

Geological Constraint-Driven Directional Relational Graph Networks with Contrastive Pretraining for Explainable Porphyry-Gold Prospectivity Mapping Under Scarce Labels

Hakeem B. Ajileye^{1,2*}, Olughu Daniel Sunday³, Olugbenle Olatomide Alfred³, Olagundoye O. Omosalewa³, Dorothy Mokeira Kerage⁴, Agbailu Adejoke Adewumi⁵

¹Geomathematics Key Laboratory of Sichuan Province (Chengdu University of Technology), Chengdu 610059, China

²School of Mathematical Sciences, Chengdu University of Technology, Chengdu 610000, China

³College of Nuclear Technology and Automation Engineering, Chengdu University of Technology, Sichuan P.R., 610059, China

⁴College of Earth Sciences, Chengdu University of Technology, Sichuan P.R., Chengdu 610059, China

⁵College of Ecology and Environment, Chengdu University of Technology, Sichuan P.R., Chengdu 610059, China

DOI: <https://doi.org/10.36347/sjet.2026.v14i10.001>

| Received: 03.08.2026 | Accepted: 25.09.2026 | Published: 01.10.2026

*Corresponding author: Hakeem B. Ajileye

Geomathematics Key Laboratory of Sichuan Province (Chengdu University of Technology), Chengdu 610059, China

Abstract

Review Article

Mineral prospectivity mapping is difficult in regions where confirmed deposits are few and available geological data vary widely in type and quality. Standard machine learning and convolutional methods have improved prediction accuracy, but they ignore the directional structure of geological space, do not enforce geological logic during training, and rarely explain which factors drove a prediction. We present GeoEDL, a directional relational graph network built on geological constraints, combined with contrastive pretraining, for porphyry-gold prospectivity mapping under scarce labels. The framework provides interpretability at four levels: data, model architecture, geology, and results. Tiles are first classified as favorable, ambiguous, or barren using Mahalanobis distance and fracture density, without relying on mineralization labels, and these classes generate contrastive pairs for pretraining. A nine-edge-type graph encoder is then pretrained with geological contrastive loss and fine-tuned using evidential deep learning under structural and geochemical consistency penalties. On 9,324 tiles with 94 positive labels from Tuwu-Yandong, the model reached ROC-AUC 0.9418 and sensitivity 1.0000, with geochemical anomaly identified as the leading driver (51.6%) and epistemic uncertainty mapping flagging priority drilling targets.

Keywords: Mineral Prospectivity Mapping, Graph Neural Networks; Explainable Artificial Intelligence, Contrastive Learning, Uncertainty Quantification.

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

1. INTRODUCTION

Mineral prospectivity mapping is a widely adopted methodology in mineral exploration that aims to delineate target areas with the highest likelihood of hosting undiscovered mineral deposits by integrating and analyzing multi-source geological, geochemical, geophysical, and remote sensing data [1-4]. With continued advancements in computational power and the growing availability of geological prospecting big data, artificial intelligence algorithms have emerged as powerful tools for mining nonlinear mineralization patterns from heterogeneous datasets, demonstrating strong predictive performance across a wide range of deposit types and geological settings [5]. Early AI-driven MPM relied on supervised machine learning methods

such as weights of evidence, logistic regression, support vector machines, and random forests, which established the foundation for data-driven mineral targeting by learning statistical associations between geological features and known deposit locations [6,7]. These methods demonstrated consistent utility in integrating multi-source datasets and ranking spatial prediction units by mineralization potential. However, these approaches do not capture the complex nonlinear interactions and hierarchical feature representations embedded in high-dimensional geological datasets, limiting their capacity to fully exploit the richness of modern multi-source prospecting data.

The limitations of shallow supervised methods prompted the adoption of deep learning architectures in

Citation: Ajileye, H. B., Sunday, O. D., Alfred, O. O., Omosalewa, O. O., Kerage, D. M., & Adewumi, A. A. (2026). Geological Constraint-Driven Directional Relational Graph Networks with Contrastive Pretraining for Explainable Porphyry-Gold Prospectivity Mapping Under Scarce Labels. *Sch J Eng Tech*, 14(10), 575-598.
<https://doi.org/10.36347/sjet.2026.v14i10.001>

MPM, with convolutional neural networks emerging as the dominant approach due to their ability to automatically extract hierarchical spatial features from gridded geological rasters [8,9]. CNN-based frameworks have demonstrated substantial gains in predictive performance by treating MPM as an image classification problem, learning localized spatial patterns across geochemical, structural, and geophysical input bands without manual feature engineering [10,11]. Recurrent and hybrid architectures have further extended these capabilities to sequential and multi-modal geological data [12]. Despite these advances, however, these approaches do not account for the explicit relational structure of geological space, where mineralization potential is governed not only by local pixel values but by directional connectivity between spatial prediction units, fault orientations, and the anisotropic propagation of hydrothermal fluids along structurally controlled pathways, relationships that grid-based convolutional operations are fundamentally unable to encode.

Graph neural networks represent a particularly suitable architecture for MPM because the graph structure, composed of nodes representing spatial prediction units and edges encoding directional geological relationships, enables flexible characterization of spatial anisotropy and the complex coupling relationships between geological entities that cannot be captured by pixel-based or image-based methods [13,14]. Relational GNNs with direction-aware message passing are especially well-suited to encoding the structural anisotropy of fault and fracture systems, where mineralization potential varies with both proximity to structures and the orientation of structural connections [15]. Moreover, GNNs are amenable to self-supervised pretraining on graph structure, enabling large amounts of unlabeled geological data to be leveraged for representation learning before fine-tuning on the limited labeled set, directly addressing the scarce label problem [16,17]. However, these approaches do not address the interpretability of the resulting predictions: GNN-based MPM models remain largely opaque, producing prospectivity scores that cannot be decomposed into contributions from distinct ore-controlling factors, and offering no mechanism for quantifying the uncertainty of predictions in a geologically meaningful way [18,19].

Explainable artificial intelligence has emerged as a necessary direction for addressing the interpretability limitations of AI-driven MPM [20-22]. In the context of MPM, interpretability can be understood at multiple levels: the quality and geological relevance of input data, the degree to which geological knowledge governs model behavior during training, the ability to explain individual predictions in geological terms, and the interpretability of the final output maps for exploration decision-making. Existing XAI approaches in MPM have made progress in embedding geological knowledge into loss functions as soft constraints [23,24], applying attribution algorithms such as SHAP to rank the

importance of input variables [25], and using visualization techniques to understand feature extraction processes [26]. However, these approaches do not provide a unified framework that systematically embeds geological knowledge from data preparation through model design to output explanation, and the specific challenge of spatially decomposing predictions into contributions from distinct ore-controlling factors at the regional map scale, a distinction directly actionable for designing follow-up exploration campaigns, remains largely unaddressed.

A deeper and largely unaddressed challenge arises when these limitations compound simultaneously: when a study region offers only a handful of known deposit locations as training labels, when the geological space is best characterized by directional relational structure rather than regular grids, and when exploration decisions require not only a prospectivity score but a geologically interpretable, uncertainty-aware explanation of what drives each prediction. In such settings, the absence of geological constraints during model training means that even architecturally expressive models such as GNNs may produce predictions that are statistically plausible but geologically inconsistent, assigning high prospectivity to tiles that violate known ore-controlling criteria such as structural proximity and geochemical anomaly. Label-free geological supervision, the use of ore-controlling geological rules to organize representation learning before any mineralization labels are introduced, offers a principled route to reducing overfitting under these conditions, but has not been formalized within a GNN-based MPM framework. Furthermore, the quantification of both aleatoric and epistemic uncertainty within a geologically constrained framework and the spatial decomposition of predictions into structural, geochemical, and contextual contributions across the full study area remain absent from existing approaches. Furthermore, the quantification of both aleatoric and epistemic uncertainty within a geologically constrained framework and the spatial decomposition of predictions into structural, geochemical, and contextual contributions across the full study area remain absent from existing approaches. No existing framework simultaneously addresses the scarce label problem, directional relational graph structure, geological constraint enforcement, multi-level XAI, and principled uncertainty quantification within a single unified MPM pipeline, a gap this study directly addresses through the proposed GeoEDL framework.

The Tuwu-Yandong metallogenic belt in eastern Tianshan, Xinjiang, China, hosting porphyry copper and hydrothermal gold deposits with only 94 known deposit locations across 9,324 spatial prediction units, provides a demanding and geologically meaningful testbed for evaluating such a framework under real-world scarce-label conditions. The specific contributions are as follows:

- ❖ This study proposes GeoEDL, a four-level explainable AI framework for MPM that systematically embeds geological knowledge across data preparation, model training, prediction explanation, and uncertainty quantification within a single unified pipeline.
- ❖ A geology-guided two-criterion tile classification scheme and GeoSupCon, a geologically supervised contrastive pretraining approach for directional relational graph networks, are developed to collectively address overfitting under scarce mineralization labels.
- ❖ A two-term geological constraint loss for geologically consistent fine-tuning and geological factor decomposition, producing regional attribution maps quantifying structural, geochemical, and contextual contributions to prospectivity predictions, is introduced.
- ❖ The evidential deep learning output head simultaneously generates prospectivity and uncertainty maps whose combination constitutes a principled and geologically grounded exploration decision map.

