S. Adler-Golden and R. Sundberg, "Identifying vehicles with VNIR-SWIR hyperspectral imagery: sources of distinguishability and confusion," Proc. SPIE 9976, Imaging Spectrometry XXI, 99760K; (2016)

doi:10.1117/12.2238811

See next page.

Identifying vehicles with VNIR-SWIR hyperspectral imagery: Sources of distinguishability and confusion

Steve Adler-Golden* and Robert Sundberg
Spectral Sciences, Inc., 4 Fourth Ave., Burlington, MA 01803

ABSTRACT

Multispectral and hyperspectral imaging can facilitate vehicle tracking across a series of images by gathering spectral information that distinguishes the vehicle of interest from confusers. Developing effective algorithms for utilizing this information requires an understanding of the sources and nature of both the common and unique components in vehicle spectra, as well as the variations associated with lighting, view angle, and part of the vehicle being observed. In this study, focusing on the VNIR-SWIR spectral region, we analyze hyperspectral data from a recent field experiment at the Rochester Institute of Technology. We describe the spectra of painted vehicle surfaces in general terms, and demonstrate effective classification of automobiles based on spectra from upward facing surfaces (the roof, hood or trunk) using a method that combines the Support Vector Machine with data pre-conditioning.

1. INTRODUCTION

Hyperspectral imaging (HSI) sensors have the ability to identify objects based on the spectral signatures of their surface materials. With these signatures, moving targets can be tracked across multiple images, which may be taken over time periods of days or more and under varied lighting conditions. This capability can however be compromised by variations in the observed object signature. Such variability is common with three-dimensional objects, such as vehicles, that contain specular (non-Lambertian) surfaces and/or multiple materials and surface angles. Especially when combined with non-uniform and/or changing illumination and viewing geometries, there can potentially be significant variation in the target spectral signature at both whole-object and sub-object spatial scales. The challenge is to understand these signatures and their variability, and to use that knowledge to build algorithms that optimize the desired of objectives of detection and identification.

In this paper we present an analysis of VNIR-SWIR (visible-near infrared-shortwave infrared) HSI data acquired in a Fall, 2015 measurement campaign at Rochester Institute of Technology (RIT), focusing on the identification of automobiles across images taken hours apart. The experiment is described in a recent paper.[1] First-principles spectral image simulations[2] have been conducted and found to compare well with the data. Both the data and simulations show that spectra of vehicle surfaces with a given paint tend to be similar in shape but subject to large differences in overall brightness with different illumination and viewing angles due to their specularity. This property of the paint spectra motivates the use of either amplitude-normalized data or a normalized measure of spectral similarity, such as spectral angle, in data exploitation algorithms. The small variations in spectral shape that are physically meaningful can be represented with a subspace of only a few dimensions, typically one or two.

The purpose of this paper is to demonstrate an effective data processing scheme for identifying vehicles in one hyperspectral image using spectra from another image taken with a different sun angle, making use of our general knowledge of paint spectrum behavior. To minimize the influence of nearby objects, we take training spectra from only the vehicle top surfaces (the roof, hood or trunk), which are the most consistently illuminated. The data processing employs a Support Vector Machine (SVM) algorithm[3], operating in an all-against-one classification mode, together with data pre-processing and spatial filtering of the SVM output score. We use amplitude-normalized data in lieu of a spectral angle-based SVM.[4]

*adlergolden@spectral.com; phone 1 781 273-4770; fax 1 781 270-1161; spectral.com

The distinguishability of the vehicle spectra arises mainly from the visible part of the spectrum, where there are absorption features associated with pigments. There is some additional information in the near infrared and SWIR regions of the data, extending to around 2500 nm, that can facilitate discrimination. Here the paint spectra tend to be fairly linear, but they may differ in slope. We found that including the SWIR region can dramatically improve discrimination with a less than optimal algorithm. However, with proper algorithm tuning we were able to get very good, and nearly equivalent, classification performance using VNIR region (below 1000 nm) alone. Somewhat surprisingly, even the silver painted cars, which appear virtually identical in color, could be discriminated using this approach.

2. EXPERIMENTAL DATA

The field experiment was conducted at the RIT campus on October 8, 2015, during which 23 hyperspectral images of vehicles and other targets of interest were collected over a six-hour time span. There was little cloud cover during most of the experiment, yielding full sun illumination in all but a few of the images. The HSI system, consisting of Headwall Photonics, Inc. Hyperspec VNIR and SWIR sensors, was operated by the University of Massachusetts, Amherst from the rooftop of RIT's Center for Imaging Science. The sensors looked down on a nearby parking lot, a portion of which was reserved for the placement of spectrally characterized rental vehicles and calibration targets. The VNIR and SWIR sensors had overlapping fields of view to provide complete coverage of wavelengths from 400 nm to 2600 nm in twenty-four hyperspectral image pairs.

