COLOR OUTLIER DETECTION FOR SEARCH AND RESCUE

Travis Marshall and L. Nathan Perkins

May 1, 2015

Boston University

Department of Electrical and Computer Engineering

Technical Report No. ECE-2015-01

BOSTON
UNIVERSITY

COLOR OUTLIER DETECTION FOR SEARCH AND RESCUE

Travis Marshall and L. Nathan Perkins

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

May 1, 2015

Technical Report No. ECE-2015-01

Summary

Unmanned aerial vehicles (UAV) are a tremendous resource for search and rescue missions, as they are able to collect high resolution imagery of large swaths of land without the cost and manpower of human search parties. With the tremendous amount of imagery that can be collected, a new bottleneck in the process has emerged as the task of evaluating and identifying objects of interest, including people, clothing or other details in these images is still predominantly a manual task. Automated color anomaly detection can significantly expedite this process by identifying unexpected colors likely associated with synthetic clothing or equipment.

Hyperspectral imaging, used in satellite and geographic exploration, has a strange set of tools for anomaly detection. In this paper, we evaluate the applicability of these hyperspectral techniques to search and rescue, as well as introduce a novel algorithm that uses gaps in the principal component analysis to identify clusters of anomalous pixels. In addition, we compare performance of algorithms across multiple color spaces and across multiple types of environments to identify potential strengths and weaknesses.

Existing algorithms offer a clear tradeoff between computational performance and anomaly detection accuracy. Our proposed algorithm strikes a balance between these tradeoffs, achieving high accuracy with relatively low computational intensity.

Contents

List of Figures

List of Tables

1 Introduction

As the prevalence of unmanned aerial vehicles (UAV) increases in both search and rescue and military missions, so too does the volume of imagery collected by the onboard cameras. Human analysis of such imagery is both unreliable [1], given the need to broadly analyze large fields of view, and impractical, given the number of images produced. Instead, algorithmic analysis of images is employed to quickly identify regions of interest in the imagery.

Given the unknowns regarding the visual appearance of the target (e.g., an individual in a search and rescue mission), the automatic processing must identify anomalous regions in the image rather than implementing a more directed target detection approach. Exploiting color information is especially well suited for identifying outlying pixels in a large natural scene, as clothing and man-made objects will often have distinguishing spectral signatures. Identifying such pixels with anomalous colors can aid in both identifying the target or important information that may lead to the target of a search and rescue mission. In this project, we aim to use spatial and spectral data from a photograph to classify anomalies – that is, local pixels in which the color patterns deviate from expectations identified in the global scene.

2 Literature review

Most of the research in the area of anomaly detection has focused on hyperspectral images (used in medical and geological imaging) [2]. The Reed-Xiaoli (RX) detector, which is broadly used and often considered a benchmark [3], compares individual pixels to either a surrounding window or global statistics under the assumption that the spectral data are distributed according to a multivariate Gaussian [4]. There are many variations and adaptations of the RX detector based on the windowing, allowing better adaptation throughout the image, such as local RX and quasi-local RX [5].

Another approach to anomaly detection is the dual window eigen-separation transform (DWEST) [6], which compares spectral projection matrices for an inner and outer window to identify distinctive spectral statistics. Similar to this windowing approach, the nested spatial window-based target detection (NSWTD) [7] and the multiple-window anomaly detection (MWAD) [8] use multiple windows of comparison to achieve background suppression or target whitening relative to varying contextual scales, helping identify potential anomalies of variable size within the image.

Alternatively, the cluster-based anomaly detection (CBAD) approach segments images into regions that maximize the conformation of the spectral distribution to a Gaussian model [9]. Anomalies can then be identified based on values that deviate from the distribution of the surrounding cluster.

In the limited literature that specifically addresses search and rescue, approaches include increasing the saturation of rare hues and decreasing the saturation for common hues in order to assist human users in identifying anomalous colors [10]. More

recent work has surveyed adapting some of the hyperspectral techniques, and found that many of the detectors described above can be successfully implemented for search and rescue, but the RX algorithm proved most robust and least dependent on parameter tuning [1].

Finally, in a previous EC520 project, students extended the above techniques, which predominantly identify point anomalies, by incorporating spatial information via Markov random fields [11]. Their results show an improvement when compared to the RX detector approach. Yet their model was limited: it simply focused on color (ignoring luminance) and had a high rate of false positives in photos that include sharp transitions from one environment to another (ocean and shore, for example).

3 Problem statement

Due to the prevalence of UAV photography in search and rescue applications, there is a need to improve analysis of wide angle aerial perspective high resolution imagery to find clues that can lead to finding missing persons. Improvement in analysis of this type of imagery in search and rescue applications will be measured by being able to successfully identify anomalous objects/pixels within natural aerial perspective scenes while (a) maximizing the area under the ROC curve for anomaly detection by implementing algorithms on images with known anomalous pixels and (b) minimizing processing time for the image analysis algorithms.

In section 4, we propose our solution to this problem. Subsections 4.1- 4.4 first survey existing anomaly detection algorithms developed for hyperspectral applications. We then then introduce a new solution, PCAG, in subsection 4.5. In subsection 4.7, we describe the color spaces are considered across all algorithms to prevent biases against algorithms in one or more color spaces and to identify a color space most suitable to search and rescue. Section 5 discusses further details of implementation and evaluation across algorithms and color spaces. Section 6 presents and explains the experimental results across algorithms, color spaces, scene types and algorithm processing time performance to determine comparison data for (a) and (b) and evaluate our solution against existing solutions for (a) and (b). Lastly, conclusions and future work recommendations are in section 7.

4 Solution

In order to improve detection of anomalous colors in search and rescue photos, we survey a range of the hyperspectral anomaly detection algorithms, including some of the most recent developments that achieve high accuracy through nested, overlapping windows of analysis. In addition to simply implementing and comparing existing algorithms, we compare algorithm performance for a range of color spaces (including L*a*b, XYZ, YCbCr) as well as subsets of color spaces (such as the “a” and “b” channels of L*a*b) to identify an analysis space that is most productive to search

and rescue applications.

