March 2024

Automated identification and mapping of interesting mineral spectra in CRISM images

Arun M. Saranathan
University of Massachusetts Amherst

AUTOMATED IDENTIFICATION AND MAPPING OF INTERESTING MINERAL SPECTRA IN CRISM IMAGES

A Dissertation Presented

by

ARUN M. SARANATHAN

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

February 2024

Electrical & Computer Engineering

AUTOMATED IDENTIFICATION AND MAPPING OF INTERESTING MINERAL SPECTRA IN CRISM IMAGES

A Dissertation Presented

by

ARUN M. SARANATHAN

Approved as to style and content by:

Mario Parente, Chair

Marco Duarte, Member

Janice Bishop, Member

Christopher V. Hollot, Department Chair
Electrical & Computer Engineering

DEDICATION

To my parents- count one for us!

ACKNOWLEDGMENTS

First, I would like to express my deep gratitude to my advisor Prof. Mario Parente, for his guidance, support, and mentor-ship during throughout my Ph. D. studies. I am also thankful to my committee members Prof. Marco Duarte and Dr. Janice Bishop for their time and invaluable feedback.

I am also very thankful for my colleagues in the Remote Hyperspectral Observers Group (RHOGroup) @UMass, Amherst for creating a fun and collegial workspace. In particular, I would like to thank Yuki Itoh, Ping Fung, and Shobitha Shivakumar for the many interesting discussions and feedback at various stages of my dissertation.

I am also thankful to our various scientific collaborators including Dr. Frank Seelos, Dr. Ray Arvidson, and Dr. Darby Dyar for providing me with timely feedback and for sharing their vast experience with me.

Finally I would like to express my deep gratitude to my family and friends (both in Amherst and elsewhere) without whose support and encouragement none of this would be possible.

Thank you all! I am indeed grateful.

ABSTRACT

AUTOMATED IDENTIFICATION AND MAPPING OF INTERESTING MINERAL SPECTRA IN CRISM IMAGES

FEBRUARY 2024

ARUN M. SARANATHAN

B.E., VISVESVARAYA TECHNICAL UNIVERSITY, BELGAUM, INDIA

M.S., UNIVERSITY OF MASSACHUSETTS, AMHERST

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Mario Parente

The Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) has proven to be an invaluable tool for the mineralogical analysis of the Martian surface. It has been crucial in identifying and mapping the spatial extents of various minerals. Primarily, the identification and mapping of these mineral spectral-shapes have been performed manually. Given the size of the CRISM image dataset, manual analysis of the full dataset would be arduous/infeasible. This dissertation attempts to address this issue by describing an (machine learning based) automated processing pipeline for CRISM data that can be used to identify and map the unique mineral signatures present in a CRISM image. The pipeline leverages a highly discriminative representation learned through the use of Generative Adversarial Networks, such that in this novel representation space simple distance metrics are sufficient to discriminate between even very similar spectral shapes. The pipeline leverages this enhanced

feature space to set up an open set classification problem that labels each new pixel as either a member of a known mineral class or novel spectral shape (or outliers). Following this, a segmentation technique is used on the outliers to group them, and further, reduce them to a representative set of the novel spectral shapes present in the image. These novel spectral shapes can then be labeled based on expert analysis and used to update the open-set classifier. The performance of these tools are validated over a subset of CRISM images from different parts of the Martian surface such as Jezero Crater, North East Syrtis, and Mawrth Vallis.

TABLE OF CONTENTS

LIST OF TABLES

Table **Page**

LIST OF FIGURES

Figure **Page**

CHAPTER 1

INTRODUCTION

Somewhere, something incredible is waiting to be known.

Carl Sagan

The Martian planetary surface holds many secrets on the various processes involved in the formation and evolution of our nearest planetary neighbor. Given their extraordinary ability to remotely identify and map mineral families present on the surface of far away planetary bodies, spectral sensors have emerged as an essential tool for the study of the Martian surface. In particular, spectral sensors at various spatial and spectral resolutions were employed to achieve NASA's aim for Mars Exploration to "follow the water" [43]. Thermal sensors with low spatial resolution like Thermal Emission Spectrometer (TES) on Mars Global Surveyor [21] (several kilometers per pixel and spectral resolution between $0.5 - 1$ nm over a spectral range of $6 - 50$ μm) and THErmal Emission Imaging System (THEMIS) on the Mars 2001 Odyssey mission [22] (100 m per pixel, with 10 bands over 9 wavelengths over a spectral range of $6.78 - 14.9$ μm) were used to map and locate silicate materials that formed the igneous rocks and other crystalline materials on the Martian surface. Following this Visible and Near InfraRed (VNIR) sensors were deployed to track fine-grained hydrated and hydroxylated secondary minerals that were formed in the presence of past "aqueous activity". One such sensor is the Observatoire pour la Minéralogie, l'Eau, les Glaces et l'Activité (OMEGA)[11] (spatial resolution of a few hundred to thousand meters per pixel and spectral range of $0.4 - 5$ μm). While such sensors have proven to be valuable for global or large scale studies of the planetary surface, a new sensor capable of measurements in the VNIR of a specific targeted location at a fine scale was

required- the Compact Reconnaissance Imaging SpectroMeter (CRISM) aboard the Mars Reconnaissance Orbiter (MRO) [60] was launched to address this need.

CRISM measures the electromagnetic energy scattered by the surface in the VNIR wavelength range ($0.362 - 3.920\mu$m) at a spatial resolution of 18m/pixel (in the full resolution mode) and a spectral resolution of 6.55 nm/channel. The analysis of the CRISM spectral data has been instrumental in revealing the history of aqueous alterations on the Martian surface, thus confirming the presence of liquid water on the Martian surface in a bygone era [30, 10, 61, 85]. Such analysis has enabled researchers to identify the widespread presence of various phyllosilicate (smectite) minerals such as hydrated silica [62], a variety of Fe/Mg-smectites [15], and Al-smectites [95]. Researchers have also been able to identify a variety of other hydrated smectites and chlorites [72]. Another commonly seen component of the Martian surface are mafic minerals[1] which are a primary component of igneous rocks. Commonly seen mafic minerals on the martian surface are olivines (both magnesium-rich fosterite [83] and iron-rich fayalite [62]) and pyroxenes (again both low-calcium [62] and high-calcium [82]). Analysis has also lead to the identification of other minerals including Mg-carbonates [29] and sulfates [92], together with several other species [27]. These detections have been instrumental in understanding the physical and chemical processes that led to the formation of the Martian surface. Identification of specific groups of minerals and their spatial distributions on the Martian surface have also shed a great amount of light on the habitability[2] of Mars [88, 2, 14].

These detections have greatly improved our understanding of the various processes at play in the formation of the Martian surface; and they have also greatly aided in the selection of the landing sites for the Mars rover exploration missions [71]. Further,

---

[1] The word mafic is a portmontaeu of magnesium and iron (ferrous or ferric) minerals or rocks

[2] "The habitability zone is defined by the sum of the physical and chemical conditions, which would support the presence of liquid water." Treatise on Geophysics, 2007 [84]

in its operating lifetime ranging from September 2006 to April 2023 CRISM acquired tens of thousands of images of the Martian surface [80], providing a large amount of granular spectral information of the Martian surface. Analysis of such a large granular spectral dataset would be valuable in furthering our understanding of our nearest planetary neighbor. Such analysis is made significantly harder by the various noise/distortion processes present in the CRISM image dataset, such as (i) incomplete or imperfect atmospheric correction, due to calibration errors, atmospheric distortion artifacts persist in even the highest quality products available in the NASA CRISM Planetary Data System (PDS)[3], (ii) Cubes with increased noise due to instrument performance degradation, in the later stages due to the degradation of the instrument cooling systems many of the later cubes were measured at a higher then expected detector temperature leading to increased noise/spike clusters in the measurements [53, 45]. Many of these noise/distortion processes appear to be column-dependent [63, 62, 29].

Given, such data conditions most of the analysis of the CRISM spectral data is performed manually by expert in spectroscopy. Practitioners use techniques like "spectral ratioing", wherein spectra of interest are divided by an "unremarkable spectrum" (a spectrum with no absorption features) from the same column. While spectral ratioing can help ameliorate noise/distortion effects, ratioing also has adverse effects like changing the continuum of the spectrum (see Appendix. A.1 for more details on the continuum) and distorting the spectral features present in spectra of interest (if the denominator spectra chosen are not completely "unremarkable"). Practitioners also use techniques like spatial averaging to eliminate zero-mean noise present in the spectrum. Following this experts identify the spectral shapes and map them manually in each image.

[3] `https://pds-geosciences.wustl.edu/missions/mro/crism.htm`

Historically, the most common tool used by experts for mapping/analyzing the spatial distribution of minerals in a specific image is the set of spectral summary parameters (initially introduced in [69], and then updated and revised in [90]). The spectral summary parameters use simple mathematical formulations to test each spectrum for the presence of a specific spectral feature. Each summary parameter assigns a score to a spectrum based on the perceived strength of the spectral feature in the specific spectrum. Practitioners further combine different spectral summary products to create "browse products", each browse product is 8-bit scaled RGB color composite that display three-band combinations of thematically related spectral summary products. These browse products can be used to provide high-level products to guide the experts to identify groups of related minerals. The browse products for specific images can also be downloaded from the NASA PDS mentioned above.

Despite their widespread use, spectral summary parameters exhibit significant limitations namely, the spectral summary parameters exhibit a large number of "false positives", i.e. spectral summary parameters are known to show high values for spectra that do not contain the specific feature of interest but rather contain a different feature in a similar spectral range (see [90] for a full list of "caveats" associated with each spectral summary parameter). Another well known issue is that the summary parameter values are greatly affected by many global factors and as such it does not appear that there is a global statistical basis on whether a feature is present or not. Rather they are used locally to identify spectra that are diverse from its neighbors, and then the identifications are made manually. Also, given the nature of their mathematical formulation, which uses the reflectance values seen at specific spectral bands, they are greatly affected by the presence of noise/distortions. A more detailed description of the spectral summary parameters and the issues associated with these spectral summary parameters are shown in Section 2.2.

Given the extensive size of the CRISM image database automated algorithms capable of identifying and mapping minerals in various CRISM images would be highly valuable to the community. This quest is greatly helped by recent advances in the atmospheric correction of the CRISM spectral data (e.g. [45, 79]) making significantly cleaner spectral data available for analysis. Initial attempts to perform automated mineral identification and mapping endeavor to leverage the summary parameters, by building automated rule-based systems/pipelines to identify and map various mineral signatures in hyperspectral images. The methods attempt to build various rule-based systems on top of the spectral summary parameters to make automated identifications and maps. Carter et al. [17], built a procedure for identifying the presence of weakly hydrated minerals in CRISM images. Recently, the same authors also published a first attempt at a global high resolution map of specific aqueous minerals using data from multiple instruments in the Mars Orbital Catalog of Aqueous Alteration Signatures (MOCAAS) [18]. The CRISM portion of their work still primarily leverages spectral summary parameters to generate identifications and mappings. After a variety of preprocessing steps (such as destriping, continuum-removal, spurious pixel removal, etc.), the approach subtracts the average column spectrum from the current pixel and performs continuum removal to highlight weak mineral absorptions. Following this, spectral summary parameters are calculated, and the presence of an individual mineral is identified by using a rule-based system based on whether a specific set of summary parameters are enhanced while others appear to be suppressed. The system was developed primarily for minerals showing characteristic absorptions in the spectral range $1.7 - 2.6\ \mu m$.

Allender et al. [1] introduced another rule-based system, which combines a set of mineral specific rules along with full summary parameter representation for both mapping known spectra as well as identifying novel spectra. The proposed approach begins by applying basic preprocessing steps such as photometric and atmospheric

correction, destriping, and despiking to the chosen CRISM image cube. Following this, the full list of standard CRISM spectral summary parameters described in [90] are calculated for each CRISM-pixel. Then, the authors apply a super-pixel segmentation [31] to summary-parameter cube, to find spatially contiguous regions (referred to as "superpixels") with similar spectral summary parameter representations. In all future processing, each "superpixel" is represented by the median summary parameter vector, which significantly reduces the data complexity for future processing. The superpixel formulation is also known to significantly reduce the effect of localized noise artifacts. Following this, the procedure uses Discovery through Eigen-basis Modeling of Uninteresting Data (DEMUD) [91], which is an anomaly detection algorithm, to find the 500 most diverse spectral signatures (in terms of the estimated summary parameters) from the CRISM image, the authors then use Ordering Points To Identify the Clustering Structure (OPTICS) [3] to find the most diverse clusters from this data. The summary parameter vector corresponding to the diverse clusters are thresholded to find the enhanced summary parameter values for the cluster. The list of enhanced summary parameter are compared to list of enhanced summary parameters for known mineral identifications and a rule based system is used to map mineral shapes in the CRISM image. While this method is an interesting way of leveraging available knowledge of known mineral identification, the method is primarily dependent on summary parameters, and the method inherits all the known issues of the summary parameter technique. Further, the lack of a global statistical basis makes the extension of these mapping across images much harder as it is quite difficult to compare pixel-spectra from different images.

The techniques mentioned above are primarily based on summary parameters and as such these methods inherit all the issue of the summary parameters. Namely, they are highly susceptible to variations due to noise/distortions. Also, the absence of a global statistical basis to these summary parameters makes the creation of multi-

image mapping products quite hard as it is often hard to compare pixel-spectra from various regions based on their summary parameter values. Furthermore, these methods often require the construction of a rule-based system to identify various mineral shapes, as such the addition of novel mineral shapes to the mapping library requires significant time and effort from expert practitioners.

Another strain of mineral identification algorithms is the similarity based algorithms that attempt to measure the similarity between two spectra using various methods. Primarily these approaches attempt to match each test spectrum to a library of known spectral shapes or laboratory measurements referred to as *exemplars*. A seminal approach of this strain of algorithms was introduced by Parente et. al. [67] for the identification and mapping of mafic minerals in CRISM images. The method first applies continuum removal to eliminate illumination and other acquisition effects from the spectral data. Then the method models the absorption features using automated Gaussian curve-fitting methods. Finally, the coefficients for the fit to the target spectrum are compared to the fit-coefficients associated with the exemplar spectra to identify the closest match from the library.

Other practitioners attempt to leverage unmixing[4] based methods for mineral identification and mapping. Gilmore et al. [34] proposed a technique that first applies a superpixel segmentation algorithm [31] to a CRISM image cube and then identifies pixel-spectra which are considered the "corners" the data cloud (referred to as *image endmembers*) by applying the Sequential Maximum Angle Convex Cone (SMACC) [38]. Finally, the algorithm uses a spectral angle similarity to generate the maps corresponding to the different image endmembers identified in the previous step. The algorithm assumes that purest mineral shapes are the most extreme points of the data cloud and as such these extreme pixel-spectra are the only mineral shapes of

[4]Unmixing refers to the attempt at identifying the constituent components in each spectrum. See Appendix A for more information.

interest in the data. While such assumptions appear reasonable in scenarios that are noise free and with well known mixing models, such assumptions break-down in the presence of random noise and atmospheric distortions like the ones present in CRISM data. Rohani et al. [74] have previously shown that in scenarios where the data is afflicted by significant amounts of noise, it appears that more and more points can be considered extreme. Using certain graph-theoretic metrics they show that most of the pixel-spectra in any CRISM image can be considered extreme, making the identification of the mineral-spectra of interest using unmixing algorithms much harder.

An interesting strain of algorithms for mineral identification are the Factor Analysis and Target Transformation (FA/TT) based techniques [56]. *Factor Analysis* is a statistical method for identifying the significant independent components of a datasets, the independent components can be considered the smallest set of vectors, such that every element in the dataset can be expressed as a linear combination of these components. The technique uses an eigenvalue decomposition to express the given dataset as combination of a small number of uncorrelated and unobserved vectors. The algorithm further postulates that a specific eigenvector is significant or can be considered a *factor* if and only if the eigenvalue associated with the specific vector is above a chosen threshold. The *target transformation* portion of the algorithm states given a list of factors associated with the dataset, the factors are sufficient to not only reconstruct the spectral-vectors present in the dataset, but that they can also reconstruct the spectra associated with the pure minerals which contribute to the original dataset. Therefore the algorithm attempts to reconstruct known pure mineral shapes in a spectral library using the factors extracted from the dataset of interest. If the dataset specific factors are able to reconstruct an exemplar spectrum accurately (measured using a simple euclidean distance), then it is concluded the material corresponding to the exemplar is present in the dataset. This approach is

very well known and has been commonly applied to Thermal InfraRed (TIR) spectral data from various sources. [5, 6, 20, 35, 33, 40].

One of the first approaches which leverages FA/TT for CRISM data was described by Thomas et al. [89], wherein the authors apply FA/TT to a significantly sub-sampled version (the sub-sampling is about $1:25$) of the image dataset to identify factors. This approach has been quite successful in identifying the dominant mineral modes in the CRISM images. The main issue with the FA/TT appears to be that the process is computationally quite expensive as it requires an eigenvalue decomposition for each image. Methods leveraging this technique often includes a significant sub-sampling step to make the technique tractable for CRISM data. While sub-sampling helps with the feasibility of the method it carries inherent risk that small deposits may be missed as part of the sub-sampling. Another disadvantage of this technique is that it produces a list of spectra corresponding to the mineral components present in the CRISM image but in this form it does not produce a map of the different components.

A recent approach Dynamic Aperture Factor Analysis for Target Transformation (DAFA/TT) [55] attempts to address some of the issues for FA/TT mentioned above. DAFA/TT applies FA/TT for small moving windows of different shapes. Since DAFA/TT considers a small dataset at each step it ameliorates the computation complexity issue of the eigenvalue decomposition. Further, since the method looks for the presence of each mineral in a small window it is also able to produce maps indicating whether the mineral is present in/around each pixel. The technique also uses information on the presence of a specific mineral in windows/apertures of different shapes to further refine the shape of each deposit. While DAFA/TT addresses many of the issues endemic to FA/TT, it should be noted that DAFA/TT in some sense trades computational complexity for a large number of iterations. Not only does this algorithm require application for FA/TT for each pixel in an image, it also requires

application of multiple apertures or windows. Practical applications have shown that in spite of the ability of DAFA/TT to generate spatial extent maps the DAFA/TT maps are a bit chunky and the creation of precise mineral maps requires the use of a large number of windows/apertures. Further since DAFA/TT uses a distance function to estimate whether an independent component is well reconstructed they often inherit the issues with similarity metrics to measure spectral similarity mentioned in Section 2.1.

Leveraging recent advances in machine learning Plebani et al. [71] released a machine learning toolbox/toolkit for the identification and mapping of mineral shapes in CRISM images. The toolbox is built on "ratioed" spectra- wherein each pixel-spectrum of interest is ratioed with a neutral spectrum (i.e. a spectrum with no clear mineral absorptions) in its vicinity. For this purpose the toolkit consists of a neutral-pixel detector which both identifies spectra with mineral absorptions and the closest neutral spectrum that can be used for ratioing the mineral spectrum. Following this the toolkit has two main tools for mineral classification, namely a Hierarchical Bayesian Model (HBM) and Convolutional Neural Network (CNN). The HBM leverages bayesian inference [50] to estimate the posterior membership of the test spectrum. The HBM has a two layer structure such that: "the lower layer models spectral variations of the same class across images whereas the upper layer models spectral variations across different classes." [71]. In this scenario, the lower layer models intra-class variability whereas the upper layer models inter-class variability. The HBM starts with a prior at both levels and uses the bayesian machinery to estimate the posterior probability for each test spectrum. The other approach uses neural network models to perform a supervised classification of the test-spectra.

While the toolkit mentioned above shows impressive performance in terms of detecting a variety of minerals in CRISM images it seems to primarily function as mineral detectors. Since the toolkit is built for ratioed pixels the performance of

the models are highly dependent on the presence of neutral spectral shapes in the vicinity of a mineral spectrum (this is not always guaranteed as many images have lots of pixels with clear mineral absorptions). While ratioing is a viable and valid tool for spectral analysis it creates spectral shapes not present in the dataset and thus requires validation based on expert analysis. Further to mitigate issues like false positives etc, the tool-kit includes strict thresholds on model confidences and seems to be generally used to detect the presence of minerals in specific images rather than as a mapping algorithm. Further, since the machine learning models used in this toolkit are supervised models, they are affected by the limitations of the supervised algorithms, namely they are unsuited to process previously unseen spectral shapes. Further, the training set used for the mineral classification is highly dis-balanced (i.e. contains in-equal number of examples of the various classes) which also affects model performance. Also, currently the toolkit only contains labeled datasets for the CRISM Targeted Reduced Data Records (TRDR) images, and for analyzing other CRISM data products significant work needs to be done for re-training the machine learning models.

Past attempts at mineral mapping can be generally classified into three broad tiers. The first tier consists of methods leverage the classical spectral summary parameters to compare each test spectrum to known mineral shapes, often these techniques build an additional rule based system on top the summary parameters to identify specific spectral shapes. While such methods are quite intuitive in terms of how they work such techniques inherit the drawbacks of the spectral summary parameters. Other techniques that leverage methods to measure spectral similarity in different forms to compare test spectrum to different members of a known spectral library. The main drawback of this suite of models is that in general spectral similarity is a complicated concept that is hard to measure using classical similarity metrics. Finally, recent attempts of using supervised machine learning methods for mineral mapping is held

back as the dataset is expected to contain some novel spectral shapes that are not currently known. Overall, there is a severe lack of a comprehensive tool for identifying and mapping all the unique spectral shapes present in a CRISM image.

1.1 Research objectives and contributions

The main objective of the research presented in this dissertation is to create/design a comprehensive mapping pipeline for each CRISM image- which can create a summary of the various unique spectral shapes as well as the spatial extent of each individual mineral deposits. Based on previous research the spectral shapes associated with the minerals present on the martian surface are known *apriori*. The full processing pipeline needs to have the ability to:

1. To accurately identify and map the spatial extent of each of the known spectral (mineral) shapes (if present) in each individual CRISM image.

2. To identify the presence of spectral (mineral) shapes that are different from the mineral shapes in the known spectral library.

3. To identify the spatial extents of the novel spectral shapes for each individual CRISM image.

4. To group novel spectral mineral shapes based on spectral similarity to create products which experts can use for mineral attribution for each novel spectral shapes identified in CRISM images.

1.2 Contributions of this work

To achieve these objectives mentioned in the previous section, the pipeline presented in this dissertation leverages some machine learning tools based on unsupervised/semi-supervised learning (due to the absence of a large exhaustive database of labeled min-

eral shapes present in the CRISM image database). The full list of tools designed for this purpose are:

- **A Feature Extractor based on Generative Adversarial Network (GAN)**
  In this portion, an (unsupervised) Generative Adversarial Network (GAN) was trained on CRISM I/F spectral data. This dissertation will clearly demonstrate the features extracted from the GAN are highly discriminative of the different spectral shapes present in the CRISM image database . It will also be illustrated that these features can be used to identify clear examples of all the known mineral spectral shapes from the CRISM image database. This tool is used to create a labeled dataset with examples of all the known spectral shapes from the CRISM image database

- **An open set classifier for identifying hitherto unknown spectral shapes**
  Following this, the dissertation also presents an *open-set* classifier which is capable of known spectral shapes while simultaneously being able to also identify/flag pixels with spectral shapes that are unlike the ones present in the labeled training set, i.e. "outliers".

- **A graph-based segmentation algorithm to group "outliers" with similar spectral shapes into segments/groups**
  A modified version of the graph-based segmentation algorithm [31] that groups similar spectral shapes into groups/segments.

1.3 Organization of this dissertation

The rest of this dissertation is organized as follows: **Chapter. 2** introduces the reader to common dissimilarity metrics used to measure spectral similarity. Further, this chapter will also describe the issues with using these dissimilarity metrics for mineral discrimination. The chapter will also introduce readers to the spectral summary

parameters, commonly used by practitioners for mineral discrimination. The chapter will also attempt to highlight some common issues with these spectral summary parameters for data containing noise/outliers.

Chapter. 3, introduces the GAN based feature extractor that was designed for CRISM I/F spectral data. The chapter briefly introduces GANs and describes the various preprocessing steps used by the tool as well the training procedure for the GANs. The following sections will describe the use of the GAN as a feature extractor for the CRISM I/F data. The chapter will also attempt to clearly show the gains in terms of the mineral discrimination when using these features relative to the original data. Following this the chapter will describe the procdure to use these features to map mineral shapes in CRISM images. The chapter concludes by listing known issues and limitations with the described feature space.

Chapter. 4 will describe the full open set pipeline for mapping all known spectral shapes as well as a the functionality to identify/flag the locations of the novel spectral shapes (referred to as "outliers"). The chapter also describes a graph based clustering tool to group these outliers into clusters with similar spectral shapes. Finally, the chapter also describes a grouping scheme to create groups of the cluster means based on spectral similarity.

Finally, the conclusions are presented in **Chapter. 5**.

CHAPTER 2

WHY IS SPECTRAL DISCRIMINATION SO DIFFICULT?

The core requirement for any competent spectral mapping method is the ability to accurately match the shape of a test spectrum to known mineral spectral shapes. This matching is made significantly more difficult by the fact that mineral identification is often performed using the presence of specific absorption shapes in very narrow spectral windows[1]. As mentioned in the previous chapter the two techniques used to perform spectral similarity are similarity/dissimilarity metrics that can be used to measure the spectral similarity and the spectral summary parameters. This chapter will provide some insight into why these methods are not efficient methods for measuring spectral similarity.

2.1 Issues with classical similarity metrics for spectral discrimination

First, consider the issues present when using metrics including Spectral angle divergence (used in [34]) and euclidean distance (used as RMSE in FA/TT based methods) to measure spectral similarity. A general metric that measures either the similarity/dissimilarity between the objects in a given set M can be defined as a function $d : MXM \rightarrow \mathbb{R}^+$, (where $\mathbb{R}^+$ denotes the set of positive real numbers) which satisfies the following conditions:

- $d(x, y) \geq 0, \forall x, y \in M$, that is the similarity/dissimilarity is always positive.

[1] see Section A.1 for more information.

- $d(x, y) = d(y, x)$, i.e., the similarity/dissimilarity metrics must be symmetric.
- $d(x, z) \leq d(x, y) + d(y, z)$, this condition which is known as the triangle inequality ensures that the distance between two points is the shortest path between the two points.

In such a scenario, the choice of the metric is an important factor in being able to generate an accurate matching of a test spectrum with the library spectra. As mentioned above, some well-known metrics that have found wide-spread use are the cosine distance and the euclidean distance (note that these metrics technically measure a distance or dissimilarity that is the inverse of the similarity). Given two data-points x, y that are two real valued vectors (i.e. $x, y \in \mathbb{R}^d$), we define the euclidean distance (d_{euc}) as

$$d_{euc}(x, y) = \|x - y\| = \sqrt{\sum_{i=0}^{d} (x_i - y_i)^2} \tag{2.1}$$

where $\|.\|$ refers to simple ℓ_2-norm. Geometrically, the ℓ_2-norm is the euclidean distance can be interpreted as the length of the shortest line between the two points. While the cosine distance (d_{cosine}) is defined as:

$$d_{cosine}(x, y) = 1 - \frac{\sum_{i=0}^{d} x_i y_i}{\sqrt{\sum_{i=0}^{d} x_i^2} \sqrt{\sum_{i=0}^{d} y_i^2}} \tag{2.2}$$

Geometrically, this is 1 minus the cosine of the angle subtended by the vectors from the origin to x and the origin to y. The main issue with such metrics appears to be that they treat each dimension of the real vectors x & y equally, i.e., these metrics do not prioritize some dimensions. This does not seem to agree with how spectra are analyzed manually; generally mineral determination is performed based on spectral features that span very few spectral bands. The classical interpretation means that

often global factors have an inordinate effect on the measured similarity/dissimilarity (with respect to the spectral similarity).

