Pre-Print Manuscript of Article:
Bridgelall, R., Rafert, J. B., Tolliver, D., Lee, E., "Resolution Agile Remote Sensing for Detection of Hazardous Material Spills," *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 2547, DOI: 10.3141/2547-08, 2016.

RESOLUTION AGILE REMOTE SENSING FOR DETECTION OF HAZARDOUS MATERIAL SPILLS

Raj Bridgelall, Corresponding Author
SMARTSe[SM], Upper Great Plains Transportation Institute
North Dakota State University
P.O. Box 863676, Plano, TX 75086
Phone: (408) 607-3214; Fax: (701) 231-1945; Email: raj@bridgelall.com

James B. Rafert
Department of Physics
North Dakota State University
1536 Cole Boulevard, Suite 140, Lakewood, CO 80401
Phone: (720) 238-0080; Fax: (701) 231-1945; Email: bruce.rafert@ndsu.edu

Denver D. Tolliver
Upper Great Plains Transportation Institute
North Dakota State University
Fargo, ND 58108
Phone: (701) 231-7190; Fax: (701) 231-1945; Email: denver.tolliver@ndsu.edu

EunSu Lee
School of Business
New Jersey City University
Jersey City, NJ 07311
Phone: (701) 205-1525; Fax: (201) 200-3242; Email: elee3@njcu.edu

Word count: 5,987 + 3 Tables (250 each) + 2 Figures (250 each) = 7,237 words

TRR Paper number: 16-0110

Submission Date: February 23, 2016

ABSTRACT
Traffic carrying flammable, corrosive, poisonous, and radioactive materials continues to increase in proportion with the growth in their production and consumption. The sustained risk of accidental releases of such hazardous materials poses serious threats to public safety. The early detection of spills will potentially save lives, protect the environment, and thwart the need for expensive clean up campaigns. Ground patrols and terrestrial sensing equipment cannot scale cost-effectively to cover the entire transportation network. Remote sensing with existing airborne and spaceborne platforms has the capacity to monitor vast areas regularly but often lack the spatial resolution necessary for high accuracy detections. The emergence of unmanned aircraft systems with lightweight hyperspectral image sensors enables a resolution agile approach that can adapt both spatial and spectral resolutions in real-time. Equipment operators can exploit such a capability to enhance the resolution of potential target materials detected within a larger field-of-view to verify their identification or to perform further inspections. However, the complexity of algorithms available to classify hyperspectral scenes limits the potential for real-time target detection to support rapid decision-making. This research introduces and benchmarks the performance of a low-complexity method of hyperspectral image classification. The hybrid supervised-unsupervised technique approaches the performance of prevailing methods that are at least 30-fold more computationally complex.

INTRODUCTION

Sustained growth in industrial and commercial activities that rely on hazardous material transport increases the risks of hazardous material spills. Government studies indicate that traffic carrying flammable, corrosive, poisonous, and radioactive materials exceed 800,000 shipments per day (*1*). A case study of crude oil transport demonstrates that spills from pipelines, trains, and commercial motor vehicles pose a serious threat to public safety (*2*). The release of hazardous materials often results in environmental contamination and property damages that cost millions of dollars and years to remediate (*3*). Corrosion is a leading cause of pipeline failures (*4*). Hence, the risk of future accidents increases with pipeline aging. Currently, the age of more than 70% of the crude oil pipelines in the U.S. exceeds 45 years. The lack of pipeline capacity has led to a factor of 24 increase in rail carloads from 2008 to 2012 (*3*). This ramp in rail traffic has raised new concerns about testing and packaging. In the United States, there are more than 630,000 miles of hazardous liquid gathering and transmission pipelines, and more than 140,000 miles of railroad (*4*). Many miles of the pipeline and railroad rights-of-way pass through populated and environmentally sensitive areas.

The lack of regular ground inspections and systems to detect hazardous material releases result in missed detections. A government study found that the public generally reports pipeline leaks quicker than company installed sensors (*5*). Operators cannot easily inspect buried pipes, and rugged terrain often precludes ground inspections (*5*). Remote sensing using spaceborne and airborne platforms offer the potential to monitor large areas for hazardous material spills, quickly and regularly (*6*). Hyperspectral image sensors add a spectral dimension to enhance sensitivity and reduce missed detections (*7*). However, high-altitude platforms limit the spatial resolution needed to preclude false positives. The emergence of small, rotary and fixed-wing, unmanned aircraft systems (sUAS) offer the potential for resolution agile platforms. Their ability to trade-off mobility and accessibility while flying enables the possibility of validating potential targets in real-time. Pilots or autopilot algorithms can adapt sensor parameters and/or maneuver the aircraft to obtain higher resolution images of select target areas. However, real-time adaptation requires a method of rapid hyperspectral image classification to provide operators or algorithms with the data needed to make decisions while navigating.

The computing capacity that existing methods of hyperspectral image classification need to match the typical rate of image acquisition is impractically large. The desired combination of high computational capacities, low power consumption, and low cost is not yet available for integration with sUAS (*6*). Methods of hyperspectral image classification vary in performance and computational complexity as a function of the available spatial and spectral resolutions. An increase in the onboard processing requirement typically leads to a more rapid depletion of the energy available for flight endurance. The limited wireless bandwidth and communications range available also precludes the real-time transmission of hyperspectral data to a remote processor.

