Interest Segmentation of Large Area Spectral Imagery for Analyst Assistance

Ariel Schlamm, David Messinger, and William Basener

***Abstract*—Widely used methods of spectral clustering, target, and anomaly detection when applied to spectral imagery provide less than desirable results across sensor type, scene content, spectral and spatial resolutions due to the complex nature of the data. This results in a large burden placed on the analyst in terms of the amount of data needed to be processed and the ability to discern the difference between "interesting" and "uninteresting" regions in the imagery. For this research, a variety of data driven algorithms for spectral image analysis are applied to spatial tiles of a large area scene. A feature map is created by assigning a metric determined for each algorithm result to each spatial tile. The feature maps are organized into a tiled, multi-band feature image. Two visualization methods introduced here provide a detection map which can cue image analysts to visually inspect locations within a large area scene with a high likelihood of containing interest. Unsupervised classification is applied to this feature image such that the image is divided into segments representing either "interesting" or "not interesting" content with the tile. False-color visualization of three independent metrics is also presented as a way to indicate the type and strength of the amount of interest within a tile.**

***Index Terms*—Clustering methods, detectors, hypercubes, image analysis, image region analysis, image segmentation, remote sensing, search methods, spectral analysis.**

I. Introduction

MUCH of the automated analysis of spectral imagery is signature-based, meaning that there are specific material signatures of interest that are being located and that pixels are exploited individually in order to determine whether or not they contain these targets. Detection of these materials typically involves algorithms that place geometric and/or statistical assumptions on the data. Statistical methods describe the distribution of the data in the hyperspace and are often multivariate Gaussian and T-Elliptical assumptions [1]. These can be used in classification, target, or anomaly detection schemes. Geometric assumptions placed on the data rely on defining basis vectors, which may or may not be orthogonal, to describe how the data lie in the space. They are commonly divided into two categories: vector subspace and linear mixture models. If the data lie in a vector subspace, then a set of orthogonal basis vectors can be found to describe the entire data set where the number of basis vectors is less than the original number of bands measured. Linear mixture models (LMM) are based on the assumption that the data lie in a convex hull [2]. The data can be described by the corners of the simplex or endmembers (or non-orthogonal basis vectors) and each point contained in the convex hull can be expressed as a linear combination of the endmembers. The number of endmembers is typically less than the number of dimensions measured. Linear unmixing uses the LMM to determine the fractional amount of each endmember in each pixel. [3] Orthogonal subspace projection (OSP) is a target detector that minimizes the background contribution through a projection onto the endmembers, and then projects the pixel onto the target to measure the similarity [4]. Statistical target detection algorithms such as the spectral matched filter (SMF) [5] use the mean and covariance of the distribution to determine whether the target is present in the data. Other statistical target detectors are variations of SMF and include constrained energy minimization (CEM) [2], and adaptive coherence estimator (ACE) [6]. These, and other target detection methods or distance metrics, can be used to measure the similarity of a pixel to a particular target. In low spatial and spectral resolution imagery, each pixel is a mixture of many materials and the data distribution is often sufficiently represented by these assumptions. However, as the current generation of sensors typically have higher spatial resolution, the complexity of the data collected is increasing and these assumptions are no longer adequate. As a result, algorithms based on these assumptions provide inconsistent and unpredictable results when applied to data that is more complex [7], [8]. New algorithms for hyperspectral image analysis, including dimension estimation [9], [10], anomaly detection [11], spectral volume estimation, spectral clustering [12], [13], and complexity [14], that are inherently data-driven and reduce the number of assumptions placed on the data have been developed for improved exploitation of spectral imagery at any spatial resolution.

Wide area scenes provide another source of added complexity to spectral image analysts posed with this task. Large area search is a difficult problem for many reasons. While algorithms exist for automated target detection and identification, within large area search the analyst may not always have a specific material of interest; however, he/she must find everything in a scene that is possibly interesting. As a result, a higher false alarm rate is tolerated to guarantee a high probability of detection. Additionally, the background scene content can change significantly across the collected area, which can affect the performance of scene-wide algorithms. The amount of data collected

Manuscript received July 20, 2011; revised December 13, 2011; accepted March 04, 2012. Date of publication May 09, 2012; date of current version May 23, 2012.

A. Schlamm and D. Messinger are with the Center for Imaging Science, Rochester Institute of Technology, Rochester, NY 14624 USA (e-mail: schlamm@cis.rit.edu).

W. Basener is with the School of Mathematical Sciences, Rochester Institute of Technology, Rochester, NY 14624 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/JSTARS.2012.2195298

by sensors that can revisit the same area nearly once a day, including WorldView2, for large area search also requires a significant time investment from the analyst, who is typically tasked with visually inspecting individual pixels. However, much of the area is expected to be uninteresting with only some regions of the scene containing significant and valuable information. Currently, much of the large area search process is performed by visual inspection on high resolution imagery. Visual inspection is often a time intensive process in a field where many problem sets are time sensitive. The Wide Area Reconnaissance Hyperspectral Overhead Real-time Surveillance Experiment (WAR HORSE) was built by the Naval Research Laboratory in order to function as a real time detection and cueing operation. The imaging system consists of two line scanners, a visible hyperspectral sensor, and a high resolution visible sensor. A subspace RX detector is used as the data are collected. When an anomaly is found in the hyperspectral data, the high resolution imager is cued and a high resolution image is taken over the identified anomaly [15]. Other than the WAR HORSE system, hyperspectral imagery is not typically used for wide area search. This significantly reduces the amount of high resolution data an analyst must inspect.

Many of the algorithms used in spectral image analysis indicate whether interest is likely contained in a particular pixel, such as automatic target detection or anomaly detection. In this case, areas containing any manmade material of unknown specific signature are considered interesting and natural materials are not. For the problem of large area search, it may be more useful for an analyst if algorithms were used to indicate whether interest is contained in a particular region, in this case spatial tiles. Additionally, the processing of spatial tiles allows for the use of distribution-based metrics that are often overlooked in spectral image processing. This information can be used to divide the image into interest segments. Low or no interest tiles can be removed from the data or ignored by the analyst, significantly reducing the amount of data he/she must look at. The interest segments or individual tile scores can also be used to prioritize the task by directing the analyst to inspect high interest areas before low interest areas. Regions labelled "interesting" can be used to cue a sensor to collect additional data of the same region, either high resolution or in a completely new modality, or as input to pixel-based analysis.

In this research, a selection of data driven algorithms have been applied to hyper- and multispectral image tiles. These algorithms are briefly introduced in Section II. Each method provides at least one metric relating to the amount of interest contained within each tile. The data used for this analysis are described in Section III. A selection of individual feature results are presented in Section IV-A. These metrics are combined using k-means unsupervised classification in order to divide the image into segments representing "interest" or "no interest" in order to assist an analyst in prioritizing visual inspection of the data. False-color visualizations of the feature maps are shown to indicate the degree and type of interest contained within image tiles. The final interest segmentation results for analyst assistance are presented and discussed in Sections IV-A and IV-B.

II. Approach

A. Methodology

Large area search of imagery is a challenging task to both human analysts and automated algorithms. Unlike target detection, the goal of large area search does not generally include identifying specific targets with known spatial characteristics or spectral signatures. Instead, the goal is to "cue" an analyst to individual regions (not pixels) of interest for further analysis. The later analysis can be visual inspection, automated algorithms, or collection of higher resolution imagery. Large area search requires a very high probability of detection (as any missed detection is a failure), while accepting a moderate number of false alarms. Here, we are interested in identifying localized regions of a large multispectral or hyperspectral image that contain evidence of man-made activity. Typically, this problem has been addressed through the application of change detection methodologies, automated target detection, or anomaly detection on a pixel level, which have their own challenges. The approach taken here is based not on statistical models of the data, nor is it based on detection of changes between imagery.

