

RESEARCH

Open Access



Mineral prospectivity mapping under extreme imbalance using contrastive embeddings balanced learning and integrated uncertainty analysis

Dipak Kumar Nidhi^{1*}, Sudhir Kumar Mohapatra¹, Paavo Nevalainen¹, Jukka Heikkonen¹ and Rajeev Kanth²

*Correspondence:

Dipak Kumar Nidhi
dknidh@utu.fi

¹Department of Computing,
University of Turku, Turku
20014, Finland

²Savonia University of Applied
Sciences, 70201 Kuopio, Finland

Abstract

Mineral Prospectivity Mapping (MPM) is a pivotal methodology for identifying prospective deposits across large regions using complex geophysical datasets. The application of machine learning could significantly improve these processes. However, a critical challenge in data-driven mineral prospectivity mapping is the class imbalance between the mineralized locations and large background, which can severely limit model performance. To address this, this study systematically evaluates two machine learning workflows: a supervised Multilayer Perceptron (MLP) and a contrastive representation learning with radius classifier. The algorithm applied to a geophysical dataset from Finland included integrated data balancing ($M \approx N$), nested cross-validation, and methods for uncertainty quantification (radius distance and Shannon entropy) and interpretability (Shapley Additive exPlanations(SHAP)). The supervised MLP performed well with an Area Under the Curve (AUC) of 0.99, perfect of recall 100%, and Geometric Mean (G-mean) of 0.9937. The Shapley Additive explanations analysis showed that magnetic and pseudo-gravity anomalies are among those more significant features. Findings indicate that a well developed MLP can address significant data imbalance, successfully reducing the investigation footprint to around 1% of the total area while detecting all known deposits. The use of uncertainty maps showed that such deposits are found in high-confidence zones (low-uncertainty) along transitional corridors at geological boundaries, providing a reliable and economical framework for directing mineral exploration.

Keywords Mineral prospectivity mapping, Multilayer perceptron, Contrastive representation learning, Uncertainty estimation, Model interpretability

1 Introduction

Mineral Prospectivity Mapping (MPM) is a geospatial scientific method that focuses on identifying areas with high potential for mineral deposits [1]. MPM economically reduces exploration expenses by focusing on high-probability areas. MPM helps in economic development and understanding of Earth surface [2]. The success of MPM,



© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

however, depends on resolving intrinsic issues with data validity and quality, a process made more difficult by the nature of mineral exploration.

Simplicity and clarity of conventional Geographic Information System (GIS)-based methods—such as logistic regression and weights-of-evidence—have made them rather popular [3]. Nevertheless, these conventional techniques often struggle with imbalanced datasets found in mineral exploration; known mineral occurrences are somewhat rare relative to undiscovered regions. This leads to often erroneous forecasts either omitting important areas or producing too high false positives.

MPM is at a crucial a turning point. As near-surface deposits are becoming more and more scarce, mineral exploration is moving beyond intuition-based tools and towards data-driven, Artificial Intelligence (AI)-based tools that can interpret the complex and non-linear geological patterns in the high-dimensional data [4–7]. Recent reviews show that intelligent lithological mapping continues to be problematic in terms of data fusion and interpretability, and thus requires strong future frameworks [8]. Two major challenges compromise the reliability and adoption of machine learning in mineral prospecting and mining: extreme class imbalance, characterized by a few known deposits overshadowed by millions of background locations [9–11], and the "black box" nature of predictive models, which affects trust and operational applicability [12, 13]. This study handles these challenges using a thorough comparative approach which includes contrastive representation learning, supervised classification, uncertainty quantification, and feature interpretability to attain transparent, high-performance mineral targeting.

We focus on a real-world dataset comprising only 17 annotated mineralized points embedded within 1,843,564 raster cells, representing one of the most imbalanced scenarios in geospatial prediction [14, 15]. To address this, we develop and evaluate two complementary workflows: (1) Contrastive Representation Learning (CRL)-Radius, a novel contrastive representation learning approach that learns a 2D embedding space from multi-source geophysical data (magnetic, gravity, electromagnetic) without relying on class labels during pre-training, then applies a simple global prospectivity circle to identify high-potential zones [16, 17]; (2) a supervised Multilayer Perceptron (MLP) trained using 17-fold nested stratified cross-validation [11, 18], where each outer fold holds out one entire deposit to simulate realistic discovery conditions, and an inner 16-fold grid search optimizes hybrid resampling—Synthetic Minority Over-sampling Technique (SMOTE) oversampling of mineralized points and random undersampling of background data [19, 20];

This study is guided by three research questions:

1. Does contrastive learning create embedding space where a simple global-circle rule both recovers known belts and reveals new corridors
2. With very few positives, how should we balance oversampling (M) and undersampling (N) to get the best generalization without adding too many synthetic points or wasting compute?
3. Do our uncertainty measures (radius distance, Shannon entropy) and SHAP explanations agree, so we know both where to explore and why?

The experiment demonstrates that both CRL-Radius and the supervised MLP achieve high recall and AUC value, effectively recovering all known deposits while concentrating high-prospectivity predictions into spatially continuous, corridor-like belts

that encompass very low study area [21]. For data-balancing state the ideal value was around the diagonal $M \approx N$, which specifically prevents underfitting (low M, high N) and over-fitting (high M, low N), leading to stable F1-scores from [0.92, 1.00]. Shapley Additive exPlanations (SHAP) analysis pinpoints magnetic properties (Mag AS, Mag TD), pseudo-gravity (Pseu Grv) and electromagnetic responses (EM Inph, Em Qd) as the main drivers during both embedding and classification stages [22]. More precisely, the CRL-Radius SHAP values identify the effect of features on the embedding structure while the MLP SHAP values clarify classification decisions, revealing a consistent set of geophysically interpretable variables.

Moreover, both the radius-based and entropy-based uncertainty maps showed that the recognized mineralized points lie in low uncertainty regions enclosed with thin decision corridors, which represent high model confidence where transition regions (i.e., among various classes of rocks) exist in decision spaces, and small differences or geophysical contrast are essential [23]. These locations are coincide with different geological features trends. Used alongside with SHAP, this forms a robust feedback loop: wherever uncertainty is high, more investigation is required; while the feature identification clarify the reasons for the ambiguity or or prospective.

This paper further discusses explainable AI, which implements representation learning, balanced supervised training, and spatial prospectivity mapping. This work enhances geoscience towards being more consistent, clear, and useful at the same time. Their integration is essential in mineral exploration to foster confidence and direct field operations.

2 Related work

Mineral Prospectivity Mapping (MPM) has continuously evolved from traditional methods to advanced machine learning (ML) and deep learning techniques. These advances in machine learning can integrate geoscientific data from different source. Traditionally, mineral exploration were based on statistical techniques such as logistic regression, Weights-of-Evidence (WoE), and fuzzy logic [1, 3, 24]. The transition from classical to data-driven approaches during recent decades made exploration more accurate and reliable by integrating data from different sources, such as remote sensing, geophysics, and geochemistry[2, 9].

2.1 Supervised machine learning based approaches

Supervised machine learning algorithms are commonly used techniques for MPM in the domain of mineral exploration. Numerous machine learning techniques help in recognizing the mineralization. For example random forests, support vector machines, ensemble learning, and multilayer perceptrons can identify complex patterns from geophysical and geochemical signals [9, 11]. Here, Josso et al. deployed random forests to forecast Fe-Mn crusts in coastal environments, while Farahnakian et al. utilized ensemble learning to examine geophysical data for the prediction of minerals in Finland [2, 11]. In order to further improve these classifiers, recent works have been able to apply swarm-intelligence optimization to both Support Vector Machines (SVM) [25, 26] and nature-inspired algorithms such as Harris Hawks optimization to Random Forests [27] in geochemical anomaly detection. The above models outperform classical methods, as they can identify complicated, nonlinear correlations among geological data.

Nonetheless, the dependency of the model on labeled data, susceptibility to class imbalance, and limited capacity to quantify prediction uncertainty still continue to make them difficult to trust and implement in other unexplored domains.

2.2 Unsupervised/self-supervised representation learning

In MPM, unsupervised and self-supervised learning is becoming increasingly important as researchers consider mechanisms to use the limited amount of labelled data and better the transfer of models across regions. Research papers with methods such as Isolation Forest and Extended Isolation Forest and Bayesian Kernel Density Estimation (KDE) have demonstrated that complex geoscientific data can be mineralized with no supervision [28, 29]. On the same note, self-organizing maps, which use a clustering mechanism, are capable of creating interpretable maps to be explored with minimal manual labeling [30]. More recently, the area of representation learning has been given the attention it deserves due to its capacity to overcome label shortage and domain transfer. For example, graph-contrastive learning and Simple Framework for Contrastive Learning of Visual Representations (SimCLR) are used in geoscientific graph analysis, subsurface imaging, and seismic interpretation. An example is contrastive graph attention networks, used to achieve interpretable mapping at regional scale, and they are effective at combining multimodal data [31]. These strategies help models in extracting meaningful information from various datasets and geological contexts. Nonetheless, these techniques have seldom been used on tabular geophysical rasters such as magnetic, electromagnetic and radiometric data which are critical to MPM, which creates a significant research gap.

2.3 Geophysical data fusion & uncertainty quantification

In order to achieve effective MPM, complementary geophysical layers (magnetic, electromagnetic, radiometric) are to be fused in a single predictive signal. The recent literature focuses on multi-modal fusion using deep architects [32] or probabilistic models [33]. Zhang et al. suggested a 3D uncertainty-constrained MPM model that relies on Bayesian neural networks and measures epistemic uncertainty through Monte Carlo dropout [33]. However, the majority of uncertainty estimates do not have an interpretation to them: entropy or variance maps show the points where uncertainty is actually large, but not the cause. Post-hoc MPM explanations have been attempted with SHAP and Local Interpretable Model-agnostic Explanations (LIME) with little or no use of uncertainty layers [34]. The combination of SHAP with Shannon entropy (in MLP) and embedding-distance uncertainty (in CRL) to form a joint where and why system of exploration guidance is presented in our work.

2.4 Spatial imbalance handling

This severe spatial imbalance in MPM (i.e., 17 deposits vs. 1.8 million background points) requires special techniques. Synthetic Minority Over-sampling Technique (SMOTE) and Random Undersampling (RUS) are some of the frequently used resampling techniques to balance training data [19, 20]. More recent attempts propose spatially conscious validation (e.g. spatially k-fold CV) and hybrid resampling [11, 15]. Farahnakian et al. demonstrated that SMOTE + random undersampling has higher recall, but do not provide any recommendations to choose the oversampling (M) and undersampling (N)

ratios, instead of trial and error [11]. This study bridges the gap by mapping M-N trade-offs within nested spatial CV, with systematic M-N mapping illustrating the diagonal, representing the stability zone that neither underfits nor overfits.

2.5 Hybrid and deep-representation approaches

Recent researches leverage the traditional ML classifier with deep learning or self-supervised feature extraction methods to enhance the quality of representations. Deep auto-encoders, convolutional neural networks (CNNs) and graph-based models have been considered to learn spatial relation in geoscientific data [9, 35]. Deep learning architectures have proved to be highly effective in learning the complicated spatiotemporal behavior in various applications, including medical imaging [36, 37], which motivates their implementation to geological settings. The notable ones are the regimentation of indicator elements via convolutional algorithms [38] and the application of geologically-constrained deep learning via frequency domain filtering to mitigate uncertainty [39]. Contrastive Representation Learning (CRL) has become one of the promising paradigms that are learning strong low-dimensional embeddings through maximizing similarity between samples with similar geological origins and minimizing similarity between disparate samples [7]. Even though CRL has been effectively used in computer vision, its use on tabular geophysical data is not widely used. This study will implement a CRL-Radius workflow to learn mineralization-relevant manifolds without explicit supervision and classify them with a simple geometric decision boundary.

The development of MPM has taken place in the form of supervised learning, unsupervised anomaly detection, uncertainty modeling, and hybrid resampling. Nevertheless, the literature reviewed in Table 1 demonstrates that there are still a number of significant gaps that need to be filled. First, contrastive learning has not been experimented with tabular geophysical data, and none of the studies have examined the quality of embeddings or simple geometric decision rules in the case of extreme imbalance. Second, SMOTE with undersampling is commonly used, but there is no clear benchmark for selecting the oversampling (M) and undersampling (N) ratios at scales beyond 1:100,000, which often results in unstable models. Third, uncertainty and interpretability are normally recorded independently, and their field lacks a common strategy that combines uncertainty with feature-level explanations. To address the mentioned gaps, the research paper is providing a combined CRL-Radius and MLP model that introduces embedding-based geometric classification, systematic M-N mapping under nested spatial cross-validation, and combined uncertainty and interpretability to make mineral targeting clearer.

