

Over the last couple of decades, ML models originally developed to address multi-field classifi-

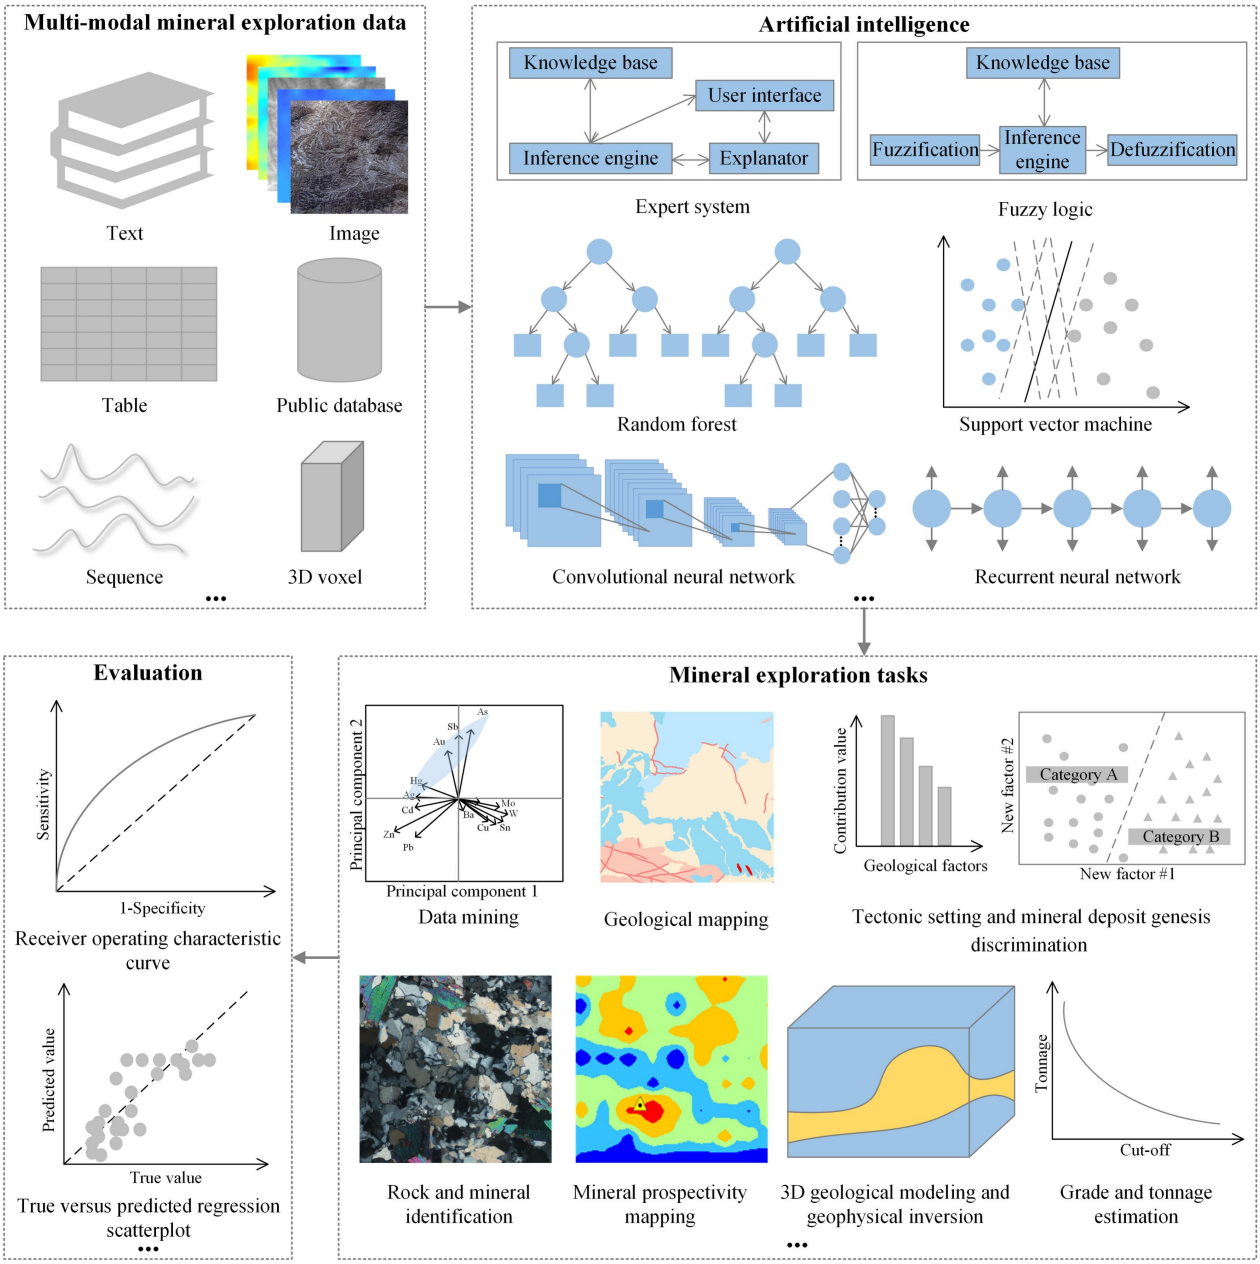


Figure 2.4 Conceptual workflow illustrating the use of AI in mineral exploration. The process usually involves acquiring geoscience datasets, constructing AI models, applying these models to specific exploration tasks such as mapping, classification, or prospectivity analysis, and evaluating their performance (Yang et al., 2024a).

cation and pattern recognition problems have shown promise in generating predictive models of mineral prospectivity (Lary et al., 2016; Li et al., 2020; Saljoughi and Hezarkhani, 2018; Sun et al., 2019; Zuo, 2017).

One of the earliest and most widely adopted ML methods was the artificial neural network

(ANN). An ANN is an adaptive, parallel, computational model with biological inspiration, made up of small processing units called neurons connected by links with weights that represent the strength of each connection (Rosenblatt, 1958). Neural networks are adaptive in that they learn by studying examples, and once the network has been trained, the previously learned information is distributed throughout it as updated weight values. Following key theoretical developments and advances in computing, Jeffrey and Rosner (1986) and Ritter et al. (1988) recognize the possibility of integrating neural networks with GIS technology. Much of the research that follows concentrates on automatic pattern identification and the classification of remotely sensed imagery (Decatur, 1989; Ritter and Hepner, 1990). Many of these experiments demonstrate how neural networks can be used to combine multiple datasets for categorization (Bruzzone et al., 1997).

Multi layer perceptrons (MLPs) or multi-layer feedforward neural networks are widely used in MPM due to their excellent performance across a wide range of problem types and strong function-approximation capabilities. A MLP is typically composed of an input layer that receives the evidence layers, one or more hidden layers, and an output layer that gives the final prospectivity score. Each neuron in the hidden and output layers performs a set of multiplication and addition operations on its inputs, followed by a non-linear transformation (an activation function), as illustrated in Figure 2.5.

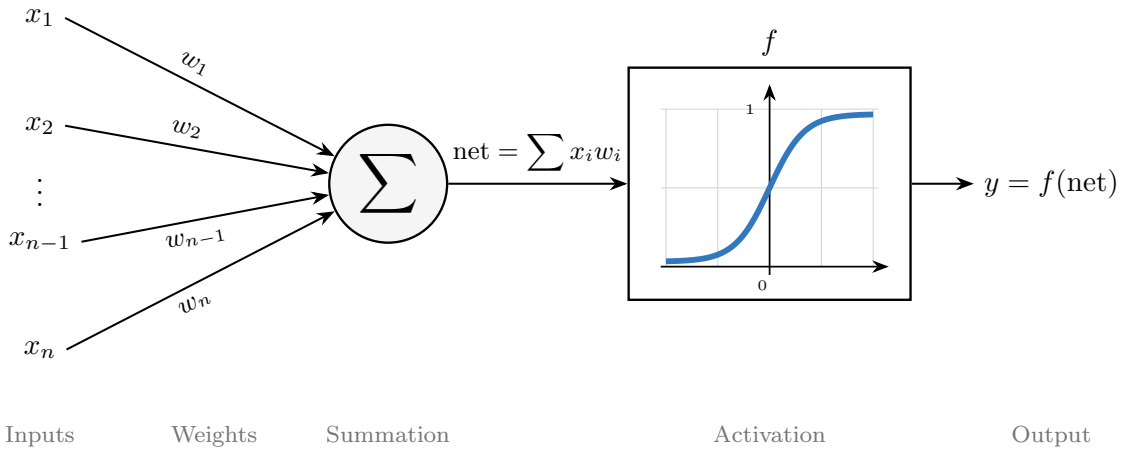


Figure 2.5 Computations performed at a single perceptron unit

When processing a set of training samples repeatedly, MLP neural networks learn by adjusting the connection weights. Because every example includes both the desired outputs and all the input parameters, this training method is known as supervised learning. The network's error is the difference between its actual output and the desired output, and this error is

minimized through the back-propagation algorithm, which iteratively adjusts the connection weights to improve the model's accuracy. Once trained, the network can generalize, or provide appropriate results for an unseen test set.

ANNs have been widely applied in MPM case studies and often demonstrate strong predictive performances (Brown et al., 2000; Rodriguez-Galiano et al., 2015; Sun et al., 2020). The primary strength of ANNs lies in their ability to model complex, non-linear, and interdependent relationships between predictor variables thanks to the universal approximation theorem (Hornik et al., 1989). They can respond to critical parameter combinations rather than simply summing favourable factors, and they can uncover hidden patterns in data that other methods miss. However, their main weakness is their black-box nature: because the learned information is distributed across thousands of weights, it can be very difficult to interpret how the model reaches its decisions, which has been a significant barrier in high-risk applications such as mineral exploration.

Next, support vector machines (SVMs) are another ML method, a kernel-based algorithm widely used for classification and regression problems (Vapnik, 1999). SVMs work by finding an optimal hyperplane that creates the best possible separation between the two classes (e.g., deposits and non-deposits in MPM applications). For non-linear data, SVMs use a kernel function, such as the popular radial basis function, to map the input features into a higher-dimensional space where a linear separation is possible. So far, SVMs have been shown to produce MPM models with high classification accuracy (e.g., Abedi et al., 2012; Ghezelbash et al., 2021; Zuo and Carranza, 2011). Notable advantages of SVMs include their effectiveness in high-dimensional feature spaces, making them particularly suitable for complex datasets. They are also relatively robust to overfitting, especially when the data are high-dimensional, and can be adapted for both linear and non-linear classification tasks through the use of kernel functions. However, SVMs also have certain limitations. Their computational complexity can be problematic when dealing with large datasets, particularly in MPM applications where geoscience data volumes can easily contain millions of points across several dozen features. In addition, the interpretability of SVM models becomes limited when non-linear kernel functions are used, making it challenging to understand the underlying decision-making process.

Another popular ML method applied in MPM is the random forest (RF) algorithm (Carranza and Laborte, 2016; Hariharan et al., 2017; Xiang et al., 2020). RF operates by constructing a large number (a forest) of individual decision trees during training, where each tree is built using a random subset of the data and features. The final prediction for a new location is determined by aggregating the predictions of all trees, typically through a majority vote for

classification tasks or averaging for regression tasks. The main strengths of RFs are their non-parametric nature, robustness to outliers, and ability to handle complex, high-dimensional datasets. They can also provide a measure of variable importance, making them useful for feature selection and model interpretation. Moreover, RFs are relatively easy to implement and generally perform well even with heterogeneous input data. Despite these advantages, RFs have several limitations. While they provide global measures of variable importance, the decision logic across numerous trees is difficult to trace for individual predictions. In addition, training large forests can be computationally expensive and may lead to over-fitting when the model becomes too complex.

While ANNs, SVMs, and RFs are among the most widely adopted ML models in MPM, other algorithms, such as k-nearest neighbours (Ghezelbash et al., 2023; Silva dos Santos et al., 2022) and gradient boosting machines (Li et al., 2024; Roshanravan and Kreuzer, 2025; Zhang et al., 2022a), have also been successfully applied. Overall, these ML algorithms demonstrated substantial improvements over earlier statistical approaches, particularly in addressing non-linear relationships. Yet as geoscience datasets expanded in volume, dimensionality, and complexity, their manual feature engineering requirements and limited representational capacity motivated a shift toward DL, which can automatically learn hierarchical features directly from raw data (LeCun et al., 2015).

2.2.2.3 Deep learning models

DL techniques are a significant subset of ML algorithms defined by their use of multiple processing layers to learn hierarchical representations of data (LeCun et al., 2015; Zhou et al., 2017). Unlike standard ML models, which often require manually engineered features, DL models can identify complex structures and latent connections within specific datasets by learning multiple levels of representation. This capability is critical as the complexity of geospatial data patterns increases with the volume of geoscience data. By leveraging deep neural network architectures, these techniques can capture the intrinsic relationships between ore-indicating characteristics and mineral deposits (Zuo et al., 2021b). Most importantly, DL algorithms have the ability to directly extract features from raw inputs. Although many cutting-edge and reliable algorithms have been presented for MPM, it is important to note that no single model consistently outperforms others in all circumstances (Porwal and Carranza, 2015).

One of the most widely used DL methods in MPM is a convolutional neural network (CNN), an ANN with a specialized architecture that falls under the DL domain and was originally designed for image processing, using local connection patterns in order to preserve spatial

information. CNNs have an input layer, one or more convolutional layers, pooling layers, fully connected layers, and an output layer as their basic architecture. The core of this architecture lies in the convolutional layers, which apply a set of filters (or kernels) to scan the input data and generate feature maps, and the pooling layers, which down-sample these feature maps (Figure 2.6). This process allows the model to automatically capture local spatial correlations, hierarchical features, and complex geo-anomalies that a standard MLP might miss.

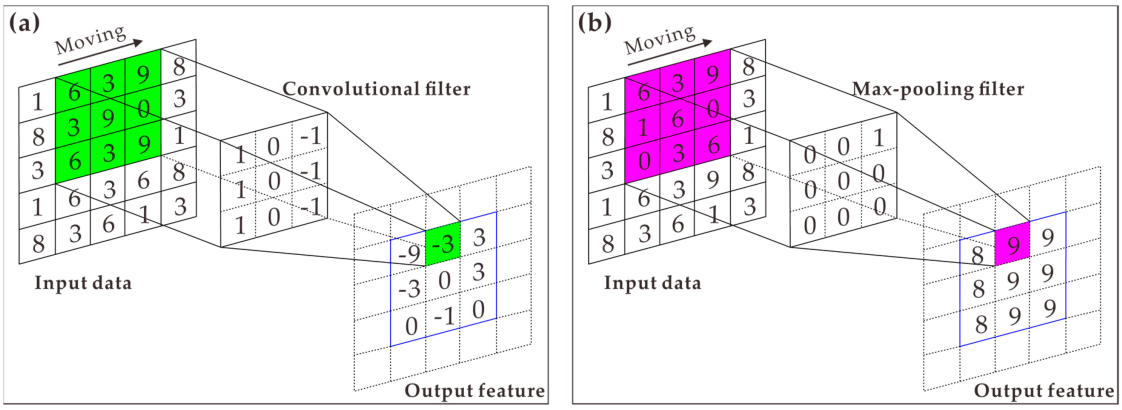


Figure 2.6 Illustrations of (a) convolution and (b) max-pooling operations (Sun et al., 2020).

CNNs offer several benefits for MPM. Their ability to learn hierarchical spatial features allows them to capture complex geo-anomalies related to mineralization. They are also robust to noise because multi-layer filter extraction helps identify stable patterns even in distorted inputs. In addition, CNNs efficiently process image-like geoscientific data via convolutional operations and achieve high predictive accuracy when working with large, mixed datasets, making them a strong option for MPM applications. Overall, CNNs show promising results in producing MPMs (He et al., 2024; Liu et al., 2024; Yang and Zuo, 2024). However, the use of big data for mineral exploration primarily depends on the high dimensionality of evidence layers (beyond spatial dimensions), specifically, a large number of channels. As a result, standard CNNs struggle to extract crucial channel properties associated with mineralization. Another drawback is the complexity of building CNNs for modelling datasets from geological prospecting in MPM. Images typically have three (RGB) or four (CMYK) channels, and standard CNN architectures such as LeNet (LeCun et al., 1998), AlexNet (Krizhevsky et al., 2017), and VGG (Simonyan and Zisserman, 2014) have been developed to handle these image formats. These architectures are designed to extract features from data that are spatially large but shallow in channel depth (typically 224 $\times$ 224 pixels with three channels). In contrast, the underlying geoscience sampling grids used in MPM are often spatially much

smaller; for example, Li et al. (2021b) reports a geochemical grid of only 85 × 76 points for a study area of approximately 25,840 km^2. At the same time, MPM inputs contain a large number of channels, commonly exceeding 40 multi-source evidence layers. This structural mismatch makes it difficult to directly apply classical CNN architectures to MPM.

Next, recurrent neural networks (RNNs) are another type of DL method that is capable of processing sequential or ordered data. In geoscience applications, several types of data exhibit inherent sequential structure, including stratigraphic columns, drill-hole lithogeochemical profiles, and depth-ordered geophysical logs, making them natural candidates for sequence-based modelling. Unlike feed-forward architectures, RNNs incorporate recurrent connections that allow information to persist across time steps, which enables the network to model temporal dependencies and contextual relationships within the input sequence. A standard RNN consists of an input layer, a recurrent hidden layer, and an output layer; however, basic RNNs often struggle to capture long-term dependencies due to vanishing gradients. To address this limitation, advanced recurrent architectures have been developed, most notably long short-term memory (LSTM) networks and gated recurrent units (GRUs). LSTMs introduce memory cells together with three gating mechanisms: the forget gate, the input gate, and the output gate. These components work together to selectively retain or discard information, enabling the network to learn long-range dependencies more effectively. GRUs use a simplified gating structure that lacks a separate memory cell, and they often achieve comparable or superior performance on smaller datasets. The design of these gated architectures makes RNNs well-suited for analyzing sequential geoscientific data. Recent studies have demonstrated the effectiveness of LSTM- and GRU-based models in capturing contextual information in evidence layers, improving the interpretation of geochemical patterns, and producing MPMs with strong predictive performance (Qaderi et al., 2025; Wang and Zuo, 2022; Yin et al., 2022; Zhao et al., 2020).

A significant limitation in training any type of MPM is the data imbalance between known mineral deposits and barren areas, where known mineral deposits are rare, which results in small positive training sets that can lead to large model errors or overfitting. Generative adversarial networks (GANs) are developed to address this specific problem (Goodfellow et al., 2014). A GAN is a deep generative model composed of two competing neural networks, including a generator that learns to create synthetic deposit samples and a discriminator that learns to distinguish the synthetic data from real data. Therefore, in MPM, their primary application is data augmentation, where the GAN is trained to capture the complex patterns of the real deposits and generate new, high-quality, synthetic positive samples to balance the training set. So far, GANs have been applied to address the data imbalance issue in MPM (Kong et al., 2025; Li et al., 2022; Wake et al., 2025). For example, Wu et al. (2023)

applies a conditional generative adversarial network (cGANs) that extends the standard GAN architecture by incorporating class-specific information into both the generator and the discriminator. In this framework, the generator receives not only random noise but also an embedded label vector which enables it to synthesize positive samples that better match the distribution of known mineral deposits. Likewise, the discriminator evaluates each input image together with its corresponding label, improving its ability to distinguish real deposit samples from synthetic ones. A diagram of the cGAN architecture used for data augmentation in MPM is shown in Figure 2.7, where label embeddings are fused with noise inputs to produce class-conditioned images, and the discriminator jointly processes real and generated samples under the same conditioning.

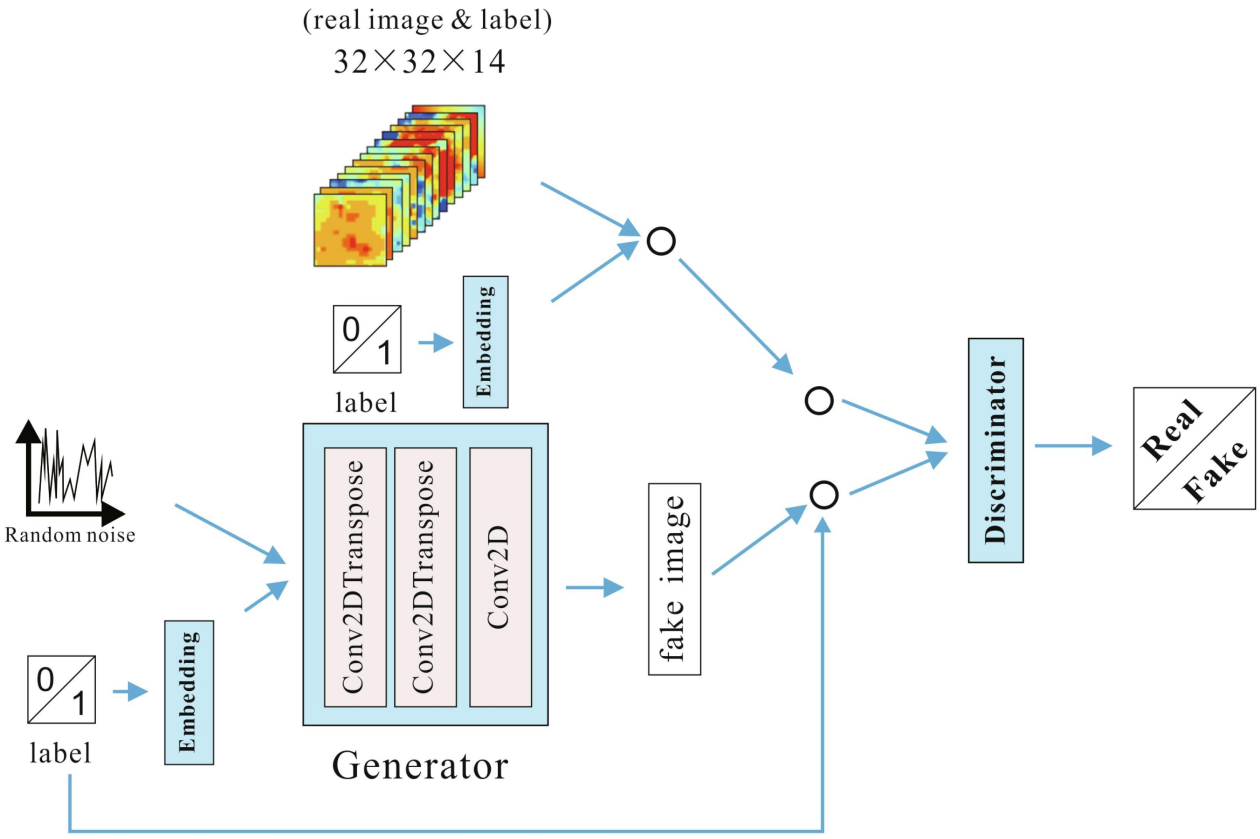


Figure 2.7 Schematic architecture of the cGAN used for data augmentation, showing label embedding, generator structure, and discriminator processing of real and synthetic samples (Wu et al., 2023).

This cGAN architecture enables the generator to learn complex spatial patterns in known deposits and produce synthetic examples that expand the positive training set while maintaining geological plausibility. The primary application of GANs and their variants in MPM is data augmentation to address class imbalance, and their effectiveness and limitations in

this role are discussed in detail in Section 2.3.1.

2.2.3 Hybrid models

Hybrid models are not a single method but are instead a framework that combines a knowledge-driven approach with a data-driven one. This is done to leverage the strengths of both techniques. An example of this category is the hybrid WofE model (Porwal et al., 2006b). First, it uses a knowledge-driven fuzzy-logic approach, where expert-defined fuzzy membership values transform the data into fuzzy predictor patterns. Next, these fuzzy predictor maps are combined using the data-driven, Bayesian-based integration function of the WofE model. The primary strength of this hybrid approach is its flexibility, which uses expert knowledge to better represent the evidence layers while still using a data-driven method for the final integration. However, one weakness of this approach is that the fuzzy WofE model remains subject to the WofE's conditional independence assumption, which can lead to artificially inflated posterior probabilities (Porwal et al., 2006b). Various hybrid methods have been successfully applied across MPM studies (Chung and Fabbri, 1993; Porwal et al., 2004; Roy et al., 2006).

2.2.4 Modelling with empirical functions

To better assign weights to evidence features in data-driven and knowledge-driven methods, researchers proposed modelling with empirical functions, which uses an analyst's judgment to define a specific mathematical equation (Yousefi et al., 2021). These functions assign weights indirectly via this expert-defined function rather than through direct scoring. However, like other knowledge-driven methods, these models are still subject to biases and uncertainties arising from the analyst's arbitrary judgments (Chung and Fabbri, 1993; Luo and Dimitrakopoulos, 2003).

2.2.5 Logistic-based modelling

The methods in this category apply logistic sigmoid (S-shaped) functions to allocate weights to evidence values. This approach is designed to avoid the primary sources of bias found in other methods, as it does not use known deposit sites for training, arbitrary expert judgments, or the categorization (or discretization) of continuous data (Yousefi et al., 2021). Instead of simplifying data into classes, this method assigns continuous weights to the exploration data, which can be applied in either greenfield or brownfield settings. Researchers have applied logistic sigmoid functions (Tsoukalas and Uhrig, 1996) to quantitatively model spatial

associations and assign these continuous weights without relying on training sites or expert opinion (e.g., Almasi et al., 2017; Yousefi and Carranza, 2016; Yousefi and Nykänen, 2016). Studies comparing this continuous weighting method with traditional categorical data have confirmed that it can achieve better predictive performance and reduce bias and uncertainty arising from the simplification of complex geological data.

Overall, the integration function in MPM has evolved from simple expert-driven overlays to complex DL models. This wide array of methodologies, from knowledge-driven to data-driven, each has distinct advantages and limitations, where no single model is known to be superior for every application. The selection of the most suitable method is therefore a critical decision for the geoscientist, depending on the specific context of the exploration problem, the geological setting, the scale of the survey, and, most importantly, the quality and quantity of the available data.

2.3 Key knowledge gaps and challenges in modern mineral prospectivity modelling

While the previous section reviewed the specific methods and their limitations of each individual model, the general MPM workflow also faces several critical challenges. The increasing adoption of geodata science, driven by big data analytics and DL, has not only advanced MPM but also highlighted and mitigated these issues. The following subsections present these key challenges.

2.3.1 Data scarcity and class imbalance

One of the major challenges in MPM is the inherent scarcity of known mineral deposit data. Since mineralization is a rare event (Cheng, 2007), it results in a problem where there are insufficient known mineral deposits to properly train a predictive model. This leads to severe class imbalance, where the majority class (barren or negative samples) vastly outnumbers the minority class (deposit or positive samples). Therefore, ML classifiers tend to be biased toward the majority class, causing poor predictive performance for the target mineralization.

Researchers try to address this issue through two main categories of techniques: algorithm-level and data-level methods (Chawla et al., 2004, 2003). Algorithm-level methods aim to adapt the classifier itself, for example, by using cost-sensitive learning to assign a higher penalty for misclassifying the rare positive events (Zhang et al., 2023; Zhou and Liu, 2006) or by using ensemble classifiers to integrate multiple models for better generalization (Seiffert et al., 2009). On the other hand, data-level methods focus on resampling the training set

to optimize the class distribution. This can be achieved by undersampling, where examples from the majority class are removed, or by oversampling, where minority-class examples are duplicated or synthesized (Zuo, 2020).

However, undersampling is often considered a suboptimal approach, especially in large-scale MPM, because deleting examples from the majority class can lead to the loss of valuable geological information about barren environments. Consequently, researchers generally prefer data augmentation techniques to artificially expand the volume and diversity of the training dataset. Noisy injection and synthetic data generation are the two main data augmentation methods employed in MPM (Wu et al., 2023; Yang et al., 2022a). Noise injection techniques introduce random noise, such as Gaussian or uniform noise, to the input features to improve model robustness. Conversely, synthetic generation aims to construct new, realistic samples that reproduce the statistical properties and spatial trends of the original dataset. Common methods in this synthetic category include the synthetic minority over-sampling technique (SMOTE, Chawla et al., 2002; Farahnakian et al., 2024; Prado et al., 2020; Tang et al., 2025), as well as more advanced generative DL models such as GANs (Farahbakhsh et al., 2023; Wu et al., 2023) and conditional variational autoencoders (CVAEs, Nagasingha et al., 2024).

Among these generative approaches, GANs and their conditional variants (cGANs) have shown promise for synthesizing realistic deposit samples (Farahbakhsh et al., 2023; Wu et al., 2023). However, their application in MPM is constrained by several practical limitations, including training instability, susceptibility to mode collapse where the generator produces limited diversity in synthetic samples, and vanishing gradients that can hinder convergence. Furthermore, GANs lack an explicit inference mechanism, which limits their utility for uncertainty quantification or deeper interpretation of the generated data (Sun et al., 2024). These limitations have motivated the exploration of alternative generative architectures, such as CVAEs, which offer more stable training through an explicit probabilistic framework and provide a mechanism for uncertainty estimation (Nagasingha et al., 2024). Beyond these, researchers have also applied positive-unlabelled (PU) learning (Hajihosseinlou et al., 2025; Xiong and Zuo, 2021; Zhang et al., 2022b) and meta-learning (Soltani and Imamalipour, 2022) in order to address the ambiguity of barren labels and the scarcity of training samples. Despite these advancements, selecting the most appropriate balancing strategy remains a critical step in the MPM workflow, as it directly influences the model's performance. This critical challenge of class imbalance and training label representativeness is addressed in Chapter 4 of this thesis.

2.3.2 Negative training sample selection

Closely linked to the challenge of class imbalance is the difficulty in selecting valid negative training samples. In supervised MPM, the positive class is readily defined by the locations of known mineral occurrences. However, defining the negative class is far more ambiguous, as it is often impossible to confirm with certainty that a location is truly barren without exhaustive drilling or investigation.

Researchers have proposed several approaches to address this problem, including selecting random points within the study area, using locations of different types of mineral deposits, or utilizing drilling sites that have historically returned no mineralization (Nykänen et al., 2015). Ideally, in order to minimize bias, these samples should be carefully selected. For instance, Carranza et al. (2008) proposes that negative samples should be randomly distributed, located at a sufficient distance from known mineralization, and possess data values for all input variables.

A recent study by Parsa and Cumani (2025) formalizes these strategies into distinct categories to systematically evaluate their impact on model performance. They define the random selection of distant sites as the traditional approach, and the selection of sites with different mineralization styles as the geological contrast approach. Crucially, the authors contrast these methods with another dissimilarity-based approach. Unlike traditional methods that rely primarily on geographic distance or geological differences, this approach focuses on the representativeness of class labels within the multi-dimensional feature space. In this method, negative labels are identified by calculating the mathematical dissimilarity, such as using cosine similarity between potential sites and known deposits. Their findings indicate that this dissimilarity-based selection can ensure that negative samples are geospatially distinct in terms of evidential features rather than merely physically distant, thereby significantly enhancing the model's spatial selectivity. This concept is visually illustrated in Figure 2.8, taken from the study which focused on Mississippi Valley-type (MVT) Zn-Pb prospectivity in Canada (Parsa and Cumani, 2025), where the map contrasts the spatial distribution of negative samples selected via the traditional approach, the geological contrast approach and those selected through the dissimilarity-based approach. Overall, results from Parsa and Cumani (2025) demonstrate that effective negative selection and dataset augmentation are governed by the principle of class label representativeness. Specifically, selecting negative labels from geospatially distinct localities is critical to reducing label noise and enhancing the model's spatial selectivity. This underscores the need for practice in assessing the representativeness of all training labels, including both augmented positive and selected negative, before using them in classifiers.

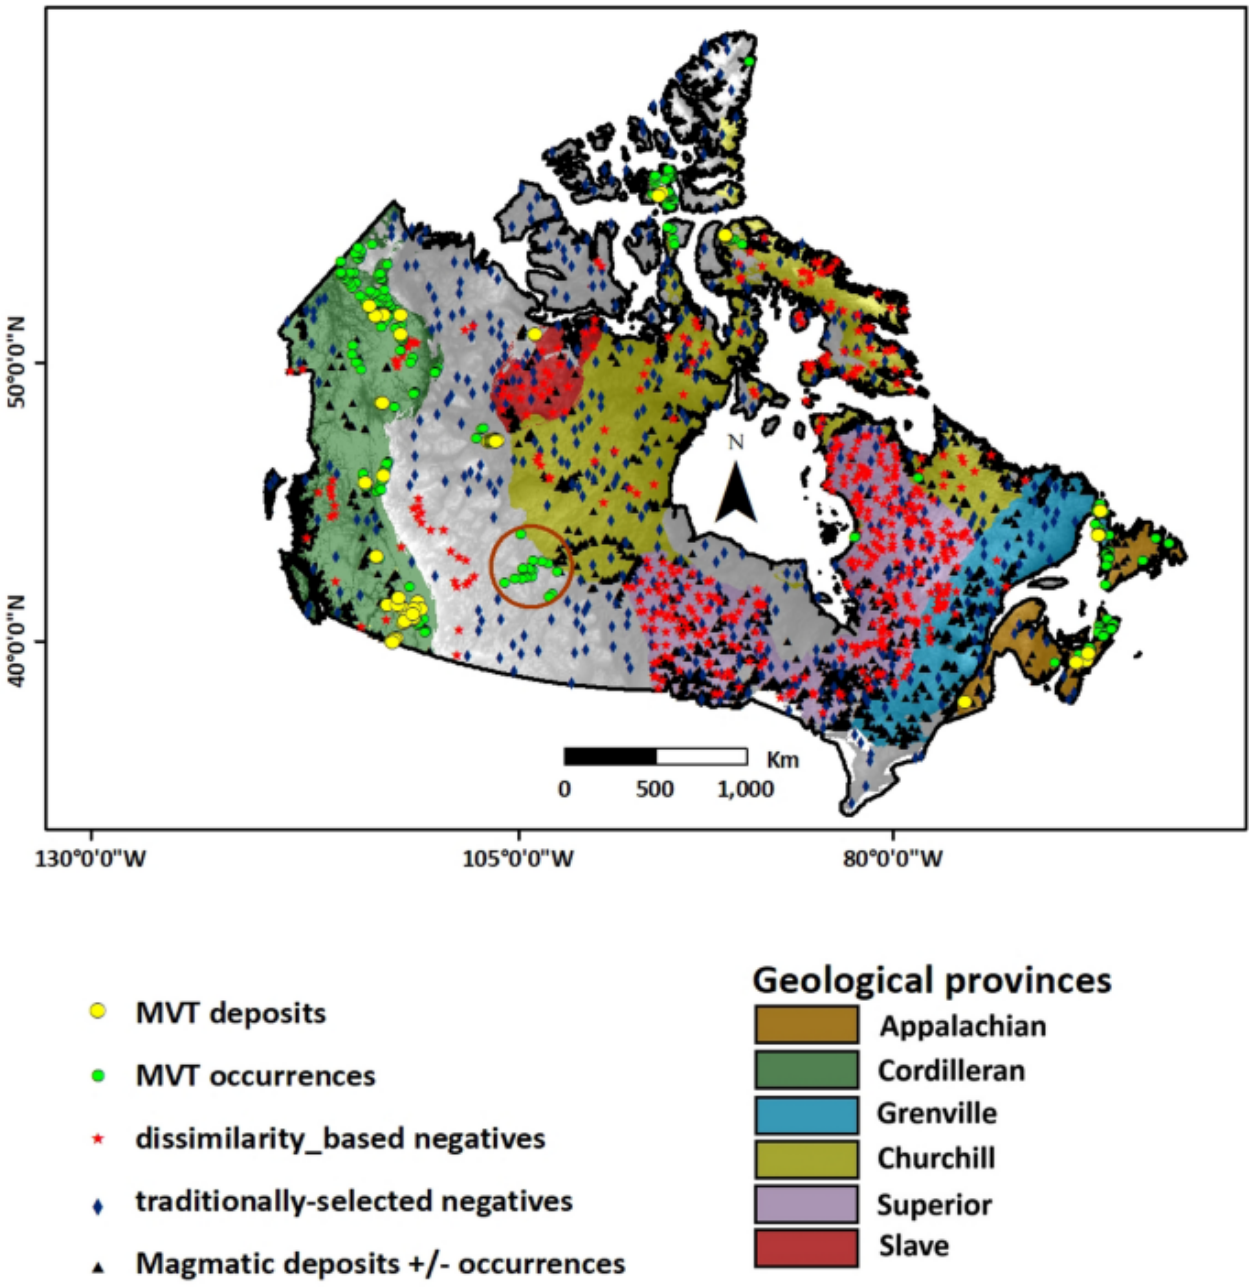


Figure 2.8 Map of Canada's geological provinces illustrating the spatial distribution of various class label groups used for MVT prospectivity modelling. The map highlights the spatial patterns of traditionally selected negatives, magmatic deposits used for geological contrast, and dissimilarity-based negatives, compared to the known MVT deposits and occurrences (Parsa and Cumani, 2025).

2.3.3 Input data quality and deep controls

The effectiveness of any MPM model is undermined by inconsistencies in the input geoscience data, often failing to objectively represent the search space of interest. Geoscientific data can be categorized into two major types. First, uniformly and objectively sampled over the area of interest, which are unbiased data sets (e.g., regional geophysical data and digital terrain models). Second, data are produced by non-uniform, biased sampling, resulting in biased datasets that include information related to known mineral occurrences, most regional geochemical data, and the known distribution of key rock units, referred to as fact geology (Hronsky and Kreuzer, 2019). The latter point is critical, as fact geology is typically only reliably determined in areas of better exposure or those that have already received considerable work due to known mineralization.

Including both biased and unbiased datasets without proper spatial weighting inevitably yields outputs that are biased toward regions of high exploration maturity. This problem often leads to prospectivity maps that simply emphasize known deposits, which is not a practically useful result for finding new large targets in poorly explored areas (Hronsky, 2009). Furthermore, if only surface-mapped geology is used, the prospectivity output is inherently biased away from areas concealed by post-mineral cover where the regions with the greatest residual prospectivity (Hronsky and Kreuzer, 2019). This problem of arbitrary surface data is compounded by the failure to include critical deep-targeting geoscientific elements. The majority of inputs usually used in MPM are derived from or close to Earth's surface. Therefore, for example, it will ignore trans-lithospheric structures and other deep controls that are fundamental to the spatial location of ore deposits (e.g., Billingsley and Locke, 1933; Crafford and Grauch, 2002; McCuaig and Hronsky, 2014; O'Driscoll, 1986). Because these deep structures are often cryptic at the surface and are not systematically compiled in public datasets, their exclusion is a common failure that leads to inaccurate prospectivity mapping.

Next, these challenges are linked to the conceptual framework used to define the model. This often leads to overlooked strong dependency between input layers, a problem that critically biases algorithms that assume independence, such as WofE algorithms. A clear example of this is seen in targeting komatiite-hosted nickel sulphide deposits (Barnes, 2006) where three important elements include major trans-lithospheric faults (pathways for magma), channelized komatiite volcanic facies (host successions), and sulphide-rich metasediments (sulphur source). However, the spatial distribution of the latter two elements is strongly controlled by the first element, as they are both related to anomalous crustal heat/magma flux, which is itself controlled by the trans-lithospheric faults. Therefore, properly recognizing and incor-

porating these intrinsic geological dependencies is critical for producing reliable and unbiased prospectivity models.

In MPM, the underlying conceptual framework for defining input elements is derived from a particular interpretation of the target mineral system, which represents the set of processes that must converge in space and time to form an ore deposit (McCuaig et al., 2010). However, the most common interpretation, the traditional investigative lens defined through the source-transport-trap paradigm, is typically flawed for predictive modelling, because this framework is characterized by intrinsic dependencies among its components, which can compromise predictive performance. This lens focuses primarily on the relationships between multi-scale processes that need to occur to produce a deposit. The greatest value in spatial prediction, however, comes from identifying areas of coincidence among multiple, independent factors. Therefore, a different framework, the predictive lens, focuses on the conjunction of totally independent, contextual factors that need to coincide for ore genesis. These key factors include architecture, which is the structural framework providing the basic plumbing for the mineral system; fertility, which comprises the basic chemical aspects of the system that give it ore-forming potential; transient geodynamics, representing the instant in time when the system actually organises to make ore; and preservation/exhumation, the post-ore processes that allow the ore deposit to be discoverable and exploitable. Qualitative representations of these two perspectives are presented in Figure 2.9. Use of this predictive lens avoids the dependence problems inherent in the investigative lens, making it better suited to support predictive exploration targeting (McCuaig and Hronsky, 2014).

2.3.4 Uncertainties in mineral prospectivity modelling

Effective decision-making in mineral exploration inherently requires robust frameworks for quantifying risk, as prospectivity mapping is characterized by high uncertainty due to geological complexity and the non-unique interpretations of exploration data, particularly geophysical data. While many factors can drive these uncertainties, in the context of data-driven MPM, they are often classified into four main categories. These include process-related uncertainties (a lack of knowledge about ore-forming mechanisms) and judgment-related uncertainties (irrational human decision-making), alongside the two most critical categories for ML models: aleatoric uncertainty (arising from irreducible noise and randomness in the data) and epistemic uncertainty (arising from ambiguity in the model's structure due to limited training data).

However, it is important to note that uncertainty quantification is not just about estimating errors but is essential for making rigorous decisions as well. Caers (2025) critiques that the

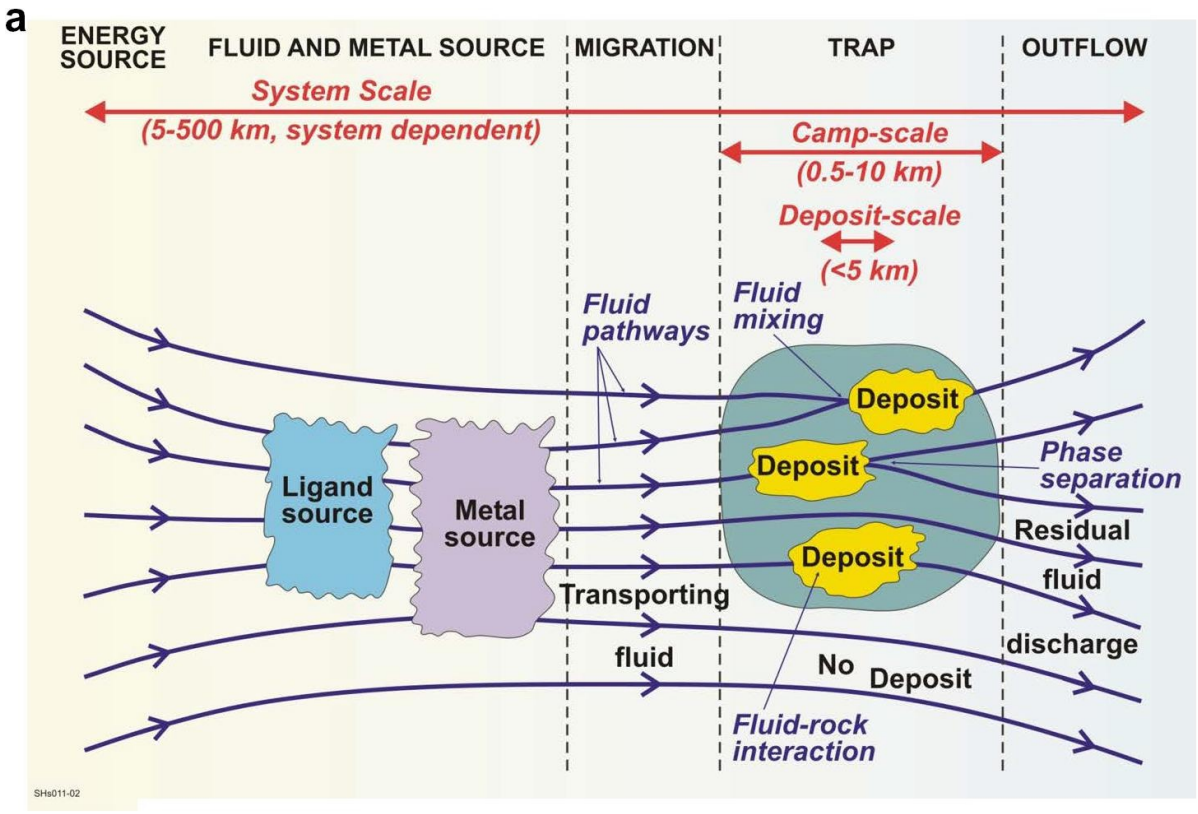


Figure 2.9 Schematic illustrating the difference between two conceptual frameworks for mineral systems. The investigative lens (a) focuses on the causal links between multi-scale component processes (Hronsky and Kreuzer, 2019; Wyborn et al., 1994), whereas the predictive lens (b) highlights the necessary spatial and temporal conjunction of independent elements needed for ore genesis (McCuaig et al., 2010).

industry's current focus on static predictions, which is just generating a map, and argues for a shift toward decision-making frameworks. In this view, the objective of an AI model applied to prospectivity modelling work is not to pinpoint a deposit immediately, but to identify which data acquisition step (e.g., drilling) will most effectively reduce uncertainty about a geological hypothesis. Therefore, it is necessary to move MPM from static classification to sequential decision planning. Despite existing frameworks, no study has simultaneously addressed aleatoric and epistemic uncertainty while providing actionable risk-return categories within such a decision-oriented approach; this gap, addressed through generative modelling for uncertainty quantification and risk-return analysis, is the focus of Chapter 5 of this thesis.

2.3.5 Model interpretability

Another key issue for data-driven MPM, which limits its industrial adoption, is the black-box nature of these models. Data-driven MPM, particularly ML and DL models, have poor interpretability because their internal mechanisms are opaque between the input and output. Despite their powerful capabilities and growing popularity, their black-box nature remains a significant concern for geologists to apply them in real-world mineral exploration projects. To address this, researchers approach interpretability at two stages.

Firstly, interpretability can be embedded during the model design phase by incorporating geological knowledge directly into the architecture or training process of data-driven models (Alvarez Melis and Jaakkola, 2018; Yang et al., 2024a). Several strategies have been proposed to achieve this integration. One method involves combining knowledge-driven and data-driven approaches, where expert geological knowledge is encoded as rules or constraints that guide the data-driven model (Porwal et al., 2004; Tahmasebi and Hezarkhani, 2012). Another approach embeds geological knowledge as mathematical formulations within the model's internal architecture or loss functions (e.g., Rasht-Behesht et al., 2022; Xiong et al., 2022). For example, Zuo et al. (2022) implements a geologically constrained autoencoder for MPM by integrating the non-linear relationship between distance to ore-controlling factors and mineral deposit density as a penalty term in the loss function. This ensures that the learned representations align with known geological principles governing mineral system architecture. Their results demonstrate that geologically constrained models outperform purely data-driven approaches, demonstrating the value of embedding domain knowledge to enhance both interpretability and predictive performance.

Secondly, post-hoc interpretability techniques are applied after model training to explain the decision-making process of trained models. These methods are important for building domain expert confidence in model predictions and for extracting geological insights from

classification models. Attribution algorithms, such as Shapley additive explanations (SHAP, Lundberg and Lee, 2017a), layer-wise relevance propagation (Bach et al., 2015), individual conditional expectation (ICE, Goldstein et al., 2015), and local interpretable model-agnostic explanations (Ribeiro et al., 2016), can identify which input features correspond to specific geological evidence layers that contribute most significantly to the model's prospectivity predictions. With these recent advancements, researchers widely implement these algorithms in MPM applications. For example, Yang et al. (2024b) applies SHAP to analyze feature importance in a DL based prospectivity model for gold deposits, which identifies distance to ore-controlling faults, proximity to favourable lithological units, and hydrothermal alteration indicators as the most influential factors in delineating prospective zones. Similarly, Liu et al. (2025b) employs multiple visualization techniques to interpret CNN decisions in tungsten prospectivity mapping. Firstly, SHAP-based analysis reveals the contribution patterns of key ore-controlling features, such as tungsten geochemical anomalies, the density of fault intersections, and proximity to intrusive rocks, while other decision-path analyses trace how individual predictions accumulate evidence across multiple geological layers. Additionally, Liu et al. (2025b) implements ICE plot analysis, partial dependence plots, and derivative ICE (d-ICE) plots to examine how variations in specific geological features influenced model predictions. Their results demonstrate that tungsten anomalies exhibit a consistently increasing trend with concentration, whereas the density of fault intersections shows distinct patterns across different sample groups.

Overall, these interpretability methods offer distinct strengths. In particular, SHAP values offer several advantages for MPM applications. They enable both global feature importance ranking to identify the most influential ore-controlling factors and local instance-level interpretations to understand individual prospectivity predictions. Furthermore, SHAP offers fast computation for tree-based models, which are commonly used in MPM. In contrast, local interpretable model-agnostic explanations (LIME) provide model-agnostic explanations that can work across different data types, including tabular, text, and image data, and generate contrastive explanations by approximating complex models with simpler, interpretable ones. Each method addresses different aspects of interpretability which is a key factor in assessing the feasibility of implementing data-driven models in MPM. Understanding the relationships between the geological knowledge embedded in model architectures, which evidence layers contribute most significantly to prospectivity predictions to evaluate model reliability, and how different model components process geological information to improve model transparency is crucial for establishing trust between geologists and data-driven models. This subject is analyzed further in Chapter 6 of this thesis.

2.4 Discussion and conclusions of the literature review

The evolution of MPM over the past several decades reflects the ongoing challenge of integrating geoscience knowledge and data with computational methods. As illustrated in this review above, the field has progressed from simple expert-driven overlays to advanced DL architectures, yet each advance has introduced new challenges alongside new capabilities. Therefore, the choice between knowledge-driven, data-driven, and hybrid approaches depends strongly on the specific exploration context, including data availability, geological complexity, and the scale of investigation.

An important observation is that increasing model complexity does not guarantee improved exploration outcomes. DL models, despite their theoretical capacity to learn complex patterns, often underperform compared to methods when training data are limited or when geological knowledge is not properly incorporated. Similarly, purely knowledge-driven approaches, while interpretable and applicable in greenfield settings, struggle to quantify spatial relationships and are vulnerable to expert bias. Given these strengths and weaknesses, the future of MPM lies not in identifying a single best method, but in developing frameworks that can be tailored to each specific exploration problem.

Another key insight is that the data challenges in MPM are interconnected. Class imbalance, negative sample selection, survival bias, and input data quality are not independent problems. Instead, they all stem from the same mismatch between the data availability and the data requirements that need to train robust predictive models. Historical exploration patterns have concentrated known deposits in accessible, well-explored terranes, creating training datasets that misrepresent the true geological search space. Addressing this requires not only methodological approaches like synthetic data generation and dissimilarity-based negative selection but also a conceptual shift in how we define and select training samples to better represent the true exploration search space.

Beyond these data-related challenges, a critical limitation lies in the gap between prediction and decision-making. The vast majority of MPM studies produce prospectivity maps that indicate where deposits are likely to occur but provide no framework for translating those predictions into exploration decisions. In practice, exploration is a sequential process in which geologists typically decide where to acquire additional data, which targets to drill, and how to allocate limited budgets under uncertainty. Therefore, this process requires not only a prediction but also an assessment of confidence, an understanding of which geological factors drive both prospectivity and uncertainty, and a framework for evaluating risk versus return, where current MPM workflows rarely provide these tools.

Additionally, the black-box problem in ML is often characterized as purely a technical challenge of extracting interpretable explanations from complex models. However, interpretability is about bridging the gap between data-driven model predictions and geological understanding. As for geologists, they will not adopt results from data-driven models regardless of their accuracy, unless they understand why the model makes specific predictions and can verify that those predictions align with geological principles. While explainable AI methods have improved model transparency, most interpretability studies focus exclusively on explaining prospectivity scores. Yet the equally important question of which geological factors control the model's uncertainty has received far less attention, despite its direct relevance to exploration decision-making and data acquisition strategies.

The challenges outlined above are significant, but they can be mitigated through methodological development and conceptual refinement. Recent advances in generative modelling, uncertainty quantification methods, and interpretability frameworks can demonstrate promising directions for improving MPM workflows. The following chapters of this thesis contribute to achieving these goals by developing probabilistic ML methods for synthetic data generation, uncertainty quantification, risk-return analysis, and interpretable prediction. These original developments are validated and deployed to model magmatic Ni sulphide mineral systems across the entire Canadian landmass.

CHAPTER 3 PRESENTATION OF THE ARTICLES

3.1 Article 1: A balanced mineral prospectivity model of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems using conditional variational autoencoders

Chapter 4 presents the first article, which develops a generative data augmentation framework based on CVAE to address severe class imbalance in MPM. Unlike traditional balancing approaches such as SMOTE and class weighting, the CVAE learns the underlying probability distribution of geoscience features associated with mineralization and generates geologically realistic synthetic training samples. The framework is evaluated against alternative balancing techniques across multiple classification algorithms.

This article was submitted to *Ore Geology Reviews* in July 2024 and accepted for publication in November 2024. *Ore Geology Reviews* covers ore geology broadly, including mathematical and computational approaches to mineral exploration, making it a suitable venue for data-driven prospectivity modelling research.

The resulting prospectivity model successfully identifies known Ni deposits while highlighting new prospective zones in underexplored regions of Canada. Geological validation confirms that these targets correspond to favourable settings such as mafic-ultramafic intrusions, while geophysical evidence is supported by gravity anomalies. A sensitivity analysis reveals how different geoscience features influence synthetic data generation, and a feature importance analysis identifies the key geological and geophysical drivers of prospectivity. These results provide data-driven insights to inform future exploration and guide target prioritization and resource allocation.

This article addresses specific objective 1, as outlined in Section 1.4.2.

3.2 Article 2: A generative neural network approach to uncertainty and risk-return analysis in mineral prospectivity modelling

Chapter 5 presents the second article, which extends the CVAE framework developed in the first study by addressing a critical limitation of deterministic ML-based prospectivity models: the quantification of both aleatoric and epistemic uncertainty in predictions.

While accurate classification is important, exploration decision-making requires understanding not only where deposits are likely to occur but also how confident the model is in those

predictions. Specifically, false positive predictions can result in millions of dollars wasted on drilling barren targets, while overlooking areas with high prospectivity but high uncertainty may lead to missing significant deposits in geologically complex or underexplored regions.

The CVAE's probabilistic nature enables the generation of diverse synthetic datasets through stochastic sampling. Therefore, by training an ensemble of 100 prospectivity models, each on a unique synthetic dataset, this research quantifies epistemic uncertainty through model disagreement while extracting aleatoric uncertainty from the CVAE's probabilistic outputs. Beyond uncertainty quantification, the study introduces a novel risk-return analysis framework that integrates these uncertainty measures with spatial clustering techniques to categorize exploration targets into four distinct groups based on both prospectivity and confidence. Finally, the practical value of this approach is demonstrated through the finding that high-return zones, comprising approximately 4 % of the total study area, contain 94.7 % of known deposits. This article was submitted to *Ore Geology Reviews* in March 2025 and accepted for publication in August 2025.

