

Research Paper

A machine learning toolkit for CRISM image analysis

Emanuele Plebani^a, Bethany L. Ehlmann^b, Ellen K. Leask^{b,c}, Valerie K. Fox^d, M. Murat Dunder^{a,*}^a Computer and Information Sciences Department, Indiana University - Purdue University, Indianapolis, 46202, IN, USA^b Div. of Geological & Planetary Sciences, California Institute of Technology, Pasadena, 91125, CA, USA^c John Hopkins University Applied Physics Laboratory, Laurel, 20723, MD, USA^d Department of Earth and Environmental Sciences, University of Minnesota, Minneapolis, 55455, MN, USA

ARTICLE INFO

Dataset link: <https://cs.iupui.edu/~mdundar/CRISM.htm>

Keywords:

Machine learning
Hyperspectral image
Mars
Hierarchical Bayesian
CRISM

ABSTRACT

Hyperspectral images collected by remote sensing have played a significant role in the discovery of aqueous alteration minerals, which in turn have important implications for our understanding of the changing habitability on Mars. Traditional spectral analyzes based on summary parameters have been helpful in converting hyperspectral cubes into readily visualizable three channel maps highlighting high-level mineral composition of the Martian terrain. These maps have been used as a starting point in the search for specific mineral phases in images. Although the amount of labor needed to verify the presence of a mineral phase in an image is quite limited for phases that emerge with high abundance, manual processing becomes laborious when the task involves determining the spatial extent of detected phases or identifying small outcrops of secondary phases that appear in only a few pixels within an image. Thanks to extensive use of remote sensing data and rover expeditions, significant domain knowledge has accumulated over the years about mineral composition of several regions of interest on Mars, which allow us to collect reliable labeled data required to train machine learning algorithms. In this study we demonstrate the utility of machine learning in two essential tasks for hyperspectral data analysis: nonlinear noise removal and mineral classification. We develop a simple yet effective hierarchical Bayesian model for estimating distributions of spectral patterns and extensively validate this model for mineral classification on several test images. Our results demonstrate that machine learning can be highly effective in exposing tiny outcrops of specific phases in orbital data that are not uncovered by traditional spectral analysis. We package implemented scripts, documentation illustrating use cases, and pixel-scale training data collected from dozens of well-characterized images into a new toolkit. We hope that this new toolkit will provide advanced and effective processing tools and improve community's ability to map compositional units in remote sensing data quickly, accurately, and at scale.

1. Introduction

The Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) on board the Mars Reconnaissance Orbiter has enabled the discovery and mapping of a broad array of aqueous minerals on the surface of Mars (Murchie, Mustard, et al. 2009; Viviano-Beck et al. 2014). Hyperspectral data collected by CRISM have revolutionized our understanding of the planet and have been instrumental in the selection of landing sites for Mars rover exploration missions. Tens of thousands of images are available for analysis to understand the geologic history and habitability of Mars by identifying minerals and mapping mineral-bearing units.

1.1. Traditional methods and limitations

Direct automated matching of minerals from spectral libraries measured on Earth to CRISM data is not possible due to the fact that the observation is of large scale geologic materials representing a mix of minerals, which possess distinctive spectral properties, and because of the complexity of the dataset's signal processing. Summary parameters (Pelkey et al., 2007; Viviano-Beck et al., 2014), derived from key wavelengths that capture spectral features characteristics of specific mineral phases, form the essence of traditional CRISM mineral identification methods. These parameters can be used for high-level mineralogical mapping of CRISM images. Different combinations of

* Corresponding author.

E-mail address: mdundar@iupui.edu (M.M. Dunder).URL: <https://cs.iupui.edu/~mdundar> (M.M. Dunder).<https://doi.org/10.1016/j.icarus.2021.114849>

Received 29 June 2021; Received in revised form 18 October 2021; Accepted 6 December 2021

Available online 1 January 2022

0019-1035/© 2021 The Author(s).

Published by Elsevier Inc.

This is an open access article under the CC BY-NC-ND license

<http://creativecommons.org/licenses/by-nc-nd/4.0/>.

summary parameters are used to characterize each image in terms of oxidized iron minerals, mafic mineralogy, hydroxylated silicates, and ice content. More specific mineral identification often requires filtering noise to enhance spectral features. This is performed by finding the ratio of the average spectrum from an area of interest (numerator) and the average spectrum derived from a spectrally homogeneous “bland” region along-track the area of interest (denominator). Spectral ratioing reduces multiplicative noise while spatial averaging filters additive noise.

Summary parameters are also useful in ratioing for identifying candidate regions for the numerator and denominator by highlighting pixels of strongest signal to contribute to the former and pixels not highlighted in any summary parameter to contribute to the latter. Finding the right pair of numerator and denominator may require multiple trials. Although this amount of labor may be acceptable when determining the presence or absence of a specific mineral in an image, manual ratioing becomes a tedious task when it has to be done at many locations in the image to determine the spatial extent of one or more mineral phases. Summary parameters become less useful during ratioing when the region of interest spans only a limited number of not necessarily contiguous pixels or when the scene is mineralogically diverse, leaving few to no options for denominator pixels. When the average spectrum is computed over a small number of pixels, key spectral features may remain suppressed under random noise, and summary parameters may end up capturing artificial features that arise due to noise. Another limitation of summary parameters appears when key absorption bands of two spatially co-occurring minerals are within few channels of one another (such as alunite and kaolinite). The 6.55 nm increments between two channels in CRISM offer enough spectral resolution to differentiate between such phases in ideal conditions. However, considering the practical limitations of CRISM data and the occurrence of phases in mixtures, such a distinction may not be possible without harnessing a large amount of data from many images to identify less obvious spectral features characterizing these phases in different locales, a task for which machine learning methods are ideally suited.

1.2. Potential utility of machine learning in CRISM image analysis

Machine learning has the potential to improve CRISM image analysis in various tasks. Mineral identification and classification at pixel scale (Dundar and Ehlmann, 2016; Dundar et al., 2019b,a; Caggiano et al., 2019), outlier detection (Dundar et al., 2013; Leask et al., 2018), spectral and spatial denoising (Parente et al., 2014; Itoh and Parente, 2021; Arun and Parente, 2017; Kreisch et al., 2017; He et al., 2019; Saranathan and Parente, 2021) are some of the areas where machine learning has already shown some promise toward automating CRISM image analysis.

The algorithms that eventually became part of this toolkit have been validated and refined in multiple projects demanding automated analysis of CRISM imagery. We demonstrated the utility of machine learning algorithms in detecting small outcrops of mineral phases in images with extensive traditional spectral analyzes by reporting new geologic discoveries from NE Syrtis area and Jezero crater (Dundar et al., 2019b,a). Jezero crater and NE Syrtis are regions of high interest as places where the Mars-2020 rover will conduct its in situ exploration. Our algorithms detected a new hydrated iron oxide phase in NE Syrtis, attributed elsewhere on Mars to akageneite, along with several new rare phases including Al clays, jarosite, chlorite/smectite, and hydrated silica in Jezero. At Jezero and NE Syrtis, small detections of rare phases are crucial for guiding the Mars-2020 rover and for contextualizing its discoveries. In another project we studied stratigraphies with Al phyllosilicates on Mars to determine the implications for near-surface habitability during the Noachian and Hesperian eras. Several new jarosite and alunite regions have been identified in Nili Fossae, Terra Sirenum, and Mawrth Vallis (Ehlmann and Dundar, 2015; Leask et al.,

2019). As part of this verification process that included an effort to map occurrences of perchlorate, a salt that had been previously reported in CRISM data and is a freezing point depressant, we have discovered a previously unidentified artifact at 2.1 μm in CRISM I/F data that mimics the appearance of certain minerals, including perchlorate. Upon further verification with CRISM radiance data we showed that none of the perchlorate detections reported in the literature remain robust in these data sets (Leask et al., 2018).

Here, we present a new machine learning toolkit for advanced CRISM processing to improve the community’s ability to map discrete compositional units in remote-sensing data to more accurately identify mineral phases on Mars. The toolkit contains Python scripts, pixel-scale training data collected from dozens of well-characterized images, and documentation illustrating use cases of the algorithms on several test images. In this toolkit we tackle two specific tasks by machine learning, namely, bland region identification and mineral classification at pixel scale. The first task requires robust estimation of the distribution of bland pixels and the second task requires a generalizable supervised classifier, both of which in turn depend on the availability of a large amount of labeled data from a representative set of well-characterized CRISM images. Collection of these labeled data is hard because of the aforementioned challenges involving manual analysis. Fortunately, ever since the first set of CRISM data became available in 2006, CRISM images have been used in numerous studies, accumulating enough domain knowledge in the literature to obtain a relatively large set of labeled data from well-characterized CRISM images. Although the current set of images studied in the literature may not be spatially or spectrally representative of all mineral phases on Mars, the amount of labeled data we can obtain based on this domain knowledge is still sufficiently large enough to train basic machine learning algorithms. This first generation machine learning models implemented with limited labeled data may become instrumental in facilitating the collection of more labeled data on a larger scale, but more importantly may help compellingly demonstrate the utility of machine learning in planetary exploration and pave the way for development of more advanced and effective second-generation algorithms.

2. Training data collection and labeling

In the first phase of data collection a small subset of relatively well-characterized CRISM images from Nili Fossae and Mawrth Vallis regions were processed by a nonparametric Bayesian clustering technique (Yerebakan et al., 2014; Dundar and Ehlmann, 2016). This method generates a few hundred spectra per processed image, which are visually inspected and classified into the mineral phases already reported in the literature in these images to create an initial spectral training library. Detections with limited spatial extent and other detections with ambiguous average spectra are ignored. Spectra with no conspicuous spectral features are added into the training set to represent the *bland* category. Selection is performed in a semi-automated way; spectra are extracted algorithmically while assignment to the bland category is performed manually.

In the second phase, this training library is used to implement two models: a bland pixel scoring function for column-wise ratioing and a classifier model that operates on the ratioed data to predict mineral phases associated with each pixel, which are described in detail in Section 5. Both the scoring function and the classifier use our two-layer Bayesian Gaussian mixture model, which is described in Section 3.

We then obtain a more representative training library using these models in an active learning scheme. After each new image is classified by the current version of the classifier, detections are visually verified at the connected-component level,¹ and unambiguous detections,

¹ A connected component in an image is defined as a group of pixels connected to each other by 8-pixel connectivity. For low quality images this restriction is relaxed to have near contiguous pixels, which are pixels that are separated by 1–3 pixels.

consistent with detections reported in the literature, are added to the training library. The updated training library is used to train a new version of the classifier, which in turn is used to classify future images. Using this active learning framework we processed over five hundred CRISM images over time, mainly from Nili Fossae, Mawrth Vallis, Terra Sirenum, Valles Marineris, Libya Montes, the polar regions, and select craters on Mars. As only detections with unambiguous spectra and those consistent with mineral detections reported in the literature are included, our current spectral training set includes labeled data from over seventy well-characterized images.

2.1. List of mineral phases and training images

The list of CRISM images, their regions, and mineral phases detected in these images is shown in Table 1. The number of labeled pixels for each phase and the number of images from which these data are collected are shown in Table 2. Images of segmentation masks and plots of average spectra for each phase detected are provided as a supplementary document. This mineral training set contains ratioed data at pixel scale, which is used to train the mineral classifier. We also maintain a separate training set with unratioed data from only bland pixels collected from over three hundred images. This data is used to estimate the distribution of the *bland* class, which is then used to evaluate the bland likelihood of a pixel when ratioing new images as described in Section 4. Both training datasets are included with the toolkit.

2.2. Dataset structure

The variables stored in the datasets are reported in Table 3. Pixel-scale spectra, class labels, and image identifiers are stored in the `pixspec`, `pixlabs`, and `pixims` variables, respectively. Spectra contain data from 350 channels, which include the ranges 1.0210 μm –2.6483 μm and 2.8070 μm –3.4769 μm . However, the channels corresponding to ranges 1.6447 μm –1.7303 μm and 1.9741 μm –2.0600 μm are excluded because of known artifacts, and channels corresponding to 2.8070 μm –3.4769 μm are excluded because data in these channels do not offer much additional information for classifying minerals of interest and the longer wavelengths show low data quality and residual artifacts (Murchie et al., 2009). Coordinates of pixels are stored in `pixcrds`. The ratioed spectra (Table 3b) also include a numeric ID in the variable `pixpat` to uniquely determine connected components for each image. The image IDs vary between one and the number of images in the dataset. The variable `im_names` maps this image ID to the last five digits of the CRISM image identifier of the source image. The coordinates are relative to the Targeted Reduced Data Record (TRDR) images, i.e., without geometric correction. Each mineral phase is associated with a unique numeric class label; Table 2 shows the correspondence between class labels and mineral phases.

3. Hierarchical Bayesian model

At the core of our machine learning approach is a hierarchical Bayesian Model (HBM), illustrated in Fig. 2, to estimate the spectral distributions of mineral classes. HBM uses a Gaussian mixture model (GMM) to model pixel-scale training data for each mineral class. The number of Gaussian distributions in the mixture model for a given mineral class is determined by the number of images in which that mineral class is detected. The model has thus a two-layer structure: the lower layer models spectral variations of the same class across images whereas the upper layer models spectral variations across different classes. This two-layer structure gives flexibility and robustness in modeling spectral distributions of mineral classes. Our approach is motivated by the observation that different instances of the same mineral class detected across different images exhibit varying spectral

properties due to differences in atmospheric effects and viewing geometry, as well as inherent differences in spectral properties of the surface materials. Fig. 1 shows the average spectra and individual pixel-scale spectra for a randomly selected ten pixels for two mineral phases (Alunite and Jarosite) to highlight the spectral variations within and between mineral phases.