The objective of this research is to develop a method of rapid hyperspectral image classification that can detect anomalous materials such as hydrocarbon spills with performance levels that approach the prevailing methods. The organization of this paper is as follows: the next section will provide a background on the existing methods of hyperspectral image classification and their computational complexities. The third section will develop the methods and models of rapid feature extraction and classification. The fourth section will assess the performance of the method using two criteria and provide case studies of each. The final section will summarize and conclude the research.

BACKGROUND

Noise and distortions from the data collection apparatus and the large path lengths through the atmosphere corrupts reflectance values. Furthermore, contamination such as dirt, vegetation, and water distorts the spectral signature of a pure target material such as hydrocarbons. Classification methods, therefore, attempt to assign every hyper-pixel to a class of materials based on some measure of similarity. Current methods of hyperspectral image classification fall into two categories: supervised or unsupervised (*8*). The former require a training set or a library of endmembers to determine each pixel assignment based on similarities of their characteristics. Unsupervised methods preclude training requirements by forming clusters of closely related hyper-pixels.

Methods of unsupervised classification, such as principle component analysis (PCA), independent component analysis (ICA) and singular value decomposition (SVD), identifies at least one orthogonal feature set in the hyperspectral scene. However, they are at least $O(PN^2+N^3)$ computationally complex (*9*). Here, P is the number of hyper-pixels, and N is the number of spectral bands. An important shortcoming is that new orthogonal features generally do not provide a clear physical meaning for data interpretation and decision-making (*10*). Algorithms such as the Iterative Self-Organizing Data Analysis Technique (ISODATA) assign hyper-pixels with similar characteristics into clusters. The iterative procedure is very sensitive to the number and types of initial features selected. Convergence depends on the heuristics of setting a threshold for the number of endmember re-assignments. Such algorithms are $O(PKN^2I)$ complex where K is the number of clusters, and I is the number of iterations (*11*). To minimize their computational complexity, analysts typically incorporate methods of feature selection to identify a minimum number of subset bands that would maintain some measure of sufficiency in class separability. However, the feature selection algorithms themselves typically have $O(PN^K)$ complexity (*12*).

Supervised methods assign each hyper-pixel to one of several user-defined classes based on a measure of ‘similarity’ to members of each class. Spectral similarity measures include both statistical and machine learning methods. The statistical methods most often used are Spectral Angle Mapper (SAM), Minimum Distance Classifier (MDC), Maximum-Likelihood Classifier (MLC), Spectral Information Divergence (SID), and Spectral Correlation Mapper (SCM). The SAM is by far the most popular method (*13*). The computationally complexities of the prevailing supervised methods range from $O(N^2)$ to $O(N^3)$. The literature lacks algorithms that are significantly less computationally complex.

METHOD OF RAPID CLASSIFICATION

The method of rapid classification described in this research is a hybrid supervised-unsupervised technique. The unsupervised aspect is a feature extraction method that operates once on every new hyper-pixel and library endmember. The supervised aspect is a comparison of feature sets that uses either radial cell or rectangular quadrant assignments in a two-dimensional (2-D) feature space. The feature extraction for library endmembers is precomputed. Hence, it could occupy a much smaller amount of digital memory within onboard computers. The reduced computational complexities of one-time feature extraction per new hyper-pixel, and the simpler similarity comparisons with endmembers enable the potential for real-time classification. For equidistant endmembers, the assignment will be the same as that of the nearest hyper-pixel neighbor that does not have an ambiguous assignment to the same set of equidistant endmembers.

Feature Extraction

The typical spectral library contains a list of endmembers represented as atmosphere corrected albedo values for each spectral band available (Figure 1). The albedo is a measure of the portion of incident solar energy reflected from a material. NASA's earth observation satellites regularly measure and report the average albedo of the earth's surface in the visible wavelength ranges. This value has been about 30% (*14*).

The typical ground cover materials of remote areas include various types of vegetation and bare soil (*15*). The overall shape of each endmember will be unique with sufficient spectral resolution and bandwidth. The *selectivity* of the approach, which is an ability to discern among different materials, improves with greater spatial and spectral resolutions. A unique advantage of sUAS is the ability to enhance spatial resolution and signal quality by moving the sensor closer to the target. This resolution agile capability will enhance the quality of images within the *sensitivity* range of the system to detect lower spill concentrations among contaminating materials.

One of the simplest features often extracted is the average albedo. Related efforts used the average albedo to estimate the age of asphalt pavements (*16*), to identify snow-cover, and to track waterways in real-time (*17*). This research modifies the average albedo and defines another statistical feature to form a simple 2-D feature space.

The Wavelength Normalized Average Albedo (AVN)

The average albedo μ_g is

$$\mu_g = \frac{1}{N}\sum_{n=1}^{N} g_n \tag{1}$$

where the albedo within spectral band n is g_n. The wavelength normalized average albedo (AVN) is

$$AVN = \frac{\mu_g}{\lambda_H - \lambda_L} \tag{2}$$

where λ_H and λ_L are the highest and lowest wavelength bands, respectively. The normalization per wavelength band facilitates comparisons between endmembers with different spectral resolutions and bandwidths, potentially from combining different libraries. Hence, normalization accommodates band selection methods that attempt to eliminate wavelength channels that do not appreciably decrease the separability between a subset of endmembers targeted to a specific application. A similarly modified feature is the normalized standard deviation (SDN) defined in the next section. The pair of extracted features {AVN, SDN} provides a low complexity feature space that establishes the separability amongst endmember classes.