3 Study area and dataset

3.1 Geographic and geological setting

The study area is located in western Finnish Lapland, which covers regions of the Kolari, Kittilä, and Muonio municipalities, encompassing a region approximately 70×75 km (Fig. 1A). The Geological Survey of Finland (GTK) provides diverse information incorporating both Archean and Paleoproterozoic domains as given in Fig. 1B. The Archean basement comprises paraschist, gneiss, Tonalite-Trondhjemite-Granodiorite (TTG) gneiss, and migmatite, as well as quartz-rich and quartz-poor volcanic rocks. The

Table 1 Summary of related works referenced in the Related Work section

References	Approach	Key contribution	Strength	Limitation	Classification
Al-Najjar et al. [15] RS	GAN synthetic sample generation	Uses GANs to generate minority-class samples for spatial susceptibility mapping.	Adds imbalance beyond classical SMOTE; preserves complex patterns.	Computationally heavy; domain transferability uncertain.	Imbalance handling (Hybrid)
Farahkhanian et al. [11] OGR	SMOTE + RUS + ensembles	Systematically evaluates SMOTE + RUS with ensembles for MPM.	Improves minority recall; thorough ablation on imbalance settings.	Spatial autocorrelation and spatial CV not fully addressed.	Supervised ML + Imbalance
Zhang et al. [33] Eng. Geology	3D modeling with UQ (Bayesian/MC Dropout)	Builds a 3D ML framework with explicit uncertainty quantification.	Uncertainty-aware predictions aid decision-making.	High compute demand; workflow complexity.	Hybrid ML + UQ
Josso et al. [2] OGR	Ensemble tree classifier	Applies RF to global Fe–Mn crust prospectivity using multi-source data.	Handles nonlinearities; robust with mixed predictors; strong baseline.	Interpretability limited; marine-focused; transfer to tabular geophysics not shown.	Supervised ML
Sun et al. [32] IEEE TGRS	CRL for hyperspectral anomaly detection	Proposes contrastive background reconstruction for hyperspectral anomaly detection.	Label-efficient; strong anomaly separation.	Primarily imagery; transfer to tabular geophysics not evaluated.	Self-supervised / Contrastive
Liu and Tsoumakas [20] KBS	Undersampling majority class	Shows random undersampling strategies in classifier chains.	Simple; reduces training time; avoids synthetic artifacts.	Information loss; can underfit minority boundaries.	Imbalance handling
Yang et al. [34] Math Geo	CNN with SHAP/LIME	Designs a dual-branch CNN with post-hoc interpretability for MPM.	End-to-end features; integrates XAI for geological insight.	Data-hungry; architecture complexity; dataset specificity.	Deep learning + XAI
Zeghouane et al. [3] AJG	Bayesian evidential overlay	Implements WoE to integrate evidential layers for Sn–W prospectivity.	Simple, probabilistic integration; transparent.	Conditional-independence assumption; weak for nonlinear effects.	Knowledge-driven GIS/Statistical
Zuo [1] NRR	Synthesis of GIS/ML methods	Synthesizes GIS- and ML-based MPM approaches and workflows.	Comprehensive scope; clear taxonomy of methods.	Limited coverage of recent self-/contrastive learning and UQ coupling.	Review
Porwal & Carranza [24] C&G	Intro to GIS-based MPM	Positions ML within MPM workflows and challenges.	Sets context; clarifies opportunities for ML.	Not a quantitative benchmark; limited technical detail.	Editorial
Zuo & Carranza [9] Math Geo	Overview of ML for exploration	Synthesizes GIS- and ML-based MPM approaches and workflows.	Comprehensive scope; clear taxonomy of methods.	Limited coverage of recent self-/contrastive learning and UQ coupling.	Editorial/Review
Chawla et al. [19] JAIR	Synthetic oversampling	Introduces the SMOTE algorithm for synthetic oversampling.	Enhances minority representation; simple and generic.	May create overlapping classes; ignores spatial autocorrelation.	Imbalance handling
Saremi et al. [6] Minerals	Isolation-based outlier detection	Applies Isolation Forest and Extended Isolation Forest for unsupervised mineral prospectivity mapping.	No labels required; robust to noise; anomaly-focused.	Hyperparameter sensitivity; limited interpretability.	Unsupervised anomaly detection

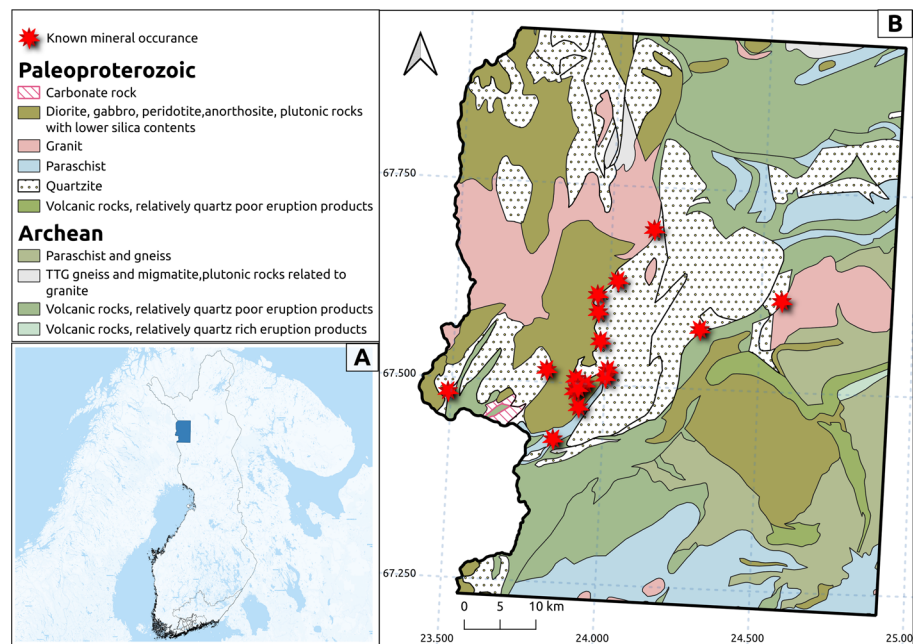


Fig. 1 Study area geology with known mineral occurrences (red stars). **A** The study area location in Finland. **B** The generalized geology of location area

Paleoproterozoic formations consist of carbonate rocks, quartzite, paraschist, granites, and mafic-to-intermediate plutonic rocks, including diorite and gabbro.

This research mainly focuses on iron oxide-copper gold (IOCG) mineralisation. Figure 1B shows 17 known IOCG mineral occurrences in the area, which are marked with red stars. Most of these deposits are found along faults where synorogenic monzonite intrusions intersect the rock that surrounds them. The information related to these known mineral occurrences was taken from the national mineral deposit database of the GTK.¹

3.2 Geophysical datasets and features

The Geological Survey of Finland (GTK) provides a diverse archive of geoscientific data, which includes airborne geophysical data, lithological maps, till geochemistry, etc. In this study, however, we intentionally used airborne geophysical data. The above strategy complies with the main goals of the Horizon Europe Exploration Information System (EIS) project that focuses on deep-penetrating indicators of the hidden mineral systems. The modern exploration of mineral deposits is more interested in deep-penetrating geophysical techniques, which play a vital role, whereas the surface information is less informative [5]. In addition, the geophysical datasets have a consistent coverage and spatial resolution (50 m x 50 m) of the entire area, which is important for minimizing sampling bias and achieving robustness in machine learning analysis [6].

Surface-based datasets, like till geochemistry and detailed lithological maps, have not been incorporated since they have an irregular spatial distribution in most instances. These datasets are primarily near-surface processes and in most cases, they have to be heavily interpolated to get continuous coverage which may impose some level of uncertainty and bias in the analysis.

¹ https://tupa.gtk.fi/paikkatieto/meta/mineral_deposits.html

3.2.1 Data generation and processing

The layer of geophysical evidence was created by The Geological Survey of Finland (GTK) based on low-altitude aerial geophysical surveys that were conducted between 1972–2007 and flown with a standard 200 m line spacing and with a nominal terrain clearance of about 30 m. The preprocessing of these survey data was carried out using common methods like leveling and micro-leveling. The operations reduce the noise of the data and improve the geographic uniformity of the study area. The procedures resulted in spatially uniform raster grids which could be utilized in the analysis at the regional scale. Local geomagnetic parameters were considered because the magnetic measurements are affected by the main magnetic field of Earth. The magnetic inclination of the area of study is 77.1° and magnetic declination is 7.7°E at the equivalent location of 67.5°N and 24.0°E respectively as per the National Oceanic and Atmospheric Administration (NOAA) / National Centers for Environment Information (NCEI) International Geomagnetic Reference Field (DGRF2000).

To ensure objectivity and reproducibility, no manual weighting of expert and/or subjective enhancement of the individual layers was performed on the individual layers. Before being used in any machine learning method, all geophysical layers were resampled to a uniform spatial resolution of $50\text{ m} \times 50\text{ m}$. The layers were then scaled to the numerical range between 0 and 1, thus ensuring that there was equivalent spatial weighting among datasets that had different physical units and dynamic ranges [6].

The final resulting dataset comprised 13 geophysical evidence layers, which are derived using magnetic, airborne electromagnetic, and radiometric data.

3.2.2 Mineral systems framework and evidence layer selection

The IOCG mineral system is viewed in mineral prospectivity mapping as a set of interacting ore-forming subsystems, which operate before, during and after mineralization [5, 40]. These subsystems are made up of metal and fluid sources, fluid transportation routes, mineral deposition, and preservation conditions. All geophysical evidence layers were thus identified and subjected to processing to reflect one or more of these subsystems using measurable physical properties.

- *Airborne Electromagnetic (AEM) Data (Traps and Deposition)*: The AEM dataset accounts for the electrical properties of the subsurface.² The features it includes are the in-phase and quadrature components of the secondary field, apparent resistivity, and the electromagnetic ratio. The above features are sensitive to variations in the electrical conductivity of the subsurface and give an indirect measure of the changes in hydrothermal, fluid circulation and lithological difference [41]. They show the process of mineral formation. Areas that have high electrical conductivities usually indicate the presence of sulfide minerals or fluid openings. These areas may either be physical or chemical traps where the metals are accumulated or deposited with time [5, 41].
- *Magnetic Data (Source and Pathways)*: The magnetic data show how the Earth's magnetic field changes because of the qualities of rocks underneath the surface.³ The magnetic data is available in the form of Total Magnetic Intensity (TMI) and its

²https://tupa.gtk.fi/paikkatiето/meta/aeroelectromagnetic_raster_data_of_finland.html

³https://tupa.gtk.fi/paikkatiето/meta/aeromagnetic_raster_data_of_finland.html

different derivatives, like the analytic signal, tilt derivative, the directional cosine of the tilt (45°), and pseudogravity. These derivatives contribute to the strengthening of the geological boundaries, structural discontinuities, and the contrast of variations between subsurface densities. In mineral systems, the magnetic layers mostly illustrate processes that happen before and during mineralization. TMI and the pseudogravity are sensitive to the intrusive rocks and are sources of heat and metal [42]. In addition, the derivatives show faults and structural corridors. These corridors are channels through which the hydrothermal fluids are transported through the crust [40].

- *Radiometric Data (Preservation and Footprint)*: Radiometric dataset involves the use of natural gamma radiation on the surface of the earth to establish the quantity of the radioelements present in the air.⁴ These data include potassium (K), uranium (U), thorium (Th), and total gamma-ray. These layers consider both the natural gamma radiations and near-surface geochemical variations [43]. Radiometric characteristics mainly represent the subsequent phases of mineralization. They are used to demonstrate the degree of preservation of a mineral system. These features determine areas of alteration e.g. potassic alteration and broader geochemical patterns in relation to IOCG mineralization [6, 43].

Figure 2 shows 13 geophysical evidence layers and positions of 17 known mineralized occurrences. These layers provide information on different physical properties of the subsurface, such as electrical conductivity (AEM data), magnetic susceptibility, magnetic derivatives structural features, and near-surface geochemical data (radiometric K, U, Th and total gamma). The layers has different features of mineral systems such as trapping events, fluid pathways and geochemical signatures. The majority of mineral occurrences

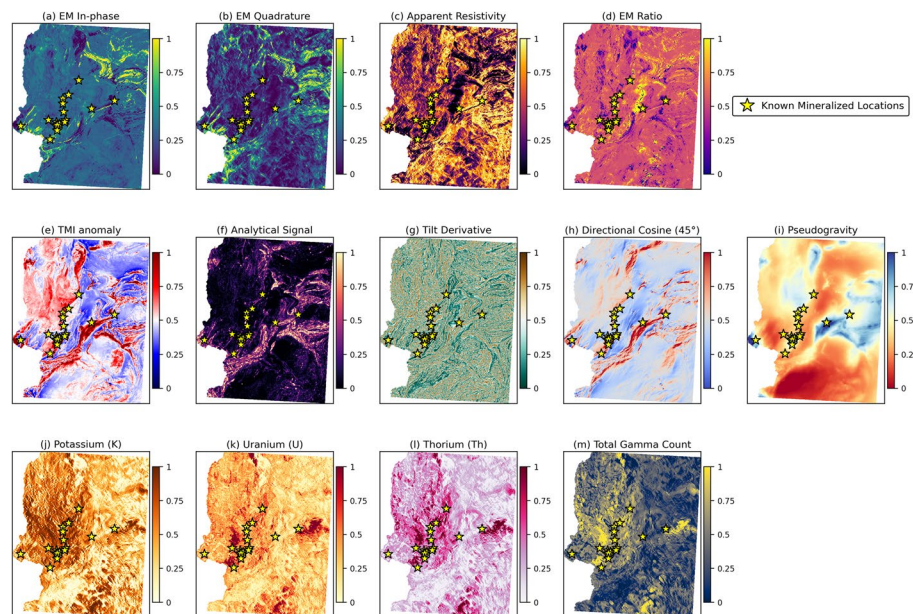


Fig. 2 Generated geophysical evidence layers used in the study area (Lapland, Finland). All 13 maps are normalized to the range between 0 and 1 for consistency. The yellow stars locate the position of the 17 known mineralized IOCG occurrences. The data includes **a–d** electromagnetic, **e–h** magnetic intensity and generated derivatives, **i** pseudogravity, and **j–m** radiometric. Data sourced from GTK

⁴https://tupa.gtk.fi/paikkatieto/meta/aeroradiometric_raster_data_of_finland.html

are related to anomalies, which can be found in several evidence layers. Therefore, this spatial relationship reveals that the geophysical evidence layers can provide valuable geological information for mineral prospectivity mapping.

