

Spectral Image Utility Prediction

Marcus S. Stefanou
United States Air Force
Rochester Institute of Technology
Rochester, New York 14623 USA
mss3679@cis.rit.edu

John P. Kerekes
Rochester Institute of Technology
Chester F. Carlson Center for Imaging Science
54 Lomb Memorial Drive
Rochester, New York 14623 USA

Abstract—The utility of an image is an attribute that describes the ability of that image to satisfy performance requirements for a particular application task. The robust ability to predict the utility of an image for a given application would facilitate sensor design trade studies, provide a basis for tasking image collection activities, and create the foundation for an image archive indexing scheme. In this paper, we examine methods for predicting the utility of spectral images for detecting sub-pixel targets using the constrained energy minimization matched filter detector.

The result of our initial work is a prediction of the likelihood of finding a synthetically implanted target in a target-free image in advance of actually applying the detector. We define image utility for the target detection application as the probability of detection at a specified probability of false alarm. We analytically predict this utility for a given image by first estimating statistical parameters directly from the image, then operating on these parameters with the matched filter detector. Three parametric statistical models are used for characterizing the image background: the global Gaussian, the multiple-class Gaussian, and the elliptically contoured t-distributions. The target models come from a library of target materials and are assumed to be multivariate Gaussian with known mean and covariance. We benchmark prediction performance by comparing predicted detection probabilities to empirical results obtained by applying the detector to the data for two HYDICE images.

Our longer term objective is to build on this initial result by developing a more general spectral image quality and utility framework, and specific metrics for the prediction of utility across many potential applications.

Keywords—spectral image utility, target detection performance prediction, hyperspectral

I. INTRODUCTION

Advances in spectral imaging – the process of collecting spatially co-registered images in multiple spectral bands – have enabled sensor systems with increasing spectral and spatial image resolution. The additional scene information inherent in these data promises improved performance in many diverse applications such as land cover classification, mineral exploration, atmospheric characterization, and unresolved object detection. Of course, performance is relative to the application task at hand, so we introduce the term “image utility” to describe the ability of the image to satisfy performance requirements for a specific application.

While many research projects have demonstrated high performance accuracies in applications for a given image, an open and interesting research topic is the general *prediction* of potential application utility for a particular image or imaging scenario. The robust ability to predict the image utility for a given application is desirable because it would facilitate sensor design trade studies, provide a basis for tasking image collection activities, and create the foundation for an image archive indexing scheme.

We propose an approach to predicting the spectral image utility for the specific application of detecting sub-pixel targets using the constrained energy minimization (CEM) detector. Our methodology predicts the likelihood of finding a synthetically implanted target in a target-free image in advance of applying the detector. It differs from the Forecasting and Analysis of Spectroradiometric System Performance (FASSP) analytical model-based approach to prediction described in [1-3] in that it begins with a real image rather than a statistical description of a notional imaging scenario, thereby capturing the effects of the scene, atmosphere, and sensor in estimated statistical parameters. These multivariate statistical parameters are then transformed, rather than the data itself, by the CEM detector into parameters that describe the scalar output test statistic. Hence, the prediction assesses detector performance in a more computationally efficient manner than actually applying the filter to the data. This paper presents initial results from the application of this methodology to two Hyperspectral Digital Imaging Collection Experiment (HYDICE) datasets.

II. APPROACH

A. Data Preparation

The two HYDICE reflectance scenes were obtained via the empirical line method (ELM) of atmospheric compensation. Feature space transformations removed atmospheric absorption bands, resulting in 144 bands for the Desert Radiance II (DR II) dataset and 145 bands for the Forest Radiance I run 05 (FR I) dataset. Both images were collected with 0.75m spatial resolution. Because we did not want to contaminate the background statistical parameter estimates with target pixels, target masks based on ground truth were applied to each image resulting in 387,149 target-free pixels in the FR I dataset and 284,581 target-free pixels in the run 03 DR II dataset. The target-free pixels in each dataset were then classified using a k-means classifier, with 12 classes for the FR I dataset and 8 classes for the DR II dataset. 18 target mean vectors and covariance matrices were drawn from the FAASP spectral

reference library. These library statistical descriptions of the targets were previously estimated from the known ground truth target pixels in HYDICE imagery as part of the FASSP development.