The model is based on the principles of Bayesian inference (Gelman et al., 1995), where the data is assumed to be produced by a *generative model*, i.e., a specific distribution (a mixture of Gaussian mixtures in our case) but the parameters of the distribution, i.e., the means and covariances of single Gaussian components, are themselves considered random variables with their own distributions. These vaguely defined distributions of the parameters (the *prior*) when combined with the training dataset by Bayes rule gives rise to a more informative distribution with lower variance and uncertainty compared to the prior, called the *posterior*.

In HBM the Gaussian distributions corresponding to the observed instances of the same mineral class across all images are regulated by a shared local prior in the lower layer, and local priors associated with each mineral class are in turn modeled by a global prior in the upper layer. In this context the local prior can be thought of as a template distribution for the distribution of the mineral classes and the global prior in turn can be interpreted as a template distribution for the distribution of the local priors. A template Gaussian component is the one that combines spectral characteristics of all mineral classes sharing the same prior in a single Gaussian distribution. Spectral signatures describing mineral phases are highly structured and continuous signals with high levels of correlation across certain ranges of bands. The prior can be considered as a template Gaussian that captures this structure and correlation in viable signals so that informative signals *a priori* can be chosen over non-informative ones (such as random noise) by the model.

We use the following generative model to fit spectral data available in our training set:

$$\text{Data model: } x_{ijk} \sim \mathcal{N}(\mu_{jk}, \Sigma_k) \quad (1)$$

$$\text{Local prior: } \mu_{jk} \sim \mathcal{N}(\mu_k, \Sigma_k \kappa_1^{-1}) \quad (2)$$

$$\text{Global prior: } \mu_k \sim \mathcal{N}(\mu_0, \Sigma_k \kappa_0^{-1}) \quad (3)$$

$$\Sigma_k \sim \text{IW}(\Sigma_0, m) \quad (4)$$

where the indices k , j , and i enumerate mineral classes, the instances (images) for class k , and individual pixels inside image j and from class k , respectively. $\mathcal{N}(\mu, \Sigma)$ denotes the Normal distribution with mean μ and covariance matrix Σ , whereas $\text{IW}(\Sigma_0, m)$ denotes the inverse Wishart distribution, with scale matrix Σ_0 and degrees of freedom m . The inverse Wishart distribution is a multivariate generalization of the inverse Gamma distribution, and it can be thought as a distribution over covariances (Nydic, 2012). The prior mean μ_0 and covariance Σ_0 are the average of all class means and covariances estimated for each training image in the dataset. The hyperparameters κ_0 and κ_1 control the amount of dispersion of the class means and the means of the intra-class Gaussian components respectively; see the supplementary document for a more detailed discussion.

The posterior distribution allows us to compute the likelihood of a new pixel belonging to a specific mineral class. For our choice of priors this step can be computed in closed form by evaluating the Bayes rule to produce a multivariate *Student-t* distribution as the posterior predictive distribution:

$$P(x|D) = T(x_{ji} | \bar{\mu}_k, \bar{\Sigma}_s, \bar{v}_s) \quad (5)$$

Now the likelihood of a new pixel represented by a spectrum x depends only on the training dataset D and hyperparameters of the model. The parameters $\bar{\mu}_k$, $\bar{\Sigma}_s$ and \bar{v}_s depend on sample means, covariances, sizes, and hyperparameters; see the supplementary document for technical details.

Table 1

List of CRISM images included in the training set. The first column shows the last five digits of the CRISM image IDs, the middle column shows the names of regions from which respective images are taken, the last column shows mineral detections included in the training set from the respective images^a. Segmentation maps and average ratioed spectra for each mineral detection are provided in the supplementary material.

ID	Region	Mineral detections
098B2	Aram Chaos	Ferric hydroxysulfate, monohydrated sulfate
09312	Baldet crater	Analcime, hydrated silica
02885	North Circumpolar dunes	H ₂ O ice, gypsum
096BF	Robert Sharp crater	Akaganeite
0634B	Claritas rise	Serpentine, illite, chlorite
062B6	Columbus crater	Jarosite, alunite, kaolinite, polyhydrated sulfate 2, monohydrated sulfate
0750A	Columbus crater	Alunite, kaolinite, polyhydrated sulfate
07D87	Columbus crater	Kaolinite, polyhydrated sulfate 1, polyhydrated sulfate 2
08565	Columbus crater	Alunite, kaolinite, polyhydrated sulfate 1, monohydrated sulfate
1212A	Columbus crater	Polyhydrated sulfate 1, polyhydrated sulfate 2
12C19	Columbus crater	Kaolinite, polyhydrated sulfate 1, polyhydrated sulfate 2
13D1F	Columbus crater	Alunite, kaolinite, polyhydrated sulfate 2
13EEF	Columbus crater	Alunite, kaolinite, polyhydrated sulfate 1
167FA	Columbus crater	Alunite, kaolinite, polyhydrated sulfate 1, polyhydrated sulfate 2
0987B	Cross crater	Alunite
0B252	Cross crater	Alunite, kaolinite
0CC44	Cross crater	Alunite, kaolinite
1187B	Cross crater	Alunite, kaolinite
12E09	Cross crater	Alunite
137C2	Cross crater	Alunite
21B59	Cross crater	Alunite, kaolinite
0B385	Eos Chasma	Monohydrated sulfate, hematite
096FE	Eridania Basin	Fe smectite, jarosite
0BABA	Gale crater	(Artifact)
09036	Huygens crater	Ca/Fe carbonate, illite, prehnite
027E2	Ius Chasma	Jarosite
0A91C	Ius Chasma	Jarosite, monohydrated sulfate
20AE1	Ius Chasma	Polyhydrated sulfate 2, monohydrated sulfate
05814	Juventae Chasma	Hydrated silica, ferricopiapite
0BAD4	Kunowsky crater	Mg Olivine
0C26C	Kunowsky crater	Mg Olivine
0285A	Mawrth Vallis	Fe smectite, jarosite, Al smectite
03BFB	Mawrth Vallis	Fe smectite, alunite, Al smectite
043EC	Mawrth Vallis	Fe smectite, Al smectite
0863E	Mawrth Vallis	Fe smectite, jarosite, Al smectite
09326	Mawrth Vallis	Fe smectite, bassanite, Al smectite
0A2C2	Mawrth Vallis	Fe smectite, jarosite, Al smectite
0A425	Mawrth Vallis	Fe smectite, jarosite, Al smectite
0C467	Mawrth Vallis	Fe smectite, alunite, Al smectite
173F4	Mawrth Vallis	Fe smectite, Al smectite
19AA0	Mawrth Vallis	Fe smectite, Al smectite
20BF9	Mawrth Vallis	Fe smectite, jarosite, Al smectite
21D02	Mawrth Vallis	Fe smectite, jarosite, alunite, Al smectite
0AD3D	Melas Chasma	Jarosite, monohydrated sulfate
13F5B	Melas Chasma	Jarosite, monohydrated sulfate
08F68	NE Syrtis	Mg smectite, Ca/Fe carbonate, illite, chlorite
19538	NE Syrtis	Akaganeite, Mg smectite, Mg CO ₃
19DAA	NE Syrtis	Mg smectite, akaganeite, Mg CO ₃
0B8C2	NE Syrtis	Jarosite, Mg CO ₃
0CBE5	NE Syrtis	Mg smectite, epidote
03E12	Nili Fossae	Mg smectite, Mg CO ₃ , Fe Olivine
03FB9	Nili Fossae	Mg smectite, Al smectite, Mg CO ₃
050F2	Nili Fossae	Mg smectite, prehnite, Ca/Fe carbonate, illite, chlorite
064D9	Nili Fossae	Mg smectite, low-Ca pyroxene
093BE	Nili Fossae	Mg smectite, hydrated silica, Mg CO ₃
09786	Nili Fossae	Prehnite, analcime, chlorite, high-Ca pyroxene
097E2	Nili Fossae	Mg smectite, kaolinite, hydrated silica, Mg CO ₃ , Fe olivine
09971	Nili Fossae	Mg smectite, jarosite, analcime, low-Ca pyroxene
0A053	Nili Fossae	Mg smectite, jarosite
0A09C	Nili Fossae	Mg smectite, Al smectite, Mg CO ₃
0A4FC	Nili Fossae	Mg smectite, Fe olivine
0AA03	Nili Fossae	Mg smectite, Al smectite, Mg CO ₃
0B438	Nili Fossae	Mg smectite, Mg CO ₃ , Mg Olivine
21C5A	Nili Fossae	Low-Ca pyroxene
0CA5C	North Polar area	Gypsum
0A546	Schroef crater	Ca/Fe carbonate, illite, chlorite
36F60	Sirenum Fossae	Kaolinite
3703F	Sirenum Fossae	Kaolinite
07E26	South pole	CO ₂ ice, H ₂ O ice
081CF	South Syrtis	Ca/Fe carbonate, illite

(continued on next page)

Table 1 (continued).

ID	Region	Mineral detections
14703	SW Syrtis	Ca/Fe carbonate, chlorite
02FC5	Syrtis Major	Prehnite, illite, chlorite
0454E	Syrtis Major	Prehnite, illite, chlorite
0B868	Syrtis Major	Prehnite, high-Ca pyroxene
0CBAC	Syrtis Major	Analcime, hydrated silica
0AB81	Terra Sirenum	Chloride, Mg smectite
0AA7D	Mawrth Vallis	Fe smectite, jarosite, alunite

^aThe list of detections reported for each image is not an exhaustive list of all detections. For most of the images only detections that were previously reported in the literature are included in the training set.

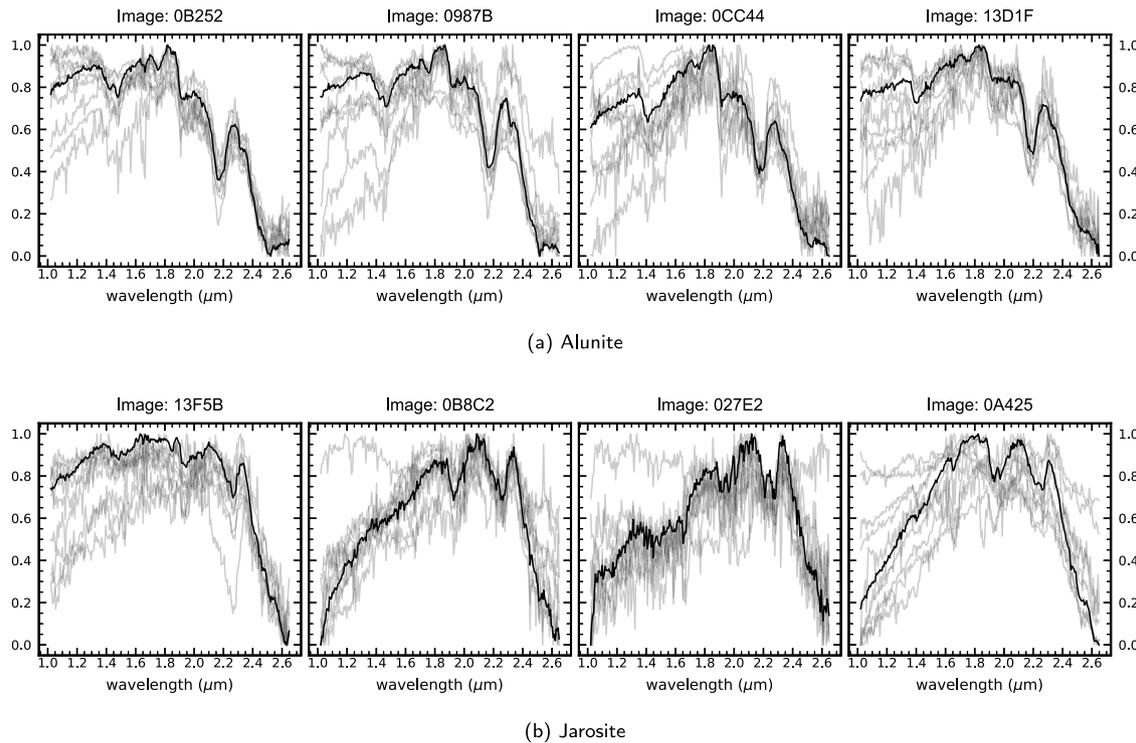


Fig. 1. Sample pixel spectra for alunite (a) and jarosite (b). Average spectra shown in solid black and spectra of individual pixels from each image are shown by gray. Both jarosite and alunite exhibit variation not only across different images but across different pixels of the same image as well.

4. Spectral ratioing

To compute the likelihood of individual pixels originating from a bland category the HBM model described in Section 3 is employed to estimate a distribution of bland pixels as part of the training phase. Unratioed training I/F data collected as described in Section 2 are used for this purpose. During the testing phase, this distribution is evaluated on pixels to generate likelihood scores, which in turn are used to identify denominator regions for use during column-wise ratioing. For a pixel of interest the denominator is obtained from the same column as the pixel of interest. The average spectrum of three pixels with the highest blandness score lying within a $\pm w$ row neighborhood of the pixel of interest is used as the denominator. These three pixels are selected based on the “relative blandness” score of the pixels in that column. As a result, even when there is no “truly bland” pixels, the scoring function can select pixels that have relatively less conspicuous spectral features compared to other pixels in the column. Although this may not constitute the ideal scenario for ratioing for reducing the nonlinear noise, as long as the spectral features in the denominator signal are weaker than those of the numerator, the ratioing may still serve for its intended purpose. The current implementation uses $w = 50$ as the default window size, but w is added as an external parameter to the ratioing algorithm as the optimum value often depends on the signal-to-noise ratio (SNR) and mineralogical composition of the image.