This methodology for large area search is presented in Fig. 1. First, the image is divided into spatial tiles. A series of algorithms are applied to each image tile. These algorithms, described in Section II-B, estimate metrics based on the spectral distribution of the data, each in a unique way. The result of each algorithm is a numeric score between zero and one, which is then assigned to that tile. The result of this step is a multi-band feature image, of equal size to the original image, where each band corresponds to the result of a particular algorithm. Each band of the image can then be presented to the user with individual tile brightnesses adjusted to the complexity of that tile relative to the rest of the image. The multi-band feature image can be put through an unsupervised clustering algorithm, such as k-Means, to cluster the image into "interesting" and "uninteresting" classes. This segmentation map can be presented to an analyst in order to help prioritize the search. Additionally, any three of the feature bands can be used to create a false color image to visualize the results. In this case, the color of the image tile corresponds to the metric(s) that scored the highest and brightness of a tile corresponds to the value of the metric(s).

B. Metrics

The estimated inherent dimension of spectral imagery is one way to characterize the "size" of a distribution in the native hyperspace. The point density approach to hyperspectral image analysis has been introduced as a method for estimating the inherent dimensionality of hyperspectral data [9], [16]. This method creates a point density plot (PDP) by counting the number of data points that fall within spheres of increasing radii around a center point and plotting the result in log-space, shown in Fig. 2. The point density dimension (PDD) is found by estimating the slope of the primary incline of this plot. As a result of making this plot, another measure indicating the amount of interest contained has been determined: the point density tail length (PDTL) [17]. These two metrics are

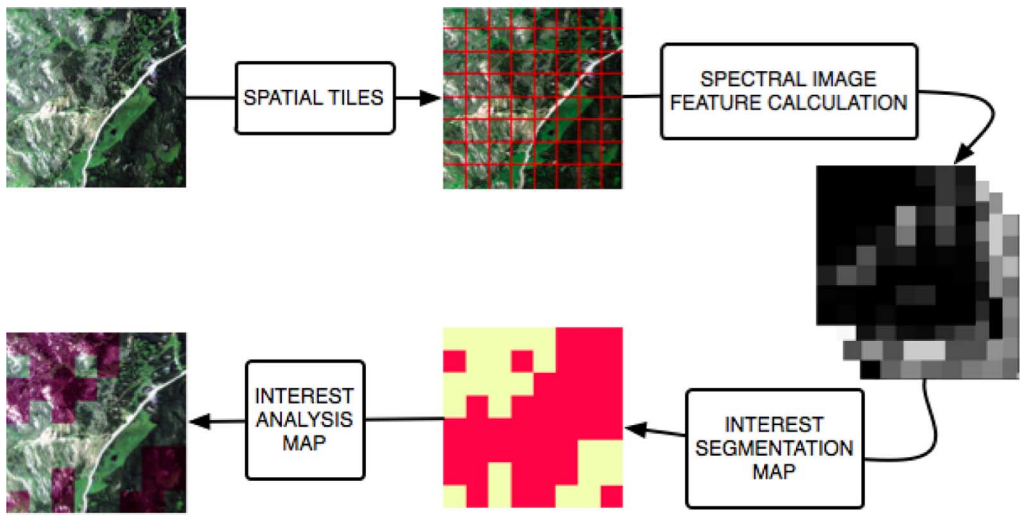

Fig. 1. The interest segmentation process.

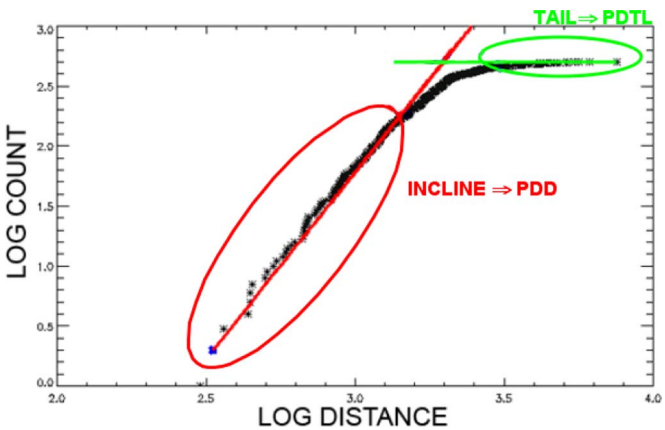

Fig. 2. Example point density plot (PDP) of real data showing the two metrics estimated from point density approach.

illustrated in Fig. 2. The short study in Schlamm, *et al.* [17], showed that, in general, as manmade spectra were implanted into natural material hyperspectral tiles, each of these metrics tended to gradually increase. Therefore, the PDD and PDTL relate to the amount of interest within an image tile. In particular, a larger PDTL indicates the presence of anomalous pixels. The PDD relates to the material type and amount of clutter within the tile.

Estimating the multivariate normality of a high dimensional data set is not an easy task. Many of the methods for accessing the univariate normality of a data set, including the Mardia, Kolmogorov-Smirnov, Shapiro-Wilk and other tests based on estimating high order statistical moments or graphical "goodness of fit" metrics, have been extended to the multivariate case but do not provide reliable results for high dimensional data, such as spectral imagery [18]. Mecklin and Mendfrom (2004) point out that though over 50 methods for estimating the multivariate normality of data sets can be found in the literature, many of them do not perform consistently when applied to high dimensional data. Point density estimation provides an empirical way to quantify and compare the multivariate normality of high dimensional data distributions [19]. For each PDP created from true image data, a second PDP can be created from random data with the same mean, covariance, and number of pixels as the true image data. The average error between these two PDPs is calculated and used as a measure of the multivariate normality (MVN) of the data. This operates under the assumption that the spectral distributions of natural materials are better approximated by a normal distribution than the distributions of imagery containing manmade materials. This is illustrated in Fig. 3 where the image-based PDP is shown in black, a PDP generated from random data with the same statistics is shown in red, and the corresponding error is displayed in green. Fig. 3 shows that natural material data (a) has a lower MVN error score than the data containing manmade materials (b), meaning that manmade materials do not follow a multivariate normal distribution.

Another way to characterize the "size" of a spectral distribution in the hyperspace is to estimate the volume of the distribution. Messinger, *et al.*, present an iterative method which uses the framework of the linear mixture model to calculate the volume of the parallelotope that encloses the data as a function of the number of dimensions used in the calculation. This is done through the use of the Max-D algorithm [20], [21] and the Gram Matrix [22], [23], which both have unique properties useful for this approach to volume estimation. The Max-D algorithm returns the endmembers in order of decreasing magnitude of separation.The determinant of the Gram Matrix, termed the Gramian, is the squared volume of the parallelotope formed by the vectors in the test set [24]. A parallelotope is an n-dimensional parallelogram. Due to the convex geometry of the data, the volume of the parallelotope that encloses the data is the same as the volume of the parallelotope defined by the endmembers. The Gramian is calculated iteratively as a function of the number of endmembers used in the calculation, producing a plot shown in Fig. 4. The three metrics estimated, shown in Fig. 4, that relate to the complexity of a spectral image tile include peak volume, area under the curve, and number of endmembers where the volume goes to zero. Additionally the pro-

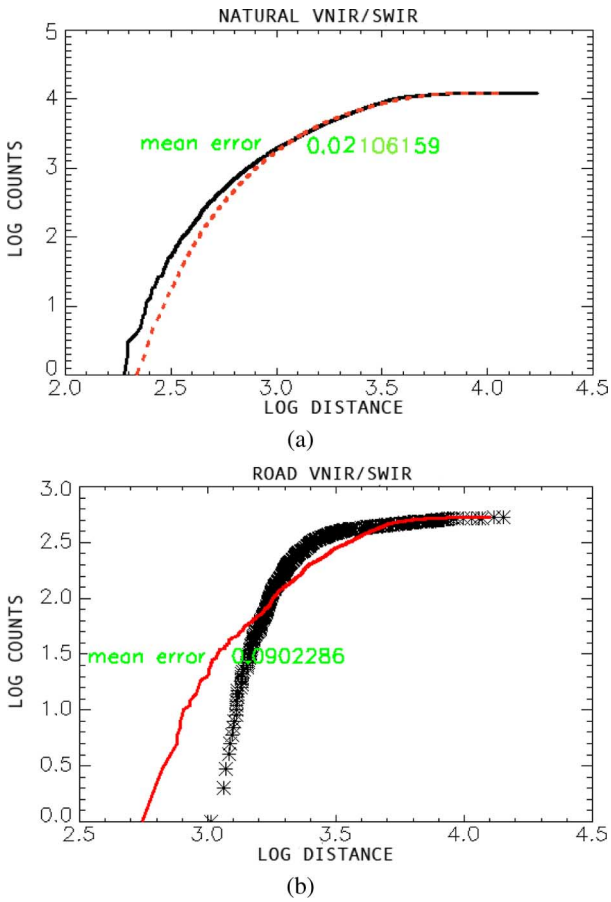

Fig. 3. Example PDPs with corresponding MVN PDPs and error scores for a region of pure natural material (a) and manmade materials (b).