The remainder of this paper is organized as follows. Section 2 describes the study area geology and the dataset. Section 3 presents the methodology of the proposed four-level XAI framework, covering the input feature multicollinearity analysis, two-criterion tile

classification, and contrastive pair construction at the data interpretability level, the GeoSupCon pretraining, geologically constrained GeoEDL fine-tuning, and model architecture at the model interpretability and architecture level, and the GNNExplainer and geological factor decomposition components at the geological interpretability level. Section 4 presents the full framework results and discussion, comprising model performance evaluation, ablation study, geological interpretability findings, and prospectivity and uncertainty map analysis with exploration decision implications. Section 5 concludes the paper.

2. Study area and Dataset

2.1 Geological Background

2.1.1 Regional Geology

The eastern Tianshan is a key segment of the Central Asian Orogenic Belt (CAOB), lying between the Tarim craton to the south and the Siberian craton to the north, and comprises a collage of Paleozoic tectonic terranes assembled through oceanic subduction, arc accretion, and continental collision [27,28]. From north to south, five major tectonic units are recognized: the Bogeda-Haerlik magmatic arc, the Dananhu-Tousuquan arc belt, the Kanggur-Huangshan ductile shear belt, the Aqishan-Yamansu arc, and the Middle Tianshan Terrane, separated by the regional Dacotan, Kanggur, Yamansu, and Aqikekuduke fault systems [29,30].




Figure 1: Geological map of the Tuwu-Yandong metallogenic belt, eastern Tianshan, Xinjiang, China

These fault systems served as long-lived crustal-scale fluid conduits that focused hydrothermal circulation and magmatic ascent during multiple metallogenic events. The Dananhu-Tousuquan arc belt hosts the Tuwu-Yandong metallogenic belt and contains several significant porphyry copper deposits, including Tuwu, Yandong, Fuxing, Linglong, and Chihu, formed during Devonian to early Carboniferous Island arc accretion driven by northward subduction of the Kanggurtag oceanic plate [31-34]. The adjacent Kanggur-Huangshan ductile shear belt hosts orogenic gold, porphyry molybdenum, and magmatic Cu-Ni sulfide deposits formed during late Carboniferous continental collision [35,36].

2.1.2 Stratigraphy and Lithology

The Tuwu-Yandong belt is located in the southern margin of the Dananhu-Tousuquan arc belt, approximately 80 km southwest of Hami City, bounded to the south by the Kanggurtag fault and to the north by the Dacatotan fault, with all major deposits clustered within a 10 km-wide corridor as seen in Figure 1 [37-39]. The dominant stratigraphic unit is the Carboniferous Qi'eshan Group, striking east-west and dipping southward at 43°-63° with well-developed schistosity, typically 600-2,000 m thick [40]. It comprises three members: a lower sequence of andesite and basalt lavas intercalated with pyroclastic tuffs representing early arc volcanism; a middle sequence of andesite and brecciated andesite lavas whose brecciated facies enhanced permeability and created favorable structural traps for hydrothermal fluid infiltration; and an upper sequence of lithic sandstone, tuffaceous siltstone, and intercalated volcanic flows recording waning arc activity [41]. The Jurassic Xishanyao Formation of sandstone, siltstone, mudstone, and conglomerate unconformably overlies the Qi'eshan Group and postdates all mineralization [42]. The ore-bearing intrusive complex, comprising plagiogranite porphyry, diorite porphyry, tonalite, quartz albite porphyry, and quartz porphyry, was emplaced into the Qi'eshan Group volcanic host rocks [43], [44]. Zircon U-Pb ages of 332.8 ± 2.5 Ma at Tuwu and 340.0 ± 3.0 to 332.2 ± 2.3 Ma at Yandong, together with Re-Os molybdenite ages of 322.7–331.8 Ma, confirm synchronous early Carboniferous magmatic and hydrothermal activity in a subduction-related arc setting [45-48].

2.1.3 Structural Controls

Three fault orientations govern mineralization localization in the belt [49], [50]. East-west-trending faults define the northern and southern boundaries of the ore-hosting corridor and controlled intrusive emplacement. Northwest-trending transtensional faults enhanced permeability at structural intersections. Northeast-trending faults, best expressed at Yandong, where they controlled tonalite dike elongation, served as the primary channelways for ascending ore fluids, evidenced by their spatial coincidence with maximum alteration intensity and highest ore grades. The

intersection of northeast- and east-west-trending faults defines the most structurally favorable sites in the belt, where converging fluid pathways created conditions necessary for metal precipitation. Fracture density derived from mapped fault and lineament traces therefore, serves as a primary ore-controlling criterion in this study.

2.1.4 Hydrothermal Alteration, Ore Mineralogy, and Resources

Alteration at both deposits follows the classic porphyry concentric zonation: a potassic core of secondary biotite, K-feldspar, and magnetite with disseminated chalcopyrite and bornite; a phyllic zone of quartz-sericite-pyrite marking maximum structural permeability and highest Cu grades; and a distal propylitic halo of chlorite, epidote, and calcite defining the outermost hydrothermal footprint [51], [52]. Ore minerals are dominated by chalcopyrite and pyrite with minor bornite, chalcocite, and magnetite in disseminations and veinlets, proceeding through four stages from early chalcopyrite-bornite-pyrite through quartz-molybdenite veins to a late carbonate-laumontite stage [53], [54]. The Tuwu deposit contains approximately 4.7 Mt Cu at 0.67% with 18 t Au, and the Yandong deposit contains 372 Mt ore at 0.58% Cu with significant Mo and Au [55].

2.1.5 Geological Basis for Joint Porphyry Copper and Hydrothermal Gold Prediction

The joint prediction of porphyry copper and hydrothermal gold mineralization reflects their shared geological controls within the Tuwu-Yandong belt rather than genetic equivalence. Both deposit types are hosted within the same Qi'eshan Group volcanic sequence, controlled by the same intersecting east-west, northeast-, and northwest-trending fault network, and generate spatially overlapping multi-element geochemical anomalies enriched in Cu, Au, Mo, As, Sb, Pb, and Zn that are inseparable at the 2,000 m prediction unit scale of the RGNR survey data used in this study. The structural intersections that localized porphyry copper emplacement are the same intersections that focused hydrothermal gold-bearing fluid flow, making the fracture density criterion simultaneously valid for both deposit types. Joint prospectivity prediction is therefore geologically justified and consistent with the known deposit distribution, where porphyry copper and hydrothermal gold occurrences cluster within the same 10 km-wide structural corridor.

2.2 Dataset

The multi-source dataset used in this study was compiled from regional geochemical survey data, structural lineament analysis, geophysical measurements, and known deposit location data covering the Tuwu-Yandong metallogenic belt in eastern Tianshan, Xinjiang, China. The study area was subdivided into a regular fishnet grid of 9,324 spatial prediction units using ArcGIS, each represented by a 16

× 16-pixel tile at 125 m resolution corresponding to a ground footprint of 2,000 m × 2,000 m, designed to preserve localized geochemical variability and structural context at the regional scale. Each tile is characterized by a 44-band feature stack composed of four data categories. The geochemical component comprises 39 elemental concentration maps derived from stream sediment sampling of the Regional Geochemistry National Reconnaissance (RGNR) program [56], [57], covering elements including copper pathfinder elements, gold pathfinder elements, and major geochemical indicators of hydrothermal alteration, together with one composite element association band. The structural component consists of one fracture density raster derived from mapped fault and lineament traces, capturing the density of structural features within each tile as a proxy for fluid pathway intensity. The geophysical component comprises one geophysical anomaly band capturing subsurface magnetic or gravity signatures associated with intrusive bodies and mineralized zones. The spatial proximity component comprises two distance-to-known-mineralization rasters encoding Euclidean distance from known porphyry copper deposit and hydrothermal gold deposit locations, respectively, serving as spatial prior knowledge inputs to the model. Global normalization statistics for all 44 bands were computed across all 9,324 tiles and saved as mean and standard deviation vectors, used to standardize tile features during both training and inference. From the full tile population, a labeled subset of 188 tiles was selected for model training and evaluation. This comprises 94 positive (mineralized) tiles defined by spatial overlap with known Au_Deposit and Gold_Deposit locations, and 94 negative (barren) tiles selected from areas lacking any evidence of mineralization and situated at maximum distance from known deposit locations and geochemical anomaly zones, consistent with established criteria for negative sample selection in MPM [58]. The labeled dataset was partitioned into a training set of 150 tiles (75 positive and 75 negative) and a held-out test set of 38 tiles (19 positive and 19 negative) using stratified random sampling to preserve class balance in both partitions. The complete set of 9,324 unlabeled tiles was used for generating the regional prospectivity map during inference.

3. METHODOLOGY

3.1 Framework Overview

The proposed framework, referred to as the Geologically Constrained Evidential Deep Learning framework (GeoEDL), addresses two interconnected problems in MPM for data-limited regions: overfitting under scarce mineralization labels and the inability to explain predictions in geological terms. The framework is organized as four sequential components: data interpretability, model interpretability and architecture, geological interpretability, and result interpretability, each building on the preceding one. The data interpretability component covers an input feature multicollinearity analysis, a two-criterion geological tile classification scheme assigning each spatial prediction

unit to a favorable, ambiguous, or barren category without mineralization labels (Stage 1), and the construction of geologically annotated contrastive training pairs (Stage 2). The model interpretability and architecture component covers the relational GNN architecture, Geological Supervised Contrastive pretraining (GeoSupCon), and geologically constrained evidential deep learning fine-tuning. The geological interpretability component delivers tile-level explanation through GNNExplainer and regional attribution mapping through geological factor decomposition. The result interpretability component produces simultaneous prospectivity, aleatoric uncertainty, and epistemic uncertainty maps through the evidential deep learning output head. Figure 2 presents a schematic overview of the complete framework.

