Universität der Bundeswehr München
Institut für
Flugsysteme

Sensor Performance Modelling for Camouflage Detection in Airborne Multispectral Imagery

Tobias Hupel

Vollständiger Abdruck der von der Fakultät für Luft- und Raumfahrttechnik der Universität der Bundeswehr München zur Erlangung des akademischen Grades eines

Doktor-Ingenieurs (Dr.-Ing.)

angenommenen Dissertation.

Gutachter:

1. Univ.-Prof. Dr.-Ing. Peter Stütz
2. Univ.-Prof. Dr.-Ing. Michael Heizmann

Diese Dissertation wurde am 04.06.2025 bei der Universität der Bundeswehr München eingereicht und durch die Fakultät für Luft- und Raumfahrttechnik am 15.10.2025 angenommen. Die mündliche Prüfung fand am 17.11.2025 statt.

Acknowledgments

I want to express my gratitude to everyone who encouraged and supported me during my time as a doctoral candidate at the Institute of Flight Systems at the University of the Bundeswehr Munich.

None of this would have been possible without my supervisor, Univ.-Prof. Dr.-Ing. Peter Stütz. I thank him for his trust, for believing in the value of my work, for backing me throughout the completion of this thesis, and for relying on me to handle tasks beyond the thesis itself.

I also want to thank my closest colleagues and friends, David Nospes and Linda Eckel, for cheering me up during difficult periods, for the good times we shared, and for the many technical discussions.

Finally, I thank my parents for always being there, as well as all the students and colleagues who contributed to this work, especially during the data collection campaigns.

Tobias Hupel
November 2025

Abstract

Multispectral imaging systems extend the capabilities of traditional sensor payloads on small tactical reconnaissance drones beyond the visual spectrum, enabling the collection of information that is inaccessible to human perception. This capability could be a significant advantage in camouflage detection, as camouflage typically relies on visual deception but may be ineffective in other spectral ranges. However, this valuable information comes at the cost of an increased processing workload that operating personnel would have to manage. Therefore, deploying multispectral sensors on tactical reconnaissance UAVs could prove more effective when paired with advanced methods to streamline information processing and support operators. In order to investigate this potential, the Institute of Flight Systems at the University of the Bundeswehr Munich conducted four studies in recent years that examine the capabilities and limitations of multispectral imaging systems for camouflage detection in tactical reconnaissance scenarios. In addition to exploring computer-aided detection methods, the studies focus on managing the large numbers of sensor bands in order to improve detection rates and minimize the associated workload. More specifically, sensor performance modelling techniques that identify the most valuable sensor bands under varying environmental conditions are conceptualized, implemented, and evaluated for their effectiveness in enhancing camouflage detection. This thesis organizes these studies into a cumulative body of research by summarizing and linking their major motivations, methodologies, experiments, and key results.

Throughout the conducted studies, several novel approaches have been proposed, and a variety of methodologies have been explored. The first study addresses the problem of camouflage detection in general. It shows that spectral anomaly detection is an effective and efficient method for identifying camouflaged targets in multispectral imagery. In order to manage the increased complexity inherent to multispectral sensors, the second study proposes and evaluates a lightweight sensor performance modelling approach. Using this novel method, the performance of each sensor band in terms of camouflage detection can be predicted based on current environmental conditions, allowing the generation of a subset of only the most valuable sensor bands. By considering only this subset in any subsequent application, the complexity associated with the multispectral imaging system can be substantially reduced. The third study builds on the first two studies by investigating a combination of their methodologies, called sensor-managed anomaly detection. In this novel approach, only the most valuable sensor bands

are processed with the detection algorithms instead of all available sensor bands. Although this method does not significantly improve detection rates, it offers considerable potential for conserving resources by reducing the amount of data to be processed while maintaining performance, which could be particularly useful given the constrained resources on small tactical drones. The final study expands upon the second study by integrating sensor performance modelling into an optimization approach that generates novel spectral indices specifically designed to expose camouflaged targets. Combined with the spectral anomaly detection methods explored in the first study, these optimized indices provide significant improvements in camouflage detection performance. Furthermore, given the inherently low computational overhead associated with such indices, the proposed approach could prove particularly valuable in tactical reconnaissance scenarios. In summary, the presented approaches demonstrate the high utility of multispectral imaging systems aboard tactical drones for camouflage detection while considering the constraints imposed by the limited computational resources.

However, while the results show the effectiveness and indicate promising future applications of the proposed methods, it is important to note that they are based on relatively limited data sources, which might affect their generalizability and reliability. Therefore, further validation in a broader context and with additional data is advised to ensure their validity in different settings. Future research directions could include exploring faster and more efficient spectral anomaly detection methods, exploiting sophisticated and compact hardware for incorporating deep neural networks, and considering hyperspectral imaging systems as sensor technology evolves.

Kurzfassung

Multispektrale Bildgebungssysteme erweitern herkömmliche Sensornutzlasten auf taktischen Aufklärungsdrohnen über das sichtbare Spektrum hinaus und ermöglichen die Erfassung von Informationen, die für das menschliche Auge nicht zugänglich sind. Diese zusätzlichen Informationen könnten erhebliche Vorteile bei der Detektion von Tarnmitteln mit sich bringen, da Tarnmittel typischerweise nur auf visueller Täuschung beruhen und daher in anderen Spektralbereichen möglicherweise ihre Wirkung verlieren. Allerdings ist mit der Zunahme verfügbarer Daten auch ein erhöhter Verarbeitungsaufwand verbunden, den das Betriebspersonal zu bewältigen hätte. Daher könnte sich der Einsatz von Multispektralsensoren in taktischen Aufklärungsdrohnen als effektiver erweisen, wenn sie mit fortschrittlichen Methoden zur Aufbereitung von Informationen und zur Unterstützung der Operateure kombiniert werden. Um dieses Potenzial zu untersuchen, hat das Institut für Flugsysteme an der Universität der Bundeswehr München in den letzten Jahren vier Studien durchgeführt, in denen die Fähigkeiten und Grenzen multispektraler Bildgebungssysteme zur Tarnungserkennung in taktischen Aufklärungsszenarien untersucht wurden. Neben der Erforschung computergestützter Detektionsmethoden konzentrieren sich die Studien auf das Management der hohen Anzahl von Spektralbändern, um die Erkennungsraten zu verbessern und die damit verbundene Arbeitsbelastung zu minimieren. Im Detail werden Methoden zur Modellierung von Sensorleistung, die die geeignetsten Spektralbänder unter verschiedenen Umgebungsbedingungen identifizieren, konzipiert, implementiert und auf ihre Effektivität bei der Detektion von Tarnmitteln untersucht. Diese Arbeit ordnet die Studien in einen kumulativen Forschungsrahmen ein, indem sie ihre grundlegenden Motivationen, Methoden, Experimente und wichtigsten Ergebnisse zusammenfasst und miteinander verknüpft.

Im Rahmen der durchgeführten Studien wurden mehrere neue Lösungsansätze entwickelt und eine Vielzahl von Methoden ausgewertet. Die erste Untersuchung beschäftigt sich mit der Detektion von Tarnmitteln im Allgemeinen. Sie zeigt auf, dass die spektrale Anomaliedetektion eine effiziente Methode zur Aufklärung getarnter Objekte in multispektralen Daten darstellt. Um die erhöhte Komplexität multispektraler Sensoren zu bewältigen, wird in der zweiten Studie ein effizienter Ansatz zur Modellierung der Sensorleistung entwickelt und evaluiert. Konkret ermöglicht diese Methode die Vorhersage der Leistung jedes Spektralbandes in Bezug auf die Detektion von Tarnmitteln und in Abhängigkeit aktueller Umgebungsbedingungen. Dadurch kann die Menge aller Spektralbänder in nachfolgenden Verarbeitungsinstanzen auf diejenigen

reduziert werden, die die höchste Leistung oder den größten Nutzen erbringen, was die Komplexität in der Verarbeitung multispektraler Daten erheblich verringern kann. Die dritte Studie baut auf den ersten beiden Studien auf und untersucht eine Kombination ihrer Methoden. Dabei werden nicht alle verfügbaren Spektralbänder von den Anomaliedetektoren verarbeitet, sondern nur die Spektralbänder mit der höchsten vorhergesagten Leistung. Dieser Ansatz führt zwar nicht zu einer signifikanten Verbesserung der Detektionsleistung von Tarnmitteln, bietet aber ein erhebliches Potenzial zur Ressourcenschonung. Denn durch die geringere Anzahl an Spektralbändern reduziert sich auch die zu verarbeitende Datenmenge, was insbesondere bei den begrenzten Ressourcen auf taktischen Aufklärungsdrohnen von Vorteil sein kann. Die vierte Studie entwickelt und untersucht eine Methode, die zur Optimierung von Spektralindizes dient. Dazu werden Konzepte aus der Leistungsmodellierung aus der zweiten Studie adaptiert, um neue Spektralindizes speziell für die Detektion von Tarnmitteln zu bestimmen. In Kombination mit den Methoden zur spektralen Anomaliedetektion aus der ersten Studie führen diese Indizes zu signifikanten Verbesserungen bei der Detektion getarnter Objekte. Darüber hinaus sind solche Indizes mit einem äußerst geringen Rechenaufwand verbunden, was den Optimierungsansatz besonders attraktiv für die taktische Aufklärung macht. Zusammenfassend zeigen die vorgestellten Ansätze deutliche Vorteile für den Einsatz multispektraler Bildsysteme zur Detektion von Tarnmitteln auf taktischen Aufklärungsdrohnen. Zudem berücksichtigen sie die Einschränkungen durch begrenzte Rechenressourcen mit dem Einsatz effizienter Algorithmen.

Allerdings ist darauf hinzuweisen, dass die erzielten Ergebnisse auf einer relativ begrenzten Datenmenge beruhen, was ihre Generalisierbarkeit einschränken kann. Um die Zuverlässigkeit der Methoden in anderen Szenarien zu gewährleisten, sollten daher weitere Experimente mit anderen oder zusätzlichen Daten durchgeführt werden. Zukünftige Forschungsarbeiten könnten sich mit der Untersuchung schnellerer und effizienterer Verfahren zur Anomalieerkennung, der Integration tiefer neuronaler Netze durch moderne Hardware oder der Erprobung miniaturisierter hyperspektraler Bildgebungssysteme befassen.

Contents

1 Introduction **1**

2 Research Scope **7**

3 Experiments **11**

3.1 Anomaly Detection for Camouflage Detection 11

3.1.1 Approach 11

3.1.2 Dataset: MUCAD 12

3.1.3 Spectral Anomaly Detectors 13

3.1.4 Results 17

3.2 Sensor Performance Modelling for Camouflage Detection 19

3.2.1 Approach 19

3.2.2 Dataset: MUDCAD-X 20

3.2.3 Measuring and Predicting Sensor Performance 20

3.2.4 Results 22

3.3 Sensor-Managed Anomaly Detection for Camouflage Detection 25

3.3.1 Approach 25

3.3.2 Results 26

3.4 Sensor Performance Modelling for Spectral Index Optimization 26

3.4.1 Approach 28

3.4.2 Spectral Index Optimization 29

3.4.3 Results 30

4 Conclusions and Future Prospects **33**

5 Publications **37**

5.1 Adopting Hyperspectral Anomaly Detection for Near Real-Time Camouflage Detection in Multispectral Imagery 37

5.2 Measuring and Predicting Sensor Performance for Camouflage Detection in Multispectral Imagery 59

5.3 Sensor-Managed Anomaly Detection for Camouflage Detection in Airborne Multispectral Imagery 81

5.4 Optimized Spectral Indices for Camouflage Detection in Multispectral Imagery . 93

References **127**

List of Abbreviations

AED	attribute and edge-preserving filter detector
AED-F	AED with additional filtering
AP	attribute profile (morphological filtering)
AUC	area under curve
BNDVI	blue normalized difference vegetation index
CRD	collaborative-representation-based detector
EIR	edge-infrared (also known as red edge)
ϵ-SVR	ϵ-Support Vector Regression
GNDVI	green normalized difference vegetation index
GSD	ground sample distance
IOU	intersection over union
LPD	local point density
LRI	linear ratio index
LRI_2	LRI based on two bands
LRI_6	LRI based on six bands
LRI_{nd}	normalized difference index expressed by a LRI
LRI_r	ratio index expressed by a LRI
LRX	local Reed-Xiaoli
LWIR	long-wave infrared
MUCAD	Multispectral Dataset for Camouflage Detection
MUDCAD-X	eXtended Multispectral Dataset for Camouflage Detection
NIR	near-infrared
NDRE	normalized difference red edge index
NDVI	normalized difference vegetation index
RF	Random Forests
RX	Reed-Xiaoli
ROC	receiver operating characteristic

SWaP-C	size, weight, power and cost
TVI	Target Visibility Index
UAV	unmanned aerial vehicle
VIS	visual
XGBoost	eXtreme Gradient Boosting

1 Introduction

Unmanned aerial vehicles (UAVs) paired with advanced imaging systems have revolutionized remote sensing operations over the past decade. Their low operational requirements compared to manned aircraft, ability to cover large areas in a short time, and high utility have opened up a variety of new research opportunities and led to unprecedented amounts of data [28, 53]. One of the key technologies that is driving this research is multispectral imaging. Multispectral sensors are sensitive to multiple different spectral regions that can be within or beyond the visual spectrum, thereby producing an image with great spectral depth. Unlike visual sensors, these spectral regions do not necessarily correspond to human perception, but to the requirements of the application for which they are designed. Therefore, when used in conjunction with UAVs, multispectral sensors allow insights into surface properties on a large scale that are inaccessible to conventional visual sensors and otherwise only accessible through tedious and potentially destructive on-site probing. For example, healthy green vegetation reflects much of the incoming light in the near-infrared spectrum between 700 nm and 900 nm, while it shows comparatively strong absorption in the blue spectrum at 450 nm and in the red spectrum at 670 nm [43]. As a result, these spectral regions are widely used in precision agriculture or vegetation monitoring applications. Figure 1.1 shows a UAV weighing about six kilograms with several multispectral sensor payloads that can be used for data collection in such use cases. It can be operated without extensive or specialized training and typically requires only an online test for a civil piloting license.

Recent research has explored the potential of multispectral imaging in a wide range of remote sensing applications, including pest [50] and invasive species detection [7], water body extraction [27], land cover classification [16, 29, 2], as well as crop yield [14, 54], chlorophyll content [14, 1], leaf area index [14, 52], and water quality [51] estimation. In these scenarios, it is common practice to combine the individual spectral regions captured by multispectral sensors using spectral indices to create images derived from multiple raw images. The resulting index images typically provide a direct indication of certain surface properties and are therefore immediately interpretable, which can also improve the prediction of more complex surface properties by machine learning models, such as canopy nitrogen content [48]. For example, the near-infrared spectrum at about 850 nm and the red spectrum at about 700 nm are used to calculate the normalized difference vegetation index (NDVI) [36], which is a common indicator of plant

(a) airborne (b) on ground

Figure 1.1: UAV with multispectral sensor payloads. The left image shows an airborne UAV equipped with multiple multispectral sensors, while the right image shows the same UAV on the ground with a single downward-facing multispectral sensor, providing six different spectral bands (annotated with an orange circle).

(a) VIS (b) NIR (c) red (d) NDVI

Figure 1.2: Multispectral capture with visual, red and NIR spectral bands of a sample scene. The NDVI indicates healthy green vegetation with green areas, while red areas most likely represent non-vegetation.

health. Figure 1.2 shows a multispectral capture consisting of the visual, red and NIR spectral bands along with the corresponding NDVI image. Note that the NDVI image is color-coded with green for high values and red for low values, indicating healthy green vegetation and other surface types, respectively. As can be observed, the green vegetation reflects much of the incoming light in the near-infrared band and much less in the red band. In addition, the NDVI shows high values for the dense green canopy of the trees and significantly lower values for the asphalt of the road and other regions with less developed green vegetation.

With the ability to capture spectra beyond the range of human perception, multispectral imaging systems may prove to be of great value in military reconnaissance scenarios. In particular, they can provide information that mitigates the effects of camouflage materials, which are typically designed to deceive the human eye. Camouflage materials are commonly employed to conceal military vehicles, personnel or positions by mimicking their immediate surroundings

(a) 2D camouflage net (b) 3D camouflage net

Figure 1.3: Two different camouflage nets that blend into their surroundings. They are commonly used to hide objects that would otherwise be easily visible.

with similar-looking but different materials, such as green camouflage nets deployed in forest areas. Examples of such materials are shown in Figure 1.3. Here, two different camouflage nets are placed in an environment where they blend in very well.

Considering currently operated small tactical UAVs in military reconnaissance missions, such as ALADIN or MIKADO of the German Armed Forces, sensor payloads usually consist of conventional visual (VIS) or long-wave infrared (LWIR) imaging systems. Although the LWIR spectrum is beyond the visual spectrum like the spectra provided by multispectral sensors, the information captured is typically limited to a single channel integrating the electromagnetic spectrum from 8 µm to 14 µm. This is the spectral region where materials emit the most radiation at regular ambient temperatures, therefore providing substantial information about the temperature of objects and materials. However, multispectral sensors extend or exceed these capabilities by capturing multiple spectral regions simultaneously, while maintaining a similar low size, weight, power and cost (SWaP-C) footprint. As a result, their use on small tactical UAVs could significantly improve the success of reconnaissance missions by exposing camouflaged targets more effectively than conventional VIS and LWIR imaging systems. For example, Figure 1.4 shows the VIS, LWIR and NIR spectral regions of the same scene with an additional label mask identifying two different camouflaged targets. While the VIS band is of little help for detection, the target in the lower left can be easily detected in the LWIR band, but not the other one. However, in the NIR band, both targets are clearly visible as very dark spots, which demonstrates the enormous potential of multispectral imaging for camouflage detection.

Published research also suggests that multispectral imaging is particularly effective in detecting camouflaged targets. For instance, successful detection has been achieved using satellite imagery [4] as well as specially designed ground-based military sensor systems [3, 12]. Furthermore, advanced camouflage detection algorithms based on transformers [46] or constrained energy

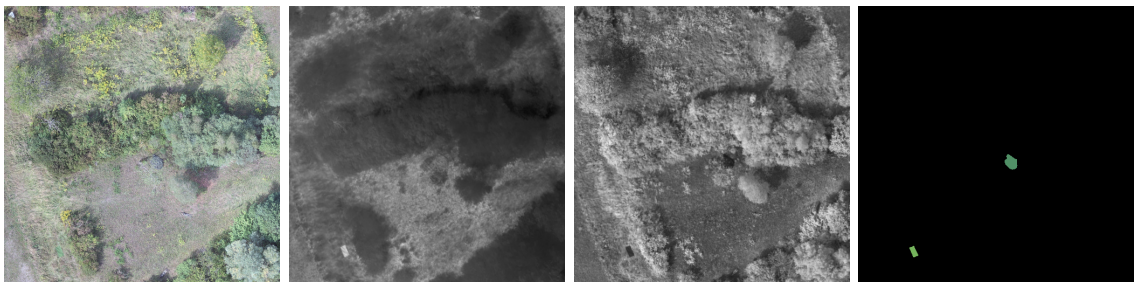

(a) VIS (b) LWIR (c) NIR (d) label

Figure 1.4: Multispectral capture with visual, LWIR and NIR spectral bands of a sample scene. The last image identifies two different camouflaged targets in the scene. Both targets are easily detected in the NIR band (two very dark spots), while the LWIR band only helps with the lower left target (one bright spot). Naturally, the visual band is of limited use for either target.

minimization combined with adaptive thresholding [39] have been explored. In the context of airborne multispectral imagery with a typical nadir perspective, camouflage detection has been investigated using anomaly detection [26] and target detection [30]. Moreover, there are indications that the incorporation of spectral indices specifically designed for camouflage detection can further enhance the performance of anomaly detection methods [26].

Although multispectral sensors promise great benefits in military reconnaissance scenarios by providing additional spectral information, their exploitation also imposes a significant workload on the operating personnel, specifically for small tactical UAVs, where the sensor operator and the drone pilot are usually the same person. Compared to single-image visual or thermal sensor streams, multispectral sensor streams are composed of multiple image streams and are therefore inherently more difficult to process, potentially overwhelming operators. Consequently, the use of multispectral sensors on tactical reconnaissance UAVs may be more effective if accompanied by higher levels of automation that can optimize the flow of information, assist operators, and prevent unmanageable workloads. As a first step, these concerns can be addressed by automating the process of detecting camouflaged targets using computer-aided detection algorithms. More advanced approaches may involve the integration of performance models that help determine the relevant and irrelevant sensor bands of a multispectral imaging system. In general, performance models represent the capabilities of particular components, such as hardware components or software algorithms, by relating current situational or environmental conditions to a measure of the performance provided by these components with respect to a specific task. In the context of tactical reconnaissance scenarios using multispectral sensors,

these models would continuously evaluate the capabilities of each sensor band. With such an approach, relevant information can be emphasized, while irrelevant information can be omitted. This may improve camouflage detection rates and reduce workload at the same time.

Over the past few years, the concept of performance modelling has been extensively studied by the Institute of Flight Systems at the University of the Bundeswehr Munich in various contexts. For example, Bayesian networks [37] and artificial neural networks [17] have been successfully explored to predict the most appropriate object detector onboard a UAV under varying environmental circumstances, thereby increasing overall detection performance over a single detector. Similarly, sensor performance models have been combined with optimal control to generate optimized UAV trajectories that maximize the performance of object detection algorithms [55]. In another application dealing with multi-target tracking using visual sensors on UAVs [23, 24], performance models estimate the probability of successful object detections given a context consisting of sensor parameters and environmental as well as aircraft conditions. This novel approach maximizes the observation times of all targets compared to traditional techniques. Furthermore, random forests have been effectively applied to predict the most appropriate sensor bands of an airborne hyperspectral sensor given an environmental context, improving the detection of camouflaged targets and unexploded ordnance [10, 9].

In view of the potential benefits as well as the significant challenges, the Institute is also researching the utility of multispectral sensors onboard small tactical UAVs for camouflage detection in reconnaissance scenarios, with a strong focus on sensor performance modelling techniques. As part of this effort, several complementary studies [18, 19, 21, 20] have been conducted in recent years. This thesis builds a framework around these studies by connecting and highlighting their motivations, methodologies, and findings.

The following chapters of this thesis are organized as follows. First, Chapter 2 presents the individual areas of research, derived research questions, and interrelationships of all studies, forming the cumulative research scope of this work. Then, Chapter 3 summarizes the conducted experiments addressing the research questions and the obtained results across all studies, including a brief introduction to the employed materials and methods. Conclusions and future prospects are subsequently presented in Chapter 4. Finally, Chapter 5 contains the list of publications forming the foundation of this thesis.

2 Research Scope

This thesis explores the utility of multispectral imaging onboard small tactical UAVs for camouflage detection in military reconnaissance scenarios through four complementary studies. Compared to current VIS and LWIR sensor payloads, these imaging systems provide much more spectral information in the form of additional sensor bands, which increases their general utility but also the associated processing workload. This increased workload risks exceeding the capacity of the operating personnel, especially in the dynamics of military reconnaissance scenarios. Therefore, the studies explore approaches that automate the camouflage detection process by incorporating detection algorithms and advanced methods for managing the additional information while exploiting it to the fullest extent possible.

Considering that all materials have unique transmission, reflection, and absorption rates across the electromagnetic spectrum, it can be assumed that camouflage materials cannot completely mimic the spectral characteristics of their surroundings. As a result, these materials may appear as anomalous spectral signatures that can be captured by multispectral imaging systems, which sample the electromagnetic spectrum at many different wavelengths. Consequently, the use of spectral anomaly detection methods has strong potential for detecting camouflaged targets in multispectral imagery. Although target detection could be a more accurate approach than anomaly detection, given the distinctive and unique signatures of camouflage materials, it requires known target signatures in the detection process. Since knowledge of target signatures cannot be guaranteed and should not be assumed in military reconnaissance scenarios, anomaly detection appears to be more appropriate than target detection in such use cases. Therefore, the first study addresses the elementary problem of camouflage detection using multispectral imaging by investigating the following research question.

Research Question 1: *How effective is spectral anomaly detection for camouflage detection in multispectral imagery?*

Although the information provided by multispectral imaging systems offers significant value through the large and diverse set of captured spectral regions, it also imposes an inherently high workload for any subsequent processing instance. For example, sensor operators of small tactical UAVs are usually confronted with a single VIS or LWIR image stream, whereas multispectral image streams introduce many more sensor bands and thereby image streams to evaluate.

This workload may drive operators beyond their processing capabilities, ultimately rendering the additional information useless if it cannot be utilized. In addition, camouflage detection algorithms stand to benefit if the underlying sensor data to be processed is composed of sensor bands that are actually beneficial for the detection task. Since sensor bands correspond to specific spectral regions, each with its own unique physical characteristics, some are likely more useful for camouflage detection than others. Similarly, some sensor bands may even be disadvantageous, given the right environmental circumstances, which are constantly changing in the course of military reconnaissance scenarios. Consequently, there appears to be a strong and learnable relationship between the suitability of a sensor band for camouflage detection and the current environmental conditions. Therefore, an approach that determines the most suitable set of sensor bands for camouflage detection by modelling this relationship may reduce workload and increase detection performance at the same time.

This problem is specifically addressed by the concept of sensor performance modelling. Here, sensor performance models continuously evaluate the value, suitability, or performance of each sensor band of a multispectral imaging system with respect to camouflage detection. This evaluation is based on a context state that describes current environmental conditions, like landscape or lighting conditions, as these influence the spectral information that is captured by each sensor band. With a rating of performance for each sensor band, a subset can be generated that contains only the most suitable sensor bands for camouflage detection. Naturally, such an approach requires methods for describing sensor performance, the environmental context itself, and methods for modelling the relationship between these two, i.e. performance models. An implementation of such methods is evaluated in the second study, which explores the concept of sensor performance modelling in the context of this thesis by addressing the following research question.

Research Question 2: *How effective is sensor performance modelling for identifying the most suitable sensor bands for camouflage detection in multispectral imagery?*

Assuming that sensor performance modelling enables the identification of the most suitable sensor bands for camouflage detection, it remains unclear how these bands affect the actual performance of detection algorithms. In particular, processing only this subset rather than all available sensor bands may improve their detection and false alarm rates, since less suitable and possibly confusing bands are ignored. This can be particularly useful on resource-constrained platforms, such as small tactical reconnaissance UAVs, as there are fewer bands to process, which reduces computational complexity. Given these potential benefits, the third study explores such an approach by combining sensor performance modelling and spectral anomaly detection,

called sensor-managed anomaly detection. In this approach, the sensor bands to be processed are dynamically selected or managed by performance models. Consequently, the third study investigates the following research question.

Research Question 3: *How effective is sensor-managed anomaly detection for camouflage detection in multispectral imagery?*

Using the concept of sensor performance, the suitability of any sensor band for camouflage detection can be evaluated. This applies to any available raw sensor band, as well as to any derived or processed sensor band. Common types of derived sensor bands are spectral indices. They are widely studied and used in various remote sensing applications, and combine multiple raw sensor bands in a particular predefined way. Therefore, there may also be a spectral index that is particularly beneficial in detecting camouflaged targets. With sensor performance modelling, the actual utility of any spectral index can be measured, allowing the determination of the most suitable one. As a result, a spectral index specifically designed for camouflage detection should be generated in such a way that it maximizes sensor performance.

In tactical reconnaissance scenarios, a spectral index geared towards camouflage detection offers clear advantages, especially in view of the typical minimal computational burden. For example, operators can focus on evaluating a single-image stream, the spectral index, thereby drastically reducing workload while benefiting from more spectral information than is available in any raw sensor band. In addition, processing such an index together with all other sensor bands may improve the performance of detection algorithms. Therefore, the fourth study of this thesis explores an approach that employs sensor performance modelling to find and optimize spectral indices specifically tailored for camouflage detection. Thus, it addresses the following fourth and final research question of this thesis.

Research Question 4: *How effective is spectral index optimization in combination with sensor performance modelling for camouflage detection in multispectral imagery?*

3 Experiments

In this section, the studies encompassing the experiments that address the four research questions of this thesis are summarized. This includes their general approach, the materials and methods used, the evaluation metrics employed, the results obtained, and a discussion of their significance in relation to the research questions. All experiments and computations were conducted on a system running Ubuntu, equipped with an AMD Ryzen 9 3950X processor (16C/32T) and 128 GB of RAM.

3.1 Anomaly Detection for Camouflage Detection

This section summarizes the study [18] addressing the research question *"How effective is spectral anomaly detection for camouflage detection in multispectral imagery?"*. The approach is outlined in Section 3.1.1 while the results are shown and discussed in Section 3.1.4. In addition, the employed dataset and anomaly detection methods are introduced in Section 3.1.2 and Section 3.1.3, respectively. A full reprint of the corresponding publication can be found in Section 5.1.

3.1.1 Approach

Spectral anomaly detection is a highly active field of research, especially in the hyperspectral domain, which is similar to multispectral imaging but with many more spectral bands [34, 42, 31]. Compared to multispectral sensors, hyperspectral sensors are much more expensive in terms of money and required computing power, making them a less suitable payload for small tactical UAVs. However, methods designed for hyperspectral data may also be effective for multispectral data given the strong similarities. Therefore, this study employs hyperspectral anomaly detection to investigate the effectiveness of anomaly detection for camouflage detection in multispectral imagery. For this purpose, the following detectors are implemented: the Reed-Xiaoli detector (RX) [35], local Reed-Xiaoli detector (LRX) [35], the local point density detector (LPD) [18], the attribute and edge-preserving filter detector (AED) [22], the AED with additional filtering detector (AED-F) [18] and the collaborative-representation-based detector (CRD) [25]. Each relies on completely different detection mechanisms, which minimizes anomalous results due

to method-specific biases or errors, thereby increasing the robustness of this investigation. In order to evaluate them for camouflage detection in multispectral imagery, they are applied to a custom multispectral dataset specifically created for this study, the Multispectral Dataset for Camouflage Detection (MUCAD). It consists of multiple different scenes, each containing at least a single camouflaged target. The effectiveness of the detection algorithms is assessed over the entire dataset using receiver operating characteristic (ROC) and area under curve (AUC). These metrics are often used to measure detection performance, as they allow evaluation of detection maps without a specific binarization threshold.

Since hyperspectral anomaly detection methods are not necessarily designed to be resource-efficient, they could be very slow given the limited computing capabilities of small tactical UAVs. Therefore, the detectors are also evaluated for their runtimes to get an indication of their computational efficiency.

In addition to the evaluation based on MUCAD, the detectors are also evaluated based on MUCAD-VI, which is an extended version of MUCAD and contains some extra bands derived from the original raw bands. These bands are the blue normalized difference vegetation index (BNDVI), the green normalized difference vegetation index (GNDVI) and the normalized difference red edge index (NDRE). During preliminary and exploratory experiments, these indices appeared to enhance the perceived visibility of the camouflaged targets contained in MUCAD. Therefore, they may also improve the performance of the detection algorithms compared to the plain MUCAD. Consequently, this additional evaluation provides an initial indication of the utility of spectral indices for camouflage detection, which is investigated in more detail in the fourth study of this thesis. The selected indices are calculated using the blue, green, edge-infrared (EIR) and near-infrared (NIR) bands, as shown below.

$$\text{BNDVI} = \frac{\text{NIR} - \text{blue}}{\text{NIR} + \text{blue}} \quad \text{GNDVI} = \frac{\text{NIR} - \text{green}}{\text{NIR} + \text{green}} \quad \text{NDRE} = \frac{\text{NIR} - \text{EIR}}{\text{NIR} + \text{EIR}} \tag{3.1}$$

3.1.2 Dataset: MUCAD

The Multispectral Dataset for Camouflage Detection (MUCAD) is intended to closely resemble footage that could have been captured by an actual small tactical reconnaissance drone equipped with a multispectral imaging system. Therefore, its collection was carried out with two modern multispectral cameras, a MicaSense Altum and a Zenmuse XT2, mounted on a DJI Matrice 210 RTK V2, a UAV whose size is close to that of the actual tactical reconnaissance UAVs, such as ALADIN or MIKADO. The data acquisition flights for MUCAD were conducted at the test site of the University of the Bundeswehr Munich. It features multiple terrain types with varying vegetation and characteristics, enabling the integration of many different camouflaged targets. Figure 3.1 shows the areas of the test site that were captured during the acquisition flights. As

can be seen, there are a number of different environments, such as grassland, woodland, gravel, concrete and roads, providing excellent conditions for creating extensive and versatile datasets. This allowed for the placement and eventual capture of multiple different camouflaged targets in areas where they blended in almost seamlessly.

In total, MUCAD features 23 different scenes and 8 different camouflaged targets. Each scene has a resolution of 512 px by 512 px, corresponding to a ground sample distance (GSD) of 10 cm px^{-1}. Additional properties of the dataset, such as the target classes it contains, are shown in the middle column of Table 3.1. Note that the right column shows the properties of another dataset, MUDCAD-X, which was compiled in a subsequent study but has much in common with MUCAD.

Each sample of MUCAD consists of a multispectral capture and an accompanying label mask that identifies all camouflaged targets in the scene. An example is given in Figure 3.2, where Figures 3.2a to 3.2g show the individual spectral bands and Figure 3.2h shows the label mask. The spectral properties of each band are given in Table 3.2. As can be observed, there are three single-channel bands in the visual spectrum and three single-channel bands in the infrared spectrum. In addition, there is a visual band, which captures the scene close to the human eye and is technically composed of three individual bands. However, since these bands are generally considered as a single unit, they are also treated as such in this thesis.

Although a comparatively small dataset, MUCAD contains a variety of camouflaged targets in many different environments. Considering that hyperspectral anomaly detection methods are typically evaluated using a few single-scene datasets [44, 45, 49, 8], these conditions should provide a sufficient basis for evaluating the effectiveness of spectral anomaly detection methods for camouflage detection.

3.1.3 Spectral Anomaly Detectors

The selection of hyperspectral anomaly detection methods is primarily motivated by their expected computational requirements and their prominence in the public literature, resulting in the implementation of six different algorithms: the RX [35], LRX [35], LPD [18], AED [22], AED-F [18] and CRD [25] detector. Each of them uses entirely different principles to characterize anomalous and non-anomalous pixels in an image. Therefore, the following brief introductions to these algorithms focus on their unique detection mechanisms. All algorithms are implemented and parallelized in `C/C++` to ensure maximum runtime performance.

As one of the most dominant and classic algorithms, the RX detector assumes that background pixels originate from a normal distribution. Consequently, pixels with very low likelihood under that distribution are considered anomalous. Based on this principle, the anomaly score of a single pixel x in an image is equal to its Mahalanobis distance d_{M}, calculated according to

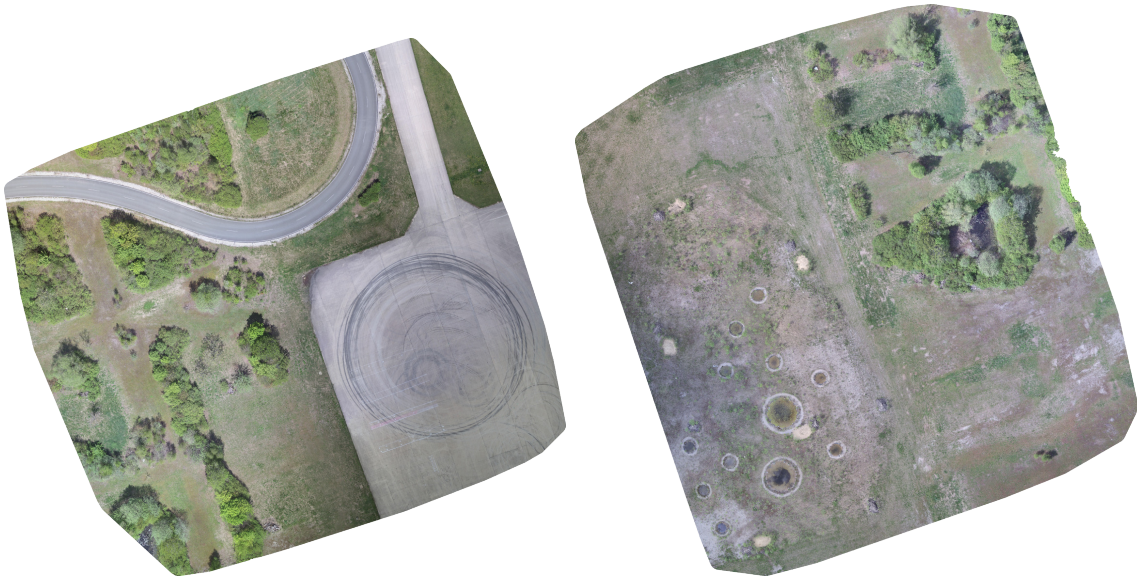

Figure 3.1: The areas of the test site of University of the Bundeswehr Munich captured in the data collection campaigns for MUCAD and MUDCAD-X.

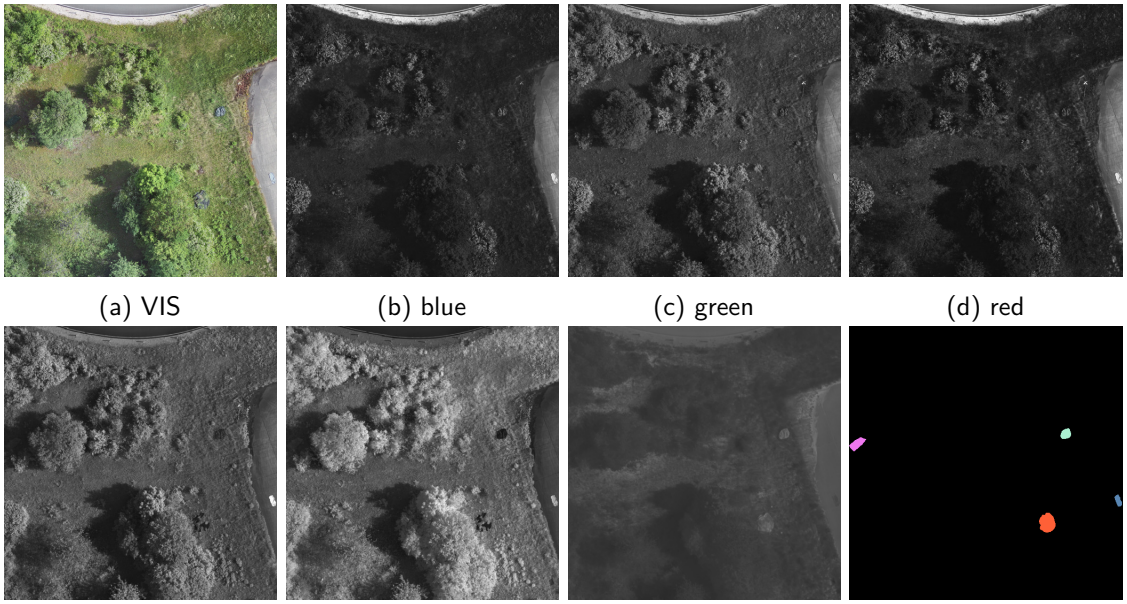

(a) VIS (b) blue (c) green (d) red

(e) EIR (f) NIR (g) LWIR (h) label

Figure 3.2: All bands from VIS (a) to LWIR (g) contained in a single sample, either from MUCAD or MUDCAD-X. The last image shows the label mask (h) of the scene, identifying four different camouflaged targets in this case.

Table 3.1: Basic properties and camouflaged targets of MUCAD and MUDCAD-X.

	MUCAD	MUDCAD-X
Properties		
# camouflaged target classes	8	10
# samples	23	853
seasons	summer	spring, summer, fall
bands	VIS, blue, green, red, EIR, NIR, LWIR	
resolution	512 px × 512 px	
ground sample distance (GSD)	10 $\frac{\text{cm}}{\text{px}}$	
Camouflaged Target Classes		
gray car	✓	✗
anthracite fleece	✗	✓
artificial turf	✓	✓
artificial hedge	✓	✓
gray tarp	✓	✓
green tarp	✓	✓
green 2D camouflage net	✓	✓
green 3D camouflage net	✓	✓
gray 3D camouflage net	✗	✓
yellow 3D camouflage net	✗	✓
person in green uniform	✓	✓
person in yellow uniform	✗	✓

Table 3.2: Spectral characteristics of the bands contained in each sample of MUCAD and MUDCAD-X.

band	center	bandwidth
visual (VIS)	-	-
blue	475 nm	32 nm
green	560 nm	27 nm
red	668 nm	14 nm
edge-infrared (EIR)	717 nm	12 nm
near-infrared (NIR)	842 nm	57 nm
long-wave infrared (LWIR)	10.5 µm	6 µm

the formula given in Equation (3.2). Here, μ is the average value over all pixels and Σ is the corresponding covariance matrix.

$$d_{\mathrm{M}} = \sqrt{(x-\mu)^{\mathrm{T}} \Sigma^{-1} (x-\mu)} \tag{3.2}$$

Following the same idea, the LRX detector computes the Mahalanobis distance using an average value and a covariance matrix based only on the local neighborhood of a pixel under test. A dual window centered on the pixel under test defines this neighborhood. It is composed of an inner and an outer square window, where the neighborhood consists of all pixels that are inside the outer window but outside the inner window. This way, the background estimation is unique to each pixel and considers only its immediate surroundings, usually resulting in a more accurate estimation and thus better detection results than the RX detector. However, the potential improvements come at the cost of increased computational complexity, since the background estimation must be performed separately for each pixel.

Similar to the LRX detector, the LPD detector uses a local neighborhood based on a dual window to calculate anomaly scores. But instead of relying on the Mahalanobis distance, it defines anomalies in terms of density. The lower the density of a pixel, the higher its abnormality and vice versa. Here, the density ρ of a pixel x is computed as shown in Equation (3.3), where m_i is the i-th element of the set M containing the pixel under test and all pixels in the dual window neighborhood. s_M is the size of M and μ_M is the average value of all pixels in M. After applying this algorithm to an image, the resulting density map is easily converted to an anomaly map by negating each element and then adding the lowest value of all elements.

$$\rho = \frac{1}{s_M} \sum_{i=1}^{s_M} \exp\left(-\frac{\|x - m_i\|_2^2}{d_{\mathrm{c}}^2}\right) \tag{3.3}$$

$$d_{\mathrm{c}} = \frac{1}{s_M} \sum_{i=1}^{s_M} \|m_i - \mu_M\|_2^2 \tag{3.4}$$

In contrast to statistical approaches, CRD works on the assumption that non-anomalous pixels can be represented by a linear sum of their neighbors, while anomalous pixels cannot. As with the LRX and LPD detectors, the neighborhood of a pixel is defined by a dual window. Fortunately, the problem of collaboratively representing a pixel by its neighborhood has a closed-form solution. Therefore, the anomaly score a, which is equal to the reconstruction error, of a pixel x is calculated as given in Equation (3.5). Here, N is a matrix containing all pixels of the dual window neighborhood, I is an identity matrix of fitting size and λ is a regularization parameter that penalizes high contributions of individual neighboring pixels to

the reconstruction.

$$a = \|x - N\hat{\alpha}\|_2 \tag{3.5}$$

$$\hat{\alpha} = \left(N^{\mathrm{T}}N + \lambda I\right)^{-1} N^{\mathrm{T}}x \tag{3.6}$$

AED uses neither statistical nor representational approaches. Instead, it is based on the assumption that anomalous areas in an image are generally small and have significantly different intensity values compared to their surroundings. This is implemented by morphological attribute filters that decompose each channel of an image under test into two morphological attribute profiles (APs). The first has all bright connected components removed and the second has all dark connected components removed that are smaller than a predefined constant κ. Subtracting these two APs produces a difference map that highlights all areas of very low or high intensity areas smaller than κ. The difference map is additionally filtered by a special operation to retain only large differences. Finally, all difference maps are summed to generate an overall anomaly map, which is further smoothed using an edge-preserving filter. With additional filtering of each difference map, AED-F is a minor modification of AED that aims to suppress further false positives while maintaining correct detections.

