OUTLIER COLOR DETECTION FOR SEARCH AND RESCUE APPLICATIONS

Philip Yuan, Matthew Brawley

Boston University
Department of Electrical and Computer Engineering
8 Saint Mary's Street
Boston, MA 02215
www.bu.edu/ece

May 5, 2017

Technical Report No. ECE-2017-02

Contents

List of Figures

List of Tables

1 Introduction

Search and rescue operations vary vastly in terrain, goals, and implementation. The operations may be completed by manned or unmanned aerial vehicles and are conducted across many different terrains and in search of a variety of objects. Regardless of the choice of aircraft, a number of high resolution cameras can be mounted on the vehicle for capturing video and images of the terrain. Due to the breadth and scope of these operations, the challenge of identifying an object of interest is usually quite significant; the object of interest, e.g. a person, typically occupies a small portion of the image or view due to the altitude from which the image is captured or from which the viewer observes the terrain.

Outlier color detection methods are particularly useful for these applications due to the tendency of search and rescue targets to be a different color than the surrounding terrain, as displayed by the blue body in the woods below in Figure 1. Automated color anomaly detection algorithms can be implemented to assist in the detection of such targets.

Figure 1: Image of 'Woods' with Anomaly

2 Literature Review

A large portion of the literature dedicated to anomaly detection are directed towards hyperspectral images. Hyperspectral imagery differs from the images used in search and rescue applications in that the use of hyperspectral sensors results in an image

composed of up to hundreds of spectral bands including the non-visible and visible [1]. However, many of the underlying concepts of these algorithms developed for use in hyperspectral imaging can be applied to imagery with many fewer visible spectral bands, such as color images.

A basic algorithm for anomalous color enhancement in aerial imagery was developed by Rasmussen, Thornton and Morse [2]. Their algorithm was centered on enhancing the image selectively by determining the frequency of hues in the image, determining the saliency of an object (i.e. some group of pixels that is of note compared to others), and then boosting the saturation of the pixels that are salient and infrequent hue-wise while reducing the saturation of the pixels that are not salient and frequent hue-wise. However, this method is not particularly robust nor automated and was designed to be more of a visual aid to an observer.

Looking at a selected assortment of hyperspectral anomaly detectors, a widely used concept is to examine and compare the statistical properties of the spectral distributions between two classes of pixels: first, the local class of pixels around some center pixel under test (PUT) that has potential to be anomalous and second, an outer class of background pixels. Other detectors compare only the pixel under test to either the statistical properties of the entire image, or to an outer class of background pixels.

A widely accepted method of hyperspectral anomaly detection is the Reed-Xiaoli (RX) algorithm [5], as well as the multiple variants of the algorithm. This normalization method statistically compares each pixel in the image in the form of the mean and covariance relationships between the spectral distributions of either the entire image or a local window surrounding the pixel, excluding a determined center of that window. An alternative set of detectors is the Dual-Window based Eigen Separation Transform (DWEST) [3] and the Dual Spatial Window Target Detector (DSWTD) [6][7] which not only uses a dual window to do locally-based comparisons, but also performs subspace projection on the pixels in order to enhance the statistical differences between the pixels of the two regions of the dual window.

3 Problem Statement

Search and rescue operations could greatly benefit from the use of outlier color detection methods performed on still images captured from cameras mounted on aerial vehicles to automate the process of finding the object of interest. The object of interest is considered to be an unknown color, as well as the background, which is of a color different to that of the object. The object is also assumed to be small compared to the search area. The purpose of our project is to implement anomaly detection algorithms and find the "best" algorithm to use, and what the best color spaces are to incorporate this algorithm. Also, we determine whether some fusion of results between algorithms and color spaces help improve detection accuracy.

4 Solutions

4.1 RX Algorithm

The RX algorithm is based on each pixel's component values being drawn from a multivariate normal distribution [4]. The original algorithm is based on the global mean and global sample covariance matrix of the image. It is represented by the following formula:

$$\delta_{RX}(\boldsymbol{r}) = (\boldsymbol{r} - \boldsymbol{\mu}_g)^T \mathbf{C}_g^{-1} (\boldsymbol{r} - \boldsymbol{\mu}_g) \quad (1)$$

In the above equation C_g represents the global covariance matrix of the image and μ_g represents the global mean. The value calculated for each pixel represents its distance from the center of the multivariate normal distribution.

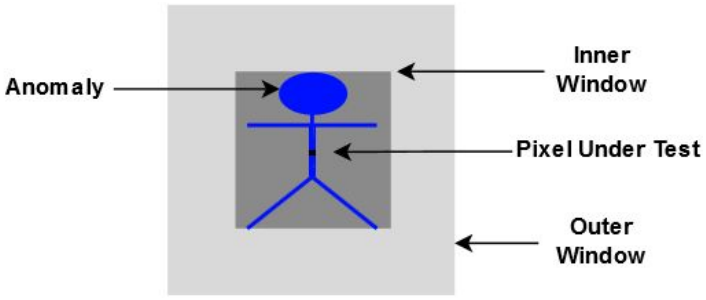

Figure 2: Inner and Outer Window Example

Variations of the RX algorithm have been developed since its conception. A

common alteration is to calculate variance across the entire image and to calculate a mean local to the pixel being evaluated, while another is to use this local mean as well as a local covariance matrix. These local statistics would be computed through two windows around the pixel in order to separate and classify two distinct regions around the PUT: an inner window would include pixels closer to the target pixel, and an outer window would be a larger box around the target pixel. The outer window, however, would exclude the pixels in the inner window. This general procedure can be seen in Figure 2. The size of these windows can highly impact the algorithm's results. The statistics of the outer window are utilized in the RX variants' calculations, so the inner window must be large enough to fit the probable maximum dimensions of the object being searched for; if the anomaly is present in the outer window, then the local statistics would reflect that. Essentially, if the PUT is an anomaly, and part of the anomaly is affecting the local statistics, the PUT would be considered more similar to the local area compared to if the anomaly was not contaminating the outer window and local statistics [1]. The outer window must be large enough to include enough background pixels to be able to compute an invertible covariance matrix for the second variation. Meanwhile, these windows must be small enough to be considered "local" to the pixel being evaluated. These RX variants result in a varying mean across the image, or a varying mean and covariance matrix.