Three deliberate design choices underpin the framework. A directional relational graph network is adopted because fault-controlled mineralization is governed by directional structural relationships that isotropic convolution cannot encode, with nine edge types capturing the anisotropic permeability of the northeast- and northwest-trending fault network. Contrastive pretraining precedes fine-tuning to decouple representation learning from classification, organizing the embedding space according to geological favorability rules before scarce mineralization labels are introduced, directly preventing overfitting. The loss function addresses three distinct failure modes: GeoSupCon enforces geological similarity in embedding space prior to labeled training; the evidential deep learning loss replaces softmax with a Dirichlet distribution for simultaneous prospectivity prediction and uncertainty quantification; and the two-term geological constraint penalty ensures predictions remain structurally and geochemically consistent with the known ore-controlling factor distributions throughout fine-tuning.

3.2 Data Interpretability

3.2.1 Input Feature Multicollinearity Analysis

Pearson correlation analysis was conducted across all 9,324 tiles using tile-level band means to characterize the multicollinearity structure of the 44-band input feature space as seen in Figure 3. The resulting correlation heatmap reveals coherent inter-element correlation clusters reflecting shared lithological affinities: ferromagnesian elements co-vary with mafic and ultramafic host rock compositions, copper-group pathfinder elements exhibit mutual enrichment consistent with hydrothermal alteration, and rock-forming elements reflect the calc-alkaline volcanic associations of the Qi'eshan Group. The presence of these correlation clusters confirms substantial redundant information across individual bands, meaning that univariate anomaly detection treating each element independently would fail to capture the multivariate geochemical signature of mineralization and would be distorted by inter-element correlations. The Mahalanobis distance statistic is therefore adopted as the geochemical

anomaly criterion in Stage 1 because it simultaneously accounts for the spatial variability of each element and the correlations between elements through the inverse covariance matrix, making it inherently suited to the spatial distribution analysis of multi-element geochemical data. Unlike univariate thresholding,

Mahalanobis distance measures how far a tile's complete geochemical signature deviates from the regional background in the full multivariate feature space, capturing coherent multi-element anomaly patterns associated with hydrothermal alteration regardless of which specific pathfinder elements are enriched.




Figure 2: Schematic overview of the four-level GeoEDL framework

3.2.2 Stage 1: Two-criterion geological tile classification

The study area is subdivided into a regular fishnet grid of spatial prediction units, each represented

as a 16×16 -pixel tile. To provide the contrastive pretraining stage with geologically meaningful positive and negative examples without requiring any mineralization labels, each tile is classified into one of

three geological categories: favorable, ambiguous, or barren, using two independent ore-controlling criteria

derived from the known mineralization system of the Tuwu-Yandong belt.




Figure 3: Pearson correlation heatmap of the 18 most intercorrelated bands from the 44-band input feature space

Criterion 1: Geochemical Anomaly Intensity

The multivariate geochemical signal of each tile is summarized using the Mahalanobis distance statistic D^2 , which measures how far a tile's geochemical signature deviates from the regional background in the full multivariate feature space. For each tile i with mean geochemical feature vector $x_i \in \mathbb{R}^p$, where p is the number of geochemical bands D^2 , is computed in Equation 1, where μ is the regional background mean vector, and Σ is the sample covariance matrix, both estimated from all 9,324 tiles without using any deposit location information.

$$D_i^2 = (x_i - \mu)^T \Sigma^{-1} (x_i - \mu) \quad (1)$$

Under the assumption of multivariate normality, D^2 follows a chi-squared distribution with p degrees of freedom. A tile is classified as geochemically anomalous if it satisfies Equation 2, where $\chi^2_{p, 0.05}$ is the chi-squared critical threshold at $p = 0.05$.

$$D_i^2 > \chi^2_{p, 0.05} \quad (2)$$

This criterion is deposit-type-agnostic, capturing coherent multi-element enrichment patterns associated with hydrothermal alteration without preselecting specific pathfinder elements, and accounts for inter-element correlations through the inverse covariance matrix Σ^{-1} .

Criterion 2: Structural Favorability through Fracture Density

Fracture density for each tile is computed as the total length of mapped fault and lineament traces per unit area within the tile boundary, following the linear density estimator as computed in Equation 3, where l_{ij} is the length of the j -th fault or lineament segment intersecting the tile i , n_i is the total number of such segments, and A is the tile area ($2,000 \text{ m} \times 2,000 \text{ m} = 4 \text{ km}^2$).

$$F_i = \frac{\sum_{j=1}^{n_i} l_{ij}}{A} \quad (3)$$

Fault and lineament traces were extracted from the regional geological map and digitized using ArcGIS. A tile is classified as structurally favorable if its fracture density exceeds the regional spatial median as defined in

Equation 4, where $\tilde{F}$ is the median fracture density across all 9,324 tiles.

$$F_i > \tilde{F} \quad (4)$$

The regional median serves as a data-driven threshold requiring no external parameter, and reflects the role of fault and fracture networks as primary fluid conduits for both porphyry copper and hydrothermal gold systems in the Tuwu-Yandong belt. A tile is

assigned to the favorable class only if it satisfies both criteria simultaneously, reflecting the co-spatial occurrence of geochemical anomaly and structural permeability that characterizes known mineralized environments in the Tuwu-Yandong system. Tiles satisfying neither criterion are assigned to the barren class. All remaining tiles satisfying one but not both criteria are classified as ambiguous, excluded from contrastive pair construction, but retained for fine-tuning and regional inference.




Figure 4: Data interpretability component of the GeoEDL framework

3.2.3 Stage 2: Contrastive pair construction

Geologically annotated contrastive pairs are constructed from the Stage 1 tile categories without using

any mineralization labels. For each favorable anchor tile, one other favorable tile is selected as the positive example, and N barren tiles are randomly sampled as

negative examples, with $N = 7$ in the current implementation. Negative tiles are resampled every epoch to prevent memorization of specific tile identities and to expose the encoder to diverse barren geological environments throughout training. Ambiguous tiles are excluded from all pairs to maintain the sharpness of the geological contrast between positive and negative examples. Together, Stage 1 and Stage 2 constitute the data interpretability component of the framework, providing geologically annotated contrastive training pairs derived entirely from ore-controlling criteria without any mineralization labels. The complete data interpretability workflow is illustrated in Figure 4.

3.3 Model Interpretability and Architecture

3.3.1 GeoSupCon - Geological supervised contrastive pretraining

The encoder is trained using the Geological Supervised Contrastive loss (GeoSupCon), which uses the geological category labels from Stage 1 as pseudo-labels to organize the embedding space according to geological favorability. The GeoSupCon loss enforces three simultaneous objectives: favorable tiles attract each other in embedding space, barren tiles attract each other in embedding space, and favorable and barren tiles repel each other. No mineralization labels are used at any point; geology drives everything. By the end of pretraining, the encoder has learned that geologically favorable tiles cluster in one region of embedding space and barren tiles cluster in another, purely from geological rules applied to the full unlabeled tile population. When fine-tuning subsequently begins with the scarce mineralization labels, the model is not required to learn geological understanding from scratch; it simply draws a decision boundary between two neighborhoods in embedding space that already exist. This is the core mechanism by which the framework reduces overfitting under scarce labels.

3.3.2 Geologically constrained evidential deep learning fine-tuning

Following pretraining, the GNN encoder weights are frozen for the first 60 epochs of fine-tuning to allow the GeoEDL classification head to stabilize before joint optimization begins. The GeoEDL head models the output as a Dirichlet distribution over class probabilities rather than a point estimate, enabling simultaneous quantification of aleatoric uncertainty arising from inherent data ambiguity and epistemic uncertainty arising from model ignorance alongside prospectivity prediction.

To ensure that predictions remain geologically consistent with the ore-controlling factor distributions of the Tuwu-Yandong system, the GeoEDL loss is

augmented with a two-term geological constraint penalty applied across all tiles simultaneously. The structural penalty penalizes high prospectivity predictions for tiles with fracture density below the structural threshold, ensuring that tiles predicted as prospective must be near fault structures. The geochemical penalty penalizes high prospectivity predictions for tiles with Mahalanobis D^2 below the chi-squared threshold, ensuring that tiles predicted as prospective must exhibit a geochemical anomaly. Both penalties are applied simultaneously during every training step, meaning that both geological ore-controlling factors govern model behavior across the entire study area throughout training. Together, the GeoSupCon pretraining and the geologically constrained evidential deep learning fine-tuning constitute the model interpretability component of the framework, ensuring that geological knowledge governs both the representation learning and prediction stages entirely without mineralization labels during pretraining. The complete workflow is illustrated in Figure 5. The formal loss formulations and all training hyperparameters are detailed in Section 3.3.3.

3.3.3 Model Architecture

The proposed GeoEDL consists of three components: a relational graph neural network encoder, a projection head used during pretraining only, and an evidential deep learning classification head used during fine-tuning.

