MINERALS IN THE WILD: A HYPERSPECTRAL–XRF DATASET FOR ELEMENTAL COMPOSITION ESTIMATION

Eleftheria Tetoula-Tsonga[1], *George Arvanitakis*[2], *Theodoros Giannakas*[1]

[1]Institute of Communication and Computer Systems, Athens, Greece
[2]Geonova, Athens, Greece

ABSTRACT

Rapid mineral characterization is essential for applications ranging from mineral exploration to industrial ore processing. To this end, Hyperspectral Imaging (HSI) has emerged as a promising sensing modality thanks to its fine spectral resolution, enabling mineral discrimination in both close-range and remote sensing settings. However, the scarcity of publicly available datasets with reliable ground-truth labels hinders the development and evaluation of HSI-based mineral identification methods. We release *Minerals in the Wild*, a multi-purpose dataset comprising 1,132 rock specimens collected across Europe [1]. For each specimen, we provide an HSI acquisition together with an elemental characterization obtained via an XRF sensor. We define the task of elemental characterization on our dataset and propose a pruning mechanism that removes distant signatures from the USGS dictionary prior to a convex optimization approach for matching HSI pixels with USGS spectral signatures. Finally, we empirically show that our approach outperforms simpler baselines.

Index Terms— Mineral Exploration, Hyperspectral Imaging, XRF, Spectral Unmixing, Datasets and Benchmarks

1. INTRODUCTION

1.1. Motivation

The access to critical raw materials is becoming a bottleneck for key technologies–from AI and advanced manufacturing to electrification and clean energy. For example, AI infrastructure relies heavily on mineral-intensive semiconductors and data centers; while clean-energy deployment depends on lithium, cobalt, nickel, copper, graphite, and rare earth elements for batteries and grids [2]. Traditionally, to *explore* the mineral content of a new area, one needed to set up costly field campaigns, often with highly uncertain reward. Beyond exploration, modern mining workflows increasingly require *automated sorting* of raw material, under stringent latency requirements (often in the order of seconds), rendering the use of manual high-precision equipment impractical [3].

This work has received funding from the European Union under Grant Agreement No. 101178897 (DEXPLORE).

To address these challenges *at scale*, the community is moving toward efficient, yet approximate, mineralogical inference using image-based sensing. In particular, the imagery employed in this setting is typically of high spectral resolution, commonly referred to as hyperspectral imaging (HSI), due to the fact that many minerals exhibit informative features, e.g., absorption bands, at wavelengths not captured by conventional RGB. Nonetheless, importantly, the camera on its own is by no means enough; it has to be paired with an identification algorithm capable of inferring the presence of minerals in the image. To devise such algorithms, we typically have to associate two (or potentially more) modalities: (a) hyperspectral image, and (b) matched mineralogical or elemental characterization. While acquiring modality (a) is relatively fast and scalable, acquiring modality (b) requires time-consuming measurements with specialized equipment. Taking the above into account, the importance of such paired datasets emerges naturally, as it enables practitioners and researchers to build a variety of applications related to remote sensing and automated material sorting.

1.2. Related Work

Datasets. Hyperspectral mineral datasets span satellite, airborne, and laboratory acquisition regimes. Satellite (e.g., EnMAP) and airborne (e.g., Cuprite/HYDICE) datasets primarily support regional-scale mineral mapping and exploration [4–6]. Laboratory datasets are closer to our setting; for example, Koirala *et al.* introduced a multisensor benchmark based on controlled mineral powder mixtures [7]. In contrast, our dataset comprises intact natural rock specimens, pairing HSI cubes with XRF-derived elemental compositions to support specimen-level HSI-to-XRF prediction.

Unmixing. The problem of pixel unmixing has been thoroughly studied in the context of satellite imaging. Classical approaches include constrained least-squares (sparse) formulations [8], with them either decomposing pixels to representative in-image endmembers or using reference spectra (e.g., USGS libraries) [9]. More recent approaches incorporate nonlinear modeling and spectral variability [10, 11]. Our pixel unmixing pipeline builds on [9] by introducing a novel library-reduction filter for improved fit during optimization.

