Machine learning for automated texture classification of Ni-Cu-PGE ores: a comparative study from Västerbotten, Sweden

Aghata Zarelli Viana[1]*, Natalia Katherine Soler-Aragon[1], Cecilia Lund[1], Mehdi Parian[1]

[1]Minerals and Metallurgical Engineering, Dept. of Civil, Environmental, and Natural Resources Engineering, Luleå University of Technology, Luleå, SE-971 87, Sweden

*Corresponding author. E-mail: aghata.viana@ltu.se

Abstract

Quantitative characterization of ore textures is important in process mineralogy, since texture influences comminution behavior, mineral liberation, and metallurgical performance. However, texture classification is still largely qualitative and subjective, which limits its reproducibility and its use in process modeling. This study evaluates supervised and unsupervised machine learning approaches for automated ore texture classification of the Lappvattnet Ni-Cu-PGE deposit, Västerbotten, Sweden, using μ-XRF mineral maps of drill core.

A total of 459 mineral maps were divided into 100 × 100 pixel tiles (5 × 5 mm), from which 73 quantitative texture descriptors (composition, morphology, diversity and entropy, and dissemination) were extracted and reduced to 46 by correlation-based feature selection. A supervised gradient-boosted tree classifier (XGBoost), trained on 2,975 tiles labelled into 14 ore texture classes adapted from the scheme of Barnes et al. (2017, 2018), and an unsupervised multi-view clustering approach were applied to this same feature set. Evaluated on core samples held out from training, the XGBoost model achieved a macro-averaged F1-score of 0.71 ± 0.03 (0.78 ± 0.01 when tiles were split at random). Performance was highest for texturally distinctive classes such as ball ore, semi-massive/massive ore and barren veins, and most errors were between neighboring classes within the same lithological domain. Mineral composition and sulfide dissemination features carried most of the discrimination (38% and 28% of the feature importance, respectively). Hierarchical multi-view clustering into 14 groups recovered the lithological domains and the separation between mineralized and non-mineralized tiles (98% and 89% of the tiles fell in clusters dominated by their own domain and mineralization group, respectively), but not the finer textural subdivisions (adjusted Rand index of 0.26 with the labelled texture classes).

The supervised model is suitable for routine texture classification but requires expert-labelled data, whereas the clustering recovers the main lithological and mineralization groups without labels and can be used to delineate textural domains and to check label consistency.

Keywords: ore texture; μ-XRF; automated mineralogy; machine learning; unsupervised clustering; Ni-Cu-PGE sulfide ores; geometallurgy

1 Introduction

Ore texture is widely used to reconstruct the evolution of magmatic Ni-Cu-PGE sulfide systems. The spatial relationship between sulfide and silicate minerals records processes such as magma differentiation, sulfide liquid segregation, and post-magmatic remobilization. These relationships are the basis of the ore-texture classification for Ni-Cu-PGE deposits proposed by Barnes et al. (2017, 2018).

Ore texture also influences particle behavior in comminution and concentration processes (Lund et al., 2015; Pérez-Barnuevo et al., 2018b). However, ore texture classification is still largely manual, qualitative, and dependent on expert interpretation, which makes it difficult to use in quantitative geometallurgical and process models.

Advances in computer vision, machine learning (ML), and automated mineralogy can help address this limitation. Automated mineralogy systems based on SEM (QEMSCAN, Gottlieb et al., 2000; MLA, Fandrich et al., 2007; TIMA, Hrstka et al., 2018) and μ-XRF (AMICS, Menzies et al., 2022) allow rapid quantification of modal mineralogy, mineral liberation, grain size, and mineral associations. ML has been applied to ore texture and drill core classification (Gu et al., 2024; Guntoro et al., 2020; Koch et al., 2019; Merrill-Cifuentes et al., 2022; Pérez-Barnuevo et al., 2018a; Tiu, 2017); to lithology and mineral classification (W. Chen et al., 2023; Seo et al., 2022; Xu et al., 2021; Zhang et al., 2024); and to mineral and grain segmentation in optical thin section images (Ferreira, 2023; Filippo et al., 2021; Maitre et al., 2019; Morales et al., 2025; Rubo et al., 2019). However, these approaches have not yet been combined into a texture classification workflow for magmatic Ni-Cu-PGE sulfide ores.

The Barnes et al. (2017, 2018) classification is well suited for automation, since its textural classes are defined by measurable criteria, such as modal sulfide abundance and textural characteristics. The classes also have genetic significance and provide a standardized framework that can be linked to processing behavior.

In this study, μ-XRF mineral maps from the Lappvattnet Ni-Cu-PGE deposit were used to test a supervised machine learning model, developed by Soler-Aragon (2026), which was built on the ore texture classification proposed by Barnes et al. (2017, 2018) and adapted to the Ni-Cu-PGE deposits in the Nickel Line. A further question was whether a comparable classification could be achieved using unsupervised models applied to the same quantified texture features, in which each textural parameter (e.g., grain size, grain shape, mineral association) is represented as its own matrix. These quantitative descriptors could also be compared with processing behavior and used as geometallurgical proxies. The objectives of this study are therefore (i) to evaluate the supervised XGBoost model on a labeled Lappvattnet dataset, (ii) to test whether unsupervised multi-view clustering of the same descriptors recovers comparable textural groupings, and (iii) to evaluate which descriptors discriminate the ore textures and how the two approaches agree.

1.1 Lappvattnet Ni-Cu-PGE deposit

The Lappvattnet deposit is one of the highest-grade nickel deposits within the Nickel Line, with inferred resources of 780 kt @1.35% Ni (Reddick & Lindholm, 2020). The deposit is located approximately 70 km southwest of Skellefteå in the Västerbotten District, northern Sweden (Figure 1), and it consists of a magmatic Ni-Cu-PGE sulfide mineralization associated with small mafic-

ultramafic intrusions emplaced into migmatitic metagreywackes and biotite-graphite gneisses. The mineralized intrusion extends for approximately 620 m along strike and varies from 1 to 20 m in width (Åkerman, 1987).

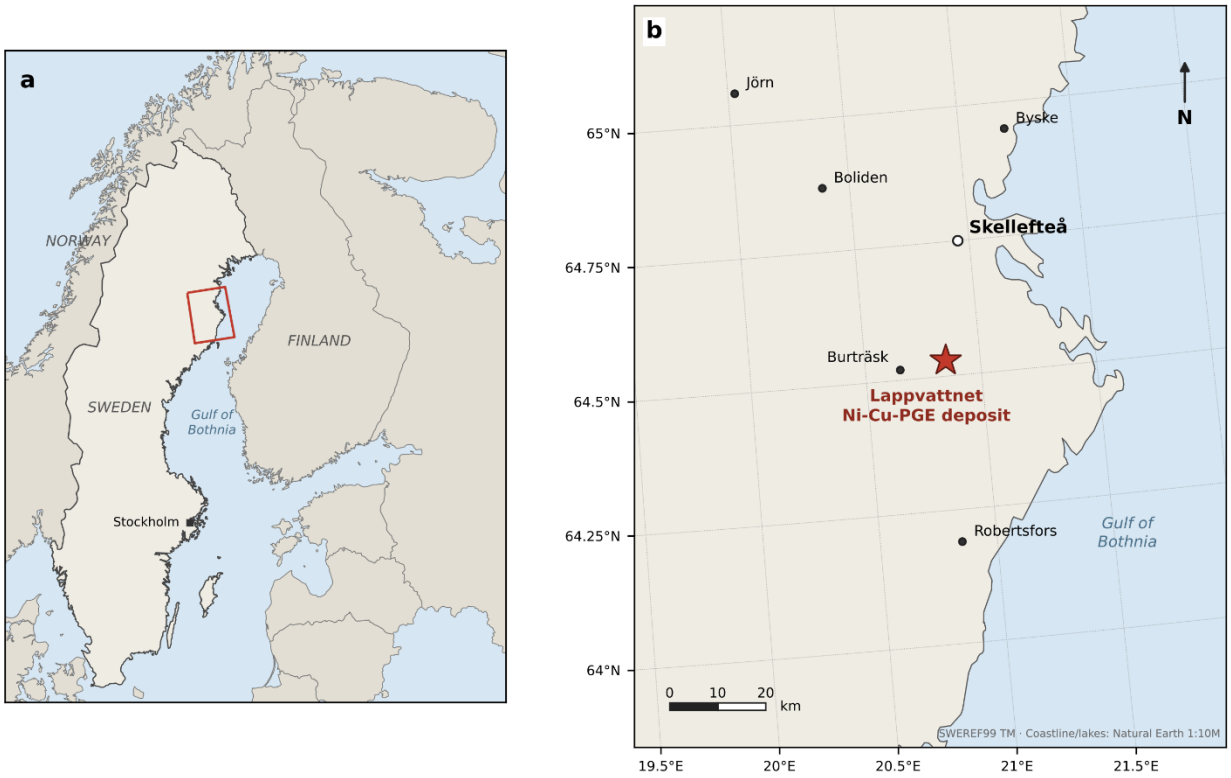

Figure 1. Location of the Lappvattnet deposit.

