

# Superpixel Endmember Detection

David R. Thompson, Lukas Mandrake, Martha S. Gilmore, and Rebecca Castaño

**Abstract**—Superpixels are homogeneous image regions comprised of multiple contiguous pixels. Superpixel representations can reduce noise in hyperspectral images by exploiting the spatial contiguity of scene features. This paper combines superpixels with endmember extraction to produce concise mineralogical summaries that assist in browsing large image catalogs. First, a graph-based agglomerative algorithm oversegments the image. We then use segments' mean spectra as input to existing statistical endmember detection algorithms such as sequential maximum angle convex cone (SMACC) and N-FINDR. Experiments compare automatically detected endmembers to target minerals in an Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) scene of Cuprite, Nevada. We also consider a planetary science data set from the Compact Reconnaissance Imaging Spectrometer (CRISM) instrument that benefits from spatial averaging due to higher noise. In both cases, superpixel representations significantly reduce the computational complexity of later processing while improving endmembers' match to the target spectra.

**Index Terms**—Airborne Visible/Infrared Imaging Spectrometer (AVIRIS), Compact Reconnaissance Imaging Spectrometer (CRISM), endmember detection, hyperspectral images, planetary geology, segmentation, superpixels.

## I. INTRODUCTION

STATISTICAL analysis of hyperspectral image data generally starts with one of two representations: 1) Independent measurements taken one per pixel; or 2) large segmented regions that correspond directly to the scene features of interest. Pixels taken independently are sensitive to instrument noise and intraclass variability. Scene segmentation can reduce noise and variability but can be difficult to automate, and segmentation may incorrectly mask spectral features. This paper explores a superpixel representation [1], [2] that combines some benefits of both extremes. Superpixels are homogeneous image regions comprised of several contiguous pixels having similar value (e.g., grayscale, color). They are generated by an intentional oversegmentation that forms a one-to-many partitioning of scene features into smaller segments (Fig. 1). Our use of the term superpixel is distinct from its occasional use to describe spatial binning in a uniform grid; we impose no morphological constraints except spatial contiguity.

Superpixels are a common preprocessing step before object segmentation in RGB and single-channel images [1]–[3].

Manuscript received November 24, 2009; revised March 29, 2010 and June 6, 2010. Date of publication September 30, 2010; date of current version October 27, 2010.

D. R. Thompson, L. Mandrake, and R. Castaño are with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 USA (e-mail: david.r.thompson@jpl.nasa.gov).

M. S. Gilmore is with Wesleyan University, Middletown, CT 06459 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TGRS.2010.2070802



Fig. 1. Segmentation of an image patch from CRISM frt00003e12, originally from [4] ( $R : 2.0 \mu\text{m}, G : 1.5 \mu\text{m}, B : 1.1 \mu\text{m}$ ). (Left) Original subimage. (Center) Coarse segmentation, minimum region size 100. (Right) Fine segmentation, minimum region size 20. Finer segmentations provide additional resolution at the cost of greater computation time and sensitivity to noise. Image courtesy NASA / John Hopkins.

They could also benefit hyperspectral images by identifying populations of adjacent spectra that serve as independent and identically distributed (IID) samples of each surface feature. This additional spatial information distinguishes real spectral features from noise artifacts and permits more robust estimation of spectral properties. Note that a “perfect” (i.e., most physically meaningful) scene segmentation would also provide these advantages, but automating such segmentation is generally quite difficult. Oversegmentation is a pragmatic alternative when features have diverse scales or when they are not known in advance. Additionally, superpixel representations can improve the computational efficiency of later processing; depending on the spatial complexity of the image, they can reduce the number of distinct spectra to analyze by one or more orders of magnitude.

This paper uses superpixel representations to identify pure materials, or endmembers, of hyperspectral images. Our motivating application is the challenge of browsing and summarizing large hyperspectral archives. Here, endmember detection can provide fast automated summaries to guide analysts' attention to regions of interest. Our primary focus is planetary science data sets with low signal-to-noise ratios (SNRs) and atmospheric or instrument variability at the pixel level. We review the endmember detection task, describe our segmentation approach, and evaluate both pixel and superpixel representations for recovering known image features.

The evaluation uses two distinct data sets. We first consider an Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) image [5] to test the technique with well-studied data of high SNR. We also examine the Compact Reconnaissance Imaging Spectrometer (CRISM) aboard the Mars Reconnaissance Orbiter [6], a planetary data set for which we anticipate superpixel methods could enhance classification. While superpixels may mask the mineralogy of very small or subpixel outcrops, they offer significantly improved recall and spectral accuracy for scenes dominated by large features. In all cases, the data reduction yields significant benefits for run time.

## II. RELATED WORK

According to the geographic mixing assumption observed, reflectances are linear combinations of a small number of endmember materials. We use the standard linear mixing model, with  $m$  endmember spectra in  $w$  wavelengths generating each image pixel independently according to the expression

$$\rho_j = \mathbf{Q}\mathbf{a}_j + \mathcal{N}(0, \sigma^2) \quad (1)$$

All the pixels in the image share a  $w \times m$  matrix  $\mathbf{Q}$  whose columns are the endmember spectra. Here,  $\mathbf{a}$  is an  $m$ -vector of nonnegative mixing coefficients whose entries give each endmember's contribution to the observation. Gaussian noise of variance  $\sigma^2$  is independent in each channel. Excluding the noise term, the spectral observations  $\rho$  then lie on the  $m$ -simplex with vertices defined by the endmembers. Endmember detection aims to identify the vertices of this simplex, which typically correspond to the purest examples of the materials in the scene. Typically, endmember extraction fits a simplex to the convex hull of the data. A wide range of algorithms exist; Plaza *et al.* offer a comprehensive review [7]. In fact, nearly any statistical endmember extraction strategy could operate on superpixel representations.

We are interested primarily in automated search, so we focus on fully unsupervised endmember detection strategies that do not require significant user intervention. This also prevents operator performance from influencing evaluation results. The following experiments consider two representative algorithms. The sequential maximum angle convex cone (SMACC) method typifies sequential endmember detection [8]. It iteratively adds endmember spectra to a basis of increasing dimensionality. At each step, it adds as an endmember of the spectrum with the largest orthogonal projection onto the current basis. The SMACC method requires little tuning to achieve reasonable performance. In contrast, the N-FINDER algorithm exemplifies simultaneous methods that choose all endmembers at once [9]. It randomly initializes the set of endmembers and then modifies it by identifying the endmember position and pixel that best improves the volume of the resulting simplex. Multiple random initializations help prevent convergence to a local minimum.

Endmember extraction based on convex geometry favors outlier spectra and is sensitive to measurement noise. Noisy observations distort the geometry of the convex hull, appear as false endmembers, and hide real ones. Spectral smoothing (through PCA, for example) can reduce this effect but it also masks subtle spectral features. This has led researchers to consider methods of utilizing spatial information to improve endmember detection.

Some endmember extraction algorithms incorporate spatial information directly. This can recognize subtle spectral features near the noise level by exploiting features' presence in multiple neighboring pixels. For instance, Rogge *et al.*'s Spatial-Spectral Endmember Extraction estimates endmember pixels' spectra by averaging similar candidates within a local spatial window [10]. The Successive Projection Algorithm (SPA) of Zhang *et al.* also uses spatial averaging [11]. It chooses candidate pixels based on their orthogonal projection onto the subspace spanned by existing endmembers. Noise is reduced by averaging spectrally

similar target pixels within a local spatial window around the candidate.

Conversely, Zortea and Plaza use spatial information for endmember selection but draw each spectrum from the original pixel [12]. During selection of candidate endmember pixels, they weight each pixel's orthogonal projection based on its similarity to neighbors inside a local square image window. This biases the endmember selection process to favor pixels from homogeneous image regions. Finally, another spatial technique, Automated Morphological Endmember Extraction, modifies the definition of endmember to include spatial considerations in addition to spectral purity [13]. It defines a local morphological eccentricity index (MEI) for each pixel through local operators. It selects endmembers by segmenting the MEI image to find spectra that are both spatially and spectrally distinctive.

There are several ways superpixel representations differ from previous endmember detection methods that use spatial information. First, existing spatial endmember detection generally assume spatial relationships are constant throughout the image; they describe spatial similarity using a constant distance metric or membership to a fixed-size window. This is certainly appropriate when spectral variation is smooth and constant over the scene. In contrast, our segmentation uses pixel connectivity; it describes similarity relationships based on pairwise differences between neighboring pixels. This permits spatially adaptive influence; superpixels warp to cover the features of interest (Fig. 1). This could benefit natural scenes where features have arbitrary morphologies or sharp boundaries. Another difference with existing spatial endmember detection algorithms is modularity. In principle, superpixel representations could pair with virtually any statistical endmember extraction algorithm.