To illustrate the inability of the classical similarity metrics to measure notions of spectral similarity accurately consider the example shown in Fig. 2.1. Fig. 2.1 (A) shows spectra of the CMS[2] standard for the mineral Nontronite at different grain sizes. There is no difference in the mineralogy or chemistry of the two samples. The two samples in Fig. 2.1 (A) display absorption features at 1.41μm and 1.91μm (indicating the presence of water in the lattice), and absorptions at 1.43 μm and 2.3 μm, which indicates the presence of ferric hydrooxide ($Fe(OH)_3$) in the mineral lattice. Fig. 2.1 (B) on the other hand compares the CMS Nontronite samples to a CMS sample of the mineral Montmorillionite. Montmorillionite is quite similar to Nontronite except that the lattice primarily contains aluminium hydrooxide ($Al(OH)_3$), instead of the ferric hydro-oxide. This results in the absence of a strong 2.3 μm absorption and the over-tone at 1.43 μm and instead shows a strong absorption feature at 2.2 μm. The spectra corresponding to these samples were extracted from the NASA Reflectance Experiment LABaratory (RELAB) [59] database.

Fig. 2.1 also shows the estimated dissimilarity values between the two pairs of spectra shown in Fig. 2.1 (A) & (B). Note, that the two Nontronite spectra that have the same mineralogy and chemical composition are shown to be more dissimilar as compared to the Nontronite and Montmorillionite spectra which are chemically and compositionally different as mentioned above. The actual value of the dissimilarity metrics appears to be more affected by global factors such as differences/similarity in the continuum etc., rather than the specific discriminatory features like the position of the oxide band. This example clearly indicates that often the dissimilarity between samples belonging to the same mineral class is significantly higher than the dissimi-

[2]CMS: Clay Mineralogical Society

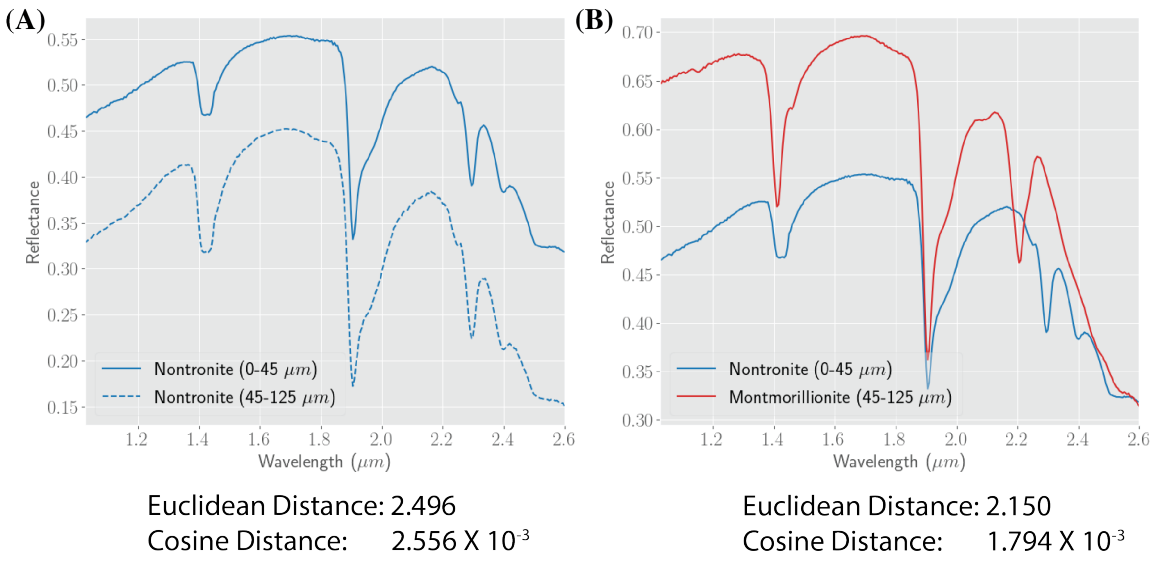

Figure 2.1. The issue with using standard similarity metrics to measure spectral similarity. (A) CMS nontronite samples. (B) CMS nontronite vs CMS montmorillionite. samples

larity (in terms of spectral angle divergence) between samples from differing classes.

To better illustrate this phenomenon, a simple experiment was carried out, wherein the largest intra-class divergence/dissimilarity (i.e. the divergence of samples which have the same mineralogy and chemical composition while having other physical differences) and the smallest inter-class divergence/dissimilarity (divergence between samples which have different chemical compositions) were estimated for some well-

Table 2.1. Maximum Intra-class Divergence for Some Mineral Families

Mineral	max intra class angle
Nontronite (NO)	1.27e-2
Illite (IL)	6.23e-3
Ripidolite (RI)	4.06e-3
Sepiolite (SE)	3.85e-3
Olivine* (FO)	2.28e-3
Kaolinite (KA)	1.89e-3
Montmorillonite (MO)	5.13e-3

*Olivine(Fo): - fosterite

Table 2.2. Minimum Inter-class Divergence for the Mineral Families in Table 2.1

	NO	IL	RI.	SE	FO	KA	MO
NO	X	2.86e-2	6.15e-1	**3.24e-3**	5.32e-2	9.7e-3	**3.43e-3**
IL	2.86e-2	X	**2.09e-3**	3.26e-2	**2.23e-3**	2.39e-2	1.52e-2
RI	6.15e-1	**2.09e-3**	X	7.29e-2	**2.88e-3**	5.73e-2	4.53e-2
SE	**3.24e-3**	3.26e-2	7.29e-2	X	6.71e-2	6.17e-3	**4.8e-4**
FO	5.32e-2	**2.23e-3**	**2.88e-3**	6.71e-2	X	4.86e-2	3.75e-2
KA	9.7e-3	2.39e-2	5.73e-2	6.17e-3	4.86e-2	X	**1.85e-3**
MO	**3.43e-3**	1.53e-2	4.53e-2	**4.8e-4**	3.75e-2	**1.85e-3**	X

known mineral classes. In this experiment the cosine distance was chosen as the divergence metric. Spectra corresponding to samples from the 7 well-known mineral classes were extracted from the RELAB and CMS spectral databases. To avoid any confusion between members of related families we limited ourselves to minerals that have a 4-level classification under the Dana classification systems for mineral species [24]. Table (2.1) reports the largest divergence observed between samples of the same mineral class (i.e. the maximum intra-class divergence). Each cell in Table (2.2) on the other hand reports the smallest difference between a pair of samples corresponding to the two mineral classes. The values highlighted in red in Table (2.2) correspond to the cases where the maximum intra-class divergence of one of the classes is larger than the minimum inter-class divergence between the class of another mineral family.

Note that for each mineral there is at least one other mineral such that intra-class divergence is often higher than the inter-class divergence to some other spectrally similar mineral class. Due to this behavior, it is almost impossible to use the similarity metric to differentiate a member of the same mineral class from members of other mineral classes. Setting a high threshold on the similarity metric carries the risk that spectra not belonging to the same class may be considered similar. On the other hand setting a low-threshold on the similarity metric would mean that diverse samples of the same class are not accurately classified.

[N.B.: Much of this analysis on the issues with using similarity metrics to measure spectral similarity has been previously presented in our previous works, [75] & [76].]

2.2 Spectral summary parameters and the associated issues

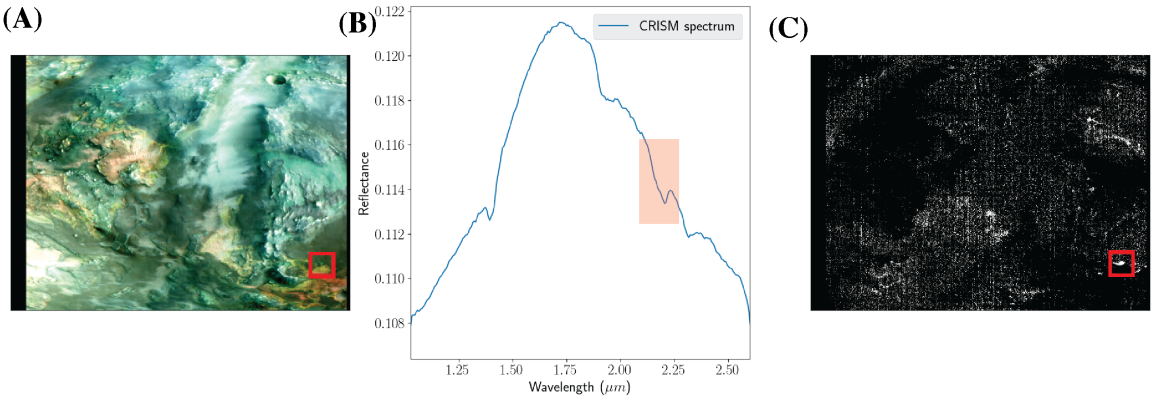

Figure 2.2. Application of summary parameters to CRISM Data. (A) RGB Composite of FRT000097E2 with the kaolinite deposit highlighted. (B) The spectrum of the kaolinite deposit. (C) Summary Parameter MIN2200H on CRISM data denoised using the method described in [45]-in spite of the image being aggressively denoised the summary parameter map appears quite noisy and displays large parameter values for many pixels other than the highlighted deposit.

In this section, let us consider the choice/use of the spectral summary parameters for mineral identification. Spectral summary parameters identify the presence of a absorption feature in a specific spectral window, by using a specific mathematical formula, that is designed to produce a large value if the spectral feature is present in the spectrum; otherwise, a small or negative value is produced. [The reader is referred to Section. 4 in [90] for the exact procedure and nature of the mathematical formulations used for various spectral shapes]. In Table 2, [90] also provides the list of the commonly used spectral summary parameters along with minerals identified by each of the parameters. The table also identifies the caveats or associated minerals for which the parameter provides spuriously high values that may cause misidentifications/errors. In addition to providing spuriously high values for other minerals

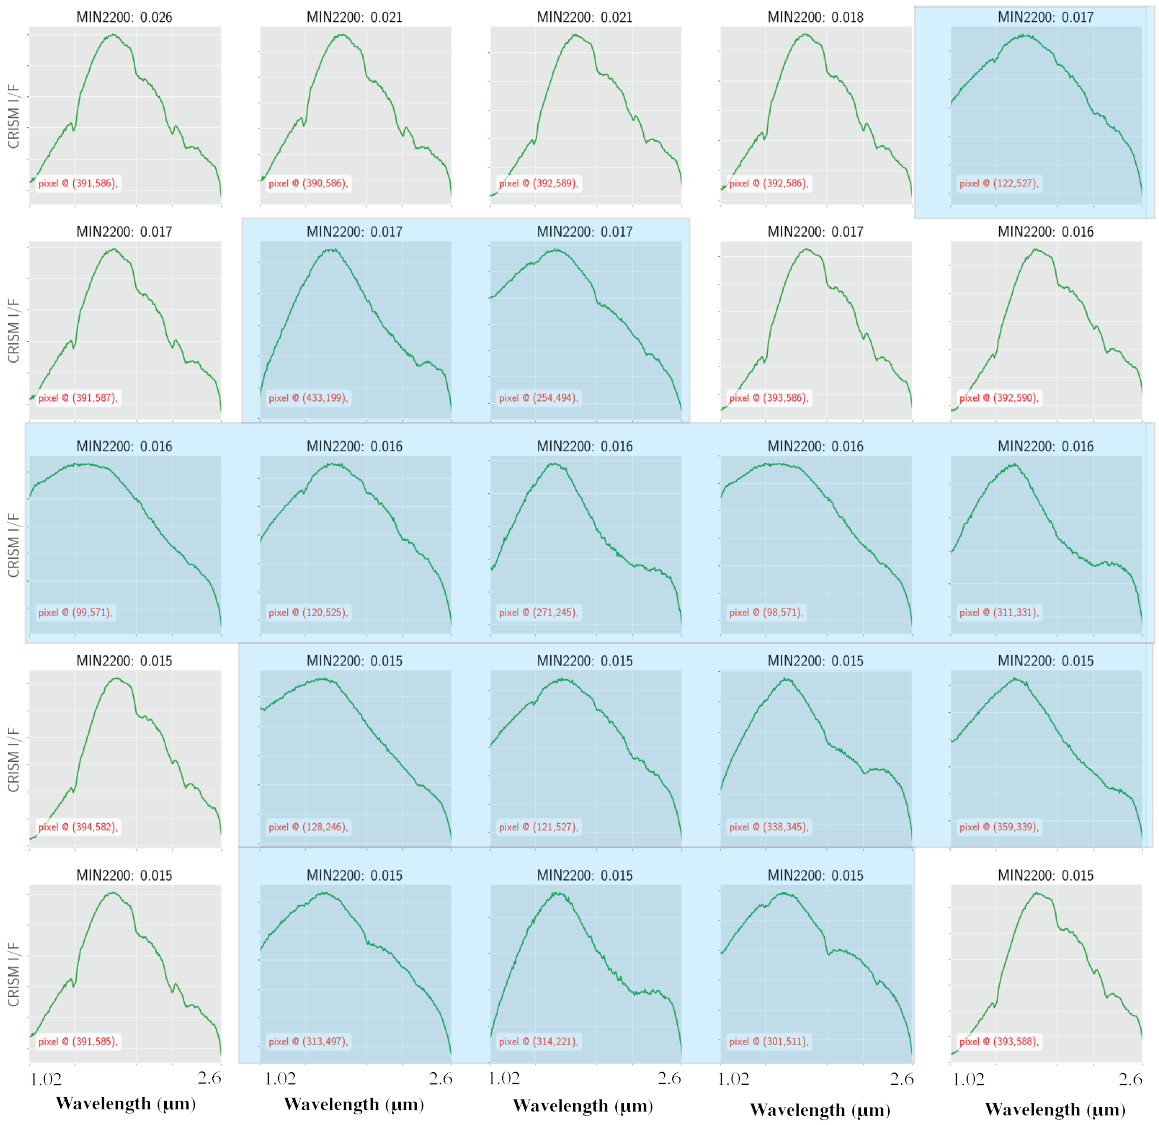

Figure 2.3. The 25 pixel-spectrum from the CRISM image FRT000097E2 with the highest score for the summary parameter MIN2200. Notice that a large number of these spectra (highlighted in blue) do not consider the characteristic 2.2 μm absorption expected to be present in kaolinite-bearing spectra like the one shown in Fig. 2.2 (B)

which may have absorption features in a neighboring spectral windows, the spectral summary parameters are highly susceptible to the presence of any noise or distortions in the spectral window of interest. Given, the rather precise mathematical formula-

tion of the spectral summary parameters even small amounts of noise in the window of interest may affect the calculation of these summary parameters.

To illustrate the effect of noise on the spectral summary parameters, consider the subset of the CRISM image $FRT000097E2$. The image has a small deposit (highlighted by the red rectangle in Fig. 2.2 (A)), which is spectrally consistent with the mineral kaolinite. An expert spectroscopist would make this determination based on the presence of the "doublet" feature with the minimum near $2.2\mu m$. The specific feature is highlighted in red in Fig. 2.2 (B). The spectral summary parameter MIN2200 was designed to identify spectra with this exact feature. Fig. 2.2 (C) shows the MIN2200 parameter map for this image. Clearly the pixels corresponding to the deposit of interest have a large value for the chosen summary parameter. It should also be noted that the summary parameter map also appears quite noisy. Additionally, when we consider the 25 pixel-spectra from $FRT000097E2$ which have the highest value of the MIN2200 (as shown in Fig. 2.3), it seems clear that many of the pixel-spectrum with a high score for the MIN2200, show no evidence of the characteristic kaolinite-like absorption at $2.2\mu m$ (such examples are highlighted in blue in Fig. 2.3).

The issues described in this chapter are intended to provide the reader a general notion on how complex spectral similarity is to define and measure. Spectral similarity is often defined based on specific features and their location in terms of the spectral window. This discrimination is made significantly more challenging by the presence of noise and atmospheric distortions. In many cases the determination of whether a feature present in a specific spectrum is noise/artifact or a characteristic mineral absorption requires expert manual analysis. This chapter illustrates the need/value of learning a representation wherein mineral discrimination can be better accomplished.

CHAPTER 3

MINERAL MAPPING USING AN ADVERSARIAL FEATURE EXTRACTOR

In the light of the issues described in Chapter 2, a representation which would enable improved mineral discrimination would be very valuable. Algorithms like Generative Adversarial Networks (GANs) are highly capable of learning a very rich and detailed representation of any dataset. A GAN is a specific neural network variant composed of two neural network models with two components and *generator* and a *discriminator*. At convergence, the *generator* portion of the network is highly capable of creating samples that are indistinguishable from real samples in the dataset. At the same stage the other portion of the model, the *discriminator*, which attempts to separate "real samples" from the "generated samples"- in this process the discriminator must have learned a rich representation of the data, which is clearly able to describe and identify samples corresponding to the training dataset. The discriminator is expected to have this property as the discriminator is the main source of the training signal for the generator in this formulation [1]. As such the discriminator portion of a GAN trained to generate CRISM pixel spectra is expected to have learned a very rich representation of the CRISM data, which is able to successfully discriminate between the various minerals present in the dataset. Given the richness of the representation learned by the discriminator, it can be surmised that a similarity analysis performed on the representation learned by the discriminator is far more potent than the similarity analysis performed in the original data space [52].

[1] More information on this model is presented in Section 3.2.1

This chapter describes a mineral mapping technique for hyperspectral CRISM data based on the features learned by the discriminator of a GAN trained on CRISM spectral data. This chapter will describe in detail all the aspects involved in the pipeline for mapping of known mineral spectra in the CRISM image database, namely, (i) describing the pre-processing steps, (ii) formulating an appropriate training dataset using the CRISM image library, (iii) building and training a GAN on the CRISM spectral dataset, (iv) demonstrating the improved discriminatory power of the adversarial features over classical metrics and summary parameters, (v) describing the complete mineral mapping pipeline for CRISM images, and (vi) illustrating the value of these mapping pipelines for a couple of locations on the Martian surface. [N.B.: Some of the material covered in this chapter has also been previously published in [76]]

3.1 Preprocessing steps for the CRISM image data

The CRISM image dataset [60] is well known to be affected by a multitude of noise and atmospheric distortion processes, even the most advanced version of the CRISM data product available publicly on the NASA CRISM Planetary Data System (PDS)[2], i.e. the Targeted Empirical Record (TER) and Map-projected Targeted Reduced Data Records (MTRDR) [45], are still affected by some amount of noise and distortions [See [45] for more information on specific distortions]. Minimizing the effect of the atmospheric distortions and noise is the most essential aspect of the preprocessing stage. Additionally, as described in the primer on hyperspectral data analysis (see Appendix A), each pixel-spectrum contains components that provide no mineralogical information such as the continuum etc.., which can also be eliminated to further highlight the mineralogical information of interest. This section will describe

[2] `https://pds-geosciences.wustl.edu/missions/mro/crism.htm`

all the preprocessing steps used for the mineral mapping pipeline to minimize the effect of these processes in the CRISM spectral data.

3.1.1 Atmospheric correction and denoising

Historically, the presence of various atmospheric and noise effects in CRISM data products is ameliorated using the so called "volcano scan" [60, 61, 78, 79] method. This is a legacy technique for atmospheric correction that has been previously also used for OMEGA data products [51, 63]. The technique leverages empirical transmission spectra derived from the ratio of I/F spectra at the base and summit of Olympus Mons (the tallest mountain/volcano on Mars). The atmospheric contribution in each CRISM pixel is removed by dividing each pixel in the CRISM image by an empirically scaled transmission spectrum. This volcano-scan correction is well known in the community and is available as part of the CRISM Analysis Toolkit which is available in the CRISM PDS and has been the most popular method for atmospheric correction and has been leveraged for the creation of PDS products including the most advanced TER/MTRDR products were created by using a version of the volcano scan corrections.

Recently, Itoh & Parente [45] have recently proposed a novel method to do a better job eliminating these atmospheric artifacts and severe noise present in the CRISM image data. This method derives a refined transmission spectrum from the observed I/F spectrum inside the image. This method works by estimating the transmission spectrum by considering the light propagation through the atmosphere and surface mixing. The method also uses this modeling to identify and correct spectral regions that are affected by severe noise. Experimental results have shown this method produces fewer column-dependent artifacts, in the spectral regions $1.1 - 1.7\ \mu m$ and $1.9{-}2.1\ \mu m$ (regions with strong atmospheric absorptions). The new method also does a better job in identifying and correcting bands suffering from severe noise compared

to the volcano scan method. Since the technique is highly data-driven it is also better able to contend with adverse scenarios including elevated detector temperatures and water ice aerosols.

3.1.2 Continuum removal and scaling

As has been mentioned in Section A.1, each spectrum appears to be a combination of smooth global shape (continuum) and some localized absorptions (which are spread over small spectral regions), referred to as absorption bands. Since mineralogical interpretation is generally performed using these absorption bands, the continuum can be generally considered a nuisance factor for the mineral identification/discrimination efforts. Continuum removal enhances the absorption bands(as shown in Figure 3.1 middle column). Since continuum spectra in the short-wave infrared (SWIR) are generally convex or flat, we can estimate the continuum by fitting a set of convex line segments (a convex hull) to the spectrum [23]. We estimated the convex hull by using the iterative algorithm described in [9]. Finally, we eliminate the continuum from the spectrum by ratioing the pixel spectrum at a specific wavelength with the continuum value at that same wavelength.

While continuum removal minimizes the confusion from factors like illumination-effects etc., even after continuum removal it is clear that the absorption bands of some end-members are much bigger than others (see Figure 3.1 middle column). Mineral identification is not based on the shape and position of various absorption features as opposed to the overall size of the bands, which depends on factors like the spectral ranges used for continuum removal, the slope in the spectrum, and grain size etc. Further, the size of the absorption band cannot be directly linked to mineral abundance due to factors including noise, other admixtures, and grain size. For this reason, we introduce a *band normalization* step. Wherein, we normalize each spectrum so that the smallest value in each continuum removed spectrum is the same, the effect of

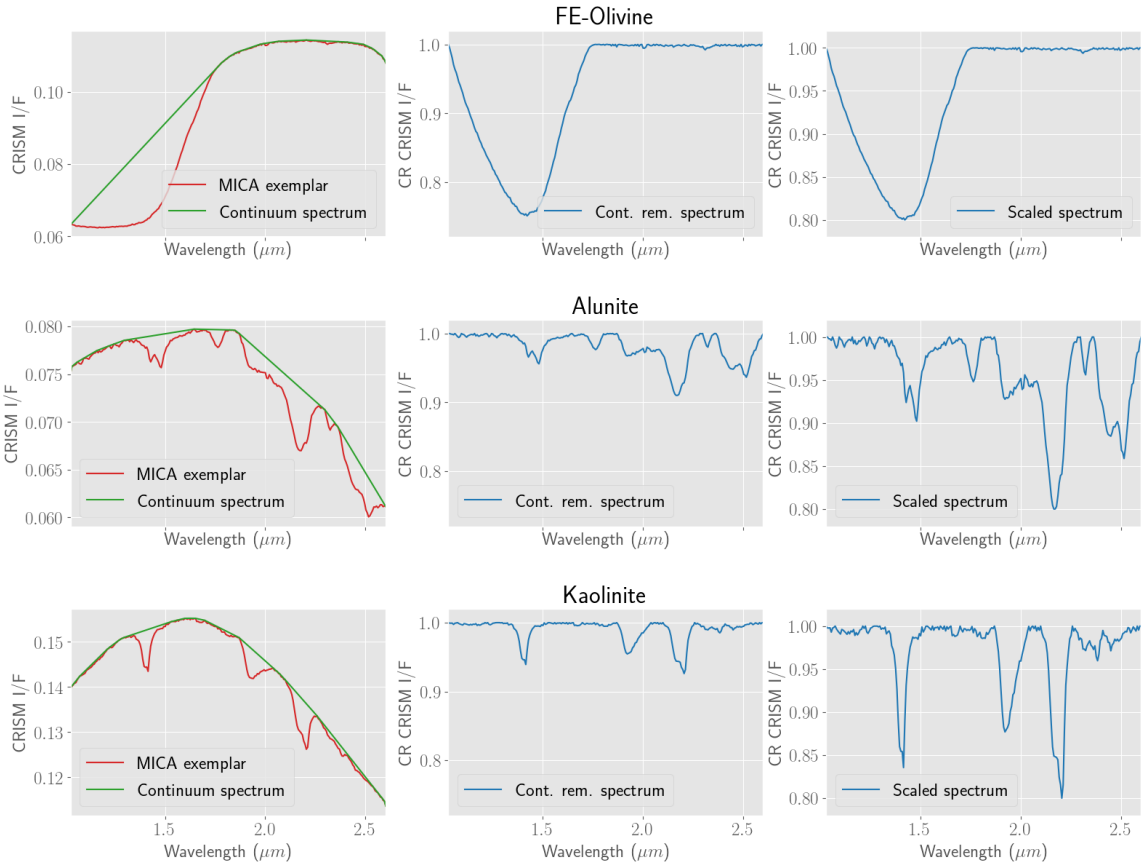

Figure 3.1. Illustration of the effect of the preprocessing steps. The original spectrum and the continuum (left column), the continuum removal (middle column) and scaling (right column) on CRISM I/F spectrum

this step is shown in Figure 3.1 right column. Notice that these steps in combination cause the bands to be more clearly visible and eliminates many of the known nuisance factors from the spectrum.

3.1.3 Eliminating uninteresting spectra

Many of the pixel spectra in the CRISM database appear to have no obvious absorption features. A spectrum is considered as "neutral" or "bland" if the continuum-removed spectrum (non-scaled) has no absorptions that are even 1% of the spectrum continuum. The 1% threshold was chosen based on the amount of noise commonly seen in CRISM pixel spectrum from multiple images. Given the amount of noise seen

in these images, absorptions that are of this size by themselves are not considered sufficient for mineral identification.

3.2 Creating the adversarial feature extractor

This section describes the various aspects of the model training to create a feature extractor using the discriminator portion of the Generative Adversarial Networks (GAN). It will describe the various aspects involved in training the feature extractor.

3.2.1 A brief description of Generative Adversarial Networks (GAN)

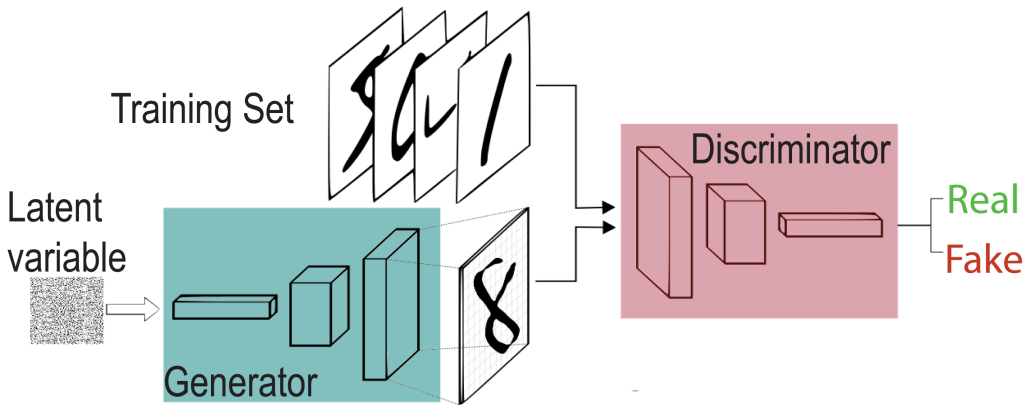

Figure 3.2. The general GAN scheme[3]

Classical machine learning algorithms referred to as *discriminative* or *conditional* models leverage the presence of labeled examples to estimate or "learn" the mapping between the output (or dependent variable) given the corresponding input (or independent variable) (e.g. logistic regression [25] or decision trees [49]). Generative models on the other hand attempt to learn a representation of the data distribution. Such models may be either *explicit* (such as variational auto-encoders [48]) or *implicit*

[3]This figure has been extracted from the blog "An intuitive introduction to Generative Adversarial Networks (GANs)", available online at `freecodecamp.org/news/an-intuitive-introduction-to-generative-adversarial-networks-gans-7a2264a81394/`

(such as GANs [36]). Explicit generative models assume a specific model for the data distribution and then attempt to estimate the parameters of the specific distribution. On the other hand the implicit generative models do not explicitly model/estimate the distribution. Rather the implicit generative models describe a stochastic process that can be used to directly generate samples from the data distribution. [N.B. For more information on implicit generative models readers are referred to Section II in Goodfellow et al. (2016) [36].]