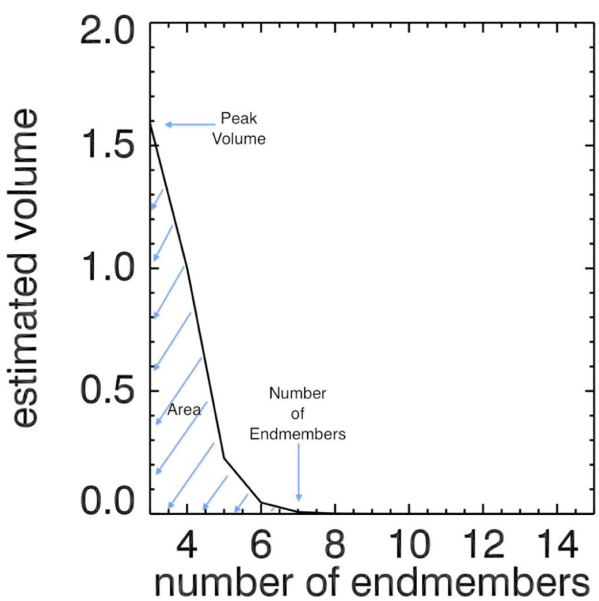

Fig. 4. Example Gram matrix plot estimated from real data showing the three metrics estimated from this approach.

portion of the total area under the volume curve that is the peak volume is used as a complexity metric.

An image tile containing a large number of anomalies or many material clusters is more complex than a tile containing only one cluster and no anomalies. Topological analysis and graph theory provide ways to make data driven measurements on the data. The Topological Anomaly Detector (TAD) [11] uses methods from these fields in order to locate spectral anomalies. This algorithm connects points in the hyperspace that are spectrally "close" to each other in order to build a simplicial complex, or a graphical representation of the background distribution based on small simplex formation. This process may find a single or many graph components of a variety of different sizes. As anomalies have a low frequency of occurrence, the small or single point components are labeled anomalies while the large components are labeled background, notionally shown in Fig. 5. How anomalous these pixels are is determined by measuring the distance between the anomaly and the nearest few pixels connected to a background component. This method has been shown to consistently perform well across a variety of scenes, sensors, ground sample distance (GSD), and image content [25]. Basener, *et al.* [12], introduce a clustering algorithm, named the gradient flow algorithm, which clusters points together based on the direction of increasing density. For each point in the hyperspace, the n nearest neighbors are found and the density of each point with respect to the n neighbors is calculated. A smoothing step locates points that are in the center of high density regions. Finally, the gradient differential is calculated; for each point, this will find the direction of the highest increasing density and "push" the point in that direction, thus assigning it to the same cluster as other points associated with this high density sink. This step is iterated until a steady state is reached and a cluster map can be displayed in image form. This algorithm is illustrated in Fig. 6. This algorithm requires no prior estimation of the number of classes in the scene and is not restricted by geometric or statistical assumptions. For this search methodology, anomaly detection and unsupervised clustering are performed image-wide. The detection and/or clustering maps are tiled. The feature is calculated by counting the number of anomalies and/or clusters per tile. Any preferred anomaly detection or spectral clustering routine may be used in place of these, such as the RX anomaly detection [26] or Isodata for clustering [2].

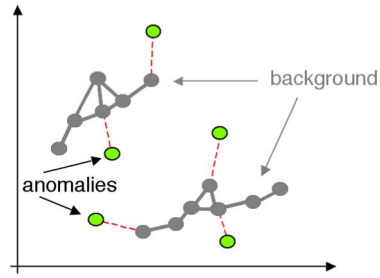

Fig. 5. Notional result of TAD on a simple data set.

From these methods, nie individual mathematical features of spectral image data will be used, including PDD, PDTL, MVN, number of anomalies (NANOM), number of spectral clusters (NCLUST), peak volume (PV), area under the Gram matrix plot (AREA), ratio of peak volume and area (RATIO), and number of endmembers for the volume to reach zero (NEM). These features, shown in Table I, characterize the distribution complexity of spectral data in the native hyperspace, each in a unique way,

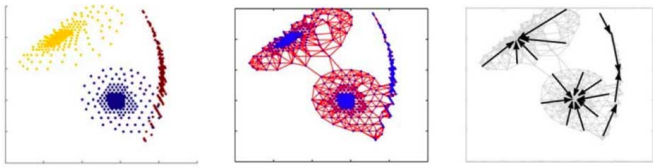

(a) (b) (c)

Fig. 6. Notional description of the Gradient Flow algorithm. The points in (a) are connected to their nearest neighbors in (b). The direction of highest density change are shown in (c).

TABLE I
INTEREST SEGMENTATION METRICS

Metric	Abbreviation	Approach
number of endmembers where volume function equals zero	NEM	simplex volume
peak of volume function	PV	simplex volume
area under volume function	AREA	simplex volume
ratio of area to peak	RATIO	simplex volume
point density dimension estimate	PDD	point density
point density tail length	PDTL	point density
multivariate normality error	MVN	point density
number of spectral clusters	NCLUST	external algorithm
number of spectral anomalies	NANOM	external algorithm

and do not attempt to identify spectral features based on material property or in terms of wavelength dependencies.

C. The Effect of Tile Size on Metrics

Some of the metrics introduced are dependent on the tile size used in the metric calculation. The distribution-based metrics, including those based on the Gram matrix and the PD approaches, have the potential to be the most effected by the chosen tile size. According to Camastra, [27] there exists a requirement on the number of data points needed to accurately estimate a fractal dimension D, such as the point density dimension. The number of data points must satisfy

$$D < 2 \log N \tag{1}$$

where N is the number of data points in the set. For a 100 by 100 pixel image, this implies accurate estimation of a dimension up to 8. Practically, the fractal dimension of hyperspectral imagery or image tiles has not beed found to be larger than 6 [9]. An image tile size of 50 by 50 satisfies this criteria. Fig. 7, shows the estimated dimension of hyperspectral data over a residential neighborhood as a function of the number of pixels used in the calculation. With the exception of the 2500 pixel sample, the estimated dimension is fairly constant. The PDTL metric is essentially a local anomaly detector. The ideal tile size for this metric has to do with how different the PDTL is when an anomaly is and is not present. The effect of tile size on the PDTL metric is shown in Fig. 8. The two plots are most distant from each other at smaller tile sizes and begin to converge at larger tile sizes. This occurs because as the tile not containing anomalies grows bigger, the distribution better approximates a normal distribution and the PDTL increases. In this methodology, the anomaly detection and spectral clustering and performed image-wide so the NANOM and NCLUST metrics are less effected by the tile size chosen. As a result, a larger tile size may result in more anomalies and/or clusters per tile. The Gram matrix based methods are similarly effected by tile size. The volume estimated and number of endmembers cannot go down as tile size increases. All of the results presented here use a tile size of 30 by 30 as this meets the number of samples criteria for the point density approach, but is small enough to provide granularity and detail across the scene.