1.3. Contributions and Outline

(C.1) We release *Minerals in the Wild*, a new multi-purpose dataset of 1,132 solid hand-sized rock specimens collected across Europe [1]. For each specimen we store: a) HSI, and b) XRF characterization, i.e., the specimen's composition in the form of percentages over chemical elements (Section 2).
(C.2) We define a case study on the dataset; in particular, estimating the elemental composition of a specimen using its HSI cube. To this end, we introduce an efficient algorithm, based on convex optimization, paired with a novel (regularizing) filter. We predict elemental composition and validate the estimates against XRF measurements (Section 3).

2. A HYPERSPECTRAL-XRF DATASET

Physically, *Minerals in the Wild* comprises a collection $\mathcal{S}$ of $S = 1,132$ hand-sized specimens gathered from multiple mines and exploration areas across Europe. Our first contribution is to convert this collection into a structured dataset. Below, we detail the acquisition protocol for the HSI and the XRF sensing modalities, as well as the data representation of each specimen in the dataset. Finally, we report dataset- and sample-level insights. In what follows, tensors of arbitrary size are denoted with bold. For example, $\mathbf{X} \in \mathbb{R}^{A\times B}$ is a 2D matrix, while $\mathbf{X}(i,j)$ is its scalar value at index (i,j).

2.1. Hyperspectral Modality: Specimens to HSI Cubes

Step (1): Scan. The setup consists of a conveyor belt and a Specim SWIR hyperspectral camera [12], mounted approximately 30 cm above the belt and oriented towards it. The spatial content of specimen s is acquired as follows. At each timestep, the push-broom sensor captures a spatial line of width $W = 384$ pixels. As s passes under the sensor, successive lines being accumulated result in an image of size $H_s \times W$, where H_s denotes the number of acquired lines, and depends on its size. Spectrally, for each pixel the camera stores $B = 273$ values, where B is the number of bands, in the 1000–2500 nm wavelength range[1]. At the end of Step (1), the data object stored for specimen s is denoted as $\mathbf{X}_s^{(1)} \in \mathbb{R}^{H_s\times W\times B}$.
Step (2): Calibrate. Using dark- and white-references, denoted by $\mathbf{I}_{\text{Dark}}$, $\mathbf{I}_{\text{White}} \in \mathbb{R}^{W\times B}$, respectively, we calibrate $\mathbf{X}_s^{(1)}$. These references vary across detector position and wavelength, but remain constant along the scan dimension. Specifically we compute:

$$\mathbf{X}_s^{(2)}(i,j,b) = \frac{\mathbf{X}_s^{(1)}(i,j,b) - \mathbf{I}_{\text{Dark}}(j,b)}{\mathbf{I}_{\text{White}}(j,b) - \mathbf{I}_{\text{Dark}}(j,b)}, \quad \forall (i,j,b). \quad (1)$$

Step (3): Mask-out. We manually create a binary spatial mask $\mathbf{M}_s \in \{0,1\}^{H_s\times W}$ for specimen s, where $\mathbf{M}_s(i,j) = 1$ indicates that pixel (i,j) visually belongs to s, and $\mathbf{M}_s(i,j) = 0$ otherwise. To obtain the final representation, denoted by $\mathbf{X}_s^{(3)} \in \mathbb{R}^{H_s\times W\times B}$, we apply the mask as follows:

$$\mathbf{X}_s^{(3)}(i,j,b) = \begin{cases} \mathbf{X}_s^{(2)}(i,j,b), & \text{if } \mathbf{M}_s(i,j) = 1, \\ \texttt{NaN}, & \text{if } \mathbf{M}_s(i,j) = 0, \end{cases} \quad \forall (i,j,b). \quad (2)$$

Effectively, $\mathbf{X}_s^{(3)}$ is the HSI cube of specimen s with the background filled with `NaN` across all B bands. For the rest of this work, we drop the step superscript and $\mathbf{X}_s \equiv \mathbf{X}_s^{(3)}\ \forall\ s \in \mathcal{S}$. For each specimen s, let Ω_s denote the set of valid pixel locations after masking; we enumerate these by $p = 1,\ldots,|\Omega_s|$, and denote the corresponding spectra by $\mathbf{x}_{s,p}$.