Restricting the search to a small window is important as the nonlinear noise component of pixels will be most similar when they are spatially closer. Using the average spectrum of pixels farther away from the pixel of interest as the denominator can produce large spikes in the ratioed spectrum and suppress key spectral features used for identification, especially in images with poor SNR.

One trivial approach to identify a column-specific bland signal for automated ratioing is to use the median of pixel spectra from each column in the image as the bland signal. In what follows we use the column-median approach as a baseline to evaluate HBM ratioing on four single-pixel mineral detections. The first pixel is obtained from a jarosite outcrop detected in Mawrth Vallis (0AA7D, $x = 201$, $y = 429$) (Ehlmann and Dundar, 2016). The second pixel is obtained from an Al clay outcrop detected in Jezero Crater (040FF, $x = 201$, $y = 428$) (Dundar et al., 2019b). The third pixel is obtained from an alunite outcrop recently detected in Eridania Basin (09E4C, $x = 451$, $y = 380$) (Leask et al., 2019). The fourth one is a pixel of Fe smectite (0A425, $x = 98$, $y = 93$), also used in the CRISM Type Spectra Library (Pelkey et al., 2007; Viviano-Beck et al., 2014). In the case of the jarosite pixel in 0AA7D, the column $x = 201$ contains a large number of Fe and Al smectite pixels in addition to jarosite. In the case of the kaolinite pixel in 040FF, the column $x = 286$ is mainly dominated by Mg carbonate pixels. In the case of the alunite pixel in 09E4C, the spatial extent of Fe smectite in column $x = 451$ is significantly

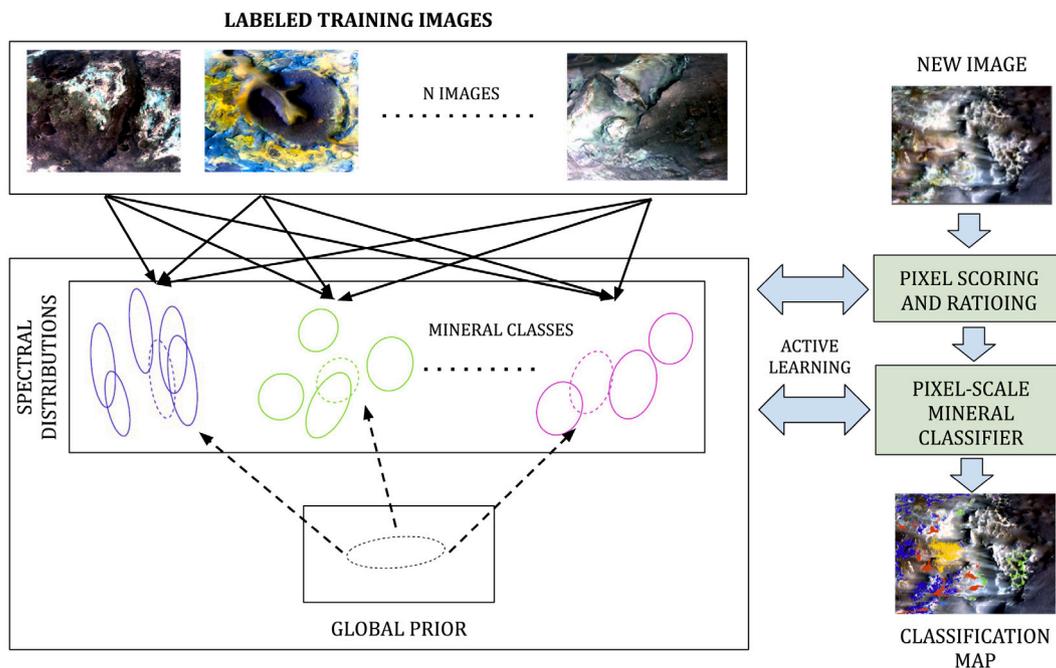


Fig. 2. Two-layer Bayesian Gaussian Mixture Model Training and Classification. Each training image contains detections from one or more mineral classes. Solid contours of the same color show the observed spectral distribution of the same mineral class detected across different images. Colored contours with dashed lines indicate unobserved spectral distributions of mineral classes from which all observed distributions assumed to originate. Black contour with dashed line is the global prior, which can be considered as the distribution of viable mineral spectral patterns. Training the model involves estimating the posterior predictive distribution (PPD) of unknown spectral distributions, i.e., colored dashed contours, for each mineral class using labeled pixel data in training images. A pixel in a new test image is classified by evaluating PPDs of each mineral class for this pixel and assigning the pixel to the mineral class that generates the maximum probability. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

larger than that of alunite. In the case of the Fe smectite pixel in 0A425, the column $x = 98$ is dominated by Al smectite pixels. The ratioed signals obtained by the column-median and HBM approaches on four pixels are shown in Fig. 3. The ratioed spectra obtained by HBM consistently enhances spectral features in all four cases making identification (even at pixel-scale) easier. The column-median approach generates deformed versions of spectra where key diagnostic features of these phases are distorted. In the cases of jarosite, kaolinite, and alunite pixels, the column-median generates a denominator with a higher albedo than the numerator. In the case of Fe smectite, the denominator generated by the column-median has lower albedo than the numerator but has strong spectral features characteristic of Al smectite, which lead to significant deformations during ratioing. For this specific pixel the unratioed spectrum already show strong spectral features of Fe smectite, which mostly disappear after ratioing by the column-median approach.

Overall this level of deformation in spectra may still be acceptable for manual analysis as the expert can still accurately identify mineral phases from distorted versions of spectral features. However, for automated analysis, significant distortions in the shape of spectra generate spatial discontinuities in the ratioed image causing automated classifiers to generate classification maps with large semantic gaps or, in more severe cases, scattered detections that appear as speckle noise. A machine-learning-based approach for ratioing can help eliminate these problems to a greater extent and help discover small outcrops of mineral phases even in images with extensive traditional analysis as has been demonstrated for the kaolinite outcrop in 040FF.

5. Mineral classification

An ensemble of $N = 15$ submodels is constructed to classify ratioed data at the pixel scale. Ensemble models mitigate noise in the data by using only a subset of the available data to construct each submodel (Breiman, 2001; Agjee et al., 2018). These subsets are created

by subsampling either the number of available samples, features or both. Training submodels with low correlation is the key to construct a diverse and robust ensemble. In our ensemble classifier, each submodel uses an HBM trained on the same set of training pixels but with channels selected from different ranges of the spectrum. Each submodel evaluates K class conditional probabilities of a pixel, where K is the number of mineral classes defined during training. These probability values are weighted differently for each class to account for the relative importance of each spectral range for each mineral. As a comparison a single model trained on the whole spectral range classifies pixels to specific classes with almost certainty, making false positive mitigation by probability-based thresholding less useful. Detailed results of this experiment is available in the supplementary document.

If we denote the probability computed by the submodel i for class j by p_{ij} and the weight assigned to the output of the submodel i for class j by w_{ij} , the ensemble probability for class j can be computed by $p_j = \sum_{i=1}^N w_{ij} p_{ij}$. Probabilities p_{ij} follows a multinomial distribution, i.e., $\sum_{i=1}^N p_{ij} = 1$, and weights assigned for each mineral class satisfies $\sum_{i=1}^N w_{ij} = 1$. Thus, these two constraints ensure that $p_j \leq 1$. The higher the p_j , the higher the confidence that the pixel is classified correctly. The ranges of wavelengths used in each submodel are shown in Table 4. The first four submodels use channels covering the entire spectrum from 1.0 μm to 2.65 μm , subsampled at every fourth channel. The sub-sampling rate was found through experimentation; increasing the number of channels makes the model slightly more susceptible to noise and other numerical issues. The remaining eleven models use ranges of wavelengths covering the key absorption features of phases detected on Mars. The weights w_{ij} assigned to each submodel/class pair is determined based on whether the spectral range the submodel is trained on contains a key absorption band of the mineral class, in which case a higher weight is assigned. For example as shown in Table 4, submodels 6, 10, and 11, which are trained on spectral ranges of 1.7–1.9 μm , 2.17–2.36 μm , and 2.23–2.43 μm , respectively, receive the highest weights for jarosite. Similarly, submodels 5 and 9, which

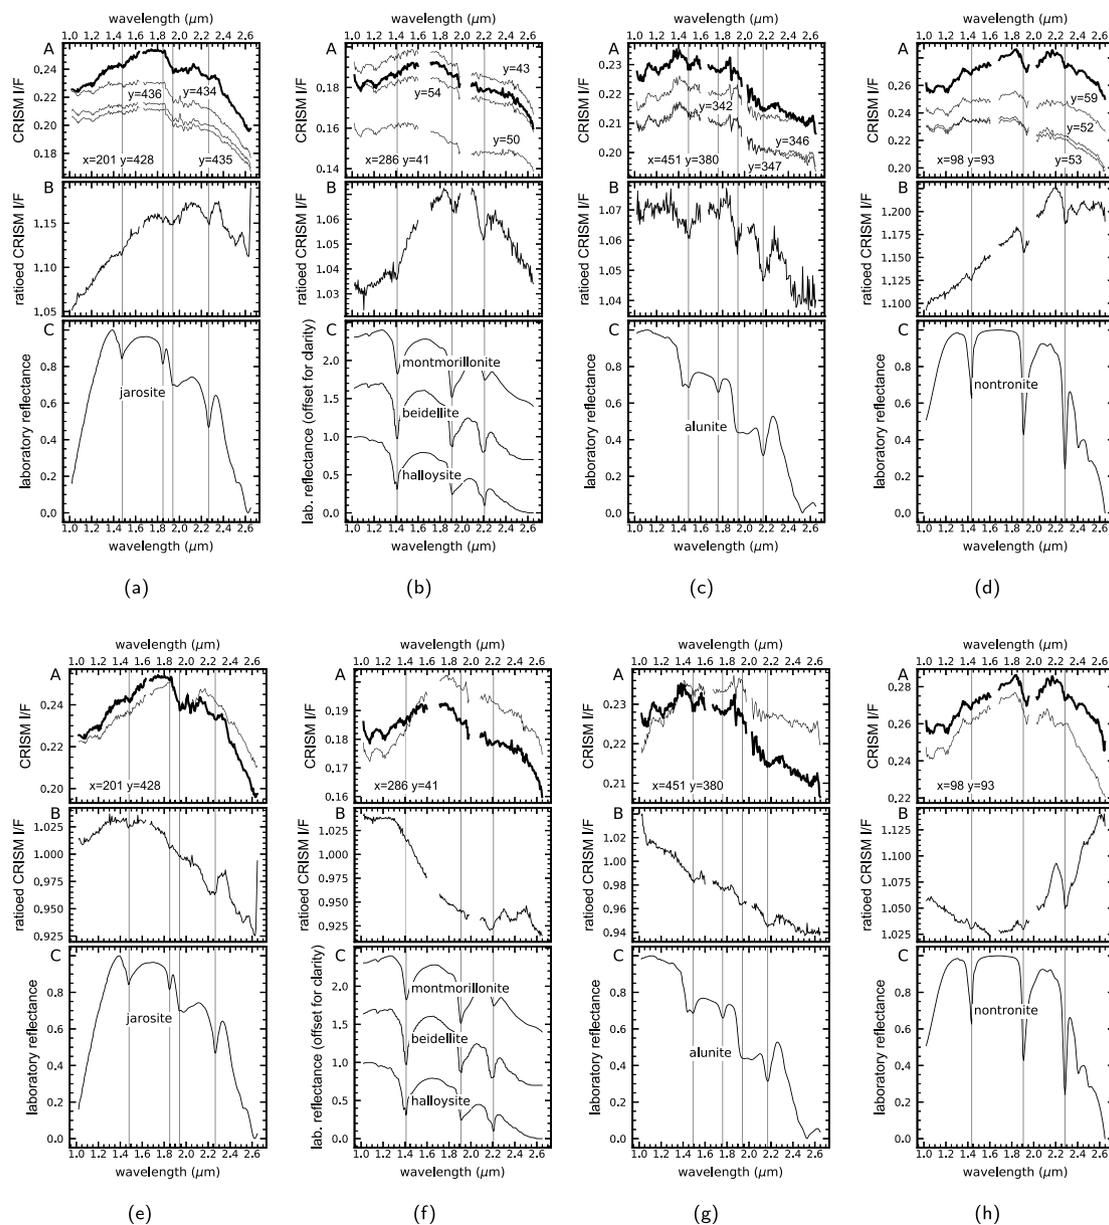


Fig. 3. Comparing ratioed spectra obtained by HBM (a–d) and column-median technique (e–h) on four pixels: jarosite in OAA7D (a, e), Al clay in O40FF (b, f), alunite in O9E4C (c, g) and Fe smectite in OAA25 (d, h). HBM automatically selects three pixels within ± 50 rows along-track the column of the numerator pixel and uses the average of these pixels as the denominator. In the top panel in figures (a)–(d) thick darker lines show the unratioed numerator pixel, thinner lines the unratioed denominator pixels. In the top panel in figures (e)–(h) thick darker lines show the unratioed numerator pixel, thinner line the column median. The median is computed using all the pixels along-track the column of the numerator pixel.

are trained on spectral ranges of 1.3–1.5 μm , 2.1–2.3 μm receives the highest weights for kaolinite. The entire weight matrix is provided as a spreadsheet file in the supplements. For every pixel, K probability values are computed, one for each mineral class, and the pixel is assigned to the class with the highest probability subject to additional filtering and validation as described next.