This work fulfills specific objective 2 listed in Section 1.4.2.

3.3 Article 3: Tractable approximate Gaussian inference for interpretable mineral prospectivity modelling

Chapter 6 presents the third article, which addresses computational and methodological limitations inherent in gradient-based Bayesian DL approaches. While the second article demonstrated the value of probabilistic ML for uncertainty quantification, most Bayesian methods rely on gradient-based optimization through iterative backpropagation. This iterative process can be computationally expensive, susceptible to convergence issues such as local minima and vanishing gradients, and may require extensive hyperparameter tuning. These limitations can be problematic for MPM projects where training datasets contain thousands to millions of samples and computational efficiency is essential for real-world implementation.

Therefore, to overcome these challenges, this study implements tractable approximate Gaussian inference (TAGI), an analytical Bayesian method that propagates probability distributions through neural networks without gradient-based optimization. Unlike conventional approaches that iteratively adjust parameters through backpropagation, TAGI analytically computes posterior distributions by propagating means and variances layer by layer using Gaussian approximations. By doing this, TAGI offers several advantages, such as faster training convergence and the ability to quantify inherent uncertainty without requiring ensemble models. The performance of TAGI is compared against two other Bayesian approximation

techniques, including Bayesian neural network (BNN) with variational inference (VI) and Monte Carlo (MC) dropout, across multiple evaluation criteria, including predictive accuracy, computational cost, and model robustness. Additionally, beyond comparing predictive performance, this study also develops an interpretability framework using explainable artificial intelligence techniques, specifically SHAP, to analyze not only which features drive prospectivity predictions but also which features control the model's uncertainty. The findings reveal that lithological and geochronological features are the dominant drivers of both prediction and confidence, which provides valuable insights for exploration targeting and data acquisition strategies.

This article was submitted to *Applied Computing and Geosciences* in November 2025 and accepted for publication in March 2026. *Applied Computing and Geosciences* is an open-access journal at the interface between computer sciences and geosciences, covering areas such as artificial intelligence, computational methods, and modelling applied to earth science problems, making it well-suited for this research.

This article addresses specific objective 3 as outlined in Section 1.4.2.

CHAPTER 4 ARTICLE 1: A BALANCED MINERAL PROSPECTIVITY MODEL OF CANADIAN MAGMATIC NI (±CU ±CO ±PGE) SULPHIDE MINERAL SYSTEMS USING CONDITIONAL VARIATIONAL AUTOENCODERS

Ore Geology Reviews
Volume 175, 16 November 2024, 106329

Lahiru M.A. Nagasingha[a], Charles L. Bérubé[a], Christopher J.M. Lawley[b]

[a] Polytechnique Montréal, P.O. Box 6079, Station Centre-Ville, Montréal, H3C 3A7, QC, Canada

[b] Natural Resources Canada, Geological Survey of Canada, 601 Booth Street, Ottawa, K1A 0E8, ON, Canada

Authorship contribution statement

Lahiru M.A. Nagasingha: Writing – original draft, review & editing, Literature research, Experimentation, Validation, Software, Methodology, Formal analysis, Visualization, Data curation, Conceptualization. **Charles L. Bérubé:** Writing – review & editing, Supervision, Validation, Resources, Investigation, Funding acquisition. **Christopher J.M. Lawley:** Writing – review & editing, Data curation.

Abstract

With the increasing demand for raw materials, innovative exploration techniques are needed to discover large mineral deposits that are accessible from the surface. In recent years, various supervised machine learning techniques have proven effective for mineral prospectivity modelling (MPM). However, the successful application of these techniques has been limited due to the scarcity of known mineral deposits compared to barren regions, which leads to a model imbalance favouring the latter. We address the data imbalance challenge in MPM by proposing a novel generative modelling approach using a conditional variational autoencoder (CVAE). We compare the proposed method with two other data balancing techniques, namely the synthetic minority oversampling technique and class weighting. Furthermore, the efficacy of the balancing strategies is evaluated for three MPM classification methods, including extreme gradient boosting machines (XGBM), random forests, and multilayer perceptrons. We implement and test the approaches by modelling the prospectivity of magmatic Ni (±

Cu ± Co ± Platinum group elements) sulphide mineral systems for the Canadian landmass. With an area under the success rate curve of 0.94 for a spatially distinct testing data set, we observe that a combination of the proposed CVAE framework with the XGBM classification model surpasses the other methods. Furthermore, the geographical representation of our XGBM-CVAE model demonstrates a strong association with known Ni mineral occurrences in Canada, along with new prospective regions in underexplored areas.

4.1 Introduction

Economic predictions show that so-called critical minerals, sought after for consumer electronics, cyber-infrastructure, electric automobiles, and renewable energy generation and storage, are especially susceptible to supply issues (Ali et al., 2017; Lawley et al., 2021; Meinert et al., 2016; Schodde, 2017; Skirrow et al., 2013; Watari et al., 2020). Over the years, deposits close to the surface have been discovered, and a significant portion of these ores and metals have been mined and recovered, respectively (Johnson, 2018). The remaining mineral deposits that can be mined economically from the surface are likely to be exceptionally rare (Cheng, 2007, 2012). Therefore, mineral discovery has declined for decades due to the exhaustion of near-surface discoveries and the increasing need to explore deeper deposits (Okada, 2021). This decline is further compounded by the survival bias concept, where areas lacking clear signs of mineralization are often overlooked despite their potential to have deposits (Yousefi et al., 2024b). According to the survival bias concept, areas with weak or no mineralization signals may exhibit special patterns or trends that require further investigation. Therefore, this pressure exacerbates technical and economic risks and highlights the urgent need to adopt new technologies and innovative practices, including null and negative exploration search space tests to ensure that no potential deposits are prematurely dismissed (Yousefi et al., 2024b). Within the context of mineral prospectivity modelling (MPM), geoscientists use complex mathematical functions to delineate potentially mineralized zones in diverse geological settings. MPM is a pivotal tool that detects characteristic anomalies in geoscience data, suggesting the occurrence of mineralized zones by integrating geological, geochemical, geophysical, and remote sensing data.

MPM techniques can be divided into at least five categories (Yousefi et al., 2021). These include (1) data-driven models, which are suitable for well-explored, data-rich areas that use established mineral prospect locations as control points and spatial statistical algorithms to estimate the weight of each data layer, allowing the automatic identification of control points; (2) knowledge-driven models, which generally apply to small-scale areas and do not systematically require a data set of known mineral prospects, occurrences, or deposits as controls since

they largely depend on geologist expertise to assess the significance of evidence maps (Carranza, 2008; Zuo and Carranza, 2011); (3) hybrid models, which combine both data-driven and knowledge-driven methods (Porwal et al., 2006b; Zuo et al., 2009); (4) modelling with empirical functions, which involves empirical functions along with expert judgements to calculate weight (Luo and Dimitrakopoulos, 2003); and (5) logistic-based modelling approaches, which apply suitable logistic functions to allocate weights to evidence values without relying on known deposit sites or expert judgements (Almasi et al., 2017; Yousefi and Carranza, 2016).

Among these, data-driven methods have seen significant advancements in recent years. A data layer represents a set of detailed geoscientific features, such as geophysical anomalies or geological maps. This approach reduces the potential for subjective geologist input or biases, requiring less human engagement in mineral prospecting. Advancements in MPM have been made during the last few decades thanks to a broad range of data-driven modelling techniques. Rapid growth in artificial intelligence and machine learning (ML) techniques has drastically improved the success of mineral resource evaluation and exploration projects, improved mineral processing efficiency, and reduced exploration time and cost (Chen and Xiao, 2023; Shirmard et al., 2022). ML methods, including logistic regression (Chung and Agterberg, 1980), support vector machines (Zuo and Carranza, 2011), neural networks (Leite and de Souza Filho, 2009), Bayesian networks (Porwal et al., 2006b) and random forests (Carranza and Laborte, 2015) have been successfully applied in data-driven MPM. With further advancements in the ML domain, deep learning (DL) is a subset of ML that employs multiple levels of representation by assembling simple, non-linear modules. DL models facilitate the metamorphosis of raw input into a sophisticated and more abstract representation by applying successive interconnected computational layers (LeCun et al., 2015; Xiong et al., 2018). Among DL architectures, convolutional neural networks (Li et al., 2023), recurrent neural networks (Yin et al., 2022), and convolutional autoencoders (Yang et al., 2022b) apply successfully to MPM. The ability of DL algorithms to automatically learn by identifying patterns and improve from prior knowledge without being explicitly programmed has significantly differentiated them from traditional ML methods, whose performance improves with increasing amounts of data and often requires human intervention.

MPM typically employs positive and negative samples to train an ML model to identify new potentially mineralized areas based on geological features associated with the positive samples. Positive samples represent known mineral occurrences, whereas negative samples indicate barren areas. Mineralization is a rare event, and in a certain area of interest, the number of negatively labelled areas is much higher than that of positively labelled areas. It is also impossible to confirm if all negative labels are true negatives. This disparity can

impact the effectiveness of MPM (Xiong and Zuo, 2017; Zuo et al., 2022). The selection of positive and negative label locations is of utmost importance and requires the identification of subtle evidence of a mineralized deposit, as well as undertaking follow-up investigations with a complete understanding of the volumetric dimensions of the target deposit (Rahimi et al., 2021). On the other hand, selecting negatively labelled areas can involve identifying areas with unsuitable geological contexts without significant exploration indicators and spatial distribution patterns of these areas, which considerably influence target generation, and variations in negatively labelled amounts can impact the prediction ability of MPM (Carranza and Hale, 2008; Nykänen et al., 2015; Yousefi et al., 2019). For example, Rahimi et al. (2021) examines how the number of negative labelled data affects the MPM and highlights the need for an optimal ratio of balancing between the negative and positive labels by implementing self-organizing maps to improve the models' prediction ability and reliability. Apart from that, many optimization techniques, including the Genetic algorithm, particle swarm optimization, simulated annealing, and Cuckoo search algorithm, were successfully applied in MPM upon searching the optimum ratio between negative and positive labels to balance the training data (see Shabani et al., 2019; Tsai et al., 2014, for examples).

Previous research on data imbalance and scarcity in MPM focuses on augmenting the original geoscience data using image processing techniques like rotation and flipping. However, these operations are known to alter the geospatial and geological significance of data layers (Li et al., 2021b). Alternatively, sliding window approaches and noise injection methods can augment scarce data; however, a greater sample number can lead to data redundancy and over-fitting in MPM applications (Li et al., 2023). More recent work leverages the unique generative properties of DL algorithms, including generative adversarial networks (GAN, Farahbakhsh et al., 2023; Wu et al., 2023), or ML algorithms such as synthetic minority oversampling technique (SMOTE, Hariharan et al., 2017; Peng et al., 2023; Prado et al., 2020) to augment geoscience data for MPM applications. These generative models aim to generate new samples that closely resemble the original geoscience data or images while preserving intrinsic geological information. Through data augmentation, the GAN-based strategies tackle the issues of data scarcity and the imbalance of labelled data in the geoscience domain. This framework allows geologists to alter the inputs within an automated processing workflow, which captures the complicated nature of complex mineral systems formed by geological and tectonic processes. However, Wu et al. (2023) pointed out the limitations of GAN-based MPM, including the model's architecture, training methods, and the size of the chosen data set, which are factors that might affect the generalization capabilities of predictive models. SMOTE-based MPM can also tackle this vast data imbalance issue by modifying the original training data set to equalize the number of positive and negative

training examples. Nevertheless, the SMOTE algorithm also has several limitations, such as oversampling of uninformative and noisy samples and difficulty in determining the number of nearest neighbours (Hong et al., 2024; Jiang et al., 2021). Furthermore, other methods for dealing with limited training data in MPM include using cosine similarities or clustering to find nearest neighbours in feature space or the geospatial domain (Lawley et al., 2023). However, this method can lack effectiveness in high-dimensional spaces when it is difficult to differentiate between various vectors due to the increased distance between points.

This study proposes an alternate approach to address MPM's ubiquitous imbalance of positive and negative labelled data. We develop a conditional DL model using an autoencoder framework and leverage its generative modelling properties to augment geoscience data characteristic of the positive label class. This data augmentation strategy requires only a small amount of realistically available labelled data. We demonstrate the proposed method's efficacy through MPM of magmatic Ni ($\pm$ Cu $\pm$ Co $\pm$ Platinum group elements (PGE)) sulphide mineral systems covering the entire Canadian territory. The contributions of this study are two-fold: (1) we develop a novel geoscience data augmentation method to balance the positive and negative labels in MPM data with a conditional variational autoencoder (CVAE), and (2) we use the balanced data set to generate improved magmatic Ni prospectivity maps for Canada and compare the results with geological and geophysical characteristics. So far, CVAEs have also been successfully applied to solve geophysical inversion problems (Cheng and Jiang, 2020; McAliley and Li, 2024) and other data imbalance issues in various real-world applications, such as detecting rare medical disease diagnoses (Abdellatif et al., 2024; Si et al., 2024), human action recognition tasks (Gong et al., 2020) and short-term building energy prediction applications (Fan et al., 2022). However, to the best of our knowledge, no other examples of CVAE are specifically applied to balance the positive and negative labels in MPM.

The following section presents the Canadian geological context and magmatic Ni ($\pm$ Cu $\pm$ Co $\pm$ PGE) sulphide mineral systems. Then, we describe the data sets used for MPM, and following this, we detail the methods used to implement the CVAE and MPM models. The results section presents the new and improved Ni mineral prospectivity maps, validates their accuracy using quantitative and qualitative approaches and analyzes the CVAE model. Finally, we discuss the advantages and disadvantages of CVAE for data augmentation compared to other generative models in MPM.

4.2 Canadian geological context and magmatic Ni (± Cu ± Co ± PGE) sulphide mineral systems

4.2.1 Geological provinces in Canada

Canada has high geological potential for critical minerals and comprises seventeen geological provinces that are divided into four primary groups: the continental shelf, platform, accretionary orogen, and shield (Wheeler, 1996). Figure 4.1 illustrates the spatial distribution of geological provinces in Canada. Composed of deformed sediments and volcanic rocks of Mesozoic and Cenozoic age, the Atlantic, Pacific, and Arctic continental shelves are the youngest provinces. The St. Lawrence, Interior, Arctic, and Hudson Platforms comprise Phanerozoic strata. Next, the Appalachian, Cordilleran, and Innuitian orogens comprise deformed sedimentary and volcanic rocks generated during distinct Phanerozoic orogenies. Finally, the Precambrian Canadian shield comprises Proterozoic orogens (Grenville, Churchill, Southern, and Bear provinces) and Archean cratons (Superior, Slave, and Nain Provinces).

4.2.2 Magmatic Ni (± Cu ± Co ± PGE) sulphide mineral systems

Wyborn et al. (1994) define a mineral system as all the geological factors that influence mineral deposit formation and preservation. Advances in mineral deposit metallogeny have resulted in unifying mineral systems models that divide deposit types into genetically similar groups. One strategy is to categorize deposits according to geological processes and mineralizing fluid types (see Herrington, 2011, for examples). Therefore, seven main mineral system groups can be defined: ortho-magmatic, magmatic-related hydrothermal, deformation and metamorphism, basin-related fluid movement, regional metasomatic, sedimentary, weathering and regolith (Hofstra et al., 2021; Lewis and Downes, 2008). This approach has the benefit of being process-oriented, which means that a region must have the following evidence of all four essential components for mineral deposits to be prospective: (1) sources which associated with ore components, transporting fluids, and energy to drive the system, (2) pathways where the metals take from their source to sink, often linked to fluid flow conduits, (3) traps, which are either physical or chemical mechanisms that cause the ore components to deposit at the sink and (4) preservation processes that allow mineralization to be held in the crust until today (Wyborn et al., 1994).

The distribution of giant magmatic Ni (± Cu ± Co ± PGE) deposits is non-uniform and limited to a few countries, including Australia, Canada, China, Russia, South Africa, and the United States. These deposits can be categorized into three types: (1) Ni-Cu-PGE, which contains Ni-Cu-Co-PGE; (2) Ni-Cu-(PGE), which has less PGE compared to Ni-Cu-Co; and

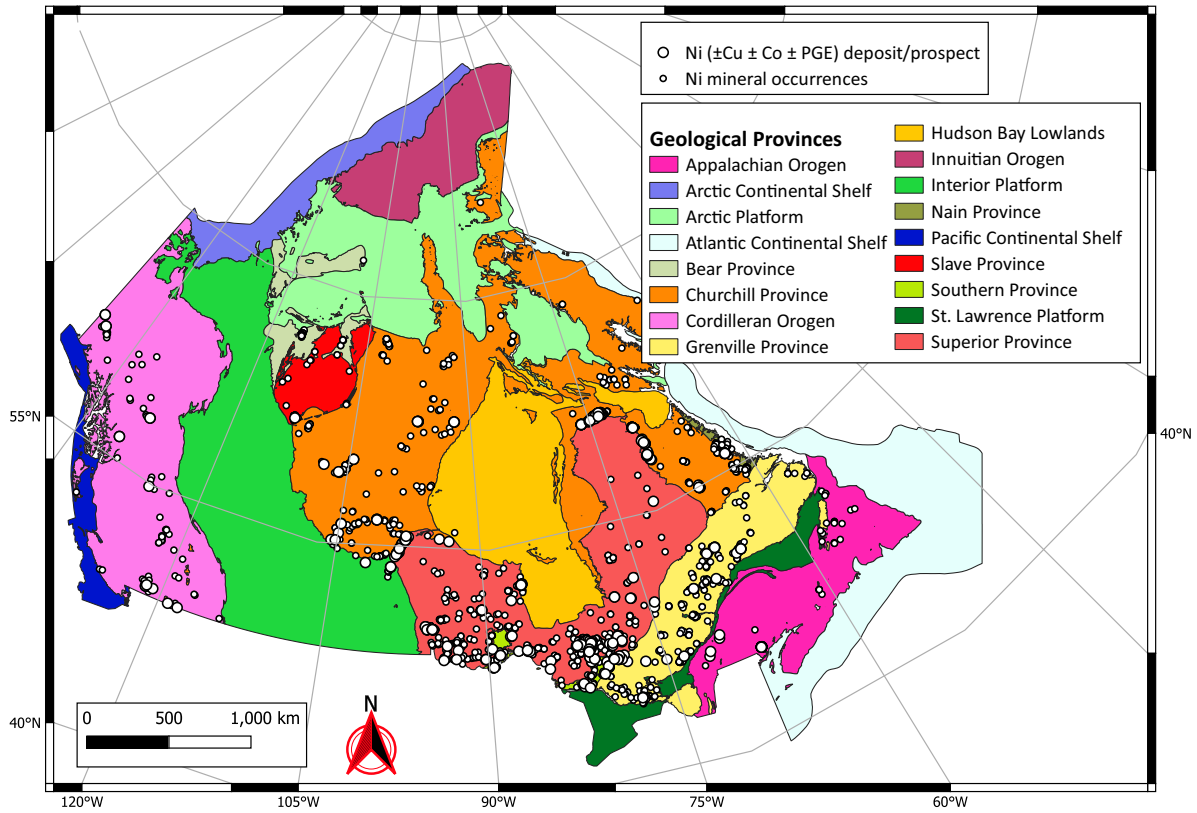


Figure 4.1 Map of the spatial distribution of geological provinces in Canada (Wheeler, 1996). Ni mineral deposits include active mining operations, past-producing mines, and advanced-stage exploration projects (Good et al., 2015; Lawley et al., 2022b), and Ni mineral occurrences include prospects, showings, and anomalous drill intercepts obtained through provincial and territorial databases.

(3) PGE-(Cu)-(Ni), which is PGE-enriched in comparison to Ni-Cu-Co (Lu et al., 2019). Ni is primarily extracted from magmatic Ni-Cu-PGE sulphide deposits (Cross and Peterson, 2023), which originates through the segregation and concentration of immiscible sulphide liquid within mafic or ultramafic magmas. The most abundant and most prospective magmatic Ni-Cu-PGE mineral systems are associated with these mafic to ultramafic igneous rocks in rift-related tectonic settings, primarily from Archean and Paleoproterozoic ages (Eckstrand and Hulbert, 2007). The formation of magmatic Ni-Cu-PGE sulphide deposits depends on three crucial factors: (1) extensive partial melting in the source region; (2) the segregation of sulphide, possibly triggered by the assimilation of sedimentary rocks; and (3) the mechanical and geochemical process for concentrating ore (Naldrett, 2011; Song et al., 2011). However, one of Canada's largest Ni deposits, the Sudbury Basin in Ontario, is a notable exception where sulphide ores formed through a unique combination of crustal melting and meteorite

impact processes, with no evidence of similar deposits worldwide (Maier et al., 2001).

The mineral system approach described above studies the magmatic Ni-Cu-PGE sulphide deposits from lithospheric-scale processes, breaking them into two sub-types based on their principal metal production, regardless of tectonic context: deposits with high sulphide content (10–90 vol. %) and those with low sulphide content (<0.1–5 vol. %) (see Maier, 2005; Nixon et al., 2020; Schulz et al., 2010, for examples). Sulphide-rich Ni-Cu-PGE deposits are the predominant subtype of mineral systems found in Canada, including notable examples like Raglan, Thompson, and Voisey's Bay, which are observable in Archean to Proterozoic greenstone belts, rift-related settings, and more rarely, magmatic arcs. Alternatively, sulphide-poor and PGE-rich deposits, located in layered intrusions and extensive igneous provinces, are associated with magma conduits and extensional settings, highlighting the high prospectivity of continental breakup tectonism (Barnes et al., 2016; Good et al., 2015; Lesher, 2019). Convergent boundaries and ophiolite complexes in the Churchill, Grenville, Cordilleran, and Appalachian geological provinces are the primary hosts for the small yet significant magmatic Ni-Cu-PGE sulphide mineral system sub-types in Canada (Lawley et al., 2021; Nixon et al., 2020; Sappin et al., 2011).

4.3 Data sets

The host rocks of the various Ni deposit types have textural, mineralogical, geochemical, and physical properties that contrast with the neighbouring barren rock units in some cases (Elitok and Tavlan, 2019; Zuccarelli et al., 2022). Because contrasts in rock physical properties are the sources of gravity, magnetic, radiometric, and electromagnetic signals that may be observable with geophysical exploration methods, MPM often relies on integrating geophysical and geological data as the most effective input data sources. We start with the MPM data set initially compiled and described by Lawley et al. (2022b). First, all vector and raster geoscience data sets were converted to the EPSG-4326 coordinate reference system and then indexed using the H3 discrete global grid system (DGGS, Uber Technologies, Inc., 2023). The H3 DGGS is a geospatial indexing approach that employs hexagonal tiling to cover the globe at multiple spatial resolutions. Each hexagon is assigned a unique address, enabling precise spatial indexing, aggregation, and disaggregation across resolutions, assuring consistency and accuracy in data processing at various scales. In this study, the data were spatially indexed based on resolution 7 of the DGGS with 1.22 km per hexagon side, resulting in about 1.8 million unique H3 addresses for the Canadian landmass. Most importantly, the higher spatial resolution of provincial and territorial geological data included in the data set relative to Lawley et al. (2022b) is critical for mapping small ultramafic to mafic intrusions, which

serve as the primary host rock for magmatic Ni-Cu-PGE sulphide deposits (Barnes et al., 2016). Mineral deposits and occurrences are used as true positives during model training, and all other H3 cells were initially assumed to be barren for the purposes of this study. All the geological and geophysical dataset descriptions and sources used in this study are included in the Appendix A, and qualitative representations of the seismic, magnetic, and gravity dataset anomalies are illustrated in Figure 4.2.

4.3.1 Geological data sets

The geological data sets of Canada are comprised of previously published provincial, territorial, and smaller-scale geological maps (Wheeler, 1996). Additionally, Lawley et al. (2022a) describes the unstructured rock descriptions in geological map data sets, which map up to 17 geological properties for sediment-hosted mineral systems. This is crucial for prospectivity modelling, as it helps identify rare prospective rock types not represented by the map unit, ensuring accurate mapping of mineral systems. The mineral system framework uses maximum geological ages of surface bedrock as proxies for mineral deposit sources and transport mechanisms, while minimum geological ages reflect the criticality for knowing whether the igneous or sedimentary rocks are coeval with major mantle melting events (Lawley et al., 2022a). Additionally, faults (Chorlton, 2007), terrane boundary (Eglington et al., 2013), passive margins (Bradley, 2008) and carbonaceous sedimentary rocks such as black shale's (Granitto et al., 2017) data sets are also used as geological data inputs.

Ni deposits are often hosted by minor intrusions that are less than a few km and absent from regional geological maps. In these circumstances, mafic dykes and intrusions with the predominant map unit are the only evidence of mantle-derived melts. The mantle sources of magmatic mineral systems are mapped using both majority and minority rock types, as well as unstructured map unit descriptions.

4.3.2 Geophysical data sets

We select the geophysical data sets based on their completeness, applicability to mineral systems, and capability to provide insights into features ranging from hundreds to thousands of metres beneath the Earth's surface (Lawley et al., 2022a). Firstly, seismic data sets from active and passive surveys reveal the spatial distribution of old, cold, and melt-depleted cratonic lithospheres due to their high seismic velocities. In contrast, slow seismic velocities are found in lithospheric regions modified by younger melts within hot orogens (Debayle et al., 2020; Schetselaar and Snyder, 2017). Therefore, seismic data are used to map abrupt lithospheric and crustal thickness changes, i.e., depth to the lithosphere-asthenosphere boundary

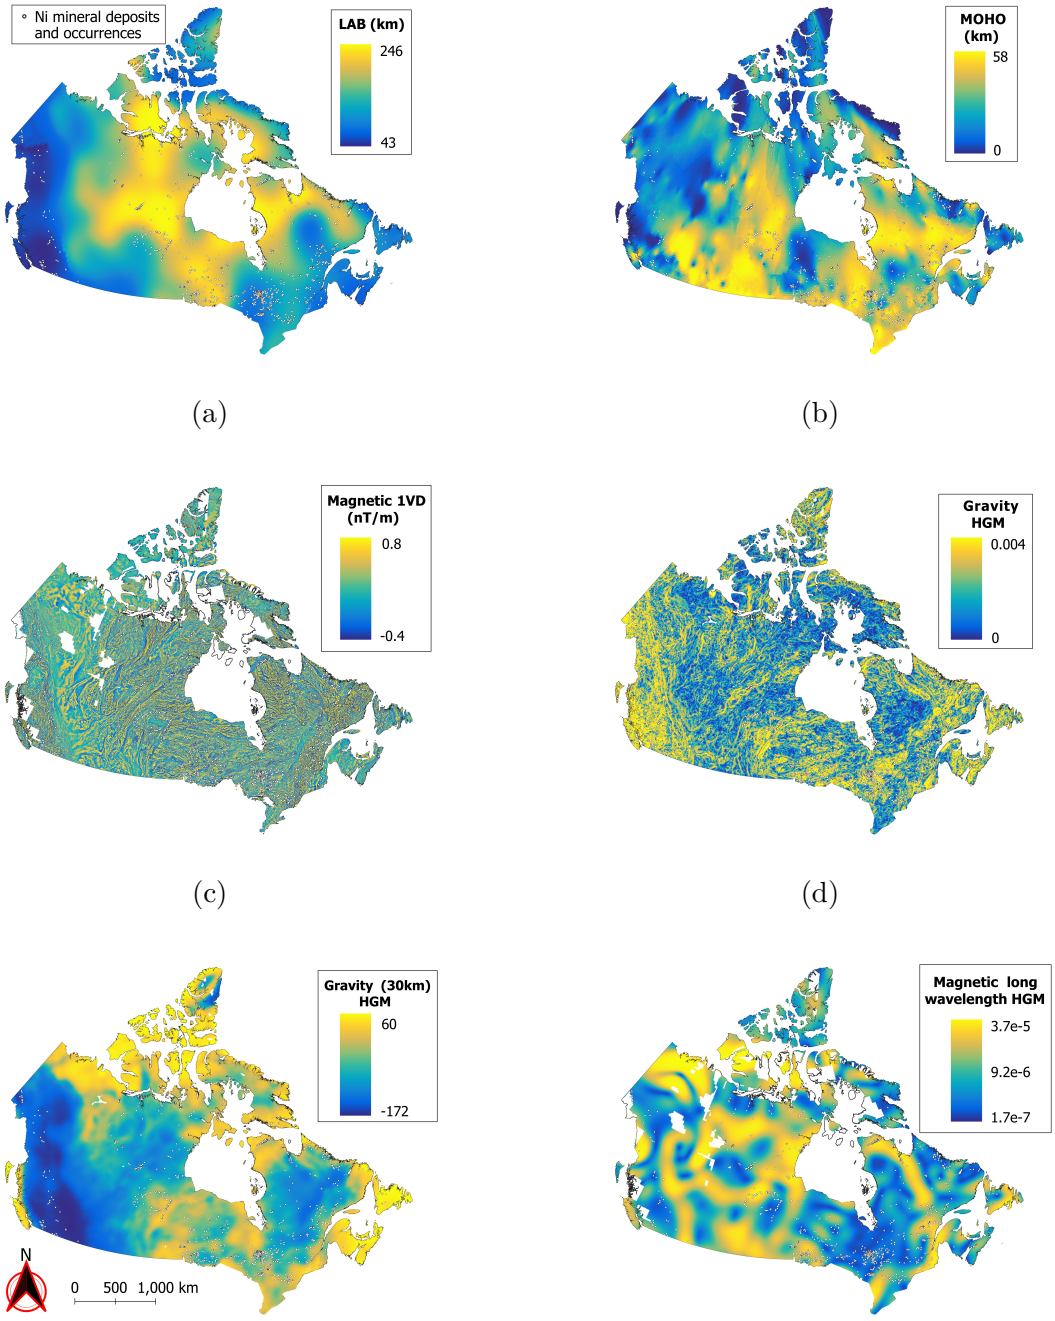


Figure 4.2 Qualitative representation of the feature anomalies (a) depth to the lithosphere-asthenosphere boundary (LAB) and (b) depth to the Mohorovičić discontinuity (MOHO) demonstrates the cratonic edges and crustal thickness respectively; (c) magnetic first vertical derivative (1VD) map demonstrates the magnetic contrasts in shallow crust; (d) gravity horizontal gradient magnitude (HGM) highlights the edges of near-surface contacts or structures; (e) upward –continued (30 km) gravity HGM maps and (f) magnetic long –wavelength HGM emphasize the variations of deep density and magnetic susceptibility within the upper and lower crust respectively.

(LAB) and depth to the Mohorovičić discontinuity (MOHO), respectively. Next, gravity data sets are collected from various surveys, including satellite, airborne, and ground-based surveys (Jobin et al., 2017). Satellite gravity data sets and derivative products (Ebbing et al., 2018) are used to map paleo-tectonic settings with reprocessed gravity grids, i.e., Bouguer and horizontal gradient magnitude (HGM) of the Bouguer gravity, upward-continued (30 km) versions, and maximum values of gravity HGM, highlighting density sources' edges at shallow to deeper crustal depths. Finally, the residual magnetic anomaly data are reduced to the magnetic north pole, and the first vertical derivative (1VD) is used to evaluate magnetic contrasts in the shallow crust while long-wavelength HGM of the magnetic data demonstrate the variations of deep magnetic susceptibility within the upper and lower crust.

4.3.3 Training and testing data sets

We divide the geoscience data compilation into training, validation, and testing subsets using the spatial cross-validation approach implemented in the *blockCV* R package (see Lawley et al., 2022b). This step is essential because the alternative approach of randomly splitting strongly clustered data into training and testing sets is known to overestimate model performance (Valavi et al., 2018). Figure 4.3 illustrates the spatial cross-validation approach.

Each block (500 x 250 km) in Figure 4.3 is sized using semivariogram analysis for a subset of geophysical data sets, including seismic, gravity, and magnetic data. There are six folds in total, including five cross-validation folds for the training set and the remaining fold for the testing set, which we randomly divide based on the limited number of known mineral deposits available for training. The data set contains 1,820,815 individual data points, divided into a training set of 1,491,227 data points and a testing set of 329,588 data points. In total, the data set contains 1452 Ni deposits and occurrences, 1056 of which are in the training set and 396 in the testing set. The total data set is filled with 29 different geological and geophysical data features.

4.4 Methods

This study uses techniques and processes, including data preparation, data augmentation, classification modelling, model evaluation, and interpretation. Figure 6.2 depicts a flowchart of the proposed framework to summarize these techniques, and the following sections describe the primary processes in more detail.

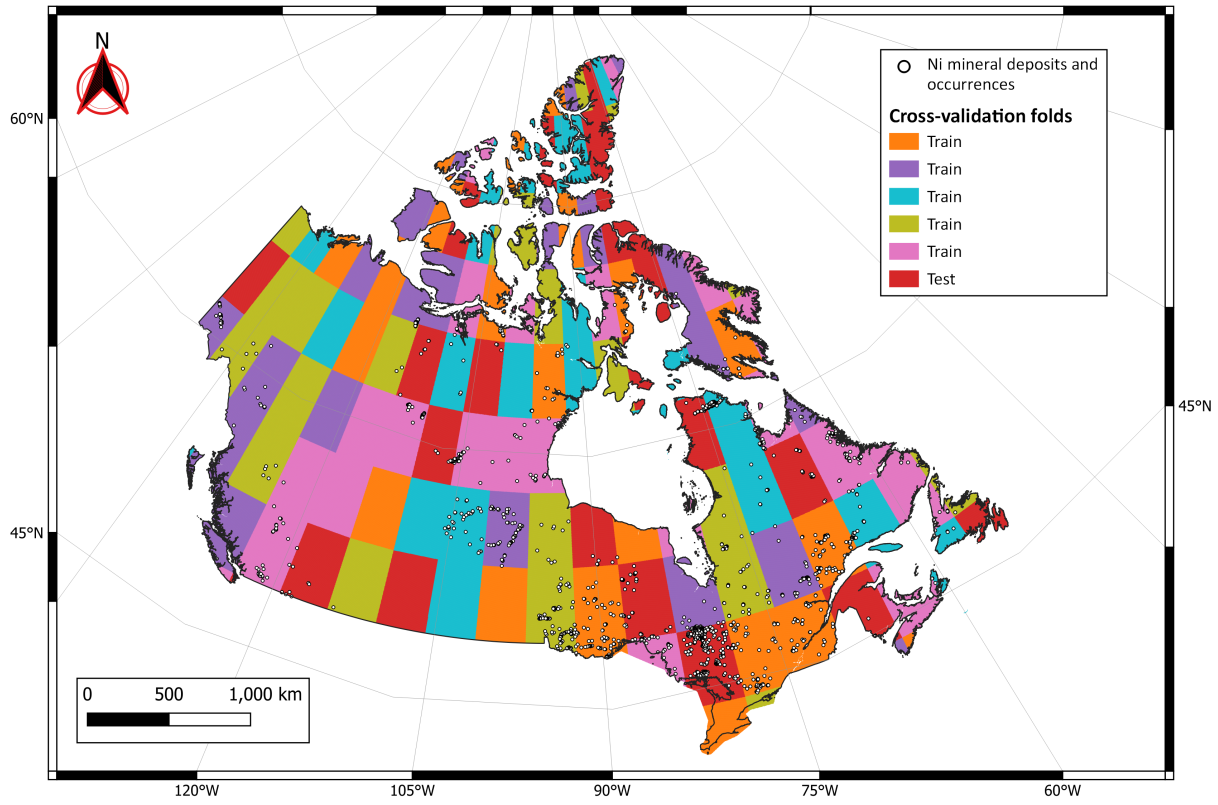


Figure 4.3 Spatial distribution of training and testing data set areas divided based on cross-validated technique. Based on the known Ni mineral deposits and occurrences, five cross-validation folds are obtained for the training data set and one cross-validated fold for the testing data set.

4.4.1 Data preprocessing

Firstly, we used a nonparametric outlier detection method to address the outliers upon calculating the interquartile range (IQR) by creating a fence boundary at a distance of 1.5 IQR, which is beyond the 1st and 3rd quartiles, and subsequently replacing them with the 95th percentile value. Specifically, geophysical data including magnetic HGM, gravity Bouguer HGM, magnetic 1VD, gravity HGM worms, and Magnetic HGM Worms demonstrate the highest outlier percentages of 4.51 %, 5.42 %, 12.79 %, 2.34 %, and 11.38 %, respectively. Secondly, we use the mean feature values to impute 6.5 % of missing magnetic data and 0.2 % of missing gravity data. Since geological data sets, including rock types and period data sets, are categorical, they are mapped into a numeric array using the *OneHotEncoder* class in the *scikit-learn* Python library. After encoding the categorical data, the total data

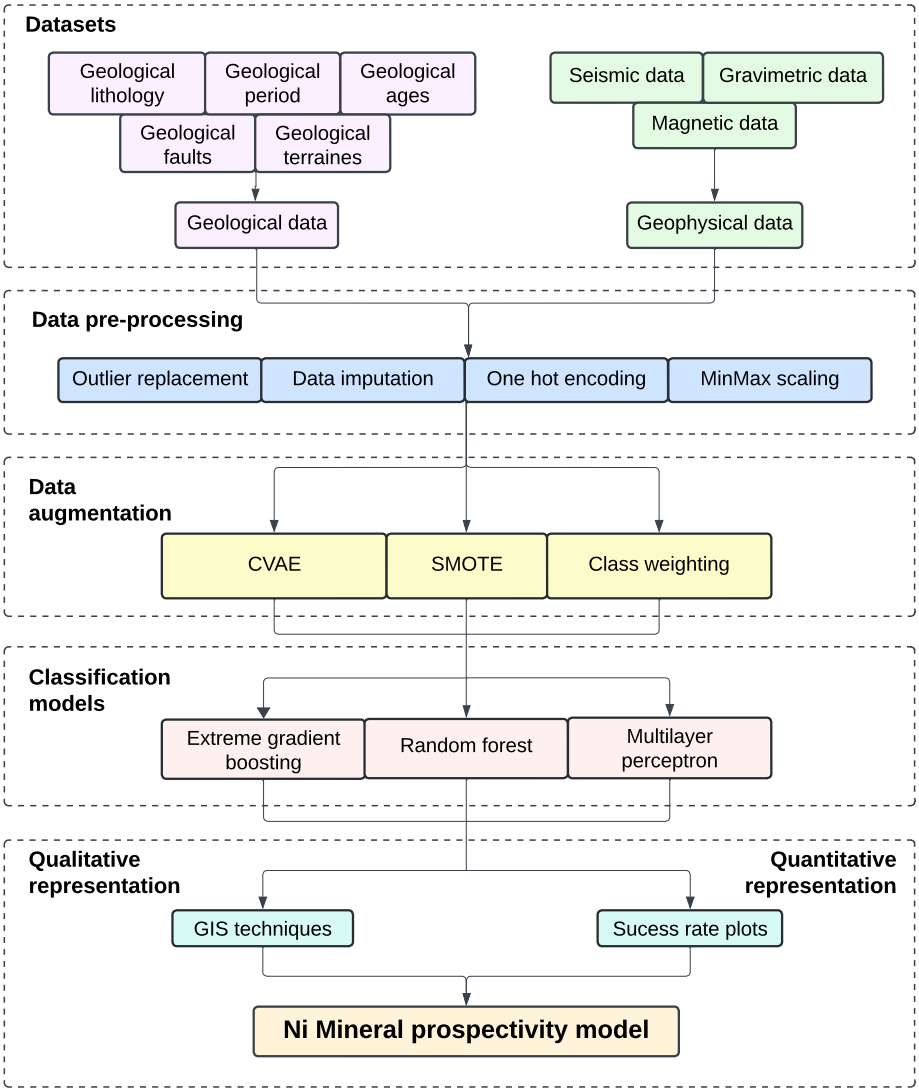


Figure 4.4 Flowchart of the proposed framework of this study.

set contains 173 features. As for the final preprocessing step, the training and testing data sets are scaled to a fixed range of [0, 1].

4.4.2 Data augmentation methods

4.4.2.1 Conditional variational autoencoder

For decades, the concept of an autoencoder has been a major component of developing neural networks (Bourlard and Kamp, 1988). The primary application of autoencoders is to achieve nonlinear dimensionality reduction of large data sets. However, thanks to recent advancements in latent variable models, they have become prominent in generative modelling

applications (Kramer, 1991; Makhzani et al., 2015). Autoencoders have two primary components: the encoder and the decoder. The encoder maps input data ($\mathbf{x}$) into a lower-dimension encoding space by yielding a latent code ($\mathbf{z}$). Then, the decoder component attempts to recreate an approximation of the input data ($\mathbf{x}'$) using $\mathbf{z}$. Autoencoder optimization typically aims to minimise the reconstruction error between the input and output data (respectively $\mathbf{x}$ and $\mathbf{x}'$) by using the backpropagation algorithm with a variant of stochastic gradient descent.

Variational autoencoders (VAE) are generative models that use deep neural networks to predict variational distribution parameters. Its advanced representation capabilities have greatly improved upon those of traditional autoencoder models (Kingma and Welling, 2013). The training objective for the VAE is achieved through a variational Bayesian approach because the true posterior distribution $p_\theta(\mathbf{z}\,|\,\mathbf{x})$ of the latent variable is assumed to be intractable because of the lack of a closed-form marginalization in the latent space. This can be denoted with Bayes' theorem by

$$p_\theta(\mathbf{z}\,|\,\mathbf{x}) = \frac{p_\theta(\mathbf{x}\,|\,\mathbf{z})\,p_\theta(\mathbf{z})}{p_\theta(\mathbf{x})}, \tag{4.1}$$

with the marginal likelihood $p_\theta(\mathbf{x})$ reading

$$p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{x}\,|\,\mathbf{z})\,p_\theta(\mathbf{z})\,\mathrm{d}\mathbf{z}, \tag{4.2}$$

and where $p_\theta(\mathbf{x}\,|\,\mathbf{z})$ is the data likelihood and $p_\theta(\mathbf{z})$ is the prior latent distribution. The VAE model must implement a tractable distribution, which is known as the variational inference model $q_\phi(\mathbf{z}\,|\,\mathbf{x})$, to approximate $p_\theta(\mathbf{z}\,|\,\mathbf{x})$. The approximate distribution is parameterized by an encoder multilayer perceptron with internal weights ϕ. Additionally, $p_\theta(\mathbf{x}\,|\,\mathbf{z})$ is a probabilistic decoder multilayer perceptron parameterized by weights θ. For data augmentation purposes, our study aims to map the input data into the parameters of its underlying generative distribution. Therefore, the input data $\mathbf{x}$ is transformed into vectors $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$, which correspond to the mean and covariance parameters of a Gaussian density function, respectively. The reparametrization trick

$$\mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}, \tag{4.3}$$

where $\odot$ denotes the element-wise product, is then used to explicitly learn $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ in the multivariate Gaussian situation, leaving the stochastic properties of an auxiliary random vector $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$, where $\mathbf{I}$ is the identity matrix, intact. Finally, the decoder reconstructs the latent vector $\mathbf{z}$ to $\mathbf{x}'$, resembling the standard autoencoder. Put simply, VAEs aim to balance the trade-off between data reconstruction loss, or likelihood, and regularization measured by

Kullback-Leibler divergence (D_{KL}) to learn significant latent features and generate novel data effectively. The reader is referred to Bérubé and Bérubé (2022) for a more detailed example of VAE implementation for generative geophysical data modelling.

Conditional variational autoencoders (CVAE) enhance the basic VAE architecture by conditioning the generative model with auxiliary variables. Figure 4.5 illustrates the proposed CVAE network architecture for geoscience data augmentation. Likewise, to the inference model, CVAEs also employ an encoder network that is augmented by the auxiliary variables, $q_\phi(\mathbf{z} \,|\, \mathbf{x}, \mathbf{c})$ where $\mathbf{c}$ are the conditions (see Bérubé and Baron, 2023, for a detailed example of VAE conditioning with petrophysical data). Our work uses the *PyTorch* DL framework to implement the CVAE model with $\mathbf{x}$ containing the pre-processed input geoscience data and $\mathbf{c}$ containing the location of barren and known Ni mineral occurrence data. Algorithm 1 outlines the training process of the proposed CVAE network. The primary objective of training the CVAE is to achieve optimal conditional evidence lower bound (ELBO), which involves maximizing the likelihood of observed data while considering the latent variables within the model (Sohn et al., 2015). The ELBO is

$$\mathrm{ELBO}(\mathbf{x}, \mathbf{z}, \mathbf{c}) = \mathbb{E}\left[\log p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c})\right] - D_{\mathrm{KL}}\left(q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c}) \,\|\, p_\theta(\mathbf{z}, \mathbf{c})\right), \tag{4.4}$$

where the first term is the expectation of the log-likelihood of observing the geoscience data, and the second term is the Kullback-Leibler divergence between the encoder distribution and latent prior. We use the negative ELBO as the CVAE loss function ($\mathcal{L}_{\mathrm{CVAE}}$) to be minimised, which, after simplification reads

$$\mathcal{L}_{\mathrm{CVAE}} \simeq \underbrace{-\frac{1}{2}\sum_{k=1}^{K}(1 + \log\sigma_k^2 - \mu_k^2 - \sigma_k^2)}_{\text{Regularization } (D_{\mathrm{KL}})} + \underbrace{\frac{1}{J}\sum_{j=1}^{J} \parallel x_j - x_j' \parallel_2^2}_{\text{Data reconstruction loss}}, \tag{4.5}$$

where $K = 16$ is the dimensionality of the latent space and $J = 173$ is the dimensionality of the data. The Adam (Kingma and Ba, 2014) extension to the stochastic gradient descent algorithm is employed to optimize the loss function. Furthermore, we use three hidden layers with 128 nodes in the encoder and decoder, a batch size of 32, and an initial learning rate of 10^{-3}. In both the encoder and decoder networks, all fully connected hidden layers employ the sigmoid linear unit (SiLU) activation function defined as

$$\mathrm{SiLU}(\cdot) = \frac{(\cdot)}{1 + e^{-(\cdot)}}. \tag{4.6}$$

Notably, the output layers corresponding to μ, σ, and $\mathbf{x}'$ use a Sigmoid activation function,

constraining their values to the range between [0, 1]. After training, we use the CVAE to generate 1,489,120 positive class data and concatenate it into the existing training data to balance the positive and negative classes.

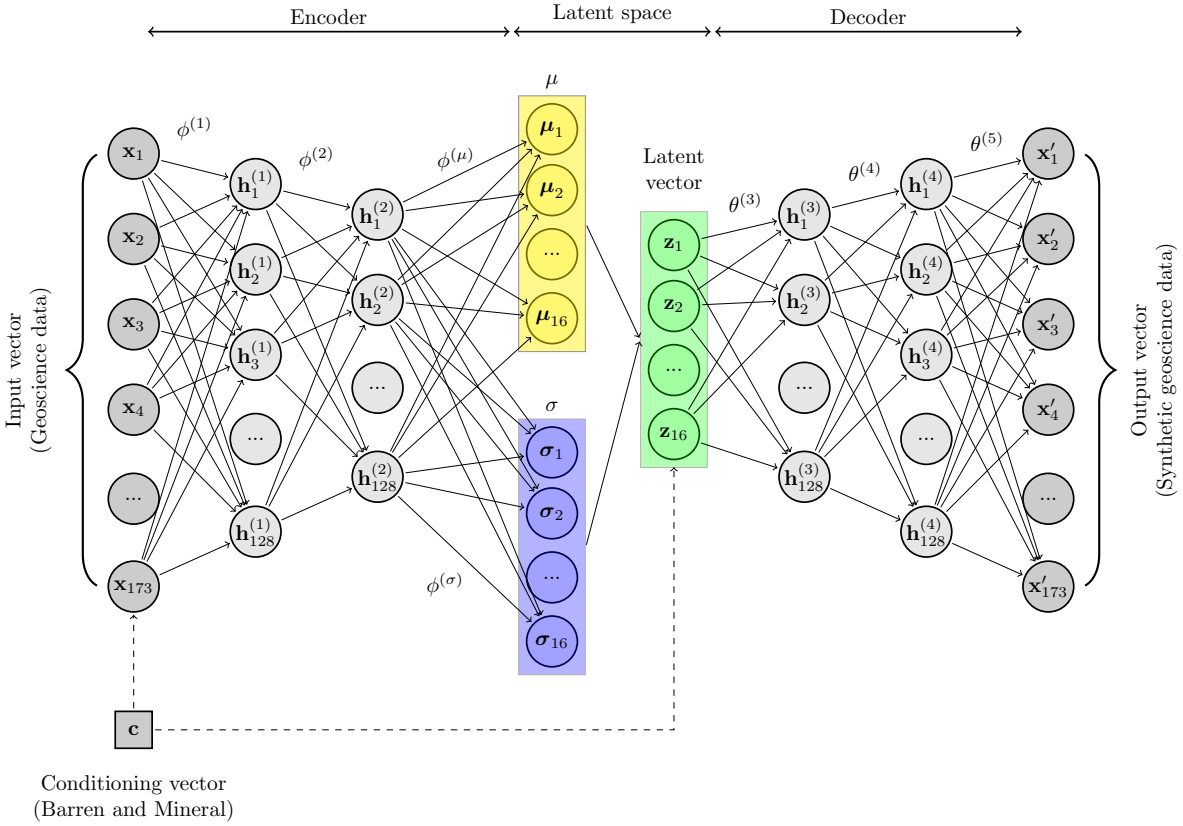


Figure 4.5 The proposed CVAE framework. In the encoder, hidden layers $\mathbf{h}^{(l)}$ are interconnected by linear transformation parameters $\phi^{(l)}$, whereas in the decoder, they are connected by $\theta^{(l)}$. The dashed lines represent vector concatenation.

4.4.2.2 SMOTE

SMOTE generates new instances for the minority class by considering the feature space formed by each instance and its nearest neighbours instead of duplicating data (Chawla et al., 2002). Therefore, the synthetic geoscience data is obtained by

$$\mathbf{x'} = \mathbf{x}_i + (\mathbf{x}_{\text{nn}} - \mathbf{x}_i) \times t, \tag{4.7}$$

where, $\mathbf{x'}$ is the newly synthesized mineral data, which is the minority class, $\mathbf{x}_i$ is the original mineral data that is used to synthesize new data, $\mathbf{x}_{\text{nn}}$ is the data point that is nearest to the data point that needs to be synthesized, and t is a uniform random number between 0 and 1.

Algorithm 1: CVAE model training algorithm.

Data: Geoscience data $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(N)}$ with conditions $\mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(N)}$ ▷ N = Data set length

Result: Probabilistic encoder q_ϕ, probabilistic decoder p_θ

$\phi, \theta \leftarrow$ Initialize parameters

repeat

  for $i = 1$ *to* N **do**

    Draw samples from $\mathbf{z}^{(i)} \sim q_\phi(\mathbf{z} \,\|\, \mathbf{x}^{(i)}, \mathbf{c}^{(i)})$

    $\boldsymbol{\mu}^{(i)}, \boldsymbol{\sigma}^{(i)} = q_\phi(\mathbf{x}^{(i)}, \mathbf{c}^{(i)})$

    $\mathbf{z}^{(i)} = \boldsymbol{\mu}^{(i)} + \boldsymbol{\sigma}^{(i)} \odot \boldsymbol{\epsilon}^{(i)}$ where $\boldsymbol{\epsilon}^{(i)} \sim \mathcal{N}(0, \mathbf{I})$ ▷ $\odot$ = Element-wise product

    $\mathcal{L}^{(i)}_{\text{CVAE}} = D_{\text{KL}}(q_\phi(\mathbf{z} \,\|\, \mathbf{x}^{(i)}, \mathbf{c}^{(i)}) \,\|\, p_\theta(\mathbf{z})) - \log p_\theta(\mathbf{x}^{(i)} \,\|\, \mathbf{z}^{(i)}, \mathbf{c}^{(i)})$

    $\mathcal{L}_{\text{CVAE}} \simeq -\frac{1}{2}\sum_{k=1}^{K}(1 + \log \sigma_k^2 - \mu_k^2 - \sigma_k^2) \; + \frac{1}{J}\sum_{j=1}^{J} \| \, x_j - x'_j \, \|_2^2$

    ▷ J = Data dimensionality

    ▷ K = Latent space dimensionality

    $\phi, \theta \leftarrow$ Update parameters with gradients of $\mathcal{L}_{\text{CVAE}}$ using backward propagation and Adam

  end

until Convergence of parameters ϕ, θ

The following steps are required to generate synthetic geoscience data samples using SMOTE. Initially, a feature vector and nearest neighbours are chosen. Next, the difference between the feature vector and its neighbours is computed. After that, this difference is multiplied by a random number between 0 and 1. The resulting value is added to the original feature vector, creating a new point along the line segment. These steps are repeated iteratively to identify additional feature vectors, thereby enhancing the diversity of the data set.

4.4.2.3 Class weighting

Positive weighting overcomes class imbalance problems by assigning a higher weight to the minority class throughout the model training phase. We have calculated the positive weight by the ratio of the sum of negative instances, which are barren areas, to the sum of positive instances, which are known mineral occurrences. In the context of MPM, this added weight emphasises the importance of the mineral occurrences in the trained model's decision-making process. Consequently, the model's predictions should become more balanced, with less bias towards the dominant class (i.e., barren occurrences).

4.4.3 MPM classification methods

We have implemented three classification algorithms, including extreme gradient boosting machines (XGBM), random forest (RF), and a multilayer perceptron (MLP), to evaluate the performance of each data augmentation method. Firstly, we generate balanced training data sets using each data augmentation method. Next, each classifier's hyperparameters are tuned by considering the balanced or unbalanced data, depending on the experiment. The tuning of hyperparameters is crucial in enhancing the model accuracy, computation speed, and ability to generalize from the geoscience data. Therefore, to make sure that the performance differences are not solely due to the hyperparameter tuning process favouring one balancing approach over another, we have conducted cross-validation grid search experiments for each balanced data set to find the hyperparameters that optimize each MPM classifier's performance.

4.4.3.1 Extreme gradient boosting machines

XGBM is a boosting ensemble ML algorithm that uses multiple decision tree iterations to integrate the predictive capacity of multiple learners into a single model. We have implemented the Python-based XGBM package (Chen and Guestrin, 2016). It employs decision trees as base learners, and the computed error is used to correct the previous predictor. In contrast, the objective function is used to assess the change in model performance. This objective function incorporates additional regularization to avoid over-fitting in ensemble models, thus penalizing complexity and improving generalization. Unlike the simple gradient boosting algorithm, which minimises the loss function by iteratively adding new models that predict the errors of the previous models, XGBM employs first- and second-order partial derivative approximations to acquire information on the direction of gradients, allowing it to minimise the loss function more efficiently. In this study, the XGBM model is optimized using grid search and trained using the following hyperparameters: $n\_estimators = 700$, $max\_depth = 6$, $min\_child\_weight = 1$, $learning\_rate = 0.3$, $subsample = 1$, $gamma = 0.9$ and $colsample\_bytree = 0.6$. All other hyperparameters are left to their default values, as the grid search did not indicate a significant improvement from adjusting them.

