

Fast automated mineral differentiation using optical hyperspectral imaging

Sebastian Harenbrock

Trinity College Dublin, the University of Dublin

January 27, 2024

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

Summary

The processing of mineral-bearing rocks in the mining industry is a highly energy and resource-intensive process. This has motivated the development of more effective ways to automatically sort valuable ore from barren material before the most expensive crushing and processing stages.

Many different sensor-based methods to perform automated ore sorting have been investigated in the literature and in the industry. The present work explores potential improvements to sensor-based sorting methods based on optical reflectance hyperspectral imaging.

These improvements are drawn from the field of remote sensing and revolve around the decomposition of reflectance spectral images into relatively small numbers of component spectra (“endmembers”) and their corresponding abundance values. Classification models using linear discriminant analysis (LDA) based on these abundances are fitted and compared against classifiers operating directly on the original spectral data or on RGB colour values. All classifiers were fitted and evaluated on lithium-bearing pegmatite (lepidolite) samples to be differentiated from waste rock (weathering granite) from Gonçalo, Portugal.

The resulting abundance-based classifiers are found to be competitive and generalise well from training data to test data without risk of overfitting, unlike direct spectral classification models. They also offer the option of manually overriding the classification significance of individual spectral components, which is demonstrated on the example of contamination consisting of labels written on the surface of the samples.

Acknowledgements

I would like to express my sincere gratitude to my advisor Prof. Igor Shvets for his long-time support and mentorship since my undergraduate days.

Special thanks to Prof. José António de Almeida of the Earth Sciences Department of the NOVA School of Science and Technology (Lisbon), and Alexandra Carolino of Pegmatítica - Sociedade Mineira de Pegmatites, Ld^a for collecting the lepidolite and granite samples used in this project, and for their hospitality during our visit to the country.

Thanks also to John Güven of the Science Foundation Ireland Research Centre in Applied Geosciences (iCrag) for advice on geological and mining engineering aspects of the problem, and to Dr. David Klindt of Stanford University for invaluable discussions on linear and non-linear classification and potential machine learning options.

Huge thanks to my parents, Dr. Marion Mensing and Dieter Harenbrock, for their love and support throughout my life, and to my mother and Dr. Adam Kelly for proofreading.

The presented work was supported by:

- Sustainable Energy Authority Ireland (SEAI), award number: RDD440,
- AMBER Directors Fund 2018, award number: 208531-12565,
- Science Foundation Ireland (SFI) Future Innovator Prize 2018, award number: 208705-15649.

Contents

1	Introduction	1
2	Literature review	3
3	Samples	9
4	Spectrometer design	15
5	Off-line analysis	21
5.1	Scanning procedure	21
5.2	Pre-processing steps	22
5.3	Post-processing steps	23
5.3.1	Cropping and flipping	24
5.3.2	Downscaling	24
5.3.3	White-reference application	25
5.3.4	Subject detection and trimming	27
5.3.5	Continuum removal	29
5.4	Analysis steps	31
5.4.1	Virtual dimensionality	32
5.4.2	Extraction of spectral endmembers	33
5.4.3	Abundance-map generation	36
5.4.4	Reconstruction	39
5.4.5	Linear discriminant analysis (LDA)	40
6	Real-time differentiation	43
6.1	Scanning procedure	43
6.2	Post-processing	43
7	Results	45
7.1	Inference on training set	47

7.1.1	RGB photo classifier	48
7.1.2	Best-performing direct spectral classifier	50
7.1.3	Best-performing abundance-based classifier	51
7.2	Inference on test set	53
7.2.1	RGB photo classifier	54
7.2.2	Best-performing direct spectral classifier from training set	55
7.2.3	Best-performing direct spectral classifier	56
7.2.4	Best-performing abundance-based classifier	58
7.2.5	Best-performing abundance-based classifier with marker labels ignored	59
7.3	Discussion	60
7.3.1	RGB models	61
7.3.2	Direct spectral models	61
7.3.3	Abundance-based models	61
8	Conclusion	63
8.1	Summary	63
8.2	Industrial applicability	64
8.3	Further work	65
	References	67
	Appendix A Inference results on training set	71
	Appendix B Inference results on test set	79

Chapter 1

Introduction

Extraction of valuable materials from mineral ores in mining involves highly energy and resource-intensive processes, including multiple stages of crushing and grinding of mined rock to a particle size in the μm range, as well as chemical and hydrometallurgical extraction methods. The input of run-of-mine (ROM) material is typically highly dilute, with concentrations of the target material often at just a fraction of a percent. Increasing this concentration by removing waste before the energy-intensive processing stages is critical for reducing the environmental and carbon footprints of a mining operation and has become an area of highly active research.

The pre-concentration concept presented in this thesis belongs to the field of ore sorting, in which mined rocks are physically separated at a macroscopic scale into multiple streams (typically two: target material and waste) according to measured physical properties.

The proposed method involves the measurement of reflectance spectral images of the surface of the input rocks. Particular attention is paid to the need to make a sorting decision for each rock in real time. This is achieved by first conducting an intensive off-line study of the reflectance characteristics of the orebody in question. Here, an example study is conducted on representative samples of lepidolite-bearing pegmatite rocks from an active hard-rock open-pit lithium extraction operation in Gonçalo, Portugal. Based on the results of this study, models are fitted to classify fast low-resolution scans into valuable rocks (lepidolite) and waste rocks (granite) in real time. The main contribution of this thesis consists of the development of classification models based on the relative abundances of different spectral components found across the samples. These abundance-based models will be compared against baseline conventional models that perform classification directly on entire spectral data or on the RGB colour values extracted from photos of the samples.

The goals of the project consist of:

- designing and tuning a spectral imaging hardware system suitable for real-time appli-

cations,

- developing a postprocessing software pipeline capable of extracting spectral images with sufficient signal-to-noise ratio (SNR),
- generating a spectral library for the given orebody,
- optimising the proposed abundance-based methods,
- investigating the benefits of a logarithmic (absorbance) view of reflectance, and of separating spectra into continua and feature spectra, and
- comparing the inference performance on test samples of abundance-based models against baseline approaches.

In [Chapter 2](#), relevant literature on automated mineral classification and remote sensing, from which several methods form the core of the proposed method, is presented.

[Chapter 3](#) introduces the training and test sets of samples analysed in this project, lithium-bearing pegmatite from Portugal.

In [Chapter 4](#), the spectral imaging hardware that was designed for this project is described in detail along with calibration procedures.

[Chapter 5](#) walks through the scanning procedure, and details the postprocessing and analysis pipelines that were applied to captured scans for the off-line study, resulting in the set of classification models to be compared.

[Chapter 6](#) describes the real-time application of the classification models to fast lower-resolution scans.

[Chapter 7](#) presents the classification performance of each model variant on the training and test sets, followed by a discussion.

The thesis concludes in [Chapter 8](#) with a summary, along with identified opportunities of further research.

Chapter 2

Literature review

The importance of sorting as a mineral processing step, i.e. discarding rocks that do not contain the desired material in sufficient quantities, is self-evident enough that it has been in use since prehistoric times when stone tools were first developed. Early industrial mines made heavy use of hand sorting, however for modern mines the associated throughput is simply insufficient to meet demand. With decreasing ore grades it also becomes harder to distinguish valuable rocks from waste by visual inspection. For these reasons, the field of sensor-based ore sorting has received intense and widespread attention.

There are numerous types of sensor that have been explored for the purposes of ore sorting. They can be grouped into the following categories (summarised across Refs. ^[26,21]):

- Optical:
 - Reflectance spectroscopy (incl. colour sensing)
 - Luminescence (incl. fluorescence)
 - Raman scattering
 - Crystal probing (laser scattering, refractive index)
 - Laser-induced breakdown spectroscopy (LIBS)
 - Laser-induced fluorescence
 - Image analysis (structure, texture)
- Other electromagnetic:
 - X-ray transmission (XRT)
 - X-ray fluorescence (XRF)
- Nuclear:
 - Radioactivity (induced or natural)

- Nuclear magnetic resonance (NMR)
- Electrical:
 - Conductivity
 - Permittivity
 - Differential heating by microwaves
 - Triboelectric effect
- Magnetic (e.g., inductivity)
- Acoustic (e.g., transmission)
- Gravity (rarely implemented as sensor, but instead uses these properties physically for separation):
 - Density (e.g., shaking tables)
 - Weight (differential time of flight separating into different receptacles)

Due to the extreme variety of minerals and the compositions of rocks they appear in, no single solution is expected to ever dominate all types of mine. Each type of sensor has its own benefits and drawbacks. For example, optical techniques tend to be fast and inexpensive, however they can only ever probe the surface of a rock, leaving the contents of the bulk as an estimation problem. This makes them most suited for orebodies where the mixture of valuable ore and waste is more heterogeneous, such that the presence of a material in the bulk of a rock is highly correlated with its presence on the surface. Few technologies actually penetrate the bulk of a rock. They mainly include X-ray transmission (XRT), which produces a density profile which can highlight lumps of dense ore inside rocks (making it particularly suitable for elements such as tungsten), and differential heating by microwaves and subsequent detection of hotspots using thermal imaging (MW/IR)^[10].

Modise^[19] reviewed the current use of different parts of the electromagnetic (EM) spectrum for ore sorting and compares and scores EM-sensor based sorting methods based on fundamental principles, such as dependence on orientation, surface cleanliness, molecular polarity and others (scores reproduced in parentheses):

- Colour (5/16)
- Gamma emission (9/16)
- Infrared transmission (7/16)
- XRF (10/16)

- XRT (12/16)
- Combined XRF/XRT (14/16)

Note that these are very general scores and that individual orebodies may well be served best by a lower-scoring method, depending on its properties and not least based on cost-benefit analysis. Reflectance spectroscopy is not evaluated by Modise^[19] directly, but should score similarly to colour sorting among these criteria due to it using a conceptually equivalent measurement technique. The scoring criteria relate to the fundamental limitations of each method and are thus largely independent of improvements to implementation, so the present work makes no attempt to overcome these. As suggested by the score of combined XRF/XRT, it is likely that such limitations will only be overcome by sensor fusion, allowing multiple types of sensors to complement each other^[21]. This further motivates the continued improvement of lower-scoring methods (such as reflectance) as they may become parts of combined systems.

Apart from the fairly common colour sensing approaches, which should be considered a reflectance spectroscopic technique in the visible spectrum (VIS, 0.4 μm to 0.7 μm) with three coarse wavelength bands, the field of optical spectroscopic sensor-based ore sorting typically involves the near-infrared (NIR, 0.7 μm to 1.4 μm)^[8,26], and short-wave infrared (SWIR, 1.4 μm to 3.0 μm) regions of the electromagnetic spectrum, as most defining molecular absorption bands are found there. The use of the visible-to-near-infrared (VNIR, 0.4 μm to 1.0 μm) spectrum for this purpose^[22] is rare, even though sensors in this region are inexpensive. The present work is an attempt to further explore the limits of the VNIR spectrum for ore sorting.

Classification in the VNIR, NIR and SWIR cases usually consists of the measurement of the depth of manually identified absorption bands^[8] or operates on pre-processed spectra (e.g., auto-scaled and transformed to a principal component space^[22]). In some cases, machine learning methods are applied directly to the spectral inputs^[20].

Reference spectral libraries for numerous types of material (incl. minerals) are available^[14,9]. In principle, these could be used to assemble a spectral library for the purposes of classification of samples of a given orebody, however this would be a completely manual process requiring significant geology expertise, specialised for each orebody in question. The spectra chosen would also need to be resampled to match the characteristics of the spectrometer to be used in the sorting application^[35].

The related field of remote sensing (see Gupta^[11] for an overview) has developed a wealth of tools and methods that are of relevance for automatically classifying mineral samples, and which appear underused in the spectroscopic ore sorting literature. Remote sensing often deals with imagery collected from airborne hyperspectral cameras in which pixels cover ground areas of at least several square metres, leading to significant mixture of reflectance

components (e.g., vegetation, rock, water). There is an equivalent problem of separating out spectral components originating from different materials in spectroscopic ore sorting, albeit at a much smaller scale and due to intimate mixtures of minerals. This suggests the application of remote sensing techniques for this problem.

The techniques in question revolve around the concept of *spectral endmembers*, of which the mixed spectra present in a scene are assumed to be a linear combination. In a geometric interpretation, these endmembers form the vertices of a simplex containing all input spectra, that is embedded in the high-dimensional space spanned by all possible spectra.