3.1.4 Results

The detection algorithms are evaluated by applying them to each capture of MUCAD and MUDCAD-VI with multiple different parameter configurations. For the final results, only those parameter configurations are considered that yield the best detection result in terms of area under curve (AUC) per camouflaged target over all captures. This leaves one parameter configuration per detector, camouflaged target and dataset. For each of these configurations, the final receiver operating characteristic (ROC) and corresponding AUC are determined by threshold averaging [11] over all occurrences of the respective camouflaged target in the dataset.

The obtained AUCs for each detector are shown in the upper and lower halves of Table 3.3 for MUCAD and MUCAD-VI, respectively. Here, the rows represent the detectors with their respective parameter configuration and the columns represent the camouflaged targets. Note that values below 0.9 are not displayed because the algorithm is considered to be unable to detect the corresponding camouflaged target. As the results reveal, most targets are found by all detectors, with some AUCs even close to perfect detection, such as those for the green and gray tarp. However, certain targets seem to be difficult to detect for any algorithm, such as the hedge and the green persons. This may be due to the inability of the sensor data to

Table 3.3: Detection results in terms of Area under Curve (AUC) of hyperspectral anomaly detectors applied to MUCAD and MUCAD-VI. Columns and rows represent camouflaged targets and detectors, respectively. The upper half of the table shows the achieved AUC for MUCAD and the lower half for MUCAD-VI. Top results per camouflaged target are highlighted in bold.

	Hedge	Turf	Grn. Tarp	Gry. Tarp	Grn. 2D	Gry. 3D	Grn. Per.	Gry. Car
				MUCAD				
RX	-	0.968	0.977	0.995	0.954	-	-	0.957
LRX	0.906	0.996	0.993	**1.000**	**0.990**	**0.960**	0.963	0.961
AED	-	0.991	0.995	0.998	0.966	-	-	0.923
AED-F	-	0.997	0.997	0.999	0.973	-	-	0.936
LPD	**0.947**	**0.998**	0.988	0.980	0.973	-	**0.989**	0.902
CRD	-	0.997	**0.998**	1.000	0.989	0.955	-	**0.978**
				MUCAD-VI				
RX	0.943	0.995	0.980	0.993	0.959	0.906	-	0.984
LRX	**0.989**	**1.000**	0.995	**1.000**	**0.991**	0.972	0.966	0.975
AED	-	0.999	0.999	0.998	0.976	-	-	0.923
AED-F	-	1.000	**1.000**	0.999	0.982	-	-	0.939
LPD	0.956	0.999	0.988	0.978	0.976	0.946	**0.991**	-
CRD	-	1.000	0.997	**1.000**	0.985	**0.975**	0.903	**0.985**

highlight these targets as anomalous as required to be detected. The additional vegetation indices in MUCAD-VI help to counteract this problem and result in increased detection rates over all camouflaged targets and detectors.

In terms of runtime, the majority of algorithms returns in less than a second for any parameter configuration. Although they are executed on a very capable machine, these results suggest sufficient runtimes on the limited hardware of a small tactical reconnaissance drone. With the exception of the CRD detector, which can take up to several minutes, making it virtually impossible to use in real-world scenarios.

Overall, the obtained results demonstrate that spectral anomaly detection algorithms are able to effectively detect camouflaged targets in a multispectral context. However, it should be noted that there may be camouflaged targets that are harder to detect than others, depending on the degree of abnormality provided by the imagery. This suggests that successful detections depend more on the underlying sensor data than on the actual detection algorithm, as the detectors collectively tend to work either well or poorly for a specific camouflaged target. Additional bands derived from the original data, such as vegetation indices, may help address this issue, as the results indicate.

3.2 Sensor Performance Modelling for Camouflage Detection

In this section, the study [19] addressing the research question *"How effective is sensor performance modelling for identifying the most suitable sensor bands for camouflage detection in multispectral imagery?"* is summarized. The approach and achieved results are described in Section 3.2.1 and Section 3.2.4, respectively. In addition, Section 3.2 introduces the utilized dataset and Section 3.2.3 explains the novel sensor performance modelling approach. For a full reprint of the corresponding publication, see Section 5.2.

3.2.1 Approach

Exploring the concept of sensor performance modelling to identify the most suitable sensor band in multispectral imagery requires a proper implementation first. For this purpose, this study introduces a novel approach for measuring and predicting sensor performance in the context of camouflage detection. In particular, for each sensor band of a multispectral imaging system, a performance model continuously predicts its performance by evaluating the current environmental context. The context is represented by a context state, a feature vector extracted from the visual band by texture and image descriptors. For measuring sensor performance, a new metric is introduced, the Target Visibility Index (TVI). It provides a simplified but efficient description of the visibility of an object in a nadir image. With the context state as input and the TVI as output, practically any machine learning model with a compatible interface can be trained to predict the performance for a particular sensor band. Consequently, the resulting machine learning models form the performance models that dynamically provide the value of each sensor band of a multispectral imaging system.

The proposed approach is evaluated using a new dataset specifically created for this purpose, the eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X). This dataset is very similar to MUCAD, but with many more samples and diversity. It is divided into a training and a test dataset, which consist of approximately 80 % and 20 % of all samples, respectively. For the performance models, several different machine learning algorithms are employed and evaluated: ϵ-Support Vector Regression (ϵ-SVR) using LIBSVM [5], Random Forests (RF) using DecisionTree.jl [38] and Gradient Boosted Trees using eXtreme Gradient Boosting (XGBoost) [6]. They are based on entirely different modelling approaches and are generally characterized by relatively fast inference times, which can be critical for small tactical UAVs with limited resources. In order to fit these models to data, the context state is first extracted from the grayscale visual band of each sample in the training dataset. Then, it is mapped to the TVIs resulting from all other bands of the same sample, leading to six different mappings per sample. As a result, six separate performance models per algorithm are trained to learn these mappings

for each band over the entire training dataset. Once trained, the predictions of the models are evaluated using the test dataset. Note that if multiple camouflaged targets are present in the same sample, the calculated TVI per band used for training and testing results from an average over the individual TVIs resulting from each camouflaged target.

The evaluation on the test dataset is based on a comparison of the predicted and actual best-performing band order, rather than evaluating the performance models separately for their TVI predictions. This allows for an evaluation based on rank accuracy, which is a more suitable perspective on the predictions for addressing the research question. In addition, best-performing band orders are easier to comprehend and interpret than TVI predictions. The predicted best-performing band order is determined by ranking the bands according to their predicted TVIs provided by the performance models, while the actual best-performing band order is based on the calculated TVIs using the label masks.

3.2.2 Dataset: MUDCAD-X

Although MUCAD offers sufficient diversity for evaluating detection algorithms, training performance models requires a more extensive and versatile data foundation. Therefore, the acquisition of the eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X) was spread over an entire year including multiple different seasons. This resulted in a much larger and diverse dataset, which is more suitable for evaluating the proposed sensor performance modelling approach than MUCAD. Otherwise, MUDCAD-X and MUDCAD are closely related. Both were captured at the University of the Bundeswehr Munich test site using the exact same setup. Table 3.1 shows the properties of each dataset, with MUCAD in the middle column and MUDCAD-X in the right column. As can be observed, the main differences are the number of samples and the captured seasons. In addition, both feature a slightly different set of camouflaged targets. However, everything else is identical, including the structure of the samples, shown in Figure 3.2, and the characteristics of each spectral band, shown in Table 3.2.

3.2.3 Measuring and Predicting Sensor Performance

In this thesis, sensor performance modelling aims to estimate the relationship between an environmental or situational context and the performance or value of a specific sensor band with respect to a particular task. Since this research is about camouflage detection, the degree to which a sensor band exposes camouflaged targets defines its performance. For the quantification of this degree, this thesis introduces the Target Visibility Index (TVI). The TVI is an efficient metric that relies solely on simple statistics to describe the visibility of an object in a single-channel nadir image. It is computed as given in Equation (3.7), where μ_T is the average over all target pixels T, μ_B is the average over all background pixels B, σ_T is the standard

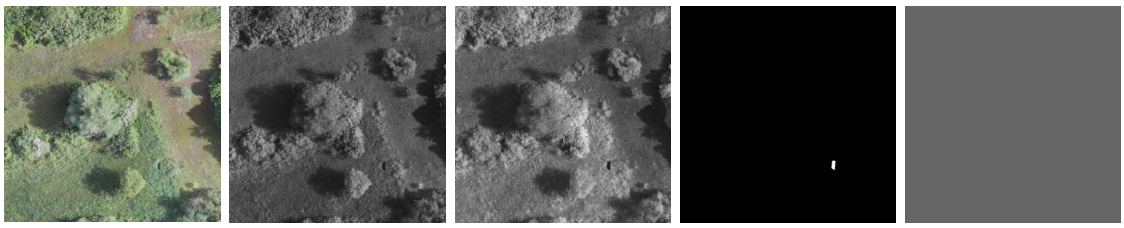

(a) VIS (7.1×10^{-6}) (b) EIR (0.0042) (c) NIR (0.013) (d) best (1.0) (e) worst (0.0)

Figure 3.3: Demonstration of the TVI, shown in parentheses, for different sensor bands of the same capture. The more a band exposes the camouflaged target in the scene, the higher the resulting TVI and vice versa.

deviation over all target pixels T, and σ_B is the standard deviation over all background pixels B.

$$\text{TVI} = (\mu_T - \mu_B)^2 (1 - 2\sigma_T)^2 (1 - 2\sigma_B)^2 \tag{3.7}$$

Note that the TVI is defined for images with a range of values between zero and one, for which it also ranges between zero and one, defining minimum and maximum visibility of the target, respectively. Figure 3.3 demonstrates the TVI for multiple different sensor bands of the same scene containing a camouflaged target. For each sensor band, the TVI is shown in parentheses. Ideal conditions are provided by a hypothetical band in Figure 3.3d, which perfectly highlights the target. The worst conditions are provided by a hypothetical band in Figure 3.3e, which shows a uniform color. Real sensor bands of the scene are given in Figures 3.3a to 3.3c. As can be seen, the TVI and therefore the performance is higher for the EIR and NIR bands, which expose the target more than the visual band, which is converted to a grayscale image before TVI calculation. Consequently, the TVI constitutes an appropriate metric for describing sensor performance in the context of camouflage detection, assuming that more exposed camouflaged targets are also easier to detect.

In addition to expressing sensor performance using the TVI, sensor performance modelling requires an approach to quantify the environmental context. Since the current environment is captured in the image produced by the sensor itself, the environmental context is derived directly from the image data. For this purpose, multiple efficient image or texture descriptors are employed, which are able to extract global features from an image: local binary patterns [32] and Haralick features [15]. Given an image, these descriptors generate a feature vector, representing an abstract encoding of the information contained in the scene, such as the structure of the depicted landscape. Although this encoding changes in different environments or scenes, it is designed to remain consistent in similar ones. In this thesis, this feature vector is extracted from the grayscale converted visual band, as these methods were originally developed for this kind of data. Consequently, the visual band and the corresponding feature vector are

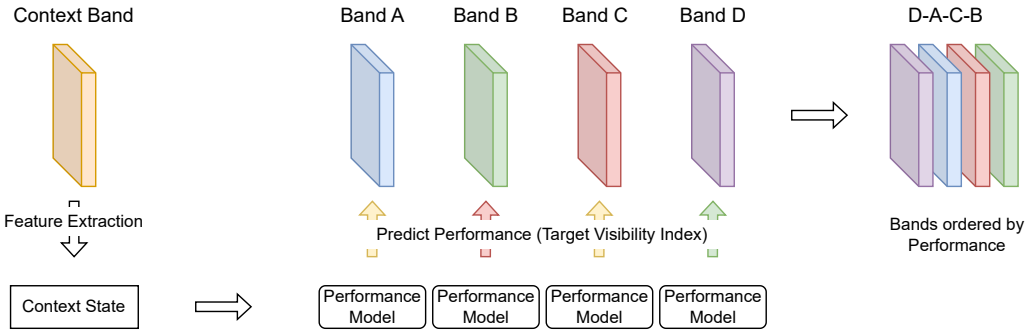

Figure 3.4: Sensor performance modelling approach. Each performance model predicts the performance from the context state for its associated sensor band. These predictions rank the sensor bands from best to worst for the current environmental conditions.

also referred to as the context band and the context state, respectively. Although there are more powerful and advanced feature descriptors, such as convolutional neural networks, the selected ones require very few resources and do not depend on a graphics processor. Considering the limited resources of tactical UAVs, computational efficiency takes therefore precedence over potential performance gains.

With environmental conditions captured as context state and sensor performance expressed as TVI, the problem of modelling sensor performance is transformed into a band-wise regression task, which can principally be resolved by any suitable machine learning algorithm. In particular, each band of a multispectral imaging system is associated with a regression model, i.e. performance model, that takes the context state extracted from the visual band and predicts the TVI for its associated band. As a result, the number of models scales with the number of sensor bands for which the performance is to be predicted. This concept is illustrated in general in Figure 3.4. Here, the performances of four different sensor bands are predicted by four individual performance models from a context state extracted from a separate context band. The resulting predictions are then used to rank the bands from best- to worst-performing. With this information, the most valuable bands for the current conditions can be easily determined.

3.2.4 Results

Comparing predicted and calculated best-performing band orders over the entire test dataset allows the evaluation of the predictions in terms of rank accuracy. These accuracies are shown in Table 3.4, where each quarter corresponds to one of the utilized machine learning algorithms. In each cell of the quarters, the value shows the accuracy that the #column predicted best-performing bands are among the #row calculated or actual best-performing bands. Consequently, the first row provides the classic Top-1-Accuracy in the first column to

the Top-5-Accuracy in the last column. The lower right quarter of the table contains a special static baseline, where the predicted best-performing band order is constant over all samples of the test dataset, regardless of the context state. It is determined by calculating the band order that performs best on average over all samples of the training dataset. As can be seen, the performance models based on machine learning algorithms consistently outperform the static baseline in all accuracy categories. However, most of the top results are achieved using ϵ-SVR with 8 out of 15 top results, while the Random Forest Regression and XGBoost account for 5 and 2 of the remaining top results, respectively. Overall, the machine learning algorithms are able to identify the best-performing band among the actual top three best-performing bands with over 80 % accuracy, which is significantly better than the static baseline, which does not even reach 70 % in this category. As a result, if only the three predicted performing bands are considered in an actual application of the sensor performance modelling approach, the true best-performing band is most likely among them, but the amount of information or bands is halved.

Overall, the performance models show robust predictions for the samples in the test dataset, indicating a strong and learnable relationship between context state and TVI. Consequently, this demonstrates that the environmental conditions have a strong influence on the performance of the sensor bands, which supports the core idea of the proposed approach. This is further supported by the poor baseline results, which are independent of the environmental context and perform consistently worse than any of the performance models. Therefore, in terms of rank accuracy, it can be concluded that the proposed sensor performance modelling approach can effectively identify the most suitable sensor band for camouflage detection. Furthermore, the approach is well-suited to the dynamic and resource-limited environments of tactical reconnaissance scenarios due to the simplicity and efficiency of its methods. However, it should be noted that this evaluation assumes that a high TVI is actually beneficial for camouflage detection, while a low TVI is not. Although this might be trivial to assess for the minimum and maximum values of zero and one, the extent of support for camouflage detection for any value in between remains difficult to determine. As a result, there could be circumstances, where sensor bands with high TVIs are not necessarily more suitable for camouflage detection than sensor bands with lower TVIs. This limitation should be considered in any subsequent application of the proposed approach.

Table 3.4: Rank accuracies of all performance models in percent based on MUDCAD-X. Each individual table shows the predictions of the respective model, where the first row corresponds to Top-1 accuracy in the first column to Top-5 accuracy in the last column. This principle continues throughout the remaining rows. For example, the value in the fourth column of the third row represents the accuracy with which the three (row number) predicted best-performing bands are among the four (column number) actual best-performing bands. The highest accuracies per column and row are highlighted in bold.

ϵ-SVR					Random Forest				
56.1	**71.9**	83.6	88.9	95.3	51.5	67.8	**84.2**	88.3	97.1
	35.7	59.1	**71.3**	**88.9**		33.3	**59.6**	70.8	87.7
		37.4	**53.8**	72.5			35.7	53.2	**75.4**
			23.4	55.0				22.2	**59.6**
				34.5					**35.7**
XGBoost					Baseline				
50.9	68.4	83.6	**92.4**	**97.7**	47.4	56.1	69.6	75.4	82.5
	29.8	57.9	70.8	87.1		24.6	48.0	57.9	71.9
		31.0	46.8	68.4			33.9	46.2	60.2
			18.1	53.8				19.9	48.5
				34.5					28.7

3.3 Sensor-Managed Anomaly Detection for Camouflage Detection

This section summarizes the study [21] addressing the research question *"How effective is sensor-managed anomaly detection for camouflage detection in multispectral imagery?"*. The approach is outlined in Section 3.3.1 and the obtained results are reviewed in Section 3.3.2. A full reprint of the corresponding publication can be found in Section 5.3.

3.3.1 Approach

Sensor-managed anomaly detection combines sensor performance modelling with spectral anomaly detection. In particular, as in the previous study, performance models predict the value of each single-channel sensor band, defining the best-performing band order. Then, instead of processing all available sensor bands, as in the first study, only the best-performing bands are concatenated with the visual band and processed by anomaly detection algorithms. This reduces the number of bands to process, which reduces the processing times of the anomaly detectors. In addition, since the considered bands are the ones that most effectively expose camouflaged targets, their detection and false alarm rates may improve, as well.

The proposed approach is evaluated using the best performance models from the previous study, the ϵ-SVR models, as well as the same test dataset split from MUDCAD-X. For each sample of the test dataset, the three best-performing single-channel bands are selected for detection, which halves the number of single-channel bands from six to three. These three bands are then concatenated with the visual band, forming the managed band stack. As in the first study, this band stack is processed by multiple anomaly detectors to ensure a certain degree of meaningfulness of the generated results. These are the RX, LRX, LPD and AED detectors. AED-F and CRD are excluded this time, since AED-F does not provide significantly higher detection rates than plain AED and CRD is too slow even on a powerful machine. In addition to processing the managed band stack, the unmanaged band stack containing all available bands is also processed for comparison. Again, the detections are evaluated based on a threshold averaged ROC and corresponding AUC per camouflaged target and detector configuration. Furthermore, the detections are assessed using intersection over union (IOU), which is also a very popular metric but requires a binarization threshold. As binarization is usually required in an actual application of detection algorithms, this metric provides a more practical perspective on the detection results. The binarization threshold can be considered as part of the detector configuration. Consequently, as with the AUCs, there is also an IOU for each camouflaged target and detector configuration.

3.3.2 Results

The results of the sensor-managed anomaly detection approach are shown in Table 3.5, where the first and second halves display the resulting AUC and IOU values, respectively. Each of these halves is divided into results obtained from the unmanaged approach and results obtained from the managed approach. The individual values show the AUC or IOU per camouflaged target and detection algorithm for the respective approach. As can be seen, unmanaged detection produces more top results in terms of AUC than sensor-managed detection. However, in terms of IOU, the pattern is quite the opposite, with more top results for the managed approach. Overall, the differences between the approaches are rather small, with no clear superiority in terms of detection performance. Nevertheless, sensor-managed anomaly detection generally achieves higher IOU, which could be considered a more practical metric than AUC. Therefore, from a strictly practical perspective, the managed approach slightly outperforms the unmanaged counterpart.

The very similar results of both approaches could be caused by too little differences in the information contained in the final sensor band stacks, even though there are half as many single-channel bands with managed detection as with unmanaged detection. However, the results could also indicate that more information, even if it is not helpful for camouflage detection, does not necessarily confuse the detection algorithms. Consequently, there is no significant difference to observe whether this irrelevant information is processed or not.

In summary, sensor-managed anomaly detection cannot generally be considered superior to unmanaged anomaly detection for camouflage detection, based on the specific sensor data and setting evaluated in this study. However, it achieves comparable results, but with fewer sensor bands and thus less computational overhead. This also suggests that a high TVI is actually beneficial for camouflage detection, as only the bands with the highest TVI predictions are processed. As a result, sensor-managed anomaly detection provides similar effectiveness for camouflage detection as unmanaged anomaly detection. In view of the limited resources on small tactical UAVs, the managed approach may be regarded as more efficient, as the computational complexity for the anomaly detectors is significantly reduced, while detection performance is maintained.

3.4 Sensor Performance Modelling for Spectral Index Optimization

In this section, the study [20] investigating the research question *"How effective is spectral index optimization in combination with sensor performance modelling for camouflage detection in multispectral imagery?"* is summarized. The approach is introduced in Section 3.4.1, for

Table 3.5: Detection results in terms of Area under Curve (AUC) and Intersection over Union (IOU) of sensor-managed anomaly detection compared to unmanaged anomaly detection based on MUDCAD-X. Rows and columns correspond to anomaly detection algorithms and camouflaged targets, respectively. Top results per camouflaged target are highlighted in bold.

	Hedge	Turf	Grn. Tarp	Grn. 2D	Grn. 3D	Grn. Per.	Fleece	Gry. Tarp	Gry. 3D	Yew. 3D	Yew. Per.
Area under Curve (AUC)											
	Unmanaged										
RX	0.972	0.988	0.984	0.892	0.893	0.909	0.979	0.999	0.987	0.949	0.930
LRX	**0.993**	0.996	**0.995**	**0.960**	0.969	0.948	0.989	**1.000**	**0.990**	**0.976**	**0.991**
LPD	0.943	0.985	0.987	0.882	**0.970**	0.896	0.982	0.992	0.932	0.946	0.965
AED	0.880	0.955	0.970	0.918	0.935	0.886	0.946	0.994	0.921	0.967	0.975
	Managed										
RX	0.960	0.987	0.979	0.888	0.889	0.895	0.986	0.999	0.983	0.917	0.890
LRX	0.990	**0.997**	0.993	0.955	0.965	**0.955**	**0.994**	1.000	0.980	0.954	0.985
LPD	0.880	0.976	0.984	0.856	0.935	0.879	0.979	0.990	0.905	0.942	0.956
AED	0.875	0.956	0.967	0.898	0.905	0.892	0.955	0.993	0.921	0.962	0.975
Intersection over Union (IOU)											
	Unmanaged										
RX	0.255	0.167	0.188	0.065	0.017	0.060	0.014	0.758	0.085	0.103	0.037
LRX	**0.297**	0.248	0.325	0.106	0.078	**0.233**	0.113	0.748	0.152	0.226	0.207
LPD	0.042	0.168	0.231	0.050	0.078	0.228	0.094	0.509	0.024	**0.307**	**0.212**
AED	0.009	0.031	0.095	0.035	0.027	0.085	0.004	0.340	0.007	0.207	0.181
	Managed										
RX	0.164	**0.299**	0.296	0.076	0.020	0.049	0.057	0.754	0.098	0.077	0.038
LRX	0.227	0.296	**0.330**	**0.124**	0.105	0.217	0.180	**0.761**	**0.188**	0.209	0.174
LPD	0.071	0.202	0.202	0.056	**0.106**	0.219	**0.195**	0.588	0.026	0.283	0.186
AED	0.009	0.112	0.161	0.045	0.061	0.123	0.005	0.543	0.009	0.151	0.167

which the results are shown and discussed in Section 3.4.3. Section 3.4.2 briefly explains the novel index optimization approach. A full reprint of the corresponding publication can be found in Section 5.4.

3.4.1 Approach

Building on the previous experiments, where sensor performance modelling was used to determine a set of the most suitable sensor bands, this study explores how that concept can be partially adopted to derive a single most suitable sensor band. Deriving sensor bands from multiple raw bands is common practice in remote sensing and is typically referred to as the creation of spectral indices. Although many spectral indices exist, most are designed for highly specific purposes, such as indicating canopy nitrogen concentration [13, 33]. Given the limited published research on camouflage detection using airborne multispectral imagery, a spectral index specifically designed for this use case has yet to be proposed. In order to address this shortcoming, this study introduces a novel spectral index, the Linear Ratio Index (LRI). The LRI is a generalization of the well-known normalized difference, ratio and difference indices commonly employed in various remote sensing applications. It is composed of a ratio of two linear functions. These functions are essentially weighted sums of all available raw bands plus a bias, where the weights and the biases are the adjustable parameters. Along with the LRI, an optimization approach is introduced for determining its optimal parameters with respect to a given criterion. Since the index is intended to provide as much sensor performance as possible, this criterion is sensor performance itself, which corresponds to the TVI in the context of this thesis. However, the optimization approach is designed in such a way, that the criterion can be easily replaced, allowing it to be applied to other use cases as well. Furthermore, it incorporates a technique that enables the optimization of Linear Ratio Indices (LRIs) that rely only on a subset of all available bands, resulting in indices with reduced complexity.

The LRI and its optimization approach are evaluated for camouflage detection using training and test datasets split from MUDCAD-X. In particular, a two-band LRI, called LRI_2, and a six-band LRI, called LRI_6, are determined with the average TVI over all samples and camouflaged targets in the training dataset as optimization criterion. Using the same criterion, an optimized normalized difference index, called $\mathrm{LRI}_{\mathrm{nd}}$, and an optimized ratio index, called $\mathrm{LRI}_{\mathrm{r}}$, are also determined by permuting all possible band combinations. Due to their generic nature, these indices are often optimized for specific use cases where predefined indices are not available, thereby providing a solid baseline for the presented novel approach. After optimization, there are a total of four different indices: LRI_6, LRI_2, $\mathrm{LRI}_{\mathrm{nd}}$ and $\mathrm{LRI}_{\mathrm{r}}$. All optimized indices are evaluated using the test dataset for their effectiveness in terms of sensor performance, i.e. TVI. In addition, their effectiveness for camouflage detection is evaluated by separately concatenating

all indices with all available raw bands, resulting in one individual band stack per optimized index. These band stacks are then processed by several spectral anomaly detectors, which are the RX, LRX, LPD and AED detectors, as in the previous study. Again, the detections are assessed based on threshold averaged ROCs, corresponding AUCs and IOUs for a fixed binarization threshold.

3.4.2 Spectral Index Optimization

The novel approach to optimizing spectral indices is based on a generalized formulation of the normalized difference, ratio and difference indices. Compared to these indices, this reformulation, the linear ratio index (LRI), is capable of modelling much more complex relationships between raw bands at the cost of a higher number of parameters. Given a multispectral capture $C = \{c_1, \ldots, c_n\}$ with n bands, the LRI is calculated as given in Equation (3.8).

$$I_{\mathrm{lr}}(C, W^\alpha, W^\beta, b^\alpha, b^\beta) = \frac{\sum_{i=1}^{n} w_i^\alpha c_i + b^\alpha}{\sum_{i=1}^{n} w_i^\beta c_i + b^\beta} \tag{3.8}$$

Here, $W^\alpha = \{w_1^\alpha, \ldots, w_n^\alpha\}$ and $W^\beta = \{w_1^\beta, \ldots, w_n^\beta\}$ are the weight parameters that determine the contributions of each band in the numerator and denominator, respectively. In addition, there is a bias parameter b^α for the numerator and b^β for the denominator. While normalized difference, ratio and difference indices are optimized by testing different raw band combinations, the LRI is optimized by finding the best-fitting parameters for W^α, W^β, b^α and b^β. With an objective function f that takes a single spectral index image and returns a cost, the optimization problem of the LRI can be formulated as given in Equation (3.9). If the LRI is to be optimized for N multispectral captures $D = \{d_1, \ldots, d_N\}$, the costs of the individual index images are summed as shown in Equation (3.10).

$$\underset{W^\alpha, W^\beta, b^\alpha, b^\beta}{\arg\min} f\left(I_{\mathrm{lr}}\left(C, W^\alpha, W^\beta, b^\alpha, b^\beta\right) = \frac{\sum_{i=1}^{n} w_i^\alpha c_i + b^\alpha}{\sum_{i=1}^{n} w_i^\beta c_i + b^\beta}\right) \tag{3.9}$$

$$\underset{W^\alpha, W^\beta, b^\alpha, b^\beta}{\arg\min} \sum_{i=1}^{N} f\left(I_{\mathrm{lr}}\left(d_i, W^\alpha, W^\beta, b^\alpha, b^\beta\right)\right) \tag{3.10}$$

The search for optimal parameters for the LRI is performed by an algorithm based on differential evolution [41] with some general improvements in parameter selection [47] and population updates [40]. Since this optimizer can principally deal with any kind of problem, it leaves much room for an arbitrarily formulated objective function that fits a specific use case.

Using the proposed approach to find optimal linear ratio indices (LRIs) for camouflage detection also requires an appropriate objective function f. Since the sensor performance modelling approach already provides a measure of the performance of a sensor band in this context, the objective is maximizing the TVI. Naturally, this also requires label masks in the optimization procedure that identify all camouflaged targets in a multispectral capture, such as those available with MUCAD or MUDCAD-X. Given a set $A = \{a_1, \ldots, a_M\}$ with M label masks and a function f_{TVI} that takes an image and a label mask and returns the corresponding TVI, the optimization in Equation (3.9) can be adapted for camouflage detection as shown in Equation (3.11). Note that the function f_{norm} normalizes its input image to ensure a range of values between zero and one, which is the only range for which the TVI is defined. Furthermore, the result is negated to transform the maximization problem into a minimization problem.

$$\underset{W^\alpha, W^\beta, b^\alpha, b^\beta}{\arg\min} \; -\sum_{i=1}^{M} f_{\mathsf{TVI}}\left(f_{\mathsf{norm}}\left(I_{\mathrm{lr}}\left(C, W^\alpha, W^\beta, b^\alpha, b^\beta\right)\right), a_i\right) \tag{3.11}$$

In addition to being able to optimize the LRI for arbitrary applications, such as camouflage detection based on the TVI, the optimization approach incorporates a parameter complexity reduction technique. It allows limiting the precision of the parameters and the number of bands actually utilized for the LRI. This is achieved by rounding the parameters and adaptively setting some of them to zero directly during the optimization procedure. As a result, human readability of the parameters is enforced and dependence on available raw bands is reduced, while ensuring as maximum performance as possible.

3.4.3 Results

The average TVIs over all samples of the test dataset for each raw band, optimized index and camouflaged target are shown in Table 3.6. Top results are marked bold, while italics show TVIs of optimized indices that are higher than all the TVIs provided by the raw bands for the respective camouflaged target. Also, the last row shows an average over all camouflaged targets, representing an overall performance of respective sensor band or optimized index. As can be observed, the optimized indices provide the highest average TVI and the most top results over all camouflaged targets. Moreover, the LRI_6 and the LRI_2 clearly stand out by outperforming all the raw bands and all the optimized baseline indices ($\mathsf{LRI}_{\mathsf{nd}}$ and $\mathsf{LRI}_{\mathsf{r}}$) in terms of both average TVI and top TVI count.

Table 3.7 shows the detection results of all band stack configurations featuring optimized indices in terms of AUC and IOU per camouflaged target averaged over all samples in the test dataset. Note that the columns represent band stack configurations rather than anomaly detectors, as opposed to Table 3.3 and Table 3.5. The displayed values are averages over all

employed detection algorithms. In addition to the band stack configurations with optimized indices, the last row of each sub-table shows the results of the raw band stack configuration without any additional band or index. As can be seen, the band stack configurations with optimized indices consistently outperform the raw configuration, with very few exceptions for some specific camouflaged targets. Considering the average over all camouflaged targets shown in the last column of each sub-table, the LRI_6 band stack configuration clearly demonstrates superior performance to all other configurations. However, the LRI_2 band stack configuration practically matches the average detection performance of the LRI_6 band stack configuration in terms of AUC, but uses only two out of six possible bands. Similarly, the LRI_r band stack configuration achieves almost comparable results to the LRI_6 band stack configuration in terms of IOU, also using only two of the six available bands. In summary, the bank stack configurations with optimized LRIs significantly increase the detection performance compared to the raw band stack configuration.

Overall, the presented approach can effectively provide spectral indices that are more suitable for camouflage detection than raw sensor bands or traditionally optimized spectral indices with respect to the optimization criterion, which is sensor performance or the TVI. In addition, combining the resulting indices with the raw sensor bands leads to superior detection performance compared to all other band stack configurations using spectral anomaly detection in terms of AUC and IOU. This also suggests that incorporating sensor bands with high TVIs actually improves detection performance, which in turn supports the TVI as an appropriate optimization criterion for spectral indices designed for camouflage detection. As a result, the proposed approach could be particularly valuable in resource-constrained tactical reconnaissance scenarios, as the LRI is very easy to compute, once the parameters are found. Furthermore, processing a single additional band adds only minor computational overhead to an anomaly detector [18]. Given the demonstrated effectiveness of the LRI for camouflage detection, this additional overhead can be considered almost negligible.

Table 3.6: Average TVI for each camouflaged target of all optimized indices and all single-channel bands over the test dataset. An average over all targets is shown in the last row. Bold values indicate target top results per camouflaged target and italic values show where indices resulted in higher values than any raw band.

	blue	green	red	EIR	NIR	LWIR	LRI_6	LRI_2	LRI_{nd}	LRI_r
anthracite fleece	0.010	0.023	0.014	0.046	0.058	0.061	*0.113*	0.051	*0.097*	**0.144**
artificial turf	0.004	0.002	0.005	0.024	0.024	0.025	**0.067**	*0.058*	*0.041*	*0.041*
artificial hedge	0.005	0.012	0.006	0.011	0.009	0.006	*0.025*	**0.043**	0.008	0.004
gray tarp	0.017	0.008	0.027	0.085	0.061	0.004	0.060	**0.089**	0.023	0.007
green tarp	0.003	0.003	0.003	0.019	0.031	0.018	**0.065**	*0.038*	*0.052*	*0.056*
green 2D camouflage net	0.007	**0.009**	0.007	0.007	0.005	0.009	0.003	0.003	0.006	0.004
green 3D camouflage net	0.001	0.005	0.001	0.024	0.026	0.004	*0.044*	**0.044**	*0.029*	0.025
gray 3D camouflage net	0.026	0.004	0.004	0.010	0.021	0.011	**0.060**	*0.038*	*0.030*	*0.027*
yellow 3D camouflage net	0.021	**0.025**	0.030	0.010	0.001	0.006	0.004	0.003	0.002	0.002
person in green uniform	0.009	0.005	0.008	0.009	0.011	**0.021**	0.005	0.009	0.005	0.002
person in yellow uniform	0.019	0.012	0.023	0.008	0.001	**0.026**	0.003	0.002	0.007	0.005
average	0.011	0.010	0.012	0.023	0.023	0.017	**0.041**	*0.034*	*0.027*	*0.029*

Table 3.7: Average detection results in terms of Area under Curve (AUC) and Intersection over Union (IOU) for all optimized indices based on MUDCAD-X. Rows and columns correspond to band stack configurations and camouflaged targets, respectively. The last column shows the average over all camouflaged targets. Top results per camouflaged target are highlighted in bold.

	Hedge	Turf	Grn. Tarp	Grn. 2D	Grn. 3D	Grn. Per.	Fleece	Gry. Tarp	Gry. 3D	Yew. 3D	Yew. Per.	Avg.
Area under Curve (AUC)												
LRI_6	0.963	**0.994**	0.993	0.929	**0.976**	0.925	0.993	0.998	**0.970**	**0.966**	0.968	**0.970**
LRI_2	**0.976**	0.993	0.991	**0.934**	0.970	0.920	0.987	**0.999**	0.966	0.962	0.968	0.970
LRI_{nd}	0.949	0.990	0.991	0.919	0.963	**0.933**	0.980	0.997	0.959	0.962	0.968	0.965
LRI_r	0.953	0.991	**0.993**	0.919	0.966	0.932	**0.996**	0.997	0.964	0.966	0.969	0.968
raw	0.956	0.985	0.987	0.916	0.933	0.916	0.974	0.997	0.957	0.963	**0.970**	0.959
Intersection over Union (IOU)												
LRI_6	0.181	**0.377**	0.456	0.109	**0.213**	0.157	0.218	0.666	**0.151**	0.213	0.182	**0.266**
LRI_2	**0.208**	0.350	0.352	0.097	0.170	0.162	0.091	**0.706**	0.103	0.222	0.173	0.240
LRI_{nd}	0.174	0.227	0.347	0.095	0.139	0.165	0.054	0.610	0.102	0.210	0.183	0.210
LRI_r	0.165	0.293	**0.487**	**0.112**	0.167	0.162	**0.344**	0.570	0.111	0.203	0.185	0.255
raw	0.199	0.221	0.287	0.078	0.063	**0.177**	0.059	0.624	0.091	**0.245**	**0.193**	0.203

4 Conclusions and Future Prospects

Considering the obtained results in the described studies, multispectral imaging systems show generally robust capabilities for camouflage detection in reconnaissance scenarios with tactical UAVs. Using relatively fast and efficient spectral anomaly detection methods, camouflaged targets are reliably detected if they appear sufficiently anomalous in the sensor data. In addition, the potentially overwhelming flood of information provided by multispectral sensors can be managed by employing lightweight sensor modelling techniques that allow only the most valuable sensor bands to be considered even under changing environmental conditions. Although the combination of these two approaches does not provide significant advantages for camouflage detection under the specific conditions and sensor data used in this thesis, it may help conserve valuable resources, as the amount of data processed by the detection algorithms can be substantially reduced while performance is maintained. The experiments also demonstrated that sensor performance modelling, combined with a novel optimization approach, enables the derivation of spectral indices that significantly improve camouflage detection rates using spectral anomaly detection. Given the low computational overhead of these indices, their application is particularly beneficial for resource-constrained tactical UAVs. In summary, the proposed methods and approaches demonstrate sufficient effectiveness and efficiency for potential application in real-world reconnaissance scenarios, highlighting the value of this research.

The key scientific contributions of this thesis can be summarized as follows:

- first systematic evaluation of spectral anomaly detection for camouflage detection in tactical reconnaissance scenarios using airborne multispectral imagery
- introduction and evaluation of a novel sensor performance modelling approach to measure and predict the suitability of individual sensor bands based on environmental conditions
- development and evaluation of a sensor-managed anomaly detection method that adaptively processes the most suitable bands, reducing computational complexity while maintaining performance
- proposal and evaluation of a novel spectral index optimization approach that significantly improves camouflage detection rates with minimal processing overhead

- compilation and publication of two dedicated multispectral datasets for airborne camouflage detection (MUCAD and MUDCAD-X)

Despite the promising results, several limitations should be considered. For instance, all studies are based on relatively limited data sources, MUCAD and MUDCAD-X, as this type of data is rare and difficult to obtain. However, this may affect the degree to which they can be generalized and relied upon. In particular, the number of camouflaged targets in each dataset is comparatively small, making it difficult to infer the performance of the proposed methods when applied to different camouflaged targets. This is particularly relevant to the sensor performance prediction approach, since it is based on models fitted to data. In addition, the choice of methods for spectral anomaly detection and sensor performance modelling is inherently limited. Although they cover a wide range of different operating principles, alternative methods may produce different or possibly conflicting results even when applied to the same data. As a result, the transferability of the proposed approaches to other environments and scenarios may be limited and must be considered in any prospective real-world application. Additional studies would be necessary to address these potential limitations and to validate the proposed methods in a broader context.

In light of the findings of this thesis, multiple promising areas for future research emerge. For example, since spectral anomaly detection is a very active area of research, with new methods published frequently, faster and more resource-efficient algorithms for camouflage detection could be explored. This may further increase the utility of multispectral imaging systems on tactical UAVs through potentially faster data processing and higher detection rates. Moreover, as technology advances and embedded computing platforms become more powerful and efficient, more advanced methods such as deep neural networks may be particularly well suited for camouflage detection, context state extraction and sensor performance prediction. Especially for context state extraction, where rich semantic information must be encoded in a relatively small vector, such sophisticated methods can offer significant advantages, possibly leading to more accurate sensor performance predictions. As a result, predictions of the most valuable sensor bands may become even more reliable and thus more useful to operators in tactical reconnaissance scenarios. With sensors becoming cheaper, smaller and more powerful, the proposed approaches can also be applied to tactical UAVs equipped with hyperspectral imaging systems, which provide far more bands than multispectral imaging systems. In particular, the sensor performance modelling approach may be worth investigating in this context, as it theoretically scales with the number of bands, regardless of their origin. Combined with spectral anomaly detection, the benefit can be much greater due to the significantly larger set of sensor bands available for selective processing. Consequently, the relative runtime reduction of spectral anomaly detectors compared to multispectral sensors may also increase significantly,

potentially allowing processing speeds as fast as those required in tactical reconnaissance scenarios. Ultimately, regardless of the direction of future research, this thesis provides a solid foundation for further investigation and exploration.

5 Publications

This chapter reprints the four peer-reviewed studies on which this thesis is based. The sections are named according to the original titles, each with a full citation of the corresponding publication.

5.1 Adopting Hyperspectral Anomaly Detection for Near Real-Time Camouflage Detection in Multispectral Imagery

Tobias Hupel and Peter Stütz. „Adopting Hyperspectral Anomaly Detection for near Real-Time Camouflage Detection in Multispectral Imagery“. In: *Remote Sensing* 14.15 (2022). ISSN: 2072-4292. DOI: 10.3390/rs14153755

remote sensing

MDPI

Article

Adopting Hyperspectral Anomaly Detection for Near Real-Time Camouflage Detection in Multispectral Imagery

Tobias Hupel * and Peter Stütz

Institute of Flight Systems, Universität der Bundeswehr München, 85577 Neubiberg, Germany
* Correspondence: tobias.hupel@unibw.de

Abstract: Tactical reconnaissance using small unmanned aerial vehicles has become a common military scenario. However, since their sensor systems are usually limited to rudimentary visual or thermal imaging, the detection of camouflaged objects can be a particularly hard challenge. With respect to SWaP-C criteria, multispectral sensors represent a promising solution to increase the spectral information that could lead to unveiling camouflage. Therefore, this paper investigates and evaluates the applicability of four well-known hyperspectral anomaly detection methods (RX, LRX, CRD, and AED) and a method developed by the authors called local point density (LPD) for near real-time camouflage detection in multispectral imagery based on a specially created dataset. Results show that all targets in the dataset could successfully be detected with an AUC greater than 0.9 by multiple methods, with some methods even reaching an AUC relatively close to 1.0 for certain targets. Yet, great variations in detection performance over all targets and methods were observed. The dataset was additionally enhanced by multiple vegetation indices (BNDVI, GNDVI, and NDRE), which resulted in generally higher detection performances of all methods. Overall, the results demonstrated the general applicability of the hyperspectral anomaly detection methods for camouflage detection in multispectral imagery.