4.4.3.2 Random forest

In contrast with XGBM, RF is a bagging ensemble ML algorithm that comprises many decision trees that each reflect a set of hierarchical criteria applied from a root node to its child nodes. RF integrates each decision tree's predictions to predict a target variable with more

accuracy and stability while avoiding over-fitting (Breiman, 2001). This employs a bagging procedure, where approximately two-thirds of the training data, known as bagged samples, is randomly sampled for prediction, and the remaining one-third is utilized for validation. Additionally, random feature selection is conducted at each node in a decision tree to improve the model's generalization and prediction accuracy. We have used the *RandomForestClassifier* implementation of RF from the *scikit-learn* library. In this study, the best RF model is optimized using grid search with the following parameters: $n\_estimators = 450$, $max\_depth = 16$, $min\_sample\_split = 2$, and $max\_features = 0.2$. Any parameters not explicitly specified retain their default values, as the grid search did not yield significant enhancements in model performance upon their modification.

4.4.3.3 Multilayer perceptron

A MLP is a neural network comprising three significant layers: input, hidden, and output. The input layer receives the geoscience data, and the hidden layer transforms the input geoscience data into a form that the output layer can use by applying weights and passing it through a $LeakyReLU$ activation function with a negative slope value of 10^{-2}. This process captures non-linear relationships between input and output data and finally classified as barren or mineral. Building upon this, we have successfully developed a custom MLP using the *PyTorch* library, fine-tuning it with a trial-and-error approach to a hidden layer configuration of 64, 16, 4, a learning rate of 10^{-3}, and a batch size of 32. We employ the binary cross-entropy loss function and the Adam optimizer to train the MLP on the balanced and unbalanced datasets.

4.4.4 Evaluation metrics

Choosing the optimal exploration criteria is a crucial concern in MPM due to the ability to locate highly favourable regions with little ambiguity. Therefore, this section describes the success rate curve we used to evaluate the MPM classification results. To depict the success-rate curve, the x-axis is defined as the proportion of the study area and the y-axis displays mineral deposits that are classified accurately (Agterberg and Bonham-Carter, 2005). Therefore, the success rate can be interpreted by examining how many deposits were correctly predicted within a given area of the model space. A success-rate curve above the diagonal gauge line indicates good spatial selectivity (i.e., reducing the search space), while one under the gauge line indicates a poor relationship with known deposits.

4.4.5 Geoscience feature analysis

Geoscience feature analysis aims to understand and enhance the prospectivity model by quantifying the effect of changes in input geoscience data on the model's generated output. We divide this analysis into two categories: (1) we calculate sensitivity indices (SI) for the CVAE-generated data, and (2) we evaluate feature importance (FI) for the classification models.

4.4.5.1 Sensitivity indices

We compute the Jacobian of the CVAE to estimate the relative SI involved in the generative modelling process (Bérubé and Baron, 2023; Molamohammadi et al., 2020). Therefore, the SI are

$$\mathrm{SI} = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{\partial \mathbf{x}_i'}{\partial \mathbf{c}_i} \right), \tag{4.8}$$

where n is the number of data generated. In other words, we obtain the SI of the decoder network by computing the partial derivatives of the generated data with respect to the mineralized and barren conditions and averaging the derivatives across all augmented data.

4.4.5.2 Feature importance

When using decision tree-based algorithms, the FI calculates a score for all the input geoscience data to demonstrate the importance of each feature in the decision-making process. This is accomplished by summing up the total gains from all splits where the geoscience data has been used and averaging across all of them. The gain measures the relative improvement in accuracy caused by each geoscience data in the splits. In this work, we use the *feature_importances_* attribute in the *scikit-learn* Python library. The implementation of FI for tree-based models is straightforward. However, finding FI for neural networks can be challenging due to the complexity and non-linear behaviour (de Sá, 2019).

4.5 Results

In this study, we compare three data augmentation methods to address the data imbalance issue in MPM. Therefore, the MPM results obtained with the CVAE framework are compared with two other types of data augmentation techniques, namely SMOTE, class weighting and the base model, which uses the original unbalanced data set. Furthermore, we implement three different MPM classifiers for each data augmenting method to validate the efficacy of the CVAE with various MPM approaches.

4.5.1 Success rate curves

We start by evaluating the success rate curves of each MPM classifier combined with each data augmentation method. Figure 4.6 illustrates the success rate curves computed for the testing data set, and Table 4.1 presents the AUC values corresponding to each combination of data augmentation and classification method.

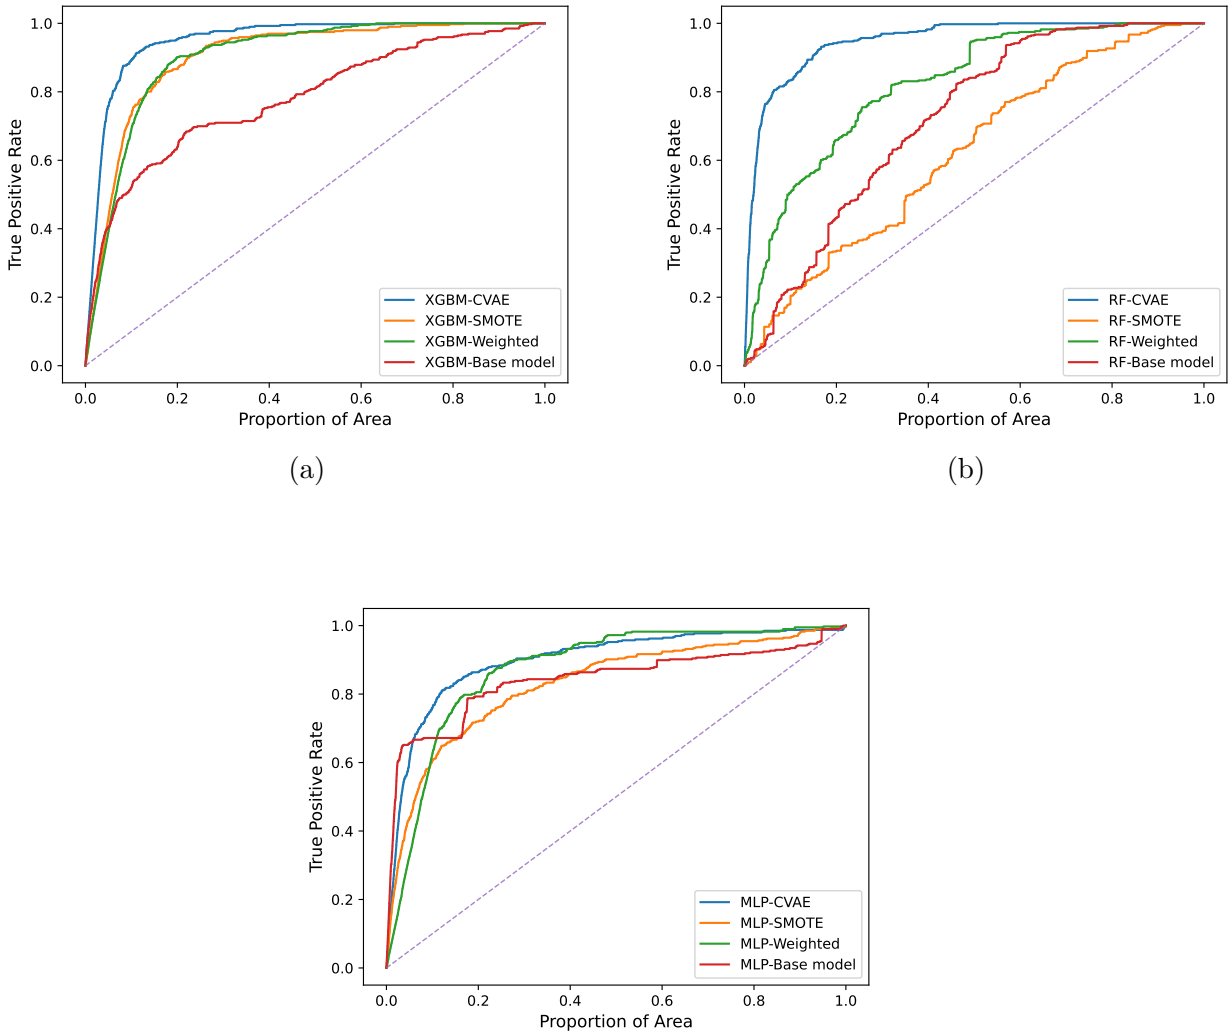


Figure 4.6 Success rate curves for the testing data set with (a) XGBM-MPM, (b) RF-MPM and (c) MLP-MPM models. The dashed lines are the success rate curves characteristic of a random classifier.

The success rate curve analysis indicates that the CVAE data augmentation method consis-

tently outperforms other methods across all classification models, achieving the highest AUC values: 0.94 for XGBM, 0.91 for RF, and 0.90 for MLP. Although SMOTE achieves an AUC of 0.90 with XGBM, its performance drops significantly with RF, resulting in an AUC of only 0.63. The weighted method exhibits stable performance, with AUC values of 0.90 for XGBM, 0.82 for RF, and 0.87 for MLP. In contrast, the base model yields the lowest AUC values across all classifiers, with scores of 0.84 for XGBM, 0.72 for RF, and 0.84 for MLP.

Table 4.1 Area under the success rate curves for the testing data set with each data augmentation and classification method.

	XGBM	RF	MLP
CVAE	**0.94**	**0.91**	**0.90**
SMOTE	0.90	0.63	0.84
Weighted	0.90	0.82	0.87
Base Model	0.84	0.72	0.84

4.5.2 Qualitative representation and geological comparison

Since the XGBM-CVAE model achieves the highest AUC value, we have utilized its results within GIS techniques for qualitative representation, as illustrated in Figure 4.7. The XGBM-CVAE model effectively identifies areas with known Ni mineral deposits, as well as some potential new under-explored regions in the Churchill province near the border of the Slave province and between the Slave and Superior provinces. As shown in Figure 4.3, the area between the Slave and Superior provinces falls within one of the blocks in the testing data set. These results qualitatively demonstrate the effectiveness of the CVAE data augmentation approach in enhancing MPM performance.

A closer representation of the Ni MPM results in the Superior Province, which is in eastern Québec and western Labrador, Canada, is illustrated in Figure 4.8. Overall, based on Figure 4.8, areas with known deposits are accurately recalled, and new prospectivity zones within the New Québec Orogen are also predicted.

We further compare the bedrock geology in the same area (Figure 4.9) because its composition plays a critical role as the source and trap of the magmatic Ni mineral system. Notably, the known Ni deposits in the selected region (Figure 4.9) are closely associated with mafic and ultramafic rocks comprising the New Québec Orogen (Labrador Trough). Additionally, our XGBM-CVAE model predicts a new Ni mineral prospectivity zone, which is depicted in Figure 4.8. Areas with the highest mineral potential correspond to Paleoproterozoic mafic

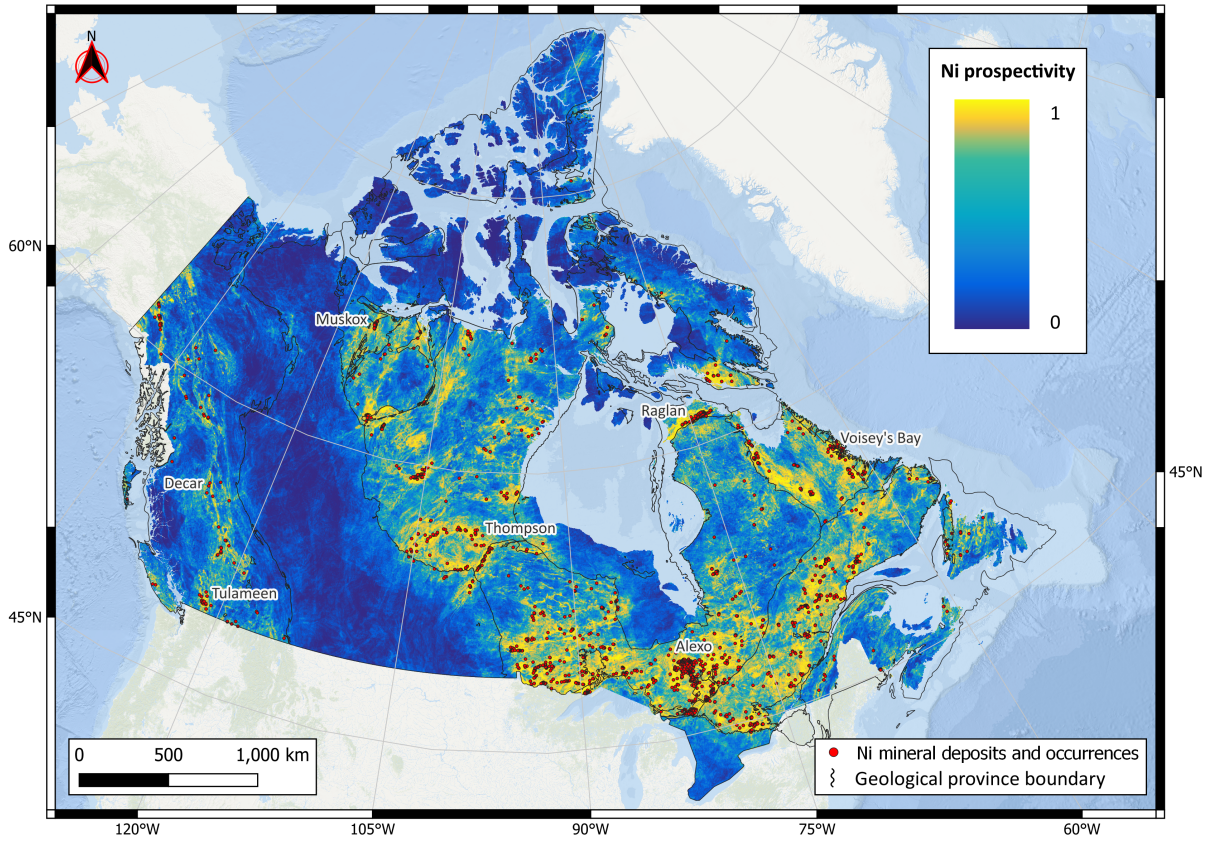


Figure 4.7 XGBM-CVAE prospectivity model results for Ni occurrences in Canada.

to ultramafic sills that have intruded into sulphide-rich sedimentary rocks. Mixing and assimilation of sedimentary rocks is one of the critical processes that can trigger sulphide saturation to form a deposit. The Montagnais Sill Complex represents an example of these prospective rocks within the New Québec Orogen (Smith et al., 2021). This congruence further validates our model's predictions.

Next, by observing the Bouguer gravity HGM in the selected area, which is illustrated in Figure 4.10, we can gain significant insights into variations of rock density in the subsurface caused by lithological contacts.

Interestingly, we observe that regions with Ni mineral deposits and the model-predicted prospective region have high horizontal gravity gradients, which are predominantly associated with contacts between low- and high-density rocks, including mafic to ultramafic lithologies. The relationship between high-density rock formations and the MPM model predictions offers additional validation for the existence of Ni-prospective areas in Figure 4.8. Therefore, this result demonstrates the importance of geophysical data as an input feature in the augmented

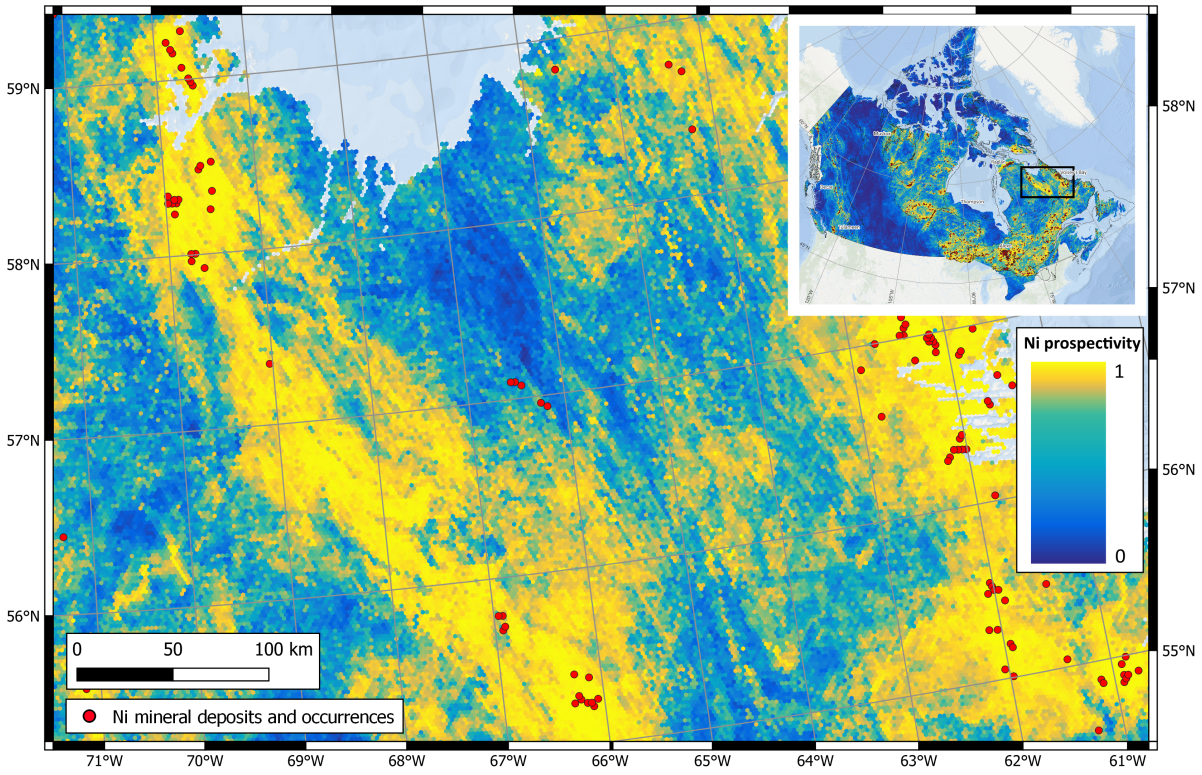


Figure 4.8 Ni mineral prospectivity around eastern Québec and western Labrador, Canada. The black rectangle on the top right side of the map denotes the exact zoomed area.

MPM classification model.

4.5.3 Model analysis

To analyze the CVAE data augmentation model, we start by calculating the root mean squared error (RMSE) between the synthetic output data and actual input data for the testing set, which yields a value of 0.0279. Subsequently, we highlight the generative capabilities of the proposed CVAE decoder. The similarity in distribution patterns between the synthetic and true geoscience data, as illustrated in Figure 4.11, demonstrates not only its capacity to learn and replicate complex geoscience data distributions but also its proficiency in preserving the intricate details of the real data. Specifically, magnetic and gravity HGM, magnetic 1VD, proximity to terrane boundary, gravity Bouguer upward continued HGM, and gravity worms synthetic data sets exhibit high similarity between the two distributions. However, other data sets, such as gravity (Bouguer) and magnetic long wavelength worms, display some

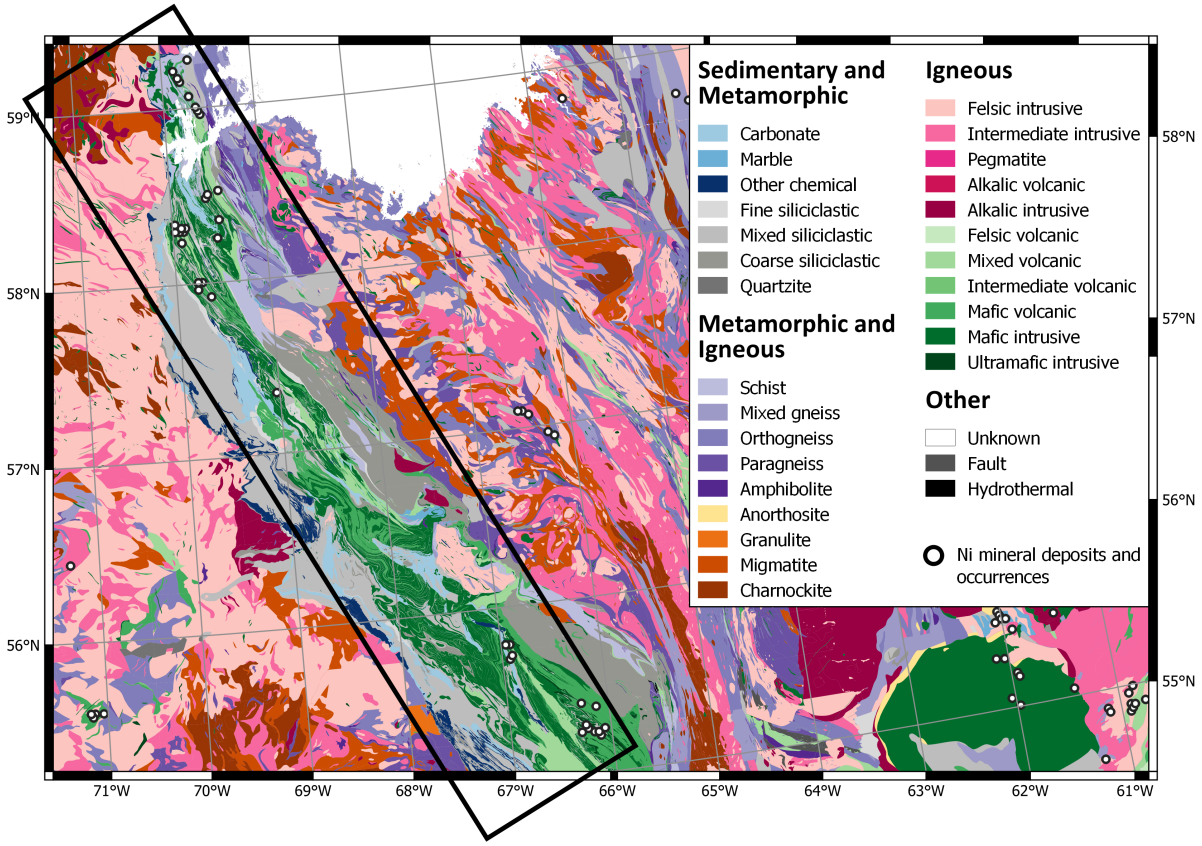


Figure 4.9 Detailed representation of the bedrock geology around eastern Québec and western Labrador, Canada. Modified from Lawley et al. (2024), which provides a comprehensive review of the area's geology.

overlap, albeit with the synthetic data exhibiting a smoother distribution compared to the real data, likely due to the inherent complexity of geoscience data. Further elaboration on this behaviour is provided in the Discussion section.

Finally, we compare the SI and FI results. Figure 4.12 presents the SI for the positive class data generated with the CVAE model. The SI are normalized so that all rows sum up to 100 %.

Interestingly, the SI analysis results align with the earlier real and synthetic data comparisons. For instance, the data sets where the synthetic data showed moderate overlap with real data and also exhibited higher sensitivity values include proximity to passive margins (7.78 %), magnetic long wavelength worms (6.76 %), satellite gravity minimum curvature derivative (6.20 %), satellite gravity shape index (6.14 %), satellite gravity mean curvature derivative (5.08 %), and LAB (5.57 %). This shows that the synthetic data in these data sets are

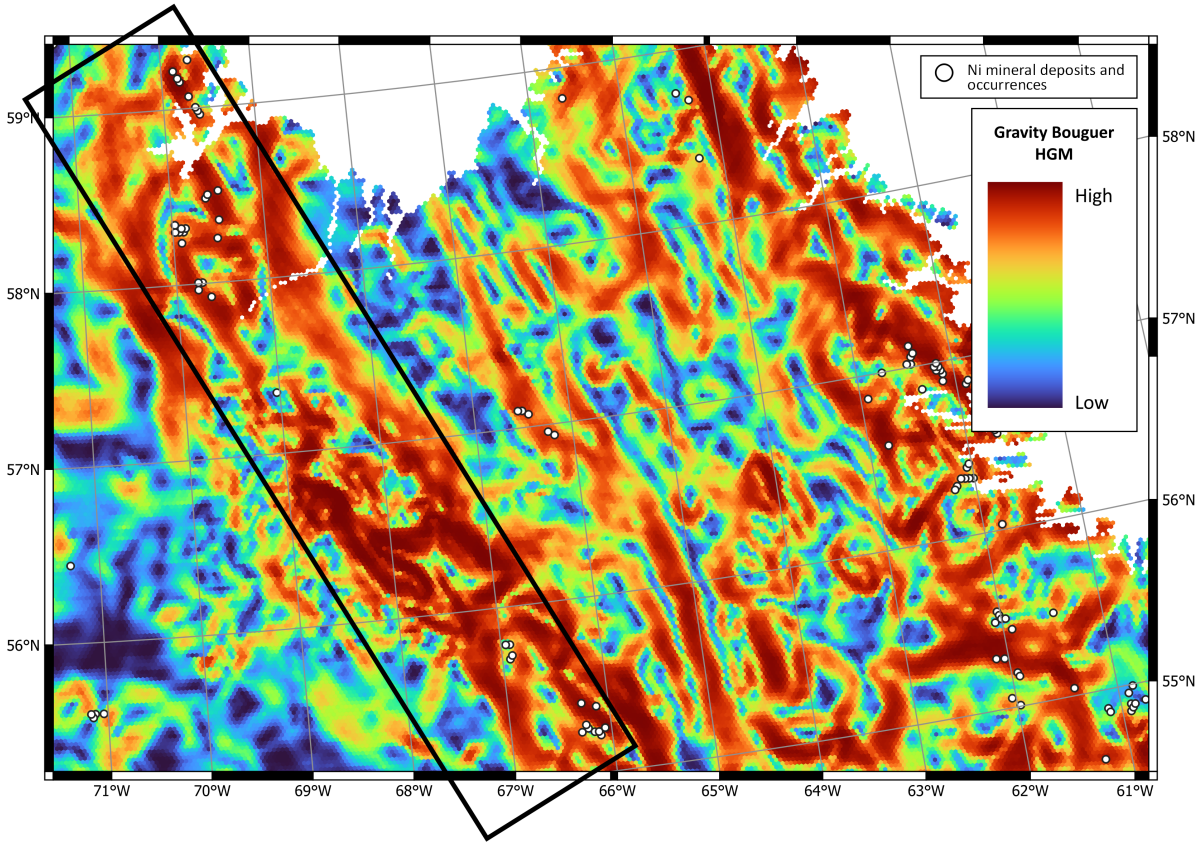


Figure 4.10 Horizontal gradient magnitude of the gravity Bouguer anomaly across the Slave and Superior provinces in Québec, Canada.

more susceptible to changes in the mineralized and barren conditions, which could clarify the differences seen with the real data.

However, for data sets where the synthetic data has a higher degree of overlap with real data, the sensitivity values are the lowest, including magnetic 1VD (0.14 %), gravity worms (0.23 %) and magnetic HGM (0.50 %), respectively. Furthermore, categorical data sets also show relatively low sensitivity values, except for the presence and absence of lithological igneous data. This indicates that the synthetic data in these data sets is less affected by changes in the mineralized and barren conditions.

The FI for the XGBM-CVAE model is illustrated in Figure 4.13. The log scale is used to plot the FI values between the input geoscience data to demonstrate the effect of FI on all the geoscience data.

Figure 4.13 emphasises the significance of each feature in the classification task performed by the XGBM model, allowing for a better understanding of their impact on the model's

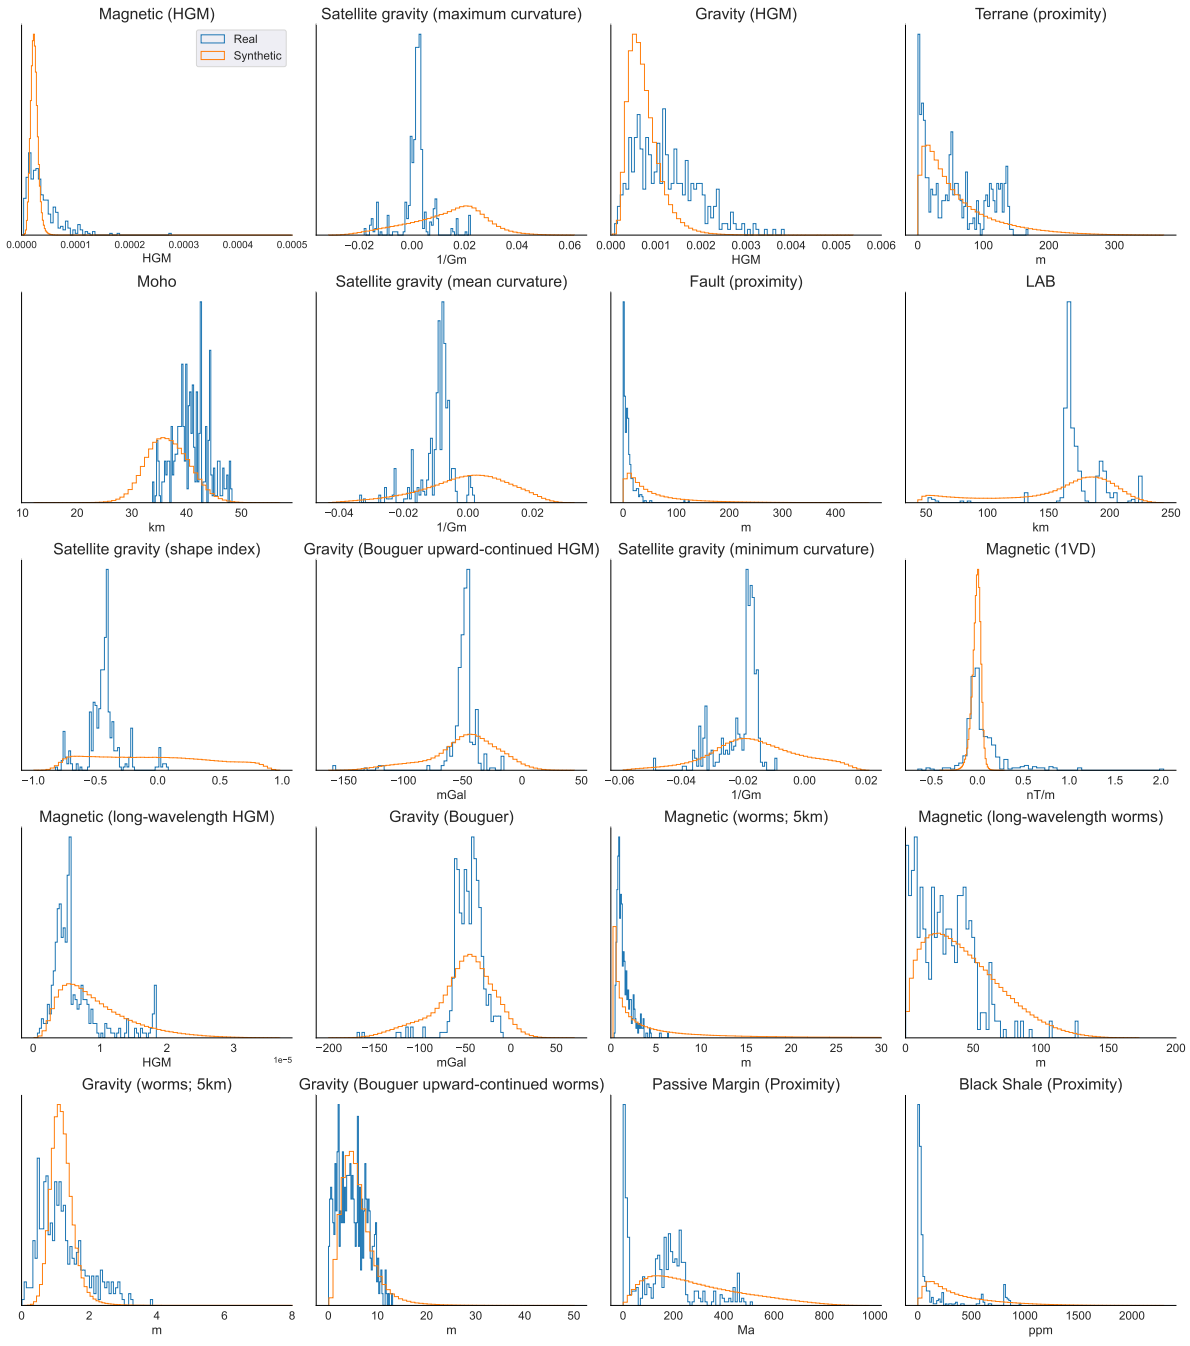


Figure 4.11 Histogram plots of the distributions between generated data from the CVAE model and real geoscience data. One hot encoded data subset is deliberately excluded from this figure due to its high dimensionality.

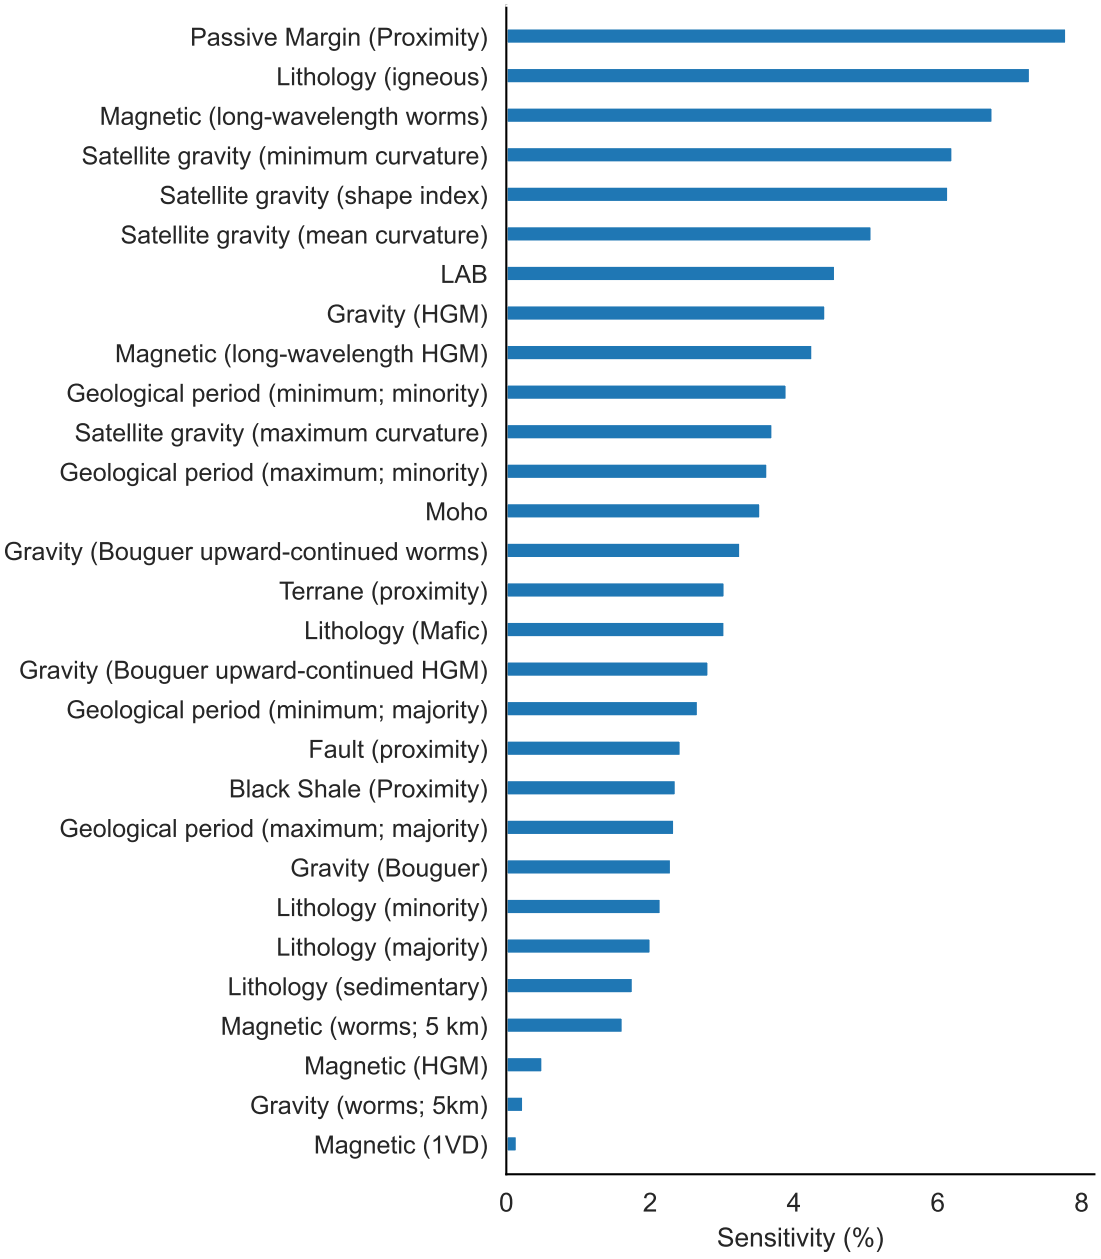


Figure 4.12 The bar chart illustrates the normalized sensitivities of the CVAE's decoded variables to different geological and geophysical data sets. The data are sorted in ascending order of sensitivity. The highest sensitivity corresponds to the augmented geoscience data set, which is most influenced by the mineralized and barren conditions.

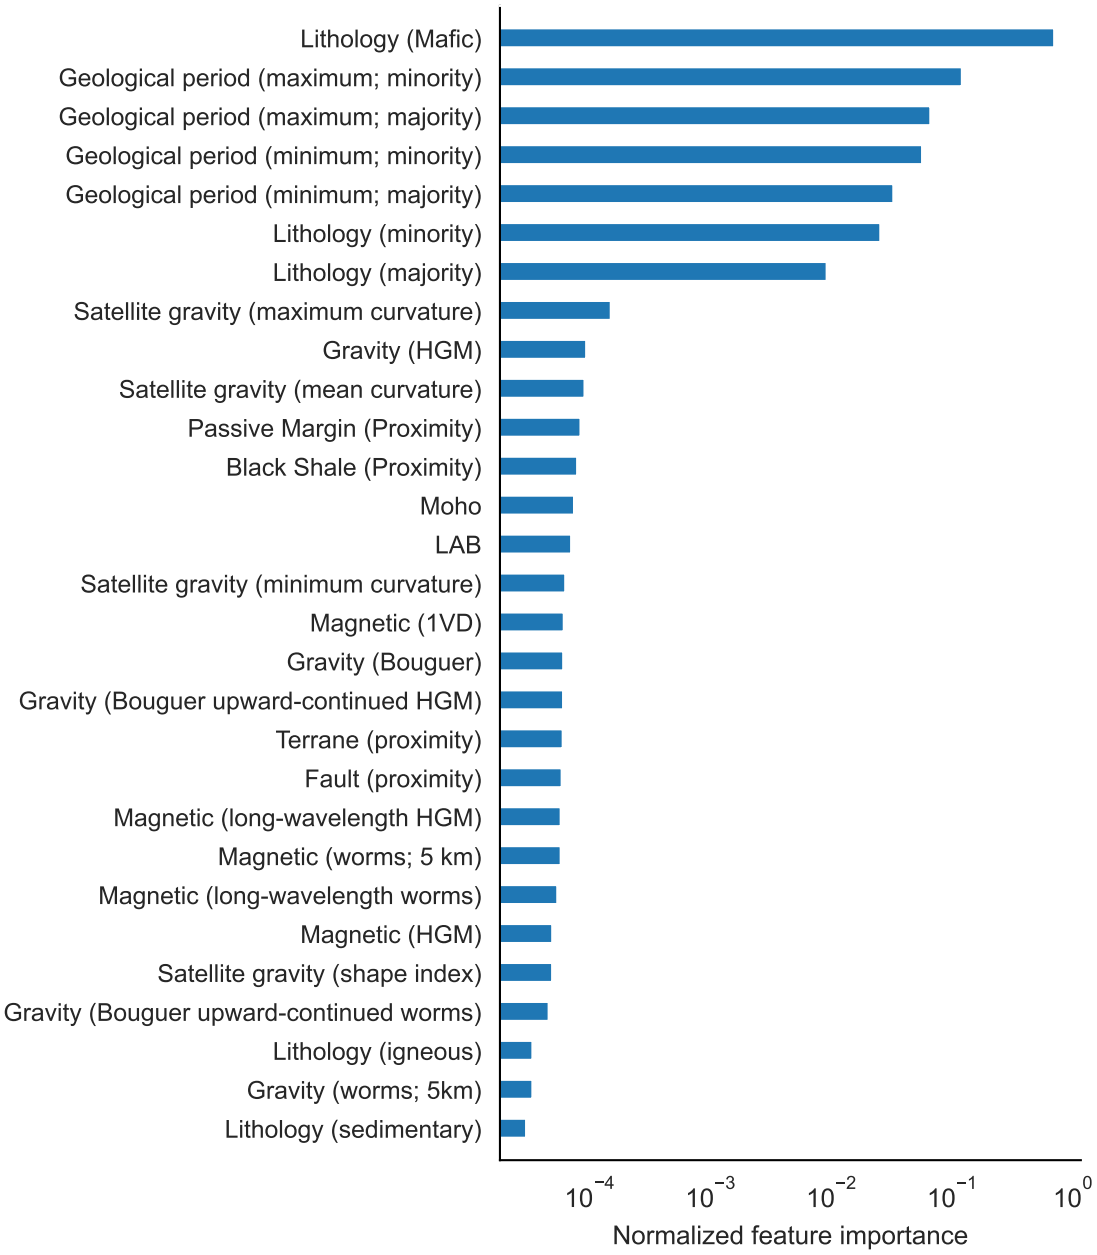


Figure 4.13 Normalized feature importance in the preferred XGBM-CVAE prospectivity model. The importance value of each data set is averaged within its category for one-hot encoded data and directly obtained for non-encoded data. The bars show the average feature importance of each feature, providing insights into their relative significance in the XGBM model's predictions.

predictions. From high to low FI values, the presence or absence of mafic intrusions and minimum and maximum geological periods are followed up with minority and majority lithology data. Even though most geophysical data have relatively low FI values, they still illustrate their effectiveness when mapping the mineral potential. For instance, mineral deposits and occurrences illustrated in Figure 4.10 are associated with the edges of gravity anomalies (i.e., high HGM values) that point to the significance of crustal pathways when forming mineral deposits. Finally, the lowest FI values are observed in the presence and absence of the igneous and sedimentary lithological data, respectively.

4.6 Discussion

The CVAE-based MPM captures complex and non-linear geoscience data relationships and generates high-quality synthetic samples, improving minority class performance. Nevertheless, the CVAE's effectiveness is highly dependent upon the quality of the input geoscience data, and we have noticed that it requires a significant amount of computer resources when handling large data sets. In comparison, SMOTE and class weighting methods are easy to implement and computationally efficient. However, the SMOTE method can introduce noise, which generates more noisy synthetic samples than real geoscience data. Our results show poor performances for some of the SMOTE models, possibly due to the highly complex and non-linear class boundaries, which fail to capture the underlying distribution. As for the class weighting technique, the MPM success rate curves are slightly improved in comparison with the base model. However, class weighting often leads to overfitting due to excessive focus on the minority class (Kumar et al., 2024).

In the context of generative modelling, tractability and flexibility are often presented as two opposing objectives. On the one hand, tractable models are amenable to analytical evaluation and cost-effective data fitting (for instance, via Gaussian or Laplace distributions), but sometimes struggle to capture the intricate structure in very diverse data sets adequately. On the other hand, flexible models, capable of accommodating any data format, typically incur higher computational costs for evaluation, training, or sampling. According to Figure 4.11, the reason some data sets have real and synthetic data partially overlap is that the decoder has learned to map the complex, high-dimensional real data to a simpler, lower-dimensional latent space, and when it maps back to the original space, some of the intricate details of the real data distribution might be lost or smoothed out. One reason for this is due to the cases where two distributions of latent codes overlap with each other, causing the optimal decoding to be the average of those two inputs with the same latent code (Janjoš et al., 2023). This common characteristic of CVAE models is often seen as a trade-off between data fidelity

and model generalization. To address this issue, it is possible to implement other generative models, including conditional generative adversarial networks (CGAN), which can produce high-fidelity samples (Chen et al., 2022). However, mode collapse is a common problem in CGAN, where the discriminator is over-trained. Alternatively, CGAN can also result in catastrophic forgetting, whereby models continue to lose performance on previously learned tasks. Dealing with geoscience data sets in MPM can be challenging for various reasons, including volume and complexity, data heterogeneity, data quality and uncertainty, and computational challenges. Therefore, each generative model has strengths and weaknesses and must be adapted to the problem at hand.

4.7 Conclusions

This study enhances the application of MPM in scenarios with limited ground truth data by introducing a novel data augmentation framework that necessitates only a small amount of realistically available labelled data. The proposed approach is implemented to model Canada's prospective areas for magmatic Ni ($\pm$ Cu $\pm$ Co $\pm$ PGE) sulphide mineral systems. The results reveal that the combination of the CVAE model with the XGBM classifier outperforms all other methods by having a success rate AUC value of 0.94 on a spatially distinct test area, highlighting the potential of this approach in MPM. Furthermore, the accuracy and effectiveness of the XGBM-CVAE model are validated through qualitative analysis. The XGBM-CVAE model not only classifies known mineral deposits accurately but also successfully identifies new prospective regions to support future mineral exploration targeting. In addition to the tests mentioned above, the results indicate high accuracy in the CVAE model's ability to generate augmented data that closely overlaps the original data set. Finally, the geoscience feature analysis concerning Ni MPM models provides valuable insights into the relationships between geophysical, geological, and mineral occurrence data, further demonstrating the reliability of the proposed data augmentation method. In future work, it is possible to leverage the CVAE's probabilistic framework to explore and analyze how data and modelling uncertainties propagate into the MPM model's predictions. This is an important step forward due to model-based risk analysis and decision-making, which depend heavily on uncertainty analysis since it provides risk assessors and decision-makers insight into the reliability of MPM predictions.

Acknowledgements

The authors acknowledge funding from Natural Resources Canada, [TGI project GC-130021S], and support of the Natural Sciences and Engineering Research Council of Canada, [funding reference number 2024-06161]. The authors thank the Editor and reviewers for constructive criticism of the manuscript.

CHAPTER 5 ARTICLE 2: A GENERATIVE NEURAL NETWORK APPROACH TO UNCERTAINTY AND RISK-RETURN ANALYSIS IN MINERAL PROSPECTIVITY MODELLING

Ore Geology Reviews
Volume 186, 23 August 2025, 106844

Lahiru M.A. Nagasingha[a], Charles L. Bérubé[a], Reza Ghanati[b]

[a] Civil, Geological and Mining Engineering Department, Polytechnique Montréal, Montréal, H3C 3A7, QC, Canada

[b] Department of Earth Physics, Institute of Geophysics, University of Tehran, 14359-44411, Tehran, Iran

Authorship contribution statement

Lahiru M.A. Nagasingha: Writing – original draft, review & editing, Literature research, Experimentation, Validation, Software, Methodology, Formal analysis, Visualization, Data curation, Conceptualization. **Charles L. Bérubé:** Writing – review & editing, Supervision, Validation, Resources, Investigation, Funding acquisition. **Reza Ghanati:** Writing – review & editing.

Abstract

The industrial adoption of machine learning techniques for mineral prospectivity modelling (MPM) remains limited due to their inability to model uncertainties and a lack of systematic frameworks for evaluating risk and return in mineral predictions. A major challenge is that most existing methods fail to simultaneously capture both epistemic uncertainty, which arises from limitations in the predictive modelling process, and aleatoric uncertainty, which stems from the inherent randomness in geoscience data. To address this, we propose a conditional variational autoencoder (CVAE) approach incorporating decoder calibration and uncertainty estimation, which we apply to Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems. Aleatoric uncertainty is quantified from the CVAE's posterior distribution, whereas epistemic uncertainty is assessed from 100 MPM realizations based on datasets generated by the CVAE. We also introduce a novel risk-return framework which integrates relative uncertainty measures with the non-parametric Getis-Ord G^* statistics spatial clustering technique to categorize exploration targets into four distinct risk-return categories. Results from the

spatial distribution and kernel density estimation analysis reveal that most known deposits are situated in low-uncertainty zones. Notably, high-return zones, which comprise approximately 4 % of the total area, account for 94.7 % of the known deposits. This research highlights the significance of incorporating uncertainty and risk-return analysis to improve decision-making in mineral prospecting.

5.1 Introduction

With the progressive exhaustion of shallow exploration search spaces, discovering new mineral resources has transitioned from surface outcroppings to identifying and examining deeper deposits (Yousefi and Kreuzer, 2024). A better understanding of ore-forming processes is therefore necessary to explore new deep-seated spaces. However, discovering new mineral deposits is a challenging and time-consuming task with low success rates (Kreuzer and Etheridge, 2010). This process is facilitated by mineral prospectivity modelling (MPM), which identifies favourable locations for specific types of undiscovered mineral deposits through relationships between known mineral deposits and relevant geoscience datasets (Bonham-Carter, 1994; Carranza, 2008; Yousefi and Carranza, 2016). The availability of extensive geoscience data compilations enables machine learning (ML) techniques to integrate and interpret complex, high-dimensional, and non-linear geospatial datasets in MPM (Porwal and Carranza, 2015; Xiong et al., 2018).

ML algorithms can be categorized into supervised, unsupervised, and semi-supervised techniques. Supervised methods require training datasets with labelled examples, where positive labels represent locations of confirmed mineral deposits and negative labels represent barren areas. These methods examine the relationships among evidence features, such as remote sensing, geological, geochemical, and geophysical data, in relation to positive or negative labels to predict prospective areas. However, supervised methods are limited because ore-forming processes are complex, poorly understood, and inaccessible to direct observation (Cheng, 2007; Yousefi et al., 2024b). This leads to inadequate positive labels for supervised models, compounded by severe class imbalance, where negative samples vastly outnumber positive samples, resulting in biased predictions that favour the majority class (Rahimi et al., 2021). Conversely, unsupervised techniques identify patterns and structures in geoscience data without requiring labelled training samples, making them suitable when exploration data is sparse. Nonetheless, these methods fail to efficiently utilize valuable information from existing positive and negative labels when such data are available. To address these challenges, Nagasingha et al. (2024) proposed a semi-supervised conditional variational autoencoder (CVAE) approach that augments positive samples to create a balanced dataset.

Then, they use extreme gradient boosting (XGBoost) to model magmatic Ni (±Cu ±Co ±platinum group elements (PGE)) sulphide mineral systems in Canada. However, a key limitation of that study is that it focuses solely on predictive accuracy, without a framework to quantify the uncertainty or risk associated with its predictions.

Uncertainties exist at every stage of data-driven modelling (Wang and Huang, 2015), and ML techniques applied to MPM often act as black box models, leading to poor generalizability, limited interpretability, and discrepancies in model parameters and outcomes (Karpatne et al., 2017; Reichstein et al., 2019). Additionally, neglecting prior and expert knowledge hinders the interpretation of MPM results for geoscientists (Zuo et al., 2024). Therefore, assessing uncertainty in MPM is crucial for ensuring the reliability of predictions and understanding the confidence levels associated with potentially undiscovered mineral occurrences. MPM aims to identify target locations that exhibit high returns, which are determined by the likelihood of success and failure in exploring for certain mineral deposits (Singer and Kouda, 1999). MPM uncertainty is closely associated with risk-return analysis, and high exploration risk can sometimes lead to high returns. However, it is essential to quantify uncertainties in MPM and find low-risk, high-return areas during the exploration process (Caers, 2011; Zuo and Wang, 2020).

Within the context of MPM, uncertainty types can be divided into four main categories. First, process-related uncertainties are arising from a lack of knowledge about the processes that control ore-forming mechanisms. These processes are particularly significant as mineral deposit formation cannot be directly observed or measured, and are primarily determined by mineral conceptual models and limited knowledge (Kreuzer et al., 2020; Zhang et al., 2024). Second, aleatoric uncertainty results from intrinsic randomness, noise, or overlapping feature distributions within geoscience data. This type of uncertainty cannot be mitigated even with perfect knowledge and additional training data (Hüllermeier and Waegeman, 2021). Aleatoric uncertainty may occur during data acquisition due to instrument and transmission errors, as well as inappropriate sampling techniques. Data storage, representation, and interpolation methods can exacerbate this uncertainty. For example, Yang et al. (2023) and Gao et al. (2024) studied aleatoric uncertainty caused by the evidence layers in MPM using the Monte Carlo simulation and a hybrid model that integrates the direct sampling technique with a convolutional neural network, respectively. Third, epistemic uncertainty results from ambiguity in a model's predictions due to a lack of knowledge regarding MPM choices based on current input geoscience data. This includes model architecture uncertainty, parameter uncertainty resulting from imperfect learning algorithms or limited training data, and uncertainties in the training sample distribution (He et al., 2023). To address epistemic MPM uncertainty, Yousefi et al. (2024a) proposed a majority voting and confidence

index-based approach, while Zhang et al. (2022b) used Bayesian logistic regression. Fourth, judgment-related uncertainty arises from irrational, inaccurate, and illogical judgments that significantly influence human thinking and decision-making (Zuo et al., 2021a). In mineral exploration, epistemic and aleatoric uncertainties can hinder rational decision-making, often leading geoscientists to rely on intuition. The impact of these uncertainties largely depends on how well ore-forming processes can be translated into mappable criteria (McCuaig et al., 2010). Overall, aleatoric uncertainty significantly impacts geoscientific models, such as 3D geomodelling (Lindsay et al., 2012; Wellmann et al., 2010), whereas epistemic uncertainty holds greater importance for regional-scale MPM (Yousefi et al., 2024a).

To date, variational autoencoders (VAEs) and CVAEs have been successfully applied in various real-world applications for quantifying uncertainties, including aerodynamic applications (Saetta et al., 2024), linear inverse problems (Böhm et al., 2019), and medical image reconstructions (Edupuganti et al., 2020). However, uncertainty quantification in MPM using generative models remains underexplored. This study investigates how decoder calibration techniques applied to a CVAE can improve the accuracy and reliability of uncertainty estimates in MPM. A framework is developed to quantify the uncertainties in the MPM results presented by Nagasingha et al. (2024) and to assess confidence in exploration targeting. Specifically, we aim to: (1) assess aleatoric uncertainty through the posterior distribution of the CVAE; (2) analyze epistemic uncertainty using an ensemble of models combined with relative uncertainty metrics; and (3) propose a novel risk-return framework that integrates uncertainty results and spatial clustering to guide decision-making in mineral exploration. To the best of our knowledge, this is the first application of an ensemble-based generative neural network approach for uncertainty quantification in MPM.

