

Automated Labeling of Materials in Hyperspectral Imagery

Brian D. Bue, *Student Member, IEEE*, Erzsébet Merényi, *Senior Member, IEEE*, and Beáta Csathó, *Member, IEEE*

Abstract—We present a technique for automatically labeling segmented hyperspectral imagery with semantically meaningful material labels. The technique compares the mean signatures of each image segment to a spectral library of known materials, and material labels are assigned to image segments according to the most similar library entry. Similarity between spectral signatures is evaluated using the $CICR_d$ similarity measure proposed in [1], designed specifically for hyperspectral imagery. This measure considers both the continuum-intact reflectance spectrum and its continuum-removed representation. We provide a thorough assessment of this measure by comparison to several commonly used similarity measures on a well-studied Low Altitude AVIRIS image of an urban area. We evaluate our results using both information-theoretic techniques and visual validation of the resulting spectral matches.

Index Terms—automatic labeling, material labeling, spectral matching, AVIRIS, hyperspectral imagery, urban

I. INTRODUCTION

IDENTIFYING surface materials is a fundamental goal of employing airborne hyperspectral sensors. The high spectral and spatial resolution of these sensors, along with many recent advances in image processing techniques, has brought fully automated material mapping close to reality. A significant remaining step towards automated material mapping is addressing the semantic gap between hyperspectral image pixels and their material interpretations. Because hyperspectral pixels are effectively unique material descriptors, it is often possible to determine which materials hyperspectral pixels represent by locating similar spectra in libraries of field or lab measured spectra; a methodology has been practiced by experts in spectroscopy for many years.

Spectral libraries such as those available from the USGS [2] and NASA (e.g., RELAB [3], ASTER [4] and CRISM [5]) have been used extensively by spectroscopists to interpret terrestrial and planetary spectral data. Their use in automated analysis, however, has been somewhat limited, because current spectral libraries often do not capture the diverse variations and types of spectral signatures that can be extracted from hyperspectral imagery. Another limiting factor is the computational expense of comparing thousands/millions of hyperspectral image pixels to each entry in a spectral library (which may, itself, consist of thousands of signatures).

A promising method to reduce computational costs is to summarize segments (clusters) of similar image pixels (spectra) that capture the most relevant spectral variations in the

image. Once an image is segmented according to spectral properties of the pixels, we can compute mean signatures (or other summary statistics) for each segment, which can be subsequently compared to spectral library signatures.

In this work, we describe a methodology for assigning semantically meaningful labels to segmented hyperspectral image data through automated comparisons of cluster signatures to field and lab measured spectra. In previous work, we proposed a novel similarity measure for hyperspectral image signatures that considers both the continuum-intact (hereafter referred to as CI) reflectance spectrum in addition to its continuum-removed (CR) representation [1]. In this paper, we provide a more thorough assessment of this new measure by comparison to several commonly used similarity measures on a well-studied Low Altitude AVIRIS image. We evaluate the measures using techniques proposed by Chang et al. in [6], and through visual inspection of the resulting spectral matches.

This paper is organized as follows: in the next section, a review of several topics related to automatic labeling and spectral matching is presented. Next, we describe our automatic labeling methodology, methods for preprocessing spectra and evaluating spectral similarity, and techniques for evaluating spectral similarity measure performance.. Section IV describes the imagery and spectral library we analyze, along with our automatic labeling results. Analysis of these results is provided in section V, and further discussed in section VI. Finally, we discuss conclusions and future work in section VII.

II. RELATED RESEARCH

Hyperspectral Image Segmentation: The goal of an image segmentation algorithm is to partition an input image into subsets of similar pixels. One input to this work is segmented hyperspectral imagery, therefore the quality of the segmentation is very important. Due to the dimensionality and complexity of hyperspectral data, segmentation is a challenging task and a subject of ongoing research in the remote sensing community. Of the numerous approaches for segmenting hyperspectral data, Self-Organizing Maps have performed demonstrably well [7], [8], [9]. Also, recent approaches which have yielded promising results are hierarchical methods (e.g [10], [11], and [12]), and combining spectral and spatial properties in clustering (for instance, [13], [14] and [15]).

Spectral Similarity Measures: Characterizing spectra in terms of their shape and the positions/widths of their absorption bands is crucial in measuring spectral similarity. Many hyperspectral analysis techniques assess similarity using a function such as the Euclidean distance (ED), Cross-Correlation (CCSM) [16], symmetrized Kullback-Leibler di-

B. Bue and E. Merényi are with the Department of Electrical and Computer Engineering, Rice University, Houston TX, 77005, USA. E-mail: {bbue, erzsebet}@rice.edu. B. Csathó is with the Department of Geology, University at Buffalo, SUNY, Buffalo NY, 14260, USA. E-mail: bcsatho@buffalo.edu. Color figures available in online version.

vergence (Spectral Information Divergence (SID)) [6], or cosine similarity (Spectral Angle Mapper) [17]). Several techniques measure similarity between CR spectra (e.g., Spectral Feature Fitting [18], Cross-Correlation Spectral Matching for Continuum Removed signatures (CCSM-CR) [19]), but these are similarly applicable to CI spectra.

Despite the considerable utility of these measures (which have been compared in several previous works, such as [20], [21], and [22]), they all consider either CI reflectance signatures, or CR absorption band characteristics — but not both. If CI signatures are used, differences in absorption band characteristics are often poorly captured because the continuum shape tends to dominate similarity comparisons. Alternatively, using CR spectra alone emphasizes differences in absorption bands, but discards essential continuum information.

Automated Hyperspectral Image Analysis with Spectral Libraries: Utilizing established lab or field-measured spectral libraries in automated analyses has only been explored in a few instances in the hyperspectral imaging domain. One of the seminal works in this area is the Tetracorder algorithm, developed by Clark et al. [18]. The Tetracorder compares continuum removed image signatures to library signatures by calculating a modified least squares fit between the signatures, constrained by the reflectance levels, continuum slopes, and presence/absence of ancillary features of interest in the spectra. Spectral matches are filtered by a rule-based system that searches for “diagnostic features” defined by spectroscopists.

Wagstaff et al. leverage spectral libraries in semi-supervised clustering of hyperspectral imagery [23]. Here, library signatures are used to define initial cluster centroids in K-means clustering. By explicitly “seeding” clusters with signatures from known species, the converged clusters often carry more definite interpretations, while also achieving faster convergence in clustering (up to a 40% decrease in convergence time, when compared to randomly initialized cluster centroids).

Keshava [24] utilizes spectral libraries in the context of hyperspectral band selection for material identification. Here, band selection is framed as an optimization problem where a set of categorized reference spectra are used to select the bands which maximize the angular separation (using the SAM similarity function) between categories. Classification accuracies produced using only the selected bands are within 5% of accuracies employing all spectral bands.

While spectral libraries have not been employed extensively to automatically match hyperspectral image signatures, their use in automated matching techniques has been more widely explored in mass spectrometry. This methodology dates back to as early as 1971 when Hertz et al. [25] developed methods to autonomously evaluate similarities between a set of known mass spectral signatures to spectra collected by a gas chromatograph. Analogous techniques are still used today (for example, in [26] and [27]), and while the measures for evaluating similarity between mass spectra differ greatly from those for hyperspectral signatures, the spectral matching methodology is essentially the same (a review of spectral matching topics for mass spectroscopy is given in [28]).

III. METHODOLOGY OF AUTOMATIC MATERIAL LABELING

After a hyperspectral image has been segmented, each segment consists of the set of pixels “most similar” to one another, according to the selected similarity measure. Because the pixels (the spectra) in each segment are similar, we can summarize each segment by its mean spectral signature (we will use “mean signature” and “cluster signature” interchangeably). To assign a material label to a segment, we calculate the similarity between its mean signature and signatures in a library of (lab or field-measured) signatures with known material labels. We assume that each spectral signature in the library is a unique descriptor for the material it represents. Therefore, if the similarity measure yields a high similarity score for a given cluster signature and a particular library signature, we can assign the material label from the library signature to the members (pixels) of the given cluster.

Because we assign material labels based on similarity scores between cluster signatures (spectral properties of unknown materials) and library signatures (spectral properties of known materials), the validity of the label assignments depends on two assumptions: (1) cluster/library signatures adequately characterize their constituent materials, and (2) the chosen similarity measure can accurately quantify relationships between signatures. If either of these assumptions break down, the resulting label assignments will be unreliable.

Additionally, due to the automated nature of the technique, the quality of the labeling is constrained by the material metadata available in the spectral library. This can be problematic because some library entries have *object* labels, rather than *material* labels. By “object” we mean some collection of one or more materials collectively described as a high level semantic concept (for instance, a tennis court). We must distinguish between material vs. object identification because we often cannot infer objects from material properties without additional context. For instance, we cannot differentiate between an asphalt rooftop and an asphalt road by their spectral signatures alone. Conversely, if we are only provided an object label, we cannot (automatically) infer the material composition of that object. A higher level of semantic context is necessary to determine the relationship between identified materials and the object(s) to which they belong, which is beyond the scope of this paper. In this work we focus on identification of materials as determined by their spectral properties.