The Wavelength Normalized Standard Deviation (SDN)

To diversify the feature space, the SDN characterizes the variability of a spectral signature. The standard deviation σ_g of the reflectance spectra is

$$\sigma_g = \sqrt{\frac{1}{N}\sum_{n=1}^{N}\left(g_n - \mu_g\right)^2} \tag{3}$$

and the SDN is

$$SDN = \frac{\sigma_g}{\lambda_H - \lambda_L}. \quad (4)$$

Division by the spectrum bandwidth normalizes the variations across the available spectral ranges of library endmembers. The physical meaning of the AVN is the reflective nature of a material per unit of wavelength across the spectral region of interest whereas the SDN is a measure of the overall variations in reflectivity of a material per unit of wavelength.

Distance Measure

The simple spectral classifier (SSC) computes a pair of {AVN, SDN} feature for each hyper-pixel of the acquired image frame and compares their distance with each target endmember. For a given spatial and spectral resolution, the level of reflectance noise and material contamination of the target will proportionately increase the feature space distance between the captured signature and the signature of the corresponding library endmember. Therefore, a feature space that exhibits relatively large separation distances among endmember combinations will accommodate a greater level of target noise and contamination while maintaining their association with the endmember class of materials. Subsequently, an ability to compare the average separability of selected endmembers in their respective classifier feature spaces establishes an effective approach to compare their relative potential for generating false positives. This approach is a first step that precludes the enormous expense of conducting extensive field experiments to compare actual false positive rates among different classifiers.

The SSC uses the Euclidian distance proportion of the maximum distance in the normalized feature space as the measure of separability. The Euclidian or radial distance proportion D_η is

$$D_\eta = \frac{1}{D_{\max}} \sqrt{(g_x - h_x)^2 + (g_y - h_y)^2} \quad (5)$$

where $D_{\max}$ is the maximum radial distance for all features in the normalized 2-D space, $\boldsymbol{g}$ and $\boldsymbol{h}$ are vectors of the extracted features for any two materials, x and y are the scaled feature space vector components for the horizontal and vertical axis, respectively. The SDN feature space is scaled to equal the dynamic range of the AVN feature space.

The SSC feature space for 15 typical ground cover materials organizes into six macro-classes of materials (Figure 2). Members of the same class have similar spectral signatures. Hence, the average separability among intra-class signatures is substantially less than the average separability among inter-class members. At small zenith angles, materials of the aquatic class are highly absorptive throughout the spectral region. This characteristic places water and ice at an extreme lower corner of the feature space. Conversely, snow of different consistency is typically highly reflective in the visible region and varies in albedo at longer wavelengths. Those features place it near the top of the feature space. The hydrocarbon endmembers shown in the feature space are relatively small concentrations of Benzene or unleaded gasoline mixed with clay (*15*). The hydrocarbons exhibit a combination of high average reflectivity and medium variability in albedo that places it at the extreme right corner of the SSC feature space.

By inspection, the SSC separates hydrocarbons and snow reasonably well from the other materials so that they appear as outliers in the feature space. Conversely, intra-class materials

such as evergreen trees and green grass exhibit less separability. Hence, applications that need to distinguish among intra-class similar materials will likely require a different type of classifier that features a higher average separability, but likely more computationally complex. This limitation of the SSC points to a trade-off in computational complexity and intra-class separability. Consequently, the rapid classification capability of the SSC will be most suitable for applications that seek to identify anomalies in a scene for further scrutiny. For example, oil spilled from a pipeline onto a vegetation or soil patch will likely exhibit obvious and abrupt changes in the SSC material classes. Subsequently, the appropriate use of sUAS will provide a resolution agile capability for closer inspection and validation.

RESULTS AND DISCUSSIONS

The two key performance measures are the average separability of endmembers in the new feature space and the computational complexity of the new classifier. The separability analysis will use the materials sampled from the ASTER Spectral Library (*15*). A case study of the separability performance will compare the relative distances between endmember combinations in the SSC and the SAM feature spaces. Comparing the actual computational resource needs of several classification methods will require a new benchmark that is appropriate for computer architectures that manufacturers optimize to process images at high speed. A case study of the computational complexity using state-of-the-art mobile image processors will quantify the trade-off in processing needs and image classification speed.

Separability Performance

The average feature-space distance between the same combinations of inter- and intra-class endmember samples provides a means to compare the relative separability of different classifiers. Table 1 shows the separability for materials in the denser cluster near the center of the feature space, as a proportion of the maximum SSC feature space distance (cell entries ‘na’ means not applicable.) This comparison excludes the outlier clusters such as hydrocarbons and snow to remove bias in the separability assessment. This combination also simplifies the table to a more meaningful set of materials for ease of visualization and clarity. Hence, these endmember samples from the large spectral library will serve as the standard to compare the SSC separability with other classifiers.

The average separability for the selected materials is 15.8%. The inter-class separability (highlighted cells) is 22.1% whereas the intra-class separability is 3.3%. Borrowing from the interpretation of chi-squared statistics goodness-of-fit testing that uses a 5% significance threshold, a candidate signature is not likely a member of the tested class if its separability is greater than 5%. Therefore, the greater than 10% separability among different classes qualifies the SSC performance to identify materials that are likely contaminants or anomalies, for example, oil spills that do not naturally occur within vegetation, top soil, or aquatic bodies. However, the less than 5% intra-class separability indicates that this method may not be as suitable for distinguishing among materials with similar signatures.