B. Methodology

The utility or target detection performance of a particular image for a given target is assessed using a receiver operating characteristic (ROC) curve that expresses the probability of detection over a range of probabilities of false alarm. The probabilities used in the ROC curve are derived from the distributions of the test statistic at the output of the target detection algorithm. The CEM filter, \mathbf{w} , is constructed with the background inverse covariance matrix, Σ^{-1} , estimated from the target-free data, and a target mean vector, \mathbf{t} , drawn from the FASSP spectral library.

$$\mathbf{w} = \frac{\Sigma^{-1}(\mathbf{t} - \boldsymbol{\mu})}{(\mathbf{t} - \boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{t} - \boldsymbol{\mu})} \quad (1)$$

Vector \mathbf{w} is the CEM filter operator and $\boldsymbol{\mu}$ is the target-free data mean vector.

In order to generate the ROC curve, the output of the matched filter is needed for both a target absent and a target present case as described in [4]. When calculating the empirical (actual) image utility, the target absent case is created by applying the filter operator to each pixel of the demeaned target-free data.

$$y_{TA} = \mathbf{w}^T (\mathbf{x} - \boldsymbol{\mu}) \quad (2)$$

The scalar y_{TA} is the output in the target absent case at one pixel location. The target present case is created by first fractionally implanting a target pixel vector at the prescribed fraction in each pixel of the target-free data (3) and then applying the matched filter to the resulting target present pixel (4). The implanted target pixel, \mathbf{x}_{TP} , is created by realizing a random vector, \mathbf{t}' , from a random process obeying the parameters of the target mean and covariance from the FASSP reference library. This target variability leads to a more realistic detection scenario.

$$\mathbf{x}_{TP} = f\mathbf{t}' + (1-f)\mathbf{x} \quad (3)$$

The random target vector is mixed fractionally with every data pixel to the specified fraction f . The calculated empirical target present output of the CEM filter at one pixel location is y_{TP} .

$$y_{TP} = \mathbf{w}^T (\mathbf{x}_{TP} - \boldsymbol{\mu}) \quad (4)$$

In calculating the predicted image utility, three statistical models are employed to describe the distribution of the scalar test statistic at the output of the filter. The assumption is that the assumed multivariate statistical distribution of the dataset is linearly transformed by the action of the filter to a scalar distribution. The first statistical model is a Gaussian having the parameters described by the transformed mean and covariance of the entire (global) target-free dataset. The second model is created by a weighted sum of the individual spectral class test statistic distributions. In the equations that follow, the global Gaussian model is treated as a special single class case of the

class sum model, in which the subscript i refers to the class number. The transformation from multivariate to scalar statistics for the target absent case of this model is:

$$\begin{aligned} \theta_{TA_i} &= \mathbf{w}^T (\boldsymbol{\mu}_i - \boldsymbol{\mu}) \\ \sigma_{TA_i}^2 &= \mathbf{w}^T \Sigma_i \mathbf{w} \end{aligned} \quad (5)$$

where θ_{TA_i} is the target absent distribution scalar mean, $\boldsymbol{\mu}_i$ is the multivariate class mean, Σ_i is the class covariance, and $\sigma_{TA_i}^2$ is the target absent distribution variance. The test statistic output distribution is then the weighted sum of Gaussian distributions, each parameterized by θ_{TA_i} and $\sigma_{TA_i}^2$. The target present situation is created by appropriately mixing the mean vectors and covariance matrices of the target and background, and then applying the filter to the target present multivariate statistics to obtain their scalar versions.

$$\begin{aligned} \boldsymbol{\mu}_{TP_i} &= f\mathbf{t} + (1-f)\boldsymbol{\mu}_i \rightarrow \theta_{TP_i} = \mathbf{w}^T (\boldsymbol{\mu}_{TP_i} - \boldsymbol{\mu}) \\ \Sigma_{TP_i} &= f^2 \Sigma_i + (1-f^2)\Sigma \rightarrow \sigma_{TP_i}^2 = \mathbf{w}^T \Sigma_{TP_i} \mathbf{w} \end{aligned} \quad (6)$$

Vector $\boldsymbol{\mu}_{TP_i}$ is the i th class mean mixed with the FASSP reference library target mean vector, \mathbf{t} , θ_{TP_i} is the target present distribution scalar mean, Σ_{TP_i} is the i th class covariance mixed with the target covariance, Σ_T , and $\sigma_{TP_i}^2$ is the target present distribution variance.