5.2 The mineral system approach and related uncertainties

The mineral system concept was initially introduced by Wyborn et al. (1994) and further refined by Knox-Robinson and Wyborn (1997) and Hronsky and Groves (2008), and is a fundamental principle in modern mineral exploration. This approach is now widely used to develop genetic models and conceptual frameworks (Hagemann et al., 2016; McCuaig et al., 2010; Wyman et al., 2016) that play a pivotal role in MPM (Groves et al., 2022; Skirrow et al., 2019; Zuo et al., 2025). The mineral system approach can help determine the likelihood of a specific mineral deposit in a particular area by breaking down the mineralizing process into their fundamental stages and estimating likelihood independently (Fallon et al., 2010; Kreuzer et al., 2008). This approach can enhance MPM by leveraging geological processes and tectonic settings rather than merely relying on pattern recognition. Consequently, the

prospectivity model integrates all available geoscience data with a thorough understanding of the mineralization processes relevant to the deposit, which can help reduce uncertainties in exploration targeting (Hronsky and Groves, 2008; Partington et al., 2024).

The evolution of a mineral system involves processes that span vast timescales, from long-term architectural development to transient mineralizing events (Blewett et al., 2010; Czarnota et al., 2010). Many critical features of the mineral system represent ground preparation, which is the development of favourable structures and rock types that occur long before the mineralizing event itself (Bruce et al., 2020). For example, the fundamental architecture of the system, such as the development of long-lived crustal magma staging chambers and crustal-scale faults that later act as fluid conduits, is the result of ground preparation established long before the actual mineralization occurs. On the other hand, exhumation and upgrading through metamorphism or supergene enrichment occur subsequently. Therefore, the mineral system approach is more practical than conventional descriptive deposit models, as it emphasizes critical geological processes necessary to form major mineral deposits (Lord et al., 2001). It can explain the spatial and temporal co-occurrence of certain deposits within a specific province and families of coeval mineral systems that may have formed in the same or adjacent terranes or provinces (Hagemann et al., 2016).

Most importantly, when developing a targeting model for a mineral system, four key factors of geological uncertainty must be considered (Bárdossy and Fodor, 2001; McCuaig et al., 2010). These factors, which range from the quality of the input data to the completeness of our conceptual understanding of ore-forming processes, are summarized in Table 5.1. As outlined in Table 5.1, the first three sources of uncertainty stem from the challenge of using measurable proxies to represent geological elements and processes, while the fourth arises from an incomplete understanding of the mineral system itself. This distinction highlights a common trade-off that must be made between the representativeness of a targeting criterion and the spatial resolution of the data used to map it.

Table 5.1 Four key factors of geological uncertainty in MPM.

Uncertainty type	Key question	Controlling factors and examples	Quantification challenge
Accuracy	How good is the raw data?	The quality, resolution, and signal-to-noise ratio of the geoscience datasets	Relatively easy to quantify using standard geographic information system (GIS) techniques
Precision	Is the mappable proxy a good representation of the target element?	The degree to which a mappable feature (e.g., a geophysical anomaly) accurately represents the geological element of interest (e.g., a specific rock type)	Challenging to quantify; often disregarded by geoscientists
Alignment	Does the target element truly relate to the critical ore-forming process?	The confidence that a mapped feature represents a true geological structure (e.g., identifying a mantle-tapping or translithospheric fault). This is challenged by the unique limitations inherent to each geophysical dataset (e.g., non-unique interpretations for gravity/magnetics; coarse resolution for magnetotellurics; difficulty in imaging steep structures with seismic)	Challenging; requires integrating multiple datasets to reduce ambiguity

Continued on next page

Table 5.1 Four key factors of geological uncertainty in MPM. (continued)

Uncertainty type	Key question	Controlling factors and examples	Quantification challenge
Knowledge	Is the conceptual model of the mineral system correct?	The completeness of our fundamental understanding of the critical mineralization processes. Models often rely on empirical proxies that may not fully capture reality	Difficult to quantify as it is a conceptual, knowledge-based uncertainty

5.3 Magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems

Concerning this study, magmatic Ni (±Cu ±Co ±PGE) sulphide deposits are formed through the segregation and concentration of liquid sulphide droplets from mafic or ultramafic magma, along with the partitioning of chalcophile elements from a silicate melt (Naldrett, 1997, 2005). Thanks to the mineral systems approach, we can investigate the genesis of magmatic Ni (±Cu ±Co ±PGE) sulphide deposits from a lithospheric-scale perspective by considering drivers, sources, pathways, and trapping processes. This includes evaluating the nature of magmatism in relation to cratonic architecture, the role of magmatic and structural controls in forming magma conduits, access to crustal sulphur sources, favourable intrusion geometry for sulphide deposition and upgrading, and the structural and erosional history that influences preservation and detectability (Barnes et al., 2016). Most importantly, these elements can be converted into mappable geological criteria. Furthermore, magmatic mineral systems can be classified into subtypes according to host rocks and tectonic environments: komatiite-hosted, flood basalt-hosted, anorthosite-hosted, ferropicrite-hosted, and Alaska-type (Naldrett, 1981). While some processes may vary between these subtypes at the depositional scale, they are typically considered together for national-scale modelling, as their fundamental, mappable criteria for sources, pathways, and traps are broadly consistent at this scale (Lawley et al., 2022b; McCuaig et al., 2010).

Firstly, in the magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system, which is illustrated in Figure 5.1, mantle melting is the primary source and driver of ore-forming elements (Lawley et al., 2021). Thus, it is necessary to have higher-resolution, accurate, and detailed bedrock geological information in order to reduce potential uncertainties and map the small ultramafic to mafic intrusions and primary host rocks for magmatic Ni (±Cu ±Co ±PGE)

sulphide deposits. Extensive partial melting often generates the most promising ore-forming magmas, highlighting specific geological settings and periods as optimal for these mineral systems (Barnes, 1990; Barnes and Fiorentini, 2012; Begg et al., 2018). The high prospectivity of hot mantle temperatures is demonstrated by numerous komatiite-associated Ni (±Cu ±Co) deposits. Secondly, structural pathways for forming a Ni deposit guide ore-forming magmas from their lithospheric source toward the mid- to upper-crust, whereas crustal pathways facilitate the concentration and refinement of these mantle-derived melts. Deposits can form either at the surface, such as with komatiites, or within the upper crust. These magmas typically move through tube or blade-shaped dykes, channelized sills, or, in some Archean and Paleoproterozoic cases, as komatiitic lava flows (Arndt et al., 2005; Naldrett, 2013). Therefore, geophysical datasets play a critical role in mapping these pathways, as mafic and ultramafic rocks that form these conduits and flows exhibit distinctive physical properties, such as higher density and magnetic susceptibility, compared to most crustal host rocks. Finally, depositional traps in these mineral systems depend on the sulphur saturation of ore-forming magmas, often driven by the assimilation of sulphur-bearing country rocks during ascent or emplacement (Eckstrand and Hulbert, 2007). However, mapping the trap component is particularly challenging at a national scale. While the potential for this geochemical trigger can be proxied, many of the specific physical trap mechanisms (e.g., changes in conduit geometry, mechanical accumulation) (Arndt et al., 2005; Barnes et al., 2016; Naldrett, 1981) are often impossible to identify with the available regional datasets (Lawley et al., 2021).

5.4 Dataset

The datasets used in this study are the same as those from our earlier publication on magmatic Ni (±Cu ±Co ±PGE) mineral systems (Nagasingha et al., 2024), which includes geological and geophysical data sourced from publicly available Canadian geoscience data compilations (Lawley et al., 2022a, 2021). Initially, we converted all vector and raster geoscience datasets into the EPSG-4326 coordinate reference system and indexed them using the H3 discrete global grid system (DGGS). The H3 DGGS technique uses hexagonal tiling to cover the globe at multiple resolutions, assigning unique addresses to each hexagon. Consequently, the Canadian landmass is spatially indexed at a resolution of 7, resulting in approximately 1.8 million unique H3 addresses.

The primary geological dataset used for this study is the new Canada Geological Map Compilation (CGMC), shown in Figure 5.2 (Lawley et al., 2024). Unlike previous single-scale national maps, the CGMC is a multi-source patchwork compilation, integrating 32 previously

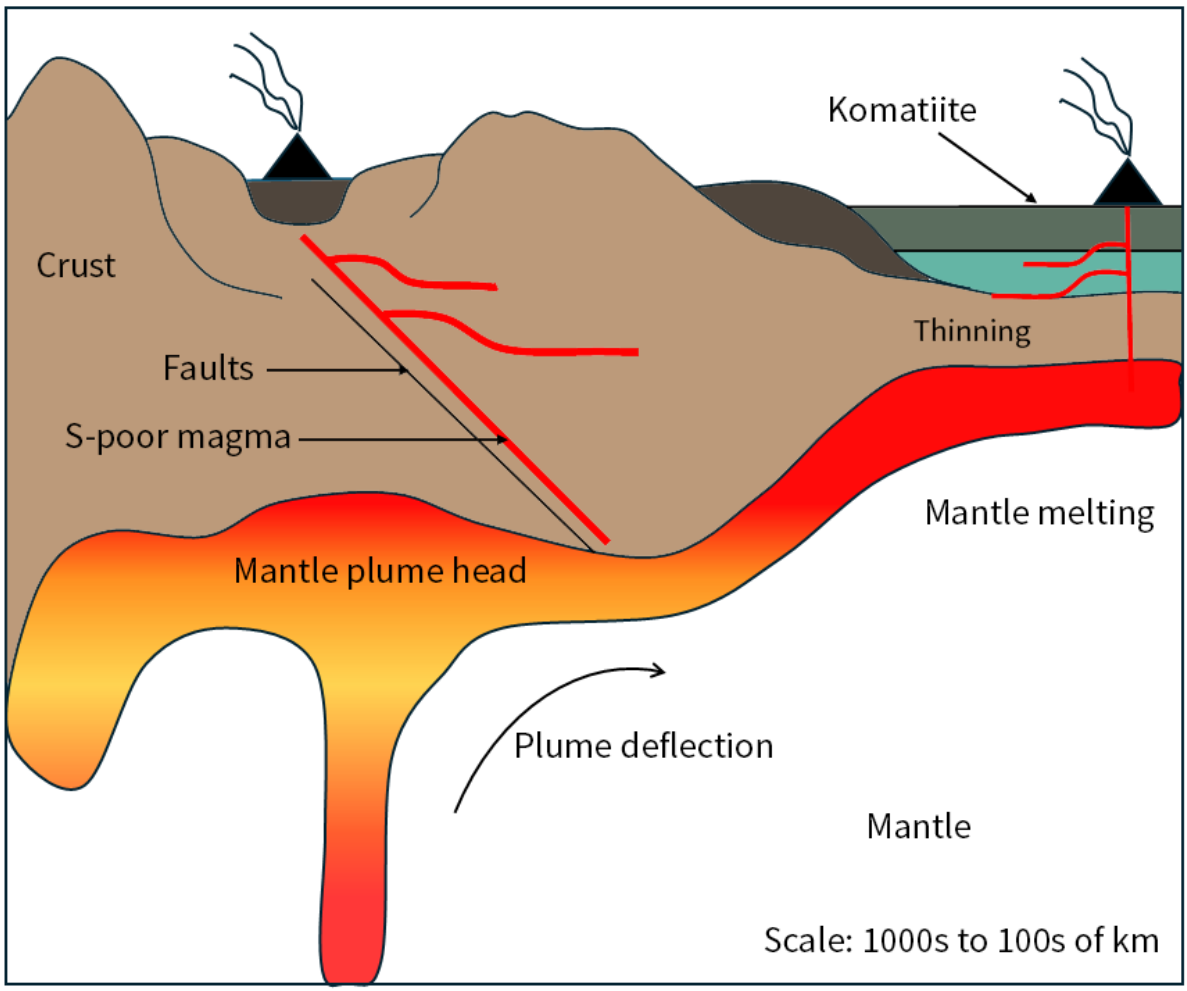


Figure 5.1 Conceptual model for magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems. The model illustrates mantle plumes as the primary drivers that trigger mantle melting and the generation of sulphur-poor (S-poor) magma that rises through pathways such as thinning crust and faults. Key attributes, including geological faults, magnetic and gravity anomalies, lithospheric thickness, and geological terrane boundaries, are used to map the movements of magma. Structural and lithological traps facilitate sulphide deposition, with Bouguer gravity providing information on these mechanisms. Modified from Lawley et al. (2021).

published bedrock maps of varying scales, from 1:30,000 to 1:5,000,000. Key enhancements include correcting geometry errors and applying a new hierarchical classification scheme, which standardizes the large and varied set of lithological descriptions from the 32 source maps into 35 consistent classes. However, it is crucial to acknowledge the dataset's inherent limitations for prospectivity targeting, such as the unaddressed overlapping map polygons and boundary artifacts between jurisdictions noted by Lawley et al. (2024). Therefore, the final map represents the most accurate and up-to-date information available for any given

location.

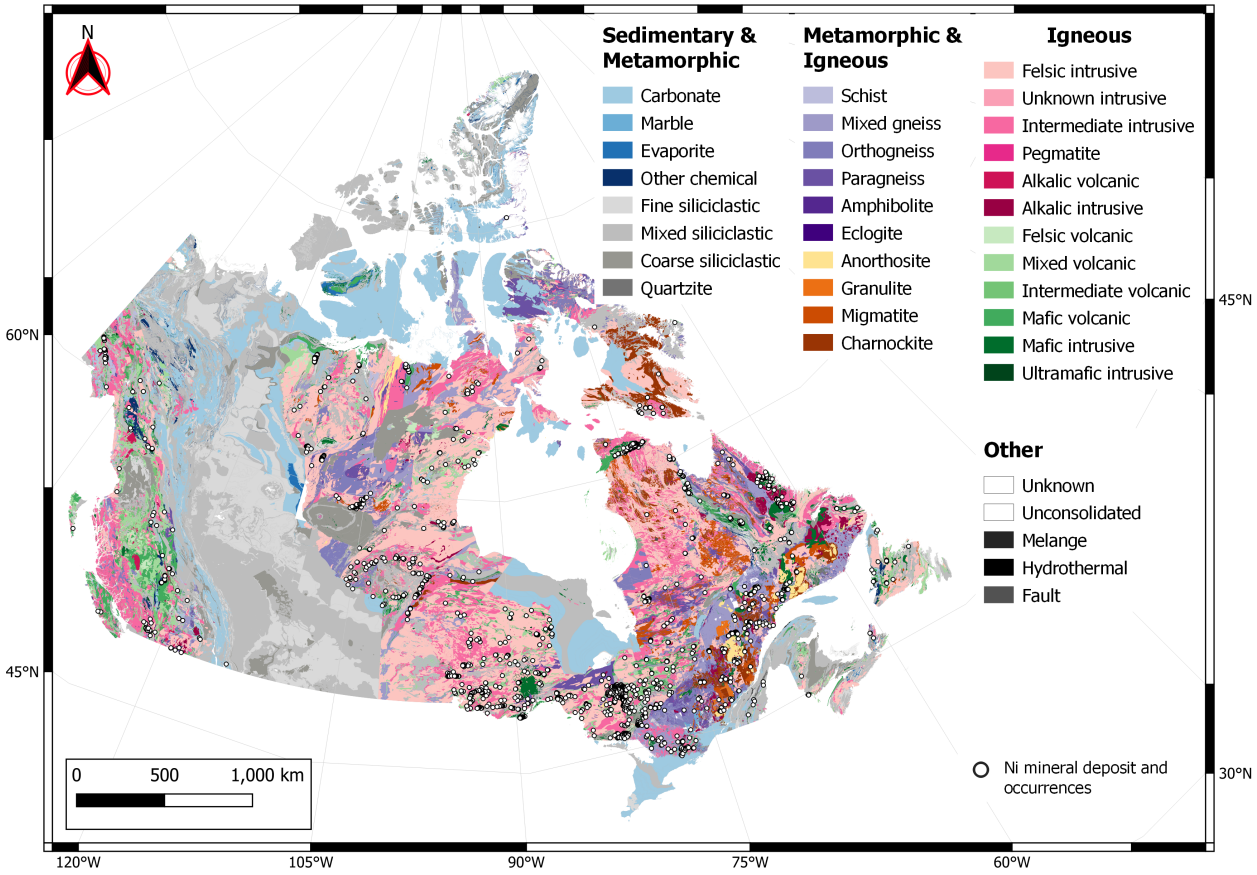


Figure 5.2 New bedrock geological map of Canada (Lawley et al., 2024). The bedrock geology shown on this map is used to define the source and trap components of the mineral systems approach. The known Ni mineral deposits and occurrences demonstrated here are used for model training.

Additionally, our mineral system framework utilizes the maximum geological ages of surface bedrock to determine sources and transport mechanisms. In contrast, minimum ages illustrated in Figure 5.3 reflect the possibility of igneous or sedimentary rocks that are coeval with major mantle melting events (Lawley et al., 2022a).

Intrusions hosting Ni deposits are often small or rare and absent from regional geological maps. Therefore, the primary evidence for mantle-derived melts often comes from small mafic dykes and intrusions, as illustrated in the descriptions of the majority map units. We use geological map data (Figure 5.2), which includes the majority and minority rock types, as well as unstructured map unit descriptions, to map the mantle sources of magmatic mineral systems. Proximity to faults, terrane boundaries, and lithological contacts is used to map mineral system pathways at shallower crustal depths. Next, gravity and magnetic survey

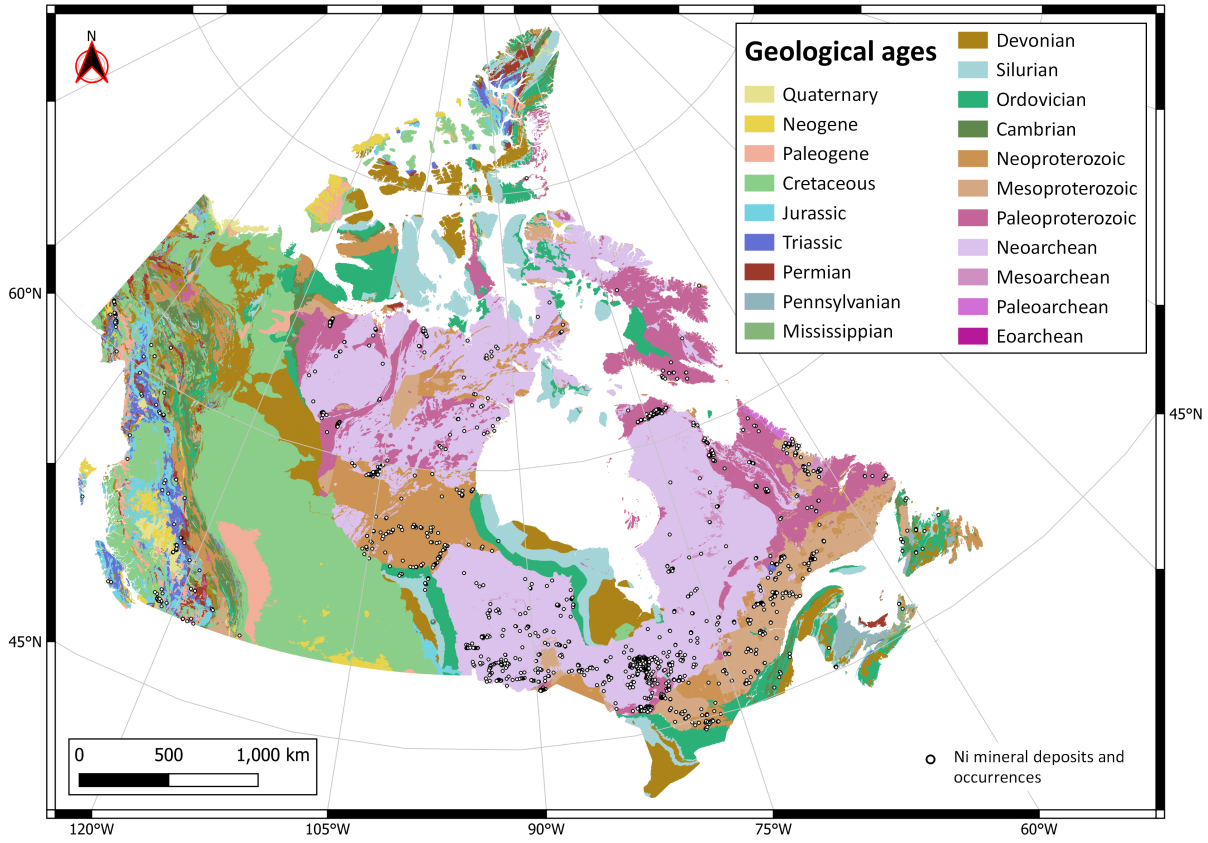


Figure 5.3 Spatial distribution of the minimum geological age map of Canada. These minimum geological ages are used to identify rocks that are coeval with major mantle melting events, which act as a transient driver in the mineral system.

compilations serve multiple purposes. These datasets, including derivative products such as the first vertical derivative (1VD), horizontal gradient magnitude (HGM), and satellite gravity dataset, are used to: 1) map mineral system pathways under surficial cover at multiple depths within the crust; 2) delineate paleotectonic settings; and 3) find density and magnetic contrasts indicative of major lithological boundaries (Ebbing et al., 2018; Lawley et al., 2022b). Additionally, seismic data compilations that include the depth to the lithosphere-asthenosphere boundary (LAB) and the depth to the Mohorovičić discontinuity (Moho) are used to map abrupt changes in the thickness of the lithosphere and crust, which are important for determining where magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system subtypes may be located (Champion and Huston, 2016; Griffin et al., 2013; Huston et al., 2016). We assess the potential of a depositional trapping mechanism across Canada by utilizing carbonaceous rock types found in geological map databases, as their assimilation by magma can be a key trigger for sulphide deposition (Lawley et al., 2022b). However, the mechanisms

recognized for segregating and accumulating sulphide liquid, such as mechanical accumulation and gravity-driven processes, can be difficult to determine from the currently available datasets (Arndt et al., 2005; Lawley et al., 2021). Figure 5.4 illustrates a combination of high and low-resolution qualitative representations of some of the feature anomalies that are described. Finally, based on the best available datasets, we have primarily focused on mapping the magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system through sources (e.g., ultramafic to mafic melts), drivers (e.g., mantle melting events) and pathways (e.g., regions of mantle and crustal weakness). Table 5.2 summarizes the mineral system components and related datasets used in this study.

Table 5.2 Mappable criteria and datasets utilized for prospectivity modelling of the magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system (Lawley et al., 2021).

Component	Mappable criteria	Dataset(s)
Drivers	Timing of melt events	Minimum and maximum geological ages of bedrock units (proxy for timing of major magmatic events)
Sources	Presence of mafic to ultramafic rocks	Lithology classifications from bedrock maps and the igneous rock dictionary
Pathways	Crustal thickness	Depth to the Mohorovičić discontinuity (Moho)
	Lithospheric thickness	Depth to the lithosphere-asthenosphere Boundary (LAB)
	Major crustal structures	Proximity to known geological faults
	Terrane architecture	Proximity to geological terrane boundaries
	Crustal density structures	Regional ground gravity products (e.g., Horizontal gradient magnitude (HGM), upward-continued HGM, worms)
	Magnetic anomalies and structural lineaments	Regional airborne magnetic products (e.g., 1st vertical derivative (1VD), HGM, long-wavelength HGM, worms)
	Long-wavelength lithospheric variations	Satellite gravity derivative products (e.g., curvature, shape index)

Continued on next page

Table 5.2 Mappable criteria and datasets utilized for prospectivity modelling of the magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system (Lawley et al., 2021). (continued)

Component	Mappable criteria	Dataset(s)
Traps	Potential for sulphur contamination	Proximity to carbonaceous rocks (derived from geological maps and dictionaries)
	Favourable conduit and chamber geometries hosted by dense mafic-ultramafic rocks	Bouguer gravity anomaly

Next, we employed a spatial cross-validation approach (Valavi et al., 2018) to split the total dataset into training and testing partitions. The total geoscience dataset comprises 1,820,815 individual data points across 29 geological and geophysical features, including 1,452 known Ni deposits and occurrences. We used a six-fold cross-validation technique to divide the total dataset into training (1,491,227 points with 1,056 known deposits) and testing (329,588 points with 396 known deposits) sets, using block sizes of 500 km × 250 km as determined through semivariogram analysis. Additionally, we have considered known Ni mineral deposits and occurrences as true positives during model training (refer to Lawley et al., 2022b; Nagasingha et al., 2024, for further details about the training and testing datasets). This spatial cross-validation approach reduces the risk of overestimating model performance associated with random data splitting (Raschka, 2018; Roberts et al., 2017; Valavi et al., 2018).

5.5 Methods

The proposed framework involves data pre-processing, generative modelling, model training and testing, uncertainty quantification, risk-return analysis, and is summarized in Figure 5.5. Detailed explanations are provided in the following sections for the role and implementation of each step within the framework.

5.5.1 Data pre-processing

We have performed several pre-processing steps, including replacing outliers, imputing missing data, converting categories into numerical data, and scaling. Firstly, Tukey fences are

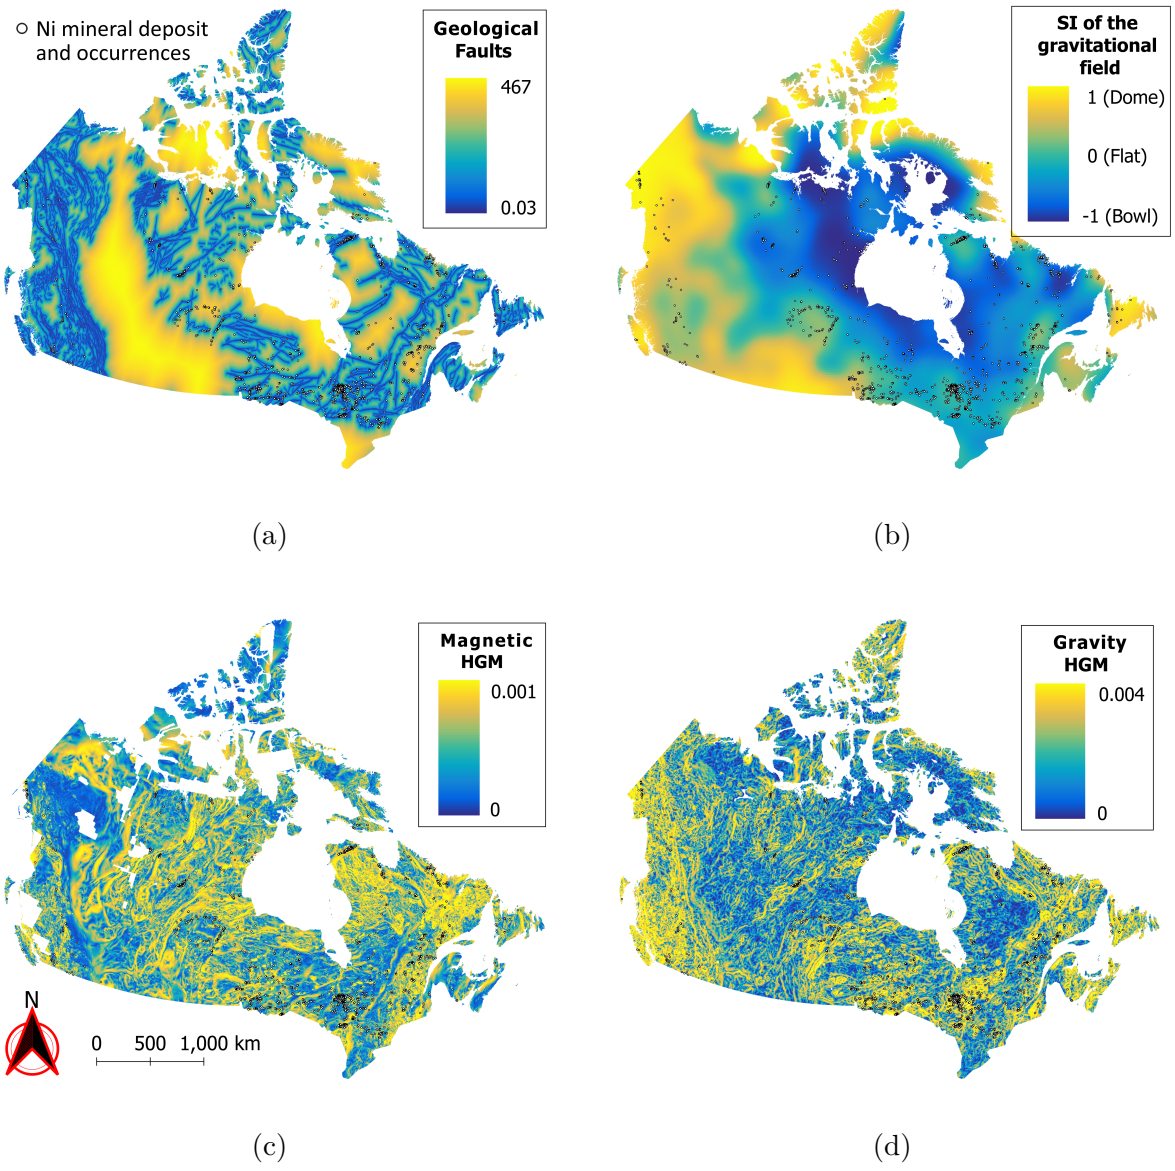


Figure 5.4 Examples of geoscience features used in this study: (a) faults dataset highlighting geological structures (Styron and Pagani, 2020), (b) shape index (SI) derived from the Gravity field and steady-state Ocean Circulation Explorer satellite gravity data representing variations in the Earth's gravitational field (Ebbing et al., 2018), (c) magnetic HGM map showing magnetic anomalies, and (d) gravity HGM map visualizing density variations (Lawley et al., 2022a).

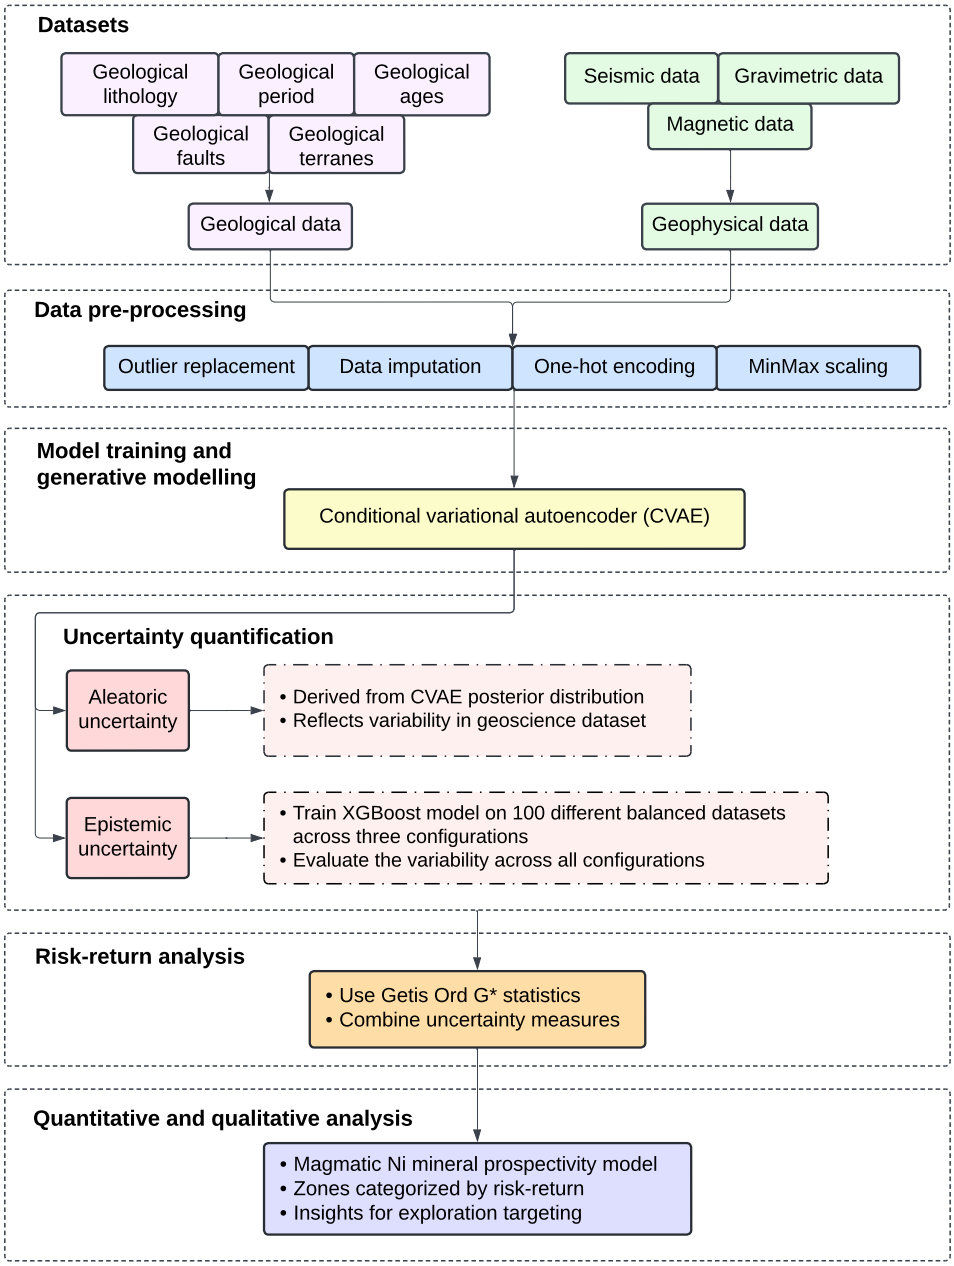


Figure 5.5 Flowchart of key steps involved in the uncertainty and risk-return analysis of prospectivity modelling.

used to replace the outliers of magnetic HGM, gravity Bouguer HGM, magnetic 1VD, gravity HGM worms, and magnetic HGM worms. Secondly, relatively few missing data points in the magnetic and gravity datasets are imputed using mean feature values. Then, categorical geological data sets, including rock types and period data sets, are mapped into a numeric array using the `OneHotEncoder` class in the scikit-learn Python library, resulting in a total

of 173 features. Finally, the training and testing data sets are scaled to a fixed range of $[0, 1]$ using scikit-learn's `MinMaxScaler` class.

5.5.2 Data augmentation and CVAE

An autoencoder is a deterministic neural network designed to learn efficient data representations by minimizing reconstruction loss. Autoencoders are effective for data compression and reconstruction tasks; however, their absence of a probabilistic framework limits their capacity to model data uncertainty and generate diverse samples. Nevertheless, their pattern recognition capabilities still make them valuable in MPM. For example, Mirzabozorg and Abedi (2023) utilized an autoencoder to identify anomaly patterns linked to specific mineralization events by training it to reconstruct input geoscience data. In this process, reconstruction errors highlighted areas that deviated from typical background patterns, which potentially indicate mineralization targets.

Kingma and Welling (2013) introduced the VAE, which incorporates a probabilistic latent space ($\mathbf{z}$) that models data distributions rather than relying on deterministic mappings. This probabilistic framework allows VAEs to learn meaningful latent representations and generate new samples by maximizing the evidence lower bound (ELBO).

Next, the CVAE extends the VAE framework by integrating specific conditional information ($\mathbf{c}$) into the encoding and decoding processes, enabling the model to generate data that reflects dependencies in the input (Sohn et al., 2015). In this study, $\mathbf{c}$ represents geological conditions, such as barren or known Ni mineral deposit data. In the CVAE framework, the input data $\mathbf{x}$, representing the geoscience feature vectors, is encoded into $\mathbf{z}$, where both the encoder and decoder networks take $\mathbf{c}$ as additional input. The CVAE employs a conditional variational inference model by introducing an approximate posterior $q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c})$ to estimate the true posterior $p_\theta(\mathbf{z} \mid \mathbf{x}, \mathbf{c})$, which is intractable due to the integration over all possible values of $\mathbf{z}$. This approximate posterior is defined as a multivariate Gaussian distribution ($\mathcal{N}$) by

$$q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c}) = \mathcal{N}(\mathbf{z}; \mu_\phi(\mathbf{x}, \mathbf{c}), \sigma^2_\phi(\mathbf{x}, \mathbf{c})), \tag{5.1}$$

where $\mu_\phi(\mathbf{x}, \mathbf{c})$ and $\sigma^2_\phi(\mathbf{x}, \mathbf{c})$ are the mean and diagonal variance outputs of the encoder network, parameterized by ϕ, respectively. The decoder then reconstructs $\mathbf{x}$ from $\mathbf{z}$, which is conditioned on $\mathbf{c}$ using the conditional likelihood $p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c})$. Here, θ denotes the parameters of the decoder network. Therefore, the generative process of the CVAE can be described in two steps: (1) sampling $\mathbf{z}$ from a prior distribution $p_\theta(\mathbf{z})$ which is typically a standard

Gaussian $\mathcal{N}(0, \mathbf{I})$, where $\mathbf{I}$ is the identity matrix, and (2) decoding $\mathbf{z}$ using $p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c})$ to generate realistic $\mathbf{x}$ samples that reflect the conditioned geological information $\mathbf{c}$.

The CVAE is optimized by maximizing the ELBO, which balances the decoder's reconstruction accuracy and the regularization of the latent space. Maximizing the ELBO is equivalent to minimizing the CVAE loss function ($\mathcal{L}_{\text{CVAE}}$) that reads

$$\mathcal{L}_{\text{CVAE}}(\mathbf{x}, \mathbf{c}; \theta, \phi) = \underbrace{\mathbb{E}_{q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c})}\Big[\log p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c})\Big]}_{\mathcal{L}_{\text{reconstruction}}} - \underbrace{D_{\text{KL}}\Big[q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c}) \,\|\, p_\theta(\mathbf{z})\Big]}_{\mathcal{L}_{\text{regularization}}}. \tag{5.2}$$

Here, the reconstruction term ensures the decoder can accurately generate $\mathbf{x}$ through $\mathbf{z}$, which is conditioned by $\mathbf{c}$. The regularization term minimizes the Kullback-Leibler divergence (D_{KL}) between the $q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c})$ and the prior $p_\theta(\mathbf{z})$, promoting a smooth and meaningful latent space.

By explicitly incorporating $\mathbf{c}$, CVAEs enable the latent space to capture both the variability of input data and the conditional dependencies inherent to geological characteristics. This capacity allows the model to effectively generate and reconstruct geoscientific data tailored to specific geological contexts, such as distinguishing barren regions from Ni mineral occurrences.

Next, the uncertainty estimation and calibration of the CVAE decoder network are further enhanced by implementing the σ-VAE approach proposed by Rybkin et al. (2021), which replaces the fixed variance assumption in the reconstruction term of the CVAE loss function with a learnable variance. Then, the Gaussian decoding distribution part becomes

$$p_\theta(\mathbf{x} \mid \mathbf{z}, \mathbf{c}) = \mathcal{N}(\mathbf{x}', \sigma^2 \mathbf{I}), \tag{5.3}$$

where $\mathbf{x}' = \mu_\theta(\mathbf{z} \mid \mathbf{c})$ represents the output of the decoder. This modification enables the decoder to represent uncertainty in data reconstruction by acquiring the optimal variance during training. Therefore, the corresponding reconstruction term in the loss function becomes the Gaussian negative log-likelihood (NLL)

$$\mathcal{L}_{\text{reconstruction}} = \frac{1}{2\sigma^2}\|\mathbf{x}' - \mathbf{x}\|^2 + J \log \sigma, \tag{5.4}$$

where $J = 173$ is the dimensionality of the data that represents input geoscience features. Building upon this, we can define the complete objective function of this σ-CVAE approach as

$$\mathcal{L}_{\sigma\text{-CVAE}}(\mathbf{x}, \mathbf{c}; \theta, \phi) = \frac{1}{2\sigma^2}\|\mathbf{x}' - \mathbf{x}\|^2 + J \log \sigma + D_{\text{KL}}\Big[q_\phi(\mathbf{z} \mid \mathbf{x}, \mathbf{c}) \,\|\, p_\theta(\mathbf{z})\Big]. \tag{5.5}$$

This modification enhances the decoder's calibration, ensuring that the D_{KL} and reconstruction loss are well-balanced in relation to the data variance, thereby improving uncertainty estimation.

This study utilizes the PyTorch Python library to develop a CVAE network featuring three hidden layers, each containing 128 nodes in both the encoder and decoder networks. We also set the latent space dimensionality to 16. All fully connected hidden layers utilize the Sigmoid Linear Unit (SiLU) activation function. In contrast, the output layers corresponding to $\mathbf{x}'$ employ a sigmoid activation function to constrain their values to [0, 1] because the input data is also normalized in that range. We train the CVAE with the Adam optimizer (Kingma and Ba, 2014), a batch size of 32 and an initial learning rate of 10^{-3}. Following training, the CVAE generates 1,489,120 positive-class samples, which are then concatenated with the existing training data to achieve a balance between the positive and negative classes, resulting in a total of 2,980,342 samples.

5.5.3 Extreme gradient boosting machines

We employ the XGBoost algorithm (Chen and Guestrin, 2016) as the classification model for MPM to differentiate between positive and negative labelled samples. XGBoost is a scalable and efficient variant of the gradient-based boosting algorithm, which is well-suited for MPM due to its ability to handle complex feature interactions (Bigdeli et al., 2024; Yin and Li, 2022). XGBoost performs classification by building an ensemble of decision trees, with each tree aimed at correcting the errors made by the previous ones. Considering our final input geoscience dataset $(\mathbf{x}_i, \mathbf{y}_i)$ where $\mathbf{x}_i$ represents the geoscience features and $\mathbf{y}_i \in \{0,1\}$ denotes the corresponding class label (positive or negative), thereby the predicted output for a sample is

$$\hat{\mathbf{y}}_i = \sum_{k=1}^{K} f_k(\mathbf{x}_i), \quad f_k \in F, \tag{5.6}$$

where $\hat{\mathbf{y}}_i$ represents the predicted result, f_k represents the k^{th} tree, F is the space of all possible classification trees and K is the total number of trees in the ensemble. Thus, the objective function L is

$$L = \sum_i l(\hat{\mathbf{y}}_i, \mathbf{y}_i) + \sum_k \Omega(f_k), \tag{5.7}$$

where $l(\hat{\mathbf{y}}_i, \mathbf{y}_i)$ is the prediction error of the i^{th} sample and $\Omega(f)$ is used to control model complexity and prevent overfitting. $\Omega(f)$ can be written as

$$\Omega(f) = \gamma T + \frac{1}{2}\lambda \|s\|^2, \tag{5.8}$$

with T being the number of leaf nodes, s represents the leaf scores, and γ and λ are the regularization parameters. Finally, the loss function can be expressed as

$$L(f_t) \approx \sum_{m=1}^{T} \left[\left(\sum_{i \in I_m} g_i \right) s_m + \frac{1}{2} \left(\sum_{i \in I_m} h_i + \lambda \right) s_m^2 \right] + \gamma T, \tag{5.9}$$

where g_i and h_i represent the first and second-order gradient statistics of the loss function, respectively, and I_m refers to the set of data instances assigned to the m^{th} leaf node.

In this study, the XGBoost model is fine-tuned using a grid search to identify optimal hyperparameters for classification. The optimal configuration uses `n_estimators` = 5000, `max_depth` = 8, `min_child_weight` = 1, `eta` = 0.1, `subsample` = 1, `colsample_bytree` = 1, and `gamma` = 0.05. The rest of the hyperparameters are retained at their default values.

5.5.4 Uncertainty quantification

5.5.4.1 Aleatoric uncertainty

After we train the CVAE model on $\mathbf{x}$ and the corresponding $\mathbf{c}$, the decoder outputs reconstructions of the input features. During the inference phase, the CVAE generates 100 realizations (n_{real}) from the $q_\phi(\mathbf{z} \,|\, \mathbf{x}, \mathbf{c})$ for each input sample $(\mathbf{x}, \mathbf{c})$, resulting in a reconstruction tensor with dimensions (N, n_{real}, d). Here, N represents the number of samples, and d denotes the 173 geoscience features, thereby effectively capturing the inherent variability in the data. We define this output variance as the aleatoric uncertainty, an approach consistent with other applications of generative models for uncertainty quantification (Rautela et al., 2025). Then, the aleatoric uncertainty ($U_{\text{aleatoric}}$) for a selected feature ($f \in \{1, \ldots, d\}$) can be determined by the variance across the posterior realizations, which reads

$$U_{\text{aleatoric}}(\mathbf{x}_f) = \frac{1}{n_{\text{real}}} \sum_{r=1}^{n_{\text{real}}} (\mathbf{x}_f^{(r)} - \overline{\mathbf{x}}_f)^2, \tag{5.10}$$

where $\mathbf{x}_f^{(r)}$ represents the r^{th} realization of feature f and $\overline{\mathbf{x}}_f$ is the mean of the posterior realizations.

5.5.4.2 Epistemic uncertainty

Using an ensemble-based approach, we quantify the epistemic uncertainty related to the final MPM. First, we train the XGBoost model on three separate balanced dataset configurations to examine their prediction variability. To clarify, we define real data as the original geoscience data; synthetic data consists of CVAE-augmented positive samples; reconstructed

data represents geoscience data that has been encoded and decoded through the CVAE. By leveraging the inherent stochastic nature of the CVAE, we generate 100 distinct datasets for each configuration. Put simply, the CVAE produces slightly different synthetic positive samples and reconstructed geoscience data during each iteration, thereby ensuring diversity across all subsequent configurations.

- **Configuration 1:** Real geoscience data combined with synthetic positive samples generated by the CVAE. This is the same configuration as Nagasingha et al. (2024).
- **Configuration 2:** Reconstructed geoscience data combined with synthetic positive samples (Fully synthetic data).
- **Configuration 3:** A hybrid dataset consisting of randomly shuffled and equally split real geoscience data and reconstructed data, along with synthetic positive samples.

Next, we train an ensemble of M=100 XGBoost models separately for each data configuration, where M is the total number of models in the ensemble. Each model in the ensemble generates a probability prediction ($p_i(\mathbf{s})$) for the positive class on a common test set ($\mathbf{s}$), where i ranges from 1 to M. For clarity, these predictions can be represented as

$$\mathbf{p} = \begin{bmatrix} p_1(\mathbf{s}_1) & p_2(\mathbf{s}_1) & \cdots & p_M(\mathbf{s}_1) \\ p_1(\mathbf{s}_2) & p_2(\mathbf{s}_2) & \cdots & p_M(\mathbf{s}_2) \\ \vdots & \vdots & \ddots & \vdots \\ p_1(\mathbf{s}_n) & p_2(\mathbf{s}_n) & \cdots & p_M(\mathbf{s}_n) \end{bmatrix},$$

where each row corresponds to a test sample (n is the number of test samples), and each column represents predictions from an individual model. To aggregate predictions for each sample $\mathbf{s}$, we compute the mean prediction $p_{\mu(\mathbf{s})}$ and the standard deviation $p_{\sigma(\mathbf{s})}$ across the ensemble, which can be expressed as

$$p_{\mu(s)} = \frac{1}{M} \sum_{i=1}^{M} p_i(\mathbf{s}), \tag{5.11}$$

$$p_{\sigma(s)} = \sqrt{\frac{1}{M} \sum_{i=1}^{M} \left(p_i(\mathbf{s}) - p_\mu(\mathbf{s})\right)^2}, \tag{5.12}$$

where the $p_\mu(\mathbf{s})$ represents the ensemble's consensus on the likelihood of a test sample being positive, while $p_\sigma(\mathbf{s})$ captures the spread of predictions that indicates model disagreement. To quantify epistemic uncertainty, we calculate the relative uncertainty (U_{relative}) metric with

$$U_{\text{relative}} = \frac{p_{\sigma(\mathbf{s})}}{p_{\mu(\mathbf{s})}}. \tag{5.13}$$

To ensure numerical stability when $p_{\mu}(\mathbf{s})$ approaches zero, we clip the denominator to a minimum threshold of 1×10^{-10} in our implementation.

5.5.5 Risk-return analysis

We developed a novel risk-return analysis framework that integrates predictive spatial modelling with quantified uncertainty metrics. We denote the risk as the U_{relative} metric that reflects confidence in the results, while the return is characterized by spatial clusters of high mineralization probability, quantified using the non-parametric Getis-Ord G^* statistic (Getis and Ord, 1992). The G^* identifies spatial clusters by quantifying the clustering intensity of mineralization probabilities in relation to their spatial context, which is obtained by

$$G_i^* = \frac{\sum_j w_{ij} p_j}{\sum_j p_j}, \tag{5.14}$$

where p_j represents the predicted mineralization probability at neighbour j, and w_{ij} denotes the spatial weight assigned to neighbour j at location i. We use the following different functions in the PySAL Python library (Rey and Anselin, 2007) to conduct the spatial clustering analysis. Firstly, the weights are computed using the `Queen` contiguity method. This defines neighbours based on shared boundaries in the spatial arrangement of geospatial data, ensuring that spatial relationships are effectively captured for each polygon. Next, the isolated polygons, also known as islands, are handled using the `attach_islands` function in conjunction with the k-nearest neighbours function. Then, the G^* is implemented through the `G_local` function from the PySAL library's exploratory spatial data analysis subpackage.

After calculating the G^* statistic for each polygon, the scores and the relative uncertainty values are rescaled using min-max normalization to fall within the range [0, 1]. This allows for a consistent comparison across various geospatial locations. Next, thresholds for risk and return are determined using the centroids of the normalized distributions. Specifically, the return threshold, G^*_{centroid}, is the mean of the normalized G^* scores, while the risk threshold, U_{centroid}, is the mean of the normalized U_{relative} values. These thresholds are then used to categorize each polygon into one of four distinct risk-return groups:

- Low risk - high return: $U_{\text{relative}} < U_{\text{centroid}}$ and $G^* \geq G^*_{\text{centroid}}$
- High risk - high return: $U_{\text{relative}} \geq U_{\text{centroid}}$ and $G^* \geq G^*_{\text{centroid}}$

- High risk - low return: $U_{\text{relative}} \geq U_{\text{centroid}}$ and $G^* < G^*_{\text{centroid}}$
- Low risk - low return: $U_{\text{relative}} < U_{\text{centroid}}$ and $G^* < G^*_{\text{centroid}}$

Finally, we validate the results of our proposed risk-return framework by comparing the overlap between high-return zones with known mineral deposits.

5.5.6 Evaluation metrics

The downstream utility of the CVAE-generated data is evaluated across all three data configurations by training a unique model for each of the 100 datasets and assessing their performance on a common, unseen test set. After that, to assess the performance of the MPM predictions, several tools have been developed over the years, including receiver operating characteristics (ROC) curves (Nykänen et al., 2015), normalized density (Mihalasky and Bonham-Carter, 2001), posterior probability values sampled at training point locations (Nykänen and Salmirinne, 2007), success-rate and prediction-rate curves (Agterberg and Bonham-Carter, 2005), and prediction-area plots (Yousefi and Carranza, 2015). In this study, we utilized the ROC curves to evaluate the performance across three data configurations. The ROC curves are assessed by plotting the true positive rate (sensitivity) against the false positive rate (1 - specificity). This technique facilitates the evaluation of the model's ability to distinguish between positive and negative classes. The area under the curve (AUC) serves as a summary metric for evaluating overall classification performance, with higher AUC values indicating superior predictive accuracy. Furthermore, a curve close to the diagonal line indicates performance equivalent to random chance, corresponding to an AUC value of 0.5.

5.6 Results

5.6.1 Model comparison and Ni MPM

Figure 5.6 presents the ROC-AUC analysis results for the three configurations. The first configuration, which incorporates original geoscience data with CVAE-generated synthetic positive samples, achieved an AUC value of 0.94 by demonstrating better predictive capability. The second configuration, which employs reconstructed geoscience data alongside synthetic positive class data, yielded a slightly lower AUC value of 0.88, indicating a decrease in model accuracy when real geoscience data is replaced with reconstructed data. The third configuration, a hybrid model that integrates real, reconstructed, and synthetic positive class data, yielded an AUC value of 0.92.

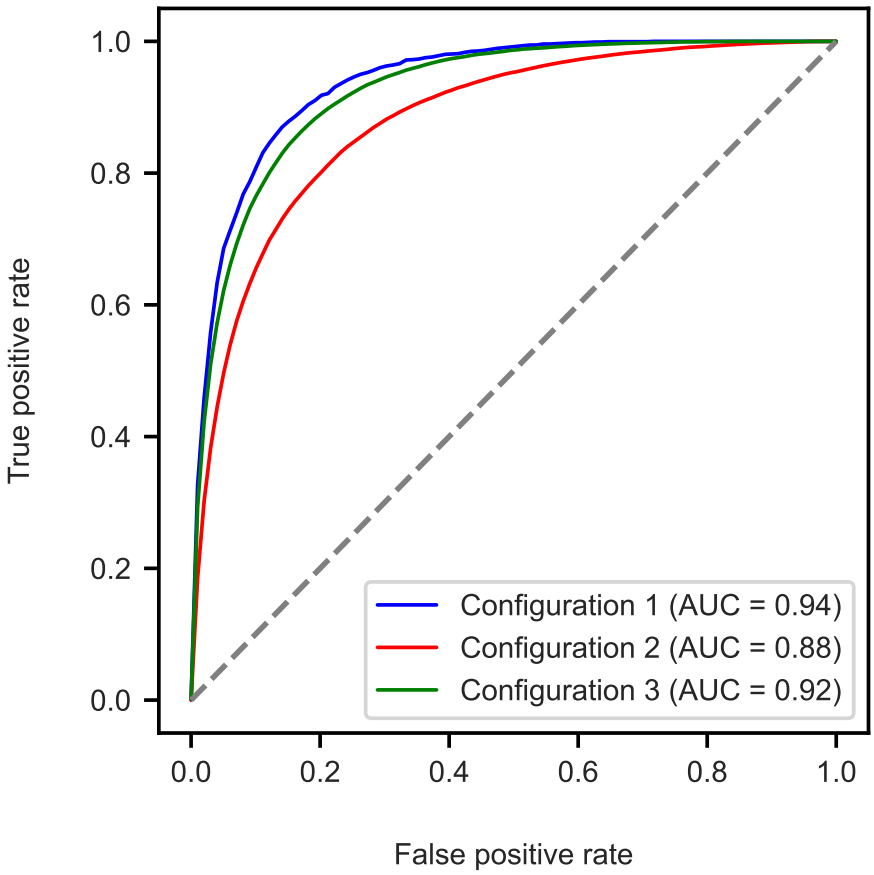


Figure 5.6 ROC curves for the testing dataset that indicate the classification performance of MPMs across three dataset configurations: (1) real geoscience data with CVAE-generated synthetic positive samples, (2) reconstructed data with synthetic positive samples, and (3) a hybrid dataset combining real, reconstructed, and synthetic positive samples.