The GANs specifically work by pitting two networks against each other in a zero sum game with competing objectives. A schematic representation of GANs is shown in Fig. 3.2. The first network referred to as the *generator* (the green network in Fig. 3.2) attempts to sample the data distribution. The generator accepts a low-dimensional random noise vector to activate the network (i.e. input) and attempts to generate a sample from the data distribution. The other network referred to as the *discriminator* (shown in red in Fig. 3.2), attempts to separate samples generated by the generator from the true data samples. GAN training consists of two steps: (i) discriminator training, which leverages a combination of equal number true data samples and samples generated by the current generator to improve the discrimination between the two classes. This can be achieved by training the network using a simple binary-cross entropy loss; (ii) generator training, in this stage the generator is trained to fool the discriminator. To achieve this the generator is chained with the discriminator, the discriminator weights are frozen and the generator output is used as the input to the discriminator. Following this the generator is trained to reduce the binary cross entropy with flipped labels. Overall, this step is designed to train the generator to create samples that the discriminator identifies as "real". The above two training updates are alternated unto convergence. At convergence, the samples from the generator should be completely indistinguishable from real data samples. Further, since the discriminator is constantly being trained to separate real samples

from fake samples, it is expected that representation for an input sample learned by the discriminator is rich and very useful for discrimination.

3.2.2 Creating the training dataset

To build a GAN that can model/recreate spectral data, a large spectral dataset representative of the spectral variability present in the CRISM spectral dataset was required. Such a dataset was created by sampling spectra from the ∼30 CRISM images. To further ensure as much spectral variability as possible is represented in the training dataset as possible, spectra were specifically drawn from the CRISM images that were used for the clearest mineral identifications in the CRISM image database. The clearest CRISM mineral identifications are referred to as the Minerals Identified through CRISM Analysis (MICA library)[90] [4].

Additionally, since the primary effort is to ensure that the model is able to learn "spectral features" we use the various preprocessing steps mentioned in Section 3.1. Also, rather than use the original noisy spectra which contain some remnant noise/ distortions we train the model by using the noiseless approximations of CRISM pixel spectra that are available as a byproduct of the denoising algorithm [45] mentioned previously. This procedure is valuable since the aim is to enable the model to discriminate mineral shapes, so even if at a pixel-level, the reconstructions are not accurate the presence of these clean spectra allows the model to focus its full descriptive power on capturing the mineral absorption features present in the CRISM data instead of any noise features. Each spectrum in this dataset contains the mineral absorptions of chosen CRISM pixels in the spectral range $1.0 - 2.6\mu$m over 240 spectral bands. In total, about a million total spectra were chosen and placed in the training dataset.

[4]The data corresponding to these observations is available at `http://crism.jhuapl.edu/data/mica/`

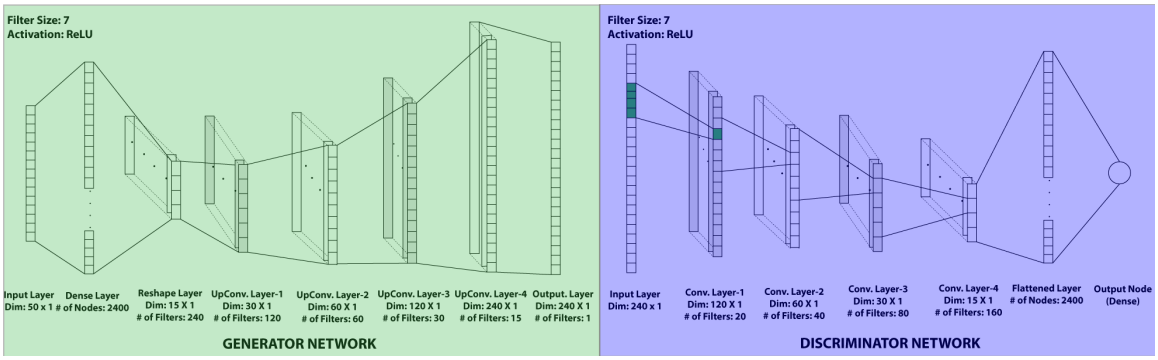

Figure 3.3. The architecture of the GAN used to create the CRISM feature extractor

3.2.3 Training the model: architecture and hyperparameters

The full model architecture is shown in Figure 3.3. The generator input is a noise vector from $\mathbb{R}^{50}$. The generator has a dense input layer followed by 4 convolutional layers, the output is spectrum (vector) $\in \mathbb{R}^{240}$. This spectrum is the input for the discriminator, which contains 4 convolutional layers with a 1 node dense output layer. Every dense or convolutional layer in the figure also followed by some common normalization layers, namely- *batch normalization* (to stabilize training by normalizing the input of each unit to have zero mean and norm-1) [44] and *dropout* (to minimize over fitting by randomly dropping a specific fraction of weights at each training step) [86]. In particular, this network is an instance of the popular Deep Convolutional Generative Adversarial Network (DCGAN) [73] adapted to 1D dimensional data. As such, convolutional networks are used for both the generator and the discriminator, with a convolutional filter size of 7 and a stride of 2. Overall, these networks have approximately 1 million parameters combined.

The network is then trained to convergence (which took about $250,000$ iterations), each training step includes alternative updates of the discriminator and generator weights. Each batch contained 256 randomly drawn samples. The stochastic Adam [47] optimization technique was used for training, with a learning rate ($\alpha = 1e^{-4}$), and the multipliers for the decay rate to the various moments set to ($\beta_1 = 0.9$, & $\beta_2 =$

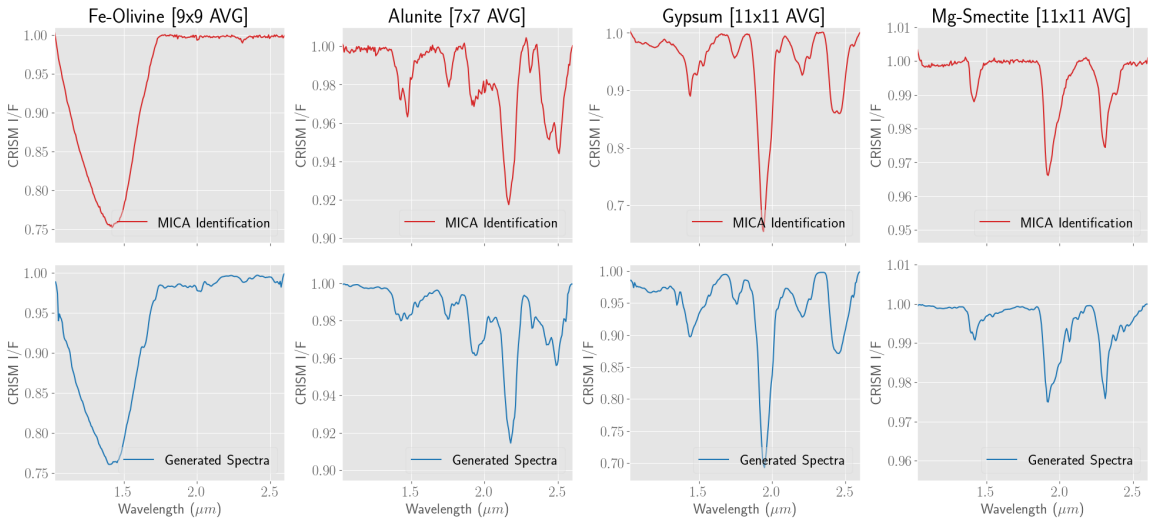

Figure 3.4. Comparison of real continuum rremoved Spectra from CRISM Images (top row shown in red) to continuum removed spectra generated by the generator (bottom row shown in blue)

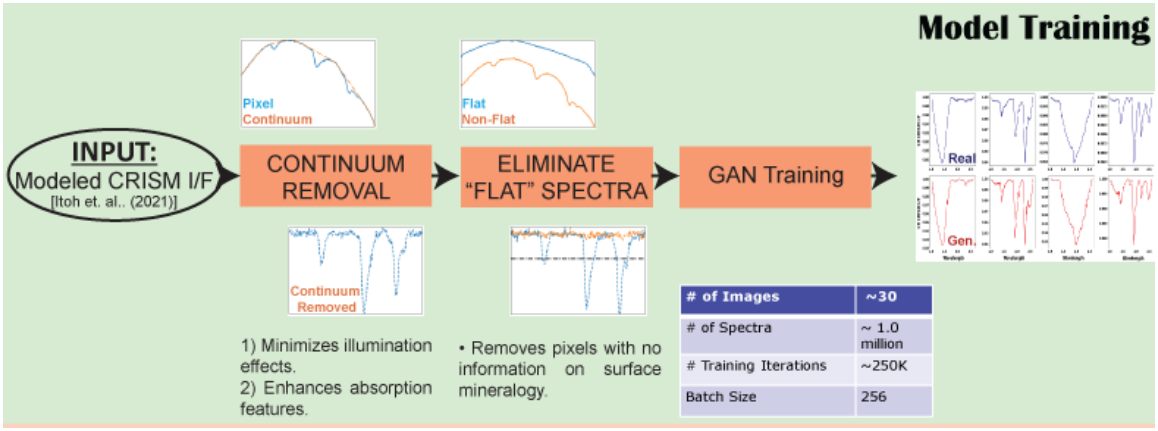

Figure 3.5. The scheme used for training a GAN for CRISM spectra.

0.99). On a machine with one TITAN Xp ® GPU [5], the full training process takes around 6 hours. At equilibrium, the GAN is able to produce continuum removed spectra, which are very similar to the real continuum removed CRISM spectra. Fig.

[5]Titan Xp® is a registered trademark of the NVIDIA Corporation. Full specs on the Titan Xp can be found at `https://www.nvidia.com/en-us/titan/titan-xp/`

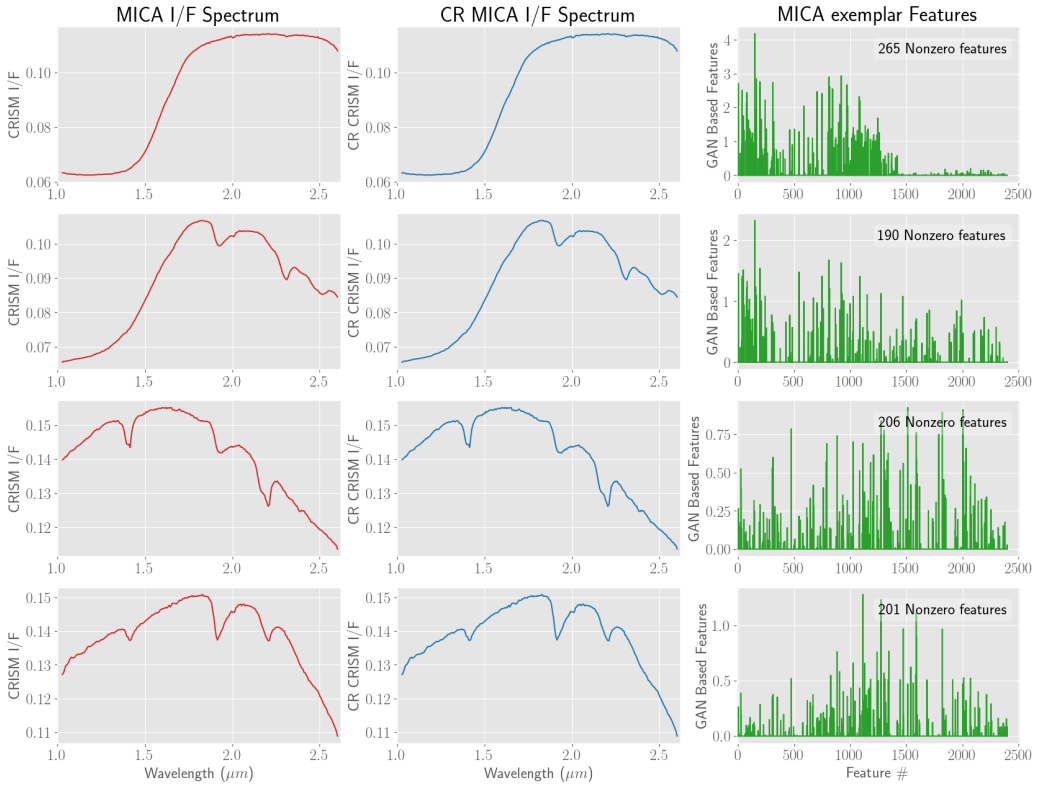

Figure 3.6. The GAN representations for various end-members from the MICA Library. (A) Fe-olivine. (B) Mg-carbonate. (C) Kaolinite/halloysite. and (D) Al-smectite.

3.4, shows a comparison of some real continuum removed MICA spectra(shown in red on the top row), to their generated analogues (shown in blue on the bottom row). Note, in particular, that the generated spectra appear quite similar to the spectrum from the training dataset, and have the same characteristic absorption features as the real spectra. The only difference appears to be a small amount of extra noise in the generated spectra.

The full procedure needed for training the GAN required for mineral mapping is shown in Fig. 3.5.

3.3 Using the trained discriminator as a feature extractor

As was discussed briefly in the introduction of this chapter since the discriminator is trained to identify spectra, the representation learned by the discriminator is highly descriptive of CRISM pixel-type spectrum, this section will describe the procedure for using the trained discriminator as a feature extractor. It will also describe the procedure to use the features in conjunction with simple similarity metrics to perform improved mineral discrimination. Finally, the next section (Sec. 3.4) will show the results of mapping some well-known mineral detections present in the CRISM image database referred to as "exemplars" on some well-studied regions of the Martian surface.

3.3.1 Feature extraction from pixel-level CRISM spectral data

Pixel-level spectral data from the relevant spectral range (i.e. $1.0 - 2.6\ \mu m$) are extracted from CRISM images of interest. The data are processed using the various preprocessing steps described in Section 3.1. The only significant difference between this stage and the training stage is that rather than using the modeled continuum-removed spectra from an unmixing model, actual continuum-removed pixel spectrum from CRISM images are considered for discrimination. These pixels are generally noisier than the model derived continuum-removed CRISM type spectra used in training the GAN model.

For each pixel (which is not considered "bland" or "neutral" as per the technique in Section 3.1.3) considered for feature extraction, feature extraction is performed by using a modified version of the discriminator network, wherein similar to previous works such as Radford et al. [73], the discriminator without the final layer is used as the feature extractor-i.e. the output of the penultimate layer of the discriminator are used as the features. Since the discriminator has been fully-trained to identify

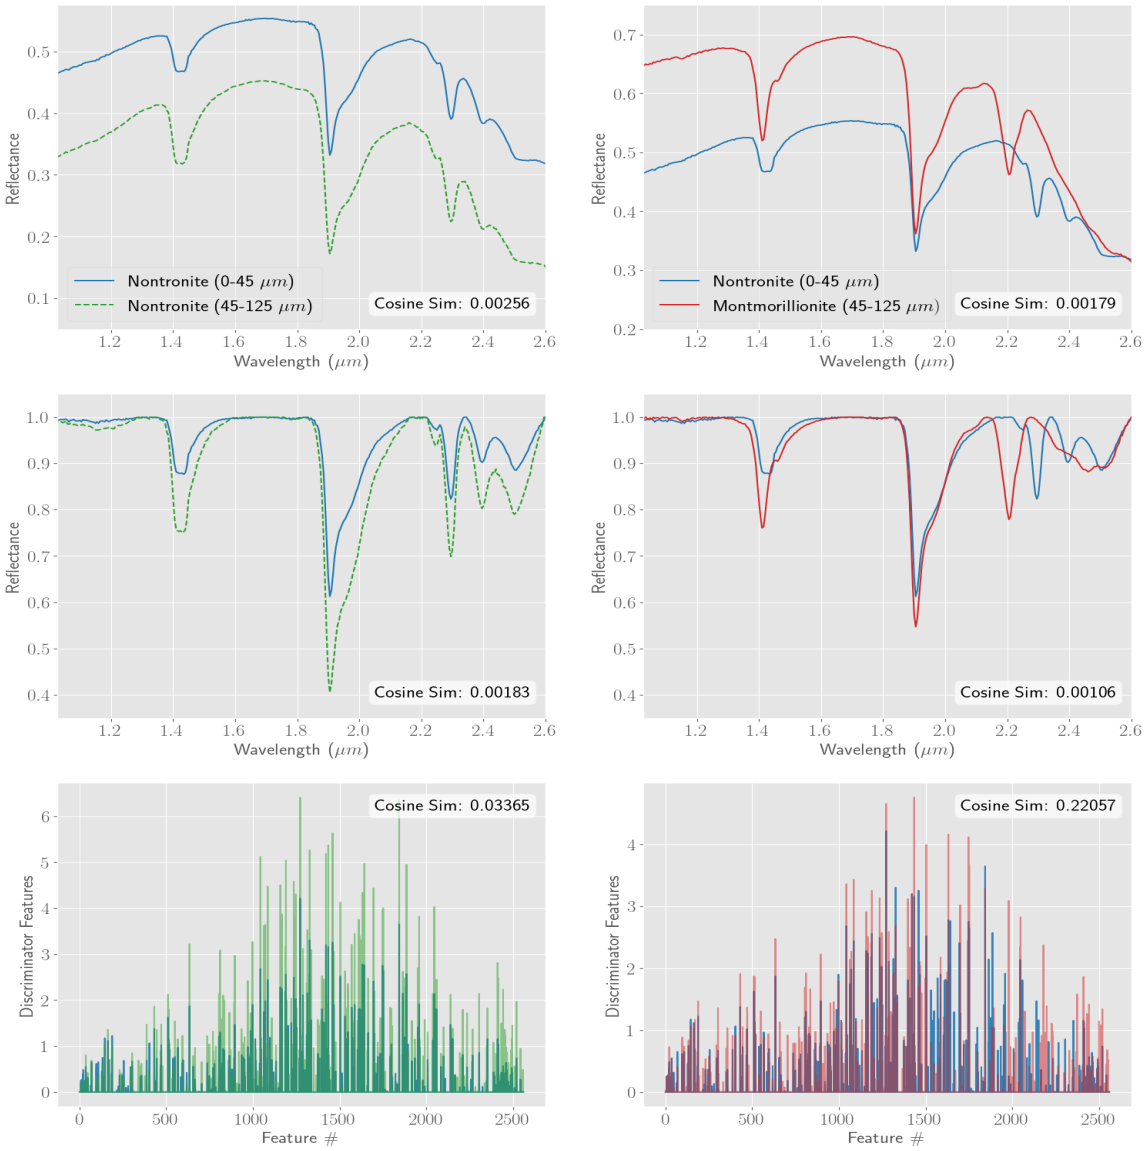

Figure 3.7. Illustration of the improved discriminative strength of the GAN feature space. Spectral Angle divergence between intra-class (2 nontronite) and inter-class (nontronite and montmorillionite) members from the RELAB database in original reflectance space (top row). Same analysis in the continuum removed reflectance space (middle row). Same analysis in the GAN based feature space (bottom row). Note that cleary, in the GAN-feature space (unlike the original and CR spaces) *in class* distance is much smaller than the inter-class divergences.

CRISM type spectra, it is expected that the mineral discrimination will be easier in this newly derived feature space.

3.3.2 Improved mineral discrimination in the feature space

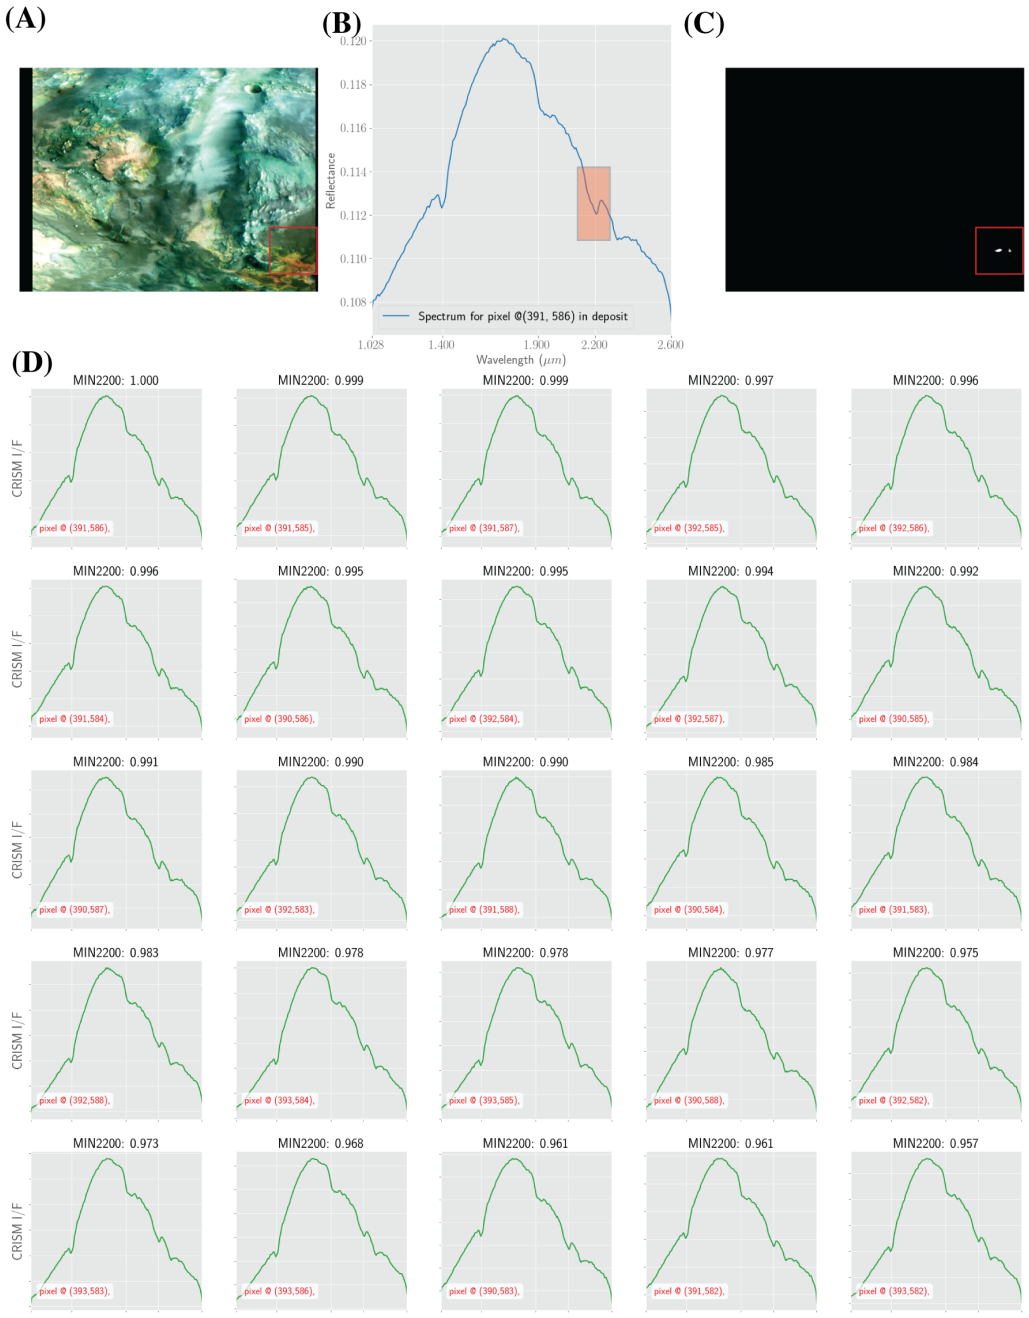

Figure 3.8. Mapping the kaolinite deposit in FRT000097E2 using the features from the GAN discriminator. (A) RGB Composite of FRT000097E2 with the kaolinite deposit highlighted. (B) the characteristic spectrum corresponding to the deposit. (C) The pixels with high similarity to the characteristic spectra in the feature space. (D) 20 most similar spectra to the characteristic spectrum in the image.

The first relevant question for mineral discrimination in this novel feature space is to prove that there is indeed increased discrimination between the various mineral classes in this novel feature space. The first prerequisite for such discrimination is that the features corresponding to the manually identified exemplar spectra are different and easily discriminable. To illustrate this Fig. 3.6 shows the representations generated by the discriminator for some MICA exemplar spectra. It is interesting to note that the features generated for the various exemplars are diverse in a variety of ways such as: (i) the number of non-zero features; (ii) the profile of where the features have large values as opposed to smaller values; and (iii) the average values of the features etc.

Even for exemplar signatures that are very similar to each other, the representation learned by the models is quite different, e.g., the MICA minerals kaolinite/halloysite (shown in Fig. 3.6 third row) and Al-smectite (shown in Fig. 3.6 bottom row). Both have characteristic absorptions at 1.4, 1.9 and 2.2μm, with minor differences in the band shapes at 1.4 and 2.2μm where kaolinite appears to have a doublet (an absorption band with two local minima) as opposed to an absorption band (with a single minimum) since the bands are in the same position cosine similarities are expected to be similar. On the other hand, the representations of two spectra have different numbers of non-zero features and completely different profiles (the Al-smectite profile appears to have large values in the middle, while the large values are more spread out for the kaolinite).

To further illustrate the improved spectral discrimination in the feature space consider the example shown in Section 2.1. The example shows the inability of classical global metrics in spectral space to accurately measure spectral similarity since spectral discrimination is based on subtle differences. The issue is further illustrated in the top and middle rows of Fig. 3.7, wherein the distance/divergence between members of the same class is bigger than the distance/divergence between members of different

classes. On the other hand note that (in the bottom row) the representations of the two nontronite spectra are almost identical, i.e. they have the peaks and valleys at the same locations. On the other hand the nontronite and montmorillionite spectra have slightly different profiles with many instances where the peak and valleys do not really match. From being farther away from each other in the reflectance I/F space the two nontronite spectra are almost an order of magnitude closer to each other in the feature space.

Similarly, even for the task of identifying the kaolin-type spectra from the CRISM image FRT000097E2, mentioned in Sec. 2.2, these discriminator features are far more capable of identifying the exact deposit and it exact boundaries. Notice the similarity map shown in Fig. 3.8 (C) exhibit almost no noise, rather only identifies the deposit of interest. Similarly, look at the 20 spectra that are closest to the characteristic spectrum (shown Fig. 3.8 (D)) in the novel feature space; they all have the same clear mineral absorption features as the characteristic spectrum. There are no noisy spectra with high scores as was seen with the summary parameters (shown in Sec. 2.2).

3.3.3 Quantifying the improved discrimination of the features learned by the discriminator

The previous subsection shows that the feature space of the discriminator is better able to discriminate between mineral spectra from different classes with similar shapes. Given these observations, it is essential to quantify the improved discrimination of the discriminator feature space relative to the original reflectance space. As has previously discussed there are not a lot of labeled CRISM spectra on which a statistical test can be performed. To get around this issue, a simulated CRISM dataset was created to ensure that a sufficiently large labeled dataset is available for analysis.