Alterations to the original RX detector result in two variants:

$$\delta_{RX_{var1}}(\boldsymbol{r}) = (\boldsymbol{r} - \boldsymbol{\mu}_{out})^T \mathbf{C}_g^{-1} (\boldsymbol{r} - \boldsymbol{\mu}_{out}) \tag{2}$$

$$\delta_{RX_{var2}}(\boldsymbol{r}) = (\boldsymbol{r} - \boldsymbol{\mu}_{out})^T \mathbf{C}_{out}^{-1} (\boldsymbol{r} - \boldsymbol{\mu}_{out}) \tag{3}$$

The first equation represents the RX algorithm evaluated using a local mean (calculated using the outer window) and a global covariance matrix. The second equation represents the algorithm using fully local statistics: locally calculated means and covariance matrices utilizing the outer window.

These RX algorithms return values at each pixel location, the distance from the center of the multivariate normal distribution. From this point a threshold must be determined which will be used to classify all values below the threshold as non-anomalies and those above the threshold as anomalies.

4.2 DWEST and DSWTD

The next set of detectors assumes a linear mixing model for each pixel, that is, each pixel is a linear combination of spectra of the many materials it may be made up of [4]. The general idea is to find a set of projection vectors to project the pixels of the two regions of the dual window into a subspace that enhances the computation of the statistical differences between the two regions. First, the DWEST uses the inner and outer dual window classification method as in the RX detector. The covariance matrices for both the inner and outer window regions, $\mathbf{C}_{in}$ and $\mathbf{C}_{out}$ respectively, are computed and the difference $\mathbf{C}_{diff}$ is found in the following equation:

$$\boldsymbol{C}_{diff} = \boldsymbol{C}_{in} - \boldsymbol{C}_{out} \tag{4}$$

Eigendecomposition is done on the matrix $\mathbf{C}_{diff}$ and the eigenvectors of only the positive eigenvalues of this difference covariance matrix are found. Given this set of i positive eigenvectors, $\mathbf{P}_i$, the difference of the mean vectors of the inner and outer windows, $\boldsymbol{\mu}_{out}$ - $\boldsymbol{\mu}_{in}$, is projected onto positive eigenvectors $\mathbf{P}_i$ as shown in equation 5 below. This value is then compared to some global threshold where the pixel is classified an outlier if it exceeds the threshold.

$$\delta_{DWEST}(\boldsymbol{r}) = \left| \sum_i \boldsymbol{P}_i^T (\boldsymbol{\mu}_{out} - \boldsymbol{\mu}_{in}) \right| \tag{5}$$

Second, the DSWTD is a variant of the DWEST that uses orthogonal projection instead. The detector result is determined by computing a measure called the orthogonal projection divergence (OPD) as defined below which can be interpreted as the distance between two orthogonal projections. This OPD measure for two windows is calculated from the means of the inner and outer window where $\mathbf{s}_i$ and $\mathbf{s}_j$ in the definition is $\boldsymbol{\mu}_{in}$ and $\boldsymbol{\mu}_{out}$.

$$\delta_{DSWTD}(r) = OPD(\mu_{in}, \mu_{out}) \tag{6}$$

$$OPD(s_i, s_j) = (s_i^T P_{s_j}^{\perp} s_i + s_j^T P_{s_i}^{\perp} s_j)^{\frac{1}{2}} \tag{7}$$

$$P_s^{\perp} = I - s(s^T s)^{-1} s^T \tag{8}$$

In essence, both the DWEST and DSWTD try to compute a measure of the difference of the means of the inner and outer regions with the DWEST finding eigenvalues and eigenvectors while the DSWTD calculates the OPD.

4.3 Color

The above five algorithms are each evaluated in five color spaces RGB, LAB, YCbCr, XYZ and HSV.

4.4 Intersection of Results

Upon determining the best detectors in the best color spaces, two methods are attempted to evaluate the intersection of the results, data fusion and decision fusion. Through data fusion, the most accurate outlier detection methods from above are identified, and the resultant matrices of the algorithms (pre-thresholding to determine outliers) are scaled by one another through element-wise multiplication, to determine if the scaling of results would improve detector accuracy. Decision fusion is used on the best detectors via logical AND and OR operators performed on the thresholded results of the best detectors to determine if either logical operation on the results would increase the number of true positives detected or decrease the number of false negatives.

5 Implementation

In order to evaluate the algorithms, images were required. There is no known data set of images, so three of the four test images needed anomalies to be synthetically introduced. The first image found is shown below in Figure 3 [9]. Artificial anomalies were placed in the image of the shore.