To further illustrate the predictive capability of the best-performing configuration, we present the Ni MPM generated using the predictions from the first configuration in Figure 5.7, which achieved the highest ROC-AUC value. This map demonstrates the model's ability to predict areas of interest, highlighting regions with significant mineral prospectivity that align with known mineralized zones and revealing new, underexplored areas with potential for mineralization. Importantly, the spatial patterns observed in this prospectivity map are consistent with the results of our previous publication (Nagasingha et al., 2024), which utilized real geoscience data alongside a single synthetic positive class dataset.

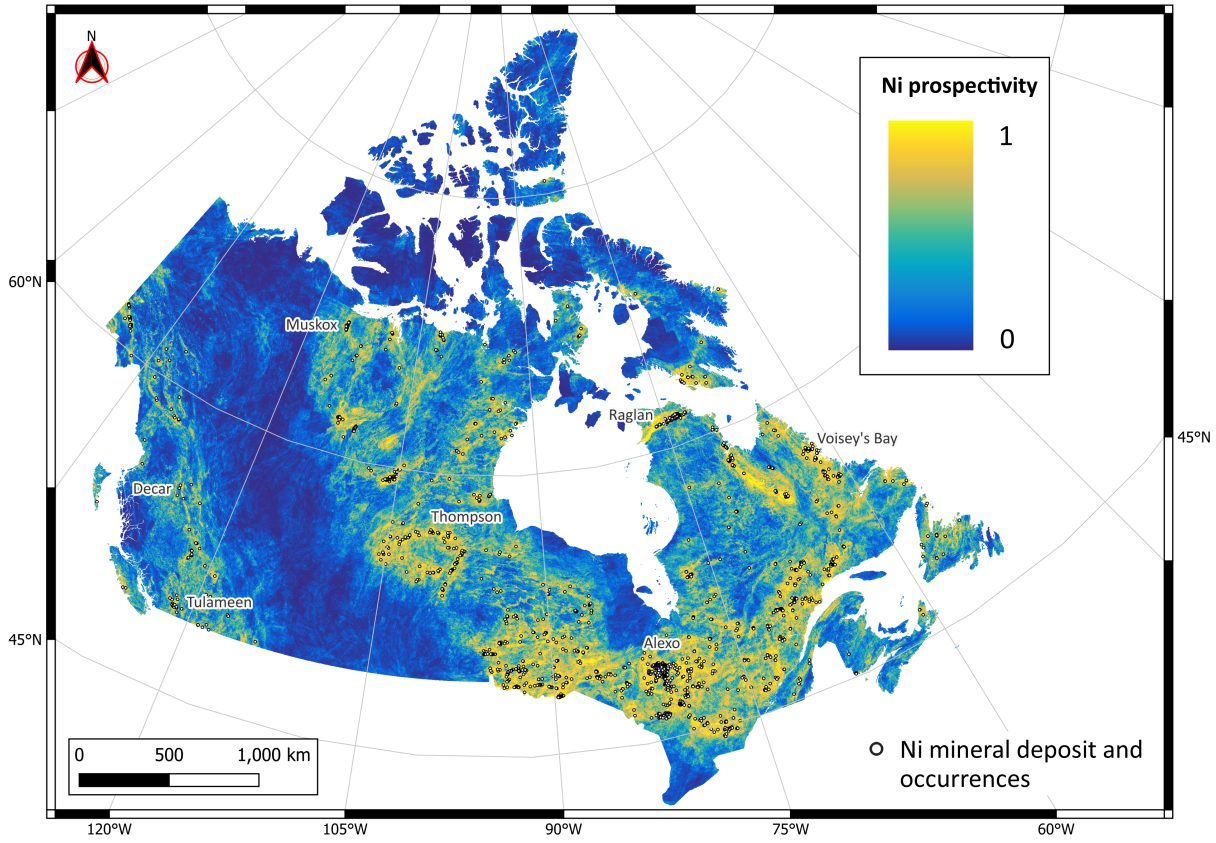


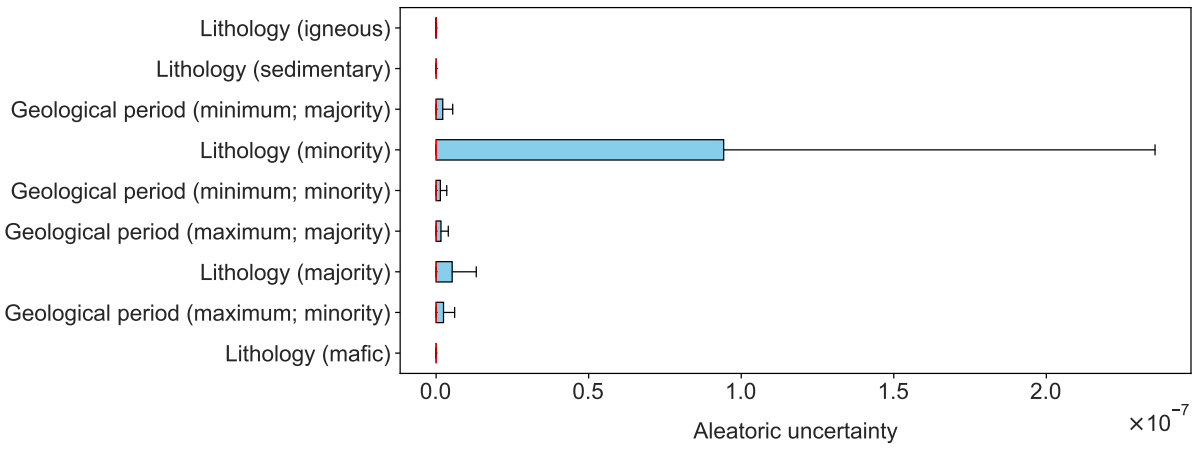
Figure 5.7 Modelled mineral prospectivity across Canada based on the results from configuration 1. The prospectivity values are quantified on a scale from 0 to 1, where high-prospectivity zones, indicating areas with a higher likelihood of mineralization, are closer to 1. Low-prospectivity zones are closer to 0.

5.6.2 Uncertainty quantification

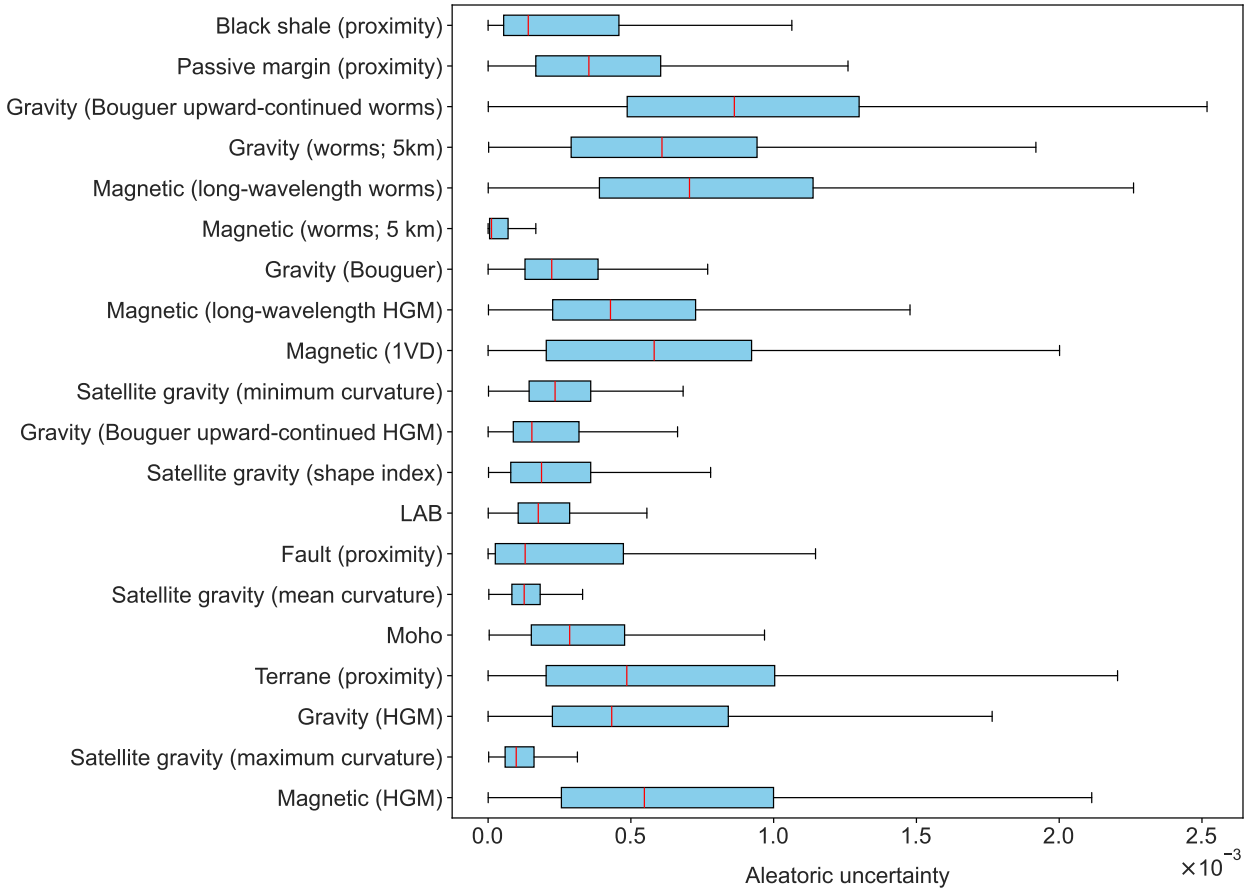
5.6.2.1 Aleatoric uncertainty

The box plots in Figure 5.8 summarize the aleatoric uncertainty values for all the scaled input features, highlighting the inherent variability and noise associated with each geoscience feature.

Figure 5.8a represents categorical geology attributes that display exceptionally low aleatoric uncertainty values, approximately 10^{-7}. We calculate the uncertainty of these features by identifying the maximum variance within specified feature ranges and assigning the highest value to each group. On the other hand, the majority of the geophysical datasets in Figure 5.8b exhibit significantly higher aleatoric uncertainty, which varies around 10^{-3}.



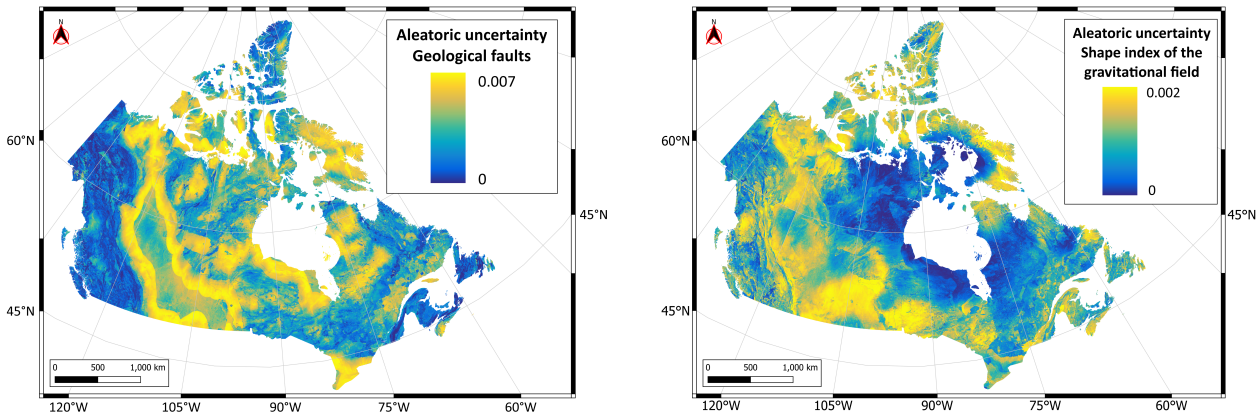
(a)



(b)

Figure 5.8 Box plots of the distribution of aleatoric uncertainty in the scaled input geoscience features. Plot (a) represents one-hot encoded features that include categorical geology data, while plot (b) comprises the rest of the non-encoded geoscience features. The red lines denote the median uncertainty, while the blue boxes illustrate the interquartile range, with whiskers representing the overall spread of uncertainty.

Figure 5.9 illustrates the aleatoric uncertainties of geological faults and the satellite gravity shape index datasets. Compared to the real data presented in Figure 5.4, areas with high geological fault values consistently exhibit low aleatoric uncertainty (Figure 5.9a). In contrast, for the satellite gravity shape index data, both high and low feature values correspond to low uncertainty (Figure 5.9b). This behaviour suggests that the aleatoric uncertainty model is more confident in areas with well-defined or extreme signal characteristics. Conversely, areas with intermediate feature values show higher uncertainty. This pattern likely arises due to the complexity or variability of geological characteristics.



(a) (b)

Figure 5.9 Spatial distribution maps of aleatoric uncertainty across Canada for two geoscience features: (a) geological faults and (b) the shape index of the gravitational field. Regions with higher uncertainty indicate areas where input features display greater variability.

5.6.2.2 Epistemic uncertainty

Figure 5.10 illustrates the spatial distribution of epistemic uncertainty, highlighting areas where the MPM predictions show the greatest confidence as well as regions with higher uncertainty. Notably, we observe distinct low-uncertainty patterns corresponding to known deposit locations, while new prospective areas exhibit both high and low uncertainty values in different locations. For example, the northwestern region of Churchill Province shows a high uncertainty pattern despite being recognized as a promising new area for exploration, as indicated in Figure 5.7. On the other hand, negative class-labelled regions exhibit moderate uncertainty across most areas.

To further support this qualitative representation, we performed a kernel density estimation (KDE) analysis using the `KernelDensity` class in the scikit-learn Python library, with a

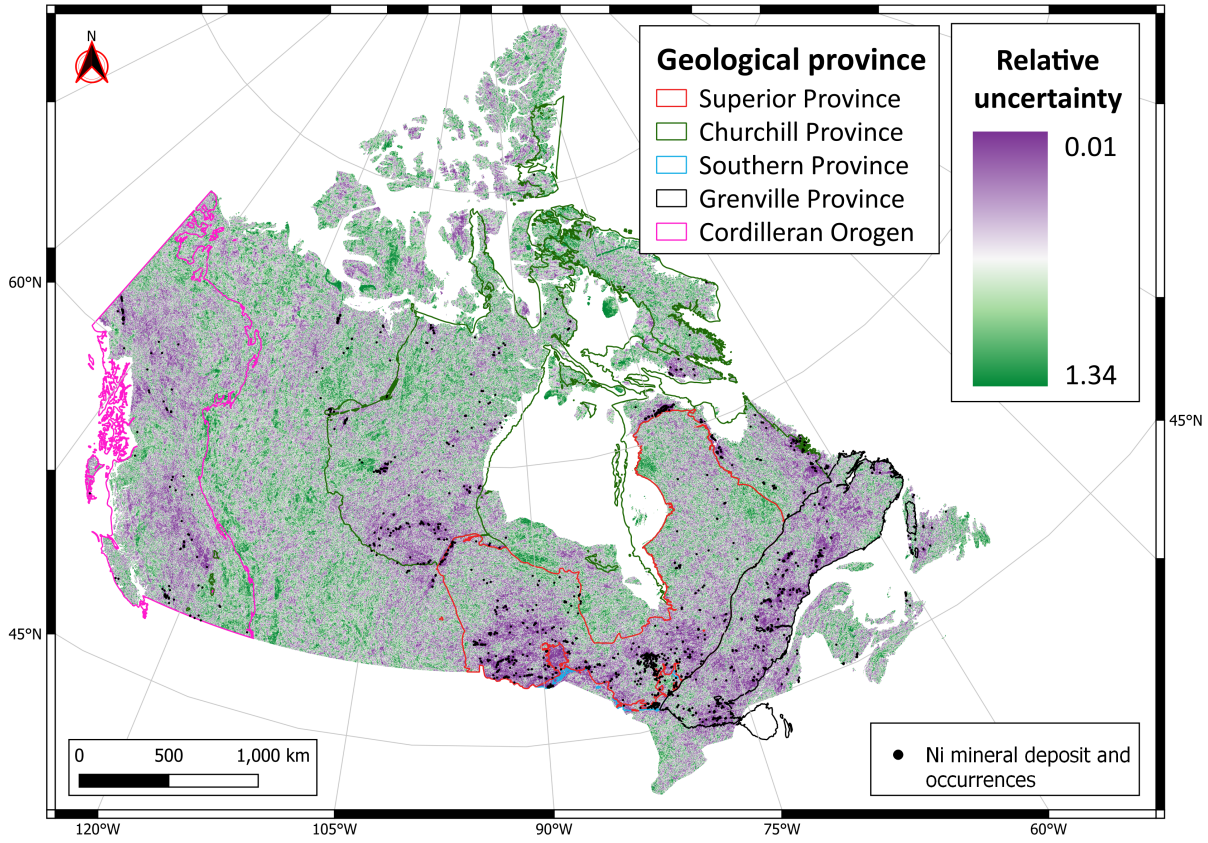


Figure 5.10 Spatial distribution of relative uncertainty in mineral prospectivity across Canada. The map also shows five key geological province boundaries that contain over 90 % of Canada's known Ni deposits.

bandwidth value of 0.5, to investigate the density distribution of relative uncertainty for positive and negative class predictions across major geological provinces in Canada. The KDE analysis is conducted on the testing set, which comprises 396 known Ni deposits out of a total of 1452. The uncertainty versus density distributions for the Superior, Churchill, and Southern provinces are illustrated in Figure 5.11. According to the KDE analysis, we observed that the density values of uncertainties in negative predictions remain relatively consistent across both the training and testing datasets, whereas positive predictions exhibit higher uncertainty in the testing dataset compared to the training dataset across all provinces. As shown in Figure 5.11, among the positive class predictions, the Superior Province shows the lowest peak uncertainty, while the Southern Province records the lowest peak uncertainty for negative class predictions. Notably, the Churchill Province shows a broader KDE distribution in the positive class compared to the other two provinces, particularly on the side of low uncertainty.

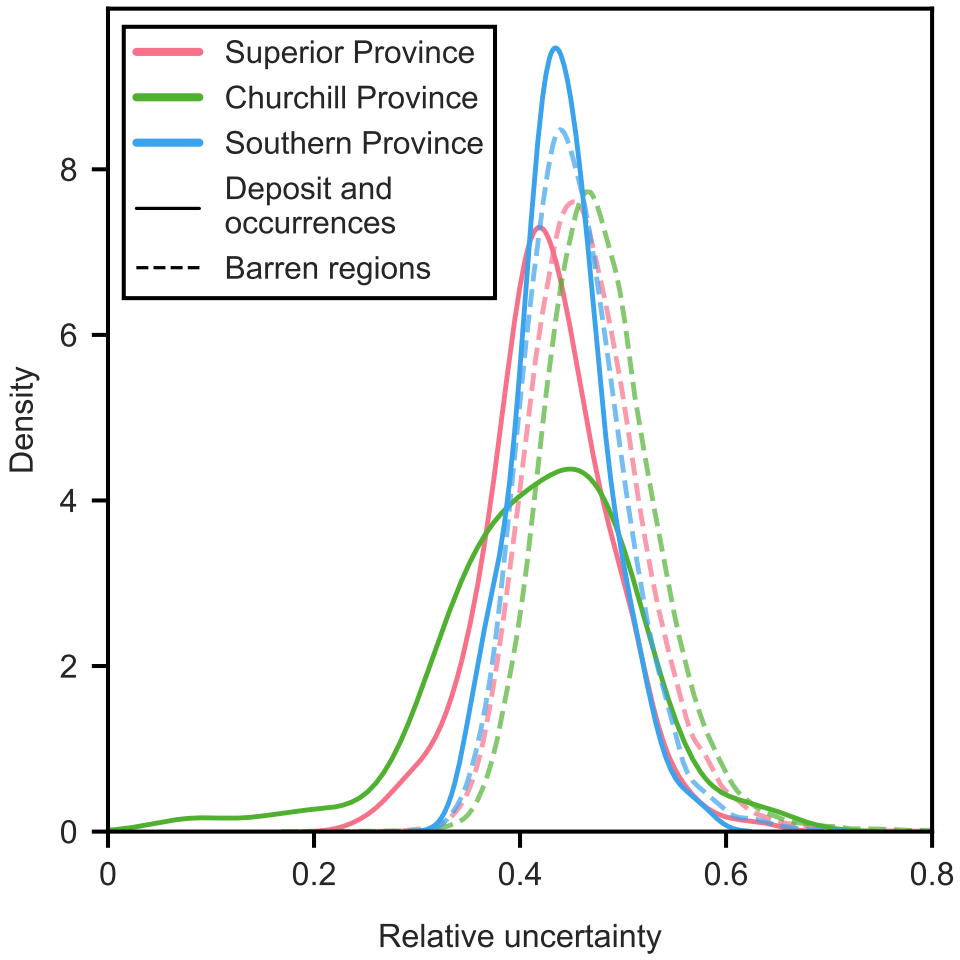


Figure 5.11 KDE plots of the relative uncertainty distributions in the testing dataset for positive and negative regions across major geological provinces in Canada. The density signifies the concentration of relative uncertainty values, whereas peaks denote the most frequent uncertainty levels. For clarity and because they contain more than 88 % of known deposits, only the Superior, Churchill, and Southern provinces are shown.

5.6.3 Risk-return analysis

Building on our new risk-return approach, we constructed a map that effectively depicts the trade-off between exploration uncertainty and mineral prospectivity across the study area. The risk-return map in Figure 5.12 represents areas of high prospectivity coupled with varying levels of uncertainty. As shown in Figure 5.12a, the majority of known Ni deposits are strongly aligned with low-risk-high-return zones. Figure 5.12b and Figure 5.12c provide detailed views of specific regions within the Churchill and Superior Provinces, respectively. These zoomed-in perspectives illustrate the distribution of known deposits and occurrences,

with a significant concentration in high-return categories.

Additionally, new prospective regions have been identified in both low-risk-high-return and high-risk-high-return categories. For example, in the Churchill Province, the new prospective regions are observed within the high-risk-high-return category, entirely separate from the outcrops of known deposits. This separation indicates untapped potential that could yield significant discoveries; however, it requires careful exploration strategies to address the associated uncertainties. Furthermore, the Superior Province offers new prospective zones within the low-risk-high-return category. The identified zones are located between established deposit sites within the New Québec Orogen (Labrador Trough), aligning with favourable geological conditions for mineralization (Nagasingha et al., 2024; Smith et al., 2021).

Next, Figure 5.13 presents the quantitative comparison of the distribution of area and known deposit percentages across the four risk-return categories. High-return zones, which encompass both high-risk and low-risk classifications, collectively represent approximately 4 % of the total area and contain 94.7 % of the known deposits. On the other hand, low-return areas cover 96 % of the total area but contain only 5.3 % of the known deposits.

5.7 Discussion

5.7.1 Geological interpretation of the results

The performance difference between the model configurations clearly demonstrates the impact of real geoscience data on predictive accuracy, as Configuration 1 achieved the highest performance (AUC = 0.94). On the other hand, Configuration 2, which utilized fully synthetic data, yielded the lowest accuracy (AUC = 0.88), likely due to the introduction of noise and the limitations of the reconstruction process in replicating complex real-world geological patterns perfectly. Nevertheless, the robust performance of both the reconstructed and hybrid models highlights their potential utility, especially in contexts where real data is limited or incomplete, which can make them valuable tools for expanding MPM with certain precision trade-offs.

A key challenge in large-scale prospectivity modelling is that differing data resolutions across regions can heighten the uncertainty associated with geoscience features, a problem compounded by the preprocessing steps involved in integrating original datasets with different resolutions using the H3 DGGS technique. This is clearly reflected in the uncertainty analysis, which reveals a dichotomy between the categorical geological data and the geophysical datasets. For example, the low uncertainty in geological features, such as lithology and geological period, stems from their foundation in well-established, expert-interpreted frameworks

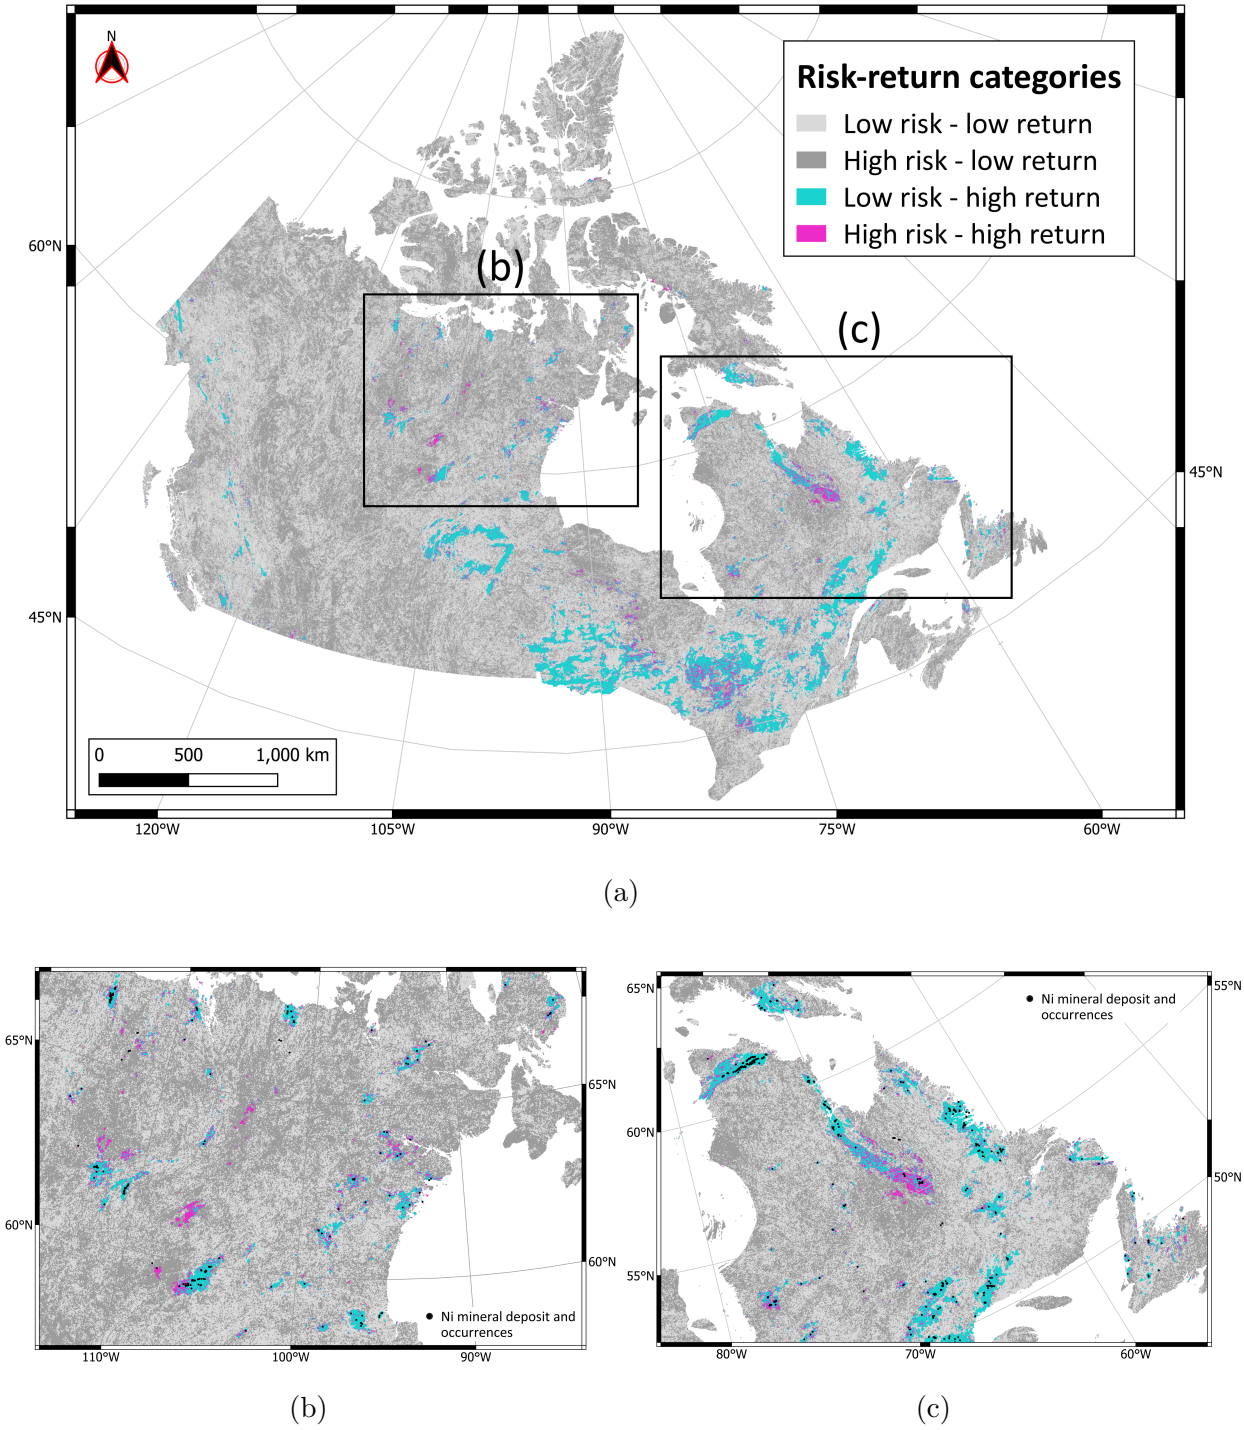


Figure 5.12 Map (a) depicts the spatial distribution of risk-return categories across Canada. Subplot (b) provides a zoomed-in view of a region within the Churchill Province, whereas subplot (c) focuses on a region within the Superior Province. Known Ni deposits and occurrences are included in the subplots as reference points to demonstrate their alignment with the risk-return categories.

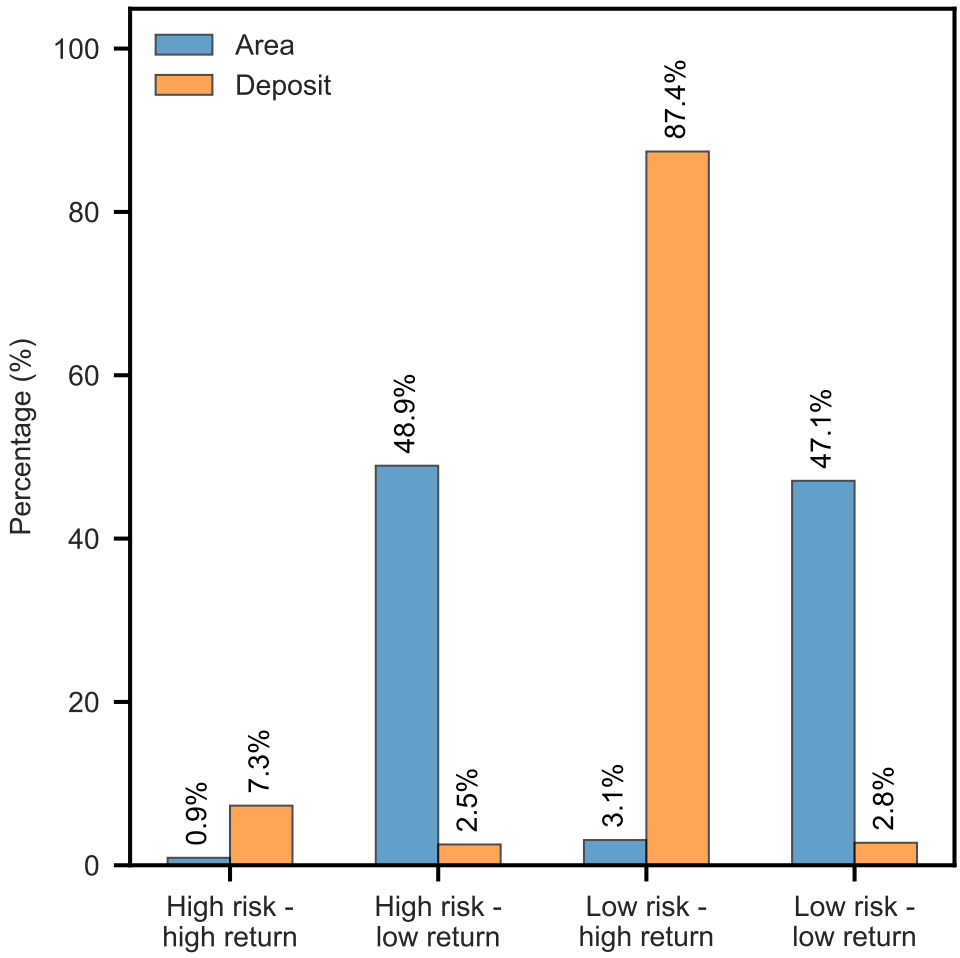


Figure 5.13 Bar chart comparing the area and known deposit percentages across risk-return categories. The percentages are calculated based on the proportion of the total area (blue bars) and the total number of known deposits (orange bars, out of a total of 1,452) in each category, respectively.

based on decades of field observations. In contrast, the higher aleatoric uncertainty inherent in the geophysical datasets reflects these challenges. Derivative products like magnetic worms and the HGM, for example, are designed to enhance subtle subsurface features, but this sensitivity also makes them susceptible to environmental noise, data resolution limitations, and ambiguity in interpretation (Hábel et al., 2020; Saibi et al., 2022). Beyond this data-inherent uncertainty, the KDE analysis of epistemic uncertainty reveals a broad distribution for positive predictions in the Churchill Province (Figure 5.11). This pattern can be attributed to the high concentration of world-class deposits within specific areas, such as the Raglan region, where the model has learned the distinct geological patterns with high confidence. For prospective areas outside of these well-defined camps, the model's confidence

decreases, leading to higher uncertainty. Understanding this behaviour is crucial, as the risk-return analysis clarifies the trade-offs associated with exploring these highly prospective but variably certain regions.

5.7.2 Advancements and implications for exploration

So far, several risk-return analysis frameworks have been introduced in the MPM domain. Wang et al. (2020) established definitions for return and risk through the application of the mean and standard deviation of log-odd ratios, which is successfully applied in MPM studies (Parsa et al., 2024; Zuo and Wang, 2020). Next, Parsa and Carranza (2021) used predictive model values to define return and risk. Overall, our proposed framework presents several key advancements over existing methods. First, training 100 distinct dataset models and aggregating their mean outputs effectively captures the stochastic uncertainty inherent in mineral predictions. Second, the integration of non-parametric spatial clustering analysis, using the Getis-Ord G^* statistic, enables us to capture localized clusters of high mineral potential effectively. Third, combining this with relative uncertainty enhances the interpretability and robustness of the risk assessment. Additionally, our implementation of k-nearest neighbours analysis for isolated polygons ensures spatial coherence across the total study area.

Furthermore, by employing the classical risk equation: $Risk = Impact \times Probability$ (Krisper, 2021), we can implement a real quantitative perspective on risk. In the MPM, we interpret the impact as the potential cost of false-positive exploration, where a company might invest significant resources in non-viable targets, while probability can be defined as relative uncertainty. Therefore, by prioritizing high-return zones and limiting exploration in low-return areas, our approach can significantly mitigate risks and improve the success of both financial investment and exploration.

5.7.3 Limitations of the proposed framework

The proposed framework, while robust, has several methodological limitations that present clear avenues for future research. Firstly, the ensemble-based approach is computationally intensive, as it requires generating and training 100 separate models. Secondly, the relative uncertainty map, as a direct consequence of the stochastic nature of the CVAE-driven ensemble, can sometimes appear spatially incoherent. Future work could explore alternative uncertainty quantification methods, such as Monte Carlo dropout or Bayesian neural networks, which are more computationally efficient but may also yield more spatially coherent uncertainty maps. Finally, the risk-return analysis relies on key methodological choices. The results are sensitive to both the thresholds used to define the four quadrants (here, the mean)

and the spatial weights matrix used for the G^* statistic (here, Queen contiguity). A valuable next step would be to conduct a sensitivity analysis, testing alternative thresholds and spatial weighting schemes to assess the robustness of the identified prospective zones.

Most importantly, a fundamental challenge in national-scale prospectivity modelling is the inherent mismatch between the coarse-resolution regional data used as input and the small, localized ore-hosting intrusions, which are frequently too small to be resolved on these maps. This challenge is particularly acute when mapping trapping mechanisms. Indeed, many of the specific physical and chemical mechanisms that accumulate sulphide liquid are impossible to identify with available datasets (Lawley et al., 2022b). This overall scale-mismatch problem, therefore, can be one of the primary drivers of uncertainty in prospectivity modelling. Future improvements in targeting depend on developing new, high-resolution geophysical and geochemical datasets, which are only available at the district- and deposit-scale and which could serve as better proxies for physical and chemical trapping mechanisms.

5.8 Conclusions

This study introduces a CVAE-based uncertainty quantification framework for Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide MPM. We evaluate and distinguish aleatoric and epistemic uncertainty through generative modelling by using CVAE's probabilistic characteristics and decoder calibration capabilities. Our aleatoric uncertainty analysis reveals that geological period and lithological data exhibit lower uncertainty compared to geophysical data. Next, through three different balanced dataset configurations, we find that the combination of real data and an augmented positive class dataset is superior, achieving a ROC-AUC value of 0.94, which is further validated by qualitative analysis. Using an ensemble of models, we quantify epistemic uncertainty as a relative measure for mineral predictions, with spatial analysis and KDE confirming that most known deposits align with low-uncertainty zones. Moreover, we combine uncertainty measures with the Getis-Ord G^* spatial clustering technique to develop an innovative risk-return framework for prospectivity modelling, enhancing decision-making tools. Our findings demonstrate the effectiveness of this approach by classifying outputs into four risk-return categories. High-return zones, comprising approximately 4 % of the total area, contain 94.7 % of the known deposits, demonstrating the framework's efficiency in prioritizing prospective targets and potentially reducing unnecessary exploration efforts. Overall, this national-scale MPM methodology advances current practices, aids in exploration decision-making and offers practical advantages for real-world mineral exploration. Future work could incorporate high-resolution geoscience data to resolve depositional-scale trapping mechanisms and utilize borehole data to identify true negatives and false positives

more accurately.

Acknowledgements

C. L. Bérubé and L. M. A. Nagasingha acknowledge funding from Natural Resources Canada [TGI project GC-130021S] and the support of the Natural Sciences and Engineering Research Council of Canada [reference number 2024-06161]. The authors thank the Editor and reviewers for their constructive criticism of the manuscript.

CHAPTER 6 ARTICLE 3: TRACTABLE APPROXIMATE GAUSSIAN INFERENCE FOR INTERPRETABLE MINERAL PROSPECTIVITY MODELLING

Applied Computing and Geosciences
Volume 29, 2 March 2026, 100332

Lahiru M.A. Nagasingha[a], Charles L. Bérubé[a], Reza Ghanati[b]

[a] Civil, Geological and Mining Engineering Department, Polytechnique Montréal, Montréal, H3C 3A7, QC, Canada

[b] Department of Earth Physics, Institute of Geophysics, University of Tehran, 14359-44411, Tehran, Iran

Authorship contribution statement

Lahiru M.A. Nagasingha: Writing – original draft, review & editing, Literature research, Experimentation, Validation, Software, Methodology, Formal analysis, Visualization, Data curation, Conceptualization. **Charles L. Bérubé:** Writing – review & editing, Supervision, Validation, Resources, Investigation, Funding acquisition. **Reza Ghanati:** Writing – review & editing.

Abstract

Bayesian deep learning methods are powerful tools for mineral prospectivity modelling (MPM) thanks to their ability to quantify uncertainty. However, their reliance on stochastic gradient-based optimization or iterative variational updates presents challenges, and tools for interpreting the sources of MPM uncertainty remain limited. We address these challenges by applying tractable approximate Gaussian inference (TAGI), an analytical Bayesian method, to model the prospectivity of magmatic Ni (±Cu ±Co ±Platinum group elements) sulphide mineral systems across Canada. First, we use the adaptive synthetic sampling technique to handle class imbalance, and second, we benchmark the TAGI approach against a Bayesian neural network (BNN) with variational inference and a Monte Carlo (MC) dropout model. Comprehensive evaluations of performance, computational cost, and robustness show that TAGI outperforms the other Bayesian approaches, achieving an area under the receiver operating characteristic curve value of 0.91 and recall of 0.70 compared to 0.88 and 0.23 for BNN and 0.87 and 0.22 for MC dropout. Additionally, TAGI also trains 1.4 times faster per

epoch and demonstrates 39 % less performance degradation under added noise compared to the second-best performing method. Furthermore, the TAGI MPM results capture 85 % of known test deposits within just 15 % of the most prospective land area. Finally, Shapley additive explanations analysis confirms that the TAGI-based model learns geologically sensible relationships and reveals that lithology and geological period data are the dominant drivers of both its prospectivity predictions and its associated uncertainty. This result highlights that, when integrated into an MPM framework, TAGI provides not only a powerful predictive tool but also an effective decision-support system, as it quantifies both prospectivity and model reliability across geological domains.

6.1 Introduction

Numerous countries have established lists of critical minerals to recognize their indispensable role in the global transition to green energy and national technological progress. However, supply chains for these materials, e.g., nickel, copper, cobalt, and platinum-group elements (PGE), face significant risks. These risks arise both from the depletion of near-surface deposits, which forces exploration toward deeper and more complex geological settings, and from an over-reliance on single mineral commodities that currently lack viable substitutes. A promising approach to mitigate these risks and guide new ore deposit discoveries is the regional-to-local exploration-scale adoption of mineral prospectivity modelling (MPM) techniques by the mining industry (Bonham-Carter, 1994; Zuo, 2020).

Despite the advancements in machine learning techniques over the past few decades, the industry implementation of these algorithms in MPM is hindered by several challenges. One challenge is the inherent class imbalance due to mineral deposits being a rare geological phenomenon (Cheng, 2007; Xiong and Zuo, 2017). Consequently, training datasets for a given area will typically be dominated by barren examples (the negative class) and feature only a few known deposits (the positive class). Studies have confirmed that finding an appropriate balance between positive and negative datasets is crucial for producing reliable exploration targets (Rahimi et al., 2021). So far, class imbalance in MPM has been addressed from two main perspectives: (1) augmenting the minority class using generative techniques such as conditional variational autoencoders (CVAE, Nagasingha et al., 2024), generative adversarial networks (GAN, Farahbakhsh et al., 2023), and the synthetic minority oversampling technique (SMOTE, Prado et al., 2020), or (2) positive-unlabelled learning, a technique that treats barren locations as unlabelled rather than negative (Hajihosseinlou et al., 2025; Xiong and Zuo, 2021).

Beyond data imbalance, many classifiers also lack uncertainty quantification, which is cru-

cial for determining model confidence. Although accurate classification is the primary goal of MPM, a simple binary prediction of high prospectivity versus low prospectivity is often insufficient for informed decision-making. To date, various uncertainty quantification techniques have been implemented in MPM, including generative ensemble modelling-based uncertainty quantification (Nagasingha et al., 2025), majority voting and confidence index-based ensemble techniques (Yousefi et al., 2024a), and fuzzy logic-based Monte Carlo (MC) simulations (Lisitsin et al., 2014). Moreover, the application of machine learning-based MPM in industry remains limited due to the apparent black box nature of artificial intelligence models. This lack of transparency can lead to models that fail to generalize to new areas, or that produce results inconsistent with known geological principles (Karpatne et al., 2017; Reichstein et al., 2019; Zuo et al., 2024). In high-risk economic activities such as mineral exploration, unreliable predictions can impact major investment decisions. To overcome this, explainable artificial intelligence principles have been successfully applied to MPM case studies (Liu et al., 2025b; Mou et al., 2025; Pradhan et al., 2022). While researchers have made significant strides in explaining MPM predictions, understanding their associated uncertainties is an equally critical yet less explored avenue of research.

Uncertainty quantification in machine learning is commonly achieved through probabilistic approaches, particularly Bayesian techniques. Unlike deterministic models that produce only point estimates, Bayesian inference treats model parameters as random variables and derives their posterior distribution to quantify model uncertainty (Goodfellow et al., 2016). One of the earliest Bayesian MPM implementations is the *prospector expert system*, which relies on knowledge-driven probabilistic relationships (Campbell et al., 1982; Duda et al., 1977). Subsequent developments have shifted toward data-driven approaches, like the weights of evidence, which calculates posterior probabilities from data but operates under the assumption of conditional independence among predictor variable patterns (Agterberg and Cheng, 2002; Bonham-Carter and Agterberg, 1990). This limitation can be addressed by modelling interdependencies with Bayesian neural network (BNN) classifiers (Porwal et al., 2006a). However, because exact posterior inference in BNNs is computationally intractable, scalable approximation methods such as variational inference (VI, Barber and Bishop, 1998; Blundell et al., 2015; Hinton and Van Camp, 1993; Kingma et al., 2015; Wu et al., 2018) and dropout as a Bayesian approximation (Gal and Ghahramani, 2016) have become standard practice, including in geoscience applications (Huang et al., 2025; Jordao et al., 2023; Puzyrev and Duuring, 2025; Zuo et al., 2022).

Most modern machine learning uncertainty quantification approaches reformulate inference as an optimization problem built around gradient backpropagation. While powerful, this process can face significant challenges, including local minima, vanishing or exploding gradients in

deep neural networks, and high computational cost due to its iterative nature. To address these limitations, Goulet et al. (2021) introduces tractable approximate Gaussian inference (TAGI), an analytical Bayesian method that achieves computational efficiency by propagating means and variances through the network without relying on gradient-based backpropagation.

While the TAGI method has shown promising results on benchmark datasets, its application to complex, real-world problems remains limited (Nguyen and Goulet, 2021; Vuong et al., 2025). This study extends the use of TAGI to large-scale MPM and develops an uncertainty-driven interpretability framework by leveraging TAGI outputs. The contributions of this study are twofold: (1) we apply the TAGI method to model magmatic Ni (±Cu ±Co ±PGE) prospectivity across Canada and benchmark its performance against two Bayesian approximation techniques: a BNN trained with VI and a deep neural network using the MC dropout method; and (2) we extend the application of Shapley additive explanations (SHAP) to not only describe the model predictions but also to provide feature-based insights about their associated uncertainties. The remainder of this paper is organized as follows: Section 6.2 describes the study area and datasets used; Section 6.3 details the methodology for the three probabilistic machine learning models; Section 6.4 presents the comparative results, discusses their geological implications and underlying mechanisms, and acknowledges framework limitations; and Section 6.5 highlights the conclusions of our numerical experiments.

6.2 Datasets

The study area for this research encompasses the entire Canadian landmass. Figure 6.1 provides a simplified geological map of the study area, showing the boundaries of Canada's major geological provinces and the spatial distribution of major rock categories, as well as the locations of the magmatic Ni (±Cu ±Co ±PGE) deposits and occurrences that are the focus of this study. The majority of these known Ni deposits are concentrated within the geological provinces listed in Figure 6.1, which reflects the highly favourable geological conditions for mineralization present within these regions.

The data cube we used in this study is built upon Uber's H3 Discrete Global Grid System (Uber Technologies, Inc., 2023). The H3 system is a geospatial indexing approach that employs hexagonal tiling to cover the globe at multiple spatial resolutions. This data cube utilizes resolution 7 of the H3 grid, which corresponds to hexagons with a side length of approximately 1.22 km. This results in approximately 1.8 million unique cells and has proven to be a versatile foundation for many national-scale prospectivity models, including those for Mississippi Valley–type and clastic–dominated Zn–Pb deposits (Lawley et al., 2022a), magmatic Ni (±Cu ±Co ±PGE) sulphide deposits (Nagasingha et al., 2024), carbonatite-hosted

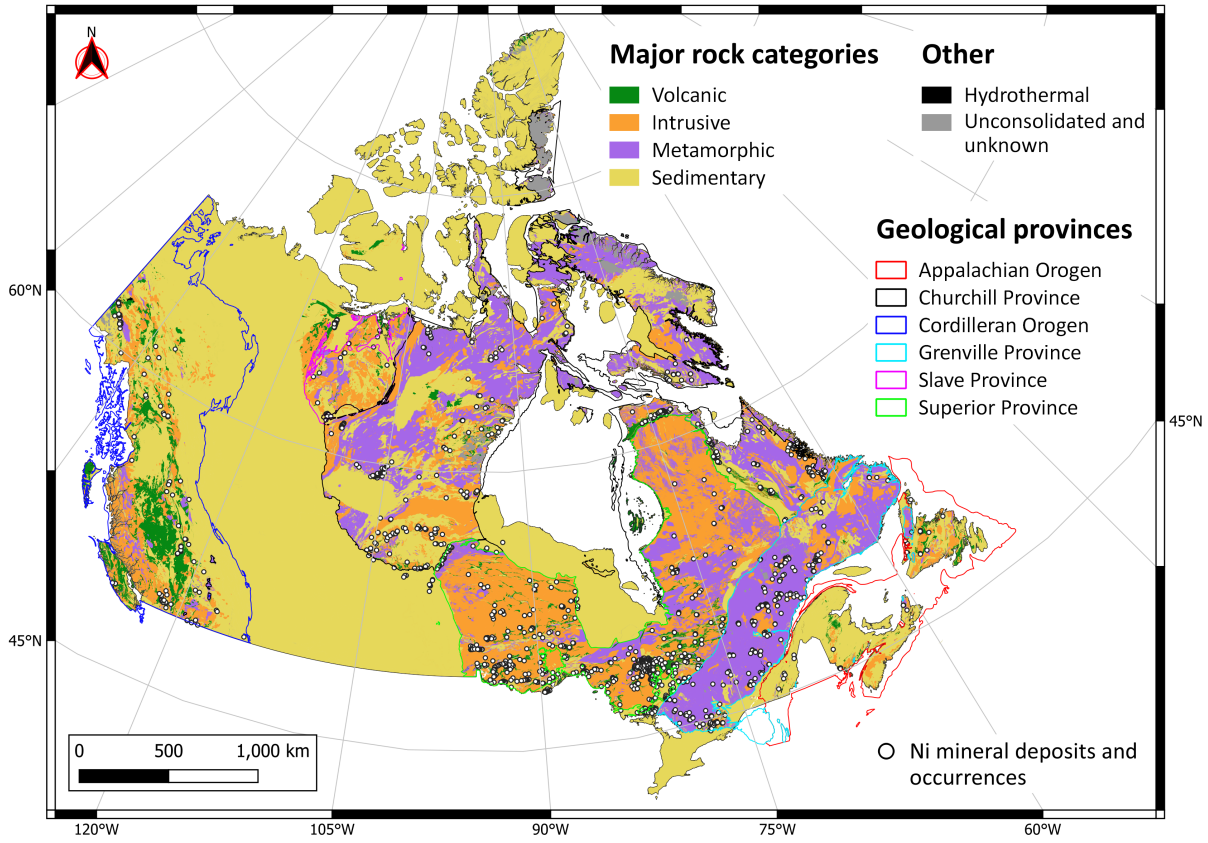


Figure 6.1 Simplified geological map showing the distribution of major rock categories across Canada. The locations of the known Ni deposits and occurrences used to train and test the models are indicated by white circles, and the boundaries of key geological provinces are highlighted to demonstrate that the majority of these deposits are situated within these specific terranes (Modified from Lawley et al. (2024)).

rare earth element deposits (Parsa et al., 2024), lithium-cesium-tantalum pegmatites (Parsa et al., 2025), and graphite deposits (Zhang et al., 2025).

Each H3 cell in the data cube is characterized by a set of geoscience evidence layers, including geochronology, geology, and geophysical layers, as detailed in Table 6.1. These features provide evidence for the key components of magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems, encompassing: (1) the drivers of mineralization, such as mantle melting events; (2) the sources of metal-rich ultramafic to mafic melts; (3) the pathways that controlled magma ascent through long-lived zones of mantle and crustal weakness; and (4) the geological traps that triggered sulphur saturation and sulphide accumulation. The reader is referred to Lawley et al. (2022a,b); Nagasingha et al. (2024) for a comprehensive description of magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems and the geoscience features that we use as

evidence layers in this study. Additionally, qualitative representations of evidence layers are included in the supplementary material (S1-S3).

Table 6.1 Summary of the geoscience evidence layers used as input features, categorized by their role (drivers, sources, pathways and traps) in the magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system

Data category	Evidence layer	Role in mineral system
Geochronology	Minimum and maximum geological ages of bedrock units	Proxy for the timing of major magmatic events (*Driver*)
Geology	Lithology classifications from bedrock maps	Identifies the presence of favourable mafic to ultramafic host rocks (*Source*)
	Proximity to carbonaceous rocks	Proxy for potential crustal sulphur contamination (*Trap*)
	Proximity to known geological faults	Represents major crustal structures that act as magma pathways (*Pathway*)
	Proximity to geological terrane boundaries	Represents deep lithospheric architecture and potential magma conduits (*Pathway*)
Geophysics	Depth to the Mohorovičić discontinuity (Moho)	Proxy for crustal thickness (*Pathway*)
	Depth to the lithosphere-asthenosphere Boundary (LAB)	Proxy for lithospheric thickness and potential magma sources (*Pathway*)
	Regional airborne magnetic products (e.g., 1st vertical derivative (1VD), Horizontal gradient magnitude (HGM), long-wavelength HGM, and worms)	Maps magnetic anomalies and structural lineaments (*Pathway*)

Continued on next page

Table 6.1 Summary of the geoscience evidence layers used as input features, categorized by their role (drivers, sources, pathways and traps) in the magmatic Ni (±Cu ±Co ±PGE) sulphide mineral system (continued)

Data category	Evidence layer	Role in mineral system
	Regional ground gravity products (e.g., HGM, upward-continued HGM, and worms)	Maps crustal density structures and potential intrusions (*Pathway*)
	Satellite gravity derivative products (e.g., curvature, shape index)	Maps long-wavelength lithospheric variations (*Pathway*)
	Bouguer gravity anomaly	Identifies dense mafic-ultramafic bodies that create favourable geometries (*Trap*)

The data cube comprises 1,820,815 data points across 29 different geoscience data features, as summarized in Table 6.1. To mitigate model performance bias caused by spatial autocorrelation in standard random splits, we divide the dataset using a six-fold spatial cross-validation approach (Valavi et al., 2018). As a result, the study area is divided into geographically distinct blocks, each measuring 500 km × 250 km in size. These blocks are then randomly distributed among one of the six folds, with the distribution stratified by the number of known deposits to ensure representativeness (Lawley et al., 2022b; Nagasingha et al., 2024). Five of these folds are designated for the training set (1,491,227 data points, including 1,056 positive samples), while the remaining fold is held out as an independent testing set (329,588 data points, containing 396 positive samples). Following Lawley et al. (2021), positive samples include deposits, advanced-stage exploration projects, and other Ni mineral occurrences, which together comprise 1,452 samples. All remaining H3 cells are labelled as negative samples, including cells containing other mineral system types and cells with no recorded mineralization. However, this negative class represents a practical limitation: the absence of documented Ni mineralization does not prove that a location is barren, as undiscovered deposits may exist at depth or in underexplored regions. While drill hole databases could improve identification of truly barren areas, such data remain largely inaccessible at the national scale.

