



Article

Evaluation of Color Anomaly Detection in Multispectral Images for Synthetic Aperture Sensing

Francis Seits, Indrajit Kurmi  and Oliver Bimber * 

Institute of Computer Graphics, Johannes Kepler University Linz, 4040 Linz, Austria

* Correspondence: oliver.bimber@jku.at; Tel.: +43-732-2468-6631

Abstract: In this article, we evaluate unsupervised anomaly detection methods in multispectral images obtained with a wavelength-independent synthetic aperture sensing technique called Airborne Optical Sectioning (AOS). With a focus on search and rescue missions that apply drones to locate missing or injured persons in dense forest and require real-time operation, we evaluate the runtime vs. quality of these methods. Furthermore, we show that color anomaly detection methods that normally operate in the visual range always benefit from an additional far infrared (thermal) channel. We also show that, even without additional thermal bands, the choice of color space in the visual range already has an impact on the detection results. Color spaces such as HSV and HLS have the potential to outperform the widely used RGB color space, especially when color anomaly detection is used for forest-like environments.

Keywords: multispectral; image processing; anomaly detection; search and rescue; unmanned aerial vehicles; airborne optical sectioning



Citation: Seits, F.; Kurmi, I.; Bimber, O. Evaluation of Color Anomaly Detection in Multispectral Images for Synthetic Aperture Sensing. *Eng* **2022**, *3*, 541–553. <https://doi.org/10.3390/eng3040038>

Academic Editor: Antonio Gil Bravo

Received: 3 November 2022

Accepted: 25 November 2022

Published: 29 November 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Color anomaly detection methods identify pixel regions in multispectral images that have a low probability of occurring in the background landscape and are therefore considered to be outliers. Such techniques are used in remote sensing applications for agriculture, wildlife observation, surveillance, or search and rescue. Occlusion caused by vegetation, however, remains a major challenge.

Airborne Optical Sectioning (AOS) [1–13] is a synthetic aperture sensing technique that computationally removes occlusion in real-time by registering and integrating multiple images captured within a large synthetic aperture area above the forest (cf. Figure 1). With the resulting shallow-depth-of-field integral images, it becomes possible to locate targets (e.g., people, animals, vehicles, wildfires, etc.) that are otherwise hidden under the forest canopy. Image pixels that correspond to the same target on the synthetic focal plane (i.e., the forest ground) are computationally aligned and enhanced, while occluders above the focal plane (i.e., trees) are suppressed in strong defocus. AOS is real-time and wavelength-independent (i.e., it can be applied to images in all spectral bands), which is beneficial for many areas of application. Thus far, AOS has been applied to the visible [1,11] and the far-infrared (thermal) spectrum [4] for various applications, such as archeology [1,2], wildlife observation [5], and search and rescue [8,9]. By employing a randomly distributed statistical model [3,10,12], the limits of AOS and its efficacy with respect to its optimal sampling parameters can be explained. Common image processing tasks, such as classification with deep neural networks [8,9] or color anomaly detection, [11] are proven to perform significantly better when applied to AOS integral images compared with conventional aerial images. We also demonstrated the real-time capability of AOS by deploying it on a fully autonomous and classification-driven adaptive search and rescue drone [9]. In [11,13], we presented the first solutions to tracking moving people through densely occluding foliage.




Figure 1. Airborne optical sectioning (AOS) is a synthetic aperture sensing technique that computationally combines multiple aerial images captured within a synthetic aperture area (a) to an integral image, which enhances targets on the synthetic focal plane while suppressing occluders above it. Right: People covered by forest canopy. Single aerial image (thermal channel) that suffers from strong occlusion (b), and corresponding integral image of the same environment with occlusion removed (c).

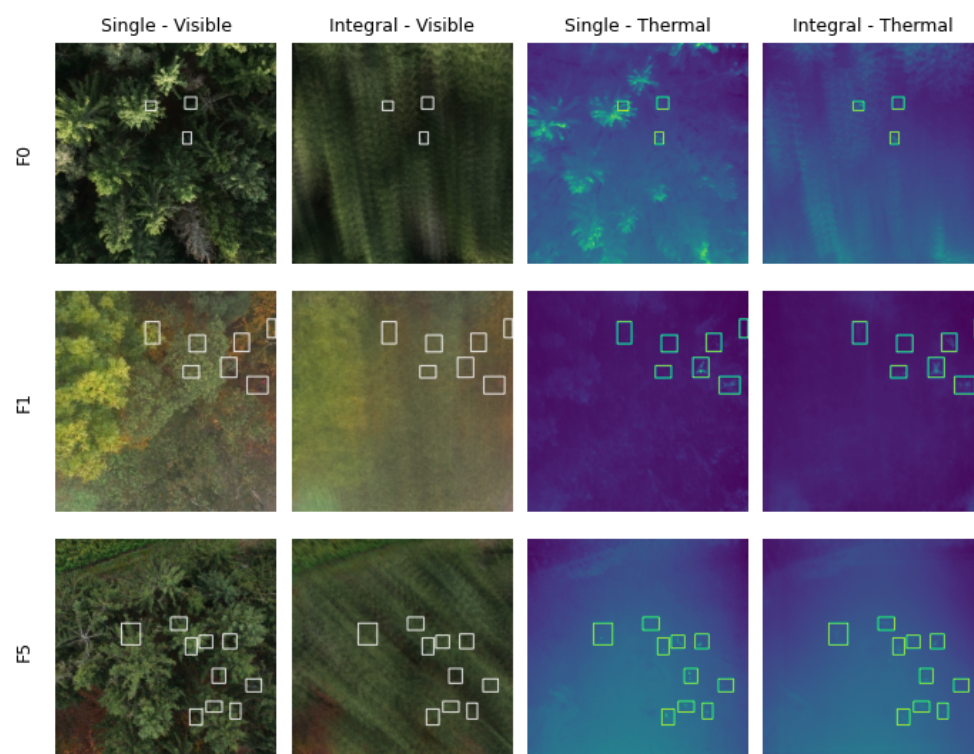
Anomaly detection methods for wilderness search and rescue have been evaluated earlier [14], and bimodal systems using a composition of visible and thermal information were already used to improve detection rates of machine learning algorithms [15,16]. However, none of the previous work considered occlusion.