5.1. Spike removal

Spectra in the ratioed data are preprocessed to remove large spikes that may affect the robustness of the distribution estimated by the HBM model. Spikes are removed by a repeated application of a moving median filter with decreasing window sizes and replacing only the spikes that exceed the median by a threshold. In each pass, the difference between the original spectrum and the median-filtered version of it is computed, and the spike is only replaced by the median if the difference

is found to be larger than five standard deviations away from the mean difference. Both mean and standard deviations are computed over all the channels and averaged over all available pixel spectra. The choice of the window size for the median filter is critical, because repeated smoothing may remove spectral features and affect accuracy of the classifier. We choose to filter with windows of 11, 7 and 3 channels in this specific order. All plots generated in Section 6 use ratioed data without despiking or any other post processing applied.

5.2. False positive mitigation

The HBM model always predicts a mineral class for a pixel as it is designed in a closed-set classification setting (Cheng et al., 2019; Dundar et al., 2012). However, in some cases there could be new mineral phases in tested images that were not included in the training data, or the spectrum may be too noisy to be identified unambiguously.

Table 2

List of mineral phases in the training set. The first column on the left shows class labels associated with each mineral phase. The second column shows names of phases. The third column shows the number of labeled pixels for each phase. The last column shows the number of images from which the labeled data are collected.

Label	Mineral	# pixels	# images
1	CO ₂ Ice	17 297	1
2	H ₂ O Ice	5 947	2
3	Gypsum	32 163	2
4	Ferric Hydroxysulfate	1 195	1
5	Hematite	666	1
6	Fe smectite	77 556	14
7	Mg smectite	110 935	18
8	Prehnite	7 715	6
9	Jarosite	13 133	17
10	Serpentine	1 893	1
11	Alunite	10 323	17
12	Akaganeite	819	3
13	Ca/Fe CO ₃	8 113	6
14	Al smectite 1	311	4
15	Kaolinite	40 555	17
16	Bassanite	130	1
17	Epidote	189	1
18	Al smectite 2	8 781	12
19	Polyhydrated sulfate	47 675	10
23	Illite	8 613	8
25	Analcime	1 308	4
26	Monohydrated sulfate	26 441	8
27	Hydrated silica	21 946	5
29	Ferricopiapite	10 559	1
30	MgCO ₃	37 473	10
31	Chlorite	57 935	8
33	Low Ca Pyroxene	32 130	3
34	Mg Olivine	80	2
35	High Ca Pyroxene	849	2
36	Fe Olivine	2 098	4
37	Chloride	1 262	1
38	2.1 μm artifact	509	1
39	Bland	5 814	24

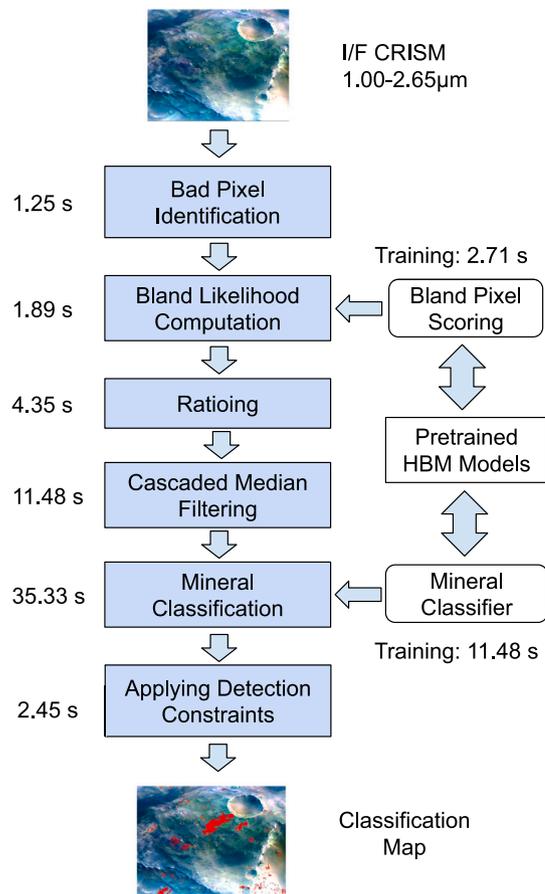


Fig. 4. Processing stages involved in the classification of an image. Numbers next to each block indicate processing time in seconds.

The latter is quite common in CRISM images as many rare phases have their distinctive spectral features within few channels of other common phases and they are often formed as part of alteration assemblages. In images with poor SNR, spatially co-occurring and spectrally mixed mineral phases may not be uniquely identified at the pixel scale. The classifier assigns such spectrally mixed pixels to the class with the highest probability, but since the total probability is shared between two or more spectrally similar classes these pixels are eventually assigned to their respective classes with low probabilities. A large number of low-confidence pixel detections can dominate the average spectrum and can have a smoothing effect on key spectral features or distort the spectrum expressed by a small number of high-confidence detections in other ways. Thus, removal of these low confidence pixel detections is crucial for more accurate identification of mineral phases.

As CRISM images exhibit different SNR levels and some mineral classes are more represented in the training set than others, a single probability threshold that works for all images and for all mineral classes may not be attainable. Pixels classified with higher probabilities produce average spectra with sharper features. However, as the probability threshold increases, the number of pixels used in computing the average decreases, and the average signal becomes noisier. For well-represented mineral classes the optimum probability threshold usually varies between 0.5 and 0.7 for good quality images and between 0.3 and 0.5 for poor quality images. We consider a mineral class well-represented if there are at least five instances of that mineral phase in the training set, where detection from each different image is considered an instance in this context. The number of instances of each mineral phase is shown in Table 2. If a mineral phase is not well-represented in the training set a probability threshold of 0.7 or higher might be needed irrespective of the image quality, to

reduce false positives. Using probability thresholds converts the closed-set classification setting into an open-set one where pixels with low confidence predictions are ignored as originating from unknown classes not represented during training.

Apart from probability thresholds, applying spatial constraints have proved to be very useful as a false positive mitigation strategy, especially for poor quality images. For such images an average spectrum computed over a large number of spatially-scattered, low-probability pixel detections can mimic the appearance of a spectral signature of a real mineral phase. To avoid such cases we seek spatial coherence in reported detections. Any real phase detection should have spatial continuity on the surface characterized by a group of contiguous or near contiguous pixel detections concentrated in certain parts of the image as opposed to being randomly distributed. More specifically, once the image is classified, we obtain a spatial map of pixel-scale detections and identify connected components. A connected component in an image is defined as a group of pixels from the same class connected to each other by 8-pixel connectivity. For low quality images this restriction is relaxed to have near contiguous pixels, which are pixels that are separated by 1–3 pixels. All connected components with less than a certain size or all pixels in the same column are considered less viable detections and are ignored from further processing. This is illustrated in Fig. 5 for pixel detections separated by one pixel and a minimum connected component size of five. In our approach probability thresholding and spatial constraints are combined to ensure that only the most viable detections are reported. In the next section we demonstrate the overall utility of the machine learning toolkit and the reliability of these specific constraints on three test images not used as part of training.

Table 3

Structure of the datasets. (a) Unratioed spectra from bland pixels. (b) Ratioed spectra from mineral classes.

(a) Variables and their descriptions in the unratioed bland training dataset. This data is used to estimate the probability density function of bland pixels.		
Name	Size	Description
pixspec	337 617 × 350	Unratioed spectra
im_names	340	List of CRISM image names, mapping them to numerical IDs
pixims	337 617	Numerical ID of the image the spectrum is from
pixcrds	337 617 × 2	(x, y) coordinates of the points in the original image
(b) Variables and their descriptions in the ratioed mineral dataset. This data is used to train the HBM model for mineral classification.		
Name	Size	Description
pixspec	592 413 × 350	Ratioed spectra
pixlabs	592 413	Mineral labels
im_names	77	List of CRISM image names, mapping them to numerical IDs
pixims	592 413	Numerical ID of the image the spectrum is from
pixpat	592 413	ID of the connected patch in the image the pixel belongs to
pixcrds	592 413 × 2	(x, y) coordinates of pixels in their respective image

Table 4

Spectral ranges employed by different HBM models in the ensemble classifier for two of the minerals classified. The first four submodels use the entire range of the spectrum subsampled at every fourth channel. The remaining eleven submodels use specific subranges associated with key spectral features of mineral phases commonly found on Mars.

Submodels	Spectral ranges	Number of channels	Weights	
			Jarosite	Kaolinite
1	1.041:0.026:1.618 μm	55	0.05	0.05
	1.737:0.026:1.948 μm			
	2.067:0.026:2.648 μm			
2	1.047:0.026:1.625 μm	54	0.05	0.05
	1.744:0.026:1.954 μm			
	2.073:0.026:2.629 μm			
3	1.054:0.026:1.632 μm	54	0.05	0.05
	1.75:0.026:1.961 μm			
	2.08:0.026:2.635 μm			
4	1.06:0.026:1.638 μm	54	0.05	0.05
	1.757:0.026:1.968 μm			
	2.086:0.026:2.642 μm			
5	1.310–1.507 μm	31	0.05	0.20
6	1.704–1.902 μm	31	0.20	0.05
7	1.770–1.968 μm	31	0.10	0.05
8	2.034–2.232 μm	31	0.05	0.05
9	2.100–2.298 μm	31	0.05	0.30
10	2.166–2.364 μm	31	0.15	0.05
11	2.232–2.430 μm	31	0.15	0.05
12	2.298–2.497 μm	31	0.05	0.05
13	2.364–2.563 μm	31	0.00	0.00
14	2.430–2.629 μm	31	0.00	0.00
15	1.408–1.605 μm	31	0.00	0.00

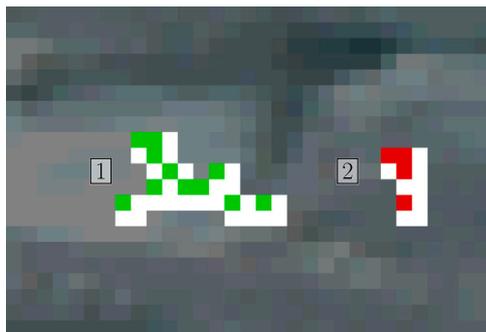


Fig. 5. Selection of connected components. The detected regions are dilated by one pixel on the right and on the bottom (white regions). Region 1 contains more than 5 detected pixels (11, in green) and thus it is selected; region 2 contains less than 5 detected pixels (4, in red) and thus it is discarded. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

6. Validation on test images

6.1. Offline training

During the offline training phase, PPDs are estimated for ratioing and mineral classification using all images in their respective training datasets. For ratioing, unratioed spectra from bland pixels are used to estimate the PPD of bland pixels. As discussed in Section 3 and the supplementary document, this PPD is obtained in the form of a *Student-t* distribution and thus it requires estimating the degree of freedom, location vector, and the scale matrix. The total computation time for estimating these three parameters is 3 s. For the mineral classifier, 15 PPDs are estimated, one for each submodel, using channels from different spectral ranges of the ratioed training data (see Table 4) collected from 39 classes (see Table 2). The total computation time for estimating these PPDs is 11 seconds. All processing is performed on a single core of an Intel i9-9900 3.1 GHz CPU.

6.2. Classification of new images

A new image is processed according to the stages outlined in Fig. 4. First, bad pixels are identified and removed. Second, the HBM model

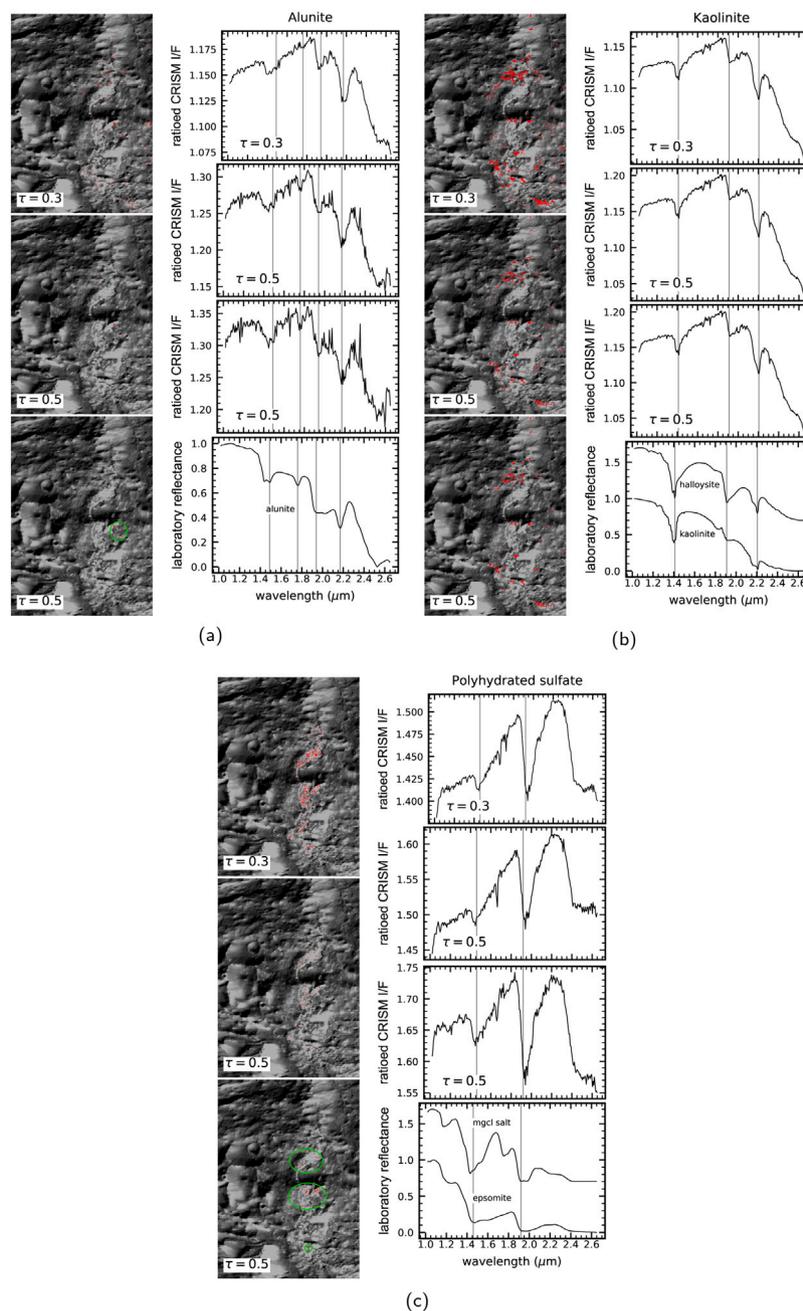


Fig. 6. Phases detected by HBM in HRL00016CFE. (a) alunite, (b) kaolinite and (c) polyhydrated sulfate. The classification maps on the left of each figure show the detected pixels for a probability threshold τ of 0.3 (top), 0.5 (middle), and after spatial constraints are applied (bottom). The top three plots on the right of each figure show the average spectra of detected pixels corresponding to each map on the left. The bottom plot shows the reference spectra of the phases most closely matching the detected phases.