Keywords: camouflage detection; anomaly detection; multispectral; hyperspectral; infrared; image processing

Citation: Hupel, T.; Stütz, P. Adopting Hyperspectral Anomaly Detection for Near Real-Time Camouflage Detection in Multispectral Imagery. *Remote Sens.* **2022**, *14*, 3755. https://doi.org/10.3390/rs14153755

Academic Editor: Qiangqiang Yuan

Received: 10 June 2022
Accepted: 26 July 2022
Published: 5 August 2022

1. Introduction

Small unmanned aerial vehicles (UAVs) equipped with imaging sensor systems such as ALADIN or MIKADO of the German Armed Forces operate at relatively low altitudes (ca. 30–150 m) and are commonly utilized for reconnaissance tasks in tactical environments, as they are quickly deployed and allow to monitor very large areas without directly risking human lives. However, military forces generally attempt to conceal themselves and their equipment in their respective environments by using camouflage, making tactical reconnaissance a particularly demanding challenge. Additionally, small tactical reconnaissance drones are usually equipped with rudimentary visual or thermal sensor systems, which do not necessarily provide enough information in order to distinguish between camouflaged objects and their surroundings. Deploying a sensor payload that provides spectral information beyond visual or thermal range could therefore be crucial for the detection of camouflaged objects, as camouflage might lose its effectiveness in certain spectral regions.

Hyperspectral imaging systems, for instance, are capable of capturing a unique and distinctive spectral signature of every physical material, which can be successfully exploited for camouflage detection [1,2]. However, considering a tactical environment with small reconnaissance drones, the size, weight, power, and cost criteria (SWaP-C) of any payload must be taken into account. As powerful as hyperspectral sensors can be, as expensive is their acquisition. In addition, raw hyperspectral sensor data requires extensive postprocessing in order to obtain a hyperspectral data cube, rendering online hyperspectral sensor data evaluation nearly impossible. Compared to hyperspectral sensor systems, multispectral

sensor systems generally come at a more favorable price, are lower in size and weight and require much less postprocessing, making them considerably more relevant for any airborne remote sensing application. Yet, their spectral resolution is significantly lower and their spectral sensitivity per band is significantly higher, which could diminish their practical value in other ways.

Nevertheless, multispectral imagery has found many applications in recent years, ranging from precision farming [3–6], vegetational monitoring [7–12] and disturbance detection [13–15] to biomass estimation [16], aquatic plant detection [17,18], bathymetry [19], and even camouflage detection [20]. In [20], constrained energy minimization (CEM) in combination with a customized version of the well-known OTSU thresholding algorithm is applied to detect multiple camouflages in complex outdoor scenes of different perspectives. In addition to the effectiveness of the proposed method, it requires prior spectral knowledge about the camouflage, which might not be available in real-world tactical reconnaissance scenarios.

Since multispectral imagery has already been proven for various applications, even camouflage detection, the Institute of Flight Systems of the Universität der Bundeswehr München is actively researching the possibilities of multispectral imaging systems in tactical reconnaissance scenarios for computer-aided near real-time camouflage detection. For this purpose, a multispectral sensor setup was composed, mounted on a commercial drone, and utilized to collect a multispectral dataset containing multiple camouflaged targets of different materials and sizes. Moreover, a set of four well-known methods for hyperspectral anomaly detection (RX [21], LRX [21], CRD [22], and AED [23]) and one new density-based method developed by the authors called local point density (LPD) were adopted and applied on the specially created multispectral dataset. In addition to the raw dataset, the methods were also applied on an extension with multiple vegetation indices (BNDVI, GNDVI, and NDRE) of it, as they appeared to have increased contrast between the targets and their surroundings. The performance of these methods for detecting the camouflaged targets in the data were evaluated using the well-known metrics receiver operating characteristic (ROC) and area under the curve (AUC) and are presented and discussed in this paper. Since detection results should be available almost immediately in a real-world tactical reconnaissance scenario, the algorithms were also assessed with respect to the imposed near real-time requirement, which corresponds to a processing time of less than one second in this paper. The specific time constraint of one second results from a trade-off between the required availability of the detection results and the high computational complexity of the detection algorithms.

Hyperspectral anomaly detection is a very active field of research in which targets are characterized by spectral signatures that appear anomalous with respect to their current context. Consequently, successful target detections do not require prior knowledge about target signatures but a general approach for separating the target from background signatures. The Reed–Xiaoli detector (RX) [21] is one of the most prominent algorithms. It estimates the global covariance and the mean value per channel in a hyperspectral image to calculate the Mahalanobis distance for each pixel, which serves as a measure of its anomaly. As RX assumes that every background pixel follows a global Gaussian distribution, which might be a very simplified assumption in some cases, several techniques built upon the original RX detector have been introduced. For example, the local RX detector (LRX) considers only a window-based neighborhood for calculating the Mahalanobis distance, allowing more accurate background estimations for each pixel. Further developments are the kernel RX (KRX) [24], the cluster kernel RX (CKRX) [25], or the weighted RX (W-RXD) [26] detector, for instance. In contrast to distribution-based background modeling, representation-based detectors [22,27–29] assume that background pixels can be represented by certain or derived parts of the original hyperspectral image such as background dictionaries or local regions while anomalous pixels cannot. Furthermore, density-based [30,31], cluster-based [32], or morphological attribute filtering-based [23] detectors have also been proposed.

As it is assumed that camouflage spectrally differentiates from its surroundings in multiple bands, although it does logically not in visual, camouflage could be treated as a target for methods of hyperspectral anomaly detection. Moreover, the methods do not require prior knowledge about spectral target signatures, which fits the needs of real-world tactical reconnaissance scenarios where spectral characteristics of hostile camouflage should generally be considered unknown. However, since the selected methods were developed for hyperspectral imagery and for detecting spectral anomalies in general, they might not necessarily work for multispectral imagery containing camouflaged targets. Hence, the research work provided by this paper by compiling the dataset and evaluating the applicability of the aforementioned methods.

In the following sections, the dataset, the extension of it, and the selected methods are introduced in detail. Subsequently, the detection performances and runtimes of the methods are presented and discussed, and possible future research focuses are pointed out.

The contributions of this work can be summarized as follows:

- Compilation of a multispectral dataset for camouflage detection (MUCAD);
- Development of a density-based hyperspectral/multispectral anomaly detector;
- Evaluation of five different methods adopted for near real-time camouflage detection in multispectral imagery.

2. Materials and Methods

2.1. Dataset: MUCAD

The multispectral dataset for camouflage detection (MUCAD) consists of 23 samples, each featuring 7 images with a resolution of 512 × 512 pixels and a corresponding ground truth mask containing target annotations. All images have a color depth of 8 bits and a ground sample distance (GSD) of 10 cm/px. Each kind of target is labeled with a different but consistent color across all samples. An example of a single sample is shown in Figure 1. (a) shows the visual band (which is technically not a single band but is treated as such in this paper), (b–g) show the blue to LWIR bands, respectively, and are ordered by ascending wavelength and (h) shows the ground truth mask with annotations for a grey tarpaulin (dark blue), a green tarpaulin (light blue) and a 2D camouflage net (red).

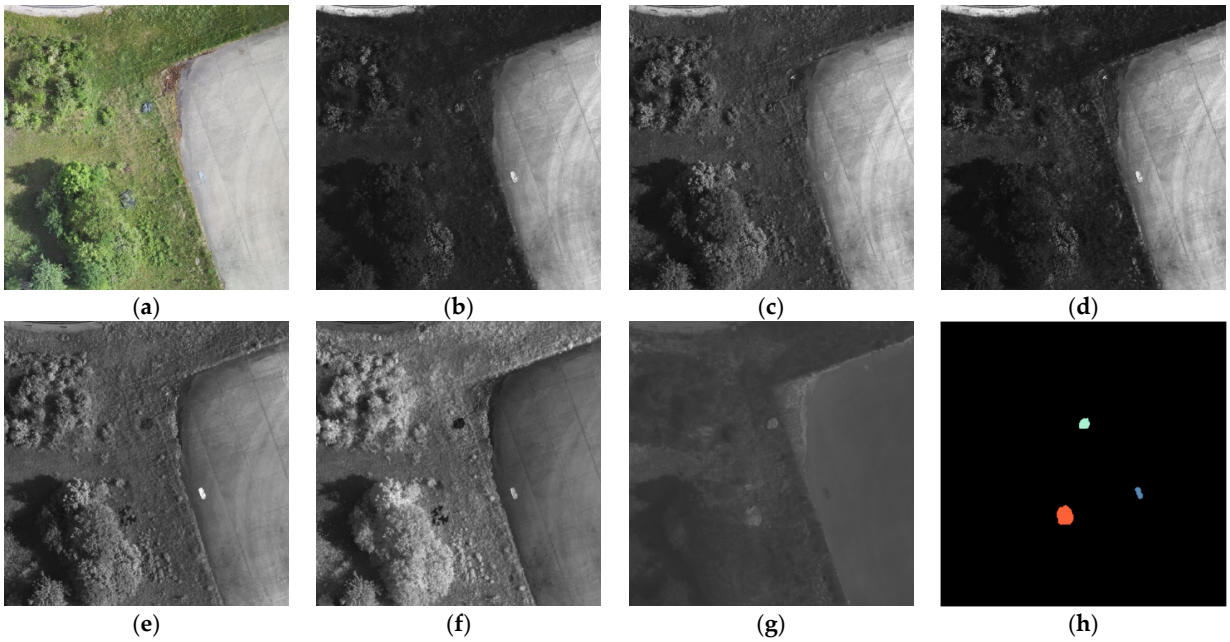

Figure 1. Single sample of MUCAD containing three different targets (grey tarpaulin, green tarpaulin, and 2D camouflage net). (**a**) VIS; (**b**) blue; (**c**) green; (**d**) red; (**e**) EIR; (**f**) NIR; (**g**) LWIR; (**h**) ground truth. The grey tarpaulin, green tarpaulin, and 2D camouflage net are labeled as dark blue, light blue, and red, respectively.

The samples of the dataset were not natively captured but cut out of a set of geographically referenced and aligned orthophotos (one for each band). In addition, the samples were pixel aligned based on the ECC criterion [33] using the implementation provided by OpenCV [34] and rescaled to fit a resolution of 512 × 512 pixels and a GSD of 10 cm/px.

The orthophotos were generated using the aerial image mapping software OpenDroneMap [35], where the raw sensor data were captured in 50 m height with 70% side and front overlap using a MicaSense Altum and a DJI Zenmuse XT2 mounted on a DJI Matrice 210 RTK V2. The complete setup, as well as a the generated visual orthophoto, are depicted in Figure 2. Visual and LWIR images were provided by the Zenmuse XT2 and blue to NIR images were provided by the MicaSense Altum. Table 1 shows further details about the captured spectral information of each band. Since both cameras capture images with 16-bit color depth except for the visual band, the intensities of all images were resampled from 16-bit to 8-bit color depth. Moreover, LWIR images were rescaled between their min and max values across all LWIR images before they were resampled down to 8-bit color depth. In contrast to common practices, the spectral bands provided by the MicaSense Altum were not converted to reflectance maps but kept in their raw form as reflectance values only matter for actual vegetational measurements which were not intended to be conducted.

Figure 2. The visual orthophoto that was used to generate the visual bands of MUCAD. The setup utilized for taking the raw images was composed of a DJI Zenmuse XT2 and a MicaSense Altum mounted on a Matrice 210 RTK V2.

The procedure for the dataset creation was chosen as it simplifies synchronization of both camera systems and allows the generation of images with arbitrary resolutions and ground sample distances only limited by the raw footage's ground sample distance. The downside of this method is the inability to capture dynamic scenes.

Table 1. Characteristics and associated sensor of each band.

Band	Sensor	Center Wavelength	Bandwidth
Visual (VIS)	XT2	-	-
Blue	Altum	475 nm	32 nm
Green	Altum	560 nm	27 nm
Red	Altum	668 nm	14 nm
Edge-infrared (EIR)	Altum	717 nm	12 nm
Near-infrared (NIR)	Altum	842 nm	57 nm
Long-wave infrared (LWIR)	XT2	10.5 µm	6 µm

Although this paper is about near real-time camouflage detection, the sensor–drone system in Figure 2 itself does not allow near real-time image processing, as both camera systems only provide simple storage interfaces (USB stick and SD card), meaning that image processing must be performed offline. An on-board system capable of online image processing would have to be composed of sensor systems that provide a proper high-speed interface. Furthermore, setting up such a system would have required a different hardware platform and additional on-board processing capabilities. However, in creating the dataset, which theoretically could have been produced by a system with online capabilities, main concerns were fast and reliable raw data generation without excessive preliminary work, hence the hardware configuration described above. In addition, this paper only considers the near real-time capabilities (processing time less than one second) of the detection algorithms and not those of an entire system.

In total, MUCAD contains 8 different kinds of targets, whose placement is shown in Figure 3. Each target was placed in a visually similar appearing environment. The first three target classes are depicted in Figure 3a: (1) a green tarpaulin that was shaped to appear like a bush instead of a simple green rectangular, (2) a 2D camouflage net that was placed near the woods to appear like an extension of it and (3) a gray tarpaulin that was placed in a concrete area and additionally shaped to avoid sharp rectangular appearing transitions between the target and its environment. The fourth target is shown in Figure 3b: (4) an artificial grass mat that was placed in very similar appearing high grass. Figure 3c contains the fifth and sixth target classes: (5) a 3D camouflage net that was placed and shaped to appear like a tree crown or a bush and (6) an artificial hedge that was placed in a shadowed area and thrown over a bush to adapt its shape.

The last two target classes are shown in Figure 3d: (7) two lying persons in a tree's shadow wearing a battle dress uniform and a German field dress, respectively, and (8) two gray cars which were initially not considered targets but became targets as their color is close to that of the road but their heat signature is completely different. Figure 4 exemplarily shows the green tarpaulin, the grey tarpaulin, and the 2D camouflage net from the ground perspective just before they were captured by the multispectral camera system.

The targets were acquired from multiple different sources. The green tarpaulin, the gray tarpaulin, the artificial grass mat, and the artificial hedge were bought at a local hardware store to cover the use case of utilizing relatively simple means and easily acquirable materials to conceal objects in a suitable environment without having to shop online. To cover a classic military use case, additional camouflage equipment was bought online. According to the store pages, the 2D camouflage net and the 3D camouflage net were supposedly original equipment of the Bundeswehr (German Armed Forces) and the Armed Forces of the Crown (British Armed Forces), respectively. The German field dress and the battle dress uniform were also authentic original military equipment, according to the store pages.

In the dataset, every target is almost evenly distributed across all samples to keep a potential target class distribution bias as small as possible. Furthermore, to obtain more meaningful detection results, the samples were cropped from the orthophotos in such a way that each target is located in different parts of the captured area and that the same target has different backgrounds considering the whole area captured by the sample. For this reason, there are more samples than targets in the dataset.

(**a**) (**b**)

(**c**) (**d**)

Figure 3. All camouflaged targets of MUCAD in their visual band. (**a**) Green tarpaulin (1), 2D camouflage net (2), and gray tarpaulin (3); (**b**) artificial grass mat (4); (**c**) 3D camouflage net (5), and artificial hedge (6); (**d**) two persons wearing a battle dress uniform and a German field dress, respectively, (7) and two gray cars (8).

(**a**) (**b**) (**c**)

Figure 4. Three targets of MUCAD from ground perspective. (**a**) Green tarpaulin (1); (**b**) grey tarpaulin (3); (**c**) 2D camouflage net (2). Numbers refer to Figure 3.

The entire dataset was captured at the end of May 2021 at the test site of the Universität der Bundeswehr München and is publicly available (See the Data Availability Statement at the end of the article).

2.2. Camouflage Detectors

For the detection of camouflage in multispectral imagery, four well-known hyperspectral anomaly detection methods were adopted: the classic Reed–Xiaoli detector (RX) and local RX detector (LRX) [21], the collaborative representation-based detector (CRD) [22] and the attribute and edge-preserving filter detector (AED) [23]. In addition to those four existing methods, this work introduces a new target detection method: local point density (LDP), which was inspired by dual window density (DWD) [31].

While there are a lot of powerful options available, the methods were primarily selected by their prominence, expected implementation effort, and expected computational requirements, although covering a great variety of distinct approaches was also a major concern. For this reason, all methods are based on entirely different target and background modeling techniques. The methods and their implementations are briefly discussed in the following sections. For a deeper understanding of the algorithms' theoretical foundations, versed readers may look up the provided references to the original publications.

2.2.1. RX and LRX

The RX detector [21] is based on the assumption that spectral target and background signatures in an image can be modeled by a single gaussian distribution, respectively. Leading to the Mahalanobis distance for calculating an anomaly score for a single pixel:

$$d_M = \sqrt{(x-\mu)^T \sum^{-1} (x-\mu)} \quad (1)$$

where d_M is the resulting Mahalanobis distance for the pixel of interest, x is the corresponding pixel value vector, μ is the mean vector consisting of the mean values for each band and $\sum$ is the covariance matrix for the image's bands. The higher the Mahalanobis distance for a single pixel the higher its anomality.

Instead of calculating the mean and covariance for the whole image, the LRX detector operates in a small local neighborhood for each pixel, defined by an inclusion and exclusion window size (W_{inc} and W_{exc}). Considering that every image possesses an arbitrary GSD and resolution and therefore also covers an arbitrary area in real-world space, it can make sense to restrict the region where mean and covariance estimation is performed to improve detection performance, as the whole image might cover multiple but entirely different backgrounds and targets. Figure 5 shows the process of local neighborhood selection for a single pixel under test (marked yellow). The outer blue window describes the inclusion area, and the orange window describes the exclusion area, meaning that the mean and covariance for calculating the anomaly score (Mahalanobis distance) of the yellow pixel are estimated with all values that are contained in the blue window but are outside of the orange window. Conclusively, the LRX detector is computational more expensive than the plain RX detector but might achieve higher detection rates because of its local nature.

Both algorithms were implemented in C++ using Eigen 3 [36], parallelized using OpenMP [37] and not modified in any way they work. Additionally, the algorithms were interfaced for the us in Python 3 using Cython [38].

Figure 5. Dual window neighborhood principle for anomaly detection. For the pixel under test (marked yellow), only pixels that are inside the blue window but outside the orange area are considered for estimating its anomaly score.

2.2.2. CRD

Collaborative representation-based detection [22] operates like the LRX detector in a local neighborhood (see Figure 5), but works entirely different in calculating a pixelwise anomaly score. In general, it is assumed that a non-anomalous pixel can be collaboratively represented by a weighted sum of its neighboring pixels. The neighboring pixels are the pixels in the local neighborhood defined by the exclusion and inclusion area of the pixel under test. Consequently, it is also assumed that an anomalous pixel cannot be collaboratively represented by its neighboring pixels. The problem can be described as an optimization problem for a pixel vector under test x and its dual window neighborhood N:

$$\arg\min_{\alpha} \|x - N\alpha\|_2^2 + \lambda\|\alpha\|_2^2 \tag{2}$$

where $N = \{n_i\}_{i=1}^{s_N}$ is a 2D Matrix of size $s_N = W_{inc}^2 - W_{exc}^2$ containing all pixel values in the dual window neighborhood of the pixel under test. x is the value of the pixel under test, α is the weight vector whose elements describe the individual contribution of each pixel in the neighborhood and λ is a regularization parameter that controls the weight of the penalty of the weight vector's squared second norm. Fortunately, the optimization problem has a closed form solution and must not be solved iteratively (I is the identity matrix):

$$\hat{\alpha} = \left(N^T N + \lambda I\right)^{-1} N^T x \tag{3}$$

After obtaining the weight vector $\hat{\alpha}$ for a pixel under test, its anomaly score a can be calculated, which is its reconstruction error at the same time:

$$a = \|x - N\hat{\alpha}\|_2 \tag{4}$$

The authors of CRD introduced an additional distance-based regularization method and a sum to one constraint for the weight vector. Details can be found in the original publication [22]. For the implementation in this work, both concepts were applied.

Although CRD is a rather slow method, it was initially thought that it can be greatly accelerated by parallelizing it with a graphics card. Unfortunately, the matrix inversion in Equation (3) turned out to be too complex for a single graphics card core, leading to an even slower solution than a parallelized C++ implementation using Eigen 3 [36] and

OpenMP [37] that was used for the results in this paper (additionally interfaced with Cython [38] for the use in Python 3).

2.2.3. AED

In attribute and edge-preserving filtering-based anomaly detection (AED) [23], it is assumed that anomalous objects are generally small and characterized by a significantly different intensity in a lot of bands compared to their backgrounds. In order to transfer that assumption into an algorithm, the authors utilized morphological attribute filters [39] to decompose the image under test into morphological attribute profiles (APs). In detail, for each band two versions (the APs) are generated. One where all bright connected components that have an area (number of pixels) lower than κ have been removed and one where all dark connected components that have an area lower than κ have been removed. These two versions are subtracted from each other to obtain a difference map that has very large values where very bright or very dark connected components have been removed. All difference maps undergo a special binary filtering step where only the pixels with large values are retained and all others are set to zero. The filtered difference maps are finally accumulated into a single anomaly map, which is processed by an edge-preserving filter. Consequently, the final anomaly map has high values where a lot of connected components have been removed that had much lower or higher intensities compared to their backgrounds, thus indicating anomalous objects. It should be noted that the original AED algorithm performs a dimensionality reduction step on the image before it is further processed. This step was not deployed in this work, since it deals with multispectral imagery, which is of low dimensionality by nature.

In addition to the originally proposed algorithm, a version with an extra filtering step right after the binary filtering step that removes every bright connected component with an area lower than ν or a compactness (area divided by the square of its perimeter length) lower than υ was implemented. This additional filtering step is supposed to eliminate false positive detections in each individual band before they are combined into a single anomaly map. The customized version of AED is called AED-F (F for the additional filtering).

To reduce implementation effort, the edge-preserving filtering step for the final anomaly map in both algorithms was achieved by a common bilateral filter [40] instead of a domain transform recursive filter [41] utilized in the original publication [23].

AED and AED-F were implemented in Python 3 using sap [42] and OpenCV [34].

2.2.4. LPD

Inspired by DWD [31], a density-based [43] algorithm was developed that encompasses a comparatively less memory intensive and more equally weighted pixelwise density computation, called local point density (LPD). LPD operates like LRX and CRD in a local neighborhood that is defined by an inclusion and exclusion area, as it is shown in Figure 5.

For a pixel under test, its density is calculated according to:

$$\rho = \frac{1}{s_M} \sum_{i=1}^{s_M} \exp\left(-\frac{\|x - m_i\|_2^2}{d_c^2} \right) \tag{5}$$

where x is the pixel value vector, m_i is the i-th element of $M = N \cup \{x\}$ and s_M is the size of M ($s_M = s_N + 1$). $N = \{n_j\}_{j=1}^{s_N}$ is the pixel's dual window neighborhood of size $s_N = W_{inc}^2 - W_{exc}^2$. d_c is the cut-off distance, which in contrast to DWD is not determined by calculating a distance matrix containing all distances between all pixel values but by computing the average distance between all pixel values and the average pixel value:

$$d_c = \frac{1}{s_M} \sum_{i=1}^{s_M} \|m_i - \mu_M\|_2 \tag{6}$$

$$\mu_M = \frac{1}{s_M} \sum_{j=1}^{s_M} m_j \tag{7}$$

Compared to calculating a distance matrix, this method is much less memory expensive, as it is purely iterative and no matrix must be stored at all, which can prove useful in large scale parallelized implementations and on memory limited platforms.

The density calculated per pixel is additionally more evenly weighted compared to DWD, as pixel density is calculated pixel-wise with respect to a local neighborhood which is not the case in DWD, where pixel density is calculated window-wise with respect to a local neighborhood, meaning that the local neighborhood is not centered around every pixel under test as it is in LPD.

In the final density map, anomalous and background pixels have a low and high density, respectively. The density map is converted into an anomaly map by negating it and adding the lowest value to all pixels afterwards to avoid negative values.

LPD was implemented using C++ and Eigen 3 [36], parallelized using OpenMP [37], and interfaced for use in Python 3 using Cython [38].

2.3. Vegetation Indices

As stated in the introduction, the methods described in Section 2.2 were applied to two different configurations of MUCAD. First, on raw MUCAD, and second, on MUCAD-VI, which is an extension of MUCAD with the vegetation indices BNDVI, GNDVI, and NDRE, as it is described in Table 2. The indices were chosen as they appeared to be increasing the visibility of some targets with respect to their environments (see Figure 6), which could support successful target detections. All indices were computed online, meaning that they were not statically saved as files but calculated right before the detection algorithms were applied.

Table 2. Vegetation indices for online extension of MUCAD.

Index	Formula
BNDVI	$\frac{NIR-blue}{NIR+blue}$
GNDVI	$\frac{NIR-green}{NIR+green}$
NDRE	$\frac{NIR-EIR}{NIR+EIR}$

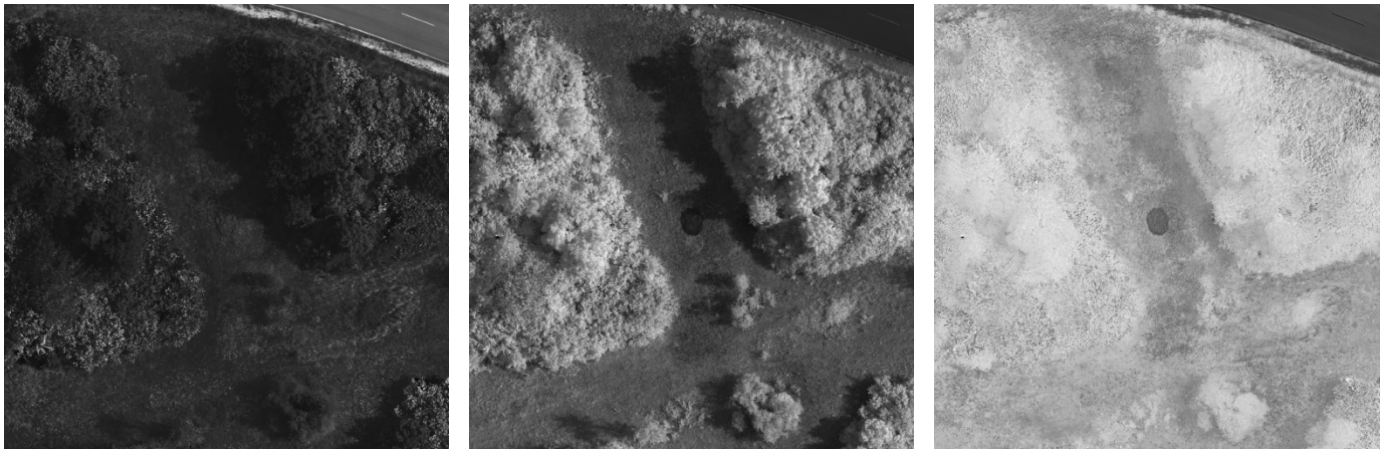

(**a**) (**b**) (**c**)

Figure 6. The vegetation index BNDVI calculated with NIR and blue bands increases visibility of 3D camouflage net in its environment compared to both raw bands. (**a**) blue; (**b**) NIR; (**c**) BNDVI.

3. Results

For the evaluation of each target detector with MUCAD, an image resolution of 256×256 pixels was chosen to conserve computing capabilities, resulting in an effective ground sample distance of 20 cm/px for each sample. It is the lowest resolution that still allows for visually detecting the smallest targets of the dataset (persons in uniforms). Before the samples were processed by the algorithms, every band was z-normalized (mean of zero and standard deviation of one). Additionally, the visual band was weighted by one-third, since it was treated as a single band but technically consists of three individual bands. As a postprocessing step, all bright connected components that had an area smaller than nine pixels were removed in the final binarized anomaly map of each target detector, as it was assumed that anomalous objects occupy an area of at least a third square meter ($9 \times 400\ \text{cm}^2 = 0.36\ \text{m}^2$).

To find a near to optimal configuration for each target detector, multiple parameters sets were applied (except for RX, which is parameterless). For LRX and LPD, W_{inc} and W_{exc} were set to: $\{(5, 15), (11, 31), (21, 61), (31, 91), (41, 121)\}$. Because of its expensive computational requirements, the window sizes of CRD were selected slightly differently: $\{(5, 15), (11, 21), (21, 31), (31, 41), (41, 51)\}$. As for the AED-based methods, κ was set to: $\{25, 50, 100, 200, 300\}$. For AED-F, ν was also set to nine and υ (compactness threshold) was set to 0.15.

The target detectors were applied on two different dataset configurations: the raw version of MUCAD and an extended version with three additional bands: BNDVI, GNDVI, and NDRE. The evaluation of each method was performed based on its receiver operating characteristic (ROC) and the corresponding area under the curve (AUC). ROCs were calculated for each target class and each capture individually, which were then combined by threshold averaging [44] afterward to obtain a single ROC for each target class over the whole dataset. ROCs with an AUC lower than 0.9 had such bad detection rates (true positive rate) and high false alarm rates (false positive rate), that they were considered to have not detected the target at all. Therefore, those weak results are not included in the following tables and graphs for better clarity of the remaining results.

All experiments were conducted on a machine running Ubuntu 20.04 LTS with an AMD Ryzen 9 3950X (16C/32T) and 128 GB Memory.

3.1. MUCAD

First, all target detectors were applied on the raw version of MUCAD. The resulting receiver operating characteristics are displayed in Figure 7. Note that for each method only the parameter configuration that achieved the highest AUC over all parameter configurations of the method is shown. All other parameter configurations are not further considered in the following evaluation. Each tile shows the results for a specific target class and contains a color-coded ROC for every successful algorithm (note that each plot has individual limits for the abscissa). The legend of every tile shows AUCs for each ROC in parentheses.

As it can be seen, detection performances vary greatly across target detectors and classes. The artificial hedge could only be detected by LRX and LPD. Still, their AUCs are relatively low, and ROCs show that false alarm rates are already considerably high before detection rates even reach 50%, which signalizes high difficulty for the methods to detect the target. In contrast to that, the artificial grass mat, the green tarpaulin and the gray tarpaulin were detected by all algorithms with high detection rates at low false alarm rates. The 2D camouflage net, as well as the cars, were also detected by all methods but with significantly lower performance, although LRX and CRD stand clearly out in AUC and ROC. For the 3D camouflage net, only LRX and CRD could achieve considerable results. Yet, the detection performances are almost as bad as for the artificial hedge, making it the second target that could barely be detected. The persons were also detected only by two algorithms: LRX and LPD. Only this time, LPD achieved a comparatively strong performance with high detection rates at low false alarm rates.

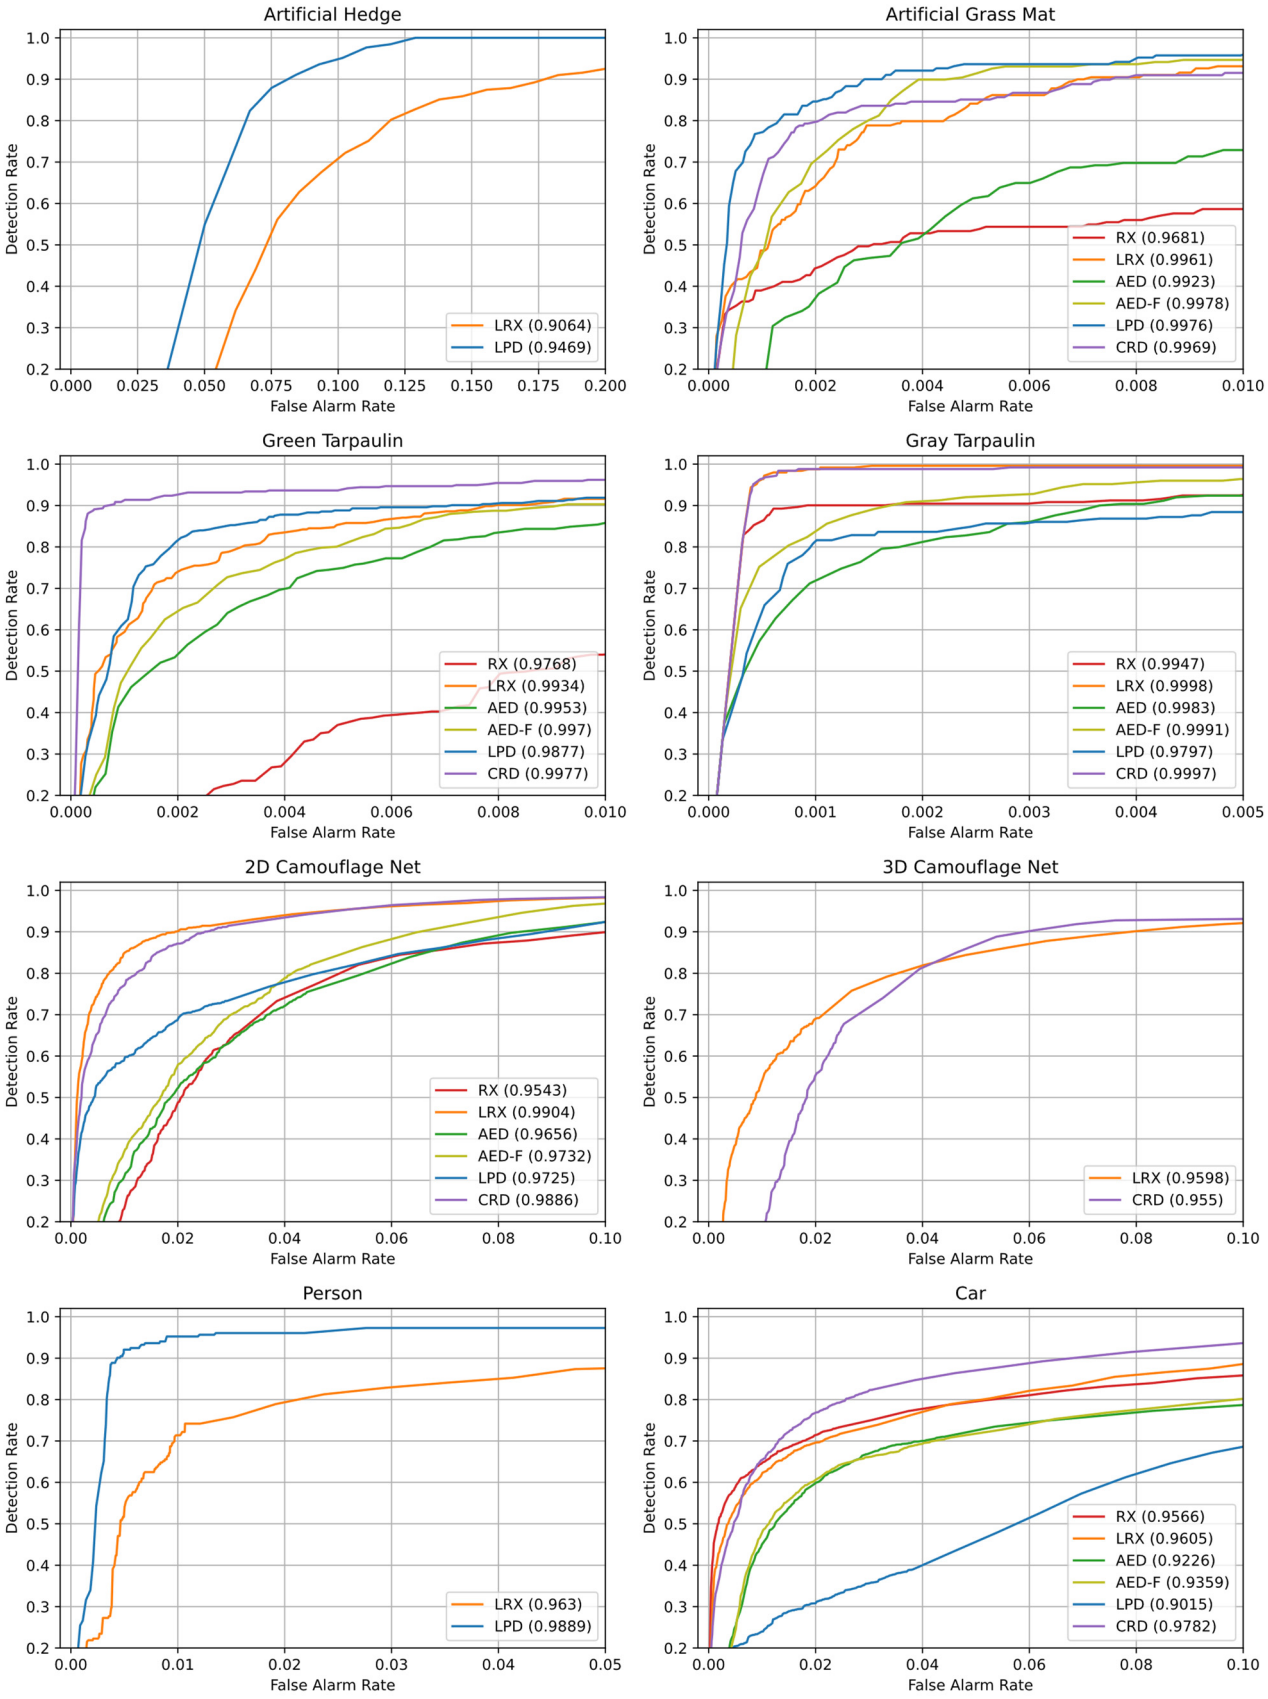

Figure 7. Receiver operating characteristics (ROCs) and areas under the curve (AUCs) for individual target classes of every method applied on the raw version of MUCAD. Algorithms that achieved an AUC less than 0.9 are not included. Note that each plot has individual limits for the abscissa.

Table 3 contains the AUCs, parameter configurations, and runtimes for the ROCs in Figure 7. The upper third of the table shows the AUC values of each algorithm for every target class where the highest values are marked bold. As mentioned before, AUC values lower than 0.9 are not included and were replaced by a "-". Top scores are almost evenly spread across LRX, LPD and CRD, showing that the detection task benefits from local target and background assumptions, while the plain RX detector stays noticeably behind the top performances in all target classes except for the gray tarpaulin. AED-based methods fail in as many target class detections as RX but achieve close to best performance for the artificial grass mat, the green tarpaulin and the gray tarpaulin. Moreover, AED-F achieves consistently higher AUCs compared to its unmodified counterpart AED, showing the effectiveness of the additional filtering. As already pointed out based on the ROCs, the artificial hedge, the 3D camouflage net and the persons cannot be detected by a lot of algorithms, though LPD achieves a strong detection performance for the person target class. Those bad detection performances are likely due to weak spectral differences of the artificial hedge (which even lies in the shadows) and the 3D camouflage net, and the small spatial size of the persons, making them particularly hard to differentiate from normal background clutter.

Table 3. Class-wise detection performances, parameter configurations and runtimes of each method applied on the raw version of MUCAD. Detection results with an AUC less than 0.9 were replaced with a "-" for better clarity.

	Art. Hedg.	**Art. Gr. Mat**	**Green Tarp**	**Gray Tarp**	**2D Net**	**3D Net**	**Person**	**Car**
Detection Performance (AUC)								
RX	-	0.9681	0.9768	0.9947	0.9543	-	-	0.9566
LRX	0.9064	0.9961	0.9934	**0.9998**	**0.9904**	**0.9598**	0.9630	0.9605
AED	-	0.9914	0.9953	0.9983	0.9656	-	-	0.9226
AED-F	-	0.9971	0.9970	0.9991	0.9732	-	-	0.9359
LPD	**0.9469**	**0.9976**	0.9877	0.9797	0.9725	-	**0.9889**	0.9015
CRD	-	0.9969	**0.9977**	0.9997	0.9886	0.9550	-	**0.9782**
Parameter Configuration (W_{inc} and W_{exc} or κ)								
RX	-	-	-	-	-	-	-	-
LRX	(41, 121)	(11, 31)	(21, 61)	(21, 61)	(31, 91)	(41, 121)	(5, 15)	(21, 61)
AED	-	100	200	100	300	-	-	300
AED-F	-	100	200	100	300	-	-	300
LPD	(41, 121)	(11, 31)	(11, 31)	(11, 31)	(21, 61)	-	(5, 15)	(21, 61)
CRD	-	(21, 31)	(21, 31)	(21, 31)	(31, 41)	(31, 41)	-	(41, 51)
Runtime								
RX				6 ms				
LRX	3700 ms	114 ms	350 ms	350 ms	894 ms	3700 ms	43 ms	350 ms
AED				295 ms				
AED-F				525 ms				
LPD	953 ms	74 ms	74 ms	74 ms	268 ms	-	24 ms	268 ms
CRD	-	2.8 min	2.8 min	2.8 min	8.2 min	8.2 min	0.5 min	17 min

The middle third of the table describes the parameter configuration for each detector whose AUC is given in the first third. Depending on the algorithm, values stand either for the inclusion and exclusion window (W_{inc} and W_{exc}) or for the area of connected components to be removed (κ). RX is parameterless. For the local neighborhood-based algorithms, window sizes logically rise and fall with the targets' sizes. As does the connected component area for the AED-based methods.

The runtimes of the algorithms are outlined in the lower third of the table. The RX detector does not have any parameters which results in a consistent runtime for each target class. LRX, LPD and CRD depend heavily on their parameters, which translates into increasing runtimes for larger window sizes. The runtimes of the AED-based methods

are not influenced by their parameters and are therefore constant for each target class, although the additional filtering in AED-F almost doubles its time consumption. CRD has striking high runtimes due to its computational complexity, making it unsuitable for any near real-time application regardless of its strong detection performance. All other methods keep their time consumption below one second except LRX, which exceeds it when window sizes are comparatively large. Note that the runtimes only contain the raw method execution time and no time consumption caused by pre- or postprocessing.

Figure 8 shows the detection maps of all methods for an exemplary capture of MUCAD containing the green tarpaulin, the grey tarpaulin, and the 2D camouflage net. The first two images show the visual band (a) and the associated ground truth (b) of the capture, followed by the detection maps (c–h) of RX, LRX, AED, AED-F, LPD, and CRD, respectively. The parameter configurations of each method correspond to the second column in Table 3. As it can be qualitatively observed, AED-F (f) produced much fewer false positives in the detection map compared to AED (e). LRX (d) also generated much fewer false positive detections than RX (c), which is consistent with their corresponding ROCs in Figure 7 and AUCs in Table 3. LPD (g) produced a significantly different detection map with generally higher anomaly scores compared to all other methods, but with no apparent negative impact on its ROC or AUC. The detection map of CRD (h) appears to be very similar to the detection map of LRX (d), although there are generally considerable differences in detection performance, as indicated by their ROCs and AUCs.