At Lappvattnet, the nickel mineralization occurs both in small, lens-shaped ultramafic bodies and in biotite-graphite gneisses (Ekström, 1978). The ultramafic bodies consist mostly of meta-peridotite, composed mainly of orthopyroxene, serpentine, and olivine, surrounded by a thin meta-pyroxenite margin, composed mainly of amphibole. The associated metasediments typically occur as veined gneiss or migmatite, composed of quartz, plagioclase, and biotite, with garnet and microcline as common accessory phases (Ekström, 1978).

Although the ultramafic bodies are the primary host of the Ni-Cu sulfide mineralization, local remobilization of sulfides has occurred, resulting in sulfide mineralization within the surrounding gneiss as fragmental and breccia-hosted zones (Reddick & Lindholm, 2020).

Mineralogical studies conducted by Ekström (1978) characterized the Lappvattnet mineralization as a complex nickel ore dominated by pyrrhotite, pentlandite, chalcopyrite, pyrite, and magnetite, with marcasite, mackinawite, melnikovite-pyrite, and violarite as secondary phases, gersdorffite, niccolite, and maucherite as notable minor minerals, and ilmenite, sphalerite, titanite, chromite, and bismuth tellurides present as accessories. Based on this mineralogical characterization, Ekström (1978) grouped the mineralization into three main ore types (Table 1).

Table 1. Mineralization and textural classification identified in each host rock in the Lappvattnet deposit.

Mineralization types	Host rock	Textural characteristic
Disseminated ore	Meta-ultramafic	Spot-type (or nest-type)
	Gneiss and meta-ultramafic	Vein and network type
	Gneiss and meta-ultramafic	Fine-grained type
Breccia ore	Gneiss	"Ball ore" – breccia ore
Compact ore	Gneiss	Massive

2 Methods and samples

The material used for this study consisted of 55 m of drill core from the Lappvattnet deposit, sampled from seven drill holes along the deposit to represent the different mineralization styles. These samples were logged and scanned using µ-XRF, resulting in 459 mineral maps (Table 2) used for the texture classification.

Table 2. Number of mineral maps for each lithology group.

Host Rock	Number of images
Gneiss	128
Meta-ultramafic	331
Total	**459**

The classification workflow was divided into five steps: (1) core logging and sampling, (2) image acquisition, (3) image processing, (4) feature extraction, and (5) automated classification.

2.1 Image acquisition

The 459 scans were obtained by µ-XRF analysis using a Bruker M4 Tornado Plus equipped with a Rh X-ray tube and two XFlash silicon drift detectors (SDDs), with the acquisition parameters listed in Table 3.

Table 3. Acquisition parameters for the Bruker M4 Tornado Plus µ-XRF analysis.

Parameters	Value
Vacuum	2 mbar
Acceleration voltage	50 kV
Pixel size	50 µm
Anode current	300 or 600 µA
Dwell time per pixel	10 or 20 ms

The scans were processed using the M4 software, which gives a grayscale X-ray intensity map of the total counts per pixel. Brighter regions correspond mainly to Fe-, Ni- and Cu-bearing sulfides and oxides and darker regions to silicates, a contrast similar to that of backscattered electron (BSE) images.

The acquired spectral data were then processed using the Advanced Mineral Identification and Characterization System (AMICS) software, which classifies minerals on a pixel-by-pixel basis by comparing each pixel's spectrum against its built-in mineral database (Figure 2).

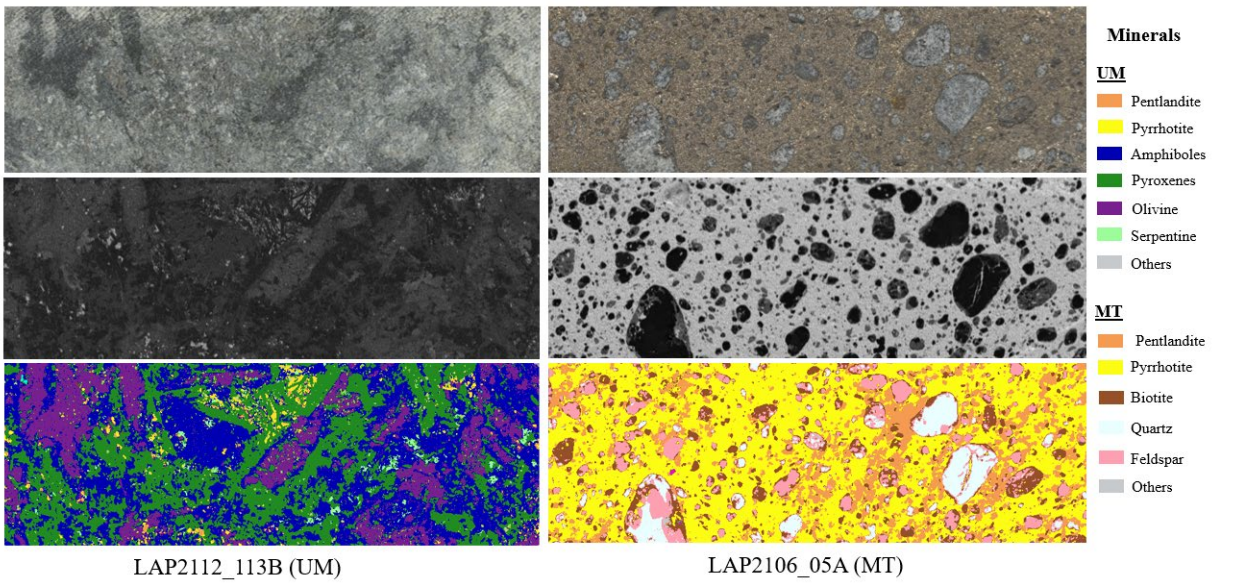

Figure 2. Example of the image products obtained for two drill-core scans (columns): RGB photograph (top), µ-XRF X-ray intensity map (center) and AMICS mineral map (bottom), with the mineral legend used throughout this study. Left column: meta-ultramafic host with disseminated sulfides; right column: meta-sedimentary host with ball ore.

2.2 Image processing

The 459 mineral maps generated by AMICS were preprocessed in Python 3.14.6, where white or black background stripes were removed from the borders of the images.

The images were then cut into 100 × 100 pixel tiles (5 × 5 mm at the 50 µm pixel size), a tile size selected in preliminary tests. Each tile was converted from the AMICS color image into an integer image in which each pixel value identifies a mineral phase. The last step consisted of labelling the tiles used in the supervised classification with the hierarchical classification key described in the Supervised Learning section, using the drill core logging as reference. Each tile was labelled according to the texture that dominates its 5 × 5 mm area.

2.3 Feature Extraction

The feature extraction step converts the µ-XRF mineral maps into quantitative descriptors for each tile. The 73 features were computed per tile and grouped into four descriptor families (Table 4): composition, morphology, diversity and entropy, and dissemination. The morphological descriptors were calculated on three inputs: the whole tile (general morphology), and dilated masks of the sulfide (chalcopyrite, pyrite, pyrrhotite, pentlandite and other Ni minerals) and carbonate phases (disk, r = 2 px). The dilation connects nearby pixels belonging to the same vein. The dissemination descriptors were calculated on the non-dilated sulfide mask.

Soler-Aragon (2026) tested different feature schemes to remove redundant features, and the best-performing scheme used correlation-based feature reduction. The absolute Pearson correlation coefficients between all feature pairs were calculated using only the training dataset, and for each pair with a coefficient greater than 0.85, one of the two features was removed. This reduced the 73 features to 46, spanning six groups: mineral composition, mineral diversity and entropy, general morphology, sulfide morphology, carbonate morphology, and sulfide dissemination, which were used in this study for both the supervised and the unsupervised classification.

Table 4. Quantitative descriptors extracted from the mineral map tiles by feature group, and features retained after correlation-based reduction (Soler-Aragon, 2026).