A. Signature Normalization

In order to compare a cluster signature to a library signature, it is necessary that both signatures have the same spectral range and resolution. We assume atmospheric effects are accounted for by converting image radiance to surface reflectance using appropriate atmospheric correction methods. After atmospheric correction is performed, we still must account for scaling effects caused by illumination conditions. To mitigate illumination effects, we scale each cluster and library signature by its Euclidean norm. This scales signatures to unit length while preserving spectral angles, but discards

geometric albedo, which can be recovered in post-processing by selecting from multiple same-spectrum matches.

B. Combining Spectral Representations: $CICR_d$

We measure spectral similarity using the $CICR_d$ distance proposed in [1], which accounts for differences in both continuum shape and absorption bands. $CICR_d$ is defined as:

$$CICR_d(\mathbf{s}_i, \mathbf{s}_j) = \frac{d(\mathbf{s}_i, \mathbf{s}_j)}{v_{CI}} + \alpha \frac{d(CR(\mathbf{s}_i), CR(\mathbf{s}_j))}{v_{CR}} \quad (1)$$

where \mathbf{s}_i and \mathbf{s}_j are two CI spectral signatures, $d(\cdot, \cdot)$ is a distance measure, α is a weighting factor that determines the contribution of the CR term, v_{CI} and v_{CR} are scaling factors (described below). The output of the $CR(\cdot)$ function is a vector with components in the range $[0, 1]$, where values of one lie on the estimated continuum and values less than one indicate the depth of absorptions relative to the estimated continuum. Scaling factors v_{CI} and v_{CR} are the variances of all pairwise distances between library and cluster spectra, for CI and CR versions, respectively. Scaling each term by its respective variance is necessary since $d(\cdot, \cdot)$ and $d(CR(\cdot), CR(\cdot))$ are not (in general) in the same range. For as straightforward comparison to earlier works as possible, we set $\alpha = 1.0$. In this work, $d(\cdot, \cdot)$ will be either the Euclidean distance (ED) or Spectral Information Divergence (SID).

C. Comparing Spectral Similarity Measures

We discuss evaluation measures proposed by Chang [6], which characterize the performance of a distance measure $d(\cdot, \cdot)$: the Spectral Discriminatory Probability (hereafter referred to as SDP^d), Spectral Discriminatory Entropy (SDE^d), and the Power of Spectral Discrimination (PW^d).

The Spectral Discriminatory Probability calculates the likelihood that a cluster signature \mathbf{c} will be identified as library signature \mathbf{l}_k using distance measure $d(\cdot, \cdot)$.

$$SDP^d(\mathbf{c}, \mathbf{l}_k) = \frac{d(\mathbf{c}, \mathbf{l}_k)}{\sum_{j=1}^m d(\mathbf{c}, \mathbf{l}_j)} \quad (2)$$

given a set of m library signatures $\mathbf{L}^c = \{\mathbf{l}_1, \dots, \mathbf{l}_m\}$. A small SDP^d value indicates the probability of distinguishing the cluster signature and library signature is low, within the context of the given library. Thus, the “best” matches, according to measure $d(\cdot, \cdot)$, are those with the smallest SDP^d values. In this work, we inspect the best three ($m = 3$) matches to balance the amount of manual validation while providing a satisfactory demonstration of the technique. We do not threshold the spectral distances when selecting these three matches, therefore we will get three candidates per signature, regardless of their similarity.

The Spectral Discriminatory Entropy quantifies the uncertainty in identifying cluster signature \mathbf{c} amongst the library matches in \mathbf{L}^c .

$$SDE^d(\mathbf{c}, \mathbf{L}^c) = - \sum_{j=1}^m SDP^d(\mathbf{c}, \mathbf{l}_j) \log SDP^d(\mathbf{c}, \mathbf{l}_j) \quad (3)$$

The SDE^d takes values in the range $0 < SDE^d \leq \log \frac{1}{m}$, reaching its maximum when all m distance values are equal.




Fig. 1: The functional forms of the PW^d (a) and SDE^d (b) for $\mathbf{L}^c = \{\mathbf{l}_1, \mathbf{l}_2\}$, according to equations (3), (4).

A smaller value indicates a better chance of identifying \mathbf{c} amongst the library signatures in \mathbf{L}^c .

The Power of Spectral Discrimination estimates the power, for a given distance measure $d(\cdot, \cdot)$, of distinguishing library signatures $\{\mathbf{l}_i, \mathbf{l}_j\} \in \mathbf{L}^c$ from one another, with respect to a reference signature \mathbf{c} , and is defined as:

$$PW^d(\mathbf{c}, \mathbf{l}_i, \mathbf{l}_j) = \max \left\{ \frac{d(\mathbf{c}, \mathbf{l}_i)}{d(\mathbf{c}, \mathbf{l}_j)}, \frac{d(\mathbf{c}, \mathbf{l}_j)}{d(\mathbf{c}, \mathbf{l}_i)} \right\} \quad (4)$$

$$= \max \left\{ \frac{SDP^d(\mathbf{c}, \mathbf{l}_i)}{SDP^d(\mathbf{c}, \mathbf{l}_j)}, \frac{SDP^d(\mathbf{c}, \mathbf{l}_j)}{SDP^d(\mathbf{c}, \mathbf{l}_i)} \right\} \quad (5)$$

PW^d values near one indicate that the library signatures are “indistinguishable” with respect to cluster signature \mathbf{c} . For each cluster signature \mathbf{c} , we calculate the mean Power of Spectral Discrimination for the corresponding m best library matches. The mean PW^d is defined as:

$$\overline{PW^d(\mathbf{c}, \mathbf{L}^c)} = \frac{2}{m(m-1)} \sum_{i=1}^m \sum_{j=i+1}^m PW^d(\mathbf{c}, \mathbf{l}_i, \mathbf{l}_j) \quad (6)$$

Both the $\overline{PW^d}$ and SDE^d estimate the uncertainty in distinguishing cluster signatures from library signatures. In fact, the PW^d and SDE^d produce rankings that are order isomorphic, i.e., given a set of distances between library and cluster signatures, the ordering produced by ranking the distances by the PW^d is equivalent to the ranking generated by the $-SDE^d$. Fig. 1 shows the functional behavior of the PW^d and SDE^d for $m = 2$. Both the PW^d and $-SDE^d$ are convex functions with minima at the same location ($\frac{1}{2}$, specifically). Thus, both are monotonically decreasing before, and monotonically increasing after that location, and are therefore order isomorphic.

This isomorphism also holds for $m > 2$. In this case, the $\overline{PW^d}$ is used (Eq. 6). Clearly, when $m = 2$, the PW^d is a convex function (as illustrated in Fig. 1). Since the sum of convex functions is convex, $\overline{PW^d}$ is also convex, with minimum at $\frac{1}{m}$ (when all outcomes are equiprobable) with value 1. Because entropy is a concave function, $-SDE^d$ is a convex function, with minimum similarly attained when all outcomes are equiprobable (at value $\log \frac{1}{m}$). Once again, we have two convex functions with minima at the same location, both monotonically decreasing (increasing) before (after) the minimum. Thus, the order isomorphism holds when $m > 2$.

Given a set of similarity values, the PW^d is better suited to discriminate values at the extreme ends of the distribution, whereas the SDE^d gives better separation across the mid-range. However, both measures produce precisely the same ordering of the similarity values, and thus, we focus on only the PW^d in this work.

The mean PW^d for library signatures in L^c characterizes how “tightly packed” the distances are between the library signatures and the cluster signature c . Intuitively, we want this value to approach 1 for similar signatures, and to be large for dissimilar signatures. But the PW^d may become skewed if distances between c and its best matching library signatures (L^c) are relatively far apart (as demonstrated in section V). This often indicates that a representative signature does not exist in the spectral library, since at least one of the m matches may be a different spectral species than the other matches.

Additionally, the PW^d is sensitive to spectral representation, and in many cases does not capture visually strong matches, particularly with the CR_d measures (see section V-A). To address this issue, we visually inspect the top m library matches for each cluster signature c , obtained using ED and SID. We assign a “visual score” (“VS” in Table I) to each $l_k \in L^c$ in the range of $[0, 3]$, based on the overall spectral shape and absorption band positions/widths. A visual score of zero indicates poor quality for all m matches. A score of one indicates the majority (but not all) of the m matches are of poor quality, while a score of two indicates the majority of the matches are of good quality. Finally, if all m signatures strongly match the cluster signature, we assign a score of three. Because these scores are subjective, multiple observers are necessary to corroborate them confidently. In this work, the matches produced by the CI_d , CR_d and $CICR_d$ measures have been judged by four independent observers. Additionally, each observer has assigned visual scores to 210 $CICR_{ED}$ vs. $CICR_{SID}$ spectral matches using a web-based form, and we have observed similar trends in the per-user rankings, in comparison to the visual scores provided in Table I.