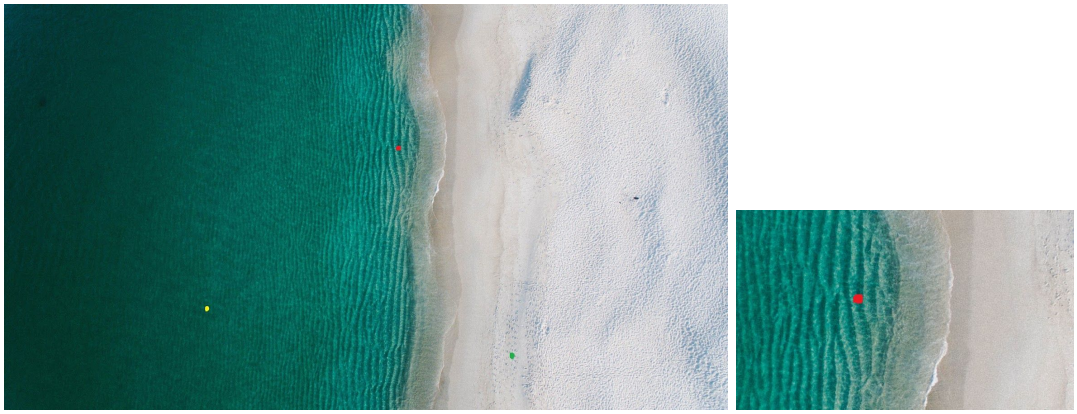

(a) *(b)*

Figure 3: 3992x2242 Image of 'Shore' with Anomalies (a) Full Image (b) Anomaly Zoom

The second image evaluated was found online and does not contain any synthetically introduced anomaly; it is of a wooded area and contains a blue body anomaly. It is shown in Figure 4 [11].

(a) *(b)*

Figure 4: 2562x1920 Image of 'Woods' with Anomaly (a) Full Image (b) Anomaly Zoom

Figures 5 and 6 below show two images shared by another team researching outlier color detection for search and rescue. Figure 5 shows a desert with three artificial anomalies introduced, and Figure 6 shows a mountain with three artificial anomalies.

Figure 5: 1200x800 Image of 'Desert' with Three Anomalies

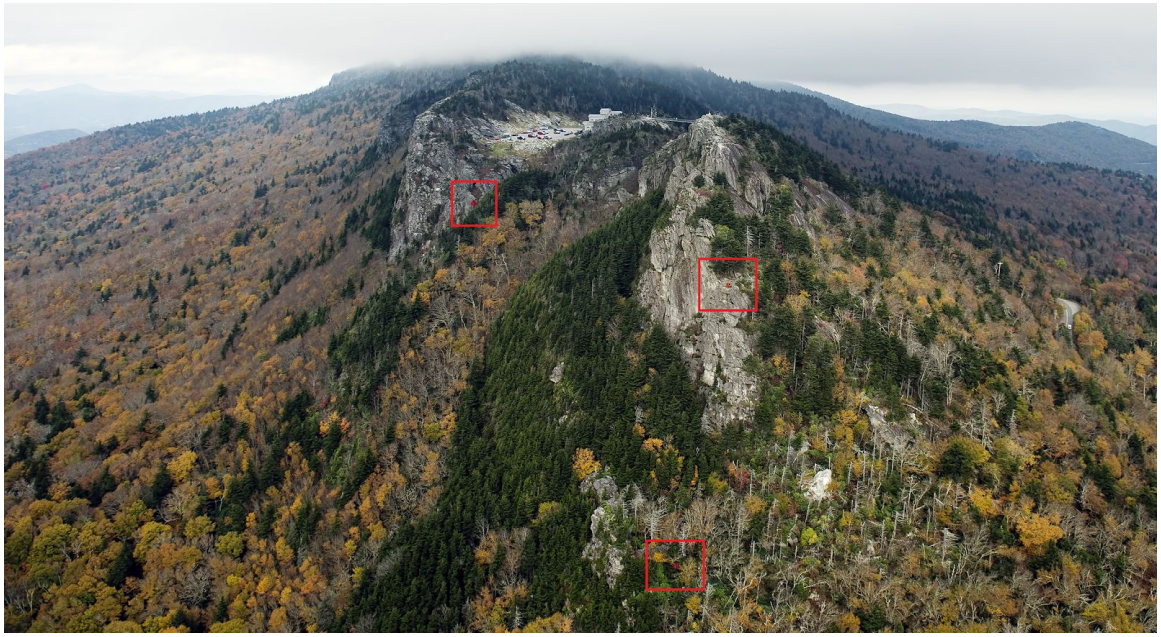

Figure 6: 2500x1368 Image of 'Mountain' with Anomalies Specified

Save for the original RX algorithm, the four other algorithms require inner and outer window usage to calculate statistics local to the pixel of interest. As was discussed previously, the sizing of these windows is very important for the success of these

algorithms. It was determined that an ideal window size can be used for this application as opposed to random sizes, due to the ability of a search and rescue team to calculate the expected pixel size of a target, and to size the window appropriately. An expected pixel dimension can be calculated using the following formula [11]:

$$height_{obj}(pixels) = \frac{flength(mm) \times height_{obj}(mm) \times height_{img}(pixels)}{dist(mm) \times height_{sensor}(mm)} \quad (9)$$

Where "flength" is the focal length of the camera and "dist" is the distance from the object. Using knowledge of the camera used to capture the imagery and equipment such as an altitude meter, the expected pixel height of the object can be obtained and inner and outer windows can be sized appropriately.

MATLAB was chosen as the software platform to evaluate the anomaly detection algorithms. In order to successfully create the required windows in the sides and corners of the full image, a buffer of NaN values is placed around the read-in image matrix which are disregarded by Matlab functions such as nanmean and nancov. This method allows the other pixels in the area to be considered as included in the windows.