pretrained on unratioed spectra from bland pixels is used to evaluate the likelihood of a bland spectrum for each pixel. Third, these likelihood scores are used to ratio each and every pixel in the image as described in Section 4. Fourth, ratioed spectra are processed to replace large spikes by a one dimensional cascaded median filtering applied at decreasing window sizes. Fifth, mineral classification is rendered on the ratioed pixel-scale spectra and each pixel is assigned to the mineral class whose distribution generates the maximum probability for that pixel. The classifier returns the predicted class and the probability of the prediction for each pixel. Finally, probability thresholds and spatial constraints are applied, and average spectra for detected phases and their respective binary classification maps are produced. The total processing time for a full resolution FRT image is around one minute. The toolkit uses targeted, reduced data record (TRDR) images and reports pixel x , y coordinates to facilitate reproducibility for

investigators seeking to replicate the findings presented here. Users can convert binary classification maps into a map-projected product for use in GIS programs by converting x , y coordinates in classification maps to latitude and longitude by accessing the *latitude* and *longitude* bands for those x , y pixels in the ancillary derived data record file (DDR), which accompanies each CRISM image.

6.3. Validation

In this section we use three images not used as part of the training set to demonstrate the usefulness of the developed machine learning processing pipeline. For all images we use the same set of probability thresholds and spatial constraints. More specifically, for mineral phases represented with fewer than five instances in the training set we use

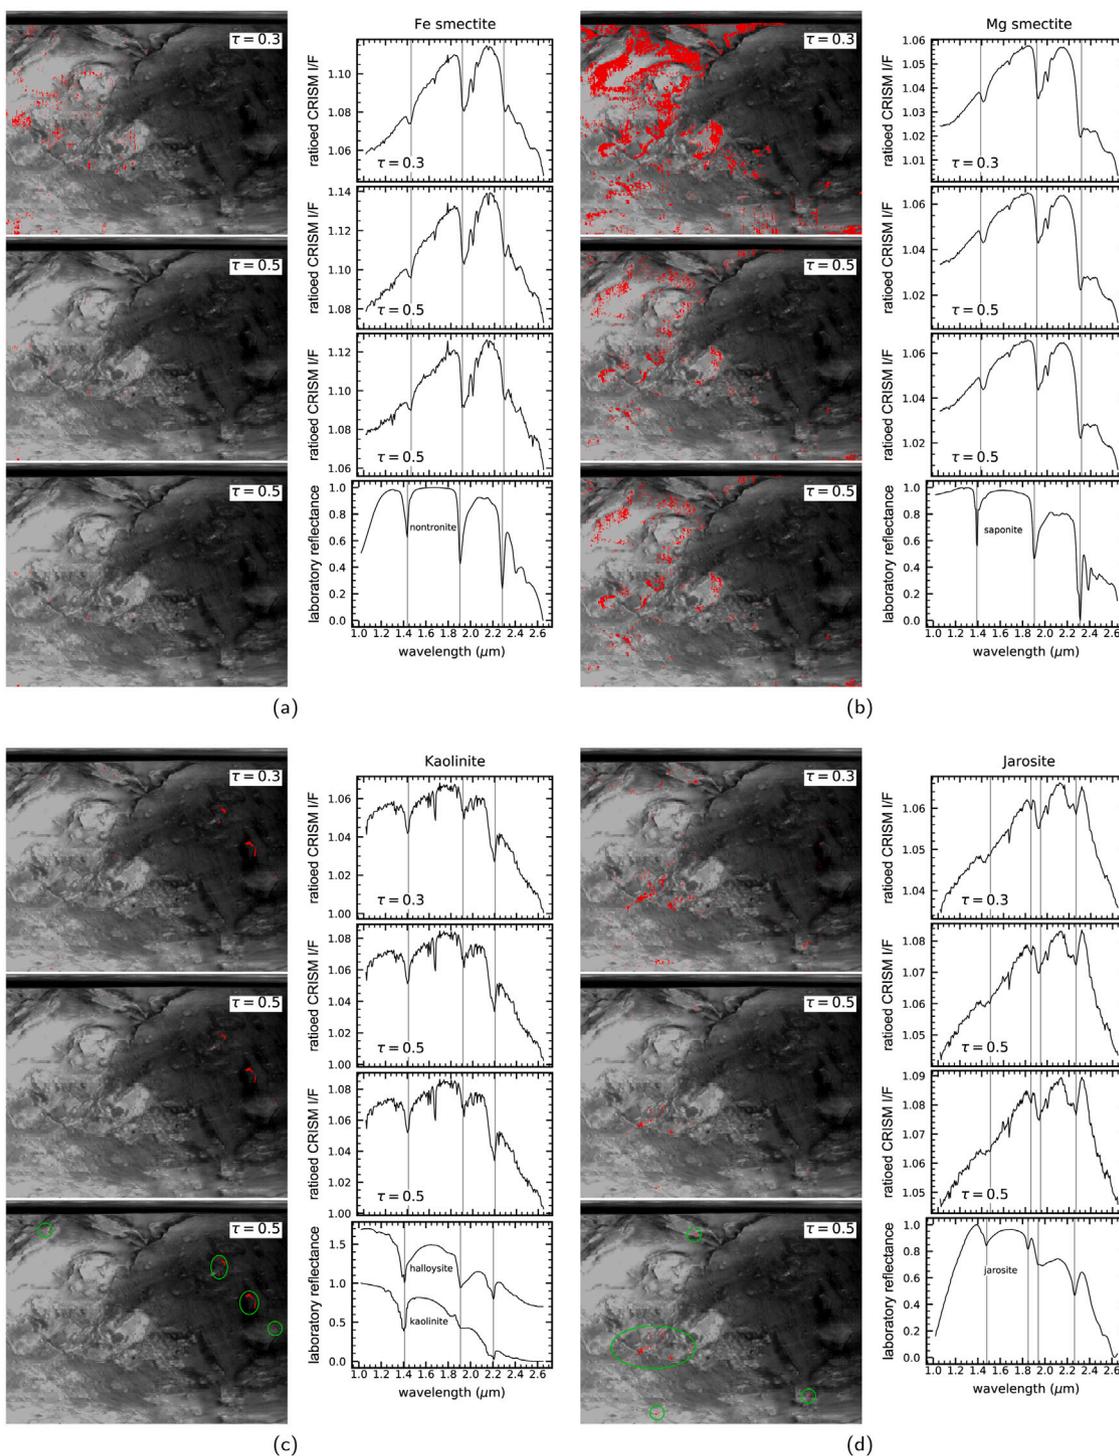


Fig. 7. Phases detected by HBM in FRT0001FD76. (a) Fe smectite, (b) Mg Smectite, (c) kaolinite, and (d) jarosite. The classification maps on the left of each figure show the detected pixels for a probability threshold τ of 0.3 (top), 0.5 (middle), and after spatial constraints are applied (bottom). The top three plots on the right of each figure show the average spectra of detected pixels corresponding to each map on the left. The bottom plot shows the reference spectra of the phases most closely matching the detected phases.

a probability threshold of 0.7 and for those represented with five or more instances we use a probability threshold of 0.5. Connected components are identified after dilating the binary classification map by one pixel to the right and to the bottom to group nearby pixels in the same connected component (see Fig. 5). Only connected components containing more than five detected pixels are considered viable and all other detections are ignored. The pixels added by the dilation (shown as white pixels in Fig. 5) are not counted toward the size threshold.

The first image is HRL00016CFE, taken from the east crater rim of Columbus Crater. Although alunite was previously reported in Columbus crater, it was detected on the west crater wall and crater floor (Ehlmann et al., 2016; Chaves et al., 2018). In this image our algorithm detects small outcrops of alunite and polyhydrated sulfate along with kaolinite. Segmentation maps and average spectra for all three phases before and after applying probability and spatial constraints are shown in Fig. 6. In images where kaolinite and alunite cooccur, distinguishing

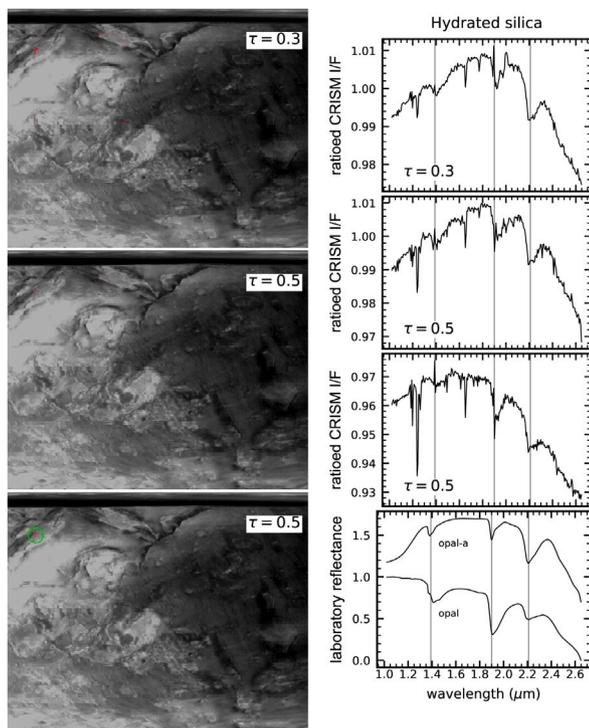


Fig. 8. Hydrated silica detected by HBM in FRT0001FD76. The classification maps on the left show detected pixels for a probability threshold τ of 0.3 (top), 0.5 (middle), and after spatial constraints are applied (bottom). The top three plots on the right show the average spectra of detected pixels corresponding to each map on the left. The bottom plot shows the reference spectra of the phases most closely matching the detected phases.

alunite from kaolinite becomes challenging because of the spectral similarity between the two. The validation with this image demonstrates that our algorithm can effectively detect subtle spectral differences between spatially cooccurring, spectrally similar phases such as the $1.75 \mu\text{m}$ feature in alunite .

The second image is FRT0001FD76, taken from the central peak of a crater in Nili Fossae. Our algorithm detects Fe/Mg smectite, kaolinite, jarosite, pyroxene, and hydrated silica. Fe/Mg smectites were first reported in [Ehlmann et al. \(2009\)](#), albeit in a different image (FRT0000A053). Kaolinite and pyroxene are visible in the 3-channel image derived from summary parameters BD2300, BD2210 and BD1900. Jarosite in this crater was first detected by our algorithm in image FRT0000A053 and was reported in [Ehlmann and Dundar \(2016\)](#). The small outcrop of hydrated silica detected in this image may require additional vetting but the broad feature at $2.2 \mu\text{m}$ suggests a phase different than kaolinite detected in this image. Segmentation maps and average spectra for all detected phases before and after applying probability and spatial constraints are shown in [Figs. 7 and 8](#).

The third image is FRT00009036, taken from the southern rim of the Huygens crater. In this image our algorithm detects small outcrops of prehnite, illite, and chlorite in addition to olivine. The 3-channel image derived from summary parameters BD2300, BD2210 and BD1900 shows an Al phyllosilicate phase. Our algorithm more precisely detects that phase as illite and prehnite with relatively well-defined spectral features. In addition, our algorithm detects two small outcrops as chlorite, one of which has a conspicuous $2.25 \mu\text{m}$ feature that appears at higher thresholds in addition to the $2.35 \mu\text{m}$ one, which may indeed suggest a chlorite phase. The other outcrop has only the $2.35 \mu\text{m}$ feature, which may suggest chlorite or Ca/Fe carbonate. A key point is that the algorithm can track and classify across images any characteristic recurring spectral shape, even if the expert user cannot

identify it. In this case, whether the phase we are tracking is a chlorite or an Ca/Fe carbonate is ambiguous. This actually highlights a key advantage of the machine learning approach: the unknown spectral type can be tracked across Mars and the geographic distribution and geologic settings of occurrence used to determine which phase(s) are most plausible. Segmentation maps and average spectra are shown in [Fig. 9](#).

