

AUTOMATED NEUTRAL REGION SELECTION USING SUPERPIXELS

Lukas Mandrake¹, David R. Thompson¹, Martha Gilmore², Rebecca Castano¹, Eldar Z. N. Dobrea³

¹JPL / California Inst. Of Tech.
4800 Oak Grove Ave.
Pasadena, CA 91109
firstname.lastname@jpl.nasa.gov

²Wesleyan University
265 Church St.
Middletown, CT 06459
mgilmore@wesleyan.edu

³Planetary & Space Science Institute
1700 E. Fort Lowell, #106
Tucson, AZ 85719
eldar@psi.edu

ABSTRACT

This work presents an automated approach utilizing superpixel segmentation for detecting spectrally Neutral Regions (NR) in hyperspectral images. NRs are often used in planetary geology as spectral divisors to Regions of Interest (ROI), both to enhance key mineralogical signatures and correct for systematic errors such as residual atmospheric distortion. We compare automated NR selections to handpicked examples with mineralogical summary products used in analysis of data from the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM). We also present a new summary product to quantify the level of atmospheric distortion in a CRISM spectrum. We find that the automated algorithm matches manual NR detection with regards to mineral spectral contrast and outperforms manual selection for reducing atmospheric distortion.

Index Terms— hyperspectral, CRISM, superpixel, neutral region, toolbox, ATMO

1. INTRODUCTION

While hyperspectral images are fairly directly interpretable given pure sample pixels and a prior known library, planetary science applications present complications including impure mineral constituents even at the sub-pixel level, dust and debris lying above the substrate, and contributions from unknown minerals [1]. Techniques to reduce noise, correct systematic errors, and enhance faint spectral features thus become vital to useful interpretation.

One common strategy for enhancing interesting mineral features in reflectance spectra is the division of a feature target by a neutral spectrum [2][3][4][5]. This neutral spectrum is generated by averaging over a Neutral Region (NR) drawn from the same image. The NR is not precisely defined: should be spectrally “bland,” i.e. lacking the spectral features of interest. Instead, it is comprised of the background continuum and any noise artifacts or atmospheric absorption features that are common across the entire image (Figure 1). The neutral spectrum then serves as a divisor to a target spectrum:

$$S_{\text{ratioed}}(b) = \frac{S_{\text{original}}(b)}{S_{\text{neutral}}(b)} \quad \forall b \in \text{bands}$$

This yields a spectral reflectance curve relative to the surrounding spectral background as characterized by the NR. Ideally the ratioed spectrum enhances important features unique to the target and reduces incidental features common to both. The quality of the result hinges on finding an appropriate NR.

Currently NR selection requires a time consuming manual trial and error search across many pixels, with the final NR chosen by the expert judgment of the analyst. Often geologists prefer regions that are visually dark (absorptive, not shadowed) and have fairly spatially uniform texture. Dusty, sandy or granular regions are often chosen due to the nonlinear convolution properties of multiple scattering sources. If these small scattering bodies are dark (absorptive), it can be shown that distinctive mineral features are generally suppressed relative to the continuum background making them ideal for NR selection. Ultimately, however, any measure of the acceptability of the NR is left to intuition and observed desirable effects. Suboptimal NR selection can introduce unwanted artifacts since those features unique to the NR appear as inverted distortions in the ratio. A NR that contains some trace of a mineral of interest could reduce rather than enhance a desired spectral feature. In some scenes, features of interest are present in every NR forcing suboptimal choices.

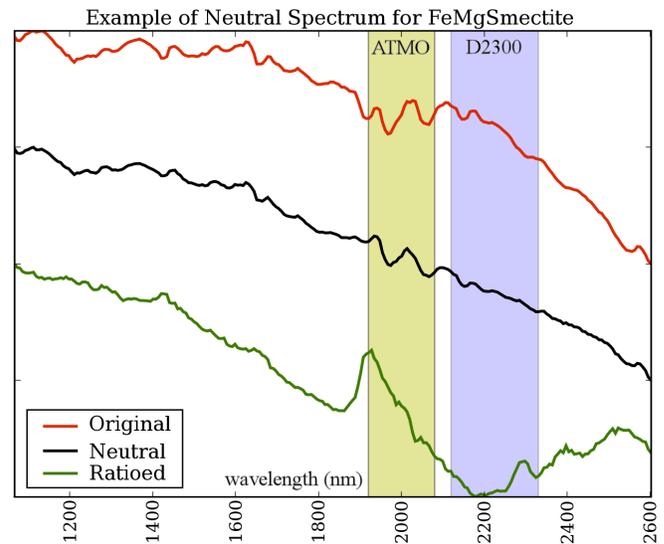


Figure 1 Original, Neutral, and Ratioed spectra for ROI in CRISM scene 863e. Yellow region shows characteristic residual CO₂ atmospheric absorption features. Absorption within blue region identified by expert as diagnostic for FeMgSmectite.