III. DATA

Multispectral data from DigitalGlobe's [28] WorldView2 (WV2) and Quickbird sensors are used for this analysis. WV2 is an 8-band VNIR/SWIR multispectral sensor with 2 m GSD launched late in 2009. Sample imagery of a coastal city, shown in Fig. 9(a), was made freely available by Digital Globe. QuickBird is a 4-band VNIR sensor with 2.5 m GSD [28]. A portion of a large QuickBird dataset, shown in Fig. 16(a), of the Esperanza forest fire in California is used for this analysis. This data set covers the area near the fire origin and contains, a dry wash and a suburban area. Hyperspectral data from the HyMap sensor [29], [30] collected over Cooke City, MT area on July 4, 2006 is also used in this analysis, shown in Fig. 12(a). HyMap is a commercially flown airborne hyperspectral sensor with 126 usable bands (after the atmospheric absorption bands are removed) between 450 nm and 2500 nm with 2.5 m ground sample distances (GSD). The scene is predominately natural material, including forest, farmland, grass, foothills, rock, and dirt roads. The manmade material in the scene includes downtown Cooke City, a residential neighborhood, and a construction site, all shown in red boxes.

IV. RESULTS

A. Feature Results

The individual feature maps for the WV2 sample imagery are shown in Figs. 9 and 10. In each image, tile brightness corresponds to the magnitude of the metric score. The NEM feature map, shown in Fig. 9(b), only shows two tile brightness, or scores of either zero (black) or one (white). In this case, a score of zero corresponds to the algorithm finding 4 endmembers while a score of one corresponds to 5 endmembers. Regions of manmade activity contain more varying types of materials than a natural regions. As a result, tiles with more endmembers contain more unique materials and are more interesting. Because this is 8 spectral band data, finding more endmembers than this is not expected. The black regions of the image correspond to the water and the tiles that contain only natural materials (trees) while the white tiles are over the manmade regions, which would likely contain more material variability and therefore more endmembers, than the natural material regions. The PV, Area, and Ratio feature maps in Fig. 9(c)–(e) are very similar to each other because the metrics are highly correlated. The PV metric estimates the peak volume from Gram matrix plot in Fig. 4 and the Area metric measures the area under the entire

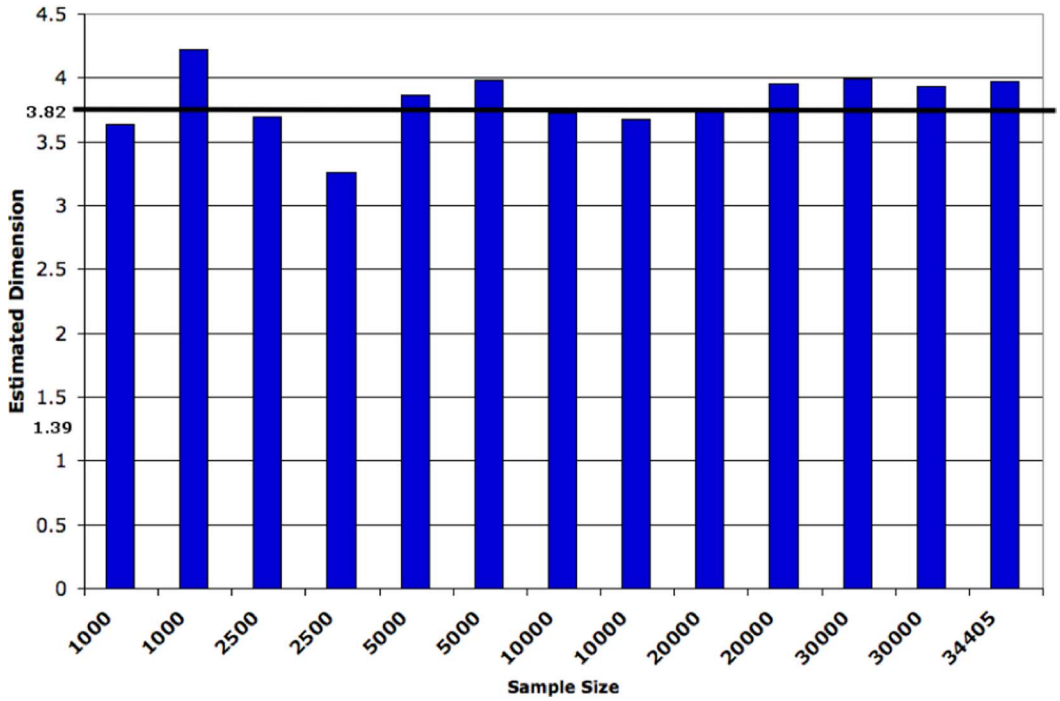

Fig. 7. Bar chart of point density dimension estimates as a function of tile size for hyperspectral data containing a residential neighborhood.

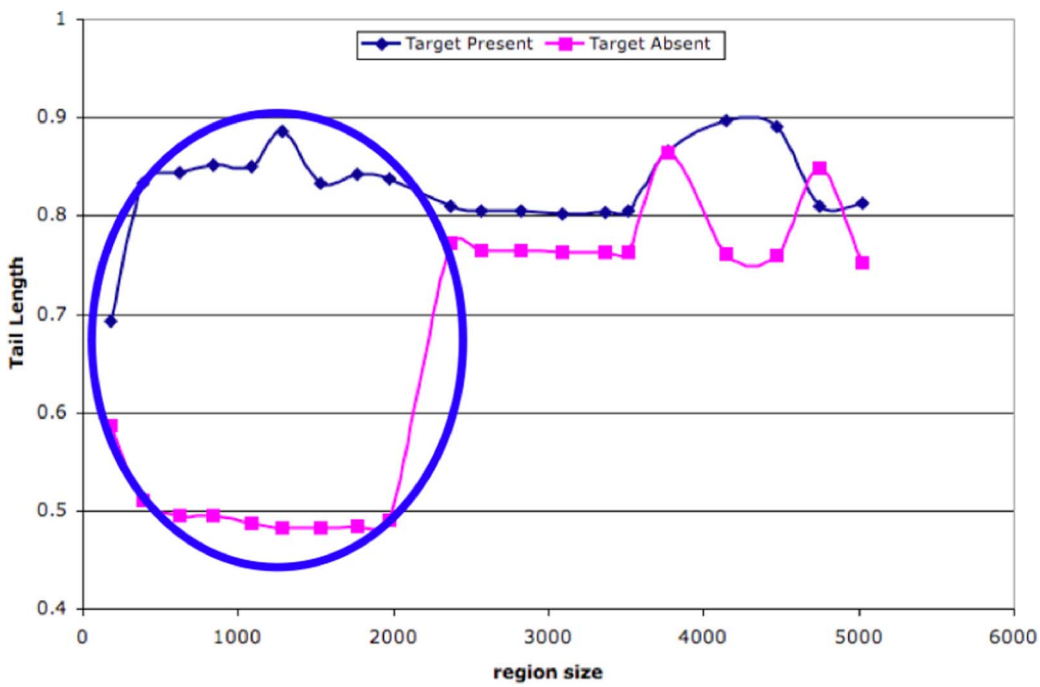

Fig. 8. The estimated PDTL as a function of tile size for an uninteresting tile and a tile containing anomalies.