3.3 Final dataset composition

A dataset with two classes, known deposit locations and geophysical background points, presents an extreme imbalance for mineral prospectivity mapping (MPM). The dataset has 17 annotated (mineralized) samples, which are the places where positive (the known IOCG deposits) are located. The other 1,843,564 raster (background) points, which are not known to have large deposits, are the negative (non-mineralized) samples. The 13 geophysical features listed above describe each sample, whether it is positive or negative. The ratio to the mineralized to the non-mineralized points is approximately 1 to 108,445.

All background points considered as negative represent a simplifying assumption. Carranza et al. discusses that background locations in MPM should be unknown because they might contain mineral deposits [44]. However, the simplifying assumption is commonly used when working with large and imbalanced data in MPM. In particular Farahnakian et al. used the above mentioned approach on IOCG data, treating the background data as non-mineralized for machine learning modeling [11].

The above assumption is valid in this research due to two reasons. To begin with, the fact that the frequency of mineral deposits is very low (approximately 1:108,445) implies that the chance of accidental sampling and discovering a latent deposit and declaring it negative is very low. Second, random undersampling of samples eliminates this potential further, given that the likelihood of such sparse latent deposits taking over the training data is minimized. Therefore, the backgrounds are mostly regarded as non-mineralized, yet they play the critical role in mineral exploration. False positive is not necessarily inappropriate in MPM. It is typically a sign of the location of a new deposit. The model is made to identify these high-probability locations so they can be further explored.

4 Methodology

This section details the computational methodology to develop and evaluate mineral prospectivity mapping (MPM) in extremely imbalanced data environment. The main issue is the presence of very few amount of known mineral locations compared to the whole background area. In order to handle this class imbalance, robust validation and two different pipelines namely, (1) a geometric (radius) classifier based on Contrastive Representation Learning (CRL-Radius), and (2) a supervised Multilayer Perceptron (MLP), were implemented.

This methodological framework aims at addressing significant issues in MPM through the integration of state-of-the-art techniques. The problem of severe data imbalance is handled by implementing effective tactics such as data resampling and classification threshold optimization— methods supported in recent MPM research [10, 11].

4.1 Nested cross-validation

A powerful validation techniques called Nested cross-validation was implemented to obtain an unbiased estimate of the model performance and tune hyperparameters. This two-level approach is important for the prevention of information leakage from the model tuning process to the final performance evaluation.

The stratification in Stratified K-Fold Cross-Validation (SKCV) is vital since it makes sure that the very small number of positive samples persists the same across all folds, such that no fold has zero positive cases.

- *Outer Loop (Model Evaluation)*: A 17-fold Stratified K-Fold Cross-Validation (SKCV) was deployed for the final model assessment. Given the 17 annotated samples, this procedure is equivalent to a Leave-One-Out Cross-Validation (LOOCV) for the positive class. During each of the 17 iterations, an annotated point is held out as the test set, and the model is trained on the remaining 16 annotated points along with the corresponding background raster data. The performance metrics from each of the 17 test sets are collected and averaged to produce the final reported model performance.
- *Inner Loop (Hyperparameter Tuning)*: For each of the 17 outer folds, a different 16-fold SKCV was run on the 16 training data. The only thing this inner loop does is use Grid Search to get the best hyperparameters. A model is trained on 15 annotated points and tested on the last one for each set of hyperparameters. The hyperparameter set that gives the best average performance over these 16 inner folds is chosen and used to train the final model for that outer loop.

4.2 Class imbalance handling

A major challenge in mineral prospectivity is the class imbalance, as mineralized locations are scarce compared to the background points. To fix this:

- *Synthetic Minority Over-sampling Technique (SMOTE)*: This method makes fabricate samples of the minority class in the training.
- *Random Undersampling (RUS)*: Under samples negative examples to get a balanced class distribution, thereby preventing the model from being biased towards majority.
- *Fold-wise application*: To prevent data leakage, Both oversampling and undersampling are implemented on the training split of each fold to prevent from data leakage.

This hybrid approach makes sure that classifiers get required minority and majority samples during training without altering the validation or test data.

4.3 Foundational methodologies

4.3.1 Multilayer perceptron

A multilayer perceptron (MLP) is a type of supervised deep neural network that contains multiple layers of nodes, encompassing an input layer, one or more hidden layers, and an output layer. Except for those in the input layer, every node is a neuron that makes use of a non-linear activation function. This structure has been specifically created to replicate the behavior of neurons in the human brain.

Figure 3 shows a picture of the fundamental components of a multilayer perceptron (MLP). It shows a general block diagram that shows how the MLP architecture functions when put together. Input Layer: This layer gets the input signal. Hidden layer: This is the name for the levels that are in between the input and output layers of a multilayer perceptron (MLP). These layers are mostly in charge of doing the real math that the MLP does. Output Layer: This layer makes the network's final prediction or output.

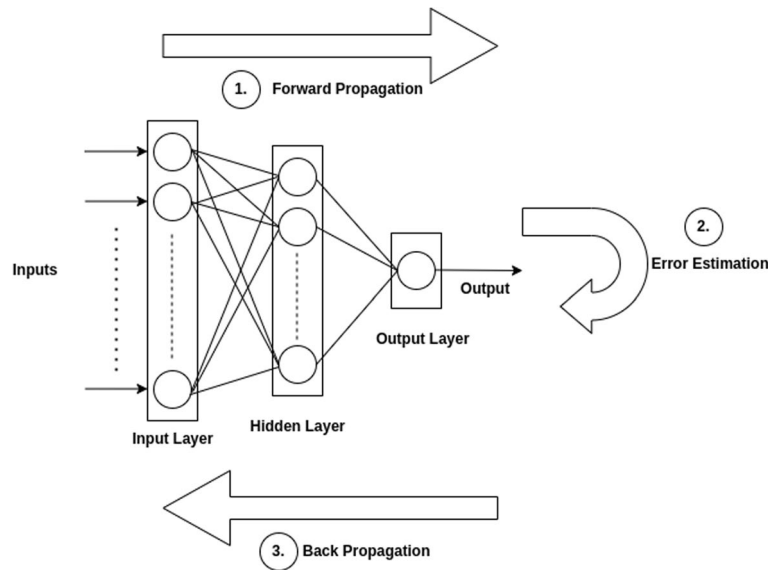


Fig. 3 Representation of multilayer perceptron neural network




Fig. 4 Schematic representation of a siamese network for contrastive representation learning

4.3.2 Contrastive representation learning

Siamese Networks have emerged as a fundamental architecture in the domain of contrastive representation learning. Their objective is to acquire a feature space in which similar samples are grouped together, while dissimilar samples are spread apart. The Siamese Network has been widely adjusted for several representation learning tasks, especially in situations when labeled data are scarce but an abundance of unlabeled data.

As illustrated in the Fig. 4, a Siamese Network consists of two or more identical subnetworks, usually neural networks, that share similar architecture and weights. Following the processing of one of the two input samples by each subnetwork, the outputs of these subnetworks are subsequently compared using a contrastive loss function.

Twin Networks At the heart of a Siamese Network, there lie two identical subnetworks. Typically, these subnetworks consist of deep neural networks such as multilayer perceptron (MLP) or convolutional neural networks (CNNs) designed for processing picture input. The primary characteristic is that both subnetworks possess identical parameters (weights and biases), which means that the procedure of extracting features is the same for both inputs.

- *Input Layer:* The Siamese Network is capable of processing two inputs, which may include pictures, text, or any other form of statistical information. Typically, these inputs are denoted as x_1 (positive pair) and x_2 (negative pair).
- *Shared Network Architecture:* Every input passes through a network f_θ , which is parameterized by θ . The function $f_\theta(x)$ transforms the input x into a feature space

that allows for the quantification of the similarity or dissimilarity between the input features. Given that the networks are identical, we have $f_{\theta}(x_1)$ and $f_{\theta}(x_2)$.

- *Embedding Layer:* The embedding layer can use the results from the subnetworks to make feature vectors with fewer dimensions. That means these vectors can be used to estimate the similarity between different inputs.
- *Distance Measure:* Distance based measure is a measure which is used to measure the outputs of the two networks. The Euclidean distance is the most frequent metric, however depending on the application, other metrics like cosine similarity can also be used. The Euclidean distance between the outputs $f_{\theta}(x_{1i})$ and $f_{\theta}(x_{2i})$ is calculated via the following formula, where ϵ (epsilon) is a small constant (e.g. 10^{-7}) added within the square root to ensure numerical stability and prevent the undefined gradient in case the distance is extremely small:

$$D_i(f_{\theta}(x_{1i}), f_{\theta}(x_{2i})) = \sqrt{\|f_{\theta}(x_{1i}) - f_{\theta}(x_{2i})\|_2^2 + \epsilon} \quad (1)$$

- *Contrastive Loss Function:* In contrastive learning, the objective is to maximize the distance between different pairings (negative pairs) and reduce the distance between similar pairs (positive pairs). A contrastive loss function is employed to do this, consolidating the embeddings of similar pairs and separating the embeddings of dissimilar pairs.
- For a batch of N sample pairs, the contrastive loss \mathcal{L} is given by:

$$\mathcal{L} = \frac{1}{2N} \sum_{i=1}^N \left[y_i D_i^2 + (1 - y_i) \max(0, m - D_i)^2 \right] \quad (2)$$

Where:

The Euclidean distance between the feature vectors is denoted by D_i , as defined in Eq. (1).

y_i is a binary label; a value of 0 denotes a dissimilar pair of inputs and a value of 1 indicates a similar pair of inputs.

A margin parameter called m indicates a minimum amount that the network will attempt to enforce between dissimilar pairings.

The distance between similar pairs is kept to a minimum by the first term, $y_i \cdot D_i^2$.

This second term, $(1 - y_i) \cdot \max(0, m - D_i)^2$, encourages dissimilar pairs to be at least m distance units apart by penalizing negative couples only if their distance is smaller than the margin, m .

By minimizing this loss function during training, the network is effectively tuned to generate representations that separate dissimilar inputs from similar ones and cluster similar inputs together.

4.4 Proposed modeling pipelines

Two different workflows were applied using methods described above of MLP and CRL, to assess their performances for producing mineral prospectivity mapping. Each pipeline is unique in its approach to learning features and classification. The next parts illustrate

a detailed, step-by-step description of each workflow, which aligns with the processes shown in the diagrams.

4.4.1 Contrastive representation learning with radius classifier(CRL-radius)

The CRL-Radius pipeline shown in Fig. 5 uses Contrastive Representation Learning to transform high-dimensional geophysical data into a compact 2D embedding space where mineralized points form dense clusters, making classification possible through a simple geometric proximity rule.

The compression of 13 geophysical features into embedding space is not intended to capture all the variations of the feature, as linear methods such as Principal Component Analysis (PCA) do. The contrastive technique enhances the embedding by closely mapping samples related to mineral deposits while separating different samples. In this context, the method acts like a discriminative compression tool. Therefore, it boosts the geophysical signals that help in identifying minerals and suppresses the ones that do not contribute to this task.

The process begins with inner-fold contrastive training. In each inner fold, 15 annotated deposits together with associated raster points are used to train a Siamese encoder f_{θ} using the contrastive loss described in Sect. 4.3.2. The encoder learns an embedding function that brings similar samples closer and pushes dissimilar ones apart. Once trained, the encoder is applied to all relevant data in that fold—including the held-out validation deposit and background points—to generate 2D embeddings for every sample.