We assess significance of comparisons between similarity measures using the Wilcoxon Signed-Rank Test (WSRT) [29]. The WSRT is a non-parametric statistical hypothesis test for paired measurements (similarity values) on a single sample (cluster signature). Three quantities define the WSRT: the number of trials performed, N_t , the sum of positive differences in paired measurements, W^+ , and the sum of negative differences in paired measurements, W^- . Equal measurements are handled by adding their mean to both W^+ and W^- . The significance of the performance is based on N_t and $\max(W^+, W^-)$ [30]. Using the WSRT to test significance of spectral similarity measure comparisons has several advantages. First, it makes no assumptions on the underlying distribution of the measurements. Second, greater emphasis is placed on larger differences in measurements than on smaller ones. Third, because the statistic for the signed rank test is resistant (i.e., unaffected by changes in a few observations), outliers are naturally suppressed (if the number of outliers is not particularly large with respect to the size of the sample). For a detailed discussion on the WSRT, see [31].

IV. DATA AND CASE STUDY

A. AVIRIS Image of Ocean City, MD

We demonstrate our automatic labeling technique on an urban hyperspectral image of Ocean City, MD [32]. This image was acquired by a low-altitude AVIRIS flight on Nov

5, 1998, with spatial resolution of 4m/pixel. Data preprocessing, segmentation, the resulting image clusters and their signatures are presented in [33]. This image is an example of the complexity in a real urban study, with many material classes of interest, and additionally has a good segmentation, with clusters that have been verified to correspond to known objects and materials. In this paper, we use the segmentation produced with a Self-Organizing Map (SOM) in [33]. The high spatial and spectral resolution of AVIRIS imagery, along with the sensitive segmentation technique allowed discrimination of 35 clusters with varied characteristics including (very) small, and spectrally similar ones. As verified from field data, most of these clusters represent objects associated with distinct material types. Examples of these are: water tower, buildings, roads, boardwalks, parking lots, mini golf course, coast guard lookout tower, and landscape units. However, for some clusters which can be recognized on the functional level (i.e., tennis court), we do not have a corresponding material identification. Discontinuities in the spectra, seen in all spectral plots in the following discussions, are due to the removal of saturated atmospheric water vapor bands. We also observe a characteristic dip near $1.96\mu m$, which may be due to imperfect removal of saturated water vapor bands, but may also occur due to shadow effects (which are common in urban areas) or multiple reflections caused by vegetation.

B. Spectral Library of Urban Materials

The spectral library used in our spectral matching procedure consists of 1250 spectral signatures from three sources:

- 1164 field-measured spectra of mostly urban materials acquired in 1075 wavelengths in the 0.35 to $2.4\mu m$ range (described in [34])
- 17 lab-measured vegetation spectra from the USGS splib06a spectral library [2]
- 21 AVIRIS image spectra (mostly vegetation and soil types) from training regions described in [35]

All library signatures are tagged with metadata indicating the objects measured, and most of the entries include a corresponding material label. Library signatures are convolved to appropriate AVIRIS wavelengths. We exclude wavelengths outside of the range $[0.42, 2.39]\mu m$ due to low signal to noise ratios in some of the library signatures, and exclude bands removed from the Ocean City image. The remaining 165 of the original 224 AVIRIS bands are used for spectral matching.

V. EVALUATION OF SPECTRAL MATCHING AND AUTOMATIC LABELING ON OCEAN CITY CLUSTERS

A. Spectral Matching Performance

Table I gives the set of all visual scores for the matches of the Ocean City clusters, along with summary statistics for the visual and PW^d scores for all cluster signatures (“(All)” in Table I) and signatures which are adequately represented in the library (“(Selected)” in Table I). Since precise material interpretations are not known for all clusters, we make the simplifying assumption that if the mean visual score (for all measures) for a cluster signature is zero, then that signature does not have a representative library signature. From Table I,

we see that clusters with mean visual scores equalling zero include **P**, **S**, **X**, **a**, and **c**. **C**, **F** and **d** are also excluded because they could not be interpreted with adequate confidence (**C** is a green tennis court shown in Fig. 3, discussed below, **F** consists of street/sidewalk materials, and **d** is likely a mixture of water and nearby building materials).

While the best possible $\overline{PW^d}$ score is 1, high $\overline{PW^d}$ scores do not necessarily indicate a measure is performing poorly. It may suggest the average of the top m matches is skewed because the similarity values are on varying scales. This is evident in Fig. 2, where PW^d scores for each cluster are provided. Of particular interest are the atypical scores for signatures **M** (vegetation) and, to a lesser degree, **T** (asphalt) and **Y** (sand). Table II gives similarity and pairwise PW^d scores for the best three matches for signature **M**. In this case (similarly with signatures **T** and **Y**), a single similarity score is relatively distant from the remaining two scores, resulting in a relatively high PW^d value. Note that each similarity measure returns a different set of library matches, yet we observe the same effect. In fact, the PW^d may be used to indicate the library has less than m suitable match candidates.

In terms of visual scores, matches using the $CICR_d$ measures score higher than matches made with their CI_d and CR_d counterparts, with significant improvements over the CR_d measures. It may seem counter-intuitive that simply including a CR term in the spectral similarity measure will improve spectral matching performance, considering the poor performance of the CR_d measures, but such performance gains are actually quite easy to explain: CI_d measures produce spectral matches that correspond well in terms of spectral shape, but fail to capture characteristic absorption features. Conversely, using CR spectra alone will often yield matches that differ greatly in spectral shape. Fig. 3 compares the CI_d and $CICR_d$ spectral matching results for Ocean City cluster **C**. The pixels in this cluster belong to a green tennis court, and while we do not have ground truth data on the material composition of this signature, based on the character of the spectrum, it is probably composed of asphalt material (and green pigment). The signature has several significant absorptions at ~ 0.45 , 0.64 and $2.22 \mu m$ that are captured by both $CICR_d$ measures, but poorly captured by CI_d measures. The $CICR_d$ matches are not only more visually agreeable, but better represent spectroscopic similarities between signatures. Fig. 4 gives the CR_d matches for cluster signature **E** (a metal rooftop). These signatures have nearly indistinguishable CR representations, but differ significantly in terms of continua, resulting in unsatisfactory matches with the CR_d measures. Even the best match differs significantly at the shorter wavelengths, due to diffusion effects commonly observed in metallic materials.

Another example where the PW^d may not reliably measure similarity measure performance can be given through a case of spectral match described by Van der Meer in [20], Fig. 7. In this work, he concludes that the SID is more discriminatory than ED based on analysis of a synthetic data set consisting of 601-band field-measured spectra, and on AVIRIS imagery consisting of 50 bands in the 2.0 to $2.5 \mu m$ range, for material signatures montmorillonite (mont), kaolinite (kaol), quartz and

alunite (alun). While we do see (visually) a slight improvement over the ED when matching the Ocean City CI signatures, the SID performance is worse than the ED with the CR_d measures, and nearly equivalent with the $CICR_d$ measures. Furthermore, in terms of PW^d scores, we see that the SID-based measures appear less discriminatory than CI_{ED} , CR_{ED} , and $CICR_{ED}$. The previous conclusion that the SID is a better measure is based partly on the ED vs. SID PW^d values for alun-kaol, alun-mont, and kaol-mont, with quartz as the reference signature. However, this is a somewhat pathological case for the Euclidean Distance (particularly for the AVIRIS signatures given in [20], Fig. 7). On further inspection of the AVIRIS signatures, the quartz reference signature (a signature lacking significant absorption features in the selected wavelengths) will yield high PW^d scores in comparison to kaol and alun signatures (both of which are approximately bisected by the quartz signature) and the mont. signature (which is similar to the quartz signature, but darker). Consider the following: Let the sum of the polygon areas defined by the upper and lower portions of the alun signature bisected by the quartz signature be $A_{a,q}$, and let $A_{k,q}$ be the summed area of the polygons defined by the same bisector (quartz) with the kaol signature. Next, let $A_{m,q}$ be the area defined by the polygon with boundaries defined by the mont and quartz signatures. It is trivial to show that if $A_{a,q} = A_{k,q}$ then the alun-quartz and kaol-quartz Euclidean distances are equal. If $A_{m,q}$ also has the same area, then these three signatures are equidistant from the quartz signature. Granted, for the signatures in [20], Fig. 7, these areas are not precisely equal, but they do show strong similarities which would result in similar Euclidean distances.

Table III gives the WSRT p -scores for the PW^d for each similarity measure, evaluated on the 35 SOM clusters. Values near 1.0 (bold) indicate low confidence in statistical significance. The scores indicate the distributions of CR_d and $CICR_d$ similarity values differ, and therefore should not be (directly) compared. The p -scores are also relatively high between CI_d and CR_d measures. These findings are not surprising, since spectra that are very different in terms of continuum shape can be identical after continuum removal (see signatures **h** and **i** in Fig. 6, for instance). Because similar signatures produce PW^d scores near 1.0, the CR_d measures appear more discriminatory than the other measures, which is contrary to our findings in Figures 2, 3, and 4.