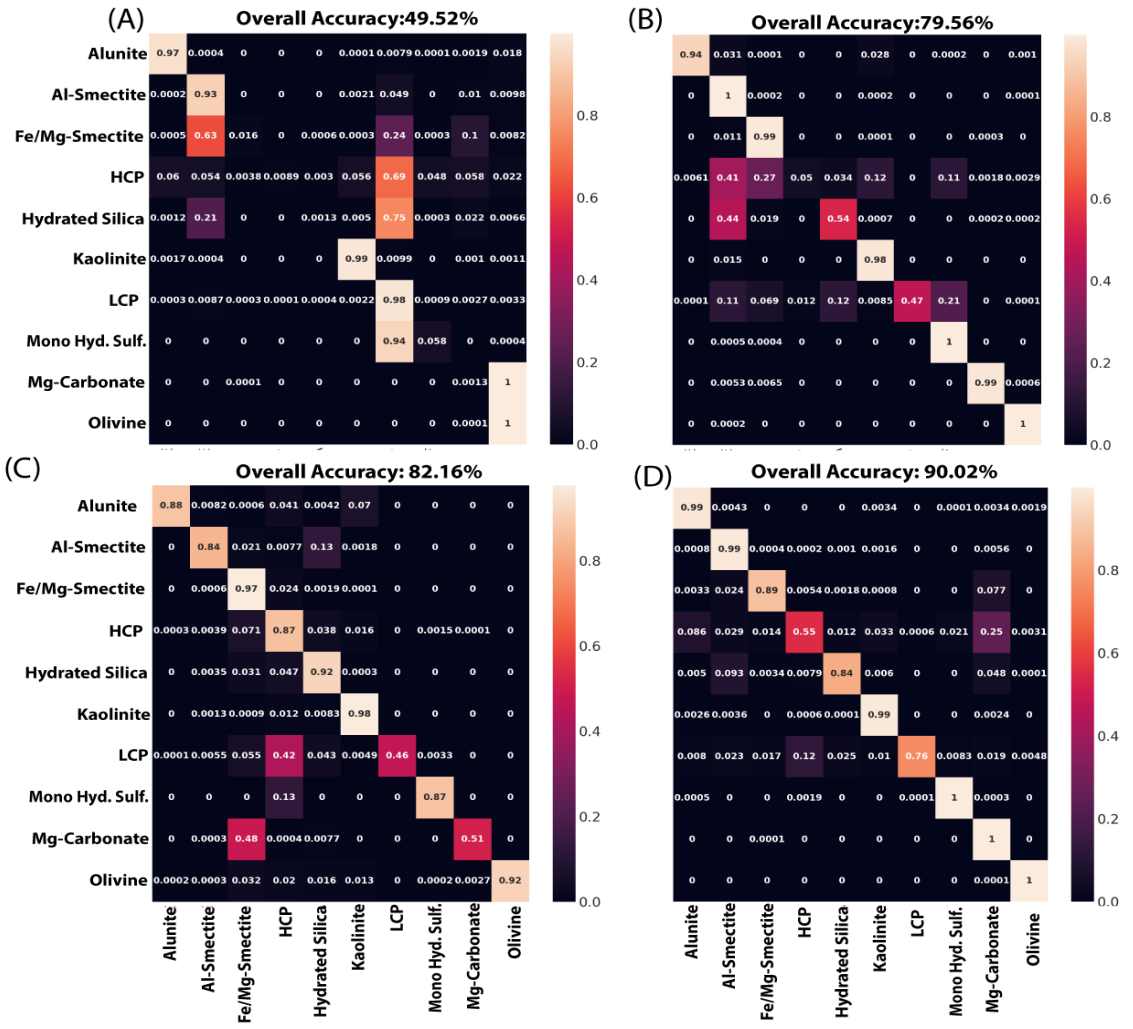

Figure 3.9. The Classification performance on the simulated dataset under the various regimes. (A) Regime-1. (CR + SAD) (B) Regime-2 (CR + Norm. + SAD). (C) Regime-3 (CR + GAN Rep. + SAD). (D) Regime-4 (CR + Norm + GAN Rep. + SAD).

This simulated dataset attempts to mimic the spectra in the CRISM image dataset as much as possible. The dataset contains spectra with 10 constituent endmembers referred to as "exemplars" (these are chosen from some of the most commonly seen minerals in the CRISM image database). The mineral classes used in creating the

mixtures are shown in the labels of the various confusion matrices in Fig. 3.9. Each spectrum is modeled as mixture of up to 5 endmembers, with a single dominant endmember. The dominant endmember is more than 70% by weight in each mixed sample. Spectra with the same dominant endmember are considered to be in the same class. There are $10,000$ spectrum corresponding to each class. Following this various atmospheric effects, atmospheric distortions, and noise commonly seen in martian data are simulated and added to the mixed spectrum. The presence of dust was modeled by randomly mixing each spectrum with a small fraction (between $0-30\%$) of a "bland" spectrum (i.e., a flat spectrum with no clear absorption features). Finally to model atmospheric correction/noise seen in CRISM pixel, we add a randomly selected residual from the atmospheric correction method[45]. (The residual is estimated as the difference between the CRISM spectrum and the model spectrum estimated by the atmospheric correction method).

This simulated dataset contains CRISM-like spectral data wherein each sample is associated with a corresponding label (based on the dominant mineral endmember). Using this labeled dataset one can setup a classification task to provide a quantitative measure of improved discrimination of the GAN discriminator feature space. Consider for example a simple classification task where each samples is assigned the same class label as the closest "exemplar" (closest is measured using the cosine distance or Spectral Angle Divergence (SAD)). To illustrate the impact of the continuum removal (CR) and normalization/feature extraction techniques (Norm.) we consider the classification performance under four regimes (i) CR + SAD; (ii) CR + Norm. + SAD; (iii) CR + GAN Rep. + SAD; (iv) CR + Norm. + GAN Rep. + SAD. In the first regime (CR+SAD), the spectra were processed using a simple continuum removal followed by using the SAD based classification scheme described above. In the second regime (CR + Norm. + SAD), both continuum removal and the band normalization (these normalizations are described in Sec. 3.1.2) before performing the

classification using the SAD. In the third regime (CR + GAN Rep. + SAD), following continuum removal a GAN based feature representation of both the pure mineral library and the simulated dataset are extracted, and the SAD-based classification was performed in this feature space. In the final regime (CR + Norm. + GAN Rep. + SAD), both continuum removal and band normalization before extracting the GAN representation, and the SAD-based classification was performed in this representation space. Finally, the estimated labels were compared to the true labels to estimate the various regimes' performance.

The results of each of these regimes in terms of confusion matrices are shown in Fig. 3.9 (A)-(D) (where the entry in row i and column j of each sub-figure is the fraction of the samples which belongs to class i that were predicted as being in class j). Notice that regime-1 (CR + SAD) performed rather poorly (see Fig. 3.9 (A)). On the other hand, when we added the normalization in regime-2, the performance improved quite significantly (see Fig. 3.9 (B)). In these regimes the classification scheme confuses endmembers that are similar to each other, for example Fig. 3.9 (A) & (B) shows poor discrimination for endmembers with absorptions in similar spectral regions like Al-smectite and Hydrated Silica. The other issues appear to be for minerals without sharp absorptions such as LCP/HCP. Using the GAN representation in Regime-3, the performance is already superior to the previous two regimes (see Fig. 3.9 (C), but the best results are obtained in Regime-4, which includes both the normalization and the GAN representation (see Fig. 3.9 (D)). This experiment clearly underlines the value of both the GAN based representation and the various normalization steps in improving mineral discrimination.

3.4 Mapping minerals in CRISM images

The previous section illustrates that the GAN-based features have significantly enhanced discrimination compared to the data in the original reflectance space and

Mg-Olivine	Fe-Olivine	Mg-Carbonate	Low-Ca Pyroxene	Monohydrated Sulfate
Jarosite	Gypsum	Kaolinite	Al-smectite	Fe/Mg-Smectite
Prehnite	Hydrated Silica	Polyhydrated Sulfate	Illite	Talc
Serpentine	Epidote	Analcime	Hydroxylated Fe-Sulfate	High-Ca Pyroxene
Chloride	Margarite	Hematite		
Chlorite	Ca/Fe-Carbonate	H_2O-Ice	CO_2-Ice	Bassanite

Table 3.1. List of mineral spectra from the MICA Library chosen as exemplars. The minerals colored in pink are excluded as they are not suitable, and the minerals colored in red have been left for future use

simple algorithms like spectral summary parameters. This section illustrates a the use of this GAN based feature space for the mapping of CRISM images.

3.4.1 Leveraging the GAN-discriminator feature for mapping spectral shapes in CRISM images

In the initial mapping approach, the mapping pipeline attempts to map some of the previously identified mineral shapes from CRISM data. The mineral shapes were extracted from the Minerals Identified by CRISM Analysis (MICA) [90], these mineral detections have been verified by consensus in the community and are not subject to interpretation. The various minerals in the CRISM MICA library are shown in Table. 3.4. In the initial implementation, only the minerals shown in black are used for mapping the current images. Some of the MICA library minerals were excluded as they were not suitable for our technique; such minerals are colored in pink in Table. 3.4. For example, the Hematite spectrum is excluded as its characteristic absorption bands are outside the spectral range $1.0 - 2.6$ μm, which is the region of interest in this study. Further, the minerals Margarite and Chlorite have been excluded as the numerator pixel spectra associated with the detections appear neutral and do not show any absorptions in the spectral range $1.0 - 2.6$ μm. The other minerals marked in red were not considered at this stage. In addition to the known MICA spectrum,

some mixed spectra that have been identified in various scenes by expert analysis, such as the hydrated silica-bearing spectrum identified by [87] were also included. The mapping method was also used in community efforts such as "the CRISM Fandango" [66, 65] to generate improved mineral mapping products for Jezero crater. These products were validated by consensus across various mineral identification/mapping techniques and expert analysis (Improved mapping products from these efforts will be described in a future publication).

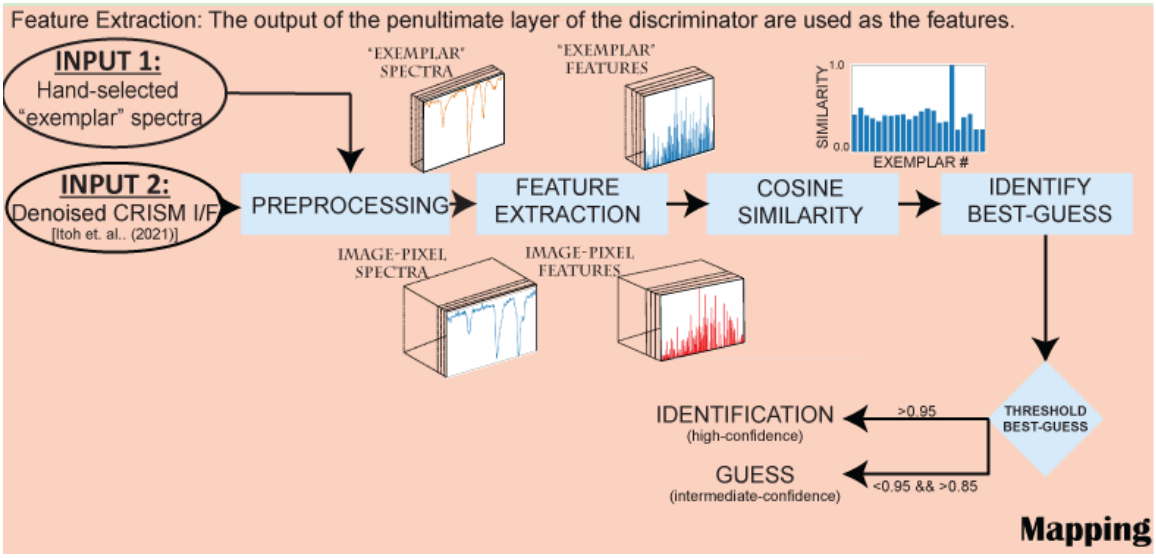

Figure 3.10. The scheme using the features from the GAN discriminator to map mineral shapes in CRISM images.

A summary of the steps involved in mapping a single CRISM pixel is shown in Fig. 3.10. The various steps in mapping the mineral shapes are:

1. **Preprocessing:** In the pre-processing stage each pixel is processed though the various procedures such as *denoising*, *continuum removal*, and *neutral-spectrum* removal, described in Section. 3.1.

2. **Feature Extraction and Similarity Analysis:** In this stage, for each pixel, a feature space representation is extracted using the GAN discriminator. Fol-

lowing, this the pixel representation is compared to the representation of known mineral shapes or "exemplars". The similarity is measured using a simple distance metric, in this case the cosine similarity. Following this step each pixel is represented by a vector containing the similarity scores of each pixel to the various exemplars in the mapping library.

3. **Identify Best Guess:** First, we threshold the similarity scores such that all scores below 0.707 are set to 0 (as a score below this level indicates that the pixel representation subtends an angle $> 45°$ from the exemplar). Then, we also zero out all but the highest score in the similarity vector for each pixel. Any non-zero score in the similarity vector identifies the spectrum as closest in shape to a specific member in the exemplar library.

4. **Separating Best Guesses into Different Confidence levels:** The cosine similarity calculated over the previous stages provides us a degree of confidence in our detection over various classes. Pixel-spectra with very high similarity (i.e., cosine similarity > 0.95) to the exemplars are referred to as the *identifications*. Empirically we have verified that the pixels labeled as identifications have the same absorption features (in terms of band positions, band shapes, band minima, etc.) as the exemplar, such that based on visual analysis, an expert would also classify the spectrum as being similar to the exemplar. The next level are pixels, which we refer to as *guesses*, whereby guess pixels have an intermediate similarity to the exemplar ($0.85 <$ cosine similarity < 0.95). The spectra also show moderately high similarity to the exemplar in terms of the mineral absorption features, but there are certain clear differences. In most cases, these differences are not significant enough to produce a changed class assignment from an expert (differences such as a change in the relative size of the hydration band, etc. will affect spectral similarity but will not affect an

expert's labeling of the pixel spectrum). In some instances, the changes in relative band size, etc. may affect the labeling, and as such, the guesses need to be manually investigated to ensure perfect accuracy.

Following this process, it is necessary to create composite mineral maps that enable the user to simultaneously identify the various minerals present in the scene. To enable this a dictionary is created wherein each mineral shape in the "exemplar library" is assigned a single color or RGB code. Pixels identified with high confidence or "identifications" are marked in the mineral map using the exact color corresponding to the specific exemplar, while pixel identified with intermediate confidence or "guesses" are marked using subdued versions of the same color (i.e. 0.7 $\times$ mineral RGB code).

3.4.2 Illustration of the mineral mapping on a single CRISM image

To illustrate the mapping capability of the GAN discriminator features, consider the task of mapping the spectral shapes in a single CRISM image. FRT000093BE is an image in the Jezero Watershed (N:19.2363, E:76.3465) that shows significant spectral variability, and provides an excellent test case for the mapping abilities of the GAN based pipeline. Fig. 3.11 (A) shows an RGB composite of the region covered by the specific image. Fig. 3.11 (C) shows a browse product made with three well known summary parameters, wherein the summary parameters corresponding to the various channels are R: OLINDEX, G: BD2500, B: D2300. In these browse products, pixels that appear white (with a greenish tinge) represent Mg-Carbonate, pixels that are blue identify Fe/Mg-Smectite or Serpentine and red pixels would identify olivine. Note that the summary parameter map appears quite noisy, and identifying the presence of small deposits in this map would be quite challenging. Note that the olivine bearing units are not clearly visible in this browse product. Similarly, Fig. 3.11 (B), shows the mapping of well known mineral spectra using the GAN based discriminator features. The mineral identification maps are overlaid on the CRISM image RGB composite to

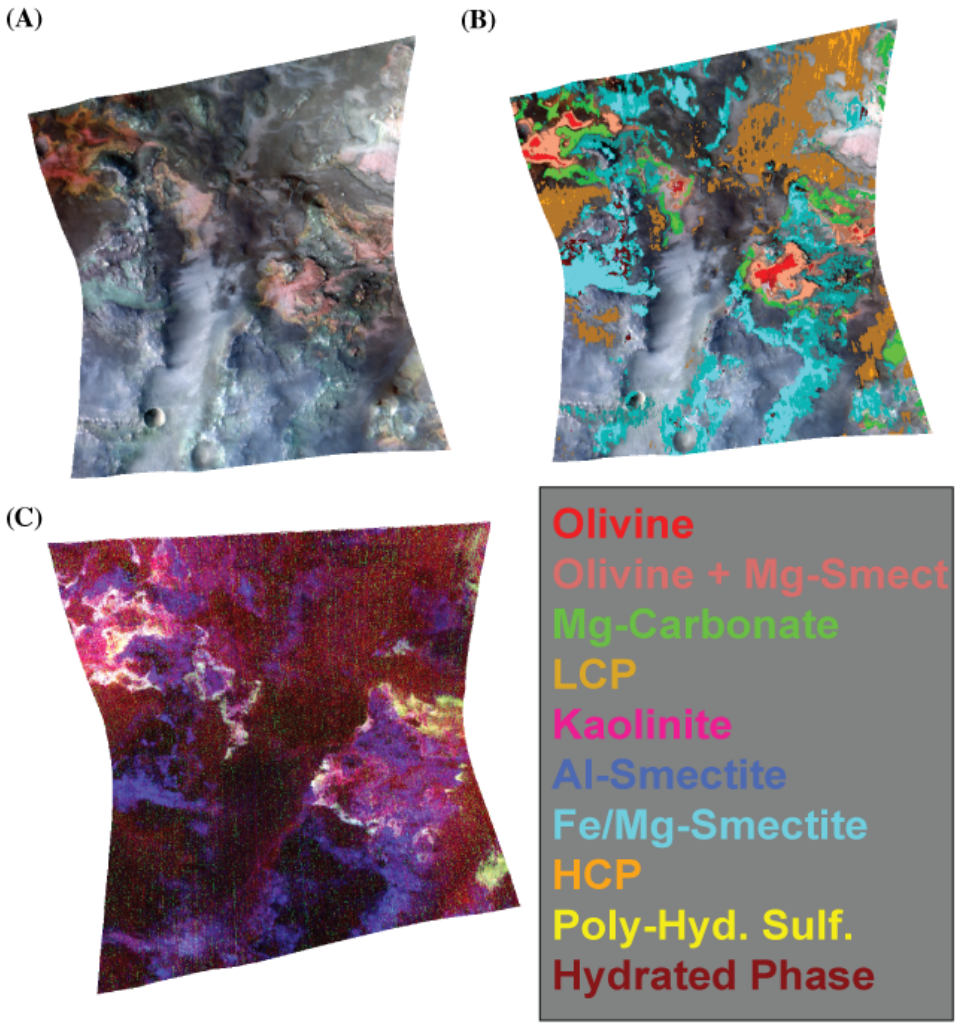

Figure 3.11. GAN based Mapping of the CRISM image. (A) RGB Composite of the image FRT000093BE. (B) Mineral Map using the GAN pipeline for FRT000093BE. (C) CRISM parameter browse product for FRT000093BE. (R: OLINDEX, G: BD2500, B: D2300).

allow the reader to appreciate the correlation between mineral detections and natural geological features present in the image. Pixels that have been assigned no color correspond to pixels not considered sufficiently similar to any of the mineral shapes in the exemplar library (or pixels eliminated from consideration in the preprocessing). The colors corresponding to the various mineral identifications are shown in the legend on the bottom right of Fig. 3.11.

The boldly colored pixels in Fig. 3.11 (B) represent **identifications**, while the subdued colors indicate **guesses**. It is also observed that the regions of high confidence *identifications* are surrounded by areas containing *guesses* or mixed classes, which agrees with our expectations of optical mixing at the margins of deposits, which leads to the presence of mixed classes or variations at the edges of various deposits. The regions identified as Mg-Carbonate and Fe/Mg-Smectite by the GAN representation (shown in Fig. 3.11 (B)) have quite a good correspondence to the appropriate regions in the browse product (shown in Fig. 3.11 (C)) while showing clearer margins and separations between the various classes. The detection map can further identify a variety of other minerals as well without needing to generate/analyze more browse products.

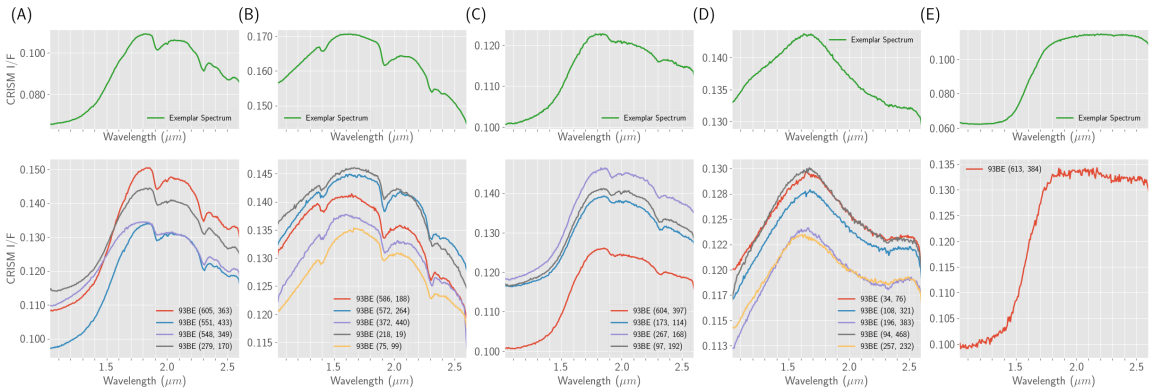

Figure 3.12. Comparison of mineral spectra for the image $FRT000093BE$ **identified** using the GAN representation to the exemplar spectrum for the minerals. (A) Mg-Carbonate. (B) Fe/Mg-Smectite. (C) Olivine/Mg-Smectite Mix. (D) High-Ca Pyroxene. (E) Fe-Olivine.

Fig. 3.12 illustrates that the pixel spectra from the regions marked by the GAN representations as *identifications* have the same absorptions as the "exemplar". Comparing the pixel spectrum *identified* in the CRISM image FRT000093BE as Mg-Carbonate, Fe/Mg-Smectite, olivine/Mg-Smectite Mix, High-Ca pyroxene, and Fe-olivine are shown in Fig. 3.12. We see that the spectra have very similar absorption

features across different spatial deposits to the "exemplars". The only differences are in factors like brightness, slope, etc., which have no bearing on the mineralogy. On the other hand, if we look at the spectra from the **guess** classes shown in Fig. 3.13, there are some differences in the absorption features, which reduces the model confidence. For example, while the detected spectra in Fig. 3.13 (A) and (B) have the bands at the same spectral location as the "exemplars" the characteristic absorptions are not as clear/deep. The detected spectrum in Fig. 3.13 (C) shows an additional absorption feature at 2.3 μm compared to the exemplar. Similarly, in Fig. 3.13 (D), the bands are not as deep as the exemplar (especially the band at 1.9 μm), and in Fig. 3.13 (E) the hydration band (at 1.9 μm) is not as deep, and the faint shape at the end of the spectrum around 2.5 μm is not as clear. In general, the guesses have many of the same absorption features as the associated "exemplar", but these spectra also exhibit some differences as described above. As such, the "exemplar" chosen is interpreted as the spectrum in the library that is closest to the given spectrum.

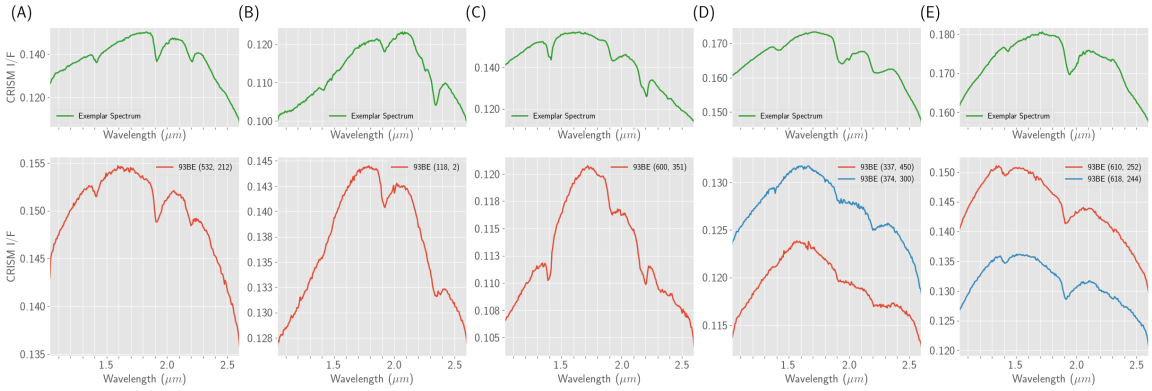

Figure 3.13. Comparison of mineral spectra for the image $FRT000093BE$ **guessed** using the GAN representation to the exemplar spectrum for the minerals. (A) Al-smectite. (B) Epidote. (C) Kaolinite/Halloysite. (D) Hydrated Silica. (E) Poly-Hydrated Sulfate.

It is also interesting to note that the the mineral maps for overlapping images in similar regions have very similar mineral maps. For examples, consider the CRISM image $FRT000097E2$, which almost completely overlaps with the image $FRT000093BE$.

The maps from the two images are shown in Fig. 3.14. Note that the maps are very similar in terms of the margins of the different exemplars and the corresponding geological units where these minerals are found.

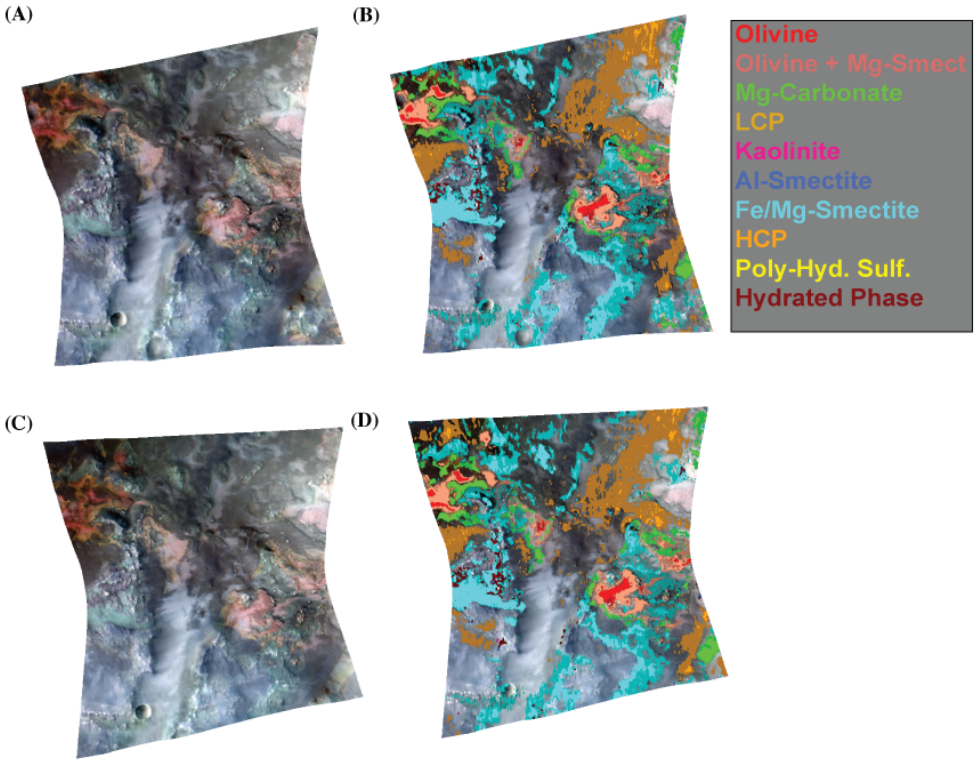

Figure 3.14. Comparison of mineral maps for overlapping CRISM images (A) RGB Composite of the image $FRT000093BE$ (B) Mineral Map using the GAN pipeline for $FRT000093BE$ (C) RGB Composite of the image $FRT000097E2$ (D) Mineral Map using the GAN pipeline for $FRT000097E2$

3.4.3 Mineral map composites by using multiple CRISM images over well known Martian regions

This mapping method can be used over multiple CRISM images from a specific spatial region on the Martian surface to create composite maps that enable the study of the various trends in the minerals seen on the surface. To illustrate this, mineral detection maps were generated for some of the CRISM images near Jezero Crater (E:18.4463, W:77.4565) and North East Syrtis (E:17.8899, W:77.1599). These regions

are well known and have been well studied in the literature as they were candidate landing sites for the "**NASA Perseverance Rover**" (Jezero crater was finally chosen as the landing spot for this mission).