Feature group	Key metrics	Method	Extracted (n)	Retained (n)
Compositional	Modal mineralogy: area fraction per mineral group	Pixel count per mineral group divided by the total tile pixels	20	20 (all mineral fractions)
Morphological (general, sulfide and carbonate inputs)	Edge content	Canny edge detection ($\sigma = 2$): edge pixel count and edge density (edge pixels / total tile pixels)	6	3 (edge pixel count, all three inputs)
	Skeleton network	Skeletonization of the detected edges: skeleton pixel count and skeleton density	6	0
	Line segment morphology	Connected skeleton segments longer than a minimum length (5 px general; 30 px masks): number of segments, mean and maximum line length	9	4 (mean line length, all inputs; maximum line length, general)
	Line orientation	Circular statistics of the segment orientations: mean, standard deviation and anisotropy (mean resultant length, 0–1)	9	8 (mean and standard deviation, all inputs; anisotropy, sulfide and carbonate)
Diversity and entropy	Mineral heterogeneity	Mineral richness (S); Shannon entropy $H = -\Sigma\, p_i \log_2 p_i$ and normalized entropy $J = H / \log_2 S$; Simpson index $D = \Sigma\, p_i^2$ and Gini-Simpson index $1 - D$	5	2 (mineral richness, Shannon entropy)
Dissemination	Overall abundance	Sulfide pixel count (N_S) and sulfide area fraction (A_S = N_S/N_T)	2	0

Grain frequency	Number of connected sulfide grains (K, $a \geq a_min = 2$ px) and grain density (G_den = K/N_T)	2	0
Grain size distribution	Grains per sulfide area (G_norm = K/N_S); mean, median, maximum and standard deviation of grain areas	5	3 (G_norm, mean and standard deviation of grain area)
Size-class partition	Small ($a \leq a_small = 20$ px) and large ($a \geq a_large = 100$ px) grains: counts (K_small, K_large), area fractions (F_small, F_large), count fractions (C_small, C_large), and largest-component fraction (F_max = a_max/N_S)	7	4 (K_large, F_large, C_small, F_max)
Composite features	Dissemination indices: $DI_1 = F_small \times (1 - F_max)$ and $DI_2 = (F_small \times G_norm)/(F_max + \varepsilon)$	2	2
Total		**73**	**46**

2.4 Ore textures in the Ni-Cu-PGE Lappvattnet Deposit

The ore textures present in the Lappvattnet deposit were classified based on the nomenclature proposed by Barnes et al. (2017, 2018), which distinguishes disseminated, net-textured, semi-massive, massive, and sulfide-matrix breccia ores based on the textural relationships between sulfide and silicate minerals.

Based on the examination of magmatic Ni-Cu-PGE sulfide deposits from around the world, the authors identified common textural characteristics and proposed a generic framework relating ore textures to the evolution of sulfide and silicate melts. This framework was deliberately built on examples taken, as far as possible, from undeformed deposits that have undergone little or no post-magmatic alteration or metamorphic modification. The classification recognizes three principal textural groups based on the modal sulfide abundance within the host rock: (1) disseminated ores, containing up to about 5 modal % sulfides (most typically 0.5–2.5%) occurring as isolated grains or small aggregates; (2) net-textured ores, characterized by continuous sulfide networks enclosing cumulus silicate minerals and containing about 30–70 modal % sulfides; and (3) massive ores, composed predominantly of sulfides and typically containing 80–100% sulfides, or more than 90% by volume where semi-massive ores are distinguished.

The authors also recognize transitional ore types with intermediate sulfide abundances between disseminated and net-textured ores. These ores commonly consist of centimeter-scale domains of disseminated and net-textured mineralization occurring together within the same sample, referred

to as “patchy net-textured ores”. Associated with these domains it is common to observe a “globular” texture, consisting of sulfide aggregates with a sub-spherical morphology at a scale from tens of micrometers to a few centimeters.

The Lappvattnet deposit, as other deposits in the Nickel Line, was affected by metamorphism and sulfide remobilization, which modified the primary magmatic textures and produced textures that are not part of the original Barnes et al. (2017, 2018) framework, such as ball ore. For this reason, this classification was adapted to the Lappvattnet context (Figure 3), both in the definition of the classes and in the sulfide thresholds used to separate them. In the meta-ultramafic domain, mineralized tiles (≥ 1% sulfides) were classified as disseminated (1–10%), intermediate (10–30%), net-textured (30–70%) and semi-massive/massive (> 70%); in the meta-sedimentary domain, they were classified as disseminated or ball ore according to the continuity of the sulfide matrix. The class names are used in a purely descriptive sense and do not carry any genetic interpretation: for instance, a texture classified as disseminated describes isolated sulfide grains or aggregates in a silicate matrix, without interpreting whether it is primary magmatic or the result of remobilization. This follows the recommendation of Barnes et al. (2018) that purely descriptive terminology should be used where evidence for tectonic remobilization is lacking or ambiguous.

The 14 textural classes identified in the Lappvattnet deposit, their main sulfide and gangue minerals and their defining characteristics are summarized in Table 5, together with a representative drill-core image of each class.

Table 5. Textural classes of the Lappvattnet deposit with their main sulfides, major gangue minerals, defining characteristics and a representative drill-core image. po = pyrrhotite, pn = pentlandite, ccp = chalcopyrite, ol = olivine, opx = orthopyroxene, cpx = clinopyroxene, amp = amphibole, bt = biotite, srp = serpentine, car = carbonate, gr = graphite, fsd = feldspar, qz = quartz.

Lithological Domain	Texture type	Main Sulfides	Major gangue	Characteristics	Rock image
Meta-Ultramafic	00 UM NM Fine grained	-	ol, opx, cpx, amp, bt, srp	Sulfides <1%; fine-grained silicate matrix (grain size <2 mm)	
	01 UM NM Coarse grained	-	ol, opx, cpx, amp, bt, srp	Sulfides <1%; coarse-grained silicate matrix (grain size >2 mm)	
	02 UM NM Barren veins	-	ol, opx, cpx, amp, bt, srp, car	Sulfides <1%; barren veins/veinlets (e.g. carbonate, serpentine) cutting the host	
	03 UM M Vein/veinlets	po>pn>ccp	ol, opx, cpx, amp, bt, srp	Sulfides ≥1% hosted in veins and veinlets cutting the host rock	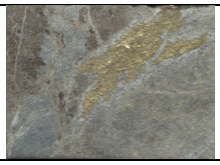

Lithological Domain	Texture type	Main Sulfides	Major gangue	Characteristics	Rock image
	04 UM M Disseminated	po>pn>ccp	ol, opx, cpx, amp, bt, srp	1–10% sulfide as isolated interstitial (along grain boundaries) or globular blebs in a silicate matrix	
	05 UM M Intermediate (Patchy/net-textured)	po>pn>ccp	ol, opx, cpx, amp, bt, srp	10–30% sulfide; net-textured patches alternating with sulfide-poor or disseminated domains	
	06 UM M Net-textured	po>pn>ccp	ol, opx, cpx, amp, bt, srp	30–70% sulfide forming a continuous interstitial network around cumulus silicate grains (olivine/pyroxene pseudomorphs)	
	07 UM M Semi-massive/massive	po>pn>ccp	ol, opx, cpx, amp, bt, srp	>70% sulfide; continuous sulfide matrix enclosing isolated silicate grains or inclusions, up to near-pure sulfide	
Meta-sediments	08 MT NM Fine grained layered	-	bt, gr, fsd, qz, amp	Sulfides <1%; foliated; fine-grained layering (layers <2 mm)	
	09 MT NM Coarse-grained layered	-	bt, gr, fsd, qz, amp	Sulfides <1%; foliated; coarse-grained layering (layers 0.2–1 cm)	
	10 MT NM Coarse grained massive	-	bt, gr, fsd, qz, amp	Sulfides <1%; not foliated; coarse, massive texture with porphyroblasts >1 cm	
	11 MT M Vein/veinlets	po>pn>ccp	bt, gr, fsd, qz, amp	Sulfide veins and veinlets cutting the metasedimentary country rock	
	12 MT M Disseminated	po>pn>ccp	bt, gr, fsd, qz, amp	Sulfides ≥1% as dispersed grains in the country rock	

Lithological Domain	Texture type	Main Sulfides	Major gangue	Characteristics	Rock image
	13 MT M Ball Ore	po>pn>ccp	bt, gr, fsd, qz, amp	Continuous sulfide matrix enclosing rounded to subangular country-rock inclusions	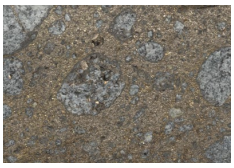

2.5 Texture classification

2.5.1.1 Supervised Learning

Following the proposed workflow, each image tile was assigned to one of 14 ore texture classes (8 meta-ultramafic and 6 meta-sedimentary) using a hierarchical classification key (Figure 3). Tiles were first separated by host lithology (meta-ultramafic or meta-sedimentary) and by sulfide content (mineralized, above ~1% sulfides, or non-mineralized). Non-mineralized tiles were further subdivided by vein presence and grain size (meta-ultramafic) or by foliation and layer size (meta-sedimentary). Mineralized tiles were first separated by the presence of sulfide veins/veinlets and then, in the meta-ultramafic domain, by sulfide proportion into disseminated (~1–10%), intermediate (~10–30%), net-textured (~30–70%) and semi-massive/massive (above ~70%) ore; in the meta-sedimentary domain, they were separated by the continuity of the sulfide matrix into disseminated and ball ore.