Graph construction: Each 16×16 spatial prediction tile is converted into a graph $G = (V, E, R)$ where each of the 256 pixels becomes a node $v \in V$ carrying the full 44-dimensional feature vector, and spatial adjacency defines directed edges $(u, v) \in E$ with relation type $r \in R$. Nine directional relation types are defined corresponding to the eight compass directions (N, S, E, W, NE, NW, SE, SW) and a self-loop, indexed $R = \{0, 1, \dots, 8\}$, enabling direction-aware message passing that captures structural anisotropy in fault and fracture systems.

GNN encoder: The encoder consists of two stacked relational graph convolutional (RGCNConv) layers with a hidden dimension of 64. Each layer applies relation-specific linear transformations to aggregate messages from directional neighbors, followed by batch normalization and ReLU activation, as given in Equation 5, where $h_v^{(l)}$ is the node representation at layer l , $W_r^{(l)}$ is the relation-specific weight matrix for relation r , $\mathcal{N}_r(v)$ is the set of neighbors of node v connected by relation r , and σ is the ReLU activation function. Dropout with a rate of 0.2 is applied during pretraining and 0.4 during fine-tuning.

$$h_v^{(l+1)} = \sigma \left(W_0^{(l)} h_v^{(l)} + \sum_{r \in R} \sum_{u \in \mathcal{N}_r(v)} \frac{1}{|\mathcal{N}_r(v)|} W_r^{(l)} h_u^{(l)} \right) \quad (5)$$




Figure 5: Model interpretability of the GeoEDL framework

A tile-level graph embedding is obtained by concatenating global mean pooling and global max pooling over all node representations at the final layer, as given in Equation 6, producing a 128-dimensional tile embedding vector. The encoder has a total of 90,400 trainable parameters.

$$z_G = \left[\frac{1}{|V|} \sum_{v \in V} h_v^{(L)} \parallel \max_{v \in V} h_v^{(L)} \right] \quad (6)$$

Projection head: During Stage 3 pretraining only, a two-layer MLP projection head maps the 128-dimensional encoder output into a 32-dimensional

contrastive space: Linear (128 → 128) → BatchNorm1d → ReLU → Linear (128 → 32). The projection head is discarded after pretraining, and only the encoder weights are retained.

GeoEDL classification head: During Stage 4 fine-tuning, an evidential deep learning head is attached to the frozen encoder output: Linear (128 → 128) → ReLU → Dropout (0.4) → Linear (128 → 2). The head outputs raw evidence values $e_k \geq 0$ for each class k , from which Dirichlet distribution parameters are computed as given in Equation 7.

$$A_k = e_k + 1, S = \sum_{k=1}^K \alpha_k \quad (7)$$

The prospectivity probability, aleatoric uncertainty, and epistemic uncertainty for each tile are then derived analytically from the Dirichlet parameters as given in Equation 8, where $\hat{p}_k = \alpha_k/S$ and $K = 2$ is the number of classes. High total evidence S indicates confident predictions; low S indicates high epistemic uncertainty.

$$\hat{P} = \frac{\alpha_1}{S}, u_{\text{aleat}} = \sum_k \frac{\hat{p}_k(1 - \hat{p}_k)}{S + 1}, u_{\text{epist}} = \frac{K}{S} \quad (8)$$

GeoSupCon loss: The encoder is pretrained using the Geological Supervised Contrastive loss, which treats geological category labels from Stage 1 as pseudo-labels. For a batch containing both favorable (label 1) and barren (label 0) tiles, the loss maximizes agreement among tiles of the same geological category while repelling tiles of different categories, as defined in

Equation 9, where z_i is the ℓ_2 -normalized projection of tile i , $P(i)$ is the set of tiles in the batch sharing the same geological category as tile i , and $\tau = 0.3$ is the temperature parameter controlling the sharpness of the similarity distribution.

$$\mathcal{L}_{\text{GeoSupCon}} = -\frac{1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(z_i \cdot z_p / \tau)}{\sum_{a \neq i} \exp(z_i \cdot z_a / \tau)} \quad (9)$$

EDL loss with label smoothing: The fine-tuning objective combines an evidential MSE loss with KL divergence regularization and label smoothing of $\epsilon = 0.1$ as formulated in Equation 10, where $y_k^{(i)} = y_k(1 -$

$\epsilon) + \epsilon/K$ is the label-smoothed target, and β is a KL weight annealed from 0 to 0.05 over the first 30% of training epochs.

$$\mathcal{L}_{\text{EDL}} = \frac{1}{N} \sum_{i=1}^N \left[\sum_k \left(y_k^{(i)} - \frac{\alpha_k^{(i)}}{S^{(i)}} \right)^2 + \frac{\alpha_k^{(i)} / S^{(i)} \cdot (1 - \alpha_k^{(i)} / S^{(i)})}{S^{(i)} + 1} \right] + \beta \cdot \mathcal{L}_{\text{KL}} \quad (10)$$

Geological constraint loss: A two-term penalty enforces geological consistency across all tiles simultaneously, as defined in Equation 11, where $\hat{p}_i$ is the predicted prospectivity probability of the tile i ,

$\tau_{\text{frac}} = 23.3246$ is the structural threshold, $\chi^2 = 59.3035$ is the geochemical threshold, frac_i is the fracture density of the tile i , and D_i^2 is its Mahalanobis distance score.

$$\mathcal{L}_{\text{constraint}} = \frac{1}{N} \sum_{i=1}^N \hat{p}_i \left[\frac{\max(0, \tau_{\text{frac}} - \text{frac}_i)}{\tau_{\text{frac}}} + \frac{\max(0, \chi^2 - D_i^2)}{\chi^2} \right] \quad (11)$$

The penalty grows when tiles with low fracture density or low geochemical anomaly receive high prospectivity predictions. The total training objective combining the EDL loss and geological constraint penalty is given in Equation 12, where λ is annealed from 0 to $\lambda_{\text{max}} = 0.15$ over the first 40% of training to allow the EDL head to stabilize before geological enforcement begins.

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{EDL}} + \lambda \cdot \mathcal{L}_{\text{constraint}} \quad (12)$$

Optimization: During the encoder-frozen phase (epochs 1–60), only the GeoEDL head parameters are optimized using AdamW with a learning rate 5×10^{-4} and weight decay 10^{-3} . After epoch 60, the encoder is unfrozen, and the full model is optimized with a reduced learning rate of 5×10^{-5} under cosine annealing with warm restarts ($T_0 = 50$, $T_{\text{mult}} = 2$, $\eta_{\text{min}} = 10^{-6}$). Gradient norms are clipped to 1.0 throughout training. Early stopping with patience of 40 epochs is applied after the encoder is unfrozen. Five-fold stratified cross-validation is applied across the labeled dataset of 188

tiles to ensure stable performance estimation, with class balance preserved in each fold throughout all training runs. The complete model architecture specifications, training configuration, loss function parameters, and data

configuration are summarized in Tables 1, 2, 3, and 4, respectively. The complete GeoEDL model architecture is illustrated in Figure 6.

Table 1: Model architecture specifications

Component	Pretraining	Fine-tuning
Input feature dimension	44	44
GNN type	RGCNConv	RGCNConv
Number of GNN layers	2	2
Hidden dimension	64	64
Graph embedding dimension	128 (mean pool + max pool)	128 (mean pool + max pool)
Number of edge relation types	9	9
Projection head	Linear (128-32)	-
GeoEDL head	-	Linear(128-128)-ReLU-Dropout-Linear (128-2)
Total encoder parameters	90,400	90,400




Figure 6: Model architecture component of the GeoEDL framework

Table 2: Training configuration

Component	Pretraining	Fine-tuning
Optimizer	AdamW	AdamW
Learning rate	5×10^{-4}	5×10^{-4} (head) 5×10^{-5} (full)
Weight decay	1×10^{-3}	1×10^{-3}
Batch size	16 favorable + 16 barrens	16
Epochs	200	400 (max)
Early stopping patience	-	40 epochs
Encoder freeze epochs	-	60
LR scheduler	Cosine annealing ($\eta_{\min}=1 \times 10^{-5}$)	Cosine annealing with warm restarts ($T_0=50$, $T_{\text{mult}}=2$, $\eta_{\min}=1 \times 10^{-6}$)
Gradient clipping	1.0	1.0
Dropout rate	0.2	0.4

Table 3: Loss function parameters

Component	Pretraining	Fine-tuning
Temperature (τ)	0.3	-
Label smoothing (ϵ)	-	0.1
KL weight (β)	-	0 $\rightarrow$ 0.05 over the first 30% of epochs
Constraint weight ($\lambda_{\max}$)	-	0.15
Constraint annealing	-	0 $\rightarrow$ $\lambda_{\max}$ over the first 40% of training
Constraint batch size	-	64

Table 4: Data configuration

Component	Pretraining	Fine-tuning	
Negative per anchor (N)	7	-	
Cross-validation	-	5-fold stratified	
Random seed	42	42	

3.4 Geological Interpretability

3.4.1 GNNExplainer

Following model fine-tuning, GNNExplainer is applied to each positively predicted tile to identify which pixel locations within each tile were most critical in driving that specific prediction. For each tile, the explainer learns a continuous soft mask over the node features of the tile's local computational graph that maximally preserves the model's prediction while using the minimum number of features. The resulting node importance scores identify which spatial locations within the tile carry the highest geochemical importance for the prospectivity prediction, providing tile-level geological reasoning for individual predictions that cannot be obtained from global feature importance methods.