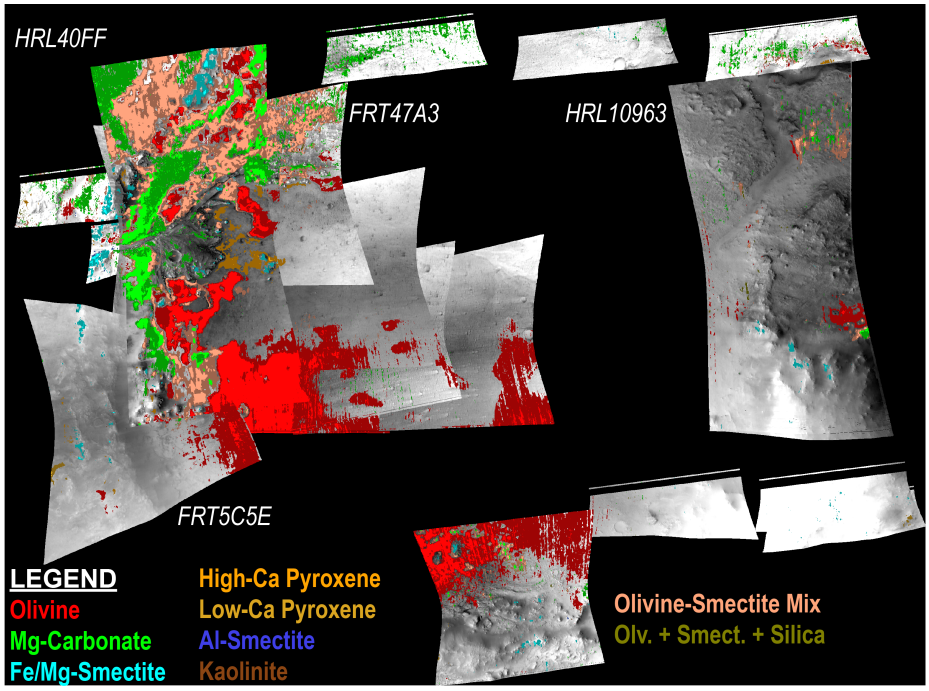

Figure 3.15. Mineral Identification Maps for the region around Jezero Crater. The pixels identified as minerals are shown by the colored pixels, and the proposed landing site is marked in yellow.

Figs. 3.15 & 3.17 shows these composite detection maps overlain on the CRISM images themselves and allows one to easily appreciate the trends of the various minerals present in the region. These maps show that the dominant mineralogy in these sites are the Mg-Carbonates (shown in green in Fig. 3.15 and 3.17) and Fe/Mg-smectites (shown in cyan in Figs. 3.15 and 3.17). There are also some olivine-smectite mixtures in both sites (shown in salmon in both figures). The various mineral phases detected and mapped in the Jezero region are shown in Fig. 3.16. The dominant

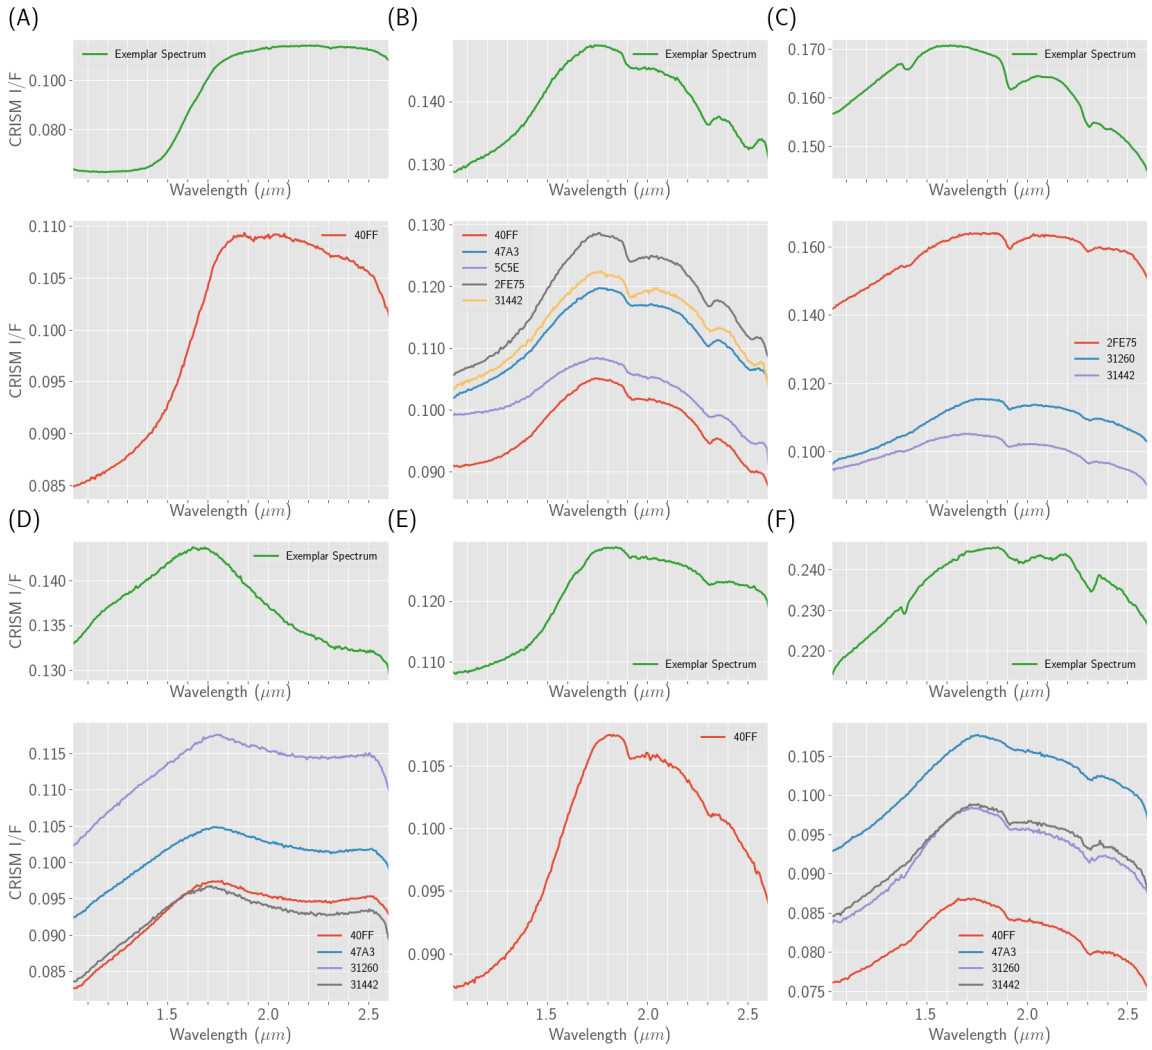

Figure 3.16. Identified Mineral Spectra from/near Jezero Crater. (A) Fe-Olivine. (B) Mg-Carbonate. (C) Fe/Mg-Smectite. (D) High-Ca Pyroxene. (E) Olivine-Smectite Mixture. (F) Serpentine.

phases in this region and Olivine (shown in red), Mg-Carbonate (shown in green), and olivine+Smectite Mix (shown in salmon).

The North East Syrtis region appears to have a significant amount of high-Ca pyroxene (shown in orange). The spectra for high-Ca pyroxene shown in Fig. 3.16 (D) and Fig. 3.18 (D) appear to generally have the same shape as the "exemplar", but there are some detections that appear to have a significantly different slope.

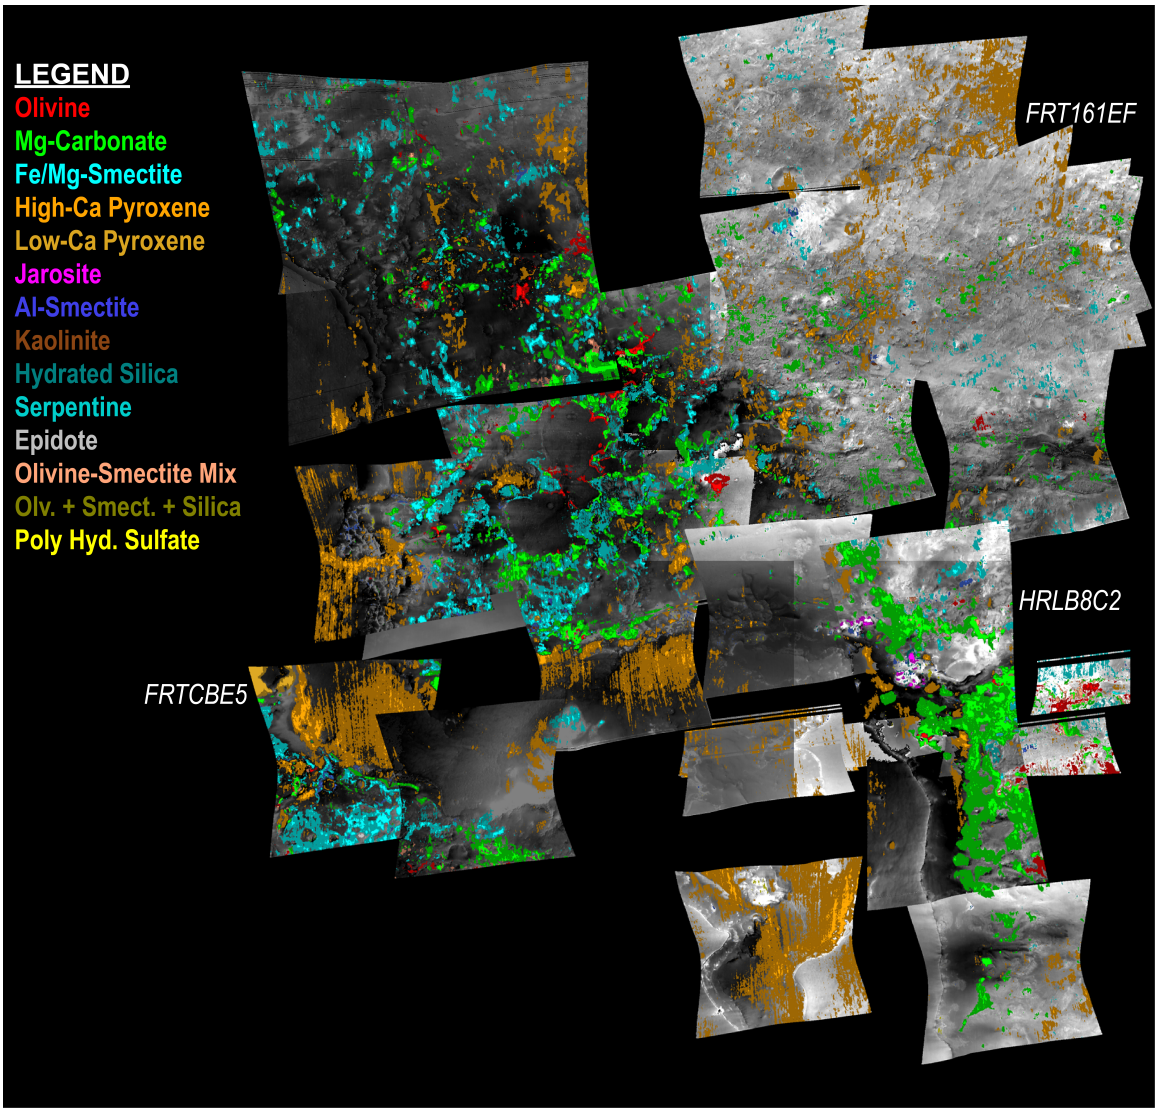

Figure 3.17. Mineral identification maps for the region around north east Syrtis.

However, these spectra have the same shape in terms of the continuum removed spectra. Additionally, there are jarosite deposits (shown in pink in Fig. 3.17) in the CRISM image HRL0000B8C2. The jarosite spectra from the image HRL0000B8C2 shown in Fig. 3.19 (E) (This detection agrees with previous detections mentioned in [28]). There is also evidence of serpentine in a few images, and the spectra associated with the mineral shown in Fig. 3.19 (A). There is also some Low-Ca pyroxene (shown in dull yellow in Fig. 3.17) in some images in the North East Syrtis regions. The

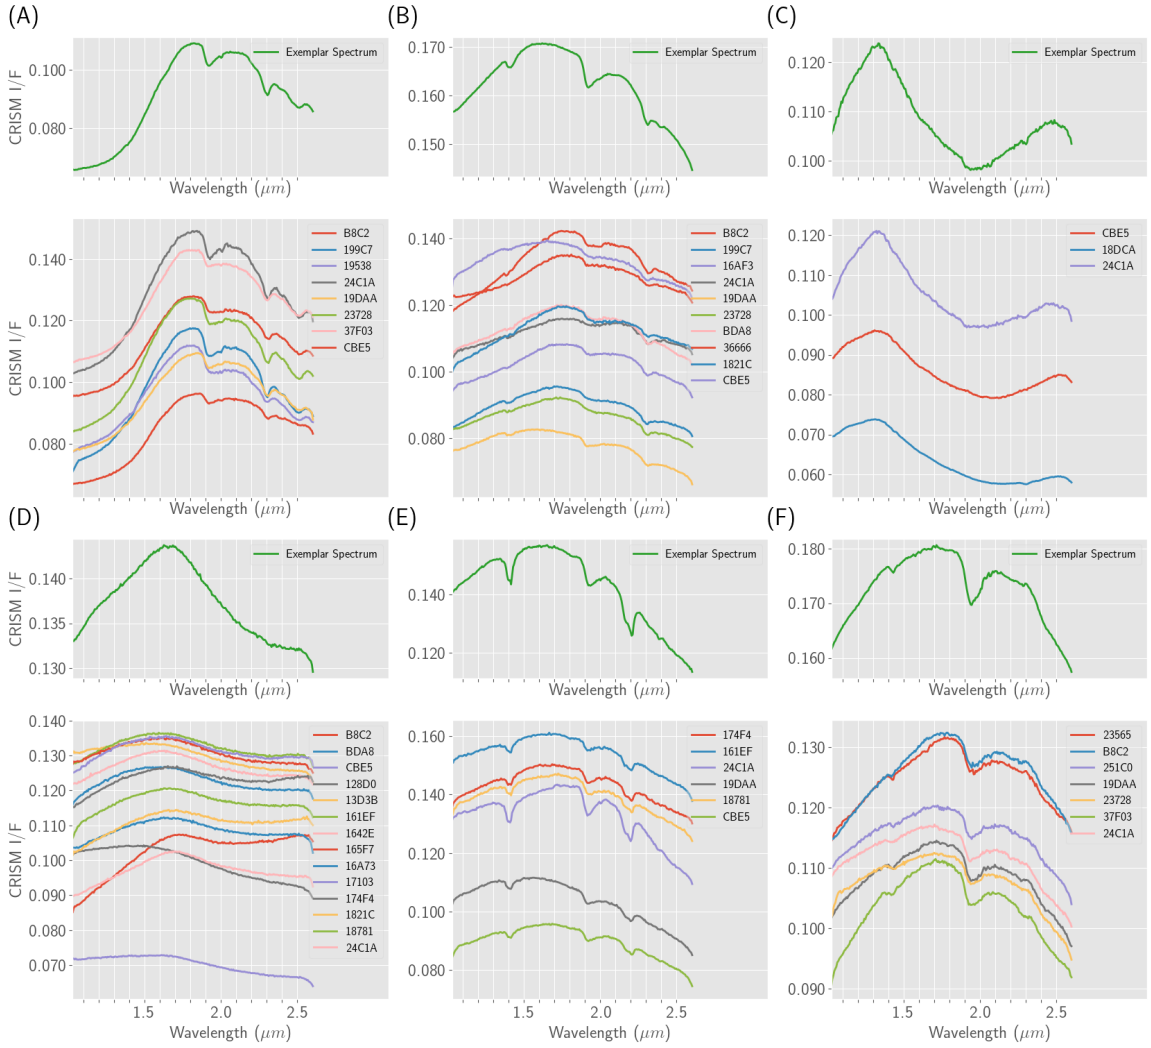

Figure 3.18. Spectra of minerals seen across many images from/near north east Syrtis. (A) Mg-Carbonate. (B) Fe/Mg-Smectite. (C) low-Ca Pyroxene. (D) high-Ca Pyroxene. (E) Kaolinite/halloysite. (F) poly hydrated sulfate.

detected low-Ca pyroxene spectra are shown in Fig. 3.18 (C). Various other minerals such as epidote, Al-smectite, kaolinite, etc. are found in much smaller deposits, but the spectra show clear similarities to the specific "exemplar" spectra.

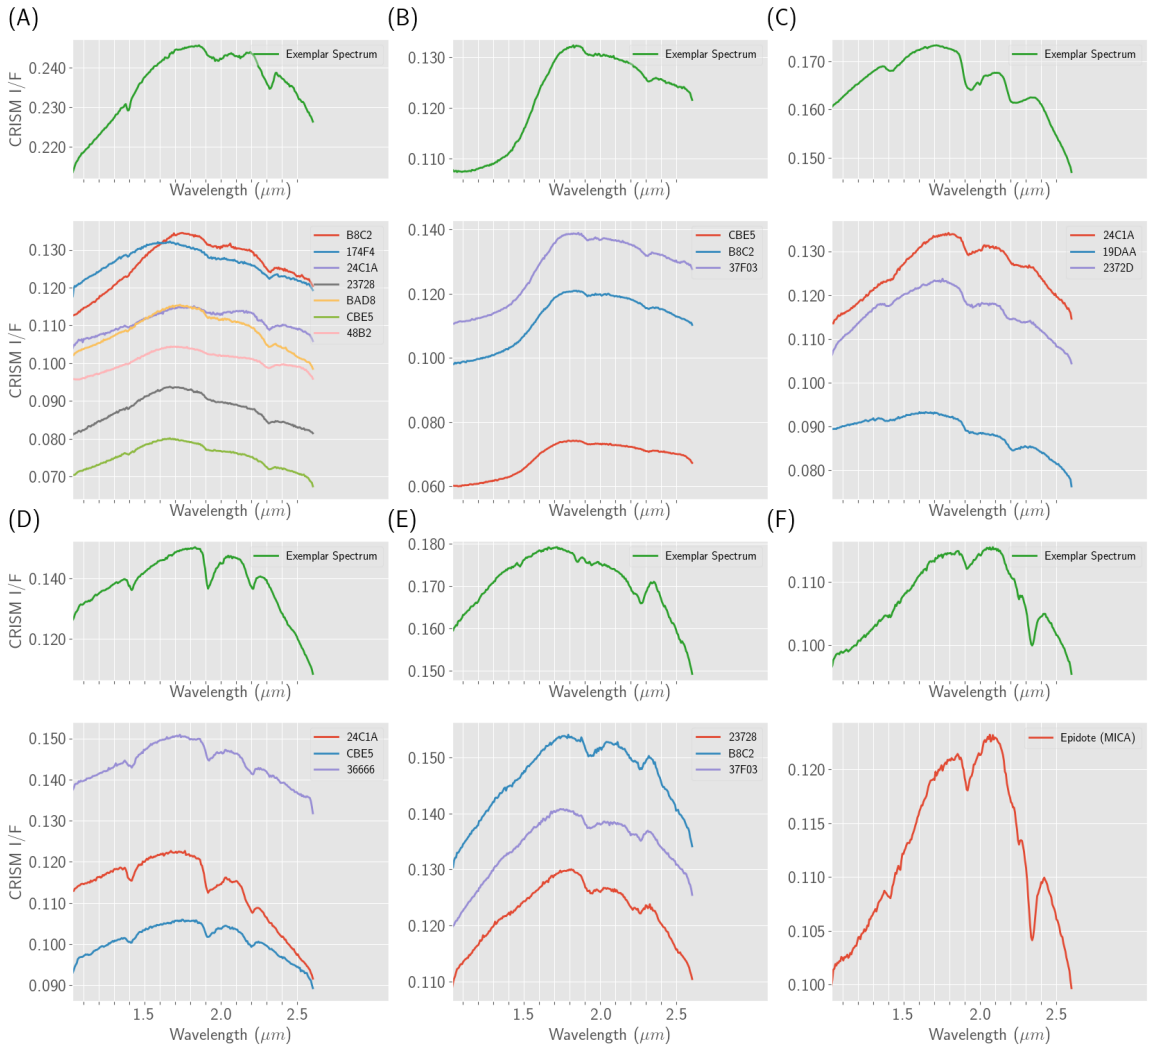

Figure 3.19. Spectra of minerals seen across a few images from/near northeast Syrtis. (a) Serpentine. (b) Olivine + Mg-smect mixture. (C) Hydrated silica. (D) Al-Smectite. (E) Jarosite. (F) Epidote.

3.5 Known issues and limitations of the mapping method

The main limitation of the mapping technique is that it is a ***spectrum mapping*** technique instead of a **mineral mapping** technique. This is a subtle difference, as mineral identification is generally made based on the presence of specific absorption features, there may/may not be some other absorption features that indicate some alteration of the mineral in question that is not sufficient to change the mineral iden-

tification, e.g., some spectra tend to have an additional feature at 1.9 μm based on whether the mineral in that region was hydrated. While this feature appears quite prominently in the mineral, the presence/absence of this feature is not generally used by the experts to make a specific mineral identification. Unlike the expert, the mapping method has no means of identifying specific absorption features as unimportant, so does not immediately identify alterations of known mineral spectra. An example of this is shown for the mineral Mg-Carbonates, (shown in Fig. 3.20 (A)), where the identification is made based on the presence of the bands at 2.31 and 2.51 μm (bands at the black lines in Fig. 3.20 (A)), but the spectra often show some variation in the other spectral regions in Fig. 3.20 (A); for example, see the bands at 1.34 and 1.9 μm. These differences are further enhanced in the case of the continuum removed spectra (shown on the bottom row), and as such, the model is unable to identify the green spectra as "similar" to the blue or "exemplar spectra" due to the differences in the additional bands. In this specific case these slope differences are neglected as slope artifacts in the range $1.0 - 1.4$ μm are known artifacts of the atmospheric correction.

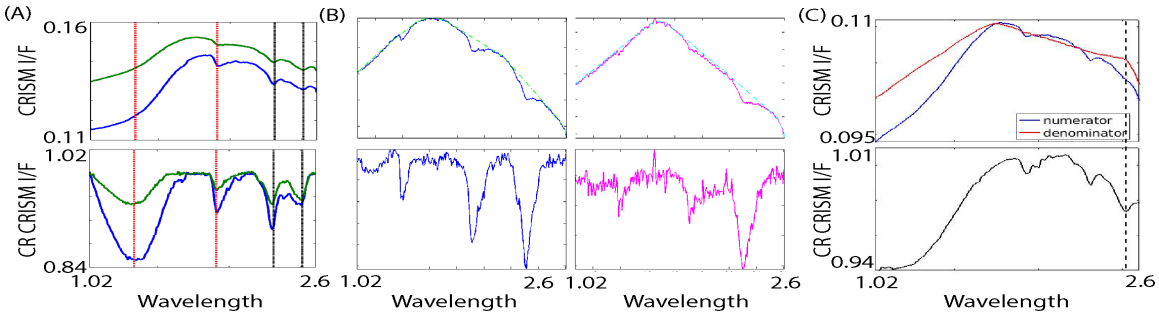

Figure 3.20. Known issues with the GAN based automatic identification. (A) Inter Class variability. (B) Continuum Issues. (C) Ratio problems.

Another associated problem with the method is since all decisions are made based on the continuum removed spectra; any errors in the continuum removal process is detrimental to the performance. This is mainly a problem for the CRISM spectra having some concavity at low and high wavelength ranges. These concavities appear

as small bands in the continuum removed spectra, making mineral identifications more difficult, particularly when the absorption features of interest are also small. An example of this is shown in Fig. 3.20 (B), where we have two spectra that appear to have the same absorptions in the CRISM I/F space, but in the continuum removed space we see a band in the low wavelength range of the blue spectrum. Also, note that the "fake" band is similar in size to the smallest absorption bands in the spectra, and this makes mineral discrimination hard in these cases. This problem can be thought of as artificially increasing the spectral variability of a class. This issue is further exacerbated by the fact that an artificial edge is created at 1.0 μm because of the spectral properties of the sensor and the exact location at which the continuum is fitted to the spectrum.

Note that both these above problems cause the model to "miss" out on identifying some possible mineral spectra (i.e., these are *false negatives*) as opposed to misidentifying a pixel as belonging to a different class (or *false positives*). Additionally, one can work around these issues by augmenting the exemplar library to contain examples relevant to each class's spectral variants. This would ensure that the model can identify all the different mineral phases present in the data. {N.B., the exemplar library used to create the maps shown in the previous section has been enhanced by including variants for the different MICA classes to improve mineral identification}.

Another limitation of the technique is that it can only make detections if the absorption features are present in the CRISM I/F data. A common procedure used in manual mineral identification is ***ratioing***, wherein the pixel spectrum of interest is ratioed with a *neutral* spectrum (i.e., a spectrum with no absorption features) to remove instrument based distortions and increase the spectral contrast of the data. In some cases, the ratioing may induce/eliminate certain spectral features in the ***ratioed*** spectra, leading to mineral identifications. An example of this is shown in Fig. 3.20 (C), where both the numerator (blue) and denominator (red) spectra are shown in

the top row while the ratio spectrum (black) is shown in the bottom row. Note that the band at 2.5 μm is present in the ratioed spectrum and appears to be in line with the convexity in the denominator spectraum at that wavelength. The original spectrum (and its convex continuum removed spectrum), does not contain such a large and clear band, and the model is therefore, unable to find such spectra in an automated manner. It should be mentioned that this is not to argue that the ratioing is incorrect (the ratio is similar to subtracting a spectrum in the log-space, so this operation can be considered to be analogous to unmixing), but rather to highlight that the GAN-based mapping is unable to identify minerals that are only highlighted after the removal of some fractional components.

CHAPTER 4

AUTOMATED IDENTIFICATION OF NOVEL MINERAL SHAPES IN CRISM IMAGES

The previous chapter describes a promising method for mapping known mineral shapes in CRISM images which leverages features learned by an unsupervised machine learning algorithm–namely, the GANs. One capability the existing suite of tools lack, is that they are currently unable to identify variants/alterations of known mineral shapes or previously unknown/novel spectral shapes present in CRISM images. A traditional classification technique, such as the GAN feature based classifier described in Section. 3.3.3 also does not include this capability (the full pipeline is schematically shown in Fig. 4.1). This technique (classifier) simply assigns each test sample to the same class as the exemplar to which it is closest to in the GAN feature space, i.e. it would simply classify each test pixel as a member of one of the known classes. In terms of machine learning, the GAN feature based classifier described above is known as a *closed set classifier*, wherein each test sample is classified as a member of one of the known classes. Such an approach, does not include the ability to check whether a test spectrum is an example of a novel class (i.e., the spectral shape is unknown *apriori*).