B. Automatic Labeling of the Ocean City Image

We evaluate the performance of the automatic labeling procedure by categorizing the spectral library into ten distinct material groups (loosely based on the taxonomy of urban materials given in [34]). We group materials for convenience in these groups: Concrete materials, Asphalts, Composites (which largely consist of shingle materials), Metals, Vegetation, Coatings (i.e., paint), miscellaneous roofing materials (e.g., tile and wood shingles), Soil/Dirt, Water, and “Other” (“Other” refers to library signatures for which material information is not provided. In our library, this includes only tennis and basketball court signatures). We select the best match produced by the $CICR_{ED}$ measure (which yielded the

	Visual Scores for Individual Clusters																															Mean Visual Score		Mean PW ^d						
	A	C	D	E	F	G	I	J	K	L	M	O	P	Q	R	S	T	U	V	W	X	Y	Z	a	b	c	d	e	f	g	h	i	j	l	m	All	Selected	All	Selected	
CI _{ED}	2	1	1	1	1	1	2	1	2	3	1	1	0	1	1	0	1	2	1	2	0	2	2	0	1	0	0	2	2	0	1	1	1	1	1	1.1143	1.1053	1.0614	1.0708	
CR _{ED}	0	2	2	0	1	0	2	0	2	3	2	1	0	1	1	0	1	2	2	1	0	1	1	0	1	0	1	2	2	1	0	1	1	1	0	1.0000	0.8947	1.0441	1.0503	
CICR _{ED}	2	2	2	1	1	1	2	1	1	3	2	1	0	1	1	0	1	2	1	2	0	2	1	0	1	1	1	2	2	1	1	1	1	1	2	1	1.2571	1.3684	1.0395	1.0439
CI _{SID}	2	1	2	1	1	2	2	1	2	3	1	1	0	1	1	0	1	2	1	2	0	2	2	0	1	0	0	2	2	0	1	1	1	1	1	1.1714	1.2105	1.1354	1.1586	
CR _{SID}	0	2	1	0	1	0	2	0	2	3	2	1	0	1	1	0	1	2	2	1	0	1	1	0	1	0	1	2	2	1	0	1	1	1	0	0.9714	0.8421	1.0908	1.1029	
CICR _{SID}	1	2	2	1	1	1	2	1	2	3	2	1	0	1	1	0	1	2	2	1	0	1	1	0	1	1	1	2	2	1	1	1	1	2	1	1.2286	1.3158	1.0898	1.1047	
Average	1	2	2	1	1	1	2	1	2	3	2	1	0	1	1	0	1	2	2	0	2	1	0	1	0	1	0	1	2	2	1	1	1	1	1					

TABLE I: Visual scores and averages for Ocean City SOM clusters. We assume clusters with average visual scores of zero (indicated by italics) are not represented in the library. Average scores are given for all clusters (All) and clusters represented in the library (Selected). The best scores are given in bold text, and the worst scores are underlined. $CICR_d$ measures generally outperform both CI_d and CR_d measures. Visual scores are in the range [0,3] (0 = worst match, 3 = best match) and PW^d scores are in the range [1,∞] (1.0 = indistinguishable w.r.t. the cluster signature).




Fig. 2: Top: PW^d scores for each cluster signature according to the CI_{ED} (solid line, circle marker), CR_{ED} (dashed line, square marker), and $CICR_{ED}$ (dotted line, diamond marker) measures ($m=3$). Bottom: PW^d scores for each cluster signature using the CI_{SID} (solid line, circle marker), CR_{SID} (dashed line, square marker), and $CICR_{SID}$ (dotted line, diamond marker) measures. On average, ED-based measures yield better PW^d scores than SID-based measures. Cluster M (similarity scores given in table II) is the most spectrally ambiguous according to the employed measures.

best overall matching performance experimentally), and if the material group of the matching library signature corresponds well to the material group of the cluster signature, we consider the label assignment a success. For some cases, determining this correspondence requires the translation of an object label (for instance, “rooftop”) to a material group (“asphalt”) (based on manual inspection of the cluster signature and expert interpretation), since the expert interpretations are sometimes given on the object, rather than on the material level.

Of the 25 clusters with adequate library representation, 21 are successfully labeled using the $CICR_{ED}$ measure. These 21 clusters comprise 67.6% of the image pixels with known material labels available in the library. The automatic labeling

$d(\cdot, \cdot)$	M, l_1	M, l_2	M, l_3
CI_{ED}	27.097 (shaded concrete)	51.683 (green paint)	51.683 (grass)
CR_{ED}	94.866 (grass)	137.969 (tall grass)	138.628 (gray shingle)
$CICR_{ED}$	165.857 (grass)	217.640 (shaded concrete)	222.888 (sage brush)
CI_{SID}	0.187 (shaded concrete)	0.5817 (green paint)	0.610 (green paint)
CR_{SID}	0.052 (grass)	0.107 (palm tree)	0.110 (green paint)
$CICR_{SID}$	0.309 (shaded concrete)	0.793 (green paint)	0.844 (green paint)
$PW^d(\cdot, \cdot, \cdot)$	M, l_1, l_2	M, l_1, l_3	M, l_2, l_3
CI_{ED}	1.907	1.954	1.025
CR_{ED}	1.454	1.461	1.004
$CICR_{ED}$	1.312	1.343	1.024
CI_{SID}	3.114	3.266	1.049
CR_{SID}	2.026	2.080	1.026
$CICR_{SID}$	2.565	2.731	1.065

TABLE II: Similarity and PW^d values for cluster signature M. Significantly higher PW^d scores are due to spectral ambiguity between cluster signature M and library signatures l_2 and l_3 (bottom, bold), combined with a strong match to l_1 .

	CI_{ED}	CR_{ED}	$CICR_{ED}$	CI_{SID}	CR_{SID}
CR_{ED}	0.1542				
$CICR_{ED}$	0.0885	0.7064			
CI_{SID}	0.0000	0.0007	0.0001		
CR_{SID}	0.0200	0.0034	0.0004	0.1957	
$CICR_{SID}$	0.0797	0.0238	0.0036	0.0769	0.6465

TABLE III: WSRT-based p -values for the PW^d using CI_d , CR_d and $CICR_d$ similarity measures for the 35 Ocean City SOM cluster signatures. Significantly higher p -values between CI and CR-based similarity measures indicate that the similarity values produced by these measures do not follow the same distribution, and should not be compared directly.

results for each Ocean City cluster using the $CICR_{ED}$ measure are given in Fig. 5. Expert interpretations of clusters are given in plain text (column 2), and the $CICR_{ED}$ library matches are given in colored text (column 3). If the match is considered a success (according to our library categorization), the text is colored green, while mismatches are colored red. Clusters without a clear expert interpretation for their materials are displayed in black text. Clusters with an asterisk by the expert interpretation (in column 2) lack representative material signatures in the spectral library; therefore these matches should be

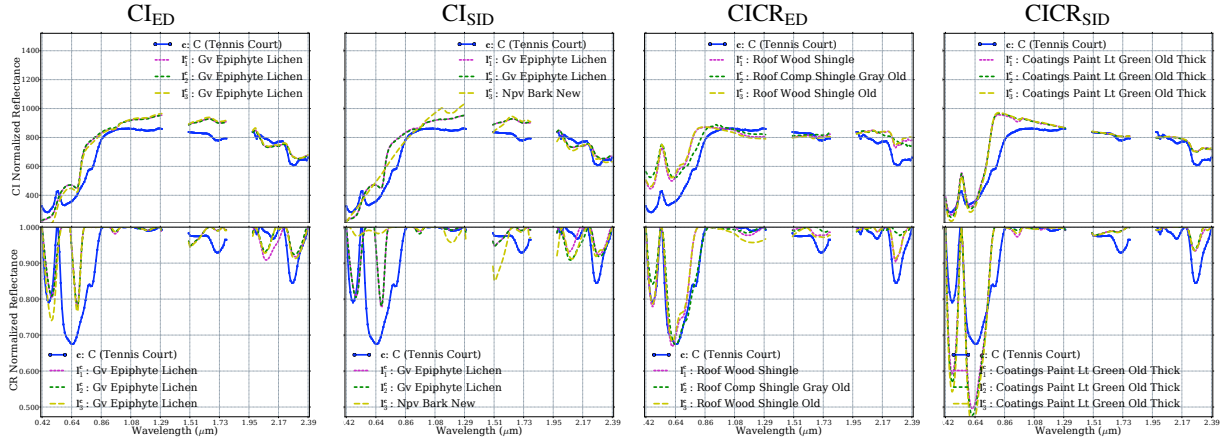


Fig. 3: Top: Best three matches for cluster **C** using CI_{ED} , CI_{SID} , $CICR_{ED}$ and $CICR_{SID}$ measures. Bottom: Corresponding CR spectra. Using CI signatures alone results in poor matches of absorption bands (particularly at the shorter wavelengths). The $CICR_d$ measure can exploit differences in absorption band characteristics, and thus, achieves improved matching performance.