Within each inner fold, the cluster center is computed as the mean of the 15 training-deposit embeddings. A candidate classification radius is then set as the Euclidean distance from this center to the inner validation annotated embedding. This candidate

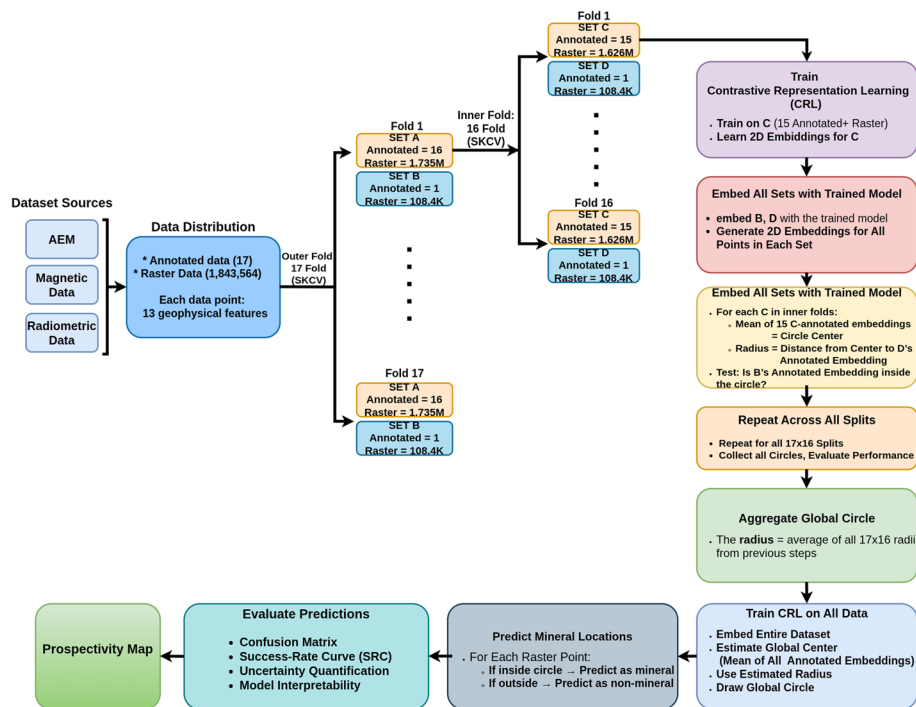


Fig. 5 Mineral prospectivity workflow using contrastive representation learning and radius-based classification

radius represents the decision boundary for that fold. Its effectiveness is assessed by verifying that the outer validation annotated point lies inside or outside the circle.

This procedure is repeated across all 17×16 outer-inner fold combinations, resulting in multiple circles. These are aggregated to form a global classification circle and the radius is taken as the average of the optimal radii obtained from all folds.

After completing cross-validation, the Contrastive Representation Learning model is retrained on the entire dataset, using all annotated and raster points. The trained encoder embeds the full dataset into the 2D space, from which the global center is calculated as the mean embedding of all annotated points. Using the previously estimated global radius, a global prospectivity circle is drawn in the embedding space. Each raster point is then evaluated: points falling inside the global circle are predicted as probable mineralized locations, while those outside are predicted as non-mineralized.

Predictions are evaluated using confusion matrices, Success-Rate Curves (SRC), uncertainty quantification, and SHAP-based interpretability applied to the embeddings and classifications. The final output is a prospectivity map that highlights regions with the highest potential for mineralization.

4.4.2 Supervised multilayer perceptron

This pipeline uses a nested spatial cross-validation method to make sure that model evaluation is robust along with the optimal hyperparameter tuning as shown in Fig. 6. The dataset is first divided into 17 outer folds (spatial K-fold cross-validation, SKCV), where each fold holds out one annotated deposit and its associated raster points as the test set. Within each outer training set, a 16-fold inner SKCV is used for hyperparameter optimization and decision threshold selection.

During each inner fold, the training data is balanced using a hybrid resampling strategy. Annotated (positive samples) points are oversampled with SMOTE to synthetically expand the minority class from 15 points to 1,626,659, while negative samples (raster points) are randomly undersampled to match the same count, resulting in a balanced

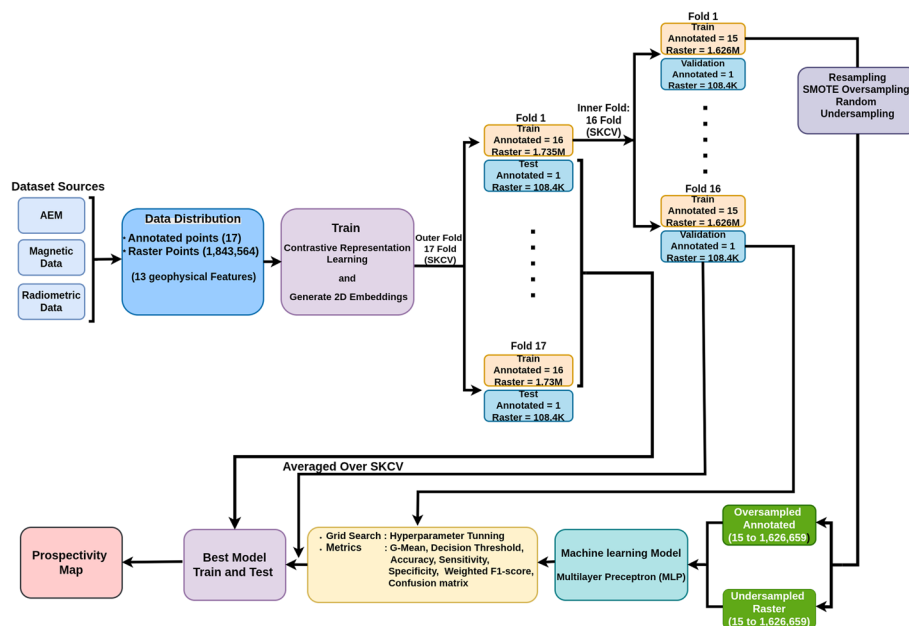


Fig. 6 Mineral prospectivity mapping via multilayer perceptron with nested spatial cross-validation

dataset. These resampled data is then used to train the MLP classifier one various sets of hyperparameters. Hyperparameters tuning is performed on various sets of hyperparameters in the inner loop.

The model selection is done with the help of a crucial metric called Geometric mean (G-mean) used in class imbalance scenario. We consider one of various evaluation metrics such as accuracy, sensitivity, specificity, weighted F1-score besides G-mean along with some confusion matrices to have an extensive understanding of the model. Once you have found the best one, you retrain the model on the entire outer training set using these best parameters. Next, it make predictions on the outer test set that corresponds with this.

This process is repeated across all 17 outer folds, and metrics are averaged to produce an unbiased performance estimate. After cross-validation, the best-performing MLP configuration is retrained on the entire dataset and used to score every raster point, producing both probability maps and binary prospectivity maps indicating predicted mineral potential.

4.5 Model evaluation and selection

To ensure that we could obtain an unbiased performance estimate and perform robust model selection, we conducted nested cross-validation in the model evaluation. The candidate model was trained and evaluated multiple times with different resampling settings to deal with severe class imbalance in the inner loop. The geometric mean (G-mean) was selected as the main criterion because it strikes a balance between sensitivity and specificity, thus being quite fair when one of the classes is much smaller than the other. The Additional metrics to calculate include the accuracy, sensitivity, specificity, weighted F1-score, and confusion matrices.

Receiver Operating Characteristic (ROC) curve is used to test a classifier's ability to distinguish between the classes. This graph shows the true positive rate in function of the false positive rate. The curve provides a more detailed picture showing model performance at different thresholds, while the Area Under the Curve (AUC) summarizes this data into one number. Perfect discrimination is represented by an AUC of 1.0 for the model. Since a random classifier will have an AUC of 0.5; any higher value of AUC suggests that the classifier is better, in terms of discrimination.

A success-rate curve, or SRC, is a plot of abundance (as the cumulative percentage of known deposits captured; y-axis) versus area (as the cumulative percentage of area inspected; x-axis), which ranks all cells in order by model scores [45]. The steeper the SRC, the better the priority; the diagonal indicates random selection. In conjunction with the SRC, the Prediction-Area (P-A) plot is also utilized to determine the practical utility of the model to exploration targeting [46].

The inner loop identified the optimal configuration, the corresponding model was then retrained on the entire outer-training set and evaluated on held-out outer-test set. This separation of model tuning and evaluation avoided information leakage, leading to performance estimates that are transferable to new spatial locations. At last, the variance estimation including uncertainty maps was implemented measuring model stability among folds in outputs of prospectivity.

4.6 Uncertainty quantification and feature interpretability

There are two ways of measuring uncertainty. In the case of the MLP-based model, we use the Shannon entropy of the projected probability p as:

$$H(p) = - [p \log p + (1 - p) \log(1 - p)] \quad (3)$$

in which the outcome is normalized into $[0, 1]$.

Entropy is strongest at $p=0.5$ and lowest near 0 or 1, therefore high-entropy zones are decision borders and suitable follow-up spots. The distance to the global prospectivity circle in the 2-D embedding determines CRL-Radius. Values near zero imply severe uncertainty, whereas large positive or negative distances indicate confidence.

SHAP uses Shapley values to assign each feature a local contribution to a prediction. Positive values push a sample toward the mineralized class, and negative values pull it away; averaging absolute values ranks global importance. Applied to CRL-Radius, SHAP shows which features move points toward the deposit-centered region of the embedding. Applied to the MLP, it explains which features raise or lower predicted probability, giving a clear link between geophysics and model output.

4.7 Final prospectivity map generation

The generated models were ultimately used to predict the probability of mineralization for all raster cells in the study area and yielded the final prospectivity maps. The model generated non-binary probabilities that were displayed as a prospectivity probability map representing the gradient of mineral potential throughout the region. Using the optimal decision threshold determined during the inner-loop cross-validation, a binary prospectivity map was generated.

5 Implementation

All experiments were conducted in a controlled computational environment to ensure reproducibility and scalability of the proposed method. We used several statistical, machine learning, and deep learning frameworks to clean, train models and for visualization. The computational pipeline was executed on a high-performance computing (HPC) system with specialized resources for CPU-bound and GPU-accelerated work loads, as appropriate based on model complexity. Systematic log of each experiment was maintained that included splits, model configuration, hyperparameters, and generated outputs as probability maps or the classification label.

Model evaluation was conducted based on nested cross-validation. The outer loop had 17 folds, in each of which one annotated deposit was left out with the rasterized neighborhood for testing, and the rest constituted the training set. In every outer training set, a 16-folds inner loop was performed for hyperparameter optimization and model selection. We performed preprocessing, resampling, and parameter tuning steps only inside the inner folds so data leakage was not an issue. Performance metrics were calculated on the outer test sets and averaged over the 17 folds and bootstrapped confidence intervals.

For imbalance handling, SMOTE was applied to oversample the minority class (annotated deposits) and random undersampling was used to reduce the majority class (background raster points), achieving approximately a 1:1 class ratio. Resampling was applied only to training splits in the inner loop; validation and test splits preserved the original class distributions. All features were standardized using a standard scaler fit on the

inner-training data, with separate fitting for each outer fold. For contrastive models, features were standardized prior to encoder training, but the resulting 2D embeddings were left unscaled.

Two pipelines were implemented. The first, CRL-Radius (Fig. 5), employed a Siamese MLP encoder trained with a contrastive loss to produce 2D embeddings where mineralized deposits formed compact clusters. In each inner fold, the center of the mineralized cluster was defined as the mean of embeddings from the 15 training deposits, and the classification radius was calculated as the Euclidean distance to the embedding of the single validation deposit. This radius was validated against background intrusion and averaged across all inner folds to form a global decision boundary. After cross-validation, the encoder was retrained on all data, the full raster was embedded, and points were classified as mineralized if they fell inside the global circle.

The second pipeline, supervised multilayer perceptron (MLP, Fig. 6). The custom tabular dataset uses these 13 geophysical features as input, fed into a Rectified Linear Unit (ReLU) feedforward architecture with batch normalization and dropout. Inner loop hyperparameter tuning: hidden layer widths, dropout rates, learning rate, batch size, weight decay and class weighting. Models were trained using binary cross-entropy loss and Adaptive Moment Estimation (Adam) optimization with early-stopping on inner-validation G-mean.

To guarantee reproducibility and transparency in the form of the model architecture, a thorough grid search was used to determine the best hyperparameters. Table 2 provides the search-space that was actually searched in the inner loop cross-validation and the optimal values found to be used in the final MLP model. Since the Geometric Mean (G-mean) is robust to extreme class imbalance, it was given priority during the selection process.

In second pipelines, the primary model selection metric in the inner loop was the geometric mean (G-mean) of sensitivity and specificity, which provides a balanced measure under severe class imbalance. Outer-loop evaluation included G-mean, sensitivity, specificity, weighted F1-score, accuracy, ROC-AUC curve, and confusion matrices.

Model predictions uncertainty was measured by the variability in classification outcomes between outer folds. An uncertainty map was created to represent the variance of the predicted probabilities, which was displayed alongside the corresponding prospectivity probability map. The final outputs consisted of probability maps and uncertainty layers.