6 Experimental Results

6.1 Observational Results

Figure 7 below shows the results of the RX algorithm using a local mean in the RGB color space utilizing varying thresholds.

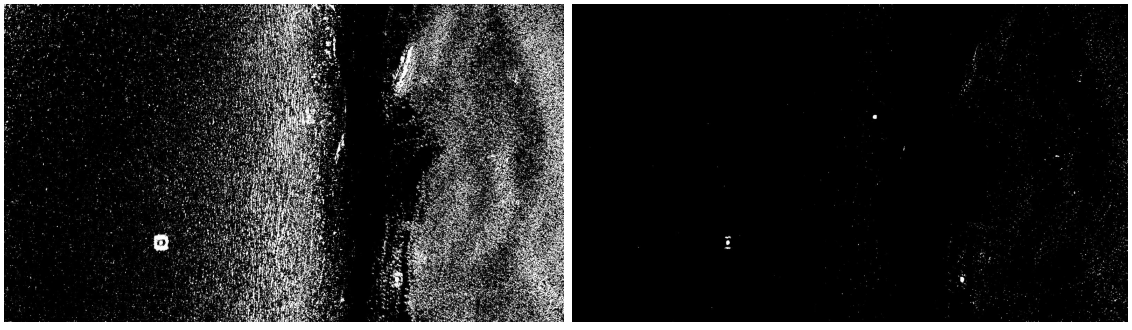

(a) *(b)*

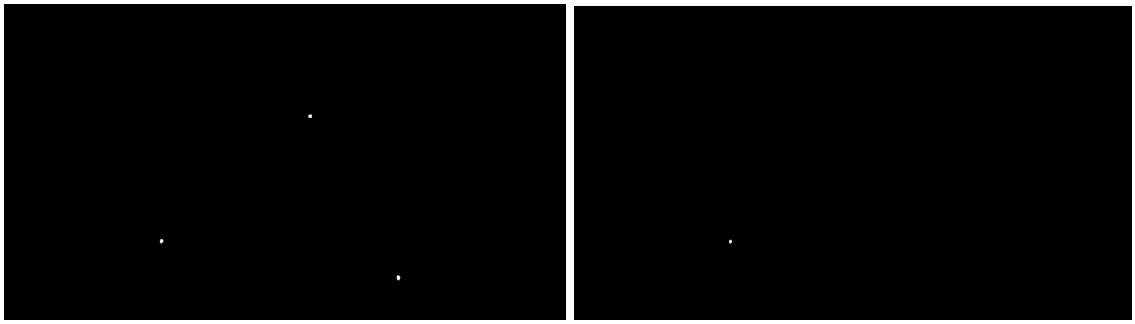

(c) *(d)*

Figure 7: Results of the RGB Local Mean RX on 'Shore' Image with Varying Thresholds of (a) 1 (b) 10 (c) 100 and (d) 1000

It is apparent from this figure that the chosen threshold is extremely important in determining the pixels which are classified as anomalies. Of the thresholds shown above the one that produces the fewest false positives and still detects all three anomalies is (c) which displays the result with a threshold of 100. Each of the five algorithms in the five color spaces produce images such as this when thresholded, displaying the detected locations of anomalies.

Another example is shown in Figure 8, below, of how an incorrect threshold could cause detection of the many bushes in the 'desert' image, and how a more accurate threshold would detect the color anomalies much more accurately.

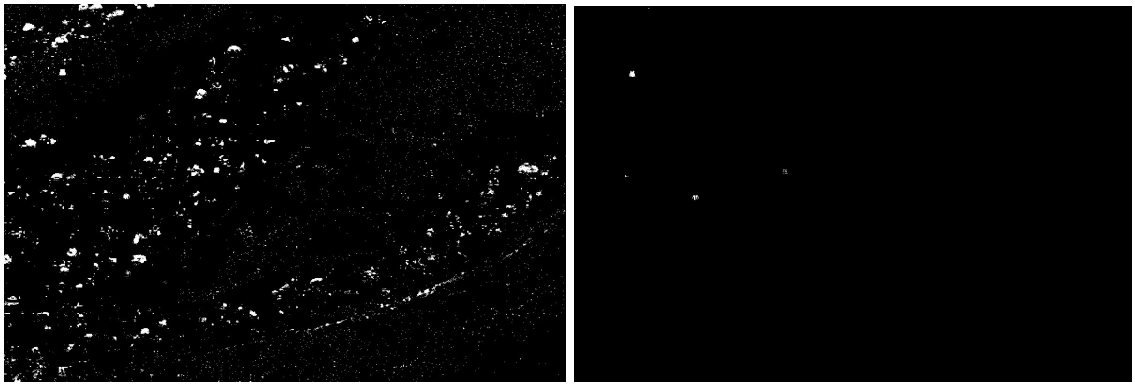

(a) *(b)*

Figure 8: Results of the RGB Local Mean RX on 'Desert' Image with Thresholds of (a) 10 and (b) 80

6.2 Receiver Operating Characteristic Curve Results

To quantify our results, a receiver operating characteristic curve (ROC) is

modeled for each of the detectors and the area under this curve (AUROC) is calculated. The ROC curve highlights the tradeoff between the sensitivity and specificity where sensitivity is defined as the true positive rate (TPR) while specificity is defined as 1 minus the false positive rate (1-FPR). The TPR is the proportion of pixels that are true anomalies that are successfully detected as true while the FPR is the proportion of pixels that are not anomalies but are successfully detected as such. The ROC curve itself is determined by using many thresholds and calculating the FPR and TPR results of using that threshold; the resultant point with FPR measured on the x-axis and TPR measured on the y-axis is just one of many calculated similarly on the ROC curve. Ultimately, the measure we are interested in is the area under this curve: the AUROC. The AUROC can be seen as a sort of measure for accuracy since it is defined in our case as the probability that the detector gives a value higher to some randomly selected pixel that is a true anomaly than it gives to some randomly selected pixel that is not an anomaly.