With AOS, we are able to combine multispectral recordings into a single integral image. Our previous work has shown that image processing tasks, such as person classification with deep neural networks [8–10], perform significantly better on integral images when compared to single images. These classifiers are based on supervised architectures, which have the disadvantage that training data must be collected and labeled in a time-consuming manner and that the trained neural networks do not generalize well into other domains. It was also shown in [11] that the image integration process of AOS decreases variance and covariance, which allows better separation of target and background pixels when applying the Reed–Xiaoli (RX) unsupervised anomaly detection [17].

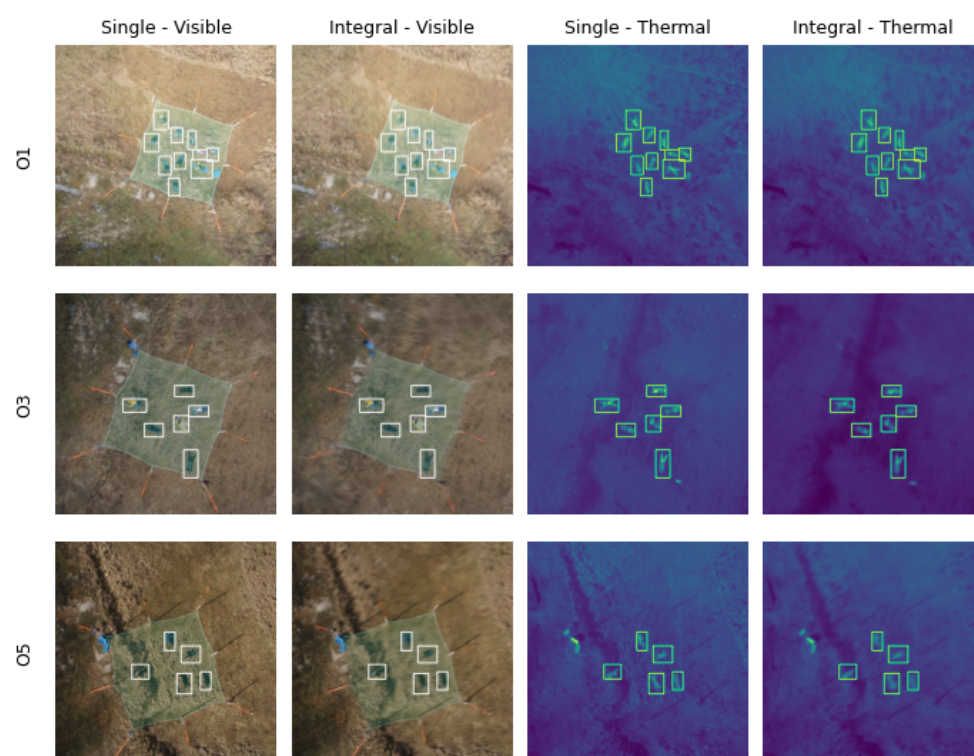
In this article, we evaluate several common unsupervised anomaly detection methods being applied to multispectral integral images that are captured from a drone when flying over open and occluded (forest) landscapes. We show that their performance can significantly be improved by the right combination of spectral bands and choice of color space input format. Especially for forest-like environments, detection rates of occluded people can be consistently increased if visible and thermal bands are combined and if HSV or HLS color spaces are used for the visible bands instead of common RGB. Furthermore, we also evaluate the runtime behavior of these methods when considered for time-critical applications, such as search and rescue.

2. Materials and Methods

For our evaluation, we applied the dataset from [8], which was used to prove that integral images improve people classification under occluded conditions. It consists of RGB and thermal images (pairwise simultaneously) captured with a drone prototype over multiple forest types (broadleaf, conifer, mixed) and open landscapes, as shown in Figure 2. In all images, targets (persons laying on the ground) are manually labeled. Additional telemetry data (GPS and IMU sensor values) of the drone during capturing are also provided for each image.



(a) Forest Landscapes



(b) Open Landscapes

Figure 2. Our evaluation dataset consists of several forest (a) and open (b) landscape images captured with a drone from an altitude of about 35m AGL. Each scenery (F0, F1, F5, O1, O3, O5) contains about 20 consecutive single images taken in the visible (RGB) and thermal spectrum, which are combined into integral images. Rectangles indicate manually labeled persons lying on the ground.

While the visible bands were converted from RGB to other color spaces (HLS, HSV, LAB, LUV, XYZ, and YUV), the thermal data were optionally added as a fourth (alpha) channel, resulting in additional input options (RGB-T, HLS-T, HSV-T, LAB-T, LUV-T, XYZ-T, and YUV-T).

