



openaccess

Rapid hyperspectral image classification to enable autonomous search systems

Raj Bridgelall,^{*} J. Bruce Rafert,^a Denver D. Tolliver^b and EunSu Lee^c

Upper Great Plains Transportation Institute, North Dakota State University, Fargo, USA

^{*}PO Box 863676, Plano, TX 75086, USA. E-mail: raj@bridgelall.com

^a1536 Cole Boulevard, Suite 140, Lakewood, CO 80401, USA. E-mail: bruce.rafert@ndsu.edu

^bPO Box 6050, Fargo, ND 58108, USA. E-mail: denver.tolliver@ndsu.edu

^cManagement Department, School of Business, New Jersey City University, NJ 07311, USA, USA. E-mail: elee3@njcu.edu

The emergence of lightweight full-frame hyperspectral cameras is destined to enable autonomous search vehicles in the air, on the ground and in water. Self-contained and long-endurance systems will yield important new applications, for example, in emergency response and the timely identification of environmental hazards. One missing capability is rapid classification of hyperspectral scenes so that search vehicles can immediately take actions to verify potential targets. Onsite verifications minimise false positives and preclude the expense of repeat missions. Verifications will require enhanced image quality, which is achievable by either moving closer to the potential target or by adjusting the optical system. Such a solution, however, is currently impractical for small mobile platforms with finite energy sources. Rapid classifications with current methods demand large computing capacity that will quickly deplete the on-board battery or fuel. To develop the missing capability, the authors propose a low-complexity hyperspectral image classifier that approaches the performance of prevalent classifiers. This research determines that the new method will require at least 19-fold less computing capacity than the prevalent classifier. To assess relative performances, the authors developed a benchmark that compares a statistic of library endmember separability in their respective feature spaces.

Keywords: acquisition systems, autonomous vehicles, classification accuracy, endmember separability, hazardous material detection, real-time spectrometer, resolution agility, simple spectral classifier, unmanned aircraft systems, video imaging

Introduction

The recent emergence of golf-ball size hyperspectral cameras that weigh less than a tennis ball will create new capabilities that shift the paradigm in remote sensing. These full-frame cameras¹ also have relatively low power consumption. Hence, they are suitable for integration into small mobile platforms that could navigate autonomously on the surface, in the air or in water.² For a given spatial resolution, hyperspectral imaging offers the benefit of high spectral resolution to minimise missed detections and maximise the accuracy of target identifications.³ The possibility of adjusting spatial resolution to verify target detections while in operation will preclude the

expense of repeat missions to recheck potential targets after offline image processing. The ability to achieve resolution agility will also compensate for trading off some theoretical accuracy for lower classification complexity. Organisations can use such autonomous search systems for many new applications. A few of the most notable are emergency response, post-disaster damage assessment, environmental hazard (e.g. oil spill or chemical release) detection and the performance evaluation of transportation systems.⁴

To conduct verifications *in situ*, the system must quickly classify large-scale hyperspectral scenes to identify poten-



tial targets, and then be capable of enhancing the spatial resolution of selected areas by moving closer or adapting the optical system. The achievable rate of hyperspectral image classification will dictate the maximum search speed and coverage. Adding computing capacity could accelerate the execution of existing classifiers, but doing so would deplete on-board energy supplies more quickly, reducing ground coverage. Hence, the main idea of this paper is to develop a rapid hyperspectral classifier to enable small, agile and adaptive autonomous search systems. Subsequently, the following are the main objectives of this research:

- 1) to develop a simple hyperspectral image classifier that requires substantially less computing capacity than existing high-performance classifiers.
- 2) to identify a method of benchmarking the theoretical accuracy of the simple classifier relative to the prevalent classifiers that avoids the common and costly approach of immediately conducting expensive ground truth experiments.
- 3) to identify a method of estimating and comparing the absolute and relative computational complexities of the proposed empirical and the prevalent classifier.