In addition to implementing existing algorithms, we propose a potential algorithm called "principal component analysis gaps" (PCAG) that looks at gaps between clusters along the vector of greatest variability for individual non-overlapping windows.

Each of the considered anomaly detection algorithms is described below.

4.1 RX

The benchmark in hyperspectral anomaly detection is the RX algorithm, developed by Reed and Yu [4]. The detector takes the following form:

$$\delta_{\mathrm{RX}}(\mathbf{u}[\mathbf{x}]) = (\mathbf{u}[\mathbf{x}] - \mu)^T \mathbf{K}^{-1} (\mathbf{u}[\mathbf{x}] - \mu) \tag{1}$$

Where μ is the global sample mean and $\mathbf{K}$ is the global sample covariance matrix (across spectra, or in this case, color components). This equation is equivalent to the square of the Mahalanobis distance (a measure of distance between a single point and a distribution). In its application to anomaly detection, it can be thought of as a whitening process that suppresses the background.

We implement two versions of the RX algorithm. **Global RX** calculates the global mean and global covariance for the full image, and applies equation 1 based on these global measures.

The **Local RX** evaluates each pixel individually, calculating a local sample mean and sample variance based on an outer window of pixels excluding a guard window around the pixel being evaluated. The windows used in calculating the local RX algorithm are illustrated in figure 1.

4.2 DWEST

The dual window-based eigen separation transform (DWEST) is an adaptive algorithm that, like the RX detector, uses covariance to compare material differences between an inner and outer window [6]. Unlike the RX detector, the differential covariance matrix is calculated $\mathbf{K}_{\mathrm{diff}} = \mathbf{K}_{\mathrm{inner}} - \mathbf{K}_{\mathrm{outer}}$, where $\mathbf{K}_{\mathrm{inner}}$ is the covariance matrix for the inner window and $\mathbf{K}_{\mathrm{outer}}$ is the covariance matrix for the outer window. A small subset of the eigenvalues of $\mathbf{K}_{\mathrm{diff}}$ will have positive values, and the corresponding eigenvectors ($\mathbf{V}_i$). These eigenvectors correspond with distinctive, differential data structures and can be used to project the mean difference between the windows ($\mu_{\mathrm{diff}} = \mu_{\mathrm{outer}} - \mu_{\mathrm{inner}}$); this, in effect, achieves adaptive whitening by suppressing background patterns from the outer window.

As an anomaly detector, this is written:

$$\delta_{\mathrm{DWEST}}(\mathbf{u}[\mathbf{x}]) = \left| \sum_i \mathbf{v}_i^T \mu_{\mathrm{diff}}[\mathbf{x}] \right| \tag{2}$$

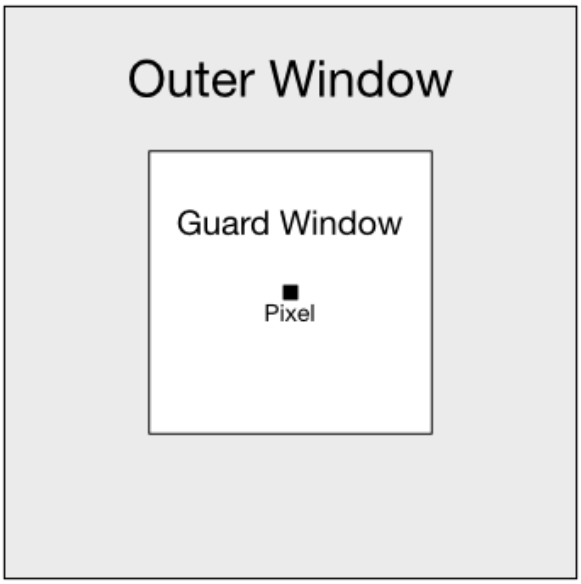

Figure 1: Windowing used for the Local RX algorithm. Sample mean and covariance are based on the outer window and used to evaluate the center pixel.

Where $\mathbf{v}_i$ are the eigenvectors corresponding with positive eigenvalues of $\mathbf{K}_{\text{diff}}$.

We specifically implement a multi-window version of the DWEST algorithm [8], where individual detectors are calculated according to equation 2 for multiple inner window sizes. Figure 2 illustrates a range of inner window sizes. The final detector represents the maximum of these individual detectors:

$$\delta_{\text{MW-DWEST}} = \max_{i=1,\ldots,N} \left\{ \delta^{(i)}_{\text{DWEST}}(\mathbf{u}[\mathbf{x}]) \right\} \tag{3}$$

Specifically, our implementation uses an outer window of 11×11 and inner windows of 1×1, 3×3, 5×5 and 7×7.

4.3 NSWTD

A relatively new hyperspectral algorithm is the nested spatial window-based target detection (NSWTD), which builds on ideas inherent in the DWEST detector [7]. The technique employees the orthogonal subspace projection [12]:

$$\mathbf{P}^{\perp}_{\mathbf{s}} = \mathbf{I} - \mathbf{s}(\mathbf{s}^T\mathbf{s})^{-1}\mathbf{s}^T \tag{4}$$

The orthogonal subspace projection matrix allows projecting a vector into a new space that optimally minimizes the influence of the background signal $\mathbf{s}$. This can be seen by considering the image $\mathbf{u}[\mathbf{x}]$ as a linear combination of desired signal $\mathbf{v}[\mathbf{x}]$ (the anomaly) and background $\mathbf{s}[\mathbf{x}]$.