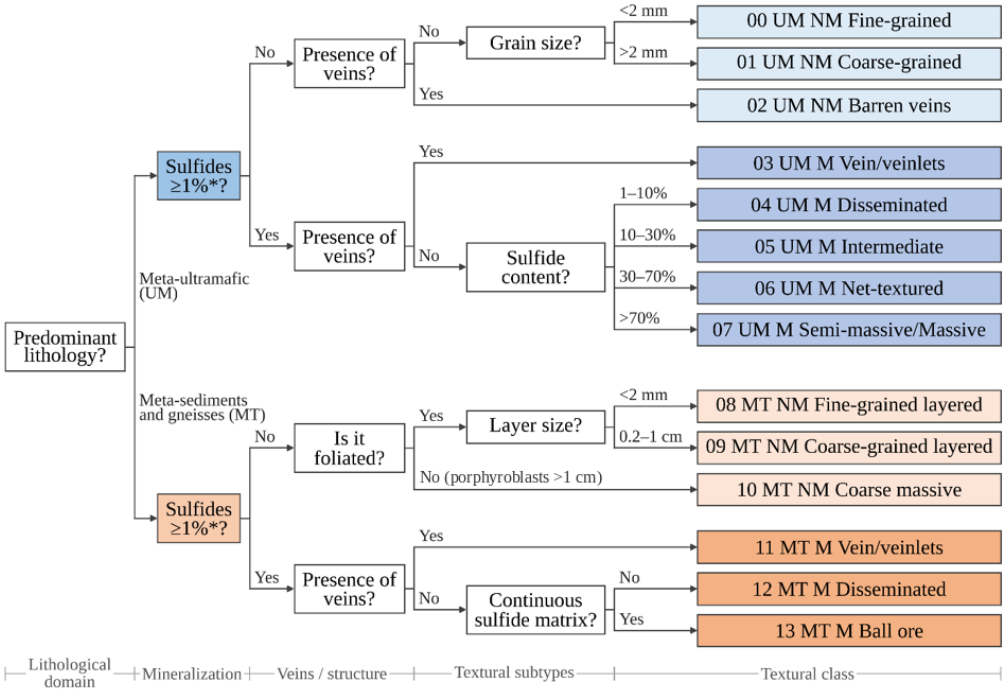

Figure 3. Hierarchical ore texture classification key used to assign the 14 training labels, adapted from Barnes et al. (2017, 2018) to the Lappvattnet deposit. UM = meta-ultramafic, MT = meta-sedimentary and gneissic, M = mineralized, NM = non-mineralized. Asterisks mark sulfide proportions visually estimated during core logging and validated by the µ-XRF results.

A minimum of 200 labelled tiles per class was targeted to limit class imbalance, resulting in 2,975 labelled tiles from 342 of the 459 µ-XRF mineral maps, 1,670 meta-ultramafic and 1,305 meta-sedimentary. Class sizes range from 200 to 296 tiles, except semi-massive/massive meta-ultramafic ore (85 tiles from 11 maps), which remained below the target. Of the 2,975 labels, 702 (24%) were accepted from a proposal computed automatically from the measured sulfide proportion, grain size and sulfide morphology following the same key; the remaining 2,273 were assigned manually.

Soler-Aragon (2026) tested several supervised classifiers: decision tree (Quinlan, 1986), random forest (Breiman, 2001), linear support vector machine (Cortes & Vapnik, 1995), LightGBM (Ke et al., 2017), and XGBoost (T. Chen & Guestrin, 2016), each combined with different feature schemes. XGBoost trained on the 46-feature set achieved test performance comparable to the best classifiers with a reduced feature set and was therefore adopted as the classification model in this study. The model was configured as a multiclass gradient-boosted tree classifier (100 trees, learning rate 0.05, maximum depth 5, subsample and column-subsample ratios of 0.8, L1/L2 regularization of 0.1/2.0), with balanced sample weights to compensate for class imbalance. The configuration was adopted unchanged; no hyperparameters were tuned on the dataset of this study.

Two evaluation protocols were used. In the *random protocol*, each class was capped at 200 randomly drawn tiles and the tiles were split into a training set (80%) and a test set (20%) by stratified random sampling. In the *grouped protocol*, all tiles from the same core sample were kept on the same side of the split, using stratified group 5-fold cross-validation repeated with five random seeds (25 folds in total). The grouped protocol prevents tiles cut from the same sample, which share mineralogy and fabric, from appearing in both training and test sets, and therefore estimates performance on unseen core samples (Roberts et al., 2017).

Performance was summarized by accuracy, macro-averaged F1-score and weighted F1-score; the macro-averaged F1-score weights all texture classes equally regardless of their size and is therefore more informative than accuracy under class imbalance. Per-class precision, recall and F1-score and a confusion matrix were computed from the pooled out-of-fold predictions of the grouped protocol. Gain-based feature importances were extracted from a model fitted to all labelled tiles to identify the texture descriptors that contributed most to the classification.

2.5.2 Unsupervised Learning

Since the nickel ore texture classes proposed by Barnes et al. (2017, 2018) describe a continuum rather than strictly discrete categories, an unsupervised multi-view clustering approach (Chao et al., 2021) was applied in parallel with the supervised classification. In this approach, the texture groups are formed directly from the data instead of from expert-assigned labels, and the two approaches can be compared with each other.

The clustering used the same feature set applied to the supervised classification. No class labels, sample identifiers or lithological information were used as input. The resulting 46 variables were organized into the six feature groups (views): mineral composition (20 variables), mineral diversity (2), general morphology (5), sulfide morphology (5), carbonate morphology (5) and sulfide

dissemination (9). Each view was scaled independently by a quantile transformation to a normal distribution, so that bounded, skewed or heavy-tailed variables contribute on a comparable scale.

Two multi-view strategies were tested, differing in the level at which the six views are combined. In hierarchical clustering with feature-level fusion, each scaled view was reduced to its first five principal components (or to all components where a view has fewer variables), and the reduced views were concatenated, giving each view approximately equal weight regardless of its number of variables; the concatenated matrix was clustered by agglomerative hierarchical clustering with Ward linkage (Ward, 1963). In multi-view spectral clustering with similarity-level fusion, a similarity matrix was computed for each view with a radial basis function kernel, with the bandwidth set to the median squared distance between tiles in that view; the six matrices were averaged and the result was partitioned by spectral clustering (von Luxburg, 2007) with k-means assignment of the spectral embedding. A fixed random seed was used throughout.

The number of clusters (k) was evaluated on the 2,975 labelled tiles, a subset small enough for both methods to be computed exactly. For k = 2 to 15, the silhouette coefficient (Rousseeuw, 1987), the Davies–Bouldin index (Davies & Bouldin, 1979) and the Calinski–Harabasz index (Caliński & Harabasz, 1974) were calculated. The reproducibility of the clustering was assessed by applying each method, for k = 3, 6, 9, 12 and 14, to eight random subsets of 80% of the tiles and measuring the agreement between every pair of solutions on their shared tiles with the adjusted Rand index (ARI; Hubert & Arabie, 1985). The final number of clusters was set to k = 14, equal to the number of texture classes, so that clusters and classes could be compared directly.

Because pairwise methods require an n × n similarity or distance matrix, which is impractical for 25,452 tiles, the final hierarchical clustering was fitted on 3,000 randomly selected landmark tiles (X. Chen & Cai, 2011). Each remaining tile was assigned to the cluster receiving the largest affinity-weighted vote among its 10 most similar landmarks, measured with the same per-view kernels averaged across views, and the share of the vote won by the assigned cluster was recorded as an assignment confidence. On a 4,000-tile test subset, fitting the hierarchical clustering on 2,000 landmarks and extending it to the remaining tiles reproduced the partition obtained from all tiles with an ARI of 0.93.

Clusters and texture classes were compared on the 2,975 labelled tiles by the ARI and the adjusted mutual information (AMI; Vinh et al., 2010), by contingency tables of clusters against classes, and by the accuracy and macro-averaged F1-score obtained after matching each cluster one-to-one to a texture class (Hungarian algorithm; Kuhn, 1955). The supervised reference was the out-of-fold prediction of the grouped protocol, so that both classifications are evaluated on tiles not used to define the corresponding classes.

3 Results

3.1 Supervised ore texture classification

On the Lappvattnet dataset, the random protocol gave a macro-averaged F1-score of 0.776 ± 0.014. When all tiles from a core sample were held out together, the macro-averaged F1-score decreased to 0.711 ± 0.027, and accuracy from 0.773 to 0.716 (Table 6). The grouped folds ranged from 0.665 to 0.766.

Table 6. Performance of the adopted XGBoost model on the Lappvattnet dataset of this study (14 classes) under the two evaluation protocols. Values are mean ± standard deviation over 25 random splits (random protocol) or 25 cross-validation folds (grouped protocol), with the range in brackets.