Numerous different algorithms^[3,5] have been developed for the extraction of endmembers from a spectral image. These include NFINDR^[32], which is used in the present work, and ATGP^[24], which is used here to initialise NFINDR, which has been shown to improve extraction results^[23].

These endmember extraction algorithms require the input of the number of spectral components present. This is also known as the *virtual dimension* (VD) of an input, which can be estimated using methods such as the Harsanyi-Farrand-Chang (HFC)^[4] algorithm.

Endmember extraction is a “local” process, in that extracting endmembers separately from two or more different spectral scans generally produces very different sets of endmembers, without an easily established common reference. A technique called *endmember bundling* (introduced by Somers^[27] as a way to address endmember variability), has proven useful in the present work as way to generate a common reference spectral library for the orebody under study. It essentially consists of k -means clustering of a set of n observations of endmembers across many scans, which partitions this set into k disjoint subsets while minimising the sum of intra-cluster variances.

The benefit of such a spectral library that results is that each pixel of a spectral scan can be unmixed into *abundance values* corresponding to each spectral endmember. This results in a lower-dimensional *abundance space* where individual components can often be mapped directly to the presence of specific materials. The unmixing procedure is known in the remote sensing literature as *abundance map generation*. As with endmember extraction, multiple different algorithms are available. Most commonly used are least-squares optimal methods which differ in the constraints that are applied to the abundance values: unconstrained least-squares (UCLS, essentially multiplication by a pseudoinverse matrix), non-negative least-squares (NNLS)^[1] and fully-constrained least-squares (FCLS)^[12], which ensures non-negative abundances that sum to one.

These methods assume linear (additive) mixing, which is a fair assumption for pure first-surface reflectance with negligible light penetration into the material. With significant light penetration, absorbance (which mixes multiplicatively) becomes dominant, which necessitates

non-linear unmixing for best results. This is most easily achieved by transforming data to the log-domain (applying the negative log to transform reflectance into *apparent absorbance*) and then performing linear unmixing^[34].

Ore sorting typically consists of the classification of samples into two categories (ore and gangue) and therefore lends itself to the use of established binary regression methods. In this work, an approach known as *linear discriminant analysis* (LDA)^[25] will be used for classification. LDA projects input observations of known classes using linear combination into a lower-dimensional space where the separation of the classes (in terms of Euclidean distance) is maximised. Once the parameters of this transformation are fitted using training observations, they can be used to project unknown test observations into this same space, enabling the prediction of the most likely class using the distances to each class.

This thesis will focus on incorporating these techniques in the construction of real-time-capable ore classification models, explaining how suitable spectral libraries can be generated for this purpose and comparing the performance of abundance-based methods against direct spectral-based classification as well as colour sensor-based classification. It will also include the well-established spectral processing technique known as *continuum removal*^[7], which aims to separate local absorption features from a larger smooth background continuum. In the context of continuum removal, the possibility of using information contained in the continua themselves (as opposed to the extracted features) for classification will for the first time also be investigated and compared against the use of unmodified spectra and feature spectra.

Chapter 3

Samples




Figure 3.1: Geographic and geological setting of the lithium Alvarrões-Gonçalo deposit (from Sousa^[28] under Creative Commons 4.0). G1: biotite granite (Guarda granite), G2, G3, G4: biotite-muscovite granites.

The samples used in this study originate from an active open-pit lithium mining operation (“Mine C-57”) in the Gonçalo Li-pegmatite field in Portugal (see [Figure 3.1](#)). They were collected (see [Figure 3.2](#)) with the goal of being representative and classified by an expert geologist familiar with this orebody (see [Acknowledgements](#)) into:

- Valuable material (lepidolite, 0.6%+ Li)
- Waste/host rock (weathering granite, <0.1% Li)



Figure 3.2: Sample collection site and process.

The rocks were then labelled according to these classes (PL for lepidolite samples, PG for granite). For the analysis portion of this study, six medium-sized samples (2 cm to 10 cm across, see [Figure 3.3](#)) per class were analysed and used to fit classification models. Remaining medium-sized samples (four per class) were used as a test set (see [Figure 3.4](#)). Each sample was analysed from two opposite sides, referred to by A and B.

Due to happenstance, most of the lepidolite samples in the training set were labelled with a different marker (pink) than the granite samples (purple). It is important to be aware of this, as it could lead a classification system to use the spectral signature of the label as basis of its classification, which is undesirable. Since the affected area is relatively small, a majority-based classification system is unlikely to be significantly affected by this. Nevertheless, [Section 5.4.5](#) will describe a way to suppress such individual spectral components in the classification. Also, all samples in the test set were labelled with the purple marker to “punish” any inadvertent fit to the pink labels.

Detail views of some representative granite and lepidolite samples along with visual descriptions can be seen in [Figures 3.5](#) and [3.6](#), respectively. Typical spectra of the minerals making up this orebody are shown in [Figure 3.7](#).

It should be noted that some samples classed as lepidolite have a side that contains little to no lepidolite, seen in: PL8-A (right edge), PL9-B, PL13-B and PL16-B (middle). These samples were likely collected from the edge of the lepidolite layer. While containing sufficient quantities of the desired lepidolite to be classified as such, samples like these (depending on their orientation) are likely to confuse most surface-based sorting methods such as the one presented in this work and are therefore expected to cause false-negative errors.




Figure 3.3: Set of training samples used for fitting the classifiers in this study. The top two rows are granite samples, the bottom two rows lepidolite. Note that the perspective of the photos may differ from the perspective of the corresponding scans. Photos taken with a Nikon D7200 DSLR with 24 mm f/1.4 lens under diffuse white LED lighting.




Figure 3.4: Set of test samples used for validating the classifiers. Top two rows: granite, bottom two rows: lepidolite. Note that the perspective of the photos may differ from the perspective of the corresponding scans. Photos taken with a Nikon D7200 DSLR with 24 mm f/1.4 lens under diffuse white LED lighting.




Figure 3.5: Detail of two granite samples. Left: PG11-A from the training set, consisting of albite (white-brown) and oligoclase (grey) containing very small clear quartz crystals. Right: PG13-B from the test set, showing reddish K-feldspar and black biotite crystals in albite.




Figure 3.6: Detail of two lepidolite samples. Left: PL8-A from the training set, demonstrating the lavender colour of lepidolite crystals interspersed with clear quartz, embedded in albite. The brown K-feldspar on its right-hand side indicates that this sample originated from a contact zone with the granitic layer. Right: PL16-B from the test set, showing mainly yellow-to-white albite with some lepidolite and quartz.




Figure 3.7: Example spectra published by the US Geological Survey^[14] for the main constituent minerals of the studied samples. Circle fills indicate the approximate colour of each material and the RGB sampling points used to estimate this. The shaded region designates the (cropped) wavelength range of the imaging spectrometer described in Chapter 4.

Chapter 4

Spectrometer design

Due to the eventual goal of applying the method developed in this project to a moving stream of mined material, an imaging spectrometer was designed that allows the incremental scanning of a 2-d sample surface by imaging a narrow line through which the sample moves orthogonally. This concept is known as push-broom scanning.




Figure 4.1: Optical diagram of the PGP-type imaging spectrometer design. The blue and red rays are used to illustrate the light paths of the shortest and longest imaged wavelengths, respectively (not the corresponding colours as perceived by a human). The way that the spectrometer was configured in this project, the centre wavelength actually corresponds to red light, with the longest wavelength in the near-infrared.

The imaging of this line is achieved by narrowing the field of view using a vertical slit as a field stop located at the focal point of the front lens. The light is then collimated by another lens and diffracted by a diffraction grating which is inserted between two opposite beam-deflecting wedge prisms. This design is known as a prism-grating-prism (PGP) spectrometer (see [Figure 4.1](#)). In this design, the outgoing first-order diffraction mode at the centre wavelength exits the second prism along the same axis that the collimated light enters, with the zeroth order being deflected off-axis and absorbed. This allows the re-imaging system (a lens and a sensor) to be situated on this same axis. This greatly simplifies the alignment of all op-

tical components, as they can all be mounted along this axis using standard optomechanical parts.




Figure 4.2: Render of the imaging spectrometer including optical components used. Some parts are rendered translucent to show the components inside.

The optical components used (shown assembled along with mechanical components in [Figure 4.2](#)) consist of:

- MVL12M23: front lens (12 mm f/1.4 C-mount). This was chosen to achieve a relatively large field of view, to allow for the scanning of larger samples.
- VA100C/M: variable-width slit as field stop, positioned at the focal point of both the front lens and the collimating lens. This limits the field of view to the vertical axis.
- AC254-050-AB: achromatic doublet (50 mm) as collimating lens.
- LTFW6: filter wheel (useful with notch filters for fluorescence or Raman spectroscopy experiments; not used in this work).
- 2x PS812: wedge prism (6° beam deviation).
- GT25-03: visible transmission grating (300 grooves/mm with 17.5° groove angle).
- MVL35M23: re-imaging lens (35 mm f/2.0 C-mount).
- BFS-U3-32S4M: imaging sensor (2048x1536 monochrome with 12 bit ADC).

All part numbers refer to products from Thorlabs except for the camera sensor, which is produced by FLIR.

The vertical axis on the sensor (1536 pixels tall) corresponds to the vertical axis imaged on the sample, and the light is dispersed by the diffraction grating perpendicular to this, making the horizontal sensor axis (2048 pixels wide) the wavelength axis.

The Bragg condition for a transmission grating is given by:

$$d(\sin \theta_m - \sin \theta_i) = m\lambda \quad (4.1)$$

where d is the grating spacing, θ_i is the angle at which incident light of wavelength λ impinges on the grating, and θ_m is the angle at which light of order m exits it.

Given two identical prisms deflecting incident and outgoing light (at the centre wavelength λ_{centre}) onto the optical axis by beam deflection angle δ , (4.1) can be rewritten for the dominant order $m = 1$ as:

$$d(\sin \delta + \sin \delta) = \lambda_{\text{centre}} \quad (4.2)$$

Given $\delta = 6^\circ$ and $d = 1 \text{ mm}/300$, the centre wavelength can be calculated as $\lambda_{\text{centre}} \approx 697 \text{ nm}$. The horizontal extent of the sensor results in an imaged wavelength range of 362 nm to 1027 nm. As a result, the $m = 1$ and $m = 2$ orders overlap from around 700 nm. However, the transmission grating's groove angle of 17.5° results in most of the transmitted optical power being concentrated in the $m = 1$ order, with $m = 2$ having negligible intensity. It was therefore not deemed necessary to implement any order sorting.

Due to drop-off towards the ultraviolet of the emissivity of the light source, the transmissivity in the optical components and the sensitivity of the sensor, in the presented experiments the first 256 of the 2048 spectral bands were discarded due to excessive noise. The resulting wavelength range is 446 nm to 1027 nm with around 0.325 nm between bands before down-sampling (see [Section 5.3.2](#)). The best spectral resolution was measured in the centre of the frame to be around 1.23 nm FWHM given a monochromatic 532 nm light source, with resolution worsening towards the edges of the frame.

This spectral resolution is likely excessive for the orebody under examination here, since the sample classes are plausibly distinguished by colour (see [Figure 3.3](#) and [Figure 3.4](#)) and the constituent minerals present no significant sharp spectral features in this wavelength range (see [Figure 3.7](#)). However, the method described here is intended to be generally applicable to many orebodies, and as such the full spectral capabilities of the spectrometer along with the resulting data processing challenges will be explored, as well as compared in performance against equivalent approaches using RGB data only.




Figure 4.3: Left: imaging system with halogen lamp and cylindrical lens mounted above. Middle: 532 nm laser light (spread into a line with a cylindrical lens) illuminating the white reference for the spectral resolution measurement. Right: light column from the halogen lamp illuminating the white reference.

The spectrometer can be configured for reflectance as well as fluorescence measurements by choosing appropriate light sources and filters. For the reflectance measurements performed in this project, a 10 W quartz-tungsten-halogen lamp (Thorlabs QTH10) was used due to its wideband continuous emission spectrum across VIS and NIR. To raise the illumination level of the part of the sample visible to the spectrometer, the collimated output beam of the lamp was focused down to a narrow column concentric with the field of view of the spectrometer using a cylindrical lens. No filters were used.