The raw VNIR and SWIR data cubes were provisionally calibrated to spectral radiance, georegistered using ground control points, and combined into full-length VNIR-SWIR data cubes. For this paper, the data were downsampled by 2x2 pixel aggregation and spectrally resampled to 259 evenly spaced spectral bands to make the data files smaller and faster to manipulate. As part of the resampling process, problems with the absolute and spectral radiometric VNIR calibrations and the presence of a baseline offset were identified and partially corrected. Recalibration of the sensors is in progress and will be incorporated in an anticipated future release of the HSI data. The final step was atmospheric correction processing on selected images using the FLAASH code[5] in ENVI 5.0 (Harris Corp.), which converts the data to reflectance units. This last step amounts to a spectrally and solar angle-dependent linear rescaling, as scattering contributions (i.e., path radiance and adjacency components) are negligible at the close range of these measurements.

3. TEST CASES

Our objective is to identify a target vehicle in one hyperspectral image (the training image) and find that same vehicle in another image (the test image) acquired at a different time, when the vehicle may have moved or the sun angle changed. We present results here for vehicles identified in image number 4 (i.e., the training image), which was taken at 11:35 AM, and searched for in image number 14 (the test image), taken at 2:32 PM. While measurements continued for close to two more hours, the later images are narrower and do not show as many vehicles in common with image 4.

Figure 1 shows RGB color renderings of the parking lot in images 4 and 14. The rental cars, which are the first three cars in the left hand row, have changed location and/or orientation. Many of the vehicles in the right hand row have remained in their place. However, the sun angle variation over the three-hour period between the images is sufficient to cause up to factor-of-two changes in the brightness of the tops of these vehicles (for example, see the gold truck) as well as dramatic changes in the lighting of the facing sides of the vehicles, which are in shadow in image 14.

Besides the white and bright blue rental cars, we are interested in vehicles that appear in both images behind the rental cars and in the right hand row. Their colors include silver, gold, dark blue, and dark green. The VNIR and SWIR images are poorly registered for the more distant vehicles, resulting in corruption of their full-length spectra, so those vehicles could be analyzed with only one of the two spectral ranges at a time. We have also run a few cases with different training and test images, including one taken in the late afternoon. Those results are incomplete, but thus far they are similar to those reported here.