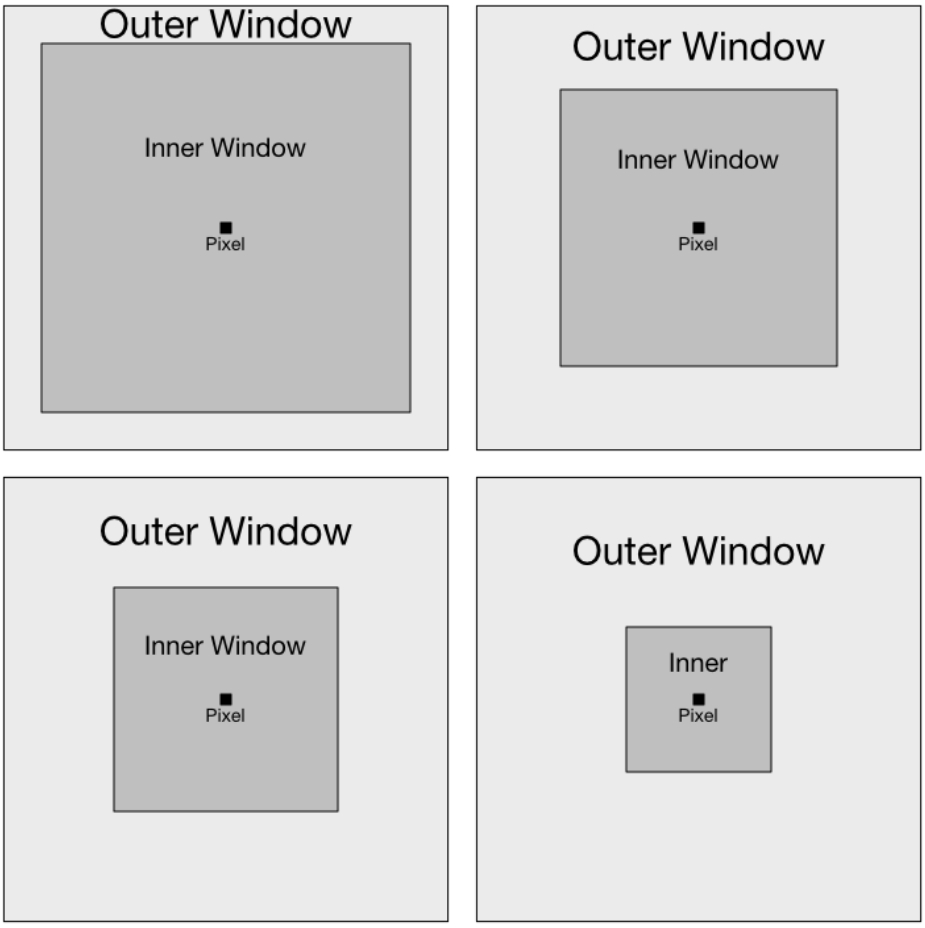

Figure 2: Windowing used for the DWEST algorithm. The maximum anomaly detector is calculated over multiple sets of inner and outer windows.

$$\mathbf{u}[\mathbf{x}] = \mathbf{v}[\mathbf{x}] + \mathbf{s}[\mathbf{x}] \tag{5}$$
$$\mathbf{P}_{\mathbf{s}}^{\perp}\mathbf{u} = \mathbf{P}_{\mathbf{s}}^{\perp}\mathbf{v} + \mathbf{P}_{\mathbf{s}}^{\perp}\mathbf{s} \tag{6}$$
$$= \mathbf{P}_{\mathbf{s}}^{\perp}\mathbf{v} + (\mathbf{I} - \mathbf{s}(\mathbf{s}^T\mathbf{s})^{-1}\mathbf{s}^T)\mathbf{s} \tag{7}$$
$$= \mathbf{P}_{\mathbf{s}}^{\perp}\mathbf{v} + \mathbf{s} - \mathbf{s} \tag{8}$$
$$= \mathbf{P}_{\mathbf{s}}^{\perp}\mathbf{v} \tag{9}$$

The orthogonal subpsace projections are used in calculating a measure called the orthogonal projection divergence, which measures the distance between two orthogonal projections:

$$\mathrm{OPD}(\mathbf{s}_i, \mathbf{s}_j) = \sqrt{\mathbf{s}_i^T P_{\mathbf{s}_j}^{\perp}\mathbf{s}_i + \mathbf{s}_j^T P_{\mathbf{s}_i}^{\perp}\mathbf{s}_j} \tag{10}$$

The OPD measure can be used to compare spectral differences between two windows:

$$\delta_{\text{NSWTD}}(\mathbf{u}[\mathbf{x}]) = \text{OPD}(\mu_{\text{inner}}[\mathbf{x}], \mu_{\text{outer}}[\mathbf{x}]) \tag{11}$$

Again, like the DWEST algorithm, this is specifically implemented via a multi-window approach where individual detectors are calculated according to equation 11 for multiple inner window sizes. The aggregate detector is the maximum of these individual calculations for a given pixel. The same window sizes from the DWEST algorithm are used in this implementation.

4.4 MW-NSWTD

A further variation on the NSWTD algorithm is the multiple window nested window-based target detection (MW-NSWTD) algorithm proposed in [8]. Like NSWTD, this relies on the orthogonal projection divergence, but rather than simply comparing an inner and outer window, uses three windows. The outer most window is used for whitening (background suppression), and the algorithm classifies anomalies based on the divergence in the projected (whitened) versions of the middle and inner windows:

$$\delta_{\text{MW-NSWTD}}(\mathbf{u}[\mathbf{x}]) = \sqrt{\mu_{\text{inner}}[\mathbf{x}]^T P^{\perp}_{\mu_{\text{outer}}[\mathbf{x}]} \mu_{\text{inner}}[\mathbf{x}] + \mu_{\text{middle}}[\mathbf{x}]^T P^{\perp}_{\mu_{\text{outer}}[\mathbf{x}]} \mu_{\text{middle}}[\mathbf{x}]} \tag{12}$$

The equation 12 closely resembles the previously defined function for OPD (equation 10), but uses the orthogonal subspace projection from the outer window to measure the divergence between the inner and middle window.

Again, as with the last two algorithms, the detectors are calculated according to equation 12 for multiple middle and outer window sizes (the inner window remains fixed at one pixel). The aggregate detector is the maximum of these individual detectors. Figure 3 illustrates the multiple versions of the windowing configuration used for the detector.