It is worth noting a related class of spatial smoothing strategies. Anisotropic diffusion [14]–[17] could also reduce image noise as a preprocessing step prior to endmember extraction. It obtains a smoothed image through a partial differential equation where the smoothing scale varies according to a locally varying diffusion conductance function of the base image. By basing diffusion conductance on the image gradient, one can denoise the image while preserving edges. Diffusion methods and superpixels have similar disadvantages and advantages for endmember extraction. Both can reduce spatially uncorrelated noise at the cost of intraclass variability and potentially spectral purity. However, they also differ in their treatment of spatial similarity. Most diffusion filters use a smooth differentiable function. This contrasts with superpixels' arbitrary morphologies and discrete all-or-nothing memberships.

A final benefit of superpixel representations *vis a vis* other spatial processing is that it decreases the number of spectra in the image. This diversity-preserving data reduction improves computational complexity both for the endmember detection stage or any additional analysis that follows. The segmentation operation is  $\mathcal{O}(n \log n)$  in the number of pixels [18] and time required for segmentation is generally far less than for later processing. In the experiments of Section IV, the speed benefit is sufficient to make an offline N-FINDER algorithm suitable for interactive use.




Fig. 2. Graph representation of pixels in a hyperspectral image. (Left) An example edge  $(v_i, v_j)$  connects vertices  $v_i$  and  $v_j$ . (Right) We consider joining segments  $S_a$  and  $S_b$ . We compare the minimum internal distance in their spanning trees  $\text{MInt}(S_a, S_b)$  with the minimum connecting distance  $\text{Dif}(S_a, S_b)$ . Here, the candidate edge with minimum distance is denoted by a dashed line. All vertices are part of some segment, although for clarity we label only  $S_a$  and  $S_b$ .

### III. SUPERPIXEL SEGMENTATION APPROACH

Superpixel representations split an image into a set of segments  $S$ , each containing one or more contiguous image pixels. The physical features also partition the image, and a valid superpixel segmentation should be a refinement of this partition. In other words, each superpixel should intersect at most a single scene feature. We will call this the refinement criterion. Note that this standard is impossible to satisfy for subpixel features. Otherwise, it is trivially satisfied whenever superpixels are reduced to the size of single-image pixels. The refinement criterion becomes more difficult to enforce as superpixel size increases.

Any segmentation of an image into contiguous regions could produce superpixels; candidate algorithms include watershed methods [19] or Markov Random Field Models [20]. Here, we use a graph-based approach based on the work of Felzenszwalb and Huttenlocher [18]. It has several advantages for our application:

- The resulting segments are spatially contiguous and relatively uniform in area.
- Computation is efficient and scales with  $n \log(n)$  for images of  $n$  pixels. Thus the algorithm is suitable for hyperspectral images of megapixel size or larger.
- It can accommodate any desired spectral divergence function (e.g., Euclidean distance, spectral angle).

The algorithm has been used for superpixels in three-channel images [21] and translates easily to the hyperspectral domain.

We represent the image of  $n$  pixels as an eight-connected graph of vertices  $V = \{v_i\}_{i=1}^n$ . Each vertex  $v_i$  is associated with a spectrum  $\rho_i$ , and edges  $(v_i, v_j)$  that connect the vertex to its neighbors (Fig. 2). These edges carry weights corresponding to a function  $d(v_i, v_j)$ , describing the divergence between the two spectra. An agglomerative clustering algorithm partitions the graph into segments  $S$ . Each pixel starts as a separate segment which trivially satisfies the refinement criterion. We merge neighboring subgraphs whenever there is no evidence of a boundary between the two regions. The Felzenszwalb algorithm [18] determines the boundary evidence by comparing the weight associated with segments' smallest connecting edge to the maximum internal weight in either subgraph. Following this previous work, we represent the maximum internal difference

of a segment  $S$  by the largest edge weight in the segment's minimum spanning tree  $\text{MST}(S)$ .

$$\text{Int}(S) = \max_{v_i, v_j} d(v_i, v_j) \quad \forall v_i \in S, \quad v_j \in S, \\ (v_i, v_j) \in \text{MST}(S) \quad (2)$$

For the set of all graph edges  $E$ , any adjacent superpixels  $S_a$  and  $S_b$  have a minimum connecting weight given by the minimum distance between candidate connecting edges

$$\text{Dif}(S_a, S_b) = \min_{v_i, v_j} d(v_i, v_j) \quad \forall v_i \in S_a, \quad v_j \in S_b, \\ (v_i, v_j) \in E. \quad (3)$$

A boundary between neighboring regions is preserved when this intersegment distance is larger than the minimum of either internal weight  $\text{Int}(S)$ , where  $\text{Int}(S)$  is the largest distance in the minimum spanning tree of a segment. This threshold is biased by a constant  $k$  and inversely proportional to a superpixel's area  $|S|$

$$\text{MInt}(S_a, S_b) = \min \left( \text{Int}(S_a) + \frac{k}{|S_a|}, \text{Int}(S_b) + \frac{k}{|S_b|} \right). \quad (4)$$

The bias  $k$  inflates the internal variability of the smallest regions, controlling the size of the resulting superpixels. A final step cleans any remaining small noise regions by merging segments less than a minimum size with their spectrally closest neighbors. Designers can trade speed for accuracy by changing this merging threshold to produce coarser or finer segmentations as desired (Fig. 1).

We tried two spectral divergence measures: 1) The Euclidean distance  $d_e$  across all wavelengths  $\lambda$ ; and (2) the spectral angle distance (SAD) given by  $d_{sa}$

$$d_e(v_i, v_j) = \sqrt{\sum_{\lambda} (\rho_{i\lambda} - \rho_{j\lambda})^2} \quad (5)$$

$$d_{sa}(v_i, v_j) = \sum_{\lambda} \frac{\rho_{i\lambda} \rho_{j\lambda}}{|\rho_i| |\rho_j|}. \quad (6)$$

After segmentation, we exclude any pixels with obvious bad data or instrument errors and compute the mean spectrum for each superpixel. This represents the population of member pixels for future processing. Assume that a feature is associated with a single vector of mixing coefficients  $a$ . If the refinement criterion is satisfied, the feature's superpixels are internally homogeneous and its member pixels share mixing coefficients. The component image pixels constitute IID measurements from a common distribution. For a pixel  $v_j \in S_i$ , this produces a revised sampling distribution

$$\rho'_j = \mathbf{X}_i = \frac{1}{|S_i|} \sum_{\rho_k \in S_i} \rho_k = \mathbf{Q} \mathbf{a}_j + \mathcal{N} \left( 0, \frac{\sigma^2}{|S_i|} \right). \quad (7)$$

We have transformed the original endmember detection task of (1) into an equivalent problem with a new noise variance reduced proportionally to the superpixel area. This limits outliers' influence on endmember extraction, and can improve spectral

fidelity since the recovered spectrum will be closer to the real noise-free reflectances.

This noise reduction comes with a risk. If the refinement criterion is not satisfied, or scene features are not homogeneous, then the new representation will degrade endmembers' spectral purity with respect to the original data set. This tradeoff is probably justified in planetary geology where instrument and atmospheric anomalies are pervasive. Planetary geologists report spectral averages from small contiguous regions, accepting a reduction in purity to improve detection confidence and interpretability [22]–[25].

Violations of the refinement criterion would also impact fractional abundance estimation for raw pixels because superpixel endmembers would be less “convex” in the original image. For this reason, we advise caution in using superpixel endmembers for unmixing raw pixels, particularly when the targets of interest are likely to be small. Note that superpixel endmembers are still valid endmembers within the superpixel image, so abundance estimation within this representation is still meaningful.

#### IV. EVALUATION METHODOLOGY

We evaluate pixel- and superpixel-based endmember detection for a guided search task. Here, an analyst browsing a large image database desires a fast summary of scene constituents to identify outliers and trends. A good automated system would detect the same features that would have been found from an exhaustive manual search. Our tests use a list of manually selected target regions as a performance standard. We evaluate superpixel-augmented endmember detection using an AVIRIS image of Cuprite [9], as well as images from the Compact Reconnaissance Imaging Spectrometer (CRISM) aboard the Mars Reconnaissance Orbiter [6].