6.3 Methodology

This study employs a probabilistic frameworks for MPM, which integrates data preprocessing, synthetic oversampling, and a comparative analysis of three Bayesian deep learning methods. Following model evaluation, the superior model is used to generate prospectivity and uncertainty maps, which are further interpreted using SHAP analysis to quantify the drivers of both prediction and uncertainty. The overall workflow adopted for this research is illustrated in Figure 6.2.

6.3.1 Data preprocessing

The total dataset undergoes key preprocessing steps, including outlier handling, missing value imputation, categorical encoding, and feature scaling. Firstly, we address outliers using Tukey's fences, where geophysical data, such as magnetic and gravity products, exhibit the highest outlier percentages, ranging from 2.34 % to 12.79 %, and are replaced with the 95^{th} percentile value. Secondly, we impute relatively few missing data points in the magnetic (6.5 %) and gravity (0.2 %) datasets using the mean feature values. Next, the original 29 features, which comprise 20 continuous and nine categorical features (e.g., lithology and geological period), are processed by transforming the categorical features into a numeric array using the `OneHotEncoder` class in the scikit-learn Python library (Pedregosa et al., 2011), thereby expanding the total feature space to 173. Finally, all features are scaled to a [0,1] range using scikit-learn's `MinMaxScaler` class, which is fit only on the training data before being applied to both the training and test sets.

6.3.2 Data augmentation

To address the MPM class imbalance, we employ the adaptive synthetic technique (ADASYN, He et al., 2008), an oversampling technique designed to generate synthetic samples for the minority positive class and overcome the bias towards the majority negative class. Crucially, the ADASYN method is applied only to the scaled training dataset after all preprocessing steps, ensuring no data leakage into the validation or testing sets. Unlike other simple oversampling techniques, such as SMOTE, ADASYN adaptively shifts the decision boundary by focusing more on generating synthetic data for minority class instances that are harder to learn. Therefore, in the context of MPM, these harder-to-learn instances often represent mineral deposits located in subtle, atypical, or isolated geologically complex environments where the distinction from negatively labelled regions is less clear.

The ADASYN procedure begins by calculating the total number of synthetic samples to

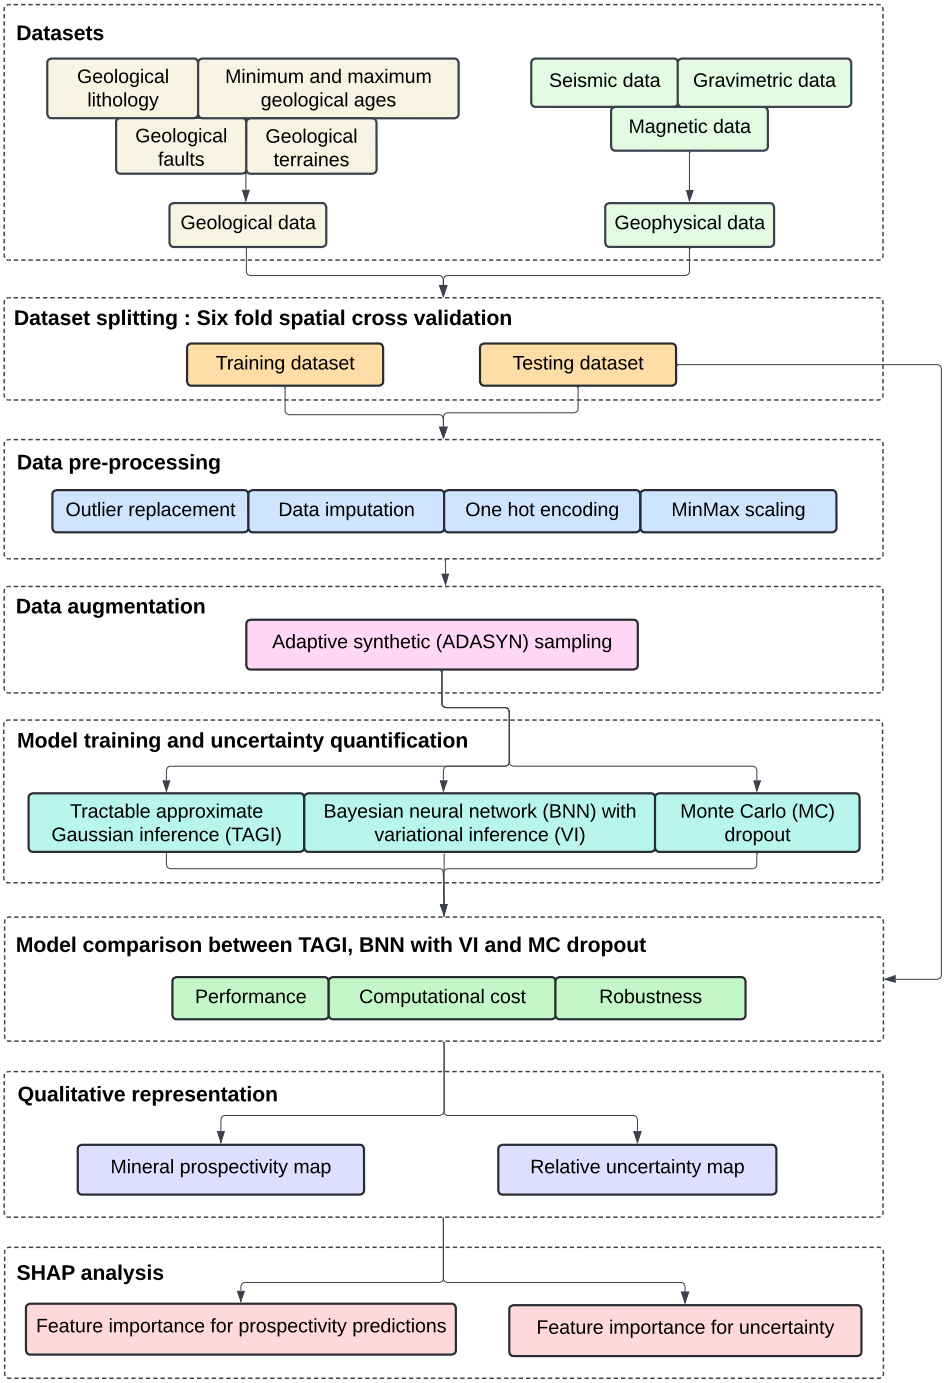


Figure 6.2 Schematic workflow of the proposed MPM framework comparing three Bayesian deep learning methods.

generate, G, which reads

$$G = \beta(N_{\text{maj}} - N_{\text{min}}), \tag{6.1}$$

where N_{maj} and N_{min} are respectively the number of samples in the majority negative and minority positive classes from the training geoscience dataset $\mathcal{D}$, and $\beta \in [0, 1]$ is the desired balance level. For each minority sample, represented by its feature vector $\mathbf{x}_i$, a difficulty ratio r_i is calculated as the proportion of majority-class samples within its k-nearest neighbours. These ratios are then normalized to form a density distribution that allocates a greater number of synthetic samples to minority instances with higher difficulty ratios. Finally, each new synthetic feature vector $\mathbf{s}_i$ is created by linearly interpolating between a minority sample and one of its randomly selected minority-class neighbours $\mathbf{x}_{ni}$, where n is the index of the neighbour, through

$$\mathbf{s}_i = \mathbf{x}_i + (\mathbf{x}_{ni} - \mathbf{x}_i) \times \lambda, \tag{6.2}$$

where λ is a random value drawn from a uniform distribution between $[0, 1]$. We have used the ADASYN implementation available in the imbalanced-learn Python library (Lemaître et al., 2017), where we set the balance level $\beta = 1$ to create a fully balanced training set and use the default value of $k = 5$ for the number of nearest neighbours.

6.3.3 Uncertainty quantification methods

We train three distinct Bayesian deep learning methods for MPM, including TAGI, a BNN with VI and MC dropout. To ensure a fair comparison, all models share an identical base architecture, consisting of an input layer, two hidden layers with 32 neurons each, using the rectified linear unit (ReLU) activation function. Furthermore, all models are trained for 100 epochs with a consistent batch size of 32. All other method-specific hyperparameters are detailed in the following subsections. All experiments are conducted on the Apple M1 chip.

6.3.3.1 Bayesian neural network with tractable approximate Gaussian inference

Unlike the gradient-based VI and MC dropout methods used in this study, TAGI is an analytical Bayesian approach that approximates the posterior distribution of a model's parameters by propagating means and covariances through the network, thereby avoiding gradient-based iterative optimization (Goulet et al., 2021). The overall architecture of the TAGI model is illustrated in Figure 6.3a. The fundamental principle of TAGI is to propagate probability distributions instead of single point values throughout the neural network.

Rather than explicitly characterizing the full joint distribution between the network parameters $\boldsymbol{\theta}$ (weights $\mathbf{W}$ and biases $\mathbf{B}$), hidden units $\mathbf{z}$, and observations $\mathbf{y}$, which would be computationally intractable, TAGI propagates only their first- and second-order statistics (means, variances, and covariances) through the network in a layer-wise manner. These local

Gaussian assumptions enable tractable analytical inference while maintaining computational efficiency.

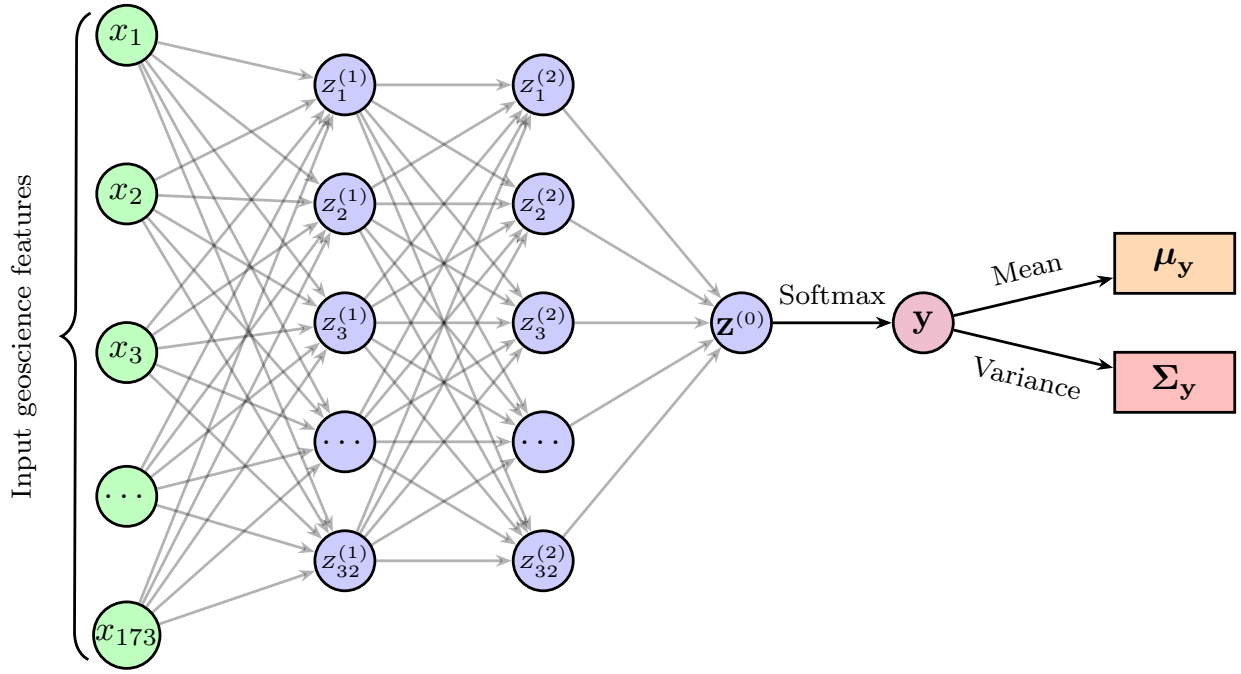
The training process using TAGI is then divided into two stages for each batch of data, as illustrated in Figure 6.3b. The forward pass constitutes the uncertainty propagation by computing the joint probability distributions between adjacent layers (e.g., $f(\mathbf{x}, \mathbf{z}^{(1)})$) while the backward pass performs the analytical inference by computing the conditional posterior distributions (e.g., $f(\mathbf{z}^{(1)}|\mathbf{y})$) for each layer and its associated parameters. For full mathematical derivation of these distributions, the reader is referred to Goulet et al. (2021).

The forward pass is applied to the preprocessed and augmented training data, where each input feature is represented as a Gaussian distribution. The transition begins with the activation units of a given layer j, which are themselves described by a multivariate Gaussian distribution $\mathbf{A}^{(j)} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{A}}^{(j)}, \boldsymbol{\Sigma}_{\mathbf{A}}^{(j)})$. After that, this distribution is propagated to the i^{th} hidden unit of the next layer $j+1$, which can be described by

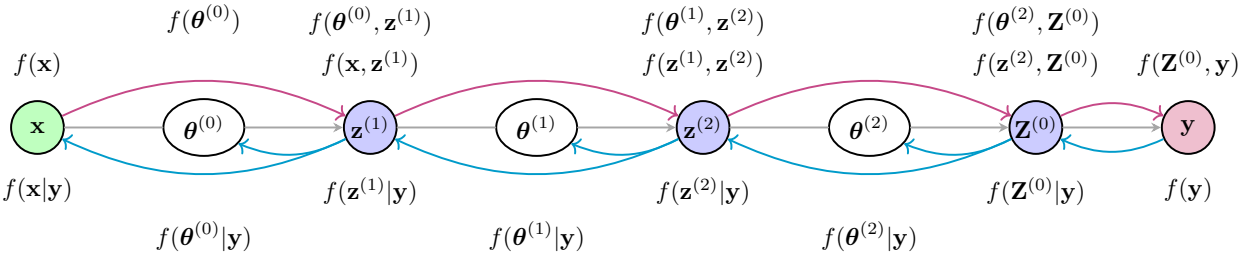
$$z_i^{(j+1)} = \sum_{k=1}^{N} W_{i,k}^{(j)} A_k^{(j)} + B_i^{(j)}, \tag{6.3}$$

where the parameters $W_{i,k}^{(j)}$ and $B_i^{(j)}$, which learn to represent the relationships between different geological evidence layers, are also treated as Gaussian random variables, and N is the number of activation units in the preceding layer. To perform this propagation analytically, TAGI employs two essential approximations. First, the Gaussian multiplication approximation (GMA) is used to handle the product of Gaussian variables, including the weight and an activation. Since this product does not produce a true Gaussian, TAGI calculates the exact first two moments of the resulting distribution analytically to approximate it as a new Gaussian using the GMA. This approximation is supported by the central limit theorem, which ensures that the sum of many such independent products converges to a Gaussian distribution. Second, local linearization is employed to tackle the non-linear activation functions used for the network. In order to maintain analytical tractability, this technique employs a 1^{st} order Taylor series approximation for each input data point. This approach is similar to the extended Kalman filter (Haykin, 2004), which facilitates the analytical propagation of mean and variance through non-linearity. Finally, during this forward pass, all layer-wise statistics, including the means, variances, and cross-covariances, are computed and stored for use in the subsequent backward pass.

Next, the backward pass analytically updates the model parameters using the statistics computed during the forward pass. This is achieved not with gradient descent, but with a recursive algorithm from state-space modelling called the Rauch-Tung-Striebel (RTS, Rauch



(a) TAGI-MPM architecture illustrating the 173 input geoscience feature vector ($\mathbf{x} = [x_1, x_2, \ldots, x_{173}]$) that are processed through two hidden layers ($\mathbf{z}^{(1)}$ and $\mathbf{z}^{(2)}$) of 32 neurons each, followed by an output layer ($\mathbf{Z}^{(0)}$) which contains logits for binary prospectivity classification (barren vs. mineral). The hierarchical softmax activation produces a single probabilistic observation output $\mathbf{y}$. From this output distribution, we can then derive the prospectivity score $\boldsymbol{\mu}_{\mathbf{y}}$ and its associated uncertainty $\boldsymbol{\Sigma}_{\mathbf{y}}$.



(b) Compact representation of the TAGI model, where $\boldsymbol{\theta} = \{\mathbf{W}, \mathbf{B}\}$ represents the parameters connecting the layers. The forward pass (magenta arrows) propagates uncertainty through the network layers, while the backward pass (cyan arrows) performs layer-wise analytical parameter inference.

Figure 6.3 Figures showing TAGI models' (a) network architecture and (b) the compact representation of the analytical information flow (Modified from Goulet et al. (2021)).

et al., 1965) smoother, which works backward from the output layer to the input layer. As illustrated in Figure 6.3b, the process begins by updating the posterior distribution of the output units $\mathbf{Z}^{(0)}$ based on the ground-truth observation vector $\mathbf{y}$. This update uses the Gaussian conditional equations, where $\mathbf{Y}$ is the random variable for the model's predictive output, with its prior mean $\boldsymbol{\mu}_{\mathbf{Y}}$, prior covariance $\boldsymbol{\Sigma}_{\mathbf{Y}}$ and $\boldsymbol{\Sigma}_{\mathbf{YZ}^{(0)}}$ is the prior cross-covariance between $\mathbf{Y}$ and $\mathbf{Z}^{(0)}$. The equations can be denoted as

$$\boldsymbol{\mu}_{\mathbf{Z}^{(0)}|\mathbf{y}} = \boldsymbol{\mu}_{\mathbf{Z}^{(0)}} + \boldsymbol{\Sigma}_{\mathbf{YZ}^{(0)}}^{\mathsf{T}} \boldsymbol{\Sigma}_{\mathbf{Y}}^{-1} \left(\mathbf{y} - \boldsymbol{\mu}_{\mathbf{Y}}\right), \tag{6.4}$$

$$\boldsymbol{\Sigma}_{\mathbf{Z}^{(0)}|\mathbf{y}} = \boldsymbol{\Sigma}_{\mathbf{Z}^{(0)}} - \boldsymbol{\Sigma}_{\mathbf{YZ}^{(0)}}^{\mathsf{T}} \boldsymbol{\Sigma}_{\mathbf{Y}}^{-1} \boldsymbol{\Sigma}_{\mathbf{YZ}^{(0)}}. \tag{6.5}$$

Since MPM is a binary classification problem, it is important to consider the choice of the output layer. A traditional softmax function does not have a closed-form analytical solution for propagating Gaussian uncertainty; therefore, to maintain analytical tractability, TAGI employs a hierarchical softmax output layer instead. Once the output layer is updated, the RTS smoother recursively propagates this new information backward through the network, layer by layer, to infer the posterior for all parameters $\boldsymbol{\theta}$ and hidden units $\mathbf{z}$. After processing a batch of data, this posterior becomes the prior for the next batch.

From the resulting trained model, the mean of the output distribution for a given location is taken as its final prospectivity score, while the variance serves as a direct, analytical measure of the model's epistemic uncertainty. To maintain a linear computational complexity, the TAGI implementation we use relies on several key assumptions, including a diagonal covariance structure for the parameters and hidden units, a layer-wise inference that uses the conditional independence between non-adjacent layers, and the assumption that the observation error variance (σ_V^2) is homoscedastic and must be set as a fixed hyperparameter (Nguyen and Goulet, 2022). Therefore, the optimal value for the primary hyperparameter σ_V^2 is selected empirically. We evaluated numerous standard deviation (σ_V) values in the range of [0, 10] based on performance on a validation set, which yielded an optimal value of $\sigma_V = 5$.

6.3.3.2 Bayesian neural networks with variational inference

BNNs treat their weights as random variables rather than fixed parameters, with each weight represented by a probability distribution. The training objective of a BNN is to learn the posterior distribution over the parameters $\boldsymbol{\theta}$, conditioned on the training dataset $\mathcal{D}$, which is denoted as $p(\boldsymbol{\theta}|\mathcal{D})$. In this study, we employ VI to approximate this posterior by introducing a tractable variational distribution, $q(\boldsymbol{\theta}|\boldsymbol{\phi})$, which is itself defined by a set of variational

Algorithm 2: TAGI model training for MPM

Data: MPM training dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{S}$, where S is the total number of training samples, $\mathbf{x}_i$ is a vector of geoscience features and y_i is its label (positive or negative);
Validation dataset $\mathcal{D}_{\text{val}}$

Result: A trained MPM model defined by the posterior parameters $(\boldsymbol{\mu}_{\boldsymbol{\theta}}^*, \boldsymbol{\Sigma}_{\boldsymbol{\theta}}^*)$

▷ `Training the TAGI prospectivity model`

Initialize: Model parameter prior $(\boldsymbol{\mu}_{\boldsymbol{\theta}}, \boldsymbol{\Sigma}_{\boldsymbol{\theta}})$ with a diagonal $\boldsymbol{\Sigma}_{\boldsymbol{\theta}}$
Observation noise standard deviation σ_V

for $e = 1$ **to** number of epochs **do**
  for each batch of locations $\mathcal{B} = \{\mathbf{x}_b, \mathbf{y}_b\} \subset \mathcal{D}$ **do**
    ▷ `1. Forward pass: Propagate feature distributions`
      ⋆ Propagate geoscience feature distributions layer-by-layer using:
        (a) The Gaussian multiplication approximation (GMA)
        (b) Local linearization of non-linear activation functions
      ⋆ Store all layer-wise statistics for the backward pass
    ▷ `2. Backward pass: Update model parameters`
      ⋆ Update the output layer using the hierarchical softmax encoding of the labels $\mathbf{y}_b$
      ⋆ Recursively update each hidden layer's parameters using the Rauch-Tung-Striebel (RTS) smoother
      ⋆ Update the model's parameter prior with the calculated batch posterior
  end
  ▷ `Monitor training and validation losses for convergence`
    ⋆ Evaluate and record losses on $\mathcal{D}$ and $\mathcal{D}_{\text{val}}$
end

return The trained model posterior $(\boldsymbol{\mu}_{\boldsymbol{\theta}}^*, \boldsymbol{\Sigma}_{\boldsymbol{\theta}}^*)$

parameters $\boldsymbol{\phi}$. The optimal variational parameters are found by minimizing the Kullback-Leibler divergence (D_{KL}) between the approximate and true posteriors. This is equivalent to maximizing the evidence lower bound (ELBO), which can be expressed as

$$\mathcal{L}_{\text{ELBO}}(\boldsymbol{\phi}) = \mathbb{E}_{q(\boldsymbol{\theta}|\boldsymbol{\phi})}[\log p(\mathcal{D}|\boldsymbol{\theta})] - D_{\text{KL}}(q(\boldsymbol{\theta}|\boldsymbol{\phi})||p(\boldsymbol{\theta})). \tag{6.6}$$

The first term is the expected log-likelihood, which encourages the model to fit the data, while the second term is the D_{KL} between the variational distribution and the prior, which acts as a regularization term. A standard Gaussian distribution is used as prior, $p(\boldsymbol{\theta})$, which is a common choice for BNNs (Jospin et al., 2022).

Our BNN is built upon PyTorch using flipout layers from the Bayesian-torch library (Krishnan et al., 2022). Flipout is an efficient VI technique that reduces the variance of the gradient estimates by decorrelating the gradients within each mini-batch (Wen et al., 2018). Since MPM is a binary classification task, the ELBO's likelihood is calculated using the `BCEWithLogitsLoss` (binary cross entropy with logits loss) function from the PyTorch library. The network optimization is performed using the Adam optimizer (Kingma and Ba, 2014) with an initial learning rate of 10^{-3}. Predictions and their associated uncertainties are obtained for the unseen test dataset by aggregating the results of 100 forward passes. This number of samples is chosen as it represents a conventional balance between achieving a stable estimate of the predictive mean and variance, and maintaining a reasonable computational cost for inference.

6.3.3.3 Monte Carlo dropout

MC dropout is a technique for approximate Bayesian inference that can be applied to deep neural network architectures with the addition of dropout layers (Gal and Ghahramani, 2016). Typically, dropout is used only as a regularization method during training, where a random subset of weights are deactivated to prevent over-fitting. In a standard network, these layers are then deactivated during the inference phase to produce a single, deterministic output. The main difference of MC dropout is that it keeps the dropout layers active during inference as well. By performing multiple stochastic forward passes on the same input data, the network effectively becomes a MC sampler and generates a distribution of varying predictions as different sets of neurons are randomly dropped out in each pass.

We have implemented our MC dropout model using functions in the PyTorch library and trained with the Adam optimizer, using an initial learning rate of 10^{-3} and the `BCEWithLogitLoss` loss function. The dropout rate, which is applied after each hidden layer, is selected empirically by evaluating rates of 0.1, 0.2, and 0.3 on a validation set, with the final value of 0.1 yielding the best performance. Finally, predictions and their associated uncertainties are obtained by aggregating the results of 100 MC forward passes on the test dataset.

6.3.4 Evaluation metrics

To provide a comprehensive comparison between the three Bayesian approximation methods, all evaluations are performed on the unseen testing dataset. The models are assessed based on four categories of metrics: predictive performance, computational cost, model robustness, and model interpretation. The model interpretation analysis is conducted only on the best-

performing model.

6.3.4.1 Prediction performance

We assess classification performance using metrics derived from the four outcomes of a confusion matrix: (1) true positives (TP), where a deposit is correctly classified as a deposit; (2) true negatives (TN), where a barren location is correctly classified as barren; (3) false positives (FP), where a barren location is incorrectly classified as a deposit; and (4) false negatives (FN), where a deposit is incorrectly classified as barren. Uncertainty in the classification metrics is quantified using bootstrap resampling with 1000 iterations of the test set to compute 95 % confidence intervals.

Area under the receiver operating characteristic curve (ROC-AUC): This evaluates the model's performance across classification thresholds by using its probabilistic outputs. The ROC curve is plotted with the true positive rate on the y-axis against the false positive rate on the x-axis. The final AUC score measures the model's ability to discriminate between the positive and negative classes, where an AUC score of 1 represents a perfect classifier and a score of 0.5 represents performance equivalent to random guessing.

Recall: This measures the proportion of actual deposits that are correctly identified by the model. Recall score also ranges from 0 to 1, where a higher value indicates better performance. A score of 1 means the model successfully identified every true deposit. It is calculated as

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}. \tag{6.7}$$

Geometric mean score (GMS): This metric measures the balance of performance on both the positive and negative classes. It is calculated as the geometric mean of sensitivity and specificity, which is given by

$$\text{GMS} = \sqrt{\text{Sensitivity} \times \text{Specificity}} = \sqrt{\frac{\text{TP}}{\text{TP} + \text{FN}} \times \frac{\text{TN}}{\text{TN} + \text{FP}}}. \tag{6.8}$$

The GMS ranges from 0 to 1, where 1 indicates perfect classification of both classes.

Prediction-area (P-A) plot: This curve evaluates the model's effectiveness at prioritizing prospective areas. The PA plot consists of two curves where the prediction rate curve shows

the cumulative percentage of deposits found, while the cumulative area curve represents the percentage of total land area (Yousefi and Carranza, 2015). The intersection point between the two curves provides a key measure of model performance, as a higher intersection signifies a stronger predictive ability.

6.3.4.2 Computational cost

To compare the efficiency and complexity of the models, we measured three computational cost metrics during a representative 100-epoch training run.

Number of parameters: This is the total number of trainable parameters, such as means and variances of weights and biases in the model.

Peak memory: This metric indicates the hardware requirement, which measures the maximum physical RAM in megabytes (MB) used by the training loop, excluding initial data loading overhead. To capture peak memory usage, we use the psutil Python library to monitor RAM usage after each training batch throughout the entire experimental run. This entire experiment is then repeated three times, and the mean and standard deviation of the three recorded peak values are calculated.

Average time per epoch: This reflects training efficiency and is the mean time in seconds required to complete one training epoch. We also report the standard deviation to quantify the variability in training time per epoch.

6.3.4.3 Robustness

To evaluate model robustness against noisy input data, we performed a sensitivity analysis using the ROC-AUC drop. This is the absolute difference between the ROC-AUC score on the original test set and the score on the same test set after adding a moderate level of Gaussian noise with a standard deviation of $\sigma = 0.1$ to the preprocessed data features. A smaller AUC drop indicates a more robust model.

6.3.5 Model interpretation

To understand the drivers behind the model's predictions and uncertainties, we assess the importance and contribution of each geoscience data feature using the SHAP method (Lundberg and Lee, 2017b). SHAP is developed based on game theory and explains how each

geoscience feature contributes to a specific prediction. A positive SHAP value indicates that the feature contributes to the model's prediction being more prospective, while a negative value contributes to it being less prospective. Because the SHAP library's `KernelExplainer` class can be applied to any function that maps features to a single scalar output, we perform two separate analyses to explain both the mean (prospectivity) and the variance (uncertainty) of the TAGI model's output. This dual analysis allows us to identify which features are the primary drivers of not only the prospectivity score but also the model's predictive uncertainty. When using `KernelExplainer`, a representative background dataset is required to establish a baseline for explanations. Therefore, the data is carefully sampled to handle the inherent class imbalance. A background dataset of 1,000 samples is created from the training data, consisting of 650 negative and 350 positive samples, where the majority class is sampled using k-means clustering to ensure diversity. Next, an explanation set of 3,000 samples is drawn from the test set, ensuring the inclusion of all 396 positive samples and randomly selecting the remaining 2,604 samples from the negative class. The final analysis is performed with the number of perturbations for each explanation, where `nsamples` parameter is set to 1,000. Since `KernelExplainer` is computationally expensive, all the above sample sizes are selected to maximize representativeness given our available computational resources.

6.4 Results and discussion

6.4.1 Model performance and selection

A comparison between the three models is conducted on the unseen testing dataset, and their performance, computational cost, and robustness are evaluated using the metrics summarized in Table 6.2. To ensure a comprehensive evaluation, we tested all models across multiple classification thresholds (0.3-0.7). Table 6.2 shows results at threshold 0.5, with the threshold selection criteria discussed in Section 6.4.2 and complete threshold analysis provided in Appendix B.

Table 6.2 Performance comparison between the TAGI, BNN with VI, and MC dropout models evaluated on the unseen test set at classification threshold 0.5. The best result in each row is highlighted in bold. See Appendix B for performance at multiple thresholds.

Metric	TAGI	BNN (VI)	MC dropout
Performance (at threshold 0.5)			
ROC-AUC	**0.91 ± 0.009**	0.88 ± 0.007	0.87 ± 0.008
Recall	**0.70 ± 0.023**	0.29 ± 0.022	0.30 ± 0.023
GMS	**0.81 ± 0.013**	0.53 ± 0.020	0.54 ± 0.021
Computational cost			
Number of parameters	13,314	13,314	**6,657**
Peak memory (MB)	1607.30 ± 51.10	1197.69 ± 26.69	**971.46 ± 17.15**
Epoch time (s)	**139.50 ± 1.27**	390.12 ± 5.24	194.40 ± 9.87
Robustness			
AUC drop (Noise $\sigma = 0.1$)	**0.17**	0.28	0.31

In terms of predictive performance, all three models demonstrate a strong ability to discriminate between barren and deposit locations. As illustrated in Figure 6.4, the TAGI model has the best ROC-AUC value of 0.91, followed by the BNN and MC dropout models with AUC values of 0.88 and 0.87. Next, the TAGI model's ability to identify the positive deposits is highlighted by achieving a recall value of 0.70, which substantially outperforms both the BNN, which scores 0.29, and the MC dropout model, which scores 0.30. Furthermore, the TAGI model's ability to categorize deposit and barren areas is supported by a high GMS of 0.81, surpassing the BNN and MC dropout scores of 0.53 and 0.54, respectively.

The analysis of computational costs indicates that the MC dropout model has the highest parameter efficiency, requiring only 6,657 parameters because it does not require storing distributional information in its weights. On the other hand, both the TAGI and BNN models require 13,314 parameters to represent the mean and variance for each weight. This efficiency extends to memory usage, where the MC dropout model requires the least peak memory of 971.46 ± 17.15 MB. Conversely, the TAGI model requires the highest peak memory of 1607.30 ± 51.10 MB, which is probably a consequence of storing the layer-wise statistics needed for its analytical updates. However, when considering training speed, the TAGI model is significantly faster, completing an epoch in an average of 139.50 ± 1.27 seconds, whereas the BNN is the slowest, with an average of 390.12 ± 5.24 seconds.

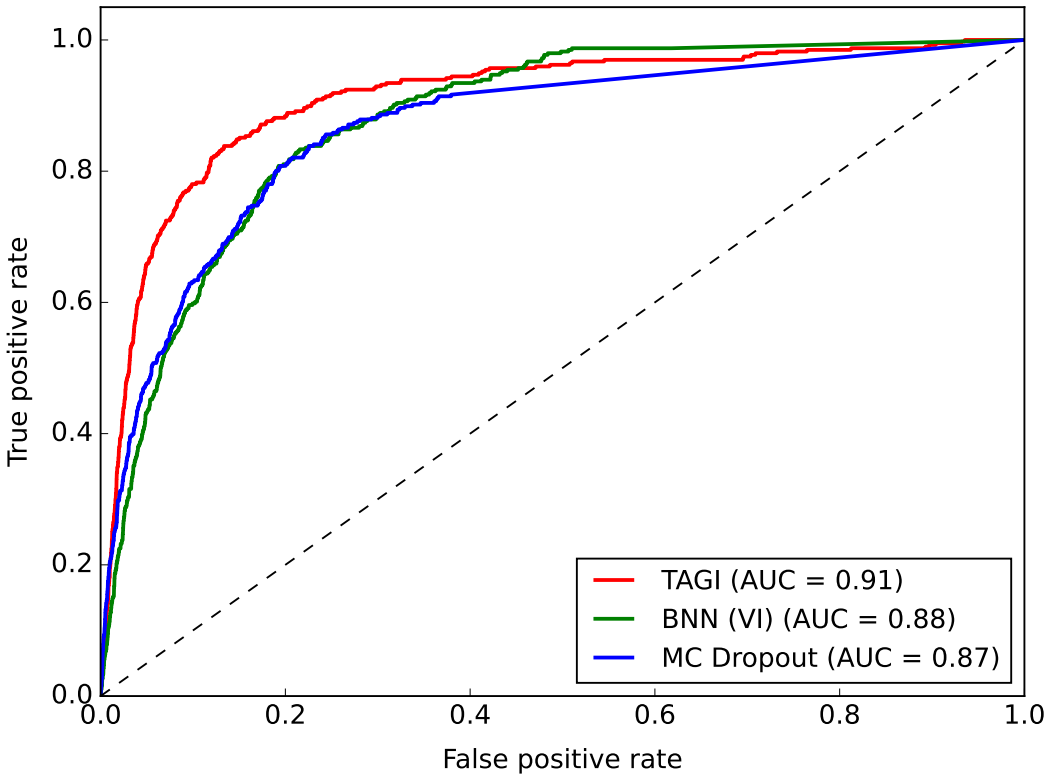


Figure 6.4 Mean ROC curves and AUC values for the TAGI, BNN with VI, and MC dropout models on the unseen test set. The dashed diagonal line represents the performance of a random classifier with an AUC of 0.5.

Finally, the robustness of the three models is evaluated by adding Gaussian noise to the geoscience features of the test set and measuring the resulting degradation in performance. With the smallest AUC value drop of 0.17, the TAGI model showed the best performance, while the BNN and MC dropout models indicate that they are more sensitive to the noise, with higher AUC drops of 0.28 and 0.31, respectively. Therefore, despite the MC dropout model's higher efficiency in parameter count and peak memory usage, TAGI's superior combination of predictive performance across all metrics, faster training speed, and greater robustness establishes it as the best-performing model.

Furthermore, to provide a more detailed view of each model's classification of the rare positive class, we examined their respective confusion matrices, which are illustrated in Table 6.3. The TAGI model correctly identifies 275 of the 396 test deposits as TPs, while the BNN and MC dropout models find only 114 and 117 TPs, respectively. This significant difference in TPs directly explains TAGI's higher recall score, which is necessary for successfully identifying the few existing mineral deposits across a vast barren area.

Table 6.3 Summary of confusion matrix results for (a) TAGI, (b) BNN with VI, and (c) MC dropout models on the test set. TP, FN, TN, and FP denote true positives, false negatives, true negatives, and false positives, respectively.

Model	TP	FN	TN	FP
TAGI	275	121	309,221	19,951
BNN (VI)	114	282	324,091	5,081
MC dropout	117	279	323,699	5,473

6.4.2 TAGI prospectivity and uncertainty maps

Having established TAGI as the superior model based on its predictive performance and robustness, the following analysis focuses on its primary outputs. The final mineral prospectivity map generated by the TAGI model is presented in Figure 6.5. The map successfully delineates high-prospective regions that coincide with the locations of known mineral deposits while also highlighting several new potential exploration targets. Furthermore, the overall spatial patterns are consistent with previously published national-scale magmatic Ni prospectivity maps for Canada (Lawley et al., 2021; Nagasingha et al., 2024), particularly across the Raglan and New Québec Orogen (Labrador Trough) regions. In this area, the model correctly prioritizes the Paleoproterozoic mafic-ultramafic sills of the Montagnais Sill Complex, which are spatially associated with strong horizontal gravity gradients (Smith et al., 2021). For detailed qualitative representations of the geological and geophysical features in this specific region, the reader is referred to Figures 9 and 10 in Nagasingha et al. (2024).

The intersecting point in the P-A curve of Figure 6.6 highlights the model's ability to significantly reduce the search space, showing that 85 % of the test set's known deposits are concentrated in just 15 % of the most prospective land area. Moreover, at a probability threshold of 0.5, the model still retains more than 70 % of known deposits, suggesting that many deposit-bearing polygons receive confident predictions.

To translate these results into actionable exploration targets, we generate a binary prospectivity map (Figure 6.7) using a classification threshold of 0.5 which we selected based on GMS values. The binary map shows that, beyond the 275 correctly identified test deposits, the model predicts an additional 19,951 hexagons as prospective. However, these should not be interpreted simply as prediction errors. At this spatial resolution, each known deposit occupies only a single hexagonal cell. However, geologically favourable conditions typically extend across multiple surrounding hexagonal cells, meaning that prospective cells near known

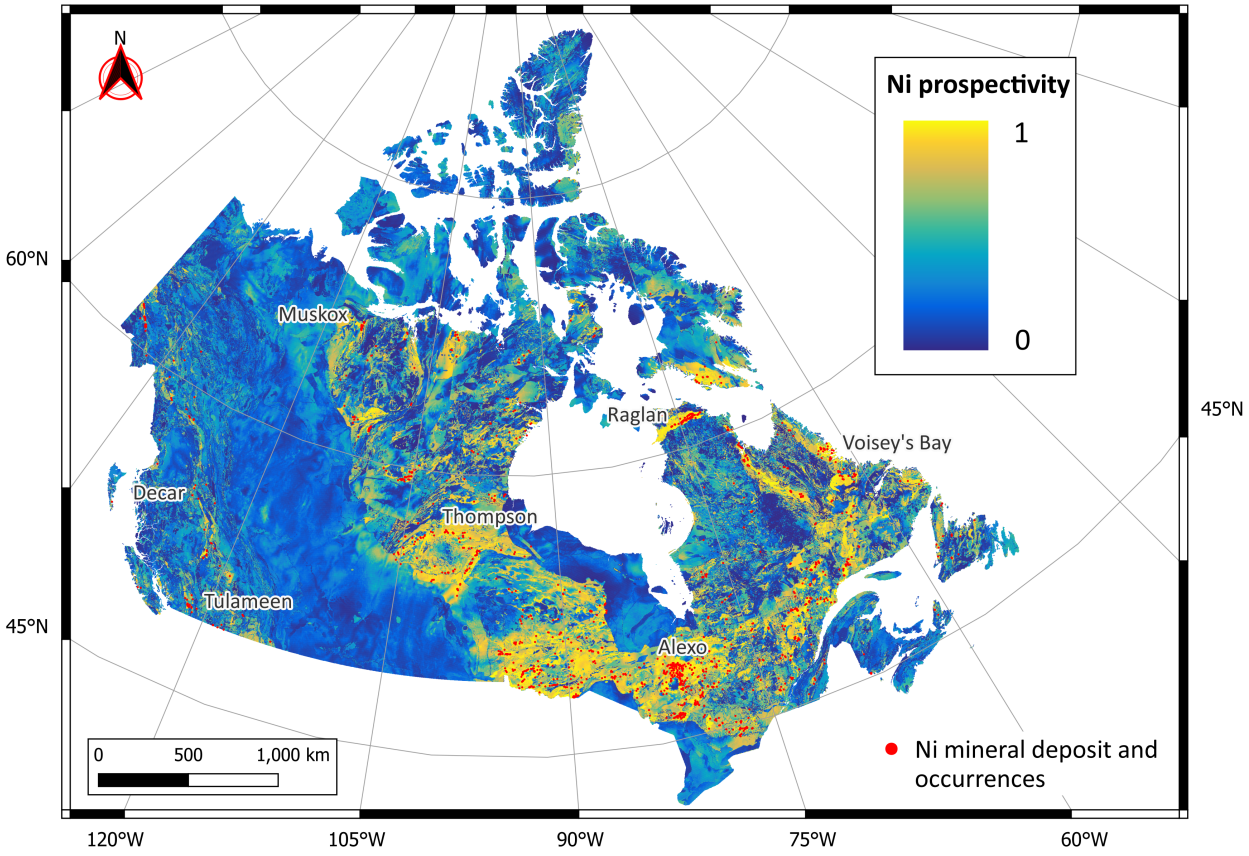


Figure 6.5 MPM for magmatic Ni (±Cu ±Co ±PGE) sulphide deposits across Canada, generated by the TAGI model. High prospectivity scores, approaching one, are shown in yellow, while the lowest scores, approaching zero, are shown in blue.

deposits often capture the broader extent of favourable geological conditions while those in underexplored regions may include prospective targets that have not yet been validated through field surveys or drilling. These spatial characteristics make GMS a more appropriate performance metric than precision-focused measures for national-scale prospectivity modelling, as GMS balances both sensitivity and specificity without being dominated by the overwhelming negative class. Consequently, this binary map serves as a decision-support tool for prioritizing specific areas for detailed follow-up investigations.

These predicted targets can be further interpreted through the lens of survival bias, which is a form of selection bias that results from neglecting areas due to their lack of visibility and the absence of obvious mineralization signals at the surface (Saremi et al., 2025; Yousefi et al., 2024b). In MPM, areas lacking documented mineral occurrences are usually excluded from exploration programs, even though the absence of known deposits does not confirm that

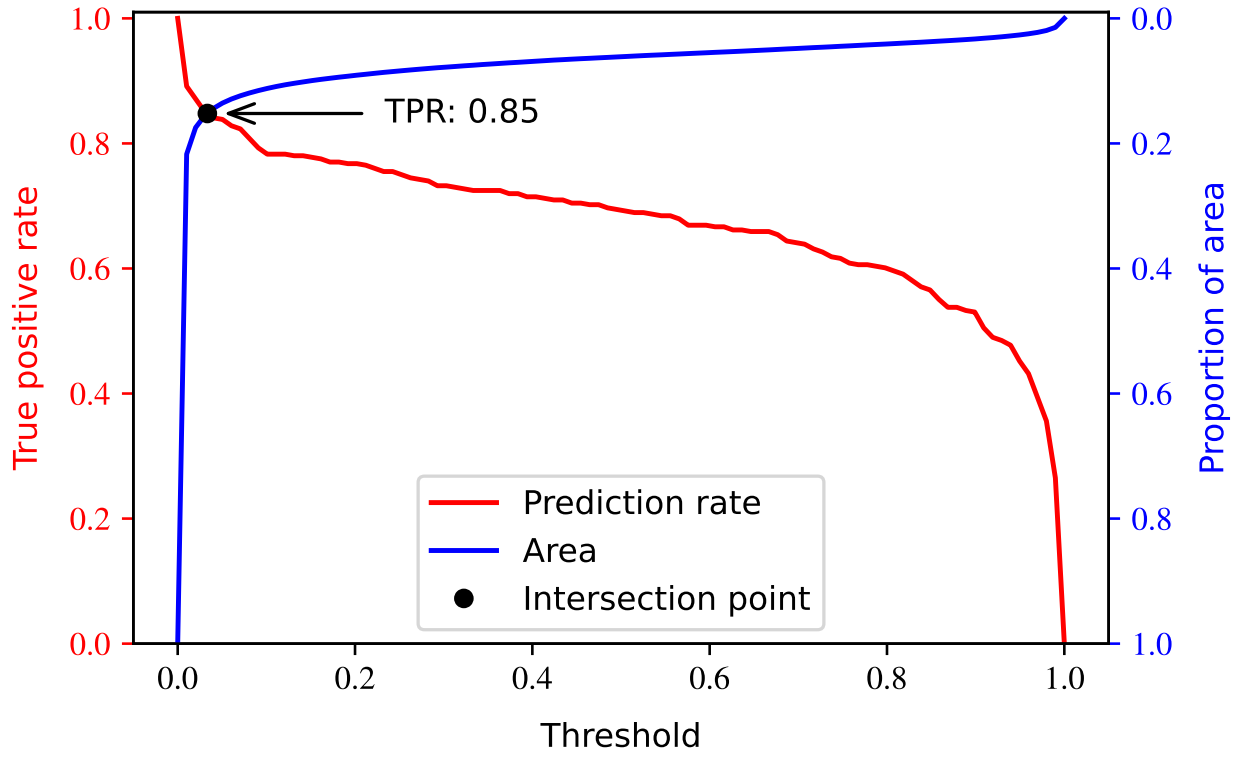


Figure 6.6 P-A curve for the TAGI model, showing the relationship between the true positive rate and the proportion of total land area flagged as prospective at different classification thresholds.

such areas are barren. The regions between clusters of known deposits can represent critical areas that warrant careful examination rather than exclusion. Our binary prospectivity map (Figure 6.7) illustrates this at two scales. In the well-explored New Québec Orogen and Raglan region, the model validates known prospective areas, which confirms its ability to capture established Ni district signatures. Conversely, the model also identifies prospective targets in less-explored regions, such as the corridor between the Muskox intrusion and the Thompson Nickel Belt within the Churchill Province (Figure 6.1), which, despite being situated between two major Ni deposit clusters, contains very few documented occurrences. The TAGI model flags these targets based on geological and geophysical signatures that are sensitive to subsurface conditions. According to the null and negative test concept (Yousefi et al., 2024b), these targets likely satisfy a null test, where available data cannot exclude the presence of mineralization, rather than a negative test where evidence confidently rules out a mineralized system.

Next, the relative uncertainty map illustrated in Figure 6.8 provides further insights into the model's confidence and complements the survival bias analysis discussed above. This met-

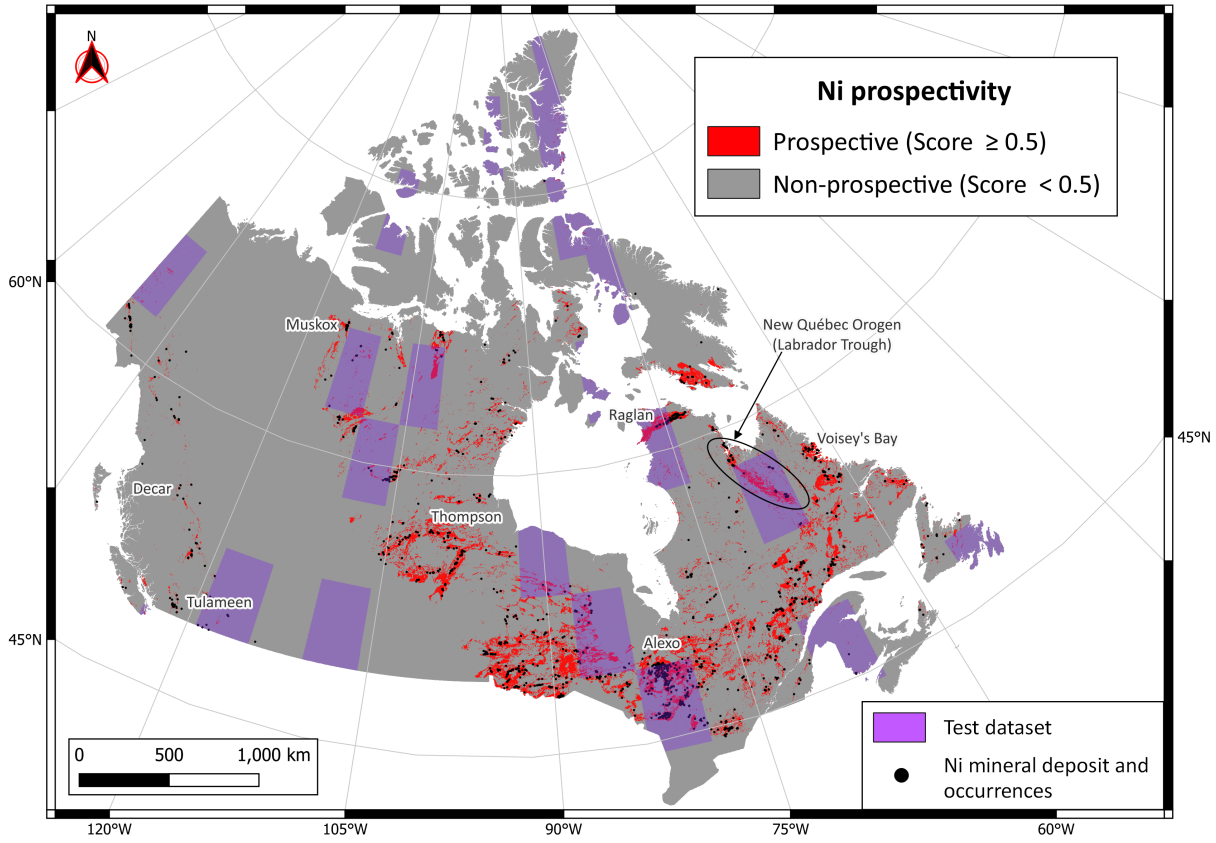


Figure 6.7 Binary MPM for magmatic Ni (±Cu ±Co ±PGE) sulphide deposits in Canada, derived from the TAGI model using an optimal classification threshold of 0.5. Prospective target areas are shown in red, while non-prospective regions are shown in gray. Black dots represent known Ni deposit locations. Additionally, the testing dataset used to evaluate model performance is illustrated in purple regions.

ric normalizes the predictive standard deviation by the predicted prospectivity score, where higher values indicate greater model uncertainty. While training data predominantly reflects easily discoverable, near-surface deposits, the uncertainty map enables dual interpretation: low relative uncertainty in regions around known mineral deposits demonstrates the model's highest confidence where it has been constrained by ground-truth data, while higher uncertainty across Canada's northern regions reflects both sparser data coverage and potential for deeper or unconventional deposits that differ from historical exploration patterns. Additionally, this relative uncertainty map is a notable improvement over our previous ensemble-based approach on the same dataset (Nagasingha et al., 2025), which appeared noisy due to pixel-level model disagreements, whereas the TAGI-based map is spatially coherent and therefore more easily interpretable. When used jointly with the binary prospectivity map, newly iden-

tified prospective areas with low relative uncertainty are more likely to satisfy a null test, while those with high relative uncertainty indicate where additional data acquisition is needed before a null or negative determination can be made.

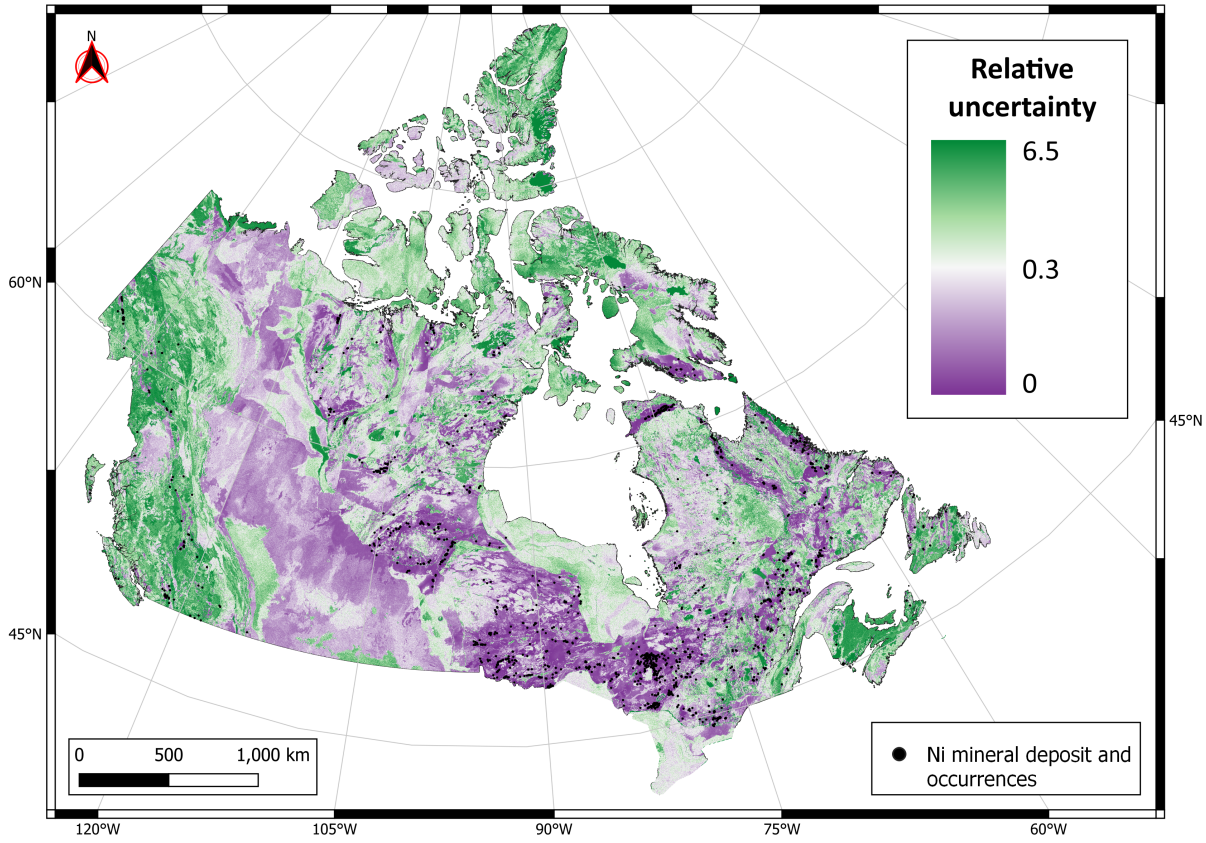


Figure 6.8 Relative uncertainty map for the TAGI model across Canada, where purple areas indicate low relative uncertainty and high model confidence, while green areas indicate high relative uncertainty and low model confidence. The locations of known Ni mineral deposits and occurrences are shown as black dots.