The ROC and AUROC results are measured in two ways: first we find the values at a per image basis to analyze the varying difficulty of determining anomalies in each image, and second we combine a number of the images to compute the values at a per detector basis to determine the best-performing detector.

Tables 1, 2, 3 and 4 below show the AUROC results of the algorithms in varying color spaces for each individual image. Figure 9 shows the same data in chart form.

"Shore" AUROC Results						
		Color Space				
		RGB	LAB	YCbCr	HSV	XYZ
Detector	RX: Global μ Global C	1.0000	1.0000	1.0000	0.9460	0.9999
	RX: Local μ Global C	1.0000	1.0000	1.0000	0.9942	1.0000
	RX: Local μ Local C	0.9999	0.9999	0.9999	0.9998	0.9998
	DWEST	1.0000	0.9999	0.9999	0.9986	0.9984
	DSWTD	0.9999	0.9999	0.9956	0.9981	0.9998

Table 1: 'Shore' Image AUROC Values

"Woods" AUROC Results						
		Color Space				
		RGB	LAB	YCbCr	HSV	XYZ
Detector	RX: Global μ Global C	0.9902	0.9872	0.9902	0.9486	0.9444
	RX: Local μ Global C	0.9988	0.9951	0.9988	0.9843	0.9739
	RX: Local μ Local C	0.9997	0.9996	0.9997	0.9939	0.9869
	DWEST	0.6869	0.7696	0.7199	0.6652	0.4898
	DSWTD	0.7698	0.6666	0.3963	0.4406	0.7591

Table 2: 'Woods' Image AUROC Values

"Desert" AUROC Results						
		Color Space				
		RGB	LAB	YCbCr	HSV	XYZ
Detector	RX: Global μ Global C	0.9778	0.9777	0.9778	0.9606	0.9193
	RX: Local μ Global C	0.9869	0.9864	0.9869	0.9707	0.9487
	RX: Local μ Local C	0.9661	0.9865	0.9861	0.9645	0.9423
	DWEST	0.7857	0.8898	0.8062	0.8649	0.6895
	DSWTD	0.9674	0.9733	0.7327	0.7294	0.7992

Table 3: 'Desert' Image Auroc Values

"Mountain" AUROC Results						
		Color Space				
		RGB	LAB	YCbCr	HSV	XYZ
Detector	RX: Global μ Global C	0.9836	0.9825	0.9836	0.9474	0.9473
	RX: Local μ Global C	0.9938	0.9936	0.9938	0.9750	0.9814
	RX: Local μ Local C	0.9885	0.9890	0.9885	0.9846	0.9669
	DWEST	0.8254	0.9088	0.8577	0.9299	0.7799
	DSWTD	0.9777	0.9871	0.8173	0.7030	0.8824

Table 4: 'Mountain' Image Auroc Values

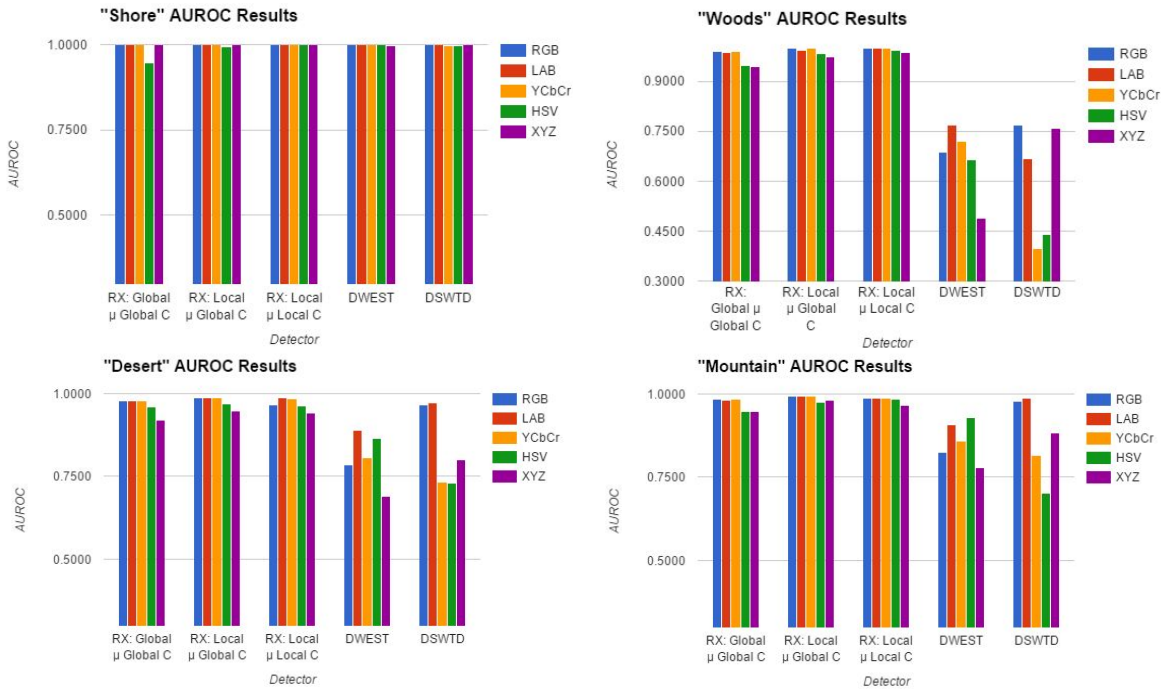

Figure 9: Detection Results for Individual Images

As mentioned, these results are shown to display the varying difficulties of the images under test. AUROC values for the 'Shore' picture indicate the detector is incredibly accurate (in fact, near-perfect) when dealing with an artificially exaggerated scenario of solid, bright, and contrasting colors. The anomalies introduced are very apparent in the picture. Meanwhile, the other three images analyzed produce much more relevant and interesting results, as the images are more complex and the outliers are more difficult to encounter. Therefore, only the 'Woods,' 'Desert' and 'Mountain' images are included in the subsequent analysis of detector accuracy as they represent slightly more realistic scenarios.