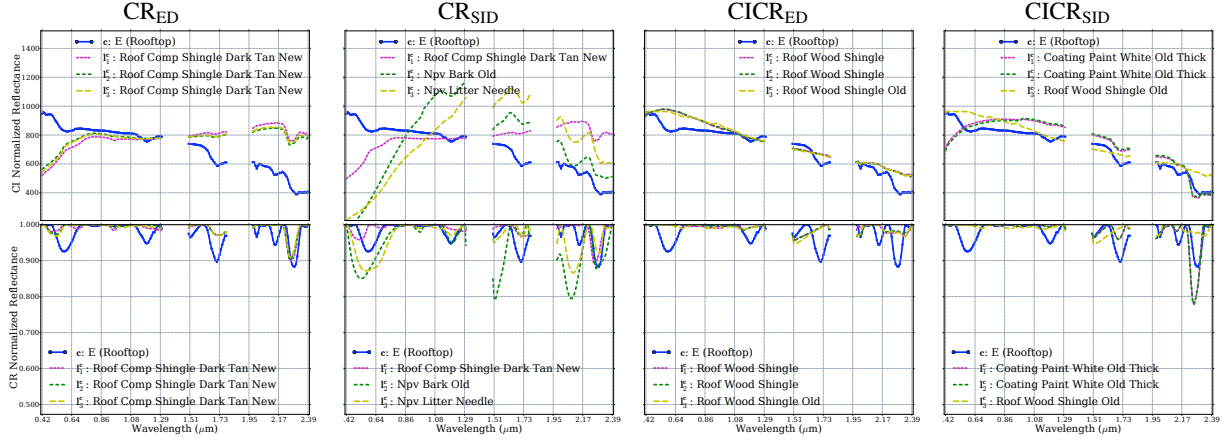


Fig. 4: Top: Best three matches for cluster **E** using CR_{ED} , CR_{SID} , $CICR_{ED}$ and $CICR_{SID}$ measures. Bottom: Corresponding CR spectra. Due to the fact that the CR representation discards information on the shape of the continuum in favor of absorption band characteristics, spectral matching with CR_{ED} and CR_{SID} is poor. The matches using the $CICR_d$ measure yield improved matches since both the continuum and the absorption features are considered.

disregarded. Selected spectral matches, grouped according to their best matching library material label, are given in Figures 6, 7, 8, 9 and 10. Even within these categories, there are often significant differences in spectral shape for similar materials, but since our library is sufficiently diverse, we find relevant matches in almost all cases.

Not surprisingly, including CR signatures does not improve discrimination between materials without significant absorption features. Fig. 6 gives the CI_{ED} and $CICR_{ED}$ matches for several asphalt cluster signatures. The best library matches using both measures are the same, with only slight changes in the ranking order. Also, visual scores (in Table I) for many of the asphalt (**h**, **i**, **T**) and composite (**G**, **I**) signatures remain the same for both the CI_d and $CICR_d$ measures.

Two of the concrete matches are of particular interest. First, cluster signature **L** (Fig. 7) is matched to a “shaded concrete” library signature. This library signature is described in detail in [34]), and is an example of an “intimate” mixture [36] of concrete and a tree canopy. The mixture of the flat concrete library signature does not cause significant perturbation of the vegetation library signature and thus appears representative of

vegetation, and therefore, matches well to cluster signature **L** (trees). The other concrete signature, **U**, corresponds well to several gray/dark gray-colored rooftop material signatures. According to recent aerial photographs, the smaller **U** signature (Fig. 5, right image) is a viewing tower, with a small enclosed building on top, that likely is composed of a concrete roof and concrete base. The larger **U** signature (Fig. 5, left image) appears to contain concrete roof tiles. Also, the match to the gravel rooftop is expected since concrete materials generally consist of a mixture of cement, gravel and water.

VI. DISCUSSION

A. Spectral Representation: CI vs. CR

While absorption characteristics are of great importance in identifying materials from spectral signatures, using the CR representation alone in spectral matching is unreliable since continuum shape information is ignored. Furthermore, the CR signatures are only useful in discriminating spectral signatures if the signatures have significant absorption features. Many urban materials such as concrete and asphalt often do not have such features, so the performance of CR_d measures will be

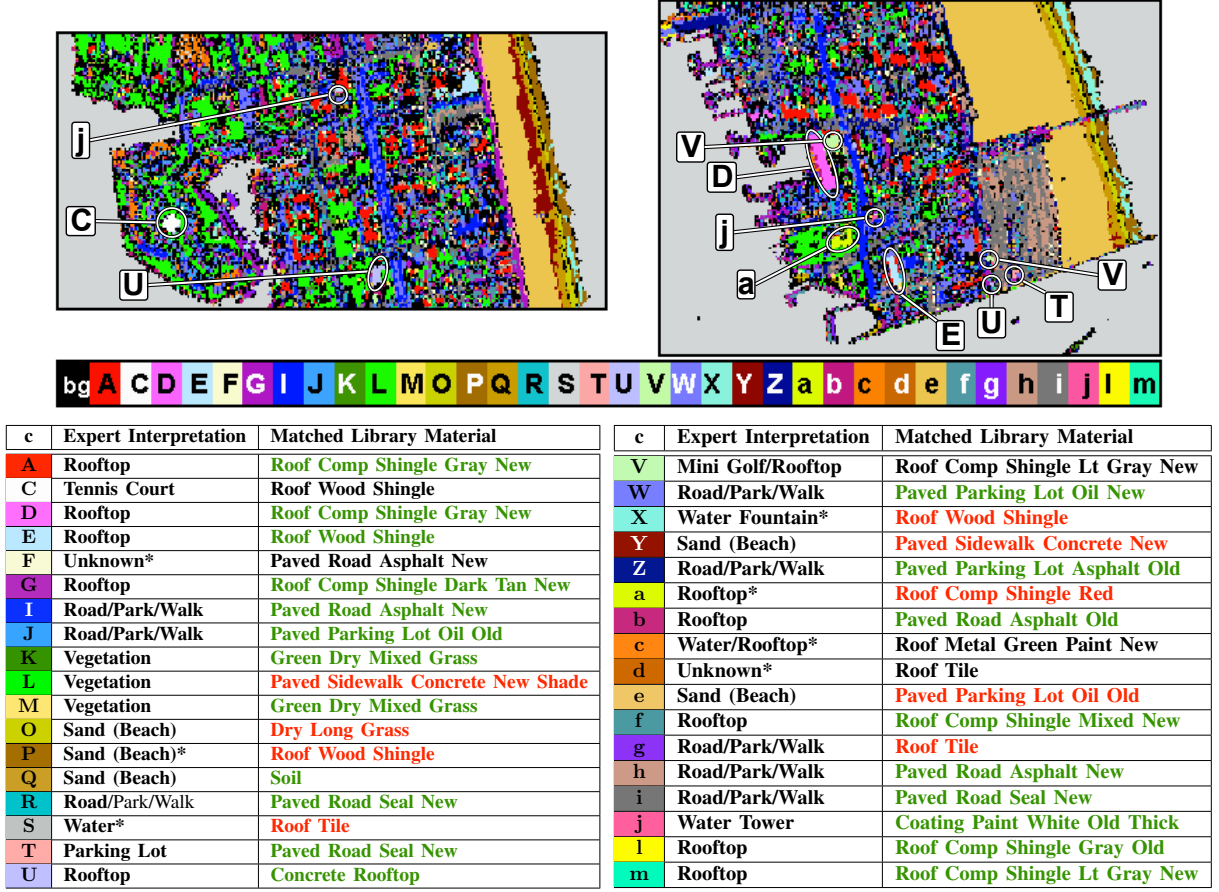


Fig. 5: Automatic labeling results for all Ocean City cluster signatures. Cluster interpretations (from field knowledge) are given in black text (column 2) and the corresponding best match using the $CICR_{ED}$ measure is given in column 3 (colored text). Cluster interpretations marked with an asterisk do not have representative material signatures in the spectral library, and are not included in the “Selected” measurements in Table I. Matches in green text (in column 3) indicate that the material of the best library match corresponds well to the expert interpretation, red text indicates a mismatch, and black text indicates that the material composition for the cluster signature is unknown. Labels are determined on the basis of the spectral shape similarities. Spectral matches are discussed in detail in figures 6, 7, 8 9 and 10.

poor, while the performance of the CI_d and $CICR_d$ measures will be approximately the same on such material signatures.

The continuum removal technique employed may have a significant effect on the $CICR_d$ measure. In this work, we approximate the continuum by first connecting a set of the most significant local maxima in a spectral signature via straight-line segments (this procedure is analogous to the technique described in [37]). This may lead to slight distortions in the CR signatures, due to the piecewise linear nature of the continuum approximation. A continuum approximation method using higher-order basis functions (for instance, using the technique described in [38]) may improve the continuum approximation and subsequent matching performance.