3.4.2 Geological factor decomposition

At the regional level, an occlusion-based geological factor decomposition is applied to all tiles to produce three spatially continuous attribution maps showing the relative contribution of each ore-controlling factor to prospectivity predictions across the entire study area.

For each tile, three factor contributions are estimated independently: the structural contribution is

measured by masking fault-parallel directional edges and recording the resulting drop in prediction confidence; the geochemical contribution is measured by replacing geochemical node features with regional background values and recording the drop; and the contextual contribution is the remaining prediction score after structural and geochemical contributions are accounted for, reflecting the influence of spatial relationships with neighboring tiles. Each contribution is expressed as a percentage of the total feature-driven prediction change, producing three regional maps: structural attribution, geochemical attribution, and contextual attribution that together decompose the model's prospectivity prediction into its geological components across the entire study area. The complete geological interpretability workflow, encompassing both GNNExplainer tile-level explanation and geological factor decomposition, is illustrated in Figure 7.

3.5 Result Interpretability

The result interpretability level is delivered through three simultaneous output maps generated by the GeoEDL evidential deep learning head for all prediction units across the study area. The prospectivity map encodes the predicted mineralization probability for each tile and constitutes the primary exploration targeting output.




Figure 7: The geological interpretability component of the GeoEDL framework




Figure 8: Result interpretability component of the GeoEDL framework

The aleatoric uncertainty map encodes inherent data ambiguity. High aleatoric uncertainty reflects locations where the geochemical and structural signal is contradictory or where sampling density is insufficient, a condition that cannot be reduced through additional model training and requires improved field sampling. The epistemic uncertainty map encodes model ignorance. High epistemic uncertainty reflects locations where the tile's geological character is unlike anything encountered during training, identifying areas where the model's prediction is least reliable.

The combination of the prospectivity map and the epistemic uncertainty map constitutes the exploration decision map: tiles with simultaneously high prospectivity and high epistemic uncertainty represent the highest-priority targets for future drilling and sampling campaigns. This uncertainty-guided prioritization provides a quantitative and geologically principled basis for allocating exploration resources in the Tuwu-Yandong belt, converting model uncertainty from a limitation into an actionable exploration signal. The result interpretability workflow is illustrated in Figure 8.

4. RESULTS AND DISCUSSION

4.1 Model Performance

The geology-guided tile classification in Stage 1 identified 344 geologically favorable tiles (3.7%), reflecting the restricted spatial extent of co-occurring geochemical anomaly and structural permeability within the Tuwu-Yandong belt, consistent with the narrow 10 km-wide ore-hosting corridor bounded by the Kanggurtag and Dacotan faults. The remaining 4,318 ambiguous (46.3%) and 4,662 barren (50.0%) tiles correspond predominantly to the structurally quiet zones and barren Qi'eshan Group volcanic country rocks distal from the porphyry intrusive complex. The 344 favorable tiles were used to construct contrastive pairs in Stage 2, and the pretrained encoder achieved a separation ratio of 0.9570 with a within-favorable similarity of 0.8674 and within-barren similarity of 0.9109, confirming that the encoder successfully learned to distinguish geologically favorable from barren environments before any mineralization labels were introduced. The GeoSupCon pretraining loss decreased consistently from 3.45 to 2.89

over 200 epochs (Figure 9; C), validating stable convergence of the geological supervision signal.

Evaluation on the held-out test set of 38 tiles demonstrated strong predictive performance, with a ROC-AUC of 0.9418, PR-AUC of 0.9119, F1 score of 0.9268, overall accuracy of 0.9211, sensitivity of 1.0000, specificity of 0.8421, and precision of 0.8636 at the optimal threshold of 0.4681 (Table 6, Figure 9; A, B, D). The perfect sensitivity of 1.0000 zero false negatives is geologically critical in the Tuwu-Yandong system, where the known porphyry copper and hydrothermal gold deposits are confined to a structurally constrained corridor, and any missed positive tile corresponds to an undetected mineralized zone within an already limited target window. The three false positives are geologically interpretable as tiles within the Qi'eshan Group volcanic sequence that satisfy both the geochemical anomaly and fracture density criteria but do not correspond to confirmed deposit locations, likely reflecting the broad propylitic alteration halo extending beyond the economic ore shell at Tuwu and Yandong. Five-fold cross-validation yielded a mean ROC-AUC of 0.8866 ± 0.0913 , confirming stable generalization across data partitions despite the scarcity of positive labels.

Regional prospectivity prediction across all 9,324 tiles produced probabilities ranging from 0.0365 to 0.9623, with 3,469 tiles (37.2%) exceeding the optimal threshold. The high-prospectivity belt (Figure 13) is spatially concentrated in the central portion of the study area, coinciding with the Carboniferous Qi'eshan Group host rock distribution and the east-west structural corridor between the Kanggurtag and Dacotan faults where the Tuwu and Yandong intrusive complexes are emplaced. Low-prospectivity zones correspond to the Jurassic Xishanyao Formation cover sequence to the north and structurally quiet areas distal from the northeast- and northwest-trending fault intersections that governed hydrothermal fluid focusing in the belt. Positive tiles exhibit systematically higher epistemic uncertainty than negative tiles (Figure 9; E, F), reflecting that mineralized zones in the Tuwu-Yandong belt represent geologically rare and spatially restricted environments that were underrepresented in the training population relative to the broad barren volcanic country rock background.

Table 6: Model performance metrics

Metric	Value / Description
Optimal Classification Threshold	0.4681
Accuracy	0.9211 (92.11%)
ROC-AUC	0.9418 (94.18%)
PR-AUC	0.9119 (91.19%)
F1-score	0.9268 (92.68%)
Sensitivity	1.0000 (100.00%)
Specificity	0.8421 (84.21%)
Precision	0.8636 (86.36%)
CV ROC-AUC	0.8866 ± 0.0913 (88.66% $\pm$ 9.13%)




Figure 9: Model performance evaluation: A) Confusion matrix, B) ROC curve, C) GeoSupCon Pretraining Loss Curve, D) Precision-Recall curve, E) Aleatoric Uncertainty Distribution (positive vs. negative tiles), F) Epistemic Uncertainty Distribution (positive vs. negative tiles)

4.2 Ablation study

To evaluate the contribution of each component of the proposed GeoEDL framework, six ablation variants were systematically compared against the full model using the fixed classification threshold of 0.4681 across all variants. The results are presented in Table 7.

Removing the GeoSupCon pretraining stage produced a marginal improvement in ROC-AUC from 0.9418 to 0.9446 but reduced sensitivity from 1.0000 to 0.9474, meaning the model without pretraining missed one known mineralized tile that the full model correctly identified. This demonstrates that while pretraining does

not always improve discriminative performance measured by AUC, it produces a more geologically conservative model that avoids missing true deposit locations, a critical requirement in mineral exploration.

Removing the geological constraint loss reduced ROC-AUC to 0.9391, F1 to 0.8649, and sensitivity to 0.8421, confirming that the two-term structural and geochemical penalty plays a meaningful role in maintaining geologically consistent predictions during fine-tuning. Without the constraint loss the model produces predictions that are less aligned with known ore-controlling factor distributions, resulting in more false negatives.

Removing the fracture criterion from Stage 1 so that tile classification relies solely on the geochemical anomaly criterion produced a ROC-AUC of 0.9446 and maintained perfect sensitivity of 1.0000 but reduced specificity from 0.8421 to 0.7895, generating more false positives. This indicates that the fracture criterion contributes to reducing false alarms by excluding geochemically anomalous but structurally unfavorable tiles from the favorable category during contrastive pretraining.

Replacing the evidential deep learning head with a standard softmax classifier produced the largest drop in ROC-AUC at 0.9224 and reduced PR-AUC from

0.9119 to 0.8789, confirming that the Dirichlet-based evidential formulation provides measurable discriminative benefit beyond simply quantifying uncertainty. The cross-validation variability also increased with the softmax variant, suggesting that the GeoEDL head contributes to training stability under scarce label conditions.

Removing the directional relational edges, reducing the graph to self-loops only with no message passing between neighboring pixels, produced the same ROC-AUC as the no-constraint variant at 0.9391 but showed the largest drop in cross-validation stability at 0.8534 ± 0.0825 . This confirms that direction-aware message passing across the nine structural orientations contributes meaningfully to generalization across different data partitions, capturing spatial anisotropy in the geochemical and structural signal that cannot be recovered from node features alone.

Overall, the ablation study shows that different components contribute to different aspects of model reliability: the EDL head improves discrimination and uncertainty-aware stability, the geological constraint loss improves geological consistency and sensitivity, GeoSupCon pretraining reduces missed mineralized tiles, and directional relational edges improve cross-validation stability.