To set the scene for achieving these objectives, the next section will review the scope of existing methods to classify hyperspectral images and to characterise their overall and relative computational complexities. The third section will achieve the first objective by developing the new classifier using an empirical formulation that hinges on the energy of wavelength rate changes across the spectrum. The fourth section will achieve the second objective by utilising separability analysis. This method uses typical ground samples from existing spectral libraries to compare the theoretical false positive potential of the new classifier with that of the prevailing classifier that has known performance. The fifth section will achieve the third objective listed above by introducing a new method of benchmarking computational complexity that is most appropriate for digital image processing architectures. The sixth section will use the benchmarks to compare the theoretical performance of the proposed and the prevalent classifier, within the context of a case study. The final section will summarise and conclude the research, demonstrating that the theoretical approach has successfully achieved the performance benchmarking objectives with existing endmember libraries as a first vetting step to avoid the common approach of incurring upfront expenses to conduct extensive field data collection to evaluate a new classifier.

Review of existing methods

Existing hyperspectral libraries⁵ contain the spectral signature of target materials obtained from ground truth data. These signatures are so-called library endmembers because of their careful measurements and spectral purity. Noise and distortions from the data collection equipment and the long path lengths from remote sensing distorts the spectral signature

of captured images. Therefore, the main purpose of a hyperspectral image classifier is to associate each noisy hyper-pixel of the image to an endmember (supervised methods) or into clusters (unsupervised methods), based on some measure of similarity.

The Euclidean n -space distance is one of the most popular measures of similarity.⁶ Therefore, the average separation distance of endmembers (separability) in a given feature space is a predictor of the potential for a candidate classifier to produce false positives at some level of noise. Hence, comparing the proportional separability of a fixed sample of library endmembers within the feature space of the candidate classifier and the feature space of a classifier with known performance offers a first step in vetting their potential performance with actual field data. The common approach to evaluate a new classifier is to collect field data immediately to assess its accuracy. However, the alternative of benchmarking the theoretical performance of the new classifier against the performance of a prevalent classifier with known field accuracy characteristics will minimise the risks of incurring the high cost of data collection to vet a potentially useless classifier. This approach is analogous to the classic methods of evaluating the theoretical bit-error-rate (BER) performance of information encoding schemes as a function of signal-to-noise ratio (SNR). The separability analysis of this research uses the same finite sample of library endmembers that typifies a scene containing the target material, for example spilled crude oil, and the contaminated materials such as soil, water, snow and vegetation.

Methods of image classification vary in performance and computational complexity as a function of both the number of hyper-pixels P and the number of spectral bands N . Methods of supervised classification use statistical and machine learning techniques to establish their measures of similarity. The most popular methods 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 the prevalent method. The computational complexities of present supervised methods range from $O(N^2)$ to $O(N^3)$. Methods of unsupervised classification use techniques such as principle component analysis (PCA), independent component analysis (ICA) and singular value decomposition (SVD); they are at least $O(PN^2 + N^3)$ computationally complex.⁷ Algorithms such as the iterative self-organising data analysis technique (ISODATA) assign hyper-pixels with similar characteristics into clusters. 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.⁸ To minimise 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.⁹

The proposed rapid classifier

We propose a new method of rapid classification based on empirical influences from both supervised and unsupervised techniques. The unsupervised aspect is a feature extraction that operates once per new hyper-pixel acquired and once per library endmember. The supervised influence is a radial or a rectangular distance threshold comparison in a two-dimensional (2D) feature space. The method precomputes the features for each available library endmember for comparison with features of captured hyper-pixels. Hence, the extracted endmember features could occupy a much smaller amount of digital memory than the entire library of endmember signatures. The reduced computational complexities of one-time feature extraction per new hyper-pixel, and the simpler similarity comparisons with endmember features enable the potential for real-time classification.