The ROC curves for each detector in each color space are plotted below.

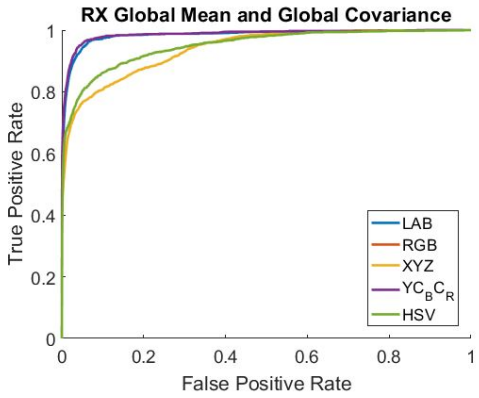

Figure 10: ROC Curve for Global RX

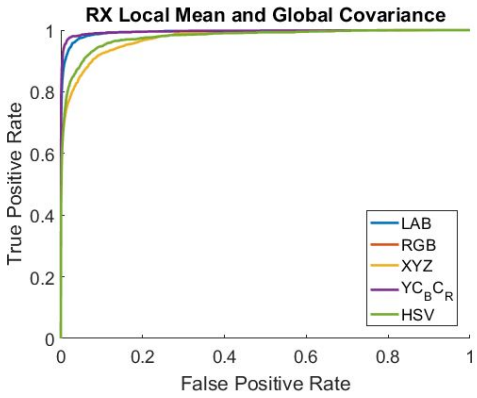

Figure 11: ROC Curve for Local Mean RX

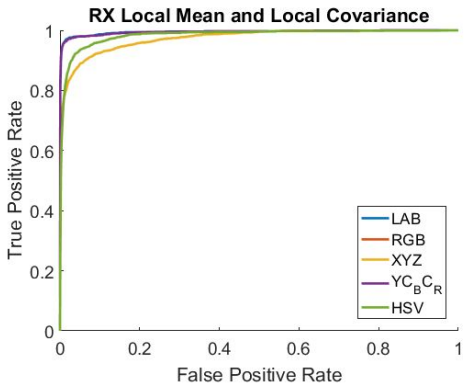

Figure 12: ROC Curve for Local RX

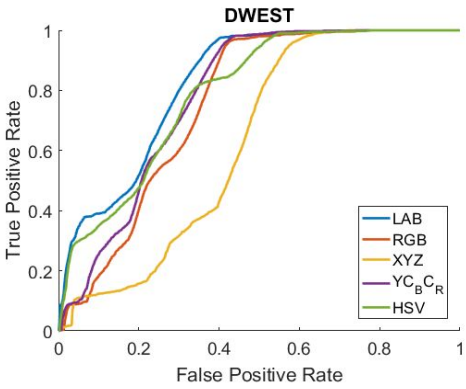

Figure 13: ROC Curve for DWEST

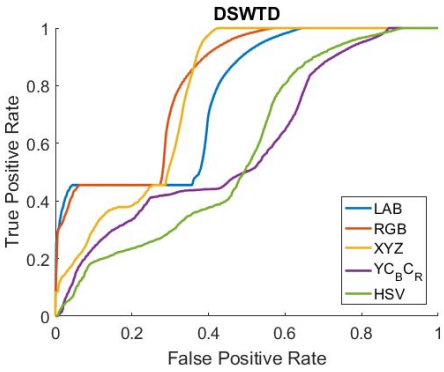

Figure 14: ROC for DSWTD

Below in Table 5 are the AUROC results for each detector in each color space, evaluated over the determined three images. The results are also displayed in Figure 15.

Detector & Color Space AUROC Values						
		Colorspace				
		RGB	LAB	YCbCr	HSV	XYZ
Detector	RX: Global μ Global C	0.9865	0.9844	0.9865	0.9493	0.9391
	RX: Local μ Global C	0.9953	0.9935	0.9953	0.9773	0.9740
	RX: Local μ Local C	0.9942	0.9944	0.9942	0.9836	0.9741
	DWEST	0.7565	0.8258	0.7835	0.7913	0.6290
	DSWTD	0.8119	0.7605	0.5949	0.5790	0.7778