Similarly even classical neural network classifiers, are also designed under the closed set paradigm, wherein each sample is classified as a member of one of the k known classes. These models are supervised machine learning models, i.e. the models use some examples of the k-known classes to learn the relationship between the input sample and class memberships. Finally, these classifiers leverage a *SoftMax* layer [32] as the final layer to produce a classification result. The *SoftMax* layer converts a

$k-$dimensional output vector of the neural network into a probability distribution over the $k-$dimensions, i.e., classifier output is a distribution across the k classes. The test sample is assigned the label corresponding to the most probable class in the output distribution.

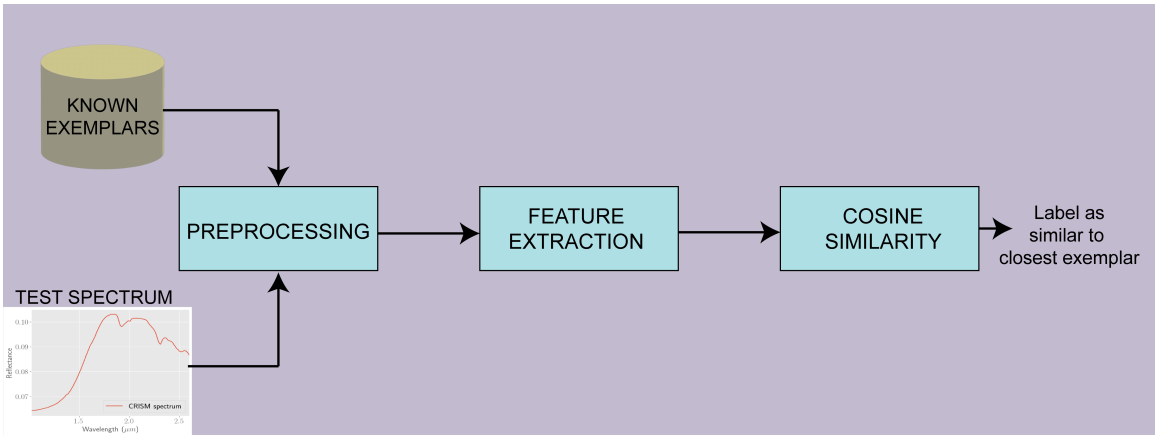

Figure 4.1. Flow chart describing the CRISM pixel classification using the GAN features

The mapping tool described in the previous chapter (in Sec. 3.4) is only able to assign high confidence labels to very few pixel-spectra (i.e., pixel spectra with a cosine similarity $\geq$ 0.95 to a specific exemplar in the GAN feature space). Most images contain a significant number of pixel spectra which are considered "non-bland" or "interesting" (as per the criterion in Sec. 3.1.3) which cannot be identified as either "identifications" or "guesses" corresponding to any of the known exemplars. Even for the of pixels marked as "guesses" in the previous chapter too, it is essential to render a decision on whether the spectral shape is the same as the exemplar or a variant that needs to be mapped separately. It should be noted that strict thresholds were set for the "identifications" in the previous chapter to ensure uniformity between the spectral shapes identified as similar to a specific exemplar. Given these conditions, the ideal tool for application to CRISM data can instead be modeled as *open-set* classifier.

Under the *open-set* scheme, the classifier considers two possible scenarios for each test sample, that it is either a (i) Known class sample, or an (ii) Out-of distribution sample. In the first scenario the model decides that the specific test sample is a member of a known class (i.e. is similar to some of the samples encountered during training). In this case the classifier also renders a decision on which of the known classes the specific test sample belongs to. In the second scenario, the model decides that the test-sample is unlike any of the spectral shapes encountered during training. Instead, the model marks/labels the test-spectrum as an "outlier" to the spectral shapes the models has been exposed to previously or and novel spectral shape. While, the machine learning literature contains a variety of techniques and approaches for open-set classification [96], there two reasonably well known architectures which allow for open set classification using classical deep learning networks, namely the *OpenMax* [8] and the *Deep Open Classifier (DOC)* [81]. Both the open set classification architectures include modified output layers (to replace the closed set *SoftMax*) designed for an open set scenario.

The *OpenMax* [8] modifies the scores/output for a traditional closed set classifier based on Extreme Value Theory. The eponymous *OpenMax* layer operates on the outputs of penultimate layer of the traditional network (referred to as an *Activation Vector (AV)*). In this formulation, each class is considered to be represented by the Mean Activation Vector (MAV), which the mean computed over the AV of correctly classified examples of that class. For each class the technique further calculates the distance between the MAV and the AV of the correctly classified examples. Following this for each class, the technique fits a Weibull distribution to the largest distances between the correctly classified AVs and the class MAV (The choice of Weibull distributions for the distances is motivated by previous analysis of classifier scores which found that such scores follow a Weibull Distribution [77]). For a test-sample the activation vector is revised based on the Weibull scores from the estimated distribution

functions. The technique also calculates a pseudo-activation for the "outlier" class a the sum of difference between the original and revised activations. The full activation vector for this layer is now the combination of the "outlier" activation and the revised class activation vector. This revised activation vector further normalized using a *SoftMax*. A test samples is marked as "outlier" if either, the (pseudo-)activation of the "outlier" class is highest or if the activation of the largest class is below a user-defined threshold.

The *DOC* attempts to generate a similar output, but instead replaces the final layer of the classifier by a K node layer with simple sigmoid activations. Shu et al. (2017) show that this layer behaves as a set of K 1-*vs-rest* classifier wherein each node decides whether the test sample is member of a specific class. Following this, the DOC identifies thresholds for each class based on the differences between the scores of in class members and scores of non-class members. Finally, any test-sample whose output in the DOC layer is below the identified thresholds for all the K classes is considered "outliers". This chapter describes a pipeline that leverages the DOC layer for setting up of open-set classifier that can be used for processing CRISM images. This pipeline is able to simultaneously identify both known and novel spectral shapes present in the image and map their spatial extents . In this implementation, both approaches were tested on a model open set problem for CRISM data (more details on the exact experiment is presented in Sec. 4.1.1) and for this specific test both approached showed very similar performance. Given this information the DOC architecture was preferred as it simpler to train and implement and does not require an additional Weibull fitting. It should be noted that the main motivation at this stage is to ensure that open set tool capable of identifying novel spectral shapes from the CRISM image database rather than optimizing for the best tool of this nature. Future works can delve into the ideal open-set tool for CRISM data.

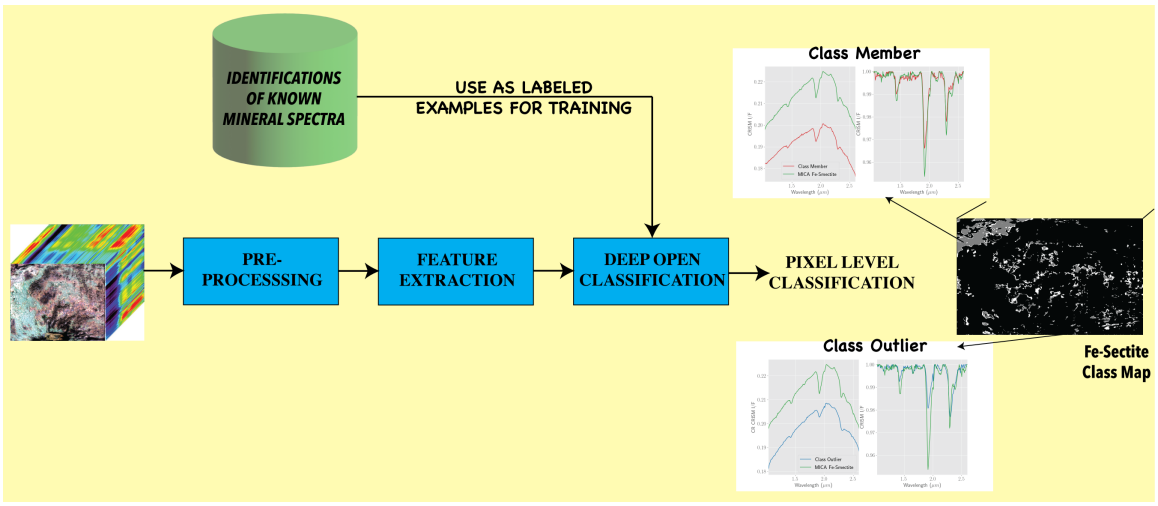

Figure 4.2. Flow chart describing the novel spectral shape identification using a DOC for CRISM images

Table 4.1. Architectural details of the CRISM-based Deep Open Classifier designed for this project

Deep Open Classifier (DOC) Architecture				
Layer	Filters/ Nodes	Stride	Activation	o/p size
Input				240 X 1
Conv-1	20	2	Leaky ReLU	120 X 1
Conv-2	40	2	Leaky ReLU	60 X 1
Conv-3	80	2	Leaky ReLU	30 X 1
Conv-4	160	2	Leaky ReLU	15 X 1
Flatten				2400 X 1
Dense-1	100	n/a	Sigmoid	100
Dense-2	50	n/a	Sigmoid	50
Output (Dense)	22	n/a	Sigmoid	22

4.1 Deep Open set Classifier (DOC) for CRISM images

This section describes the process of setting up and training an open-set classifier for the CRISM spectral data. The full DOC-based scheme for CRISM images is shown in Fig. 4.2. The first step leverages the preprocessing methods described previously (see Chapter 3.1). The DOC model is also designed to leverage the feature extractor (i.e. the discriminator) described in the previous chapter. This choice is

made as the previous chapter has shown the GAN-based feature extractor to be highly discriminative in terms of spectral shape. The overall architecture used for the DOC on CRISM spectra is described in Table 4.1. The initial layers of the model (shown in green in Table 4.1) are layers corresponding to the GAN-based feature extractor and were initialized under the "warm-start" [4] scheme, wherein the weights of a pre-trained model (in this case the initial layers of the discriminator from the previous chapter) are used to initialize a new network. These initial layers corresponding to the GAN-based feature extractor are held frozen during the training phase.

The DOC models have dropout layers with a dropout of 0.3 after each computational layer. The models use an Adam optimizer[47] with a learning rate ($\alpha = 1e^{-4}$), and the multipliers for the decay rate to the various moments set to ($\beta_1 = 0.9$, & $\beta_2 = 0.99$). All DOC models are trained with a batch size of 256 for 200 epochs, with early stopping using a validation dataset. The model is trained using the summation of the log-losses of the $k-$sigmoid functions (assuming k is the number of classes) on the training data (with N samples) as shown in Eqn. 4.1 [81]:

$$Loss = \sum_{j=1}^{N}\sum_{i=1}^{k} -\mathbb{I}(y_j == l_i)log(p(y_j) == l_i) - \mathbb{I}(y_j \neq l_i)log(1 - p(y_j) == l_i) \quad (4.1)$$

Following the training the algorithm includes a thresholding condition for each node's output for the rejection of samples as "outliers". In the original implementation, a specific threshold is evaluated for each class/node is evaluated, with any test-samples with maximal output activation below the threshold are considered outliers (assigned a label of -1). For each node/class the algorithm considers each member of the class as positive example and every member of other classes as a negative example and identifies a threshold which can best separate the two. With this modification the neural network model is expected to be able to not only identify members of the known classes but is also able to identify any test samples that the model was not exposed to in training.

4.1.1 Verifying the performance of the DOC on CRISM spectral data

Prior to the application to CRISM image data the chosen open-set classifier has to be tested under the open-set paradigm for spectral data. For this purpose a toy open set problem was created with labeled spectral data. The training dataset of this problem contains 4500 labeled examples per class from 6 mineral classes known to be present in CRISM data, these mineral classes are Fe-olivine, Mg-carbonate, low Ca-pyroxene, mono hydrated Sulfate, Alunite, Al-smectite, and Fe/Mg-smectite. To simulate the open set paradigm the test set also included samples from 4 additional classes found in Martian data, namely, Jarosite, Gypsum, hydrated Silica, and Prehnite. This open-test set contains 1500 test-samples corresponding to each of the test classes mentioned. Following this, a DOC classifier with the architecture shown in Table. 4.1 is trained with the spectra in the training set described here. Post the training with the labeled dataset the outputs of each node for the different classes are analyzed. In this implementation, the threshold to identify "outliers" is estimated as 2 times the standard deviation of the activation for the training samples of the specific class. This threshold was chosen as across all the classes in the training there were no negative samples included in the specific class with this threshold. Following this the DOC model is tested on the open-test set described above. For the samples corresponding to the novel classes the models' classification is considered successful if they are labeled as "outliers" and not members of a known class. The confidence matrix showing the performance of DOC on the open-test set is shown in Fig. 4.3.

From the results shown in Fig. 4.3, it is clear that for the simple open-test set described here, the model is able to not only classify the samples of the known classes, the DOC is also able to identify the outliers successfully. In analyzing the results for the open-test set described above it is essential to note that the open-dataset described here is significantly simpler than the CRISM image dataset, as the classes are disparate by construction, i.e. the samples in each class are drawn from a clearly

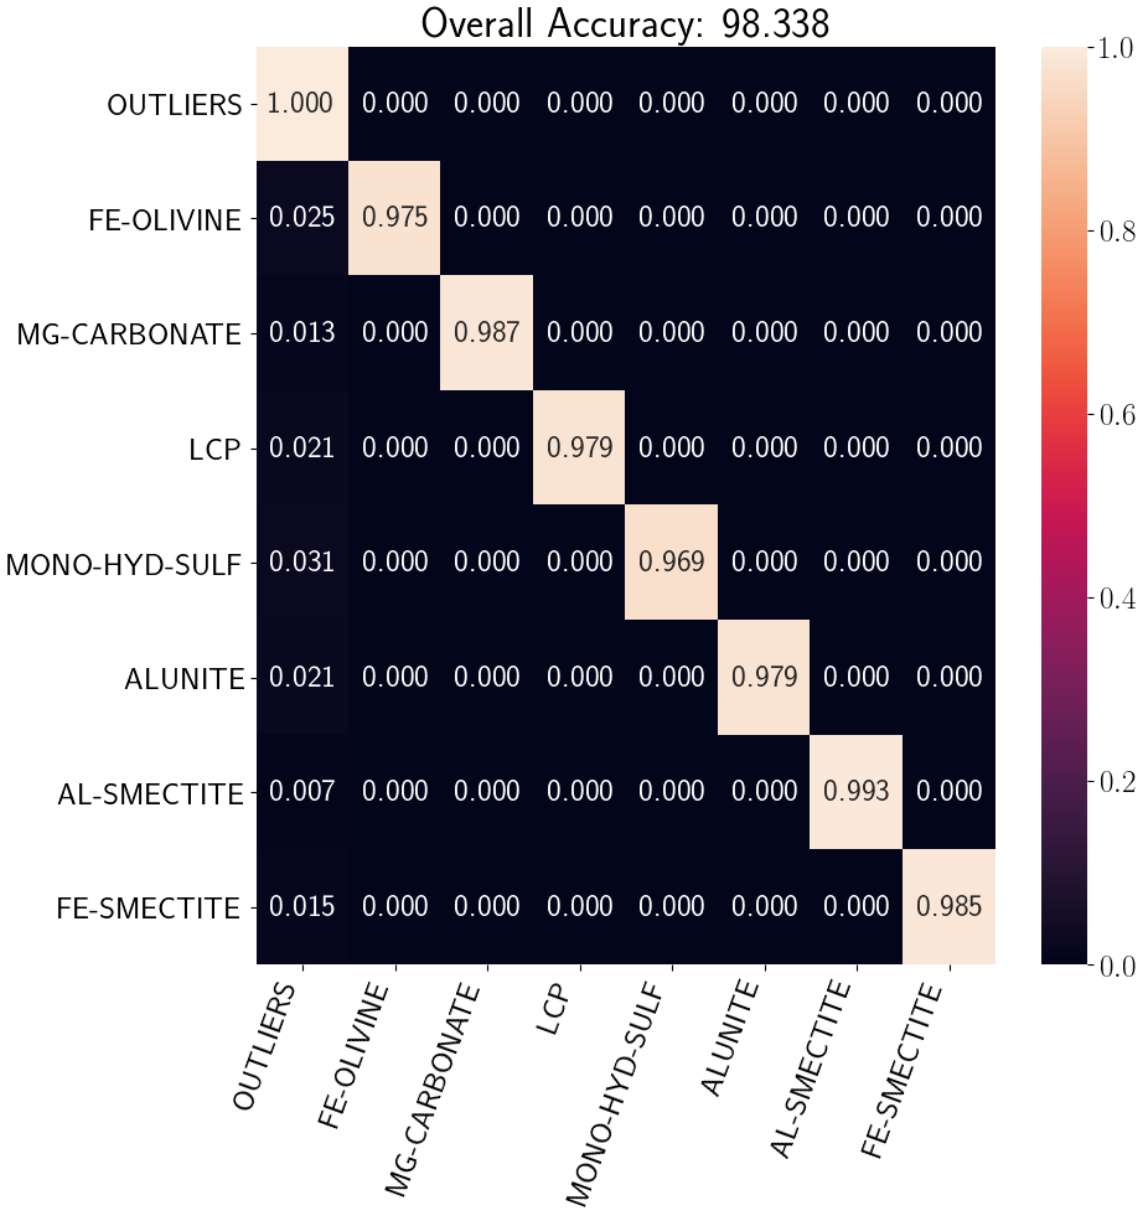

Figure 4.3. Confidence matrix showing the performance of the DOC for a simple experiment with labeled samples

separable regions in the GAN representation space. Rather, this experiment only provides a baseline to indicate that the DOC able to clearly separate the unknown samples from the known samples, as indicated by the 100% accuracy in identifying the "outliers" (see top row of Fig. 4.3). Interestingly there are some cases when class-members are marked as "outliers", this is because of the strict thresholds placed on the definition of the class, and a continuation of our preference to miss a detection, rather than mislabeling a spectrum. Further, based on this trend that model is more likely to miss a specific detection rather mislabel one, it can be surmised that samples labeled as a class by the model are spectrally consistent. This mode of labeling can be considered identifying region is the GAN-based feature space that is consistent with the spectral class.

4.1.2 Creating a training dataset for CRISM DOC classifier

In the initial implementation the DOC for CRISM images were initially trained on the 18 major well-known mineral classes (shown in black in Table. 3.4) from the CRISM MICA database. While primarily 18 classes were considered, some of the classes encompassed multiple spectral shapes, for example the expert analysis has shown two spectral shapes that are both considered Mg-Carbonate (see Fig. 3.20 (A) for specific shapes). Following this we search a training database of around 40 CRISM images for spectra that are very similar to exemplars corresponding to the various classes. Similar to the efforts in the previous section, spectra corresponding to the known classes are found by leveraging a similarity analysis in the GAN-based feature space. In this case, the spectra that can be considered *identifications* are extracted and placed in a labeled spectral database (to ensure strict similarity with the exemplar, the identification threshold (in terms of cosine similarity) was further tightened to 0.97). This search enabled the creation of a labeled CRISM spectral database that can be used for training ML models.

The one issue identified with the labeled dataset is that it is significantly imbalanced, i.e. there are more examples of certain prevalent minerals such as Fe-Olivine and Fe/Mg-smectite that have identifications ranging in the tens of thousands, while rare minerals like Jarosite have examples ranging in only the tens or hundreds. To address this issue a simple data augmentation techniques is used, convex combinations of existing members of the class are created to create more examples of the sparse class. Following this a simple training dataset is created with 10000 samples of each class to train the DOC model. Following this is a DOC model is trained with this dataset using the same architecture and hyperparameters as the model described in Sec. 4.1.1. Further, a similar scheme is used to identify the outlier threshold for each of the known classes.

4.1.3 Training the DOC classifier for CRISM data

Finally to use the DOC classifier to generate maps for various minerals/outliers present in a CRISM image, each image is first processed using the preprocessing steps mentioned in Chapter 3.1. Then each pixel with significant absorptions is applied to the trained DOC model which produces a 18−dimensional vector as on output for each pixel. All activations except the highest score in the output are zeroed out. Then for the maximal class activation if the score is above the threshold for the specific mineral the pixel is marked as "class-member" (denoted by label 1). If the score is below the class threshold the pixel is marked as an "outlier" (denoted by label 2). Thus, each non-bland "pixel-sepctrum" is represented by an 18−dimensional vector with one non-zero value. For the pixels marked "outlier" (i,e., pixels for which the non-zero value in the label vector is 2), the label vector also identifies the library exemplar to which the outlier spectrum is closest to in shape.

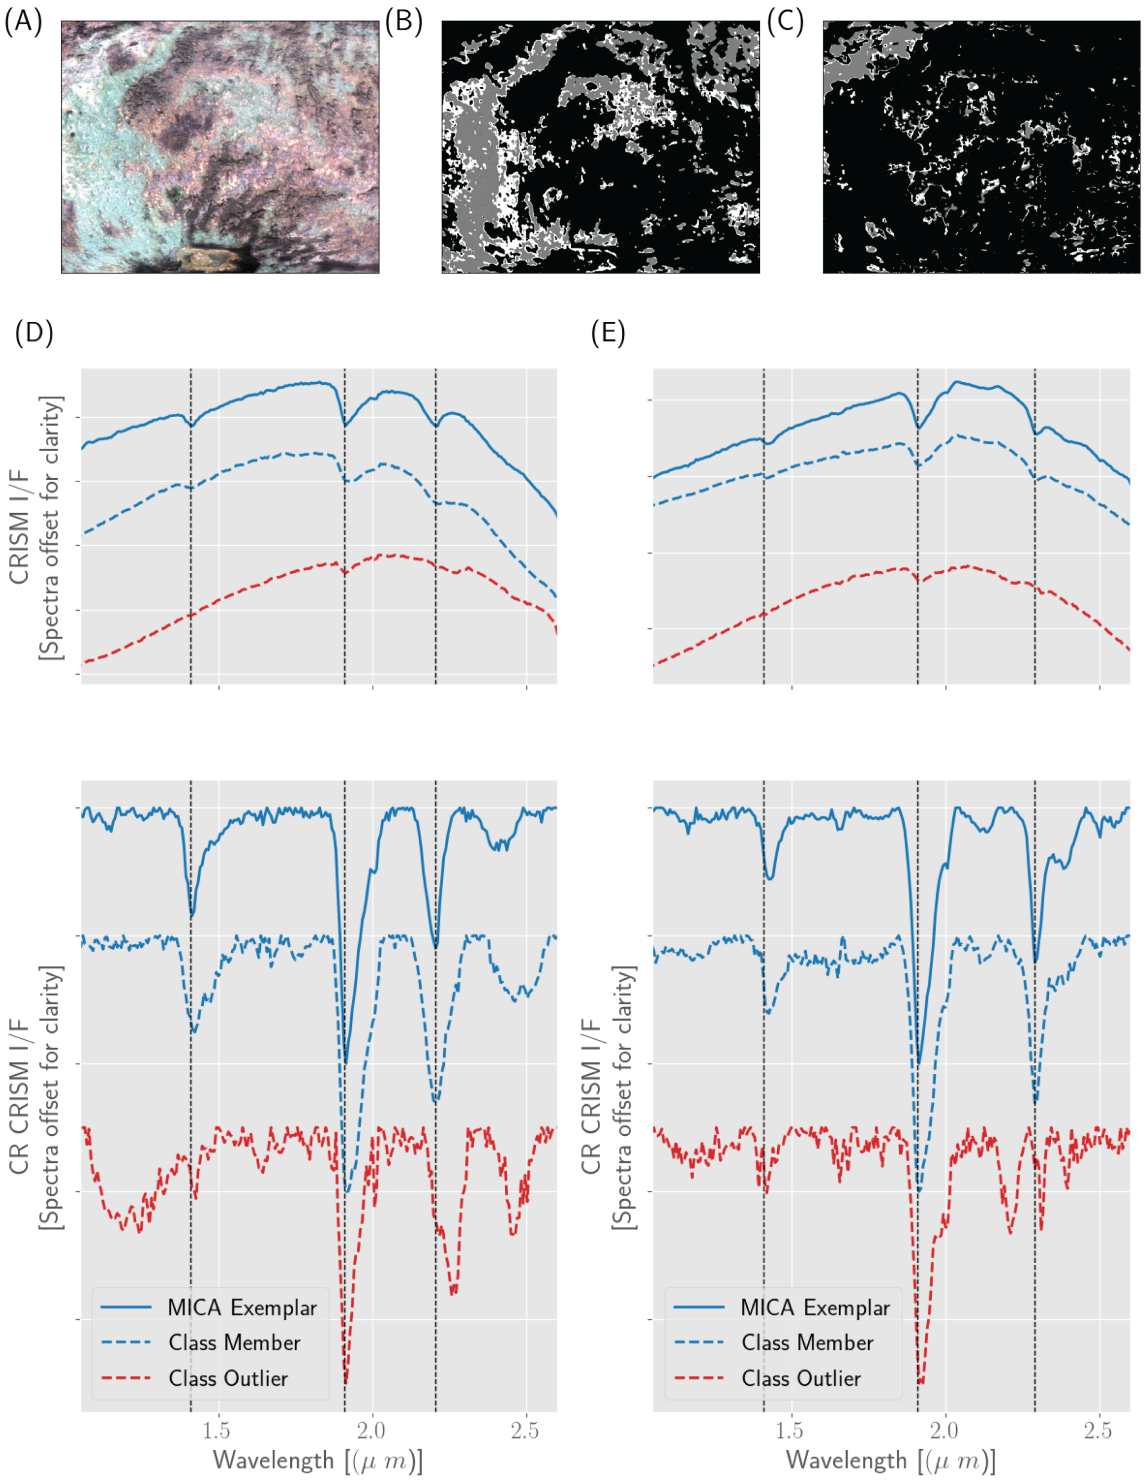

Figure 4.4. Maps showing the DOC performance of the CRISM images $FRT0000A425$ with: (A) RGB composite of the CRISM images. (B) DOC maps for Al-smectite. (C) DOC maps for Fe-smectite (D) DOC class members and outliers for Al-smectite and (E) DOC class members and outliers for Fe-smectite

4.1.4 Illustration of the DOC Mapping on a single CRISM image

To illustrate the performance of the DOC classifier we consider the CRISM image $FRT0000A425$ is an image from Mawrth Vallis (N:25.248355, W:20.254675). The image is well-known and is in fact the scene from which the MICA exemplar for Fe-Smectite was identified. The image has been well-studied and contains a wide variety of spectral shapes and provides an excellent test case for the open-set classification. The main mineral detections of the DOC classifier in this image are smectites namely Al-smectite and Fe/Mg-Smectite respectively. The maps of the two minerals from the DOC classifier are shown in Fig. 4.4 (B) & (C). The pixels marked in grey are class members of the class Fe-smectite and Al-smectite. The pixels marked in white are the outlier pixels in the image who's closest library match is the either Al-Smectite or Fe-Smectite respectively.The plots in Fig. 4.4 (D) & (E) and illustrate the comparison of the class members and outliers relative to the class exemplar.