Gram matrix curve. The Area metric is driven by the PV. As a result, when the PV is high, the Area is also high. However, it is possible to have two plots with the comparable area metrics with two significantly different volumes if tile has a large number of endmembers. In each case, the ocean and vegetation-covered areas are mostly dark while the urban areas are bright. The brightness of the tiles over land correspond to the spatial density of manmade structure within the urban area. Near the center and the upper left side of the image are two areas that are very densely populated and have no visible vegetation. An example of this area is shown in Fig. 11(a). Between these areas is a strip of land that shows up as much brighter in all three of these metrics. This is a less dense urban area containing small parks and lawns, shown in Fig. 11(b). The contrast between these areas is best seen in the Ratio feature map. The addition of vegetation causes this distribution to occupy a larger volume in the hyperspace, and therefore increases the Gram matrix-based metrics. Some tiles in the water also have large metrics values from

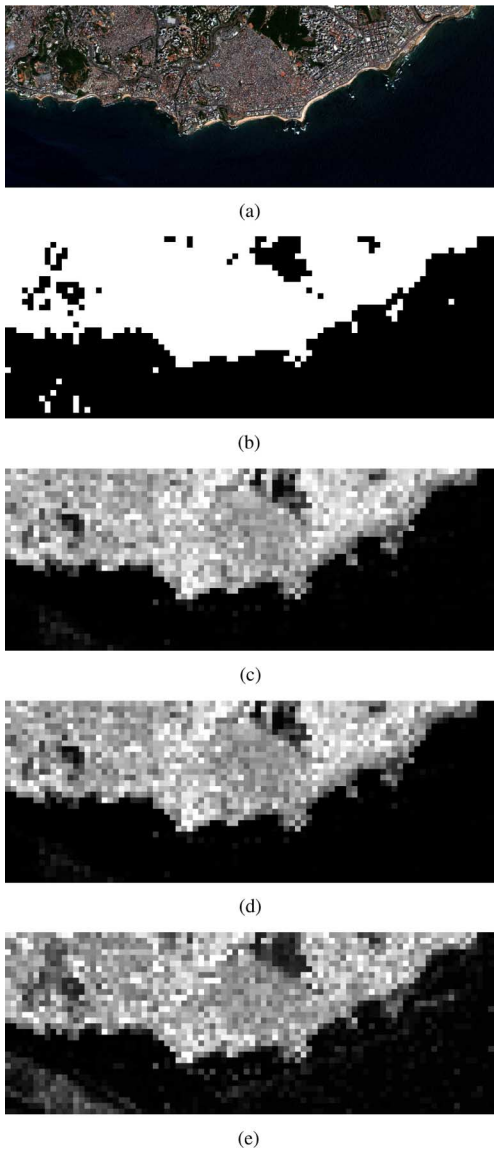

Fig. 9. True color image from the WV2 sample image (a) and corresponding gram matrix-based feature maps (a) RGB (b) NEM (c) PV (d) Area (e) Ratio.

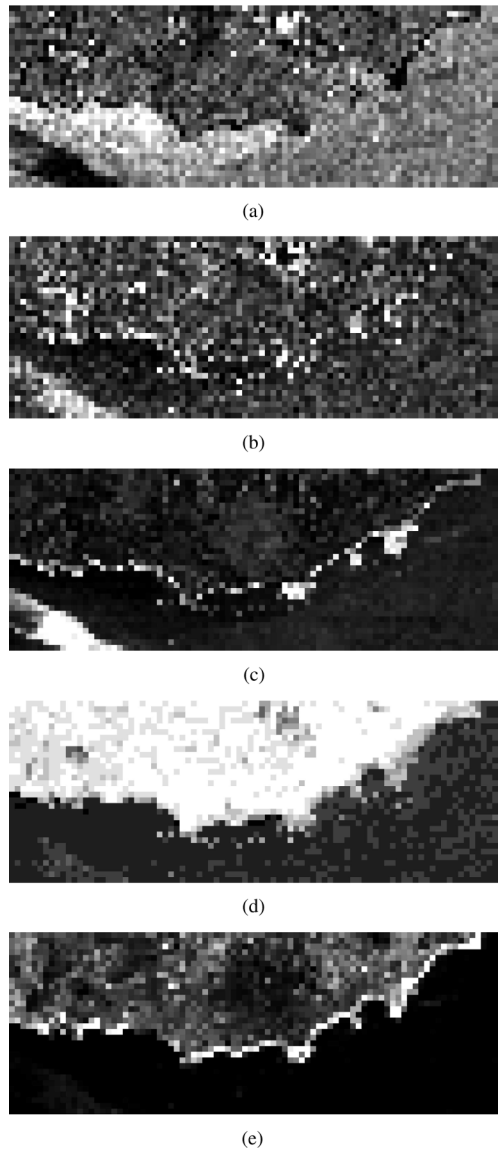

Fig. 10. Point density, cluster, and anomaly-based feature maps for the WV2 sample imagery (a) PDD (b) PDTL (c) MVN (d) NCLUST (e) NANOM.

these algorithms. In the lower left corner, the stripe of bright tiles corresponds to an area of colored noise, possibly an artifact from the registration procedure. The bright tiles near the coastline correspond to boats in the water. These areas are shown in Fig. 11(c) and (d).

For most metrics used, a bright tile corresponds to an interesting tile, the exception being PDD, shown in Fig. 10(a). In this case, the tiles with higher inherent dimension estimates are those of pure materials, including the water and vegetation regions. Tiles containing material mixtures generally have significant empty space in the hyperspace and therefore a lower dimension. The PDTL and MVN metrics are brightest over the coastline and the noise artifact in the lower left, as these deviate the most from the ideal distribution on which the PD approach is based. The MVN metric also finds higher scores for the very dense urban areas visible in the Gram-based feature bands, meaning these areas are less multivariate normal that areas containing some natural materials. Both the PDTL and MVN metrics also find the boats just off the coastline. The NCLUST feature map segments the image into tiles with between one and eight spectral clusters. Tiles with more clusters are over the urban areas while tiles with fewer clusters are predominately over the water and vegetation areas. The NANOM map highlights the coastline, similar to the PDTL and MVM feature bands; however, it provides better discriminability over the dense urban areas versus the less dense urban areas with vegetation.

The corresponding feature maps are shown for HyMap data in Figs. 12 and 13. While the NEM feature map for WV2 data only shows either 4 or 5 endmembers per tile, when measured for hyperspectral data many more endmembers can be found due to the increased number of bands. This is visible in Fig. 12(b). The areas with the largest number of endmembers are over downtown Cooke City, the residential neighborhood, and the construction site, shown in Fig. 14. The PV, Area, and Ratio feature maps produce similar results, highlighting the same, populated areas as well as the entire road system within the image.

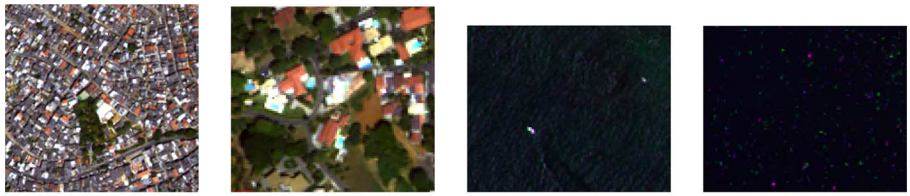

Fig. 11. True color images of the WV2 sample data over two urban areas (a–b) and two water areas (c–d) (a) Dense Urban (b) Suburban (c) Boats (d) Noise.

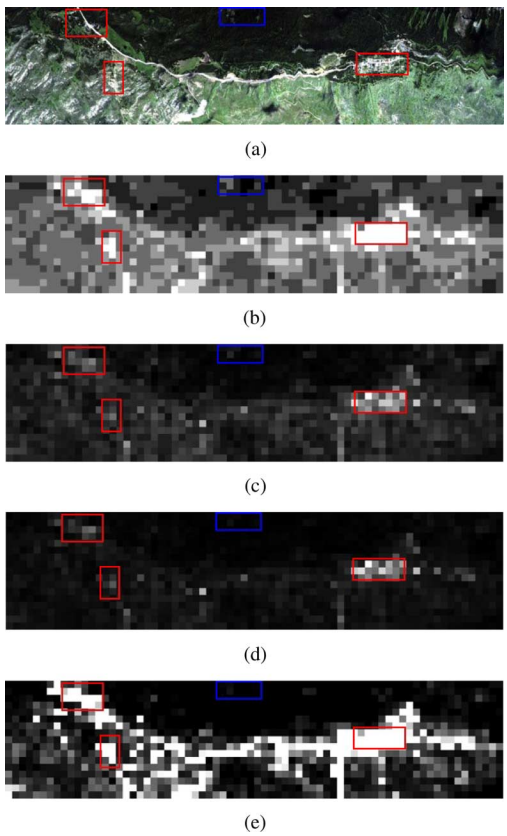

Fig. 12. True color image from the HyMap data (a) and corresponding gram matrix-based feature maps (a) RGB (b) NEM (c) PV (d) Area (e) Ratio.