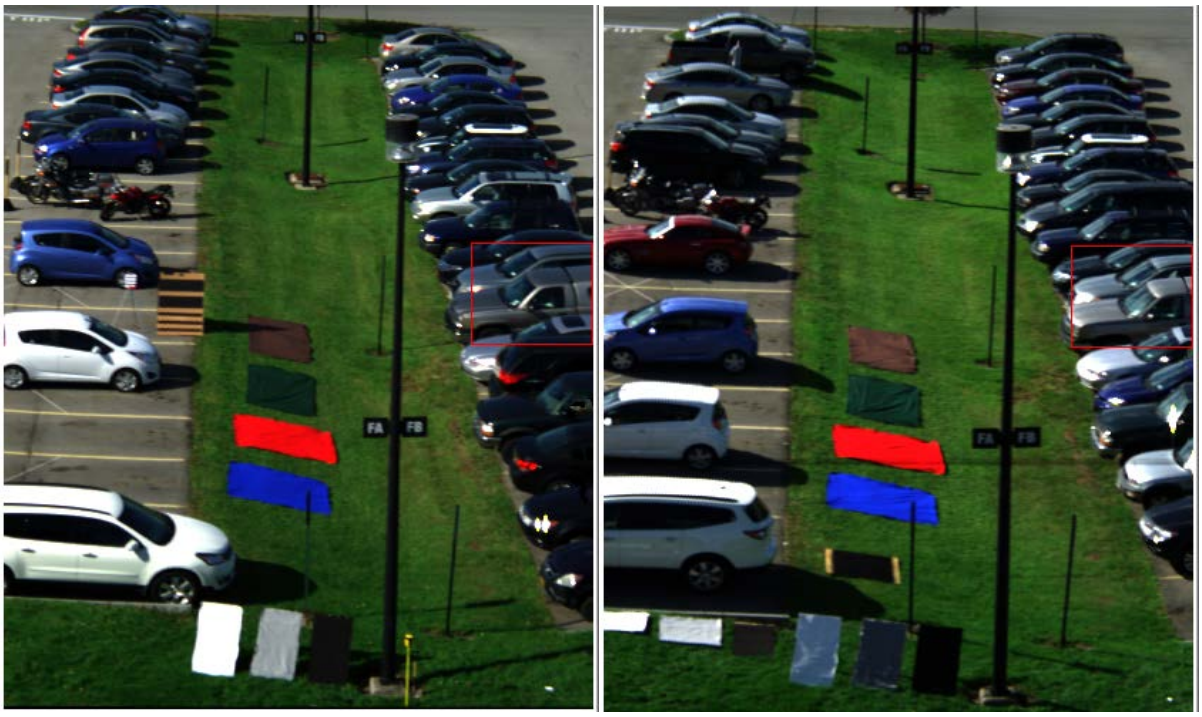

Figure 1. Vehicles of interest in image 4 (left) and test image 14 (right). Red squares are centered on a gold truck and silver car in common.

ENVI and IDL software (Harris Corp.) was used to define a region of interest (ROI) for each target vehicle in the training image, perform preprocessing on the target and test image spectra, run the ROI-trained SVM algorithm on the preprocessed data, and convert the output score image into individual vehicle detections, which were then manually identified as true or false detections. The detections were tabulated as a function of preprocessing and SVM options and parameter values. We manually chose the training ROIs on the vehicle roof, hood and/or trunk, typically numbering between 150-250 pixels. These areas are well illuminated by the sun, and, unlike the sides of the vehicles, are unaffected by illumination from the immediate surroundings.

4. ALGORITHMS

4.1 Support vector machine

The SVM classifier algorithm was selected for its typically good performance with diverse target training spectra. All-against-one classification of the test data was performed using the ENVI SVM algorithm with the radial basis function (RBF) kernel. Since the ENVI implementation is not designed to apply training data from one image to another image, we developed a work-around that uses a hybrid image containing both the training data and the test data. We found that by far the best SVM results were obtained with normalized spectra, in which the reflectance data are divided by their RMS value. With those data, the best value of gamma, which controls the width of the RBF, was in the range of 500 to 2000 times the default value, which is the reciprocal of the number of bands. The lower end of this range turned out to be somewhat more robust.

The SVM output score image provides a qualitative measure of the likelihood of target presence in each pixel. However, what is needed is a way to define detections of whole objects (vehicles). To do this, we start by selecting some number, M, of top-scoring pixels. To discriminate against isolated false positives, such as object edges, we then require that only contiguous clusters containing a certain minimum number of pixels, N, be counted as detections. With very rare exception, these clusters are all on vehicles. Finally, we use visual inspection to identify the vehicles as target or non-target. With the current HSI data we obtained nearly optimal results using M=500 and N=30 for the entire set of target vehicles.

4.2 Data whitening preprocessing

The SVM algorithm automatically normalizes each spectral band intensity to its variance in the training data—i.e., it performs a diagonal "whitening" transform. This provides a better match of the data distribution to the shape of the RBF kernel, which is isotropic in spectral space. It also makes the results independent of a linear or affine transform of the data that is diagonal in that space.

By extension, one might expect better performance of the SVM with a more thorough whitening transform, derived from a data covariance matrix, that also decorrelates the data. Such a transform may be either vehicle-specific, designed for all-against-one classification, or based on the collective distributions of all the vehicle data, as would be used in simultaneous multiclass classification. We consider the latter, "universal" type of whitening transform in this paper. A single covariance matrix is built from all the normalized ROI spectra by subtracting the corresponding ROI means; there is a sufficient total number of training pixels to insure good statistical sampling. An initial whitening matrix is calculated as the square root of the covariance matrix inverse. The user may specify a covariance regularization, or "shrinkage," step, implemented here by setting a lower bound on the covariance matrix eigenvalues, i.e., an upper bound on the whitening matrix elements in diagonalized (principal component) space. We refer to this whitening procedure as "UCW," for Universal Covariance Whitening.

5. RESULTS

Figure 2 illustrates typical outputs from the target cluster identification and size selection step, which shows the pixel clusters (in magenta) corresponding to potential targets in the test image. The left-hand image represents a case in which there are no false positives, while the right-hand image represents a more difficult identification case in which there is a false positive. These examples were run with full-length spectra and no UCW. Note that, as in the training data, the correctly identified parts of the vehicle are the upward facing surfaces.