7. Comparison with deep convolutional neural networks

Deep learning models pretrained on millions of images and vast of amount of text and speech data have radically changed the way we interact with data over the last few years. The field of hyperspectral image analysis is no exception to this change as deep convolutional neural networks (CNNs) that can jointly process spectral and spatial dimensions offer a full-fledged solution to many classification and clustering problems in hyperspectral image analysis ([Li et al., 2017](#); [Mou et al., 2017](#); [Min et al., 2018](#)). Although some issues, such as class imbalance, outliers, and missing classes are yet to be addressed, deep CNNs are considered the state-of-the-art approach in hyperspectral image classification ([Audebert et al., 2019](#)). To test the feasibility of these deep learning approaches for CRISM image analysis we have evaluated several recently proposed deep learning techniques implemented by the *DeepHyperX* toolbox ([Audebert et al., 2019](#)) on a subset of our dataset containing 14 images for training and 11 for testing. Based on this preliminary evaluation we have selected the network architecture by [Ben Hamida et al. \(2018\)](#) as the model with the most promising results and put this to further testing in the same way we have tested the proposed HBM. The model is a 3D convolutional neural network based on $3 \times 3 \times 3$ kernels and takes 5×5 image patches as input to classify pixels. The network architecture is built from two 3D CNN layers (two dimensions for the spatial image and one dimension for the spectral data), each followed by a 1D CNN layer on the channel dimension, two 1D CNN layers, and a final dense layer (each node in the layer receives input from all nodes in the previous layer) for classification. The model is applied in a sliding window manner on the input image to generate pixel-level class predictions, and a special “unknown” class is reserved for pixels without labels. More technical details can be found in [Ben Hamida et al. \(2018\)](#).

The model is trained with the full set of 77 I/F images in [Table 1](#), each ratioed following the same steps outlined in [Section 4](#). Only labeled patches are used for training and a validation set is obtained by sampling 5% of the patches stratified according to class sizes. The training is performed for 50 epochs using the training parameters from the original paper. The training time was 1 h 26 min and the final validation accuracy, weighted by class size, was 94%. Detections are filtered using the same post-processing steps applied for the HBM model (see [Section 5.2](#)). As deep learning algorithms in general produce overconfident predictions, i.e., most probabilities tend to be close to 1, the probability thresholds for well represented and under-represented images are adjusted differently for each image to have at most 300 high-confidence connected component detections per image. Here we briefly evaluate the classification results.

In image HRL00016CFE CNN detects Fe/Mg smectite, jarosite, alunite, kaolinite, polyhydrated sulfate, kieserite, and low Ca pyroxene ([Fig. 10](#)). Out of these detections only alunite, kaolinite and polyhydrated sulfate are correct. However, the spectra for alunite and polyhydrated sulfate are visibly noisier than those obtained by HBM, a sign that the areas detected by HBM are under-segmented by CNN. The kaolinite detection is oversegmented and shows vertical patterns that indicate a dependence of the classifier on column-specific noise. The other detections are all considered false positives. We have investigated all the connected components from these detections to see if there were any connected components that were correctly classified, yet missed by the HBM model, but there were none. The average spectra for all the false positives were very noisy with no recognizable spectral features.

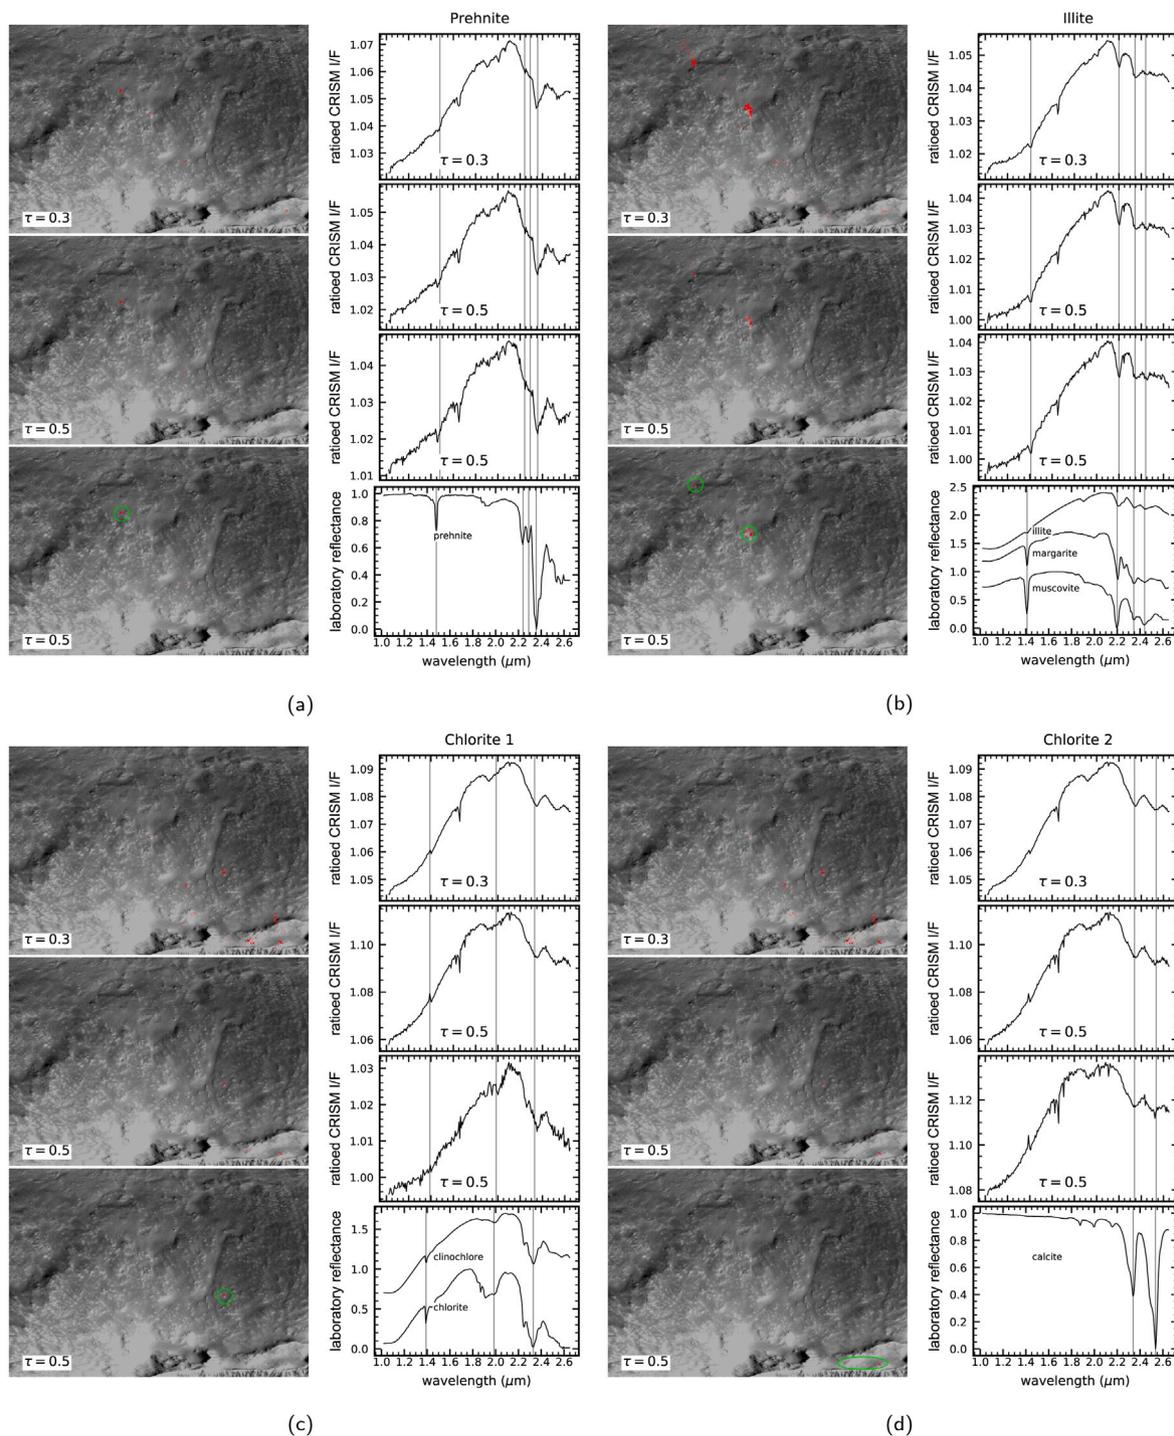


Fig. 9. Phases detected by HBM in FRT00009036. (a) prehnite, (b) illite, (c) chlorite 1, and (d) chlorite 2. The classification maps on the left of each figure show detected pixels for a probability threshold τ of 0.3 (top), 0.5 (middle), and after spatial constraints are applied (bottom). The top three plots on the right of each figure show the average spectra of detected pixels corresponding to each map on the left. The bottom plot shows the reference spectra of the phases most closely matching the detected phases.

In image FRT0001FD76 CNN detects Fe/Mg smectite, jarosite, kaolinite, monohydrated sulfate, Mg CO_3 , and pyroxene (Fig. 11). Out of these detections Fe/Mg smectite, jarosite, kaolinite, and pyroxene are correct while the monohydrated sulfate and Mg CO_3 are false positives. However, Fe/Mg smectite detections contained a large number of pyroxene outcrops. Jarosite contained several connected components with nondiagnostic spectra causing the key $1.85 \mu\text{m}$ in jarosite spectrum to vanish. CNN does not detect any hydrated silica in this image, which was detected by HBM. In image FRT00009036 CNN detects Fe/Mg smectite, prehnite, Ca/Fe carbonate, polyhydrated sulfate, illite,

hydrated silica, ferricopiapite, and chlorite (Fig. 12). Out of these detections prehnite, illite, Ca/Fe carbonate, and chlorite detections can be considered correct with some reservations. Ca/Fe and chlorite detections contain several connected components with nondiagnostic spectra causing the $2.25 \mu\text{m}$ feature in chlorite to disappear. CNN is also repeating the HBM mistake of mixing chlorite and Ca/Fe carbonate. Illite detections exhibit sharp spectral features extending beyond areas detected by HBM suggesting HBM misses some of the illite outcrops in this image at high probability thresholds. Remaining detections that involve Fe/Mg smectites, polyhydrated sulfate, hydrated silica, and

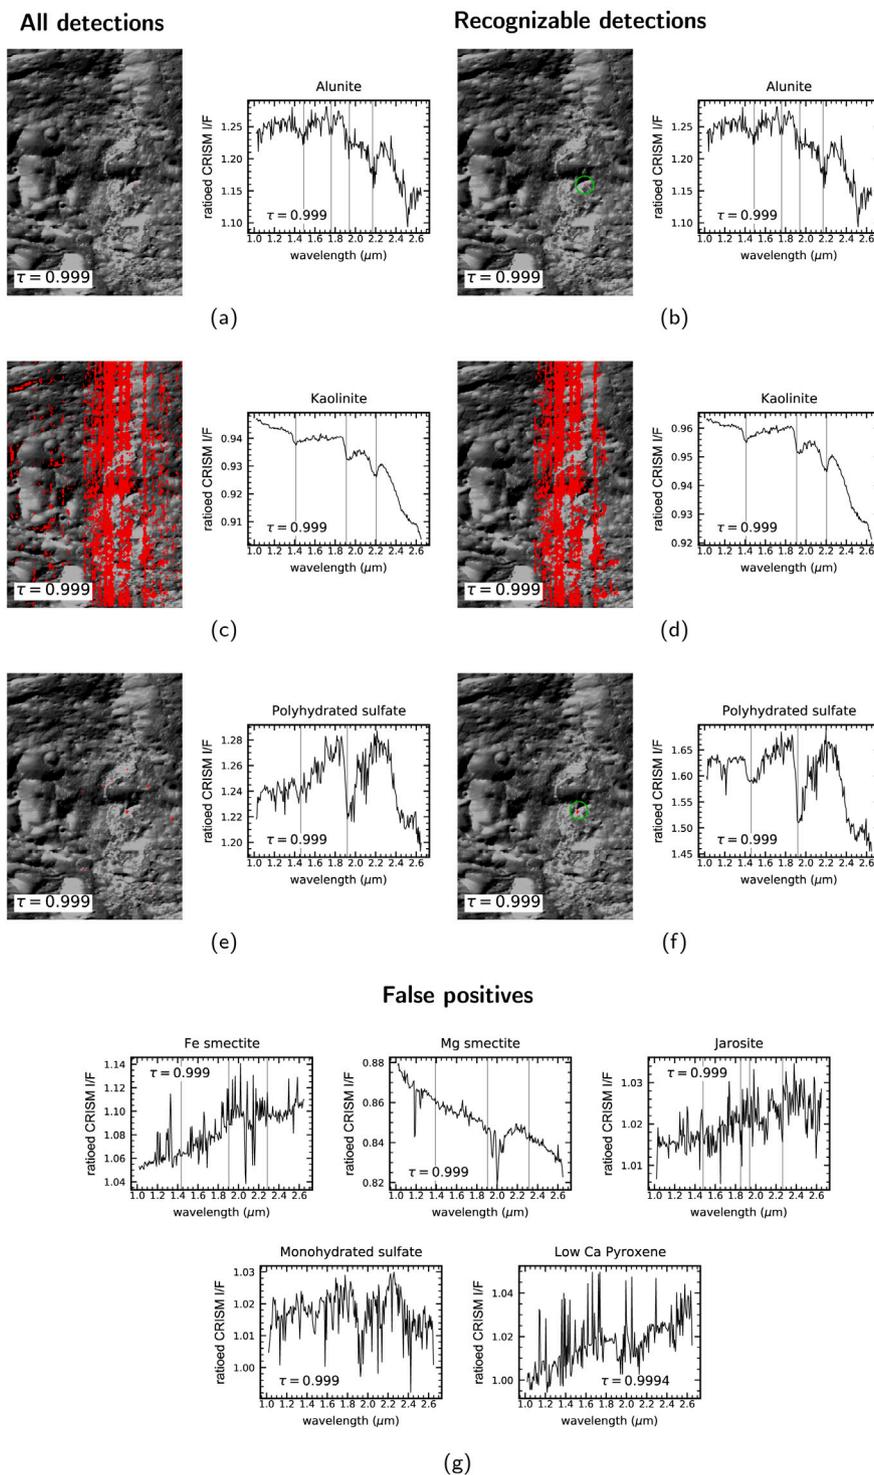


Fig. 10. Phases detected by CNN in HRL00016CFE. Figures (a), (c), and (e) on the left panel show all the connected components detected by the model and the matching average spectra for true positives. Figures (b), (d), and (f) on the right panel show the subset of connected components with recognizable spectra for true positives. Figures in the bottom panel (g) show the average spectra for false positive detections. Probability threshold is denoted by τ .

ferricopiapite exhibit flat spectra with no recognizable features, and are all false positives.