Figure 4.4: Left: scanning setup with Thorlabs ELL20/M linear piezoelectric stage and scene background lined with Thorlabs BFP1 black flocked paper. Right: reflectance of the BFP1 material as given by Thorlabs.

The background of the scene was lined with black flocked paper with low reflectivity in the visible spectrum (Figure 4.4). Note that there is significant reflectance in the NIR spectrum; for this reason, the background will show up as a spectral component in Section 5.4.3. A




Figure 4.5: Calibrated reflectance of 99% diffuse reflectance standard as supplied by Spectralon. The shaded region designates the (cropped) wavelength range of the imaging spectrometer.

piezoelectric motorised linear stage with 60 mm travel (Thorlabs ELL20/M, see [Figure 4.4](#), left) was used to incrementally scan successive slices of each sample’s surface. The stage was controlled and programmed with the Elliptec ELLO Piezoelectric Resonant Motor Control Software GUI, whereas the FLIR SpinView application was used to record images from the image sensor. No synchronisation between stage movement and frame acquisition was employed; instead, the timing of the movement sequence was set to produce the desired number of frames per step in accordance with the frame rate of the sensor. Note that this likely caused a small increase in the average horizontal positioning error due to occasional overlap of frames between steps.

To calibrate the spectral response of the spectrometer, a Spectralon 99% white diffuse reflectance standard with relatively flat response across the target wavelength range was used (see [Figure 4.5](#)).

Chapter 5

Off-line analysis

5.1 Scanning procedure

To allow enough room for the arrangement of the light source and spectrometer optics in a shared plane perpendicular to the motion of the stage, an imaging distance of 24 cm was chosen. To optimise for horizontal spatial resolution, the slit width was set to 40 μm , as narrow as repeatably possible. Given the imaging focal length of 12 mm at a distance of 24 cm, this results in a width of the imaged line of 800 μm . The distance advanced by the stage at each step was matched to this value. At the maximum stage travel of 60 mm this implies 75 steps per scan. The bit depth of the imaging sensor was set to its maximum value of 12 bits, stored as 16-bit unsigned integers at a resolution of 2048 by 1536 pixels (implying exactly 6 MiB per frame). The 99% reflectance target was used to identify the upper limit of the exposure time that avoids sensor saturation (note that specular reflection off a sample can still cause saturation). An exposure time of 250 ms at a gain of 10 was chosen, which gave about 100% headroom above the maximum signal produced using the reflectance standard. To improve SNR, a number of these exposures were captured at each step for later stacking. To keep the time taken to complete each scan reasonable, it was decided to take 160 exposures for each step. This implies 40 s per step for a total of 12000 frames (“lights”), with each complete scan taking 50 minutes and generating around 75.5 GB of raw data.

Before each scan, a set of 160 frames (“whites”) of the reflectance standard, illuminated by the halogen light source, was captured at the same settings to allow for calibration of the spectral response of the system.

A set of 160 dark frames (with a lens cap on the front lens) for reduction of fixed-pattern noise was also captured at the same settings. This was done at the end of each scan, to ensure a consistent camera temperature had been reached and thus a more representative dark noise pattern could be captured.

5.2 Pre-processing steps




Figure 5.1: Pre-processing flowchart.

The initial processing of the data consisted of summing (“stacking”) each successive set of 160 light frames into a new high-bit-depth image, leading to 75 “stacks,” each resulting from imaging a different vertical slice of the sample. These were then combined into a “data cube,” with the first two axes representing the y-x spatial directions of the imaged sample surface, and the third being the spectral axis.

The sets of dark and white frames were each stacked in the same way. The dark stack was then subtracted from all light stacks and from the white stack, largely cancelling out the fixed-pattern noise caused by dark current in the sensor. Any negative values potentially introduced in the light stacks by the subtraction were replaced with zeroes.

5.3 Post-processing steps




Figure 5.2: Post-processing flowchart. Branch points separating reflectance and absorbance pipelines are labelled with “Refl.” and “Abs.”, respectively, with absorbance-only steps shaded in grey.

To clean up the raw signals and turn them into useable spectra, the data cubes were passed through a post-processing pipeline (see Figure 5.2). This pipeline performs the application of the white-reference image stack (Section 5.3.3) in addition to simple operations such as cropping, downscaling, flipping and masking, and concludes with continuum extraction/removal (Section 5.3.5) which extracts continuum and feature spectra.

Reflectance data is sometimes better modelled as apparent absorbance (in the following plainly termed “absorbance” for simplicity), depending on the contribution of first-surface reflection compared to the level of penetration of incident light into the material. This matters in the processing of the data when multiple different material components are covered by the same pixel, as in the reflectance model their contributions combine additively, whereas in the absorbance model they combine multiplicatively.

To test which of these models is more accurate for the given data, two separate post-processing pipelines were assembled, with one operating on reflectance spectra directly and the other on absorbance spectra, calculated as the negative log of the reflectance values. Operations that are additive in the reflectance pipeline were performed in the log-domain for the absorbance pipeline (resulting in a multiplicative version of the operation). Otherwise, the two pipelines were equivalent.

Due to the three-dimensionality of a data cube, there are many ways to visually represent slices of it. The following sections will make frequent use of reconstructed RGB images, which display three spectral slices (650 nm, 550 nm and 450 nm for simplicity) of a data cube in the RGB channels of a colour image. This allows a natural visualisation of the effect of a spectral transformation on the basis of these three spectral bands.

5.3.1 Cropping and flipping

None of the samples were of sufficient height to fill the entire field of view of the spectrometer, so the y-axis was cropped above, as well as below to remove the sample stage from the view. This resulted in an initial region of interest (ROI) for all scans from pixel row 144 to 912, with its height (768) chosen to be divisible by 8, to allow even downscaling by this factor later (see [Section 5.3.2](#)). The y-axis of all data cubes was also flipped to ensure the scans were right-side-up, as the spectral imaging system described in [Chapter 4](#) captures images upside-down. The resulting dimensions (height by width by bands) of each data cube after this step were $768 \times 75 \times 1792$.

5.3.2 Downscaling




Figure 5.3: RGB representation of raw reflectance data cube of PG10-B before (a) and after (b) vertical downscaling by a factor of 8.

The relatively high pixel density of the sensor resulted in a vertical spatial resolution of $100\text{ }\mu\text{m}$, much finer than the horizontal spatial resolution of $800\text{ }\mu\text{m}$. This level of vertical resolution was deemed unnecessary, and so the vertical axis was downsampled (by local averaging) by a factor of 8 which resulted in square pixels and also helped to reduce noise levels.

Given the effective spectral resolution of at best 1.23 nm FWHM, the spectral grid defined by the sensor was also found to be unnecessarily fine, with a mean spacing of 0.325 nm . Hence the spectral axis was also downsampled by a factor of 4, resulting in 448 spectral channels with a mean spacing of 1.3 nm (and further reducing noise due to averaging). This led to a manageable size of $96 \times 75 \times 448$ for each data cube.

Note that in the absorbance pipeline, the above downscaling steps are applied in the log-domain.

5.3.3 White-reference application




Figure 5.4: Illustration of white-reference application. a)+d) White-reference stack and signal along the highlighted pixel row. b)+e) Reconstructed RGB image of PL1-A scan before division by white reference and signal at the highlighted pixel. c)+f) Reconstructed RGB image of PL1-A scan after division by white reference and signal at the highlighted pixel.

A critical postprocessing step is the white-reference application. This involves the division of the data cube by the white stack (and multiplication by the calibrated reflectance curve provided with the imaged standard), transforming raw sensor signals into reflectance spectra. This is equivalent to white-balancing in a colour camera (i.e., ensuring that a perfectly reflecting white subject results in all spectral channels having a value of unity) as well as flat-field correction (correcting for uneven illumination and vignetting effects, which can be clearly seen in [Figure 5.4 a](#)).

This step also corrects for artifacts of unclear origin in the imaging system, which can be seen as dark horizontal lines in the white-reference image and in the pre-white-balanced RGB image in [Figure 5.4 a](#)) and b), which are effectively cancelled out in c).

Division by small values being an operation particularly sensitive to noise, this step benefits from noise reduction through downscaling and is therefore done after that step.

If the white reference stack contains any negative or zero values after dark-stack subtraction, the corresponding division of those pixels is not performed. One might consider removing the corresponding row or spectral bands from analysis, but in practice this situation is easily avoided when the total exposure times of white and dark stacks are long enough.

Note that while more than 100% diffuse reflectance is impossible, reflectance values larger than one may still occur. When part of a continuous spectrum, these values typically indicate the presence of specular reflection on an exposed face of a crystal, while when discontinuous they are usually caused by noise.

5.3.4 Subject detection and trimming




Figure 5.5: Illustration of the subject masking and trimming process on the example of PL11-A. a) Reconstructed RGB image with background pixels not meeting the luminance threshold highlighted in red. Note the dark border around the subject, indicating mixture with background pixels. b) Reconstructed RGB image with subject mask contracted by one pixel. c) Reconstructed RGB image with the bounding box of the subject highlighted in green. After the trimming step, the data cube will be cropped to this region.

After evening out the brightness of the data cube across the frame, it is much easier to separate the subject from the background based on simple thresholding. This process consists of:

- Selecting the largest group of connected pixels with a mean reflectance above a chosen threshold (Figure 5.5 a).
- Contracting this subject mask inwards by one pixel to reduce mixture with background pixels (Figure 5.5 b). This is achieved by removing all pixels that have less than eight neighbouring pixels in the subject mask.
- Zeroing out non-subject pixels.
- Trimming the region of interest (ROI) to the bounding box of all non-zero pixels (Figure 5.5 c).

This allows subsequent operations to be easily applied to the subject only, by filtering out pixels with a reflectance spectrum of zero.

Note that this procedure can sometimes introduce “holes” in the subject mask, see for example Figure 5.6. In this case, the intensity of a few pixels in the interior of the subject did not meet the chosen threshold. This led to a small non-subject region inside the subject,

which grew after the mask contraction step. This could be dealt with by more finely tuning the threshold, or by detecting such holes algorithmically and assigning them to the subject mask. (This is left as an exercise to the reader.) In this work, the loss of a few pixels of information was not considered to be of significance for the development of the method and was thus accepted.




Figure 5.6: Illustration of the subject masking and trimming process on the example of PG10-A. a) Reconstructed RGB image with background pixels not meeting the luminance threshold highlighted in red. Note a small square hole in the interior of the subject. b) Reconstructed RGB image subject mask contracted by one pixel. Note that the hole in the subject mask has expanded by one pixel in each direction. c) Reconstructed RGB image with the bounding box of the subject highlighted in green.

5.3.5 Continuum removal




Figure 5.7: Continuum removal in the reflectance pipeline on the example of PL9-A. a) Reconstructed RGB image of original data cube. b) Reconstructed RGB image of data cube with absorption features removed from each spectrum by division. c) Reconstructed RGB image of data cube with continua removed from each spectrum by division. Note the enhancement of the label written on the sample. d)-f) Spectra of pixels highlighted in a)-c).

A reflectance spectrum can be modelled as a multiplicative combination of a continuum and a spectrum of absorption features caused by electronic transitions. The continuum often consists of the tails of larger absorption features centred outside the spectral region of interest. It can be helpful in the analysis of smaller absorption features to computationally remove the continuum, counteracting shifts from their usually consistent positions. A common approach to compare spectra with differing continuum-parts is to estimate the continuum of each spectrum by fitting a convex hull over it and dividing the original spectrum by this continuum. This process is called continuum removal (CR) and is illustrated in Figure 5.7. The calculation of the convex hull was performed using the *SciPy*^[30] Python package.

For the pipelines based on absorbance—the negative log of the reflectance—the convex hull is fitted from below. Also, in this log domain the continua are removed by subtraction instead of division (see Figure 5.8; note that RGB images here are scaled by a factor of 0.5 to avoid clipping, as absorbance is often larger than 1).




Figure 5.8: Continuum removal in the absorbance pipeline on the example of PL9-A. a) Reconstructed RGB image of original data cube. b) Reconstructed RGB image of data cube with absorption features removed from each spectrum by subtraction. c) Reconstructed RGB image of data cube with continua removed from each spectrum by subtraction. d)-f) Spectra of pixels highlighted in a)-c).