All images had a resolution of 512x512 pixels, so the input dimensions were either (512, 512, 3) or (512, 512, 4). Methods that do not require spatial information used flattened images with (262144, 3) or (262144, 4) dimensions.

The publicly available C/C++ implementation of AOS Source Code: <https://github.com/JKU-ICG/AOS> (accessed on 28 November 2022) was used to compute integral images from single images.

2.1. Color Anomaly Detectors

Unsupervised color anomaly detectors have been widely used in the past [17–22], with the Reed–Xiaoli (RX) detector [17] being commonly considered as a benchmark. Several variations of RX exist, where the standard implementation calculates global background statistics (over the entire image) and then compares individual pixels based on the Mahalanobis distance. In the further course of this article, we will refer to this particular RX detector as Reed–Xiaoli Global (RXG).

The following briefly summarizes the considered color anomaly detectors, while details can be found through the provided references:

The Reed–Xiaoli Global (RXG) detector [17] computes a $K_{n \times n}$ covariance matrix of the image, where n is the number of input channels (e.g., for RGB, $n = 3$ and for RGB-T, $n = 4$). The pixel under test is the n -dimensional vector r , and the mean is given by the n -dimensional vector μ :

$$\alpha_{RXG}(r) = (r - \mu)^T K_{n \times n}^{-1} (r - \mu).$$

The Reed–Xiaoli Modified (RXM) detector [18] is a variation of RXG, where an additional constant $\kappa = ||r - \mu||^{-1}$ is used for normalization:

$$\alpha_{RXM}(r) = \kappa \cdot \alpha_{RXG}(r) = \left(\frac{r - \mu}{||r - \mu||} \right)^T K_{n \times n}^{-1} (r - \mu).$$

The Reed–Xiaoli Local (RXL) detector computes covariance and mean over smaller local areas and, therefore, does not use global background statistics. The areas are defined by an inner window (*guard_win*) and an outer window (*bg_win*). The mean μ and covariance K are calculated based on the outer window but excludes the inner window. Window sizes were chosen to be *guard_win* = 33 and *bg_win* = 55, based on the projected pixel sizes of the targets in the forest landscape.

The principal component analysis (PCA) [19] uses singular value decomposition for a linear dimensionality reduction. The covariance matrix of the image is decomposed into eigenvectors and their corresponding eigenvalues. A low-dimensional hyperplane is constructed by selected (*n_components*) eigenvectors. Outlier scores for each sample are then obtained by their euclidean distance to the constructed hyperplane. The number of eigenvectors to use was chosen to be *n_components* = n , where n is the number of input channels (e.g., for RGB, $n = 3$ and for RGB-T, $n = 4$).

The Gaussian mixture model (GMM) [20] is a clustering approach, where multiple Gaussian distributions are used to characterize the data. The data are fit to each of the single Gaussians (*n_components*), which are considered as a representation of clusters. For each sample, the algorithm calculates the probability of belonging to each cluster, where low probabilities are an indication of being an anomaly. The number of Gaussians to use was chosen to be *n_components* = 2.

The cluster based anomaly detection (CBAD) [21] estimates background statistics over clusters instead of sliding windows. The image background is partitioned (using any clustering algorithm) into clusters (*n_cluster*), where each cluster can be modeled as a Gaussian distribution. Similar to GMM, anomalies have values that deviate significantly

from the cluster distributions. Samples are each assigned to the nearest background cluster, becoming an anomaly if their value deviates farther from the mean than background pixel values in that cluster. The number of clusters to use was chosen to be $n_{cluster} = 2$.

The local outlier factor (LOF) [22] uses a distance metric (e.g., Minkowski distance) to determine the distances between neighboring ($n_{neighbors}$) data points. Based on the inverse of those average distances, the local density is calculated. This is then compared to the local densities of their surrounding neighborhood. Samples that have significantly lower densities than their neighbors are considered isolated and, therefore, become outliers. The number of neighbors to use was chosen to be $n_{neighbors} = 200$.

2.2. Evaluation

The evaluation of the methods summarized above was carried out on a consumer PC (Intel Core i9-11900H @ 2.50GHz) for the landscapes shown in Figure 2.

Precision (Equation (1)) vs. recall (Equation (2)) was used as metrics for performance comparisons.

The task can be formulated as a binary classification problem, where positive predictions are considered anomalous pixels. The data we want to classify (image pixels) is highly unbalanced, as most of the pixels are considered as background (majority class) and only some of the pixels are considered anomalies (minority class).

The true positive (TP) pixels are determined by checking whether they lie within one of the labeled rectangles, as shown in Figure 2. Pixels detected outside these rectangles are considered false positives (FP) and pixels inside the rectangle but not classified anomalously are considered false negatives (FN). Since the dataset only provides rectangles for labels and not perfect masks around the persons, the recall results are biased (in general, not as good as expected). As we are mainly interested in the performance difference between individual methods and the errors introduced are always constant (rectangle area—real person mask), the conclusions drawn from the results should be the same, even if perfect masks were used instead.

The precision (Equation (1)) quantifies the number of correct positive predictions made, and recall (Equation (2)) quantifies the number of correct positive predictions made out of all positive predictions that could have been made:

$$Precision = \frac{TP}{TP + FP}, \quad (1)$$

$$Recall = \frac{TP}{FN + TP}. \quad (2)$$

Precision and recall both focus on the minority class (anomalous pixels) and are therefore less concerned with the majority class (background pixels), which is important for our unbalanced dataset.