Table 2 Detailed hyperparameter search space and optimal values for the MLP model

Hyperparameter	Search space	Optimal value
Hidden layer sizes	[(2), (4), (6), (8), (10), (16), (32), (2, 2), (4, 2), (8, 2), (16, 2), (16, 4)]	(16, 2)
Activation function	[relu, tanh, logistic]	relu
Initial learning rate	[0.0001, 0.001, 0.01]	0.01
Solver	[adam, sgd]	adam
L2 penalty (Alpha)	[0.0001, 0.001, 0.01, 0.1]	0.0001
Batch size	[32, 64, 128, 256]	256
Max iterations	[1000]	1000

6 Results and discussion

6.1 Contrastive representation learning with radius classifier(CRL-radius)

The CRL-Radius pipeline constructs a low-dimensional embedding where mineralized points form a compact cluster, and classification reduces to a geometric rule: points are predicted “mineralized” if their embedding lies inside a global prospectivity circle. The circle’s center is the mean embedding of annotated deposits and its radius is aggregated from 17×16 inner–outer validation splits.

6.1.1 Inner–outer validation behavior

In each outer fold, we ran a 16-fold inner loop. For a given inner fold, the center of the prospectivity circle is the mean embedding of the 15 training deposits (blue points; black star). The radius is set by the inner-validation deposit (purple point) as its Euclidean distance from that center. We then test the outer-held-out deposit (green point) by checking whether its embedding lies inside the dotted circle.

As the example in Fig. 7 shows, the inner-validation deposit typically lies near the edge of the training cluster, so the radius envelopes the cluster without being overly large. The outer-held-out deposit falls inside the circle, illustrating that the geometry learned from the inner loop generalizes to the unseen outer test case.

Aggregating all splits yielded a global radius of 0.1400 and a global center of (0.4636, 0.5643) in embedding coordinates, which is shown in Fig. 8.

6.1.2 Classifier outcomes

After all the dataset was embedded and the global prospectively circle was generated, 5,776 raster points were classified as being inside the circle, with 1,837,888 considered outside. Figure 8 shows the confusion matrix with all annotated deposits as positives ($n = 17$).

- *True Negatives (TN)*= 1,837,888 (99.69%) Background points were correctly predicted as non-mineralized.
- *False Positives (FP)*= 5,776 (0.31%) background points that we mistakenly classified as mineralized.

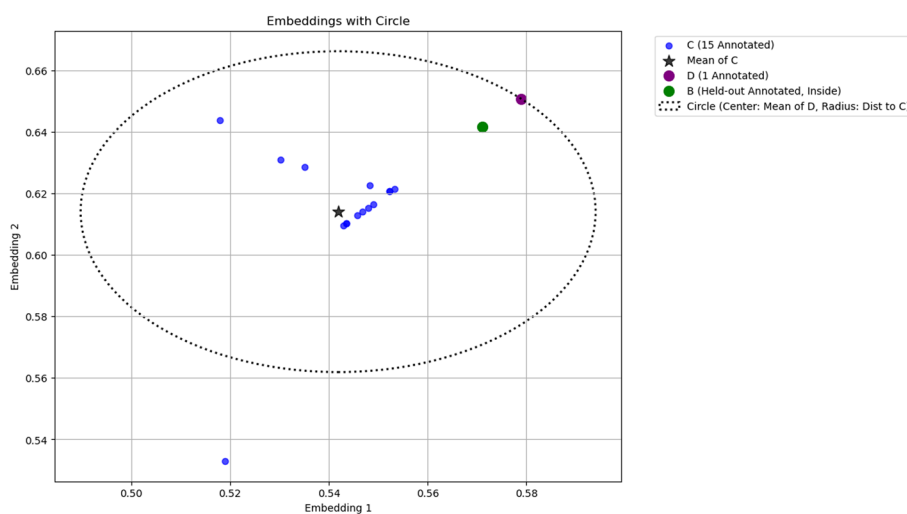


Fig. 7 One outer/inner split; training deposits (blue), mean center (black star), inner-validation deposit (purple) defining the radius, outer-held-out deposit (green), and the resulting circle

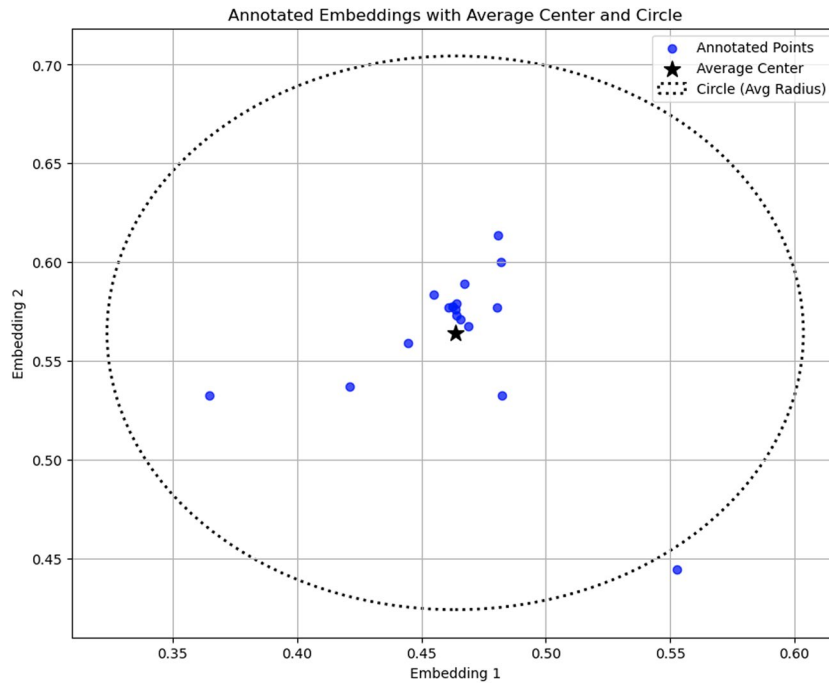


Fig. 8 All annotated deposits in the embedding space with the global circle

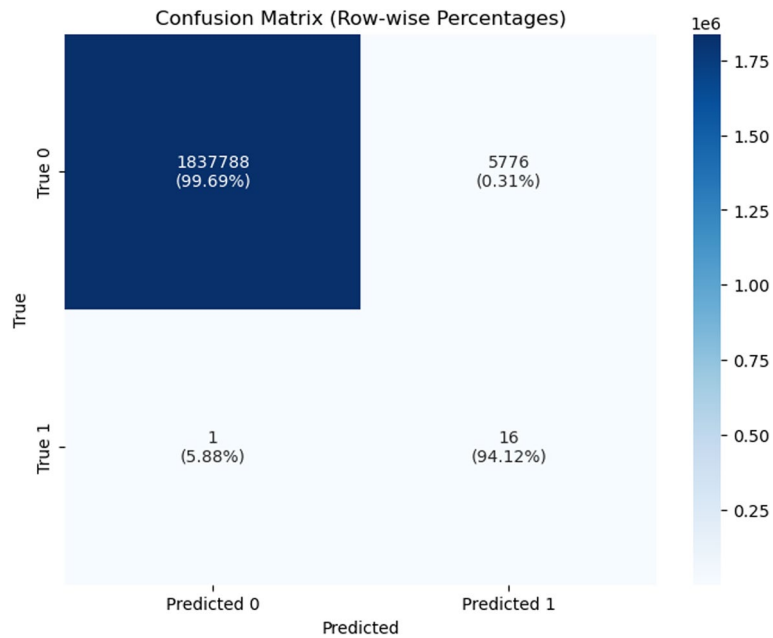


Fig. 9 Confusion matrix for CRL-Radius classifier

- *False Negatives (TN)*= 1 (5.88%) A known deposit miss-classified by the classifier.
- *True Positives (TP)*= 16 (94.12%) → Known deposits classified as mineralized correctly (Fig. 9).

Effective Calculation and Interpretation of Key Metrics within Mineral Prospectivity Mapping (MPM):

Whereas the last metric is almost "perfect", the accuracy of 99.69% sounds very good, but in MPM this metric does not give much information because of extreme class imbalance.

- *Recall/Sensitivity (94.12%)*: Nearly all known deposits are identified. In MPM, this metric was significant because missing a deposit (FN) could mean overlooking a valuable target.
- *Specificity (99.69%)*: Almost all of the background points are excluded from the mineralized class, providing an effective control against large-scale overprediction.
- *Precision (0.28%)*: A small ratio of predicted mineralized points are actual known minerals. At early stage exploration precision will inevitably be quite low as the objective is to include all true deposits (at the cost of additional false positives).
- *F1-score (0.55%)*: This value is low due to low precision value, which is acceptable for general targeting.
- *Weighted-F1 (99.84%)*: High because the dominant back-ground class is almost perfectly classified.
- *G-mean (0.9686)*: Good balance between correctly detected deposits and avoiding false alarms.

In the context of mineral prospectivity mapping(MPM), the outcome reflects a deliberate trade-off between capturing all true mineralized locations and limiting false alarms. The model achieves very high recall, meaning it successfully identifies nearly every known deposit, which is critical in exploration since missing a deposit (false negative) could mean overlooking a valuable target area. This comes at the cost of low precision—many of the points predicted as mineralized are actually background—but in early-stage targeting, this is acceptable. A small false positive rate applied to a very large background dataset naturally yields thousands of candidate locations, yet these extra points are manageable because they serve as a broad initial search space.

6.1.3 Ranking power and threshold

Test performance (from higher to lower) and threshold behavior demonstrate the strength of the model in CRL-Radius. That is, Fig. 10 demonstrates an AUC of 0.9996, demonstrating nearly excellent separability between points related to mineralization and those background in the learned embedding space forever (predicated on the inverse-distance score).

For exploration decision-making, the Success-Rate Curve (SRC) [45] offers a more meaningful perspective as shown in Fig. 11. The SRC shows that a small top-ranked fraction of the study area contains the majority of known deposits, confirming the model's ability to concentrate exploration effort into high-yield zones.

The global circle radius works as a human-understandable decision-threshold. Reducing the radius helps to eliminate false positives and increase precision but might result in some deposits being missed out (lower recall), suitable to high-confidence, focused drilling campaigns. This in turn raises recall (i.e., most of the deposits are captured) but at the cost of more false positives which is theoretically an ideal condition for a wider regional early-stage screening work-up process.




Fig. 10 ROC curve for the radius workflow

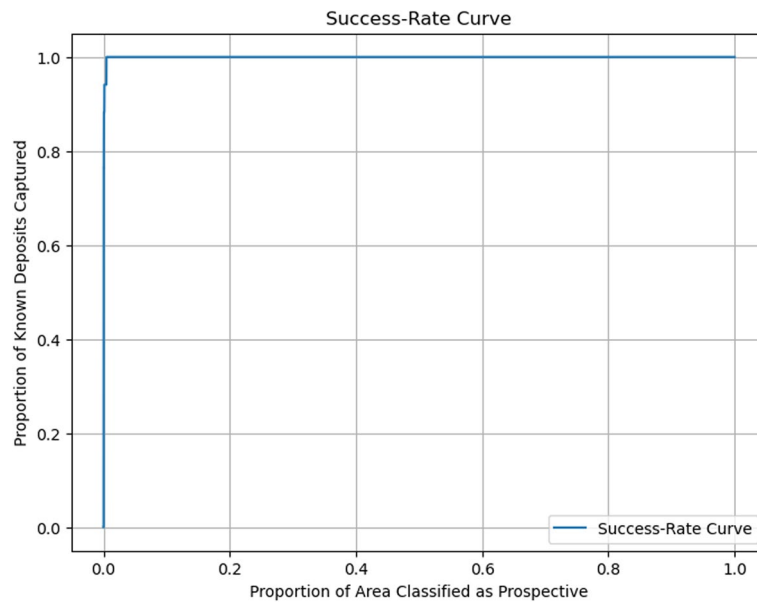


Fig. 11 Success-rate curve (SRC) for the radius workflow

6.1.4 Uncertainty quantification and feature interpretability

The uncertainty analysis of the CRL-radius classifier is based on the Euclidean distance between the decision radius and the point. To interpret this, these distances were transformed into estimated probabilities with a sigmoid function to enable us to obtain the Shannon entropy of each raster point (scaled between 0 and 1). The points on the boundary are those with a maximum entropy of 1.0 (highly uncertain) and the points deep inside or deep outside the radius are close to an entropy of 0.0 (high confidence).

Figure 12 demonstrates the spatial distribution of model uncertainty in a simple and clear manner. Dark purple areas indicate the mineralized and non-mineralized prediction with high model certainty. The thin Yellow corridors indicate the areas of highest




Fig. 12 Uncertainty map for CRL-Radius classifier

uncertainty of the model. The 17 identified mineralized locations (red stars) are generally located in high-confidence areas, with a mean probability of 0.9600 and a low mean entropy of 0.1215 (Table 3). The uncertainty is highest in deposit ID 2 ($H=0.9999$) because the deposit is slightly beyond the decision boundary but very close. This agrees with its location in Fig. 8.