Empirical feature extraction

The typical spectral library contains a list of endmembers represented as albedo values for each of the available spectral bands. The albedo is a measure of the portion of incident solar energy reflected from a material. This simple statistic is still powerful, relevant and very important. NASA's earth observation satellites regularly measure and report the average albedo of the earth's surface in the visible wavelength ranges. Figure 1 plots the albedo as a function of the spectral band for typical ground cover materials.⁵

We modify the albedo to improve its effectiveness when using signatures of different signature lengths and spectral resolutions from the same or multiple libraries. We also designed a second feature based on a heuristic that summarises the waviness of the signature. Together, these two features form a 2D feature space.

The wavelength normalised average albedo

The average albedo μ_g of a spectral signature g is

$$\mu_g = \frac{1}{N} \sum_{n=1}^N g_n \quad (1)$$

where the albedo in spectral band n is g_n . The wavelength normalised average albedo (AVN) is

$$AVN = \frac{\mu_g}{\lambda_H - \lambda_L} \quad (2)$$

where λ_H and λ_L are the highest and lowest wavelength bands, respectively. The normalisation per wavelength band facilitates comparisons between endmembers with different spectral resolutions and bandwidths, potentially from combining different libraries. Hence, normalisation accommodates band selection methods that attempt to eliminate wavelength channels that do not appreciably decrease the separability between selected endmembers.

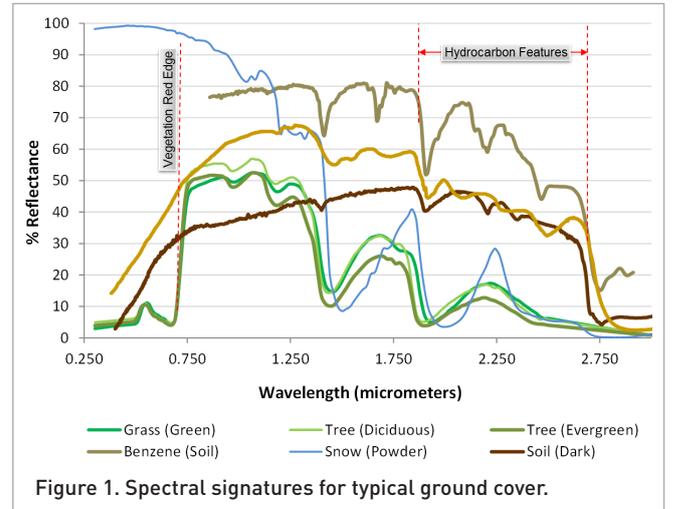


Figure 1. Spectral signatures for typical ground cover.

The wavelength sensitivity index

We call the new feature the wavelength sensitivity index (WSI) because it characterises the *shape* or waviness of a spectral signature. We define the corresponding wavelength sensitivity transform as

$$W_s = \sqrt{\frac{1}{\lambda_H - \lambda_L} \sum_{n=1}^N \left(\frac{g_n - g_{n-1}}{\lambda_n - \lambda_{n-1}} \lambda_n \right)^2} \quad (3)$$

where W_s is the WSI, and λ_n is the centre of each available wavelength band in units of micrometres. The origin of the WSI is purely empirical. Its formulation is fundamentally a measure of the band-normalised energy of the wavelength slope signature. Conceivably, this definition could include higher order derivatives instead of or in addition to the wavelength slope, but at the expense of increasing computational complexity. The weight λ_n of the wavelength slope maximises the separability of materials that might have similar wavelength slope energies in different portions of the spectrum. The fact that the weight tends to emphasise the wavelength slopes at the higher end of the spectrum is inconsequential because of the feature space normalisation. The authors have previously reported⁴ on other measures of waviness such as the normalised standard deviation, which is less effective and entropy, which is substantially more computationally complex.

A potential limitation of the WSI method could be the reduced separability of materials that have slope signatures in different portions of the spectrum and with just the right magnitudes to equalise their single-dimensional WSI feature. However, for such a hypothetical case, it is also possible that the {AVN, WSI} feature pair will compensate to increase the two-dimensional separability. Therefore, without an exhaustive separability analysis that involves all materials known to man, the authors recommend using this method to test application specific targets, for example spilled crude oil among common contaminated materials such as soil, snow, water and vegetation.