##### A. Data Sets

Our first data set consists of AVIRIS f970619t01p02, an overflight of Cuprite Nevada. This region is well studied both geologically and through hyperspectral remote imaging [26], [27]. We use wavelengths from 1.0 to 2.5  $\mu\text{m}$ , excluding absorption bands at 1.9 and 1.35  $\mu\text{m}$ . SNR is estimated to vary between 1:400 and 1:1000 in this spectral range [28]. The scene typifies terrestrial high-resolution low-noise hyperspectral imagers. An expert analyst performed a manual study to find the purest examples of minerals identified previously by Rowan *et al.* [29]. The analyst labeled regions of interest comprised of several tens to hundreds of pixels using techniques documented by Gilmore [25]. These classes often included multiple spatially separated regions.

We used the mean spectrum of each class as a target for automatic endmember detection. Fig. 3 shows an example—the “strong calcite” class. Red pixels in the insert show the manually identified region of interest. Blue pixels show the corresponding endmember superpixel identified automatically by N-FINDR. A wide range of minerals have been found at Cuprite; we limited our evaluation to the dominant mineral classes in the Rowan *et al.*'s study and with distinguishing features apparent in the 1.0–2.5  $\mu\text{m}$  window.

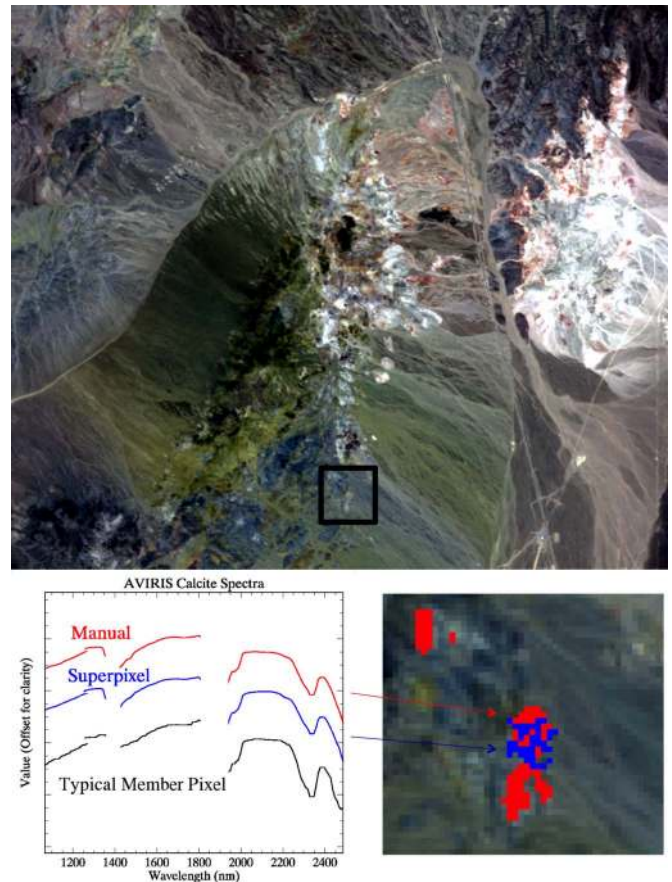


Fig. 3. (Top) AVIRIS image f970619t01p02\_r02\_sc04 of Cuprite, Nevada, covering a region 10 km at 19 m/pixel ( $R$ : 636.07 nm,  $G$ : 557.07 nm,  $B$ : 468.31 nm). (Bottom) The “Calcite Strong” region from the geologist’s target list, and an endmember superpixel detected by N-FINDR. Holes in the region are due to the eight-connectivity of the segmentation graph. A spectrum drawn from a typical member pixel appears in black.

We also consider images from the CRISM instrument which is typical of planetary science data having high noise and relatively low spatial resolution. CRISM offers multispectral survey and hyperspectral targeted modes. The hyperspectral mode yields a megapixel-scale image covering 1.0 to 4.0  $\mu\text{m}$  with over 400 measurement channels at approximately 18 m per pixel. Full Width at Half Maximum is 7.9–10.1 nm across the VNIR range [6]. This spatial and spectral resolution is inferior to the best terrestrial instruments but unprecedented for the Mars surface. We examine three well-studied full resolution CRISM scenes: 1) 3e12; 2) 3fb9; and 3) 863e (omitting the FRT0000 catalog prefix). These contain diverse spectra consistent with olivine, phyllosilicate, and carbonate minerals [22]–[24]. We also consider scene 8158 where recent studies have found spectra consistent with sulfate deposits [25].

Several characteristics of CRISM data make it a good test case for guided search. The actual minerals present on the surface are unknown which favors an unsupervised approach. Also, CRISM’s data volumes outpace scientists’ capacity for exhaustive manual study. CRISM will return over a terabyte of data to earth over the course of its mission, and automated summary analysis could benefit investigations by identifying novel mineralogy. Note that CRISM exhibits high noise relative to terrestrial instruments. Estimated instrument SNR from VNIR

laboratory calibration data, for average material at a  $30^\circ$  phase angle, is 425. However, in practice, spectral features are quite subtle or obscured by instrument noise, atmospheric noise, or surface dust. They may be present in just a few contiguous pixels; these smallest features are easily masked and challenge the limits of spatial noise reduction. In the evaluations that follow, we will reference AVIRIS and CRISM as representative examples of high and low noise extremes.

The CRISM spectra cover infrared wavelengths from 1.0 to  $2.5 \mu\text{m}$ , avoiding thermal emission present in the spectrum longward of  $2.5 \mu\text{m}$ . The Brown CRISM Analysis Tool applies radiometric and atmospheric corrections to each image [30]. Atmospheric absorption correction uses a standard approach based on spectral observations of the Olympus Mons volcano at different altitudes [31]. In all data sets, we apply a radius 1 median filter in the spectral domain to reduce single-shot noise.

After atmospheric correction, photometric correction, and preprocessing, the CRISM labeling procedure is the same as the AVIRIS data set. The geologist identified the primary constituents in each image along with the image pixels containing the purest examples of each mineral. Sometimes, constituent spectra could be identified as a single mineral, while others are mineral mixtures with a dominant mineral. Such mixed constituents are likely to be common in moderate resolution data on planetary surfaces. Identification of the dominant mineralogy within these pixels is geologically important and an optimal result given the limitations of the instrument. Fig. 4 shows CRISM image 3fb9 along with manually labeled “kaolinite” regions. High noise in the single-pixel spectrum underscores the ambiguity of these nonaveraged observations.

## B. Methods

We choose segmentation parameters to minimize spectral angle error for the CRISM 3fb9 image and apply these parameters universally for all tests. Appendix discusses the tests in greater detail. The designer can affect superpixels through the choice of distance metric, the minimum region size, and the bias  $k$ . Of these, only the minimum region size has a clear consistent effect on performance. The bias parameter appears quite stable; one can still vary it far enough to cause catastrophic failure, but short of this, it has little effect on segmentation size or morphology. In the experiments that follow, we use  $k = 0.001$  for all images. Additionally, the different spectral divergence functions yield no obvious performance difference and we use the SAD measure for all segmentations. We found that the best minimum size parameter was different for N-FINDR and SMACC. We use a minimum region size of 15 pixels for N-FINDR which produced superpixel areas from 15 to over 100, similar to the fine segmentation of Fig. 1. SMACC favors slightly larger regions, and we use a minimum size parameter of 50 (Appendix).

After segmentation, we test both SMACC and N-FINDR algorithms on pixel and superpixel representations. All preprocessing is the same for both data sets aside from the segmentation step. To avoid local minima with N-FINDR, we use ten random restarts and keep the best scoring result. Fig. 5 shows N-FINDR endmembers from the 3fb9 image. The first