Since the anomaly detection methods provide probabilistic scores on the likelihood of a pixel being considered anomalous, a threshold value must be chosen to obtain a final binary result.

The precision–recall curve (PRC) in Figure 3 shows the relationship between precision and recall for every possible threshold value that could be chosen. Thus, a method performing well would have high precision and high recall over different threshold values. We use the area under the precision–recall curve (AUPRC), which is simply the integral of the PRC, as the final evaluation metric.




Figure 3. The area under the precision–recall curve (AUPRC) is used as a metric for comparing the performance of the evaluated anomaly detection methods. We consider true positives (TP), false positives (FP), and false negatives (FN) pixels for each image and calculate precision and recall. The example above illustrates precision–recall curves of all landscapes for RXG and with RGB-T input.

The AUPRC metric provides comparable results on the overall performance of a method but is not well suited when it comes to finding the best threshold for a single image. To obtain the best threshold value for a single image, we use the F_β -score (Equation (3)), which is also calculated from precision and recall:

$$F_\beta = \left(1 + \beta^2\right) \cdot \frac{\text{Precision} \cdot \text{Recall}}{(\beta^2 \cdot \text{Precision}) + \text{Recall}}, \quad (3)$$

where β is used as a weighting factor that can be chosen such that recall is considered β -times more important than precision.

The balanced F_1 -score is the harmonic mean of precision and recall and is widely used. However, as we care more about minimizing false positives than minimizing false negatives, we would select a $\beta < 1$. A grid search of β values has given the best results for $\beta = \frac{1}{2}$. With this setting, precision is weighted more heavily than recall. The F_β metric is only used to threshold the image scores for comparison purposes, as shown in Figure 6.

3. Results

Figure 4 shows the AUPRC values across different color spaces and methods. The methods are evaluated on each color space, once with three channels (visible spectrum only) and once with four channels (visible and thermal spectrum). The results of the forest landscape are average values over F0, F1 and F5, and the results of the open landscape are average values over O1, O3 and O5.

As expected (and as we have also seen in Figure 3), the overall AUPRC of the open landscapes is much higher than the AUPRC values of the more challenging forest landscapes. The reason is an occlusion in the presence of forests.

The AUPRC values of the four-channel (color + thermal) and three-channel (color only) inputs are overlayed in the same bar. The slightly lighter colored four-channel results are always higher than the three-channel results—regardless of the method or the color space used. However, the difference is more pronounced for the forest landscapes than for the open landscapes. This shows that regardless of the scenery and regardless of the method and the color space used, the additional thermal information always improves the performance of anomaly detection.

With a look at the AUPRC values in the forest landscapes, we can observe that RXL gives the overall best results and outperforms all other methods. Utilizing the additional thermal information gives, in this case, even a $2\times$ gain. This can also be observed visually in the anomaly detection scores shown in Figure 6, where FP's detections highly decrease and TP's detections highly increase if the thermal channel is added (e.g., F1, in the visible spectral band, many background pixels are considered anomalous, with the additional thermal information those misclassified pixels are eliminated).

Looking at the AUPRC values in the open landscapes, we can observe that the difference between the methods is not as pronounced as in the forest landscapes. An obvious outlier, however, seems to be LOF, which nevertheless performs very well (second best) in the forest landscapes. This can be explained by the fact that hyper-parameters of the methods were specifically chosen for the forest landscape. In the case of LOF, the $n_neighbors$ parameter was set to be 200, which seems suboptimal for the open landscapes. The same holds for RXL (window sizes), CBAD (number of clusters), GMM (number of components) and PCA (number of components). All other methods do not require hyper-parametrization.

Another observation that can be made is that some color spaces consistently give better results than others. In the forest landscapes, HSV(-T) usually gives the best results, regardless of the methods being used. In the open landscapes, it is not as clear which color space performs best, but HSV(-T) still gives overall good results. In general, and especially for RXM, the improvements achieved by choosing HSV(-T) over other color spaces are clearly noticeable.




Figure 4. Results of area under the precision–recall curve (AUPRC) values for multiple color spaces and color anomaly detection methods. The results of the forest landscape are average values over F0, F1, and F5, and the results of the open landscape are average values over O1, O3, and O5. The stacked bar charts highlight the improvement gains caused by the additional thermal channel.

The individual results plotted in Figure 4 are also shown in Tables 1 and 2, where the mean values over all color spaces (last row) may give a useful estimate of the method's overall performance. The highest AUPRC value for the forest and open scenery is highlighted in bold.

Since anomaly detection for time-critical applications should deliver reliable results in real-time, we have also measured their runtimes, as shown in Table 3. The best-performing methods on the forest landscapes in terms of AUPRC values are RXL and LOF. In terms of runtime, both are found to be very slow, as they consume 20 to 35 s for computations, where all other algorithms provide anomaly scores in under a second (cf. Figure 5).



Figure 5. Performance in AUPRC (left bars) vs. runtime in *ms* (right bars): Color anomaly detection methods that produce results in less than a second. Reed–Xiaoli local (RXL) and local outlier factor (LOF) performed well but needed more than 20 seconds and are, therefore, not practicable for applications with real-time demands.