Distance measure

The wavelength sensitivity classifier (WSC) computes {AVN, WSI} pairs for each hyper-pixel of the acquired image frame and compares their proximity to target endmembers. Although other distance measures are possible, we elected to use the Euclidian distance because of its simplicity. The Euclidian or radial distance D_E is

$$D_E = \sqrt{(g_x - h_x)^2 + (g_y - h_y)^2} \quad (4)$$

where g and h are the extracted feature sets; the x and y components are the {AVN, WSI} features for any two materials.

The endmember samples for separability analysis are 15 typical ground cover materials from the ASTER Spectral Library.⁵ Figure 2 shows the result of applying the WSC to the endmember sample set. The WSC feature space assigns the normalised AVN and WSI to the horizontal and vertical axis, respectively. 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 feature combinations place it near the centre of the feature space. Materials of the hydrocarbon class exhibit a combination of high average reflectivity and high wavelength sensitivity that places it at the extreme upper-right corner of the feature space.

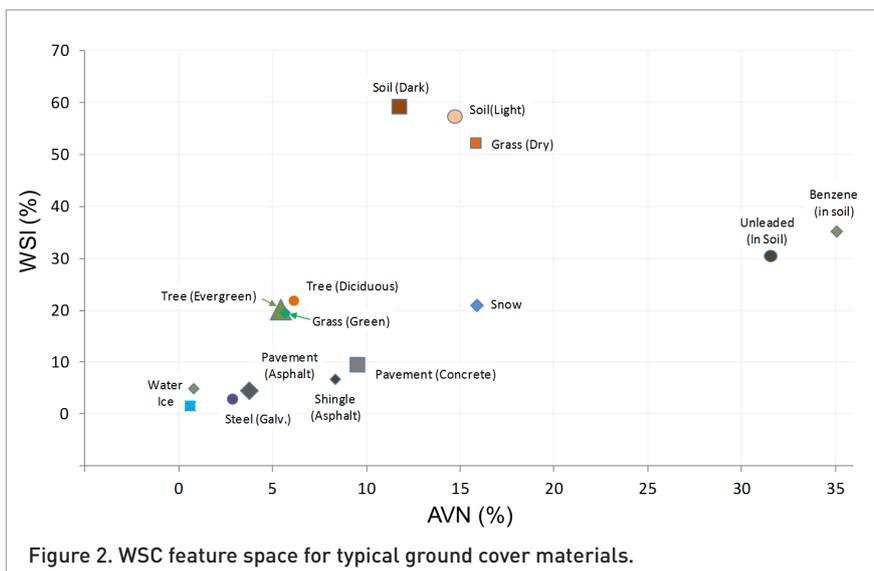
By inspection, the WSC appears to separate hydrocarbon target materials from soil and snow reasonably well. Conversely, materials within the same macro-class, such as evergreen trees and green grass, exhibit less separability. Hence, applications that need to distinguish among similar materials should use a different type of classifier that is likely more computationally or dimensionally complex. This limitation of the WSC points to a trade-off in computational complexity and intra-class separability. This scenario analysis indicates that the rapid classification capability of the WSC

will be best suited for custom applications that search for high contrast targets within the scene. Hence, in addition to oil spills (on snow, water, vegetation or soil), the WSC would be appropriate for tracking vehicles on paved or unpaved roads, and for tracking vessels on water. The primary strength of the WSC is that it allows for immediate repositioning of robotic vehicles to obtain higher spatial or directional resolution for target verification. This resolution agility will likely compensate for any loss of classification accuracy relative to the more computationally complex approaches such as SAM or SVD.

Proposed performance benchmark

This research defines the *separability* of a classifier as the average of the normalised separation distance in its feature space, for all target and contaminated material combinations of the selected sample of library endmembers. Normalising the feature space distance provides a fair means of benchmarking the potential accuracy of classifications relative to the anticipated distance errors in hyper-pixel assignments. Figure 3 graphically illustrates the concept. It shows a hypothetical distribution of the normalised distances between all combination of endmembers in the respective feature spaces of two different classifiers, A and B.