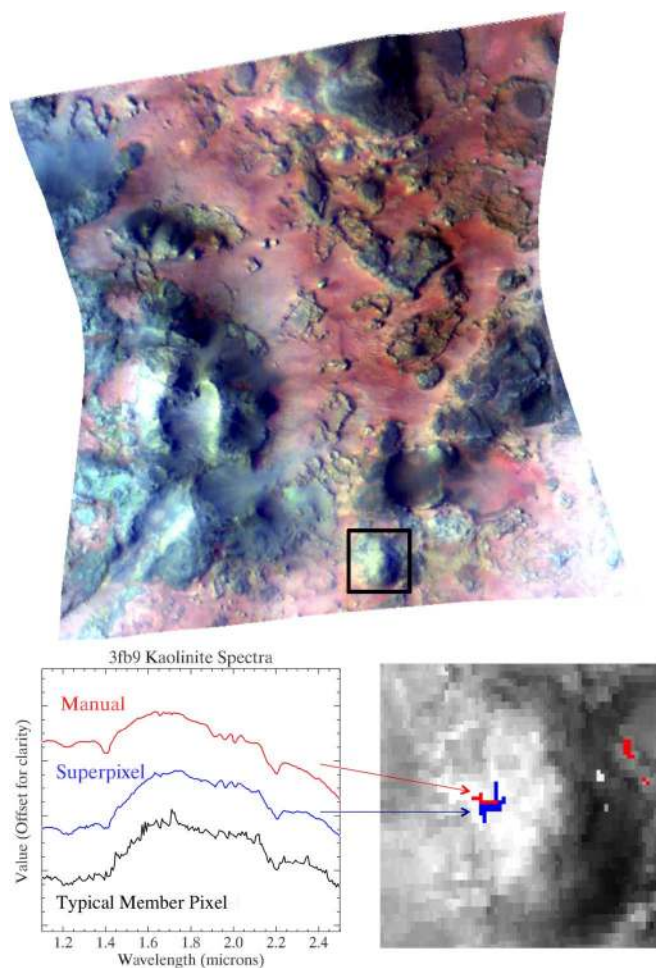


Fig. 4. CRISM image 3fb9, covering 10 km at 18 m/pixel ( $R : 2.0 \mu\text{m}$ ,  $G : 1.5 \mu\text{m}$ ,  $B : 1.1 \mu\text{m}$ ). The insert shows specific pixels identified by the geologist as belonging to the “kaolinite” class (in red). N-FINDR identifies the blue superpixel as an endmember; this region partially intersects the geologist’s pixel selections and the resulting spectrum is a close match. Pixels taken individually have ambiguous classifications due to noise.

column displays target mineralogical classes reported in the literature and corroborated with the exhaustive manual study. We use general descriptive terms where exact classifications are uncertain. The targets include carbonate, kaolinite, olivine, and a phyllosilicate. We also include a neutral (i.e., mineralogically bland) lava spectrum. The center and right columns show the best matching endmembers obtained from both superpixel and (median-filtered) pixel representations. Noise reduction benefits from segmentation are particularly evident in the detected phyllosilicate, kaolinite, and lava classes.

## C. Performance Metrics

The number of targets differs between images raises the question of how many endmembers to request. Analysts often estimate this number directly from using the intrinsic dimensionality or similar measures [32]. In guided search, the primary goal is not to produce spectra for abundance estimation but rather to detect interpretable features of interest for manual followup. Thus, the appropriate number of detections also depends on our tolerance for recall error and the cost

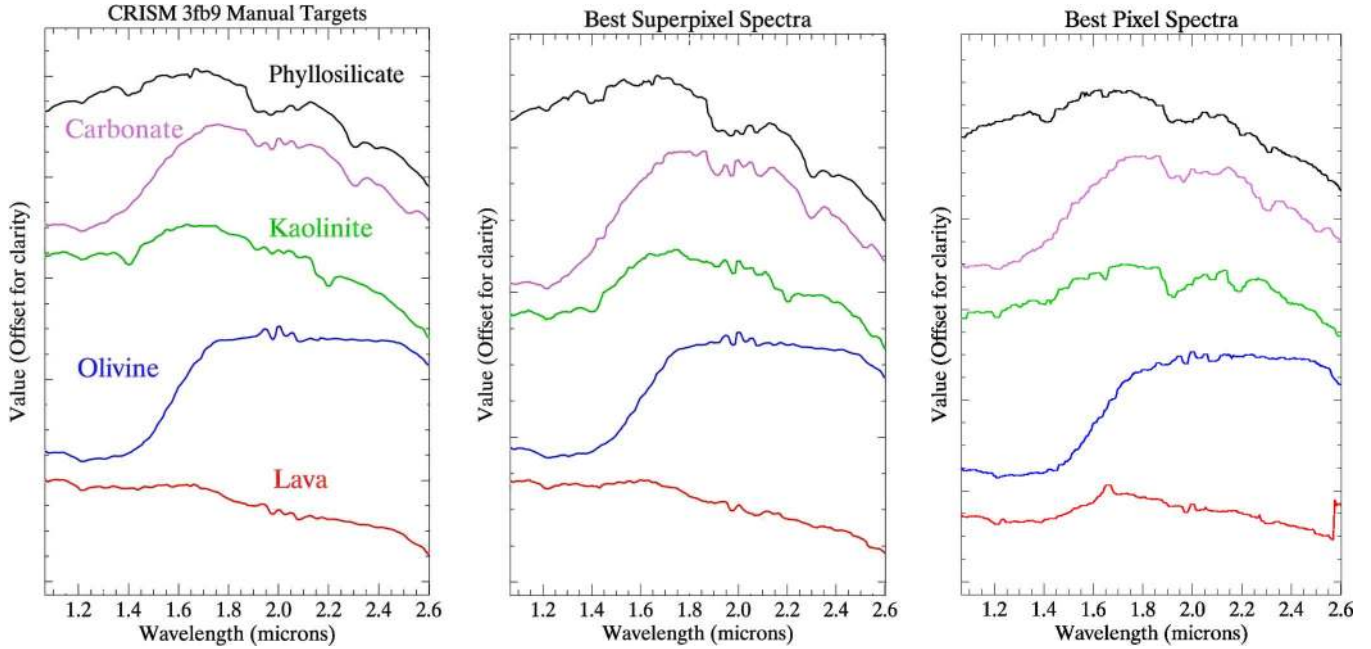


Fig. 5. Spectral classes derived from the CRISM 3fb9 image, with wavelengths in micrometers. (Left) Manual spectra identified by an expert. (Center) Superpixel-based endmember extraction. (Right) Endmembers from traditional pixel-based endmember detection. Note the loss of fidelity in the Phyllosilicate, Kaolinite, and Lava classes.

of the analysts' time to investigate each result. Analysts with little time available would opt for smaller sets that minimize redundancy. Conversely, an analyst free to investigate many detections would prefer larger sets that are more certain to include all relevant materials. For a set of endmembers  $\mathcal{E}$  and target list  $\mathcal{T}$ , we hypothesize a total score  $f(\mathcal{E}, \mathcal{T})$  that is a weighted combination of terms corresponding to spectral accuracy  $Acc(\mathcal{E}, \mathcal{T})$ , recall rate  $Rec(\mathcal{E}, \mathcal{T})$ , and the time cost of followup analysis  $Tim(\mathcal{E})$ . Here, time cost is related to the total number of targets detected but might also include computation time.

$$f(\mathcal{E}, \mathcal{T}) = \alpha_1 Acc(\mathcal{E}, \mathcal{T}) + \alpha_2 Rec(\mathcal{E}, \mathcal{T}) + \alpha_3 Tim(\mathcal{E}) \quad (8)$$

Rather than impose specific weighting values, we will compare recall and accuracy scores at many endmember list sizes. This reveals the envelope of possibilities for different accuracy/time tradeoffs. Our geologist identified 20 endmembers as the ceiling for a mineralogical search task, so we evaluate the detection process for list sizes ranging from 3–19 endmembers.

We measure spectrum accuracy by computing each target mineral's spectral angle distance to each spectrum in the retrieved endmember list. We report the mean spectral angle of the best matches. We also computed RMS errors but these differences were virtually identical to the spectral angle results across all images and provide no extra information. We report SAD exclusively for clarity.

We measure the recall score with a metric that matches each detection against all possible targets in the detection list. To count as having been "detected," a target must have an associated endmember that is more similar in spectral angle to that target than to any other. The recall improves as more endmembers are requested, so we report the first endmember list size at which each target signal appears. Together, the recall




Fig. 6. SAD error of best matching endmembers for different numbers of requested endmembers. The mean score over all images and targets is shown. Translucent bands show 95% confidence intervals assuming Gaussian-distributed errors.

and accuracy scores describe both the ability of the endmember detection to find all the scene's constituents and the fidelity of retrieved spectra to these targets.

## V. RESULTS

Fig. 6 shows the mean spectral accuracy score for all images with different numbers of requested endmembers. Translucent bars show 95% confidence intervals for the mean, presuming Gaussian error. This is a tenuous assumption but it provides some intuition about the spread of the data. We chart performance for both N-FINDER and SMACC; neither offers a

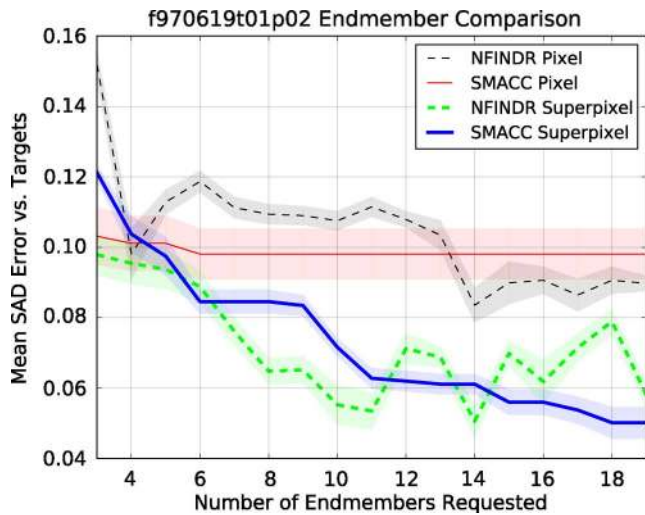


Fig. 7. SAD error of best matching endmembers for different numbers of requested endmembers (AVIRIS Cuprite image).