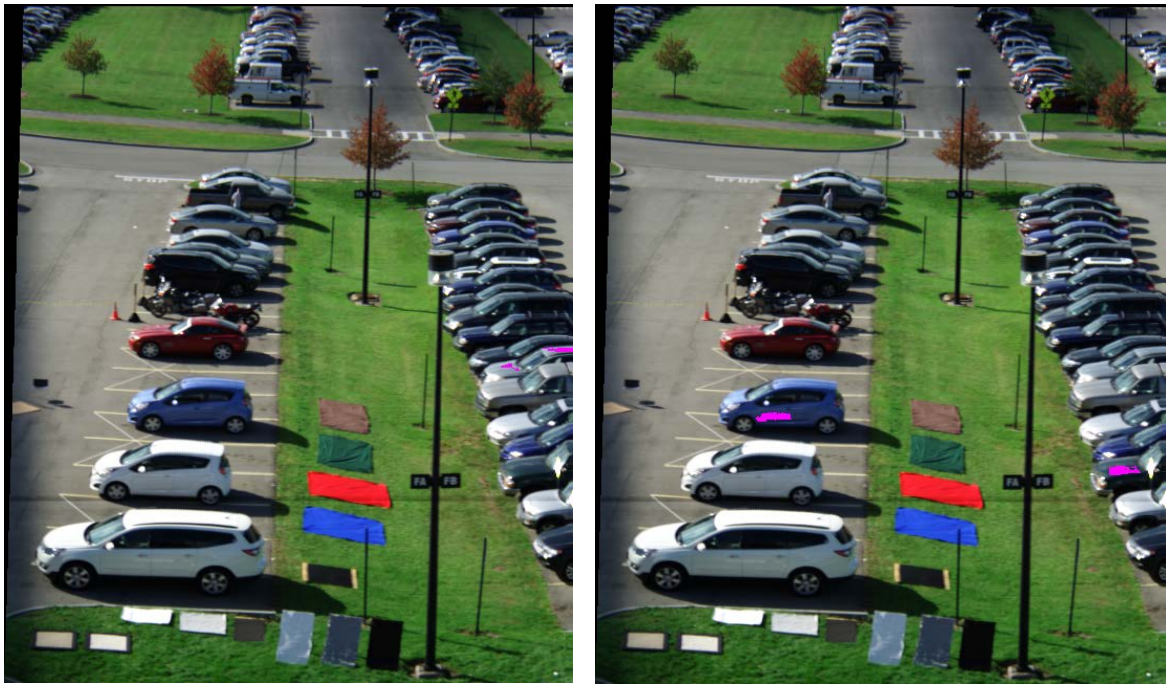

Figure 2. Example vehicle identifications. At left, silver car; at right, dark green car.

Summarized results for five vehicles without UCW are shown in Table 1 for a single, fixed set of parameters gamma, *M* and *N*. The table also includes results using only the VNIR portion of the spectrum (up to 1000 nm). We observe that the VNIR alone is inadequate for identifying the three silver and gold vehicles but improves identification of the green car by eliminating the false positive on the side of the blue rental car. The white compact car is from the same manufacturer (Ford) as the white station wagon, and may have the same paint, which would explain the presence of the false detection, which is on the station wagon.

Table 1. Vehicle identification results without UCW.

Vehicle	True detection	# False detections	True detection (VNIR only)	# False detections (VNIR only)
Gold truck	yes	0	no	0
Bright silver car	yes	1	no	0
Silver car	yes	0	no	1
Dark green car	yes	1	yes	0
White compact	yes	1	yes	1

Table 2 presents results with UCW. As in Table 1, the same processing parameters are used throughout. The UCW transform has an adjustable parameter, "rank," which specifies a covariance eigenvalue ordinal number that defines the eigenvalues' lower limit. With an appropriately tuned gamma scale factor (around 2000) and the full spectrum, UCW processing was found to yield somewhat larger identified target areas, which allowed a slightly higher size threshold of N=40 to be used.

In addition to the five vehicles in Table 1, Table 2 includes a third silver car, located behind a dark gray car in the rental car row, as well as the blue rental car, which is very easy to identify. The silver car spectrum is corrupted by poor spatial registration between the VNIR and SWIR images, so we show just the VNIR-only result, which compares favorably with those from the other vehicles.

The full spectrum (VNIR+SWIR) results with UCW in Table 2 are slightly better than those given in Table 1, since the dark green false positive is removed. However, a dramatic improvement is found in the VNIR-only results, where now all vehicles are identified, and discrimination against false positives is as good as with the full spectrum when rank=9. Except for the bright silver car, the results are insensitive to the rank value.