The interval of uncertainty represents the expected normalised distance deviations of candidate hyper-pixel distances from their endmembers, for example with less than 5% significance. This normalised deviation threshold is analogous to an acceptable noise interval for hyper-pixel assignments. For the same set of library endmembers, the mean is closer to the noise interval for Classifier A than for Classifier B. Therefore, the former has a greater potential for misclassifications than the latter. Alternatively, Classifier B could tolerate a higher noise level than Classifier A could. Consequently, Classifier B



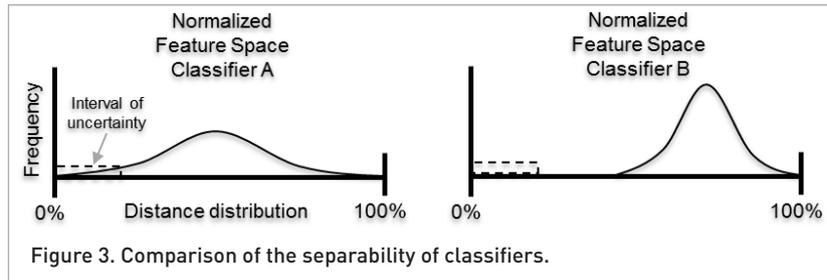


Figure 3. Comparison of the separability of classifiers.

has the potential to classify more of the hyperspectral scene than Classifier A could. In either case, a larger ratio of average normalised separation distance to an arbitrary interval of uncertainty is desirable.

As mentioned earlier, it is important to consider this benchmarking approach in analogy to the classic approach of comparing the theoretical SNR requirements of different data encoding methods to achieve some desired level of decoding accuracy. That is, this theoretical performance bound is a first indicator of the potential performance of the candidate classifier, and it does not replace the eventual need to conduct extensive field studies with ground truth data. The primary benefit of this first-step vetting is to benchmark the theoretical performance of new methods relative to those of existing methods of known performance levels to assess the value of later conducting expensive field data collection for final performance characterisations.

Proposed measure of computational complexity

This section details new benchmark for computational complexity that is appropriate for computer architectures that manufacturers optimise to process images at high speed.

The multiply-accumulate complexity

We define a unit of computational capacity $\Pi[D]$ as the multiply-accumulate complexity (MACC), 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 optimised 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.¹⁰ The MACC notation is more convenient than the Big-O notation to benchmark the computing time on processors optimised 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 pre-compute operations that involve only library endmembers. Furthermore, the MACC excludes computations that operations can store from previous cycles of an iteration, for example, when computing a series expansion.

The SAM complexity

The SAM is the most popular classifier. It represents spectra as a vector in N -dimensional space and computes the “angle” between vectors as the measure of similarity.¹¹ The SAM maps the separation of two vectors in multidimensional space to an angle α_s in degrees such that

$$\alpha_s = \arccos \left(\frac{\sum_{n=1}^N \mathbf{f}_n \mathbf{g}_n}{\sqrt{\sum_{n=1}^N \mathbf{f}_n^2 \sum_{n=1}^N \mathbf{g}_n^2}} \right) \quad (5)$$

where \mathbf{f} is the spectrum of a hyper-pixel, \mathbf{g} is the reference or endmember spectrum and n is the index of the wavelength band.

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 operation¹² provides the baseline for estimating the number of MAC operations where:

$$\sqrt{1+z} = 1 + \sum_{k=1}^C \frac{(-1)^k (2k)!}{4^k (k!)^2 (1-2k)} z^k \quad (6)$$

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 intermediate results for future iterations. For instance, the exponent of the argument z requires $\Pi[C]$ operations. 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 arccosine¹² is

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

In a manner that is similar to evaluating 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 MAC. Hence, the MACC of the arccosine operation is $2\Pi[2C+1]$. Subsequently, the total MAC complexity of the SAM classifier per image frame of P hyper-pixels is

$$\Pi_{\text{SAM}} = P \times K \times \{\Pi[3N] + \Pi[8C] + \Pi[44]\}. \quad (8)$$