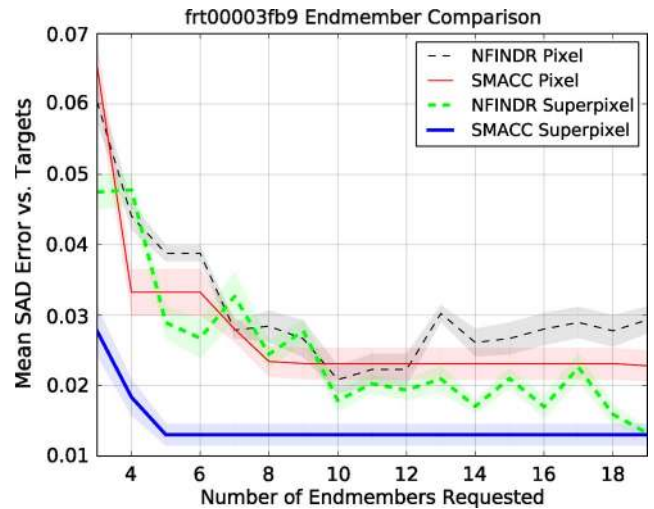


Fig. 9. SAD error of best matching endmembers for different numbers of requested endmembers (CRISM image 3fb9).

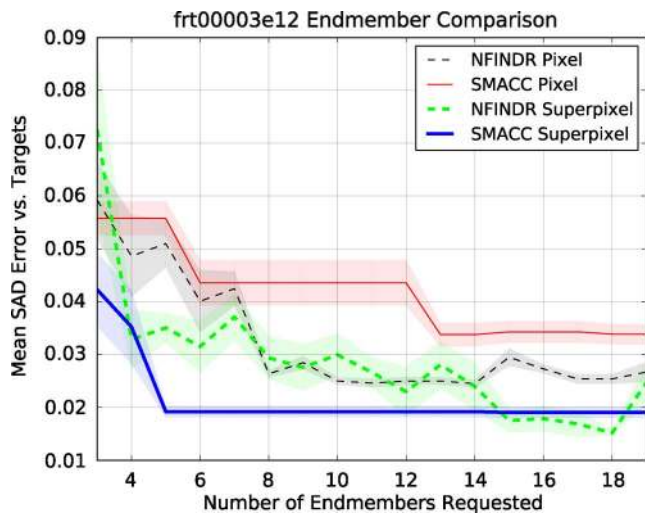


Fig. 8. SAD error of best matching endmembers for different numbers of requested endmembers (CRISM image 3e12).

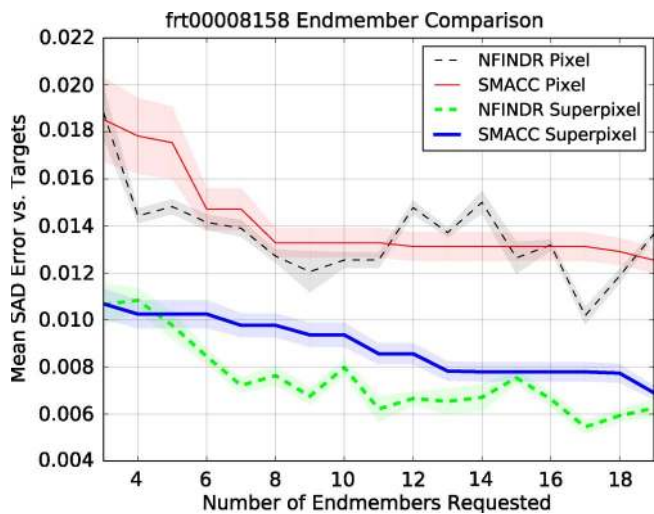


Fig. 10. SAD error of best matching endmembers for different numbers of requested endmembers (CRISM image 8158).

conclusive benefit over the other for this data set. However, on average, superpixel representations improve both methods and benefit the SAD accuracy score by factors of 20–25%. Note that the SMACC score is monotonic decreasing because it is sequential and each additional spectrum can only improve the best matches. In contrast, the simultaneous N-FINDR algorithm is more unstable and increasing the set size occasionally hurts accuracy by dropping previous detections. In general, accuracy improves as more endmembers are requested.

Scores from individual AVIRIS and CRISM scenes appear in Figs. 7–11. Benefits appear most significant for those scenes with larger targets, such as the AVIRIS image and CRISM 8158. CRISM 3fb9 and 3e12 have several small and disjoint features for which the refinement criterion is more difficult to satisfy. On these images, there is a range of list sizes for which superpixels provide no performance benefit over pixel representations for the N-FINDR method. In general, however, superpixel combinations perform at least as well as, or better than, the pixel representations.

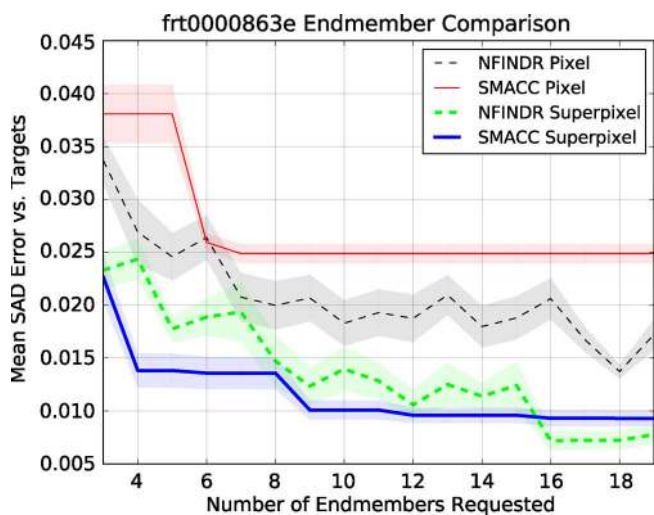


Fig. 11. SAD error of best matching endmembers for different numbers of requested endmembers (CRISM image 863e).

TABLE I

SPECTRAL ANGLE DISTANCE BETWEEN THE TARGET AND ITS CLOSEST DETECTED ENDMEMBER. THE PIXEL SIZE OF EACH TARGET APPEARS IN PARENTHESIS. THIS CHART SHOWS THE RESULT AFTER REQUESTING THE MAXIMUM NUMBER OF ENDMEMBERS (19 FOR OUR TESTS)

Target (size)	SMACC		N-FINDR	
	Pixel	Superpixel	Pixel	Superpixel
<b>AVIRIS</b>				
Alunite (437)	<b>0.0137</b>	0.0874	0.0732	0.0430
Calcite Strong (572)	0.0710	<b>0.0153</b>	0.0917	0.0490
Calcite Weaker Bands (504)	0.0697	<b>0.0398</b>	0.1022	0.0535
Jarosite/Alunite (55)	0.1069	<b>0.0440</b>	0.0964	0.0715
Kaolinite (438)	0.1350	0.0825	0.0689	<b>0.0640</b>
Muscovite (425)	0.1391	<b>0.0560</b>	0.1011	0.0996
Neutral Region (322)	0.0273	0.0259	0.0944	<b>0.0252</b>
<b>CRISM 3e12</b>				
Lava (13241)	0.0438	0.0228	0.0227	<b>0.0158</b>
Magnesite (106)	0.0222	<b>0.0215</b>	0.0222	0.0538
Olivine (1806)	0.0314	0.0202	0.0406	<b>0.0147</b>
Phyllosilicate (2290)	0.0379	<b>0.0114</b>	0.0211	0.0142
<b>CRISM 3fb9</b>				
Carbonate (183)	0.0093	<b>0.0058</b>	0.0339	0.0129
Kaolinite (54)	0.0260	0.0235	0.0205	<b>0.0189</b>
Neutral Region (3518)	0.0334	<b>0.0113</b>	0.0418	<b>0.0113</b>
Olivine (1403)	0.0330	0.0191	0.0327	<b>0.0134</b>
Phyllosilicate (333)	0.0120	<b>0.0052</b>	0.0177	0.0095
<b>CRISM 8158</b>				
Neutral Region (7822)	0.0135	<b>0.0069</b>	0.0129	0.0102
Polyhydrated Sulfate Bright (4479)	<b>0.0029</b>	0.0044	0.0153	<b>0.0029</b>
Kieserite Area A (183)	0.0100	0.0077	0.0145	<b>0.0076</b>
Phyllosilicate (15257)	0.0169	0.0088	0.0122	<b>0.0051</b>
Kieserite Area B (328)	0.0121	0.0061	0.0143	<b>0.0058</b>
Kieserite Area C (738)	0.0120	0.0086	0.0126	<b>0.0074</b>
Polyhydrated Sulfate Dark (17247)	0.0075	0.0097	0.0129	<b>0.0032</b>
Shadow (1096)	0.0139	<b>0.0061</b>	0.0143	0.0074
<b>CRISM 863e</b>				
FeMgSmectite (640)	0.0259	0.0078	0.0126	<b>0.0070</b>
Kaolinite (13)	0.0225	<b>0.0061</b>	0.0207	0.0152
Montmorillonite (133)	0.0202	<b>0.0052</b>	0.0122	0.0077
Neutral Region(4468)	0.0232	0.0065	0.0113	<b>0.0046</b>
Nontronite (98)	0.0326	0.0119	0.0292	<b>0.0040</b>