Feature interpretability, assessed using SHAP values for the encoder's first embedding coordinate (Fig. 13), identifies magnetic variables (Mag_AS, Mag_TD) along with gravity and electromagnetic attributes (Pseu_Grv, EM_Inph, Em_Qd) as the most influential in shaping the learned embedding space. SHAP features ranking is based on the average absolute SHAP value. The SHAP values of Mag_AS are rank first due to the high absolute values of the SHAP values that are mostly within the range of -0.3 to -0.6 . PseuGrv (Rank 2) is larger in shap values which are around -0.2 to 0.8 but the mean value of the absolute ones is a little lesser. These SHAP values show what feature characteristics have an impact on the shape of the learned embedding. Features with bigger SHAP values decide how the data points are distributed and clustered around the deposit area. In simple terms, CRL-SHAP shows variables that determine geological similarity and lithological continuity. This explains why gravity, along with electromagnetic variables, is so important, as shown in Fig. 13. Such properties characterize the embedding space so that the mineralized belts appear to be clear and connected.

Table 3 Estimated probability of mineralization (P) and Shannon entropy (H) of the 17 known mineralized occurrences of the CRL-Radius classifier

ID	Probability (P)	Entropy (H)
1	0.9928	0.0613
2	0.4935	0.9999
3	0.9984	0.0174
4	0.9983	0.0185
5	0.9979	0.0220
6	0.9984	0.0167
7	0.9972	0.0279
8	0.9969	0.0299
9	0.9980	0.0208
10	0.9977	0.0234
11	0.9969	0.0299
12	0.9964	0.0345
13	0.9888	0.0885
14	0.9938	0.0540
15	0.9878	0.0949
16	0.9979	0.0217
17	0.8888	0.5035
<i>Summary Statistics (N = 17)</i>		
Mean	0.9600	0.1215
Std. Dev.	0.1230	0.2540
Min	0.4935	0.0167
Max	0.9984	0.9999

The smaller values of the entropy mean that the model is more confident in the prediction of the deposits




Fig. 13 SHAP summary for the radius workflow

Together, the uncertainty map guides where to focus further exploration, while the SHAP plot explains why predictions are made, enhancing both the operational utility and transparency of the CRL-Radius approach.

6.1.5 Prospectivity mapping

The final prospectivity map as shown in Fig. 14 categorizes each raster point according to its location in relation to the global prospectivity circle. The points inside the circle (orange) show possible mineralized zones, while the points outside (blue) show areas that are not mineralized. The red stars on the annotated deposits mostly line up with the orange zones. This shows that the classifier does a good job of finding known mineralized belts.

The orange points show connected belts and clusters that line up with known mineralized deposits. They also show new areas that have not been explored enough and need more exploration. This pattern is quite similar to what the CRL-Radius classifier is meant to do: make concentrated, corridor-like areas where the embedded data is most like known deposits.




Fig. 14 Prospectivity map derived from the global-circle decision rule in the 2D contrastive embedding space

6.2 Supervised multilayer perceptron

The supervised MLP workflow depicted in Fig. 6 utilized the 17-fold nested stratified K-fold cross-validation (SKCV) methodology to evaluate prediction performance under significant class imbalance. For testing, each outer fold kept one annotated deposit and its corresponding background raster points. The rest of the data went through 16-fold inner SKCV to fine-tune the hyperparameters. To address the class imbalance, a hybrid resampling method was used. This includes SMOTE oversampling of minority (mineralized) instances and random undersampling of majority (background) instances. This assured that the model trained from a balanced collection of classes.

6.2.1 Effectiveness of data balancing

The dataset included 17 annotated positive (mineralized) points and 1,843,564 raster points. We used a 17-fold stratified outer cross-validation (SKCV) method. Each outer fold had 16 annotated positives and 1,735,119 raster points in the training set and 1 annotated and 108,445 raster points in the test set. We used an inner 16-fold SKCV for hyperparameter tuning within each outer training set. This gave us training folds with 15 annotated positives and 1,626,659 raster points, and validation folds with 1 annotated positive and 108,445 raster points.

To explore a range of M (over-sampling) and N (under-sampling) combinations, we generated 20 logarithmically spaced values between 15 and 1,626,659. This produced the sequence in Fig. 15, ensuring denser sampling for small values—where changes have the greatest relative effect—and coarser spacing for large values—where performance is expected to plateau. These values formed the grid over which we evaluated model performance.

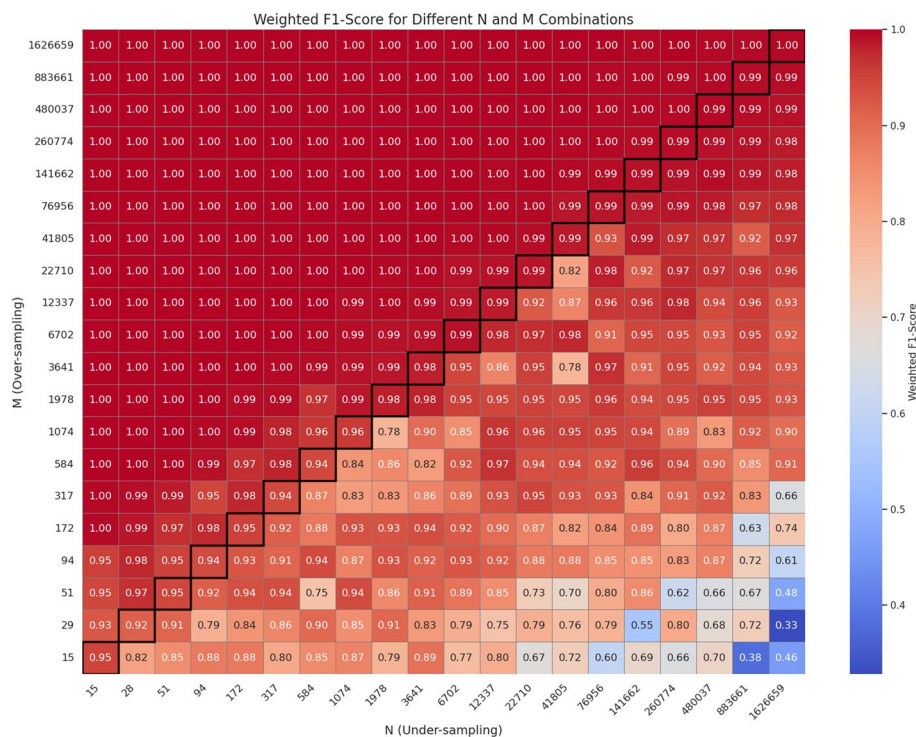


Fig. 15 Weighted F1-scores for different oversampling (M) and undersampling (N) combinations, showing underfitting in the bottom-right, overfitting in the top-left, and optimal balanced performance along the diagonal

The results, shown in Fig. 15, reveal three distinct regions. The area around bottom-right corner (low M , high N) suffers from underfitting, as the minority class is poorly represented; weighted F1-scores drop as low as 0.33 at ($M=29$, $N=1626659$) and 0.38 at ($M=15$, $N=883661$). The top-left corner (high M , low N) indicates overfitting: while scores are near-perfect, they are driven by an excessive number of synthetic positives combined with a small set of negatives, increasing the risk of memorizing artifacts. The diagonal where $M \approx N$, which means balanced sampling, seems to have the most stable and generalizable performance. Here, weighted F1-scores consistently range from 0.92 to 1.00 without relying on extreme sampling.

Figure 15 shows that the model's performance improves and stabilizes when the amounts of oversampled positive samples (M) and undersampled negative samples (N) are kept about balanced. The structure, such as $M = N = 6702$, is a large oversampling of the 15 mineralized deposits, but that does not imply that the geological evidence is artificially overestimated. This approach involves the resampling method in the geophysical features and is evaluated under the context of nested cross-validation. The findings of the results, in these settings, indicate that the model generalizes to previously unseen folds (data) rather than just memorizing the mineralized deposits. In fact, the mineralized locations are very sparse in the 13-dimensional feature space. Dense resampling gives a smooth approximation of the actual geophysical pattern of the mineral deposits, and thus the MLP is able to learn the pattern smoothly.

This behavior is similar to the findings of Prado et al., who showed that the balanced training strategy results in stable and efficient prospectivity models [47]. Although Prado et al. investigated small resampling ratios, this study applies their balanced-learning concept to larger resampling scales using an MLP model. This approach is also justified by the study by Farahnakian et al., who have confirmed that the hybrid resampling of oversampling and undersampling is effective in dealing with the issue of class imbalance in Finnish mineral exploration data [11]. The success-rate curve further confirms these results by showing that the predicted known deposits are highly concentrated in small regions.

Only one configuration, (1626659, 1626659), achieved a perfect F1-score of 1.00 in this grid search. While this setting maximizes performance in terms of the metric, it also requires the largest possible synthetic oversampling and undersampling, resulting in very high computational cost and potential over-smoothing of decision boundaries. Several slightly smaller diagonal configurations—such as (480037, 480037), (141662, 141662), and (76956, 76956) produced near perfect scores (0.98–0.99) while being far more efficient and less prone to synthetic-data artifacts. For most applications, these smaller balanced settings in the good-fit zone are the practical choice, reserving (1626659, 1626659) for final models only when computational resources are abundant and marginal performance gains are critical.

6.2.2 Predictive efficiency and confusion matrix

The performance of the MLP classifier was evaluated using a confusion matrix as shown in Fig. 16. The model can distinguish between annotated (mineralized) points and non-mineralized background points. From the confusion matrix, the results can be seen as $TN = 1,820,572$ (98.75%), $FP = 22,992$ (1.25%), $FN = 0$ (0.00%), and $TP = 17$ (100.00%). The model employs the following performance metrics:

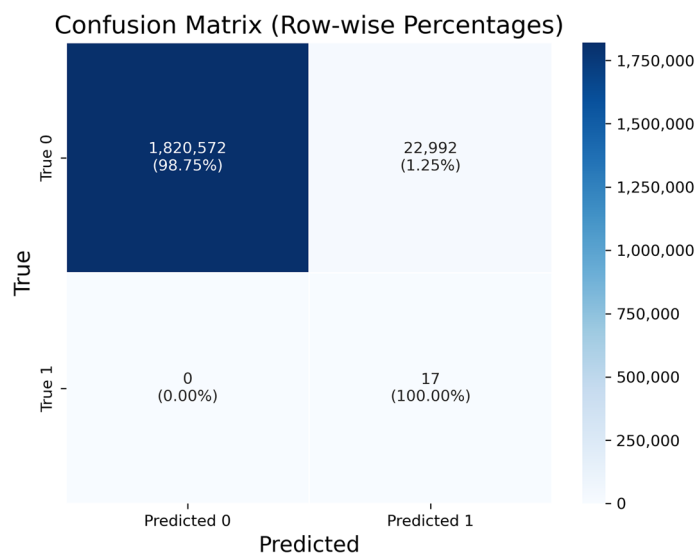


Fig. 16 MLP confusion matrix

- *Recall(100.00%)*: All known deposits are identified. In MPM, this metric was significant because missing a deposit (FN) could mean overlooking a valuable target.
- *Specificity (98.75%)*: Almost all of the background points are excluded from the mineralized class, providing an effective control against large-scale overprediction.
- *Precision (0.07%)*: A small ratio of predicted mineralized points are actual known minerals. At early stage exploration precision will inevitably be quite low as the objective is to include all true deposits (at the cost of additional false positives).
- *F1-score (0.14%)*: This value is low due to low precision value, which is acceptable for general targeting.
- *Weighted-F1 (99.51%)*: High because the dominant back-ground class is almost perfectly classified.
- *G-mean (.9937)*: Good balance between correctly detected deposits and avoiding false alarms.

In MPM, the objective is to attain perfect recall by allowing certain false positives to ensure that no deposit is missed. A total of 22,992 background points were identified as probable mineral points, thereby narrowing the search to very promising regions.

6.2.3 Ranking power and threshold

The performance of the MLP model as a predictor was re-evaluated rigorously with three complementary metrics to overcome possible evaluation bias: the Receiver Operating Characteristic (ROC) curve, the revised Success-Rate Curve (SRC), and the Prediction-Area (P-A) plot.

The AUC of the ROC curve (Fig. 17) is 0.9986 indicating that the mineralized and the background samples are very discriminated.