Case Study of Relative Separability

The SAM is one of the most popular techniques for quantifying the separability of spectra in feature space (*18*). The model represents spectra as a vector in N-dimensional space and computes the “angle” between vectors as the measure of similarity. The SAM maps the separation of two vectors in multidimensional space to an angle α_s in degrees such that

$$\alpha_s = \arccos\left(\frac{\sum_{n=1}^{N} f_n g_n}{\sqrt{\sum_{n=1}^{N} f_n^2 \sum_{n=1}^{N} g_n^2}}\right) \quad (6)$$

where $\boldsymbol{f}$ is the spectrum of a hyper-pixel, $\boldsymbol{g}$ is the reference spectrum, and n is the index of the wavelength band (*19*). It is evident that for identical hyper-pixels where $\boldsymbol{f} = \boldsymbol{g}$ the expression evaluates to zero degrees. This approach requires that the compared spectra have matching spectral bands. Of the material combinations analyzed in this study, only six were comparable using the SAM. It is possible to re-sample spectra to equalize their wavelength bands but resampling introduces errors that distort the results of the feature extraction methods.

The proportional SAM distances for the available combinations are the angle of separation as a percentage of the maximum separation angle in the feature space, which is 90 degrees (Table 2). The average SAM separability improvement over the SSC for this sample of library endmembers is 7.2%. This comparison focused on the materials in the center cluster of the SSC space and excluded the outliers such as hydrocarbons and snow. Therefore, the less than 10% improvement of the SAM over the SSC for materials with signatures that are more similar indicates that both would be similarly effective. Moreover, this result also indicates the potential effectiveness of the SSC classifier in identifying anomalies such as hydrocarbons.

Computational Complexity

This study defines the multiply-accumulate complexity (MACC), denoted $\Pi[D]$ where D is the number of clock cycles that a model requires when implemented on processors capable of single-cycle multiple-accumulate (MAC) operations. The typical digital signal processor (DSP) and some alternative architectures optimized for mobile devices implement a MAC operation within a single instruction cycle. However, they implement divisions using a series of bit shifting and comparison operations that amount to approximately 42 clock cycles for a 32-bit signed division (*20*). The MACC notation is more convenient than the Big-O notation to benchmark the computing time on processors optimized for signal and image processing. As is customary with the Big-O notation, the MACC ignores operations that do not include multiplications, such as additions or comparisons (subtractions). The MACC also excludes divisions and multiplications by integer constants that are powers of two because DSPs can calculate those using single-cycle bit-shifting operations that consume negligible resources. Additionally, the MACC excludes operations that algorithms can pre-compute and store in memory for later use. For instance, algorithms can precompute operations that involve only library endmembers. Furthermore, the MACC excludes computations that operations can store from previous cycles of an iteration.

Spectral Angle Mapper (SAM)

The SAM has a MAC complexity of $3\Pi[N]$ operations plus one square root, one division, and one arccosine operation. The Taylor series expansion for a square root operation provides the baseline for estimating the number of MAC operations where (*21*):

$$\sqrt{1+z} = 1 + \sum_{k=1}^{C} \frac{(-1)^k (2k)!}{4^k (k!)^2 (1-2k)} z^k \quad (7)$$

The selection of C provides the desired precision. The exponential and factorial operations of each iteration can use extra memory to pre-compute and store the results for future iterations. The exponent of the argument z requires $\Pi[C]$ operations by storing the results from previous iterations. Multiplication with the pre-computed constants of each iteration requires one additional MAC. Therefore, the MACC of the square root operation is $2\Pi[C]$.

The Maclaurin series expansion for the arccosine of the argument z is (*21*):

$$\arccos(z) = \frac{\pi}{2} - \sum_{k=1}^{C} \frac{(2k)!}{4^k (k!)^2 (2k+1)} z^{2k+1} \tag{8}$$

In a manner similar to the square root operation, pre-computing the constants will reduce the iterative computational requirements. The exponential operation requires $\Pi[2C + 1]$ and multiplication by the constant in each iteration will require one additional. Hence, the MACC of the arccosine operation is $2\Pi[2C + 1]$. Therefore, the total MAC complexity of the SAM classifier per image frame of P hyper-pixels is

$$\Pi_{SAM} = P \times K \times \{\Pi[3N] + \Pi[8C] + \Pi[44]\}. \tag{9}$$

Bhattacharya Distance (B-Distance)

The Bhattacharya Distance (B-Distance) and the Maximum Likelihood Classifier (MLC) are two other frequently utilized classifiers. The B-distance is an index that is proportional to the amount of overlap between two probability density functions $f(x)$ and $g(x)$ of hyper-pixel vector $\boldsymbol{x}$ (*22*). For Gaussian density functions, the model is

$$B = \frac{1}{4}\ln\left[\frac{1}{4}\left(\frac{\sigma_f^2}{\sigma_g^2} + \frac{\sigma_g^2}{\sigma_f^2} + 2\right)\right] + \frac{1}{4}\left(\frac{(\mu_f - \mu_g)^2}{\sigma_f^2 + \sigma_g^2}\right) \tag{10}$$

where σ_f and σ_g are the standard deviations of the first and second spectra, respectively, and μ_f and μ_g are the means of the first and second spectra, respectively. Each variance requires $\Pi[N + 1]$ operations. The series expansion for a logarithm operation is (*21*):

$$\ln(z) = \sum_{n=0}^{C} \frac{(-1)^k}{k+1} (z-1)^{k+1} \tag{11}$$

By inspection, the logarithm operation requires $2\Pi[C + 1]$ MAC cycles. Therefore, the total MACC of the B-Distance per image frame of P hyper-pixels is

$$\Pi\text{B-dist} = \text{P} \times \text{K} \times \{2\Pi[\text{N} + 1] + 2\Pi[\text{C} + 1] + \Pi[172]\}. \tag{12}$$