6.4.3 Geological interpretation through SHAP analysis

To understand the key drivers of the TAGI model's prospectivity predictions, we evaluated the results obtained from the SHAP analysis, which is illustrated in Figure 6.9. The SHAP analysis reveals that the model's predictions are significantly driven by geological features, with lithology and geological period data overwhelmingly dominating the feature ranking. The plot for lithology (mafic), which is a critical feature representing the primary source of metals in a magmatic Ni sulphide mineral system, demonstrates a geologically realistic relationship that extends beyond a simple positive correlation. While the long red tail on the

right side of the plot confirms that a high presence of mafic rocks can be a strong positive predictor, there is also a dense cluster where high values for this same feature result in negative SHAP values. This indicates that the model has learned a crucial geological insight: the mere presence of mafic host rocks is not sufficient for prospectivity. When other key components of a mineral system, such as pathways or sulphur sources, are absent, the model correctly identifies these as barren intrusions and downweights their importance. Additionally, the model correctly incorporates key geophysical features into its predictions, where gravity (Bouguer), which serves as a geophysical proxy for the dense mafic-ultramafic intrusions that host these deposits, is ranked as a highly important predictor. The SHAP plot reveals a clear positive correlation, where high gravity values, indicated by red points, consistently yield positive SHAP values, confirming that the model has learned to target dense rock bodies. Furthermore, the model captures the structural pathways required for magma transport, which is evidenced by fault (proximity) and terrane (proximity) plots. Both features exhibit a clear inverse relationship, where the blue dots represent closer proximity to these major structures, which increases the prospectivity score. This is further supported by magnetic (long-wavelength HGM) and gravity (HGM), where the model is used to map the structural lineaments that define these pathways.

The SHAP analysis of the model's uncertainty provides insights that are distinct from the drivers of prospectivity (Figure 6.10). The results show that, similar to the prospectivity predictions, the model's uncertainty is overwhelmingly driven by geological map features, with lithology and geological period data dominating the top of the feature ranking. While these dominant geological features demonstrate mixed feature values for the models' confidence, lithology (mafic), which exhibits the presence of some favourable host rocks, tends to slightly increase the model's uncertainty. This behaviour is probably due to the model having learned that not all mafic intrusions are mineralized, and their presence alone is not a guarantee of a deposit, thus introducing a small amount of ambiguity. On the other hand, the key geophysical features, including gravity (Bouguer), indicate that the model's confidence increases when the gravity signal is strong and clear. For example, a strong positive gravity anomaly, which is a clear indicator of a prospective dense body, results in negative SHAP values, meaning it actually decreases the model's uncertainty. However, the plot also shows that conflicting evidence, such as a low gravity value in an otherwise geologically favourable setting, can significantly increase the model's uncertainty. These SHAP-based insights enable exploration geoscientists to prioritize resource allocation by targeting areas where high prospectivity coincides with reasonable uncertainty, while identifying regions where additional data acquisition, such as targeted geophysical surveys, could probably reduce model uncertainty and improve decision confidence.

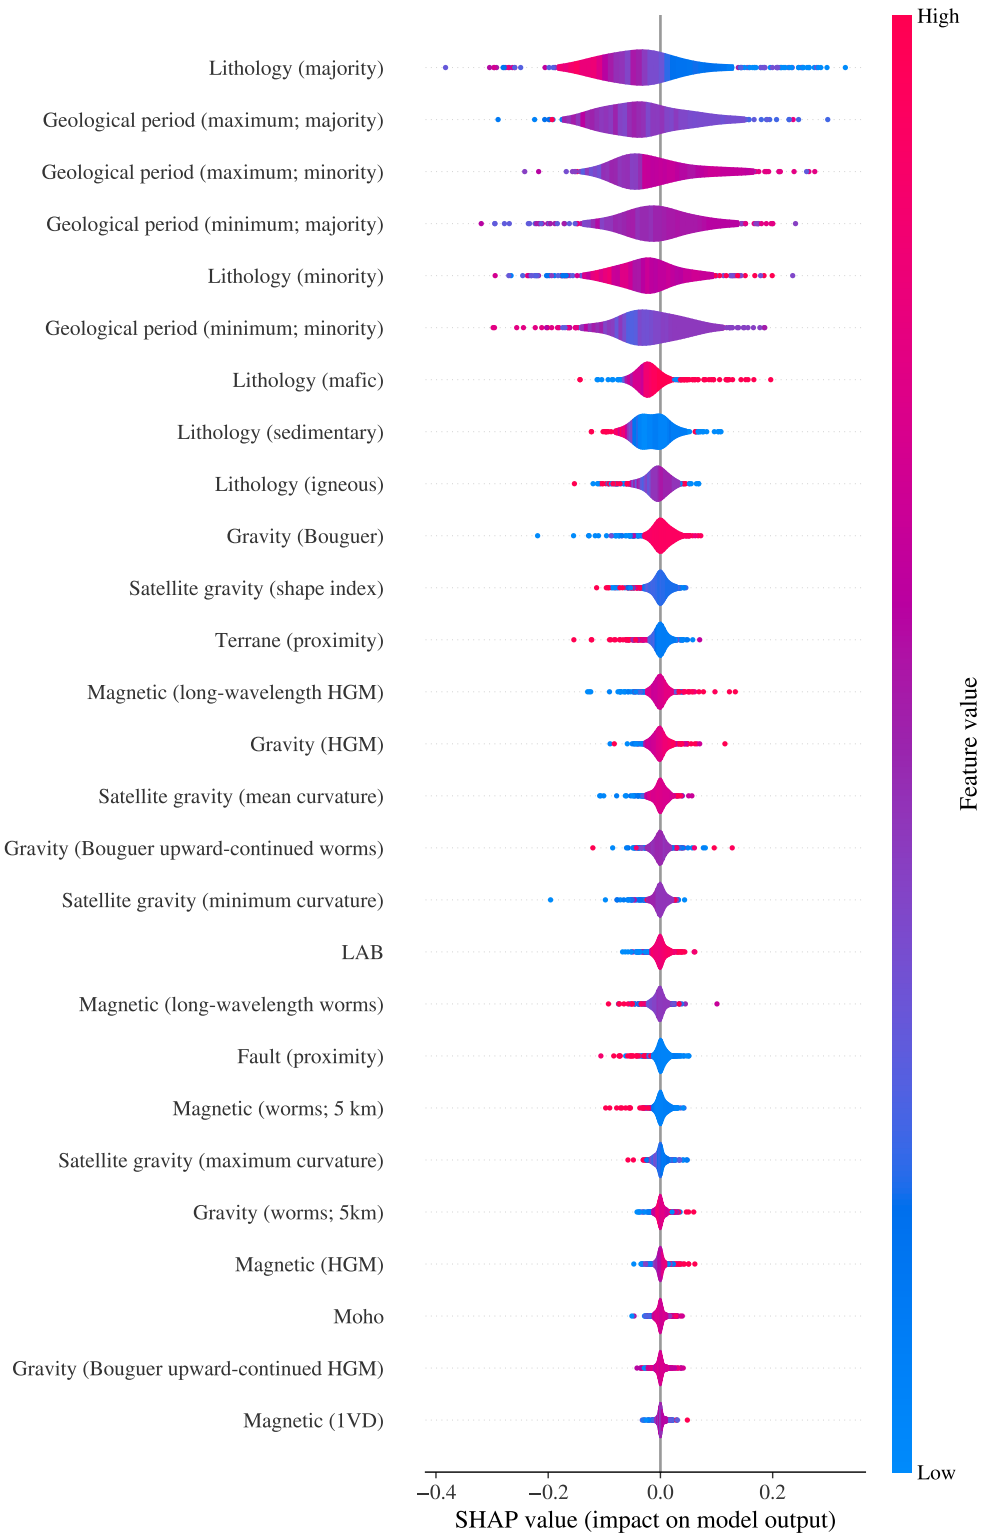


Figure 6.9 SHAP summary plot showing the feature importance for the TAGI model's prospectivity predictions. Each feature's violin shows the distribution of its SHAP values across the test set. The position on the x-axis indicates the impact on the model's output, while the colour represents the feature's value from high (red) to low (blue). The features are ranked from most to least important.

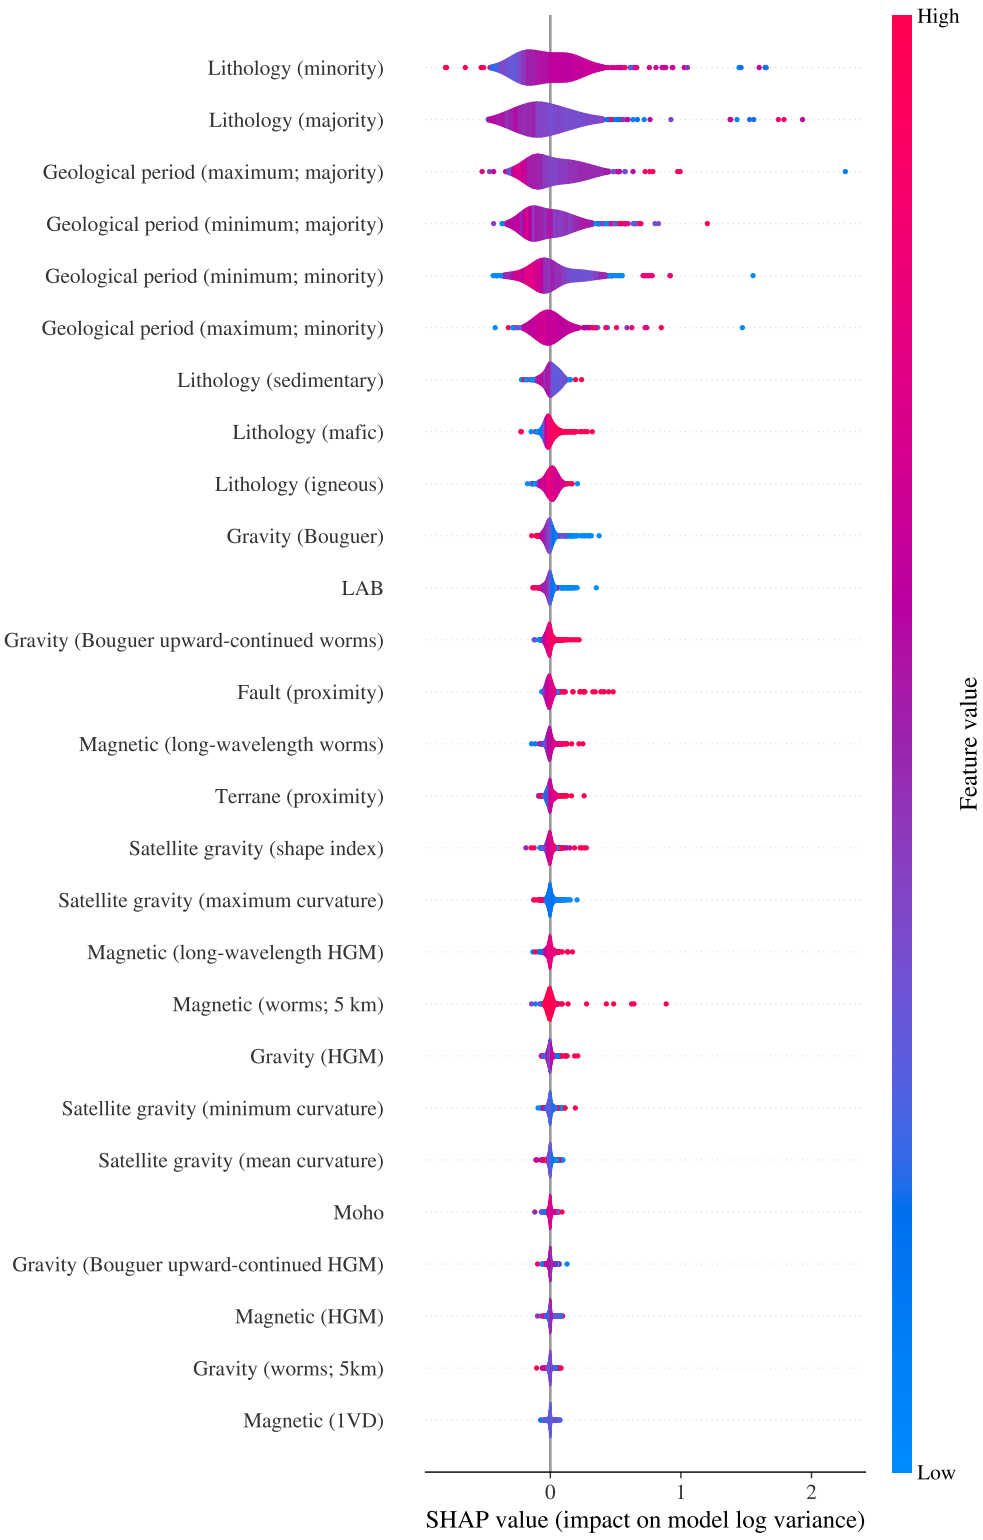


Figure 6.10 SHAP violin plot showing the feature importance for the TAGI model's uncertainty across the test set. The features are ranked by importance, where a positive SHAP value on the x-axis corresponds to an increase in output variance. The colour bar represents the feature's value, with high values shown in red and low values in blue.

6.4.4 Understanding TAGI's performance advantages

TAGI's outperformance of BNN with VI and MC dropout can be attributed to several underlying mechanisms that distinguish analytical inference from gradient-based Bayesian approximations. Unlike gradient-based methods that rely on iterative optimization procedures susceptible to local minima and gradient instability, TAGI updates parameters directly through analytical inference using the RTS smoother, thereby eliminating the risk of suboptimal convergence that can occur when gradient-based methods become trapped in local minima during training. Next, TAGI's computational efficiency stems from its diagonal covariance structure and layer-wise recursive updates, which achieve linear computational complexity with network depth (Goulet et al., 2021). Additionally, the use of ADASYN may introduce noise or synthetic artifacts into the generated positive samples (Li et al., 2021a). TAGI's analytical inference approach through GMA and local linearization may be less sensitive to such artifacts compared to gradient-based methods, which is evidenced by its 39 % lower performance degradation under added noise and consistent superior performance across all tested thresholds (Appendix B).

6.4.5 Limitations of the proposed framework

Overall, the results not only establish TAGI as the superior predictive model among the three Bayesian methods evaluated in this study but also demonstrate that a SHAP analysis of its unique uncertainty measure can yield critical insights into the geological factors that govern the model's confidence. However, it is also important to acknowledge the limitations of this proposed framework. Firstly, the assumption of homoscedastic σ_V^2 is a known limitation of the original TAGI algorithm we implemented, as the true data noise in geology is often heteroscedastic. This may lead to a uniform treatment of uncertainty across heterogeneous geological settings, potentially underestimating uncertainty in data-poor regions and overestimating it in well-constrained areas. Secondly, preliminary experiments without oversampling yielded a recall below 3 % as models predicted the majority class almost everywhere, which demonstrates that data augmentation is essential for addressing class imbalance in this national-scale prospectivity modelling. As discussed in the previous section, TAGI's analytical propagation may be more robust to such noise than gradient-based methods, though this hypothesis requires dedicated investigation in future research. Finally, since this study is conducted on a national scale, the model's performance is significantly dependent on the quality and resolution of the input geoscience data. This is evident in the SHAP analysis, where lithology and geological period showed the dominant influence on the model's outputs, consistently exhibiting the widest spread of SHAP values for both prediction and uncertainty

compared to other geological and geophysical features.

6.5 Conclusion

This study presents a comparative analysis of three Bayesian deep learning methods: TAGI, which is based on an analytical inference method, conventional gradient-based BNN with VI, and MC dropout, for modelling the prospectivity of magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems in Canada. To ensure a fair comparison, the ADASYN sampling technique is used to address the inherent class imbalance in the dataset. A comprehensive comparison has been carried out between the models, and the results show that TAGI has several advantages over the gradient-based Bayesian approximations. Specifically, TAGI yields the highest predictive accuracy, the most efficient training speed, and the greatest resilience to noisy input data, although the MC dropout model uses less trainable parameters and memory. By leveraging these advantages, the prospectivity map generated by the TAGI model proved highly effective. It accurately classified the majority of known deposits, correctly identified 275 of the 396 positive samples in the test set, and its spatial patterns are consistent with previous national-scale models. Furthermore, the P-A curve demonstrates the map's practical utility, showing that 85 % of these deposits are concentrated within just 15 % of the most prospective land area, significantly reducing the exploration search space. Next, thanks to the approximate analytical uncertainty measure produced by the TAGI model, the relative uncertainty map generated by the TAGI model is a notable improvement over the uncertainty map from our previous ensemble-based approach on the same dataset. Finally, the SHAP analysis confirmed that the model learned geologically realistic relationships within the mineral system, revealing that lithology and geological period data are the dominant drivers of prospectivity predictions, as well as the model's confidence, which is likely due to the quality and resolution of the geological map data. Future work could leverage extensions of the TAGI framework to model heteroscedastic uncertainty, thereby allowing for more realistic uncertainty estimates.

Data and code availability

The paper uses Python as the primary programming language to develop the prospectivity models. The dataset and the codes are available at: `https://github.com/Laish93/probabilistic_mpm`

Acknowledgements

C. L. Bérubé and L. M. A. Nagasingha acknowledge funding from Natural Resources Canada [TGI project GC-130021S] and the support of the Natural Sciences and Engineering Research Council of Canada [reference number 2024-06161]. The authors thank J. A. Goulet for constructive criticism of the manuscript.

CHAPTER 7 GENERAL DISCUSSION

Data-driven approaches to MPM have advanced considerably over the past two decades, yet most existing models still exhibit methodological limitations that constrain their application to real-world mineral exploration. These limitations are largely attributable to the severe class imbalance between known deposits and barren locations in training datasets (Farahbakhsh et al., 2023; Farahnakian et al., 2024; Zhang et al., 2023), the inability of most predictive models to quantify the uncertainty inherent in their predictions (Lisitsin et al., 2014; Zuo et al., 2021a), and the limited interpretability of ML algorithms, which restricts the ability to validate both predictions and their associated uncertainties against geological knowledge (Zuo et al., 2024). Although some of these challenges have been addressed individually using various methods, no studies have integrated probabilistic frameworks to address them collectively. The studies presented in Chapter 4 (Nagasingha et al., 2024), Chapter 5 (Nagasingha et al., 2025), and Chapter 6 (Nagasingha et al., 2026) demonstrate that probabilistic ML frameworks not only improve predictive performance but also expand the scope of MPM from deterministic prospectivity maps to decision-support tools that quantify where the model is confident, where it is uncertain, and which geological factors control both. This chapter discusses the key insights that emerge from examining these contributions together, their geological and practical implications, and their limitations.

7.1 Cross-article methodological insights

The national-scale prospectivity model of Lawley et al. (2021) applied gradient boosting machines to Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems without data augmentation, achieving AUC values ranging from 0.846 to 0.923. To address the class imbalance in training data, this thesis employs two different augmentation strategies: CVAE in Chapters 4 and 5, and ADASYN in Chapter 6. While the CVAE generates synthetic samples by modelling the conditional distribution of geoscience features given mineralization labels, ADASYN oversamples adaptively in regions of the feature space where positive instances are sparse. Despite these different approaches, both achieve strong predictive performance, with CVAE combined with XGBoost reaching an AUC of 0.94 and ADASYN combined with TAGI reaching 0.91. The resulting prospectivity maps capture known mineral deposits more effectively and identify additional prospective regions not delineated by Lawley et al. (2021). Although the evaluation approaches differ, both aim to quantify the spatial efficiency of the prospectivity model in narrowing the exploration search space. While Lawley et al. (2021) re-

ported that 8 % of the study area contains more than 80 % of known deposits, the high-return zones identified by the novel risk-return analysis framework presented in Chapter 5 concentrate 94.7 % of known deposits within approximately 4 % of the study area by integrating prospectivity with uncertainty estimates.

The thesis also develops two approaches to uncertainty quantification, an ensemble-based approach in Chapter 5 and an analytical approach using TAGI in Chapter 6. The ensemble approach derives uncertainty from the disagreement among independently trained models, which requires training and storing multiple networks, whereas TAGI estimates uncertainty analytically by propagating Gaussian distributions over the network parameters with low computational cost. While both approaches produce low uncertainty around known deposit locations, the ensemble-derived uncertainty maps exhibit greater spatial noise in regions away from known deposits, whereas the TAGI framework produces smoother, more geologically interpretable uncertainty estimates. Regardless of the approach used, most existing MPM evaluation metrics, such as AUC and P-A plots, focus primarily on predictive accuracy with limited consideration of uncertainty (Zhang et al., 2024). While previous risk-return frameworks for MPM have been proposed using log-odd ratios (Wang et al., 2020) and predictive model values (Parsa and Carranza, 2021), the framework developed in Chapter 5 advances these approaches by deriving relative uncertainty measures from model disagreement and integrating non-parametric spatial clustering to produce spatially coherent risk-return categories while filtering out noisy, isolated predictions. The practical importance of uncertainty quantification is further demonstrated by the concept of survival bias (Yousefi et al., 2024b), in which regions lacking documented deposits are usually excluded from exploration, even though undiscovered mineralization may still exist at depth or beneath cover. The uncertainty maps produced in this thesis highlight this by distinguishing between areas where the model is confident in its predictions and areas where uncertainty remains high, regardless of the predicted prospectivity.

Additionally, the three articles employ different evaluation metrics, from success rate curves in Chapter 4, to ROC-AUC analysis in Chapter 5, to a combination of ROC-AUC, recall, GMS, and prediction-area plots in Chapter 6. This progression reflects an important finding: reliance on a single metric can mask critical weaknesses in model performance. Precision is sensitive to the resolution mismatch between national-scale geoscience features, which span many cells, and deposit locations that occupy single cells, leading to elevated false-positive rates, whereas recall is constrained by the small number of known deposit locations, which may not fully represent the true distribution of mineralization. Therefore, rather than relying on a single metric, evaluating model performance through multiple complementary metrics, such as GMS, which balances sensitivity and specificity, provides a more comprehensive

understanding of model behaviour.

Despite these methodological differences, the prospectivity maps produced across all three studies exhibit broadly similar spatial patterns, identifying largely the same major prospective regions across Canada. From a practical perspective, the CVAE-based model can be considered the most reliable predictive map due to its higher AUC and more geologically realistic data augmentation, whereas the TAGI framework provides the most informative uncertainty and interpretability maps that are supported by dual SHAP analysis of both predictions and model confidence. Together, these outputs provide a practitioner with both reliable prospectivity targets and the information needed to prioritize where additional data acquisition would be most valuable.

7.2 Geological implications

The results across all three studies consistently identify lithological and geochronological features as the dominant drivers of both model predictions and model uncertainty for Canadian magmatic Ni sulphide prospectivity, which is consistent with the variable importance analysis of Lawley et al. (2021) on the same mineral system. While attribution algorithms such as SHAP have been increasingly applied in MPM to explain which features drive prospectivity predictions (Zuo et al., 2024), these applications have been limited to explaining prospectivity scores without examining what drives model uncertainty. The dual SHAP analysis developed in Chapter 6 extends this by applying SHAP to both predictions and their associated uncertainties. Specifically, the SHAP analysis of model uncertainty indicates that the presence of mafic host rocks alone does not result in confident predictions. When other key mineral system components, such as structural pathways or sulphur sources, are absent, the model's uncertainty increases, even in the presence of favourable host lithologies. This reinforces the fundamental geological principle that mineralization is a rare occurrence which requires the convergence of all critical mineral system components, and the absence of even a single component significantly reduces the likelihood of mineralization.

7.3 Limitations

While the individual limitations of each study are discussed within their respective chapters, two common limitations affect all three articles. First, all models assume that non-deposit H3 hexagonal cells are barren. At the national scale, comprehensive drill-hole databases are unavailable to confirm whether a given cell is truly barren or simply unexplored. This introduces label noise into the training data that can affect model performance, particularly

in underexplored regions where undiscovered deposits may exist at depth or beneath cover. The uncertainty maps produced in Chapters 5 and 6 partially mitigate this by flagging regions with low model confidence, but the assumption itself remains a source of bias across all models.

Second, most data-driven MPM studies operate at local-to-regional scales, where data density is higher, geological variability is more constrained, and input features more closely capture the geological signatures associated with mineralization. At the national scale, models must perform across diverse geological provinces with varying levels of data coverage, where features strongly associated with mineralization in one province may not generalize to another. Additionally, a scale mismatch exists between the national-scale input data and the deposit-scale geological processes that control mineralization. While the mineral system's sources and pathways can be mapped at regional scales using geological and geophysical data, the trapping mechanisms that control where sulphide-rich fluids accumulate operate at scales below the resolution of available datasets. This remains a primary driver of model uncertainty for the trap component of the mineral system, which is the most difficult to map at the national scale.

7.4 Original contributions of the thesis

As highlighted by Caers (2025), risk-informed decision-making is essential for effective mineral exploration. The risk-return framework developed in Chapter 5 addresses this by categorizing exploration targets based on both prospectivity and model confidence. Additionally, the SHAP analysis from Chapter 6 can be applied to specific prospective target clusters to identify which geological features drive both high prospectivity and model uncertainty, thereby helping prioritize which additional data to acquire at each target area. These frameworks, while applied to Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide systems, are adaptable to other deposit types and geological settings. Furthermore, the TAGI and SHAP uncertainty implementations from Chapter 6 are available as open-source Python packages, with the source code and documentation accessible through the repository listed in Chapter 6.

This thesis also demonstrates that probabilistic ML frameworks can be successfully applied at the national scale, covering the entire Canadian landmass. Despite the challenges of national-scale modelling discussed above, the results from all three studies confirm that model generalization, class imbalance, and uncertainty characterization can be effectively addressed even across across varying geological settings. This opens pathways for similar large-scale prospectivity assessments in other countries and across different mineral systems.

CHAPTER 8 CONCLUSION

8.1 Summary

This thesis aimed to develop and validate probabilistic ML frameworks for national-scale MPM of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide systems. To achieve this, three progressive steps form the core of this research. First, a generative data augmentation framework addressed class imbalance in training datasets. Next, this framework extended to quantify prediction uncertainty and introduced a risk-return analysis for exploration targeting. Finally, an analytical Bayesian inference method provided both uncertainty estimation and an interpretability framework that explains predictions and their associated uncertainties.

The first objective asks whether generative models can effectively address class imbalance in MPM. The thesis demonstrates that data balancing can yield improved prospectivity predictions and that the CVAE framework outperforms conventional balancing techniques such as SMOTE and class weighting across multiple classification algorithms. The resulting prospectivity maps capture known Ni deposits and identify new prospective regions in underexplored parts of Canada. While data augmentation consistently improves predictive performance, partial overlap between real and synthetic distributions for some geological features suggests that further refinement of the generative process could yield additional improvements.

The next objective addresses the quantification of uncertainty in MPM predictions. First, aleatoric uncertainty is assessed using the CVAE's probabilistic outputs, while epistemic uncertainty is captured by disagreement among an ensemble of models trained on distinct augmented datasets. A novel risk-return framework is then proposed to prioritize exploration targets by combining these uncertainty measures with spatial clustering to categorize targets based on both prospectivity and confidence. While this demonstrates the value of jointly evaluating prospectivity and uncertainty, the ensemble-based uncertainty maps exhibited spatially noisy patterns, which motivates the analytical approach explored in the third study.

Finally, we implement an analytical Bayesian inference method that propagates probability distributions through the network without gradient-based optimization. TAGI outperforms BNN with VI and MC dropout across all evaluation criteria, including predictive accuracy, recall, training speed, and robustness to noise, while producing spatially coherent uncertainty maps that address the noise limitation of the ensemble approach. A dual SHAP analysis, applied to both predictions and uncertainty, reveals that lithological and geochronological

features drive both prospectivity and model confidence, and that uncertainty increases at locations where only some mineral system components are present. Therefore, this interpretability framework enables validation of model outputs against geological knowledge and guides data acquisition in specific target areas.

Collectively, these three studies fulfil the general objective of this thesis by expanding MPM from deterministic prospectivity maps to probabilistic decision-support tools that quantify prediction confidence, prioritize exploration targets, and identify which geological factors control both predictions and their associated uncertainties. The findings and frameworks presented in this thesis open several directions for future research and broader applications beyond the scope of the current work.

8.2 Recommendations for future work

Based on the experience gained through this research, several directions for future work are identified. First, careful data preprocessing and quality assessment remain critical for any MPM application. Increasing the number of input features does not necessarily improve model performance, and meaningful prospectivity predictions depend on the careful selection and integration of relevant geoscience evidence layers.

A key limitation identified across all three studies is the assumption that non-deposit locations are barren. Unsupervised clustering conditioned on geological features could address this by grouping locations into geologically distinct domains prior to supervised modelling. Based on the results obtained through this thesis for the national scale, lithology and geochronology data dominate prospectivity predictions. Therefore, geological domains that do not satisfy the source requirements of the mineral system, such as regions where mafic-ultramafic lithologies are absent, could be confidently labelled as non-prospective, reducing label noise and improving negative sample selection.

Next, integrating large language models (LLMs) into prospectivity workflows is another promising direction. Geological map databases contain rich textual descriptions of rock types, lithostratigraphic relationships, and field observations that are currently reduced to categorical codes, losing significant geological information. LLMs trained on geological text could extract richer, context-aware representations from these descriptions, including geological relationships relevant to ore deposit formation, and integrate them as input features for prospectivity models, potentially capturing geological nuances that categorical encoding cannot represent.

Additionally, this thesis quantified aleatoric and epistemic uncertainty but did not address

conceptual uncertainty arising from an incomplete understanding of the ore-forming process itself. One possible approach could involve integrating expert knowledge through weighted mineral system components, in which the confidence assigned to each process, such as source, pathway, and trap, reflects the current level of geological understanding. However, how to incorporate such process-level uncertainty into MPM workflows remains an open challenge.

Finally, mineral deposit formation is governed by physicochemical processes such as magma transport, sulphide segregation, and structural controls on fluid flow, all of which operate under known physical and chemical laws. Traditional ML models learn patterns solely by minimizing the discrepancy between predictions and observed data, without any constraints imposed by the underlying processes. Physics-informed neural networks (PINNs) can be used to address this by embedding physical laws directly into the training process, which ensures that predictions remain consistent with the science of ore formation. While PINNs have not yet been applied to MPM, integrating such constraints could improve the geological plausibility of prospectivity predictions, particularly in regions where the model would otherwise rely entirely on data-driven patterns. However, physical modelling of ore deposit formation requires complex mechanical, thermodynamic, and hydrological modelling, and therefore, how to effectively implement PINNs into MPM remains an open question.

REFERENCES

Abdellatif, A., Mubarak, H., Abdellatef, H., Kanesan, J., Abdelltif, Y., Chow, C.-O., Chuah, J. H., Gheni, H. M., and Kendall, G. (2024). Computational detection and interpretation of heart disease based on conditional variational auto-encoder and stacked ensemble-learning framework. *Biomedical Signal Processing and Control*, 88:105644.

Abedi, M., Norouzi, G.-H., and Bahroudi, A. (2012). Support vector machine for multi-classification of mineral prospectivity areas. *Computers & Geosciences*, 46:272–283.

Agterberg, F., Bonham-Carter, G., Cheng, Q.-m., and Wright, D. (1993). Weights of evidence modelling and weighted logistic regression for mineral potential mapping. *Computers in Geology*, 25:13–32.

Agterberg, F. P. (1989). Systematic approach to dealing with uncertainty of geoscience information in mineral exploration. *APCO 89*, pages 165–178.

Agterberg, F. P. and Bonham-Carter, G. F. (2005). Measuring the performance of mineral-potential maps. *Natural Resources Research*, 14:1–17.

Agterberg, F. P. and Cheng, Q. (2002). Conditional independence test for weights-of-evidence modeling. *Natural Resources Research*, 11(4):249–255.

Ali, S. H., Giurco, D., Arndt, N., Nickless, E., Brown, G., Demetriades, A., Durrheim, R., Enriquez, M. A., Kinnaird, J., Littleboy, A., et al. (2017). Mineral supply for sustainable development requires resource governance. *Nature*, 543(7645):367–372.

Almasi, A., Yousefi, M., and Carranza, E. J. M. (2017). Prospectivity analysis of orogenic gold deposits in Saqez-Sardasht goldfield, Zagros Orogen, Iran. *Ore Geology Reviews*, 91:1066–1080.

Alvarez Melis, D. and Jaakkola, T. (2018). Towards robust interpretability with self-explaining neural networks. *Advances in Neural Information Processing Systems*, 31.

Arndt, N., Lesher, C., and Czamanske, G. (2005). Mantle-derived magmas and magmatic Ni-Cu-(PGE) deposits. *Economic Geology 100th Anniversary Volume*.

Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., and Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140.

Barber, D. and Bishop, C. M. (1998). Ensemble learning in Bayesian neural networks. *Nato ASI Series F Computer and Systems Sciences*, 168:215–238.

Bárdossy, G. and Fodor, J. (2001). Traditional and new ways to handle uncertainty in geology. *Natural Resources Research*, 10:179–187.

Barnes, S.-J. (1990). The use of metal ratios in prospecting for platinum-group element deposits in mafic and ultramafic intrusions. *Journal of Geochemical Exploration*, 37(1):91–99.

Barnes, S. J. (2006). Komatiite-hosted nickel sulfide deposits: Geology, geochemistry, and genesis.

Barnes, S. J., Cruden, A. R., Arndt, N., and Saumur, B. M. (2016). The mineral system approach applied to magmatic Ni–Cu–PGE sulphide deposits. *Ore Geology Reviews*, 76:296–316.

Barnes, S. J. and Fiorentini, M. L. (2012). Komatiite magmas and sulfide nickel deposits: a comparison of variably endowed Archean terranes. *Economic Geology*, 107(5):755–780.

Bazilian, M., Bradshaw, M., Goldthau, A., and Westphal, K. (2019). Model and manage the changing geopolitics of energy. *Nature*, 569(7754):29–31.

Begg, G. C., Hronsky, J. M., Griffin, W. L., and O'Reilly, S. Y. (2018). Global-to deposit-scale controls on orthomagmatic Ni-Cu (-PGE) and PGE reef ore formation. In *Processes and ore deposits of ultramafic-mafic magmas through space and time*, pages 1–46. Elsevier.

Bérubé, C. L. and Baron, F. (2023). Bayesian inference of petrophysical properties with generative spectral induced polarization models. *Geophysics*, 88(3):E79–E90.

Bérubé, C. L. and Bérubé, P. (2022). Data-driven modeling of time-domain induced polarization. *Geophysics*, 87(3):E135–E146.

Bigdeli, A., Maghsoudi, A., and Ghezelbash, R. (2024). A comparative study of the XGBoost ensemble learning and multilayer perceptron in mineral prospectivity modeling: A case study of the Torud-Chahshirin belt, NE Iran. *Earth Science Informatics*, 17(1):483–499.

Billingsley, P. and Locke, A. (1933). Tectonic position of ore districts in the Rocky Mountain region. American Institute of Mining and Metallurgical Engineers.

Bleeker, W. and Kamo, S. (2020). Structural-stratigraphic Setting and U-Pb Geochronology of Ni-Cu-Co-PGE Ore Environments in the Central Cape Smith Belt, Circum-Superior Belt. Technical report.

Blewett, R., Henson, P., Roy, I., Champion, D., and Cassidy, K. (2010). Scale-integrated architecture of a world-class gold mineral system: the Archaean eastern Yilgarn Craton, Western Australia. *Precambrian Research*, 183(2):230–250.

Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. (2015). Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR.

Böhm, V., Lanusse, F., and Seljak, U. (2019). Uncertainty quantification with generative models. *arXiv preprint arXiv:1910.10046*.

Bonham-Carter, G. (1994). Geographic information systems for geoscientists: Modelling with GIS. Number 13. Elsevier.

Bonham-Carter, G. and Agterberg, F. (1990). Application of a microcomputer-based geographic information system to mineral-potential mapping. In *Microcomputer Applications in Geology 2*, pages 49–74. Elsevier.

Bonham-Carter, G. F., Agterberg, F. P., and Wright, D. F. (1989). Integration of geological datasets for gold exploration in Nova Scotia. *Digital Geologic and Geographic Information Systems*, 10:15–23.

Bourlard, H. and Kamp, Y. (1988). Auto-association by multilayer perceptrons and singular value decomposition. *Biological Cybernetics*, 59(4-5):291–294.

Bradley, D. C. (2008). Passive margins through Earth history. *Earth-Science Reviews*, 91(1-4):1–26.

Breiman, L. (2001). Random forests. *Machine Learning*, 45:5–32.

Brodaric, B., Gahegan, M., and Harrap, R. (2004). The art and science of mapping: computing geological categories from field data. *Computers & Geosciences*, 30(7):719–740.

Brown, W. M., Gedeon, T., Groves, D., and Barnes, R. (2000). Artificial neural networks: A new method for mineral prospectivity mapping. *Australian Journal of Earth Sciences*, 47(4):757–770.

Bruce, M., Fairclough, M., Jareith, S., Kreuzer, O., Mihalasky, M., and Poliakovska, K. (2020). Descriptive uranium deposit and mineral system models. International Atomic Energy Agency.

Bruzzone, L., Conese, C., Maselli, F., Roli, F., et al. (1997). Multisource classification of complex rural areas by statistical and neural-network approaches. *Photogrammetric Engineering and Remote Sensing*, 63(5):523–532.

Caers, J. (2011). Modeling Uncertainty in the Earth Sciences. John Wiley & Sons.

Caers, J. (2025). The future of ai in critical mineral exploration.

Campbell, A., Hollister, V., Duda, R., and Hart, P. (1982). Recognition of a hidden mineral deposit by an artificial intelligence program. *Science*, 217(4563):927–929.

Carranza, E. and Hale, M. (2008). Handbook of exploration and environmental geochemistry. *Geochemical Anomaly and Mineral Prospectivity Mapping in GIS; Elsevier: Amsterdam, The Netherlands*, 11.

Carranza, E., Hale, M., and Faassen, C. (2008). Selection of coherent deposit-type locations and their application in data-driven mineral prospectivity mapping. *Ore Geology Reviews*, 33(3-4):536–558.

Carranza, E. J. M. (2008). Geochemical anomaly and mineral prospectivity mapping in GIS, volume 11. Elsevier.

Carranza, E. J. M. (2017). Natural resources research publications on geochemical anomaly and mineral potential mapping, and introduction to the special issue of papers in these fields. *Natural Resources Research*, 26(4):379–410.

Carranza, E. J. M. and Laborte, A. G. (2015). Random forest predictive modeling of mineral prospectivity with a small number of prospects and data with missing values in Abra (Philippines). *Computers & Geosciences*, 74:60–70.

Carranza, E. J. M. and Laborte, A. G. (2016). Data-driven predictive modeling of mineral prospectivity using random forests: A case study in Catanduanes Island (Philippines). *Natural Resources Research*, 25(1):35–50.

Champion, D. C. and Huston, D. L. (2016). Radiogenic isotopes, ore deposits and metallogenic terranes: Novel approaches based on regional isotopic maps and the mineral systems concept. *Ore Geology Reviews*, 76:229–256.

Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. (2002). Smote: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357.

Chawla, N. V., Japkowicz, N., and Kotcz, A. (2004). Special issue on learning from imbalanced data sets. *ACM SIGKDD explorations newsletter*, 6(1):1–6.

Chawla, N. V., Lazarevic, A., Hall, L. O., and Bowyer, K. W. (2003). SMOTEBoost: Improving prediction of the minority class in boosting. In *European conference on principles of data mining and knowledge discovery*, pages 107–119. Springer.

Chen, M. and Xiao, F. (2023). Projection pursuit random forest for mineral prospectivity mapping. *Mathematical Geosciences*, pages 1–25.

Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794.

Chen, Y., Kempton, D. J., Ahmadzadeh, A., Wen, J., Ji, A., and Angryk, R. A. (2022). CGAN-based synthetic multivariate time-series generation: A solution to data scarcity in solar flare forecasting. *Neural Computing and Applications*, 34(16):13339–13353.

Cheng, Q. (2007). Mapping singularities with stream sediment geochemical data for prediction of undiscovered mineral deposits in Gejiu, Yunnan Province, China. *Ore Geology Reviews*, 32(1-2):314–324.

Cheng, Q. (2012). Singularity theory and methods for mapping geochemical anomalies caused by buried sources and for predicting undiscovered mineral deposits in covered areas. *Journal of Geochemical Exploration*, 122:55–70.

Cheng, Q. and Agterberg, F. (1999). Fuzzy weights of evidence method and its application in mineral potential mapping. *Natural Resources Research*, 8(1):27–35.

Cheng, X. and Jiang, K. (2020). Crustal model in eastern Qinghai-Tibet plateau and western Yangtze craton based on conditional variational autoencoder. *Physics of the Earth and Planetary Interiors*, 309:106584.

China Geological Survey (2021). Global mining development report 2021.

Chorlton, L. (2007). Generalized geology of the world: bedrock domains and major faults in GIS format. *Geological Survey of Canada, Open File*, 5529(1).

Chung, C. and Agterberg, F. (1980). Regression models for estimating mineral resources from geological map data. *Journal of the International Association for Mathematical Geology*, 12(5):473–488.

Chung, C.-J. F. and Fabbri, A. G. (1993). The representation of geoscience information for data integration. *Nonrenewable Resources*, 2(2):122–139.

Crafford, A. and Grauch, V. (2002). Geologic and geophysical evidence for the influence of deep crustal structures on Paleozoic tectonics and the alignment of world-class gold deposits, north-central Nevada, USA. *Ore Geology Reviews*, 21(3-4):157–184.

Cross, B. and Peterson, H. (2023). A comparison of airborne geophysical data over two magmatic nickel deposits. *The Leading Edge*, 42(4):237–244.

Czarnota, K., Blewett, R., and Goscombe, B. (2010). Predictive mineral discovery in the eastern Yilgarn Craton, Western Australia: An example of district scale targeting of an orogenic gold mineral system. *Precambrian Research*, 183(2):356–377.

de Araújo, C. C. and Macedo, A. B. (2002). Multicriteria geologic data analysis for mineral favorability mapping: Application to a metal sulphide mineralized area, Ribeira Valley Metallogenic Province, Brazil. *Natural Resources Research*, 11(1):29–43.

de Quadros, T. F., Koppe, J. C., Strieder, A. J., and Costa, J. F. (2006). Mineral-potential mapping: a comparison of weights-of-evidence and fuzzy methods. *Natural Resources Research*, 15(1):49–65.

de Sá, C. R. (2019). Variance-based feature importance in neural networks. In *International Conference on Discovery Science*, pages 306–315. Springer.

Debayle, E., Bodin, T., Durand, S., and Ricard, Y. (2020). Seismic evidence for partial melt below tectonic plates. *Nature*, 586(7830):555–559.

Decatur, S. E. (1989). Application of neural networks to terrain classification. *Proceedings of IJCNN (Washington, DC, USA)*, 1:283–288.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.

Duda, R. O., Hart, P. E., Nilsson, N. J., and Sutherland, G. L. (1977). Semantic network representations in rule-based inference systems. *ACM SIGART Bulletin*, (63):18–18.

Ebbing, J., Haas, P., Ferraccioli, F., Pappa, F., Szwillus, W., and Bouman, J. (2018). Earth tectonics as seen by GOCE-Enhanced satellite gravity gradient imaging. *Scientific Reports*, 8(1):16356.

Eckstrand, O. R. and Hulbert, L. J. (2007). Magmatic nickel-copper-platinum group element deposits. *Mineral deposits of Canada: a synthesis of major deposit types, district metallogeny, the evolution of geological provinces, and exploration methods: Geological Association of Canada, Mineral Deposits Division, Special Publication*, 5:205–222.

Edupuganti, V., Mardani, M., Vasanawala, S., and Pauly, J. (2020). Uncertainty quantification in deep MRI reconstruction. *IEEE Transactions on Medical Imaging*, 40(1):239–250.

Eglington, B., Pehrsson, S., Ansdell, K., Lescuyer, J.-L., Quirt, D., Milesi, J.-P., and Brown, P. (2013). A domain-based digital summary of the evolution of the Palaeoproterozoic of North America and Greenland and associated unconformity-related Uranium mineralization. *Precambrian Research*, 232:4–26.

Elitok, Ö. and Tavlan, M. (2019). Geology and economic potential of Ni deposits. *Mineral Resources of Turkey*, pages 635–653.

Environment and Climate Change Canada (2021). National Inventory Report 1990–2019: Greenhouse Gas Sources and Sinks in Canada, Part 1. Technical report.

Fallon, M., Porwal, A., and Guj, P. (2010). Prospectivity analysis of the plutonic Marymia greenstone belt, Western Australia. *Ore Geology Reviews*, 38(3):208–218.

Fan, C., Chen, M., Tang, R., and Wang, J. (2022). A novel deep generative modeling-based data augmentation strategy for improving short-term building energy predictions. In *Building Simulation*, volume 15, pages 197–211. Springer.

Farahbakhsh, E., Maughan, J., and Müller, R. D. (2023). Prospectivity modelling of critical mineral deposits using a generative adversarial network with oversampling and positive-unlabelled bagging. *Ore Geology Reviews*, 162:105665.

Farahnakian, F., Sheikh, J., Zelioli, L., Nidhi, D., Seppä, I., Ilo, R., Nevalainen, P., and Heikkonen, J. (2024). Addressing imbalanced data for machine learning based mineral prospectivity mapping. *Ore Geology Reviews*, 174:106270.

Feltrin, L. (2008). Predictive modelling of prospectivity for Pb–Zn deposits in the Lawn Hill Region, Queensland, Australia. *Ore Geology Reviews*, 34(3):399–427.

Ford, A., Miller, J. M., and Mol, A. G. (2016). A comparative analysis of weights of evidence, evidential belief functions, and fuzzy logic for mineral potential mapping using incomplete data at the scale of investigation. *Natural Resources Research*, 25:19–33.

Ford, A., Peters, K., Partington, G., Blevin, P., Downes, P., Fitzherbert, J., and Greenfield, J. (2019). Translating expressions of intrusion-related mineral systems into mappable spatial proxies for mineral potential mapping: Case studies from the Southern New England Orogen, Australia. *Ore Geology Reviews*, 111:102943.

Fu, C., Chen, K., Yang, Q., Chen, J., Wang, J., Liu, J., Xiang, Y., Li, Y., and Rajesh, H. (2021). Mapping gold mineral prospectivity based on the weights of evidence method in southeast Asmara, Eritrea. *Journal of African Earth Sciences*, 176:104143.

Fuentes, I., Padarian, J., Iwanaga, T., and Vervoort, R. W. (2020). 3d lithological mapping of borehole descriptions using word embeddings. *Computers & Geosciences*, 141:104516.

Gadd, M. G., Lawley, C. J., Corriveau, L., Houlé, M., Peter, J. M., Plouffe, A., Potter, E., Sappin, A.-A., Pilote, J.-L., Marquis, G., et al. (2023). Public geoscience solutions for diversifying Canada's critical mineral production. *Geological Society, London, Special Publications*, 526(1):SP526–2021.

Gal, Y. and Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059. PMLR.

Gao, M., Wang, G., Carranza, E. J. M., Qi, S., Zhang, W., Pang, Z., Li, X., and Xiao, F. (2024). 3d Au targeting using machine learning with different sample combinations and return-risk analysis in the Sanshandao-Cangshang District, Shandong Province, China. *Natural Resources Research*, 33(1):51–74.

Getis, A. and Ord, J. K. (1992). The analysis of spatial association by use of distance statistics. *Geographical Analysis*, 24(3):189–206.

Ghezelbash, R., Maghsoudi, A., Bigdeli, A., and Carranza, E. J. M. (2021). Regional-scale mineral prospectivity mapping: Support vector machines and an improved data-driven multi-criteria decision-making technique. *Natural Resources Research*, 30(3):1977–2005.

Ghezelbash, R., Maghsoudi, A., Shamekhi, M., Pradhan, B., and Daviran, M. (2023). Genetic algorithm to optimize the SVM and K-means algorithms for mapping of mineral prospectivity. *Neural Computing and Applications*, 35(1):719–733.

Goldstein, A., Kapelner, A., Bleich, J., and Pitkin, E. (2015). Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *Journal of Computational and Graphical Statistics*, 24(1):44–65.

Gong, X., Tang, B., Zhu, R., Liao, W., and Song, L. (2020). Data augmentation for electricity theft detection using conditional variational auto-encoder. *Energies*, 13(17):4291.

Good, D. J., Eckstrand, O. R., Yakubchuk, A., and Gall, Q. (2015). World Ni-Cu-PGE-Cr deposit database. In *World Ni-Cu-PGE-Cr deposit database*.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. `http://www.deeplearningbook.org`.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. *Advances in neural information processing systems*, 27.

Goulet, J.-A., Nguyen, L. H., and Amiri, S. (2021). Tractable approximate Gaussian inference for Bayesian neural networks. *Journal of Machine Learning Research*, 22(251):1–23.

Government of Canada (2021). Canadian Net-Zero Emissions Accountability Act.

Government of Canada (2024). Canadian Critical Minerals Strategy Annual Report 2024.

Granitto, M., Giles, S., and Kelley, K. (2017). Global geochemical database for critical metals in black shales. *US Geological Survey data release. https://doi. org/10.5066/F71G0K7X*.

Grieve, R. A., Stoeffler, D., and Deutsch, A. (1991). The Sudbury structure: Controversial or misunderstood? *Journal of Geophysical Research: Planets*, 96(E5):22753–22764.

Griffin, W., Begg, G., and O'Reilly, S. Y. (2013). Continental-root control on the genesis of magmatic ore deposits. *Nature Geoscience*, 6(11):905–910.

Groves, D. I., Santosh, M., Müller, D., Zhang, L., Deng, J., Yang, L.-Q., and Wang, Q.-F. (2022). Mineral systems: Their advantages in terms of developing holistic genetic models and for target generation in global mineral exploration. *Geosystems and Geoenvironment*, 1(1):100001.

Hábel, B., Janák, J., Papčo, J., and Val'ko, M. (2020). Impact of environmental phenomena on continuous relative gravity measurements performed in urban area. *Studia Geophysica et Geodaetica*, 64:330–348.

Hagemann, S. G., Lisitsin, V., and Huston, D. (2016). Mineral system analysis: Quo vadis. *Ore Geology Reviews*, 76:504–522.

Hajihosseinlou, M., Maghsoudi, A., and Ghezelbash, R. (2024). Stacking: A novel data-driven ensemble machine learning strategy for prediction and mapping of Pb-Zn prospectivity in Varcheh district, west Iran. *Expert Systems with Applications*, 237:121668.

Hajihosseinlou, M., Maghsoudi, A., and Ghezelbash, R. (2025). A semi-supervised approach for mineral prospectivity mapping via weighted positive-unlabeled learning and tree-structured parzen estimator for hyperparameter optimization. *Ore Geology Reviews*, page 106783.

Hariharan, S., Tirodkar, S., Porwal, A., Bhattacharya, A., and Joly, A. (2017). Random forest-based prospectivity modelling of greenfield terrains using sparse deposit data: An example from the Tanami Region, Western Australia. *Natural Resources Research*, 26(4):489–507.

Harris, D., Zurcher, L., Stanley, M., Marlow, J., and Pan, G. (2003). A comparative analysis of favorability mappings by weights of evidence, probabilistic neural networks, discriminant analysis, and logistic regression. *Natural Resources Research*, 12:241–255.

Haykin, S. (2004). *Kalman filtering and neural networks*. John Wiley & Sons.

He, H., Bai, Y., Garcia, E. A., and Li, S. (2008). ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence)*, pages 1322–1328. IEEE.

He, H., Zhu, H., Yang, X., Zhang, W., and Wang, J. (2024). Mineral prospectivity prediction based on convolutional neural network and ensemble learning. *Scientific Reports*, 14(1):22654.

He, W., Jiang, Z., Xiao, T., Xu, Z., and Li, Y. (2023). A survey on uncertainty quantification methods for deep learning. *arXiv preprint arXiv:2302.13425*.

Herrington, R. J. (2011). Geological features and genetic models of mineral deposits. In *SME Mining Engineering Handbook, (3rd edition)*, pages 83–104.

Hinton, G. E. and Van Camp, D. (1993). Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 5–13.

Hofstra, A. H., Lisitsin, V., Corriveau, L., Paradis, S., Peter, J., Lauzière, K., Lawley, C., Gadd, M., Pilote, J.-L., Honsberger, I., et al. (2021). Deposit classification scheme for

the critical minerals mapping initiative global geochemical database. Technical report, US Geological Survey.

Hong, S., An, S., and Jeon, J.-J. (2024). Improving SMOTE via Fusing Conditional VAE for Data-adaptive Noise Filtering. *arXiv preprint arXiv:2405.19757*.

Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366.

Hronsky, J. (2011). Self-organized critical systems and ore formation: The key to spatial targeting. *Society of Economic Geology Newsletter*, 84:14–16.

Hronsky, J. M. (2009). The exploration search space concept: key to a successful exploration strategy. *CET Q. Newsl*, 8:14–15.

Hronsky, J. M. and Groves, D. I. (2008). Science of targeting: definition, strategies, targeting and performance measurement. *Australian Journal of Earth Sciences*, 55(1):3–12.

Hronsky, J. M. and Kreuzer, O. P. (2019). Applying spatial prospectivity mapping to exploration targeting: Fundamental practical issues and suggested solutions for the future. *Ore Geology Reviews*, 107:647–653.

Huang, D., Zuo, R., Wang, J., and Tolosana-Delgado, R. (2025). Combining sequential Gaussian co-simulation and Monte Carlo dropout-based deep learning models for geochemical anomaly detection and uncertainty assessment. *Applied Geochemistry*, 184:106385.

Hüllermeier, E. and Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3):457–506.

Huston, D. L., Mernagh, T. P., Hagemann, S. G., Doublier, M. P., Fiorentini, M., Champion, D. C., Jaques, A. L., Czarnota, K., Cayley, R., Skirrow, R., et al. (2016). Tectonometallogenic systems — The place of mineral systems within tectonic evolution, with an emphasis on Australian examples. *Ore Geology Reviews*, 76:168–210.

Imasu, R., Chang, L.-H., and Tsai, F. (2024). Multispectral, hyperspectral, and ultraspectral remote sensing technology, techniques, and applications viii. In *Proc. of SPIE Vol*, volume 13266, pages 1326601–1.