Continuum removal presents a branch point in the set of candidate models for automatic classification of samples. In [Chapter 7](#), models based on the original spectra, the extracted continua or the continuum-removed (i.e., feature) spectra will be compared. Working with the extracted continua themselves is a novel approach, and can be seen as a way to probe spectral features outside the spectral region of the spectrometer used.

In the remainder of this thesis, the resulting three reflectance-based pipelines will be referred to as:

- Reflectance
- Reflectance continua
- Reflectance features

The corresponding three absorbance pipelines will be referred to as:

- Absorbance
- Absorbance continua
- Absorbance features

Note that the continuum removal procedure is extremely sensitive to “hot pixels,” i.e. pixels with abnormally large signals caused by noise, as the convex hull would grow to accommodate these values, leading to incorrect feature extraction. For this reason, dark-stack subtraction (see [Section 5.2](#)) is an essential prerequisite for this step.

5.4 Analysis steps




Figure 5.9: Spectral library optimisation flowchart. The shown combinations of q (number of endmembers per scan) and k (number of endmember bundles across training set) were compared on the basis of the reconstruction error they produced, leading to an optimum at $q = 6, k = 30$.

After post-processing, each data cube contains a relatively clean set of reflectance or absorbance spectra which are isolated from the background. An example of 24 reflectance spectral slices of a lepidolite scan can be seen in [Figure 5.10](#). Apart from the distinct spectral characteristics of the label written on the sample, not much useable information is immediately apparent. To extract such information, a further analysis pipeline was applied that transforms the data in a way that it can be used to fit a classification model suitable for real-time inference. Briefly, this consists of the identification of a common set of spectral “endmembers” (basis spectra which are assumed to linearly combine to form all spectra found in a scene) along with corresponding abundance fractions of each endmember in each subject pixel of each scan. This was achieved with an optimisation procedure outlined in [Figure 5.9](#) and detailed in the following sections. The resulting endmember abundances could

then be used along with the known class labels to fit a linear model separating the classes.




Figure 5.10: Slices of 24 (out of 448) evenly spaced spectral bands of the PL1-A reflectance data cube.

5.4.1 Virtual dimensionality




Figure 5.11: Boxplots of estimated virtual dimension of each class of sample in the reflectance pipeline for a) reflectance spectra and b) reflectance-continuum spectra. No virtual dimension could be calculated for the feature spectra using available algorithms.

Most endmember extraction algorithms found in the literature require as input the number q of expected spectral components. A natural choice of this parameter can be the “virtual

dimension” (VD) of the given spectral image. This is an estimate of the number of distinct spectral components present in a scene. The virtual dimension of each scan was estimated using the HFC algorithm^[4] implemented in the *pysptools*^[29] Python package (see Figure 5.11 for ranges of virtual dimensions for each class of sample). This implementation failed to produce a result on feature spectra, and it was unclear what kind of data clean-up procedure would allow it to proceed. However, it is a reasonable assumption that the virtual dimension of a feature-extracted scan lies between the virtual dimension of the corresponding original scan and its continuum-extracted scan.

The virtual dimension estimated in this way was not used directly as the parameter q for endmember extraction, but rather as a starting point for an optimisation procedure (see Figure 5.9) to find the most appropriate set of endmember “bundles” (see Section 5.4.2.1).

5.4.2 Extraction of spectral endmembers




Figure 5.12: Endmembers extracted for the PL10-A data cube with $q=6$ for a) reflectance, b) reflectance continua and c) reflectance features.

To extract the spectral endmembers (essentially a set of basis vectors onto which each mixed spectrum will be projected) of a given scan, the N-FINDR^[32] algorithm was used, with the starting point initialised^[23] by ATGP^[24]. The implementation used was also part of the *pysptools*^[29] Python package.

Given the mean value of approx. 10 virtual dimensions found in Section 5.4.1, the endmember extraction was repeated for a range of values from 6 to 16. Based on the relative reconstruction error after random sampling from endmember bundles (see Sections 5.4.2.1 and 5.4.4), the optimal number of endmembers in this range was found to be 6 for all pipelines, so this value was used for the generation of each classification model. Example endmember sets are shown in Figure 5.12 (for a lepidolite scan) and Figure 5.13 (for a granite scan).




Figure 5.13: Endmembers extracted for PG3-B reflectance data cube with $q=6$ for a) reflectance, b) reflectance continua, c) reflectance features.

5.4.2.1 Endmember bundling

Note that endmembers extracted separately for two or more different scans will generally lack a common reference. For example, it is not obvious how to match up spectral endmembers between [Figure 5.12](#) and [Figure 5.13](#). For this reason, a technique called endmember bundling^[27] was applied to all endmembers extracted from all scans in a given pipeline to identify a common set of “bundles” of endmembers. Endmember bundling is typically used to address the fact of endmember variability, which means that a given material never presents only with a single reflectance spectrum. It is achieved using simple k -means clustering, which groups all endmembers into k different clusters. Here, the technique is used to establish a common reference set of endmember bundles across all samples, i.e., a spectral library.

Due to the use of k -means clustering, the number of desired endmember bundles is a required input to this process, so this number was also optimised using the relative reconstruction error of scans based on random sampling from the generated sets of endmember bundles. The optimal number of clusters minimising the reconstruction error for all pipelines was found to be 30. The resulting set of endmember bundles was then used as a spectral library for abundance map generation (see [Section 5.4.3](#)).

One prominent endmember bundle (EM20) found in the reflectance pipeline is shown in [Figure 5.14](#) and compared with lepidolite reflectance spectra published by the US Geological Survey^[14] and ECOSTRESS^[9]. There is a good correlation between the spectra, in particular the absorption feature due to Fe^{3+} or Mn^{3+} around 555 nm^[2]. This suggests that the presence of this endmember could be used as a direct indicator of the presence of lepidolite, the target lithium-bearing mineral.




Figure 5.14: EM20 (mean of endmember bundle 20) compared with 10 closest-matching lepidolite spectra from datasets published by the US Geological Survey^[14] (three samples of specimen HS167) and the ECOSTRESS^[9] project (sample PS-13B). Spectra were offset vertically to facilitate visual distinction. No spectral resampling was applied to match the spectrometer types.

Note that while it can be valuable to confirm that extracted endmembers correspond to specific known materials, this is not actually required for the presented method, as no prior information about materials or their corresponding spectra needs to be incorporated. This allows the procedure to be implemented in a relatively automated fashion, requiring only a set of known valuable samples and a set of known waste samples.

5.4.3 Abundance-map generation




Figure 5.15: Extracted reflectance endmembers for PL1-A and associated false-colour representation with each endmember’s abundance mapped to red, green, blue, cyan, magenta and yellow. a)+d): Reflectance endmembers with corresponding abundance map. b)+e): Reflectance continua endmembers with corresponding abundance map. c)+f): Reflectance features endmembers with corresponding abundance map.

Once endmembers have been identified, either for a single scan or as a common set of bundles over multiple scans, it is important to be able to estimate how much of each endmember’s corresponding material is present in a given pixel. This process is called abundance-map generation and involves the unmixing (decomposition) of a mixed spectrum into a linear combination of the given endmembers. Two particular unmixing algorithms were compared for this purpose:

- UCLS (unconstrained least squares), which consists of multiplication of the unknown spectrum by the (Moore-Penrose) pseudo-inverse of the endmember matrix. It is a fast and relatively robust procedure, but may result in abundances that are negative or do not sum to one. It is a suitable candidate for estimating abundances in real time due to its low complexity and vectorizability.
- FCLS (fully-constrained least squares)^[12], which imposes the physicality conditions that individual abundances must be non-negative and sum to one for each pixel. This is a much slower process than UCLS as it involves an optimisation procedure for each

pixel. It is therefore currently unsuitable for on-line abundance estimation, but was included nevertheless as a baseline for comparison. Note also that the constraints may induce incorrect results in the presence of noise.

As with the endmember extraction step, *pysptools*^[29] implementations of both of these unmixing algorithms were used.

The resulting abundance maps can be used for visual inspection when identification of individual endmembers is required. For example, as stated in [Chapter 3](#), the sample surfaces contain a written label which is a potential source of bias for any classification model. Abundance maps can be used to identify the endmembers that correspond to such contamination, so that they can be later explicitly ignored. See [Figure 5.16](#), where the spectral contribution of the label written in pink marker is clearly concentrated in EM12, and [Figure 5.17](#), where the label written in purple marker can be found in EM11.




Figure 5.16: FCLS-generated abundance maps for PL1-A of each reflectance endmember of the common spectral library. Note the pink label endmember EM12, the scene background endmember EM16 (see [Figure 4.4](#), right) and the potential lepidolite endmember EM20.




Figure 5.17: Abundance maps for PG1-A of each reflectance endmember. Note the purple label endmember EM11 and the lack of abundance of potential lepidolite endmember EM20.

With a common spectral library established across all samples, it was now possible to compare abundance maps for a given endmember between different scans. See [Figure 5.18](#) for abundance maps of the likely lepidolite endmember EM20 across eight scans of four different samples. Note that the lepidolite samples show significant abundance of this endmember, whereas the granite samples (and side B of PL9, which appears to come from a boundary between lepidolite and granite, see [Figure 3.3](#)) show virtually no such abundance.




Figure 5.18: Abundance maps for reflectance endmember bundle 20 (indicating presence of lepidolite) for four different samples from two sides each. Note sample PL9, with side A showing significant lepidolite content and side B virtually none.

5.4.4 Reconstruction

To assess the quality of a given spectral library (i.e., set of endmember bundles), it was used in combination with the generated abundance maps to reconstruct each original spectral image. The RMS difference between the reconstruction and the original, normalised by the mean of the original, gives the relative reconstruction error.

In this process, random sampling was used to select the endmember from each bundle that would reconstruct a given pixel. Spectral libraries with more homogeneous clusters should be more resistant to the error introduced by this random sampling and therefore should achieve lower reconstruction errors. As stated in [Section 5.4.2](#), the parameters that minimised the reconstruction error were $q = 6$ (endmembers per scan) and $k = 30$ (number of bundles) for all pipelines. These parameters were used to generate the final spectral libraries for each pipeline.

5.4.5 Linear discriminant analysis (LDA)




Figure 5.19: Endmember score visualisation for LDA model fitted on all endmembers. a) Top 5 positive indicator reflectance endmembers (line width indicates ranking). d) Top 5 negative indicator reflectance endmembers. b)+c)+e)+f) Abundance maps of four different scans with top 5 positive endmember abundances in the green channel and top 5 negative endmember abundances in the red channel. Note dominance of pink marker endmember EM12 and purple marker endmember EM11, and that lepidolite endmember EM20 is not found in the top 5 positive endmembers here.

Once the spectral libraries were finalised, they were subjected to linear discriminant analysis in order to grade each endmember's significance in distinguishing the two classes of sample (lepidolite vs granite).

Linear discriminant analysis (LDA) projects the endmember abundance values such that the two classes are maximally separated, and then determines a linear combination of the transformed abundances that best separates the classes, leading to a coefficient for each endmember and an intercept. A fitted LDA model is capable of inferring the class of a given input abundance map pixel by projecting it to either side of the separating hyperplane. In addition to this, the fitted coefficients can also be inspected to determine the relative importance of each component.

Endmembers that are predominantly present in the valuable class will have relatively large positive values (their presence is a positive indicator of the valuable class), and those present more in the waste class will have relatively large negative values (their presence is a negative

indicator of the valuable class). Some endmembers may be present in similar amounts in both classes, and so their coefficients will have small absolute values (their presence does not aid significantly in discrimination). See Figure 5.19 for the top 5 positive and top 5 negative endmembers, along with corresponding example abundance maps generated using FCLS unmixing mapped to the green and red channels, respectively. It is clear that the labels written in pink and purple marker were (correctly) identified to carry high significance for classification.

To ignore the presence of known contamination, the abundance of corresponding endmembers (e.g. EM12, see Section 5.4.3) can be set to zero (rather than removing the corresponding dimensions) before fitting the LDA model. This will ensure that the model will not be biased by them, while maintaining the same dimensionality of all inputs, such that an LDA model fitted on all endmembers can easily be compared against an LDA model fitted with some endmembers ignored. See Figure 5.20 for the top 5 positive and negative endmembers and associated abundance maps for an LDA model with the pink and purple marker endmembers ignored.