2.2. XRF Modality: Specimens to Compositions

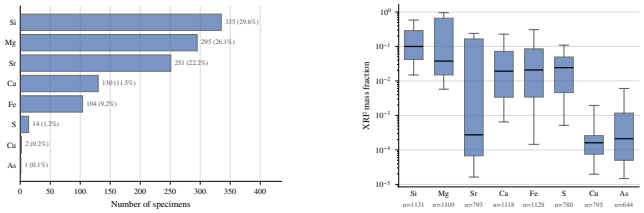
Step (1): Scan. To characterize specimen s, we use X-ray fluorescence (XRF) [14] by performing $D = 5$ point measurements at visually different locations of the surface, each with an 8 mm analysis spot. Each measurement returns a vector of N normalized compositions over a mixed set of column labels, including elements (e.g., Mn) and compounds (e.g., MgO). The resulting matrix is denoted by $\mathbf{y}_s^{(1)} \in \mathbb{R}^{D\times N}$, where each row corresponds to one location.
Step (2): Average. To obtain a specimen-level composition, we aggregate by averaging across the D measurements, which yields vector $\mathbf{y}_s^{(2)} \in \mathbb{R}^N$.
Step (3): Harmonize. This step turns the mixed (elements and compounds) label space to element-only space. Compound compositions are decomposed into their elemental contributions using stoichiometric mass fractions[2]. The contributions from all sources are then accumulated at element level. The resulting vector is filtered by removing light elements (e.g., O and C) and renormalized to obtain a valid elemental composition (summing to 1). Thus, compound coordinates (e.g., MgO) are replaced by their corresponding elemental coordinates (e.g., Mg), while the vector dimensionality is preserved. The final vector is denoted by $\mathbf{y}_s^{(3)} \in \mathbb{R}^N$ ($N = 51$ in our case). For the rest of this work, $\mathbf{y}_s \equiv \mathbf{y}_s^{(3)}\ \forall\ s \in \mathcal{S}$.

2.3. Discussion and Insights

With the dataset $\mathcal{D} = \{\mathbf{X}_s, \mathbf{y}_s\}_{s\in\mathcal{S}}$ in place, we attempt to export interesting insights, using the two modalities.
Elemental diversity (XRF). To characterize the elemental diversity of $\mathcal{S}$, we focus on our XRF vectors. In Fig. 1(a), we see a histogram showing the number of specimens in which an element dominates. For example, around 30% of the specimens have Si as their top element. We observe that the absence of a single dominant element suggests that $\mathcal{D}$

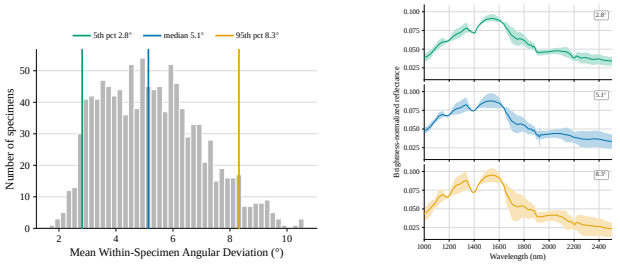
[1] This range is known to contain diagnostic absorption features that facilitate mineral discrimination [13].

[2] For example, suppose Mn=0.9 and MgO=0.1. The MgO splits to: Mg as $24.3/(24.3{+}16.0)$ and O as $16.0/(24.3{+}16.0)$. These are weighted by 0.1 to yield Mn=0.9, Mg=0.06, and O=0.04. After removing O and renormalizing, the final representation becomes Mn=0.937 and Mg=0.063.



(a) Dominant elements (b) Elemental mass fractions

Fig. 1: XRF diversity: (a) Number of specimens for which an element is dominant. (b) Distribution of non-zero elemental mass fractions across specimens.