B. Evaluating Material Mismatches

Sometimes translation between the expert interpretations of image segments and the labels provided in the spectral library is nontrivial. Fig 11 illustrates this issue. Here, the expert interpretation of cluster C, “tennis court” material, matches well to several of “wood shingle” library signatures, even though several tennis court material signatures exist in the

library. Since the precise material composition of signature C is unknown, and the library metadata, in this case, does not provide a material label for the tennis court signatures, it is difficult to assess the accuracy of this labeling. Here, determining the correct labeling for C requires additional contextual information, since the wood shingle signatures are clearly stronger matches than the tennis court signatures (both in terms of spectral shape and absorption bands). These ambiguities are best resolved by employing more diverse spectral libraries with extensive metadata, complete with material descriptions.

C. Segmentation Sensitivity

The performance of our proposed labeling technique is clearly dependent on the quality of the segmentation to be labeled. Using a poor segmentation will result in poor spectral matches (and subsequently, labels without clear meaning) since cluster signatures would not accurately capture distinctions between spectral species. To demonstrate this, we compare the SOM and ISODATA [39] segmentations of our AVIRIS image discussed in [33]. The SOM clustering (as mentioned in section IV-A) yielded 35 clusters, while ISO-

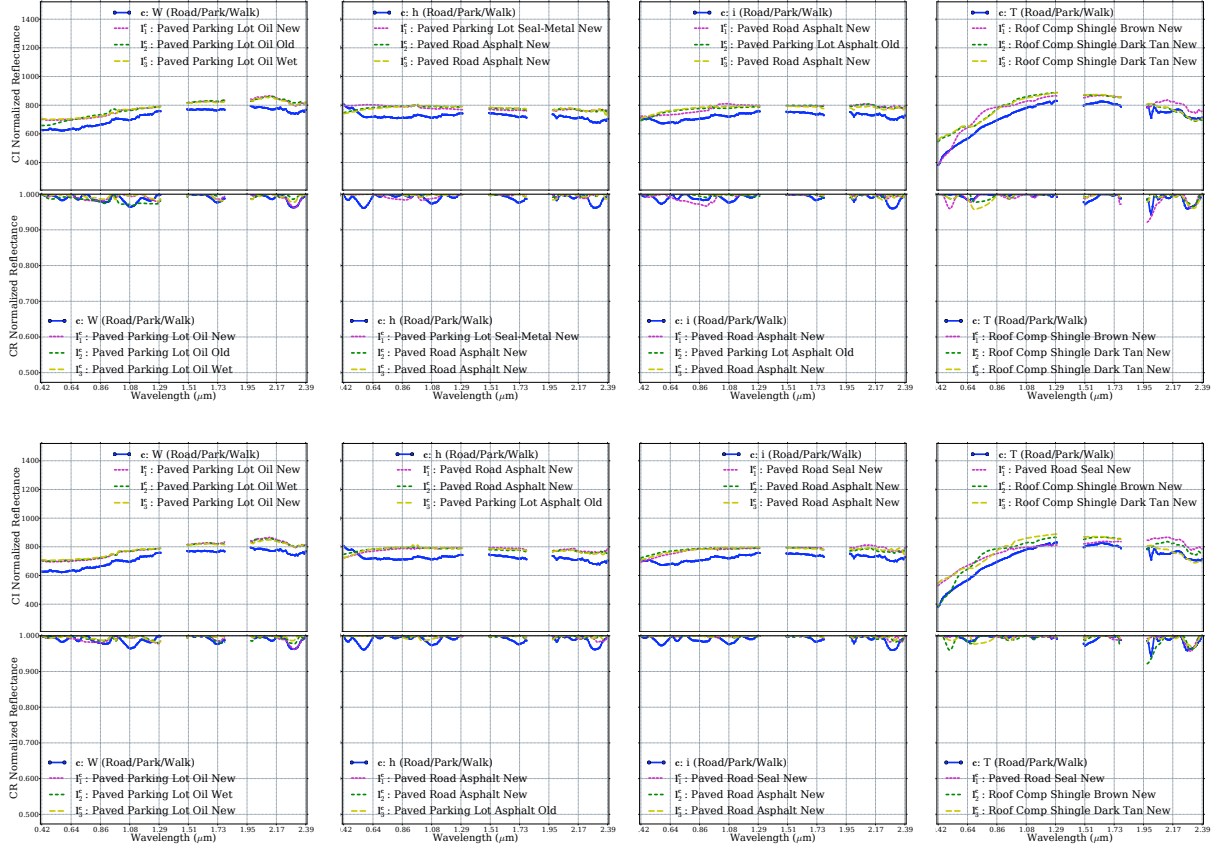


Fig. 6: Top: CI and corresponding CR library matches for category “Asphalt” using the CI_{ED} measure. Bottom: CI and corresponding CR library matches for the same clusters using the $CICR_{ED}$ measure. Both measures yield nearly the same matches due to the lack of prominent absorption features in these signatures.

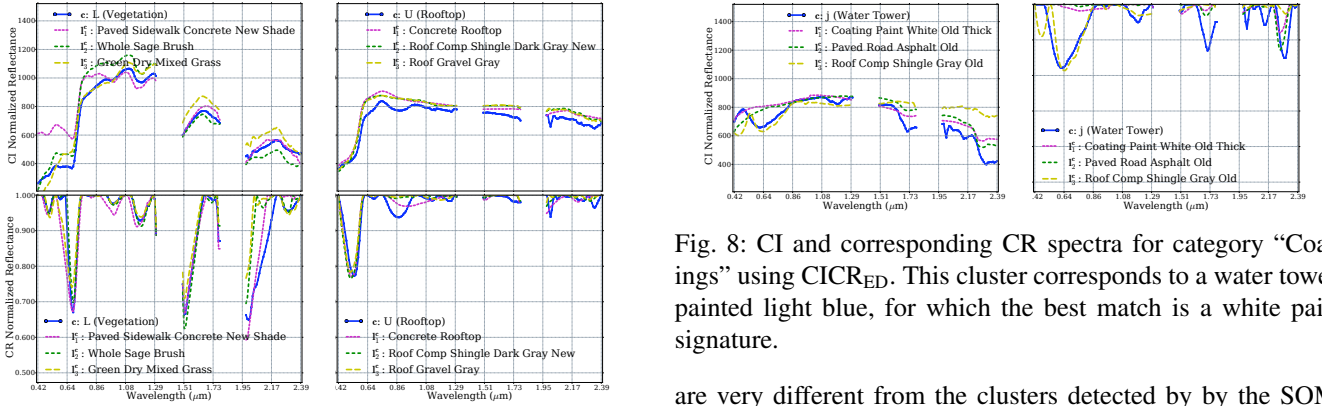


Fig. 7: CI and corresponding CR library matches for category “Concrete” using $CICR_{ED}$. Due to intimate mixing effects caused by the shadow of a tree canopy on the concrete material, the “shaded concrete” library signature has spectral shape typical to a vegetation signature and closely matches cluster signature **L** (grass).

DATA yielded 20 clusters. Both clusterings, and corresponding cluster signatures, are discussed in detail in [33]. Note that the cluster labels (colors) in these two segmentations are not consistent with each other because reconciling clusters is nontrivial or impossible since there is not a one-to-one (or even a clean one-to-many) correspondence between the two clusterings. In fact, clusters in the ISODATA segmentation

Fig. 8: CI and corresponding CR spectra for category “Coatings” using $CICR_{ED}$. This cluster corresponds to a water tower, painted light blue, for which the best match is a white paint signature.

are very different from the clusters detected by the SOM. Fig. 12 illustrates the problem of using a poor segmentation in spectral matching. In this case, ISODATA assigns pixels, corresponding to a clearly recognizable building (SOM cluster **D** in Fig. 5), into three separate clusters (**K**, **L** and **M**, not to be confused with the SOM clusters with the same labels), none of which represents the true signature of the building. The ISODATA cluster map is not shown here, but the signatures of ISODATA clusters **K**, **L** and **M** are compared to SOM cluster **D** in Fig. 12. There are two related issues here: (1) ISODATA fails to detect an area of a unique signature clearly delineated by the SOM, and (2) a number of spectrally similar materials correctly grouped together by the SOM are incorrectly assigned to several quite different ISODATA clusters. Consequently, these ISODATA clusters give no clear,