Using the same approach, the complexity of the Bhattacharya distance [B-distance] is

$$\Pi_{\text{B-dist}} = P \times K \times \{2\Pi[N+1] + 2\Pi[C+1] + \Pi[172]\} \quad (9)$$

and the complexity of the MLC is:

$$\Pi_{\text{MLC}} = P \times K \times \{\Pi[2N] + \Pi[P(N+1)] + 2\Pi[C+1]\}. \quad (10)$$

Wavelength sensitivity index complexity

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

$$\Pi_{\text{WSC}} = P \times K \times \{2\Pi[C] + \Pi[3]\} + P \times \{\Pi[2N] + 2\Pi[C] + \Pi[3]\}. \quad (11)$$

Assigning WSC 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 WSC-rectangular (WSC-R) classifier that has a complexity of

$$\Pi_{\text{WSC-R}} = P \times 1 \times \{\Pi[2N] + 2\Pi[C] + \Pi[3]\}. \quad (12)$$

Analytical results and discussions

This section will quantify the two key performance measures of endmember separability and the computational complexity.

Separability analysis

Table 1 summarises the normalised separation distances for materials in the denser cluster near the centre of the feature WSN space. This comparison excludes the outlier clusters such as hydrocarbons and snow to remove comparison bias. The selected combinations also simplify the table to a more meaningful set of materials for ease of visualisation and clarity. Hence, these endmember samples from the large spectral library will serve as the standard for comparing the separability of candidate and prevailing classifiers for a specific application. The average separability for the selected materials is 23.7%. The average inter-class separability (emphasised in bold font) is 33.9% whereas the intra-class separability is 3.4%. 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%. Hence, these results indicate that the WSC will be effective in identifying specific contaminants such as oil spills or non-native materials that disrupts the homogeneity of a hyperspectral scene.

Case study of relative separability

The SAM requires that the compared spectra have the same bands and bandwidths. Of the available material combinations in the endmember sample set, only six were comparable. It is possible to re-sample spectra to equalise their wavelength bands but resampling introduces errors that distort the results of the feature extraction methods. Table 2 compares the separability of the SAM combinations available from the sample set.

The SAM separability advantage over the WSC is 8.2%. The average inter-class and intra-class improvements are 8.9% and 7.9%, respectively. The relatively small improvement of the SAM over the WSC indicates that the latter has the potential to approach the performance levels of prevalent classifiers for a small improvement in image quality that would reduce the interval of uncertainty.

Case study of computational complexity

The case studies will use parameters for an existing airborne remote sensing platform and a state-of-the-art processor to benchmark the computational requirements. At the time of this publication, the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) sensor is still the most popular platform for airborne hyperspectral image acquisition.¹³ It provides $N=224$ spectral channels that range from 360 nm to

Table 1. WSC separability matrix for typical ground cover.

	Soil (dark)	Tree (con)	Tree (dec)	Concrete	Ice
Soil (light)	6.4%	47.6%	45.2%	57.3%	71.6%
Grass (green)		0.9%	3.0%	14.1%	23.4%
Tree (conifer)			2.4%	16.1%	26.4%
Shingle (asphalt)				4.0%	16.8%
Pavement (concrete)					20.3%

Table 2. Separability comparison of the SAM and the WSC.

Class separability	SAM	WSC	Δ
Soil (light)–Soil (dark)	17.2%	6.4%	10.7%
Grass (green)–Tree (deciduous)	8.0%	0.9%	7.1%
Tree (evergreen)–Tree (deciduous)	4.1%	2.4%	1.7%
Shingle (asphalt)–Concrete	16.2%	4.0%	12.2%
Shingle (asphalt)–Ice	25.8%	16.8%	9.0%
Concrete–Ice	29.1%	20.3%	8.8%
Average	16.7%	8.5%	8.2%

2500 nm. When the host aircraft is a Twin-Otter flying at an altitude of 4 km, the AVIRIS provides a spatial resolution of 4 m. Hence, there will be $P=62,500$ hyper-pixels per square-kilometre 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 WSC 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.¹⁴ To summarise, the parameters for the case study are $P=62,500$, $N=224$, $C=3$ and $K=15$.