Maximum Likelihood Classifier (MLC)

The MLC estimates the probability $g_i(x)$ that an observed hyper-pixel $\boldsymbol{x}$ belongs to a pre-determined spectral class ω_i of the i^{th} class in the scene such that

$$g_i(x) = \ln p(\omega_i) - \frac{1}{2}\ln|\Omega_i| - \frac{1}{2}[x - \mu_i]^T \Omega_i^{-1} [x - \mu_i] \tag{13}$$

where $p(\omega_i)$ is the probability that spectral class ω_i is present in the scene (*14*); $|\Omega_i|$ is the determinant of the covariance matrix for the group of spectra in class ω_i; Ω_i^{-1} is the inverse of the

covariance matrix; μ_i is the mean vector for the group of spectra in class ω_i. MLC implementations may precompute the matrix determinants and inversions for each target class; hence, the MACC does not include them. Furthermore, the first term to determine the probability of class presence in a scene is done only once for each class endmember. This procedure requires calculating at least a variance for each hyper-pixel that will require at least $\Pi[(N + 1)P]$ operations, plus the final natural logarithm. The final set of matrix operations will require at least $\Pi[2N]$ operations for a single endmember representing a class. The division by 2 is a right-shift operation so the MACC does not included it. Therefore, the total MACC of the MLC per image frame of P hyper-pixels is

$$\Pi_{MLC} = P \times K \times \Pi[2N] + \Pi[\mathrm{P}(N + 1)] + 2\Pi[C + 1]. \quad (14)$$

Wavelength Normalized Standard Deviation

Computing SDN requires $\Pi[N] + \Pi[1] + 2\Pi[C]$ MAC cycles. The wavelength ratios are pre-computed. The AVN requires $\Pi[2]$. The SSC operates on each of the P hyper-pixels only once to determine their {AVN, SDN} coordinate. The SSC assigns each coordinate to the class having the minimum Euclidian distance. There are $P \times K$ Euclidian distance calculations that require $2\Pi[C] + \Pi[3]$ MAC cycles. Therefore, the one-time SSC computation per hyper-pixel and the assignment to a class requires $P \times \{\Pi[N] + 2\Pi[C] + \Pi[3]\ \}$ and $P \times K \times \{2\Pi[C] + \Pi[3]\}$ operations, respectively. Therefore, the total MACC of the SSC classifier is

$$\Pi_{SSC} = P \times K \times \{2\Pi[C] + \Pi[3]\} + P \times \{\Pi[N] + 2\Pi[C] + \Pi[3]\}. \quad (15)$$

Assigning SSC features to a rectangular quadrant of the feature space would reduce the complexity further by requiring only $P \times K$ subtraction operations. This would yield a SSC-Rectangular (SSC-R) classifier that has a complexity of

$$\Pi_{SSC\text{-}R} = P \times 1 \times \{\Pi[N] + 2\Pi[C] + \Pi[3]\}. \quad (16)$$

Case Study of Computational Complexity

As of 2015, the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) sensor has been the most popular platform for airborne hyperspectral image acquisition. It provides $N = 224$ spectral channels that range from 360 to 2500 nanometers (*23*). When aboard a Twin-Otter aircraft at an altitude of 4 km, the AVIRIS provides a spatial resolution of 4 meters. Hence, there will be $P =$ 62,500 hyper-pixels per square-kilometer of the scene. Although a typical application will tend to classify materials into dozens of classes, this case study will use the $K = 15$ material types shown for the SSC as prototype endmembers for a class. The highest exponent of the polynomial in the series expansion should be at least $C = 3$ when computing the arccosine, logarithm, and square root functions with at least one significant digit of accuracy (*24*). To summarize, the parameters for the case study are $P = 62{,}500$, $N = 224$, $C = 3$, and $K = 15$.

The processing requirements computed are per square-kilometer of hyperspectral scenes collected with the AVIRIS Twin Otter system (Table 3). For this scenario, the number of classifications per frame is PK, which totals 937,500. The third and fourth columns list the number of MAC operations per classification (Πs/PK) and the total MACs per frame (Total Πs), respectively. It is evident that the SAM requires 30 and 48 times more processing capacity than the SSC and the SSC-R, respectively.

The last column of Table 3 lists the execution time for each method for a processor that can allocate 13 million multiply-accumulate cycles per second (MMACS) of capacity. The latest generation of mobile computers has approximately 400 MMACS of total processing capacity (*25*). Hence, the SSC will consume 3% of that capacity whereas the SAM would require 98% of it to classify scenes at the same rate. The SSC and the SSC-R processing speeds shown will support image acquisition rates greater than 0.5 square-kilometers per second. The AVIRIS Twin-Otter can capture hyperspectral images at a maximum rate of approximately 0.4 square-kilometer per second (*23*). Therefore, UAS platforms capable of similar image acquisition rates can perform real-time classification of hyperspectral images using the SSC and SSC-R and only 3% of the processing capacity of state-of-the-art mobile computing platforms.