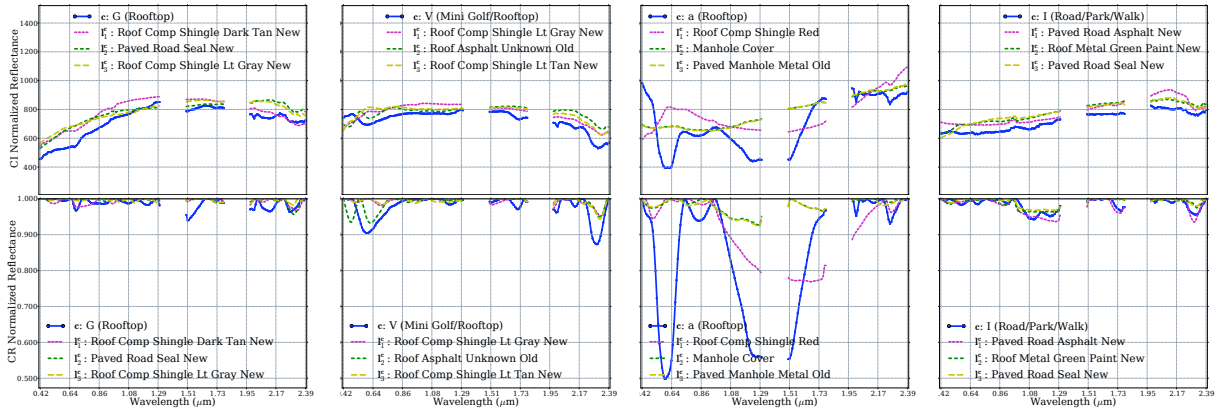


Fig. 9: CI and corresponding CR library matches for category “Composites” using CICR_{ED} . As observed in [34], considerable spectral confusion exists between dark asphalt road and composite shingle rooftop signatures (since the composite shingles often have a strong asphalt component), so material matches such as those observed in signature **G** are expected. Signature **a** (a building rooftop consisting of a mixture of metal alloy and aluminum, painted blue) is a mismatch due to both signatures having dramatically different spectral shapes (which indicates that there is not a representative signature present in the library). The material content of cluster signature **V** (mini golf/rooftop) is unknown, but the marked similarity to other asphalt signatures suggests it may be dominated by asphalt as well.




Fig. 10: CI and corresponding CR spectra for category “Vegetation” using CICR_{ED} . The second and third “green paint” matches for cluster **K** are due to strong similarities in absorption features common to vegetation species, as observed in the CR signatures. As a result of these similarities, the measure would incorrectly label the vegetation spectra as green paint if the first match, “grass,” had not been present.




Fig. 11: Left plots: Tennis court cluster signature **C** matched to library signatures of tennis court materials. Right plots: Spectral matching results for signature **C** using CICR_{ED} . Both CI and CR spectra of the tennis court better match several library signatures of “shingle” materials in comparison to the tennis court library signatures.




Fig. 12: Left block: Library matches for ISODATA clusters **K** (comprising SOM clusters verified as shingle rooftops, roads and parking areas, and a mini golf course), **L** (various rooftop materials) and **M** (various rooftop and road materials). Right: Library matches for SOM cluster **D** (a shingled rooftop). In this case, pixels that are delineated well by the SOM cluster **D** are misclustered by ISODATA into 3 of its clusters — **K**, **L** and **M** — none of which represent the true signature. Note that the spatial distributions, as well as the labels of the ISODATA clusters are different from the SOM clusters (as shown in [33]).

or worse, misleading interpretations. Matches from a library — while they may be good matches to the mean cluster

signatures — may not represent the species at the locations of the incorrectly delineated ISODATA clusters.

VII. CONCLUSIONS AND FUTURE WORK

In this work, we demonstrated a technique for automated labeling of segmented hyperspectral imagery using a library of known material signatures. Labels derived by our proposed labeling technique are determined by the contents of the library, the quality of the segmentation, and the similarity measure used to compare spectral signatures. A similarity measure designed specifically to characterize hyperspectral signatures was evaluated and shown to outperform the Euclidean distance and Spectral Information Divergence measures in matching hyperspectral signatures. This measure was used to select material labels from a library of urban material signatures for each cluster in a verified, high quality segmentation of an AVIRIS image. The technique we present successfully labeled 21 of the 25 clusters with known material interpretations and representative library signatures. The remaining clusters could not be labeled because either their material interpretations were unknown, or the library lacked representative material signatures for those clusters. Both of these issues could potentially be mitigated by augmenting the spectral library with additional, detailed metadata describing the exact material composition of all library spectra, or by including additional library spectra which include such metadata.


The lack of exhaustive and detailed ground-truth data makes the objective evaluation of automated labeling methods challenging. Even in cases where ground-truth data exists, it is often given for objects, rather than materials. Since it is currently impossible to acquire exhaustive material labels for large remote sensing surveys, synthetically-generated hyperspectral imagery may be of significant help. We are currently analyzing several images generated with the RIT DIRSIG [40] image synthesis algorithm to determine their value in the verification of automated labeling techniques.

As observed in [1], allowing the CI and CR terms to carry equal weight in the $CICR_d$ measure may be suboptimal. In this work, we set α to 1.0 for straightforward comparison to earlier works, but we expect to achieve improved spectral matching performance by allowing other α values. An optimization procedure may be necessary to learn the best α -value, dependent on the data. Another extension can be to treat α as a vector of the same length as the number of spectral bands, and then learn an optimal weighting for each band.

It is interesting to note that the $CICR_d$ measures outperform the CI_d and CR_d measures regardless of whether we use the ED or SID distance measure. However, the choice of these distance measures was primarily motivated by convenient comparisons to existing works, and one could easily substitute other measures in this calculation. Distance measures that consider global covariances (e.g., the Bhattacharyya distance [41]) or functional measures (e.g., the Sobolev distance [42]) may improve spectral matching performance.

Acknowledgments: We thank Dar Roberts for the use of their library of urban field spectra, and for his gracious assistance with it; and Bill Farrand for sharing his expertise on several technical and software-related details. We also wish to thank both reviewers for their excellent comments. B. Bue

is supported by NASA GSRP Fellowship NNX0AR79H. E. Merényi is partially supported by grant NNG05GA94G from the Applied Information Systems Research Program, NASA, Science Mission Directorate.



imagery (in particular, hyperspectral imagery), for applications in Earth and planetary science.

Brian D. Bue is a Ph.D. student in the department of Electrical and Computer Engineering at Rice University, and an associate member of the technical staff in the Machine Learning and Instrument Autonomy group at the NASA Jet Propulsion Laboratory. He earned a M.S. in computer science from Purdue University, West Lafayette, IN, and B.S./B.A. degrees in computer science and mathematics, respectively, from Augsburg College, Minneapolis, MN. His research involves developing machine learning and pattern recognition techniques for remotely sensed



Erzsébet Merényi earned an M.Sc. in mathematics (1975) and Ph.D. in computational science (1980) at Szeged (Attila József) University, Hungary. From 1980 through 1996 she held parallel appointments at the Central Research Institute for Physics of the Hungarian Academy of Science and at the Lunar and Planetary Laboratory of the University of Arizona, Tucson, Arizona, where she was a staff scientist until 2000. She is currently a research professor in the Electrical and Computer Engineering Department of Rice University, Houston, Texas, U.S. A. Her interests include self-organized machine learning, artificial neural networks, structure discovery in complicated high-dimensional manifolds, segmentation and classification of high-dimensional, complex patterns, data fusion, data mining, knowledge discovery. She has been analyzing data from various space missions and terrestrial remote sensing projects for over 20 years, including resource mapping and diagnostics from multi- and hyperspectral imagery of Mars, Earth, the Moon, asteroids; and biological tissue analysis from microscopic hyperspectral imagery. Prof. Merényi's work has been funded by the Applied Information Systems Research, Mars Data Analysis, and Solid Earth and Natural Hazards Programs of NASA's Science Mission Directorate, as well as by the Baylor College of Medicine. She is a senior member of the IEEE, a member of the IEEE Computational Intelligence Society, the International Neural Network Society, and the Division of Planetary Science of the American Geophysical Union.



Beáta Csathó earned her MS degree in mathematics at the Eötvös Loránd University, Budapest, Hungary in 1989 and a PhD degree in geophysics from the University of Miskolc, Miskolc, Hungary in 1993. From 1981 to 1994 she was a Research Scientist of the Eötvös Loránd Geophysical Institute, Budapest, Hungary. From 1994 to 2006 she was a Research Scientist of the Byrd Polar Research Center of The Ohio State University, Columbus, Ohio. In 2006 she joined the faculty of the Department of Geology at the University at Buffalo, where she is currently an Associate Professor. Her research interest spans a large number of disciplines, including geophysics, remote sensing, photogrammetry, geodesy, spatial statistics, GIS, visualization, digital image processing, pattern recognition, and data fusion. She serves as a science team member of NASA's ICESat and ICESat-2 satellite missions. Prof. Csathó is a member of ASPRS, AGU, AAAS, IGS, SEG, and she has held numerous offices at the International Society of Photogrammetry and Remote Sensing (ISPRS). She was awarded a Fulbright Fellowship in 1992, the President's Honorary Citation and a Certificate of Appreciation from ISPRS in 2000 and 2001, and NASA's Group Achievement Award in 2004.