Figure 5.20: Endmember score visualisation for LDA model fitted ignoring the sample label endmembers. a) Top 5 positive indicator reflectance endmembers (line width indicates ranking). d) Top 5 negative indicator reflectance endmembers. b)+c)+e)+f) Abundance maps of four different scans with top 5 positive endmember abundances in the green channel and top 5 negative endmember abundances in the red channel. Note that pink and purple marker endmembers EM12 and EM11 have been suppressed and that lepidolite endmember EM20 is now found among the top 5 positive endmembers.

Chapter 6

Real-time differentiation

6.1 Scanning procedure

To test the classification performance of the fitted models in a more real-time sorting setting, all samples of the test set were scanned with a total exposure time per scan reduced from 50 minutes to 10 seconds (a 300-fold reduction) compared to the training scans. At the maximum stage travel of 6 cm, this corresponds to a scanning speed of 6 mm/s. To compare this against the required speed of an industrial ore sorting system, a number of assumptions must be made. For an open-pit mine such as the one from which the rock samples under investigation were sourced, a conveyor throughput on the order of 100 m³/h is typical^[31]. Assuming the majority of the throughput consists of particles in a sortable size range around 10 cm (and modelling particles as cubes) results in a conveyor belt speed of approx. 3 m/s, with the scan speed of the laboratory prototype being 3 orders of magnitude below this. See [Section 8.2](#) for possible ways to close this gap.

For these real-time scans, as for the off-line training scans, dark and white frames were captured in order to reduce fixed-pattern noise and perform white-balancing. Note that these can still be captured and prepared off-line, so large numbers of frames can be used for this purpose to avoid amplification of the higher noise levels resulting from the short exposure times. For the fast scans of the test set of samples, 1600 frames were used each for the dark and light stacks.

6.2 Post-processing

Due to the shorter scan time, each fast scan consisted of 40 steps (compared to 75 steps for the off-line scans). To achieve approximately square pixels, vertical downscaling of a factor of 16 was applied (compared to a factor of 8 off-line, [Section 5.3.2](#)). This resulted in a worse spatial resolution of 1.5 mm (horizontal) by 1.6 mm (vertical) with the benefit of

reduced noise. The wavelength axis was cropped and downscaled by a factor of 4 for a mean spacing of spectral bands of 1.3 nm, as before.

In addition to dark subtraction and white division, for each variation of the main method some additional processing is required:

- For absorbance models, the negative natural logarithm of each spectrum is taken.
- For models fitted on continuum or feature spectra, convex hull fitting is performed on each pixel. This is not an expensive operation as it involves only linear interpolation.

The abundance-based discrimination process for each spectral pixel consists of performing a low-complexity estimate of the abundance of each endmember using UCLS unmixing and applying the LDA model fitted in [Section 5.4.5](#). This is a fast computation that can easily be performed within milliseconds. This results in an estimate of the probability of the pixel belonging to the target class.

All of these processing steps as well as the inference step can be applied to each column of pixels individually as it becomes available during the scan. Even without special optimisation of the algorithms, they can be easily completed in real time.

For each scan, the application of a classification model resulted in a map of estimated probabilities of each pixel belonging to the target (lepidolite) class. Pixels with an estimated probability of 0.5 or higher were assigned to lepidolite, and the rest to granite. In the following, the error of this inferred class is evaluated for each subject pixel. To assign a class to an entire sample, the same threshold of 0.5 was applied to the mean of the estimated probabilities over all subject pixels. Note that in a given application one may decide to set a different threshold, e.g. a lower threshold could be appropriate in cases where there is a large cost to false negatives.

Chapter 7

Results

In a mining context, a classification system should be evaluated on the basis of how many samples were misclassified, either as false positives or false negatives (potentially favouring one type of error over the other). It is also common to evaluate on the basis of the mass fraction of valuable material recovered. However, to accurately compare the performance of multiple models against each other in this sense would require large numbers of samples. So in the following, the evaluation will be carried out on the pixel level, i.e. comparing the class assigned to each pixel against the actual class that the sample belongs to. This allows for finer comparisons, as most high-performing models perform almost identically at the sample level given the small number of samples. Here, the positive class is the valuable mineral (lepidolite) and the negative class is the waste (granite).

To score the performance of each classifier, a confusion matrix was constructed with true positives (TP), false negatives (FN), false positives (FP) and true negatives (TN). From this, the following measures were calculated for each model across all scans (n being the total number of pixels classified):

- The portion of correctly classified pixels:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{n} \quad (7.1)$$

- The portion of true positive pixels among all positively classified pixels:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (7.2)$$

- The portion of correctly classified pixels among all positive pixels:

$$\text{Sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (7.3)$$

- The portion of correctly classified pixels among all negative pixels:

$$\text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}} \quad (7.4)$$

- The geometric mean of precision and sensitivity (F_1 score). This is a popular measure of the performance of binary classifiers, however it has been criticised in recent times^[6] for the potential for misleading results:

$$F_1 = 2 \times \frac{\text{Precision} \times \text{Sensitivity}}{\text{Precision} + \text{Sensitivity}} \quad (7.5)$$

- The Matthews correlation coefficient, considered to be a more reliable indicator of binary classification performance^[6]. Ranges from -1 to 1, unlike the other measures:

$$\text{MCC} = \frac{(\text{TP} \times \text{TN}) - (\text{FP} \times \text{FN})}{\sqrt{(\text{TP} + \text{FP}) \times (\text{TP} + \text{FN}) \times (\text{TN} + \text{FP}) \times (\text{TN} + \text{FN})}} \quad (7.6)$$

Tables with confusion matrices, accuracy, precision, sensitivity, specificity, F_1 score and MCC can be found for the training set in [Appendix A](#) and for the test set in [Appendix B](#). In the following, MCC will be used as an overall score to judge relative performance of all classifiers.

The performance of the RGB photo model and of the best-performing spectral and abundance-based models will be shown in more detail, with a visual representation of the class assigned to each pixel (green for lepidolite, red for granite), as well as a violin plot with the class probabilities assigned to each pixel broken down by scan.

The following abbreviations were used to identify each LDA model:

- RGB (p): fitted on high-resolution RGB photos shown in [Figure 3.3](#), as a benchmark representing conventional colour sensors
- RGB (s): fitted on RGB images reconstructed from the hyperspectral scans, as an RGB benchmark at the same resolution as the other spectral/abundance-based models, to allow correction for resolution as a confounding variable
- Spectral: fitted directly on the full spectra of the pipeline, as a benchmark of typical approaches found in the literature
- Abundance-based methods:
 - U/U: fitted on UCLS-generated abundances, with UCLS unmixing during inference, the main proposed real-time method
 - F/F: fitted on FCLS-generated abundances, with FCLS unmixing during inference, as a baseline to investigate whether there could be a benefit to improving the performance of FCLS unmixing to the point of being able to use it in real time
 - F/U: fitted on FCLS-generated abundances, with UCLS unmixing during inference, to investigate whether FCLS is worth using for fitting a model even when using UCLS during inference

- U/F: fitted on UCLS-generated abundances, with FCLS unmixing during inference, unlikely to be useful, but investigated for symmetry
- A star (“*”) appended to a model name means that this model was fitted while ignoring label marker endmembers (see [Section 5.4.5](#)).

7.1 Inference on training set




Figure 7.1: Matthew correlation coefficient for all classification models across all pipelines under consideration for the training set of samples. Spectral methods are shown in purple, RGB methods in orange, abundance-based methods in green.

On the training set (i.e., the set of samples used to fit all classifiers), most models performed well across the board and across most pipelines (see [Figure 7.1](#)). This is unsurprising, as it is mostly a test of how well the training data can be separated into two classes, instead of a truly informative comparison of different classifiers. Only the features pipelines apparently presented more of a challenge to fit some models. The clear best performer across all pipelines is the direct spectral model, i.e. the model introduced as a literature baseline.

The RGB models performed relatively poorly in comparison to spectral and abundance-based models, with the high-resolution photo-based RGB model performing slightly worse than the reconstructed RGB-based model.

Among abundance-based models, certain trends are immediately apparent:

- UCLS-based models perform slightly better than FCLS-based models.
- Models ignoring the marker endmembers perform slightly worse than their standard counterparts.
- Mixing UCLS and FCLS between training and inference performs worse than both pure FCLS and pure UCLS-models.

These will be discussed in [Section 7.3.3](#).

7.1.1 RGB photo classifier




Figure 7.2: Inference results across the training set for the RGB photo classifier. The green channel of each pixel contains the assigned probability of the lepidolite class, and the red channel the probability of granite.




Figure 7.3: Violin plots of the class probabilities across all training pixels broken down by scan for the RGB photo classifier. The circular marker represents the mean probability and is coloured green/red for correct/incorrect overall assignment.

The RGB photo-based classifier produces very high-resolution inference results (see [Figure 7.2](#)), however none of them have a high confidence. Most pixels are classified with a probability not very far from the threshold of 0.5. Also, there are clear signs of it overfitting on the marker labels, assigning both the pink and the purple marker to the lepidolite class. These are the most confidently classified pixels, though as can be seen from the upper tips of the violin plots (see [Figure 7.3](#)), they only make up a very small portion of each scan.

7.1.2 Best-performing direct spectral classifier




Figure 7.4: Inference results across the training set for the reflectance spectral classifier. The green channel of each pixel contains the assigned probability of the lepidolite class, and the red channel the probability of granite.




Figure 7.5: Violin plots of the class probabilities across all training pixels broken down by scan for the reflectance spectral classifier. The circular marker represents the mean probability and is coloured green/red for correct/incorrect overall assignment.

The best-performing direct spectral classifier (fitted on the reflectance pipeline) is also the best-performing classifier overall on the training set. Figures 7.4 and 7.5 show that all

scans apart from PL9-B were correctly classified, and typically with a high confidence. The misclassification of PL9-B is unsurprising given that this side of this sample is expected to cause a false-negative error, as noted in [Chapter 3](#) and visible in [Figure 3.3](#); this will be discussed in [Section 7.3](#).

7.1.3 Best-performing abundance-based classifier




Figure 7.6: Inference results across the training set for the reflectance U/U classifier. The green channel of each pixel contains the assigned probability of the lepidolite class, and the red channel the probability of granite.




Figure 7.7: Violin plots of the class probabilities across all training pixels broken down by scan for the reflectance U/U classifier. The circular marker represents the mean probability and is coloured green/red for correct/incorrect overall assignment.

The best-performing abundance-based classifier is the U/U model of the reflectance pipeline, i.e. a model fitted on UCLS-generated abundances where UCLS was also used for unmixing during inference. As in the reflectance spectral classifier ([Section 7.1.2](#)), all scans apart from PL9-B were classified correctly (see [Figure 7.7](#)). However, some scans (especially PG7-A, PG7-B, PL11-B) show a little more ambiguity than with the direct spectral model (see [Figure 7.6](#)). This may be due to these sample surfaces being dominated by feldspar minerals that are common to both classes and therefore carry less discriminative power.

7.2 Inference on test set




Figure 7.8: Matthew correlation coefficient for all classification models across all pipelines under consideration for the test set of samples. Spectral methods are shown in purple, RGB methods in orange, abundance-based methods in green.

On the test set (i.e., the set of samples unknown to the fitted models), inference performance fluctuated a lot more than on the training set (see [Figure 7.8](#)). The previous best performer (the spectral model of the reflectance pipeline) was now the worst performer in its pipeline. Only in the reflectance continua pipeline did the direct spectral model maintain a relatively high MCC score. The two best-performing models overall on the test set were U/U (reflectance) and U/U (reflectance continua). The RGB models retained relatively moderate scores.

All models fitted on the absorbance pipelines performed fairly poorly here, with some models even producing inferences with a negative correlation coefficient. The reflectance features pipeline also did not produce very good inference results on the test set.

Among the abundance-based models, the same general trends found in the training set results tend to hold here.

7.2.1 RGB photo classifier




Figure 7.9: Inference results across the test set for the RGB photo classifier. The green channel of each pixel contains the assigned probability of the lepidolite class, and the red channel the probability of granite.




Figure 7.10: Violin plots of the class probabilities across all test pixels broken down by scan for the RGB photo classifier. The circular marker represents the mean probability and is coloured green/red for correct/incorrect overall assignment.

The inference results for the RGB photo classifier on the test set (see Figures 7.10 and 7.9) shows similarly unconfident but often-correct results as on the training set. All granite

samples were classified correctly, however four lepidolite scans (PL13-B, PL14-B, PL16-A, PL16-B) were incorrectly classified as granite.