Table 1. Area under the precision–recall curve (AUPRC) values for each color space and color anomaly detection method. The scores are obtained from integral images and are averaged over forest landscapes. The last row is the mean AUPRC value over all color spaces. The best performance is highlighted in bold.

	CBAD	GMM	LOF	PCA	RXG	RXL	RXM
HLS	0.047	0.029	0.037	0.028	0.029	0.093	0.044
HSV	0.050	0.033	0.040	0.028	0.033	0.096	0.048
LAB	0.029	0.028	0.038	0.023	0.028	0.077	0.031
LUV	0.032	0.032	0.038	0.024	0.032	0.098	0.034
RGB	0.029	0.027	0.036	0.026	0.027	0.084	0.030
XYZ	0.027	0.027	0.039	0.027	0.027	0.088	0.027
YUV	0.029	0.026	0.038	0.023	0.026	0.077	0.032
AUPRC	0.035	0.029	0.038	0.025	0.029	0.088	0.035
	CBAD	GMM	LOF	PCA	RXG	RXL	RXM
HLS-T	0.081	0.061	0.098	0.043	0.062	0.218	0.099
HSV-T	0.093	0.076	0.099	0.053	0.077	0.223	0.108
LAB-T	0.083	0.064	0.084	0.052	0.065	0.196	0.047
LUV-T	0.076	0.056	0.083	0.047	0.056	0.186	0.052
RGB-T	0.084	0.057	0.092	0.032	0.057	0.191	0.055
XYZ-T	0.082	0.057	0.083	0.032	0.058	0.196	0.037
YUV-T	0.078	0.055	0.088	0.042	0.056	0.190	0.050
AUPRC	0.082	0.061	0.090	0.043	0.062	0.200	0.064

Table 2. Area under the precision–recall curve (AUPRC) values for each color space and anomaly detection method. The scores are obtained from integral images and are averaged over open landscapes. The last row is the mean AUPRC value over all color spaces. Best performance is highlighted in bold.

	CBAD	GMM	LOF	PCA	RXG	RXL	RXM
HLS	0.219	0.255	0.022	0.239	0.256	0.215	0.275
HSV	0.211	0.256	0.044	0.246	0.258	0.217	0.291
LAB	0.243	0.242	0.021	0.276	0.241	0.144	0.189
LUV	0.234	0.230	0.020	0.229	0.228	0.136	0.185
RGB	0.261	0.257	0.078	0.094	0.256	0.153	0.175
XYZ	0.254	0.254	0.019	0.078	0.253	0.155	0.144
YUV	0.257	0.256	0.019	0.249	0.255	0.153	0.162
AUPRC	0.240	0.250	0.032	0.202	0.249	0.168	0.203
	CBAD	GMM	LOF	PCA	RXG	RXL	RXM
HLS-T	0.258	0.293	0.077	0.292	0.294	0.268	0.296
HSV-T	0.259	0.303	0.082	0.302	0.304	0.268	0.316
LAB-T	0.305	0.309	0.082	0.306	0.310	0.222	0.250
LUV-T	0.295	0.302	0.082	0.267	0.300	0.212	0.249
RGB-T	0.312	0.312	0.086	0.187	0.312	0.225	0.260
XYZ-T	0.308	0.310	0.123	0.169	0.310	0.229	0.219
YUV-T	0.314	0.312	0.075	0.297	0.312	0.225	0.250
AUPRC	0.293	0.306	0.087	0.260	0.306	0.236	0.263

Table 3. Runtime for each input format and method in milliseconds. The input format (color spaces) does not have an influence on the runtime, but addition channels (thermal) may increase the runtime for some algorithms. The last row is the mean runtime of an algorithm. Best performance is highlighted in bold.

	CBAD	GMM	LOF	PCA	RXG	RXL	RXM
HLS	887	219	18,440	98	39	36,647	28
HSV	883	225	18,094	102	40	36,346	27
LAB	877	219	19,833	99	36	36,495	29
LUV	888	212	19,975	96	41	36,166	29
RGB	925	216	19,367	96	40	35,732	29
XYZ	872	224	19,107	100	40	36,367	27
YUV	874	228	20,485	100	42	36,409	27
Runtime	887	221	19,329	99	40	36,309	28
	CBAD	GMM	LOF	PCA	RXG	RXL	RXM
HLS-T	878	243	28,054	104	43	36,002	30
HSV-T	882	237	28,235	107	40	35,801	30
LAB-T	879	230	28,005	108	32	35,876	31
LUV-T	893	229	27,506	105	45	36,059	32
RGB-T	904	230	27,323	98	37	35,618	32
XYZ-T	880	241	24,395	108	43	36,047	29
YUV-T	877	238	27,146	107	40	36,180	29
Runtime	885	236	27,238	105	40	35,940	31

4. Discussion

The AUPRC results in Figure 4 show that all color anomaly detection methods benefit from additional thermal information, but especially in combination with the forest landscapes.

In challenging environments, where the distribution of colors has a much higher variance (e.g., F1 in Figure 6, due to bright sunlight), the additional thermal information improves results significantly. If the temperature difference between targets and the surrounding is large enough, the thermal spectral band may add spatial information (e.g., distinct clusters of persons), which is beneficial for methods that calculate results based on locality properties (e.g., RXL, LOF).