Table I compares the spectrum accuracy for each target class. Columns show scores for pixel- and superpixel-based methods requesting the maximum number of endmembers. The leftmost column lists a tentative classification for each of the geologist's targets. We use general descriptive terms where the classification is uncertain. The total pixel size of each labeled region appears in parenthesis; it ranges widely with the smallest deposits having only a few tens of pixels. The top-performing combination for each row appears in bold. Here, again, results suggest an advantage to superpixel methods; a pixel-based method scores best for just one of the targets in the data set and ties for another. There are no obvious patterns across scenes or region sizes that predict whether SMACC or N-FINDR will outperform.

We also evaluate run time for each method. SMACC is quite fast and computation speed is not a significant factor to analysis time. In contrast, the N-FINDR routine requires multiple runs of a computationally demanding combinatorial optimization. Fig. 12 shows how N-FINDR run times vary with the number of endmembers requested, using a modern single-core desktop processor. A typical run of N-FINDR on each pixel representation took several hours. Superpixel representations execute in several minutes including the cost of segmentation, which makes the N-FINDR strategy feasible for use in an interactive session.

We compare spectral accuracy for different SNR levels. Fig. 13 shows retrieved spectrum accuracy for an image corrupted by synthetic instrument noise. We start with the AVIRIS Cuprite scene (which has a high intrinsic SNR) and add additional zero-




Fig. 12. Runtime for the N-FINDR algorithm increases with the number of requested endmembers. This image shows run times for CRISM 3fb9. (Black line) Superpixel. (Gray line) Pixel.

f970619t01p02 Endmember Comparison for Ubiquitous Noise

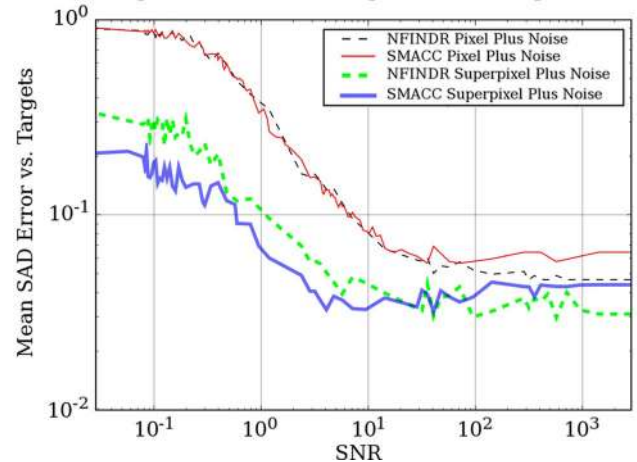


Fig. 13. SAD error of best matching endmembers for the AVIRIS Cuprite image with various simulated SNRs. We apply Gaussian noise across all hyperspectral channels. Note the log scale.

mean Gaussian noise of increasing variance to all channels. This simulates more difficult imaging conditions. The graph shows SAD error score for a midrange list size of ten requested endmembers and the standard minimum region sizes. Superpixels shift the curve both downward and to the left. This suggests relative improvements in overall accuracy as well as improved margin of resistance to small amounts of noise.

Fig. 14 shows a similar test for impulse noise. Here, each pixel has a small probability of corruption in a single random spectral channel. The alteration is additive noise with a Gaussian distribution of mean zero and standard deviation equal to three times the mean intensity of the image. This noise profile is more typical of CRISM and is particularly damaging to pixel-based representations. Relative benefits of spatial averaging become more pronounced as the pulse probability increases. Note that the SMACC version, which uses a minimum region size of 50, yields slightly larger pixels than the N-FINDR runs. The larger averaging regions are the top performers at very low SNR for both impulse and uniform noise, corroborating our intuition from (7).

Our hypothetical cost function includes a term describing recall, or the ability to find all the significant targets in a scene. We compare recall scores using the pairwise matching method defined in Section IV-C. Table II reports the smallest end-member set size for which each target spectrum was detected.



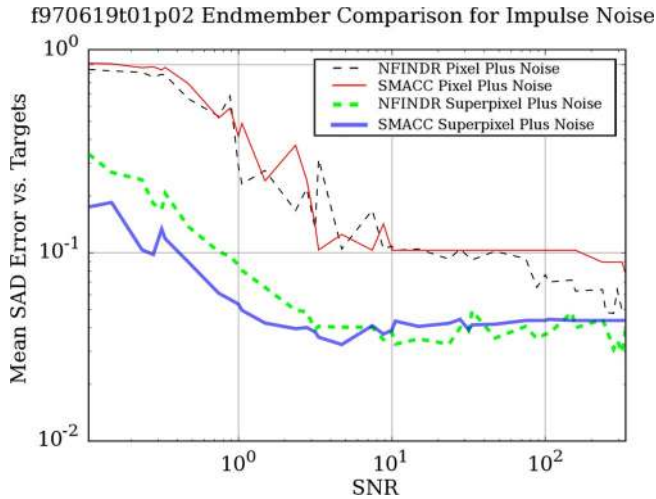


Fig. 14. SAD error of best matching endmembers for the AVIRIS Cuprite image with various simulated SNRs. Each pixel has a small probability of corruption by impulse noise of Gaussian-distributed amplitude. Note the log scale.

TABLE II

TARGET RECALL PERFORMANCE USING IMAGE PIXELS AND SUPERPIXELS. WE REPORT THE LOWEST ENDMEMBER LIST SIZE AT WHICH EACH TARGET APPEARS AS A DETECTION (LOW SCORES ARE BETTER). A SUCCESSFUL DETECTION REQUIRES THAT SOME ENDMEMBER BE A BETTER MATCH TO THAT TARGET THAN ANY OTHER TARGET IN THE LIST. A DASH '-' INDICATES TARGETS THAT WERE NOT DISCOVERED

Target (size)	SMACC		N-FINDR	
	Pixel	Superpixel	Pixel	Superpixel
<b>AVIRIS</b>				
Alunite (437)	3	3	3	6
Calcite Strong (572)	-	17	7	3
Calcite Weaker Bands (504)	14	4	13	8
Jarosite/Alunite (55)	13	10	6	4
Kaolinite (438)	11	-	5	5
Muscovite (425)	3	4	3	5
Neutral Region (322)	3	3	3	3
<b>Image 3e12</b>				
Lava (13241)	3	4	3	3
Magnesite (106)	6	5	7	5
Olivine (1806)	3	3	3	4
Phyllosilicate (2290)	3	3	3	3
<b>Image 3fb9</b>				
Carbonate (183)	3	3	3	3
Kaolinite (54)	10	10	4	3
Neutral Region (3518)	5	5	7	3
Olivine (1403)	3	3	3	5
Phyllosilicate (333)	4	3	3	18
<b>Image 8158</b>				
Neutral Region (7822)	3	5	3	5
Polyhydrated Sulfate Bright (4479)	3	3	3	3
Kieserite Area A (183)	-	3	3	3
Phyllosilicate (15257)	17	5	6	5
Kieserite Area B (328)	3	4	7	4
Kieserite Area C (738)	-	9	-	9
Polyhydrated Sulfate Dark (17247)	6	3	4	3
Shadow (1096)	4	9	5	9
<b>Image 863e</b>				
FeMgSmectite (640)	3	5	3	5
Kaolinite (13)	11	3	6	3
Montmorillonite (133)	6	8	4	8
Neutral Region(4468)	3	3	3	3
Nontronite (98)	3	3	3	3

For example, for image 3e12, the superpixel/SMACC approach first detects the magnesite class when four endmembers are requested. Five endmembers are required to reveal the lava class. Every trial requests at least three endmembers, so this is the minimum number in the table. The top-performing combination

for each target appears in bold. A missed target, denoted by a dash, does not preclude a good spectral angle score (in Table I) since a detected endmember can be spectrally close to a target that is not its best pairwise match.