Considerations for Transfer to Practice

The first step in transfer to practice requires the regulatory approval for a suitable sUAS platform and its hyperspectral camera payload. Subsequently, practitioners must implement the SSC classifier on the existing single-board computer of the sUAS, or design the algorithms to execute on a separate computing module dedicated to image classification tasks. The pre-computed SSC feature space of the library endmembers must include at least their pair of extracted features {AVN, SDN}, using Equations (2) and (4). The dedicated computing module or memory space of the single on-board computer may store the compressed endmember library of the SSC feature space. The classification tasks must communicate the results of real-time classification to the control module that implements operating and navigational response rules. Alternatively, the classification tasks may transmit classification results for real-time display to the remote pilot who is controlling the sUAS. Therefore, the sUAS may adjust its navigation autonomously or under remote pilot control to approach potential targets for more detailed scrutiny and verification.

SUMMARY AND CONCLUSIONS

Remote sensing using small and agile UAS has the capacity to scan large ground areas rapidly. The ability to adapt flight altitudes and speeds continuously enables real-time searches by trading off area coverage for higher image resolution while flying. Remote pilots or onboard algorithms can navigate to potential targets discovered within large swaths, by swooping down or zooming in to obtain images at higher spatial resolution for further scrutiny. Such a resolution agile system will enable verifications of hazardous material spills in real-time to reduce or eliminate false positives. Hyperspectral imaging adds high spectral resolution to enhance the sensitivity of target detections. However, hyperspectral imaging comes at the price of large data cubes per image frame. The high processing capacity needed to classify hyperspectral images poses significant challenges when using low-power mobile computing platforms. This limitation often precludes their use aboard small and agile unmanned aircraft systems (UAS).

This research developed a method of rapid hyperspectral image classification that enables real-time navigational guidance based on target materials detected in the scene. The approach is a hybrid supervised-unsupervised technique that extracts simple statistical features of the spectra for comparison with target endmembers. The statistical features are a wavelength-normalized average albedo (AVN) and a wavelength normalized standard deviation (SDN). Together, the pair of extracted features establishes a simple two-dimensional (2D) feature space. This enables the simple spectral classifier (SSC) to perform a Euclidian distance or quadrant comparison between each hyper-pixel and target endmember. The simple features facilitate pre-compression

of the spectral library of target materials to result in several-fold reduction of onboard computer memory requirements.

Separability analysis demonstrates that the SSC provides approximately 16% separation among library endmembers that comprises a majority of ground cover materials, including hydrocarbons. Prevailing algorithms such as the spectral angle mapper (SAM) provide a modest improvement in separability of 7.2% for materials in tight clusters of the SSC feature space. However, the SSC is less capable of separating materials with similar signatures.

This research developed a method to benchmark the computational speed of classifiers on computer architectures that manufacturers optimize for mobile image processing. The analysis reveals that the SAM requires at least 30 times more processing capacity than the SSC to perform image classifications at the same rate. The case study used optical specifications for a system that has capabilities similar to the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) aboard a Twin-Otter aircraft. The results indicate that the SSC will require a processing capacity of 13 million multiply-accumulate cycles per second (MMACS) to classify hyperspectral images at a rate that exceeds the image capture capacity of the case study system. This requirement represents only 3% of the processing capacity available from state-of-the-art mobile computing platforms, including smartphones. UAS utilize most of the available computing capacity for navigational controls and communications. The SAM will require 98% of the available processing capacity to provide hyperspectral image classifications at the same rate of the SSC. Hence, the reduced complexity of the SSC will enable longer flight endurance by trading off excess capacity for lower power consumption.

The results of this research motivate the need for future studies to characterize the trade-off in sensitivity and selectivity of schemes for rapid hyperspectral image classification. In particular, the researchers will conduct field studies to establish a performance baseline using the SAM to assess classification accuracy as a function of spatial resolution. The ability to use agile UAS and the steady emergence of higher performance hyperspectral sensors will continue to enhance the sensitivity of detectors. Therefore, the authors will investigate the relationship between the achievable spatial-spectral resolutions and the sensitivity of spill detection in terms of the concentration levels detectable.

ACKNOWLEDGEMENTS

The University Transportation Center, a program of the United States Department of Transportation, sponsored this research through its Mountain Plains Consortium (MPC).

REFERENCES

1. Craft, R. *Crashes Involving Trucks Carrying Hazardous Materials*. Publication FMCSA-RI-04-024. FMCSA, U.S. Department of Transportation, 2004.
2. *DOT Agencies Take Coordinated Actions to Increase the Safe Transportation of Energy Products*. Bulletin DOT 36-15. U.S. Department of Transportation, 2015.
3. *Oil and Gas Transportation: Department of Transportation is Taking Action to Address Rail Safety, But Additional Actions are Needed to Improve Pipeline Safety*. GAO-14-667. Government Accountability Office, 2014.
4. *The State of the National Pipeline Infrastructure*. PHMSA, U.S. Department of Transportation, 2011.
5. Kiefner & Associates, Inc. *Leak Detection Study*. DTPH56-11-D-000001. PHMSA, U.S. Department of Transportation, 2012.