Localized instrument artifacts such as “push-broom” column striping can also complicate NR identification [1]. In the case of the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) [9], each column of an unprojected / ungeorectified image represents a single detector with a unique spectral response

curve. Analysts typically correct for calibration error through division by a planetary standard spectrum of Olympus Mons [6], but this can leave residual errors due to local atmospheric phenomena or time-varying instrument artifacts. To ideally exercise the NR technique, a NR should share precisely the same columns (detectors) as a target Region of Interest (ROI). Additionally, if the target is the average spectrum of a region of interest, the associated NR should have a similar shape and size in order to appropriately weigh contributions from each column / sensor. In practice, this last constraint is difficult to satisfy and almost never used. The more general shared column constraint, however, is commonly obeyed. It should be noted that the NR technique does not represent a means to reduce shot noise and actually increases the effects of random errors. We moderate this effect through a superpixel noise reduction strategy [7]. Superpixel segmentation is an edge-preserving method for reducing spatially uncorrelated noise. It splits the image into small contiguous, homogeneous fragments and draws representative spectra from each. This preserves column relationships for NR selection [8].

In an attempt to formalize and automate the NR selection process, we have developed an algorithm for identifying NRs. We evaluated the method with a new summary product that estimates the residual atmospheric distortion in a spectrum. Correction for this distortion along with target mineral spectral feature contrast was taken as a numerical estimate of overall acceptability for a given NR. We have demonstrated our automated NR technique with CRISM images and compared the result to manual NR selection by a planetary geology expert. CRISM exhibits higher noise levels than most Earth-based hyperspectral images, and the superpixel segmentation strategy significantly enhances the selection of useful NR.

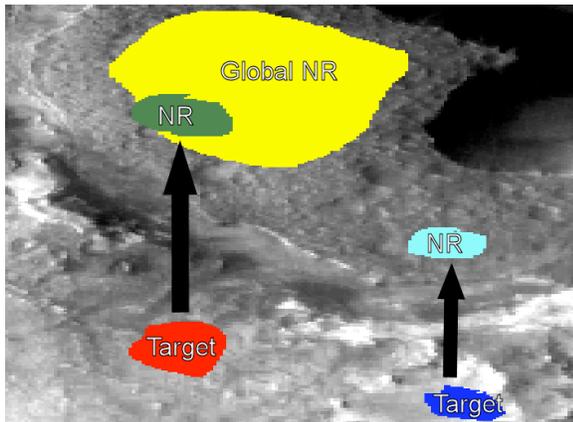


Figure 2 Subsection of CRISM image FRT00003E12. The yellow region is the manually selected image-wide NR. Red and blue are the target ROIs. Green and cyan show manually selected NRs.

2. AUTOMATED NEUTRAL REGION DETECTION

We first utilize the relative simple Felzenswab image segmentation method [7] to generate regions of spectral similarity. Two hyperparameters control this segmentation: a tolerance parameter for internal spectral variability within a segment, and a minimum permissible segment size. We then consider each segment as a

potential NR. Recalling the definition of a NR above, we recognize that a properly representative NR should share the columns of the target region and lack the key diagnostic features of the target ROI. However, as we are constructing an automated process that requires no user input, we attempt to minimize all spectral features within the NR. We simply fit a line to each L1 normalized mean super-pixel segment spectrum in the image using the range 1000 – 2600 nm. The best-fit line minimizes RMS error, and the lowest RMS is interpreted as the most neutral spectrum. For other images, instruments, or spectral regions possessing large-scale spectral structure it might be preferable to fit a spline or higher order polynomial, but in our case a line was sufficient and better characterizes CRISM spectral continuum.