4.5 PCAG

We introduce and evaluate a new algorithm that we call "principal component analysis gaps" (PCAG). The anomaly detector looks at gaps between clusters along the vector of greatest variability for individual non-overlapping windows. The algorithm works by performing a principle component analysis across spectral components – a widely used step in many anomaly detection algorithms [13] – on block of 15×15 pixels (treated as a 1D vector, $\mathbf{u}[j]$). The PCA yields eigenvectors – basis vectors describing the direction of variability across the components or spectra. The algorithm identifies the largest delta (hence the PCA "Gap" name) along each eigenvector that isolates a potential set of anomalous pixels. Figure 4 shows a hypothetical, two-dimensional

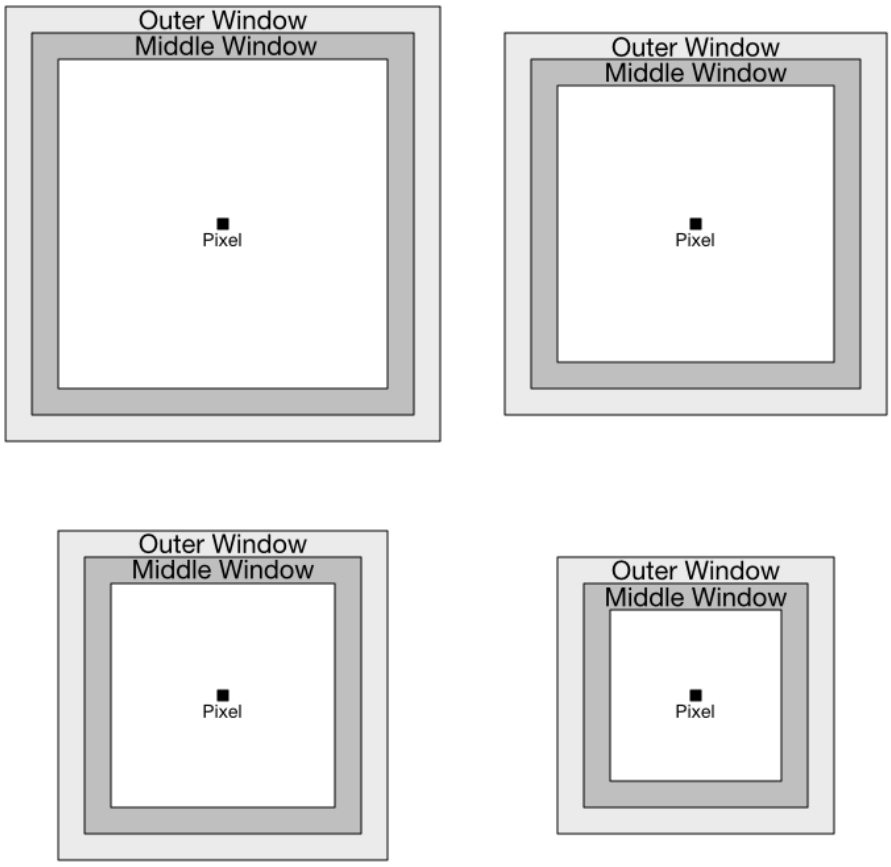

Figure 3: Windowing used for the MW-NSWTD algorithm. Many size outer and middle windows are evaluated, while the inner window remains fixed (1×1).

data set with eigenvectors superimposed, where each access represents one spectral component. As anomalous pixels will inherently be offset from the majority of pixels, there will be gaps along the dominant eigenvectors. The algorithm looks for the largest delta that isolates between $a_{\min}$ pixels and $a_{\max}$ pixels, a computationally inexpensive calculation. The isolated pixels are scored according to the sum of the deltas along each eigenvector.

Where $\mathbf{v}_i$ is an eigenvector of $\text{Cov}(\mathbf{u}[j])$ and a column of the PCA decomposition matrix $\mathbf{W} = \begin{bmatrix}\mathbf{v}_1 & \dots & \mathbf{v}_N\end{bmatrix}$, a sorted list of principle component scores is constructed $a_{i,j} = \mathbf{u}[j]^T\mathbf{v}_i$ such that $a_{i,j} \leq a_{i,j+1}$. From this sorted list, boundaries are calculated:

$$b_i^* = \arg \max_{j \in [a_{\min}, a_{\max}]} a_{i,j+1} - a_{i,j} \tag{13}$$

$$\overline{\Delta}_i = a_{i,b_i^*+1} - a_{i,b_i^*} \tag{14}$$

Equation 13 considers one boundary. Similar calculation must be conducted for a second potential boundary in the range $[N - a_{\max}, N - a_{\min}]$. Based on the identified

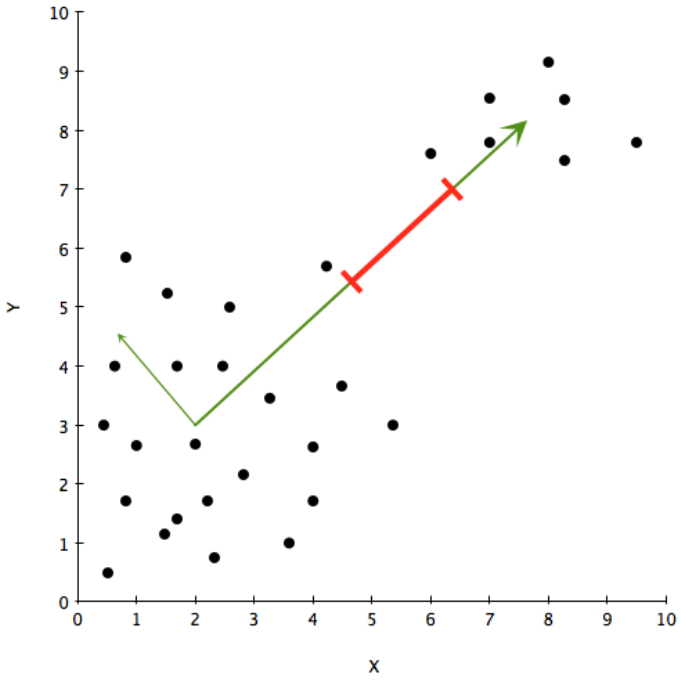

Figure 4: Example of gap identification in a hypothetical two dimensional data set (representing two spectral components). Gaps are identified along each vector.

boundary, pixels on the other side of the boundary will receive an anomaly score represented as the sum of the deltas:

$$\delta_{\text{PCAG}}(\mathbf{u}[\mathbf{x}]) = \begin{cases} \sum_i \overline{\Delta}_i & \text{if } \mathbf{x} \notin \mathcal{W}_{b_i^*} \\ 0 & \text{otherwise} \end{cases} \tag{15}$$

Where $\mathcal{W}_{b_i^*}$ represents those pixels within the boundaries identified by the gap maximization process.

