Robert K. McConnell
President
WAY-2C Color Machine Vision
14 Stevens Terrace
Arlington, MA 02476
tel: 781-641-0605
fax: 781-648-5807
rkm@way2c.com

Relevance Spectroscopy: Optimizing Interpretation Efficiency of Hyperspectral, Multispectral, and Other Data When Time, Bandwidth and/or Computing Power is Limited.

Despite a flood of visual, sonic and other sensory data, we humans are able to quickly discard the mundane and the irrelevant and focus our attention only on that which is relevant. The process may be regarded as one of data reduction with the useful output equivalent to a few bits of information of momentary interest.

Modern robotic devices and vehicles are being equipped with ever more sophisticated imaging and other sensory devices with the potential to return information of interest for human or automated interpretation. Usually each sensor is designed to handle a wide variety of as yet undetermined scenes in order to detect objects or conditions that may or may not be known much in advance of the mission. By necessity, much of the data captured for any single mission is irrelevant.

The irrelevance can arise from such factors as unused data precision, redundant channels and excessive spatial resolution. It can become a problem when time, bandwidth and/or computing power is limited. For example, a hyperspectral imager capturing two bytes of data for each of 240 spectral bands would generate 480 bytes of data per pixel. A typical instrument that could be carried on a light aircraft or robotic vehicle might have 640 spatial pixels and generate 145 image frames per second. If all this data must be transmitted to a central location for immediate interpretation, a bandwidth of the order of 300 megabits per second would be required. While this can be handled by present day wired connections, it can be a serious issue with wireless communications.

Typically the bandwidth can be reduced by at least one, and possibly two orders of magnitude if only the relevant information, the data necessary to differentiate the classes of interest for the current mission, are returned for interpretation and/or display. Often this relevant information is available from previous knowledge of the spectra of the classes of interest. Our concern here is what can be done when it's not available and must be determined on the fly.

One well-known approach to reducing the bandwidth is to perform principal component analysis (PCA) of the scene. PCA seeks to find a spectral coordinate system in which the data set can best be represented using only a small number of components. Once these principal components are determined, the on-board system might transform each pixel into the principal component space and then return the coefficients of only the most important.

PCA has two obvious drawbacks: a high computational cost because of its traditional implementation which requires calculation of a covariance matrix across the data[1] and, as Landgrebe [2] has pointed out "*the principal component transformation provides a convenient* *means for ordering the features and selecting those that will make the greatest contribution to representation of the data (but not necessarily to classification!)*"

Another way to accomplish the data reduction task is what we call relevance spectroscopy. This differs from PCA in that it assumes the identity of the classes of interest are known and that images containing these classes are available. Using supervised classification and information theory, relevance spectroscopy provides a direct method of determining which combination of spectral bands is most suitable for differentiating the specified classes. It makes no assumptions about the modality of the classes of interest and therefore can accept such things as multicolor camouflage, water bodies with glints, and multicomponent desert scrub each as a single class. Typical computation times are under a minute on a PC.

Our experiments on a variety of hyperspectral images have shown that once the classes of interest have been chosen, for practical purposes enough necessary to reliably differentiate these classes is usually available in three or fewer spectral bands. If the necessary data subset can be determined before, or even during, the mission, the transmission bandwidth requirements can typically be reduced between one and two orders of magnitude.

Eliminating irrelevant data is only one step towards optimizing interpretation efficiency. The ultimate goal is to determine the few bits of information of momentary interest such as the class associated with a region in a scene, the next word in a conversation or the state of a person or process being monitored. Once the data set has been reduced so it contains only the most relevant information, classification becomes much easier. When images can be reduced to three or fewer spectral bands they can be displayed directly in RGB false color suitable for interpretation by a human. When real-time and/or unattended automated interpretation is required, this same reduced data set is well suited to the information theoretic classification methods discussed and demonstrated at an earlier conference [3]. Classification results in a further bandwidth reduction, typically to less than 8 bits per pixel; it can typically be performed real time on-board the vehicle or at the base station with the same modest PC class computer and the same software used for the optimum band determination.

While this presentation focuses on image interpretation, the methods discussed are equally applicable to, and have been successfully demonstrated on, a variety of sensor data, both primary and derived.

References:

[1] Qian Du and James E. Fowler, "Low-Complexity Principal Component Analysis for Hyperspectral Image Compression", *International Journal of High Performance Computing* Applications *22*, 438-448, November 2008.

[2] David Landgrebe, *Signal Theory Methods in Multispectral Remote Sensing,* New York: John Wiley & Sons, 2003.

[3] Robert K. McConnell, ""Off-the-Shelf System for Color and Multispectral Based Recognition and Control", *2008 IEEE International Conference on Technologies for Practical Robot Applications*, Woburn Massachusetts, November 10-11, 2008