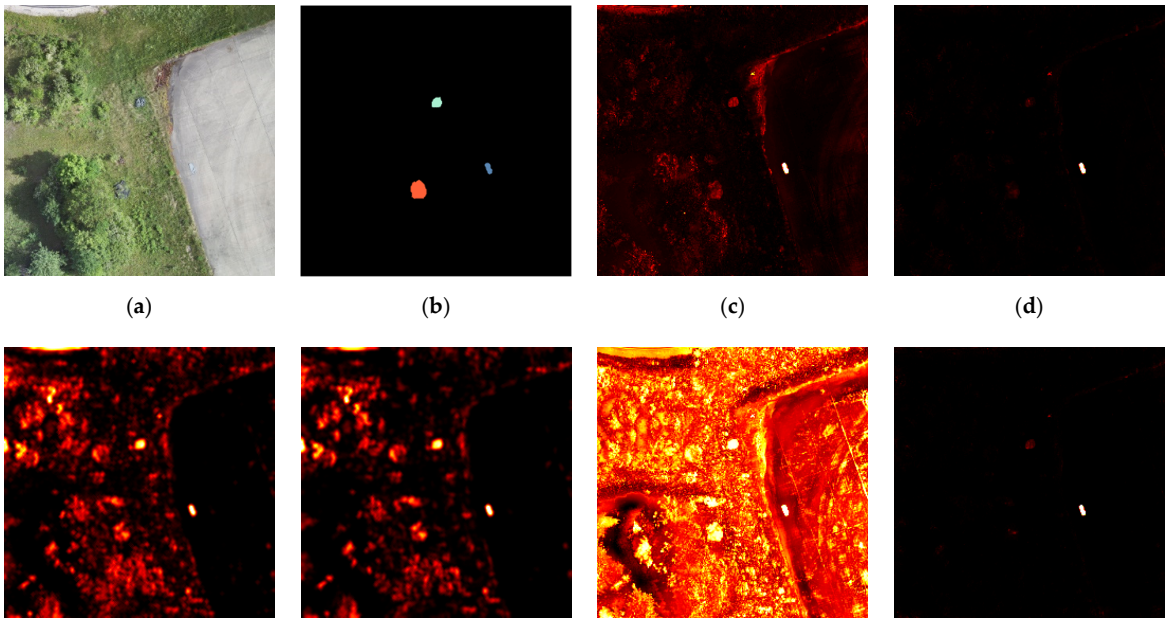

Figure 8. Detection maps of all evaluated methods applied on an exemplary capture of the raw version of MUCAD. (**a**) VIS; (**b**) ground truth; (**c**) RX; (**d**) LRX (21, 61); (**e**) AED (200); (**f**) AED-F (200); (**g**) LPD (11, 31); (**h**) CRD (21, 31). The parameter configurations (in parentheses) correspond either to the inclusion and exclusion window or to the area of connected components to be removed.

3.2. MUCAD-VI

For the extended version of MUCAD (MUCAD-VI), receiver operating characteristics for every target class are shown in Figure 9. The tiles and their contents follow the same pattern as in Figure 7. For the artificial hedge, only a subset of algorithms (RX, LRX, and LPD) could achieve more than 0.9 AUC, but LRX obtained much higher detection rates at equal false alarm rates compared to RX and LPD. The artificial grass mat, the green tarpaulin, and the grey tarpaulin were detected by all algorithms, with LRX, CRD, AED, and AED-F being close to an ideal detector for certain targets. All methods could detect the 2D camouflage net, although LRX clearly outperformed all other detectors. For the 3D

camouflage net, only LRX and CRD could achieve comparatively high detection rates at low false alarm rates. LPD and RX obtained indeed more than 0.9 AUC but performed significantly lower than their competitors. The persons were detected solely by LRX, LPD, and CRD, with LPD clearly dominating in high detection rates at low false alarm rates. For the detection of the cars, RX, LRX, and CRD obtained very similar results and outperformed all other methods, with LPD not even reaching 0.9 AUC.

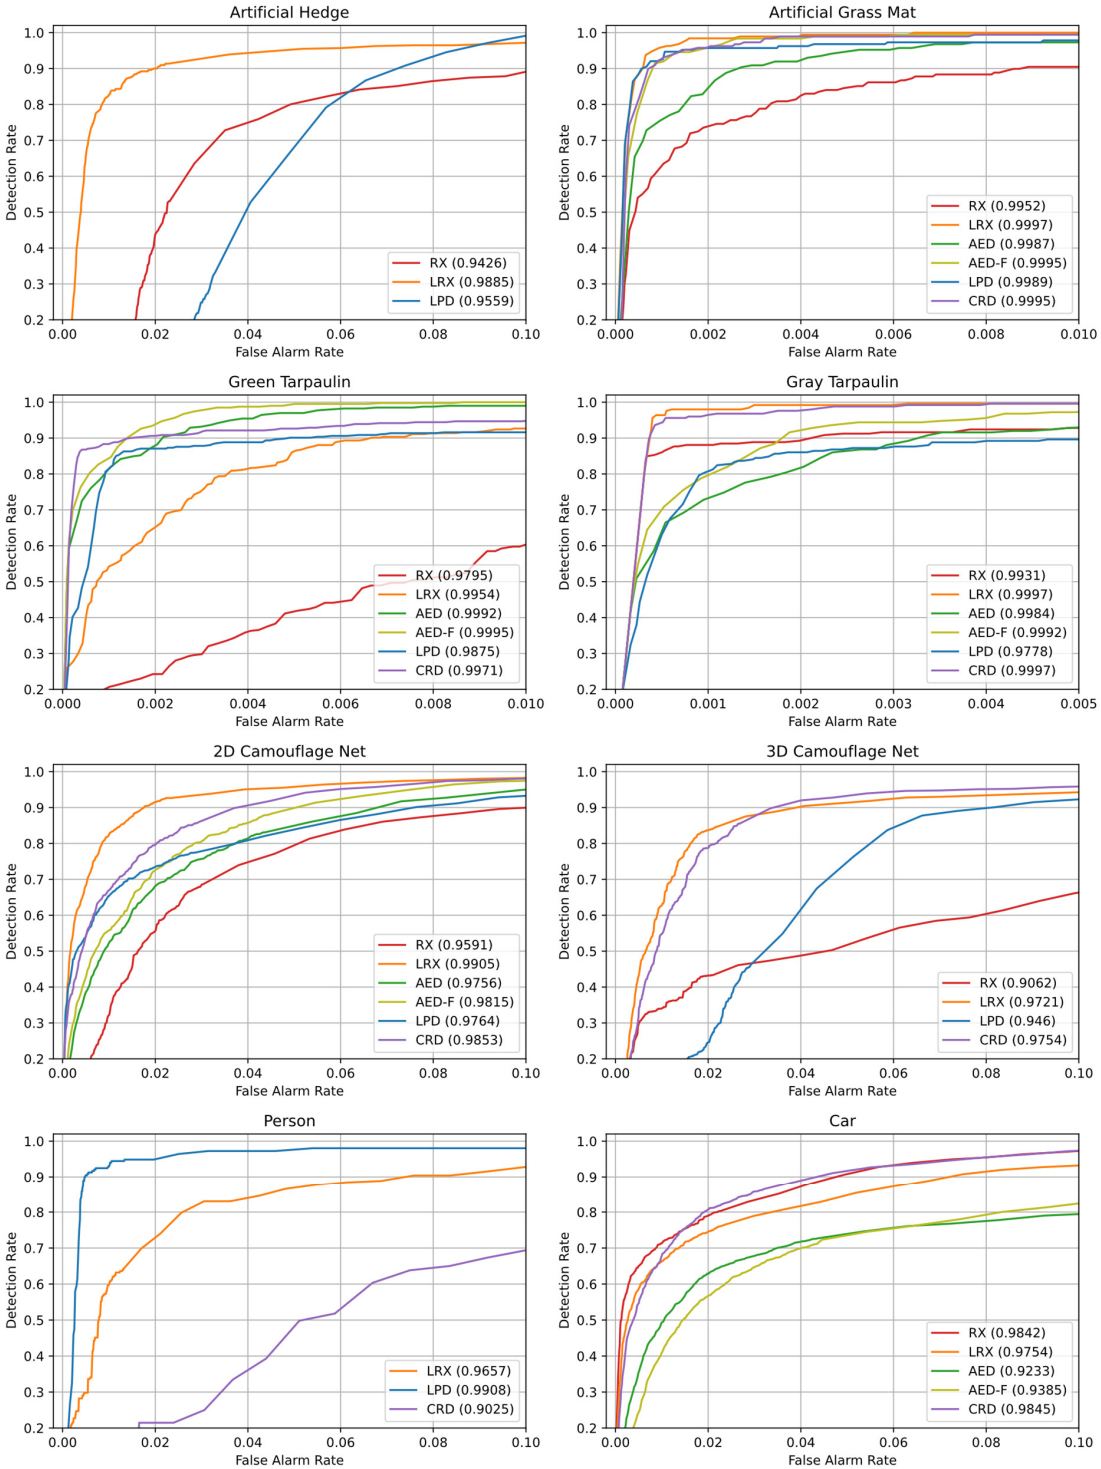

Figure 9. Receiver operating characteristics (ROCs) and areas under the curve (AUCs) for individual target classes of every method applied on MUCAD-VI. Algorithms that achieved an AUC less than 0.9 are not included. Note that each plot has individual limits for the abscissa.

Compared to the ROCs in Figure 7 (raw MUCAD), detection rates at low false alarms generally improved across all methods and target classes. Most noticeable, the artificial hedge (which was barely detected by LRX applied on the raw version of MUCAD) could be detected with much higher detection rates at much lower false alarm rates. Additionally, targets that were detected only by two methods in MUCAD, were detected by three to four methods in MUCAD-VI. Although LPD previously managed to achieve more than 0.9 AUC for the cars, its detection performance fell below that threshold in its application on MUCAD-VI.

Table 4 contains AUCs, parameter configurations, and runtimes for the methods and ROCs in Figure 9. In the first third of the table, the detection performances in AUC are displayed, with the highest value per target class marked in bold. LRX and CRD are clearly dominating in terms of the number of highest AUC per class. AED-F and LPD achieved a single top result, respectively, while RX and AED did not achieve any of the top results. Additionally, LRX is the only algorithm that detected all targets (more than 0.9 AUC for all targets). RX, LPD, and CRD could not detect the persons, the cars, and the artificial hedge, respectively. In addition, not having accomplished top results in each class, it is notable that AED-F achieved results relatively close to an ideal detector for the artificial grass mat, the green tarpaulin, and the gray tarpaulin.

Table 4. Class-wise detection performances, parameter configurations, and runtimes of each method applied on MUCAD extended by vegetation indices (MUCAD-VI). Detection results with an AUC less than 0.9 were replaced with a "-" for better clarity.

	Art. Hedg.	**Art. Gr. Mat**	**Green Tarp**	**Gray Tarp**	**2D Net**	**3D Net**	**Person**	**Car**
				Detection Performance (AUC)				
RX	0.9426	0.9952	0.9795	0.9931	0.9591	0.9062	-	0.9842
LRX	**0.9885**	**0.9997**	0.9954	**0.9997**	**0.9905**	0.9721	0.9657	0.9754
AED	-	0.9987	0.9992	0.9984	0.9756	-	-	0.9233
AED-F	-	0.9995	**0.9995**	0.9992	0.9815	-	-	0.9385
LPD	0.9559	0.9989	0.9875	0.9778	0.9764	0.9460	**0.9908**	-
CRD	-	0.9995	0.9971	**0.9997**	0.9853	**0.9754**	0.9025	**0.9845**
				Parameter Configuration (W_{inc} and W_{exc} or κ)				
RX	-	-	-	-	-	-	-	-
LRX	(41, 121)	(21, 61)	(11, 31)	(21, 61)	(31, 91)	(31, 91)	(5, 15)	(21, 61)
AED	-	100	100	100	300	-	-	300
AED-F	-	100	100	100	300	-	-	200
LPD	(41, 121)	(11, 31)	(11, 31)	(11, 31)	(21, 61)	(21, 61)	(5, 15)	-
CRD	-	(21, 31)	(21, 31)	(21, 31)	(31, 41)	(31, 41)	(5, 15)	(41, 51)
				Runtime				
RX				6 ms				
LRX	7300 ms	476 ms	148 ms	476 ms	2000 ms	2000 ms	53 ms	476 ms
AED				397 ms				
AED-F				691 ms				
LPD	1000 ms	78 ms	78 ms	78 ms	282 ms	282 ms	24 ms	-
CRD	-	2.8 min	2.8 min	2.8 min	8.2 min	8.2 min	0.1 min	17 min

As it has already been pointed out in the evaluation of the ROCs in Figure 9, compared to the results obtained with raw MUCAD, detection performances for MUCAD-VI consistently increased or were at least equally high. This is most noticeable for the artificial hedge and the 3D camouflage net, which were most difficult to detect by all algorithms (lowest AUC values). Furthermore, targets that were detected only by a small subset of methods before, were detected by one to two methods more.

The parameter configurations in the middle third of table, indicating either the inclusion and exclusion window sizes or the area of connected components to be removed, do

not show any significant changes compared to the parameter configurations of the methods when applied on raw MUCAD.

The opposite situation prevails for the runtimes in the lower third of the table. While RX and CRD seem to be almost unaffected by the increased number of bands in MUCAD-VI, all other methods show considerable higher runtimes. LRX needs roughly twice as much time and the AED-based methods apparently scale linearly with the number of bands to process. LPD is less sensitive to the number of bands, resulting only in a marginally higher time consumption for the same parameter configurations.

Figure 10 shows the detection maps of all methods for an exemplary capture of MUCAD-VI containing the green tarpaulin, the grey tarpaulin and the 2D camouflage net. The images (a–h) follow the same structure as the images in Figure 8. The parameter configurations of each method correspond to the second column in Table 4. Compared to the detection maps in Figure 8, the anomaly scores in the target areas are considerably higher, particularly in the detection maps of AED (e) and AED-F (f), which qualitatively confirms the generally higher detection performance of all methods observed in the ROCs of Figure 9 and the AUCs of Table 4. AED-F (f) again generated much fewer false positives than AED (e), as did LRX (d) compared to RX (c). In contrast to Figure 8, CRD (h) produced a detection map with significantly fewer false positives compared to LRX (d).

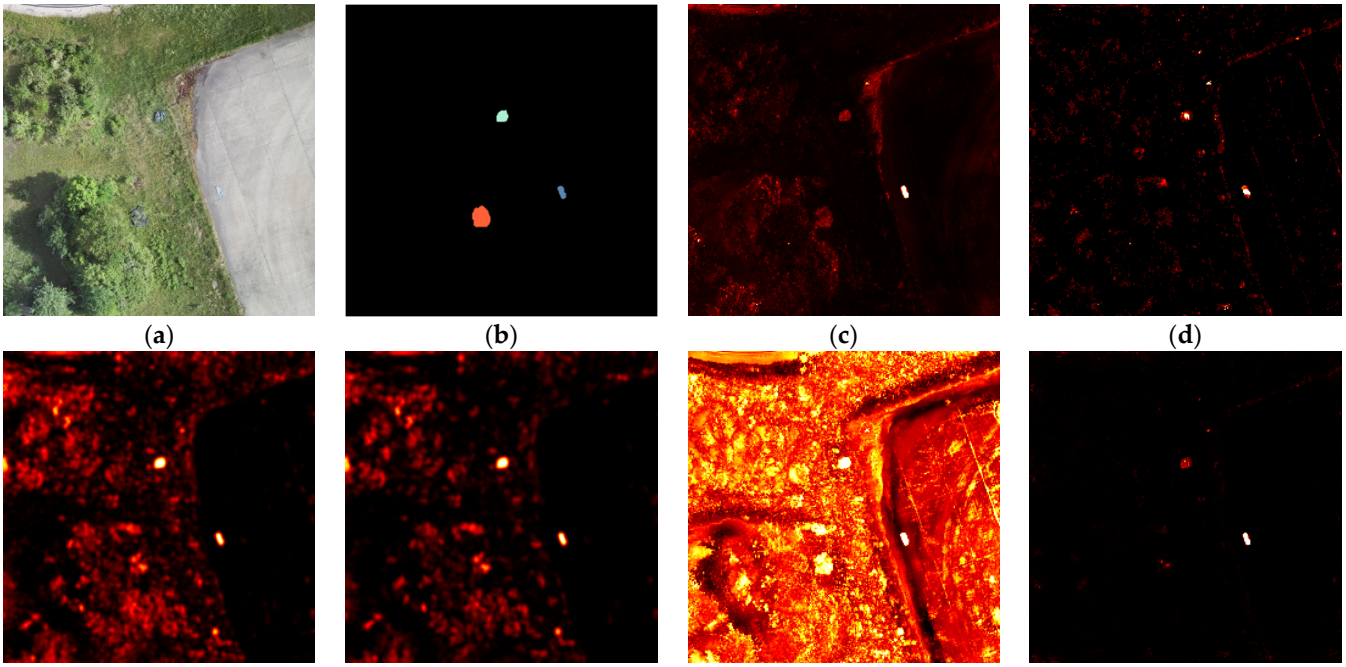

Figure 10. Detection maps of all evaluated methods applied on an exemplary capture of the extended version of MUCAD (MUCAD-VI). (**a**) VIS; (**b**) ground truth; (**c**) RX; (**d**) LRX (21, 61); (**e**) AED (200); (**f**) AED-F (200); (**g**) LPD (11, 31); (**h**) CRD (21, 31). The parameter configurations (in parentheses) corresponds either to the inclusion and exclusion window or to the area of connected components to be removed.

4. Discussions

The results showed that all adopted anomaly detection methods were principally capable of detecting multiple targets in MUCAD and MUCAD-VI, clearly indicating the applicability of hyperspectral anomaly detection methods for camouflage detection in multispectral imagery. However, none of the detectors could achieve strong and stable results over all targets, although LRX was the only method that obtained more than 0.9 AUC for all targets. Yet, the detection performances of all methods including LRX heavily fluctuated across all targets. For instance, LPD outperformed all other methods for the relatively

small person targets but achieved considerably worse results than LRX for all other targets, which indeed works entirely different but operates in a local neighborhood very similar to LPD. In addition, the AED-based methods achieved top detection performances for the artificial grass mat, the green tarpaulin, and the gray tarpaulin but far worse results for all other targets. The high detection rates for those specific targets might be induced by the very homogenous surfaces and strong spectral distinctiveness compared to the other targets. These observations indicate that each algorithm heavily depends on the target's properties for successful detections. As already indicated in the introduction, it is the nature of camouflage that its spatial and spectral properties are usually unknown in real-world reconnaissance scenarios, which makes it very difficult to account for these strong dependencies. Fusing anomaly maps of different (or differently configured) target detectors that were empirically evaluated to work for targets featuring specific properties with very low false alarm rates could be an applicable approach to mitigate that issue to some extent.

It must also be noted that all detectors generally depend on strong spectral differences between targets and their surroundings, which is a logical requirement for the detectors to work but must be considered for the evaluation of each detector. For example, the artificial hedge lacks strong spectral distinctiveness in its environment as it is occluded by shadows, but all bands except LWIR of MUCAD and MUCAD-VI depend on the target's reflectance properties which are naturally less prominent in low light environments. Therefore, the algorithms had to detect targets with very little anomalous spectral characteristics, which is of course more difficult than the detection of targets with strong anomalous properties. The artificial grass mat, the green tarpaulin, and the grey tarpaulin possess very distinctive and unique spectral signatures, which resulted in significantly high detection rates at low false alarm rates for all algorithms compared to all other targets. In this context, it becomes apparent that the successful detection of certain targets in certain environments depends not only on the detection algorithms but on the underlying multispectral sensor setup or dataset, as well. Since the environments and targets captured in MUCAD are limited, the results account only for a very limited scope of possibilities. For different targets and different environments, the multispectral sensor setup might be not suitable to provide distinctive spectral signatures for each target, which in turn might then not be detected by any algorithm.

Although the selection of vegetation indices added to MUCAD in MUCAD-VI may seem arbitrary, the results indicate that a combination of the existing bands can have a positive effect on the performance of the target detectors. Since the sensor setup or the dataset is usually fixed in terms of bands and spectral raw information, an extension of more deliberately chosen combinations of bands might even further improve the detection performances compared to a selection of already predefined vegetation indices.

In terms of computational complexity and runtimes, all target detectors except CRD were able to work at near real-time conditions under most circumstances. Only a few parameter configurations led to runtimes slower than one second, although the number of bands to process had some influence, as well. In real-world applications with near real-time requirements, parameter configurations and the number of bands under consideration must be properly selected to avoid excessive runtimes of sensitive algorithms. However, it must also be noted that all runtimes were measured on a currently very powerful CPU, which might not be available in real-world scenarios. Despite its strong detection performance, the implementation of CRD in this work always requires much more time than a single second and does therefore not fulfill near real-time requirements, making it currently unsuitable for any real-world application.

Future research will focus on further improving the detection results by collecting more multispectral data at different seasons and by investigating how multispectral raw bands can be effectively enhanced, (i.e., by combining them) and selected in order to maximize the visibility of camouflage in different environments.

5. Conclusions

From the evaluation of the detection performances, it can be concluded that the hyperspectral anomaly detection methods investigated in this paper can principally be used for camouflage detection in multispectral imagery. However, since all methods already showed strong variations in detection performance for the limited number of different targets and environments in MUCAD, it must be assumed that even the best detection algorithm could miss some targets in real-world applications. As failing to detect a camouflaged hostile unit could prove fatal in a real-world reconnaissance scenario, additional research must be conducted to further improve detection results.

In addition to the individual detector performances, the limited spectral information provided by any multispectral sensor setup or dataset must also be taken into account. Some spectral bands may not contain enough information for certain targets and environments in order to lead to successful target detections. As shown in the results, enhancing the data with carefully selected derived bands, such as vegetation indices, can successfully mitigate that issue and result in higher detection rates of the algorithms.

Since multispectral imagery is not as complex as its hyperspectral counterpart, nearly all detectors except CRD performed at near real-time requirements. However, some algorithms quickly slowed down as the number of bands to process increased or their parameters changed, which must be considered in a tactical reconnaissance scenario where time is of the essence. The number of bands to process should therefore be kept as low as possible and the parameter configuration carefully selected to minimize computation time.

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

Funding: This research was funded by the Federal Office of Bundeswehr Equipment, Information Technology, and In-Service Support (BAAINBw). The APC was funded by the Universität der Bundeswehr München (UniBwM).

Data Availability Statement: The multispectral dataset for camouflage detection (MUCAD) is publicly available on GitHub: https://github.com/Tobias-UniBwM/MUCAD (accessed on 29 April 2022).

Acknowledgments: The authors sincerely thank Tobias Kreutz for his support in creating the multispectral dataset for camouflage detection (MUCAD).

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

References

1. Chen, Y.; Chen, X.; Zhou, J.; Ji, Y.; Shen, W. Camouflage target detection via hyperspectral imaging plus information divergence measurement. In Proceedings of the International Conference on Optoelectronics and Microelectronics Technology and Application, Nanjing, China, 20–22 October 2020; Su, Y., Xie, C., Yu, S., Zhang, C., Lu, W., Capmany, J., Luo, Y., Nakano, Y., Hao, Y., Yoshikawa, A., et al., Eds.; SPIE: Bellingham, WA, USA, 2017; Volume 10244, pp. 80–92.
2. Kumar, V.; Ghosh, J.K. Camouflage Detection Using MWIR Hyperspectral Images. *J. Indian Soc. Remote Sens.* **2017**, *45*, 139–145. [CrossRef]
3. Candiago, S.; Remondino, F.; De Giglio, M.; Dubbini, M.; Gattelli, M. Evaluating Multispectral Images and Vegetation Indices for Precision Farming Applications from UAV Images. *Remote Sens.* **2015**, *7*, 4026–4047. [CrossRef]
4. Zhou, X.; Zheng, H.B.; Xu, X.Q.; He, J.Y.; Ge, X.K.; Yao, X.; Cheng, T.; Zhu, Y.; Cao, W.X.; Tian, Y.C. Predicting grain yield in rice using multi-temporal vegetation indices from UAV-based multispectral and digital imagery. *ISPRS J. Photogramm. Remote Sens.* **2017**, *130*, 246–255. [CrossRef]
5. Zhou, J.; Yungbluth, D.; Vong, C.N.; Scaboo, A.; Zhou, J. Estimation of the Maturity Date of Soybean Breeding Lines Using UAV-Based Multispectral Imagery. *Remote Sens.* **2019**, *11*, 2075. [CrossRef]
6. Yu, J.; Wang, J.; Leblon, B. Evaluation of Soil Properties, Topographic Metrics, Plant Height, and Unmanned Aerial Vehicle Multispectral Imagery Using Machine Learning Methods to Estimate Canopy Nitrogen Weight in Corn. *Remote Sens.* **2021**, *13*, 3105. [CrossRef]
7. Adam, E.; Mutanga, O.; Rugege, D. Multispectral and hyperspectral remote sensing for identification and mapping of wetland vegetation: A review. *Wetl. Ecol. Manag.* **2010**, *18*, 281–296. [CrossRef]

8. Davidson, S.J.; Santos, M.J.; Sloan, V.L.; Watts, J.D.; Phoenix, G.K.; Oechel, W.C.; Zona, D. Mapping Arctic Tundra Vegetation Communities Using Field Spectroscopy and Multispectral Satellite Data in North Alaska, USA. *Remote Sens.* **2016**, *8*, 978. [CrossRef]
9. Boon, M.A.; Tesfamichael, S. Wetland vegetation integrity assessment with low altitude multispectral uav imagery. *Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* **2017**, *XLII-2/W6*, 55–62. [CrossRef]
10. Erinjery, J.J.; Singh, M.; Kent, R. Mapping and assessment of vegetation types in the tropical rainforests of the Western Ghats using multispectral Sentinel-2 and SAR Sentinel-1 satellite imagery. *Remote Sens. Environ.* **2018**, *216*, 345–354. [CrossRef]
11. Tait, L.; Bind, J.; Charan-Dixon, H.; Hawes, I.; Pirker, J.; Schiel, D. Unmanned Aerial Vehicles (UAVs) for Monitoring Macroalgal Biodiversity: Comparison of RGB and Multispectral Imaging Sensors for Biodiversity Assessments. *Remote Sens.* **2019**, *11*, 2332. [CrossRef]
12. Žížala, D.; Minarík, R.; Zádorová, T. Soil Organic Carbon Mapping Using Multispectral Remote Sensing Data: Prediction Ability of Data with Different Spatial and Spectral Resolutions. *Remote Sens.* **2019**, *11*, 2947. [CrossRef]
13. Ozigis, M.S.; Kaduk, J.D.; Jarvis, C.H.; da Conceição Bispo, P.; Balzter, H. Detection of oil pollution impacts on vegetation using multifrequency SAR, multispectral images with fuzzy forest and random forest methods. *Environ. Pollut.* **2020**, *256*, 113360. [CrossRef] [PubMed]
14. Quan, Y.; Zhong, X.; Feng, W.; Dauphin, G.; Gao, L.; Xing, M. A Novel Feature Extension Method for the Forest Disaster Monitoring Using Multispectral Data. *Remote Sens.* **2020**, *12*, 2261. [CrossRef]
15. Minařík, R.; Langhammer, J.; Lendzioch, T.; Alvarez Taboada, F.; Govedarica, M. Detection of Bark Beetle Disturbance at Tree Level Using UAS Multispectral Imagery and Deep Learning. *Remote Sens.* **2021**, *13*, 4768. [CrossRef]
16. Naik, P.; Dalponte, M.; Bruzzone, L. Prediction of Forest Aboveground Biomass Using Multitemporal Multispectral Remote Sensing Data. *Remote Sens.* **2021**, *13*, 1282. [CrossRef]
17. Chabot, D.; Dillon, C.; Shemrock, A.; Weissflog, N.; Sager, E.P.S. An Object-Based Image Analysis Workflow for Monitoring Shallow-Water Aquatic Vegetation in Multispectral Drone Imagery. *ISPRS Int. J. Geo-Inf.* **2018**, *7*, 294. [CrossRef]
18. Song, B.; Park, K. Detection of Aquatic Plants Using Multispectral UAV Imagery and Vegetation Index. *Remote Sens.* **2020**, *12*, 387. [CrossRef]
19. Rossi, L.; Mammi, I.; Pelliccia, F. UAV-Derived Multispectral Bathymetry. *Remote Sens.* **2020**, *12*, 3897. [CrossRef]
20. Shen, Y.; Li, J.; Lin, W.; Chen, L.; Huang, F.; Wang, S. Camouflaged target detection based on snapshot multispectral imaging. *Remote Sens.* **2021**, *13*, 3949. [CrossRef]
21. Reed, I.S.; Yu, X. Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution. *IEEE Trans. Acoust.* **1990**, *38*, 1760–1770. [CrossRef]
22. Li, W.; Du, Q. Collaborative representation for hyperspectral anomaly detection. *IEEE Trans. Geosci. Remote Sens.* **2015**, *53*, 1463–1474. [CrossRef]
23. Kang, X.; Zhang, X.; Li, S.; Li, K.; Li, J.; Benediktsson, J.A. Hyperspectral Anomaly Detection with Attribute and Edge-Preserving Filters. *IEEE Trans. Geosci. Remote Sens.* **2017**, *55*, 5600–5611. [CrossRef]
24. Kwon, H.; Nasrabadi, N.M. Kernel RX-algorithm: A nonlinear anomaly detector for hyperspectral imagery. *IEEE Trans. Geosci. Remote Sens.* **2005**, *43*, 388–397. [CrossRef]
25. Zhou, J.; Kwan, C.; Ayhan, B.; Eismann, M.T. A Novel Cluster Kernel RX Algorithm for Anomaly and Change Detection Using Hyperspectral Images. *IEEE Trans. Geosci. Remote Sens.* **2016**, *54*, 6497–6504. [CrossRef]
26. Guo, Q.; Zhang, B.; Ran, Q.; Gao, L.; Li, J.; Plaza, A. Weighted-RXD and linear filter-based RXD: Improving background statistics estimation for anomaly detection in hyperspectral imagery. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2014**, *7*, 2351–2366. [CrossRef]
27. Chen, Y.; Nasrabadi, N.M.; Tran, T.D. Sparse representation for target detection in hyperspectral imagery. *IEEE J. Sel. Top. Signal Processing* **2011**, *5*, 629–640. [CrossRef]
28. Li, J.; Zhang, H.; Zhang, L.; Ma, L. Hyperspectral anomaly detection by the use of background joint sparse representation. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2015**, *8*, 2523–2533. [CrossRef]
29. Xu, Y.; Wu, Z.; Li, J.; Plaza, A.; Wei, Z. Anomaly detection in hyperspectral images based on low-rank and sparse representation. *IEEE Trans. Geosci. Remote Sens.* **2016**, *54*, 1990–2000. [CrossRef]
30. Tu, B.; Yang, X.; Li, N.; Zhou, C.; He, D. Hyperspectral anomaly detection via density peak clustering. *Pattern Recognit. Lett.* **2020**, *129*, 144–149. [CrossRef]
31. Tu, B.; Yang, X.; Zhou, C.; He, D.; Plaza, A. Hyperspectral Anomaly Detection Using Dual Window Density. *IEEE Trans. Geosci. Remote Sens.* **2020**, *58*, 8503–8517. [CrossRef]
32. Carlotto, M.J. A Cluster-Based Approach for Detecting Man-Made Objects and Changes in Imagery. *IEEE Trans. Geosci. Remote Sens.* **2005**, *43*, 374–387. [CrossRef]
33. Evangelidis, G.D.; Psarakis, E.Z. Parametric Image Alignment Using Enhanced Correlation Coefficient Maximization. *IEEE Trans. Pattern Anal. Mach. Intell.* **2008**, *30*, 1858–1865. [CrossRef]
34. Bradski, G. The OpenCV Library. *Dr. Dobb's J. Softw. Tools* **2000**, *25*, 120–123.
35. OpenDroneMap Authors ODM—A Command Line Toolkit to Generate Maps, Point Clouds, 3D Models and DEMs from Drone, Balloon or Kite Images. Available online: https://github.com/OpenDroneMap/ODM (accessed on 3 March 2022).
36. Guennebaud, G.; Jacob, B. Others Eigen v3. Available online: http://eigen.tuxfamily.org (accessed on 3 March 2022).

37. Chandra, R.; Dagum, L.; Kohr, D.; Menon, R.; Maydan, D.; McDonald, J. *Parallel Programming in OpenMP*; Morgan Kaufmann: Burlington, MA, USA, 2001.
38. Behnel, S.; Bradshaw, R.; Citro, C.; Dalcin, L.; Seljebotn, D.S.; Smith, K. Cython: The best of both worlds. *Comput. Sci. Eng.* **2011**, *13*, 31–39. [CrossRef]
39. Breen, E.J.; Jones, R. Attribute Openings, Thinnings, and Granulometries. *Comput. Vis. Image Underst.* **1996**, *64*, 377–389. [CrossRef]
40. Tomasi, C.; Manduchi, R. Bilateral filtering for gray and color images. In Proceedings of the Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271), Bombay, India, 7 January 1998; pp. 839–846. [CrossRef]
41. Gastal, E.S.L.; Oliveira, M.M. Domain transform for edge-aware image and video processing. *ACM Trans. Graph.* **2011**, *30*, 1–12. [CrossRef]
42. Guiotte, F. Sap: Python Package to Easily Compute Morphological Attribute Profiles (AP) of Images. Available online: https://github.com/fguiotte/sap (accessed on 10 March 2022).
43. Rodriguez, A.; Laio, A. Clustering by fast search and find of density peaks. *Science* **2014**, *344*, 1492–1496. [CrossRef]
44. Fawcett, T. An introduction to ROC analysis. *Pattern Recognit. Lett.* **2006**, *27*, 861–874. [CrossRef]

5.2 Measuring and Predicting Sensor Performance for Camouflage Detection in Multispectral Imagery

Tobias Hupel and Peter Stütz. „Measuring and Predicting Sensor Performance for Camouflage Detection in Multispectral Imagery“. In: *Sensors* 23.19 (2023). ISSN: 1424-8220. DOI: 10.3390/s23198025

sensors

MDPI

Article

Measuring and Predicting Sensor Performance for Camouflage Detection in Multispectral Imagery

Tobias Hupel * and Peter Stütz

Institute of Flight Systems, University of the Bundeswehr Munich, 85577 Neubiberg, Germany; peter.stuetz@unibw.de
* Correspondence: tobias.hupel@unibw.de

Abstract: To improve the management of multispectral sensor systems on small reconnaissance drones, this paper proposes an approach to predict the performance of a sensor band with respect to its ability to expose camouflaged targets under a given environmental context. As a reference for sensor performance, a new metric is introduced that quantifies the visibility of camouflaged targets in a particular sensor band: the Target Visibility Index (TVI). For the sensor performance prediction, several machine learning models are trained to learn the relationship between the TVI for a specific sensor band and an environmental context state extracted from the visual band by multiple image descriptors. Using a predicted measure of performance, the sensor bands are ranked according to their significance. For the training and evaluation of the performance prediction approach, a dataset featuring 853 multispectral captures and numerous camouflaged targets in different environments was created and has been made publicly available for download. The results show that the proposed approach can successfully determine the most informative sensor bands in most cases. Therefore, this performance prediction approach has great potential to improve camouflage detection performance in real-world reconnaissance scenarios by increasing the utility of each sensor band and reducing the associated workload of complex multispectral sensor systems.

Keywords: multispectral; infrared; camouflage detection; target visibility; sensor performance; sensor management; performance modelling

Citation: Hupel, T.; Stütz, P. Measuring and Predicting Sensor Performance for Camouflage Detection in Multispectral Imagery. *Sensors* **2023**, *23*, 8025. https://doi.org/10.3390/s23198025

Academic Editors: Chuan-Ming Liu and Wei-Shinn Ku

Received: 11 August 2023
Revised: 18 September 2023
Accepted: 20 September 2023
Published: 22 September 2023

1. Introduction

Multispectral sensor systems have become quite popular for various remote sensing applications, ranging from precision agriculture [1–3], land cover classification [4,5], detection of weeds [6], and plant disease monitoring [7–9] to shoreline extraction [10], water body detection [11], bathymetry [12], and disaster evaluation [13]. Their relatively low cost, size, weight, and power consumption make them suitable for use even on small reconnaissance drones, where the rich spectral information they provide can be utilized to detect camouflaged targets [14]. However, compared to the visual or thermal infrared sensors commonly used in reconnaissance scenarios, multispectral sensors provide a much larger number of bands, including derivatives such as vegetation indices (e.g., NDVI and NDRE). This additional information introduces a substantially heavier workload that must be managed by a sensor operator and possibly by any subsequent computer-aided processing system. For this reason, the Institute of Flight Systems at the University of the Bundeswehr, Munich, Germany is actively researching the use of multispectral sensor systems on small tactical drones in military reconnaissance scenarios.

Because each material has unique spectral characteristics, sensor bands that expose camouflaged targets in one environment, such as grassland, may not expose camouflaged targets in another environment, such as gravel. Knowing when to utilize which sensor band under given environmental conditions is usually based on experience and empirical experimentation. The large number of possible bands provided by multispectral sensor

systems makes the selection of the most useful sensor bands an even more complex task, especially in time-critical military reconnaissance scenarios. Therefore, this work presents an approach to address this issue by predicting the performance of a sensor band with respect to its ability to expose camouflaged targets. More specifically, each sensor band is linked to a performance model that predicts its performance by assessing the current environmental situation. Having a measure of performance for each sensor band of a multispectral sensor system at flight time, the sensor bands can be ranked from those providing the most performance to those providing the least performance. Moreover, the sensor bands can be reduced to the most meaningful ones, leaving the sensor operator or any subsequent processing instance with a mere subset of all sensor bands. This subset is processed more quickly and is more likely to contain the information needed to detect camouflaged targets.

In order to quantify the performance of a sensor band, this work introduces the Target Visibility Index (TVI). The TVI is a metric that provides a measure of the extent to which a given sensor band exposes a camouflaged target. Using the TVI as a reference for sensor performance, machine learning models can be trained to learn the relationship between the current environmental situation and the corresponding performance of a given sensor band. After training, these machine learning models can be employed as performance models for the sensor bands of a multispectral sensor system, where they dynamically assess the environmental situation and predict the performance of their associated sensor band. Here, the environmental situation is represented in an abstract way by a so-called context state. The context state is a feature vector extracted by multiple feature descriptors from a preselected sensor band of the multispectral sensor system. In conclusion, the performance models technically learn the relationship between the context state and the TVI of their associated sensor band.

For the training of the performance models and the evaluation of their predictions, a custom dataset featuring 853 multispectral captures containing several different camouflaged targets in various environments at different seasons was compiled. To support reproducibility and enable further research, the dataset has been made publicly available for download (see the Data Availability Statement at the end of this manuscript).

In summary, this work makes the following scientific contributions:

- Proposition and evaluation of a method for predicting sensor performance with respect to the exposure of camouflaged targets.
- Introduction of a metric for measuring sensor performance with respect to the exposure of camouflaged targets.
- Provision of an extensive multispectral dataset containing multiple camouflaged targets: the eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X).

1.1. Related Work

Although research related to sensor performance modeling and prediction is scarce, there have been a number of recently conducted relevant studies. In [15], sensor performance models were used to map selected environmental states to the detection performance of object detection algorithms for flight trajectory optimization using an optimal control approach. Incorporating these sensor performance models into the optimization procedure allowed the computation of flight trajectories that maximized the detection performance of the object detection algorithms. In [16], object detection models were used to support the sensor scheduling algorithm by predicting the probability of successful object detections given the current environmental and UAV conditions for UAV-based multi-object tracking applications with limited sensor capabilities, leading to significantly improved object observation times. In other research, the most suitable detection algorithm has been dynamically selected aboard a sensor-equipped UAV under given environmental conditions through modelling and predicting the performance of several object detection algorithms using Bayesian Networks [17] or artificial neural networks and fuzzy inference [18]. Both of these approaches were able to substantially increase overall object detection performance.

However, the prediction of sensor performance with respect to the exposure of camouflaged targets has not yet been explored, motivating the work presented in this manuscript.

The measurement of visibility or exposure of targets in a dynamic environment is a highly active field of research, especially in the automotive area, where traffic lights and signs have to be designed in such way that they cannot be overlooked by any road user. Visibility metrics based on luminance measurements and psychological behavior, such as the target visibility level [19] and the relative visual performance [20], have been proposed and evaluated in various scenarios [21–24]. The determination of visibility in terms of the distance at which objects can be identified from visual [25,26] and near-infrared [25] camera footage has been studied as well. For detecting the most salient regions and objects in an image according to human perception, a number of approaches have been introduced [27–29]. These methods generate a saliency map from an input image, highlighting those regions that the human eye would naturally focus on first. However, visibility metrics based on human perception in real-world scenes or laboratory environments are not applicable to the use case considered in this work, nor are visibility metrics in the form of viewing distances. Furthermore, saliency maps are expensive to compute and difficult to translate into a single sensor performance score. Therefore, in this paper we introduce a computationally inexpensive metric based on contrast [30] and statistical properties that have already been used for other image metrics [31,32].

1.2. Outline

In the following section (Section 2), the dataset, Target Visibility Index, and sensor performance prediction approach are introduced and explained in detail. The next section (Section 3) covers the evaluation and comparison of the machine learning models and their different training procedures with respect to their ability to determine the most informative bands given the context state. Finally, the results and their significance are discussed in Section 4, and summarized conclusions are drawn in Section 5.

2. Methods and Materials

This section introduces the dataset used to train and evaluate the proposed sensor performance prediction method in Section 2.1, the metric for target visibility in Section 2.2, and the proposed method itself in Section 2.3.

2.1. Dataset

The data used to train and evaluate the proposed sensor performance prediction approach were collected in two different areas of the test site at the University of the Bundeswehr, Munich. The areas shown in Figure 1 provided a variety of different environments, such as grassland, gravel and graveled soil, various bushes and trees, and both concrete and asphalt roads. This diversity constitutes an excellent foundation for a rich and comprehensive dataset. For the camouflaged targets, thirteen different objects were placed in visually similar environments in each of these areas: a piece of artificial turf, an artificial hedge, a green tarp, a green 2D camouflage net, a green 3D camouflage net, a gray tarp, an anthracite fleece, a gray 3D camouflage net, a yellow 3D camouflage net, and four persons, two wearing green uniforms and two wearing yellow uniforms. All targets are listed in Table 1 along with the corresponding target group indicating the type of environment in which the target was placed (i.e., green targets in green environments). In addition, the table shows the percentage distribution of the targets and target groups. As can be seen, the green targets dominate the data, making up almost two thirds, while the gray and yellow targets occupy about one and two ninths, respectively. This is due to the greater number of green targets compared to the number of yellow and gray targets as well as to the nature of the captured areas, which are dominated by green environments. Figure 2 shows the artificial hedge, the green 2D camouflage net, the gray 3D camouflage net, the green 3D camouflage net, the yellow 3D camouflage net, and the artificial turf in their corresponding environments from the ground perspective. Note that 3D camouflage nets have a more

irregularly structured surface than 2D camouflage nets, which are mostly flat and similar to a tarp.