In essence, this detector scores a small subset of pixels based on the gap that separates them from the block background (that is, the majority of the pixels in the block) based on an orthogonal projection achieved through PCA. The two parameters in this algorithm are simply the minimum (a_{min}) and maximum (a_{max}) anomaly sizes in number of pixels. Note that the block must be larger than two times the biggest anomaly.

In comparison with the existing algorithms, this approach is both fast and less dependent on having roughly square anomalies (most hyperspectral algorithms depend on a surrounding background window for comparison), but is less able to handle anomalies in noisy regions (where there is a less clear gap along the vectors of variability).

We also implement a **multi-window** version of the PCAG algorithm (MW-PCAG) that uses multiple iterations with varying window sizes to identify anomalies

of different size. That is, the traditional PCAG algorithm is performed with 11×11 windows to identify anomalies up to 50, 15×15 windows to identify anomalies up to 100 pixels, etc. This allows smaller anomalies to be better identified relative to local context. The maximum score for each pixel is then used as the overall detector output.

4.6 PCAD

In addition to the above algorithm, we propose a variation that we call the "principal component analysis distance" (PCAD). Like the previous algorithm, it relies on projecting blocks – again, 15×15 pixels – into a new space determined by the PCA of the spectral components. Rather than looking at gaps along the principal components, this approach is built on the intuition that the likelihood a pixel is anomalous is proportional to the distance to the k^{th} nearest pixel along each PCA vector; that is, a pixel is much more likely to be anomalous if there is a large distance between it and the k^{th} nearest pixel. Previous anomaly detectors have been built around this approach by measuring the distance to the k^{th} nearest pixel in multidimensional, spectral space using a k-nearest neighbors algorithm [11]. By first projecting the spectral components via PCA, this problem can be simplified and each dimension can be treated independently, eliminating the need for the more computationally intensive k-nearest neighbors computation.

Where $\mathbf{v}_i$ is an eigenvector of $\text{Cov}(\mathbf{u}[j])$ (the spectral covariance) and a column of the PCA decomposition matrix $\mathbf{W} = \begin{bmatrix}\mathbf{v}_1 & \dots & \mathbf{v}_N\end{bmatrix}$, a sorted list of principle component scores is constructed $a_{i,j} = \mathbf{u}[j]^T\mathbf{v}_i$ such that $a_{i,j} \leq a_{i,j+1}$. From this sorted list, the distances to the k^{th} nearest neighbor along each vector of variability is calculated:

$$\delta_{\text{PCAD}}(\mathbf{u}[j]) = \sqrt{\sum_{i=1}^{N} \max(a_{i,j} - a_{i,(j-k)}, a_{i,j} - a_{i,(j+k)})^2} \tag{16}$$

The calculation above is the Euclidian distance between pixel j and either nearest pixel, either $j + k$ or $j - k$, along each PCA vector of variability. The one parameter of this detector is k, which must be selected such that it is greater than the largest expected anomaly size ($k > a_{\max}$).

Much like PCAG, this algorithm will identify pixels that are substantially separated from the majority of the pixels in the block (and hence, less likely to come from the same underlying distribution). But unlike PCAG, this approach is more robust to noise as it is not dependent on having substantial gaps along the vectors of variability.

4.7 Color spaces

In addition to considering multiple hyperspectral algorithms, as well as our novel algorithm, to perform anomaly detection in search and rescue photos, we also evaluate algorithm performance across multiple color spaces. Existing literature on search and rescue image processing has focused on color space manipulation to make anomalous colors more identifiable [10]. Based on this idea, we evaluate algorithm performance on a number of color spaces that may better emphasize color anomalies. Color spaces considered include RGB, L*a*b, YCbCr, XYZ and uvL, as well as slight perceptual variations on these color spaces such as xyY and u'v'L. In addition, we consider smaller color spaces where luminance information is discarded, including the *a*b channels from L*a*b, X and Z from XYZ, x and y from xyY, u and v from uvL and u' and v' from u'v'L.

5 Implementation

Because there is no publicly available library of search and rescue scenes, we automated the generation of novel scenes by randomly superimposing color anomalies (clothing) into natural scenes. Using a range of natural scenes (such as beach, desert, mountain and forest photos), the process provides a large collection of suitable inputs for comparing and contrasting algorithm performance. The set of scenes specifically include photos with clear boundaries between regions of distinct texture. These scenes were selected based on previous reports that some anomaly detectors struggle in such an environment with sharp boundaries [11]. All anomalies are photos of clothing manually edited to remove any pixels besides the article of clothing. Each anomaly is resized to a total area of 45-90 pixels and rotated randomly. Anomalies are luminance adjusted to match the surrounding region of the scene more realistically. The boundaries of anomalies are blended with the surrounding scene to more realistically reflect likely search and rescue images. An example output of the anomaly generator is excerpted in figure 5.

The considered algorithms and color spaces were evaluated over a range of generated scenes. The performance of each algorithm is visualized via a receiver operating characteristic (ROC) curve. The area under the ROC curve (AUC) is used as a metric to compare across algorithms and color spaces. Each algorithm was applied across different natural scenes and different color spaces to assess relative advantages and disadvantages. In this way, we hope to bring current hyperspectral anomaly detection algorithms to bear in search and rescue.