Table 5: Detector & Color Space AUROC Results

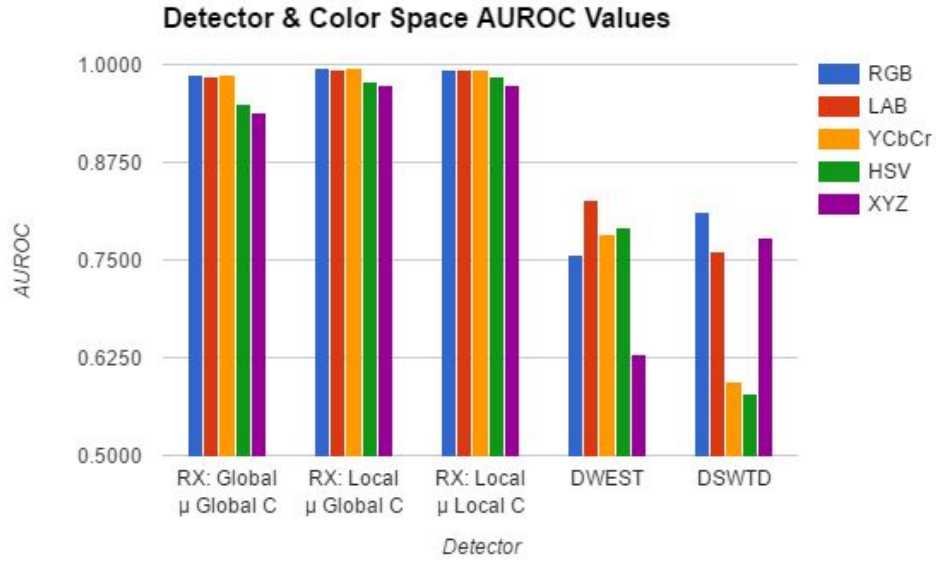

Figure 15: Detector & Color Space AUROC Results

It is quite apparent from the ROC curves that there is a significant difference between the three versions of the RX algorithm and the two subspace projection algorithms DWEST/DSWTD. This is reflected in the AUROC values specified in the table and shown in the chart in Figure 15. All variants of the RX algorithm in all color spaces had an AUROC above 0.92, while the best subspace projection-based detector had an AUROC lower than 0.83. In addition, it was observed that the subspace projection-based detectors were consistently much slower than the RX algorithm variants because of the general increase in calculative intensity and the required eigendecomposition.

The RX algorithm utilizing a global covariance matrix and a varying local mean provided consistently the best compromise for computational efficiency and quality results. The results are very comparable to those of the RX algorithm using local

covariance matrices, however utilizing a local covariance matrix resulted in a quite substantial increase in computation time. The original RX algorithm using global statistics also produced good results, and while using a local mean provides better results and a higher AUROC, this global method does provide an AUROC above 0.98 and is very quick to compute, so in situations where this slight tradeoff is acceptable and a faster computation time is desired, this global RX algorithm would be the best to use.

There is a variation of results across color spaces tested. The HSV and XYZ color spaces provide consistently poor results. The XYZ color space provides alternate R, G and B primaries within the CIE XYZ chromaticity diagram, thus reducing the variance compared to the color space in which the anomaly was introduced, so worse results are unsurprising. The HSV color space operates using hue, saturation and value (brightness). Generally, an object of interest in a search and rescue operation is expected to be of a different color than the background terrain, which is what the anomalies in the images tested reflected. This HSV color space reduced the main variation in color to hue, solely one of the three components of the color space, which caused poorer results than utilizing color spaces that focus more on color variation across the three descriptive components.

Meanwhile, the variations of the RX algorithm performed at consistently high levels in the RGB, LAB and YCbCr color spaces. RGB and LAB highly consider the variety of color in their space and it was expected that the algorithms would best detect color in these spaces. The algorithms performing well in the YCbCr color space suggest that there was a significant variation in anomaly vectors and background vectors in luminance and chrominance.

6.3 Fusion Results

To evaluate the intersection of results, data fusion first was used to scale the distance calculated from the center of the multivariate normal distribution by an algorithm in one color space with the distance calculated by the algorithm in another color space. The results for fusing the global and local mean RX algorithms in RGB and LAB color spaces are shown below in Table 6 and Figure 16.

RGB-LAB Data Fusion ROC Values				
		Color Space		
		RGB	LAB	RGB-LAB Fused
Detector	RX: Global μ Global C	0.9865	0.9844	0.9857
	RX: Local μ Global C	0.9953	0.9935	0.9951

Table 6: Data Fusion of RGB and LAB Spaces

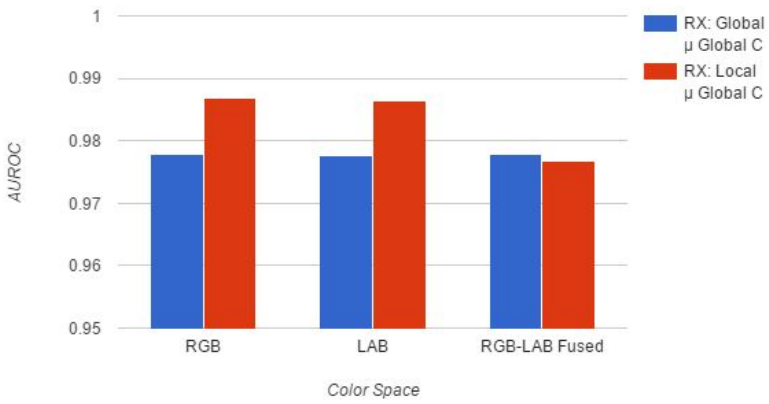

Figure 16: Data Fusion of RGB and LAB Color Spaces

All color space data fusions tested returned similar results to the above: the fusion of color spaces resulted in either a value between the original AUROC values in both color spaces, or it resulted in poorer results than the original color spaces. It was determined that the shorter distance from the center of the multivariate distribution calculated for the same pixel in one color space as another had a more negative effect on the results than the larger distance calculated by one of the algorithms.

Although the DWEST performed much worse than all variants of the RX algorithm, their fusion in the LAB color space was tested and the results are shown below in Table 7.