6. Bridgelall, R., J. B. Rafert and D. Tolliver. Hyperspectral Imaging Utility for Transportation Systems. In *Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems*, SPIE, San Diego, 2015.
7. Bridgelall, R. Remote Sensing of Oilfield Logistics with Unmanned Aircraft Systems. In *Sensors Expo and Conference*, Rosemont, IL, 2014.
8. Mather P., and B. Tso. *Classification Methods for Remotely Sensed Data*. CRC Press, Boca Raton, 2003.
9. Du, Q., and J. E. Fowler. Low-complexity Principal Component Analysis for Hyperspectral Image Compression. *International Journal of High Performance Computing Applications*, Vol. 22, No. 4, 2008, pp. 438-448.
10. Prats-Montalbán, J. M., A. D. Juan, and A. Ferrer. Multivariate Image Analysis: A Review with Applications. *Chemometrics and Intelligent Laboratory Systems*, Vol. 107, No. 1, 2011, pp. 1-23.
11. Tarabalka, Y., J. A. 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, 2009, pp. 2973-2987.
12. Bajcsy P., and P. Groves. Methodology for Hyperspectral Band Selection. *Photogrammetric Engineering and Remote Sensing*, Vol. 70, No. 7, 2004, pp. 793-802.
13. Homayouni S., and M. Roux. Hyperspectral Image Analysis for Material Mapping Using Spectral Matching. In *ISPRS Congress Proceedings*, 2004.
14. Richards, J. *Remote Sensing Digital Image Analysis*. Springer-Verlag, Berlin, 1999.
15. Baldridge, A. M., S. Hook, C. Grove, and G. Rivera. The ASTER Spectral Library Version 2.0. *Remote Sensing of Environment*, Vol. 113, 2009, pp. 711-715.
16. Herold, M., D. Roberts, V. Noronha, and O. Smadi. Imaging Spectrometry and Asphalt Road Surveys. *Transportation Research Part C*, Vol. 16, 2008, pp. 153-166.
17. Clark, R., G. Swayze, I. Leifer, K. Livo, R. Kokaly, T. Hoefen, S. Lundeen, M. Eastwood, R. Green, N. Pearson, C. Sarture, I. McCubbin, D. Roberts, E. Bradley, D. Steele, T. Ryan, and R. Dominguez. *A Method for Quantitative Mapping of Thick Oil Spills Using Imaging Spectroscopy*. U.S. Geological Survey, Reston, 2010.
18. Meer, v. d. F. D, H. M. A. v. d. Werff, F. J. A. v. Ruitenbeek, C. A. Hecker, W. H. Bakker, M. F. Noomen, M. v. d. Meijde, J. M. Carranza, J. Smeth, and T. Woldai. Multi-and Hyperspectral Geologic Remote Sensing: A Review. *International Journal of Applied Earth Observation and Geoinformation*, Vol. 14, No. 1, 2012, pp. 112-128.
19. Sohn, Y., and N. Rebello. Supervised and Unsupervised Spectral Angle Classifiers. *Photogrammetric Engineering & Remote Sensing*, Vol. 68, No. 12, 2002, pp. 1271-1280.
20. Chen, Y. T., *TMS320C6000 Integer Division*. Texas Instruments, Inc., Richardson, 2000.
21. Thomas, G. B, and R. L. Finney. *Calculus and Analytic Geometry*, 9th ed. Addison Wesley, Boston, 1995.
22. Landgrebe, D. Hyperspectral Image Data Analysis. *Signal Processing Magazine*, Vol. 19, No. 1, 2002, pp. 17-28.
23. Coulter, D., P. L. Hauff, and W. L. Kerby. Airborne Hyperspectral Remote Sensing. In *Proceedings of the 5th Decennial International Conference on Mineral Exploration*, Toronto, 2007.
24. Muller, J. M. *Elementary Functions: Algorithms and Implementation*. Springer Science & Business Media, Heidelberg, 2006.

25. Cole, B. STMicro, ARM Do a Double Whammy with New Cortex-M7 Core. *Embedded.com*, September 2014. www.embedded.com/electronics-news. Accessed June 8, 2015.

LIST OF TABLES

LIST OF FIGURES

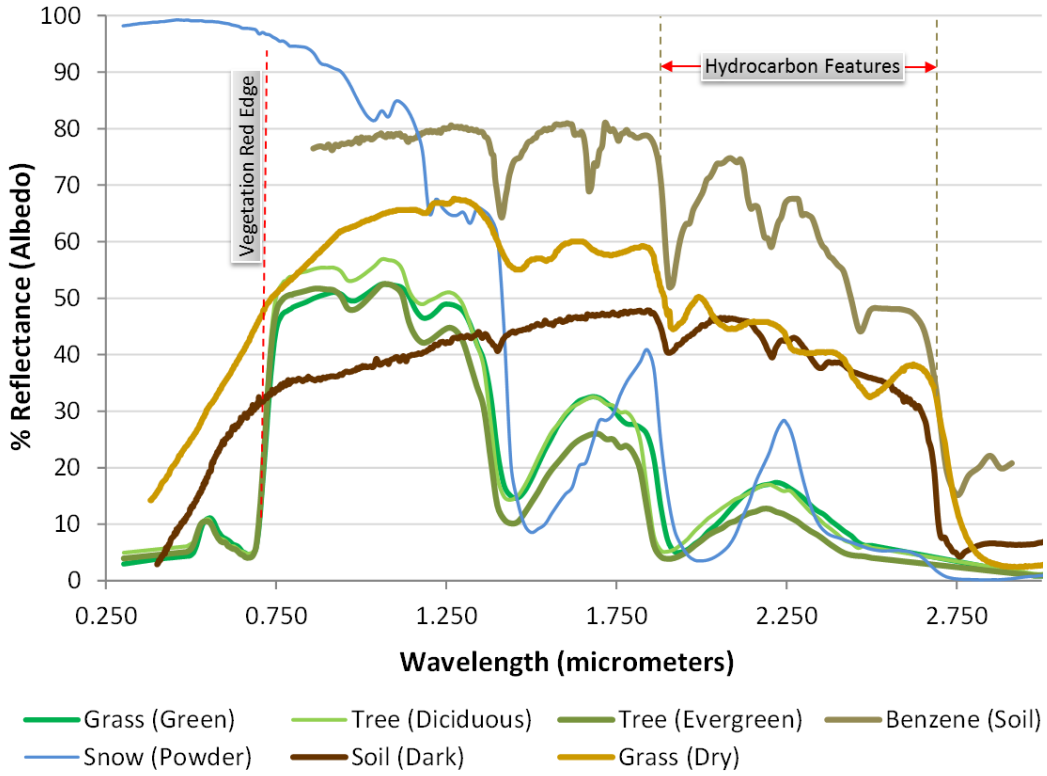

FIGURE 1 Spectral signatures for typical ground cover.