In addition to evaluating the tradeoffs in accuracy for each algorithm, there was substantial range in the complexity and, as a result, relative performance of each algorithm. Many of the newer hyperspectral algorithms evaluate each pixel individually by running five ore more comparisons with the surrounding nested windows. Alternatively, algorithms like global RX, PCAG and PCAD are able to run on large block or full scenes, and hence run much more quickly. Our implementations of the

Figure 5: Excerpt from larger scene showing two anomalies superimposed on a beach photo. The full image used in the algorithm is approximately four times larger.

each algorithms have a run time range of between 15 seconds and 8 minutes. Given the large range in performance, we also evaluate and compare relative tradeoffs of algorithm detection in terms of computational complexity.

6 Experimental results

In the subsequent sections, we discuss the results, compare the seven considered algorithms as well as evaluate relative performance across different color spaces and environments.

6.1 Algorithms

Table 1 and figure 6 compare the AUC metric for the seven algorithms across a three representative color spaces that achieved high performance based on twelve natural scenes with randomly superimposed anomalies. These AUC values are calculated across all twelve scenes by evaluating different global anomaly thresholds for each algorithm and color scheme instead of averaging AUC values for individual scenes. As a result, these results represent the expected performance on a range of backgrounds, but results could potentially be improved by tuning thresholds to the scene or environment.

The results show strong performance for many of the considered algorithms. Global RX, which is often used as a benchmark for evaluating anomaly detection

in hyperspectral imagery [3] outperforms the windowed Local RX and DWEST, but performs slightly worse than NSWTD or MW-NSWTD. Our proposed PCAG algorithm performs well, especially in the L*a*b color space. This aligns with how it separates clusters of pixels in the three dimensional color space. PCAG does not match RX or NSWTD in performance, but the added benefit of the multi-window approach (MW-PCAG) further improves its performance relative to all algorithms. MW-PCAG outperforms RX, but does not outperform NSWTD or MW-NSWTD when only considering the AUC metric as seen in figure 6.

Our second proposed algorithm, PCAD, shows strong performance in RGB and L*a*b color spaces, similar to the MW-PCAG detector. Surprisingly though, it's performance is compromised by the XYZ color space. Potentially, a MW-PCAD implementation could achieve stronger results, but has not yet been implemented.

	RGB	L*a*b	XYZ
Global RX	0.9890	0.9854	0.9665
Local RX	0.8280	0.9287	0.6506
DWEST	0.9477	0.9703	0.8278
NSWTD	0.9991	0.9980	0.9690
MW-NSWTD	0.9982	0.9974	0.9648
PCAG	0.9659	0.9733	0.9439
MW-PCAG	0.9928	0.9936	0.9525
PCAD	0.9845	0.9939	0.8329

Table 1: AUC results across all considered algorithms for three representative color spaces.

6.2 Color spaces

Figure 7 show results of Global RX, NSWTD, and MW-PCAG across seven different color spaces. As seen in the figure, RGB and L*a*b most consistently produce reliable results for these three algorithms, and Global RX and MW-PCAG work well on YCbCr, with a slight bump in performance of MW-PCAG in YCbCr. This suggests that in YCbCr, there are larger gaps along the primary component vectors between anomalous pixels in an image relative to the distribution of pixels from natural backgrounds. Other color spaces, even those tuned to human perception, such as xyY and uvL, do not perform as well.

The initial results also suggest that simple RGB or L*a*b color spaces produce consistently good results; other color spaces, even those tuned to human perception, such as uvL, do not perform as well.

Previous anomaly detection implementations for search and rescue have ignored the luminance component, suggesting it plays a minor or negligible role in such anomaly detection [11]. In general, as seen in table 2, using only the channels dedi-

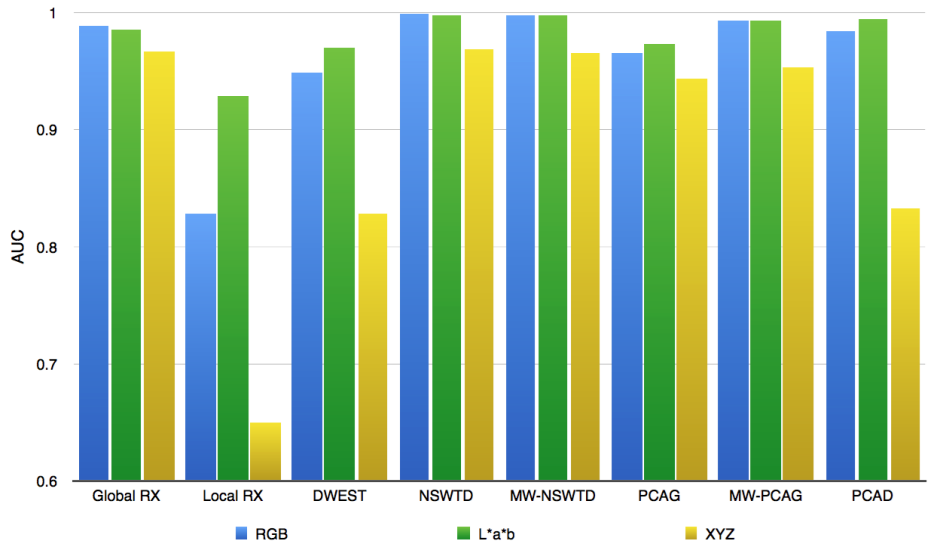

Figure 6: AUC results across all considered algorithms for three representative color spaces.

	Global RX	NSWTD	MW-PCAG
RGB	0.9890	0.9991	0.9928
L*a*b	0.9854	0.9980	0.9936
XYZ	0.9665	0.9690	0.9525
YCbCr	0.9890	0.9592	0.9950
uvL	0.9903	0.8700	0.9698
xyY	0.9894	0.9022	0.9919
u'v'L	0.9903	0.8753	0.9857
*a*b	0.9775	0.9898	0.9922
XZ	0.9082	0.6808	0.8977
CbCr	0.9801	0.9703	0.9894
uv	0.9804	0.9851	0.9952
xy	0.8421	0.9472	0.9257
u'v'	0.9804	0.9852	0.9948

Table 2: AUC results across all considered color spaces for three representative algorithms.