Table 3 lists the processing requirements per square-kilometre of hyperspectral scenes collected with the AVIRIS Twin-Otter system. For this scenario, the number of classifications per frame is $P \times K$, which totals 937,500. The third and fourth columns list the number of MAC operations per classification (IIs/ PK) and the total MACs per frame (Total IIs), respectively. It is evident that the SAM requires 19 and 24 times more processing capacity than the WSC and the WSC-R, respectively.

The last column of Table 3 lists the execution time for each method when using a processor that can allocate 20 million multiply-accumulate cycles per second (MMACS) of capacity. The latest generation of mobile computers has approximately 400 MMACS of total processing capacity.¹⁵ Hence, the WSC will consume 5% of that capacity whereas the SAM would require 94% of it to classify scenes at the same rate. The WSC and the WSC-R processing speeds shown will support image acquisition rates greater than 0.5 square-kilometres per second. The AVIRIS Twin-Otter can capture hyperspectral images at a maximum rate of approximately 0.4 square-kilometres per second.¹³ This result indicates that a hypothetical unmanned

aircraft system (UAS) platform with a similar capture rate can classify hyperspectral scenes in real-time by using the WSC and WSC-R.

Summary and conclusions

The search for dynamic targets with remote sensing platforms demands a rapid detection ability so that the system can take action to enhance the spatial resolution of the target area for immediate verification. Small unmanned and autonomous vehicles are emerging, and so are tiny hyperspectral imagers that are suitable payloads. However, the missing capability is rapid hyperspectral image classification. The WSC is a low-complexity method of hyperspectral image classification that would enable small autonomous vehicles to perform rapid searches.

The new technique extracts simple statistical and shape features of the spectra for comparison with target endmembers. The features are the wavelength normalised average albedo (AVN) and the WSI. Together, these features establish the simple two-dimensional (2D) feature space of the WSC. We further define two new measures of performance. They are the separability of the feature space and the MACC of the classifier. The former is analogous to comparing their relative SNR requirements for a given level of classification accuracy desired.

The separability analysis demonstrates that the WSC provides approximately 24% separation among library endmembers that comprise a majority of typical ground cover materials. Prevailing algorithms such as the SAM provide a modest improvement in separability of 8.2% for those mate-

Table 3. Relative complexities of the classifiers.

Model	Computational cost model	IIs/ PK	Total IIs	Time (s)
SAM	$P \times K \times \{\text{II}[3N] + \text{II}[8C] + \text{II}[44]\}$	740	694M	34.7
B-distance	$P \times K \times \{\text{II}[2(N+1)] + \text{II}[2(C+1)] + \text{II}[172]\}$	630	591M	29.5
MLC	$P \times K \times \{\text{II}[2N] + \text{II}[P(N+1)] + \text{II}[2(C+1)]\}$	463	434M	21.7
WSC	$P \times K \times \{\text{II}[2C] + \text{II}[3]\} + P \times \{\text{II}[2N] + \text{II}[2C] + \text{II}[3]\}$	39	37M	1.9
WSC-R	$P \times 1 \times \{\text{II}[2N] + \text{II}[2C] + \text{II}[3]\}$	30	29M	1.4

rials that form tighter clusters in the WSC feature space. The complexity benchmark revealed that the SAM requires at least 19 times more processing capacity than the WSC to perform image classifications at the same rate.