3. EVALUATION CRITERIA

We examined three well-studied CRISM targeted images (FRT0000)3e12, 863e, and 3fb9 focusing on the 1000-2600 nm spectral region. Each image is roughly 640 x 480 with 234 included bands. This spectral region is diagnostic of many Martian minerals of interest (Table 1). However, it is also plagued by a systematic CO₂ atmospheric absorption feature at 1920-2080 nm [10]. A planetary geologist expert provided manually constructed target ROIs that exhibited novel spectral absorption features indicative of significant mineral deposits. These potential deposits include minerals such as phyllosilicates and other products of aqueous alteration. These interesting features were often extremely faint compared to the background spectral features. Each ROI is associated with one or more CRISM summary products [11] from the CRISM Analysis Tool (CAT) toolbox [12]. These scalar metrics measure absorption depth at discriminating spectral locations useful for identifying Martian mineralogy (e.g. D2300's focus region is shown in Figure 1). We paired each ROI with a manually-selected NR from the same column. The expert also provided a generic NR for the entire image for situations in which the target region is unknown. A total of 11 target ROIs were included in our analysis across the three images along with their associated manually selected NRs.

Table 1 – Identified ROIs & Contrast Metric

Image	Mineral ID	CAT Summary Product [11]
3e12	Magnesite	BDCARB
3e12	Olivine	OLINDEX
3e12	Phyllosilicate	BD1900, D2300
3fb9	Carbonate	BDCARB
3fb9	Kaolinite	BD2210
3fb9	Olivine	OLINDEX
3fb9	Phyllosilicate	BD1900, D2300
863e	FeMgSmectite	BD1900, D2300
863e	Kaolinite	BD2210
863e	Montmorillonite	BD1900
863e	Nontronite	BD1900, D2300

To compare the automated NR detection algorithm with the expert-provided manual NRs, we evaluated the ratioed spectrum for 1) the strength of target spectral features and 2) the

degree of atmospheric distortion remaining. Each image contained several expert-identified ROIs associated with a particular mineral blend of interest. CRISM summary products [11] as implemented in the CAT toolbox [12] were then associated with each mineral per the expert's identifications. We then compared these summary products before and after ratioing using a variety of potential NRs. It should be noted that NR selection is constrained by the available image and further by the column so that a mineral's signal might actually be reduced by any candidate NRs. Our task in this case is to minimize the contrast loss from the NR technique.

Motivated by the utility of the CRISM summary products and lacking an existing metric, we created a new measure (ATMO) to quantify residual atmospheric absorption. Shown in Figure 1 as a yellow shaded region, the band of atmospheric absorption manifests several characteristic oscillations between 1920-2080 nm. ATMO is defined by fitting a 2nd order polynomial to these bands. The reported value is one minus the coefficient of determination (R^2) producing

$$ATMO \equiv 1.0 - \frac{\left[\sum (S_{orig} - \bar{S}_{orig})(S_{fit} - \bar{S}_{fit}) \right]^2}{\left[\sum (S_{orig} - \bar{S}_{orig})^2 \right] \left[\sum (S_{fit} - \bar{S}_{fit})^2 \right]}$$

where S is the reflectance intensity. Thus, a value of 0 represents a perfectly smooth 2nd order surface (no residual atmospheric absorption), while a value of 1 represents strong oscillatory departure from baseline (high residual atmospheric absorption). Note that this is a separate model unrelated to the linear fit over the wider spectrum used to detect potential NRs. As a properly chosen NR should reduce the systematic atmospheric distortion residual in the ratioed spectrum, we monitor the ATMO measure both before and after neutral spectrum ratioing as a second gauge of successful NR choice.

4. EXPERIMENT & RESULTS

We evaluated the associated summary product for each of the eleven target ROIs both before and after ratioing using the following neutral regions:

- 1) None (original superpixel spectrum),
- 2) Manual whole-image NR
- 3) Manual column-constrained NR associated with each ROI,
- 4) Automated whole-image NR selection
- 5) Automated column-constrained NR.

We evaluated the linear fit criterion using various values for the segmentation parameter defining the smallest permissible superpixel. We found a minimum size of 40-60 image pixels was consistently optimal across all metrics, and we report results for minimum size equal to 60 pixels. A second parameter that controls the superpixel internal variability tolerance (k) was fixed to 0.0001; previous studies have found segmentation results to be robust despite significant alteration of this value.

Figure 3 shows resulting changes in mineral contrast as measured by the associated CRISM summary products. The manually selected NRs often reduced relative mineral contrast in the target ROIs indicating NR contamination with features of interest. The automated method is competitive with the manual method to within error. Column constraints do not improve NR selection to within our sampling error. This could be due to the

column constraint limiting the selection of candidate NRs even as it ensures noise characteristics are similar to the target.