Table 7: Ablation study

Variant	ROC-AUC	Accuracy	F1	PR-AUC	Sensitivity	Specificity	CV ROC-AUC
Full model	0.9418	0.9211	0.9268	0.9119	1.0000	0.8421	0.8866 ± 0.0913
No pretraining	0.9446	0.8947	0.9000	0.9381	0.9474	0.8421	0.8889 ± 0.0688
No constraint loss	0.9391	0.8684	0.8649	0.9154	0.8421	0.8947	0.8829 ± 0.0915
No fracture criterion	0.9446	0.8947	0.9048	0.9215	1.0000	0.7895	0.8869 ± 0.0878
No EDL (softmax)	0.9224	0.8947	0.9000	0.8789	0.9474	0.8421	0.8784 ± 0.1061
No relational edges	0.9391	0.8684	0.8649	0.9129	0.8421	0.8947	0.8534 ± 0.0825

4.3 Geological Interpretability Findings

4.3.1 GNNExplainer findings

GNNExplainer was applied to all 3,469 positively predicted tiles to identify which pixel locations were most critical to each prospectivity prediction. The aggregated pixel importance analysis reveals that high-importance pixels are consistently concentrated in the central region of each 16×16 tile, indicating that the model assigns geochemical feature importance to the core of each spatial prediction unit rather than its periphery. This pattern directly reflects the concentric hydrothermal alteration zonation documented at both Tuwu and Yandong: the potassic core and phyllic zone, which carry the highest Cu grades and the most intense multivariate geochemical anomaly in the Cu-Mo-Au-As-Sb pathfinder element suite, are spatially confined to the intrusive center of each deposit, with anomaly intensity diminishing outward through the propylitic halo toward the barren Qi'eshan Group volcanic country rocks. The model therefore assigns highest geochemical importance to tile pixels overlying

the intrusive core and proximal alteration zones rather than the distal propylitic margins, consistent with the known spatial distribution of ore-grade mineralization at Tuwu and Yandong.

4.3.2 Geological factor decomposition findings

Regional geological factor decomposition across all 9,324 tiles produced mean contributions of 51.6% geochemical, 46.2% contextual, and 2.2% structural across the study area.

The geochemical attribution map (Figure 10) reveals that geochemically driven prospectivity is concentrated in a spatially coherent belt through the central study area, directly overlying the known porphyry copper and hydrothermal gold deposit locations. This distribution is consistent with the strong multivariate Cu-Mo-Au-As-Sb-Pb-Zn hydrothermal footprint generated by the potassic-phyllic-propylitic alteration system of the Tuwu-Yandong porphyry complex, which produces coherent multi-element

anomalies detectable at the 2,000 m tile scale in the RGNR stream sediment survey data.

The contextual attribution map (Figure 11) shows that spatial neighborhood context contributes 46.2% to prospectivity predictions, broadly distributed across the study area. This substantial contextual contribution reflects the GNN's capacity to propagate geochemical information across neighboring tiles through directional edge connections, capturing the broad geochemical dispersion halos surrounding the Tuwu and Yandong deposit centers that extend beyond individual tile boundaries into the surrounding Qi'eshan Group volcanic host rocks. In zones where individual tile geochemical signals are attenuated by dilution or cover from the Jurassic Xishanyao Formation, proximity to anomalous neighbors within the ore-hosting corridor

sustains elevated prospectivity predictions that a pixel-based model operating on isolated tiles would miss.

The structural attribution map (Figure 12) shows a mean contribution of only 2.2%, which does not reflect geological importance of the northeast- and northwest-trending fault network that controlled porphyry emplacement and hydrothermal fluid focusing in the Tuwu-Yandong belt. Rather, it reflects a resolution limitation inherent to the dataset: regional fracture density rasters derived from mapped fault and lineament traces at the 2,000 m tile scale cannot resolve the fine-scale fault intersections, vein arrays, and stockwork zones that localize ore-grade mineralization at Tuwu and Yandong at the deposit scale. Higher-resolution structural mapping incorporating drone-based lineament extraction or high-resolution satellite imagery would be expected to substantially increase the structural attribution signal in future applications.




Figure 10: Geochemical attribution map




Figure 11. Contextual attribution map




Figure 12: Structural attribution map

4.4 Prospectivity, Uncertainty Maps, and Exploration Decision (Result interpretability findings)

The GeoEDL output head generated three simultaneous spatial maps across all 9,324 tiles: a prospectivity map (Figure 13), an epistemic uncertainty map (Figure 14), and an aleatoric uncertainty map (Figure 15). All three maps share the same spatial extent and are overlaid with known deposit locations and mapped fault traces to enable direct geological interpretation.

The prospectivity map (Figure 13) reveals a spatially coherent high-prospectivity belt concentrated along the central east-west corridor of the study area, spatially coincident with the Carboniferous Qi'eshan Group volcanic sequence, the plagiogranite porphyry and tonalite intrusive corridor between the Kanggurtag and Dacotan faults, and the known Tuwu and Yandong deposit locations marked in Figure 13. A total of 3,469 tiles (37.2%) exceeded the optimal classification threshold of 0.4681, of which 1,908 were identified as priority exploration targets. The geometry of this high-prospectivity zone reflects the northeast- and northwest-

trending fault intersections that governed hydrothermal fluid focusing throughout the belt. Low-prospectivity zones flanking the central corridor correspond to the Jurassic Xishanyao Formation cover sequence in the northern portion of the study area and structurally distal barren volcanic country rocks of the Qi'eshan Group to the east and west, where neither the geochemical anomaly nor the fracture density criteria are satisfied simultaneously.

The epistemic uncertainty map (Figure 14) shows elevated uncertainty distributed along the margins of the central high-prospectivity corridor, particularly in zones where the Qi'eshan Group volcanic sequence transitions to Jurassic sedimentary cover and where the geochemical signature of the hydrothermal system attenuates away from the intrusive centers of Tuwu and Yandong. These peripheral zones were underrepresented in the 94-label training population relative to the well-characterized deposit cores, resulting in higher model ignorance at the corridor margins. Zones where high prospectivity in Figure 13 and high epistemic uncertainty in Figure 14 spatially coincide identifiable by comparing the two maps at the same location represent the highest-priority targets for follow-up field investigation, where the model predicts mineralization consistent with the

Tuwu-Yandong ore-controlling criteria, but confidence is limited by insufficient training coverage.

The aleatoric uncertainty map (Figure 15) delineates zones of inherent data ambiguity where geochemical and structural signals are internally contradictory, concentrated in areas where stream sediment sampling density in the RGNR survey is insufficient to produce a spatially consistent geochemical signature at the 2,000 m tile scale. Unlike epistemic uncertainty, aleatoric uncertainty cannot be reduced through additional model training or more deposit labels. These zones, identifiable as high-value areas in Figure 15, require denser geochemical sampling coverage, higher-resolution structural mapping, or both, providing a spatially explicit criterion for directing future field data acquisition in the Tuwu-Yandong belt independent of any modelling decisions.

The combination of Figures 13 and 14 constitutes the exploration decision output of the framework: tiles showing simultaneously high prospectivity and high epistemic uncertainty identify where mineralization is predicted, but model confidence is lowest, converting uncertainty from a limitation into a geologically grounded and spatially explicit guide for allocating drilling and sampling resources in the Tuwu-Yandong belt.




Figure 13: Prospectivity map




Figure 14. Epistemic uncertainty map




Figure 15: Aleatoric uncertainty map

4.5 DISCUSSION

The GeoEDL framework achieved a ROC-AUC of 0.9418 and perfect sensitivity of 1.0000 using only 94 positive labels across 9,324 tiles, confirming that GeoSupCon pretraining and the geological constraint loss together effectively address overfitting under scarce label conditions. The geological factor decomposition reveals that geochemical anomaly drives 51.6% of prospectivity predictions, consistent with the strong Cu-

Mo-Au-As-Sb hydrothermal footprint of the Tuwu-Yandong porphyry-gold system at the 2,000 m tile scale. The 46.2% contextual contribution reflects the GNN's capacity to propagate geochemical information across neighboring tiles, capturing dispersion halos beyond individual tile boundaries that pixel-based models would miss. The minimal 2.2% structural contribution reflects a dataset resolution limitation rather than geological unimportance. Regional fracture density rasters cannot

resolve the fine-scale NE- and NW-trending fault intersections that controlled fluid focusing at Tuwu and Yandong, and higher-resolution structural data would be expected to increase this attribution substantially.

The ablation study confirms that the geological constraint loss contributes most critically to sensitivity, its removal reducing sensitivity from 1.0000 to 0.8421, validating that geological penalties prevent statistically plausible but geologically inconsistent predictions under scarce label conditions. The evidential deep learning head produced the largest improvement in ROC-AUC stability, while directional relational edges contributed most to cross-validation consistency by capturing structural anisotropy irrecoverable from node features alone. Together, the GNNExplainer tile-level explanations and regional geological factor decomposition constitute a two-scale interpretability system bridging regional prospectivity mapping and deposit-scale exploration targeting an output that single probability maps from standard MPM frameworks cannot provide.

5. CONCLUSION

This study proposed GeoEDL, a four-level explainable AI framework for mineral prospectivity mapping with geological constraints, applied to the Tuwu-Yandong porphyry copper and hydrothermal gold belt in eastern Tianshan, Xinjiang, China. The framework addresses two persistent problems in AI-driven MPM for data-limited regions overfitting under scarce mineralization labels and the inability to explain predictions in geological terms through geology-guided contrastive pretraining, geologically constrained evidential deep learning fine-tuning, and multi-level XAI output. Applied to 9,324 tiles with only 94 positive labels, GeoEDL achieved a ROC-AUC of 0.9418, PR-AUC of 0.9119, F1 score of 0.9268, and perfect sensitivity of 1.0000, with no known mineralized tile misclassified as barren. Geological factor decomposition confirmed that geochemical anomaly drives 51.6% of prospectivity predictions and spatial context drives 46.2%, consistent with the Cu-Mo-Au-As-Sb hydrothermal footprint and broad dispersion halos of the Tuwu-Yandong system, while the 2.2% structural contribution reflects a dataset resolution limitation at the 2,000 m tile scale rather than geological unimportance of the NE- and NW-trending fault network. The simultaneous prospectivity, epistemic uncertainty, and aleatoric uncertainty maps provide a geologically principled exploration decision output that converts model uncertainty into spatially explicit guidance for field data acquisition. Future work should evaluate GeoEDL generalizability across other deposit types and geological settings and incorporate higher-resolution structural and geochemical data to improve tile classification specificity and structural attribution.