The third statistical model uses the elliptically contoured T-distribution to represent a scalar test statistic distribution of the form:

$$t(\mathbf{v}; \theta, \sigma^2, M) = \frac{\Gamma(\frac{1+M}{2})}{\sigma \sqrt{\pi} M \Gamma(\frac{M}{2})} \left[1 + \frac{1}{M\sigma^2} (\mathbf{v} - \theta)^2 \right]^{-\frac{1+M}{2}} \quad (7)$$

in which $\Gamma(\cdot)$ is the gamma function, M is the degrees of freedom (DOF) parameter, θ and σ^2 represent either the target absent or target present parameters from (5) and (6), and \mathbf{v} is the test output threshold variable vector. The T-distribution DOF parameter is varied with 12 values ranging between 1 and 50.

For both the actual and predicted utility cases, the target absent and present distributions (whether estimated or analytically specified) at the filter output are transformed into cumulative distribution functions (CDFs). The appropriate CDFs are used to calculate the probabilities of false alarm (PFA) and the probabilities of detection (PD), which lead to a ROC curve when plotted against each other. A separate ROC curve is generated for each target and target fraction. In this paper, 18 targets and 15 fractions are considered, leading to 270 ROC curves for each dataset.

A simple form of predicting the degree of difficulty associated with detecting a target against a particular background is the signal-to-clutter ratio (SCR) adapted from [5].

$$SCR = f \sqrt{(\mathbf{t} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{t} - \boldsymbol{\mu})} \quad (8)$$

Vector \mathbf{t} is the FASSP reference library mean of the target, $\boldsymbol{\mu}$ is the mean of the background, Σ^{-1} is the inverse covariance matrix of the background, and f is the fractional implant of the target. Because the SCR is similar in form to the CEM (1) operator, we would expect the SCR to be a good predictor of the degree of difficulty in detecting a particular target, and we will employ SCR as a benchmark with which to compare our image-derived utility predictions.

III. RESULTS

We first examine the behavior of the same target in both datasets. The target selected (#4) is one of the more difficult to detect, having the lowest SCR among all targets for both datasets. Fig. 1 shows the actual (empirical) and predicted results for detecting target #4 in each of the datasets using two perspectives. The first is the ROC curve at a specific target fraction implant value, in this case 0.10. This shows detection behavior over all possible PFAs, but at a single target fraction. The second perspective plots the value of the probability of detection (PD) at a specified PFA as a function of the target fraction. This potentially offers a more realistic characterization of detection performance by describing PD at a constant PFA over a range of target fractions. In most target detection problems, the operator is willing to accept a certain number of false alarms, and it is the detection at this PFA that is of interest. The specified false alarm probability for this example

is 5×10^{-4} , which translates to 194 pixels in the FR I dataset and 142 pixels for the DR II dataset.

In comparing the FR I and DR II empirical results (solid lines in Fig 1), we see that the target is slightly more detectable in the DR II than in the FR I dataset. This is apparent in the ROC curves at a single fraction, where a PD of 0.6 is attained in DR II at a PFA of 0.01 as opposed to a PD of 0.4 in the FR I dataset. It is also clear across multiple fractions, where perfect detection (1.0) is attained for the DR II dataset at a target fraction of 0.3 whereas perfect detection is attained at a fraction greater than 0.5 for FR I.