cated to color in color spaces that separate luminance and color into separate channels is not inherently beneficial to anomaly detection for search and rescue applications. The Global RX algorithm benefits from having luminance information included in the covariance matrix calculations in all color spaces. NSWTD performs poorly in color spaces tuned to the sensitivities of the human visual system (and YCbCr) and benefits from discarding luminance and only using color channels in these spaces. The

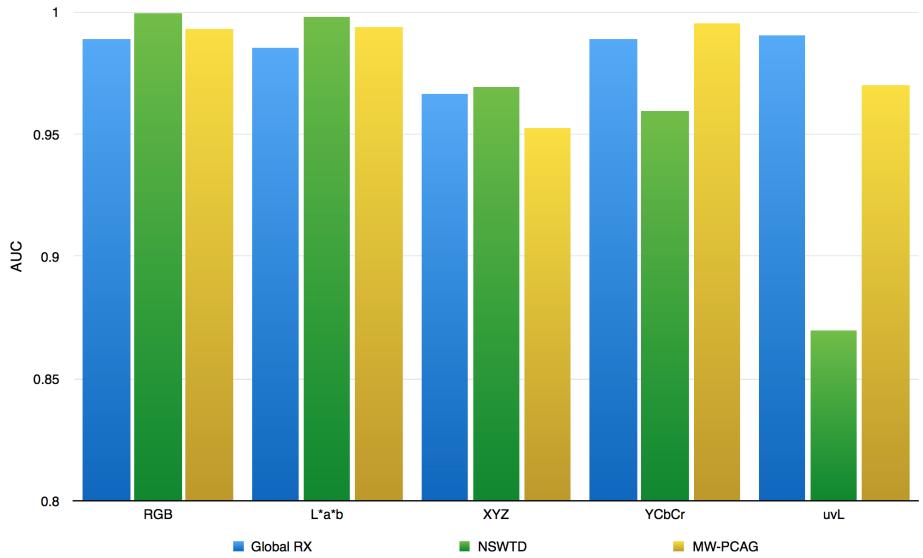

Figure 7: AUC results across color spaces for three representative algorithms.

performance of NSWTD in RGB or L*a*b space achieves the best results, suggesting that preprocessing the image or transforming the color space is unnecessary.

The higher performance of modern hyperspectral techniques in the RGB space is indicative that the algorithms benefit from images where each channel contains distinct, but statistically similar information that is equally relevant to the determination of anomaly versus non-anomaly.

MW-PCAG benefits from using only color data in uvL and u'v'L spaces, but not in any other similar color spaces. In table 2, this benefit is shown to be a 0.0024 increase in AUC from RGB and a 0.0002 increase in AUC from YCbCr for using only the uv channels of uvL color space. For such a minimal benefit, this only makes sense to pursue if the image is given in uvL or u'v'L color space.

6.3 Scene types

Table 1 shows the general performance of each algorithm for an average search and rescue image. In contrast, figure 8 breaks out the AUC performance in L*a*b space for different types of images (examples are shown in figure 10 in the appendix). As can be seen by the Global RX performance on the beach scenes, the global detector does not perform as well on images with multiple regions with distinct background statistics, but does perform very well on images with consistent backgrounds, such as the ocean image. The NSWTD and MW-PCAG both perform consistently across all of the images, with the NSWTD outperforming the PCAG in the AUC metric, as discussed in subsection 6.1.

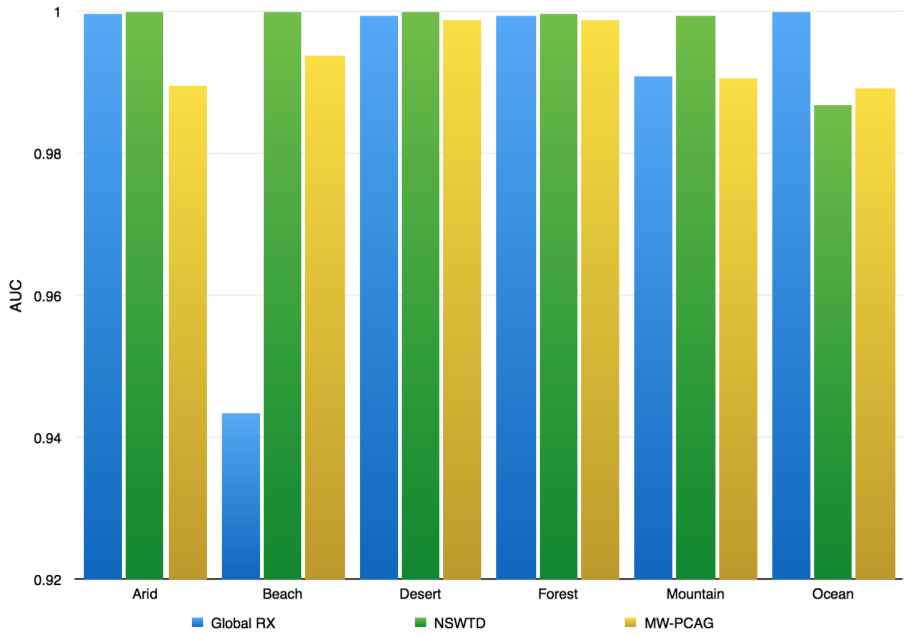

Figure 8: AUC results across different scene types for three representative algorithms using L*a*b color space.

6.4 Computation time

Figure 9 shows the average run time of all seven algorithms considered for a fixed image size. The global RX detector, the PCAG algorithm and the PCAD algorithm all have fast processing times. The desire for many applications will be to analyze many images from an aerial sweep of an area to determine whether or not a person or clues to a person's whereabouts are in an image, which can be very time sensitive information. This makes the time and performance considerations an important tradeoff for some algorithms. The NSWTD outperforms the RX, PCAG and PCAD detectors in the AUC metric, as does the MW-NSWTD, which performs slightly worse (-0.0009 in AUC metric) than NSWTD, but 1.5 times faster. MW-PCAG performs 7.6 times faster for a drop of only 0.0041 in the AUC metric from NSWTD. It can also be seen why global RX is seen as a benchmark for comparison here as it achieves performance extremely close to NSWTD (a difference in AUC of 0.0088 in their respective best color spaces), yet the RX algorithm performs the calculation 162 times faster on average.