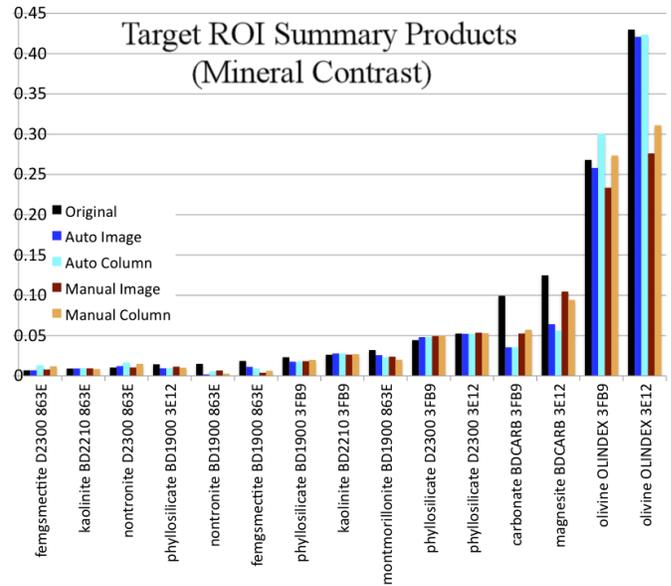


Figure 3 The CRISM summary products associated with each target ROI and the effects of NR ratioing. Higher is better. The automated method is competitive with manual NR selection.

Figure 4 replicates this same comparison using the ATMO parameter measure for residual atmospheric distortion. Here, it is seen that all NR techniques significantly reduce ATMO by 80-90% with our automated method performing equally well as manual selection to within error.

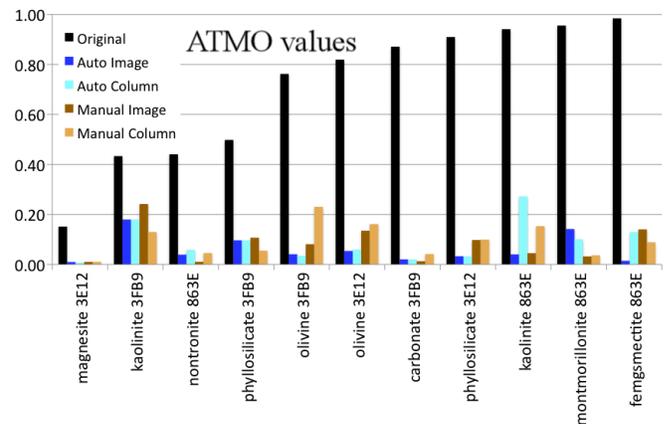


Figure 4 ATMO product of target ROIs before and after various NR selections. Lower is better. Automated method is competitive and very slightly superior to manual selection, but both greatly improve the unratioed result.

While the strong decrease in ATMO for both automated and manual NR selection is promising, it does not in itself demonstrate the value of the linear fit criterion. Any improvement must be significant compared to chance (i.e. the random selection of a NR). To validate the automated method we completed an

exhaustive comparison of ATMO variation using a global image-wide NR (Figure 5 & Table 2). We first measured the raw ATMO of each segment's mean spectrum ("Original", black) in a given image. We then paired each segment with every other segment as a potential divisor, generating an exhaustive $N \times (N-1)$ set of ATMO values ("Ratio Exhaustive", blue). This nearly uniform distribution shows that it is possible to significantly harm the ATMO measure by arbitrary NR selection. We then took the ratio of the segments with the manual global NR ("Ratio Manual Picks", magenta).

While manual selection improved the ATMO distribution, a significant number of segments retained high atmospheric distortion. The linear fit criterion was then used to select an image-wide NR ("Ratio Linear Picks", green), resulting in a more dramatic reduction in ATMO. Finally, for completeness we show the distribution of the best possible ATMO for each segment given exhaustive knowledge of all possible ratios ("Ratio Best Possible", red). This distribution is not truly fair, since it results from an exhaustive search that seeks to minimize the exact evaluation criterion. A true NR must not only decrease ATMO but also maintain mineral contrast. Atmospheric absorption can be difficult to model explicitly so the NR criterion should identify NRs that can correct atmospheric artifacts other than those within the NR selection criterion. However, this best possible ratio does show the limit of ATMO reduction ignoring other considerations. Figure 5 and Table 2 both represent the same data, resulting from the union of all three image scene distributions.