The case study used optical specifications for a system that has capabilities that are similar to the AVIRIS aboard a Twin-Otter aircraft. The results indicate that the WSC will require a processing capacity of 20 MMACS to classify hyperspectral images at a rate that exceeds the image capture capacity of the AVIRIS platform. This requirement represents only 5% of the processing capacity available from state-of-the-art mobile computing platforms, including smartphones. Mobile sensing platforms utilise most of the available computing capacity for navigational controls, communications and sensor operations. The SAM will require 94% of the available processing capacity to provide hyperspectral image classifications at the same rate of the WSC. Hence, the reduced complexity of the WSC will enable longer flight endurance by trading off excess computational capacity for lower power consumption. Hyperspectral data of the transportation infrastructure is not currently available in the public domain. Hence, the authors will use the results of this research to enable an extensive evaluation of the classification accuracy of the WSC relative to the SAM by acquiring field data with a recently purchased hyperspectral camera and UAS.

Acknowledgement

The University Transportation Centre, a program of the United States Department of Transportation, sponsored this research through its Mountain Plains Consortium (MPC). A grant from the United States Department of Transportation supported this research under Grant DTRT13-G-UTC38.

References

1. A. Lambrechts, P. Gonzalez, B. Geelen, P. Soussan, K. Tack and M. Jayapala, "A CMOS-compatible, integrated approach to hyper- and multispectral imaging", *Electron Devices Meeting (IEDM), 2014 IEEE International*. Institute of Electrical and Electronic Engineers (IEEE), Piscataway, New Jersey (2014).
2. Volpe Center, *Unmanned Aircraft System (UAS) Service Demand 2015-2035: Literature Review & Projections of Future Usage*. United States Department of Transportation, Research and Innovative Technology Administration, Washington, DC (2013).
3. J.A. Richards, *Remote Sensing Digital Image Analysis*. Springer-Verlag, Berlin (1999).
4. R. Bridgelall, J.B. Rafert and D.D. Tolliver, "Hyperspectral applications in the global transportation infrastructure", *Proc. European Association for Signal Processing*. EUSIPCO, Nice (2015).
5. A.M. Baldridge, S.J. Hook, C.I. Grove and G. Rivera, "The ASTER Spectral Library Version 2.0", *Remote Sens. Environ.* **113**, 4 (2009). doi: <http://dx.doi.org/10.1016/j.rse.2008.11.007>
6. P. Mather and B. Tso, *Classification Methods for Remotely Sensed Data*. CRC Press, Boca Raton, Florida (2003).
7. Q. Du and J.E. Fowler, "Low-complexity principal component analysis for hyperspectral image compression", *Int. J. High Perform. C.* **22**, 4 (2008). doi: <http://dx.doi.org/10.1177/1094342007088380>
8. Y. Tarabalka, J.A. Benediktsson and J. Chanussot, "Spectral-spatial classification of hyperspectral imagery based on partitional clustering techniques", *IEEE T. Geosci. Remote* **47**, 8 (2009). doi: <http://dx.doi.org/10.1109/TGRS.2009.2016214>
9. P. Bajcsy and P. Groves, "Methodology for hyperspectral band selection", *Photogramm. Eng. Rem. S.* **70**, 7 (2004).
10. Y-T Chen, *TMS320C6000 Integer Division. Application Report SPRA707*. Texas Instruments, Inc., Richardson, TX (2000).
11. E.A. Cloutis, "Review article: Hyperspectral geological remote sensing: evaluation of analytical techniques", *Int. J. Remote Sens.* **17**, 12 (1996). doi: <http://dx.doi.org/10.1080/01431169608948770>
12. G.B. Thomas and R.L. Finney, *Calculus and Analytic Geometry*, 9th Edition. Addison Wesley, Boston, Massachusetts (1995).
13. D. Coulter, P.L. Hauff and W.L. Kerby, "Airborne hyperspectral remote sensing", *Proc. 5th Decennial International Conference on Mineral Exploration*. Decennial Mineral Exploration Conferences, Toronto, Canada (2007).
14. J-M Muller, *Elementary Functions: Algorithms and Implementation*. Springer Science & Business Media, Berlin (2006).
15. B. Cole, *STMicro, ARM Do a Double Whammy with New Cortex-M7 Core*. Embedded.com, (24 September 2014). www.embedded.com, [Accessed 25 April 2015]