7.2.2 Best-performing direct spectral classifier from training set




Figure 7.11: Inference results across the test set for the reflectance spectral classifier. The green channel of each pixel contains the assigned probability of the lepidolite class, and the red channel the probability of granite.




Figure 7.12: Violin plots of the class probabilities across all test pixels broken down by scan for the reflectance spectral classifier. The circular marker represents the mean probability and is coloured green/red for correct/incorrect overall assignment.

The reflectance spectral model, which performed the best on the training set, shows unconvincing performance on the test set (see [Figure 7.12](#)). While only two samples were incorrectly classified (PL13-B, PL16-B), a lot of pixels were incorrectly classified with high confidence, resulting in speckled, noisy-looking inference images (see [Figure 7.11](#)).

7.2.3 Best-performing direct spectral classifier




Figure 7.13: Inference results across the test set for the reflectance continua spectral classifier. The green channel of each pixel contains the assigned probability of the lepidolite class, and the red channel the probability of granite.




Figure 7.14: Violin plots of the class probabilities across all test pixels broken down by scan for the reflectance continua spectral classifier. The circular marker represents the mean probability and is coloured green/red for correct/incorrect overall assignment.

The best-performing spectral classifier was the one fitted on the reflectance continua pipeline. Like the reflectance spectral classifier, it misclassified two lepidolite scans (PL13-B and PL16-B, see [Figure 7.14](#)), however the inference results were a lot more consistent and even across the pixels (see [Figure 7.13](#)).

7.2.4 Best-performing abundance-based classifier




Figure 7.15: Inference results across the test set for the reflectance U/U classifier. The green channel of each pixel contains the assigned probability of the lepidolite class, and the red channel the probability of granite.




Figure 7.16: Violin plots of the class probabilities across all test pixels broken down by scan for the reflectance U/U classifier. The circular marker represents the mean probability and is coloured green/red for correct/incorrect overall assignment.

The overall highest-performing model (by MCC score) on the test set was the reflectance U/U classifier. It performed similarly well on the test set as it did on the training set, with two

misclassifications (PL13-B and PL16-B, see Figure 7.16) and one somewhat ambiguous scan (PL14-B, see Figure 7.15). The slightly higher MCC score compared to the reflectance continua spectral model is likely explained by better classification confidence across the granite samples.

7.2.5 Best-performing abundance-based classifier with marker labels ignored




Figure 7.17: Inference results across the test set for the reflectance U/U^* classifier. The green channel of each pixel contains the assigned probability of the lepidolite class, and the red channel the probability of granite.




Figure 7.18: Violin plots of the class probabilities across all test pixels broken down by scan for the reflectance U/U^* classifier. The circular marker represents the mean probability and is coloured green/red for correct/incorrect overall assignment.

The U/U^* model, which was fitted while excluding marker endmembers (see [Section 5.4.5](#)) did not perform significantly worse than the top-performing U/U model (see [Figures 7.17](#) and [7.18](#)). The main impact was on increasing ambiguity in the inference of the PL14-A and PL14-B scans.

7.3 Discussion

Noticeably, all of the best-performing classifiers consistently misclassified the same lepidolite scans (PL9-B, PL13-B, PL16-B) as granite. As discussed in [Chapter 3](#), these scans are expected to cause such false-negatives errors. In fact, visual inspection of these lepidolite samples reveals that their “B” sides contain fairly little recognisable lepidolite, as can be seen in the corresponding photos in [Figures 3.3](#) and [3.4](#). PL9 and PL13 likely originate from a boundary between lepidolite and granite layers, whereas the low abundance of lepidolite in the middle of PL16-B appears to be due to localised displacement by another mineral (albite). In this context, these misclassifications should not be viewed as faults of the classification models themselves, but rather as a result of the limitations of a surface-based imaging system that can only view one surface of a sample at a time. Scanning samples from multiple angles simultaneously should reduce errors of this kind, however there remains the fundamental limitation of optical measurements that they are unable to probe the bulk of the rock, so some misclassifications will always be unavoidable with such methods.

7.3.1 RGB models

The RGB models generalised fairly well from the training set to the test set, however they suffered overall from low confidence inference. This is not surprising, given the only three-dimensional input.

The additional spatial resolution of the RGB photo model did not result in an edge over the low-resolution spectrally-sampled RGB model. In fact, the photo model performed slightly worse on both the training set and the test set.

7.3.2 Direct spectral models

It is striking that the best-performing classifier on the training set, the reflectance spectral model, performed so much worse on the test set. This is a sign of overfitting, which LDA can be prone to with high-dimensional input^[16]. Given that the reflectance continua spectral model performed better, it is likely that with regularisation or a different classification algorithm altogether, the performance of the reflectance spectral model could be increased.

7.3.3 Abundance-based models

The abundance-based methods generalised well from the training set to the test set. Due to the much lower-dimensional component space, they do not require any regularisation and are unlikely to overfit.

The ability to exclude arbitrary endmembers from being used for fitting, which is a unique benefit of abundance-based techniques, did not have a significant negative impact on the performance of these models here.

It is interesting that the UCLS-based models generally performed better than the FCLS-based models, even though FCLS incorporates more physically correct constraints. This may be partly due to the unconstrained method being more robust to noise, but also due to LDA assuming normally-distributed input. The output of UCLS follows a multivariate normal distribution, whereas FCLS generates abundances that are much better modelled by a Dirichlet distribution. This also results in models that mix UCLS and FCLS performing badly compared to their pure counterparts.

Chapter 8

Conclusion

8.1 Summary

The proposed method for the first time applies the established remote sensing techniques of virtual dimensionality estimation, endmember extraction and abundance map generation in conjunction with linear discriminant analysis to produce abundance-based classification models capable of real-time differentiation between two classes of mineral samples for the purposes of ore sorting. The models were tested on lithium-bearing lepidolite samples and barren granite samples from Gonçalo, Portugal.

A push-broom spectral imaging system was designed to capture reflectance spectral images of samples in a scanning fashion intended to be compatible with typical conveyor setups found in the mining industry. Slow (50 min) high-SNR reflectance scans were captured of a set of samples designated for fitting models, and faster (10s) scans were captured of another set of samples for inference validation.

Several variations of the proposed classification method were explored on the basis of transforming reflectance into absorbance, and by performing continuum extraction to split spectra into continuum spectra and feature spectra. Models fitted on absorbance data did not perform well on test samples. Continuum spectra (which are commonly disregarded in the literature) produced fairly reliable models, whereas feature spectra (which are commonly used for analysis) did not.

The most straightforward variation using unmodified reflectance spectra exhibited the most reliable inference performance, and outperformed baseline LDA models fitted directly on reflectance spectra or on RGB images. In particular, it generalised more reliably from the training samples to the test samples than the reflectance spectral model (which suffered from overfitting), and provided higher-confidence inferences than the RGB photo model.

As an additional benefit of transforming the spectral data into a lower-dimensional space of spectral components, individual endmembers can be explicitly suppressed when fitting a

model. This was demonstrated on the example of contamination by markers which were used to write labels on the samples. This kind of manual intervention is entirely optional, as the fitting process can in principle be completely automated, only requiring two known sets of samples to be distinguished.

Two common unmixing (abundance map generation) algorithms were compared in the context of fitting binary classification models: unconstrained least squares (UCLS), which may produce unphysical abundance values but is certainly fast enough for real-time use, and fully-constrained least squares (FCLS), which imposes realistic constraints but is slower. UCLS-based models consistently performed slightly better in inference, which may be due to the normally-distributed statistics of their output matching the assumptions made by LDA on its input.

8.2 Industrial applicability

As described in [Section 6.1](#), as a rough estimate, an industrially valuable ore sorting system should be able to process a material feed of approximately 3 m/s (assuming a particle size of 10 cm), whereas the “real-time” laboratory scans were captured at a scan speed of 6 mm/s. Possible ways to bridge this discrepancy include:

- replacing the 10 W halogen light source with a higher intensity bulb (e.g. 100 W or 1000 W are commonly available) for correspondingly shorter exposure times and hence faster scans,
- acceptance of higher noise levels caused by further reductions in exposure time,
- use of a more sensitive sensor,
- installation of multiple sorting systems in parallel.

A combination of these options should be sufficient to make the throughput of this sorting system viable for a given industrial application.

[Section 7.3](#) highlighted an issue of potential misclassification of samples that present different minerals on different sides. In sorting applications where false negatives are particularly undesirable, this could be mitigated by simply scanning samples with multiple separate spectrometers pointed at different sides, and potentially inducing the samples to roll while being conveyed (which is best left to engineers).

8.3 Further work

While the imaging spectrometer was custom-designed for the purposes of prototyping ore sorting methods, it was far from ideal. Its optics exhibited chromatic distortion which tended to produce curved projections of spectra (known as “smiles”), in particular near the edges of the frame. The spectral resolution was also not constant across the frame. Low sensitivity of the system in the ultraviolet and violet parts of the spectrum resulted in unacceptably low SNR there. The imaging sensor had unnecessarily high resolution and no cooling. Unidentified dark horizontal line artifacts were also prominent in raw captures, although they were easily corrected for. A professionally designed and calibrated push-broom spectral imaging system would surely yield better results. It would also be interesting to apply the proposed method to other wavelength ranges such as NIR or SWIR, and to fluorescence in place of (or in addition to) reflectance.

The method was evaluated on the basis of samples from only one orebody. It would be highly important to test its generalisability by applying it to more diverse types of ore. Of particular interest would be orebodies containing minerals with sharper spectral features in the wavelength range of the sensor.

There are newer techniques developed in recent years in the field of remote sensing which are attractive for ore sorting applications. More advanced endmember extraction techniques could be applied, such as ACEE^[33], which takes spatial information into account. Multi-sensor data could be combined such as in Ref.^[17] to improve spatial resolution of spectral images.

Further incorporation of external information into abundance-based classification models could be attempted. Perhaps the contribution of individual high-value endmembers could be boosted, analogously as contaminant endmembers were suppressed in this work.

Here, classification based on original spectra, continuum spectra and feature spectra was compared in separate pipelines. It may be worth merging two or all of these types of spectra into one higher-dimensional feature space and perform classification on this basis.

Apart from continuum extraction/removal and reflectance transformation to absorbance, no other typical spectral processing techniques were investigated in the context of abundance-based classification. For example, multiplicative scatter correction^[13] could be applied to compensate for variability in particle size distributions.

In this study, only linear discriminant analysis was used for simplicity. Many other binary classification methods are available and should be compared rigorously in this context given the statistics of abundance data.

It may be possible to relax the labelling condition imposed on the input, which is that

two separate known classes of samples are prepared for fitting. Using the presented methods of creating a spectral library on a single unlabelled set of samples and generating abundance maps for them, the resulting abundance values (or their fitted distribution parameters) could be automatically clustered into two (or more) distinct groups which could then be used to fit a classifier.

The binary classification of pixels/samples could be replaced with a regression of the concentration of the desired material in a sample. This would require the scanning of samples of known (and varying) concentrations, most likely to be determined using destructive methods.

Abundances can be reasonably seen to follow multivariate normal distributions (in the case of UCLS unmixing) or Dirichlet distributions (in the case of FCLS unmixing). Fitting the parameters of these distributions (for estimation of Dirichlet parameters see Minka^[18] and for goodness-of-fit tests see Li^[15]) would allow the generation of unlimited amounts of synthetic pixels, which could be used to train more complicated machine learning systems such as deep neural networks.

The presented abundance-based classification methods can be thought of as implementing a dimensionality reduction technique (from 448 spectral channels to 30 spectral components), similar to the way that neural networks typically incorporate a lower-dimensional latent space representation of the input data. It might be an interesting question whether a deep neural network could be trained on simulated arbitrary mineral mixtures to perform this dimensionality reduction in an unsupervised way, replacing the step of endmember extraction and automatically generating a spectral library for any given orebody.