Overall CNN classifier significantly underperforms HBM classifier in terms of accuracy and specificity of detections. As CNN produces overconfident predictions even for noisy and spectrally mixed patterns not represented in training, thresholding the probabilities was not as useful as HBM in reducing the number of false positive detections. The ensemble aspect of HBM that weights outputs of submodels differently

for each mineral class seem to be effective in generating more balanced probability values and thus significantly reducing the number of false positive detections.

8. Limitations

Although our training set contains pixel-scale labeled spectra from 77 well characterized images, we do not believe the current set of

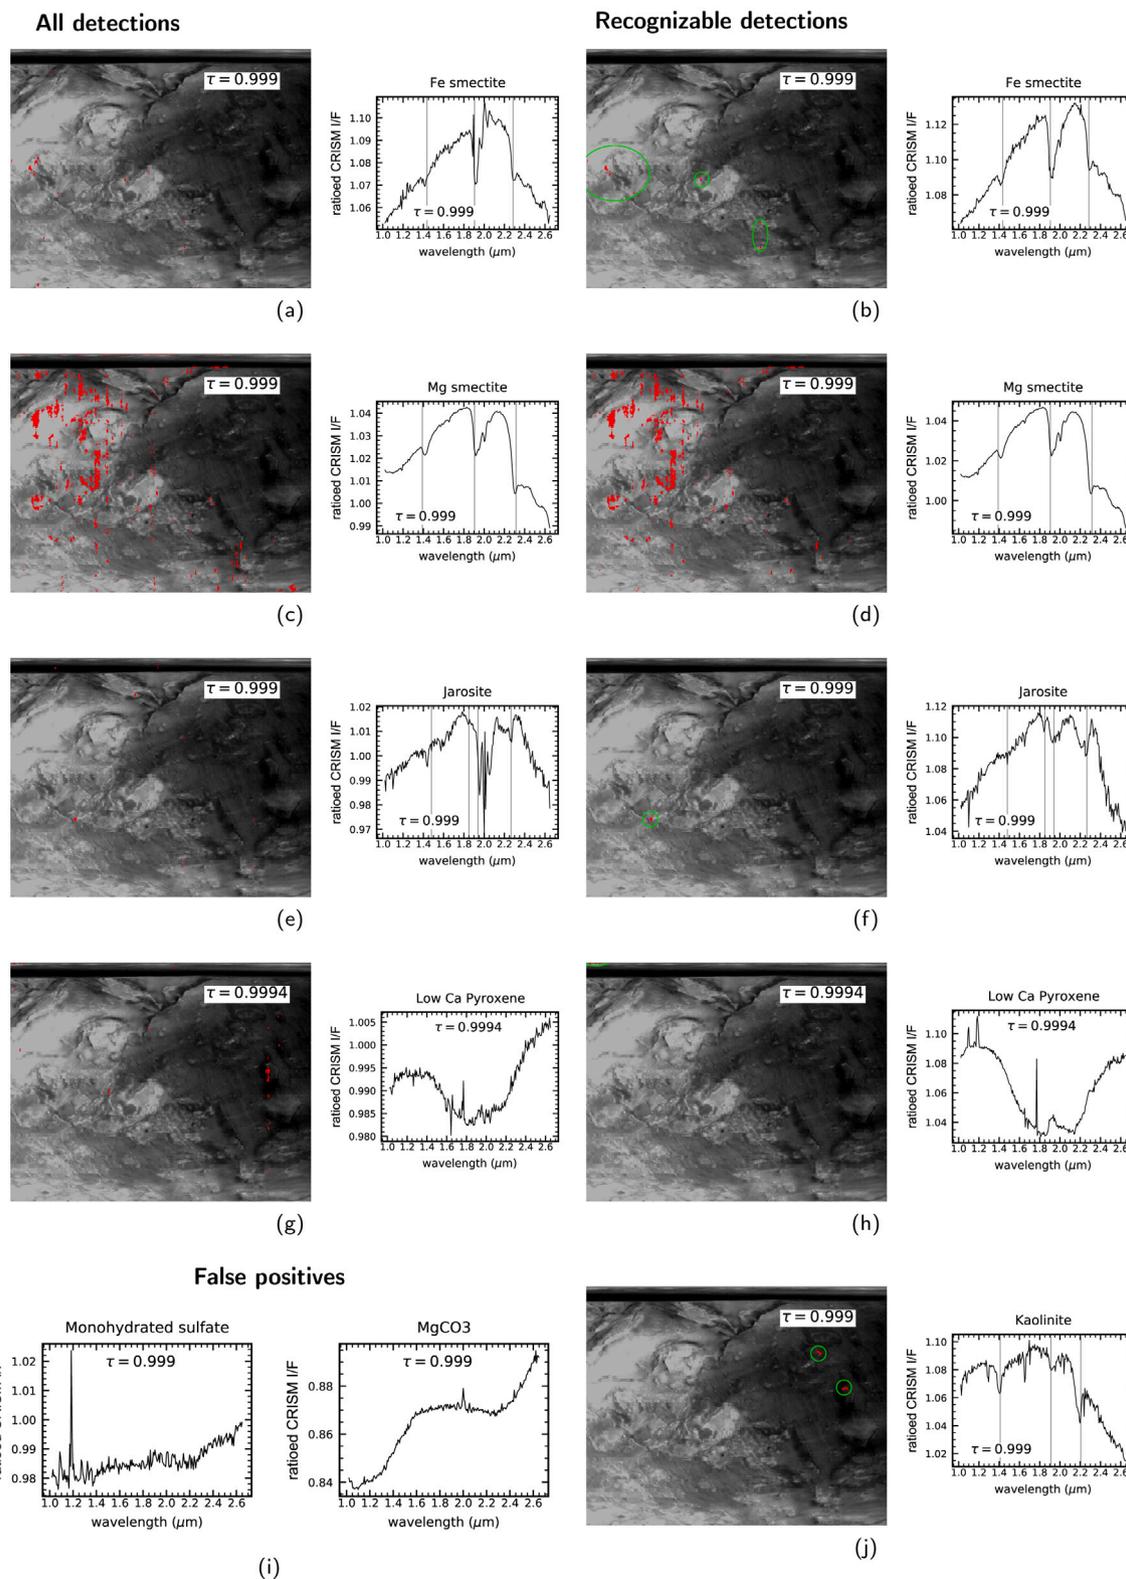


Fig. 11. Phases detected by CNN in FRT0001FD76. Figures (a), (c), (e), and (g) on the left panel show all the connected components detected by the model and their matching average spectra for true positives. Figures (b), (d), (f) and (h) on the right panel show the subset of connected components with recognizable spectra for true positives. Figures in the bottom left panel (i) show the average spectra for false positive detections. We report kaolinite once in (j) as all detected regions are valid. Probability threshold is denoted by τ .

training images exhaustively represents all spectral patterns. Certain spectral patterns detected in new images processed but not represented in the training set will be incorrectly classified to a class that is most similar to these patterns. One example of this is the 2.1 μm spectral

artifact that we reported earlier in Leask et al. (2018) and initially misclassified as a monohydrated sulfate and then identified by expert analysis as perchlorate. We could only determine that it was an overlooked CRISM pipeline processing artifact with spatial coherence after

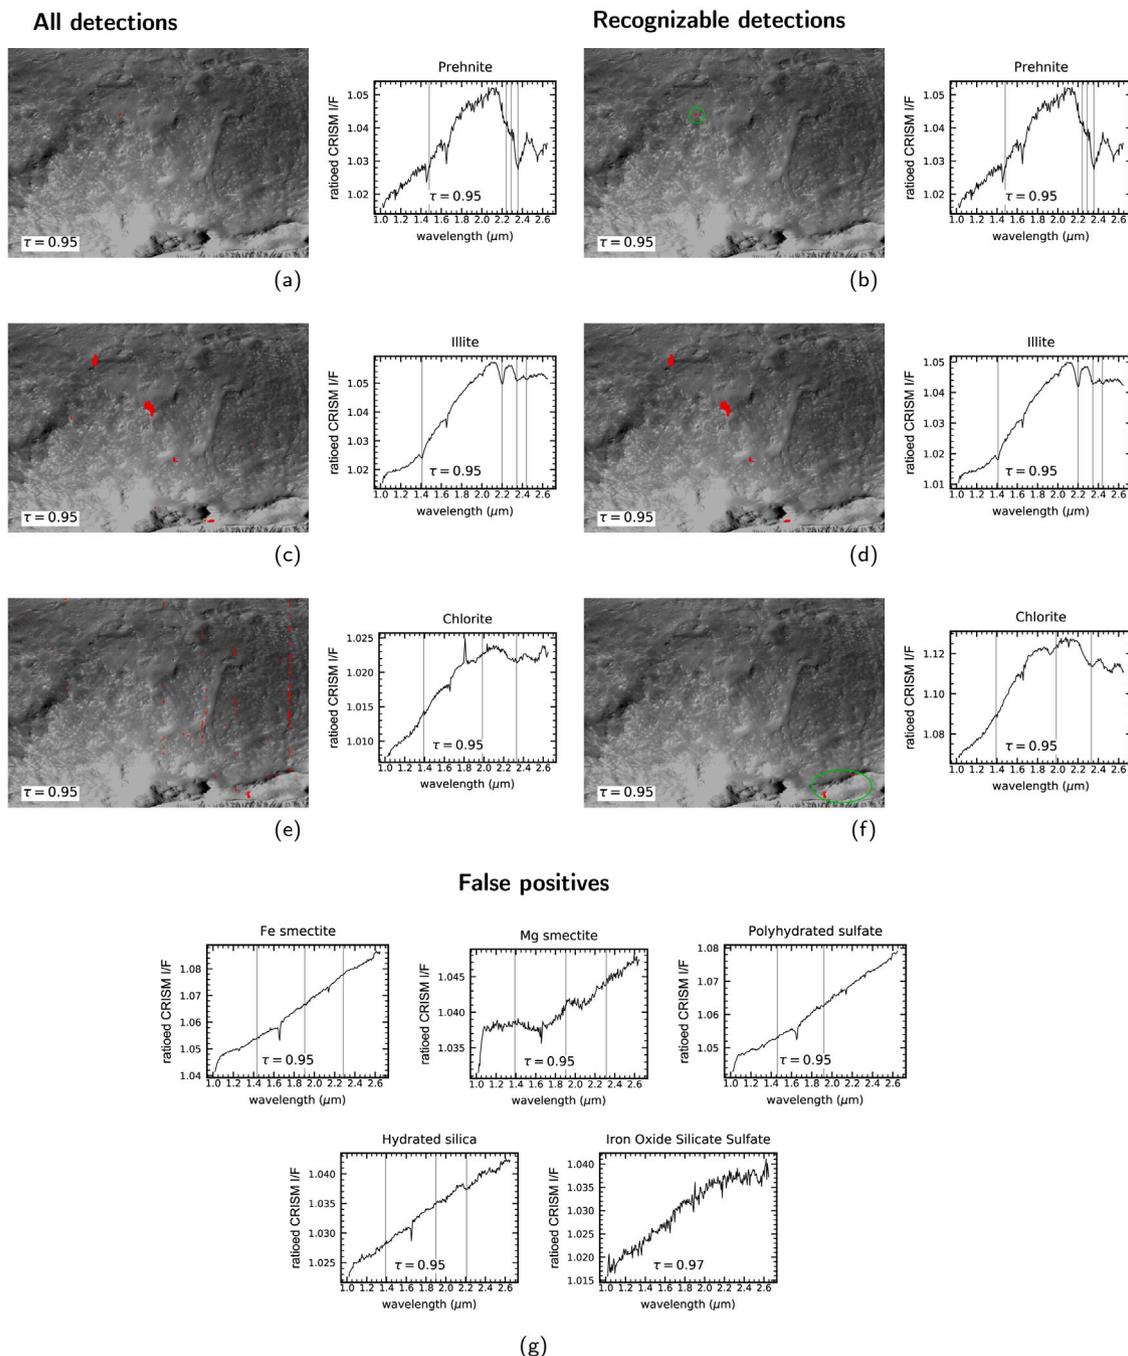


Fig. 12. Phases detected by CNN in FRT00009036. Figures (a), (c), and (e) on the left panel show all the connected components detected by the model and their matching average spectra for true positives. Figures (b), (d), and (f) on the right panel show the subset of connected components with recognizable spectra for true positives. Figures in the bottom panel (g) show the average spectra for false positive detections. Probability threshold is denoted by τ .

extensive manual vetting of the detections associated with this pattern. A new class representing this artifact has been added to the training set, so that when new instances of this pattern are detected in new images they are not classified as sulfate but correctly identified as an artifact. As the I/F data has an artifact in the 2.1 μm range, we recommend to confirm the 2.1 μm feature with radiance data in small outcrops.