REFERENCES

1. E. J. M. Carranza, *Geochemical anomaly and mineral prospectivity mapping in GIS*. Elsevier, 2009.
2. R. Zuo, "Geodata Science-Based Mineral Prospectivity Mapping: A Review," *Natural Resources Research*, vol. 29, no. 6, pp. 3415–3424, Dec. 2020, doi: 10.1007/s11053-020-09700-9.
3. R. Zuo, Y. Xiong, Z. Wang, J. Wang, and O. P. Kreuzer, "A New Generation of Artificial Intelligence Algorithms for Mineral Prospectivity Mapping," *Natural Resources Research*, vol. 32, no. 5, pp. 1859–1869, Oct. 2023, doi: 10.1007/s11053-023-10237-w.
4. R. Zuo, "Geodata Science-Based Mineral Prospectivity Mapping: A Review," *Natural Resources Research*, vol. 29, no. 6, pp. 3415–3424, Dec. 2020, doi: 10.1007/s11053-020-09700-9.
5. R. Zuo and Y. Xu, "Graph Deep Learning Model for Mapping Mineral Prospectivity," *Math. Geosci.*, vol. 55, no. 1, pp. 1–21, Jan. 2023, doi: 10.1007/s11004-022-10015-z.
6. T. Sun, F. Chen, L. Zhong, W. Liu, and Y. Wang, "GIS-based mineral prospectivity mapping using machine learning methods: A case study from Tongling ore district, eastern China," *Ore Geol. Rev.*, vol. 109, pp. 26–49, Jun. 2019, doi: 10.1016/j.oregeorev.2019.04.003.
7. M. Parsa and A. Maghsoudi, "Assessing the effects of mineral systems-derived exploration targeting criteria for random Forests-based predictive mapping of mineral prospectivity in Ahar-Arasbaran area, Iran," *Ore Geol. Rev.*, vol. 138, p. 104399, Nov. 2021, doi: 10.1016/j.oregeorev.2021.104399.
8. T. Li, R. Zuo, Y. Xiong, and Y. Peng, "Random-Drop Data Augmentation of Deep Convolutional Neural Network for Mineral Prospectivity Mapping," *Natural Resources Research*, vol. 30, no. 1, pp. 27–38, Feb. 2021, doi: 10.1007/s11053-020-09742-z.
9. K. Sun *et al.*, "A Review of Mineral Prospectivity Mapping Using Deep Learning," *Minerals*, vol. 14, no. 10, p. 1021, Oct. 2024, doi: 10.3390/min14101021.
10. S. Li, J. Chen, C. Liu, and Y. Wang, "Mineral Prospectivity Prediction via Convolutional Neural Networks Based on Geological Big Data," *Journal of Earth Science*, vol. 32, no. 2, pp. 327–347, Apr. 2021, doi: 10.1007/s12583-020-1365-z.
11. Q. Li, G. Chen, and L. Luo, "Mineral prospectivity mapping using attention-based convolutional neural network," *Ore Geol. Rev.*, vol. 156, p. 105381, May 2023, doi: 10.1016/j.oregeorev.2023.105381.
12. Z. Wang and R. Zuo, "Mineral prospectivity mapping using a joint singularity-based weighting method and long short-term memory network," *Comput. Geosci.*, vol. 158, p. 104974, Jan. 2022, doi: 10.1016/j.cageo.2021.104974.
13. Y. Xu and R. Zuo, "An Interpretable Graph Attention Network for Mineral Prospectivity Mapping," *Math. Geosci.*, vol. 56, no. 2, pp. 169–190, Feb. 2024, doi: 10.1007/s11004-023-10076-8.
14. Z. Meng and R. Zuo, "Self-Supervised Graph Contrastive Learning for Mineral Prospectivity Mapping," *Math. Geosci.*, vol. 57, no. 7, pp. 1265–1282, Oct. 2025, doi: 10.1007/s11004-025-10191-8.

15. Y. Xu, R. Zuo, and G. Zhang, "The graph attention network and its post-hoc explanation for recognizing mineralization-related geochemical anomalies," *Applied Geochemistry*, vol. 155, p. 105722, Aug. 2023, doi: 10.1016/j.apgeochem.2023.105722.
16. Z. Chen and R. Zuo, "Geological-knowledge-guided graph self-supervised pretraining framework for identifying mineralization-related geochemical anomalies," *Comput. Geosci.*, vol. 199, p. 105913, May 2025, doi: 10.1016/j.cageo.2025.105913.
17. A. Lachaud, M. Adam, and I. Mišković, "Comparative Study of Random Forest and Support Vector Machine Algorithms in Mineral Prospectivity Mapping with Limited Training Data," *Minerals*, vol. 13, no. 8, p. 1073, Aug. 2023, doi: 10.3390/min13081073.
18. Z. Zhang *et al.*, "An Uncertainty-Quantification Machine Learning Framework for Data-Driven Three-Dimensional Mineral Prospectivity Mapping," *Natural Resources Research*, vol. 33, no. 4, pp. 1393–1411, Aug. 2024, doi: 10.1007/s11053-024-10349-x.
19. Y. Liu, D. Zhang, Z. Li, H. Fan, and W. Peng, "Dirichlet-Based Uncertainty-Aware Deep Learning for Explainable Mineral Prospectivity Mapping," *Natural Resources Research*, vol. 35, no. 2, pp. 823–840, Apr. 2026, doi: 10.1007/s11053-025-10604-9.
20. R. Zuo *et al.*, "Explainable artificial intelligence models for mineral prospectivity mapping," *Sci. China Earth Sci.*, vol. 67, no. 9, pp. 2864–2875, Sep. 2024, doi: 10.1007/s11430-024-1309-9.
21. H. Cheng *et al.*, "GIS-based mineral prospectivity mapping using machine learning methods: A case study from Zhuonuo ore district, Tibet," *Ore Geol. Rev.*, vol. 161, p. 105627, Oct. 2023, doi: 10.1016/j.oregeorev.2023.105627.
22. X. Lv and G. Wang, "GIS-based mineral prospectivity mapping using machine learning methods: A case study from Duobaoshan ore district, northeastern China," *Ore Geol. Rev.*, vol. 175, p. 106352, Dec. 2024, doi: 10.1016/j.oregeorev.2024.106352.
23. Y. Xiong, R. Zuo, Z. Luo, and X. Wang, "A Physically Constrained Variational Autoencoder for Geochemical Pattern Recognition," *Math. Geosci.*, vol. 54, no. 4, pp. 783–806, May 2022, doi: 10.1007/s11004-021-09979-1.
24. R. Zuo, Z. Luo, Y. Xiong, and B. Yin, "A Geologically Constrained Variational Autoencoder for Mineral Prospectivity Mapping," *Natural Resources Research*, vol. 31, no. 3, pp. 1121–1133, Jun. 2022, doi: 10.1007/s11053-022-10050-x.
25. F. Yang, R. Zuo, Y. Xiong, Y. Xu, J. Nie, and G. Zhang, "Dual-Branch Convolutional Neural Network and Its Post Hoc Interpretability for Mapping Mineral Prospectivity," *Math. Geosci.*, vol. 56, no. 7, pp. 1487–1515, Oct. 2024, doi: 10.1007/s11004-024-10137-6.
26. Z. Luo, R. Zuo, and Y. Xiong, "Visual Interpretable Deep Learning Algorithm for Geochemical Anomaly Recognition," *Natural Resources Research*, vol. 31, no. 5, pp. 2211–2223, Oct. 2022, doi: 10.1007/s11053-022-10080-5.
27. W.-J. Xiao, "Paleozoic accretionary and collisional tectonics of the eastern Tianshan (China): Implications for the continental growth of central Asia," *Am. J. Sci.*, vol. 304, no. 4, pp. 370–395, Apr. 2004, doi: 10.2475/ajs.304.4.370.
28. Y. J. Chen, F. Pirajno, G. Wu, J. P. Qi, and X. L. Xiong, "Epithermal deposits in North Xinjiang, NW China," *International Journal of Earth Sciences*, vol. 101, no. 4, pp. 889–917, Jun. 2012, doi: 10.1007/s00531-011-0689-4.
29. B. F. Windley, D. Alexeiev, W. Xiao, A. Kröner, and G. Badarch, "Tectonic models for accretion of the Central Asian Orogenic Belt," *J. Geol. Soc. London.*, vol. 164, no. 1, pp. 31–47, Jan. 2007, doi: 10.1144/0016-76492006-022.
30. F. M. H. Sihombing, R. M. Palin, H. S. R. Hughes, and L. J. Robb, "Improved mineral prospectivity mapping using graph neural networks," *Ore Geol. Rev.*, vol. 172, p. 106215, Sep. 2024, doi: 10.1016/j.oregeorev.2024.106215.
31. C. Han *et al.*, "Geological characteristics and genesis of the Tuwu porphyry copper deposit, Hami, Xinjiang, Central Asia," *Ore Geol. Rev.*, vol. 29, no. 1, pp. 77–94, Jul. 2006, doi: 10.1016/j.oregeorev.2005.07.032.
32. L.-C. Zhang, W.-J. Xiao, K.-Z. Qin, J. Ji, and X. Yang, "Types, geological features and geodynamic significances of gold-copper deposits in the Kanggurtag metallogenic belt, eastern Tianshan, NW China," *International Journal of Earth Sciences*, vol. 93, no. 2, pp. 224–240, Apr. 2004, doi: 10.1007/s00531-004-0383-x.
33. J. Mao, R. J. Goldfarb, Y. Wang, C. J. Hart, Z. Wang, and J. Yang, "Late Paleozoic base and precious metal deposits, East Tianshan, Xinjiang, China: Characteristics and geodynamic setting," *Episodes*, vol. 28, no. 1, pp. 23–36, Mar. 2005, doi: 10.18814/epiiugs/2005/v28i1/003.
34. N. Yang, Z. Zhang, J. Yang, and Z. Hong, "Applications of data augmentation in mineral prospectivity prediction based on convolutional neural networks," *Comput. Geosci.*, vol. 161, p. 105075, Apr. 2022, doi: 10.1016/j.cageo.2022.105075.
35. J. W. Mao *et al.*, "A review of the Cu–Ni sulphide deposits in the Chinese Tianshan and Altay orogens (Xinjiang Autonomous Region, NW China): Principal characteristics and ore-forming processes," *J. Asian Earth Sci.*, vol. 32, no. 2–4, pp. 184–203, Mar. 2008, doi: 10.1016/j.jseaes.2007.10.006.
36. F. Pirajno, R. Seltmann, and Y. Yang, "A review of mineral systems and associated tectonic settings of northern Xinjiang, NW China," *Geoscience Frontiers*, vol. 2, no. 2, pp. 157–185, Apr. 2011, doi: 10.1016/j.gsf.2011.03.006.
37. L. Zhang, W. Xiao, K. Qin, and Q. Zhang, "The adakite connection of the Tuwu–Yandong copper porphyry belt, eastern Tianshan, NW China: trace element and Sr–Nd–Pb isotope geochemistry," *Miner. Depos.*, vol. 41, no. 2, pp. 188–200, May 2006, doi: 10.1007/s00126-006-0058-6.
38. Y.-H. Wang, F.-F. Zhang, J.-J. Liu, and C.-Y. Que, "Genesis of the Fuxing porphyry Cu deposit in Eastern Tianshan, China: Evidence from fluid inclusions and C–H–O–S–Pb isotope systematics," *Ore Geol. Rev.*, vol. 79, pp. 46–61, Dec. 2016, doi: 10.1016/j.oregeorev.2016.04.022.

