ADVERSARIAL FEATURE LEARNING FOR IMPROVED MINERAL MAPPING IN CRISM IMAGES

Arun M. Saranathan and Mario Parente, Department of Electrical & Computer Engineering, University of Massachusetts Amherst, Amherst, MA 01003

Introduction: The Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) [1] has been a valuable source of information on the mineralogy of the Martian surface. In spectroscopic data depending on its chemical composition and crystal structure each mineral exhibits an unique absorption pattern [2], making it possible to identify them. In fact, expert analysis of the CRISM data has lead to the identification of various minerals on the martian surface (see [3] and references therein). Given, that there are more than 10000 CRISM images containing millions of pixel spectra automated mineral identification is an important task for the community. The task of measuring similarity/dissimilarity between various spectra is made harder by the fact that, in addition to the surface information the measured spectra are also affected by various physical parameters (such as acquisition geometry, surface properties etc.), as well noise/distortion processes (such as atmospheric effects and instrument artifacts etc.).

Generally, practitioners attempt to identify the presence of each mineral absorption feature by representing it by a single number, generated by applying "an algorithm using combinations of spectral bands" [4]. These are referred to as *spectral summary parameters* [3] and have proven to be very useful guides for manual analysis. Attempts have been made to generate a rule based system using the summary parameters for automated mineral identification [5]. While the summary parameters are quite successful in identifying the presence of specific features, they are adversely affected by the presence of the noise/distortions in the CRISM images making such automation harder. Other approaches based on curve-fitting [6], image segmentation and unmixing [7], target transformation [8] and spectral signal processing [9] have also been tried. While these methods have shown some promise in identifying various minerals they are affected by some common issues (1) their performance suffers in the presence of the noise/distortion (2) they are often designed to identify specific minerals and may include processing steps which affect the identification of other minerals and (3) they are computation-intensive. We propose a new model based on Generative Adversarial Networks (GANs) to learn features which are capable of identifying the various mineral spectra. We will show that using a simple similarity metric like Spectral Angle Divergence (SAD) in this feature space is sufficient to identify the different minerals.

Data Preprocessing: Due to the presence of the various distortion effects present in CRISM data, preprocessing is essential to minimize their adverse effects. To minimize the effects of noise/atmospheric distortions we process the data using the novel image denoising algorithms designed by Itoh et al. [10][11]. We also use continuum-removal to minimize various illumination and scaling effects. Also, since mineral identification is based on absorption features, we also eliminate flat looking spectra, i.e. spectra with less than a 2% absorption.

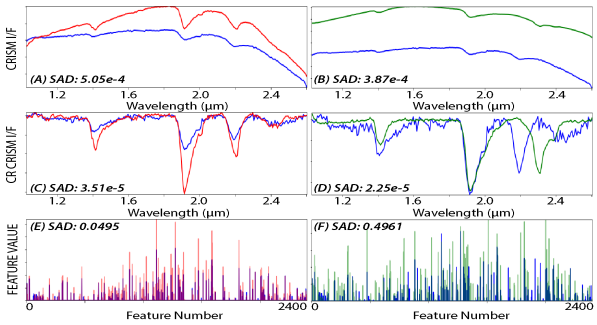

Figure 1: Discriminative power of the GAN representations (A) SAD between target CRISM I/F spectra (blue) and MICA Al-Smectite (red) (B) SAD between target CRISM I/F spectra (blue) and MICA Fe-Smectite (green) (C) SAD between target CRISM Continuum Removed spectra (blue) and MICA Al-Smectite (red) (D) SAD between target Continuum Removed spectra (blue) and MICA Fe-Smectite (green) (E) SAD between the GAN representation of the target spectra (blue) and MICA Al-Smectite (red) (F) SAD between the GAN representation of the target spectra (blue) and MICA Fe-Smectite (green)

Adversarial Feature Learning: GANs are a neural network based architecture with two networks engaged in a zero-sum game. One network the *generator* takes a random vector as input and attempts to *generate* samples that are indistinguishable from the real samples (in this case CRISM spectra), while the other the *discriminator* attempts to differentiate between real samples and the *generated* samples. In training, we first train the discriminator to differentiate between *generated* samples and real world samples. In the next step, the two networks are chained and the generator is trained to *generate* samples which would fool the discriminator (the discriminator weights are frozen for this step). These steps are alternated until the discriminator is unable to tell the difference between the *generated* samples and the real samples (convergence). At convergence, the *generated* samples appear "real", i.e. the generator has *"learned"* to generate mineral spectra. Since the generator is only trained to "fool" the discriminator, the only way the generator "learns" to generate mineral spectra

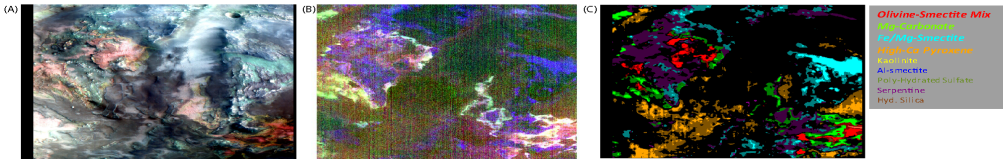

Figure 2: Mapping of the CRISM image $FRT93BE$ using the GAN representation (A) RGB Composite of the image (B) A CRISM parameter browse product with (R: OLINDEX, G: BD2500, B: D2300) (C) Mineral detection map using the GAN features