Note that there are opportunities to optimize these algorithms further. We did not do an in-depth investigation of the fastest implementations, although we did try to streamline calculations. The differences in computation time accurately reflect the relative complexity, given that DWEST, NSWTD and MW-NSWTD all involve a number of calculations for each pixel based on several surrounding windows of comparison for which eigenvalues or covariance matrices need to be calculated. Performance improvements may be gained by implementation of integral images, but the

initial calculations and memory requirements of integral images would likely minimize the benefits of such an implementation.

Other potential avenues for improving performance include preprocessing the images. Research in the hyperspectral field has identified potential benefits in performance by first using wavelet transforms prior to implementing anomaly detection algorithms such as NWSTD [14].

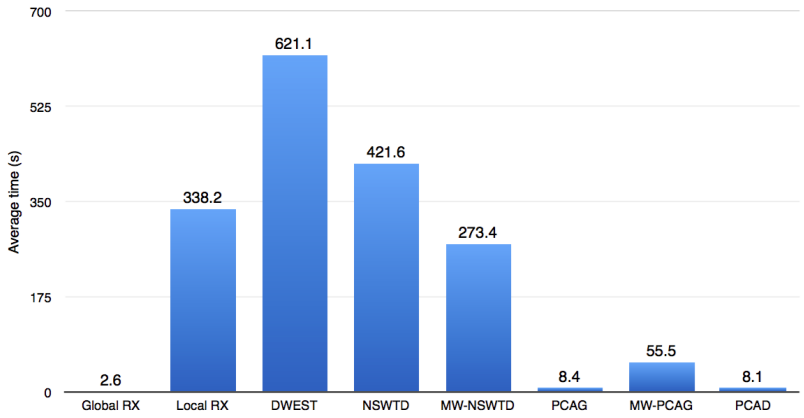

Figure 9: Average algorithm execution time on a single CPU for images of size 1536×1152.

7 Conclusions

Although results show that the NSWTD algorithm performs incredibly well in the AUC metric, there is a clear tradeoff between processing speed and accuracy. As mentioned before, the implementation of the NSWTD and MW-NSWTD here could benefit from further optimization, but inherently these algorithms require a lot more performance for individual pixel classification. Our analysis suggest that the PCAD and MW-PCAG approaches strike a very good balance between RX and NWSTD in AUC performance while maintaining an algorithm speed that approaches that of the RX detector. Future work can be done in evaluating more optimized versions of the NSWTD and MW-NSWTD algorithms as well as looking into clustering algorithms such as k-means or EM clustering for anomaly detection.

References

[1] B. S. Morse, D. Thornton, and M. A. Goodrich, *Color anomaly detection and suggestion for wilderness search and rescue*. New York, New York, USA: ACM, Mar. 2012.

[2] D. Manolakis and D. Marden, "Hyperspectral image processing for automatic target detection applications," *Lincoln Laboratory ...*, 2003.

[3] T. E. Smetek and K. W. Bauer, "Finding hyperspectral anomalies using multivariate outlier detection," *Aerospace Conference*, 2007.

[4] 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.

[5] D. C. Borghys, V. Achard, S. R. Rotman, N. Gorelik, C. Perneel, Scwheicher, and Emile, "Hyperspectral anomaly detection: a comparative evaluation of methods," *General Assembly and Scientific Symposium, XXXth URSI*, pp. 1–4, Feb. 2011.

[6] H. Kwon, "Adaptive anomaly detection using subspace separation for hyperspectral imagery," *Optical Engineering*, vol. 42, pp. 3342–10, Nov. 2003.

[7] W. Liu and C.-I. Chang, *A nested spatial window-based approach to target detection for hyperspectral imagery*, vol. 1. IEEE, 2004.

[8] W.-M. Liu and C.-I. Chang, "Multiple-Window Anomaly Detection for Hyperspectral Imagery," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 6, pp. 644–658, May 2013.

[9] M. J. Carlotto, "A cluster-based approach for detecting man-made objects and changes in imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, pp. 374–387, Jan. 2005.

[10] N. D. Rasmussen, D. R. Thornton, and B. S. Morse, "Enhancement of unusual color in aerial video sequences for assisting wilderness search and rescue," *2008 15th IEEE International Conference on Image Processing*, pp. 1356–1359, 2008.

[11] T. Bolukbasi and P. Tran, "Outlier color identification for search and rescue," *BU Department of Electrical and Computer Engineering Technical Reports*, pp. 1–25, Dec. 2012.

[12] C.-I. Chang, "Orthogonal subspace projection (OSP) revisited: a comprehensive study and analysis," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, pp. 502–518, Feb. 2005.

[13] M.-L. Shyu, S.-C. Chen, K. Sarinnapakorn, and L. Chang, "A Novel Anomaly Detection Scheme Based on Principal Component Classifier," pp. 353–365, 2003.

[14] M. Z. Baghbidi, K. Jamshidi, and A. Nilchi, "Improvement of Anomoly Detection Algorithms in Hyperspectral Images using Discrete Wavelet Transform," *arXiv.org*, 2012.

A Example scenes

The following image shows example natural scenes used in the training process as well as the rough "classification" used in comparison of algorithms across environments.

(a) Arid

(b) Beach

(c) Desert

(d) Forest

(e) Mountain

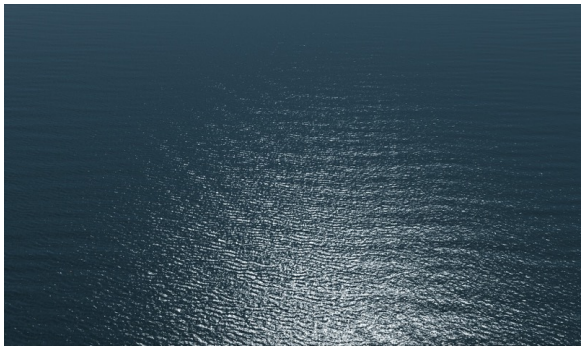

(f) Ocean

Figure 10: Representative images from each class of scene.