In forest-like environments, the RXL anomaly detector performs best regardless of the input color space. This could be explained by the specific characteristics of an integral image. In the case of occlusion, the integration process produces highly blurred images caused by defocused occluders (forest canopy) above the ground, which results in a much more uniformly distributed background. Since target pixels on the ground stay in focus, anomaly detection methods such as RXL, which calculate background statistics on a smaller window around the target, benefit from the uniform distributed (local) background. The same is true for LOF, where the local density in the blurred background regions is much higher than the local density in the focused target region, resulting in overall better outlier detection rates. Since most objects in open landscapes are located near the focal plane (i.e., at nearly the same altitude above the ground), there is no out-of-focus effect caused by the integration process. Thus, these methods do not produce similarly good results for open landscapes.

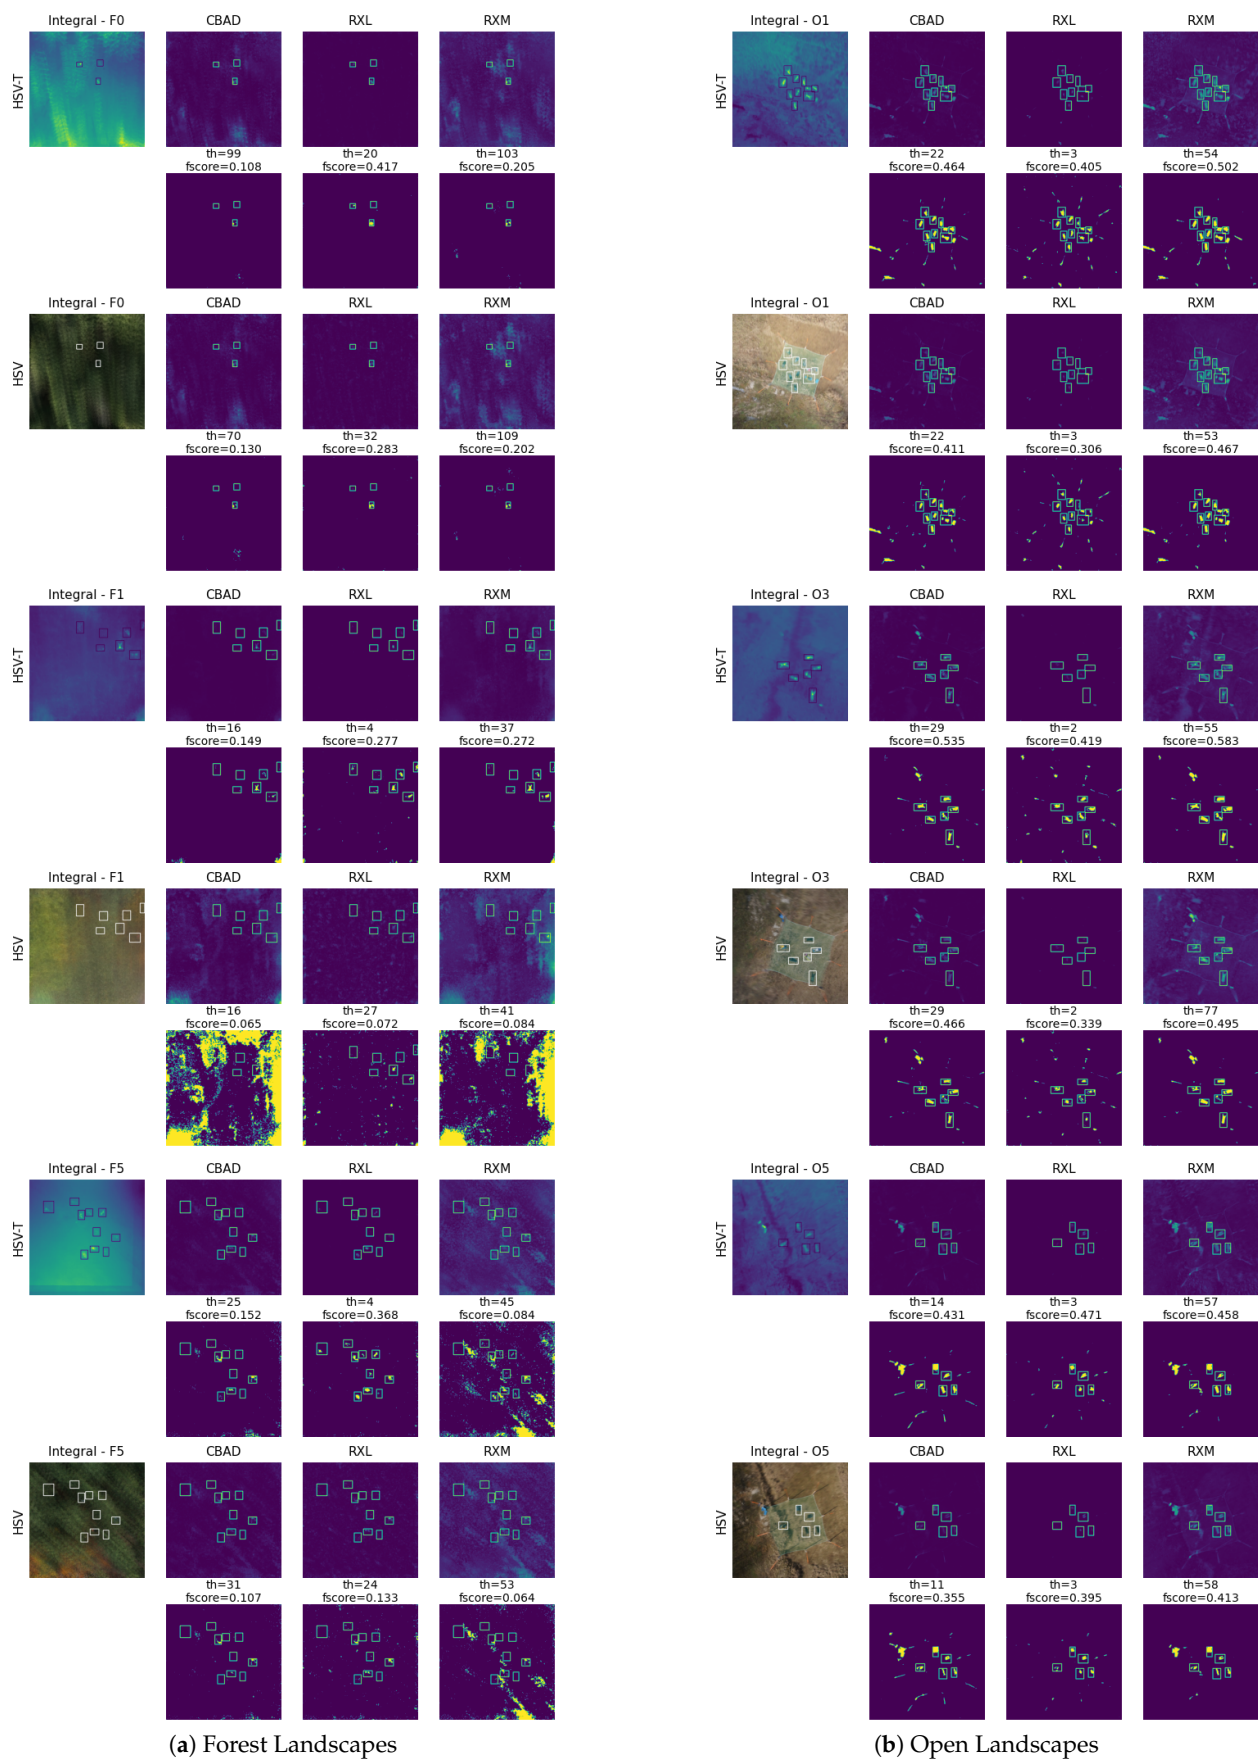


Figure 6. Color anomaly detection scores for forest (a) and open (b) landscapes, comparing the overall good performing HSV(-T) inputs. The first rows per scenery show anomaly scores of the best

(RXL) algorithm without considering runtime and the best (RXM) and second-best (CBAD) method when considering runtime. The second rows per scenery are the anomaly scores after thresholding.

For the forest landscapes, the HSV(-T) and HSL(-T) color spaces consistently give better results than others. The color spaces HSV (hue, saturation, value) and HSL (hue, saturation, lightness) are both based on cylindrical color space geometries and differ mainly in their last dimension (brightness/lightness). The first two dimensions (hue, saturation) can be considered more important when distinguishing colors, as the last dimension only describes the value (brightness) or lightness of a color. We assume that the more uniform background resulting from the integration process also has a positive effect on the distance metric calculations when those two color spaces are used, especially if the background mainly consists of a very similar color tone. This is again more pronounced for the forest landscapes than for the open landscapes.

Although the AUPRC results obtained from RXL and LOF are best for forest landscapes, the high runtime indicates that these methods are impractical for real-time applications. A trade-off must be made between good anomaly detection results and fast runtime; therefore, we consider the top-performing methods that provide reliable results within milliseconds further.