REFERENCES

- [1] B. Bue, B. Csathó, and E. Merényi, "Automated labeling of segmented hyperspectral imagery via spectral matching," *Proceedings IEEE WHISPERS 2009*, Aug 2009.
- [2] R. Clark, G. Swayze, R. Wise, E. Livo, T. Hoefen, R. Kokaly, and S. Sutley, "USGS digital spectral library splib06a: U.S. Geological Survey, digital data series 231," 2007.
- [3] "RELAB spectral library," 2008, <http://www.planetary.brown.edu/relab/>.
- [4] A. M. Baldridge, S. J. Hook, C. I. Grove, and G. Rivera, "The ASTER spectral library version 2.0," *Remote Sensing of Environment*, vol. 113, no. 4, pp. 711–715, Jan 2009.
- [5] S. Slavney and S. Murchie, "CRISM spectral library," 2006.
- [6] C.-I. Chang, "An information-theoretic approach to spectral variability, similarity, and discrimination for hyperspectral image analysis," *IEEE Trans. Info. Theory*, vol. 46, no. 5, pp. 1927–1932, 2000.
- [7] E. Merényi, R. Singer, and W. Farrand, "Classification of the LCVF AVIRIS test site with a Kohonen artificial neural network," in *Proc. Fourth Airborne Geoscience Workshop*, Washington, D.C., Oct. 25–29 1993, pp. 117–120.
- [8] E. Merényi and T. Villmann, "Self-organizing neural network approaches for hyperspectral images," *Proc. International Conference on Intelligent Computing and Information Systems*, pp. 33–41, May 2002.
- [9] T. Villmann, E. Merényi, and B. Hammer, "Neural maps in remote sensing image analysis," *Neural Networks*, vol. 16, no. 3–4, pp. 389–403, 2003.
- [10] A. Plaza and J. Tilton, "Automated selection of results in hierarchical segmentations of remotely sensed hyperspectral images," *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 7, pp. 4946–4949, Jan 2005.
- [11] S. Lee and M. Crawford, "Unsupervised multistage image classification using hierarchical clustering with a bayesian similarity measure," *IEEE Transactions on Image Processing*, vol. 14, no. 3, pp. 312–320, 2005.
- [12] W. Kim, Y. Chen, M. Crawford, J. Tilton, and J. Ghosh, "Multiresolution manifold learning for classification of hyperspectral data," *Geoscience and Remote Sensing Symposium, 2007. IGARSS 2007. IEEE International*, pp. 3785 – 3788, Jun 2007.
- [13] Y. Tarabalka, J. Chanussot, and J. Benediktsson, "Segmentation and classification of hyperspectral images using minimum spanning forest grown from automatically selected markers," *IEEE Transactions on Systems, Man, and Cybernetics - Part B: Cybernetics*, Jan 2009.
- [14] Y. Tarabalka, J. Benediktsson, and J. Chanussot, "Spectral-spatial classification of hyperspectral imagery based on partitional clustering techniques," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 8, pp. 2973–2987, Jan 2009.
- [15] G. Camps-Valls, T. Marheva, and D. Zhou, "Semi-supervised graph-based hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 10, 2007.
- [16] F. van der Meer and W. Bakker, "CCSM: Cross correlogram spectral matching," *Rem. Sens.*, vol. 18, no. 5, pp. 1197–1201, 1997.
- [17] R. Yuhas, A. Goetz, and J. Boardman, "Discrimination among semi-arid landscape endmembers using the spectral angle mapper (SAM) algorithm," in *Summaries of the 3rd Annual JPL Airborne Geoscience Workshop*, ser. AVIRIS Workshop, vol. 1, 1992, pp. 147–149.
- [18] R. Clark, G. Swayze, K. Livo, R. Kokaly, S. Sutley, J. Dalton, R. McDougal, and C. Gent, "Imaging spectroscopy: Earth and planetary remote sensing with the USGS tetracorder and expert systems," *Journal of Geophysical Research-Planets*, vol. 108, no. E12, p. 5131, 2003.
- [19] F. van der Meer, "Spectral curve shape matching with a continuum removed CCSM algorithm," *Rem. Sens.*, vol. 21, no. 16, pp. 3179–3185, 2000.
- [20] —, "The effectiveness of spectral similarity measures for the analysis of hyperspectral imagery," *Applied Earth Observations and Geoinformation*, vol. 8, no. 1, pp. 3–17, 2006.
- [21] S. Robila and A. Gershman, "Spectral matching accuracy in processing hyperspectral data," *Signals, Circuits and Systems, 2005. ISSCS 2005. International Symposium on*, vol. 1, 2005.
- [22] J. Li, D. Hibbert, S. Fuller, and G. Vaughn, "A comparative study of point-to-point algorithms for matching spectra," *Chemometrics and Intelligent Laboratory Systems*, Jan 2006.
- [23] K. Wagstaff, H. Shu, D. Mazzoni, and R. Castano, "Semi-supervised data summarization: Using spectral libraries to improve hyperspectral clustering," *Interplanetary Network Progress Report*, Jan 2005.
- [24] N. Keshava, "Distance metrics and band selection in hyperspectral processing with applications to material identification and spectral libraries," *IEEE Transactions on Geoscience and Remote Sensing*, Jan 2004.
- [25] H. Hertz, R. Hites, and K. Biemann, "Identification of mass spectra by computer-searching a file of known spectra," *Analytical Chemistry*, vol. 43, no. 6, pp. 681–691, 1971.
- [26] H. Lam, E. Deutsch, J. Eddes, J. Eng, N. King, S. Stein, and R. Aebersold, "Development and validation of a spectral library searching method for peptide identification from ms/ms," *Proteomics*, vol. 7, no. 5, pp. 655–667, 2007.
- [27] B. Bornstein, B. Bue, S. Lee, and L. Mandrake, "Autonomous identification and quantification of chemical species with the vehicle cabin atmosphere monitor (VCAM) for use onboard the international space station (ISS)," in *IEEE Aerospace Conference*, March 2008.
- [28] P. Ausloos, C. Clifton, S. Lias, A. Mikaya, S. Stein, D. Tchekhovskoi, O. Sparkman, V. Zaikin, and D. Zhu, "The critical evaluation of a comprehensive mass spectral library," *Journal of the American Society for Mass Spectrometry*, vol. 10, no. 4, pp. 287–299, 1999.
- [29] F. Wilcoxon, "Individual comparisons by ranking methods," *Biometrics Bulletin*, pp. 80–83, 1945.
- [30] —, "Probability tables for individual comparisons by ranking methods," *Biometrics*, pp. 119–122, 1947.
- [31] J. Demšar, "Statistical comparisons of classifiers over multiple data sets," *The Journal of Machine Learning Research*, vol. 7, pp. 1–30, 2006.
- [32] B. Csathó, W. Krabill, J. Lucas, and T. Schenk, "A multisensor data set of an urban and coastal scene," *International Archives of Photogrammetry and Remote Sensing*, vol. XXXII (3/2), pp. 26–31, Jan 1998.
- [33] E. Merényi, B. Csathó, and K. Tasdemir, "Knowledge discovery in urban environments from fused multi-dimensional imagery," *Proc. 4th IEEE GRSS/ISPRS Joint Workshop on Rem. Sens. Data Fusion over Urban Areas*, pp. 1–13, 2007.
- [34] M. Herold, D. Roberts, M. Gardner, and P. Dennison, "Spectrometry for urban area remote sensing—development and analysis of a spectral library from 350 to 2400 nm," *Rem. Sens. of Environ.*, Jan 2004.
- [35] E. Merényi, W. Farrand, L. Stevens, T. Melis, and K. Chhibber, "Mapping colorado river ecosystem resources in Glen Canyon: Analysis of hyperspectral low-altitude AVIRIS imagery," *Proc. of ERIM, 14th Int'l Conf. and Workshops on Applied Geologic Rem. Sens.*, 2000.
- [36] R. N. Clark and T. Rousch, "Reflectance spectroscopy: quantitative analysis techniques for remote sensing applications," *Journal of Geophysical Research-Solid Earth*, vol. 89, no. B7, 1984.
- [37] R. Clark, T. King, and N. Gorelick, "Automatic continuum analysis of reflectance spectra," *JPL Proceedings of the 3rd Airborne Imaging Spectrometer Data Analysis Workshop p 138-142(SEE N 88-13755 05-42)*, 1987.
- [38] M. Craig, "Nonconvex hulls for mineral reflectance spectra," *Applied optics*, vol. 33, no. 5, pp. 849–856, 1994.
- [39] J. Tou and R. Gonzalez, *Pattern Recognition Principles*. Addison-Wesley Publishing Company, 1974.
- [40] J. Schott, S. Brown, R. Raqueno, H. Gross, and G. Robinson, "Advanced synthetic image generation models and their application to multi/hyperspectral algorithm development," *Proceedings of SPIE*, vol. 3584, p. 211, 1999.
- [41] T. Kailath, "The divergence and Bhattacharyya distance measures in signal selection," *IEEE Transactions on Communication Technology*, vol. 15, no. 1, pp. 52–60, 1967.
- [42] S. Sobolev, "On a theorem of functional analysis," *Transl. Amer. Math. Soc.* 34 (2): 3968, vol. 34, no. 2, pp. 39–68, 1963, translation of Mat. Sb. , 4 (1938) pp. 471497.