(a) (b)

Fig. 2: Pixel-level spectral uniformity: (a) Distribution of the Mean Within-Specimen Angular Deviation (MWAD). (b) Mean spectra with corresponding spectral envelopes for representative specimens.

spans a broad range of compositions, providing a challenging benchmark for evaluating hyperspectral unmixing methods. In addition, Fig. 1(b) shows a boxplot (in logarithmic scale) of the non-zero observed fractions across the S specimens, strengthening the case that our data exhibit interesting variability when it comes to the observed XRF fractions.

Spectral variability (HSI). Since the camera acquires complete HSI cubes, we examine spectral the variability of s using angular differences between its spectra; hence, we define the Mean Within-Specimen Angular Deviation (MWAD) of s:

$$\text{MWAD}_s = \frac{1}{|\Omega_s|} \sum_{p \in \Omega_s} \theta(\mathbf{x}_{s,p}, \bar{\mathbf{x}}_s) \tag{3}$$

where $\bar{\mathbf{x}}_s$ is the mean spectrum of specimen s, and $\theta(\cdot, \cdot)$ denotes the angular deviation between two spectra. Figure 2(a) summarizes its distribution, while Fig. 2(b) shows representative $\bar{\mathbf{x}}_s$ together with their envelopes. Overall, the within-specimen variation is relatively small, but as it gets larger, the spectra seem to fluctuate more around the mean.

Cross-modal structure (HSI–XRF). We investigate how both modalities behave jointly. In Fig. 3(a), we scatterplot the S XRF vectors in 2D, and cluster them in 5 groups (see 5 groups of colored dots). Keeping the color of each cluster, in Fig. 3(b) we scatterplot all the $\mathbf{X}_s$. We can see that there is in fact some underlying structure as $\mathbf{X}_s$ are somewhat similarly (not trivially though) clustered as their XRF vectors, suggesting that learning a function that differentiates them is an interesting avenue of research. Moreover, in Fig. 3(c), we investigate if there is correlation between: (i) (x-axis) the XRF vector entropy, defined as $H(\mathbf{y}_s) = -\sum_{j=1}^{N} y_{s,j} \log(y_{s,j})$, which quantifies whether the composition is concentrated around few elements (low H) or is more uniformly distributed across many (high H), and (ii) (y-axis) spectral variability, measured using MWAD_s (defined in (3). We can observe that although there is a positive trend, it is not definitive; however, one can argue that when XRF is more "random", pixels exhibit increased variability (upper right part of figure).

3. CASE STUDY: CHARACTERIZING SPECIMENS FROM HYPERSPECTRAL IMAGERY

In this section, we define the composition estimation task on the *Minerals in the Wild* dataset, along with two metrics to assess the performance of an algorithm on the task. We will propose a solution and compare it against simpler baselines.

3.1. Task Formulation

Task. We are interested in a function $f(\cdot)$ that receives as input $\mathbf{X}_s \in \mathbb{R}^{H_s \times W \times B}$, and computes its elemental composition:

$$f(\cdot) : \mathbb{R}^{H \times W \times B} \to \mathbb{R}^N, \text{ with } f(\mathbf{X}_s) = \hat{\mathbf{y}}_s, \tag{4}$$

where $\hat{\mathbf{y}}_s$ is the elemental composition of s over N elements.
Metrics. We evaluate the performance of $f(\cdot)$ by comparing the predicted composition $\hat{\mathbf{y}}_s$ against the XRF-derived ground truth $\mathbf{y}_s$ using the L2 loss, which measures compositional error (accounting for direction and magnitude), and Cos, which measures cosine similarity and hence agreement in direction.