International Energy Agency (2021). The Role of Critical Minerals in Clean Energy Transitions. Technical report, IEA, Paris.

International Energy Agency (2025). Global Critical Minerals Outlook 2025. https://www.iea.org/reports/global-critical-minerals-outlook-2025.

Janjoš, F., Hallgarten, M., Knittel, A., Dolgov, M., Zell, A., and Zöllner, J. M. (2023). Conditional unscented autoencoders for trajectory prediction. *arXiv preprint arXiv:2310.19944*.

Jeffrey, W. and Rosner, R. (1986). Optimization algorithms-simulated annealing and neural network processing. *The Astrophysical Journal*, 310:473–481.

Jiang, Z., Pan, T., Zhang, C., and Yang, J. (2021). A new oversampling method based on the classification contribution degree. *Symmetry*, 13(2):194.

Jobin, D. M., Miles, W., and Véronneau, M. (2017). *Isostatic Residual Gravity Anomaly Map, Canada: Carte Des Anomalies Isostatiques Résiduelles Du Champ de Gravité, Canada*. Natural Resources Canada= Ressources naturelles Canada.

Johnson, D. B. (2018). The evolution, current status, and future prospects of using biotechnologies in the mineral extraction and metal recovery sectors. *Minerals*, 8(8):343.

Jordao, H., Sousa, A. J., and Soares, A. (2023). Using Bayesian neural networks for uncertainty assessment of ore type boundaries in complex geological models. *Natural Resources Research*, 32(6):2495–2514.

Jospin, L. V., Laga, H., Boussaid, F., Buntine, W., and Bennamoun, M. (2022). Hands-on Bayesian neural networks—A tutorial for deep learning users. *IEEE Computational Intelligence Magazine*, 17(2):29–48.

Karpatne, A., Atluri, G., Faghmous, J. H., Steinbach, M., Banerjee, A., Ganguly, A., Shekhar, S., Samatova, N., and Kumar, V. (2017). Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, 29(10):2318–2331.

Katiyar, N., Kulshreshtha, A., and Singh, P. (2024). A review of integration techniques of multi-geoscience data-sets in mineral prospectivity mapping. *Earth Sciences*, 13(4):127–140.

Kingma, D. and Welling, M. (2013). Auto-encoding variational bayes. In *The 2nd International Conference on Learning Representations (ICLR)*.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kingma, D. P., Salimans, T., and Welling, M. (2015). Variational dropout and the local reparameterization trick. *Advances in neural information processing systems*, 28.

Knox-Robinson, C. and Wyborn, L. (1997). Towards a holistic exploration strategy: Using geographic information systems as a tool to enhance exploration. *Australian Journal of Earth Sciences*, 44(4):453–463.

Kong, C., Zhao, J., Li, B., Wu, C., and Xu, K. (2025). Manganese mineral prospectivity mapping based on semi-supervised learning and multi-source geoscientific-sample Wasserstein generative adversarial network (Geo-WGAN) in Songtao of Guizhou, South China. *Ore Geology Reviews*, page 106933.

Kramer, M. A. (1991). Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, 37(2):233–243.

Kreuzer, O., Blenkinsop, T. G., Morrison, R., and Peters, S. (2007). Ore controls in the Charters Towers goldfield, NE Australia: Constraints from geological, geophysical and numerical analyses. *Ore Geology Reviews*, 32(1-2):37–80.

Kreuzer, O. P. and Etheridge, M. A. (2010). Risk and uncertainty in mineral exploration: implications for valuing mineral exploration properties. *AIG News*, 100:20–28.

Kreuzer, O. P., Etheridge, M. A., Guj, P., McMahon, M. E., and Holden, D. J. (2008). Linking mineral deposit models to quantitative risk analysis and decision-making in exploration. *Economic Geology*, 103(4):829–850.

Kreuzer, O. P., Miller, A. V., Peters, K. J., Payne, C., Wildman, C., Partington, G. A., Puccioni, E., McMahon, M. E., and Etheridge, M. A. (2015). Comparing prospectivity modelling results and past exploration data: A case study of porphyry Cu–Au mineral systems in the Macquarie Arc, Lachlan Fold Belt, New South Wales. *Ore Geology Reviews*, 71:516–544.

Kreuzer, O. P., Yousefi, M., and Nykänen, V. (2020). Introduction to the special issue on spatial modelling and analysis of ore-forming processes in mineral exploration targeting. *Ore Geology Reviews*, 119:103391.

Krishnan, R., Esposito, P., and Subedar, M. (2022). Bayesian-Torch: Bayesian neural network layers for uncertainty estimation. `https://github.com/IntelLabs/bayesian-torch`.

Krisper, M. (2021). Problems with risk matrices using ordinal scales. *arXiv preprint arXiv:2103.05440*.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2017). Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90.

Kumar, A., Singh, D., and Shankar Yadav, R. (2024). Class overlap handling methods in imbalanced domain: A comprehensive survey. *Multimedia Tools and Applications*, pages 1–48.

Lary, D. J., Alavi, A. H., Gandomi, A. H., and Walker, A. L. (2016). Machine learning in geosciences and remote sensing. *Geoscience Frontiers*, 7(1):3–10.

Latypov, R., Chistyakova, S., Grieve, R., and Huhma, H. (2019). Evidence for igneous differentiation in sudbury igneous complex and impact-driven evolution of terrestrial planet proto-crusts. *Nature communications*, 10(1):508.

Lauzon, D. and Gloaguen, E. (2024). Quantifying uncertainty and improving prospectivity mapping in mineral belts using transfer learning and Random Forest: A case study of copper mineralization in the Superior Craton Province, Quebec, Canada. *Ore Geology Reviews*, 166:105918.

Lawley, C., Giddy, P., Katz, L., Chu, N., Francis, A., Carvajal, J., Pinheiro, M., Tirona, K., Dettman, A., Chen, H., et al. (2024). Canada geological map compilation. *Geological Survey of Canada*.

Lawley, C. J., Gadd, M. G., Parsa, M., Lederer, G. W., Graham, G. E., and Ford, A. (2023). Applications of natural language processing to geoscience text data and prospectivity modeling. *Natural Resources Research*, 32(4):1503–1527.

Lawley, C. J., McCafferty, A. E., Graham, G. E., Huston, D. L., Kelley, K. D., Czarnota, K., Paradis, S., Peter, J. M., Hayward, N., Barlow, M., et al. (2022a). Data–driven prospectivity modelling of sediment–hosted Zn–Pb mineral systems and their critical raw materials. *Ore Geology Reviews*, 141:104635.

Lawley, C. J., Mitchell, M. G., Stralberg, D., Schuster, R., McIntire, E., and Bennett, J. R. (2022b). Mapping Canada's green economic pathways for battery minerals: Balancing prospectivity modelling with conservation and biodiversity values. *Earth Science, Systems and Society*, 2:10064.

Lawley, C. J., Raimondo, S., Chen, T., Brin, L., Zakharov, A., Kur, D., Hui, J., Newton, G., Burgoyne, S. L., and Marquis, G. (2022c). Geoscience language models and their intrinsic evaluation. *Applied Computing and Geosciences*, 14:100084.

Lawley, C. J., Tschirhart, V., Smith, J. W., Pehrsson, S. J., Schetselaar, E. M., Schaeffer, A. J., Houlé, M. G., and Eglington, B. M. (2021). Prospectivity modelling of Canadian magmatic Ni (±Cu±Co±PGE) sulphide mineral systems. *Ore Geology Reviews*, 132:103985.

Laxton, J. (2017). Geological map fusion: Onegeology-europe and inspire. *Special Publications*, 408(1):147–160.

Layton-Matthews, D., Lesher, C., Burnham, O., Hulbert, L., Peck, D., Golightly, J., and Keays, R. (2010). Exploration for komatiite-associated Ni-Cu-(PGE) mineralization in the Thompson nickel belt, Manitoba.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, 521(7553):436–444.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.

Leite, E. P. and de Souza Filho, C. R. (2009). Probabilistic neural networks applied to mineral potential mapping for platinum group elements in the Serra Leste region, Carajás Mineral Province, Brazil. *Computers & Geosciences*, 35(3):675–687.

Lemaître, G., Nogueira, F., and Aridas, C. K. (2017). Imbalanced-learn: A Python toolbox to tackle the curse of imbalanced datasets in machine learning. *Journal of Machine Learning Research*, 18(17):1–5.

Lesher, C. (2019). Up, down, or sideways: emplacement of magmatic Fe–Ni–Cu–PGE sulfide melts in large igneous provinces. *Canadian Journal of Earth Sciences*, 56(7):756–773.

Lewis, P. and Downes, P. (2008). Mineral systems and processes in new south wales: a project to enhance understanding and assist exploration. *NSW Geol. Surv. Quarterly Notes*, 128:1–15.

Li, C. and Naldrett, A. (1999). Geology and petrology of the voisey's bay intrusion: reaction of olivine with sulfide and silicate liquids. *Lithos*, 47(1-2):1–31.

Li, J., Zhu, Q., Wu, Q., Zhang, Z., Gong, Y., He, Z., and Zhu, F. (2021a). SMOTE-NaN-DE: Addressing the noisy and borderline examples problem in imbalanced classification by natural neighbors and differential evolution. *Knowledge-Based Systems*, 223:107056.

Li, Q., Chen, G., and Luo, L. (2023). Mineral Prospectivity Mapping Using Attention–based Convolutional Neural Network. *Ore Geology Reviews*, page 105381.

Li, T., Xia, Q., Ouyang, Y., Zeng, R., Liu, Q., and Li, T. (2024). Prospectivity and uncertainty analysis of tungsten polymetallogenic mineral resources in the Nanling metallogenic belt, South China: a comparative study of AdaBoost, GBDT, and XgBoost algorithms. *Natural Resources Research*, 33(3):1049–1071.

Li, T., Xia, Q., Zhao, M., Gui, Z., and Leng, S. (2020). Prospectivity mapping for tungsten polymetallic mineral resources, Nanling metallogenic belt, south China: Use of random forest algorithm from a perspective of data imbalance. *Natural Resources Research*, 29:203–227.

Li, T., Zuo, R., Xiong, Y., and Peng, Y. (2021b). Random-drop data augmentation of deep convolutional neural network for mineral prospectivity mapping. *Natural Resources Research*, 30(1):27–38.

Li, T., Zuo, R., Zhao, X., and Zhao, K. (2022). Mapping prospectivity for regolith-hosted REE deposits via convolutional neural network with generative adversarial network augmented data. *Ore Geology Reviews*, 142:104693.

Li, X., Yuan, F., Zhang, M., Jia, C., Jowitt, S. M., Ord, A., Zheng, T., Hu, X., and Li, Y. (2015). Three-dimensional mineral prospectivity modelling for targeting of concealed mineralization within the Zhonggu iron orefield, Ningwu Basin, China. *Ore Geology Reviews*, 71:633–654.

Lindsay, M. D., Aillères, L., Jessell, M. W., de Kemp, E. A., and Betts, P. G. (2012). Locating and quantifying geological uncertainty in three-dimensional models: Analysis of the Gippsland Basin, southeastern Australia. *Tectonophysics*, 546:10–27.

Lisitsin, V. A., Porwal, A., and McCuaig, T. C. (2014). Probabilistic fuzzy logic modeling: Quantifying uncertainty of mineral prospectivity models using Monte Carlo simulations. *Mathematical Geosciences*, 46:747–769.

Liu, S. Q., Liu, L., Kozan, E., Corry, P., Masoud, M., Chung, S.-H., and Li, X. (2025a). Machine learning for open-pit mining: a systematic review. *International Journal of Mining, Reclamation and Environment*, 39(1):1–39.

Liu, Y., Sun, T., Wu, K., Xiang, W., Zhang, J., Zhang, H., and Feng, M. (2025b). Interpretability analysis of data augmented convolutional neural network in mineral prospectivity mapping using black-box visualization tools. *Natural Resources Research*, 34(2):759–783.

Liu, Z., Yu, S., Deng, H., Jiang, G., Wang, R., Yang, X., Song, J., Chen, J., and Mao, X. (2024). 3D mineral prospectivity modeling in the Sanshandao goldfield, China using the convolutional neural network with attention mechanism. *Ore Geology Reviews*, 164:105861.

Lord, D., Etheridge, M., Willson, M., Hall, G., and Uttley, P. (2001). Measuring exploration success: an alternative to the discovery-cost-per-ounce method of quantifying exploration effectiveness. *SEG Newsletter*, (45):1–16.

Lu, Y., Lesher, C. M., and Deng, J. (2019). Geochemistry and genesis of magmatic Ni-Cu-(PGE) and PGE-(Cu)-(Ni) deposits in China. *Ore Geology Reviews*, 107:863–887.

Lundberg, S. M. and Lee, S.-I. (2017a). A unified approach to interpreting model predictions. Curran Associates, Inc.

Lundberg, S. M. and Lee, S.-I. (2017b). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.

Luo, X. and Dimitrakopoulos, R. (2003). Data-driven fuzzy analysis in quantitative mineral resource assessment. *Computers & Geosciences*, 29(1):3–13.

Lv, X. and Wang, G. (2024). GIS-based mineral prospectivity mapping using machine learning methods: A case study from Duobaoshan ore district, northeastern China. *Ore Geology Reviews*, 175:106352.

Ma, K., Tian, M., Tan, Y., Xie, X., and Qiu, Q. (2022). What is this article about? generative summarization with the bert model in the geosciences domain. *Earth Science Informatics*, 15(1):21–36.

Maier, W. D. (2005). Platinum-group element (PGE) deposits and occurrences: Mineralization styles, genetic concepts, and exploration criteria. *Journal of African Earth Sciences*, 41(3):165–191.

Maier, W. D., Li, C., and De Waal, S. A. (2001). Why are there no major Ni–Cu sulfide deposits in large layered mafic-ultramafic intrusions? *The Canadian Mineralogist*, 39(2):547–556.

Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., and Frey, B. (2015). Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*.

Mantilla-Dulcey, A., Goyes-Peñafiel, P., Báez-Rodríguez, R., and Khurama, S. (2024). Porphyry-type mineral prospectivity mapping with imbalanced data via prior geological transfer learning. *Gondwana Research*, 136:236–250.

McAliley, W. A. and Li, Y. (2024). Stochastic inversion of geophysical data by a conditional variational autoencoder. *Geophysics*, 89(1):WA219–WA232.

McCuaig, T. C., Beresford, S., and Hronsky, J. (2010). Translating the mineral systems approach into an effective exploration targeting system. *Ore Geology Reviews*, 38(3):128–138.

McCuaig, T. C. and Hronsky, J. M. (2014). The mineral system concept: the key to exploration targeting.

Meinert, L. D., Robinson Jr, G. R., and Nassar, N. T. (2016). Mineral resources: Reserves, peak production and the future. *Resources*, 5(1):14.

Mihalasky, M. J. and Bonham-Carter, G. F. (2001). Lithodiversity and its spatial association with metallic mineral sites, Great Basin of Nevada. *Natural Resources Research*, 10:209–226.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013a). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013b). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 26.

Mirzabozorg, S. A. A. S. and Abedi, M. (2023). Recognition of mineralization-related anomaly patterns through an autoencoder neural network for mineral exploration targeting. *Applied Geochemistry*, 158:105807.

Molamohammadi, M., Rezaei-Shoshtari, S., and Quitoriano, N. J. (2020). Jacobian of generative models for sensitivity analysis of photovoltaic device processes. In *Machine Learning for Engineering Workshop at NeurIPS*, volume 2020.

Montsion, R., Perrouty, S., Lindsay, M., Jessell, M., and Sherlock, R. (2024). Development and application of feature engineered geological layers for ranking magmatic, volcanogenic, and orogenic system components in archean greenstone belts. *Geoscience Frontiers*, 15(2):101759.

Mou, N., Carranza, E. J. M., Xue, J., Zhang, S., Wang, G., Song, H., Chen, Y., and Ren, X. (2025). Interpretable machine learning for mineral prospectivity mapping in the Qulong–Jiama district, Tibet, China. *Ore Geology Reviews*, page 106659.

Nagasingha, L. M., Bérubé, C. L., and Ghanati, R. (2025). A generative neural network approach to uncertainty and risk-return analysis in mineral prospectivity modelling. *Ore Geology Reviews*, page 106844.

Nagasingha, L. M., Bérubé, C. L., and Ghanati, R. (2026). Tractable approximate gaussian inference for interpretable mineral prospectivity modelling. *Applied Computing and Geosciences*, page 100332.

Nagasingha, L. M., Bérubé, C. L., and Lawley, C. J. (2024). A balanced mineral prospectivity model of Canadian magmatic Ni (±Cu±Co±PGE) sulphide mineral systems using conditional variational autoencoders. *Ore Geology Reviews*, 175:106329.

Naldrett, A. (1981). Nickel sulfide deposits: classification, composition, and genesis. *Economic Geology*, pages 628–685.

Naldrett, A. (1997). Key factors in the genesis of Noril'sk, Sudbury, Jinchuan, Voisey's Bay and other world-class Ni-Cu-PGE deposits: Implications for exploration. *Australian Journal of Earth Sciences*, 44(3):283–315.

Naldrett, A. J. (2005). A history of our understanding of magmatic Ni–Cu sulfide deposits. *The Canadian Mineralogist*, 43(6):2069–2098.

Naldrett, A. J. (2011). *Fundamentals of Magmatic Sulfide Deposits*. GeoScienceWorld.

Naldrett, A. J. (2013). *Magmatic sulfide deposits: geology, geochemistry and exploration*. Springer Science & Business Media.

Natural Resources Canada (2023). Nickel facts. `https://natural-resources.canada.ca/minerals-mining/mining-data-statistics-analysis/minerals-metals-facts/nickel-facts`.

Nguyen, L.-H. and Goulet, J.-A. (2021). Analytically tractable inference in deep neural networks. *arXiv preprint arXiv:2103.05461*.

Nguyen, L.-H. and Goulet, J.-A. (2022). cuTAGI: a CUDA library for Bayesian neural networks with Tractable Approximate Gaussian Inference. https://github.com/lhnguyen102/cuTAGI.

Nixon, G., Scoates, J., Milidragovic, D., Nott, J., Moerhuis, N., Ver Hoeve, T., Manor, M., Kjarsgaard, I., and File, O. (2020). Convergent margin Ni-Cu-PGE-Cr ore systems: U-Pb petrochronology and environments of Cu-PGE vs Cr-PGE mineralization in Alaskan-type intrusions. *Geological Survey of Canada, Open File*, 8722:197–218.

Nykänen, V., Lahti, I., Niiranen, T., and Korhonen, K. (2015). Receiver operating characteristics (ROC) as validation tool for prospectivity models—A magmatic Ni–Cu case study from the Central Lapland Greenstone Belt, Northern Finland. *Ore Geology Reviews*, 71:853–860.

Nykänen, V. and Ojala, V. J. (2007). Spatial analysis techniques as successful mineral-potential mapping tools for orogenic gold deposits in the Northern Fennoscandian Shield, Finland. *Natural Resources Research*, 16(2):85–92.

Nykänen, V. and Salmirinne, H. (2007). Prospectivity analysis of gold using regional geophysical and geochemical data from the Central Lapland Greenstone Belt, Finland. *Geological Survey of Finland*, 44:251–269.

O'Driscoll, E. (1986). Observations of the lineament-ore relation. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 317(1539):195–218.

Okada, K. (2021). A historical overview of the past three decades of mineral exploration technology. *Natural Resources Research*, 30:2839–2860.

Padarian, J. and Fuentes, I. (2019). Word embeddings for application in geosciences: development, evaluation, and examples of soil-related concepts. *Soil*, 5(2):177–187.

Parsa, M. and Carranza, E. J. M. (2021). Modulating the impacts of stochastic uncertainties linked to deposit locations in data-driven predictive mapping of mineral prospectivity. *Natural Resources Research*, 30(5):3081–3097.

Parsa, M. and Cumani, R. (2025). Class label representativeness in machine learning-based mineral prospectivity mapping. *Natural Resources Research*, pages 1–25.

Parsa, M., Lawley, C. J., Cawood, T., Martins, T., Cumani, R., Zhang, S. E., Thompson, A., Schetselaar, E., Beyer, S., Lentz, D. R., et al. (2025). Pan-Canadian predictive modeling of lithium–cesium–tantalum pegmatites with deep learning and natural language processing. *Natural Resources Research*, 34(2):639–668.

Parsa, M., Lawley, C. J., Cumani, R., Schetselaar, E., Harris, J., Lentz, D. R., Zhang, S. E., and Bourdeau, J. E. (2024). Predictive modeling of Canadian carbonatite-hosted REE+/-Nb deposits. *Natural Resources Research*, 33(5):1941–1965.

Partington, G. and Sale, M. (2004). Prospectivity mapping using gis with publicly available earth science data—a new targeting tool being successfully used for exploration in new zealand. In *Pacrim 2004 congress volume, Adelaide*, pages 239–250.

Partington, G. A., Peters, K. J., Czertowicz, T. A., Greville, P. A., Blevin, P. L., and Bahiru, E. A. (2024). Ranking mineral exploration targets in support of commercial decision making: A key component for inclusion in an exploration information system. *Applied Geochemistry*, 168:106010.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Peng, Q., Wang, Z., Wang, G., Zhang, W., Chen, Z., and Liu, X. (2023). 3d mineral prospectivity mapping from 3d geological models using return–risk analysis and machine learning on imbalance data. *Minerals*, 13(11):1384.

Pennington, J., Socher, R., and Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.

Porwal, A., Carranza, E., and Hale, M. (2003a). Artificial neural networks for mineral-potential mapping: a case study from Aravalli Province, Western India. *Natural resources research*, 12(3):155–171.

Porwal, A., Carranza, E., and Hale, M. (2003b). Knowledge-driven and data-driven fuzzy models for predictive mineral potential mapping. *Natural Resources Research*, 12(1):1–25.

Porwal, A., Carranza, E., and Hale, M. (2004). A hybrid neuro-fuzzy model for mineral potential mapping. *Mathematical geology*, 36(7):803–826.

Porwal, A. and Carranza, E. J. M. (2015). Introduction to the Special Issue: GIS-based mineral potential modelling and geological data analyses for mineral exploration.

Porwal, A., Carranza, E. J. M., and Hale, M. (2006a). Bayesian network classifiers for mineral potential mapping. *Computers & Geosciences*, 32(1):1–16.

Porwal, A., Carranza, E. J. M., and Hale, M. (2006b). A hybrid fuzzy weights-of-evidence model for mineral potential mapping. *Natural Resources Research*, 15(1):1–14.

Pradhan, B., Jena, R., Talukdar, D., Mohanty, M., Sahu, B. K., Raul, A. K., and Abdul Maulud, K. N. (2022). A new method to evaluate gold mineralisation-potential mapping using deep learning and an explainable artificial intelligence (XAI) model. *Remote Sensing*, 14(18):4486.

Prado, E. M. G., de Souza Filho, C. R., Carranza, E. J. M., and Motta, J. G. (2020). Modeling of Cu-Au prospectivity in the Carajás mineral province (Brazil) through machine learning: Dealing with imbalanced training data. *Ore Geology Reviews*, 124:103611.

Puzyrev, V. and Duuring, P. (2025). Uncertainty quantification of geochemical data imputation using Monte Carlo dropout. *Journal of Geochemical Exploration*, 272:107695.

Qaderi, S., Maghsoudi, A., Yousefi, M., and Pour, A. B. (2025). Translation of mineral system components into time step-based ore-forming events and evidence maps for mineral exploration: Intelligent mineral prospectivity mapping through adaptation of recurrent neural networks and random forest algorithm. *Ore Geology Reviews*, page 106537.

Rahimi, H., Abedi, M., Yousefi, M., Bahroudi, A., and Elyasi, G.-R. (2021). Supervised mineral exploration targeting and the challenges with the selection of deposit and non-deposit sites thereof. *Applied Geochemistry*, 128:104940.

Raschka, S. (2018). Model evaluation, model selection, and algorithm selection in machine learning. *arXiv preprint arXiv:1811.12808*.

Rasht-Behesht, M., Huber, C., Shukla, K., and Karniadakis, G. E. (2022). Physics-informed neural networks (pinns) for wave propagation and full waveform inversions. *Journal of Geophysical Research: Solid Earth*, 127(5):e2021JB023120.

Rauch, H. E., Tung, F., and Striebel, C. T. (1965). Maximum likelihood estimates of linear dynamic systems. *AIAA Journal*, 3(8):1445–1450.

Rautela, M., Williams, A., and Scheinker, A. (2025). Time-inversion of spatiotemporal beam dynamics using uncertainty-aware latent evolution reversal. *Physical Review E*, 111(2):025307.

Reddy, R. and Bonham-Carter, G. (1991). A decision-tree approach to mineral potential mapping in Snow Lake area, Manitoba. *Canadian Journal of Remote Sensing*, 17(2):191–200.

Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., and Prabhat, F. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743):195–204.

Rey, S. J. and Anselin, L. (2007). PySAL: A Python Library of Spatial Analytical Methods. *The Review of Regional Studies*, 37(1):5–27.

Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.

Ritter, N., Logan, T., and Bryant, N. (1988). Integration of neural network technologies with geographic information systems. In *Proceedings of the GIS symposium: integrating technology and geoscience applications*, pages 102–103.

Ritter, N. D. and Hepner, G. F. (1990). Application of an artificial neural network to land-cover classification of thematic mapper imagery. *Computers & Geosciences*, 16(6):873–880.

Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., et al. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8):913–929.

Rodriguez-Galiano, V., Sanchez-Castillo, M., Chica-Olmo, M., and Chica-Rivas, M. (2015). Machine learning predictive models for mineral prospectivity: An evaluation of neural networks, random forest, regression trees and support vector machines. *Ore Geology Reviews*, 71:804–818.

Rogge, D., Halden, N., and Beaumont-Smith, C. (2006). Application of data integration for shear-hosted Au potential modelling: Lynn Lake Greenstone Belt, Northwestern Manitoba, Canada. *GIS for the Earth Sciences*, 44:191–210.

Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386.

Roshanravan, B. and Kreuzer, O. (2025). An Optimized XGBoost Approach to Prospectivity Modelling of Orogenic Gold Systems in the Granites-Tanami Orogen, Australia. In *86th EAGE Annual Conference & Exhibition*, volume 2025, pages 1–5. European Association of Geoscientists & Engineers.

Roy, R., Cassard, D., Cobbold, P., Rossello, E. A., Billa, M., Bailly, L., and Lips, A. (2006). Predictive mapping for copper–gold magmatic-hydrothermal systems in NW Argentina: Use of a regional-scale GIS, application of an expert-guided data-driven approach, and comparison with results from a continental-scale GIS. *Ore Geology Reviews*, 29(3-4):260–286.

Rybkin, O., Daniilidis, K., and Levine, S. (2021). Simple and effective VAE training with calibrated decoders. In *International Conference on Machine Learning*, pages 9179–9189. PMLR.

Sadeghi, B. and Khalajmasoumi, M. (2015). A futuristic review for evaluation of geothermal potentials using fuzzy logic and binary index overlay in GIS environment. *Renewable and Sustainable Energy Reviews*, 43:818–831.

Saetta, E., Tognaccini, R., and Iaccarino, G. (2024). Uncertainty quantification in autoencoders predictions: Applications in aerodynamics. *Journal of Computational Physics*, 506:112951.

Saibi, H., Amrouche, M., Batir, J., Gabr, A., and Fowler, A.-R. (2022). Magnetic and gravity modeling and subsurface structure of two geothermal fields in the UAE. *Geothermal Energy*, 10(1):28.

Saljoughi, B. S. and Hezarkhani, A. (2018). A comparative analysis of artificial neural network (ANN), wavelet neural network (WNN), and support vector machine (SVM) data-driven models to mineral potential mapping for copper mineralizations in the Shahr-e-Babak region, Kerman, Iran. *Applied Geomatics*, 10:229–256.

Sappin, A.-A., Constantin, M., and Clark, T. (2011). Origin of magmatic sulfides in a Proterozoic island arc—an example from the Portneuf–Mauricie Domain, Grenville Province, Canada. *Mineralium Deposita*, 46:211–237.

Saremi, M., Hoseinzade, Z., and Yousefi, M. (2025). A deep embedded clustering algorithm in conjunction with an ensemble technique for mineral prospectivity mapping. *Scientific Reports*, 15(1):38086.

Schetselaar, E. and Snyder, D. (2017). *National database of MOHO depth estimates from seismic refraction and teleseismic surveys*. Geological Survey of Canada.

Schodde, R. (2017). Long term trends in global exploration–Are we finding enough metals? In *Fennoscandian Exploration and Mining Conference*, pages 4–55.

Schodde, R. (2025). Mineral deposit exploration—discovery trends: 1900–2023. *SEG Discovery*, (142):19–35.

Schulz, K. J., Chandler, V. W., Nicholson, S. W., Piatak, N., Seal, R. R., Woodruff, L. G., Zientek, M. L., et al. (2010). *Magmatic sulfide-rich nickel-copper deposits related to picrite*

and (or) tholeiitic basalt dike-sill complexes: A preliminary deposit model. US Department of the Interior, US Geological Survey.

Seiffert, C., Khoshgoftaar, T. M., Van Hulse, J., and Napolitano, A. (2009). Rusboost: A hybrid approach to alleviating class imbalance. *IEEE transactions on systems, man, and cybernetics-part A: systems and humans*, 40(1):185–197.

Shabani, A., Asgarian, B., Gharebaghi, S. A., Salido, M. A., and Giret, A. (2019). A new optimization algorithm based on search and rescue operations. *Mathematical Problems in Engineering*, 2019(1):2482543.

Shirmard, H., Farahbakhsh, E., Müller, R. D., and Chandra, R. (2022). A review of machine learning in processing remote sensing data for mineral exploration. *Remote Sensing of Environment*, 268:112750.

Si, X., Fu, Y., Liu, X., Wang, R., Xu, W., and Wang, L. (2024). CVAE-Based Hybrid Sampling Data Augmentation Method and Interpretation for Imbalanced Classification of Gout Disease. In *International Conference on Intelligent Computing*, pages 49–60. Springer.

Sihombing, F. M., Palin, R. M., Hughes, H. S., and Robb, L. J. (2024). Improved mineral prospectivity mapping using graph neural networks. *Ore Geology Reviews*, 172:106215.

Silva dos Santos, V., Gloaguen, E., Hector Abud Louro, V., and Blouin, M. (2022). Machine learning methods for quantifying uncertainty in prospectivity mapping of magmatic-hydrothermal gold deposits: A case study from Juruena mineral province, northern Mato Grosso, Brazil. *Minerals*, 12(8):941.

Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Singer, D. A. and Kouda, R. (1999). Examining risk in mineral exploration. *Natural Resources Research*, 8(2):111–122.

Skirrow, R. G., Huston, D. L., Mernagh, T. P., Thorne, J. P., Duffer, H., and Senior, A. (2013). *Critical commodities for a high-tech world: Australia's potential to supply global demand*. Geoscience Australia Canberra.

Skirrow, R. G., Murr, J., Schofield, A., Huston, D. L., van der Wielen, S., Czarnota, K., Coghlan, R., Highet, L. M., Connolly, D., Doublier, M., et al. (2019). Mapping iron oxide Cu-Au (IOCG) mineral potential in Australia using a knowledge-driven mineral systems-based approach. *Ore Geology Reviews*, 113:103011.

Smith, W. D., Maier, W. D., and Bliss, I. (2021). The geology, geochemistry, and petrogenesis of the Huckleberry Cu-Ni-PGE prospect in the Labrador Trough, Canada: Perspectives for regional prospectivity. *Ore Geology Reviews*, 128:103905.

Sohn, K., Lee, H., and Yan, X. (2015). Learning structured output representation using deep conditional generative models. *Advances in Neural Information Processing Systems*, 28.

Soltani, Z. and Imamalipour, A. (2022). An improved classification of mineralized zones using particle swarm optimization: A case study from Dagh-Dali ZnPb (±Au) prospect, Northwest Iran. *Geochemistry*, 82(1):125850.

Song, X., Wang, Y., and Chen, L. (2011). Magmatic Ni-Cu-(PGE) deposits in magma plumbing systems: Features, formation and exploration. *Geoscience Frontiers*, 2(3):375–384.

Srivastava, A. and Jain, S. (2025). Remote sensing for environment assessment: Multispectral, hyperspectral, and thermal imaging applications. In *Remote Sensing for Environmental Monitoring*, pages 1–31. Springer.

Stinson, H. and Cam, I. (2025). The global energy transition and critical minerals.

Styron, R. and Pagani, M. (2020). The GEM global active faults database. *Earthquake Spectra*, 36(1_suppl):160–180.

Sun, K., Chen, Y., Geng, G., Lu, Z., Zhang, W., Song, Z., Guan, J., Zhao, Y., and Zhang, Z. (2024). A review of mineral prospectivity mapping using deep learning. *Minerals*, 14(10):1021.

Sun, T., Chen, F., Zhong, L., Liu, W., and Wang, Y. (2019). Gis-based mineral prospectivity mapping using machine learning methods: A case study from Tongling ore district, eastern China. *Ore Geology Reviews*, 109:26–49.

Sun, T., Li, H., Wu, K., Chen, F., Zhu, Z., and Hu, Z. (2020). Data-driven predictive modelling of mineral prospectivity using machine learning and deep learning methods: A case study from southern Jiangxi Province, China. *Minerals*, 10(2):102.

Tahmasebi, P. and Hezarkhani, A. (2012). A hybrid neural networks-fuzzy logic-genetic algorithm for grade estimation. *Computers & geosciences*, 42:18–27.

Tang, J., Zhang, H., Bai, R., Zhang, J., and Sun, T. (2025). Mineral prospectivity mapping for exploration targeting of porphyry cu-polymetallic deposits based on machine learning algorithms, remote sensing and multi-source geo-information. *Minerals*, 15(10):1050.

Tangestani, M. and Moore, F. (2003). Mapping porphyry copper potential with a fuzzy model, northern Shahr-e-Babak, Iran. *Australian Journal of Earth Sciences*, 50(3):311–317.

Transport Canada (2021). Building a green economy: Government of Canada to require 100% of car and passenger truck sales be zero-emission by 2035 in Canada.

Tsai, J.-F., Carlsson, J. G., Ge, D., Yi-Chung, H., and Shi, J. (2014). Optimization theory, methods, and applications in engineering 2013. *Mathematical Problems in Engineering*, 2014.

Tsoukalas, L. H. and Uhrig, R. E. (1996). *Fuzzy and neural approaches in engineering*. John Wiley & Sons, Inc.

Uber Technologies, Inc. (2023). H3: A hexagon-based spatial indexing system. `https://github.com/uber/h3`. (Accessed on 09/19/2024).

Ustundag, A., Sivri, M. S., and Menguc, K. (2022). Feature engineering. In *Business Analytics for Professionals*, pages 153–169. Springer.

Valavi, R., Elith, J., Lahoz-Monfort, J. J., and Guillera-Arroita, G. (2018). blockCV: An R package for generating spatially or environmentally separated folds for k-fold cross-validation of species distribution models. *Biorxiv*, page 357798.

Vapnik, V. (1999). *The nature of statistical learning theory*. Springer.

VRIFY (2025). Cartier resources unlocks new potential in a historic gold mining region. `https://vrify.com/case-studies/cartier-resources-unlocks-new-potential-in-a-historic-gold-mining-region`.

Vuong, V.-D., Nguyen, L.-H., and Goulet, J.-A. (2025). Coupling LSTM neural networks and state-space models through analytically tractable inference. *International Journal of Forecasting*, 41(1):128–140.

Wake, N., Farahbakhsh, E., and Müller, R. D. (2025). Lateritic Ni–Co Prospectivity Modeling in Eastern Australia Using an Enhanced Generative Adversarial Network and Positive-Unlabeled Bagging. *Natural Resources Research*, 34(1):61–96.

Wang, B., Ma, K., Wu, L., Qiu, Q., Xie, Z., and Tao, L. (2022). Visual analytics and information extraction of geological content for text-based mineral exploration reports. *Ore Geology Reviews*, 144:104818.

Wang, X. and Huang, J. Z. (2015). Uncertainty in learning from big data. *Fuzzy Sets and Systems*, 258:1–4.

Wang, Z., Yin, Z., Caers, J., and Zuo, R. (2020). A Monte Carlo-based framework for risk-return analysis in mineral prospectivity mapping. *Geoscience Frontiers*, 11(6):2297–2308.

Wang, Z. and Zuo, R. (2022). Mineral prospectivity mapping using a joint singularity-based weighting method and long short-term memory network. *Computers & Geosciences*, 158:104974.

Watari, T., Nansai, K., and Nakajima, K. (2020). Review of critical metal dynamics to 2050 for 48 elements. *Resources, Conservation and Recycling*, 155:104669.

Wellmann, J. F., Horowitz, F. G., Schill, E., and Regenauer-Lieb, K. (2010). Towards incorporating uncertainty of structural data in 3D geological inversion. *Tectonophysics*, 490(3-4):141–151.

Wen, Y., Vicol, P., Ba, J., Tran, D., and Grosse, R. (2018). Flipout: Efficient pseudo-independent weight perturbations on mini-batches. *arXiv preprint arXiv:1803.04386*.

Wheeler, J. O. (1996). *Geological map of Canada: Map 1860A*. Geological Survey of Canada.

Wu, A., Nowozin, S., Meeds, E., Turner, R. E., Hernandez-Lobato, J. M., and Gaunt, A. L. (2018). Deterministic variational inference for robust Bayesian neural networks. *arXiv preprint arXiv:1810.03958*.

Wu, Y., Liu, B., Gao, Y., Li, C., Tang, R., Kong, Y., Xie, M., Li, K., Dan, S., Qi, K., et al. (2023). Mineral prospecting mapping with conditional generative adversarial network augmented data. *Ore Geology Reviews*, 163:105787.

Wyborn, L., Heinrich, C., and Jaques, A. (1994). Australian Proterozoic mineral systems: essential ingredients and mappable criteria. In *1994 AusIMM Annual Conference: Australian Mining Looks North-the Challenges and Choices*, pages 109–115. Australasian Institute of Mining and Metallurgy.

Wyman, D. A., Cassidy, K. F., and Hollings, P. (2016). Orogenic gold and the mineral systems approach: Resolving fact, fiction and fantasy. *Ore Geology Reviews*, 78:322–335.

Xiang, J., Xiao, K., Carranza, E. J. M., Chen, J., and Li, S. (2020). 3D mineral prospectivity mapping with random forests: A case study of Tongling, Anhui, China. *Natural Resources Research*, 29(1):395–414.

Xiong, Y. and Zuo, R. (2017). Effects of misclassification costs on mapping mineral prospectivity. *Ore Geology Reviews*, 82:1–9.

Xiong, Y. and Zuo, R. (2021). A positive and unlabeled learning algorithm for mineral prospectivity mapping. *Computers & Geosciences*, 147:104667.

Xiong, Y., Zuo, R., and Carranza, E. J. M. (2018). Mapping mineral prospectivity through big data analytics and a deep learning algorithm. *Ore Geology Reviews*, 102:811–817.

Xiong, Y., Zuo, R., Luo, Z., and Wang, X. (2022). A physically constrained variational autoencoder for geochemical pattern recognition. *Mathematical geosciences*, 54(4):783–806.

Yang, F., Wang, Z., Zuo, R., Sun, S., and Zhou, B. (2023). Quantification of uncertainty associated with evidence layers in mineral prospectivity mapping using direct sampling and convolutional neural network. *Natural Resources Research*, 32(1):79–98.

Yang, F. and Zuo, R. (2024). Geologically constrained convolutional neural network for mineral prospectivity mapping. *Mathematical Geosciences*, 56(8):1605–1628.

Yang, F., Zuo, R., and Kreuzer, O. P. (2024a). Artificial intelligence for mineral exploration: A review and perspectives on future directions from data science. *Earth-Science Reviews*, 258:104941.

Yang, F., Zuo, R., Xiong, Y., Xu, Y., Nie, J., and Zhang, G. (2024b). Dual-branch convolutional neural network and its post hoc interpretability for mapping mineral prospectivity. *Mathematical geosciences*, 56(7):1487–1515.

Yang, N., Zhang, Z., Yang, J., and Hong, Z. (2022a). Applications of data augmentation in mineral prospectivity prediction based on convolutional neural networks. *Computers & geosciences*, 161:105075.

Yang, N., Zhang, Z., Yang, J., and Hong, Z. (2022b). Mineral prospectivity prediction by integration of convolutional autoencoder network and random forest. *Natural Resources Research*, 31(3):1103–1119.

Yin, B., Zuo, R., and Xiong, Y. (2022). Mineral prospectivity mapping via gated recurrent unit model. *Natural Resources Research*, 31(4):2065–2079.

Yin, J. and Li, N. (2022). Ensemble learning models with a Bayesian optimization algorithm for mineral prospectivity mapping. *Ore Geology Reviews*, 145:104916.

Yousefi, M. and Carranza, E. J. M. (2015). Prediction–area (P–A) plot and C–A fractal analysis to classify and evaluate evidential maps for mineral prospectivity modeling. *Computers & Geosciences*, 79:69–81.

Yousefi, M. and Carranza, E. J. M. (2016). Data-driven index overlay and boolean logic mineral prospectivity modeling in greenfields exploration. *Natural Resources Research*, 25(1):3–18.

Yousefi, M., Carranza, E. J. M., Kreuzer, O. P., Nykänen, V., Hronsky, J. M., and Mihalasky, M. J. (2021). Data analysis methods for prospectivity modelling as applied to mineral exploration targeting: State-of-the-art and outlook. *Journal of Geochemical Exploration*, 229:106839.

Yousefi, M. and Kreuzer, O. (2024). Towards an effective exploration information system–new concepts and ideas aimed at improving mineral exploration targeting.

Yousefi, M., Kreuzer, O. P., Nykänen, V., and Hronsky, J. M. (2019). Exploration information systems–A proposal for the future use of GIS in mineral exploration targeting. *Ore Geology Reviews*, 111:103005.

Yousefi, M., Lindsay, M. D., and Kreuzer, O. (2024a). Mitigating uncertainties in mineral exploration targeting: Majority voting and confidence index approaches in the context of an exploration information system (EIS). *Ore Geology Reviews*, 165:105930.

Yousefi, M. and Nykänen, V. (2016). Data-driven logistic-based weighting of geochemical and geological evidence layers in mineral prospectivity mapping. *Journal of Geochemical Exploration*, 164:94–106.

Yousefi, M., Nykänen, V., Harris, J., Hronsky, J. M., Kreuzer, O. P., Bertrand, G., and Lindsay, M. (2024b). Overcoming survival bias in targeting mineral deposits of the future: Towards null and negative tests of the exploration search space, accounting for lack of visibility. *Ore Geology Reviews*, 172:106214.

Zadeh, J. (2025). The global shift to undercover exploration: Challenges and opportunities.

Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, 8(3):338–353.

Zeghouane, H., Allek, K., and Kesraoui, M. (2016). Gis-based weights of evidence modeling applied to mineral prospectivity mapping of Sn-W and rare metals in Laouni area, Central Hoggar, Algeria. *Arabian Journal of Geosciences*, 9(5):373.

Zekang, F., Xiaojun, Z., Yongfeng, Y., Xu, X., Fanchao, Z., Li, X., Quantong, Z., and Weikun, M. (2025). The Evolution of Machine Learning in Large-Scale Mineral Prospectivity Prediction: A Decade of Innovation (2016–2025). *Minerals*, 15(10):1042.

Zhang, H. and Sun, T. (2025). Bibliometric analysis of mineral prospectivity mapping research from 2006 to 2024: Zhang and Sun. *Natural Resources Research*, pages 1–27.

Zhang, N. and Zhou, K. (2015). Mineral prospectivity mapping with weights of evidence and fuzzy logic methods. *Journal of Intelligent & Fuzzy Systems*, 29(6):2639–2651.

Zhang, Q., Chen, J., Xu, H., Jia, Y., Chen, X., Jia, Z., and Liu, H. (2022a). Three-dimensional mineral prospectivity mapping by XGBoost modeling: A case study of the Lannigou gold deposit, China. *Natural Resources Research*, 31(3):1135–1156.

Zhang, S. E., Lawley, C. J., Bourdeau, J. E., Parsa, M., Cumani, R., and Thompson, A. (2025). Mineral Prospectivity Modeling of Graphite Deposits and Occurrences in Canada. *Natural Resources Research*, 34(3):1243–1270.

Zhang, Z., Li, Y., Wang, G., Carranza, E. J. M., Yang, S., Sha, D., Fan, J., Zhang, X., and Dong, Y. (2023). Supervised mineral prospectivity mapping via class-balanced focal loss function on imbalanced geoscience datasets. *Mathematical Geosciences*, 55(7):989–1010.

Zhang, Z., Wang, G., Carranza, E. J. M., Du, J., Li, Y., Liu, X., and Su, Y. (2024). An uncertainty-quantification machine learning framework for data-driven three-dimensional mineral prospectivity mapping. *Natural Resources Research*, pages 1–19.

Zhang, Z., Wang, G., Carranza, E. J. M., Fan, J., Liu, X., Zhang, X., Dong, Y., Chang, X., and Sha, D. (2022b). An integrated framework for data-driven mineral prospectivity mapping using bagging-based positive-unlabeled learning and Bayesian cost-sensitive logistic regression. *Natural Resources Research*, 31(6):3041–3060.

Zhao, H., Deng, K., Li, N., Wang, Z., and Wei, W. (2020). Hierarchical spatial-spectral feature extraction with long short term memory (LSTM) for mineral identification using hyperspectral imagery. *Sensors*, 20(23):6854.

Zhou, F.-Y., Jin, L.-P., Dong, J., et al. (2017). Review of convolutional neural network. *Chinese Journal of Computers*.

Zhou, G., Chen, J., An, W., Liu, C., and Li, W. (2024). Three-dimensional mineral prospectivity mapping based on natural language processing and random forests: A case study of the Xiyu diamond deposit, China. *Ore Geology Reviews*, 169:106082.

Zhou, Z.-H. and Liu, X.-Y. (2006). Training cost-sensitive neural networks with methods addressing the class imbalance problem. *IEEE Transactions on Knowledge and Data Engineering*, 18(1):63–77.

Zuccarelli, N., Lesher, C. M., Houlé, M. G., Weston, R., and Barnes, S. J. (2022). Diversity of net-textured sulfides in magmatic sulfide deposits: Insights from the Eagle's Nest Ni-Cu-(platinum group element) deposit, McFaulds Lake greenstone belt, Superior province, Canada. *Economic Geology*, 117(8):1731–1759.

Zuo, R. (2017). Machine learning of mineralization-related geochemical anomalies: A review of potential methods. *Natural Resources Research*, 26:457–464.

Zuo, R. (2020). Geodata science-based mineral prospectivity mapping: A review. *Natural Resources Research*, 29(6):3415–3424.

Zuo, R. and Carranza, E. J. M. (2011). Support vector machine: A tool for mapping mineral prospectivity. *Computers & Geosciences*, 37(12):1967–1975.

Zuo, R., Cheng, Q., and Agterberg, F. P. (2009). Application of a hybrid method combining multilevel fuzzy comprehensive evaluation with asymmetric fuzzy relation analysis to mapping prospectivity. *Ore Geology Reviews*, 35(1):101–108.

Zuo, R., Cheng, Q., Xu, Y., Yang, F., Xiong, Y., Wang, Z., and Kreuzer, O. P. (2024). Explainable artificial intelligence models for mineral prospectivity mapping. *Science China Earth Sciences*, 67(9):2864–2875.

Zuo, R., Kreuzer, O. P., Wang, J., Xiong, Y., Zhang, Z., and Wang, Z. (2021a). Uncertainties in GIS-based mineral prospectivity mapping: Key types, potential impacts and possible solutions. *Natural Resources Research*, 30:3059–3079.

Zuo, R., Luo, Z., Xiong, Y., and Yin, B. (2022). A geologically constrained variational autoencoder for mineral prospectivity mapping. *Natural Resources Research*, 31(3):1121–1133.

Zuo, R., Peng, Y., Li, T., and Xiong, Y. (2021b). Challenges of geological prospecting big data mining and integration using deep learning algorithms. *Earth Science*, 46(1):350–358.

Zuo, R. and Wang, Z. (2020). Effects of random negative training samples on mineral prospectivity mapping. *Natural Resources Research*, 29:3443–3455.

Zuo, R., Xiong, Y., Wang, Z., Wang, J., and Kreuzer, O. P. (2023). A new generation of artificial intelligence algorithms for mineral prospectivity mapping. *Natural Resources Research*, 32(5):1859–1869.

Zuo, R., Yang, F., Cheng, Q., and Kreuzer, O. P. (2025). A novel data-knowledge dual-driven model coupling artificial intelligence with a mineral systems approach for mineral prospectivity mapping. *Geology*, 53(3):284–288.

APPENDIX A DESCRIPTION OF GEOSCIENCE EVIDENCE LAYERS

Table A.1 Description of geoscience evidence layers used as input features for mineral prospectivity modelling of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems.

Data type	Dataset	Dataset description	Data source
Geology	Geology period (Maximum majority)	Maximum Period name (raster value with the largest sum of coverage fraction)	Lawley et al. (2022a)
Geology	Geology period (Maximum minority)	Maximum Period name (raster value with the smallest sum of coverage fraction)	Lawley et al. (2022a)
Geology	Geology period (Minimum majority)	Minimum Period name (raster value with the largest sum of coverage fraction)	Lawley et al. (2022a)
Geology	Geology period (Minimum minority)	Minimum Period name (raster value with the smallest sum of coverage fraction)	Lawley et al. (2022a)
Geology	Lithology majority	Lithology sub-type (raster value with the largest sum of coverage fraction)	Lawley et al. (2022a)
Geology	Lithology minority	Lithology sub-type (raster value with the smallest sum of coverage fraction)	Lawley et al. (2022a)
Geology	Lithology mafic	Presence or absence of geological property	Lawley et al. (2022a)
Geology	Igneous dictionary	Presence or absence of geological property	Lawley et al. (2022a)
Geology	Sedimentary dictionary	Presence or absence of geological property	Lawley et al. (2022a)

Continued on next page

Table A.1 Description of geoscience evidence layers used as input features for mineral prospectivity modelling of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems. (continued)

Data type	**Dataset**	**Dataset description**	**Data source**
Geology	Faults (proximity)	Proximity to fault (pixels; calculated from raster with 0.01 degree grid)	Lawley et al. (2022a)
Geology	Terrane (proximity)	Proximity to terrane boundary (pixels; calculated from raster with 0.01 degree grid)	Eglington et al. (2013)
Geology	Passive margin	Proximity to passive margin (pixels; calculated from raster with 0.01 degree grid)	Bradley (2008)
Geology	Black shale	Proximity to black shale (pixels; calculated from raster with 0.01 degree grid)	Granitto et al. (2017)
Geophysics	Gravity Bouguer	Bouguer gravity anomaly (mGal)	Lawley et al. (2022a)
Geophysics	Gravity (HGM)	Horizontal gradient magnitude (HGM)	Lawley et al. (2022a)
Geophysics	Gravity (upward-continued HGM)	Upward-continued (30 km) horizontal gradient magnitude (HGM)	Lawley et al. (2022a)
Geophysics	Gravity (worms)	Proximity to gravity worms (pixels; calculated from raster with 0.01 degree grid)	Lawley et al. (2022a)
Geophysics	Gravity Bouguer UpCont30km HGM Worms	Proximity to upward-continued (30 km) worms (pixels; calculated from raster with 0.01 degree grid)	Lawley et al. (2022a)
Geophysics	Magnetic (1VD)	1st vertical derivative (1VD; nT/m)	Lawley et al. (2022a)
Geophysics	Magnetic (HGM)	Horizontal gradient magnitude (HGM)	Lawley et al. (2022a)

Continued on next page

Table A.1 Description of geoscience evidence layers used as input features for mineral prospectivity modelling of Canadian magmatic Ni (±Cu ±Co ±PGE) sulphide mineral systems. (continued)

Data type	**Dataset**	**Dataset description**	**Data source**
Geophysics	Magnetic (worms)	Proximity to magnetic worms (pixels; calculated from raster with 0.01 degree grid)	Lawley et al. (2022a)
Geophysics	Magnetic Long-Wavelength HGM	Long-wavelength horizontal gradient magnitude (HGM)	Lawley et al. (2022a)
Geophysics	Magnetic Long-Wavelength HGM Worms	Proximity to long-wavelength worms (pixels; calculated from raster with 0.01 degree grid)	Lawley et al. (2022a)
Geophysics	Seismic (LAB)	Depth to lithosphere-asthenosphere boundary (LAB; km)	Lawley et al. (2022a)
Geophysics	Seismic (Moho)	Depth to Mohorovičić discontinuity or Moho (km)	Lawley et al. (2022a)
Geophysics	Satellite Gravity Maximum Curvature	GOCE maximum curvature derivative	Ebbing et al. (2018)
Geophysics	Satellite Gravity Minimum Curvature	GOCE minimum curvature derivative	Ebbing et al. (2018)
Geophysics	Satellite Gravity Mean Curvature	GOCE mean curvature derivative	Ebbing et al. (2018)
Geophysics	Satellite Gravity Shape Index	GOCE shape index	Ebbing et al. (2018)
Modelling	Ni deposits and occurrences	Presence of mineral deposit or absence (barren)	Lawley et al. (2021)

APPENDIX B THRESHOLD ANALYSIS FOR BAYESIAN MODEL COMPARISON

Table B.1 Performance comparison at multiple classification thresholds on the test dataset. Optimal thresholds are determined by maximizing the geometric mean score (GMS), which balances sensitivity and specificity. TP - true positives, FP - false positives, FN - false negatives, TN - true negatives.

Model	**Threshold**	**Recall**	**GMS**	**TP**	**FP**	**FN**	**TN**
TAGI							
	0.3	0.732	0.822	290	25,721	106	303,451
	0.4	0.715	0.816	283	22,462	113	306,710
	0.5	0.697	0.808	275	19,951	121	309,221
	0.6	0.669	0.796	265	17,771	131	311,401
	0.7	0.641	0.782	254	15,511	142	313,661
BNN with variational inference							
	0.3	0.321	0.561	127	6,063	269	323,109
	0.4	0.303	0.546	120	5,560	276	323,612
	0.5	0.288	0.532	114	5,081	282	324,091
	0.6	0.263	0.509	104	4,621	292	324,551
	0.7	0.240	0.487	95	4,155	301	325,017
MC Dropout							
	0.3	0.331	0.569	131	6,680	265	322,492
	0.4	0.303	0.545	120	5,995	276	323,177
	0.5	0.295	0.539	117	5,473	279	323,699
	0.6	0.288	0.532	114	5,013	282	324,159
	0.7	0.280	0.526	111	4,565	285	324,607