Neither superpixel nor pixel representations provide a clear advantage for recall scores with either data set. Only the N-FINDR/Superpixel combination finds all of the targets in every image, but the pixel-based N-FINDR also performs well and the recall difference may not be significant. While we had anticipated that small or highly disjoint regions would be the most difficult to find, spectral distinctiveness appears to have a larger effect. The Kieserite C class of image 8158 contains over 738 pixels, but it is quite similar to other endmembers and only the superpixel combinations detect it.

## VI. CONCLUSION

Superpixel endmember detection holds promise for fast summary analysis of large hyperspectral image catalogs. Automatic detection methods to find subtle mineralogical signatures can alleviate a main bottleneck to the science return of planetary missions and enable new and more comprehensive investigations. The reduction in data set size could permit more sophisticated processing on the ground, or automated data analysis by remote spacecraft.

No planetary data set yet approaches the SNR, atmospheric control, or lack of dust evidenced in the AVIRIS Cuprite scene. But these tests suggest that superpixel representations could benefit both domains in several aspects:

- spectral interpretability due to a better match with human-selected features;
- detection confidence due to combining multiple pixels;
- computational complexity for endmember detection and additional processing that follows.

The main disadvantage is that if the refinement criterion is not satisfied, small spectral features may be masked. In this case, the detected endmembers will not be suitable for abundance estimation of the raw pixels.

Future evaluations could quantify drawbacks for the smallest single-pixel or subpixel targets where spatial information is not expected to be useful. It would also be helpful to test the proposed approach against alternative smoothing methods (e.g., diffusion) and other spatially-sensitive endmember detection methods. There is also scope for improving algorithm performance. Learned or application-specific distance metrics might produce superior superpixels. Entirely new segmentation strategies such as a normalized cuts approach [13] might yield different or better results. Finally, alternatives to the sample mean such as robust estimators might improve recovery of features' true spectra.

## APPENDIX

### SETTING THE SEGMENTATION PARAMETER

Superpixels that satisfy the refinement criterion must be at least as large as the scene features. Actual superpixel sizes vary over the image; they are dependent both on segmentation parameters and the data. Here, we determine optimal

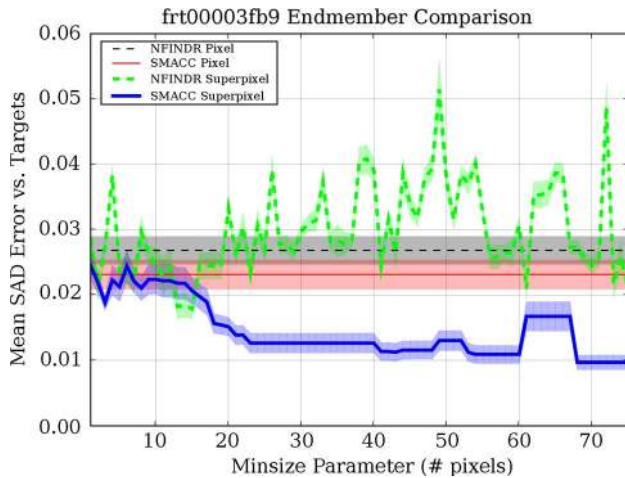


Fig. 15. SAD error of best matching endmembers across different minimum sizes (CRISM image 3fb9).

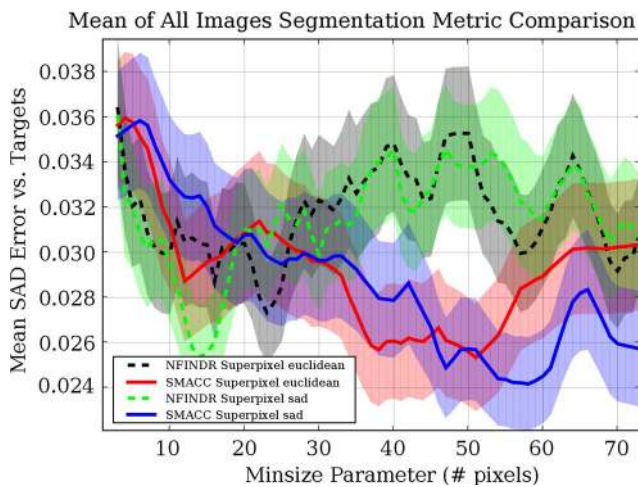


Fig. 16. SAD error of best matching endmembers across different values of the minimum size parameter (average over all images). We show the results of segmentations using both SAD and Euclidean distance metrics.

segmentation parameters empirically. There was no consistent difference in segmentation quality or overall performance between the two distance measures (Euclidean and SAD), so we use the SAD measure exclusively in the experiments. The segmentation is also robust to different choices of  $k$ , and a value of  $k = 0.001$  works well for all images we have analyzed. This leaves the minimum region size as the significant free parameter. This parameter is the dominant factor controlling the resolution of the final segmentation and accounts for all the differences between the three panels of Fig. 1.

We establish an optimal minimum size parameter to maximize performance on image 3fb9, the first in our study. Fig. 15 charts mean SAD scores as a function of the minimum size parameter for this image. The best minimum size appears different for N-FINDR and SMACC algorithms. For N-FINDR, maximum performance was achieved at a value of 15 where it outperformed the pixel-based representations. SMACC performed better for slightly larger regions, reflecting conventional wisdom that this method is more sensitive to noise. The local minimum appeared to be near 50 pixels. We used these values exclusively in all the experiments of Section IV.

Fig. 16 shows the overall average over the entire data set. We apply a local running average window of 5 to smooth the graph for clarity. While N-FINDR performance is somewhat unstable, it consistently shows a local performance peak near a minimum size near 15. A minimum size value of 50–60 is more appropriate for SMACC. Once the analyst has chosen an endmember detection algorithm, they could use a single minimum size value across different scenes during automated guided search. Fig. 16 also shows segmentation performance with both Euclidean and SAD divergence metrics.

#### ACKNOWLEDGMENT

The authors would like to thank the CRISM team and Brown University for the use of the CRISM Analysis Tool (CAT) and their software contributions to the community. The work described in this paper was carried out at the Jet Propulsion Laboratory with support from the NASA AMMOS Multimission Ground Systems and Services office.

#### REFERENCES

- [1] X. Ren and J. Malik, "Learning a classification model for segmentation," in *Proc. Int. Conf. Comput. Vis.*, 2003, pp. 10–17.
- [2] A. Moore, S. Prince, J. Warrell, U. Mohammed, and G. Jones, "Superpixel lattices," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2008, pp. 1–8.
- [3] G. Mori, X. Ren, A. Efros, and J. Malik, "Recovering human body configurations: Combining segmentation and recognition," in *Proc. CVPR*, 2004, vol. 2, pp. II-326–II-333.
- [4] D. R. Thompson, R. Castano, and M. S. Gilmore, "Sparse superpixel unmixing for exploratory analysis of crism hyperspectral images," in *Proc. WHISPERS*, 2009, pp. 1–4.
- [5] R. Green, M. Eastwood, C. Sarture, T. Chrien, M. Aronsson, B. Chippendale, J. Faust, B. Pavri, C. Chovit, M. Solis, M. Olah, and O. Williams, "Imaging spectroscopy and the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)," *Remote Sens. Environ.*, vol. 65, no. 3, pp. 227–248, Sep. 1998.
- [6] S. Murchie, R. Arvidson, P. Bedini, K. Beisser, J.-P. Bibring, J. Bishop, J. Boldt, P. Cavender, T. Choo, R. T. Clancy, E. H. Darlington, D. Des Marais, R. Espiritu, D. Fort, R. Green, E. Guinness, J. Hayes, C. Hash, K. Heffernan, J. Hemmler, G. Heyler, D. Humm, J. Hutcheson, N. Izenberg, R. Lee, J. Lees, D. Lohr, E. Malaret, T. Martin, J. A. McGovern, P. McGuire, R. Morris, J. Mustard, S. Pelkey, E. Rhodes, M. Robinson, T. Roush, E. Schaefer, G. Seagrave, F. Seelos, P. Silverglate, S. Slavney, M. Smith, W.-J. Shyong, K. Strohhenn, H. Taylor, P. Thompson, B. Tossman, M. Wirzburger, and M. Wolff, "Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) on Mars Reconnaissance Orbiter," *J. Geophys. Res.*, vol. 112, no. E5, p. E05 S03, 2007.
- [7] A. Plaza, P. Martinez, R. Pérez, and J. Plaza, "A quantitative and comparative analysis of endmember extraction algorithms from hyperspectral data," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 3, pp. 650–663, Mar. 2004.
- [8] J. Gruninger, A. Ratkowski, and M. Hoke, "The sequential maximum angle convex cone (SMACC) endmember model," in *Proc. SPIE—Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery X*, 2004, vol. 5425, pp. 1–14.
- [9] M. Winter, "N-FINDR: An algorithm for fast autonomous spectral endmember determination in hyperspectral data," *Proc. SPIE*, vol. 3753, pp. 266–275, 1999.
- [10] D. M. Rogge, B. Rivard, J. Zhang, J. Harris, A. Sanchez, J. Harris, and J. Feng, "Integration of spatial-spectral information for the improved extraction of endmembers," *Remote Sens. Environ.*, vol. 110, no. 3, pp. 287–303, Oct. 2007.
- [11] J. Zhang, B. Rivard, and D. Rogge, "The Successive Projection Algorithm (SPA), an algorithm with a spatial constraint for the automatic search of endmembers in hyperspectral data," *Sensors*, vol. 8, no. 2, pp. 1321–1342, 2008.