(**a**) (**b**)

Figure 1. The two areas of the test site at the University of the Bundeswehr Munich (in May) where the camouflaged targets were placed: (**a**) area A; (**b**) area B.

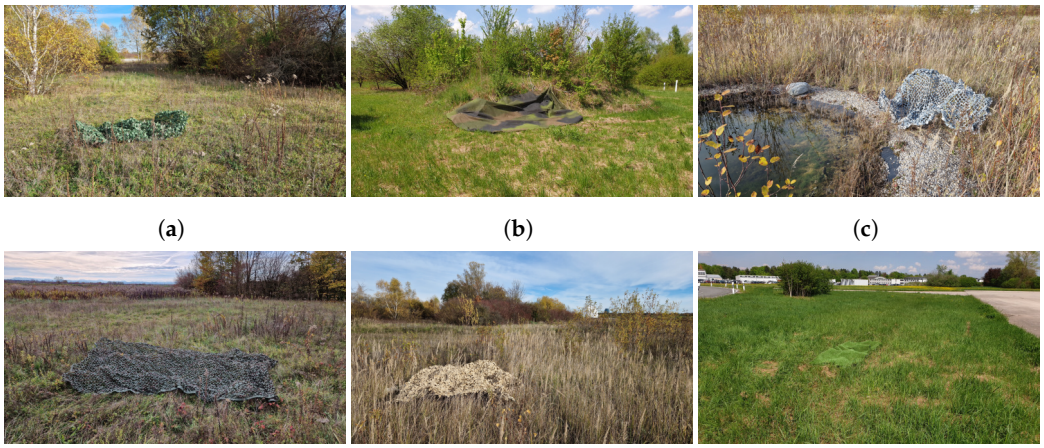

(**a**) (**b**) (**c**)

(**d**) (**e**) (**f**)

Figure 2. Multiple different camouflaged targets of the dataset from the ground perspective; all targets were placed in environments where they easily integrate: (**a**) artificial hedge, (**b**) green 2D camouflage net, (**c**) gray 3D camouflage net, (**d**) green 3D camouflage net, (**e**) yellow 3D camouflage net, (**f**) artificial turf.

For acquisition of the multispectral data, the camouflaged targets were placed in one of the areas and captured from the nadir perspective by an unmanned aerial vehicle (UAV). The UAV was equipped with a multispectral sensor system providing the bands described in Table 2. After each capture flight, the objects were placed in a different environment of the same area and captured again, resulting in seven different locations for each target in both areas (the battery life of the UAV limited the number of capture flights per area to seven). When the first area had been captured seven times, the same process was repeated for the second area. The entire capture process was conducted on three different days in three different seasons: spring in May, summer in August, and autumn in November. This was done to provide variety in the data in order to make the results as meaningful as possible.

Consequently, all camouflaged targets were captured in fourteen different environments (seven locations in two areas) in three different seasons, with only a few exceptions:

- The yellow 3D camouflage net was not used in area A in summer or autumn, as the environment was all green and no appropriate spot could be found for it.
- Only four capture flights over area B were conducted in summer, as the UAV broke during the experiments and could not be repaired in time.
- The yellow 3D camouflage net was left in the same place on all four summer capture flights in area B, as it had been overlooked when the camouflaged targets were rearranged.

Table 1. All thirteen camouflaged targets and their corresponding target group captured in the dataset. The percentages show the proportion of each target or target group among the annotations in the dataset.

Camouflaged Target		Group
artificial turf	9.3%	green 65.8%
artificial hedge	9.4%	
green tarp	9.2%	
green 2D camouflage net	9.9%	
green 3D camouflage net	9.6%	
2 persons in green uniforms	18.4%	
gray tarp	3.1%	gray 11.4%
anthracite fleece	2.2%	
gray 3D camouflage net	6.2%	
yellow 3D camouflage net	5.9%	yellow 22.8%
2 persons in yellow uniforms	16.9%	

Table 2. The bands and their associated properties provided by each capture of the dataset.

Band	Center Wavelength	Bandwidth
visual (VIS)	-	-
blue	475 nm	32 nm
green	560 nm	27 nm
red	668 nm	14 nm
edge-infrared (EIR)	717 nm	12 nm
near-infrared (NIR)	842 nm	57 nm
long-wave infrared (LWIR)	10.5 µm	6 µm

The final dataset, called eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X), was not derived directly from the acquired data, instead being derived from orthophotos generated separately for each sensor band. Two of these orthophotos, computed from the visual band images, have already been shown in Figure 1. The orthophotos were generated with a ground sample distance (GSD) of $10\frac{\text{cm}}{\text{px}}$ using the command line toolkit *Open Drone Map* [33]. Using a sliding window with a resolution of 512 by 512 pixels, the captures of the final dataset were cropped from the orthophotos to ensure that each capture contained at least a single camouflaged target. In addition, the individual sensor bands of each capture were pixel-aligned using Enhanced Correlation Coefficient Maximization [34] provided by the computer vision library *OpenCV* [35]. Figure 3 shows a sample capture of the dataset with all bands from VIS to LWIR (Figure 3a–g) and multiple different camouflaged targets in the scene that are identified by the ground truth mask in Figure 3h. In total, the final dataset contained 853 annotated and pixel aligned multispectral captures, each with a resolution of 512 by 512 pixels, a GSD of $10\frac{\text{cm}}{\text{px}}$, and containing at least a single camouflaged target. The ground truth masks were created using the Computer Vision Annotation Tool v2.3.0 [36].

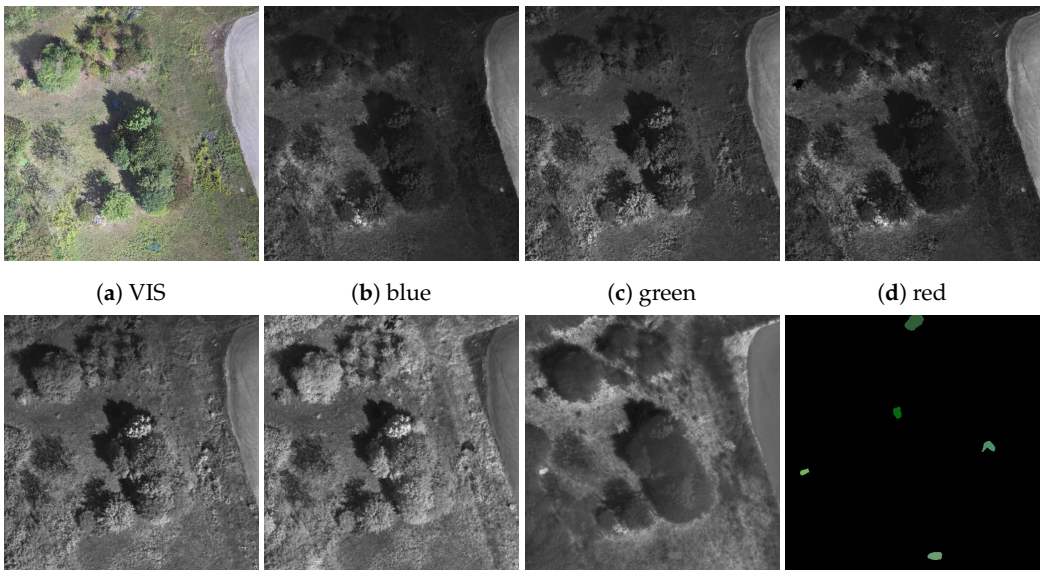

(**a**) VIS (**b**) blue (**c**) green (**d**) red

(**e**) EIR (**f**) NIR (**g**) LWIR (**h**) label

Figure 3. Sample capture of the dataset, with bands from VIS (**a**) to LWIR (**g**) and a ground truth mask (**h**) identifying all captured camouflaged targets. Note that each camouflaged target is denoted in a different color, making for five different targets in the scene.

2.2. Measuring Sensor Performance

In order to train a machine learning model to predict the extent to which camouflaged targets are exposed in a given sensor band, a metric describing that extent is first required. Because the prediction is made for the entire sensor band as a single unit, this metric must consider the entire sensor band. For this purpose, in this paper we introduce the Target Visibility Index (TVI), provided in Equation (1), where μ_T is the mean over all pixel values belonging to the camouflaged target, μ_B is the mean over all pixel values belonging to the background, σ_T is the standard deviation of all pixel values belonging to the camouflaged target, and σ_B is the standard deviation of all pixel values belonging to the background. The mean and standard deviation are computationally efficient and commonly used statistical properties in well-established image metrics for a wide range of problems [31,32]. As such, they were employed for the TVI.

$$\text{TVI} = (\mu_T - \mu_B)^2(1 - 2\sigma_T)^2(1 - 2\sigma_B)^2 \tag{1}$$

In general, the TVI is based on the idea that an ideal sensor band exposes a camouflaged target as much as possible, which is illustrated in Figure 4. The visual image in Figure 4a shows a scene containing a single camouflaged target, the green 3D camouflage net. According to the TVI, the corresponding ideal sensor band for that exact same scene is depicted in Figure 4b. All pixel values belonging to the camouflaged target differ as much as possible from all pixel values belonging to the background, resulting in the highest possible value of the TVI (1.0). The first factor of the TVI ($(\mu_T - \mu_B)^2$) serves as an approximate measure of fulfillment of that property. It is zero when the difference between the camouflaged target pixel values and the background pixel values is zero and one when the difference between the camouflaged target pixel values and the background pixel values is maximal. Thus, it can be interpreted as a measure of contrast [30] between the target and the background. However, the difference between the mean values does not sufficiently describe the extent to which the target is exposed, as can be seen in Figure 4c,d. In both bands, the respective means over all pixel values belonging to the camouflaged target are identical, as are the respective means over all pixel values belonging to the background. Consequently, the first factor of the TVI ($(\mu_T - \mu_B)^2$) yields the same result (0.04) in both cases. However, the difference in means in Figure 4d results from two different distributions

of two pixel values, while in Figure 4c it results from a difference of two constant pixel values. Considering the exemplary case that the difference of mean values is already at its maximum, a band such as that in Figure 4c would most likely be preferable over the one in Figure 4d in an actual reconnaissance scenario. Therefore, the metric must take into account the distribution of the pixel values belonging to the background and the distribution of the pixel values belonging to the camouflaged target. To ensure that small spreads in the pixel value distributions are preferable over large spreads, the TVI implements the second $((1-2\sigma_T)^2)$ and third $((1-2\sigma_B)^2)$ factors, which penalize large spreads of pixel values and favor small spreads. In essence, the greater the TVI, the closer the camouflaged target pixel values and the background pixel values are to each other, respectively. Each factor equals zero if the spread of the respective pixel values is maximal and one if there is no spread at all. Because there is no spread in both distributions in Figure 4c, both factors are one and the first factor determines the final TVI. In contrast, the spread in both distributions in Figure 4d is close to the maximum, resulting in a TVI close to zero. However, there are limits, as shown in Figure 4e. Although the standard deviations for the target and background pixels are zero, their means are equal. As a result, the first factor of the TVI equals zero, leading to a TVI of zero as well. Eventually, the TVI can only be maximal when there is minimal spread in both camouflaged target pixel values and background pixel values, and when the difference between their mean values is maximal.

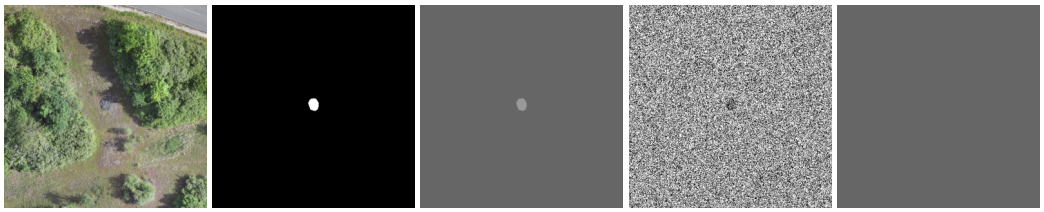

(**a**) VIS (**b**) ideal (1.0) (**c**) good (0.04) (**d**) bad (≈0.0) (**e**) bad (0.0)

Figure 4. Demonstration of camouflaged target visibility, where (**b**) corresponds to an ideal band for the scene depicted in (**a**). Likewise, (**c**) corresponds to a band with good visibility of the target, while (**d**,**e**) correspond to a band with poor or no visibility of the target. The associated TVIs are shown in parentheses.

The TVI is designed for single-channel images and a range of values from zero to one. Other ranges must be normalized, or the TVI will produce inconclusive results. Each individual factor of the TVI ranges between zero and one. If one of the means is zero and the other is one, then the first factor of the TVI is one. If the means are equal, then the first factor is zero regardless of their actual values. Because the theoretical maximum of the standard deviation is 0.5 for a range of values from zero to one [37], the second and third factors of the TVI are zero for the maximum standard deviation and one for zero standard deviation. With all factors ranging between one and zero, the TVI range is between zero and one. The factors are multiplicatively combined to prevent a strong individual factor from outweighing a weak individual factor, which would be possible in an additive combination, for instance. Additionally, each factor is squared to avoid negative values and retain the differentiability of the TVI, which might be useful if the TVI is used for numerical optimization problems.

Figure 5 shows real-world examples of the TVI. As can be observed, the TVI is very low for the blue band (Figure 5a) and relatively high for the EIR band (Figure 5b) and NIR band (Figure 5c). This is consistent with the expected behavior of the metric, as the target appears to be much more exposed in the EIR and NIR bands than in the blue band. It can be seen that the TVI generally produces relatively low values for real-world images, even when the camouflaged target is easily distinguishable from its surroundings. At this point, it is important to note that the design of the TVI is based on those edge cases where it is equal to either one, as shown in Figure 4b, or zero, as shown in Figure 4d,e. The closer the sensor band is to one of these edge cases, the closer the TVI is to zero or one, where

closer is mathematically defined by the factors provided in Equation (1). For any TVI value in between these edge cases, its true meaning in terms of target visibility is difficult to determine and does not necessarily correspond to human perception. For example, if the target in Figure 5c were placed in the shadows of the tree line immediately next to it, it would be much less visible to the human eye; however, in terms of the TVI, the visibility of the target would be roughly the same in both cases, as the change in the mean and standard deviation of the background pixels would be negligible. Notably, actual human perception of visibility is currently the subject of active research, which is beyond the scope of this work except for those trivial edge cases in which the TVI generates predefined values of zero and one. The TVI quantifies target visibility as a single value in a computationally efficient and comparable way, which naturally involves approximation; ultimately, this is necessary in order to measure and compare the extent to which a target is exposed in different sensor bands.

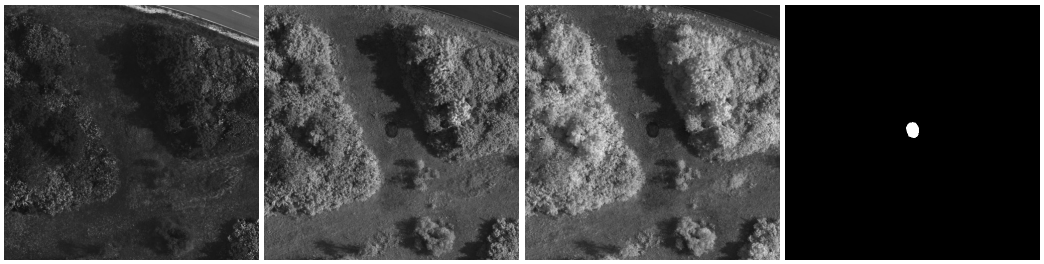

(**a**) blue (≈0.0) (**b**) EIR (0.013) (**c**) NIR (0.025) (**d**) ideal (1.0)

Figure 5. Demonstration of the Target Visibility Index (TVI), producing relatively low values for bad visibility of the target in (**a**) and relatively high values for good target visibility in (**b**,**c**). The ideal band in (**d**) results in a TVI of 1.

2.3. Predicting Sensor Performance

To predict sensor performance, in this paper we introduce the concept illustrated in Figure 6. Considering a multispectral sensor system with multiple different bands, a preselected context band is used to extract a context state that provides abstract information about the current environment and scenery. From the context state, the individual performance models predict the performance of their associated sensor bands. In the illustrated example, the predicted performance is high for band D, medium for bands A and C, and low for band B. Finally, the sensor bands are ranked by their performance predictions in order to obtain the subset of sensor bands that is most likely to provide the highest visibility of camouflaged targets. This greatly reduces the amount of information that must be processed in any subsequent evaluation instance.

In this work, the context state is extracted from the gray-level converted visual band using 16 bit rotation-invariant uniform local binary patterns (LBP_{16}^{riu2} operator) [38] and the fourteen Haralick features [39]. Both are computationally inexpensive and common choices for feature extractors in image classification problems [40–42], where an abstract representation of the scene to be classified is required as well. Table 3 shows the final composition of the context state. The LBPs were extracted with a radius of 2 px and a resolution of 16 using the implementation of the ImageFeatures.jl package of the Julia Programming Language [43]. To obtain the first eighteen values of the context state, the histogram over the extracted rotation-invariant uniform patterns was computed. Each value of the histogram represents the number of occurrences of each individual pattern. Because there are exactly seventeen rotation-invariant uniform patterns for a bit size of 16, the resulting histogram holds eighteen values, where the last one accounts for the occurrences of all non-uniform patterns. In the last step, the histogram is normalized to ensure that the values of the histogram sum to one. The remaining Haralick features of the context state were computed using the Python package Mahotas [44]. Each value of the first fourteen Haralick features is an average of four individual features values produced by four different gray-level co-occurrence matrices, each generated for a radius of 1 px and

the directions left, right, up, and down. The second fourteen Haralick features contain the differences between the maximum and minimum values generated by each of the four individual gray-level co-occurrence matrices. In total, the context state consists of 46 features that abstractly describe the environmental situation based on the preselected context band.

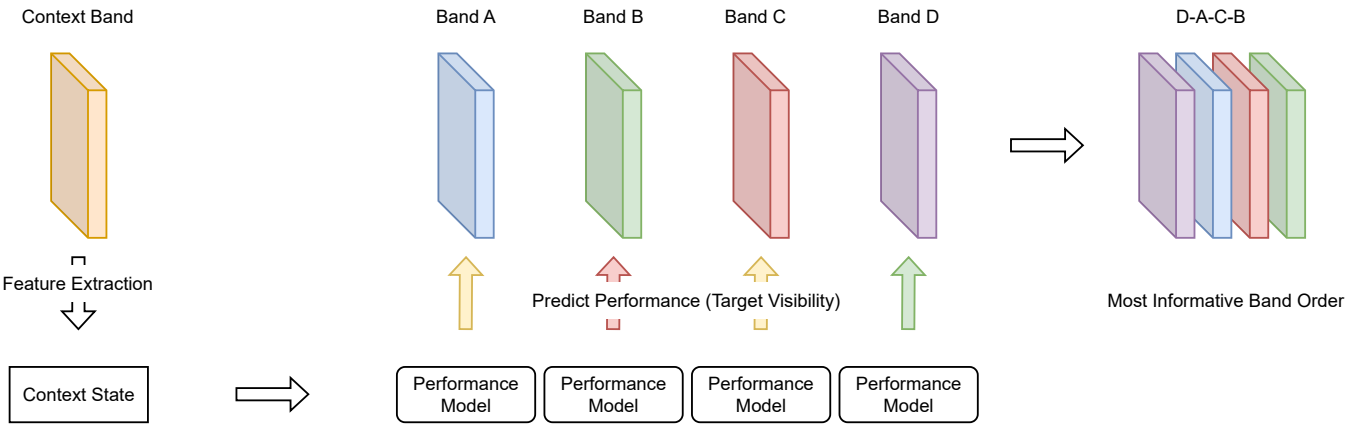

Figure 6. Conceptual basis of sensor performance prediction (target visibility). First, the context state is extracted by image descriptors from a preselected context band. Based on the context state, each performance model then predicts the target visibility for its associated band. The bands can be sorted after the predictions are made, allowing them to be ordered from the most informative to the least informative band. The green, yellow, and red prediction arrows indicate good, medium, and bad performance, respectively.

Table 3. The structure of the context state extracted from the context band (VIS) using local binary patterns and Haralick features. The numbers correspond to the feature value positions of the context state.

LBP		**Haralick**	
uniform	non-uniform	mean	min-max
1–17	18	19–32	33–46

For the performance models that predict the sensor band performance from the context state, three machine learning methods for regression tasks were applied: ϵ-Support Vector Regression (ϵ-SVR), Random Forests (RFs), and Gradient Boosted Trees (GBTs). All are based on different concepts, have been thoroughly studied, and are commonly used for complex regression tasks. In addition, their training is efficient and a robust implementation is usually available for the most popular programming languages. Therefore, they were chosen for the regression task in this work. The ϵ-SVR, RFs, and GBTs were trained and evaluated using the common interface of the Machine Learning Framework for Julia [45], where the models relied on the LIBSVM [46], DecisionTree.jl [47], and XGBoost [48] backends, respectively. The parameter selections used to train the models are introduced in Section 3.1.

3. Experiments and Results

This section first introduces the parameters and data used to train the machine learning models (i.e., performance models) in Section 3.1. Afterwards, the evaluation of the prediction performances of all models is presented in Section 3.2.

3.1. Training

In order to divide the dataset introduced in Section 2.1 into training and test data, 80% of the captures were randomly selected as training data and the remaining 20% were

used to evaluate the models. In each capture, for each band except VIS the context state was extracted from the gray-level converted visual band as described in Section 2.3 and mapped to a single TVI. Because there are six bands (blue, green, red, EIR, NIR and LWIR), each context state maps to six different TVIs per capture. In the case of multiple camouflaged targets in the capture, the resulting TVIs had to be reduced to a single value. This was achieved by averaging all of the individual TVIs calculated separately for each camouflaged target. Thus, a camouflaged target belongs to the background of every other camouflaged target in the scene. Although averaging could dilute the mappings from the context state to the TVI, it is able to consider all targets in the scene equally for the single sensor performance value. The context states were additionally z-normalized, leading to a mean value of zero and a standard deviation of one for each feature value over all context states. The means and standard deviations required for the normalization were calculated for the training data only, then applied to both the training and test data.

Considering a reconnaissance scenario in which a priori knowledge on the camouflaged targets is available, it could be beneficial to employ performance models that are able to account for this additional knowledge. For example, if the camouflaged targets are known to be located in green environments such as woods, bushes, and grass, a performance model trained only on targets commonly used in these environments may outperform a model trained on additional kinds of targets. Therefore, the models were additionally trained on data for which the resulting TVI for each sensor band was not calculated over all camouflaged targets in the scene, only over those belonging to specific target group, as has been already introduced above in Table 1. This reduces the potential dilution caused by averaging over the TVIs of targets in different target groups. Because not every capture in the dataset contains a camouflaged target of each target group, the training and testing splits and feature normalization for the specialized models were performed only on the number of captures that actually contained a target of the respective target group. With three different target groups, the models were trained on a total of four different data variations: one for each of the target groups, and one in which the target groups were ignored. With six bands for each capture, three different machine learning models, and four different data variations, a total of 72 models were trained. After training, the models use a normalized feature vector extracted from a gray-level converted visual band to predict the TVI for their associated sensor band. For the models trained on data where the TVI was calculated only for a specific target group, their predictions consider only the targets belonging to that specific target group. In contrast, the predictions of the models trained on data containing all camouflaged targets consider all target groups.

The optimal parameters for each model were searched by a simple grid scan based on cross-validation over the training data with five folds and utilizing the root mean square error (RSME). The four-part Table 4 shows the final training parameter configurations for the models considering all target groups, only targets belonging to the green target group, only targets belonging to the gray target group, and only targets belonging to the yellow target group, respectively. Note that the predictions of the individual models were not evaluated further; instead, the evaluation of the models is based on comparisons of the predicted most informative band orders with the actual most informative band orders, which is explained in detail in Section 3.2.

Table 4. The final training parameter configurations for each machine learning model (i.e., performance model), including the models that considered targets of any target group and the models that considering only targets belonging to one of the green, gray, or yellow target group. Unmentioned parameters were retained at the default values provided by their respective implementations. The trees, leaves, split, features, and fraction parameters of the Random Forest model specify the number of trees, minimum number of samples belonging to a single leaf node, minimum number of samples required for further splitting, number of random subfeatures for each tree, and fraction of random training samples for each tree, respectively. The rounds and depth parameters of the XGBoost model represent the maximum depth of each tree and the number of boosting rounds, respectively.

	ϵ-SVR		**Random Forest**					**XGBoost**			
	ϵ	**C**	**Trees**	**Leaves**	**Split**	**Features**	**Fraction**	**Rounds**	η	**Depth**	λ
any target models											
blue	0.00412	0.044	46	3	15	$\sqrt{46}$	1.0	200	0.05	6	2.5
green	0.00162	0.06	46	3	7	$\sqrt{46}$	1.0	400	0.025	2	0.001
red	0.0041	0.041	28	2	16	$\sqrt{46}$	0.8	450	0.1	2	30
EIR	0.0018	0.022	14	1	8	$\sqrt{46}$	0.9	275	0.04	6	0.1
NIR	0.00112	0.026	28	1	4	$\sqrt{46}$	1.0	125	0.05	4	0.001
LWIR	0.00747	0.1	23	1	3	$\sqrt{46}$	0.8	475	0.1	2	25
green target models											
blue	0.00068	0.038	41	3	5	$\sqrt{46}$	1.0	125	0.065	4	0.1
green	0.00202	0.038	14	1	5	$\sqrt{46}$	0.8	100	0.075	4	0.25
red	0.00163	0.1	28	1	2	46	1.0	150	0.06	4	0.001
EIR	0.00538	0.041	23	2	16	$\sqrt{46}$	0.6	75	0.08	2	0.001
NIR	0.00748	0.014	10	13	19	46	0.5	75	0.095	2	0.001
LWIR	0.00835	0.089	10	4	11	$\sqrt{46}$	0.8	300	0.1	2	27.5
gray target models											
blue	0.01	0.093	32	5	8	46	0.9	200	0.1	2	1.0
green	0.00689	0.03	23	4	6	$\sqrt{46}$	1.0	150	0.05	4	0.5
red	0.00996	0.018	19	11	2	46	0.5	100	0.095	2	50
EIR	0.00989	0.086	23	2	2	46	0.6	100	0.06	2	5.0
NIR	0.00985	0.028	19	9	11	$\sqrt{46}$	0.9	125	0.045	2	0.5
LWIR	0.00428	0.021	10	14	7	46	0.5	375	0.1	2	27.5
yellow target models											
blue	0.00144	0.08	14	5	9	$\sqrt{46}$	1.0	100	0.1	6	0.1
green	0.00705	0.081	14	4	14	$\sqrt{46}$	0.5	225	0.075	8	0.25
red	0.00847	0.096	41	3	2	$\sqrt{46}$	0.8	75	0.095	6	0.1
EIR	0.00299	0.006	10	1	7	46	1.0	125	0.06	4	0.1
NIR	0.00138	0.001	28	2	2	$\sqrt{46}$	0.6	125	0.095	2	50
LWIR	0.00668	0.1	37	4	12	46	1.0	200	0.09	12	0.1

3.2. Evaluation

With the models were already trained on 80% randomly selected data, their evaluation was performed on the remaining 20%. For each capture of the test data, the models had to predict the TVI for their associated sensor band. Afterwards, the bands were sorted from the band with the highest TVI prediction to the band with the lowest TVI prediction. The result of this sorting is called the predicted most informative band order. Because the targets in the test data were known, the actual TVI for each sensor band could be calculated, as was done for the training data during the training procedure. From the calculated TVIs for each sensor band, the bands were then sorted from the band with the highest calculated TVI to the band with the lowest calculated TVI. The result of this sorting is called the actual most informative band order. With the actual and predicted most informative band for each capture in the dataset, the band orders could then be compared for accuracy. For

example, the Top-1 accuracy is the proportion of captures in the test data where the first band of the predicted most informative band order is the same as the first band of the actual most informative band order. The same principle applies to the Top-3 accuracy, which is the proportion of captures in the test data in which the first band of the predicted most informative band order is one of the first three bands of the actual most informative band order.

Tp compare the predicted most informative band orders generated by the performance models using a static approach, a static baseline was introduced. The static baseline provides only the single most informative band order over all captures (the static most informative band order). It is motivated by the idea that it is not worth utilizing performance models if a simple static most informative band order already performs better than the predicted most informative band orders over all captures in the test data. The static most informative band order was obtained by penalizing each band using its position in the actual most informative band order over all captures in the training data. For example, as there are six bands, the most informative band is penalized by one and the least informative band is penalized by six. By accumulating the penalties of each band over all captures, the bands can be sorted from the band with the lowest accumulated penalty to the band with the highest accumulated penalty. This sorting results in the static most informative band order. Because the models were trained on four different sets of training data (one with all target groups, one with only green targets, one with only gray targets, and one with only yellow targets), there are four separate static most informative band orders:

- LWIR, NIR, EIR, red, blue, and green for any camouflaged target.
- NIR, LWIR, EIR, green, blue, and red for green camouflaged targets.
- NIR, blue, EIR, red, green, and LWIR for gray camouflaged targets.
- Red, blue, LWIR, green, EIR, and NIR for yellow camouflaged targets.

Table 5 shows the prediction accuracies of each model as a percentage. The individual quarters of the table contain the results of the general models trained on any camouflaged target and the specialized models trained only on green, gray, and yellow camouflaged targets, respectively. Each individual quarter provides four tables showing the prediction accuracies of the different machine learning models along with the prediction accuracies of the static baselines. Here, each cell contains the proportion of captures in the test data where the `<row number>` predicted most informative bands were among the `<column number>` actual most informative bands. For example, the value of the second column in the first row is the proportion of captures in which the first band of the predicted most informative band order was among the first two bands of the actual most informative band order (Top-2 accuracy). Likewise, the value of the third column and the second row represents the proportion of captures in which the first two bands of the predicted most informative band order were among the first three bands of the actual most informative band order. Therefore, the value of the first column and first row stands for the proportion in which the predicted most informative band was the actual most informative band (Top-1 accuracy). Although there are six different bands, the individual tables consist of only five columns and rows. This is for better clarity, as the last column would contain ones for each model and for the baseline. In all of the results for the different training procedures, the best predictions for each accuracy category are shown in bold.

As can be seen for the general models, the ϵ-SVR and Random Forest models are superior to the XGBoost models in terms of prediction accuracy. The ϵ-SVR models provide slightly higher prediction accuracy than the Random Forest models, with eight top results compared to six top results. While the XGBoost model achieves only two top results, the static baseline is inferior to all machine learning models, without a single top result. In addition, all models reach over 50% Top-1 accuracy and over 80% Top-3 accuracy. This means that the predicted most informative band is the actual most informative band more than half of the time, while most of the time the predicted most informative band is at least one of the three actual most informative bands. Similar results are shown by the models specializing in targets belonging to the green target group. However, the ϵ-SVR

model significantly outperforms all other models, achieving ten of the top fifteen results. In addition, the prediction accuracies are generally slightly higher than for the general case. The same applies to the results of the specialized gray target models, where only the number of top results is evenly split between the ϵ-SVR and the Random Forest models. In contrast, the specialized yellow target models are inferior to the static baseline in terms of the number of best results. Nonetheless, the first row, which represents the prediction accuracy of the single most informative band, is dominated by the top results of the ϵ-SVR model. Overall, the ϵ-SVR models perform the best, with the Random Forest and XGBoost models performing only slightly worse.

Table 5. The prediction accuracies of all models (in percentages). The upper left quarter contains the results of the general models, while the others contain the results of the specialized models. Each individual table shows the prediction accuracy of the respective model, where first row of the table corresponds to the Top-1 accuracy in the first column and the Top-5 accuracy in last column. In the second row, the value in the second column represents the accuracy of predicting the two most informative bands, regardless of their order. Similarly, the third column represents the accuracy of predicting two bands out of the three actual most informative bands. The same pattern applies to all other cells as well. The best results within each target group are shown in bold.

Any Target Models										**Green Target Models**									
ϵ-SVR					Random Forest					ϵ-SVR					Random Forest				
56.1	**71.9**	83.6	88.9	95.3	51.5	67.8	**84.2**	88.3	97.1	57.3	**74.5**	86.0	**93.0**	**98.7**	58.0	72.6	84.7	91.7	97.5
	35.7	59.1	**71.3**	**88.9**		33.3	**59.6**	70.8	87.7		**42.7**	**73.2**	**84.7**	**93.0**		36.9	67.5	79.6	91.1
		37.4	**53.8**	72.5			35.7	53.2	**75.4**			47.8	68.2	**86.0**			49.7	**69.4**	84.7
			23.4	55.0				22.2	**59.6**				**47.1**	**71.3**				38.9	68.8
				34.5					**35.7**					38.9					**42.0**
XGBoost					Baseline					XGBoost					Baseline				
50.9	68.4	83.6	**92.4**	**97.7**	47.4	56.1	69.6	75.4	82.5	**58.6**	70.1	**86.6**	91.7	96.8	31.2	68.2	81.5	89.2	96.8
	29.8	57.9	70.8	87.1		24.6	48.0	57.9	71.9		35.7	68.8	82.8	90.4		28.7	60.5	73.2	83.4
		31.0	46.8	68.4			33.9	46.2	60.2			**51.0**	68.8	83.4			47.8	65.0	80.3
			18.1	53.8				19.9	48.5				44.6	65.6				35.0	54.1
				34.5					28.7					39.5					25.5
Gray Target Models										**Yellow Target Models**									
ϵ-SVR					Random Forest					ϵ-SVR					Random Forest				
57.4	**77.0**	**86.9**	**93.4**	**98.4**	**59.0**	**77.0**	**86.9**	91.8	**98.4**	**49.5**	**75.7**	91.3	**95.1**	99.0	47.6	74.8	85.4	94.2	**100.0**
	34.4	59.0	**82.0**	**93.4**		**44.3**	**72.1**	77.0	90.2		35.9	66.0	82.5	93.2		**37.9**	66.0	83.5	93.2
		24.6	57.4	72.1			29.5	52.5	78.7			39.8	**74.8**	90.3			34.0	72.8	90.3
			34.4	62.3				24.6	**63.9**				51.5	**84.5**				54.4	80.6
				29.5					34.4					74.8					72.8
XGBoost					Baseline					XGBoost					Baseline				
50.8	**77.0**	85.2	88.5	96.7	24.6	63.9	83.6	91.8	98.4	43.7	69.9	83.5	91.3	99.0	31.1	71.8	**92.2**	94.2	96.1
	36.1	65.6	77.0	88.5		24.6	44.3	70.5	88.5		36.9	69.9	84.5	**96.1**		31.1	**75.7**	**89.3**	95.1
		34.4	**60.7**	**82.0**			18.0	41.0	59.0			**45.6**	71.8	90.3			38.8	65.0	**92.2**
			26.2	62.3				24.6	47.5				**57.3**	**84.5**				53.4	**84.5**
				39.3					34.4					74.8					**75.7**

To highlight the effectiveness of the performance models, Table 6 shows the relative accuracy improvements of each model compared to the static baseline, for which the prediction results are shown in the lower right of each quarter in Table 5. Apart from this, Table 6 has the same structure as Table 5. As can be seen, the general models achieve significant improvements over the static baseline, peaking at 45.2%, 35.7%, and 22.5% for the ϵ-SVR, Random Forest, and XGBoost models, respectively. Although the XGBoost model is slightly worse in one of the accuracy categories, it generally achieves much higher accuracy than the static baseline. Nonetheless, its improvements are not as great as those of

the ϵ-SVR and Random Forest models. The green target models achieve similar, though slightly lower overall improvements over the static baseline, with a notable very high increase in Top-1 accuracy. Likewise, the gray target models significantly outperform the static baseline, especially in Top-1 accuracy, where the prediction accuracy is more than doubled with improvements of up to 140%. Despite strong improvements in Top-1 accuracy, the yellow target models fail to improve in many of the accuracy categories. However, as noted above, the single most informative band is predicted the best by the ϵ-SVR model. In general, when considering both general and specialized models, all of the models are superior to the static baseline.

Table 7 shows the benefits of the performance models trained only on camouflaged targets belonging to a specific target group. Again, the structure of the table is the same as Table 5. The cells show the relative improvements in each accuracy category of the specialized models compared to the general models for only that target group on which the models were specialized. For example, the improvement of the green target models was obtained by comparing their prediction accuracy with that of the general models, while, the prediction accuracy of the general models was obtained by considering only green targets rather than of all targets. The same approach was applied to obtain the improvements of the gray and yellow target models. In this way, the benefits of the specialized models were quantified in an objective manner. As can be seen for the green target models, specialization leads to an overall improvement in prediction accuracy. The ϵ-SVR model achieves the most significant improvements, peaking at nearly three times the accuracy, a 191.9% improvement. The improvements are even greater for the gray target models, with the Random Forest model achieving nine top improvements and a 270% increase in Top-1 accuracy. Similar improvements can be observed for the yellow target models, with the XGBoost model showing the greatest improvements (a maximum increase in accuracy of 767.4%). Overall, the specialized models clearly outperform the general models within their respective target groups.

Table 6. The relative prediction accuracy improvements in each category of all models compared to their respective statically computed baselines (in percentages). The best results are shown in bold, even if all models predicted worse than the static baseline. This table follows the same structure as Table 5.

Any Target Models										**Green Target Models**									
ϵ-SVR					Random Forest					ϵ-SVR					Random Forest				
18.5	**28.1**	20.2	17.8	15.6	8.6	20.8	**21.0**	17.1	17.7	83.7	**9.3**	5.5	**4.3**	**2.0**	85.7	6.5	3.9	2.9	0.7
	45.2	23.2	**23.2**	**23.6**		35.7	**24.4**	22.2	22.0		**48.9**	**21.1**	**15.7**	**11.5**		28.9	11.6	8.7	9.2
		10.3	**16.5**	20.4			5.2	15.2	**25.2**			0.0	4.9	**7.1**			4.0	**6.9**	5.6
			17.6	13.3				11.8	**22.9**				**34.5**	**31.8**				10.9	27.1
				20.4					**24.5**					52.5					**65.0**
XGBoost										XGBoost									
7.4	21.9	20.2	**22.5**	**18.4**						**87.8**	2.8	**6.3**	2.9	0.0					
	21.4	20.7	22.2	21.1							24.4	13.7	13.0	8.4					
		−8.6	1.3	13.6								**6.7**	5.9	4.0					
			−8.8	10.8									27.3	21.2					
				20.4										55.0					

Table 6. *Cont.*

Gray Target Models										Yellow Target Models									
ϵ-SVR					Random Forest					ϵ-SVR					Random Forest				
133.3	**20.5**	**3.9**	**1.8**	**0.0**	**140.0**	**20.5**	**3.9**	0.0	**0.0**	**59.4**	**5.4**	**−1.1**	**1.0**	3.0	53.1	4.1	−7.4	0.0	**4.0**
	40.0	33.3	**16.3**	**5.6**		**80.0**	**63.0**	9.3	1.9		15.6	−12.8	−7.6	−2.0		**21.9**	−12.8	−6.5	−2.0
		36.4	40.0	22.2			63.6	28.0	33.3			2.5	**14.9**	**−2.1**			−12.5	11.9	**−2.1**
			40.0	31.0				0.0	**34.5**				−3.6	**0.0**				1.8	−4.6
				−14.3					0.0					**−1.3**					−3.8
XGBoost										XGBoost									
106.7	**20.5**	2.0	−3.6	−1.7						40.6	−2.7	−9.5	−3.1	3.0					
	46.7	48.1	9.3	0.0							18.8	**−7.7**	**−5.4**	**1.0**					
		90.9	**48.0**	**38.9**								**17.5**	10.4	**−2.1**					
			6.7	31.0									**7.3**	**0.0**					
				14.3										**−1.3**					

Table 7. The relative prediction accuracy improvements of the specialized models compared to the general models. Note that the improvements are based on the prediction accuracy of the general models that results when considering only those targets considered for the prediction accuracy of the respective specialized model, and not on the prediction accuracy of the general models shown in Table 5. This table follows the same structure as Table 5.

Green Target Models									
ϵ-SVR					Random Forest				
−0.8	**1.7**	1.8	3.3	**3.9**	4.9	−0.9	1.0	**4.0**	2.5
	22.7	**35.6**	**20.7**	9.3		8.1	23.5	14.5	**9.4**
		47.9	**35.5**	**28.2**			50.9	34.7	20.7
			191.9	**57.3**				172.0	38.4
				69.1					50.4
XGBoost									
7.2	−3.6	**2.5**	2.5	3.2					
	4.4	21.7	14.9	7.9					
		54.8	33.4	16.8					
			111.1	32.0					
				44.5					

Gray Target Models										Yellow Target Models									
ϵ-SVR					Random Forest					ϵ-SVR					Random Forest				
182.8	77.2	53.7	40.2	**19.1**	**270.2**	**112.7**	**71.3**	**54.5**	**19.1**	38.1	63.8	**72.7**	**36.3**	**19.3**	**48.3**	**65.1**	48.5	21.7	10.4
	13.1	40.4	**66.3**	21.7		**45.4**	**55.5**	47.7	**24.4**		246.2	133.3	98.8	**62.0**		301.4	191.6	96.7	47.5
		−0.2	80.0	42.2			45.4	72.4	**59.7**			122.1	147.6	108.1			260.2	175.7	117.5
			375.1	**79.1**				112.1	76.5				319.6	198.4				343.3	205.1
				85.1					137.5					366.1					328.8
XGBoost										XGBoost									
150.5	96.9	59.0	35.7	15.1						25.2	57.7	50.0	19.4	16.6					
	38.3	46.0	51.9	17.5							**334.5**	**208.7**	**108.2**	61.7					
		58.4	**109.3**	52.9								**706.1**	**231.1**	**139.3**					
			158.5	71.9									**767.4**	**258.1**					
				146.8										**560.4**					

4. Discussion

This section first discusses the possible implications of the proposed sensor prediction approach in Section 4.1, followed by a discussion of its limitations in Section 4.2. Finally, future research prospects are reviewed in Section 4.3.

4.1. Implications

Altogether, our results demonstrate the effectiveness of the sensor performance prediction approach presented in this paper, with the ϵ-SVR models showing the most robust performance. While not perfect, the performance models were able to learn a meaningful relationship between the context state and the corresponding TVI, which supports the utility of the extracted features and the expressiveness of the TVI. In the general case, when only the three predicted most informative bands were considered, the actual most informative band was most likely among them (around 84%), while the associated workload was reduced by a half compared to processing all six bands. Although a reduction from six to three sensor bands may seem small in absolute terms, the sensor performance prediction approach is adaptable to any number of bands. For the sake of simplicity and clarity, however, our evaluation of the proposed performance prediction approach focuses on the raw bands of the multispectral sensor system employed in this study. While not explicitly explored here, the nature of the proposed methodology suggests similar results for a smaller or larger number of bands. Therefore, the proposed method could significantly increase the utility of multispectral sensor systems in real-world applications. For example, reconnaissance drones could be equipped with much more powerful multispectral sensor systems, as the increased number of sensor bands would not result in an equally increased workload. In this case, the sensor performance prediction approach would determine the most informative bands and ignore the least informative bands. The resulting increased meaningfulness of each sensor band and the additional spectral information due to the larger number of bands could greatly improve camouflage detection performance in reconnaissance scenarios.