Detector AUROC Values		
Detector / Fused-Detector		
DWEST	RX Local μ	Dwest-RX
0.8258	0.9935	0.9875

Table 7: Data Fusion of Local Mean RX and DWEST

Fusing the poor results of the DWEST algorithm with the local mean RX resulted in a large increase in the performance compared to solely the DWEST, however this did not

overcome the performance of using solely the RX algorithm.

To perform decision fusion, the intersection of results post-ideal thresholding between algorithms was considered. Logical AND and OR operations were performed on the results of the RX algorithm using local mean in RGB color space and LAB color space across the three more realistic images. These results are compared against the individual color spaces below in Table 8 in the form of confusion matrices.

RGB RX Local μ		
	Pred No	Pred Yes
True No	9297131	157
True Yes	896	856

(a)

LAB RX Local μ		
	Pred No	Pred Yes
True No	9297171	117
True Yes	974	778

(b)

RGB & LAB RX Local μ		
	Pred No	Pred Yes
True No	9297183	105
True Yes	994	758

(c)

RGB | LAB RX Local μ		
	Pred No	Pred Yes
True No	9297119	169
True Yes	876	876

(d)

Table 8: Confusion Matrices (a) RGB RX Local Mean (b) LAB RX Local Mean (C) Decision Fusion RGB&LAB RX Local Mean (d) Decision Fusion RGB|LAB RX Local Mean

As the total number of predicted and true “no” values is magnitudes larger than the other components of these matrixes, accuracy is a poor measure of the results because it will be overly influenced by the correct predictions of true negatives.

As logic would lead one to expect, the AND operation results in values shifted from the “predicted yes” column to the “predicted no” column. This shift caused a slight decrease in false positives, however also pushed true positive results to the “predicted no” column. The total true positives as a result of the AND operation was lesser than the individual results of the original algorithms and didn’t significantly reduce the false positive rate.

The OR operation shifted results from predicted “nos” to predicted “yes’s”. This shift resulted in a higher number of true positives than either of the original algorithms,

but also resulted in a larger increase in the number of false positives. Overall, the results of the AND or OR decision fusion would not be worth the increase in computational requirement to evaluate the algorithms in multiple color spaces.

7 Conclusions

In conclusion, the applications of hyperspectral anomaly detectors to outlier color detection in search and rescue applications are quite evident and useful. The RX algorithm using local mean computed in RGB, LAB and YCbCr color spaces produced the best results consistently with consideration for computation speed. However, if fast computation is truly necessary, or computation tools are limited, the global RX algorithm in the same color spaces provides quality results as well. Fusion of data and decisions between algorithms/color spaces does not provide enough increase in results quality to justify the increase in computational requirement in evaluating the captured images in multiple color spaces or algorithms.

Image quality tested is very important in getting accurate results of algorithm performance. The original “shore” image showed quality results for algorithms and color spaces that were later proved to be poor anomaly detectors in more difficult terrain and outlier situations. Evaluation of these algorithms on more difficult and realistic images is important for further development in this area. Also, varying the window sizes in algorithms using local statistics would be a further area of research, as the window sizes utilized here were somewhat ideal under the assumption that the expected size of the anomaly could be determined. If a search and rescue operation anticipated this to be a more difficult task, evaluating the results on varying window sizes could prove to be beneficial. Finally, a further possible improvement would be the implementation of a Markov Random Field on the results of the RX algorithm. This would potentially decrease the number of scattered false positives detected in the image, and increase the number of true positives in the area where the anomaly exists.

8 References

[1] Matteoli, S., Diani, M., Corsini, G., “A Tutorial Overview of Anomaly Detection in Hyperspectral Images” *IEEE A&E Systems Magazine*. Vol. 25 No. 7, July 2010.

[2] N. Rasmussen, D. Thornton, and B. Morse, "Enhancement of unusual color in aerial video sequences for assisting wilderness search and rescue," in Proc. IEEE ICIP, Oct. 2008, pp. 1356-1359.

[3] A. F. El-Rewainy, E. E. Farouk, and A. E. Fouda, "A comparison study of dual window-based anomaly detection algorithms for hyperspectral imagery," Proc. 13th International Conference on Aerospace Sciences & Aviation Technology, 2009.

[4] David W. J. Stein, Scott G. Beaven, Lawrence E. Hoff, Edwin M. Winter, Alan P. Schaum, and Alan D. Stocker, “Anomaly Detection from Hyperspectral Imagery” IEEE *Signal Processing Magazine* January 2002.

[5] I. S. Reed and X. Yu, “Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 38, pp. 1760-1770, Oct. 1990.

[6] H. Kwon, S. Z. Der, and N. M. Nasrabadi, “Adaptive anomaly detection using subspace separation for hyperspectral imagery,” Opt. Eng., vol. 42, no. 11, pp. 3342–3351, Nov. 2003.

[7] L. Wei-Min, I. C. Chein, "Multiple-window anomaly detection for hyperspectral imagery", *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 6, no. 2, pp. 644-658, Apr. 2013.

[8] S. Matteoli, M. Diani, G. Corsini, "A tutorial overview of anomaly detection in hyperspectral images", *IEEE Aerosp. Electron. Syst. Mag.*, vol. 25, no. 7, pp. 5-27, Jul. 2010.

[9] Matthew Kane Photo Collection
Retrieved from https://unsplash.com/photos/9EM7s13H2I0.

[10] Search and Rescue Applications for the UAV, Victoria Air Photos and Survey
Retrieved from http://victoriaairphotos.com/search_and_rescue.html.

[11] Stack Exchange Photography: How do I calculate the distance of an object in a photo? Retrieved from https://photo.stackexchange.com/questions/12434/how-do-i-calculate-the-distance-of-an-object-in-a-photo