Protocol	Tiles	Accuracy	Macro-averaged F1-score	Weighted F1-score
Random 80/20 split, ≤ 200 tiles per class	2,685	0.773 ± 0.014 (0.741–0.801)	0.776 ± 0.014 (0.744–0.804)	0.771 ± 0.014 (0.738–0.800)
Grouped 5-fold cross-validation by core sample, all tiles	2,975	0.716 ± 0.023 (0.670–0.768)	0.711 ± 0.027 (0.665–0.766)	0.713 ± 0.024 (0.665–0.766)

Under the grouped protocol, F1-scores ranged from 0.45 to 0.96 (Table 7). Meta-sedimentary ball ore (0.96), meta-ultramafic semi-massive/massive ore (0.88), barren veins (0.86) and disseminated ore (0.84) were classified most reliably. The lowest scores were obtained for the non-mineralized meta-sedimentary classes, coarse layered (0.45) and coarse massive (0.49), which were mainly confused with each other and with fine layered tiles (Figure 4). Among mineralized meta-ultramafic classes, vein/veinlet tiles (0.53) were most often classified as net-textured (48 of 206 tiles), intermediate (33) or disseminated (20) ore. Confusions between the meta-ultramafic and meta-sedimentary domains were rare (< 1% of tiles).

Table 7. Per-class precision, recall and F1-score of the grouped protocol (pooled out-of-fold predictions, one repetition of 5-fold cross-validation, 2,975 tiles). The macro average therefore differs slightly from the mean over folds in Table 6.

Class	Tiles	Precision	Recall	F1-score
00 UM NM fine	220	0.69	0.66	0.67
01 UM NM coarse	219	0.71	0.68	0.70
02 UM NM barren veins	216	0.85	0.88	0.86
03 UM M vein/veinlets	206	0.59	0.48	0.53
04 UM M disseminated	296	0.84	0.84	0.84
05 UM M intermediate	228	0.76	0.79	0.78
06 UM M net-textured	200	0.68	0.81	0.74
07 UM M semi-massive/massive	85	0.87	0.88	0.88
08 MT NM fine layered	234	0.75	0.62	0.68
09 MT NM coarse layered	204	0.43	0.47	0.45
10 MT NM coarse massive	209	0.47	0.51	0.49

11 MT M vein/veinlets	213	0.73	0.76	0.74
12 MT M disseminated	215	0.71	0.71	0.71
13 MT M ball ore	230	0.96	0.97	0.96
Macro average	2,975	0.72	0.72	0.72

The confusion matrix (Figure 4) shows that misclassifications were almost entirely confined to the lithological domain of the true class: only 14 of the 2,975 tiles (0.5%) were assigned to a class of the other domain. When the 14 classes are collapsed into the four first-order groups of the classification key (meta-ultramafic or meta-sedimentary, mineralized or non-mineralized), 94% of the tiles were assigned to the correct group. In the meta-ultramafic domain, 40 tiles (2.4%) crossed the boundary between mineralized and non-mineralized classes, almost all of them between the non-mineralized classes and disseminated ore (21 non-mineralized tiles predicted as disseminated and 15 disseminated tiles predicted as non-mineralized), i.e. tiles close to the ~1% sulfide threshold. In the meta-sedimentary domain this boundary was less distinct, with 115 tiles (8.8%) crossing it, mainly coarse massive tiles predicted as disseminated (24) and mineralized vein/veinlet and disseminated tiles predicted as non-mineralized (51).

Within the mineralized meta-ultramafic classes, errors followed the sulfide-proportion sequence of the key. Among the four sulfide-proportion classes (disseminated, intermediate, net-textured and semi-massive/massive), errors occurred only between adjacent classes (e.g. 14 disseminated tiles as intermediate, 17 intermediate tiles as net-textured and 10 semi-massive/massive tiles as net-textured), and no tile was assigned to a class two or more steps away. The vein/veinlet class was the main exception: 108 of its 206 tiles (52%) were assigned to the sulfide-proportion classes, mostly net-textured (48), intermediate (33) and disseminated (20) ore, while 65 tiles of these classes were predicted as vein/veinlets. Among the non-mineralized meta-ultramafic classes, fine- and coarse-grained tiles were confused symmetrically (52 fine tiles predicted as coarse and 51 coarse tiles predicted as fine, ~24% of each class).

The non-mineralized meta-sedimentary classes formed the least separable group. Of their 647 tiles, 580 (90%) were predicted as one of the three non-mineralized meta-sedimentary classes but only 347 (54%) as the correct one. The largest confusions were between coarse layered and coarse massive tiles (62 and 58 tiles in each direction) and between fine layered and coarse layered tiles (45 and 31), placing the coarse layered class between its two neighbors. Among the mineralized meta-sedimentary classes, vein/veinlet and disseminated tiles were confused with each other (23 and 28 tiles), whereas ball ore remained distinct, with 222 of 230 tiles correctly classified and only 10 vein/veinlet tiles predicted as ball ore.

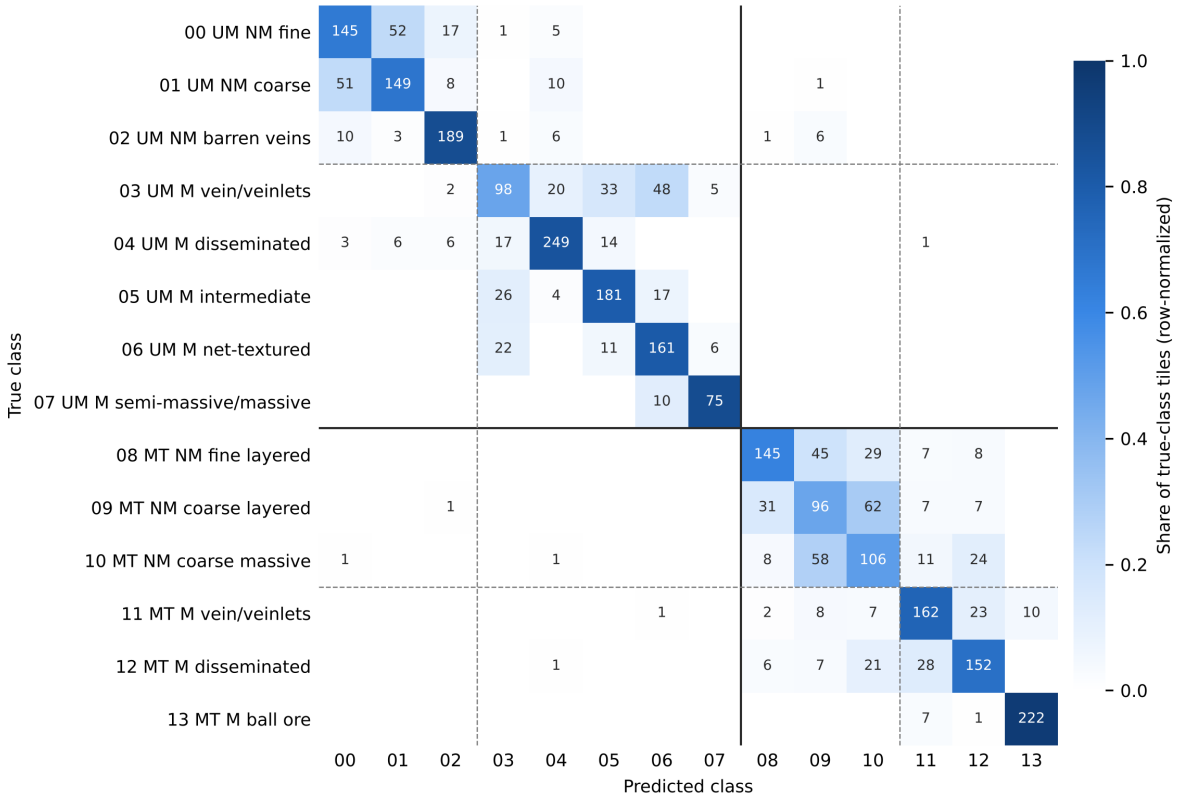

Figure 4. Confusion matrix of the adopted XGBoost model under the grouped protocol (pooled out-of-fold predictions of one repetition of 5-fold cross-validation grouped by core sample; 2,975 tiles). Numbers are tile counts; shading shows the fraction of each true class (row-normalized), so the diagonal corresponds to recall. Solid lines separate the meta-ultramafic (00–07) and meta-sedimentary (08–13) domains; dashed lines separate non-mineralized from mineralized classes. UM = meta-ultramafic, MT = meta-sedimentary, M = mineralized, NM = non-mineralized.

In terms of feature importance, mineral composition (38% of the total gain) and sulfide dissemination (28%) contributed most to the classification, followed by sulfide morphology (12%), carbonate morphology (12%) and general morphology (9%); mineral diversity contributed 2%. The most important individual features were the mean sulfide grain area (9.2%), mean carbonate line length (8.5%), the fraction of large sulfide grains (7.1%) and the plagioclase fraction (7.0%).