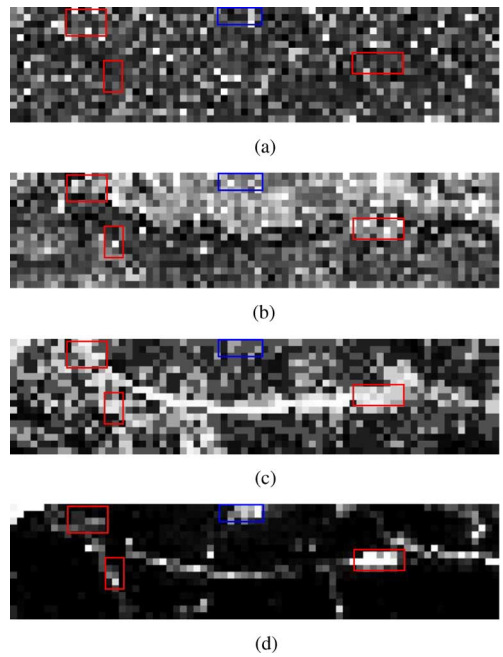

Fig. 13. Point density, cluster, and anomaly-based feature maps for the HyMap data (a) PDD (b) PDTL (c) NCLUST (d) NANOM.

Fig. 14. True color images of the HyMap data over downtown Cooke City (a), a construction site (b), and a residential neighborhood (c) (a) City (b) Construction (c) Residential.

Similar to the WV2 result, the Ratio metric produces a higher contrast feature map, exaggerating the differences between the manmade and natural material regions of the image. The PDD map for WV2 is useful for discriminating between the urban and water regions of the scene. Due do the high dimensional nature of hyperspectral imagery, the estimated dimension per tile is much less consistent across neighboring tiles and the entire image. As a result, the PDD map in Fig. 13(a) shows this variability and in this case is a less useful metric. The PDTL, NCLUST, and NANOM feature maps all detect the manmade, complex regions of the image.

B. Interest Segmentation Results

The interest segmentation results for the sample WV2 imagery multi-band feature maps from Section IV-A are shown in Fig. 15. The multi-band feature image is used as input for k-Means unsupervised clustering for $k = 2$ (Fig. 15(a)). In this case, because of the extreme difference between the water and land areas, the two segments correspond roughly to urban and natural areas. The two areas of pure vegetation fall into the same segment as the water because they are both natural materials, which manifest in the hyperspace similar to each other and therefore have similar scores from all of the metrics. Dividing the image into more than two segments provides more discriminability between interest levels on and off the land. Fig. 15(b) shows the interest segmentation result where $k = 4$. In this case, the water segment is divided into two regions: the main water body and water tiles which contain point anomalies, mainly the boats and noise artifacts from Fig. 11(c) and (d). The landmass is divided into three regions. Tiles only containing vegetation are assigned into the same segment as the water, similar to the 2-class result in (a). The coastal regions and dense urban areas

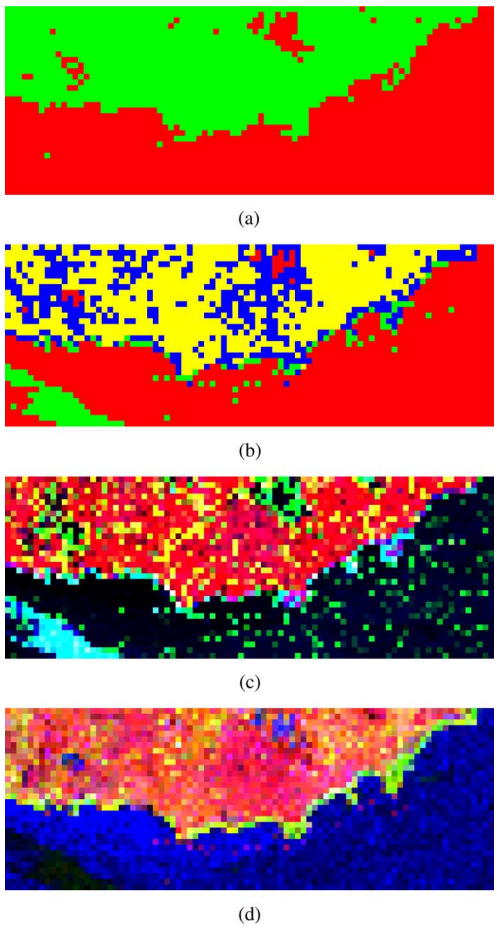

Fig. 15. Two-class (a), four-class (b), and two false color (c–d) interest segmentation maps for the WV2 sample imagery (a) INTEREST SEGMENTATION: 2 classes (b) INTEREST SEGMENTATION: 4 classes (c) R: Ratio, G: PDTL, B: MVN (d) R: NCLUST, G: NANOM, B: PDD.

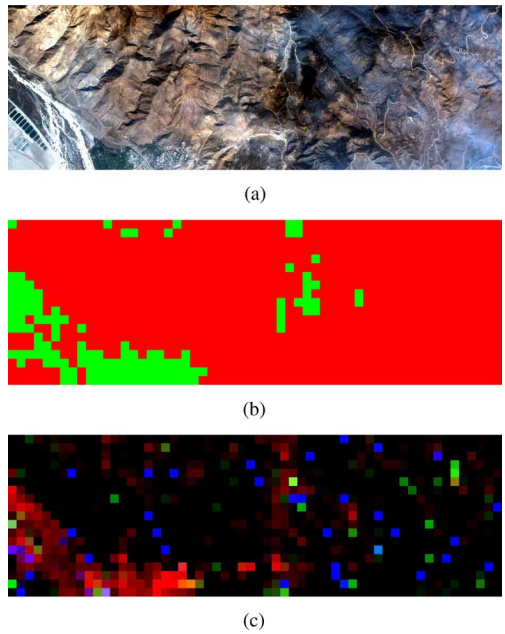

Fig. 16. RGB image (a), interest segmentation map (b), and false color visualization (c) for Quickbird imagery near the origin of the Esperanza forest fire containing multiple regions of interest (a) RGB IMAGE (b) INTEREST SEGMENTATION (c) R: PV, G: PDTL, B: MVN.

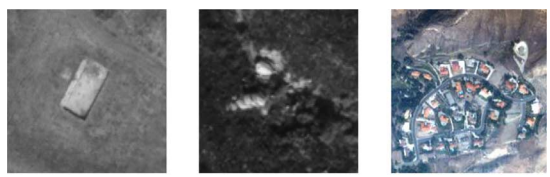

Fig. 17. High resolution tiles corresponding to "interesting" tiles in the interest segmentation map in Fig. 16(b).

are assigned into the blue segment, as they have similar spectral complexity levels. The less dense urban areas are assigned separately into a yellow segment.

Fig. 16(a) shows a region of a Quickbird image near the location of the Esperanza wild fire, demonstrated by the black burn scar and smoke on the right side of the image. The left side of this image contains a small town area and a dry wash. Upon broad inspection, most of this image contains exposed soils and nothing of interest. Approximately 70% of this image was assigned into the "no interest" (red) segment and 30% of the image is labeled "interesting" (green) in Fig. 16(b). All of the tiles over the small town, factory, and dry wash area are assigned into the interest class because they are spectrally very different from the background of the image. Twenty additional tiles, or around 10% of the image, are also assigned to the interest segment. Example high spatial resolution data of a selection of these "interesting" tiles are shown in Fig. 17. These tiles contain small buildings and point anomalies below resolution of the sensor.