Based on the AUPRC and runtime results shown in Figure 5, one could suggest that the RXM method may be used. The AUPRC results combined with HSV-T are the best among methods that run under one second, regardless of the landscape. Since this method does not require a-priori settings to be chosen (only the final thresholding value) and the runtime is one of the fastest, it would be well suited for usage in forests and open landscapes. The second-best algorithm based on the AUPRC values would be CBAD, with the disadvantage that it requires a hyper-parameter setting and does not generalize well for open landscapes.

5. Conclusions

In this article, we have shown that the performance of unsupervised color anomaly detection methods applied to multispectral integral images can be further improved by an additional thermal channel. Each of the evaluated methods performs significantly better when thermal information is utilized in addition, regardless of the landscape (forest or open). Another finding is that even without the additional thermal band, the choice of input color space (for the visible channels) already has an influence on the results. Color spaces such as HSV and HLS can outperform the widely used RGB color space, especially in forest-like landscapes.

These findings might guard decisions on the choice of color anomaly detection method, input format, and applied spectral band, depending on individual use cases. Occlusion cause by vegetation, such as forests, remains challenging for many of them. In the future, we will investigate anomalies caused by motion in the context of synthetic aperture sensing. In combination with color and thermal anomaly detection, motion anomaly detection has the potential to further improve detection results for moving targets, such as people, animals, or vehicles.

Author Contributions: Conceptualization, O.B. and F.S.; methodology, F.S.; software, F.S. and I.K.; validation, F.S., I.K. and O.B.; formal analysis, F.S.; investigation, F.S.; resources, I.K.; data curation, I.K.; writing—original draft preparation, F.S. and O.B.; writing—review and editing, F.S. and O.B.; visualization, F.S.; supervision, O.B.; project administration, O.B.; funding acquisition, O.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Austrian Science Fund (FWF) under grant number P32185-NBL and by the State of Upper Austria and the Austrian Federal Ministry of Education, Science and Research via the LIT-Linz Institute of Technology under grant number LIT-2019-8-SEE114.