3.2 Unsupervised ore texture classification

Internal validity indices did not identify a preferred number of clusters. The silhouette coefficient of the multi-view spectral clustering remained below 0.17 for all k between 2 and 15, and that of the final hierarchical clustering was 0.21, indicating that the texture descriptors form a continuum rather than well-separated groups. Both methods were highly reproducible at low k (resampling ARI $\geq$ 0.94 at k = 3 and 6), but at higher k the hierarchical clustering was more stable than the multi-view spectral clustering (ARI of 0.90 $\pm$ 0.04 against 0.73 $\pm$ 0.07 at k = 14; Figure 5a) and agreed better with the texture classes at every k tested (Figure 5b). At k = 3, the hierarchical clusters mainly separated the host lithology (ARI of 0.80 with the meta-ultramafic/meta-sedimentary

division) rather than the textures; agreement with the 14 classes increased to 0.29 at k = 6–9 and was 0.26 at k = 14.

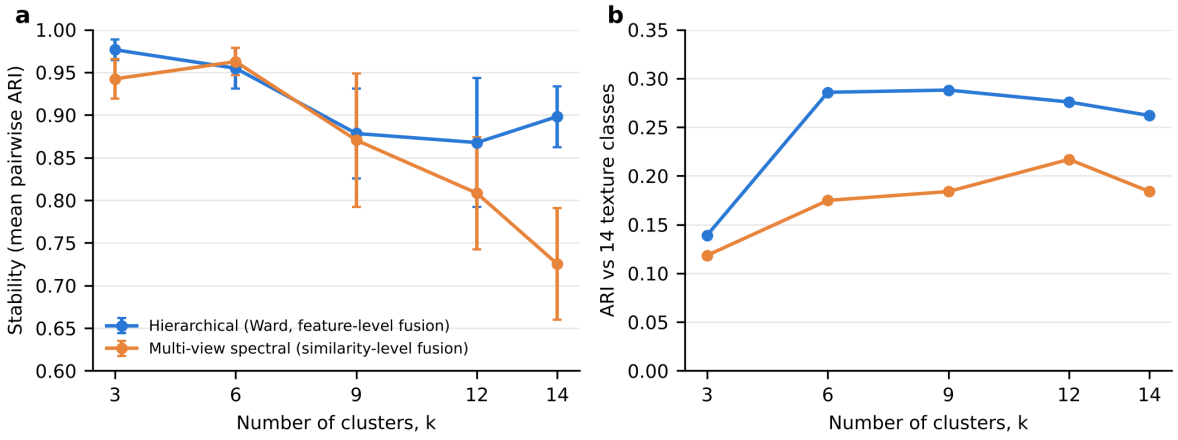

Figure 5. Selection of the number of clusters on the 2,975 labelled tiles. (a) Reproducibility of the clustering under resampling (mean ± standard deviation of the pairwise adjusted Rand index between eight solutions on random 80% subsets). (b) Agreement (adjusted Rand index) between the clusters and the 14 texture classes.

At k = 14, the hierarchical clustering of the labelled tiles gave an ARI of 0.26 and an AMI of 0.49 with the texture classes, against 0.18 and 0.34 for the multi-view spectral clustering and 0.54 and 0.67 for the supervised model (Table 8). After one-to-one matching of clusters to classes, 38% of the tiles were assigned to the correct class (macro-averaged F1-score of 0.35), compared with 72% for the supervised model under the grouped protocol.

Table 8. Agreement of the unsupervised clusters (k = 14) and of the supervised model with the 14 texture classes on the 2,975 labelled tiles. Matched accuracy and F1-score are obtained after one-to-one matching of clusters to classes; stability is the mean pairwise ARI under resampling.

Method	k	Tiles clustered	ARI	AMI	Matched accuracy	Matched macro F1	Stability (ARI)
Hierarchical, feature-level fusion	14	2,975	0.26	0.49	0.38	0.35	0.90 ± 0.04
Multi-view spectral, similarity-level fusion	14	2,975	0.18	0.34	0.33	0.31	0.73 ± 0.07
Hierarchical, all tiles via landmarks	14	25,452	0.26	0.46	0.37	0.33	–
Supervised XGBoost, grouped protocol	–	–	0.54	0.67	0.72	0.71	–

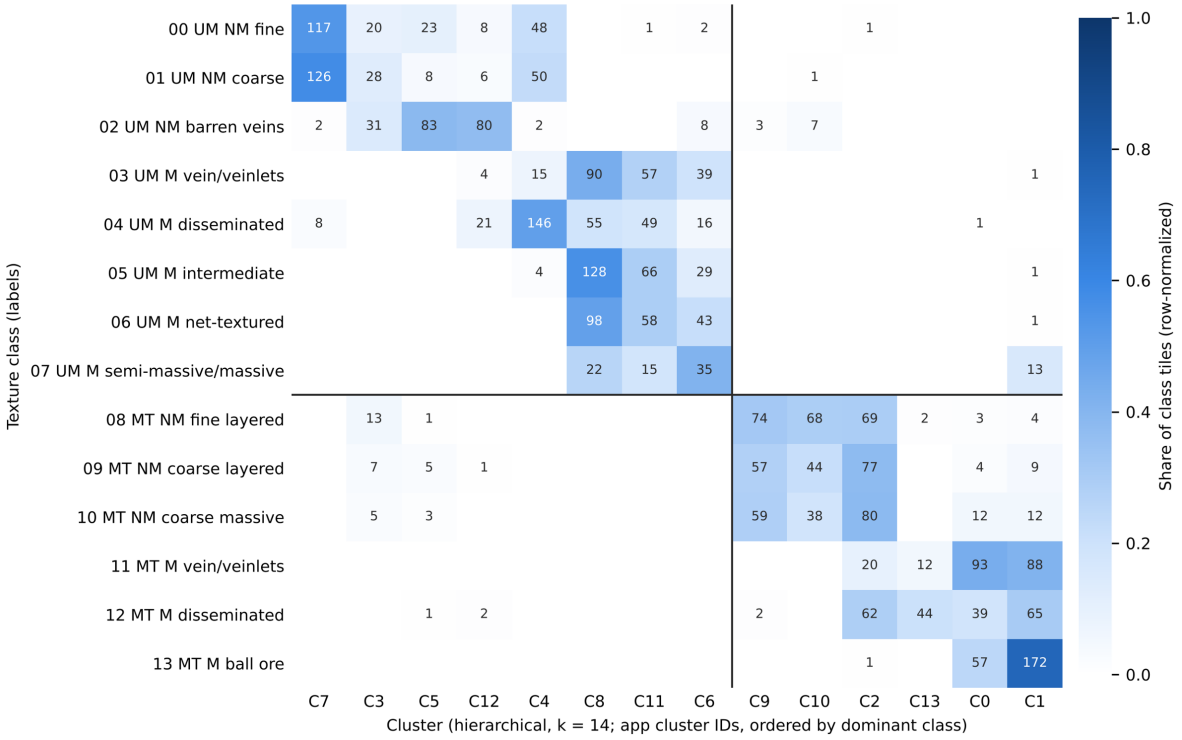

Figure 6. Contingency table of the 14 texture classes (rows) against the 14 hierarchical clusters (columns) for the 2,975 labelled tiles. Numbers are tile counts; shading shows the fraction of each class (row-normalized). Cluster numbers are arbitrary identifiers, ordered by their dominant class; the solid lines separate the meta-ultramafic and meta-sedimentary domains.

The clusters respected the first-order structure of the classification key (Figure 6). 97.7% of the labelled tiles fell in clusters dominated by their own lithological domain and 88.5% in clusters dominated by their own domain and mineralization group, against 99.5% and 94.3% of tiles correctly assigned by the supervised model. Within these groups, however, the clusters did not reproduce the textural classes. Non-mineralized meta-ultramafic tiles formed one cluster combining fine- and coarse-grained tiles (C7) and two clusters dominated by barren veins (C5 and C12), while part of the disseminated ore was grouped with non-mineralized tiles (C4). The intermediate, net-textured and vein/veinlet classes were divided between two clusters of nearly identical class composition (C8 and C11), and most semi-massive/massive tiles were assigned, together with net-textured and vein tiles, to C6. In the meta-sedimentary domain, the three non-mineralized classes were distributed almost evenly across three clusters (C2, C9 and C10); the mineralized tiles formed two clusters dominated by ball ore and veins (C1 and C0) and a small cluster dominated by disseminated ore (C13, 76%).

Applied to all 25,452 tiles, the landmark-based hierarchical clustering produced clusters of 337 to 4,049 tiles, with a mean assignment confidence of 0.95 and only 3.5% of the tiles assigned with a vote share below 0.6. Restricted to the labelled tiles, its agreement with the texture classes (ARI = 0.26) was the same as that of the clustering of the labelled tiles alone, showing that the extension to the full population did not degrade the result.

Figure 7 compares both classifications on a meta-sedimentary mineral map (LAP2106-001A). The supervised model separated the tiles crossed by sulfide veins (class 11) from disseminated (12) and non-mineralized layered tiles (08 and 09), following the sulfide veins visible in the mineral map. The clustering placed most of the vein-bearing tiles in C0 and C1 and the remaining tiles in C2, C9 and C10, with boundaries that follow the proportion of light felsic layers and biotite-rich bands rather than the supervised classes.