In addition, our evaluation shows that specializing the performance models for certain target groups can significantly increase prediction performance. This could potentially increase reconnaissance performance for scenarios in which camouflaged targets are known to be present in a specific kind of environment, as the specialized models are able to focus on the environment associated with their specific target group even when the reconnaissance area consists of different kinds of environments. In contrast, the general models consider all relevant environments even when camouflaged targets are known to be present in only one environment. Thus, the predictions of the specialized models are more closely tailored to the environment in which the camouflaged targets are located, resulting in greater camouflage detection performance. Naturally, the specialized models cannot generalize to environments that are not associated with their specific target group. For this reason, they can only be of use if this specific kind of prior knowledge is available.

Comparing the performance models to a static baseline further highlights the benefits of their application. Although the use of static most informative band orders is computationally less expensive than the use of performance models, the former approach is not able to achieve the same level of prediction accuracy. Therefore, the comparatively low computational overhead of the performance models is preferable to the lower performance of the static baselines. However, it should be noted that the yellow target models did not significantly outperform their associated static baselines. This could be due to the relatively small amount of training data, as yellow targets were not as common as green targets in our dataset. A lack of training data may have prevented the performance models from sufficiently learning the complex relationship between the context state and the TVI, resulting in more limited generalization capabilities. On the other hand, even though the dominance of gray targets in the dataset was even lower than that of yellow targets, the performance models for gray targets were far superior to the static baseline. This could be due to the relationship between the context state and the TVI being less complex for the gray target models than for the yellow target models. Unfortunately, the causes of the relatively poor performance of the yellow target models could not be further explored in this study.

Because the idea behind the performance prediction approach is not strictly bound to the camouflage detection task, it can be generalized and applied to other multispectral

sensing problems. In the present work, sensor performance corresponds to the TVI; however, this particular metric could be replaced with any other metric that fits the problem at hand. For example, such a metric could describe the ability of a sensor band to detect invasive species. In this case, the performance models simply had to learn the relationship between the context state and the new metric instead of the TVI. Even the context state is not specifically tailored to camouflage detection, as it is generated by general image descriptors. Therefore, it could be of equal utility in other use cases. Considering the positive results we obtained when applying the proposed concept to multispectral camouflage detection, it could be equally successful when applied to other tasks.

4.2. Limitations

Although the value of the proposed performance prediction approach has been confirmed, it should be noted that all of our results are based on the Target Visibility Index metric introduced in this paper. As has already been discussed in Section 2.2, the TVI defines visibility using its mathematical formula, which does not necessarily correspond to human perception. Therefore, certain targets that may actually have poor visibility to the human eye can result in a relatively high TVI, and vice versa. This behavior may have led to predictions of the performance models that were correct with respect to the TVI and incorrect with respect to the human eye. As a result, the benefits of the proposed performance prediction approach may be limited in a real-world application involving humans.

Furthermore, because each camouflaged target possesses unique spectral characteristics, the mappings from the context state to the TVI may have been diluted in the data. This may have limited the achieved prediction accuracy of the performance models. For example, while a target in one environment will result in a different TVI than another target in the same environment, the context state will not have changed in either case, as the context state is mainly determined by the scenery and not by the targets. This leads to mappings from one context state being applied to different TVIs for the same sensor band, which could have confused the training of the performance models. The averaging of all TVIs in the same capture could have further amplified this potential issue, as already mentioned in Section 3.1. However, the environment, and consequently the context state, already provide an indication of the properties of the camouflaged targets, as green targets, for example, are usually found in green environments. Therefore, the TVI may follow a certain distribution for a given environment, which could have limited the potential negative effects on the training of the performance models.

In addition, it is important to note that the performance models were not evaluated for their ability to generalize to unknown camouflaged targets. Although the data were split into training and test data, all of the camouflaged targets were part of both datasets. However, the test data contained captures that were completely unknown to the performance models, on which they showed high prediction accuracy. This suggests that the proposed sensor prediction approach has great potential for generalization.

Another limiting factor on the prediction accuracy could have been the meaningfulness of the context state extracted from the visual band. Because the context state results from relatively simple feature extractors, the performance models may not have learned every aspect of the complex relationship between the environmental situation and the TVI. More sophisticated and computationally expensive feature extraction methods, such as convolutional neural networks, might have provided even more meaningful context states. With more information about the environmental situation available in the context state, the performance models may have achieved even higher prediction accuracy. However, the computational resources on a small reconnaissance drone in a real-world application are usually limited. This requires computationally inexpensive methods for both the feature extraction process and the performance models, which have been successfully implemented and demonstrated in this paper.

4.3. Future Research

Our future research will primarily focus on predicting sensor performance for multispectral sensor systems with an even higher number of bands. In addition, the sensor performance prediction approach proposed in this paper will be included in a larger framework in which the most informative bands will be incorporated into a computer-aided camouflage detection system. As noted above, richer features in the context state may improve the prediction accuracy of the performance models, which will be another subject of future research.

5. Conclusions

The sensor performance prediction approach presented in this paper has been shown to be a successful method for obtaining those sensor bands that best expose camouflaged targets. This increases the meaningfulness of each individual sensor band, allowing for the use of more powerful multispectral sensor systems. As a result, camouflage detection performance may be significantly increased in real-world reconnaissance scenarios.

In addition, specialized training of the performance models showed promising improvements in prediction accuracy. This may further increase camouflage detection performance in real-world reconnaissance scenarios, provided that the necessary prior knowledge of the camouflaged targets to be exposed is available.

Moreover, it has been shown that the performance models are superior to the statically computed most informative band order. This indicates the existence of a complex relationship between the environmental situation and the TVI that can be successfully exploited and learned by performance models. Therefore, the benefits of the proposed performance prediction approach outweigh its computational overhead compared to a static baseline and motivate its application in real-world reconnaissance scenarios.

However, it should be noted that all results are based on the TVI, which is an experimental metric of sensor performance in the context of camouflaged target detection. Because the TVI does not necessarily correspond to human perception and is difficult to apply to multiple targets in the same scene, the range of applications of the proposed sensor performance prediction approach may be limited. In addition, the context state may not be as informative as it might have been with more sophisticated feature extraction methods, which in turn may have limited the prediction accuracy of the performance models.

Future research will address the integration of the proposed sensor prediction approach into an automated camouflaged target detection system and the generation of a richer context state.

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

Funding: This research was funded by the Federal Office of Bundeswehr Equipment, Information Technology, and In-Service Support (BAAINBw). The APC was funded by the University of the Bundeswehr Munich (UniBwM).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X) dataset is publicly available on GitHub: https://github.com/Tobias-UniBwM/MUDCAD-X (accessed on 11 August 2023).

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

Abbreviations

The following abbreviations are used in this manuscript:

EIR	Edge Infra-red/Red-edge
GBT	Gradient Boosted Tree
LBP	Local Binary Pattern
LWIR	Long-Wave Infra-Red
MUDCAD-X	eXtendend Multispectral Dataset for Camouflage Detection
NDRE	Normalized Difference Red-Edge index
NDVI	Normalized Difference Vegetation Index
NIR	Near Infra-Red
RF	Random Forest
RMSE	Root Mean Square Error
TVI	Target Visibility Index
UAV	Unmanned Aerial Vehicle
UniBwM	University of the Bundeswehr Munich
VIS	Visual
XGBoost	eXtreme Gradient Boosting

References

1. Ampatzidis, Y.; Partel, V. UAV-Based High Throughput Phenotyping in Citrus Utilizing Multispectral Imaging and Artificial Intelligence. *Remote Sens.* **2019**, *11*, 410. [CrossRef]
2. Khaliq, A.; Comba, L.; Biglia, A.; Ricauda Aimonino, D.; Chiaberge, M.; Gay, P. Comparison of Satellite and UAV-Based Multispectral Imagery for Vineyard Variability Assessment. *Remote Sens.* **2019**, *11*, 436. [CrossRef]
3. Osco, L.P.; Arruda, M.d.S.d.; Marcato Junior, J.; da Silva, N.B.; Ramos, A.P.M.; Moryia, É.A.S.; Imai, N.N.; Pereira, D.R.; Creste, J.E.; Matsubara, E.T.; et al. A convolutional neural network approach for counting and geolocating citrus-trees in UAV multispectral imagery. *ISPRS J. Photogramm. Remote Sens.* **2020**, *160*, 97–106. [CrossRef]
4. Vali, A.; Comai, S.; Matteucci, M. Deep Learning for Land Use and Land Cover Classification Based on Hyperspectral and Multispectral Earth Observation Data: A Review. *Remote Sens.* **2020**, *12*, 2495. [CrossRef]
5. Gao, Y.; Li, W.; Zhang, M.; Wang, J.; Sun, W.; Tao, R.; Du, Q. Hyperspectral and Multispectral Classification for Coastal Wetland Using Depthwise Feature Interaction Network. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 1–15. [CrossRef]
6. Osorio, K.; Puerto, A.; Pedraza, C.; Jamaica, D.; Rodríguez, L. A Deep Learning Approach for Weed Detection in Lettuce Crops Using Multispectral Images. *AgriEngineering* **2020**, *2*, 471–488. [CrossRef]
7. Yu, R.; Luo, Y.; Zhou, Q.; Zhang, X.; Wu, D.; Ren, L. Early detection of pine wilt disease using deep learning algorithms and UAV-based multispectral imagery. *For. Ecol. Manag.* **2021**, *497*, 119493. [CrossRef]
8. Lan, Y.; Huang, Z.; Deng, X.; Zhu, Z.; Huang, H.; Zheng, Z.; Lian, B.; Zeng, G.; Tong, Z. Comparison of machine learning methods for citrus greening detection on UAV multispectral images. *Comput. Electron. Agric.* **2020**, *171*, 105234. [CrossRef]
9. Kerkech, M.; Hafiane, A.; Canals, R. Vine disease detection in UAV multispectral images using optimized image registration and deep learning segmentation approach. *Comput. Electron. Agric.* **2020**, *174*, 105446. [CrossRef]
10. McAllister, E.; Payo, A.; Novellino, A.; Dolphin, T.; Medina-Lopez, E. Multispectral satellite imagery and machine learning for the extraction of shoreline indicators. *Coast. Eng.* **2022**, *174*, 104102. [CrossRef]
11. Yuan, K.; Zhuang, X.; Schaefer, G.; Feng, J.; Guan, L.; Fang, H. Deep-Learning-Based Multispectral Satellite Image Segmentation for Water Body Detection. *IEEE J. Sel. Top. Appl. Earth Observ. Remote Sens.* **2021**, *14*, 7422–7434. [CrossRef]
12. Geyman, E.C.; Maloof, A.C. A Simple Method for Extracting Water Depth From Multispectral Satellite Imagery in Regions of Variable Bottom Type. *Earth Space Sci.* **2019**, *6*, 527–537. [CrossRef]
13. Shin, J.i.; Seo, W.w.; Kim, T.; Park, J.; Woo, C.s. Using UAV Multispectral Images for Classification of Forest Burn Severity—A Case Study of the 2019 Gangneung Forest Fire. *Forests* **2019**, *10*, 1025. [CrossRef]
14. Hupel, T.; Stütz, P. Adopting Hyperspectral Anomaly Detection for Near Real-Time Camouflage Detection in Multispectral Imagery. *Remote Sens.* **2022**, *14*, 3755. [CrossRef]
15. Zwick, M.; Gerdts, M.; Stütz, P. Sensor-Model-Based Trajectory Optimization for UAVs to Enhance Detection Performance: An Optimal Control Approach and Experimental Results. *Sensors* **2023**, *23*, 664. [CrossRef] [PubMed]
16. Koch, S.; Krach, B.; Katsilieris, F.; Stütz, P. Sensor Scheduling Strategies for 1-to-N Multi-Object Tracking. In Proceedings of the 2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC), Portsmouth, VA, USA, 18–22 September 2022; pp. 1–9. [CrossRef]
17. Ruß, M.; Stütz, P. Airborne Sensor and Perception Management. In *Proceedings of the Modelling and Simulation for Autonomous Systems*; Mazal, J., Fagiolini, A., Vašík, P., Bruzzone, A., Pickl, S., Neumann, V., Stodola, P., Lo Storto, S., Eds.; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2023; pp. 182–206. [CrossRef]

18. Hellert, C.; Koch, S.; Stütz, P. Using Algorithm Selection for Adaptive Vehicle Perception Aboard UAV. In Proceedings of the 2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), Taipei, Taiwan, 18–21 September 2019; pp. 1–8. [CrossRef]
19. Adrian, W. Visibility of targets: Model for calculation. *Light. Res. Technol.* **1989**, *21*, 181–188. [CrossRef]
20. Rea, M.S.; Ouellette, M.J. Relative visual performance: A basis for application. *Light. Res. Technol.* **1991**, *23*, 135–144. [CrossRef]
21. Brémond, R.; Dumont, E.; Ledoux, V.; Mayeur, A. Photometric measurements for visibility level computations. *Light. Res. Technol.* **2011**, *43*, 119–128. [CrossRef]
22. Brémond, R.; Bodard, V.; Dumont, E.; Nouailles-Mayeur, A. Target visibility level and detection distance on a driving simulator. *Light. Res. Technol.* **2013**, *45*, 76–89. [CrossRef]
23. Chen, Z.; Tu, Y.; Wang, Z.; Liu, L.; Wang, L.; Lou, D.; Zhu, X.; Teunissen, K. Target visibility under mesopic vision using a driving simulator. *Light. Res. Technol.* **2019**, *51*, 883–899. [CrossRef]
24. Lebouc, L.; Boucher, V.; Greffier, F.; Liandrat, S.; Nicolaï, A.; Richard, P. Influence of road surfaces on the calculation of a target visibility taking into account the direct and indirect lighting. In Proceedings of the CIE 2021 Conference, Online, 27–29 September 2021. [CrossRef]
25. Kwon, T.M. *Atmospheric Visibility Measurements Using Video Cameras: Relative Visibility*; University of Minnesota Duluth: Duluth, MN, USA, 2004.
26. Yin, X.C.; He, T.T.; Hao, H.W.; Xu, X.; Cao, X.Z.; Li, Q. Learning Based Visibility Measuring with Images. In *Proceedings of the Neural Information Processing*; Lu, B.L., Zhang, L., Kwok, J., Eds.; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2011; pp. 711–718. [CrossRef]
27. Hou, X.; Zhang, L. Saliency Detection: A Spectral Residual Approach. In Proceedings of the 2007 IEEE Conference on Computer Vision and Pattern Recognition, Minneapolis, MN, USA, 17–22 June 2007; IEEE: Piscataway, NJ, USA, 2007; pp. 1–8. [CrossRef]
28. Montabone, S.; Soto, A. Human detection using a mobile platform and novel features derived from a visual saliency mechanism. *Image Vis. Comput.* **2010**, *28*, 391–402. [CrossRef]
29. Veksler, O. Test Time Adaptation with Regularized Loss for Weakly Supervised Salient Object Detection. In Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, BC, Canada, 18–22 June 2023; pp. 7360–7369. [CrossRef]
30. Peli, E. Contrast in complex images. *J. Opt. Soc. Am. A* **1990**, *7*, 2032–2040. [CrossRef]
31. Wang, Z.; Bovik, A.; Sheikh, H.; Simoncelli, E. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Process.* **2004**, *13*, 600–612. [CrossRef] [PubMed]
32. Wang, Z.; Bovik, A. A universal image quality index. *IEEE Signal Process. Lett.* **2002**, *9*, 81–84. [CrossRef]
33. Authors, O. ODM—A Command Line Toolkit to Generate Maps, Point Clouds, 3D Models and DEMs from Drone, Balloon or Kite Images. 2020. Available online: https://github.com/OpenDroneMap/ODM (accessed on 3 March 2023).
34. Evangelidis, G.D.; Psarakis, E.Z. Parametric Image Alignment Using Enhanced Correlation Coefficient Maximization. *IEEE Trans. Pattern Anal. Mach. Intell.* **2008**, *30*, 1858–1865. [CrossRef]
35. Bradski, G. The OpenCV Library. *Dr. Dobb's J. Softw. Tools* **2000**. Available online: http://www.drdobbs.com/open-source/the-opencv-library/184404319 (accessed on 3 March 2023).
36. CVAT.ai Corporation. Computer Vision Annotation Tool (CVAT). 2022. Available online: https://github.com/opencv/cvat (accessed on 3 March 2023).
37. Shiffler, R.E.; Harsha, P.D. Upper and Lower Bounds for the Sample Standard Deviation. *Teach. Stat.* **1980**, *2*, 84–86. [CrossRef]
38. Ojala, T.; Pietikäinen, M.; Mäenpää, T. Gray Scale and Rotation Invariant Texture Classification with Local Binary Patterns. In Proceedings of the Computer Vision-ECCV 2000, Dublin, Ireland, 26 June–1 July 2000; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2000; pp. 404–420. [CrossRef]
39. Haralick, R.M.; Dinstein, I.; Shanmugam, K. Textural Features for Image Classification. *IEEE Trans. Syst. Man Cybern.* **1973**, *SMC-3*, 610–621. [CrossRef]
40. Spanhol, F.A.; Oliveira, L.S.; Petitjean, C.; Heutte, L. A Dataset for Breast Cancer Histopathological Image Classification. *IEEE Trans. Biomed. Eng.* **2016**, *63*, 1455–1462. [CrossRef] [PubMed]
41. Ahmad, M.; Shabbir, S.; Roy, S.K.; Hong, D.; Wu, X.; Yao, J.; Khan, A.M.; Mazzara, M.; Distefano, S.; Chanussot, J. Hyperspectral Image Classification-Traditional to Deep Models: A Survey for Future Prospects. *IEEE J. Sel. Top. Appl. Earth Observ. Remote Sens.* **2022**, *15*, 968–999. [CrossRef]
42. Lu, D.; Weng, Q. A survey of image classification methods and techniques for improving classification performance. *Int. J. Remote Sens.* **2007**, *28*, 823–870. [CrossRef]
43. Bezanson, J.; Edelman, A.; Karpinski, S.; Shah, V.B. Julia: A fresh approach to numerical computing. *SIAM Rev.* **2017**, *59*, 65–98. [CrossRef]
44. Coelho, L.P. Mahotas: Open source software for scriptable computer vision. *J. Open Res. Softw.* **2013**, *1*, e3. [CrossRef]
45. Blaom, A.D.; Kiraly, F.; Lienart, T.; Simillides, Y.; Arenas, D.; Vollmer, S.J. MLJ: A Julia package for composable machine learning. *J. Open Source Softw.* **2020**, *5*, 2704. [CrossRef]
46. Chang, C.C.; Lin, C.J. LIBSVM: A library for support vector machines. *ACM Trans. Intell. Syst. Technol.* **2011**, *2*, 27:1–27:27. Available online: http://www.csie.ntu.edu.tw/~cjlin/libsvm (accessed on 3 March 2023). [CrossRef]

47. Sadeghi, B.; Chiarawongse, P.; Squire, K.; Jones, D.C.; Noack, A.; St-Jean, C.; Huijzer, R.; Schätzle, R.; Butterworth, I.; Peng, Y.F.; et al. DecisionTree.jl—A Julia Implementation of the CART Decision Tree and Random Forest Algorithms, Zenodo, 2022. [CrossRef]
48. Chen, T.; Guestrin, C. XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, 13–17 August 2016; KDD '16; pp. 785–794. [CrossRef]

5.3 Sensor-Managed Anomaly Detection for Camouflage Detection in Airborne Multispectral Imagery

Tobias Hupel and Peter Stütz. „Sensor-Managed Anomaly Detection for Camouflage Detection in Airborne Multispectral Imagery“. In: *2024 IEEE Aerospace Conference*. ISSN: 1095-323X. Mar. 2024, pp. 1–11. DOI: 10.1109/AERO58975.2024.10521231

Sensor-Managed Anomaly Detection for Camouflage Detection in Airborne Multispectral Imagery

Tobias Hupel
Institute of Flight Systems
University of the Bundeswehr Munich
85577 Neubiberg, Germany
tobias.hupel@unibw.de

Peter Stütz
Institute of Flight Systems
University of the Bundeswehr Munich
85577 Neubiberg, Germany
peter.stuetz@unibw.de

Abstract— **Multispectral sensors provide valuable spectral information suitable for a variety of remote sensing applications, especially when cost-effectiveness, portability, and energy efficiency are essential, such as on small tactical drones. Since all materials have unique spectral signatures, this spectral information can be utilized to detect camouflaged targets in military reconnaissance scenarios. However, varying environmental conditions like vegetation and weather in real-world reconnaissance scenarios, and the computational complexity introduced by the large number of bands, make successful camouflage detections difficult. To address these issues, a novel sensor-managed anomaly detection method is proposed. This approach determines a subset of sensor bands of a multispectral imaging system based on current environmental conditions and their ability to expose camouflaged targets. Instead of processing all available bands, a spectral anomaly detector then processes only this subset of bands. To evaluate the robustness of the proposed approach, four different spectral anomaly detection methods are employed and compared: the RX detector, LRX detector, AED detector, and LPD detector. Experiments on an extensive multispectral dataset for camouflage detection show nearly equivalent detection performance when spectral anomaly detectors are applied to the managed subset compared to all sensor bands. This demonstrates the effectiveness of the proposed approach in reducing computational complexity while maintaining detection performance, which is particularly important for small tactical drones in military reconnaissance scenarios.**

TABLE OF CONTENTS

1. INTRODUCTION 1
2. RELATED WORK 2
3. METHODOLOGY 2
4. EXPERIMENTS 5
5. RESULTS 6
6. CONCLUSIONS 8
REFERENCES 8
BIOGRAPHY 11

1. INTRODUCTION

Small drones such as MIKADO and ALADIN of the Bundeswehr are commonly deployed in close-range reconnaissance scenarios. Although their imaging systems typically cover only visual or thermal spectral ranges, technological advances already allow the use of sophisticated multispectral imaging systems. Compared to visual or thermal sensors, multispectral sensors provide a much larger number of bands and therefore much more valuable spectral information. The high utility and comparatively low cost, size, weight and power consumption of multispectral imaging systems has led to immense research interest in recent years. Applications range from vegetational monitoring [1, 2], precision agriculture [3–5] and land cover classification [6, 7] to biomass estimation [8, 9], bathymetry [10, 11], disaster evaluation [12] and urban heat vulnerability determination [13]. In addition, multispectral sensors proved to be sufficient for computer-aided detection of camouflaged targets [14]. The detection of camouflaged targets in complex military reconnaissance scenarios is an inherently challenging and time-consuming task. Therefore, computer-aided camouflage detection system are of great use by reducing the workload of the operating personnel while increasing the successful detection rates of camouflaged targets.

Although multispectral imagery is suitable for detecting camouflaged targets through spectral anomaly detection, the targets do not consistently appear anomalous in all sensor bands. For example, a sensor band may expose camouflaged targets in a certain environment, but it may not expose camouflaged targets in other environments. In addition, the high number of bands to be processed introduces an increased computational complexity. Therefore, this paper proposes a novel sensor-managed anomaly detection approach to detect camouflaged targets in multispectral imagery. In this approach, only the bands most likely to expose camouflaged targets are processed by a spectral anomaly detector. These sensor bands are determined by a performance prediction method [15] that evaluates the current environmental conditions and predicts the performance of each sensor band. In this context, performance refers to the extent to which the particular sensor band exposes camouflaged targets. The predicted performance measure dynamically ranks the sensor bands from those performing best to those performing worst. After ranking the sensor bands based on their current performance, the spectral anomaly detector processes only the best-performing subset of all bands, ignoring all others. This is intended to make camouflaged targets more susceptible to detection, on the assumption that more exposed targets will also appear more anomalous to an anomaly detector. At the same time, the smaller number of bands to process reduces the computational complexity associated with the spectral anomaly detector. The evaluation of the proposed approach is based on the eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X) and four different anomaly detectors: the Reed-Xiaoli detector (RX), the Local Reed-Xiaoli detector (LRX), the Local Point Density detector (LPD) and the Attribute and Edge-preserving Filter detector (AED). In addition, the evaluation primarily focuses on a comparison with an unmanaged approach, where all bands are processed by a spectral anomaly detector instead of the most performant ones.

979-8-3503-0462-6/24/$31.00 ©2024 IEEE

Overall, this paper makes the following scientific contributions:

- Proposition of a sensor-managed anomaly detection method for camouflage detection in multispectral imagery
- Evaluation of the proposed approach using multiple different anomaly detection methods

The following sections are organized as follows. Related work is discussed in Section 2, while the proposed approach is introduced in detail in Section 3. The obtained results are presented and discussed in Section 5. Final conclusion of the presented work are drawn in Section 6.

2. RELATED WORK

The management of sensors and algorithms using performance models has been investigated primarily in the context of object detection [16–18]. In [16], artificial neural networks and fuzzy inference are utilized to predict a performance index of perception chains (sensor combined with data processing algorithms) using an environmental state vector. Similarly, in [17], the performance of perception chains is predicted using Bayesian Networks. The performance prediction enables dynamic selection of the most suitable perception chain based on environmental conditions, resulting in improved overall object detection performance. In [18], the performance models are incorporated into a flight trajectory optimization approach using optimal control algorithms. This allows flight paths to be determined that maximize the utility of the available sensors and object detection algorithms. In addition to object detection, [19] applies performance modelling in the context of sensor scheduling and object tracking. Here, the performance models support the sensor scheduling algorithm by predicting the probability of successful object observations, resulting in increased object observation times. In [15], sensor performance models predict the visibility of camouflaged targets in all sensor bands of a multispectral camera system given an environmental context state. With this information, the most informative bands for camouflage detection are dynamically determined, while all other bands can be ignored. Under the assumption that high visibility of camouflaged targets benefits their detection by an anomaly detection algorithm, this approach is also used in this work.

Anomaly detection in hyperspectral and multispectral imagery is a very active area of research. The Reed-Xiaoli (RX) detector [20] is one of the most prominent approaches, followed by further developments such as the Kernel RX detector [21], the Weighted RX detector [22] and the Cluster Kernel RX detector [23]. While these methods are based on pixel representation by statistical distributions, other successful approaches rely on collaborative [24, 25] or sparse [26–28] pixel representation. In addition, methods using clustering [29], morphological filtering [30], isolation forests [31] or pixel density analysis [14, 32, 33] have also been proposed. Deep learning based detectors [34] mostly exploit the reconstruction error of autoencoders [35–37] for anomaly detection. Since the RX [20], LRX [20], AED [30] and LPD [14] detectors have already been used for camouflage detection in multispectral imagery [14], they are also employed in this work.

3. METHODOLOGY

This section first shortly reviews the dataset on which this work is based. Then, the general concept, the sensor performance prediction method, and the anomaly detectors are briefly introduced.

Table 1: **The bands and their associated properties provided by each capture of MUDCAD-X.**

band	center	width
visual (VIS)	-	-
blue	475nm	32nm
green	560nm	27nm
red	668nm	14nm
edge-infrared (EIR)	717nm	12nm
near-infrared (NIR)	842nm	57nm
long-wave infrared (LWIR)	10.5µm	6µm

Table 2: **The structure of the context state extracted from the gray-level converted visual band.**

LBP		Haralick	
uniform	non-uniform	mean	min-max
1 - 17	18	19 - 32	33 - 46

MUDCAD-X

The work presented in this paper relies on MUDCAD-X [15], an extensive and comprehensive multispectral dataset for camouflage detection. It features 853 multispectral captures, each containing at least one of the following camouflaged targets: artificial hedge, artificial turf, green tarp, green 2D camouflage net, green 3D camouflage net, person in green military uniform, anthracite fleece, gray tarp, gray 3D camouflage net, yellow 3D camouflage net and person in yellow military uniform. In addition, the footage was acquired at three different seasons: Spring, Summer and Fall. A sample capture along with the corresponding label mask is shown in Figure 1. As can be seen, the bands included in each capture are visual (VIS), blue, green, red, edge-infrared (EIR), near-infrared and long-wave infrared (LWIR). Their center wavelengths and bandwidths are shown in Table 1. More details about the dataset can be found in the original publication [15].

Sensor-Managed Anomaly Detection

The concept of sensor-managed anomaly detection for camouflage detection is illustrated in Figure 2. First, the context state is extracted by multiple image descriptors from the visual band. The performance models use the context state to predict the performance for each of the other bands. Next, the bands are sorted from the band with the highest predicted performance to the band with the lowest predicted performance. This is the predicted most informative band order. The first N bands of the predicted most informative band order and the visual band are concatenated to form the detection base, where N is a predefined constant. Finally, an anomaly detector processes the detection base and generates a detection map showing the most anomalous regions in the capture. In this case, these regions correspond to camouflaged target detections. The sensor performance prediction method and the anomaly detection process are described in more detail below.

(a) VIS (b) blue (c) green (d) red

(e) EIR (f) NIR (g) LWIR (h) label

Figure 1: **A sample capture of MUDCAD-X with bands from VIS (a) to LWIR (g) and the label mask (h).**

Sensor Performance Prediction

The entire sensor performance prediction method is adopted from [15] and is based on machine learning models, i.e. performance models, that learn the relationship between environmental context and sensor performance. Each performance model is linked to a single sensor band of a multispectral sensor or capture and uses an abstract representation of the environmental context to predict the performance of its associated sensor band. The predicted performances rank the sensor bands from the best performing band to the worst performing band. Note that the following explanations cover only the basic methodology of the performance prediction approach. For more insights, see the original publication [15].

Context State Extraction—In order to efficiently obtain information about the current environmental situation, local binary patterns (LBPs) [38] and Haralick's image statistics [39] are utilized to extract global features from the gray-level converted visual band, as shown in Figure 2. These features provide an abstract description of the visual band and thus an abstract representation of the environmental context. Therefore, the extracted features are called context state. Its structure is shown in Table 2.

The LBP features are generated using the LBP_{16}^{riu2} operator [38], which extracts uniform rotation-invariant patterns with a radius of 2px and size of 16 for each pixel. For the LBP feature values in the context state, the histogram over all patterns is computed. Since the LBP_{16}^{riu2} operator produces exactly 17 different uniform rotation-invariant patterns, the histogram consists of 18 values. The first 17 feature values are the number of occurrences of each individual uniform rotation-invariant pattern and the last value is the number of occurrences of all non-uniform patterns. Finally, the histogram is normalized so that all values sum up to one. The implementation for extracting the LBPs is provided by the ImageFeatures.jl package of the Julia Programming Language [40].

The remaining 28 feature values of the context state correspond to the Haralick image statistics [39]. Here, four gray-level co-occurrence matrices with a radius of 1px and the directions left, up, right and down are generated from the gray-level converted visual band. Each of the four co-occurrence matrices produces a separate set of Haralick features, each consisting of 14 different statistical properties. The first 14 values of the Haralick feature values of the context state correspond to the average of each statistical property over these four sets. Similarly, the last 14 values correspond to the difference between the maximum and minimum of each statistical property across these four sets. The implementation for extracting the Haralick features is provided by the Mahotas Python package [41].

Performance Prediction—For the prediction of the performances of each sensor band, machine learning models are utilized. They are trained to learn the relationship between the context state and the performance of their associated sensor band. Here, sensor performance refers to the Target Visibility Index (TVI) [15], a metric that quantifies the extent to which a camouflaged target is exposed in a given sensor band. It is given in Equation (1).

$$\text{TVI} = (\mu_T - \mu_B)^2(1 - 2\sigma_T)^2(1 - 2\sigma_B)^2 \qquad (1)$$

μ_T is the mean of all camouflaged target pixel values, μ_B is the mean of all background pixel values, σ_T is the standard deviation of all camouflaged target pixel values, and σ_B is the standard deviation of all background pixel values. The more a camouflaged target is exposed in a sensor band, the closer the TVI is to one. Likewise, the less a camouflaged target is exposed in a sensor band, the closer the TVI is to zero. A demonstration of the TVI is shown in Figure 3. The

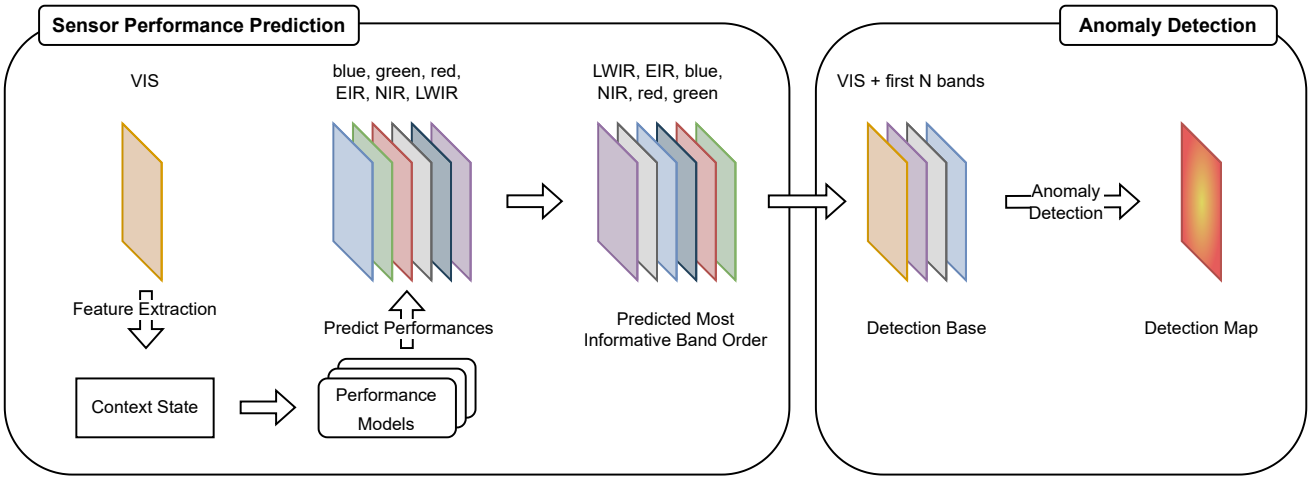

Figure 2: **The concept of sensor-managed anomaly detection for camouflage detection using the bands provided by the captures of MUDCAD-X.**

corresponding TVI values for each sensor band are shown in parentheses. As can be seen, the camouflaged target in the scene is barely visible in the visual and blue bands, resulting in a comparatively low TVI. In contrast, it is much exposed in the EIR and NIR bands, where the TVI gives much higher values. According to the TVI, the last image shows ideal visibility of the target. Note that the TVI is not based on human perception, but on predefined edge cases where it is either zero or one. Using the TVI as a quantification of sensor performance, the performance models are trained to learn the mapping from context state to TVI for their respective sensor band. After training, the models use the context state to predict the TVI for their associated sensor band. Ranking the sensor bands by their TVI, i.e. performance, prediction results in the predicted most informative band order. The detailed training procedures, machine learning model selection and parameter configurations are introduced in Section 4.

Anomaly Detection

As shown in Figure 2, the visual band is concatenated with the first N predicted most informative bands and then processed by an anomaly detector. N is a predefined constant that is set to three in the example illustrated in Figure 2. Since camouflaged targets are usually minimally exposed in the visual band, but ideally are the only objects maximally exposed in the N predicted most informative bands, the concatenation is intended to support their detection by the anomaly detector. The generated detection map constitutes the end result of the sensor-managed anomaly detection. In order to investigate the robustness of the sensor performance prediction approach for camouflage detection using anomaly detection, multiple detectors are employed in this work: the RX [20], the LRX [20], the LPD [14] and the AED [30] detector. They are compared to each other in Section 5 and briefly introduced below. For more details on the anomaly detection algorithms, see their original publications.

For the RX detector [20], the anomaly score for each pixel corresponds to its Mahalanobis distance. The larger the Mahalanobis distance, the higher the abnormality of the pixel of interest. It is calculated as stated in Equation (2), where d_M is the Mahalanobis distance of the pixel, x is the value vector of the pixel, μ is the mean vector over all bands, and Σ is the covariance matrix over all bands.

$$d_M = \sqrt{(x-\mu)^T \Sigma^{-1} (x-\mu)} \tag{2}$$

The LRX detector [20] works similar to the RX detector by calculating the Mahalanobis distance for each pixel according to Equation (2). But instead of considering all pixels in the capture, only a comparatively small neighborhood of pixels is considered for the calculation of μ and Σ. The neighborhood is defined by a dual window that spans around the pixel of interest with the inner window size w_i and the outer window size w_o. For the calculation of μ and Σ, only pixel value vectors inside the outer window but outside the inner window are considered. Compared to the RX detector, this allows for more accurate background estimates and therefore typically higher detection rates.

The LPD detector [14] uses local background estimation based on a sliding dual window similar to the LRX detector. But instead of calculating the Mahalanobis distance for each pixel as a measure of its abnormality, its density is computed as in Equation (3).

$$\rho = \frac{1}{s_M} \sum_{i=1}^{s_M} \exp\left(-\frac{\|x - m_i\|_2^2}{d_c^2}\right) \tag{3}$$

ρ is the density of the pixel of interest, x is its value vector, m_i is the ith element of M, s_m is the size of M, and d_c is the cut-off distance. $M = N \cup x$ is the neighborhood $N = \{n_i\}_{i=1}^{s_N}$ of the pixel of interest defined by the dual window appended by x. The cut-off distance is calculated as in Equation (4) where μ_M is calculated as in Equation (5).

$$d_c = \frac{1}{s_M} \sum_{i=1}^{s_M} \|m_i - \mu_M\|_2^2 \tag{4}$$

$$\mu_M = \frac{1}{s_M} \sum_{i=1}^{s_M} m_i \tag{5}$$

In the final density map, the low density pixels are more anomalous than the high density pixels.

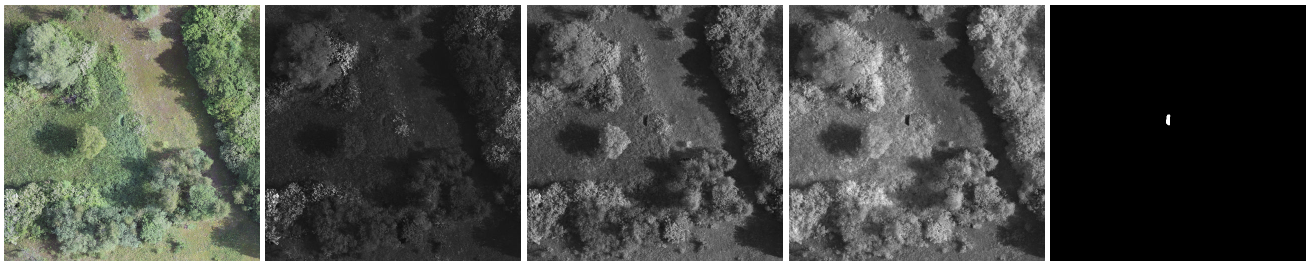

(a) VIS (0.0002) (b) blue (0.0002) (c) EIR (0.008) (d) NIR (0.016) (e) ideal (1.0)

Figure 3: **Demonstration of the Target Visibility Index (TVI).**

In AED [30] based detection, each band of the capture of interest is morphologically filtered using morphological attribute profiles. The filter removes bright and dark connected components that are smaller than a predefined area size defined by κ. In the final detection map, the areas where many connected components were removed are more anomalous than areas where no or few connected components were removed. This is motivated by the idea that anomalous objects result in many either bright or dark small continuous regions across all bands. As a post-processing step, the detection map is refined by an edge-preserving filter.

4. EXPERIMENTS

This section introduces the performance models, their training procedure, as well as the configurations of the anomaly detectors and the sensor-managed anomaly detection approach.

First, the captures of MUDCAD-X are split into training and test data. Since the dataset is divided into three different days, two different areas, and seven groups of captures, the fourth group of captures of each area and day is selected as test data, while the remaining observations are selected as training data. This leaves 724 captures for training and 129 captures for testing. The training data is used to train the performance models, while the test data is used to evaluate sensor-managed anomaly prediction approach.

Performance Model Training

For the training of the performance models, the context state is extracted from each gray-level converted visual band in the training data as described in Section 3. Furthermore, the feature values of the extracted context states are z-normalized over all training samples. This prevents features with relatively large value ranges from dominating feature values with relatively small value ranges during performance model training. In addition, the Target Visibility Index for each band is calculated for each capture in the training data. Since the TVI can only be calculated for a single camouflaged target, it is calculated separately for each camouflaged target in the scene. The resulting TVIs for each sensor band are then averaged to produce a single index value for each sensor band. With six managed sensor bands per capture, six separate performance models are trained. After training, the performance models predict the TVI for their respective sensor band given a z-normalized extracted context state. Since the ϵ-SVR models provided the most robust predictions in [15], they are also used as the performance models in this work. The optimal parameters of the models are obtained using a simple grid search with cross-validation over five folds and the root-mean-square error. Table 3 shows the final parameters for each performance model. Parameters not mentioned are left at their default values. The performance models and their training is implemented using Machine Learning Framework for Julia [42] in conjunction with the backend LIBSVM [43]. Note that the actual performance of the ϵ-SVR model predictions is not further evaluated. Instead, the evaluation is based solely on the anomaly detection results.

Table 3: **The training parameters of the ϵ-SVR performance models.**

Model	ϵ	C
blue	0.00383	0.054
green	0.00221	0.099
red	0.00481	0.065
EIR	0.00373	0.010
NIR	0.00682	0.012
LWIR	0.00533	0.076

Sensor-Managed Anomaly Detection Evaluation

With a successfully trained performance model for each sensor band, the sensor-managed anomaly detection approach can be evaluated. For each capture of the test data, the performance models predict the TVI of their respective sensor band. This results in the predicted most informative band order. The best N bands of the predicted most informative band order are concatenated with the visual band, forming the detection base. Each band of the detection base is z-normalized to avoid the influence of the value range of each band. In addition, the three bands of the visual band are weighted by one-third, since it is treated as a single band but actually consists of three individual bands. Finally, the detection base is reduced to a resolution of 256px x 256px ($20\mathrm{cm\,px^{-1}}$ ground sample distance) to lower the computational complexity. The detection base is then processed by several anomaly detectors: RX, LRX and LDP with inner and outer window sizes (w_i, w_o) of (5, 15), (11, 31), (21, 61), (31, 91) and (41, 121), and AED with a maximum area size κ of connected components to be removed of 25, 50, 100, 200 and 300. As a post-processing step, the binarized detection map is morphologically filtered to remove all detections with an area smaller than 10px, which is about one-third of a square meter.