Data Availability Statement: The data and source code used in the experiments can be downloaded from <https://doi.org/10.5281/zenodo.3894773> and <https://github.com/JKU-ICG/AOS> (accessed on 28 November 2022).

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Kurmi, I.; Schedl, D.; Bimber, O. Airborne optical sectioning. *J. Imaging* **2018**, *4*, 102. [CrossRef]
2. Bimber, O.; Kurmi, I.; Schedl, D. Synthetic aperture imaging with drones. *IEEE Comput. Graph. Appl* **2019**, *39*, 8–15. [CrossRef] [PubMed]
3. Kurmi, I.; Schedl, D.; Bimber, O. A statistical view on synthetic aperture imaging for occlusion removal. *IEEE Sens. J* **2019**, *19*, 9374–9383. [CrossRef]
4. Kurmi, I.; Schedl, D.; Bimber, O. Thermal airborne optical sectioning. *Remote Sens.* **2019**, *11*, 1668. [CrossRef]
5. Schedl, D.; Kurmi, I.; Bimber, O. Airborne optical sectioning for nesting observation. *Sci. Rep.* **2020**, *10*, 7254. [CrossRef]
6. Kurmi, I.; Schedl, D.; Bimber, O. Fast Automatic Visibility Optimization for Thermal Synthetic Aperture Visualization. *IEEE Geosci. Remote Sens. Lett.* **2021**, *18*, 836–840. [CrossRef]
7. Kurmi, I.; Schedl, D.; Bimber, O. Pose Error Reduction for Focus Enhancement in Thermal Synthetic Aperture Visualization. *IEEE Geosci. Remote. Sens. Lett.* **2021**, to be published. [CrossRef]
8. Schedl, D.; Kurmi, I.; Bimber, O. Search and rescue with airborne optical sectioning. *Nat. Mach. Intell.* **2020**, *2*, 783–790. [CrossRef]
9. Schedl, D.; Kurmi, I.; Bimber, O. An autonomous drone for search and rescue in forests using airborne optical sectioning. *Sci. Robot* **2021**, *6*, 1188. [CrossRef]
10. Kurmi, I.; Schedl, D.; Bimber, O. Combined person classification with airborne optical sectioning. *Sci. Rep.* **2022**, *12*, 3804. [CrossRef]
11. Nathan, R.; Kurmi, I.; Schedl, D.; Bimber, O. Through-Foliage Tracking with Airborne Optical Sectioning. *J. Remote Sens.* **2022**, *2022*, 9812765. [CrossRef]
12. Seits, F.; Kurmi, I.; Nathan, R.; Ortner, R.; Bimber, O. On the Role of Field of View for Occlusion Removal with Airborne Optical Sectioning. *arXiv* **2022**, arXiv:2204.13371. <https://doi.org/10.48550/ARXIV.2204.13371>.
13. Nathan, R.; Kurmi, I.; Bimber, O. Inverse Airborne Optical Sectioning. *Drones* **2022**, *6*, 231. [CrossRef]
14. Morse, B.; Thornton, D.; Goodrich, M. Color anomaly detection and suggestion for wilderness search and rescue. In Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction, Boston, MA, USA, 5–8 March 2012; Association for Computing Machinery: New York, NY, USA, 2012; pp. 455–462. [CrossRef]
15. Rudol, P.; Doherty, P. Human Body Detection and Geolocalization for UAV Search and Rescue Missions Using Color and Thermal Imagery. In Proceedings of the 2008 IEEE Aerospace Conference, Big Sky, MT, USA, 1–8 March 2008; pp. 1–8. [CrossRef]
16. Hinzmann, T.; Stegemann, T.; Cadena, C.; Siegwart, R. Deep Learning-based Human Detection for UAVs with Optical and Infrared Cameras: System and Experiments. *arXiv* **2020**, arXiv:2008.04197. <https://doi.org/10.48550/ARXIV.2008.04197>.
17. Reed, I.; Yu, X. Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution. *IEEE Trans. Acoust. Speech Signal Process.* **1990**, *38*, 1760–1770. [CrossRef]
18. Chang, C.; Chiang, S. Anomaly detection and classification for hyperspectral imagery. *Geosci. Remote Sens. IEEE Trans.* **2002**, *40*, 1314–1325. [CrossRef]
19. Shyu, M.; Chen, S.; Sarinnapakorn, K.; Chang, L. A Novel Anomaly Detection Scheme Based on Principal Component Classifier. In Proceedings of the IEEE Foundations and New Directions of Data Mining Workshop, in conjunction with the Third IEEE International Conference on Data Mining (ICDM'03), Melbourne, FL, USA, 19–22 November 2003;
20. Bishop, C.M.; Nasrabadi, N.M. Pattern Recognition and Machine Learning. *J. Electron. Imaging* **2007**, *16*, 049901.
21. Carlotto, M. A cluster-based approach for detecting man-made objects and changes in imagery. *IEEE Trans. Geosci. Remote Sens.* **2005**, *43*, 374–387. [CrossRef]
22. Breunig, M.; Kriegel, H.; Ng, R.; Sander, J. LOF: Identifying Density-Based Local Outliers. *ACM SIGMOD Rec.* **2000**, *29*, 93–104. [CrossRef]