There could be additional spectral patterns not represented by the current training set. For example, there is a vaguely characterized pattern with a doublet 2.2 μm feature widely detected in the Mawrth Vallis and Valles Marineris regions that is reminiscent of multiple different phases including jarosite, hydrated silica, and Al/Fe smectite, previously identified by [Roach et al. \(2010\)](#) (see [Fig. 13](#) for sample patterns). Our current version of the classifier attributes these detections to

classes associated with jarosite, hydrated silica, or Al smectite, as this pattern is not included in the current version of the training set because of the ambiguity surrounding it. Our training set may become more representative and the classifier trained with this data may become more robust over time once these type of ambiguous spectral patterns can be better characterized and unequivocally identified and incorporated into the training set of classes.

Several phases are represented in the training set only by a few instances. The current version of the classifier generates many false positives for some of these phases when operating at low probability thresholds, and thus a probability threshold of 0.7 and higher is recommended for these phases. Viable pixel detections will almost always have spatial continuity, and as a general rule of thumb, the spatial

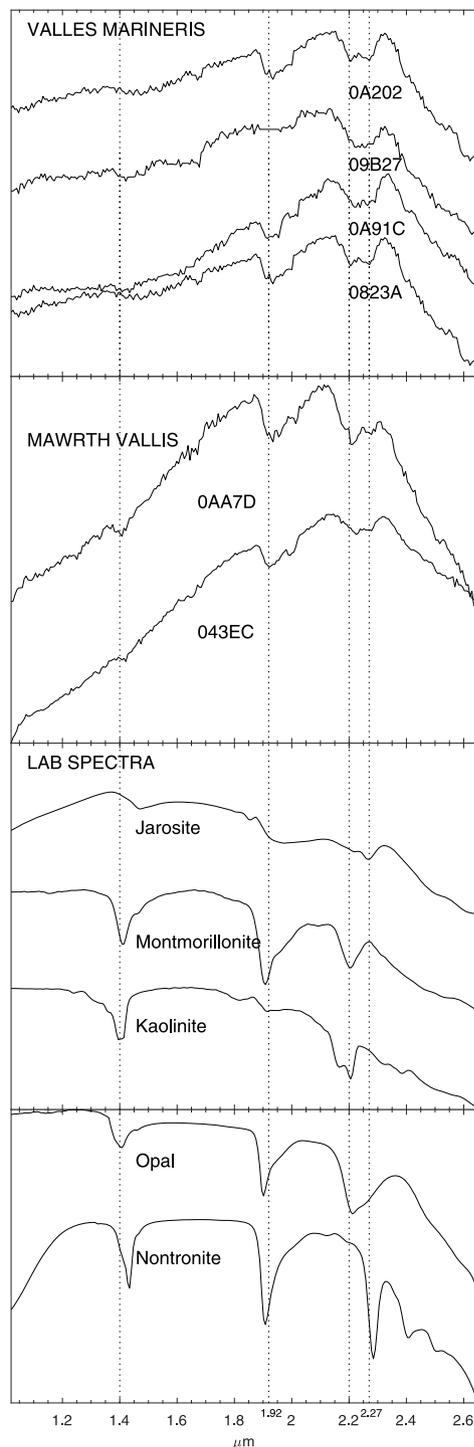


Fig. 13. CRISM detections with the doublet feature from Valles Marineris and Mawrth Vallis.

extent of viable detections grows in a continuous way as the probability threshold is gradually lowered.

The current training set includes classes with no diagnostic features in the 1.00–2.65 μm range, such as hematite and chloride. Pixels with no diagnostic features in this range, but with higher albedo compared to bland pixels, are usually classified into one of these classes. Thus, detections in these classes should only be used as a starting point for additional analyzes to confirm phases with key diagnostic features in the 0.36–1.00 μm (hematite) and 2.65–3.90 μm (chloride) ranges.

9. Conclusions

We present a new machine learning toolkit built on top of a versatile two-layer Bayesian Gaussian mixture model designed to model spectral variability within as well as between images. Using published knowledge of detected mineral phases on Mars we have collected a relatively large training dataset and demonstrated the utility of this toolkit in removing nonlinear noise and detecting both high and low abundance mineral phases in CRISM images. We believe this new toolkit will complement traditional spectral analysis techniques and will improve community's ability for fast, precise, and scalable analysis of remotely-sensed data. As more images are processed and the training set is augmented with new detections, especially from underrepresented mineral classes, the accuracy of the classifier will gradually improve. Although the current version of the classifier is implemented in a closed-set setting, training with data only from phases detected to date, open-set version can be implemented by extending the training dataset with synthesized spectra of potentially missing phases. This could potentially open new avenues for machine learning research toward discovering new phases.

Data availability

Training datasets and training segmentation maps can be downloaded from <https://cs.iupui.edu/~mdundar/CRISM.htm>. The toolkit with a user friendly Python notebook, source codes, and documentation is hosted at https://github.com/Banus/crism_ml.

Acknowledgments

Thanks to the CRISM science and operations teams for their work to collect and process these datasets. Thanks to Abigail A. Fraeman for her feedback on a subset of detections. B.L.E and M.D. were sponsored by NASA, USA under Grant Number 80NSSC19K1594. M.D. was sponsored by the National Science Foundation (NSF), USA under Grant Number IIS-1252648 (CAREER).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.icarus.2021.114849>.

References

- Agjee, N.H., Mutanga, O., Peerbhay, K., Ismail, R., 2018. The impact of simulated spectral noise on random forest and oblique random forest classification performance. *J. Spectrosc.* 2018.
- Arun, S., Parente, M., 2017. Active classification of neutral spectra for CRISM images. In: *Lunar and Planetary Science Conference*. p. 2866.
- Audebert, N., Le Saux, B., Lefevre, S., 2019. Deep learning for classification of hyperspectral data: A comparative review. *IEEE Geosci. Remote Sens. Mag.* 7, 159–173. <http://dx.doi.org/10.1109/MGRS.2019.2912563>.
- Ben Hamida, A., Benoit, A., Lambert, P., Ben Amar, C., 2018. 3-d deep learning approach for remote sensing image classification. *IEEE Trans. Geosci. Remote Sens.* 56, 4420–4434. <http://dx.doi.org/10.1109/TGRS.2018.2818945>.
- Breiman, L., 2001. Random forests. *Mach. Learn.* 45, 5–32.
- Caggiano, J., Sessa, A., Wray, J., Paty, C., 2019. Application of machine learning to identify surface minerals in CRISM imagery. *Lunar Planet. Sci.* 50.
- Chaves, L., Horgan, B., Lynch, K., Kimbrough, L., Hanley, J., Wray, J., 2018. Acidic environments in columbus crater, Mars: Implications for habitability. In: *Lunar and Planetary Science Conference*. p. 1744.
- Cheng, Y., Rajwa, B., Dundar, M., 2019. Bayesian nonparametrics for non-exhaustive learning. arXiv preprint [arXiv:1908.09736](https://arxiv.org/abs/1908.09736).
- Dundar, M., Akova, F., Qi, Y., Rajwa, B., 2012. Bayesian nonexhaustive learning for online discovery and modeling of emerging classes. In: *Proceedings of the International Conference on Machine Learning (ICML'12)*. pp. 99–106.
- Dundar, M., Ehlmann, B., 2016. Rare jarosite detection in CRISM imagery by non-parametric bayesian clustering. In: *Proceedings of Workshop on Hyperspectral Image and Signal Processing: Evolutions in Remote Sensing (WHISPERS)*. pp. 1–5.

- Dundar, M., Ehlmann, B., Leask, E., 2019a. Rare phase detections in CRISM data at pixel-scale by machine learning generate new discoveries about geology at Mars rover landing sites: Jezero and NE Syrtis. *Earth Space Sci. Open Arch.* 23. <http://dx.doi.org/10.1002/essoar.10501294.1>.
- Dundar, M., Ehlmann, B.L., Leask, E., 2019b. Machine-learning-driven new geologic discoveries at Mars rover landing sites: Jezero Crater and NE Syrtis. *Earth Space Sci. Open Arch.* 23. <http://dx.doi.org/10.1002/essoar.10501294.1>.
- Dundar, M., Rajwa, B., Li, L., 2013. Partially-observed models for classifying minerals on Mars. In: *Proceedings of Workshop on Hyperspectral Image and Signal Processing: Evolutions in Remote Sensing (WHISPERS)*. pp. 1–4.
- Ehlmann, B., Dundar, M., 2015. Are noachian/hesperian acidic waters key to generating Mars' regional-scale aluminum phyllosilicates? The importance of jarosite co-occurrences with al-phyllosilicate units. In: *Lunar and Planetary Science Conference*. p. 1635.
- Ehlmann, B., Dundar, M., 2016. Acidic conditions during open system weathering on late noachian/early hesperian Mars? Newly identified outcrops of alunite and jarosite from orbital CRISM data.
- Ehlmann, B.L., Mustard, J.F., Swayze, G.A., Clark, R.N., Bishop, J.L., Poulet, F., Des Marais, D.J., Roach, L.H., Milliken, R.E., Wray, J.J.e.a., 2009. Identification of hydrated silicate minerals on Mars using mro-CRISM: Geologic context near nili fossae and implications for aqueous alteration. *J. Geophys. Res. Planets* (1991–2012), 114.
- Ehlmann, B.L., Swayze, G.A., Milliken, R.E., Mustard, J.F., Clark, R.N., Murchie, S.L., Breit, G.N., Wray, J.J., Gondet, B., Poulet, F., et al., 2016. Discovery of alunite in cross crater, terra sirenum, Mars: Evidence for acidic, sulfurous waters. *Am. Mineral.* 101, 1527–1542.
- Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 1995. *Bayesian Data Analysis*. Chapman and Hall/CRC.
- He, L., O'Sullivan, J.A., Politte, D.V., Powell, K.E., Arvidson, R.E., 2019. Quantitative reconstruction and denoising method hyber for hyperspectral image data and its application to crism. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*
- Itoh, Y., Parente, M., 2021. A new method for atmospheric correction and de-noising of crism hyperspectral data. *Icarus* 354, 114024.
- Kreisch, C., O'Sullivan, J., Arvidson, R., Politte, D., He, L., Stein, N., Finkel, J., Guinness, E., Wolff, M., Lapôte, M., 2017. Regularization of Mars reconnaissance orbiter CRISM along-track oversampled hyperspectral imaging observations of Mars. *Icarus* 282, 136–151.
- Leask, E., Ehlmann, B., Dundar, M., 2019. Evidence for chemically distinct waters forming sulphates and chlorides in terra sirenum, Mars.
- Leask, E., Ehlmann, B., Dundar, M., Murchie, S., Seelos, F., 2018. Challenges in the search for perchlorate and other hydrated minerals with 2.1- μm absorptions on Mars. *Geophys. Res. Lett.* 45, 12–180.
- Li, Y., Zhang, H., Shen, Q., 2017. Spectral-spatial classification of hyperspectral imagery with 3d convolutional neural network. *Remote Sens.* 9. <http://dx.doi.org/10.3390/rs9010067>.
- Min, E., Guo, X., Liu, Q., Zhang, G., Cui, J., Long, J., 2018. A survey of clustering with deep learning: From the perspective of network architecture. *IEEE Access* 6, 39501–39514. <http://dx.doi.org/10.1109/ACCESS.2018.2855437>.
- Mou, L., Ghamisi, P., Zhu, X.X., 2017. Deep recurrent neural networks for hyperspectral image classification. *IEEE Trans. Geosci. Remote Sens.* 55, 3639–3655. <http://dx.doi.org/10.1109/TGRS.2016.2636241>.
- Murchie, S.L., Seelos, F.P., Hash, C.D., Humm, D.C., Malaret, E., McGovern, J.A., Choo, T.H., Seelos, K.D., Buczkowski, D.L., Morgan, M.F.e.a., 2009. Compact reconnaissance imaging spectrometer for Mars investigation and data set from the Mars reconnaissance orbiter's primary science phase. *J. Geophys. Res. Planets* (1991–2012), 114.
- Nydick, S.W., 2012. The wishart and inverse wishart distributions. *Electron. J. Stat.* 6, 1–19.
- Parente, M., Saranathan, A., Wiseman, S., Ehlmann, B., Pan, L., 2014. Denoising CRISM images: A new look. In: *Lunar and Planetary Science Conference*. p. 2900.
- Pelkey, S., Mustard, J., Murchie, S., Clancy, R., Wolff, M., Smith, M., Milliken, R., Bibring, J.P., Gendrin, A., Poulet, F.e.a., 2007. CRISM multispectral summary products: Parameterizing mineral diversity on Mars from reflectance. *J. Geophys. Res. Planets* (1991–2012), 112.
- Roach, L.H., Mustard, J.F., Swayze, G., Milliken, R.E., Bishop, J.L., Murchie, S.L., Lichtenberg, K., 2010. Hydrated mineral stratigraphy of ius chasma, valles marineris. *Icarus* 206, 253–268.
- Saranathan, A.M., Parente, M., 2021. Adversarial feature learning for improved mineral mapping of CRISM data. *Icarus* 355, 114107.
- Viviano-Beck, C.E., Seelos, F.P., Murchie, S.L., Kahn, E.G., Seelos, K.D., Taylor, H.W., Taylor, K., Ehlmann, B.L., Wiseman, S.M., Mustard, J.F., Morgan, M.F., 2014. Revised CRISM spectral parameters and summary products based on the currently detected mineral diversity on Mars. *J. Geophys. Res. Planets* 119, 1403–1431. <http://dx.doi.org/10.1002/2014JE004627>, 2014JE004627.
- Yerebakan, H.Z., Rajwa, B., Dundar, M., 2014. The infinite mixture of infinite Gaussian mixtures. In: *Advances in Neural Information Processing Systems (NIPS)*. pp. 28–36.