To identify the practical utility of the model to exploration targeting we used Prediction-Area (P-A) plot proposed by the recent reviews of imbalanced data in mineral prospectivity [46]. First, the continuous predictive probabilities of the model were reclassified with the aid of the data-driven Concentration-Area (C-A) fractal model in order to form the plot. The Prediction-Area plot (Fig. 18) compares the rate of prediction with

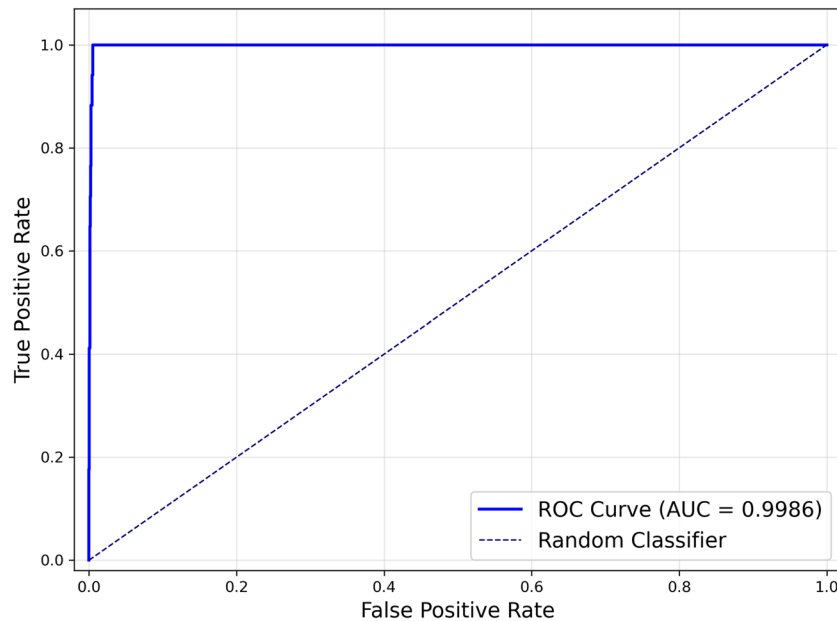


Fig. 17 MLP ROC-AUC curve

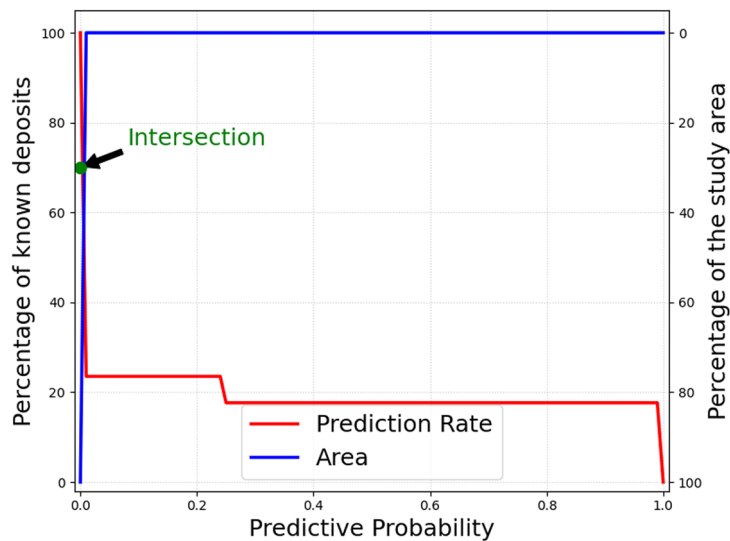


Fig. 18 MLP prediction–area curve

the area of study that is being occupied at different levels of probability. The model is able to predict, at the intersection of the two curves, 70% of the known deposits with only 30% of the total area. The intersection has a normalized density of 2.33 and the objective weight of 0.85. Since prospectivity models with normalized density above 1 is a good predictor of a mineral site [48], the high-prospectivity areas indicated by the MLP model give good targets to be used in further exploration.

To further validate the model, the Success-Rate Curve (Fig. 19) shows that the model is capable of ranking the most likely locations of minerals in the correct order [45]. The curve is steep and takes a huge percentage of known incidences in the highest percentile of the study area. This confirms mineralization is strongly correlated with the high probability scores which effectively narrows the search space during follow-up exploration.




Fig. 19 MLP cumulative success rate curve

These findings are consistent with those of Farahnakian et al. (2024), who found that MLP worked more effectively as an imbalanced data classifier [11].

6.2.4 Uncertainty quantification and feature interpretability

We applied an Shannon entropy formula based on information theory to predicate uncertainty of the MLP classifier. The uncertainty of model is high (i.e.1) when its predicted probability is 0.5 and low (0) when the prediction is 0 or 1.

The uncertainty map based on entropy (Fig. 20) shows that MLP model offers high confidence in most of the study area. To measure this observation quantitatively, we computed the Shannon entropy of the predicted probabilities of the 17 known mineralized sites (Table 4).

The quantitative analysis presented in Table 4 proves the presence of the known deposits in low uncertainty (high confidence) zones. The model indicates high mineralization probabilities (mean $P \approx 0.97$) of the known occurrences giving a very low mean entropy of 0.099. This means that the model is assured that it identified the following specific targets.

Spatially, the deposits themselves are often located in high-confidence areas, although often enclosed between thinly correlated "corridors" of increased entropy. These corridors are used as the point of transition: presumably geological contacts or structural boundaries, where the decision surface of the model passes between mineralized and non-mineralized classes. This review is also consistent with the concepts of effective exploration targeting presented by Yousefi et al. (2024), in which robust targets are characterized by a high prospectivity and low uncertainty.

Feature interpretability, assessed using SHAP values for MLP (Fig. 21), identifies magnetic variables (Mag_AS, Mag_TD) along with gravity and electromagnetic attributes (Pseu_Grv, EM_Inph, Em_Qd) as the most influential in shaping the learned embedding space. In contrast to CRL-SHAP, these SHAP values show how individual features affect the prediction of mineralization in the supervised MLP model. The SHAP analysis of the MLP model indicates that Mag_AS is the most influential feature, with SHAP values

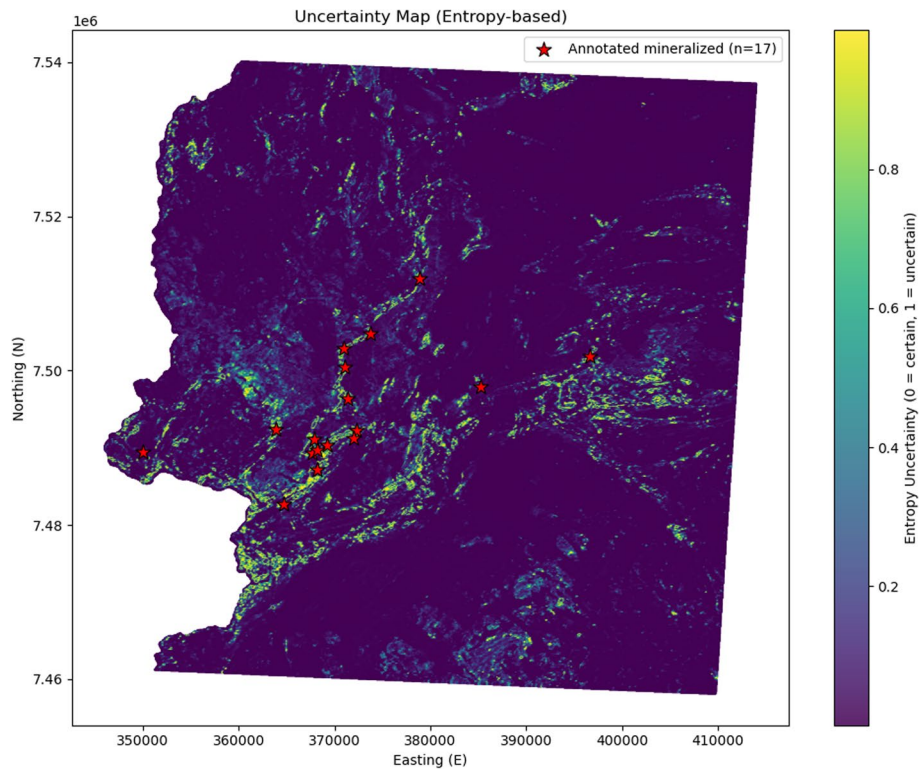


Fig. 20 Spatial distribution of uncertainty of epistemic model using Shannon entropy. Large values (yellow) show transition zones at which predictions are sensitive to minor changes in geophysical features. Small values (purple) denote high model confidence. It is important to note that the locations of known mineralized points (red stars) overlap with regions of low uncertainty, and are bounded by transitional corridors

reaching approximately 0.9–1.0, followed by *Em_Qd* (Rank 2, SHAP > 0.6) and *Rd_TC* (Rank 3, SHAP \approx 0.30). In contrast, *Pseu_Grv* is ranked last (Rank 13), with most SHAP values clustered around zero, indicating a negligible impact on the model output. Here, the model focuses on the variables that separate mineralized areas from the large background under extreme class imbalance. Because of this MLP model, magnetic derivative and other high-contrast signals look more important, as shown in Fig. 19, and help to decide the final mineralization prediction. Here, these SHAP values reflect influence on classification directly.

The uncertainty map and SHAP plot combined can give direction to subsequent exploration, and the rationale behind prediction can be obtained, as well as enhance the operational utility and interpretability of the model. The differences between Figs. 13 and 21 indicate that the CRL-radius and the MLP play different but complementary roles. CRL-SHAP illustrates the effects of the features on the embedding into an valuable geological pattern and the MLP-SHAP explains the same features being used to make the distinction between the mineralized regions and the background. The consistency in the characteristics of magnetic, pseudogravity, and electromagnetic attributes suggests that these are real indicators of IOCG mineralization.

6.2.5 Prospectivity mapping

This plot, Fig. 22, illustrates squeezed, high-probability areas (warm colors) contrasted with a low-probability background (dark blue). These high-prospectivity regions show up as distinct curvilinear belts that are spatially continuous.

Table 4 Estimated probability of mineralization (P) and Shannon entropy (H) of the 17 known mineralized occurrences

ID	Probability (P)	Entropy (H)
1	0.9975	0.0175
2	0.9869	0.0696
3	0.8707	0.3850
4	0.9995	0.0041
5	0.9818	0.0909
6	0.9879	0.0653
7	0.9493	0.2005
8	0.9972	0.0194
9	0.9994	0.0048
10	0.9966	0.0229
11	0.9657	0.1494
12	0.9977	0.0162
13	0.9888	0.0615
14	0.9866	0.0711
15	0.9599	0.1683
16	0.9987	0.0098
17	0.9015	0.3218
<i>Summary Statistics (N = 17)</i>		
Mean	0.9744	0.0987
Std. Dev.	0.0384	0.1131
Min	0.8707	0.0041
Max	0.9995	0.3850

Lower entropy values indicate higher model confidence in the predicted deposits




Fig. 21 SHAP feature importance plot

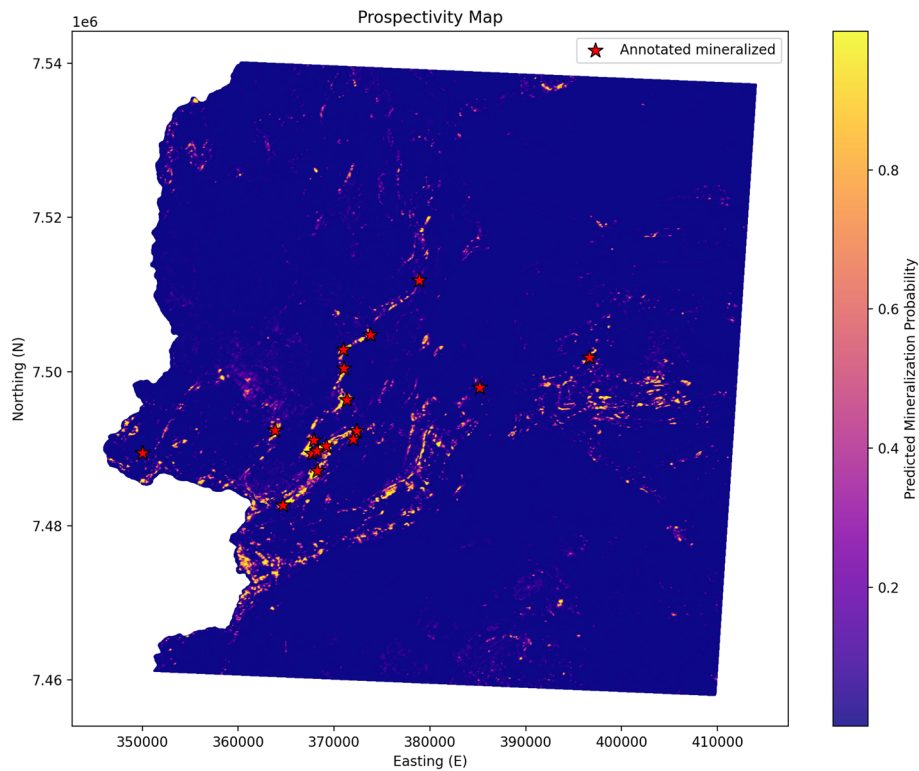


Fig. 22 This map shows the spatial distribution of predicted mineralization probabilities, with warmer colors indicating higher likelihoods. Annotated mineralized points (red stars) are overlaid to validate the model's predictions

From the figure, we can see that annotated points (mineralized) are located within or near the most probable zones. This validates that there is a robust correlation between model predictions and established deposits. By targeting these clearly identified areas, the map highly narrows the exploratory search space for MPM.