To illustrate the differences between the "class-members" and "outliers" Fig. 4.4 shows the comparison of the class exemplar with some manually chosen examples of either class. Note that the class members (the dashed-blue plots) shown in Fig. 4.4 (D) & (E) have all the same absorptions as the class exemplars (the solid-blue plots) in the continuum removed space. The major bands in general have the same approximate shape, width and depth also. On the other hand, the outlier spectra (shown in dashed-red plots) are significantly different in terms of the absorptions present. The Al-smectite outlier (shown in red in Fig. 4.4 (D)) has an absorption at 2.27 μm instead of at 2.205 μm like the exemplar of the Al-smectite class. Also there are differences in the shape of the hydration band at 1.4 μm. Similarly, for the outlier spectrum of the Fe-Smectite class (shown in red in Fig. 4.4 (E)), the specific pixel spectrum has bands at 2.205 & 2.31 μm instead of the characteristic 2.29 μm absorption typical of the Fe-Smectite. Clearly, the outlier show significant variations from the class exemplars and their identification as outliers appears reasonable.

Based on this DOC classification, the model classifies a total of $82,475$ of the significant pixels as members of the known classes, while about 73415 pixels were labeled as outliers and needed to be analyzed further.

4.2 Parsing the list of outlier spectra in CRISM images

Post-processing of a CRISM image by the DOC classifier as described in the previous section, a significant portion of the non-bland pixels are marked as "outliers". The outlier count runs in the thousands or tens of thousand for most CRISM images. Parsing such a large list of spectra manually and identifying the unique spectral shapes of interest and their spatial extents will be highly tedious and intractable on a large scale. Further, given the expectation that neighboring pixels in CRISM images have very similar spectral shapes it would be useful if an automated processes is able to group the identified outlier spectra into clusters/segments such that all the spectra in a given cluster/segment have a uniform spectral shape. In pursuit of this objective, a graph-based segmentation approach [75] was leveraged to group the outlier spectra from a given image into spectrally uniform clusters.

4.2.1 Distance based spectral clustering

The graph based clustering algorithm begins by describing a graph $\mathcal{G}$ over the elements of interest. In mathematics, a graph is a set of objects that are related in some sense. The objects are modeled as nodes/vertices with the relationships between them being modeled as edges [93]. In this initial graph definition each element in the input set (which can be represented as a vector $v \in \mathbb{R}^D$) is modeled as node. Further, the graph connects each node to its $k-$nearest neighbors with edges, each edge can be weighted according to a mathematical relationship between the nodes it connects. The full set of edges is referred to as $\mathcal{E}$.Thus a graph $\mathcal{G}$ is defined by a set of

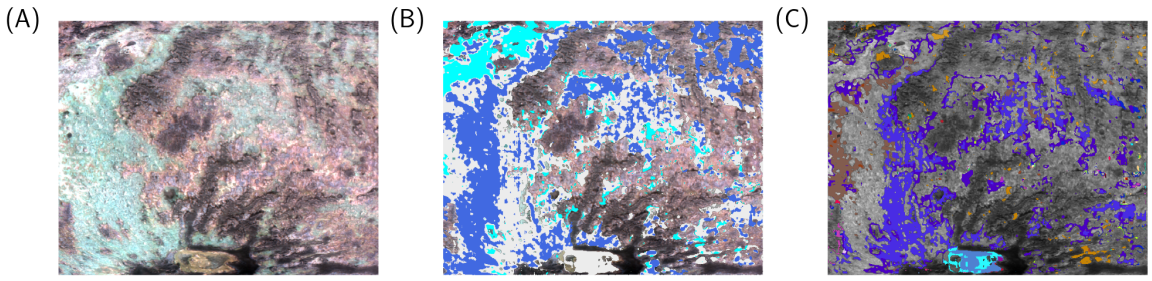

Figure 4.5. Clustering identified Outlier pixels into segments. (A) RGB composite of the CRISM images. (B) DOC maps for Al-smectite. (C) DOC maps with mineral identifications marked as colored pixels (**Cyan**:$Fe-smect$, **Blue**:$Al-smect$) and outlier pixels marked in white. (C) effect of segmenting/clustering the outlier pixels-each colored cluster identifies pixel spectra that the segmentation algorithms considers similar to each other.

nodes/vertices $\mathcal{V}$ and a set of relationships between the nodes, i.e. edges $\mathcal{E}$, together the graph is expressed as $\mathcal{G}(\mathcal{V}, \mathcal{E})$.

For clustering, the algorithm begins by considering each node/vertex as a cluster/segment on its own. Following this the algorithm arranges all the edges in $\mathcal{E}$ in increasing order of the edge weights. Each edge is then processed if it connects nodes in different segments, say S_a and S_b. In this case, the algorithm then estimates the minimal similarity ($MSim(S_a,\ S_b)$) between the two segments as:

$$\begin{aligned} MSim(S_a,\ S_b) = \min_{v_i, v_j} \quad & 1 - d_{cosine}(v_i, v_j) \\ \text{s.t.} \qquad & v_i \in S_a; \quad v_j \in S_b \end{aligned} \tag{4.2}$$

Following this, two segments are merged if the minimal similarity between them is above a user-defined threshold. Following the merging of these nodes into one cluster/segment, the segment id for all the members are updated to be the same. Once the algorithm has completed processing all the edges in the set $\mathcal{E}$, the algorithm would have grouped all elements with other similar nodes in the input dataset.

4.2.2 Grouping outlier into similarity based clusters

The various pixel spectrum in an image identified as "outliers" by the DOC in an image are extracted. Following the spectra are passed through the GAN based feature extractor- the feature extraction is performed as this representation is shown to be more discriminative than the original spectral space. Following this, the graph-based clustering algorithm (described in Sec. 4.2.1) is used to cluster the image outliers into segments/clusters of similar spectral shapes. For CRISM image set, each edge is weighted by the cosine distance between the two spectral shapes in the GAN feature space. In this application the similarity threshold for merging was chosen to be a spectral similarity of 0.95. This threshold was chosen based on heuristic manual analysis that has indicated that spectra with cosine similarity greater than this threshold are spectrally consistent and can be considered members of the class. Post clustering, each cluster is only represented by the average spectrum of each cluster. Such averaging very useful in eliminating/suppressing the noise/distortions present in an individual pixel. For the CRISM image $FRT0000A425$, such processing reduces the number of unique spectral shapes to be analyzed from 73415 down to 135. The shapes of outlier clusters/segments is shown in Fig. 4.5.

4.2.3 Grouping Cluster Medians for manual analysis

While the previous step significantly reduces the data complexity from thousands of spectral shapes to be considered to a few hundred, manually analyzing a few hundred shapes will still be physically taxing. For this purpose, an additional clustering step is proposed wherein the clustering is reapplied to the cluster medians identified in the previous step, but the threshold is loosened down to the value 0.707. This threshold now ensures that each median is grouped with spectra it is most similar to but no longer guarantees spectral uniformity inside the group. This step is useful as it allows a human expert to consider each spectrum relative to spectra that they

are most similar to. This will allow the expert to easily identify spectra that are distinct and should be mapped separately. For the CRISM image $FRT0000A425$, this approach groups the 135 unique spectral shapes into 17 distinct groups. Illustration of what the spectra in the different groups looks like is shown in Fig. 4.6

Consider for example the shapes in Fig. 4.6 (B) all the members of this group have clear absorptions at 1.4, 1.9, 2.2 & 2.3 μm. The spectra show differences in the exact depth of the Al-OH band (2.2 μm) relative to the hydration band (1.9 μm). They also show a weaker absorption at 2.3 μm. The blue spectrum in the group also shows a continuum variation (easily seen in Fig. 4.6 (C)) in the $1.0 - 1.8\mu m$ spectral range that may be interesting to an expert. Similarly all the spectra in 4.6 (E) also show a clear absorption at 2.27 μm indicative of sulfates like jarosite. Some of the spectra (specifically the blue and red ones) also show an absorption at 2.4 μm that may indicate an alteration to an expert. The spectral groups and associated maps (both group maps and segment shapes are shown in Fig. 4.6).

4.3 A brief description of the products from the automated mineral identifications

This section will briefly describe the various products that are generated by the automated mineral identification pipeline and provide some insight on how they are meant to be used. The pipeline produces three main products:

[1] The first product generated by the pipeline will be an ENVI style $3 - D$ image ('.img'), where each band contains the maps for one of the pre-known classes. Pixels identified marked as 1 can be considered as members of the specific class. Some pixels may be marked 0 for all the classes, this can either be because the mapping pipeline considers the spectrum "insignificant" (i.e. it contains no clear absorptions), or that

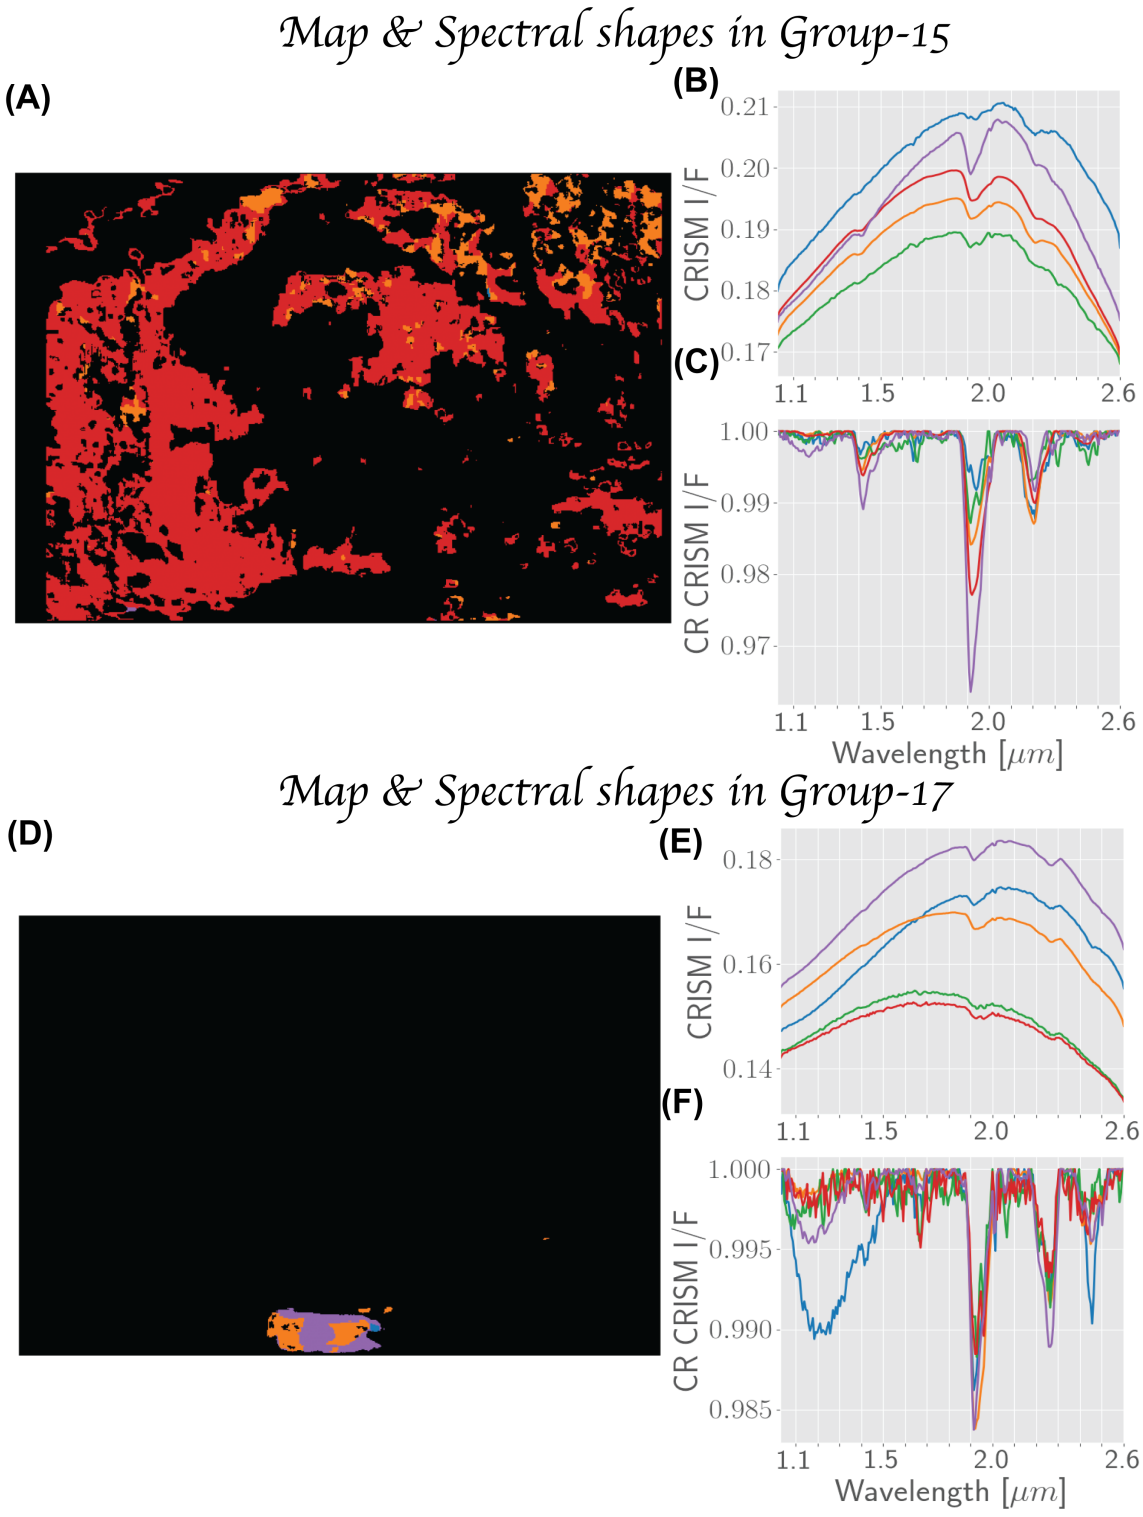

Figure 4.6. Effect of grouping unique outlier spectral shapes. (A) Maps of pixels corresponding to various shapes in Group 15, (B) & (C) Spectral shapes and Continuum removed spectral shapes in Group 15. (D) Maps of pixels corresponding to various shapes in Group 17. (E) & (F) Spectral shapes and Continuum removed spectral shapes in in Group 17.

the spectra are considered "outliers" by the mapping pipeline.

[2] The second product of the outlier detection pipeline is an ENVI style spectral library (a .sli file) containing all the unique spectral shapes identified by the mapping pipeline. As a general rule, this file contains between $30-100$ unique spectral shapes for each CRISM image. Often the differences between the members of this file are rather subtle and may not raise to the level of a spectral difference in the opinion of an expert spectroscopist, but in the interest of displaying all the shapes identified by the pipeline, this product presents all the spectral shapes identified. The names of the product will indicate the group memberships of each spectral shape and should help guide a user to compare spectral shapes that are similar.

[3] The first of the outlier-map products generated by the mapping pipeline will be a group level mapping product. This product is an ENVI style $3-D$ image ('.img'), wherein each band corresponds to one of the groups identified by the pipeline. This is a global mapping product and identifies where spectra of specific shapes are present in the image. If an expert determines that the differences inside a group are not "mineralogical", these bands will provide the maps for the specific class.

[4] The final product produced by the pipeline is a map for each spectral shape identified. As mentioned above in most cases the individual spectral shapes in each cluster are similar and may not be mineralogically relevant. The product is provided to account for cases where the grouping does lump in diverse spectral shapes and allow a fine scale analysis of the spectral differences.

CHAPTER 5

SUMMARY AND CONCLUSIONS

Given its long and fruitful operational life the CRISM instrument(2006-2023) the instrument has acquired close to 10000 images with infra-red spectral information. Given the size and scope of the dataset tools that can automatically analyze and summarize the spectral information present in the images would be a valuable asset to the planetary remote sensing community. Attempts to directly create automatic tools to analyze such images are hampered by the fact that spectral similarity is difficult to define and measure. Experimental effort such as those shown in Sec. 2.1 have shown that classical similarity metrics struggle in measuring spectral similarity, i.e., there does not seem to be a clear separation between the classes. Even metrics like spectral summary parameters which are designed for the express purpose of identifying specific spectral families (see Sec. 2.2 for details).

Given the inability of classical methods such as similarity/dissimilarity metrics and spectral summary parameters to measure spectral similarity described by human experts, automated mineral identification and mapping has proven to be an important challenge for the Martian remote sensing community. Given the issues mentioned above, efforts to leverage simple rule based systems to create mineral identification maps have encountered only partial success. This dissertation attempts to address this gap by creating/defining a suite of machine learning tools for mineral discrimination and identification. The first tool is a *GAN-based spectral feature extractor* that provides significantly improved mineral discrimination (see Chapter 3). The GAN is described as an approximate sampler of the actual CRISM spectral data distribution.

Given the ability of the discriminator portion of the GAN to separate CRISM spectral shapes from non-CRISM spectral shapes, the features learned by this discriminator are expected to be far more discriminative between different spectral shapes than the original representation in the CRISM I/F space. In this specific implementation, the features used are the output of the penultimate layers the discriminator is used as the feature space representation. Experiments have shown that this novel feature space shows a significant improvement in terms of the spectral discrimination (see Sec. 3.3.2 for details). Following this, Sec. 3.4 shows some experimental results of mapping CRISM images using this feature extractor. These results have been validated manually and used to create maps for candidate Perseverance Landing sites as part of the CRISM data "Fandango" project.

While the GAN-based space can be leveraged to create significantly improved mineral maps (or more accurately spectral maps) for mineral spectral-shapes, a missing functionality is the ability to identify hitherto unknown spectral shapes. For this purpose an open-set classifier is defined, which is designed under the specific paradigm that the test data contains classes that are previously unknown. For this purpose we define a deep open classifier that describes each class as a small compact region in the activation space of the penultimate layer of such classifiers. Such models are then capable of identifying spectral shapes that are unlike the constant shapes in a specific region. Once identified such models are able to identify novel spectral shapes present in images. The spectral shapes are then grouped based on similarity and provided to experts for further validation.

While the tools described in this dissertation address some of the issues in the identification and mapping of spectral shapes in the CRISM spectral data, they are by no means a final solution. The noise and other distortions often present in CRISM spectra can change the appearance of these spectra, making them resemble spectra of other minerals. This for the mapping and identification stages as either a spectrum

of the same class is not identified or spectrum of the same mineral is class is again identified as an “outlier. Additionally, even physical factors like mixing may affect spectral shapes which may lead to an over segmentation of the pixel space. Future work should focus on addressing these issues. In addition to this further attempts can also focus on combining the information from the VNIR and SWIR spectral data available from the CRISM instrument. Finally, while these tools show massive promise, a significant amount of work needs to be conducted in validating maps and identifications of these images,

APPENDIX

A BRIEF INTRODUCTION TO HYPERSPECTRAL IMAGING

Remote Sensing is the measurement of the properties of a surface from *a distance*, rather than *in situ* (at the original location), i.e. without any direct contact. One such remote sensing technique is *hyperspectral imaging* which studies the surface of interest using optical signals. Hyperspectral imagers capture/measure the electromagnetic energy (light) scattered in their field of view in very narrow spectral windows (of width around 10 nm) in the Visible to Near Infra-Red (NIR) wavelength range ($0.4 - 4\mu m$) [19], over a $2-$dimensional spatial grid. Hyper-Spectral Images (HSIs) take the form a $3D$ data-cube, with each plane corresponding to the measure of electromagnetic energy scattered off the surface over this narrow spectral intervals. These intervals are generally referred to as a spectral channels or band. Consequently each spatial location in the field of view is associated with a *radiance* spectrum, which represents the amount of radiation that is reflected from the surface at that location. Further, the radiance at each pixel/location is divided by the amount of total energy (in the specific spectral channel) incident at the location to create a *hyperspectral reflectance cube* [16]. The measured reflectance is a property unique to the materials present on the surface being observed, and as such can be used identify the materials present on the observable surfaces [54]. This ability to identify the materials present in the observed scenes from a distance/without any direct contact makes HSI a valuable resources for various tasks such as for agricultural management and food safety[68], monitoring aquatic ecosystems [64], forensic analysis [57], and various geological applications [7].

Geological applications leverage the ability of hyperspectral images to identify and map the spatial patterns of mineral assemblages present on planetary surfaces [16]. Since such studies generally focus on mapping specific mineral shapes on a geographic scale they focus on leveraging hyperspectral data from airborne and satellite sensors such as the Airborne Visual/Infrared Imaging Spectrometer (AVIRIS) [37] for earth observation, the Moon Mineralogy Mapper (M^3) [70] to investigate the lunar surface, and the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) [60] for studying the martian surface. Since these sensors are air or satellite-borne: each scene captured by these sensors correspond to a relatively large area on the ground (for example, each CRISM pixel corresponds to approximately 18 m on the ground, an AVIRIS pixel (in EO-2 mode) to 20 m on the ground). As a consequence, it is quite probable that this area comprises different materials. In this scenario, the spectrum associated with each pixel is expected to be some combination (mixture) of the spectra of the different materials present in the area. In remote sensing image analysis the task of identifying the constituent materials present in a scene and their fractional abundances is referred to as the Mixing/Unmixing problem [46]. Fig. A.1 shows a schematic representation of airborne/satellite borne hyperspectral imaging.

Based on the prevalent conditions different models maybe required to explain the mixing present in hyperspectral images, the mixing models proposed in the literature fall broadly into two categories: the Linear Mixing Model (LMM) and nonlinear mixing models. The LMM states that the spectrum associated with a pixel is a weighted combination of the spectra of the materials (endmembers) present in the area covered by the pixel. The weight associated with each material is proportional to the abundance of the particular material in the region. Under the geometric optics approximation, the LMM assumes that each ray of light interacts with a single material and that the optical mixing occurs at the sensor [54, 13]. In general the LMM is a feasible approximation in the case of *coarse*, *checker-board* mixtures. However

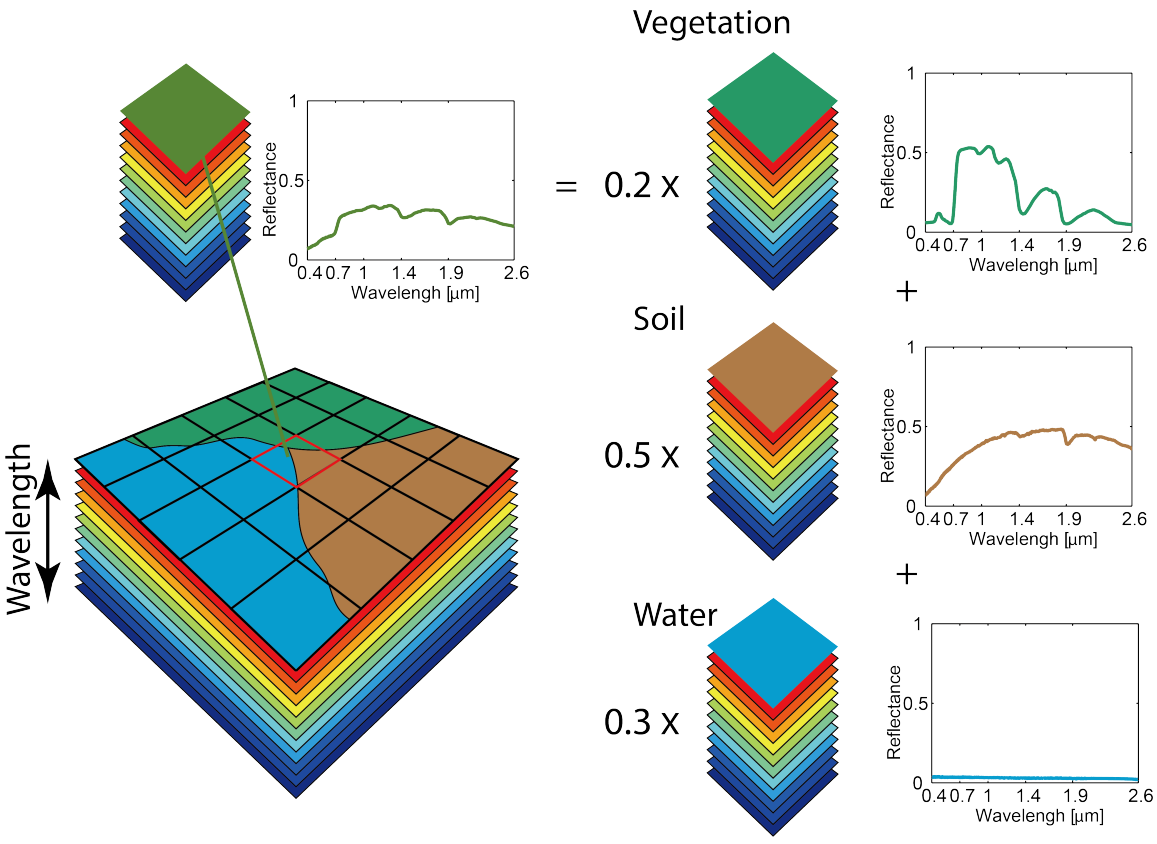

Figure A.1. The scheme for Air/Satellite-borne hyperspectral imaging. [Source: Dr. Yuki Itoh-*Private Communications*]

in the presence of *non-trivial spatial structures* or *intimate* or *microscopic* mixtures the assumptions of the LMM, i.e. light rays do not undergo multiple reflection en route to the sensor, are not valid [46, 42]. In such cases the mixing can no longer be considered linear and is modeled using non linear representations [13, 42, 26].

Most existing unmixing algorithms that have been developed assume an underlying linear mixing model to solve the mixing problem. They are based on ideas such as Non-Negative Matrix Factorization (NNMF) [58], least-squares regression [12], minimum volume simplices [94] and convex cones [38]. Non-linear methods taking into account secondary scattering have been proposed such as bilinear models (e.g. [39]). Some models explicitly describe the multiple scattering properties of the media and result in nonlinear formulations such as the ones developed by Hapke [41]. Other techniques attempt to identify endmembers in intimate mixing scenarios by non-parametric graph-based algorithms [74]. Such systems only make some general assumptions on the geometrical properties of the data such as, endmembers are always on the boundaries of such a data cloud.

Further, planetary hyperspectral datasets are affected by complicating factors such as the presence of noise, instrument distortions and outliers [45]. Also, for sensors studying planetary surfaces the atmospheric calibration is complicated due to the scarcity of usable ground-truth information. The leads to spectral distortions caused by the incorrect correction of the atmospheric effects. The presence of such spectral distortions can significantly affect the spectral shape of data obtained from the planetary instruments. In addition, these instruments are rife with significant amounts of noise. The noise further makes the application of geometric and physical methods for data processing harder as the presence of noise distorts/eliminates the presence of the geometrical structure in the data-cloud from these datasets [74]. Given these data conditions atmospheric correction and denoising are an essential pre-processing step

for hyperspectral data-processing (see Itoh and Parente (2021) [45] for an example of an atmospheric correction pipeline for CRISM data and the issues therein).

In spite of these complicating factors very often such remote sensing data sources such as hyperspectral images are the best sources of information of large swathes of often inaccessible surfaces at a great distance with which there is minimal contact.