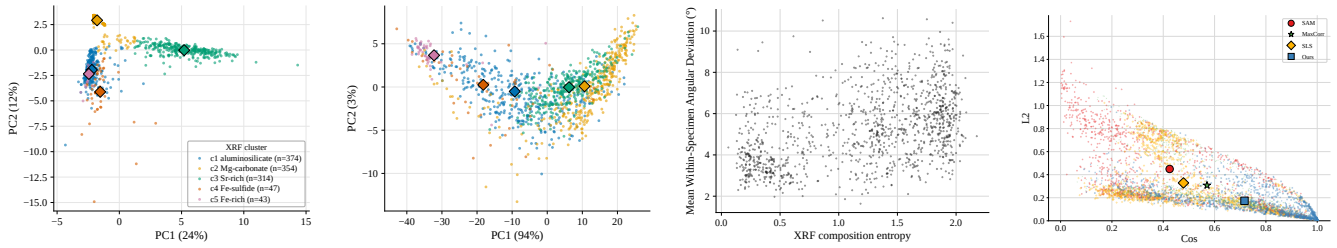
$$\boxed{\text{L2}(\hat{\mathbf{y}}_s, \mathbf{y}_s) = \|\hat{\mathbf{y}}_s - \mathbf{y}_s\|_2^2, \ \text{Cos}(\hat{\mathbf{y}}_s, \mathbf{y}_s) = \frac{\hat{\mathbf{y}}_s^\top \mathbf{y}_s}{\|\hat{\mathbf{y}}_s\|_2 \|\mathbf{y}_s\|_2}} \tag{5}$$

Preliminaries. Throughout this section, matrix $\mathbf{U} \in \mathbb{R}^{L \times B}$ encodes the L USGS mineral signatures [15], resampled at the B bands of our camera[3]. The input $\mathbf{X}$ and USGS library $\mathbf{U}$ are individually preprocessed using continuum removal [17]. To simplify notation, we denote the resulting representations also as $\mathbf{X}$ and $\mathbf{U}$. In what follows, we drop subscript s.

3.2. Proposed Pipeline

The pipeline consists of three stages, and is applied per specimen. *First*, the specimen is decomposed into its pixels and, for each pixel, we estimate its composition under the USGS spectral basis, i.e., as a combination of mineral signatures

[3]The original USGS signatures are sampled at a finer spectral resolution than our camera, so we spectrally resample them to our camera bands, following the Gaussian spectral response function approach [16].



(a) (b) (c) (d)

Fig. 3: **(a)** PCA projection of XRF composition vectors, grouped into five chemically coherent clusters. **(b)** PCA projection of the mean hyperspectral spectrum of each specimen, coloured according to the XRF-derived clusters in (a) (diamonds indicate cluster centroids of (a)). **(c)** MWAD and XRF composition entropy scatterplot. **(d)** Performance of our method and the baselines across all specimens, evaluated using the XRF measurements.

from the USGS library. *Second*, the pixel-level USGS compositions are spatially aggregated to get a specimen-level composition under the USGS basis. *Third*, the specimen-level USGS composition is transformed to elemental composition[4], enabling direct comparison with the XRF measurements.

Pixel-level USGS compositions. This is two procedures:

1) Candidates. For each pixel $p \in \Omega$, we construct a reduced USGS library $\mathbf{U}_p = g(\mathbf{U}, \mathbf{x}_p, t)$ by retaining signatures whose first- and second-order derivatives match those of $\mathbf{x}_p$. Specifically, a library signature is selected if the cosine similarities between the corresponding derivatives exceed the threshold t (excluding pairs with both negative similarities). The resulting library $\mathbf{U}_p \in \mathbb{R}^{L_p \times B}$ contains the L_p candidate signatures used in the optimization.

2) Optimization. The pixel-level USGS composition of specimen at pixel p is represented by $\mathbf{z}_p^\star \in \mathbb{R}^{L_p}$, whose entries correspond to the weights assigned to the L_p candidates; it is obtained by solving the convex optimization problem:

$$\textbf{(OP)} \qquad \mathbf{z}_p^\star = \arg\min_{\mathbf{z}\in\mathbb{R}^{L_p}} \left\|\mathbf{x}_p - \mathbf{z}^\top \mathbf{U}_p\right\|_2^2 \qquad (6)$$
$$\text{subject to} \quad \mathbf{z} \geq \mathbf{0},\ \mathbf{1}^\top \mathbf{z} = 1.$$

Then, $\mathbf{z}_p^\star$ is embedded into the full USGS basis by assigning zero to all signatures excluded during filtering. By a slight abuse of notation, we denote this vector also by $\mathbf{z}_p^\star \in \mathbb{R}^L$.