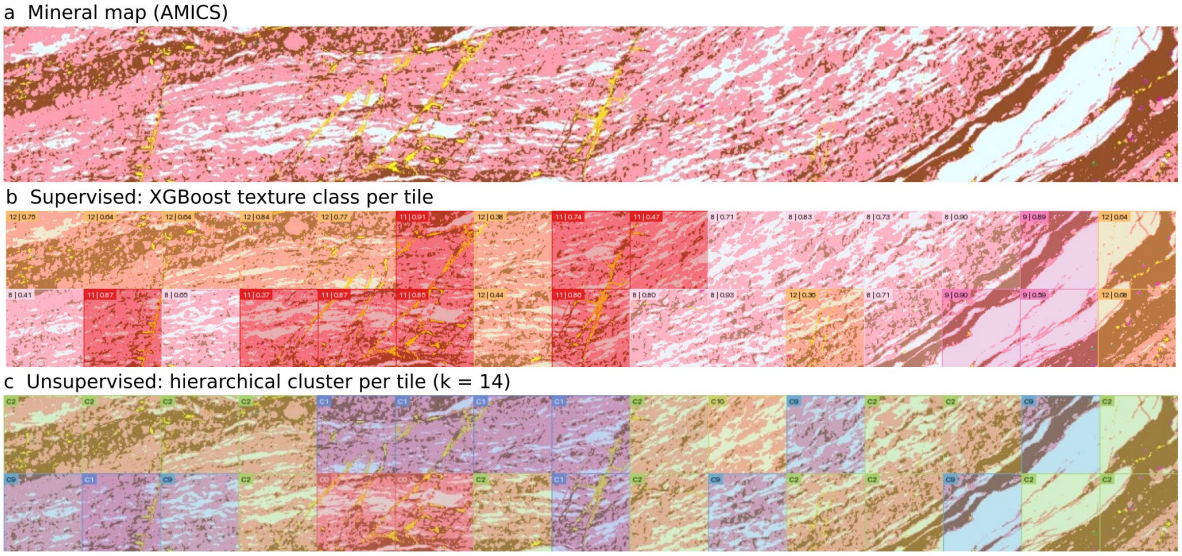

Figure 7. Classification of a meta-sedimentary mineral map (LAP2106-001A). (a) AMICS mineral map. (b) Texture class predicted by the XGBoost model for each 100 × 100 pixel tile (class number | prediction confidence). (c) Cluster assigned to each tile by the hierarchical clustering (k = 14; cluster identifiers as in Figure 6).

4 Discussion

4.1 Supervised Classification

Holding out complete core samples lowered the macro-averaged F1-score by about 6.5 percentage points (0.776 to 0.711). Tiles cut from the same µ-XRF map share mineralogy, grain size and fabric, so in a random split the test set contains close neighbors of the training tiles, and the score partly measures recognition of samples already seen (Roberts et al., 2017). The grouped estimate (macro-averaged F1-score of 0.71 ± 0.03) is therefore the more realistic measure of performance on new core samples, and it is the supervised reference against which the unsupervised classifications are compared. The spread between random splits (0.744–0.804; Table 6) also shows that single-split scores on datasets of this size should be reported with their variability rather than as point values.

Misclassifications were concentrated within lithological domains and between adjacent classes of the key. Meta-sedimentary non-mineralized tiles were confused across fine layered, coarse layered and coarse massive classes, which differ in grain size and foliation intensity rather than in mineralogy, and therefore in properties that the composition-dominated feature set captures only indirectly. In addition, several classes are defined by features that are similar to or larger than the 5 × 5 mm tile, such as grains >2 mm, layers 0.2–1 cm thick and porphyroblasts >1 cm. A single

tile therefore contains only part of a layer or a porphyroblast, and coarse layered and coarse massive tiles, or fine- and coarse-grained meta-ultramafic tiles, can look alike at this scale. This probably explains why these four classes obtained some of the lowest F1-scores (0.45–0.70), and larger tiles or features computed over neighboring tiles could improve their separation. Meta-ultramafic vein/veinlet tiles were mostly assigned to the sulfide-proportion classes (disseminated, intermediate, net-textured), consistent with a vein crossing a 5 × 5 mm tile occupying a similar sulfide proportion to these textures. These confusions follow the gradational nature of magmatic sulfide textures described by Barnes et al. (2017), where the boundaries between textural types are defined by thresholds on a continuum. Conversely, texturally distinctive classes such as ball ore, semi-massive/massive ore and barren veins were classified reliably. The dominance of composition and dissemination features in the feature importance indicates that the classification relies primarily on sulfide abundance and grain-size distribution, while sulfide vein geometry and foliation are less well represented; dedicated orientation or connectivity descriptors could improve the separation of vein and layered classes.

A macro-averaged F1-score of ~0.71 on unseen core samples, with errors largely confined to neighboring classes in the same lithological domain, indicates that the supervised model is suitable for domaining at the level of lithology and ore type, while the finest subdivisions (layer size, vein versus sulfide proportion) should be interpreted with caution.

4.2 Unsupervised Classification

The unsupervised classification shows which parts of the texture classification are expressed in the descriptor space on their own. Host lithology and the presence of sulfides, both defined by mineral composition, were recovered without labels: almost all clusters were confined to one lithological domain, and at k = 3 the clustering essentially reproduced the meta-ultramafic/meta-sedimentary division. The finer classes, in contrast, are defined by thresholds on continuous properties (sulfide proportion, grain size, layer thickness) or by geometric criteria (veins, foliation) that the 46 features capture only partially. Along such continua a distance-based method has no reason to place cluster boundaries at the project-defined thresholds (e.g. 10% or 30% sulfides), so the intermediate, net-textured and vein/veinlet tiles were divided into clusters that cut across the classes. This is consistent with the low silhouette values at all k and with the gradational nature of magmatic sulfide textures described by Barnes et al. (2017). The clusters should therefore be interpreted as data-driven textural domains rather than as a substitute for the expert-defined classes.

The two multi-view strategies differed in how the six views are weighted. Averaging one similarity matrix per view gives the two-variable diversity view and the carbonate morphology view, which is empty for most tiles, the same weight as mineral composition, diluting the compositional signal that carries most of the lithological and mineralization information. Feature-level fusion kept the views balanced by reducing each to the same number of components, and produced clusters that were more stable and more consistent with the texture classes.

4.3 Comparison between methods

Both approaches agree at the top of the classification hierarchy: lithology and mineralization were recovered by the supervised model (99.5% and 94.3% of tiles correctly assigned) and by the clustering (97.7% and 88.5% of tiles in clusters dominated by their own domain and mineralization group). They diverge below it. The supervised model reaches a macro-averaged F1-score of 0.71 on unseen core samples, with errors mostly between adjacent classes, whereas the clusters matched one-to-one to the classes reach 0.35. This difference indicates how much of the texture classification depends on expert-defined thresholds and geometric criteria rather than on structure inherent in the descriptors. The meta-sedimentary non-mineralized classes, which were poorly separated by both methods, are the clearest candidates for merging, or for redefinition using dedicated descriptors of layering and foliation.

For geometallurgical applications the two approaches are complementary. The supervised model provides classes with an established geological meaning, comparable with the Barnes et al. (2017, 2018) framework, but depends on labelled data and on the consistency of the labeler. The clustering requires no labels, was applied to all 25,452 tiles, and can be used to delineate textural domains, to check the consistency of the labels and, together with the underlying descriptor matrices, as a source of quantitative proxies to be linked with processing behavior.

5 Conclusions

On the labeled Lappvattnet dataset of 2,975 tiles and 14 texture classes, the model reached a macro-averaged F1-score of 0.78 ± 0.01 with random splits and 0.71 ± 0.03 when complete core samples were held out; the latter is the more realistic estimate for new core samples. Errors were concentrated between adjacent classes within the same lithological domain.

Mineral composition and sulfide dissemination features accounted for two thirds of the feature importance, whereas sulfide vein geometry and foliation were less well represented by the current descriptors.

Unsupervised multi-view clustering of the same 46 features recovered the lithological domains and the separation between mineralized and non-mineralized tiles, but not the finer textural subdivisions (ARI of 0.26 with the texture classes). Hierarchical clustering with feature-level fusion was more stable and more consistent with the classes than multi-view spectral clustering, and was extended to the full population of 25,452 tiles with a landmark-based approach.

Future work will include descriptors of vein geometry and foliation, larger tiles or neighborhood features for the classes defined at the centimeter scale, and will link the texture classes, the clusters and the underlying descriptor matrices to processing behavior, using modal mineralogy and liberation data and comminution tests.

References

Åkerman, C. (1987). Summary of results from nickel prospecting.