References

- [1] R. Bro and S. D. Jong. A fast non-negativity-constrained least squares algorithm. *Journal of Chemometrics*, 11(5):393–401, 1997. ISSN 0886-9383. doi:[10.1002/\(sici\)1099-128x\(199709/10\)11:5<393::aid-cem483>3.0.co;2-l](https://doi.org/10.1002/(sici)1099-128x(199709/10)11:5<393::aid-cem483>3.0.co;2-l).
- [2] J. Cardoso-Fernandes, J. Silva, M. M. Perrotta, A. Lima, A. C. Teodoro, M. A. Ribeiro, F. Dias, O. Barrès, J. Cauzid, and E. Roda-Robles. Interpretation of the Reflectance Spectra of Lithium (Li) Minerals and Pegmatites: A Case Study for Mineralogical and Lithological Identification in the Fregeneda-Almendra Area. *Remote Sensing*, 13(18):3688, 2021. doi:[10.3390/rs13183688](https://doi.org/10.3390/rs13183688).
- [3] T.-H. Chan, W.-K. Ma, A. Ambikapathi, and C.-Y. Chi. A Simplex Volume Maximization Framework for Hyperspectral Endmember Extraction. *IEEE Transactions on Geoscience and Remote Sensing*, 49(11):4177–4193, 2011. ISSN 0196-2892. doi:[10.1109/tgrs.2011.2141672](https://doi.org/10.1109/tgrs.2011.2141672).
- [4] C.-I. Chang and Q. Du. Estimation of Number of Spectrally Distinct Signal Sources in Hyperspectral Imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 42(3):608–619, 2004. ISSN 0196-2892. doi:[10.1109/tgrs.2003.819189](https://doi.org/10.1109/tgrs.2003.819189).
- [5] C.-I. Chang, C.-C. Wu, W.-M. Liu, and Y.-C. Ouyang. A New Growing Method for Simplex-Based Endmember Extraction Algorithm. *IEEE Transactions on Geoscience and Remote Sensing*, 44(10):2804–2819, 2006. ISSN 0196-2892. doi:[10.1109/tgrs.2006.881803](https://doi.org/10.1109/tgrs.2006.881803).
- [6] D. Chicco and G. Jurman. The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, 21(1):6, 2020. doi:[10.1186/s12864-019-6413-7](https://doi.org/10.1186/s12864-019-6413-7).
- [7] R. N. Clark and T. L. Roush. Reflectance spectroscopy: Quantitative analysis techniques for remote sensing applications. *Journal of Geophysical Research: Solid Earth*, 89(B7):6329–6340, 1984. ISSN 0148-0227. doi:[10.1029/jb089ib07p06329](https://doi.org/10.1029/jb089ib07p06329).

- [8] M. Dalm, M. W. Buxton, F. J. v. Ruitenbeek, and J. H. Voncken. Application of near-infrared spectroscopy to sensor based sorting of a porphyry copper ore. *Minerals Engineering*, 58:7–16, 2014. ISSN 0892-6875. doi:[10.1016/j.mineng.2013.12.016](https://doi.org/10.1016/j.mineng.2013.12.016).
- [9] J. B. Fisher, B. Lee, A. J. Purdy, G. H. Halverson, M. B. Dohlen, K. Cawse-Nicholson, A. Wang, R. G. Anderson, B. Aragon, M. A. Arain, D. D. Baldocchi, J. M. Baker, H. Barral, C. J. Bernacchi, C. Bernhofer, S. C. Biraud, G. Bohrer, N. Brunsell, B. Cappelaere, S. Castro-Contreras, J. Chun, B. J. Conrad, E. Cremonese, J. Demarty, A. R. Desai, A. D. Ligne, L. Foltýnová, M. L. Goulden, T. J. Griffis, T. Grünwald, M. S. Johnson, M. Kang, D. Kelbe, N. Kowalska, J. Lim, I. Mainassara, M. F. McCabe, J. E. Missik, B. P. Mohanty, C. E. Moore, L. Morillas, R. Morrison, J. W. Munger, G. Posse, A. D. Richardson, E. S. Russell, Y. Ryu, A. Sanchez-Azofeifa, M. Schmidt, E. Schwartz, I. Sharp, L. Šigut, Y. Tang, G. Hulley, M. Anderson, C. Hain, A. French, E. Wood, and S. Hook. ECOSTRESS: NASA’s Next Generation Mission to Measure Evapotranspiration From the International Space Station. *Water Resources Research*, 56(4), 2020. ISSN 0043-1397. doi:[10.1029/2019wr026058](https://doi.org/10.1029/2019wr026058).
- [10] G. L. Genn. *Novel Techniques in Ore Characterisation and Sorting*. PhD thesis, The University of Queensland, Australia, 09 2013.
- [11] R. P. Gupta. *Remote Sensing Geology*. Springer, 2018. ISBN 9783662558744. doi:[10.1007/978-3-662-55876-8](https://doi.org/10.1007/978-3-662-55876-8).
- [12] M. He and S. Mei. Fully constrained oblique projection approach to mixed pixel linear unmixing. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2008.
- [13] T. Isaksson and T. Næs. The Effect of Multiplicative Scatter Correction (MSC) and Linearity Improvement in NIR Spectroscopy. *Applied Spectroscopy*, 42(7):1273–1284, 1988. ISSN 0003-7028. doi:[10.1366/0003702884429869](https://doi.org/10.1366/0003702884429869).
- [14] R. F. Kokaly, R. N. Clark, G. A. Swayze, K. E. Livo, and T. M. Hoefen. USGS spectral library version 7. *U.S. Geological Survey Data Series*, 2017.
- [15] Y. Li. *Goodness-of-fit Tests for Dirichlet Distributions with Applications*. PhD thesis, Bowling Green State University, 2015.
- [16] R. Liu and D. F. Gillies. Overfitting in linear feature extraction for classification of high-dimensional image data. *Pattern Recognition*, 53:73–86, 2016. ISSN 0031-3203. doi:[10.1016/j.patcog.2015.11.015](https://doi.org/10.1016/j.patcog.2015.11.015).

- [17] S. Lorenz, P. Seidel, P. Ghamisi, R. Zimmermann, L. Tusa, M. Khodadadzadeh, I. C. Contreras, and R. Gloaguen. Multi-Sensor Spectral Imaging of Geological Samples: A Data Fusion Approach Using Spatio-Spectral Feature Extraction. *Sensors*, 19(12):2787, 2019. doi:[10.3390/s19122787](https://doi.org/10.3390/s19122787).
- [18] T. P. Minka. Estimating a Dirichlet distribution. Technical report, Massachusetts Institute of Technology, 2000.
- [19] E. G. Modise, A. M. Zungeru, B. Mtengi, and A. U. Ude. Sensor-Based Ore Sorting—A Review of Current Use of Electromagnetic Spectrum in Sorting. *IEEE Access*, 10:112307–112326, 2022. ISSN 2169-3536. doi:[10.1109/access.2022.3216296](https://doi.org/10.1109/access.2022.3216296).
- [20] H. Petersson and D. Gustafsson. Hyperspectral image analysis using deep learning—a review. *ieeexplore.ieee.org*, pages 1 – 6, 2016. doi:[10.1109/ipta.2016.7820963](https://doi.org/10.1109/ipta.2016.7820963).
- [21] D. Peukert, C. Xu, and P. Dowd. A Review of Sensor-Based Sorting in Mineral Processing: The Potential Benefits of Sensor Fusion. *Minerals*, 12(11):1364, 2022. doi:[10.3390/min12111364](https://doi.org/10.3390/min12111364).
- [22] J. Pietilä and O. Haavisto. Reflectance Spectrum Based Classification of Ore. *IFAC Proceedings Volumes*, 42(23):173–178, 2009. ISSN 1474-6670. doi:[10.3182/20091014-3-cl-4011.00032](https://doi.org/10.3182/20091014-3-cl-4011.00032).
- [23] A. Plaza and C.-I. Chang. Impact of Initialization on Design of Endmember Extraction Algorithms. *IEEE Transactions on Geoscience and Remote Sensing*, 44(11):3397–3407, 2006. ISSN 0196-2892. doi:[10.1109/tgrs.2006.879538](https://doi.org/10.1109/tgrs.2006.879538).
- [24] H. Ren and C.-I. Chang. Automatic spectral target recognition in hyperspectral imagery. *IEEE Transactions on Aerospace and Electronic Systems*, 39(4):1232–1249, 2003. ISSN 0018-9251. doi:[10.1109/taes.2003.1261124](https://doi.org/10.1109/taes.2003.1261124).
- [25] R. H. Riffenburgh. *Linear discriminant analysis*. PhD thesis, Virginia Polytechnic Institute, 1957.
- [26] C. Robben and H. Wotrub. Sensor-Based Ore Sorting Technology in Mining—Past, Present and Future. *Minerals*, 9(9):523, 2019. doi:[10.3390/min9090523](https://doi.org/10.3390/min9090523).
- [27] B. Somers, M. Zortea, A. Plaza, and G. P. Asner. Automated Extraction of Image-Based Endmember Bundles for Improved Spectral Unmixing. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2):396–408, 2012. ISSN 1939-1404. doi:[10.1109/jstars.2011.2181340](https://doi.org/10.1109/jstars.2011.2181340).

- [28] R. Sousa, V. Ramos, A. Guedes, F. Noronha, A. B. d. Sousa, M. M. Leite, R. Seltmann, and A. Dolgoplova. The Alvarrões-Gonçalo Li project: an example of sustainable lithium mining. *Advances in Geosciences*, 45:1–5, 2018. doi:[10.5194/adgeo-45-1-2018](https://doi.org/10.5194/adgeo-45-1-2018).
- [29] C. Thérien. PySptools 0.15.0, 2018. URL <https://pysptools.sourceforge.io>.
- [30] Various contributors. SciPy 1.5.4. URL <https://scipy.org>.
- [31] F. L. Wilke, F. X. Spachtholz, J. Q. Rogado, G. Brito, J. Carolino, A. Carolino, and B. Schimm. Operational Conditions for Continuous Mining Systems in Hard Rock Open Pit Mines. *BRITE/EURAM Programme Synthesis Report*, 1997. URL https://cordis.europa.eu/docs/projects/files/BRE/BRE20345/26840791-6_en.pdf.
- [32] M. E. Winter. N-FINDR: an algorithm for fast autonomous spectral end-member determination in hyperspectral data. *Imaging Spectrometry V*, pages 266–275, 1999. ISSN 0277-786X. doi:[10.1117/12.366289](https://doi.org/10.1117/12.366289).
- [33] M. Xu, B. Du, and L. Zhang. Spatial-Spectral Information Based Abundance-Constrained Endmember Extraction Methods. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(6):2004–2015, 2014. ISSN 1939-1404. doi:[10.1109/jstars.2013.2268661](https://doi.org/10.1109/jstars.2013.2268661).
- [34] H. Zhao and X. Zhao. Nonlinear unmixing of minerals based on the log and continuum removal model. *European Journal of Remote Sensing*, 52(1):277–293, 2019. doi:[10.1080/22797254.2019.1601999](https://doi.org/10.1080/22797254.2019.1601999).
- [35] H.-q. Zhao, X.-s. Zhao, Y. Cen, and H. Yang. Research on the Impact of Absorption Feature Extraction on Spectral Difference Between Similar Minerals. *Guang Pu Xue Yu Guang Pu Fen Xi/Spectroscopy and Spectral Analysis*, 37(3):869–74, 2017. ISSN 1000-0593.