Specimen-level USGS composition. We proceed to spatial averaging of $\mathbf{z}_p^\star$ across all pixels: $\mathbf{z} = \frac{1}{|\Omega|}\sum_{p\in\Omega} \mathbf{z}_p^\star$, where $\mathbf{z} \in \mathbb{R}^L$ is a vector showing the weights of the specimen across the L signatures, which sums to 1.

Specimen-level elemental composition. The $\mathbf{z} \in \mathbb{R}^L$ represents composition over the USGS basis, whereas the XRF ground truth $\mathbf{y}$ is over the elemental space, $\mathbb{R}^N$. So, we apply the harmonization procedure of Section 2.2 to convert USGS mineral composition into elemental. The resulting elemental vector is then projected onto the XRF elemental basis, discarding elements outside the basis and renormalizing, which yields the predicted elemental composition $\hat{\mathbf{y}} \in \mathbb{R}^N$.

[4]USGS contains compounds; we treat them exactly as the compounds appearing in the XRF exported vectors.

3.3. Experimental Evaluation

We first check the effect of threshold t to the performance of our algorithm and then compare it against simpler baselines.

Effect of the filtering threshold t. This parameter controls the strictness of the candidate-selection. For each value of t, we execute the complete pipeline (Section 3.2) and evaluate the resulting specimen-level compositions using the L2 and Cos metrics, as defined in (5). We additionally report the *coverage*, i.e., the percentage of specimens with at least one candidate USGS signature. Table 1 shows the trade-off induced by t. Small values retain many candidates, increasing ambiguity during the optimization, whereas large values discard plausible signatures, eventually leaving specimens with no candidates. Consequently, coverage decreases with t. We use $t = 0.10$ in the remainder, as it preserves full coverage while providing a balance between Cos and L2.

Performance comparison. All baselines are instantiated within the proposed pipeline to enable specimen-level elemental prediction and validation against the XRF. They differ only in the pixel-level USGS composition estimation step. We used the CR versions of HSI and USGS spectra throughout.

- *SAM.* Each pixel is assigned to the USGS signature with the minimum spectral angle.
- *MaxCorr.* Each pixel is assigned to the USGS signature with the maximum Pearson correlation.
- *SLS (simplex-constrained least squares).* For each pixel, we solve **(OP)** (6) over the complete USGS, i.e., without the proposed candidate filtering.

To compare our method against the baselines, Fig. 3(d) visualizes the specimen-level performance in 2D plane, where a point corresponds to $(\text{Cos}, \text{L2})_s^m$ for specimen s and method $m \in \{\text{Ours}, \text{SAM}, \text{MaxCorr}, \text{SLS}\}$. Hence, each method is represented by S points, one per specimen. Ideally, we would like the specimens to concentrate in the lower right; and in fact, our algorithm tends to lie at that area more frequently compared to the other baselines. Moreover, we present Table 2 which shows our method achieving the best performance across both Cos and L2. Compared with *MaxCorr*, our method improves the median Cos from 0.577 to

Table 1: Threshold selection of our method.

	Cos ↑			L2 ↓			
t	Mean	Std	Med.	Mean	Std	Med.	Cov.
0.05	0.659	0.201	0.680	0.202	0.148	0.155	100%
0.10	**0.717**	0.234	0.773	**0.173**	0.135	0.150	**100%**
0.15	0.711	0.270	0.809	0.215	0.220	0.166	97.1%
0.20	0.746	0.257	0.838	0.214	0.223	0.168	84.7%
0.30	0.767	0.255	0.861	0.226	0.223	0.163	61.9%
0.40	0.749	0.264	0.853	0.276	0.255	0.181	36.7%
0.50	0.774	0.244	0.885	0.278	0.264	0.177	17.0%
0.60	0.844	0.209	0.929	0.205	0.193	0.155	6.6%

Table 2: Performance/accuracy and runtime, for all methods.