For the evaluation of the proposed approach, N is set to three, which corresponds to half the number of possible bands. In addition, N is also set to six, which corresponds to the unmanaged reference, where performance predictions are

ignored. Higher values than three of N are not investigated because the sensor-managed anomaly detection approach is intended to reduce computational complexity by reducing the number of bands to process. Similarly, lower values of N are not examined to avoid overly complex results. The following evaluation is based on a comparison of the sensor-manged anomaly detection results with the unmanaged anomaly detection results.

5. RESULTS

This section evaluates and discusses the anomaly detection results of the proposed sensor-managed approach and the unmanaged reference.

The evaluation of the detection maps of the different anomaly detection methods is based on the receiver operating characteristic (ROC), the area under curve (AUC) and intersection over union (IoU). It focuses primarily on the comparison of the managed and unmanaged approaches, as anomaly detection for camouflage detection in general has already been investigated in [14].

Evaluation

Figure 4 shows the ROCs of the detection methods for each camouflaged target in the test data. UM stands for unmanaged anomaly detection while M3 stands for sensor-managed anomaly detection with an N set to 3. Each colored curve shows the ROC of a single anomaly detector. The unmanaged curves are solid and the managed curves are dashed. In addition, the AUC is shown in parentheses for each ROC. Since each target occurs multiple times in the test data, the ROC for each anomaly detection method is calculated by threshold-based averaging [44] over all target occurrences in the test data. Also, since there is more than one configuration for each anomaly detector, only the one that results in the highest AUC for the respective target is shown. Note that there are different ranges for the abscissa for all camouflaged targets.

Although the sensor-managed anomaly detectors perform generally slightly worse than their unmanaged counterparts, they outperform them for some targets, such as the artificial turf or the anthracite fleece. In addition, even though their AUCs are lower, the managed detectors often achieve higher detection rates with lower false alarm rates at high thresholds than the unmanaged detectors, as can be seen for the green 2D camouflage net and the green 3D camouflage net. The managed curves initially rise more steeply than the unmanaged curves, but they eventually intersect, resulting in a lower AUC. Even if this is not the case, the managed and unmanaged curves of each anomaly detector tend to be very close at higher thresholds. Overall, the anomaly detectors show mostly high detection rates with relatively low false alarm rates in both cases. However, it is worth noting that the LRX detector significantly outperforms all other detectors, while the AED detector performs the worst.

The AUCs and the parameter configurations of the anomaly detector in Figure 4 are additionally shown in Table 4. Here, the highest AUC per camouflaged target is marked bold. As already observed in Figure 4, the managed detectors perform slightly worse than the unmanaged detectors, with only three out of eleven top results. However, their AUCs are close to those of the unmanaged detectors. As can be seen from the parameters, the window sizes rise and fall mostly with the camouflaged target sizes, as does the area of the connected components to be removed. Note that the configurations are primarily affected by the camouflaged target, not the managed or unmanaged detection approach.

For a different perspective on the detection performances, Table 5 shows the highest achievable average IoU over all samples in the test data, given a fixed binarization threshold of the detection maps for each anomaly detector. The first part of the table shows the average IoUs that could be obtained by applying the thresholds given in percent in the middle part of the table. Parameter configurations of each anomaly detector are shown in the last part of the table, similar to Table 4. Again, the highest IoU per camouflaged target is marked bold. Compared to the results in Table 4, where the managed detectors are inferior to the unmanaged detectors, the managed detectors outperform the unmanaged detectors in terms of average IoU with seven out of eleven top results. This indicates that the managed detectors work more accurately at higher binarization thresholds than the unmanaged counterparts, as IoU tends to decrease rapidly with decreasing binarization thresholds, supporting the same observations already found in the ROCs in Figure 4. Although the IoU values may seem small in general, camouflage detection is an inherently difficult task, which naturally results in relatively low detection accuracies, especially when the camouflaged targets are not sufficiently exposed in any sensor band, as discussed in [14]. Similar to Table 4, the optimal parameter configurations of the detectors change primarily with the size of the camouflaged targets.

Discussion

The results show that the sensor-managed anomaly detection approach is generally slightly inferior to the unmanaged approach in terms of ROC and AUC. Although the bands processed by the managed detectors are more informative and expose the camouflaged targets more than the other bands, the detection rates are not significantly increased. This indicates that the increased spectral information in the unmanaged case does not confuse the anomaly detectors, but actually supports the detection of the camouflaged targets. However, the managed detectors outperform the unmanaged detectors in some cases. In addition, for the targets where they achieve lower AUCs, they tend to have higher detection rates with lower false alarm rates at higher thresholds. This is further confirmed by evaluating the highest achievable average IoU, given a fixed binarization threshold, where the managed detectors outperform the unmanaged detectors. Since higher thresholds are more practical in real-world applications to avoid too many false positive detections, the managed detectors are more appropriate in actual reconnaissance scenarios. Furthermore, the managed detectors generally produce comparable results to the unmanaged detectors, but with fewer bands to process, which is an improvement in terms of computational complexity. This is also an indication that more exposed camouflaged targets according to the Target Visibility Index support their detection, as processing fewer bands with high performance predictions yields very similar results as processing all bands. For a quantification of the actual reduction in computational complexity in terms of processing time, see [14], which shows the runtimes of all applied anomaly detectors for two different band stack sizes. Overall, the sensor-managed anomaly detection approach does not outperform the unmanaged reference in all respects. Yet, it appears to be more valuable in real-world applications and thus has great potential worthy of further investigation.

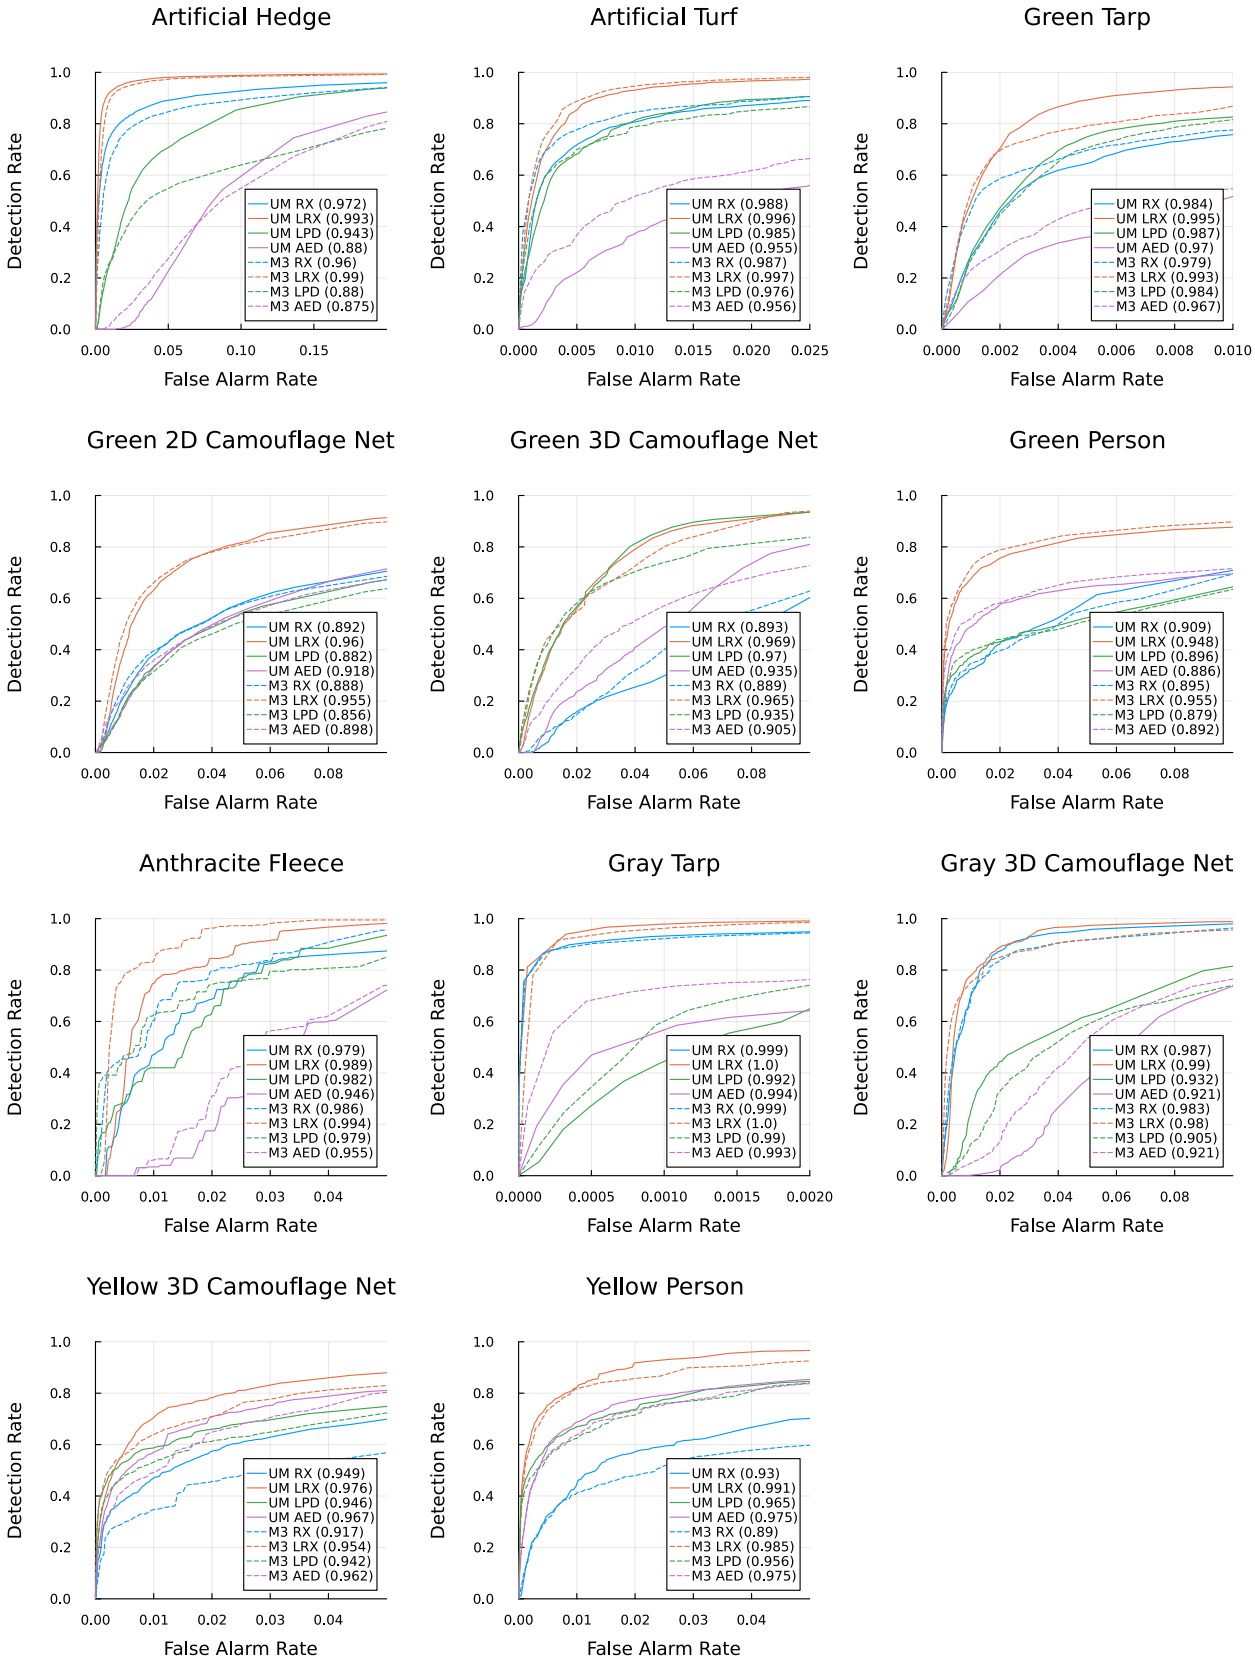

Figure 4: **Receiver operating characteristics (ROCs) of all detection methods for each camouflaged target. Managed detectors are dashed and unmanaged detectors are solid. Area under curves (AUCs) are shown in parentheses**

Table 4: **The area under curves (AUCs) and the parameter configurations for each anomaly detector and target. The results correspond to those shown in Figure 4. Top results for each target are bold. The parameters show w_i, w_o for LRX and LPD, and κ for AED. RX is parameterless.**

Area Under Curve (AUC)											
	Hedge	Turf	Grn. Tarp	Grn. 2D	Grn. 3D	Grn. Per.	Fleece	Gry. Tarp	Gry. 3D	Yew. 3D	Yew. Per.
	Unmanaged										
RX	0.972	0.988	0.984	0.892	0.893	0.909	0.979	0.999	0.987	0.949	0.930
LRX	**0.993**	0.996	**0.995**	**0.960**	0.969	0.948	0.989	**1.000**	**0.990**	**0.976**	**0.991**
LPD	0.943	0.985	0.987	0.882	**0.970**	0.896	0.982	0.992	0.932	0.946	0.965
AED	0.880	0.955	0.970	0.918	0.935	0.886	0.946	0.994	0.921	0.967	0.975
	Managed (N=3)										
RX	0.960	0.987	0.979	0.888	0.889	0.895	0.986	0.999	0.983	0.917	0.890
LRX	0.990	**0.997**	0.993	0.955	0.965	**0.955**	**0.994**	1.000	0.980	0.954	0.985
LPD	0.880	0.976	0.984	0.856	0.935	0.879	0.979	0.990	0.905	0.942	0.956
AED	0.875	0.956	0.967	0.898	0.905	0.892	0.955	0.993	0.921	0.962	0.975
Parameters											
	Hedge	Turf	Grn. Tarp	Grn. 2D	Grn. 3D	Grn. Per.	Fleece	Gry. Tarp	Gry. 3D	Yew. 3D	Yew. Per.
	Unmanaged										
RX						$--$					
LRX	$21, 61$	$21, 61$	$21, 61$	$21, 61$	$21, 61$	$5, 15$	$31, 91$	$41, 121$	$21, 61$	$21, 61$	$11, 31$
LPD	$11, 31$	$11, 31$	$11, 31$	$11, 31$	$21, 61$	$31, 91$	$41, 121$	$21, 61$	$11, 31$	$21, 61$	$11, 31$
AED	200	100	200	50	300	50	50	200	300	200	100
	Managed (N=3)										
RX						$--$					
LRX	$21, 61$	$21, 61$	$21, 61$	$21, 61$	$21, 61$	$5, 15$	$31, 91$	$41, 121$	$11, 31$	$11, 31$	$11, 31$
LPD	$11, 31$	$11, 31$	$11, 31$	$11, 31$	$21, 61$	$31, 91$	$41, 121$	$21, 61$	$41, 121$	$21, 61$	$11, 31$
AED	200	100	200	50	300	50	50	200	300	200	100

6. CONCLUSIONS

In this paper, a sensor-managed anomaly detection approach for camouflage detection is proposed. It is based on performance models that determine the bands that most expose camouflaged targets, which are then processed by an anomaly detector to produce the final detection map. While the results show that the approach does not generally outperform its unmanaged counterpart, which simply processes all bands, the difference is almost negligible. In addition, the proposed method provides more precise detections at higher binarization thresholds, making it more suitable for real-world reconnaissance scenarios. Furthermore, the performance achieved requires fewer bands to be processed by an anomaly detector, reducing computational complexity. Therefore, the sensor-managed anomaly detection approach has great potential for practical applications. In order to gain further insight into the benefits of the sensor-managed anomaly detection approach, future research will focus on applications with multispectral data containing many more bands.

REFERENCES

[1] J. A. J. Berni, P. J. Zarco-Tejada, L. Suarez, and E. Fereres, "Thermal and Narrowband Multispectral Remote Sensing for Vegetation Monitoring From an Unmanned Aerial Vehicle," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 3, pp. 722–738, Mar. 2009.

[2] S. Gxokwe, T. Dube, and D. Mazvimavi, "Multispectral Remote Sensing of Wetlands in Semi-Arid and Arid Areas: A Review on Applications, Challenges and Possible Future Research Directions," *Remote Sensing*, vol. 12, no. 24, p. 4190, Jan. 2020.

[3] N. Bagheri, H. Ahmadi, S. K. Alavipanah, and M. Omid, "Multispectral remote sensing for site-specific nitrogen fertilizer management," *Pesquisa Agropecuária Brasileira*, vol. 48, pp. 1394–1401, Oct. 2013.

[4] L. Deng, Z. Mao, X. Li, Z. Hu, F. Duan, and Y. Yan, "UAV-based multispectral remote sensing for precision agriculture: A comparison between different cameras," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 146, pp. 124–136, Dec. 2018.

[5] L. Zhang, H. Zhang, Y. Niu, and W. Han, "Mapping Maize Water Stress Based on UAV Multispectral Remote Sensing," *Remote Sensing*, vol. 11, no. 6, p. 605, Jan. 2019.

[6] B. Huang, B. Zhao, and Y. Song, "Urban land-use mapping using a deep convolutional neural network with high spatial resolution multispectral remote sensing imagery," *Remote Sensing of Environment*, vol. 214, pp. 73–86, Sep. 2018.

[7] M. Mahdianpari, B. Salehi, M. Rezaee, F. Mohammadimanesh, and Y. Zhang, "Very Deep Convolutional Neural Networks for Complex Land Cover Mapping

Table 5: **The highest achievable average intersection over union (IoU) of all detectors for each target, given a fixed binarization threshold of the detection maps. Top results for each target are bold. The thresholds for each detector are shown in percent in the middle part of the table, while their parameter configurations are shown in the last part of the table. The parameters correspond to w_i, w_o for LRX and LPD, and κ for AED. RX is parameterless.**

Intersection over Union (IoU)											
	Hedge	Turf	Grn. Tarp	Grn. 2D	Grn. 3D	Grn. Per.	Fleece	Gry. Tarp	Gry. 3D	Yew. 3D	Yew. Per.
	Unmanaged										
RX	0.255	0.167	0.188	0.065	0.017	0.060	0.014	0.758	0.085	0.103	0.037
LRX	**0.297**	0.248	0.325	0.106	0.078	**0.233**	0.113	0.748	0.152	0.226	0.207
LPD	0.042	0.168	0.231	0.050	0.078	0.228	0.094	0.509	0.024	**0.307**	**0.212**
AED	0.009	0.031	0.095	0.035	0.027	0.085	0.004	0.340	0.007	0.207	0.181
	Managed (N=3)										
RX	0.164	**0.299**	0.296	0.076	0.020	0.049	0.057	0.754	0.098	0.077	0.038
LRX	0.227	0.296	**0.330**	**0.124**	0.105	0.217	0.180	**0.761**	**0.188**	0.209	0.174
LPD	0.071	0.202	0.202	0.056	**0.106**	0.219	**0.195**	0.588	0.026	0.283	0.186
AED	0.009	0.112	0.161	0.045	0.061	0.123	0.005	0.543	0.009	0.151	0.167
Thresholds											
	Hedge	Turf	Grn. Tarp	Grn. 2D	Grn. 3D	Grn. Per.	Fleece	Gry. Tarp	Gry. 3D	Yew. 3D	Yew. Per.
	Unmanaged										
RX	99.750	99.625	99.375	97.375	96.750	99.813	98.750	99.813	99.375	99.875	99.500
LRX	99.688	99.813	99.750	98.313	98.938	99.875	99.250	99.813	99.438	99.938	99.813
LPD	99.438	99.875	99.688	98.750	99.000	99.563	99.875	99.813	96.938	99.875	99.813
AED	90.000	99.313	99.563	96.188	98.250	99.500	95.500	99.813	91.000	99.875	99.750
	Managed (N=3)										
RX	99.625	99.875	99.688	98.188	98.313	99.750	99.750	99.813	99.563	99.813	99.625
LRX	99.500	99.750	99.750	98.563	99.188	99.938	99.438	99.813	99.500	99.938	99.813
LPD	99.625	99.875	99.688	98.438	99.188	99.688	99.938	99.813	97.563	99.875	99.438
AED	92.000	99.875	99.813	98.688	99.438	99.750	98.438	99.813	95.125	99.875	99.813
Parameters											
	Hedge	Turf	Grn. Tarp	Grn. 2D	Grn. 3D	Grn. Per.	Fleece	Gry. Tarp	Gry. 3D	Yew. 3D	Yew. Per.
	Unmanaged										
RX						–					
LRX	$21, 61$	$21, 61$	$21, 61$	$31, 91$	$21, 61$	$11, 31$	$5, 15$	$41, 121$	$11, 31$	$11, 31$	$11, 31$
LPD	$21, 61$	$21, 61$	$21, 61$	$21, 61$	$21, 61$	$5, 15$	$41, 121$	$41, 121$	$11, 31$	$11, 31$	$11, 31$
AED	200	100	200	50	300	25	50	200	300	100	25
	Managed (N=3)										
RX						–					
LRX	$21, 61$	$11, 31$	$21, 61$	$31, 91$	$21, 61$	$11, 31$	$5, 15$	$41, 121$	$11, 31$	$11, 31$	$11, 31$
LPD	$21, 61$	$21, 61$	$21, 61$	$21, 61$	$21, 61$	$5, 15$	$41, 121$	$41, 121$	$11, 31$	$11, 31$	$5, 15$
AED	200	100	200	50	300	25	25	200	200	100	25

Using Multispectral Remote Sensing Imagery," *Remote Sensing*, vol. 10, no. 7, p. 1119, Jul. 2018.

[8] İ. Güneralp, A. M. Filippi, and J. Randall, "Estimation of floodplain aboveground biomass using multispectral remote sensing and nonparametric modeling," *International Journal of Applied Earth Observation and Geoinformation*, vol. 33, pp. 119–126, Dec. 2014.

[9] P. Wicaksono, P. Danoedoro, n. Hartono, and U. Nehren, "Mangrove biomass carbon stock mapping of the Karimunjawa Islands using multispectral remote sensing," *International Journal of Remote Sensing*, vol. 37, no. 1, pp. 26–52, Jan. 2016.

[10] F. Eugenio, J. Marcello, and J. Martin, "High-Resolution Maps of Bathymetry and Benthic Habitats in Shallow-Water Environments Using Multispectral Remote Sensing Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 7, pp. 3539–3549, Jul. 2015.

[11] W. J. Hernandez and R. A. Armstrong, "Deriving Bathymetry from Multispectral Remote Sensing Data," *Journal of Marine Science and Engineering*, vol. 4, no. 1, p. 8, Mar. 2016.

[12] T. Hu and R. B. Smith, "The Impact of Hurricane Maria on the Vegetation of Dominica and Puerto Rico Using Multispectral Remote Sensing," *Remote Sensing*, vol. 10, no. 6, p. 827, Jun. 2018.

[13] T. D. Mushore, O. Mutanga, J. Odindi, and T. Dube, "Determining extreme heat vulnerability of Harare Metropolitan City using multispectral remote sensing and socio-economic data," *Journal of Spatial Science*, vol. 63, no. 1, pp. 173–191, Jan. 2018.

[14] T. Hupel and P. Stütz, "Adopting hyperspectral anomaly detection for near real-time camouflage detection in multispectral imagery," *Remote Sensing*, vol. 14, no. 15, 2022.

[15] ——, "Measuring and predicting sensor performance for camouflage detection in multispectral imagery," *Sensors*, vol. 23, no. 19, 2023.

[16] C. Hellert, S. Koch, and P. Stütz, "Using algorithm selection for adaptive vehicle perception aboard uav," in *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, Sep. 2019, pp. 1–8, iSSN: 2643-6213.

[17] M. Ruß and P. Stütz, "Airborne sensor and perception management," in *Modelling and Simulation for Autonomous Systems*, ser. Lecture Notes in Computer Science, J. Mazal, A. Fagiolini, P. Vašík, A. Bruzzone, S. Pickl, V. Neumann, P. Stodola, and S. Lo Storto, Eds. Cham: Springer International Publishing, 2023, pp. 182–206.

[18] M. Zwick, M. Gerdts, and P. Stütz, "Sensor-model-based trajectory optimization for uavs to enhance detection performance: An optimal control approach and experimental results," *Sensors*, vol. 23, no. 2, 2023.

[19] S. Koch, B. Krach, F. Katsilieris, and P. Stütz, "Sensor scheduling strategies for 1-to-n multi-object tracking," in *2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC)*, Sep. 2022, pp. 1–9, iSSN: 2155-7209.

[20] I. S. Reed and X. Yu, "Adaptive multiple-band cfar detection of an optical pattern with unknown spectral distribution," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 10, pp. 1760–1770, 1990.

[21] H. Kwon and N. M. Nasrabadi, "Kernel rx-algorithm: A nonlinear anomaly detector for hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, no. 2, pp. 388–397, Feb. 2005.

[22] Q. Guo, B. Zhang, Q. Ran, L. Gao, J. Li, and A. Plaza, "Weighted-rxd and linear filter-based rxd: Improving background statistics estimation for anomaly detection in hyperspectral imagery," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 6, pp. 2351–2366, 2014.

[23] J. Zhou, C. Kwan, B. Ayhan, and M. T. Eismann, "A novel cluster kernel rx algorithm for anomaly and change detection using hyperspectral images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 11, pp. 6497–6504, Nov. 2016.

[24] W. Li and Q. Du, "Collaborative representation for hyperspectral anomaly detection," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 3, pp. 1463–1474, Mar. 2015.

[25] H. Su, Z. Wu, Q. Du, and P. Du, "Hyperspectral anomaly detection using collaborative representation with outlier removal," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 11, no. 12, pp. 5029–5038, Dec. 2018.

[26] Y. Chen, N. M. Nasrabadi, and T. D. Tran, "Sparse representation for target detection in hyperspectral imagery," *IEEE Journal on Selected Topics in Signal Processing*, vol. 5, no. 3, pp. 629–640, Jun. 2011.

[27] J. Li, H. Zhang, L. Zhang, and L. Ma, "Hyperspectral anomaly detection by the use of background joint sparse representation," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 6, pp. 2523–2533, Jun. 2015.

[28] Y. Xu, Z. Wu, J. Li, A. Plaza, and Z. Wei, "Anomaly detection in hyperspectral images based on low-rank and sparse representation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 4, pp. 1990–2000, Apr. 2016.

[29] 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, no. 2, 2005.

[30] X. Kang, X. Zhang, S. Li, K. Li, J. Li, and J. A. Benediktsson, "Hyperspectral anomaly detection with attribute and edge-preserving filters," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 10, pp. 5600–5611, Oct. 2017.

[31] R. Wang, F. Nie, Z. Wang, F. He, and X. Li, "Multiple features and isolation forest-based fast anomaly detector for hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 9, pp. 6664–6676, Sep. 2020.

[32] B. Tu, X. Yang, N. Li, C. Zhou, and D. He, "Hyperspectral anomaly detection via density peak clustering," *Pattern Recognition Letters*, vol. 129, pp. 144–149, Jan. 2020.

[33] B. Tu, X. Yang, C. Zhou, D. He, and A. Plaza, "Hyperspectral anomaly detection using dual window density," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 12, pp. 8503–8517, Dec. 2020.

[34] W. Li, G. Wu, and Q. Du, "Transferred deep learning for anomaly detection in hyperspectral imagery," *IEEE*

Geoscience and Remote Sensing Letters, vol. 14, no. 5, pp. 597–601, May 2017.

[35] J. Lei, S. Fang, W. Xie, Y. Li, and C. I. Chang, "Discriminative reconstruction for hyperspectral anomaly detection with spectral learning," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 10, 2020.

[36] S. Wang, X. Wang, L. Zhang, and Y. Zhong, "Auto-AD: Autonomous Hyperspectral Anomaly Detection Network Based on Fully Convolutional Autoencoder," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2022.

[37] G. Fan, Y. Ma, X. Mei, F. Fan, J. Huang, and J. Ma, "Hyperspectral Anomaly Detection With Robust Graph Autoencoders," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2022.

[38] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Gray scale and rotation invariant texture classification with local binary patterns," in *Computer Vision - ECCV 2000*, ser. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2000, pp. 404–420.

[39] R. M. Haralick, I. Dinstein, and K. Shanmugam, "Textural features for image classification," *IEEE Transactions on Systems, Man and Cybernetics*, vol. SMC-3, no. 6, pp. 610–621, Nov. 1973.

[40] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," *SIAM Review*, vol. 59, no. 1, pp. 65–98, 2017.

[41] L. P. Coelho, "Mahotas: Open source software for scriptable computer vision," *Journal of Open Research Software*, vol. 1, 2013.

[42] A. D. Blaom, F. Kiraly, T. Lienart, Y. Simillides, D. Arenas, and S. J. Vollmer, "Mlj: A julia package for composable machine learning," *Journal of Open Source Software*, vol. 5, no. 55, p. 2704, 2020.

[43] C.-C. Chang and C.-J. Lin, "Libsvm: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011, software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.

[44] T. Fawcett, "An introduction to roc analysis," *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, Jun. 2006.

BIOGRAPHY

Tobias Hupel *received his B.Eng. and M.Eng. in Computer Science from the Ernst-Abbe-Hochschule Jena. He currently works as a research associate and Ph.D. student at the Institute of Flight Systems at the University of the Bundeswehr Munich. His main research interests are multispectral image processing, camouflage detection, sensor performance modelling, and sensor management for reconnaissance applications.*

Peter Stütz *is a Professor of Aeronautical Engineering at the Institute of Flight Systems at the University of the Bundeswehr Munich. His scientific focus lies on platform and equipment aspects of unmanned aerial vehicles. This includes avionics, communication and sensor related issues. He considers his work on sensor integration and sensor data analysis primarily as a means to an end to enable autonomous vehicle behavior.*

5.4 Optimized Spectral Indices for Camouflage Detection in Multispectral Imagery

Tobias Hupel and Peter Stütz. „Optimized Spectral Indices for Camouflage Detection in Multispectral Imagery“. In: *GIScience & Remote Sensing* 62.1 (2025), p. 2508574. DOI: 10.1080/15481603.2025.2508574

GISCIENCE & REMOTE SENSING
2025, VOL. 62, NO. 1, 2508574
https://doi.org/10.1080/15481603.2025.2508574

Taylor & Francis
Taylor & Francis Group

Optimized spectral indices for camouflage detection in multispectral imagery

Tobias Hupel and Peter Stütz

Institute of Flight Systems, University of the Bundeswehr Munich, Neubiberg, Germany

ABSTRACT

Unmanned aerial vehicles, equipped with multispectral imaging systems, as well as publicly available datasets, have fueled various research in remote sensing applications, including precision agriculture, land cover mapping or even camouflage detection. Many of these applications make use of spectral indices, such as normalized difference or ratio indices, created by merging multiple raw bands. These indices typically provide a direct indication of certain physical surface properties, like plant health, nitrogen content or leaf area index, and are determined by studying spectral reflectance properties or using optimization techniques. Given the heavy use and utility of such indices, this work introduces a novel generalization of the normalized difference, ratio and difference indices, the linear ratio index (LRI), a ratio of two linear functions of all available bands. In addition, an optimization approach for the LRI is presented, which incorporates complexity reduction strategies that enable optimization using only a subset of all available bands and reduced parameter precision, thereby ensuring parameter readability. The LRI and its optimization are thoroughly investigated in the context of camouflage detection in tactical reconnaissance scenarios by optimizing a six-band and a two-band LRI using the eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X). For comparison with traditional spectral index optimization approaches, the resulting linear ratio indices (LRIs) are evaluated against all raw bands and an optimized normalized difference index and an optimized ratio index. The evaluation shows that the optimized LRIs provide the best overall results in terms of visibility and detectability of camouflaged targets. This could indicate a general superiority of the LRI over established indices optimized by testing band permutations, making the LRI a promising candidate for further investigation in other remote sensing applications where it could also outperform traditional index optimization approaches. Therefore, the software code for optimizing the LRI has been made publicly available for further exploitation, requiring only an adapted optimization criterion to support any other use case. Furthermore, as with most spectral indices, the LRIs obtained in this study have negligible computational overhead and, under the right conditions, can be directly integrated into any existing camouflage detection system.

ARTICLE HISTORY
Received 10 October 2024
Accepted 14 May 2025

KEYWORDS
multispectral; image processing; spectral index optimization; camouflage detection; anomaly detection

1. Introduction

Multispectral remote sensing has become a heavily studied field of research in recent years. Commercially available unmanned aerial vehicles (UAVs) and multispectral imaging systems, as well as public multispectral satellite imagery, provide easy access to valuable spectral data and have paved the way for numerous applications. Current advancements range from crop-type classification (Akbari et al. 2020; X. Wang et al. 2023; Wijayanto, Wahyu Triscowati, and Marsuhandi 2020), crop yield estimation (Anastasiou et al. 2018; Camenzind and Yu 2024; Psiroukis et al. 2022) and vineyard zoning (Gavrilovic et al. 2024) to ship detection (Wang, Ren et al. 2021), surface water mapping (Albertini et al. 2022; Cordeiro, Martinez, and Peña-Luque 2021; S. Wang et al. 2022), oil spill detection (Pérez-García et al. 2024) and rock classification (Ghrefat et al. 2021; Karimzadeh and Tangestani 2021; Li et al. 2024). In addition, recent research has shown the effectiveness of multispectral imagery for camouflage detection, using spectral anomaly (Hupel and Stütz 2022, 2024; Luo, Wang, and Deng 2023; Nandibewoor et al. 2024) or target (Nandibewoor et al. 2024, Shen et al. 2021; S. Wang et al. 2024) detection. Considering the relatively low size, weight, power and cost (SWaP-C) drawbacks of multispectral imaging systems, their use on small UAVs could provide superior performance in reconnaissance scenarios compared to traditional visual or thermal imaging systems. For this reason, the Institute of Flight Systems of the University of the Bundeswehr Munich investigates the possibilities and benefits of multispectral imaging systems on small tactical UAVs for reconnaissance applications.

Since spectral properties correspond directly to physical properties, multispectral imaging provides a non-destructive way to analyze surface conditions compared

CONTACT Tobias Hupel tobias.hupel@unibw.de

to traditional destructive sampling. This typically involves the use of spectral indices that give explicit insight into a predetermined physical surface condition. They are calculated from different spectral bands provided by the sensor or dataset. Over the years, a myriad of indices have been proposed for a wide variety of applications. For example, the normalized difference vegetation index (NDVI) (Rouse et al. 1974) is a measure of plant health, while the analytical burned area index (ABAI) (Wu et al. 2022) makes it easy to distinguish between burned and healthy forest areas. In most cases, these indices are either directly derived from spectral reflectance properties or determined through extensive field studies, experiments and numerical optimization techniques. Once an appropriate index has been found, it yields straightforward information about a particular surface property with negligible computational overhead, which is especially beneficial on resource-limited platforms, such as small tactical UAVs.

In view of the heavy use of spectral indices in multispectral remote sensing applications (Huang et al. 2021; Mouafik et al. 2024; Tran, Reef, and Zhu 2022; Zeng et al. 2020), this work introduces a novel index, the linear ratio index (LRI), a generalization of the popular normalized difference, ratio and difference indices (Jordan 1969; Rouse et al. 1974; Tucker 1979). Specifically, the generalization is achieved by extending the normalized difference index to a weighted linear sum of all available spectral bands in both the numerator and denominator. In this way, all bands can contribute to the index and normalized difference, ratio or difference indices can be considered as special instances of linear ratio indices (LRIs). Furthermore, a weighted sum of bands allows for different contributions and thus more complex relationships between spectral bands. However, taking advantage of the additional capabilities of the LRI also requires an appropriate approach to finding its parameters. For this purpose, this work employs an optimization algorithm based on differential evolution (Feldt and Stukalov 2018) that determines optimal parameters for the LRI with respect to an application-specific optimization criterion. In addition, the optimization algorithm is enhanced by complexity reduction strategies that allow certain parameters to be optimized to zero, resulting in LRIs that use fewer than all available bands. This could be useful in scenarios where the set of utilized bands needs to be reduced or the bands that contribute most to the spectral index need to be determined.

Building on previous research (Hupel and Stütz 2022, 2023, 2024), this study investigates the use of spectral indices specifically tailored to support camouflage detection using multispectral imaging in tactical reconnaissance scenarios. For this purpose, the LRI is optimized to increase the visibility and detectability of camouflaged targets in multispectral imagery using the Taget Visibility Index (TVI) (Hupel and Stütz 2023), a fast and efficient measure of the visibility of an object in a nadir image. Here, the eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X) (Hupel and Stütz 2023) serves as a data basis for optimization and evaluation by splitting it into dedicated training and test datasets. With the training dataset, two different parameter sets for the LRI are determined, one that uses all available bands and another that uses only two bands, since two-band indices are the most dominant form of spectral indices. As a baseline for these indices, a normalized difference and a ratio index are also optimized using the training dataset. Although there are countless other indices, most are designed for very specific applications that do not include camouflage detection. As a result, generic normalized difference and ratio indices are commonly optimized for other use cases (Feng et al. 2015; Guo et al. 2023; Rivera et al. 2014; Fang et al. 2015; Z. Wang et al. 2018; Y. Zhang et al. 2023; W. Zhang et al. 2023; Z. Zhang et al. 2023) and therefore provide a solid basis for comparison with the LRI. The test dataset is used to evaluate the optimized indices with respect to their ability to increase the visibility and detectability of camouflaged targets. For the evaluation of the impact on visibility, the average TVI for each index over all samples is calculated. In order to investigate the influence on detection performance, the optimized indices are incorporated into a camouflage detection pipeline based on several different spectral anomaly detection methods. The detection results are then evaluated and compared using receiver operating characteristic (ROC), area under curve (AUC) and intersection over union (IOU).

While this study focuses specifically on the application of the LRI to camouflage detection, the LRI is generally not bound to any specific application. Instead, the application of the LRI is determined by the criterion for which it is optimized, which is the TVI in this work. However, the optimization approach generally allows any optimization criterion, since it is based on differential evolution, a very versatile metaheuristic. Consequently, the LRI can be optimized and applied to any other use case, given an appropriate optimization criterion. Therefore, the software code to optimize the LRI has been made publicly available on GitHub to encourage and motivate other researchers to experiment with the proposed approach in different contexts or to replicate the results of this study. See the data availability statement at the end of this document for the specific repository URL.

Overall, this work provides the following scientific contributions:

- a novel generalization of the normalized difference, ratio and difference indices, the linear ratio index, along with an optimization approach that allows for different index complexities
- a comprehensive evaluation of the LRI and its optimization for camouflage detection using multispectral imaging in tactical reconnaissance scenarios
- application-ready linear ratio indices of different complexities improving the visibility and detectability of camouflaged targets

The remaining paper is structured as follows. Section 2 highlights previous work on spectral index optimization and camouflage detection. Then, Section 3 introduces the proposed index along with the approach for finding its optimized parameters. In addition, it provides a detailed explanation of the dataset and methods used in this work. Section 4 presents the conducted experiments and their results. The main findings, possible implications and limitations of this study as well as future prospects are subsequently discussed in Section 5. Finally, Section 6 summarizes and concludes this paper.

2. Related work

This section provides an overview of the work related to this study. First, research on spectral index optimization for various use cases is discussed in Section 2.1. Then, Section 2.2 outlines work related to camouflage detection using spectral imaging.

2.1. Spectral index optimization

Most spectral indices are based on the early normalized difference, ratio and difference indices (Jordan 1969; Rouse et al. 1974; Tucker 1979), which are calculated as shown in Equation 1-Equation 3.

$$I_{nd}(c_a, c_b) = \frac{c_a - c_b}{c_a + c_b} \tag{1}$$

$$I_r(c_a, c_b) = \frac{c_a}{c_b} \tag{2}$$

$$I_d(c_a, c_b) = c_a - c_b \tag{3}$$

$c_a \in \mathbb{R}^{m \times l}$ and $c_b \in \mathbb{R}^{m \times l}$ are the spectral input bands where m denotes the height and l denotes the width of the bands. I_{nd}, I_r and I_d correspond to the normalized difference, ratio and difference index, respectively. Note that the division operations are performed on a per-element basis, so that all resulting index images have the same size as the original input bands. Additionally, the input bands are often converted to reflectance maps before the index image is calculated. This theoretically ensures consistency of the index images regardless of lighting conditions.

Determining the appropriate spectral regions for the input bands is one of the most important tasks in developing and optimizing a spectral index. This is typically accomplished through spectral behavior analysis, extensive field studies and empirical evaluation. For example, the normalized difference vegetation index (NDVI) (Rouse et al. 1974), the most popular indicator of plant health, was developed by studying the reflectance properties of chlorophyll and general vegetation. Further developments include the soil-adjusted vegetation index (SAVI) (Huete 1988), which minimizes the influence of high soil reflectance, and the enhanced vegetation index (EVI) (Huete, Justice, and van Leeuwen 1999), which accounts for areas of high biomass and some atmospheric disturbances. Other research has led to the anthocyanin context index (ACI) (Berg and Perkins 2005) or the modified chlorophyll absorption in reflectance index (MCARI) (Daughtry et al. 2000), which provide an indicator of anthocyanin and chlorophyll content, respectively.

In contrast to tedious studies of reflectance behavior, recent research focuses on numerical optimization techniques to find optimal spectral indices. Common methods include testing all possible band combinations or testing different parameter sets for existing index formulations, followed by evaluation against a specific optimization criterion. Given the indices defined in Equation 1-Equation 3, this procedure would require testing all possible band combinations for c_a and c_b. The set of possible band combinations results naturally from the available sensor or dataset. For instance, normalized difference, ratio and difference indices have been optimized with all available bands to estimate leaf area index (Rivera et al. 2014; Wang, Sun et al. 2018), as well as anthocyanin (Z. Zhang et al. 2023), chlorophyll (Rivera et al. 2014), nitrogen (Feng et al. 2015; Guo et al. 2023; Zhang et al. 2023) and phosphorus (Y. Zhang et al. 2023) content. In addition (Feng et al. 2015), incorporated a third band and an additional scalar parameter, both of which were optimized by testing all values of a predefined set. This resulted in an even stronger indicator of nitrogen content than any of the optimized two-band indices, the modified reg-edge ratio (mRER). Furthermore (Gerstmann, Möller, and Gläßer 2016), optimized the EVI for cereal crop classification by permuting the set of possible input bands and the scalar parameters. In order to estimate the brown leaf

area index (Delegido et al. 2015), optimized a normalized difference index and proposed the Green Brown Vegetation Index (GBVI) to separate green from brown vegetation. Moreover, optimized input band selections for complex two- and three-band indices have been successfully determined to estimate biomass (Yang, Li, Wang, et al. 2021), as well as phosphorus (Wang, Shi et al. 2016), chlorophyll (Jiang et al. 2022) and nitrogen (Yang et al. 2023; Yang, Li, Hu, et al. 2021) content in various crops.

Apart from testing possible band and parameter combinations, more sophisticated optimization techniques have also been explored. For example, the analytical-burned area index (ABAI) (Wu et al. 2022) was obtained by solving an optimization problem consisting of multiple objective functions, each designed to maximize the average spectral differences between burned areas and other types of land cover. The optimization also included a sparsity constraint that allowed to control the resulting index complexity. In addition (Sameen and Pradhan 2016), developed a novel index for the mapping of built-up areas using a two-stage particle swarm optimization technique. First, the two most sensitive bands were determined using a multi-objective extension for particle swarm optimization. Then, these bands were incorporated into an index formulation similar to the normalized difference, but with an individual coefficient for each band variable. Finally, the coefficients were optimized again using particle swarm optimization.