39. Y. Xiong, R. Zuo, and E. J. M. Carranza, "Mapping mineral prospectivity through big data analytics and a deep learning algorithm," *Ore Geol. Rev.*, vol. 102, pp. 811–817, Nov. 2018, doi: 10.1016/j.oregeorev.2018.10.006.
40. S. Sheng, Y. Wang, J. Tian, X. Chen, Y. Ning, and Y. Dong, "Graph attention network-based mineral prospectivity prediction: A case study of copper exploration in eastern Tien Shan, China," *Ore Geol. Rev.*, vol. 184, p. 106766, Sep. 2025, doi: 10.1016/j.oregeorev.2025.106766.
41. B. Yin, R. Zuo, and Y. Xiong, "Mineral Prospectivity Mapping via Gated Recurrent Unit Model," *Natural Resources Research*, vol. 31, no. 4, pp. 2065–2079, Aug. 2022, doi: 10.1007/s11053-021-09979-2.
42. J. Zhao *et al.*, "Genesis and evolution of framboidal pyrite and its implications for the ore-forming process of Carlin-style gold deposits, southwestern China," *Ore Geol. Rev.*, vol. 102, pp. 426–436, Nov. 2018, doi: 10.1016/j.oregeorev.2018.09.022.
43. B. Xiao *et al.*, "Magmatic evolution of the Tuwu–Yandong porphyry Cu belt, NW China: Constraints from geochronology, geochemistry and Sr–Nd–Hf isotopes," *Gondwana Research*, vol. 43, pp. 74–91, Mar. 2017, doi: 10.1016/j.gr.2015.09.003.
44. N. Yang, Z. Zhang, J. Yang, Z. Hong, and J. Shi, "A Convolutional Neural Network of GoogLeNet Applied in Mineral Prospectivity Prediction Based on Multi-source Geoinformation," *Natural Resources Research*, vol. 30, no. 6, pp. 3905–3923, Dec. 2021, doi: 10.1007/s11053-021-09934-1.
45. W. An, C. Xue, Y. Zhao, and C. Li, "Two Periods of Porphyry Cu Mineralization and Metallogenic Implications in the Tuwu–Yandong Belt (NW China), Based on Re–Os Systematics of Molybdenite," *Minerals*, vol. 12, no. 9, p. 1127, Sep. 2022, doi: 10.3390/min12091127.
46. S. Sheng *et al.*, "Knowledge–Data Collaboration-Driven Mineral Prospectivity Prediction with Graph Attention Networks," *Minerals*, vol. 15, no. 11, p. 1164, Nov. 2025, doi: 10.3390/min15111164.
47. Y. Xu, Z. Li, Z. Xie, H. Cai, P. Niu, and H. Liu, "Mineral prospectivity mapping by deep learning method in Yawan-Daqiao area, Gansu," *Ore Geol. Rev.*, vol. 138, p. 104316, Nov. 2021, doi: 10.1016/j.oregeorev.2021.104316.
48. Y. Xiong and R. Zuo, "Recognition of geochemical anomalies using a deep autoencoder network," *Comput. Geosci.*, vol. 86, pp. 75–82, Jan. 2016, doi: 10.1016/j.cageo.2015.10.006.
49. C. Han *et al.*, "Ages and tectonic implications of the mafic–ultramafic-carbonatite intrusive rocks and associated Cu–Ni, Fe–P and apatite-vermiculite deposits from the Quruqtagh district, NW China," *Ore Geol. Rev.*, vol. 95, pp. 1106–1122, Apr. 2018, doi: 10.1016/j.oregeorev.2016.07.011.
50. R. Zuo, O. P. Kreuzer, J. Wang, Y. Xiong, Z. Zhang, and Z. Wang, "Uncertainties in GIS-Based Mineral Prospectivity Mapping: Key Types, Potential Impacts and Possible Solutions," *Natural Resources Research*, vol. 30, no. 5, pp. 3059–3079, Oct. 2021, doi: 10.1007/s11053-021-09871-z.
51. H. Yuan, P. Shen, and H. Pan, "Nature and evolution of hydrothermal fluids in the formation of the Tuwu porphyry copper deposit in the Eastern Tianshan Mountains, NW China," *J. Asian Earth Sci.*, vol. 165, pp. 210–227, Oct. 2018, doi: 10.1016/j.jseaes.2018.07.018.
52. Y. Wang *et al.*, "Overprinting mineralization in the Paleozoic Yandong porphyry copper deposit, Eastern Tianshan, NW China—Evidence from geology, fluid inclusions and geochronology," *Ore Geol. Rev.*, vol. 100, pp. 148–167, Sep. 2018, doi: 10.1016/j.oregeorev.2017.04.013.
53. X.-H. Deng *et al.*, "New ⁴⁰Ar/³⁹Ar ages from the Kalatag district in the Eastern Tianshan, NW China: Constraints on the timing of Cu mineralization and stratigraphy," *Ore Geol. Rev.*, vol. 100, pp. 250–262, Sep. 2018, doi: 10.1016/j.oregeorev.2016.08.006.
54. Y. Wang *et al.*, "Multiple mineralization events of the Paleozoic Tuwu porphyry copper deposit, Eastern Tianshan: evidence from geology, fluid inclusions, sulfur isotopes, and geochronology," *Miner. Depos.*, vol. 54, no. 7, pp. 1053–1076, Oct. 2019, doi: 10.1007/s00126-018-0859-4.
55. W. Tan, Q. Mao, M. Yu, Y. Sun, and X. Lv, "Mineralization of the Tuwu Porphyry Cu Deposit in Eastern Tianshan, NW China: Insights From In Situ Trace Elements of Chlorite and Pyrite," *Front. Earth Sci. (Lausanne)*, vol. 9, Jun. 2021, doi: 10.3389/feart.2021.648177.
56. X. Xie, X. Mu, and T. Ren, "Geochemical mapping in China," *J. Geochem. Explor.*, vol. 60, no. 1, pp. 99–113, Nov. 1997, doi: 10.1016/S0375-6742(97)00029-0.
57. X. Xie *et al.*, "Multi-scale geochemical mapping in China," *Geochemistry: Exploration, Environment, Analysis*, vol. 8, no. 3–4, pp. 333–341, Nov. 2008, doi: 10.1144/1467-7873/08-184.
58. V. Nykänen, I. Lahti, T. Niiranen, and K. Korhonen, "Receiver operating characteristics (ROC) as validation tool for prospectivity models — A magmatic Ni–Cu case study from the Central Lapland Greenstone Belt, Northern Finland," *Ore Geol. Rev.*, vol. 71, pp. 853–860, Dec. 2015, doi: 10.1016/j.oregeorev.2014.09.007.