	Cos ↑			L2 ↓			Time ↓
Method	Mean	Std	Med.	Mean	Std	Med.	s/sample
SAM	0.425	0.278	0.344	0.449	0.316	0.348	**3.83**
SLS	0.479	0.214	0.421	0.329	0.232	0.236	33.710
MaxCorr	0.571	0.211	0.577	0.308	0.227	0.221	3.753
Ours	**0.717**	0.234	**0.773**	**0.173**	0.135	**0.150**	5.521

0.773 and reduces the median L2 from 0.221 to 0.150. This suggests that restricting the optimization to pixel-dependent candidate signatures helps mitigate ambiguities introduced by the full library. Finally, in the last column, we present the average inference time needed per specimen; we can observe that our method is 6× faster compared to the vanilla *SLS* which optimizes using the whole USGS.

4. CONCLUSIONS

We introduced a new dataset of naturally occurring mineral specimens, providing paired HSI and XRF measurements that enable the study of cross-modal relationships between spectral and elemental information. Furthermore, we developed a reproducible mineral identification pipeline, leveraging the HSI as input, the USGS library as support data, and the XRF as target. Finally, we proposed a convex optimization approach with a customized pruning filter and evaluated it against other USGS-informed benchmarks.

5. REFERENCES

[1] "Minerals in the Wild," `https://github.com/ElefantheriaTtl/minerals-in-the-wild`.

[2] IEA, "Energy and AI," Tech. Rep., IEA, Paris, France, 2025, Licence: CC BY 4.0.

[3] Joseph Lessard, Jan de Bakker, and Larry McHugh, "Development of Ore Sorting and Its Impact on Mineral Processing Economics," *Minerals Engineering*, 2014.

[4] Saeid Asadzadeh and Sabine Chabrillat, "Leveraging enmap hyperspectral data for mineral exploration: Examples from different deposit types," *Ore Geology Reviews*, p. 106912, 2025.

[5] RG Resmini et al., "Mineral mapping with hyperspectral digital imagery collection experiment (hydice) sensor data at cuprite, nevada, usa," *International Journal of Remote Sensing*, 1997.

[6] Gregg A Swayze et al., "Mapping advanced argillic alteration at cuprite, nevada, using imaging spectroscopy," *Economic geology*, 2014.

[7] Bikram Koirala et al., "A multisensor hyperspectral benchmark dataset for unmixing of intimate mixtures," *IEEE Sensors Journal*, 2023.

[8] Daniel C Heinz et al., "Fully constrained least squares linear spectral mixture analysis method for material quantification in hyperspectral imagery," *IEEE TGRS*, 2001.

[9] Marian-Daniel Iordache, José M Bioucas-Dias, and Antonio Plaza, "Sparse unmixing of hyperspectral data," *IEEE TGRS*, 2011.

[10] Bikram Koirala, Zohreh Zahiri, Alfredo Lamberti, and Paul Scheunders, "Robust supervised method for nonlinear spectral unmixing accounting for endmember variability," *IEEE TGRS*, 2020.

[11] Behnood Rasti, Bikram Koirala, and Paul Scheunders, "Hapkecnn: Blind nonlinear unmixing for intimate mixtures using hapke model and convolutional neural network," *IEEE TGRS*, 2022.

[12] Specim, Spectral Imaging Ltd., "Specim swir hyperspectral camera," `https://www.specim.com/products/swir/`.

[13] Alexander F. H. Goetz, Gregg Vane, John E. Solomon, and Barrett N. Rock, *Imaging Spectrometry for Earth Remote Sensing*, Science, 1985.

[14] Bruker, "S1 TITAN Handheld XRF Analyzer for Elemental Analysis," `https://www.bruker.com/content/dam/bruger-com/literature/S1-TITAN_Overview_Brochure.pdf`.

[15] Raymond F. Kokaly et al., "Usgs spectral library version 7," *U.S. Geological Survey Data Series*, 2017.

[16] Naoto Yokoya, Akira Iwasaki, and Kazuhiro Imai, "Spectral response function estimation for hyperspectral sensors," *IEEE TGRS*, 2012.

[17] Roger N. Clark and Ted L. Roush, "Reflectance spectroscopy: Quantitative analysis techniques for remote sensing applications," *Journal of Geophysical Research*, vol. 89, no. B7, pp. 6329–6340, 1984.