7 Conclusion

In this research, two workflows of mineral prospectivity mapping (MPM) for extreme class imbalances, i.e., (1) CRL-Radius; and (2) supervised Multilayer Perceptron (MLP) were evaluated in an integrative framework.

The supervised MLP exhibited superior overall performance, achieving recall of 100%, an AUC of 0.9986, and a G-mean of 0.9937. The success-rate curve included all deposits that constituted less than 1% of the area, significantly reducing the exploration footprint. Along the diagonal line $M \approx N$, the weighted F1 score constantly ranged from 0.92 to 1.00, demonstrating no instances of underfitting or overfitting.

The CRL-Radius also has a good performance (recall 94.12%, AUC 0.9996) and obtained spatially continuous belt-like corridors due to its geometric rule.

According to SHAP, the key features were often identified as the magnetic and pseudo-gravity signatures. The enhanced MLP uncertainty map is impressive with all the 17 deposits now being detected with high confidence (low uncertainty) with narrow transitional corridors. These corridors align with geological boundaries, hence model uncertainty aligns with an exploratory objective.

In conclusion, MLP did well with the class balance $M \approx N$ strategy, nested cross-validation, uncertainty, and explainability. Although the CRL-MLP approach did not exceed

the supervised baseline, the study lays a foundation for future hybrid representation learning. The joint strength of success-rate efficiency, interpretable feature importance, and uncertainty mapping underpins trustworthy, cost-effective exploration.

Acknowledgements

The Horizon Europe research and innovation programme, Grant Agreement number 101057357, is funding the compilation of the work, EIS - Exploration Information System (<https://eis-he.eu>).

Author contributions

Dipak Kumar Nidhi led the conceptualization of the study, designed the computational framework, performed the experiments, and prepared the initial manuscript draft. Sudhir Kumar Mohapatra contributed to the methodological design, reviewed the modeling workflow, and assisted in refining the manuscript. Paavo Nevalainen supported data processing, validation of geophysical interpretations, and provided critical feedback on the results. Jukka Heikkonen supervised the research, guided the overall study design, and contributed to the interpretation of findings. Rajeev Kanth co-supervised the work, contributed to methodological discussion, and supported revisions of the manuscript. All authors reviewed and approved the final version of the manuscript.

Funding

Open Access funding provided by University of Turku (including Turku University Central Hospital).

Data availability

Data available on reasonable request.

Declarations

Ethics approval and consent to participate

Not applicable. This research exclusively utilizes archival geophysical datasets provided by the Geological Survey of Finland (GTK) and documented mineral occurrence locations. The study does not involve human participants, animals, or identifiable personal data.

Consent for publication

Not applicable.

Competing of interests

The authors declare no conflict of interest.

Received: 14 November 2025 / Accepted: 18 May 2026

Published online: 21 May 2026

References

1. Zuo R. Geodata science-based mineral prospectivity mapping: a review. *Nat Resour Res.* 2020;29(6):3415–24.
2. Josso P, Hall A, Williams C, Le Bas T, Lusty P, Murton B. Application of random-forest machine learning algorithm for mineral predictive mapping of Fe-Mn crusts in the world ocean. *Ore Geol Rev.* 2023;162:105671.
3. Zeghouane H, Allek K, Kesraoui M. GIS-based weights of evidence modeling applied to mineral prospectivity mapping of Sn-W and rare metals in laouni area, central Hoggar, Algeria. *Arab J Geosci.* 2016;9(5):373.
4. Dan WA. Transforming the business of gold exploration: adapting to deeper exploration. *SEG Newslett.* 2018;112:1–14.
5. Yousefi M, Nykänen V, Harris J, Hronsky JM, Kreuzer OP, Bertrand G, et al. 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 Geol Rev.* 2024;172:106214.
6. Saremi M, Hoseinzade Z, Yousefi M. A deep embedded clustering algorithm in conjunction with an ensemble technique for mineral prospectivity mapping. *Sci Rep.* 2025;15(1):38086.
7. Nidhi DK, Chaudhary J, Heikkonen J, Kanth R. Enhancing mineral prospectivity mapping with contrastive representation learning and dimensionality reduction techniques. In: International conference on hybrid intelligence: theories and applications, Springer, 2026;13–23.
8. Wang Z, Zuo R. Intelligent lithological mapping: challenges and future prospective. *Nat Resour Res.* 2025;1–34.
9. Zuo R, Carranza EJM. Machine learning-based mapping for mineral exploration. *Math Geosci.* 2023;55(7):891–5.
10. Nidhi DK, Seppä I, Farahnakian F, Zelioli L, Heikkonen J, Kanth R. Enhancing mineral prospects mapping with machine learning: Addressing imbalanced geophysical datasets and data visualization approaches. In: 2023 34th Conference of open innovations association (FRUCT), IEEE, 2023;125–135.
11. Farahnakian F, Sheikh J, Zelioli L, Nidhi D, Seppä I, Ilo R, et al. Addressing imbalanced data for machine learning based mineral prospectivity mapping. *Ore Geol Rev.* 2024;174:106270.
12. Maeda EE, Haapasaaari P, Helle I, Lehtikoinen A, Voinov A, Kuikka S. Black boxes and the role of modeling in environmental policy making. *Front Environ Sci.* 2021;9:629336.
13. Li H, Usman N, Coulbaly MH, Phiri R, Tang X. Does the resources curse hypothesis exist in China? What is the dynamic role of fiscal decentralization, economic policy uncertainty, and technology innovation for sustainable financial development? *Resour Policy.* 2022;79:103002.
14. Koldasbayeva D, Tregubova P, Gasanov M, Zaytsev A, Petrovskaya A, Burnaev E. Challenges in data-driven geospatial modeling for environmental research and practice. *Nat Commun.* 2024;15(1):10700.

15. Al-Najjar HA, Pradhan B, Sarkar R, Beydoun G, Alamri A. A new integrated approach for landslide data balancing and spatial prediction based on generative adversarial networks (gan). *Remote Sensing*. 2021;13(19):4011.
16. Chen X, He K. Exploring simple siamese representation learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021;15750–15758.
17. Grill J-B, Strub F, Althé F, Tallec C, Richemond P, Buchatskaya E, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Adv Neural Inf Process Syst*. 2020;33:21271–84.
18. Krstajic D, Buturovic LJ, Leahy DE, Thomas S. Cross-validation pitfalls when selecting and assessing regression and classification models. *J Cheminform*. 2014;6(1):10.
19. Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. Smote: synthetic minority over-sampling technique. *J Artif Intell Res*. 2002;16:321–57.
20. Liu B, Tsoumakas G. Dealing with class imbalance in classifier chains via random undersampling. *Knowl-Based Syst*. 2020;192:105292.
21. Zuo S, Zhu M, Xu Z, Oláh J, Lakner Z. The dynamic impact of natural resource rents, financial development, and technological innovations on environmental quality: empirical evidence from bri economies. *Int J Environ Res Public Health*. 2021;19(1):130.
22. Coughlan M, Keese A, Pinto V, Mukundan R, Marchezi JP, Adewuyi M, et al. Using machine learning explainability techniques to examine drivers of ground magnetic field localization. *Space Weather*. 2025;23(8):e2025SW004391.
23. Kendall A, Gal Y. What uncertainties do we need in bayesian deep learning for computer vision? In: *Advances in neural information processing systems*. 2017;30.
24. Porwal A, Carranza EJM. Introduction to the special issue: Gis-based mineral potential modelling and geological data analyses for mineral exploration. 2015.
25. Sabbaghi H, Tabatabaei SH, Fathianpour N. Optimization of multi-element geochemical anomaly recognition in the takht-e soleyman area of northwestern iran using swarm-intelligence support vector machine. *Front Earth Sci*. 2025;13:1352912.
26. GrobidStyleApplied="Changed". Sabbaghi H. Recognition of multi-element geochemical anomalies related to pb-zn mineralization applying upgraded support vector machine in the varchek district of Iran. *Modeling Earth Systems and Environment*. 2024;1–14.
27. Sabbaghi H, Fathianpour N. Hybrid harris hawks-optimized random forest model for detecting multi-element geochemical anomalies related to mineralization. *Sci Rep*. 2025;15(1):23662.
28. Saremi M, Hezarkhani A, Mirzabozorg SAAS, DehghanNiri R, Shirazy A, Shirazi A. Unsupervised anomaly detection for mineral prospectivity mapping using isolation forest and extended isolation forest algorithms. *Minerals*. 2025;15(4):411.
29. Nidhi DK, Nevalainen P, Chaudhary J, Heikkonen J, Kanth R. Enhancing mineral prospectivity mapping in imbalanced data environments using geophysical feature similarity and bayesian kernel density estimation. In: *2025 International conference on artificial intelligence, computer, data sciences and applications (ACDSA)*, IEEE, 2025;1–8. <https://doi.org/10.1109/ACDSA65407.2025.11166646>
30. Gajjar Y, Arora N, Sahoo NR. Self-organizing maps: concept, architecture, and use cases in engineering and finance. In: *Deep learning in engineering, energy and finance*, CRC Press, 2024;211–249.
31. Yang F, Zuo R, Kreuzer OP. Interpretable regional-scale geological mapping using a contrastive graph attention network for multimodal data fusion and recognition of controlling factors. *Eng Geol*. 2025;108508.
32. Sun X, Zhang Y, Dong Y, Du B. Contrastive self-supervised learning-based background reconstruction for hyperspectral anomaly detection. *IEEE Trans Geosci Remote Sens*. 2025;63:1–12. <https://doi.org/10.1109/TGRS.2025.3534185>.
33. Zhang Z, Wang G, Carranza EJM, Liu C, Li J, Fu C, et al. An integrated machine learning framework with uncertainty quantification for three-dimensional lithological modeling from multi-source geophysical data and drilling data. *Eng Geol*. 2023;324:107255.
34. Yang F, Zuo R, Xiong Y, Xu Y, Nie J, Zhang G. Dual-branch convolutional neural network and its post hoc interpretability for mapping mineral prospectivity. *Math Geosci*. 2024;56(7):1487–515.
35. Yu Z, Liu B, Xie M, Wu Y, Kong Y, Li C, et al. 3d mineral prospectivity mapping of Zaozigou gold deposit, west Qinling, China: deep learning-based mineral prediction. *Minerals*. 2022;12(11):1382.
36. Nidhi DK, Srivastav K, Joshi B, Joshi SR. Convolutional neural networks for the assessment of fetal echocardiography. In: *Proceedings of the 10th IOE graduate conference*, 2021;10:249–254.
37. Nidhi DK, Srivastav K, Heikkonen J, Kanth R. Evaluation of spatiotemporal fetal cardiac imaging using deep learning techniques. In: *Proceedings of the international health informatics conference: IHIC 2022*, Springer, 2023;285–298.
38. Sabbaghi H, Tabatabaei SH. Regimentation of geochemical indicator elements employing convolutional deep learning algorithm. *Front Environ Sci*. 2023;11:1076302.
39. Sabbaghi H, Tabatabaei SH, Fathianpour N. Multi-element geochemical anomaly recognition applying geologically-constrained convolutional deep learning algorithm with butterworth filtering of frequency domain information: Of frequency domain information. *Sci Rep*. 2025;15(1):43147.
40. McCuaig TC, Hronsky JM. The mineral system concept: the key to exploration targeting. 2014.
41. Dentith M, Mudge ST. *Geophysics for the mineral exploration geoscientist*. Cambridge University Press; 2014.
42. Skirrow RG. Iron oxide copper-gold (iocg) deposits-a review (part 1): Settings, mineralogy, ore geochemistry and classification. *Ore Geol Rev*. 2022;140:104569.
43. Wilford J. *Airborne gamma-ray spectrometry*, cooperative research centre for landscape environments and mineral exploration, commonwealth scientific and industrial research organization, Bertley, WA. Australia Open File Rep. 2002;144:46–52.
44. Carranza E, Hale M, Faassen C. Selection of coherent deposit-type locations and their application in data-driven mineral prospectivity mapping. *Ore Geol Rev*. 2008;33(3–4):536–58.
45. Agterberg FP, Bonham-Carter GF. Measuring the performance of mineral-potential maps. *Nat Resour Res*. 2005;14(1):1–17.
46. Yousefi M, Carranza EJM. Prediction-area (p-a) plot and c-a fractal analysis to classify and evaluate evidential maps for mineral prospectivity modeling. *Comput Geosci*. 2015;79:69–81.
47. Prado EMG, de Souza Filho CR, Carranza EJM, Motta JG. Modeling of cu-au prospectivity in the Carajás mineral province (Brazil) through machine learning: dealing with imbalanced training data. *Ore Geol Rev*. 2020;124:103611.

48. Mihalasky MJ, Bonham-Carter GF. Lithodiversity and its spatial association with metallic mineral sites, great basin of Nevada. *Nat Resour Res.* 2001;10(3):209–26.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.