Barnes, S. J., Mungall, J. E., Le Vaillant, M., Godel, B., Lesher, C. M., Holwell, D., Lightfoot, P. C., Krivolutskaya, N., & Wei, B. (2017). Sulfide-silicate textures in magmatic Ni-Cu-PGE sulfide ore deposits: Disseminated and net-textured ores. *American Mineralogist*, *102*(3), 473–506. https://doi.org/10.2138/am-2017-5754

Barnes, S. J., Staude, S., Le Vaillant, M., Piña, R., & Lightfoot, P. C. (2018). Sulfide-silicate textures in magmatic Ni-Cu-PGE sulfide ore deposits: Massive, semi-massive and sulfide-matrix breccia ores. *Ore Geology Reviews*, *101*, 629–651. https://doi.org/10.1016/j.oregeorev.2018.08.011

Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Caliński, T., & Harabasz, J. (1974). A dendrite method for cluster analysis. *Communications in Statistics*, *3*(1), 1–27. https://doi.org/10.1080/03610927408827101

Chao, G., Sun, S., & Bi, J. (2021). A survey on multiview clustering. *IEEE Transactions on Artificial Intelligence*, *2*(2), 146–168. https://doi.org/10.1109/TAI.2021.3065894

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. https://doi.org/10.1145/2939672.2939785

Chen, W., Su, L., Chen, X., & Huang, Z. (2023). Rock image classification using deep residual neural network with transfer learning. *Frontiers in Earth Science*, *10*, Article 1079447. https://doi.org/10.3389/feart.2022.1079447

Chen, X., & Cai, D. (2011). Large scale spectral clustering with landmark-based representation. *Proceedings of the AAAI Conference on Artificial Intelligence*, *25*(1), 313–318. https://doi.org/10.1609/aaai.v25i1.7900

Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*(3), 273–297. https://doi.org/10.1007/BF00994018

Davies, D. L., & Bouldin, D. W. (1979). A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *PAMI-1*(2), 224–227. https://doi.org/10.1109/TPAMI.1979.4766909

Ekström, M. (1978). Mineralogisk undersökning av nickelmineraliseringen Lappvattnet. I. Mineralisering i ultrabasit och intilliggande gnejs.

Fandrich, R., Gu, Y., Burrows, D., & Moeller, K. (2007). Modern SEM-based mineral liberation analysis. *International Journal of Mineral Processing*, *84*(1–4), 310–320. https://doi.org/10.1016/j.minpro.2006.07.018

Ferreira, I. E. (2023). *No labels needed: Self-supervised semantic segmentation of ore minerals in thin section* [Preprint]. SSRN. https://ssrn.com/abstract=4466357

Filippo, M. P., da Fonseca Martins Gomes, O., da Costa, G. A. O. P., & Mota, G. L. A. (2021). Deep learning semantic segmentation of opaque and non-opaque minerals from epoxy resin in reflected light microscopy images. *Minerals Engineering*, *170*, Article 107007. https://doi.org/10.1016/j.mineng.2021.107007

Gottlieb, P., Wilkie, G., Sutherland, D., Ho-Tun, E., Suthers, S., Perera, K., Jenkins, B., Spencer, S., Butcher, A., & Rayner, J. (2000). Using quantitative electron microscopy for process mineralogy applications. *JOM*, *52*(4), 24–25. https://doi.org/10.1007/s11837-000-0126-9

Gu, X., Cook, N. J., Metcalfe, A. V., & Aldrich, C. (2024). A machine vision approach for detecting changes in drill core textures using optical images. *Mining Technology: Transactions of the Institutions of Mining and Metallurgy*, *133*(2), 162–175. https://doi.org/10.1177/25726668241243265

Guntoro, P. I., Ghorbani, Y., Butcher, A. R., Kuva, J., & Rosenkranz, J. (2020). Textural quantification and classification of drill cores for geometallurgy: Moving toward 3D with X-ray microcomputed tomography (μCT). *Natural Resources Research*, *29*(6), 3547–3565. https://doi.org/10.1007/s11053-020-09685-5

Hrstka, T., Gottlieb, P., Skála, R., Breiter, K., & Motl, D. (2018). Automated mineralogy and petrology - applications of TESCAN Integrated Mineral Analyzer (TIMA). *Journal of Geosciences*, *63*(1), 47–63. https://doi.org/10.3190/jgeosci.250

Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, *2*(1), 193–218. https://doi.org/10.1007/BF01908075

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems 30*, 3146–3154.

Koch, P. H., Lund, C., & Rosenkranz, J. (2019). Automated drill core mineralogical characterization method for texture classification and modal mineralogy estimation for geometallurgy. *Minerals Engineering*, *136*, 99–109. https://doi.org/10.1016/j.mineng.2019.03.008

Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, *2*(1–2), 83–97. https://doi.org/10.1002/nav.3800020109

Lund, C., Lamberg, P., & Lindberg, T. (2015). Development of a geometallurgical framework to quantify mineral textures for process prediction. *Minerals Engineering*, *82*, 61–77. https://doi.org/10.1016/j.mineng.2015.04.004

Maitre, J., Bouchard, K., & Bédard, L. P. (2019). Mineral grains recognition using computer vision and machine learning. *Computers and Geosciences*, *130*, 84–93. https://doi.org/10.1016/j.cageo.2019.05.009

Menzies, A. H., Tagle, R., Reinhardt, F., Hirschle, C., Schellkopf, L. J., & Kelly, N. (2022). High-speed micro-XRF analysis of rock samples and drill cores. *Microscopy and Microanalysis*, *28*(S1), 650–651. https://doi.org/10.1017/S1431927622003117

Merrill-Cifuentes, J., Cracknell, M. J., & Escolme, A. (2022). Unsupervised textural classification of rocks in large imagery datasets. *Minerals Engineering*, *180*, Article 107496. https://doi.org/10.1016/j.mineng.2022.107496

Morales, J., Saldivia, C., Lobo, R., del Pino, M., Muñoz, M., Caniggia, G., Catalán, J., Hayde, R., & Poblete, V. (2025). Visual analysis of deep learning semantic segmentation applied to petrographic thin sections. *Scientific Reports*, *15*(1). https://doi.org/10.1038/s41598-025-99767-2

Pérez-Barnuevo, L., Lévesque, S., & Bazin, C. (2018a). Automated recognition of drill core textures: A geometallurgical tool for mineral processing prediction. *Minerals Engineering*, *118*, 87–96. https://doi.org/10.1016/j.mineng.2017.12.015

Pérez-Barnuevo, L., Lévesque, S., & Bazin, C. (2018b). Drill core texture as geometallurgical indicator for the Mont-Wright iron ore deposit (Quebec, Canada). *Minerals Engineering*, *122*, 130–141. https://doi.org/10.1016/j.mineng.2018.03.020

Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, *1*(1), 81–106. https://doi.org/10.1023/A:1022643204877

Reddick, J., & Lindholm, T. (2020). *Technical report on the Lappvattnet and Rormyrberget deposits, northern Sweden (NI 43-101 Technical Report). Gungnir Resources Inc*.

Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F., & Dormann, C. F. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, *40*(8), 913–929. https://doi.org/10.1111/ecog.02881

Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, *20*, 53–65. https://doi.org/10.1016/0377-0427(87)90125-7

Rubo, R. A., de Carvalho Carneiro, C., Michelon, M. F., & Gioria, R. dos S. (2019). Digital petrography: Mineralogy and porosity identification using machine learning algorithms in petrographic thin section images. *Journal of Petroleum Science and Engineering*, *183*, Article 106382. https://doi.org/10.1016/j.petrol.2019.106382

Seo, W., Kim, Y., Sim, H., Song, Y., & Yun, T. S. (2022). Classification of igneous rocks from petrographic thin section images using convolutional neural network. *Earth Science Informatics*, *15*(2), 1297–1307. https://doi.org/10.1007/s12145-022-00808-5

Soler-Aragon, N. K. (2026). *Ore textural characterization of Ni-Cu-PGE deposits using machine learning methods: Process relevant mineralogical and textural information* [Master's thesis, Luleå University of Technology].

Tiu, G. (2017). *Classification of drill core textures for process simulation in geometallurgy* [Master's thesis, Luleå University of Technology].

Vinh, N. X., Epps, J., & Bailey, J. (2010). Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *Journal of Machine Learning Research*, *11*, 2837–2854.

von Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and Computing*, *17*(4), 395–416. https://doi.org/10.1007/s11222-007-9033-z

Ward, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, *58*(301), 236–244.https://doi.org/10.1080/01621459.1963.10500845

Xu, Z., Ma, W., Lin, P., Shi, H., Pan, D., & Liu, T. (2021). Deep learning of rock images for intelligent lithology identification. *Computers and Geosciences*, *154*, Article 104799. https://doi.org/10.1016/j.cageo.2021.104799

Zhang, J., Liu, M., Guo, J., Wu, D., Wang, M., & Zheng, S. (2024). SHN: Rock image classification and feature visualization using multiple granularity spatial disorder hierarchical residual network. *Frontiers in Earth Science*, *12*, Article 1364209. https://doi.org/10.3389/feart.2024.1364209