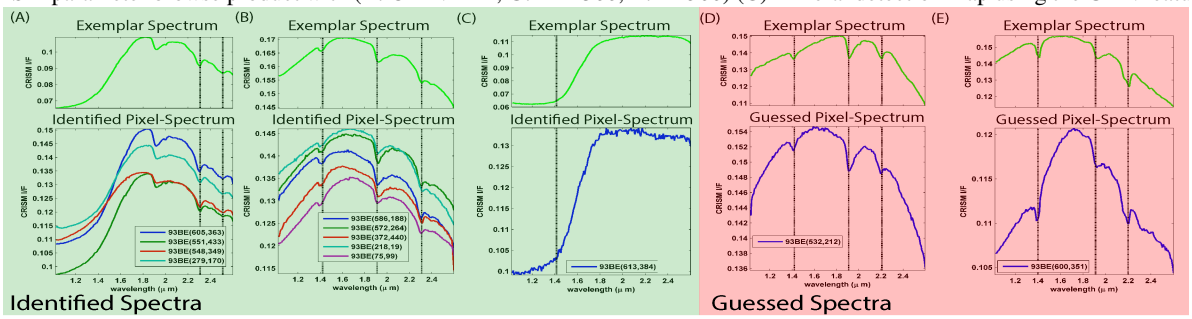

Figure 3: Comparision of mineral spectra detected in $FRT93BE$ using the GAN representation to the exemplar spectrum for the minerals identified (A-C) or guessed (D-E) (A) Mg-Carbonate (B) Fe/Mg-Smectite (C) Fe-Olivine (D) Al-Smectite (E) Kaolinite

is if the discriminator has learned a good representation for mineral spectra. Based on this we propose to use the representation learned by the discriminator as our features.

To illustrate the power of the representation we show an example in Fig. 1. The target (blue) spectrum shown in Fig. 1 (A) has absorption bands at $1.4,\ 1.9\ \&\ 2.2\ \mu m$ which is similar to a MICA Al-Smectite (red)[1] as opposed to the MICA Fe-Smectite (green). In both the CRISM I/F and continuum removed space, the SAD between the target and the Fe-Smectite is smaller than with respect to the Al-Smectite, i.e. the similarity is higher with the incorrect mineral. On the other hand, in the feature space the target is far more similar to the correct mineral (Al-Smectite), illustrating the power of the features in resolving spectral similarity.

Mineral Mapping & Results: To generate the maps, we begin by using the discriminator to extract the representation for some *exemplar spectra* (the spectrum which we wish to search the image-database for). Next we extract the representation for the test spectra. Following this we evaluate the SAD between the exemplar representation and test data representation. Since the learned representation is highly discriminative, test spectra whose representation have a very high similarity ($SAD \leq 0.1$) to the exemplar, mostly have exactly the same absorption features as the exemplar, such pixels are detected as being the same as the exemplar with high confidence and referred to as *identifications*. Test spectra with intermediate similarity (SAD between 0.1 and 0.293) have most of the same absorption features as the exemplars but exhibit some variations, these pixels are detected with a lower confidence and referred to as *guesses*.

An example of the mapping results for the CRISM image $FRT93BE$ is shown in Fig. 2. The boldly colored pixels in Fig. 2 (C) represent *identifications* while the muted coloration indicates a *guess*. Note the mineral mapping from the adversarial features (shown in Fig. 2 (C)) is less noisy than the composite maps from summary parameter maps shown in Fig. 2 (B) while still managing to identify the minerals present. In Fig. 3 we show some example of pixels either *identified* or *guessed* as being similar to the exemplars. In the case of Fig. 3 (A)–(C) we show examples of pixels *identified* as being similar to the exemplar. All these pixels have all the same absorption features as the exemplars. On the other hand Fig. 3 (D)–(E), show *guesses*, these spectra are quite similar to the exemplar but have some difference in the absorption features vis-a-vis the exemplars (for e.g. the *guessed* Kaolinite spectra in Fig. 3 (E) has an additional band at 2.3 μm). Across many CRISM images the model has been able successfully identify and map the various minerals present in the scene.

References: [1] S. Murchie, et al. (2007) *J Geophy Res: Planets* 112(E5). [2] R. N. Clark, et al. (1999) *Manual of remote sensing* 3(3-58):2. [3] C. E. Viviano-Beck, et al. (2014) *J Geophy Res: Planets* 119(6):1403. [4] S. Pelkey, et al. (2007) *J Geophy Res: Planets* 112(E8). [5] E. Allender, et al. (2017) *Icarus* 281:151. [6] M. Parente, et al. (2011) *Planetary and Space Science* 59(5-6):423. [7] M. S. Gilmore, et al. (2011) *J Geophy Res: Planets* 116(E7). [8] N. H. Thomas, et al. (2017) *Icarus* 291:124. [9] J. Carter, et al. (2013) *Planetary and Space Science* 76:53. [10] Y. Itoh, et al. (2019) *In Preparation*. [11] Y. Itoh, et al. (2019) in *50^{th} LPSC, # 2025*.

[1]MICA stands for Minerals Identified through CRISM Analysis (MICA), the data is available at `http://crismtypespectra.rsl.wustl.edu/` and is described in [3]