The detection performance for each prediction method is depicted in Fig. 1 as dashed and dotted lines. The prediction methods correspond to the three statistical models (global Gaussian, sum of Gaussian classes, and elliptically contoured T-distribution). The cases of DOF=5 and 25 are shown for the T-distribution. The ROC curves at a single fraction reveal that the prediction based on modeling the image with a global Gaussian distribution produces the worst performance predictions relative to the empirical results. The global Gaussian prediction is overly optimistic because the dataset is more complicated than a single Gaussian model can capture. The DOF=25 parameter also results in a prediction that is optimistic relative to actual performance, and confirms that a DOF>10 leads to T-distributions that tend toward Gaussian. The DOF=5 parameter produces an overly pessimistic prediction of performance, but at very low PDs appears to more accurately match the actual ROC curve than the other predictions. The best performance prediction appears to be the sum of Gaussian classes, particularly in the case of the FR I dataset. This is explained by the fact that the sum of Gaussian classes most accurately represents the reality of the distribution of a real image. We note that the unsupervised classification for the FR I dataset had 12 classes while that of the DR II dataset had 8, and although DR II seems like the simpler dataset, it is possible that the classification has caused some of this effect.

Thus far, the results have focused on the detection of one particular target. We tested the detectability of 18 targets against each dataset in order to better understand the behavior of a wide range of possible targets. For each target, there were 15 different implantation fractions ranging from 0.01 to 0.5 that were used to assess detectability. Thus, we have three parameters that we control to better understand detection performance: input image, type of target to be detected, and target fraction. The SCR (8) relates all three of these parameters; thus it would be logical to use SCR as a surrogate means of predicting performance. Fig. 2 attempts to summarize how “good” each performance prediction method is relative to the actual performance. Each scatter plot in Fig. 2 contains 270 points, corresponding to the 18 targets and 15 fractions. It uses the linear Pearson correlation coefficient (R) as the measure of closeness between the prediction and the actual detection performance at a specified PFA of 5×10^{-4} . Fig. 2 confirms the observation from Fig. 1 that the best prediction

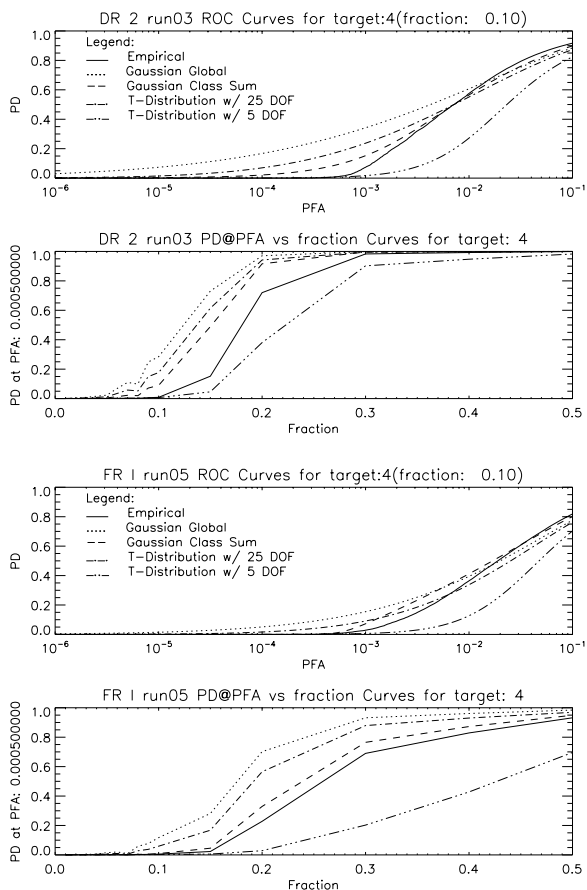


Figure 1. ROC Curves for DR II and FR I Datasets with Target #4

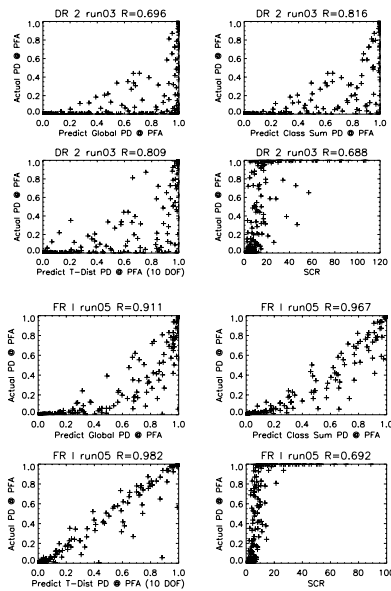


Figure 2. Scatter Plots of Predicted vs Actual Utility

is due to a good model using the sum of Gaussian classes or a properly chosen T-distribution DOF parameter. In Fig. 2, T-distributions with DOF=10 are used, which results in the closest fit to the actual prediction performance in the FR I dataset.