Table 2. Vehicle identification including the UCW transform.

Vehicle	Full spectrum true detection rank=5 or 9	Full spectrum false detections rank=5 or 9	VNIR true detection rank=5 or 9	VNIR false detections rank=5	VNIR false detections rank=9
Gold truck	yes	0	yes	1	1
Bright silver car	yes	1	yes	1	0
Silver car	yes	0	yes	0	0
Dark green car	yes	0	yes	0	0
White compact	yes	1	yes	1	1
Blue rental	yes	0	yes	0	0
Silver car behind gray	N/A	N/A	yes	1	1

6. DISCUSSION AND CONCLUSIONS

The potential of VNIR-SWIR HSI spectral signatures to assist in identifying and tracking vehicles across times and locations has been demonstrated with this case study. These results, from data taken of a parking lot in clear weather in a close range, slant view, should translate to longer ranges as long as a good atmospheric correction can be done and there are enough pixels on target. Using one image for training and another image for testing, a half dozen selected target vehicles were located in the test image, with at most one false positive over the entire parking lot. The data processing used a combination of normalized reflectance data, an optional non-diagonal whitening transform, and a Support Vector Machine (SVM) classifier. The targets included multiple vehicles with silver paint, whose colors are essentially indistinguishable by eye and whose spectra are only subtly different.

The distinguishability of the vehicle spectra arises mainly from the VNIR part of the spectrum, where there are absorption features associated with the paint pigments. Although the SWIR region contains additional spectral information, the VNIR region alone provided equally good discrimination ability when the data were processed with the whitening transform. The transform, which is derived from a training data covariance matrix, de-emphasizes the major sources of variability in the spectra, making the data distribution more isotropic in the leading principal components.

A critical ingredient in the whitening transform is regularization, which puts a floor under the covariance eigenvalues, and thus insures that the transform does not over-emphasize minor sources of variability in the training spectra. We surmise that these minor variabilities are image-dependent and not intrinsic to the paint surface, and thus represent a source of confusion in the classification process. The required degree of regularization is much greater than what is needed to compensate for covariance under-sampling.[6]

The results we have presented exclude vehicles with very dark paints, other than the green car. We did not attempt to run the processing for a black car, whose spectrum is essentially indistinguishable from that of deep, hard shadow. In addition, our processing failed for dark blue and dark gray cars. We believe this may be due to corruption of the spectra by a small background level of scattered light, making us reluctant to draw general conclusions about classification ability with dark vehicles.

Future work should include a similar investigation with better quality imagery, as well as a more thorough examination of the current imagery that includes data at later times and under cloud shadow. Another area that is ripe for additional work concerns the whitening transform, where vehicle-specific transforms and different methods of regularization should be investigated.

ACKNOWLEDGEMENTS

The authors thank the Air Force Research Laboratory, RYAT for support under contract number FA8650-15-C-1862. We also acknowledge many helpful discussions with Dr. Emmett Ientilucci and Mr. Joseph Svejkosky of the Digital Imaging and Remote Sensing Laboratory, Rochester Institute of Technology, and experimental support by Dr. Mario Parente of Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA.

REFERENCES

[1] Svejkosky, J., E. Ientilucci, S. Richtsmeier, M. Parente and C. Bachmann, “A hyperspectral vehicle BRDF sampling experiment,” Proc. SPIE 9840, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XXII, (2016).

[2] Perkins, T., S. Adler-Golden, L. Muratov, R. Sundberg, E. Ientilucci and L. Cain, “Spectral BRDF modeling of vehicle signature observations in the VNIR-SWIR,” Proc. SPIE 9840, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XXII, (2016).

[3] Hsu, C.-W., Chang, C.-C., and Lin, C.-J., “A practical guide to support vector classification,” National Taiwan University, http://ntu.csie.org/~cjlin/papers/guide/guide.pdf (2007).

[4] Honeine, P. and C. Richard, “The angular kernel in machine learning for hyperspectral data classification,” Proc. 2nd WHISPERS Conference, Reykjavic, (2010).

[5] Perkins, T., S.M. Adler-Golden, M.W. Matthew, A. Berk, L.S. Bernstein, J. Lee and M. Fox, “Speed and accuracy improvements in FLAASH atmospheric correction of hyperspectral imagery,” Opt. Eng. 51(11), 111707 (2012).

[6] Chen, Y., A. Wiesel, Y.C. Eldar and A.O. Hero, “Shrinkage algorithms for MMSE covariance estimation,” IEE Trans. Signal Processing, 58, 5016-5029 (2010).