- [12] M. Zortea and A. Plaza, "Spatial preprocessing for endmember extraction," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 8, pp. 2679–2693, Aug. 2009.
- [13] A. Plaza, P. Martinez, R. Perez, and J. Plaza, "Spatial/spectral endmember extraction by multidimensional morphological operations," *IEEE Trans. Geosci. Remote Sens.*, vol. 40, no. 9, pp. 2025–2041, Sep. 2002.
- [14] M. Lennon, G. Mercier, and L. Hubert-Moy, "Nonlinear filtering of hyperspectral images with anisotropic diffusion," in *Proc. Int. Geosci. Remote Sens. Symp.*, 2002, pp. 2477–2479.
- [15] S. Velasco-Forero and V. Manian, "Improving hyperspectral image classification using spatial preprocessing," *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 2, pp. 297–301, Apr. 2009.
- [16] J. Martin-Herrero, "Anisotropic diffusion in the hypercube," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 5, pt. 2, pp. 1386–1398, May 2007.
- [17] J. Duarte-Carvajalino, P. Castillo, and M. Velez-Reyes, "Comparative study of semi-implicit schemes for nonlinear diffusion in hyperspectral imagery," *IEEE Trans. Image Process.*, vol. 16, no. 5, pp. 1303–1314, May 2007.
- [18] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *Int. J. Comput. Vis.*, vol. 59, no. 2, pp. 167–181, Sep. 2004.
- [19] Y. Tarabalka, J. Chanussot, and J. Benediktsson, "Segmentation and classification of hyperspectral data using watershed transformation," *Pattern Recognit.*, vol. 43, no. 7, pp. 2367–2379, Jul. 2010.
- [20] A. Mohammadpour, O. Féron, and A. Mohammad-Djafari, "Bayesian segmentation of hyperspectral images," in *Proc. Bayesian Inference Maximum Entropy Methods Sci. Eng.*, 2004, vol. 735, pp. 541–548.
- [21] D. Hoiem, A. Efros, and M. Hebert, "Automatic photo pop-up," in *Proc. ACM SIGGRAPH*, 2005, vol. 24, no. 3, pp. 577–584.
- [22] J. Mustard, S. Murchie, S. Pelkey, B. Ehlmann, R. Milliken, J. Grant, J. Bibring, F. Poulet, J. Bishop, E. Dobrea, L. Roach, F. Seelos, R. Arvidson, S. Wiseman, R. Green, C. Hash, D. Humm, E. Malaret, J. McGovern, K. Seelos, T. Clancy, R. Clark, D. Marais, N. Izenberg, A. Knudson, Y. Langevin, T. Martin, P. McGuire, R. Morris, M. Robinson, T. Roush, M. Smith, G. Swayze, H. Taylor, T. Titus, and M. Wolff, "Hydrated silicate minerals on Mars observed by the Mars Reconnaissance Orbiter CRISM instrument," *Nature*, vol. 454, no. 7202, pp. 305–309, Jul. 2008.
- [23] J. L. Bishop, E. Z. Dobrea, N. K. McKeown, M. Parente, B. L. Ehlmann, J. R. Michalski, R. E. Milliken, F. Poulet, G. A. Swayze, J. F. Mustard, S. L. Murchie, and J.-P. Bibring, "Phyllosilicate diversity and past aqueous activity revealed at mawrth vallis, Mars," *Science*, vol. 321, no. 5890, pp. 830–833, Aug. 2008.
- [24] B. Ehlmann, J. Mustard, S. Murchie, F. Poulet, J. Bishop, A. Brown, W. Calvin, R. Clark, D. Marais, R. Milliken, L. Roach, T. Roush, G. Swayze, and J. Wray, "Orbital identification of carbonate-bearing rocks on Mars," *Science*, vol. 322, no. 5909, pp. 1828–1832, Dec. 2008.
- [25] M. S. Gilmore, J. P. Greenwood, and J. L. Bishop, "Sulfates in Iani Chaos, Mars," in *Proc. Lunar Planet. Sci. Conf.*, 2010, p. 2374.
- [26] G. A. Swayze, "The hydrothermal and structural history of the Cuprite mining district, southwestern Nevada: An integrated geological and geophysical approach," Ph.D. dissertation, Univ. Colorado, Boulder, CO, 1997.
- [27] J. Plaza, A. Plaza, and G. Martiacute;n, "A fast sequential endmember extraction algorithm based on unconstrained linear spectral unmixing," *Proc. SPIE*, vol. 7477, pp. 747 70L-1–747 70L-11, 2009.
- [28] R. Green, "AVIRIS and related 21st century imaging spectrometers for Earth and space science," in *High Performance Computing in Remote Sensing*. Boca Raton, FL: Chapman & Hall/CRC, 2008.
- [29] L. Rowan, S. Hook, M. Abrams, and J. Mars, "Mapping hydrothermally altered rocks at Cuprite, Nevada, using the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER), a new satellite-imaging system," *Econ. Geol.*, vol. 98, no. 5, pp. 1019–1027, Aug. 2003.

- [30] F. Morgan, F. Seelos, S. Murchie *et al.*, "CAT tutorial," in *Proc. CRISM Data User's Workshop, Lunar Planet. Sci. Conf.*, 2009.
- [31] J.-P. Bibring, Y. Langevin, A. Soufflot, C. Combes, and C. Cara, "Results from the ISM experiment," *Nature*, vol. 341, no. 6243, pp. 591–593, Oct. 1989.
- [32] C. Chang and Q. Du, "Estimation of number of spectrally distinct signal sources in hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 3, pp. 608–619, Mar. 2004.



**David R. Thompson** received the M.Sc. degree in informatics from the University of Edinburgh, Edinburgh, U.K. and the Ph.D. degree in robotics from the Carnegie Mellon Robotics Institute, Pittsburgh, PA.

He is a Researcher in the Machine Learning and Instrument Autonomy group at the Jet Propulsion Laboratory, Pasadena, CA. His work focuses on machine learning and pattern recognition for remote sensing and autonomous planetary science.



**Lukas Mandrake** received the B.A. degree in engineering physics from the University of Arizona, Tucson, in 1995, and the M.S. and Ph.D. degrees in computational plasma physics from University of California at Los Angeles, Los Angeles, in 2004.

He is a member of the technical staff in the Machine Learning and Instrument Autonomy group at the Jet Propulsion Laboratory, Pasadena, CA. He is involved in the application of machine learning techniques to both Earth-based and planetary remote sensing, hyperspectral imagery, autonomous spectral analysis, and environmental monitoring.



**Martha S. Gilmore** received the Ph.D. degree in geological sciences from Brown University, Providence, RI.

She is an Associate Professor at the Planetary Sciences Group, Department of Earth and Environmental Sciences, Wesleyan University, Middletown, CT. Her primary research interests include geomorphology and tectonics on terrestrial planets, environmental remote sensing, and mineralogical studies utilizing visible/near-infrared spectroscopy.



**Rebecca Castaño** received the Ph.D. degree in electrical engineering from the University of Illinois with her dissertation in the area of computer vision.

She is currently an Assistant Manager with the Jet Propulsion Laboratory Section for Instrument Software and Science Data Systems. She has been advancing the state of the art in autonomous planetary science and remote sensing analysis for over a decade and has been lead author on numerous publications in the field.