While the prediction based on SCR is clearly not the closest to the actual performance, it is on par with the prediction based on the global Gaussian for the DR II dataset for the specified PFA. All of the scatter plots in Fig. 2 exhibit a large number of points at low PFA and low PD. This is partially due to the relatively low PFA chosen as the operating point for PD and because of the large proportion of small target fractions, 9 of the 15 being between 0.01 and 0.1.

In an attempt to discern a relationship between target implant fraction, target type, and performance, we use a surface plot of the PD achieved at a $PFA=5 \times 10^{-4}$ for all 18 targets (ordered by their SCR) and target fractions from 0.01 to 0.1 for the FR I dataset. In Fig. 3, the x-axis is the ordered SCR of the 18 targets, the y-axis is the target fraction, and the z-axis is the PD at the specified PFA. The general trend is that greater SCR and higher target fraction correspond to easier detectability (the top right portions of the surface tend toward higher values). However, as indicated by the scatter plots of Fig. 2, SCR is not the best possible predictor of performance and thus the PDs are not monotonically increasing with increasing SCR.

An attractive feature of the prediction approaches is that they do not require the computation associated with processing each pixel of the spectral image through the detection algorithm, as actual detection performance is typically assessed. Running a prediction based on image-derived statistical parameters is roughly 11 times faster than processing the entire dataset to yield actual detection performance. The prediction based on a sum of Gaussian classes is roughly 8 times faster than the actual prediction. These estimates are based on running IDL 6.3 on a 3 GHz Pentium 4 processor.




Figure 3. Surface Plots of Targets vs Fraction

IV. SUMMARY AND DISCUSSION

This paper has presented an image-derived approach to predicting the utility of a spectral image for detection of subpixel targets. The approach is based on transforming statistical parametric models of the image into test output statistics that can be portrayed as ROC curves. The results of comparing three image-derived utility prediction approaches to the actual performance for two datasets over a range of 18 targets and 15 subpixel target fractions show that the prediction of performance is heavily dependent on how well the statistical parameters capture the complexity of the real image. In general, the predictions employing a sum of Gaussian classes and predictions based on an elliptical T-distribution with the appropriate number of degrees of freedom outperform predictions based on a global Gaussian model of the image and those based on a simple SCR. These utility prediction approaches have the advantage of running an order of magnitude faster than empirical approaches which rely on processing each data pixel. While at this point the results are dependent on the specific images and targets under investigation, we hope that by further exploration of images and targets we can generalize the results to classes of images and targets and increase prediction methodology robustness.

DISCLAIMER

The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

REFERENCES

- [1] J. P. Kerekes and J. E. Baum, "Spectral Imaging System Analytical Model for Subpixel Object Detection," *IEEE Trans. On Geosci. and Rem. Sens.*, vol. 40, no. 5, pp. 1088-1101, May 2002.
- [2] J. P. Kerekes and D. Manolakis, "Improved Modeling of background distributions in an end-to-end spectral imaging system model," *Proceedings of the 2004 IEEE Geoscience and Remote Sensing Symposium (IGARSS '04)*, Anchorage, Alaska, 2004.
- [3] J. P. Kerekes, "Parameter Studies for Spectral Imager Application Performance," *Proceedings of the 2006 IEEE Geoscience and Remote Sensing Symposium (IGARSS '06)*, Denver, Colorado, 2006.
- [4] C. E. Cafer, S. R. Rotman, J. Silverman, and P. W. Yip, "Algorithms for point target detection in hyperspectral imagery," *Proceedings of Imaging Spectrometry VIII*, SPIE Vol. 4816, pp. 242-257, Seattle, Washington, July 2002.
- [5] J. P. Kerekes, A. Cisz and R. Simmons, "A comparative evaluation of spectral quality metrics for hyperspectral imagery," *Proceedings of Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI*, SPIE Vol. 5806, pp. 469-480, April 2005.