A.1 Interpreting spectral data

As has been mentioned hyperspectral images take the form of $3D$ cube, wherein each pixel in the $2D$ spatial grid is associated with a vector, and as such each pixel in a hyperspectral image can be modeled as a spectrum as has been shown by the pixel-spectral plot in Fig. A.1. Fig. A.2 shows a selected spectrum from the CRISM image $FRT00003E12$, which observes a portion of the Martian surface in the Nili Fossae region [the image center coordinates are $22.260495°N, 77.076525°E$]. Fig. A.2 (B) shows the spectrum associated with pixel marked in red in Fig. A.2 (A). This spectrum appears to be a combination of smooth global shape and some localized absorptions (which are spread over small spectral regions), which are referred to as absorption bands. The smooth global shape is referred to as the continuum is generally modeled by the convex-hull envelope fit to the spectrum, as shown by the green-spectrum in Fig. A.2 (C). The value of the continuum depends on the global factors of the measurement such as angle of incident light to the surface normal (incidence angle), angle of the detector to the surface normal (emission angle), the atmospheric conditions, and other global factors such as cloud-cover etc.. Further, since it is known that different materials has a unique set of frequencies over which it absorbs when light is diffracted off it, these absorptions can be used identify the specific mineral components present on the chosen spatial location. Spectroscopists have used the size and shape of the absorption bands at 2.31 & $2.51\mu m$, to identify the presence of a Mg-Carbonate at that location. The absorption features are enhanced in the

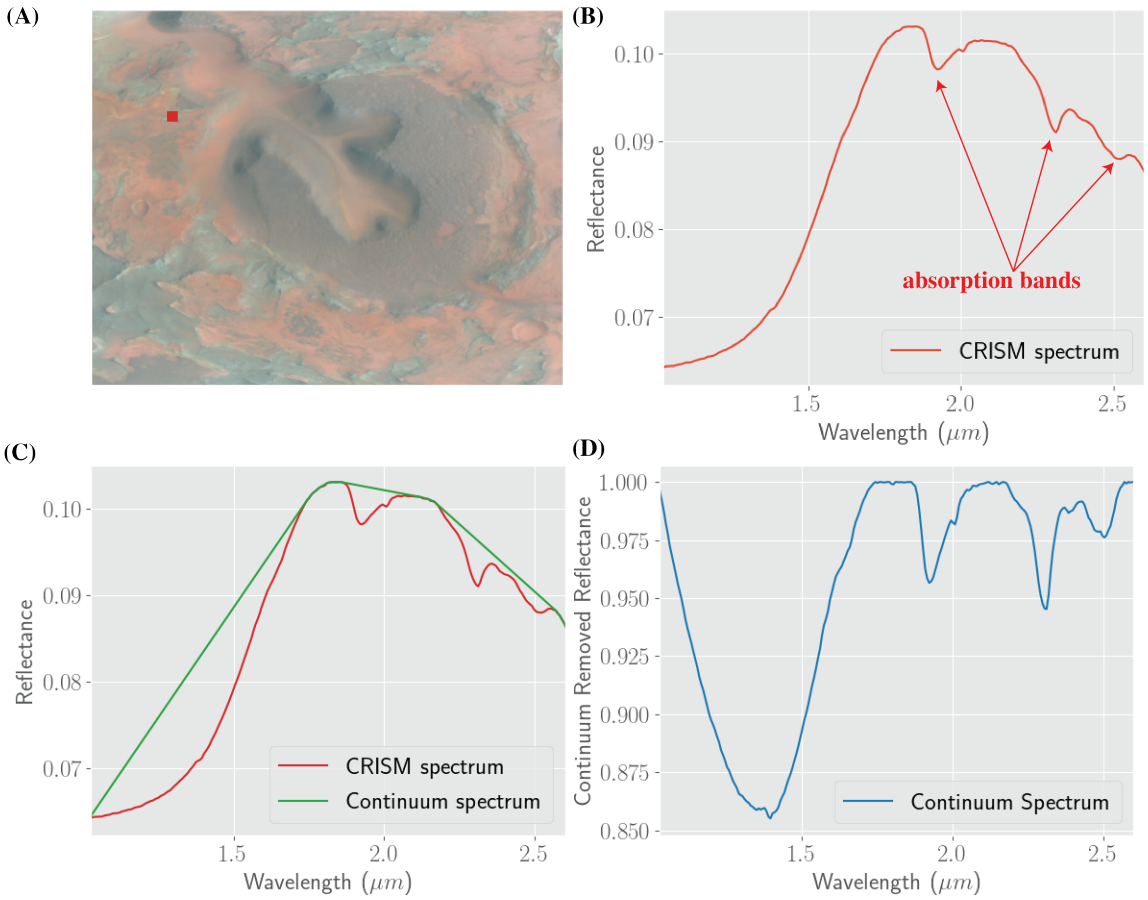

Figure A.2. Interpreting Spectral Information (A) an RGB composite of the CRISM image $FRT00003E12$ with the chosen pixel marked in red, (B) The spectrum associated with the pixel marked in (A), (C) The pixel spectrum and the associated continuum (D) the continuum-removed spectrum

continuum-removed spectrum as shown in Fig. A.2 (D). The absorption features in the continuum spectrum depends primarily upon the mineral composition of the surface at the specific spatial location as opposed to global measurement factors. These absorption features have been instrumental in successfully identifying and mapping many different minerals on various such far-away unreachable terrains.

BIBLIOGRAPHY

[1] Allender, Elyse, and Stepinski, Tomasz F. Automatic, exploratory mineralogical mapping of crism imagery using summary product signatures. *Icarus 281* (2017), 151–161.

[2] Amador, Elena S, Bandfield, Joshua L, and Thomas, Nancy H. A search for minerals associated with serpentinization across mars using crism spectral data. *Icarus 311* (2018), 113–134.

[3] Ankerst, Mihael, Breunig, Markus M, Kriegel, Hans-Peter, and Sander, Jörg. Optics: Ordering points to identify the clustering structure. *ACM Sigmod record 28*, 2 (1999), 49–60.

[4] Ash, Jordan, and Adams, Ryan P. On warm-starting neural network training. *Advances in Neural Information Processing Systems 33* (2020), 3884–3894.

[5] Bandfield, Joshua L, Christensen, Philip R, and Smith, Michael D. Spectral data set factor analysis and end-member recovery: Application to analysis of martian atmospheric particulates. *Journal of Geophysical Research: Planets 105*, E4 (2000), 9573–9587.

[6] Bandfield, Joshua L, Edgett, Kenneth S, and Christensen, Philip R. Spectroscopic study of the moses lake dune field, washington: Determination of compositional distributions and source lithologies. *Journal of Geophysical Research: Planets 107*, E11 (2002), 2–1.

[7] Bedini, Enton. The use of hyperspectral remote sensing for mineral exploration: A review. *Journal of Hyperspectral Remote Sensing 7*, 4 (2017), 189–211.

[8] Bendale, Abhijit, and Boult, Terrance E. Towards open set deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016).

[9] Berg, Mark de, Kreveld, Marc van, Overmars, Mark, and Schwarzkopf, Otfried. Computational geometry. In *Computational geometry*. Springer, 1997, pp. 1–17.

[10] Bibring, Jean-Pierre, Langevin, Yves, Gendrin, Aline, Gondet, Brigitte, Poulet, François, Berthé, Michel, Soufflot, Alain, Arvidson, Ray, Mangold, Nicolas, Mustard, John, et al. Mars surface diversity as revealed by the omega/mars express observations. *Science 307*, 5715 (2005), 1576–1581.

[11] Bibring, JP, Soufflot, A, and OMEGA, BM. Observatoire pour la minéralogie, l'eau, les glaces et l'activité. *Mars Express: The Scientific Payload. vol. ESA SP-1240* (2004), 3749.

[12] Bioucas-Dias, José M. A variable splitting augmented lagrangian approach to linear spectral unmixing. In *2009 First workshop on hyperspectral image and signal processing: Evolution in remote sensing* (2009), IEEE, pp. 1–4.

[13] Bioucas-Dias, José M, Plaza, Antonio, Dobigeon, Nicolas, Parente, Mario, Du, Qian, Gader, Paul, and Chanussot, Jocelyn. Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE journal of selected topics in applied earth observations and remote sensing 5*, 2 (2012), 354–379.

[14] Bishop, Janice L. Remote detection of phyllosilicates on mars and implications for climate and habitability. In *From habitability to life on Mars*. Elsevier, 2018, pp. 37–75.

[15] Bishop, Janice L, Dobrea, Eldar Z Noe, McKeown, Nancy K, Parente, Mario, Ehlmann, Bethany L, Michalski, Joseph R, Milliken, Ralph E, Poulet, Francois, Swayze, Gregg A, Mustard, John F, et al. Phyllosilicate diversity and past aqueous activity revealed at mawrth vallis, mars. *Science 321*, 5890 (2008), 830–833.

[16] Borengasser, Marcus, Hungate, William S, and Watkins, Russell. *Hyperspectral remote sensing: principles and applications*. CRC press, 2007.

[17] Carter, J, Poulet, F, Murchie, S, and Bibring, JP. Automated processing of planetary hyperspectral datasets for the extraction of weak mineral signatures and applications to crism observations of hydrated silicates on mars. *Planetary and Space Science 76* (2013), 53–67.

[18] Carter, John, Riu, Lucie, Poulet, François, Bibring, Jean-Pierre, Langevin, Yves, and Gondet, Brigitte. A mars orbital catalog of aqueous alteration signatures (mocaas). *Icarus 389* (2023), 115164.

[19] Chang, Chein-I. *Hyperspectral data exploitation: theory and applications*. John Wiley & Sons, 2007.

[20] Christensen, Philip R, Bandfield, JL, Clark, RN, Edgett, KS, Hamilton, VE, Hoefen, T, Kieffer, HH, Kuzmin, RO, Lane, MD, Malin, MC, et al. Detection of crystalline hematite mineralization on mars by the thermal emission spectrometer: Evidence for near-surface water. *Journal of Geophysical Research: Planets 105*, E4 (2000), 9623–9642.

[21] Christensen, Philip R, Bandfield, Joshua L, Hamilton, Vicky E, Ruff, Steve W, Kieffer, Hugh H, Titus, Timothy N, Malin, Michael C, Morris, Richard V, Lane, Melissa D, Clark, RL, et al. Mars global surveyor thermal emission spectrometer experiment: investigation description and surface science results. *Journal of Geophysical Research: Planets 106*, E10 (2001), 23823–23871.

[22] Christensen, Philip R, Jakosky, Bruce M, Kieffer, Hugh H, Malin, Michael C, McSween, Harry Y, Nealson, Kenneth, Mehall, Greg L, Silverman, Steven H, Ferry, Steven, Caplinger, Michael, et al. The thermal emission imaging system (themis) for the mars 2001 odyssey mission. *Space Science Reviews 110* (2004), 85–130.

[23] Clark, Roger N, and Roush, Ted L. Reflectance spectroscopy: Quantitative analysis techniques for remote sensing applications. *Journal of Geophysical Research: Solid Earth 89*, B7 (1984), 6329–6340.

[24] Dana, James Dwight. *A system of mineralogy*. BoD–Books on Demand, 2022.

[25] DeMaris, Alfred. A tutorial in logistic regression. *Journal of Marriage and the Family* (1995), 956–968.

[26] Dobigeon, Nicolas, Tourneret, Jean-Yves, Richard, Cédric, Bermudez, José Carlos M, McLaughlin, Stephen, and Hero, Alfred O. Nonlinear unmixing of hyperspectral images: Models and algorithms. *IEEE Signal processing magazine 31*, 1 (2013), 82–94.

[27] Ehlmann, Bethany L, and Edwards, Christopher S. Mineralogy of the martian surface. *Annual Review of Earth and Planetary Sciences 42* (2014).

[28] Ehlmann, Bethany L, and Mustard, John F. An in-situ record of major environmental transitions on early mars at northeast syrtis major. *Geophysical research letters 39*, 11 (2012).

[29] Ehlmann, Bethany L, Mustard, John F, Murchie, Scott L, Poulet, Francois, Bishop, Janice L, Brown, Adrian J, Calvin, Wendy M, Clark, Roger N, Des Marais, David J, Milliken, Ralph E, et al. Orbital identification of carbonate-bearing rocks on mars. *Science 322*, 5909 (2008), 1828–1832.

[30] Ehlmann, Bethany L, Mustard, John F, Swayze, Gregg A, Clark, Roger N, Bishop, Janice L, Poulet, Francois, Des Marais, David J, Roach, Leah H, Milliken, Ralph E, Wray, James J, et al. Identification of hydrated silicate minerals on mars using mro-crism: Geologic context near nili fossae and implications for aqueous alteration. *Journal of Geophysical Research: Planets 114*, E2 (2009).

[31] Felzenszwalb, Pedro F, and Huttenlocher, Daniel P. Efficient graph-based image segmentation. *International journal of computer vision 59*, 2 (2004), 167–181.

[32] Gao, Bolin, and Pavel, Lacra. On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805* (2017).

[33] Geminale, A, Grassi, Davide, Altieri, FRANCESCA, Serventi, G, Carli, CRISTIAN, Carrozzo, FILIPPO GIACOMO, Sgavetti, M, Orosei, ROBERTO, D'Aversa, EMILIANO, Bellucci, Giancarlo, et al. Removal of atmospheric features in near infrared spectra by means of principal component analysis and target transformation on mars: I. method. *Icarus 253* (2015), 51–65.

[34] Gilmore, Martha S, Thompson, David R, Anderson, Laura J, Karamzadeh, Nader, Mandrake, Lukas, and Castaño, Rebecca. Superpixel segmentation for analysis of hyperspectral data sets, with application to compact reconnaissance imaging spectrometer for mars data, moon mineralogy mapper data, and ariadnes chaos, mars. *Journal of Geophysical Research: Planets 116*, E7 (2011).

[35] Glotch, Timothy D, and Bandfield, Joshua L. Determination and interpretation of surface and atmospheric miniature thermal emission spectrometer spectral end-members at the meridiani planum landing site. *Journal of Geophysical Research: Planets 111*, E12 (2006).

[36] Goodfellow, Ian. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160* (2016).

[37] Green, Robert O, Eastwood, Michael L, Sarture, Charles M, Chrien, Thomas G, Aronsson, Mikael, Chippendale, Bruce J, Faust, Jessica A, Pavri, Betina E, Chovit, Christopher J, Solis, Manuel, et al. Imaging spectroscopy and the airborne visible/infrared imaging spectrometer (aviris). *Remote sensing of environment 65*, 3 (1998), 227–248.

[38] Gruninger, John H, Ratkowski, Anthony J, and Hoke, Michael L. The sequential maximum angle convex cone (smacc) endmember model. In *Algorithms and technologies for multispectral, hyperspectral, and ultraspectral imagery X* (2004), vol. 5425, SPIE, pp. 1–14.

[39] Halimi, Abderrahim, Altmann, Yoann, Dobigeon, Nicolas, and Tourneret, Jean-Yves. Nonlinear unmixing of hyperspectral images using a generalized bilinear model. *IEEE Transactions on Geoscience and Remote Sensing 49*, 11 (2011), 4153–4162.

[40] Hamilton, Victoria E, and Ruff, Steven W. Distribution and characteristics of adirondack-class basalt as observed by mini-tes in gusev crater, mars and its possible volcanic source. *Icarus 218*, 2 (2012), 917–949.

[41] Hapke, Bruce. *Theory of reflectance and emittance spectroscopy*. Cambridge university press, 2012.

[42] Heylen, Rob, Parente, Mario, and Gader, Paul. A review of nonlinear hyperspectral unmixing methods. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 7*, 6 (2014), 1844–1868.

[43] Hubbard, G Scott, Naderi, Firouz M, and Garvin, James B. Following the water, the new program for mars exploration. *Acta Astronautica 51*, 1-9 (2002), 337–350.

[44] Ioffe, Sergey, and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015).

[45] Itoh, Yuki, and Parente, Mario. A new method for atmospheric correction and de-noising of crism hyperspectral data. *Icarus 354* (2021), 114024.

[46] Keshava, Nirmal, and Mustard, John F. Spectral unmixing. *IEEE signal processing magazine 19*, 1 (2002), 44–57.

[47] Kingma, Diederik P, and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[48] Kingma, Diederik P, and Welling, Max. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).

[49] Kotsiantis, Sotiris B. Decision trees: a recent overview. *Artificial Intelligence Review 39* (2013), 261–283.

[50] Kruschke, John K. Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science 1*, 5 (2010), 658–676.

[51] Langevin, Yves, Poulet, François, Bibring, Jean-Pierre, and Gondet, Brigitte. Sulfates in the north polar region of mars detected by omega/mars express. *Science 307*, 5715 (2005), 1584–1586.

[52] Larsen, Anders Boesen Lindbo, Sønderby, Søren Kaae, Larochelle, Hugo, and Winther, Ole. Autoencoding beyond pixels using a learned similarity metric. In *International conference on machine learning* (2016), PMLR, pp. 1558–1566.

[53] Leask, EK, Ehlmann, BL, Dundar, MM, Murchie, SL, and Seelos, FP. Challenges in the search for perchlorate and other hydrated minerals with 2.1-μm absorptions on mars. *Geophysical research letters 45*, 22 (2018), 12–180.

[54] Liangrocapart, S, and Petrou, Maria. Mixed pixels classification. In *Remote Sensing* (1998), International Society for Optics and Photonics, pp. 72–83.

[55] Lin, Honglei, Tarnas, JD, Mustard, JF, Zhang, Xia, Wei, Yong, Wan, Weixing, Klein, F, and Kellner, JR. Dynamic aperture factor analysis/target transformation (dafa/tt) for mg-serpentine and mg-carbonate mapping on mars with crism near-infrared data. *Icarus 355* (2021), 114168.

[56] Malinowski, Edmund R, and Howery, Darryl G. *Factor analysis in chemistry*, vol. 3. Wiley New York, 1980.

[57] Malkoff, Donald B, and Oliver, William R. Hyperspectral imaging applied to forensic medicine. In *Spectral Imaging: Instrumentation, Applications, and Analysis* (2000), vol. 3920, International Society for Optics and Photonics, pp. 108–116.

[58] Miao, Lidan, and Qi, Hairong. Endmember extraction from highly mixed data using minimum volume constrained nonnegative matrix factorization. *IEEE Transactions on Geoscience and Remote Sensing 45*, 3 (2007), 765–777.

[59] Milliken, RE, Hiroi, T, Scholes, D, Slavney, S, and Arvidson, R. The nasa reflectance experiment laboratory (relab) facility: An online spectral database for planetary exploration. *LPI Contributions 2654* (2021).

[60] Murchie, Scott, Arvidson, R, Bedini, Peter, Beisser, K, Bibring, J-P, Bishop, J, Boldt, J, Cavender, P, Choo, T, Clancy, RT, et al. Compact reconnaissance imaging spectrometer for mars (crism) on mars reconnaissance orbiter (mro). *Journal of Geophysical Research: Planets 112*, E5 (2007).

[61] Murchie, Scott L, Mustard, John F, Ehlmann, Bethany L, Milliken, Ralph E, Bishop, Janice L, McKeown, Nancy K, Noe Dobrea, Eldar Z, Seelos, Frank P, Buczkowski, Debra L, Wiseman, Sandra M, et al. A synthesis of martian aqueous mineralogy after 1 mars year of observations from the mars reconnaissance orbiter. *Journal of Geophysical Research: Planets 114*, E2 (2009).

[62] Mustard, John F, Murchie, Scott L, Pelkey, SM, Ehlmann, BL, Milliken, RE, Grant, John A, Bibring, J-P, Poulet, Francois, Bishop, Janice, Dobrea, E Noe Dobrea, et al. Hydrated silicate minerals on mars observed by the mars reconnaissance orbiter crism instrument. *Nature 454*, 7202 (2008), 305–309.

[63] Mustard, John F, Poulet, Francois, Gendrin, Aline, Bibring, J-P, Langevin, Yves, Gondet, Brigitte, Mangold, Nicholas, Bellucci, Giancarlo, and Altieri, Francesca. Olivine and pyroxene diversity in the crust of mars. *Science 307*, 5715 (2005), 1594–1597.

[64] Pahlevan, Nima, Smith, Brandon, Binding, Caren, Gurlin, Daniela, Li, Lin, Bresciani, Mariano, and Giardino, Claudia. Hyperspectral retrievals of phytoplankton absorption and chlorophyll-a in inland and nearshore coastal waters. *Remote Sensing of Environment 253* (2021), 112200.

[65] Parente, M, Arvidson, R, Itoh, Y, Lin, H, Mustard, J. F., Saranathan, A. M., Seelos, F. P., and Tarnas, J. D. Mineral detections over jezero crater using advanced data processing techniques for crism data—the crism" fandango". *LPI Contributions 2089* (2019).

[66] Parente, M, Arvidson, RE, Itoh, Y, Lin, H, Mustard, JF, Saranathan, AM, Seelos, FP, and Tarnas, JD. Convergence on mineral detections over gale crater, ne syrtis, and jezero crater using advanced data processing techniques for crism hyperspectral imaging data. In *Lunar and Planetary Science Conference* (2019), vol. 50.

[67] Parente, Mario, Makarewicz, Heather D, and Bishop, Janice L. Decomposition of mineral absorption bands using nonlinear least squares curve fitting: Application to martian meteorites and crism data. *Planetary and Space Science 59*, 5-6 (2011), 423–442.

[68] Park, Bosoon, and Lu, Renfu. *Hyperspectral imaging technology in food and agriculture*. Springer, 2015.

[69] Pelkey, SM, Mustard, JF, Murchie, S, Clancy, RT, Wolff, M, Smith, M, Milliken, R, Bibring, J-P, Gendrin, A, Poulet, F, et al. Crism multispectral summary products: Parameterizing mineral diversity on mars from reflectance. *Journal of Geophysical Research: Planets 112*, E8 (2007).

[70] Pieters, Carle M, Boardman, Joseph, Buratti, Bonnie, Chatterjee, Alok, Clark, Roger, Glavich, Tom, Green, Robert, Head III, James, Isaacson, Peter, Malaret, Erick, et al. The moon mineralogy mapper (m^3) on chandrayaan-1. *Current Science* (2009), 500–505.

[71] Plebani, Emanuele, Ehlmann, Bethany L, Leask, Ellen K, Fox, Valerie K, and Dundar, M Murat. A machine learning toolkit for crism image analysis. *Icarus* (2022), 114849.

[72] Poulet, François, Bibring, J-P, Mustard, JF, Gendrin, A, Mangold, N, Langevin, Y, Arvidson, RE, Gondet, B, and Gomez, C. Phyllosilicates on mars and implications for early martian climate. *Nature 438*, 7068 (2005), 623–627.

[73] Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* (2015).

[74] Rohani, Neda, and Parente, Mario. Graph-based identification of boundary points for unmixing and anomaly detection. In *2013 5th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)* (2013), IEEE, pp. 1–4.

[75] Saranathan, Arun M, and Parente, Mario. Uniformity-based superpixel segmentation of hyperspectral images. *IEEE Transactions on Geoscience and Remote Sensing 54*, 3 (2015), 1419–1430.

[76] Saranathan, Arun M, and Parente, Mario. Adversarial feature learning for improved mineral mapping of crism data. *Icarus 355* (2021), 114107.

[77] Scheirer, Walter J., Rocha, Anderson, Micheals, Ross J., and Boult, Terrance E. Meta-recognition: The theory and practice of recognition score analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence 33*, 8 (2011), 1689–1695.

[78] Seelos, FP, Murchie, SL, Humm, DC, Barnouin, OS, Morgan, F, Taylor, HW, Hash, C, Team, CRISM, et al. Crism data processing and analysis products update—calibration, correction, and visualization. In *42nd Annual Lunar and Planetary Science Conference* (2011), no. 1608, pp. 1438–1439.

[79] Seelos, FP, Viviano-Beck, CE, Morgan, MF, Romeo, G, Aiello, JJ, and Murchie, SL. Crism hyperspectral targeted observation pds product sets—ters and mtrdrs. In *Lunar and Planetary Science Conference* (2016), no. 1903, pp. 1783–1784.

[80] Seelos, Frank P, Seelos, Kimberly D, Murchie, Scott L, Novak, M Alexandra Matiella, Hash, Christopher D, Morgan, M Frank, Arvidson, Raymond E, Aiello, John, Bibring, Jean-Pierre, Bishop, Janice L, et al. The crism investigation in mars orbit: Overview, history, and delivered data products. *Icarus* (2023), 115612.

[81] Shu, Lei, Xu, Hu, and Liu, Bing. Doc: Deep open classification of text documents. *arXiv preprint arXiv:1709.08716* (2017).

[82] Skok, John R, Mustard, John F, Murchie, Scott L, Wyatt, Michael B, and Ehlmann, Bethany L. Spectrally distinct ejecta in syrtis major, mars: Evidence for environmental change at the hesperian-amazonian boundary. *Journal of Geophysical Research: Planets 115*, E2 (2010).

[83] Skok, JR, Mustard, JF, Tornabene, LL, Pan, C, Rogers, D, and Murchie, SL. A spectroscopic analysis of martian crater central peaks: Formation of the ancient crust. *Journal of Geophysical Research: Planets 117*, E11 (2012).

[84] Southam, G., and Westall, F. 10.12 - geology, life and habitability. In *Treatise on Geophysics*, Gerald Schubert, Ed. Elsevier, Amsterdam, 2007, pp. 421–437.

[85] Squyres, Steven W, Grotzinger, John P, Arvidson, Raymond E, Bell III, James F, Calvin, Wendy, Christensen, Philip R, Clark, Benton C, Crisp, JA, Farrand, William H, Herkenhoff, Ken E, et al. In situ evidence for an ancient aqueous environment at meridiani planum, mars. *science 306*, 5702 (2004), 1709–1714.

[86] Srivastava, Nitish, Hinton, Geoffrey, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research 15*, 1 (2014), 1929–1958.

[87] Tarnas, J. D., Mustard, J. F., Lin, H., Goudge, T.A., Amador, E. S. Bramble, M.S., and Zhang, X. Hydrated silica in the jezero delta. In *Lunar and Planetary Science Conference* (2019), vol. 50.

[88] Tarnas, Jesse, Stack Morgan, Kathryn, Parente, Mario, Mustard, John, Koeppel, Ari, Moore, Kelsey, Horgan, Briony Noelle Heather, Seelos, Frank, Cloutis, Edward, Kelemen, Peter B, and et al. Origin of carbonate-bearing rocks in jezero crater: Implications for ancient habitability in subsurface environments. *Earth and Space Science Open Archive* (2021), 82.

[89] Thomas, Nancy H, and Bandfield, Joshua L. Identification and refinement of martian surface mineralogy using factor analysis and target transformation of near-infrared spectroscopic data. *Icarus 291* (2017), 124–135.

[90] Viviano-Beck, Christina E, Seelos, Frank P, Murchie, Scott L, Kahn, Eliezer G, Seelos, Kimberley D, Taylor, Howard W, Taylor, Kelly, Ehlmann, Bethany L, Wiseman, Sandra M, Mustard, John F, et al. Revised crism spectral parameters and summary products based on the currently detected mineral diversity on mars. *Journal of Geophysical Research: Planets 119*, 6 (2014), 1403–1431.

[91] Wagstaff, Kiri L, Lanza, Nina L, Thompson, David R, Dietterich, Thomas G, and Gilmore, Martha S. Guiding scientific discovery with explanations using demud. In *Twenty-Seventh AAAI Conference on Artificial Intelligence* (2013).

[92] Wendt, Lorenz, Gross, Christoph, Kneissl, Thomas, Sowe, Mariam, Combe, Jean-Philippe, LeDeit, Laetitia, McGuire, Patrick C, and Neukum, Gerhard. Sulfates and iron oxides in ophir chasma, mars, based on omega and crism observations. *Icarus 213*, 1 (2011), 86–103.

[93] West, Douglas Brent, et al. *Introduction to graph theory*, vol. 2. Prentice hall Upper Saddle River, 2001.

[94] Winter, Michael E. N-findr: An algorithm for fast autonomous spectral endmember determination in hyperspectral data. In *Imaging Spectrometry V* (1999), vol. 3753, International Society for Optics and Photonics, pp. 266–275.

[95] Wray, JJ, Ehlmann, BL, Squyres, SW, Mustard, JF, and Kirk, Randolph L. Compositional stratigraphy of clay-bearing layered deposits at mawrth vallis, mars. *Geophysical Research Letters 35*, 12 (2008).

[96] Yang, Jingkang, Zhou, Kaiyang, Li, Yixuan, and Liu, Ziwei. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334* (2021).