Original & Ratioed ATMO Values for Sum of Scenes

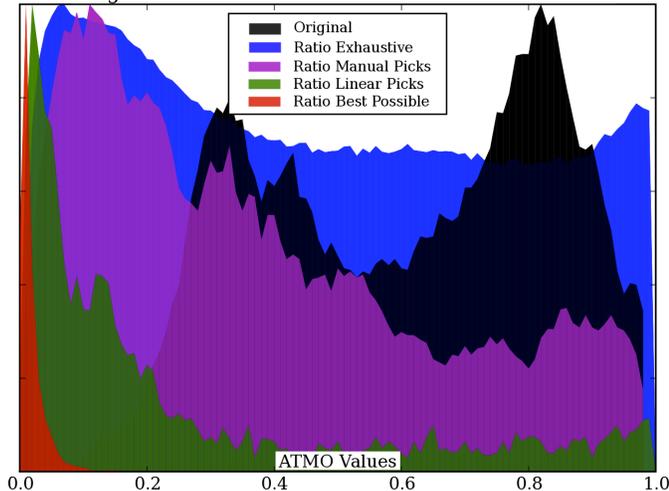


Figure 5 Distributions of ATMO values for the original segment spectra and ratioed spectra derived from various NR selection strategies. Left is better. The line fit criterion performs better than all other general techniques.

Table 2 – Original & Ratioed ATMO Values

	Mean
Original	0.61
Ratio Exhaustive	0.48
Ratio Manual Picks	0.38
Ratio Linear Picks	0.26
Ratio Best Possible	0.02

6. CONCLUSION

We have shown that a simple RMS residual to a line fit, when combined with a superpixel segmentation, can match an expert manual NR selection for mineral contrast preservation and outperform it for reducing atmospheric distortion as measured by the ATMO metric. The NR technique was also shown to occasionally reduce mineral contrast, but the automated algorithm helped to mitigate this effect.

ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2010. All rights reserved.

7. REFERENCES

- [1] T. M. Lillesand, R. W. Kiefer, and J. W. Chipman, "Remote Sensing and Image Interpretation (6th)," Wiley, New York, 2007.
- [2] E. W. Chappelle, M. S. Kim and J. E. McMurtrey, "Ratio Analysis of Reflectance Spectra," *Geoscience and Remote Sensing Symposium*, 1990.
- [3] H. B. Jensen and T. D. Glotch, "Near Infrared Spectral Analysis of Mixtures of Halite and Labradorite for Application to Putative Chloride Deposits Observed by CRISM," *Lunar and Planetary Science Conference*, 2010.
- [4] J. F. Mustard et al., "Hydrated Silicate Mineral on Mars Observed by the Mars Reconnaissance Orbiter CRISM Instrument," *Nature*, vol. 454, pp. 305-309, July 2008.
- [5] B. L. Ehlmann et al., "Orbital Identification of Carbonate-Bearing Rocks on Mars," *Science*, vol. 322, p. 1828, Dec. 2008.
- [6] Y. Langevin, F. Poulet, J. P. Bibring and B. Gondet, "Sulfates in the North Polar Region of Mars Detected by OMEGA/Mars Express," *Science*, vol. 307 (5715), pp. 1584-1586, 2005.
- [7] P.F. Felzenszwalb and D.P. Huttenlocher, "Efficient graph-based image segmentation," *International Journal of Computer Vision*, vol. 59:2, pp. 167-181, 2004.
- [8] D. R. Thompson, R. Castano and M. S. Gilmore, "Sparse Superpixel Unmixing for Exploratory Analysis of CRISM Hyperspectral Images," *IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing*, August 2009.
- [9] S. Murchie et al., "Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) on Mars Reconnaissance Orbiter," *J. Geophys. Res.*, vol. 112, no. 5, 2007.
- [10] P. E. Martin and E. F. Barker, "The Infrared Absorption Spectrum of Carbon Dioxide," *Phys. Rev.*, vol. 41, pp. 291-303, Jun 1932.
- [11] S.M. Pelkey et al., "CRISM Multispectral Summary Products: Parameterizing Mineral Diversity on Mars from Reflectance," *J. Geophys. Res.*, vol. 112, 2007.
- [12] F. Morgan et al., "CAT Tutorial," CRISM Data User's Workshop, *Lunar and Planetary Science Conf.*, 2009.