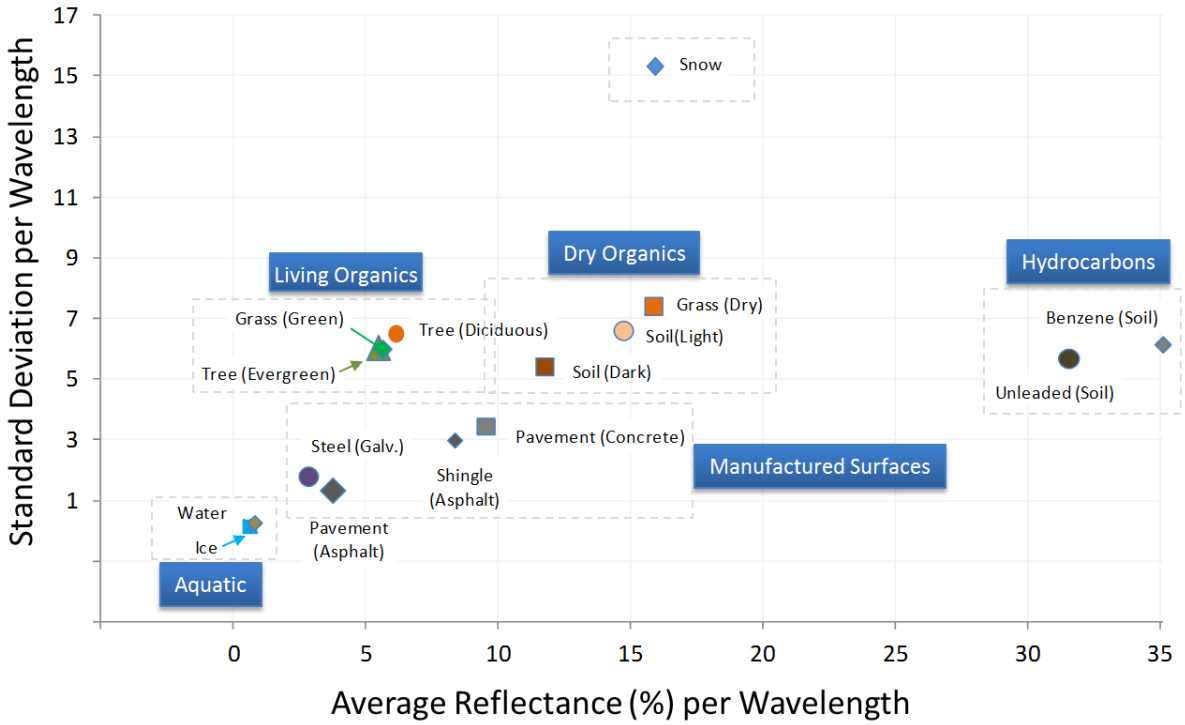

FIGURE 2 SSC separability for typical ground cover.

TABLE 1 SSC Separability Matrix for Typical Ground Cover

	Soil (Dark)	Tree (Con)	Tree (Dec)	Concrete	Ice
Soil (Light)	8.0%	18.9%	17.3%	17.4%	40.3%
Grass (Green)	na	0.4%	2.4%	13.7%	27.8%
Tree (Conifer)	na	na	2.4%	14.1%	27.8%
Shingle (Asphalt)	na	na	na	3.2%	20.0%
Pavement (Concrete)	na	na	na	na	23.2%

TABLE 2 Class Separability for SAM and WSC

Class Separability	SAM	SSC	Difference
Soil (Light) – Soil (Dark)	17.2%	8.0%	9.2%
Grass (Green) – Tree (Deciduous)	8.0%	0.4%	7.6%
Tree (Evergreen) – Tree (Deciduous)	4.1%	2.4%	1.7%
Shingle (Asphalt) – Concrete	16.2%	3.2%	13.0%
Shingle (Asphalt) – Ice	25.8%	20.0%	5.8%
Concrete – Ice	29.1%	23.2%	5.9%
AVERAGE	16.7%	9.5%	7.2%

TABLE 3 Relative Complexities of the Classifiers

Model	Computational Cost Model	Πs/PK	Total Πs	Time (s)
SAM	P×K×{Π[3N]+Π[8C]+Π[44]}	740	694M	53.4
B-Distance	P×K×{Π[2(N+1)]+Π[2(C+1)]+Π[172]}	630	591M	45.4
MLC	P×K×{Π[2N]}+Π[P(N+1)]+Π[2(C+1)]	463	434M	33.4
SSC	P×K×{Π[2C]+Π[3]}+P×{Π[N]+Π[2C]+Π[3]}	25	23M	1.8
SSC-R	P×1×{Π[N]+Π[2C]+Π[3]}	16	15M	1.1