The interest segmentation map for HyMap data over Cooke City using all of the previously introduced spectral feature metrics is shown in Fig. 18(a). In this case, K-Means is used to cluster the multi-band feature image into three clusters. Green and blue represents "interesting" while red represents "not interesting." The green segment contains the city, the main and side roads, the construction site, the neighborhood of houses, and the region in the forest that may contain something of interest. It also contains a few forest tiles that are probably not interesting, a clearing in the grass on the bottom left, and a clearing in the forest on the top right. However, though these regions do not contain something as interesting as a construction site, they are unique when compared to the rest of the scene in terms of spectral diversity and distribution size/shape. The blue region contains a few of the tiles over downtown Cooke City. In this case, approximately 85% of the image is put into the "not interesting" segment, significantly reducing the amount of area an analyst must inspect. The three very unique, manmade regions in this image shown in Fig. 14 are found to be interesting, while much of the road and near-urban areas are missing from this segmentation map. Fig. 18(b) shows the interest segmentation maps of the HyMap data over Cooke City using five of the the previously defined spectral features: NANOM, NCLUST, PDTL, PDD, and PV. In this case, using fewer metrics produces a map which includes all of the known regions of manmade activity and more of the areas road network, indicating that all

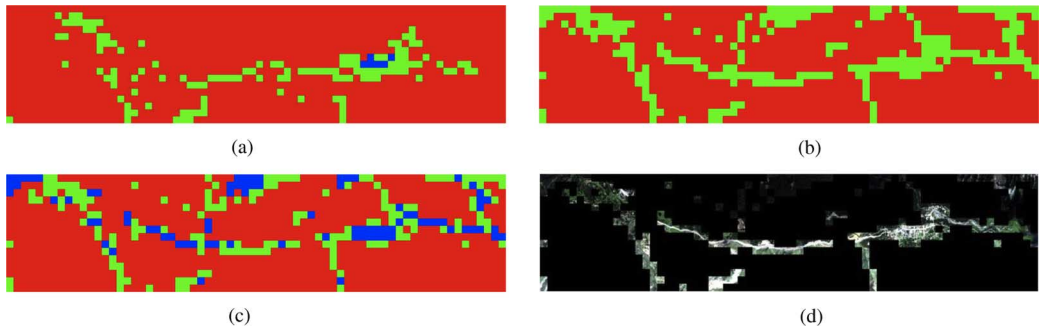

Fig. 18. Two-class (a) and three-class (b) segmentation maps using five features, three-class segmentation map using all features (b) and interest on (d) detection maps for the HyMap imagery (a) INTEREST SEGMENTATION: 3 classes, All features (b) INTEREST SEGMENTATION: 2 classes, 5 features (c) INTEREST SEGMENTATION: 3 classes, 5 features (d) INTEREST ON.

of the features are not necessary to produce a useful and accurate interest segmentation map and that in fact, using too many metrics may actually reduce the quality of the result. Fig. 18(b) shows the interest segmentation result when K-Means is used to divide the five-feature image into three clusters. Red is still "not interesting" while blue and green are both "interesting." The large blue region corresponds to the city. The blue tiles in the neighborhood are tiles that do contain at least one house. The construction site falls completely in a blue tile. In this case, the blue "interest" segment can be labeled a "high priorty" or "high clutter" interest region. The green segment contains those tiles that are more complex than the trees and grass, but perhaps less spectrally diverse than the blue segment. The blue region in the top center correspond to the region in the trees that had a high score in each of the five original features. Fig. 18(d) shows another product that may be delivered to an analyst. Based off the two interest class map, the tiles that are "not interesting" may be turned completely off, letting an analyst focus only on those areas likely to contain manmade activity.

C. False-Color Visualizations

Applying unsupervised clustering to the multi-band feature image results in a discrete segmentation map. Each tile is determined to be either "interesting" or not interesting. However, a way to visualize *how* interesting a particular tile is relative to the the rest of the image may be more useful for an analyst tasked with the search problem. The individual feature bands in Figs. 9 and 10 can be used individually to do this. However, many of these features highlight different quantities of the spectral distribution, including overall distribution volume, dimension, number of clusters, and number of anomalies. Because some of the features are independent from others, displaying multiple feature bands at once can better indicate to an analyst not only whether a particular tile is "interesting" but also in what way. Three metrics can be selected and displayed in false color. In this way, the tile brightness indicates the likelihood of the tile containing "interest" and the color indicates the particular metric(s) that detected this interest. Different metric combinations are optimal depending on sensor type, spectral sample, and scene content. A variety of combinations are presented here to show the utility of all of the metrics.

Fig. 15(c) and (d) show two different false color visualizations of the WV2 sample imagery. Fig. 15(c) shows the Ratio, PDTL, and MVN feature maps in the red, green, and blue channels, respectively. In this case, the feature bands are thresholded such that every tile that scored above 50% of the max score is displayed. The urban areas are predominantly detected by the Gram matrix-based ratio metric because the spectral diversity of the materials in the urban area cause the volume of the distribution to increase. The very dense urban area in the middle, with no vegetation, has a slight magenta cast due to the response from the MVN metric. This area, with absolutely no vegetation, is very dissimilar to a multivariate normal distribution and therefore has a high MVN score. All of the boats off the coast are detected by the PDTL metric, which is mostly sensitive to point anomalies. Many other tiles in the water are also detected by the PDTL. These are due to glint on the water surface, which manifest as anomalies in the spectral distribution. The two areas of pure vegetation are mostly black, meaning they were not detected by any of the metrics. The tiles on the edge of these vegetation areas appear green in this display due to the response from the PDTL metric. In this case, the PDTL is very high not because there are point anomalies but because the tile distribution has multiple, very dissimilar material clusters. The area of sensor artifacts in the lower left appears cyan, meaning both the PDTL and MVN metrics were very high in this area. This is because the noise appears as point anomalies, but there are enough of them in each tile to affect the multivariate normality of the distribution. However, there is not enough diversity in these artifacts to significantly affect the ratio score. In this visualization, there are likely many more false alarms detected than in the 4-class interest segmentation map. This is because the information from all three metrics are displayed. In the feature map clustering process, the information from all metrics are weighted evenly such that a very high score in only one metric may not significantly impact the clustering.

Fig. 15(d) shows a different visualization using NCLUST, NANOM, and PDD in the red, green, and blue channels, respectively. This result is not thresholded. Generally, red tiles have a large number of clusters. This corresponds to the urban areas and the boats off the coast. The coastline and the less dense urban areas appear yellow, meaning the scores were high for both the NCLUST and NANOM scores. These areas likely have the most materials in the scene and score high in NCLUST as a result. These areas are also detected by the NANOM result because these coastline tiles contain sand and vegetation, which

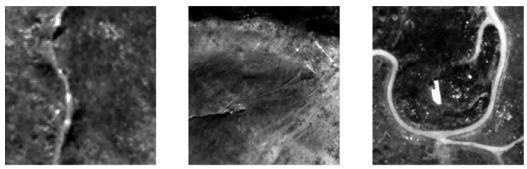

(a) (b) (c)

Fig. 19. High resolution tiles corresponding to "interesting" tiles in the false color visualization map in Fig. 16(c).

are anomalous relative to the primarily urban and water background in the scene. As stated earlier, the PDD is high over areas of a pure material and low over areas with mixed materials due to the density of the data. The water and the two vegetation areas are blue for this reason. The urban areas with more material variability, including vegetation, score high in the PDD.

Fig. 16(c) shows the interest segmentation result using PV, PDTL, and MVN in the red, green, and blue channels, respectively, using the same threshold described above, on the Quickbird data. Most of the tiles labeled interesting in the clustering process are also detected in the false color visualization; however, many more tiles are also detected that are missed in the clustering process. From this visualization, it is clear that the 2-class segmentation map is dominated by the PV results. This area includes the dry wash, the suburban area, and the area in the center with a few buildings. The tiles detected by the PDTL and MVN metrics also contain structures of interest which are missed in the 2-class interest segmentation map. Examples of these tiles, shown in Fig. 19, include cars on the road and buildings, both missed by the PV metric. Similar to the result using the WV2 sample imagery, there are more false alarms detected from these other metrics. These are typically on tiles containing high contrast due to illumination or a mix of exposed soils and burn scar.