Appendix A

Inference results on training set

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (p)	12687360	7194280	0.78	0.83	0.64	0.90	0.72	0.56
	2649808	22878708						
RGB (s)	29504	9459	0.82	0.84	0.76	0.88	0.80	0.64
	5523	39744						
Spectral	33528	5435	0.92	0.95	0.86	0.96	0.90	0.83
	1656	43611						
F/F	30705	8258	0.85	0.88	0.79	0.91	0.83	0.71
	4123	41144						
F/U	31955	7008	0.86	0.86	0.82	0.89	0.84	0.71
	5016	40251						
U/F	24336	14627	0.81	0.94	0.62	0.96	0.75	0.64
	1618	43649						
U/U	32108	6855	0.89	0.93	0.82	0.95	0.87	0.78
	2345	42922						
F/F*	29835	9128	0.81	0.82	0.77	0.85	0.79	0.62
	6612	38655						
F/U*	29698	9265	0.73	0.69	0.76	0.70	0.72	0.46
	13511	31756						
U/F*	20494	18469	0.73	0.84	0.53	0.91	0.65	0.48
	3925	41342						
U/U*	31369	7594	0.86	0.89	0.81	0.91	0.84	0.72
	4045	41222						

Table A.1: Inference results on training set in the reflectance pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	29566	9397	0.79	0.78	0.76	0.82	0.77	0.58
	8266	37001						
Spectral	32304	6659	0.88	0.91	0.83	0.93	0.87	0.76
	3307	41960						
F/F	31368	7595	0.84	0.83	0.81	0.86	0.82	0.67
	6284	38983						
F/U	30164	8799	0.81	0.80	0.77	0.83	0.79	0.61
	7551	37716						
U/F	21399	17564	0.74	0.84	0.55	0.91	0.66	0.50
	4036	41231						
U/U	31758	7205	0.88	0.91	0.82	0.93	0.86	0.75
	3303	41964						
F/F*	31009	7954	0.83	0.83	0.80	0.86	0.81	0.66
	6362	38905						
F/U*	31248	7715	0.83	0.83	0.80	0.86	0.82	0.66
	6390	38877						
U/F*	21429	17534	0.74	0.83	0.55	0.90	0.66	0.49
	4526	40741						
U/U*	31842	7121	0.87	0.89	0.82	0.91	0.85	0.74
	3940	41327						

Table A.2: Pixel-level inference results on training set in the reflectance continua pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	21218	17745	0.77	0.92	0.54	0.96	0.68	0.56
	1935	43332						
Spectral	32825	6138	0.89	0.92	0.84	0.93	0.88	0.78
	2946	42321						
F/F	22233	16730	0.76	0.87	0.57	0.93	0.69	0.54
	3177	42090						
F/U	22568	16395	0.57	0.53	0.58	0.56	0.56	0.14
	19701	25566						
U/F	12180	26783	0.67	0.94	0.31	0.98	0.47	0.41
	806	44461						
U/U	29471	9492	0.85	0.91	0.76	0.94	0.83	0.71
	2911	42356						
F/F*	22204	16759	0.76	0.85	0.57	0.92	0.68	0.52
	3838	41429						
F/U*	19970	18993	0.55	0.51	0.51	0.58	0.51	0.10
	18840	26427						
U/F*	13942	25021	0.68	0.90	0.36	0.97	0.51	0.42
	1536	43731						
U/U*	29354	9609	0.83	0.86	0.75	0.89	0.80	0.65
	4970	40297						

Table A.3: Pixel-level inference results on training set in the reflectance features pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	27107	11789	0.78	0.79	0.70	0.84	0.74	0.55
	7077	38111						
Spectral	34231	4665	0.91	0.92	0.88	0.93	0.90	0.82
	3003	42185						
F/F	32658	6238	0.84	0.82	0.84	0.85	0.83	0.69
	6928	38260						
F/U	31283	7613	0.81	0.78	0.80	0.81	0.79	0.61
	8709	36479						
U/F	32553	6343	0.83	0.80	0.84	0.82	0.82	0.66
	7911	37277						
U/U	32746	6150	0.89	0.91	0.84	0.92	0.87	0.77
	3408	41780						
F/F*	32569	6327	0.83	0.80	0.84	0.82	0.82	0.66
	7999	37189						
F/U*	32308	6588	0.73	0.67	0.83	0.64	0.74	0.48
	16130	29058						
U/F*	30008	8888	0.80	0.78	0.77	0.82	0.78	0.59
	8291	36897						
U/U*	33195	5701	0.87	0.86	0.85	0.88	0.85	0.73
	5587	39601						

Table A.4: Pixel-level inference results on training set in the absorbance pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	28814	10082	0.75	0.72	0.74	0.75	0.73	0.49
	11238	33950						
Spectral	32512	6384	0.87	0.88	0.84	0.90	0.86	0.74
	4515	40673						
F/F	30308	8588	0.77	0.74	0.78	0.77	0.76	0.55
	10463	34725						
F/U	30572	8324	0.78	0.75	0.79	0.77	0.77	0.56
	10207	34981						
U/F	23439	15457	0.71	0.71	0.60	0.79	0.65	0.40
	9347	35841						
U/U	31951	6945	0.87	0.88	0.82	0.91	0.85	0.73
	4158	41030						
F/F*	30323	8573	0.77	0.74	0.78	0.77	0.76	0.54
	10583	34605						
F/U*	31088	7808	0.79	0.76	0.80	0.78	0.78	0.58
	9825	35363						
U/F*	15098	23798	0.65	0.72	0.39	0.87	0.51	0.30
	5733	39455						
U/U*	30975	7921	0.86	0.89	0.80	0.91	0.84	0.72
	3943	41245						

Table A.5: Pixel-level inference results on training set in the absorbance continua pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	18379	20517	0.75	0.95	0.47	0.98	0.63	0.54
	888	44300						
Spectral	32388	6508	0.88	0.90	0.83	0.92	0.87	0.76
	3472	41716						
F/F	16141	22755	0.72	0.95	0.41	0.98	0.58	0.49
	886	44302						
F/U	31362	7534	0.56	0.51	0.81	0.34	0.63	0.17
	29720	15468						
U/F	4047	34849	0.58	0.93	0.10	0.99	0.19	0.22
	291	44897						
U/U	23825	15071	0.79	0.91	0.61	0.95	0.73	0.61
	2225	42963						
F/F*	16011	22885	0.71	0.91	0.41	0.96	0.57	0.46
	1600	43588						
F/U*	36971	1925	0.52	0.49	0.95	0.15	0.65	0.16
	38599	6589						
U/F*	922	37974	0.54	0.76	0.02	0.99	0.05	0.07
	286	44902						
U/U*	24611	14285	0.79	0.89	0.63	0.93	0.74	0.60
	3092	42096						

Table A.6: Pixel-level inference results on training set in the absorbance features pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Appendix B

Inference results on test set

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (p)	5170576	5827182	0.68	0.82	0.47	0.89	0.60	0.40
	1136944	9498246						
RGB (s)	1792	2667	0.73	0.85	0.40	0.95	0.54	0.44
	328	6148						
Spectral	2559	1900	0.66	0.59	0.57	0.73	0.58	0.30
	1778	4698						
F/F	2333	2126	0.80	0.96	0.52	0.99	0.68	0.60
	86	6390						
F/U	2515	1944	0.77	0.82	0.56	0.92	0.67	0.53
	542	5934						
U/F	1465	2994	0.72	0.99	0.33	1.00	0.49	0.47
	16	6460						
U/U	2789	1670	0.83	0.95	0.63	0.98	0.75	0.67
	150	6326						
F/F*	2035	2424	0.77	0.93	0.46	0.98	0.61	0.53
	144	6332						
F/U*	2721	1738	0.66	0.58	0.61	0.70	0.60	0.31
	1953	4523						
U/F*	892	3567	0.67	0.99	0.20	1.00	0.33	0.35
	11	6465						
U/U*	2574	1885	0.80	0.90	0.58	0.95	0.70	0.59
	292	6184						

Table B.1: Pixel-level inference results on test set in the reflectance pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	1849	2610	0.72	0.79	0.41	0.92	0.54	0.40
	502	5974						
Spectral	3073	1386	0.82	0.85	0.69	0.91	0.76	0.63
	561	5915						
F/F	2418	2041	0.80	0.96	0.54	0.98	0.69	0.62
	98	6378						
F/U	2489	1970	0.62	0.53	0.56	0.66	0.54	0.21
	2220	4256						
U/F	1006	3453	0.68	0.97	0.23	1.00	0.37	0.37
	29	6447						
U/U	3203	1256	0.84	0.87	0.72	0.92	0.79	0.67
	497	5979						
F/F*	2198	2261	0.79	0.96	0.49	0.99	0.65	0.58
	82	6394						
F/U*	3250	1209	0.60	0.50	0.73	0.51	0.60	0.23
	3198	3278						
U/F*	1162	3297	0.69	0.90	0.26	0.98	0.40	0.37
	130	6346						
U/U*	3095	1364	0.80	0.80	0.69	0.88	0.74	0.59
	769	5707						

Table B.2: Pixel-level inference results on test set in the reflectance continua pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	1716	2743	0.68	0.69	0.38	0.88	0.49	0.31
	783	5693						
Spectral	2116	2343	0.63	0.55	0.47	0.74	0.51	0.22
	1715	4761						
F/F	1668	2791	0.71	0.80	0.37	0.94	0.51	0.39
	419	6057						
F/U	2893	1566	0.50	0.43	0.65	0.40	0.51	0.05
	3895	2581						
U/F	833	3626	0.66	0.92	0.19	0.99	0.31	0.31
	76	6400						
U/U	958	3501	0.65	0.75	0.21	0.95	0.33	0.25
	327	6149						
F/F*	1779	2680	0.71	0.79	0.40	0.92	0.53	0.39
	486	5990						
F/U*	3033	1426	0.49	0.42	0.68	0.37	0.52	0.05
	4107	2369						
U/F*	560	3899	0.63	0.82	0.13	0.98	0.22	0.22
	123	6353						
U/U*	1654	2805	0.65	0.62	0.37	0.84	0.46	0.24
	1024	5452						

Table B.3: Pixel-level inference results on test set in the reflectance features pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	458	3831	0.63	0.87	0.11	0.99	0.19	0.22
	71	6218						
Spectral	3386	903	0.49	0.43	0.79	0.29	0.56	0.09
	4452	1837						
F/F	3219	1070	0.58	0.49	0.75	0.46	0.59	0.21
	3395	2894						
F/U	6416	2162	0.44	0.40	0.75	0.22	0.52	-0.03
	9768	2810						
U/F	4026	263	0.46	0.42	0.94	0.13	0.58	0.11
	5488	801						
U/U	8216	362	0.43	0.41	0.96	0.07	0.58	0.05
	11730	848						
F/F*	3256	1033	0.57	0.48	0.76	0.45	0.59	0.21
	3470	2819						
F/U*	1379	2910	0.53	0.40	0.32	0.67	0.35	-0.01
	2103	4186						
U/F*	2003	2286	0.66	0.61	0.47	0.79	0.53	0.28
	1292	4997						
U/U*	4061	228	0.44	0.42	0.95	0.09	0.58	0.08
	5696	593						

Table B.4: Pixel-level inference results on test set in the absorbance pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	572	3717	0.64	0.86	0.13	0.99	0.23	0.24
	93	6196						
Spectral	4096	193	0.45	0.42	0.96	0.11	0.59	0.11
	5604	685						
F/F	2851	1438	0.60	0.50	0.66	0.55	0.57	0.21
	2826	3463						
F/U	4964	3614	0.50	0.41	0.58	0.44	0.48	0.02
	7014	5564						
U/F	4229	60	0.41	0.41	0.99	0.02	0.58	0.03
	6157	132						
U/U	8338	240	0.43	0.42	0.97	0.07	0.58	0.08
	11750	828						
F/F*	2778	1511	0.60	0.51	0.65	0.57	0.57	0.22
	2692	3597						
F/U*	2228	2061	0.53	0.43	0.52	0.53	0.47	0.05
	2926	3363						
U/F*	4005	284	0.40	0.40	0.93	0.03	0.56	-0.08
	6094	195						
U/U*	4125	164	0.44	0.42	0.96	0.09	0.58	0.10
	5709	580						

Table B.5: Pixel-level inference results on test set in the absorbance continua pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.

Model	Confusion matrix		Acc.	Prec.	Sens.	Spec.	F ₁	MCC
RGB (s)	3330	959	0.50	0.43	0.78	0.31	0.56	0.10
	4334	1955						
Spectral	1892	2397	0.58	0.48	0.44	0.68	0.46	0.12
	2023	4266						
F/F	1186	3103	0.62	0.56	0.28	0.85	0.37	0.16
	941	5348						
F/U	4826	3752	0.50	0.41	0.56	0.45	0.48	0.02
	6882	5696						
U/F	1267	3022	0.50	0.36	0.30	0.64	0.32	-0.06
	2250	4039						
U/U	1842	6736	0.58	0.47	0.21	0.83	0.29	0.06
	2084	10494						
F/F*	1347	2942	0.62	0.55	0.31	0.82	0.40	0.16
	1119	5170						
F/U*	1252	3037	0.55	0.42	0.29	0.72	0.34	0.02
	1746	4543						
U/F*	1493	2796	0.47	0.35	0.35	0.56	0.35	-0.09
	2780	3509						
U/U*	330	3959	0.58	0.39	0.08	0.92	0.13	-0.01
	519	5770						

Table B.6: Pixel-level inference results on test set in the absorbance features pipeline. Confusion matrices consist of true lepidolite pixels along the top row and true granite samples along the bottom row, with predicted lepidolite/granite classes in the left/right columns.