V. Conclusion

A method for interest segmentation of spectral imagery for analyst assistance and large area search was introduced. A suite of unique metrics based on data-driven algorithms measures spectral features of the data that indicate the amount of complexity and interest contained in a tile was calculated for two multispectral images and one hyperspectral image. These algorithm results were used to create a multi-band feature image. Simple unsupervised classification techniques are used to divide the multi-band feature image into a few interest segments. In each case, at least 70% of the image was labeled as being in the "no interest" segment. The "interest" segments contain varying degrees of manmade activity. Occasional instances of manmade activity were falsely put in the "not interesting" segment, if only detected by one metric. False color visualization of the estimated feature maps is also used to display a detection plane. This approach can indicate the degree and type of interest contained within a tile based on the brightness and color of the tile in the display. This information can be used to reduce the amount of data an analyst must inspect, help prioritize the task, or cue a sensor to collect additional data. More positive detections are found using the false color visualization technique; however, more false alarms are also detected. Future work involves improving the estimation of spectral features, incorporating additional spectral and spatial features, and optimizing the process for a particular type of image. This involves determining the optimal tile size, number, and combination of features to estimate for a particular image, which can vary with scene content, sensor, and GSD. Application of this methodology to additional hyper- and multispectral datasets which vary in spatial resolution, spectral resolution, land cover, and complexity will also be performed. The interest segmentation map or individual feature maps will also be used as a cueing method for spatially adaptive processing.

References

[1] D. B. Marden and D. Manolakis, "Modeling hyperspectral imaging data using elliptical contoured distributions," *Proc. SPIE*, vol. 5093, pp. 253–262, 2003.

[2] J. R. Schott, *Remote Sensing: The Imaging Chain Approach*, 2nd ed. Oxford, UK: Oxford Univ. Press, 2007, ch. 10.2.

[3] M. O. Smith, S. L. Ustin, J. B. Adams, and A. R. Gillespie, "Vegetation in deserts: A regional measure of abundance from multispectral images," *Remote Sens. Environ.*, vol. 31, pp. 1–26, 1990.

[4] J. C. Harsanyi and C. Chang, "Hyperspectral image classification and dimension reduction: An orthogonal subspace project approach," *IEEE Trans. Geosci. Remote Sens.*, vol. 32, pp. 779–785, 1994.

[5] A. P. Schaum, "Spectral subspace matched filtering," *Proc. SPIE*, vol. 4381, pp. 1–17, 2001.

[6] L. L. Scharf and L. R. McWhorter, "Adaptive matched subspace detectors and adaptive coherence," in *Proc. 20th Asilomar Conf. Signals and Systems*, 1996, pp. 114–117.

[7] D. C. Grimm, "Comparison of hyperspectral imagery target detection algorithm chains," Master thesis, Rochester Inst. Technology, Rochester, NY, 2005.

[8] J. E. West, "Matched filter stochastic background characterization for hyperspectral target detection," Master thesis, Rochester Inst. Technology, Rochester, NY, 2005.

[9] A. Schlamm, D. W. Messinger, and B. Basener, "Geometric estimation of the inherent dimensionality of a single and multi-material clusters hyperspectral imagery," *J. Appl. Remote Sens.*, vol. 3, pp. 033527–033527, Feb. 2009.

[10] R. G. Resmini, "A tool for the nonparametric characterization of the geometry of spectra in hyperspace," *Proc. SPIE*, vol. 7334, p. 73341S, 2009, S. S. Shen and P. E. Lewis, Eds..

[11] B. Basener, D. Messinger, and E. Ientilucci, "Anomaly detection using topology," *Proc. SPIE*, vol. 6565, 2007.

[12] B. Basener, A. Castrodad, D. Messinger, J. Mahle, and P. Prue, "A dynamical systems algorithm for clustering in hyperspectral imagery," *Proc. SPIE*, vol. 6966, 2008.

[13] C. M. Bachmann, T. L. Ainsworth, and R. A. Fusina, "Exploiting manifold geometry in hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 3, pp. 441–454, Mar. 2005.

[14] A. Ziemann, D. Messinger, and W. Basener, "Iterative convex hull volume estimation in hyperspectral imagery for change detection," *Proc. SPIE*, vol. 7695, 2010.

[15] C. M. Stellman, F. M. Olchowski, and J. V. Michalowicz, "War Horse (wide area reconnaissance–Hyperspectral overhead real-time surveillance experiment)," *Automatic Target Recognition, Proc. SPIE*, vol. 4379, pp. 339–346, 2001.

[16] A. Schlamm, D. W. Messinger, and B. Basener, "Geometric estimation of the inherent dimensionality of a single material cluster in multi- and hyperspectral imagery," *Proc. SPIE*, vol. 6966, 2008.

[17] A. Schlamm, D. Messinger, and B. Basener, "Effect of manmade pixels on the inherent dimension of natural material distributions," *Proc. SPIE*, vol. 7334, 2009.

[18] C. J. Mecklin and D. J. Mundfrom, "An appraisal and bibliograph of tests for multivariate normality," *Int. Statist. Rev.*, vol. 72, no. 1, pp. 123–138, Apr. 2004.

[19] A. Schlamm, D. Messinger, and W. Basener, "An empirical estimate of the multivariate normality of spectral image data," *Proc. SPIE*, vol. 8048, 2011.

[20] J. Schott, K. Lee, R. Raqueño, G. Hoffman, and G. Healey, "A subpixel target detection technique based on the invariant approach," in *Proc. AVIRIS Workshop*, 2003, NASA JPL.

[21] P. Bajorski, "Simplex projection methods for selection of endmembers in hyperspectral imagery," in *Proc. 2004 IEEE Int. Geoscience and Remote Sensing Symp.*, Sep. 2004, vol. 5, pp. 3207–3210.
[22] D. Messinger, A. Ziemann, W. Basener, and A. Schlamm, "Spectral image complexity estimated through local convex hull volume," in *2nd Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing*, 2010.
[23] D. Messinger, A. Ziemann, B. Basener, and A. Schlamm, "A metric of spectral image complexity with application to large area search," *Opt. Eng.*, 2011.
[24] A. Kostrikin and Y. Manin, *Linear Algebra and Geometry*. New York: Gordon and Breach Science Publishers, 1997.
[25] B. Basener, D. Messinger, and E. Ientilucci, "Enhanced detection and visualization of anomalies in spectral imagery," *Proc. SPIE*, vol. 7334, 2007.
[26] I. S. Reed and X. Yo, "Comparative adaptive multiple-band CFAR detection of a optical pattern with unknown spectral distribution," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, pp. 1760–1770, 1990.
[27] F. Camastra and A. Vinciarelli, "Estimating the intrinsic dimension of data with a fractal-based method," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, pp. 1404–1407, 2002, DOI:10.1109/TPAMI.2002.1039212.
[28] Digital Globe. 2009 [Online]. Available: http://www.digitalglobe.com/
[29] HyMap. [Online]. Available: http://www.intspec.com/products/hymap/overview/
[30] D. Snyder, J. Kerekes, I. Fairweather, R. Crabtree, J. Shive, and S. Hager, "Development of a web-based application to evaluate target finding algorithms," in *IEEE Int. Geoscience and Remote Sensing Symp., IGARSS 2008*, pp. II-915–II-918, DOI: 10.1109/IGARSS.2008.4779144.

Ariel Schlamm received the Bachelor degree in imaging and photographic technology and the Ph.D. degree in imaging science, both from the Rochester Institute of Technology (RIT), Rochester, NY.

She is currently post doctoral staff in the Digital Imaging and Remote Sensing Laboratory at RIT. Her research focuses on novel mathematical approaches to spectral image processing and analysis.

David Messinger received the Bachelor degree in physics from Clarkson University, Potsdam, NY, and the Ph.D. degree in physics from Rensselaer Polytechnic Institute, Troy, NY.

He has worked as an Analyst for XonTech Inc. and on the National Missile Defense Program for Northrop Grumman. He is currently an Associate Research Professor in the Center for Imaging Science at the Rochester Institute of Technology, where he is the Director of the Digital Imaging and Remote Sensing Laboratory. His research focuses on projects related to spectral image exploitation using physics-based approaches and advanced mathematical techniques.

William Basener received the Ph.D. degree from Boston University, Boston, MA.

He is an Associate Professor in the School of Mathematical Sciences at the Rochester Institute of Technology, co-director of the Center for Applied and Computational Mathematics, and co-director of the Institute for Mathematical Methods in Counterterrorism. His research includes a textbook on Applied Topology.