However, all of these indices are limited either by the number of bands that can be considered simultaneously or by the way these bands are combined into a single index image. Therefore, this paper proposes the linear ratio index, which captures a multitude of possible index formulations mentioned in any of the previous studies in a single equation. It is introduced in detail in Section 3.2.1.

2.2. Camouflage detection

Camouflage detection is an inherently difficult task and has not been widely studied in the public literature. Most research is based on multi- and hyperspectral imagery, as it possesses the potential to break the camouflage effect. For example Brouant et al. (2023) designed and constructed a versatile ground-based multispectral imaging prototype that successfully increased the contrast between multiple camouflaged targets and their surroundings. In addition, Shen et al. (2021) presented a novel camouflage detection algorithm based on multispectral imaging using constrained energy minimization, adaptive thresholding and object region extraction. Moreover, hyperspectral imaging systems have demonstrated the ability to reveal camouflage materials in the visible Chen et al. 2017, Kriváneк, Motsch, and Bergeon 2023, Zhao et al. 2022) near-infrared (Kriváneк, Motsch, and Bergeon 2023; Zhao et al. 2022) and mid-wave infrared (Kumar and Ghosh 2017) spectrums. Although it provides limited spectral information, the detection of camouflaged targets based on traditional RGB images could also be achieved using a novel Search and Identification Network (Kumar and Singh 2024).

With the advent of UAVs and their combination with spectral imaging systems, there has also been a growing research interest in camouflage detection based on airborne spectral imagery. For instance, airborne hyperspectral imagery has been successfully combined with spectral target detection algorithms to detect dummy soldiers (Vögtli et al. 2023) and objects under camouflage materials (Gross et al. 2023). Another study (Eckel and Stütz 2024) used random forest regression on an environmental context state extracted by clustering methods to determine the most suitable bands of an airborne hyperspectral imaging system for camouflage and explosive device detection. Furthermore, Wang, Xu et al. (2024) reconstructed multispectral images from airborne RGB images to save expensive imaging equipment using a novel transformer-based network (Vaswani et al. 2017), the Spatial and Spectral Hybrid Reconstruction Transformer (S2HFormer). The reconstructed images were then processed by spectral target detection algorithms, providing detection performance comparable to processing the real multispectral images. In another study (Khangarot et al. 2023), a deep learning-based object detector was successfully investigated for camouflage detection in pure airborne RGB imagery using a limited dataset extended with data augmentation techniques.

Since camouflaged targets can be viewed as spectral anomalies, current research also focuses on their detection using anomaly detection methods (Hupel and Stütz 2022, 2024; Luo, Wang, and Deng 2023) or a combination of these with other target detection methods (Nandibewoor et al. 2024). In addition, it has been shown that the detection performance of anomaly detection methods can be improved by augmenting the original data with spectral indices (Hupel and Stütz 2022; Luo, Wang, and Deng 2023). Therefore, this study investigates the use of spectral indices specifically tailored for camouflage detection in combination with spectral anomaly detection methods to gain further insight into their utility in tactical reconnaissance scenarios.

Although deep learning has gained tremendous interest over the past decade in numerous applications, including spectral anomaly detection (Wang, Wang et al. 2022) and camouflage detection in airborne multispectral imagery (Khangarot et al. 2023; S. Wang et al. 2024), such

models typically require a significant amount of data for training and a powerful graphics processing unit for inference. Since the context of this study is tactical reconnaissance using small UAVs, where resources are generally scarce, these approaches are currently difficult to incorporate due to their resource-intensive nature. Furthermore, the multispectral data needed to study camouflage detection in tactical reconnaissance scenarios is difficult to obtain, as publicly available datasets are rare and therefore have to be generated from cumbersome data collection campaigns. As a result, the applicability of deep learning models in this study may generally be limited by the scarcity of relevant data. Considering these limitations, this study focuses on fast and efficient anomaly detection methods for camouflage detection, as presented in Section 3.3 and employed in previous research (Hupel and Stütz 2022, 2024).

3. Materials and methods

In this section, the proposed linear ratio index (LRI), its optimization problem as well as the dataset and methods used in this work are introduced. The first section, Section 3.1, presents the dataset on which all further work is based. Then, the LRI and the approach to find its optimized parameters are explained in detail in Section 3.2. Finally, Section 3.3 describes the anomaly detection methods used to evaluate the optimized indices for actual camouflage detection.

3.1. MUDCAD-X

The eXtended Multispectral Dataset for Camouflage Detection (MUDCAD-X) (Hupel and Stütz 2023) is a public dataset for training or testing camouflage detection methods based on airborne multispectral nadir imagery. It was captured at the test site of the University of the Bundeswehr Munich, which provides a diverse set of environments, such as concrete areas, graveled soil, roads and small grass- and woodland areas. Therefore, the dataset features a variety of different camouflaged targets that integrate very well into at least one of these environments. In addition, the multispectral imagery was prepared to approximate the imagery that would have been generated in a real-world reconnaissance scenario with small airborne sensor platforms, such as the MIKADO or ALADIN of the Bundeswehr.

Figure 1 shows a single sample from MUDCAD-X. Each sample contains a visual, blue, green, red, edge-infrared (EIR), near-infrared (NIR) and long-wave infrared (LWIR) band as shown in Figure 1(a–g). In addition, all samples include a label mask, as shown in Figure 1h, which marks the locations and classes for each camouflaged target with a unique color throughout the whole dataset. For example, there are five different camouflaged targets in the scene shown in Figure 1. The spectral characteristics of each band in a sample are provided in Table 1. As can be seen, each sample comes with three single-channel bands in the visual region, two single-channel bands in the near-infrared region and one single-channel band in the long-wave infrared region. Along with these single-channel bands, there is also a visual band that provides an impression of the scene as it is seen by the human eye.

All captures of MUDCAD-X were collected in two different areas of the test site on three different days in three different seasons, spring, summer and fall. On each of these days, seven capture flights were conducted per

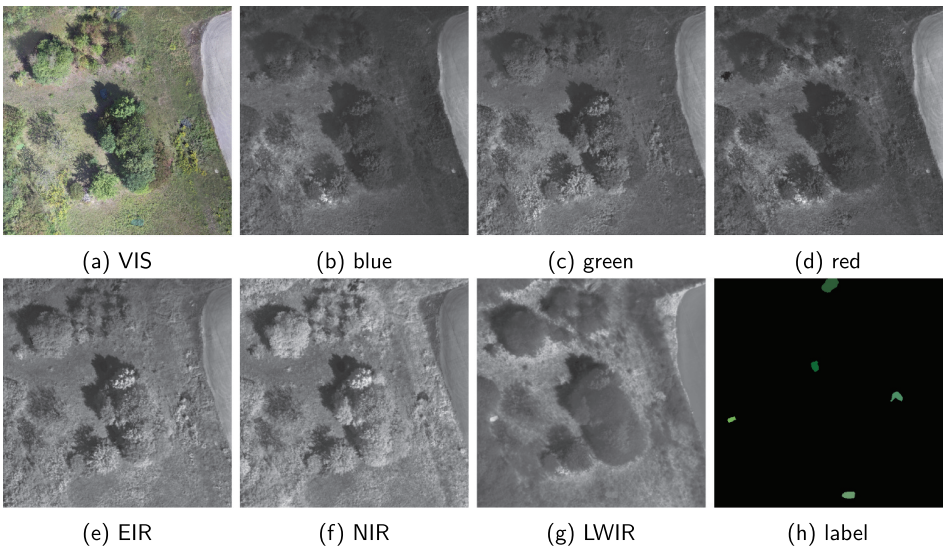

Figure 1. All bands from VIS (a) to LWIR (g) contained in a single sample from MUDCAD-X. The last image shows the label mask (h) of the scene, identifying five different camouflaged targets in this case.

Table 1. Bands included in each capture of the dataset *MUDCAD-X*.

band	center	bandwidth
visual (VIS)	–	–
blue	475 nm	32 nm
green	560 nm	27 nm
red	668 nm	14 nm
edge-infrared (EIR)	717 nm	12 nm
near-infrared (NIR)	842 nm	57 nm
long-wave infrared (LWIR)	10.5 µm	6 µm

area. After each capture flight, all camouflaged targets were moved to a new but similar location in the same area. This procedure resulted in a total of 853 samples, each containing at least one of the 10 different camouflaged target classes. The camouflaged target classes together with other basic properties of the dataset are listed in Table 2. Note that the number of targets in parentheses shows the actual number of different targets, as there were always two people in different uniforms for each person-in-uniform class. However, in the annotation process, there was no way to distinguish between these two people, so they were merged into a single class. As shown in the last two rows, all samples have a resolution of 512 by 512 and a ground sampling distance (GSD) of 10. More details about MUDCAD-X (Hupel and Stütz 2023) can be found in its original publication.

3.2. Linear ratio index

This section presents the proposed linear ratio index (LRI) and its optimization strategy. First, Section 3.2.1 introduces the formulation of the LRI, which is a generalization of the normalized difference, ratio and difference indices. Then, the optimization of the LRI and the complexity reduction strategies are explained in Section 3.2.2 and Section 3.2.3, respectively. Finally, Section 3.2.4 describes how the LRI is optimized for camouflage detection.

3.2.1. Formulation

Generalizing the normalized difference, ratio and difference indices, which are calculated as already shown in Equation 1-Equation 3, this paper introduces the linear ratio index. It is computed as shown in Equation 4.

$$I_{lr}(C, W^{\alpha}, W^{\beta}, b^{\alpha}, b^{\beta}) = \frac{\sum_{i=1}^{n} w_i^{\alpha} c_i + b^{\alpha}}{\sum_{i=1}^{n} w_i^{\beta} c_i + b^{\beta}} \tag{4}$$

Here, $C = \{c_1, \ldots, c_n \mid c_i \in \mathbb{R}^{m \times l} \; \forall i \in \{1, \ldots, n\}\}$ is the set of n available bands, each with a height of m and a width of l. $W^{\alpha} = \{w_1^{\alpha}, \ldots, w_n^{\alpha} \mid w_i^{\alpha} \in \mathbb{R} \; \forall i \in \{1, \ldots n\}\}$ is the set of numerator band weights, $W^{\beta} = \{w_1^{\beta}, \ldots, w_n^{\beta} \mid w_i^{\beta} \in \mathbb{R} \; \forall i \in \{1, \ldots, n\}\}$ is the set of denominator band weights, $b^{\alpha} \in \mathbb{R}$ is the numerator bias and $b^{\beta} \in \mathbb{R}$ is the denominator bias. In essence, all bands are multiplied by an individual weight and added together. Then, an additional bias is added to the sum of the weighted bands. This is done twice, once in the numerator and once in the denominator, using different weights and biases, leaving a single band in the numerator and in the denominator. Both the numerator and the denominator can be interpreted as linear or affine functions of C, which is why the index is referred to as linear ratio index. As with the traditional normalized difference, ratio and difference indices, the final division of the two bands in the numerator and denominator is performed element by element. Optimizing the LRI for a specific application does not require testing band combinations, as is typical with other indices, but rather finding the optimal values for W^{α}, W^{β}, b^{α} and b^{β}. Also, while normalized difference, ratio and difference indices are restricted to a range of values determined by the input bands, the LRI introduces an arbitrary range of values determined not only by the pixel values but also by the weights and biases. Furthermore, the normalized difference, ratio and

Table 2. Basic properties of the dataset *MUDCAD-X*.

Property	Value
# samples	853
# targets	10 (12)
targets	anthracite fleece, artificial turf, artificial hedge, gray tarp, green tarp, green 2D camouflage net, green 3D camouflage net, gray 3D camouflage net, yellow 3D camouflage net, person in green uniform, person in yellow uniform
seasons	spring, summer, fall
bands	VIS, blue, green, red, EIR, NIR, LWIR
perspective	nadir
resolution	512px × 512px
GSD	$10\,\frac{\text{cm}}{\text{px}}$

differences indices can be considered special cases of the LRI. For example, if there are $n = 3$ bands, the normalized difference between the first band and the third band can be expressed with $w_1^\alpha = w_1^\beta = w_3^\beta = 1$, $w_3^\alpha = -1$, $w_2^\alpha = w_2^\beta = 0$ and $b^\alpha = b^\beta = 0$. Likewise, the ratio between the first band and the third band can be expressed with $w_1^\alpha = w_3^\beta = 1$, $w_2^\alpha = w_3^\alpha = w_1^\beta = w_2^\beta = 0$ and $b^\alpha = b^\beta = 0$. The difference index between the first band and the third band can be formulated equally simple with $w_1^\alpha = 1$, $w_3^\alpha = -1$, $w_2^\alpha = w_1^\beta = w_2^\beta = w_3^\beta = 0$, $b^\alpha = 0$ and $b^\beta = 1$.

3.2.2. Optimization

As already mentioned in Section 3.2.1, the LRI is optimized by adjusting its weights $W^\alpha = \{w_1^\alpha, \ldots, w_n^\alpha\}$ and $W^\beta = \{w_1^\beta, \ldots, w_n^\beta\}$ and biases b^α and b^β. Given a single capture $C = \{c_1, \ldots, c_n\}$ with n bands and an objective function f that takes and index image and returns a cost, the LRI optimization problem can be written as in Equation 5.

$$\underset{W^\alpha, W^\beta, b^\alpha, b^\beta}{\arg\min}\ f\left(I_{lr}\left(C, W^\alpha, W^\beta, b^\alpha, b^\beta\right) = \frac{\sum_{i=1}^{n} w_i^\alpha c_i + b^\alpha}{\sum_{i=1}^{n} w_i^\beta c_i + b^\beta}\right) \tag{5}$$

For a given dataset $D = \{d_1, \ldots, d_N | d_i = \{c_{i,1}, \ldots, c_{i,n} | c_{i,j} \in \mathbb{R}^{m \times l} \forall j \in \{1, \ldots, n\}\} \forall i \in \{1, \ldots, N\}\}$ with N captures, each consisting of n bands with a resolution of l by m pixels, the cost determined by the objective function f over all captures must be reduced to a single value. For the sake of simplicity, the reduction is performed by summing over all costs per capture, as shown in Equation 6.

$$\underset{W^\alpha, W^\beta, b^\alpha, b^\beta}{\arg\min} \sum_{i=1}^{N} f\left(I_{lr}\left(d_i, W^\alpha, W^\beta, b^\alpha, b^\beta\right)\right) \tag{6}$$

For the optimization of the problem described in Equation 6, this work employs an adaptive version of differential evolution (Storn and Price 1997) with modifications for choosing the parameters (Wang et al. 2014) and updating the population (Spector and Klein 2006). The concrete implementation details are given by the *adaptive_de_rand_1_bin_radiuslimited* optimizer provided by the BlackBoxOptim.jl package (Feldt and Stukalov 2018) of the Julia Programming Language (Bezanson et al. 2017). Differential evolution is a global optimization technique that does not rely on a differentiable objective function. Both of these properties are desirable in the specific application of this work. Specifically, the use of a global optimizer eliminates batching strategies as required for global optimization using gradient descent. Thus, the optimization problem defined in Equation 6 can be optimized directly without any further modifications. In addition, this work introduces complexity reduction strategies for the optimization of the LRI that lead to non-differentiable behavior for any objective function f, making differential evolution an excellent choice for the optimizer. These complexity reduction strategies are explained in the following section.

3.2.3. Optimization with complexity reduction

In principle, the LRI is designed to make use of all the available bands for the resulting index image. However, there may be use cases where this behavior is not desirable. For example, when computational resources are extremely limited or when a multispectral imaging system needs to be designed with a minimal number of bands to cut production costs. It may also be desirable just for the sake of reducing dependencies and complexity, since most spectral indices, including normalized difference, ratio and difference indices, depend on only two to three bands. Whatever the reason, in order to support a reduced number of bands during the optimization of the LRI, a proper strategy must be found that works well with the optimizer.

Experiments have shown that the adaptive differential evolution optimizer used in this work is very robust against the following modifications of W^α and W^β, which limit the number of utilized bands to a predefined constant $k \in \mathbb{N}$. The first step is to calculate the sum of the absolute weights for each band, as shown in the equation below.

$$W^\gamma = \left\{w_1^\gamma, \ldots, w_n^\gamma \mid w_i^\gamma = \left|w_i^\alpha\right| + \left|w_i^\beta\right| \ \forall i \in \{1, \ldots, n\}\right\} \tag{7}$$

Note that w_i^α and w_i^β are both weights of the i th band, so their absolute sum can be interpreted as a measure of the contribution of the i th band to the resulting index image. In the next step, the function ϕ maps the positions of W^γ so that it returns the position of the largest element in W^γ first and the smallest last. As shown by the following equations, the set Φ_k containing the indices of the first largest k elements in W^γ can be obtained with the function ϕ.

$$\phi : \{1, \ldots, n\} \mapsto \{1, \ldots, n\} \mid w_{\phi(1)}^\gamma \geq \ldots \geq w_{\phi(n)}^\gamma \tag{8}$$

$$\Phi_k = \{\phi(1), \ldots, \phi(k)\} \tag{9}$$

Finally, all weights whose positions are not contained in Φ_k are set to zero, leading to the new weight parameters $W^{\alpha'}$ and $W^{\beta'}$.

$$w_i^{\alpha'} = \begin{cases} 0, & \text{otherwise} \\ w_i^{\alpha}, & \text{if } i \in \Phi_k \end{cases} \quad \forall i \in \{1, \ldots, n\} \tag{10}$$

$$w_i^{\beta'} = \begin{cases} 0, & \text{otherwise} \\ w_i^{\beta}, & \text{if } i \in \Phi_k \end{cases} \quad \forall i \in \{1, \ldots, n\} \tag{11}$$

$$W^{\alpha'} = \left\{ w_1^{\alpha'}, \ldots, w_n^{\alpha'} \right\} \tag{12}$$

$$W^{\beta'} = \left\{ w_1^{\beta'}, \ldots, w_n^{\beta'} \right\} \tag{13}$$

This way, only the largest k absolute weights are retained, since they contribute the most to the resulting index image. All others are set to zero, effectively nullifying the influence of the bands by which they are multiplied. The bias parameters, b^{α} and b^{β}, are not affected by this technique, as they add constant but minimal complexity to the LRI regardless of the number of bands available. It should be noted that these modifications are made each time the objective function is called, so from the optimizer's point of view, this behavior is part of the objective function. Therefore, the optimizer still tweaks the original weights W^{α} and W^{β}, which are then modified according to the equations above and passed to the objective function. This allows the optimizer to determine which bands to ignore and which bands to keep because they are important for minimizing the objective function. The following equation shows the optimization problem of Equation 5, adjusted with the weight parameters $W^{\alpha'}$ and $W^{\beta'}$.

$$\underset{W^{\alpha}, W^{\beta}, b^{\alpha}, b^{\beta}}{\arg\min} \; f\left(I_{lr}\left(C, W^{\alpha'}, W^{\beta'}, b^{\alpha}, b^{\beta} \right) \right) \tag{14}$$

Another property often found in spectral index definitions is the low precision of the weights and other parameters. For example, all normalized difference, ratio and difference indices have weights only in $\{-1, 1\}$. This may affect the accuracy of the spectral index, but it also makes it easier to remember and therefore more transferable and practical. In order to incorporate this property into the optimization of the LRI, the function r is introduced in the following Equation.

$$r(x, d) = \frac{\lfloor x \times 10^d + 0.5 \rfloor}{10^d} \tag{15}$$

It takes a value $x \in \mathbb{R}$ and a number $d \in \mathbb{N}$ and returns the value x rounded to the last d digits. For example, $r(1.8351, 2) = 1.84$. With this function, the weights and biases of the LRI can be kept at a certain precision defined by d. The following equations show how the weights $W^{\alpha'}$ and $W^{\beta'}$ and biases b^{α} and b^{β} are modified in order to obtain the weights W^{α^*} and W^{β^*} and biases b^{α^*} and $b^{\beta}*$.

$$W^{\alpha^*} = \left\{ w_1^{\alpha^*}, \ldots, w_n^{\alpha^*} \mid w_i^{\alpha^*} = r\left(w_i^{\alpha'}, d \right) \; \forall i \in \{1, \ldots, n\} \right\} \tag{16}$$

$$W^{\beta^*} = \left\{ w_1^{\beta^*}, \ldots, w_n^{\beta^*} \mid w_i^{\beta^*} = r\left(w_i^{\beta'}, d \right) \; \forall i \in \{1, \ldots, n\} \right\} \tag{17}$$

$$b^{\alpha^*} = r(b^{\alpha}, d) \tag{18}$$

$$b^{\beta^*} = r(b^{\beta}, d) \tag{19}$$

Again, these modifications are made before the parameters are passed to the objective function, so the optimizer still adjusts the original parameters W^{α}, W^{β}, b^{α} ,and b^{β}. At this point, the versed reader may notice that the parameters could grow arbitrarily large instead (e.g. 1.8351 becomes 183.51), since the index is, after all, a ratio where only relative difference matters. However, the optimizer of the BlackBoxOptim.jl package requires upper and lower bounds on all parameters, which mitigates this issue. Applying the two complexity reduction strategies described above, the optimization problem in Equation 5:

$$\underset{W^{\alpha}, W^{\beta}, b^{\alpha}, b^{\beta}}{\arg\min} \; f\left(I_{lr}\left(C, W^{\alpha^*}, W^{\beta^*}, b^{\alpha^*}, b^{\beta^*} \right) \right) \tag{20}$$

3.2.4. Optimization for camouflage detection

In order to optimize the LRI for any application, an appropriate objective function must first be defined. Since the index is intended to support the visibility and detectability of camouflaged targets in this study, the Target Visibility Index (TVI) (Hupel and Stütz 2023) is employed as the optimization criterion, which has been introduced and already applied in previous research (Hupel and Stütz 2023, 2024). The TVI is a conceptually and computationally efficient metric that describes the discriminability of an object from its background in a single-channel nadir image. Although it is an experimental metric, the simple and efficient nature of the TVI adds minimal computational overhead to the optimization problem. Alternative optimization criteria, such as an evaluation of detection maps or saliency maps, are much more expensive to compute as they require the use of sophisticated algorithms. Conclusively, they make the optimization problem much more complex than the TVI. In addition, using the TVI as an optimization criterion provides even further valuable insight into this metric and its utility in camouflage detection.

Given an image channel $c_x \in [0,1]^{m\times l}$ and an annotation mask $a_x \in \{0,1\}^{m\times l}$ denoting the camouflaged target, the TVI is calculated according to the following formula.

$$\mathrm{TVI} = (\mu_T - \mu_B)^2(1-2\sigma_T)^2(1-2\sigma_B)^2 \tag{21}$$

$$\mu_T = \frac{\sum_{i=1}^{m}\sum_{j=1}^{l} c_{xij}\cdot a_{xij}}{\sum_{i=1}^{m}\sum_{j=1}^{l} a_{xij}} \tag{22}$$

$$\sigma_T = \sqrt{\frac{\sum_{i=1}^{m}\sum_{j=1}^{l}\left(c_{xij}-\mu_T\right)^2\cdot a_{xij}}{\sum_{i=1}^{m}\sum_{j=1}^{l} a_{xij}}} \tag{23}$$

$$\mu_B = \frac{\sum_{i=1}^{m}\sum_{j=1}^{l} c_{xij}\cdot\left(1-a_{xij}\right)}{\sum_{i=1}^{m}\sum_{j=1}^{l}\left(1-a_{xij}\right)} \tag{24}$$

$$\sigma_B = \sqrt{\frac{\sum_{i=1}^{m}\sum_{j=1}^{l}\left(c_{xij}-\mu_T\right)^2\cdot\left(1-a_{xij}\right)}{\sum_{i=1}^{m}\sum_{j=1}^{l}\left(1-a_{xij}\right)}} \tag{25}$$

Here, μ_T and μ_B are the average values over all pixels belonging to the target and all pixels belonging to the background, respectively. Similarly, σ_T is the standard deviation over all target pixels and σ_B is the standard deviation over all background pixels. The TVI ranges from 0 to 1, where 0 means no visibility and 1 means maximum visibility of the target. However, these values do not necessarily correspond to human perception and should rather be interpreted as a simplified measure of contrast between target and background. Figure 2 shows the visual (VIS), blue, edge-infrared (EIR) and near-infrared (NIR) channels of a multispectral capture for which the TVIs were calculated. As can be observed, the TVI is much lower in the VIS and blue channels compared to the EIR and NIR channels. This is due to the greater difference between the target pixels and the background pixels in the infrared channels, where the target is much darker than the rest of the scene. As a result, the target is much easier to see in the infrared channels than in the visual channels. Therefore, the LRI optimized for maximum TVI should provide even higher discriminability between camouflaged targets and their backgrounds than any of the original bands. In addition, since most pixels in an image are background, a high TVI should also make the target appear more anomalous and thus more susceptible to anomaly detection algorithms. Readers interested in learning more about the TVI are encouraged to read the original publication (Hupel and Stütz 2023).

With the TVI as the criterion to be maximized, the optimization problem with complexity reduction of the LRI in Equation 20 can be adapted as follows. Let f_{TVI} be a function that, given an image channel and an annotation mask, returns the TVI computed as in Equation 21 and let $A = \{a_1,\ldots,a_M \mid a_i \in \{0,1\}^{m\times l}\ \forall i \in \{1,\ldots,M\}\}$ be a set of annotation masks, the optimization problem can be formulated as shown in the following equation.

$$\underset{W^{\alpha},W^{\beta},b^{\alpha},b^{\beta}}{\arg\min} -\sum_{i=1}^{M} f_{\mathrm{TVI}}\left(f_{\mathrm{norm}}\left(I_{lr}\left(C, W^{\alpha^*}, W^{\beta^*}, b^{\alpha^*}, b^{\beta^*}\right)\right), a_i\right) \tag{26}$$

Note that the LRI can produce a range of values beyond the definition of the TVI. Thus, the function f_{norm} normalizes the resulting index image between 0 and 1 before passing it to the function f_{TVI}. This should not cause any complications, since all that matters is the relative difference between camouflaged targets and their backgrounds, which is the property measured by the TVI. The absolute values of the resulting index image are of no concern, as they are not intended to relate to any physical property, such as plant or crop health. In addition, the TVI is negated because the optimizer minimizes the objective function by default.

3.3. Anomaly detection methods

This section briefly introduces the implemented spectral anomaly detection methods, RX (Reed and Yu 1990), LRX (Reed and Yu 1990), LPD (Hupel and Stütz 2022) and AED (Kang et al. 2017), which are used to automatically detect camouflaged targets. Using spectral anomaly

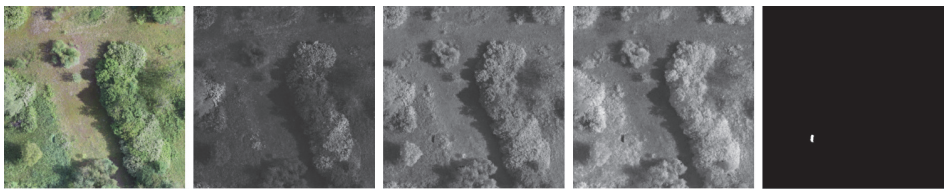

(a) VIS (0.1×10^{-3}) (b) blue (0.2×10^{-3}) (c) EIR (9.1×10^{-3}) (d) NIR (13.6×10^{-3}) (e) ideal (1.0)

Figure 2. Demonstration of the Target Visibility Index (TVI). Resulting values for each channel are shown in parentheses. Note that the visual channel has been converted to a gray-level image prior to the calculation of the TVI.

detection for camouflage detection is a common approach, as camouflaged targets are expected to show anomalous spectral behavior relative to their surroundings (Hupel and Stütz 2022, 2024; Luo, Wang, and Deng 2023; Nandibewoor et al. 2024). In addition, these methods are comparatively fast and have low resource requirements, as shown in (Hupel and Stütz 2022). This is particularly important, since this paper focuses on reconnaissance applications using small tactical drones, which by design have limited computational capabilities. Detailed information about the algorithms can be found in their original publications.

3.3.1. RX and LRX detectors

The Reed-Xiaoli (RX) (Reed and Yu 1990) detector is one of the most classic anomaly detection algorithms. It works on the hypothesis that the target pixels and the background pixels originate from two different normal distributions. Following this hypothesis, the Mahalanobis distance between a pixel and the background distribution becomes the measure of the pixel's degree of abnormality. The background distribution in a single capture with n channels is estimated by calculating its average pixel value $\mu \in \mathbb{R}^n$ and the covariance matrix $\Sigma \in \mathbb{R}^{n\times n} \mid \Sigma \geq 0,\ \Sigma = \Sigma^T$. As a result, the anomaly score $d_M \in \mathbb{R}$ of a pixel under test $x \in \mathbb{R}^n$ is calculated as shown in Equation 27.

$$d_M = \sqrt{(x-\mu)^T \Sigma^{-1} (x-\mu)} \tag{27}$$

As with the RX detector, the local RX (LRX) detector determines the abnormality score of a pixel by calculating its Mahalanobis distance to the background distribution. However, in contrast to the RX detector, the background estimation is only based on a local neighborhood of the pixel under test. The neighborhood is determined by a sliding dual-window, which consists of an inner window and an outer window centered around the pixel under test. Only pixels that are inside the outer window but outside the inner window are considered for the background estimation. This way, the background statistics are closer to the actual vicinity of the pixel under test without too much contamination of target pixels, which is suppressed by the inner window. In general, the LRX detector achieves higher detection rates traded for the introduction of additional parameters, $w_i \in \mathbb{N}$ and $w_o \in \mathbb{N} \mid w_o > w_i$, defining the inner and outer window sizes, respectively. Figure 3 shows an example of a dual window drawn into the visual band of a scene. The pixel under test is marked in blue, while the inner and outer windows are marked in red and green, respectively. As can be observed, the background estimation would only be based on the pixels immediate surroundings, ignoring pixels much further away or included in the inner window.

3.3.2. LPD detector

Local point density (LPD) (Hupel and Stütz 2022) is an anomaly detection method that relies on the soft-density of pixels (Tu, Yang, Li, et al. 2020; Tu, Yang, Zhou, et al. 2020). Unlike the RX and LRX detectors, it makes no assumptions about the distribution of background and target pixels. Instead, it estimates the density of each pixel based on its Euclidean distance from other pixels. It should be noted that this distance is based on the values of the pixels, not their locations. The higher the density of a pixel, the more pixels have similar values and are not considered anomalous. Similarly, the lower the density of a pixel, the fewer the number of pixels with similar values, and therefore the higher the abnormality of that pixel.

Like the LRX detector, the density estimation considers only a local neighborhood defined by an inner and outer window centered on the pixel under test, as illustrated in Figure 3. Consequently, the density of each pixel in a capture can be calculated as defined in Equation 28. Where $\rho \in \mathbb{R}$ is the resulting density for a pixel value $x \in \mathbb{R}^n$ with n bands and its neighborhood $M = \{m_i, \ldots, m_{s_M} \mid m_i \in \mathbb{R}^n\}$ consisting of $s_M \in \mathbb{N}$ pixels. $d_c \in \mathbb{R}$ is the cut-off distance, which is equal to the average Euclidean distance between all pixels in the

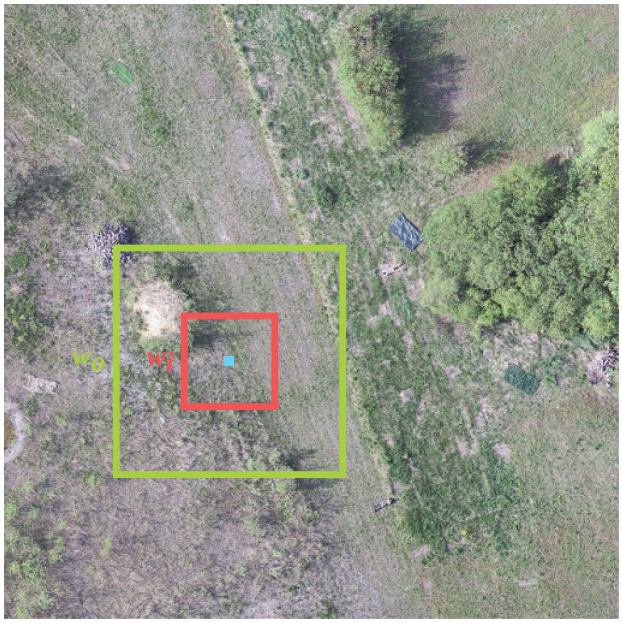

Figure 3. The dual window defining the neighborhood of pixels for the pixel under test (blue). The neighborhood consists only of pixels that are inside the outer window (green) but outside the inner window (red).

neighborhood and the average pixel value $\mu_M \in \mathbb{R}^n$ of all pixels in the neighborhood.

$$\rho = \frac{1}{s_M} \sum_{i=1}^{s_M} \exp\left(- \frac{\|x - m_i\|_2^2}{d_c^2} \right) \tag{28}$$

$$d_c = \frac{1}{s_M} \sum_{i=1}^{s_M} \|m_i - \mu_M\|_2^2 \tag{29}$$

$$\mu_M = \frac{1}{s_M} \sum_{i=1}^{s_M} m_i \tag{30}$$

Once the densities for all pixels are calculated, they are negated to create an anomaly map where higher values indicate more anomalous pixels and lower values indicate less anomalous pixels.

3.3.3. AED detector

Anomaly detection using Attribute and Edge-Preserving Filters (AED) (Kang et al. 2017) builds on the idea that anomalous objects are relatively small and either much brighter or darker than their surroundings in many bands. This idea is translated into an algorithm that generates morphological attribute profiles for each band, which are fused into a final detection map.

First, each band undergoes two morphological area filtering operations, which remove all bright and dark-connected components smaller than a predefined size $\kappa \in \mathbb{N}$. Given a capture $\{c_1, \ldots, c_n | c_i \in [0,1]^{m \times l} \forall i \in \{1, \ldots, n\}\}$ with n bands of height m and width l, this leaves a set area profile tuples $\{(\gamma_1^\kappa, \phi_1^\kappa), \ldots, (\gamma_n^\kappa, \phi_n^\kappa) \mid \gamma_i^\kappa, \phi_i^\kappa \in [0,1]^{m \times l} \forall i \in \{1, \ldots, n\}\}$, where γ_i^κ , and ϕ_i^κ are the area profiles of the ith band where all bright and dark connected components smaller than κ have been removed. Following the original idea that anomalous objects are small and either bright or dark, most anomalous objects should have been removed in at least one of these profiles. The detection map of the ith band is given by the difference $\psi_i^\kappa = \phi_i^\kappa - \gamma_i^\kappa$ of the profiles, highlighting the regions removed after at least one of the filtering operations. These detection maps are then combined into a single detection map using a special binary fusion operation, which includes additional morphological and thresholding operations. Finally, the detection map is post-processed with an edge-preserving smoothing filter. For the sake of simplicity, this is implemented by a bilateral filter in this work. The original publication used a domain transform recursive filter (Gastal and Oliveira 2011), which is computationally more efficient than a bilateral filter.

4. Experiments and results

This section presents the experiments conducted with the LRI, the obtained results and their evaluation. First, the training and test datasets used for optimization and evaluation are introduced in Section 4.1. Then, Section 4.2 describes the optimization configurations as well as the resulting LRI parameters. Finally, the obtained LRIs are evaluated for their impact on the visibility and detectability camouflaged targets in Section 4.3.

4.1. Training and test datasets

In order to optimize and evaluate the LRI for camouflage detection using the approach presented in Section 3.2, the dataset *MUDCAD-X* introduced in Section 3.1 must first be split into training and test datasets. For a ratio close to 80:20, MUDCAD-X is split by selecting all data of the 4th capture flight of each day as test data and the remaining data as training data, as done in (Hupel and Stütz 2024). After the split, the test dataset contains 129 samples and the training dataset contains the remaining 724 samples.

As a pre-processing step, the datasets are globally min-max normalized on a band-wise basis, resulting in an equal and comparable range of values for each band. This step is necessary because there are no reflectance maps available, as is usually the case with multispectral datasets. The normalization statistics are obtained only from the training dataset, but are applied to both training and test datasets. Once the minimum and maximum values for each band in the training dataset have been obtained, both datasets can be normalized. Given a pixel value $x \in \mathbb{R}$ of the ith band with a minimum value of $\min_i \in \mathbb{R}$ and maximum value of $\max_i \in \mathbb{R}$, the normalized pixel value $\hat{x} \in [0,1]$ is determined as shown below.

$$\hat{x} = \frac{x - \min_i}{\max_i - \min_i} \tag{31}$$

This leads to a range of values from 0 to 1 for each band in the training dataset. Although the bands of the raw captures already fall within this range, they do not necessarily use it to the full extent, but they do after normalization. Since the minimum and maximum values of the training dataset may not match those of the test dataset, their range of values may also be slightly different after normalization. However, this is unavoidable because the statistics of the test dataset should not influence the optimization procedure, as it is used only for evaluation. The minimum and maximum values for each band of the training dataset can be found in Table 3. As can be observed, the normalization has no effect on the blue, green and red bands. For the EIR and NIR

Table 3. Minimum and maximum values of each band over the training dataset, excluding the visual band, which is not used for index optimization. These values are used for min-max normalization of the training and test datasets.

	blue	green	red	EIR	NIR	LWIR
min	0	0	0	0.0196	0.0039	0.0118
max	1	1	1	1.0000	0.9961	0.9412

bands, the normalization effect is rather small, while it is comparatively large for the LWIR band. It should also be noted that there are no statistics for the visual band as it is not used for the optimization of the LRI.

In addition to normalizing all bands in both datasets, they are also scaled down to a resolution of 256 by 256, which corresponds to a GSD of 20 cm per pixel. This greatly reduces the computational cost of the optimization procedure while keeping the resolution high enough to detect all camouflaged targets, as explained in (Hupel and Stütz 2022). Furthermore, all annotations with an area of five pixels or less are removed from both datasets. This ensures numerical stability for any statistical calculation without removing valuable information, as it is a very small area. Finally, both datasets consist of band-wise normalized samples with a resolution of 256 by 256 pixels and no annotations smaller than 5 pixels.

4.2. Optimization

With the training dataset defined in Section 4.1, the LRI can be optimized as explained in Section 3.2.4. However, as with most datasets, there are class imbalances that need to be considered. For example, in MUDCAD-X, and therefore in the training dataset, there are many more annotations for persons in green uniforms than for the gray tarp (Hupel and Stütz 2023). This is not surprising, since there are many more environments suitable for a person in a green uniform than for a gray tarp. However, an optimized LRI should work equally well for all targets, regardless of their prevalence in the data. For this reason, the classes in the data are balanced in the optimization using the training dataset as follows. First, the equalization weight for each camouflaged target class in the training dataset is obtained according to the following equation.

$$H = \{h_1, \ldots, h_{n_H}\} \mid h_i = \frac{P}{n_H \cdot P^i} \ \forall i \in \{1, \ldots, n_H\} \quad (32)$$

Where H is the set of equalization weights, P is the total number of annotations in the training dataset, P^i is the number of annotations in the training dataset that belong to the ith camouflaged target class and n_H is the number of camouflaged target classes in the training dataset. The resulting equalization weights are listed in Table 4. As can be seen, the number of annotations for the anthracite fleece is relatively small, hence the large weight. Likewise, the weights for the persons in uniforms are relatively small because they make up a large portion of the total annotations in the training dataset. Given a set of $M \in \mathbb{N}$ tuples $A' = \{(t_1, a_1, h_{t_1}), \ldots, (t_M, a_M, h_{t_M}) \mid a_i \in \{0,1\}^{m \times l},\ t_i \in \{1, \ldots, n_H\}\ \forall i \in \{1, \ldots, M\}\}$ containing the camouflaged target class membership t_i together with the annotation mask a_i and the corresponding equalization weight h_{t_i}, the optimization problem in Equation 26 can be reformulated as follows.

$$\begin{aligned} \underset{W^\alpha, W^\beta, b^\alpha, b^\beta}{\arg\min} \ & -\sum_{i=1}^{M} h_{t_i} \\ & \cdot f_{\mathrm{TVI}}\left(f_{\mathrm{norm}}\left(I_{lr}\left(C, W^{\alpha^*}, W^{\beta^*}, b^{\alpha^*}, b^{\beta^*}\right)\right), a_i\right) \end{aligned} \quad (33)$$

By summing over all samples in the training dataset as defined in Equation 6, the LRI can finally be optimized, taking into account all data in the dataset in a single function evaluation and the imbalance of the camouflaged target classes.

For a detailed analysis of the proposed index and its optimization strategy, two different LRIs are optimized. One that uses all six single-channel bands, called LRI_6, and one that uses only two single-channel bands, called LRI_2. Both make use of the complexity reduction strategies introduced in Section 3.2.3, as their parameters are rounded to the last two digits, allowing the optimizer to directly consider their human readability. In addition, LRI_2 is optimized with only two bands, which are also determined by the optimizer. As mentioned in Section 3.2, the optimizer requires search range limits, so the allowed minimum and maximum values of each parameter of the LRI are set to −10 and 10, respectively. Also, the population size is set to 100 and the termination criterion is set to a maximum of 75,000 function evaluations. All other parameters of the optimizer are left at their default values defined by the BlackBoxOptim.jl package (Feldt and Stukalov 2018). The optimization runs of both indices are shown in Figure 4. As can be seen, both indices converge after about 40,000 iterations, which took about 24 hours on an AMD Ryzen 9 3950X with 128 memory using Julia 1.10 on Ubuntu 22.04. The total number of iterations took about 48 hours, leading to a final average TVI of 0.04045 for the LRI_6 and 0.03120 for the LRI_2.

In order to compare the LRI and the proposed optimization approach with common optimization techniques that test all possible band combinations of normalized difference and ratio indices, an optimal normalized difference and ratio are also determined. These are optimized

Table 4. Equalization weights for each camouflaged target class in the training dataset.

i	name	weight (h_i)
1	anthracite fleece	5.14
2	artificial turf	0.99
3	artificial hedge	0.99
4	gray tarp	3.11
5	green tarp	0.94
6	green 2D camouflage net	0.92
7	green 3D camouflage net	0.93
8	gray 3D camouflage net	1.46
9	yellow 3D camouflage net	1.55
10	person in green uniform	0.49
11	person in yellow uniform	0.53

using the same objective function, but with an optimizer that simply tests all possible input band combinations for the best result. The resulting normalized difference and ratio indices are called LRI_{nd} and LRI_r, respectively. Compared to the optimization times of the LRI_6 and LRI_2, which were about 2 days, the optimization times of the LRI_{nd} and LRI_r are a matter of seconds. Table 5 shows the final optimized indices, ready to be incorporated into any camouflage detection process. It is worth noting that the LRI_{nd} and the LRI_r are both based on the same bands, the NIR and LWIR bands. In contrast, the LRI_2, which is also based on two bands, relies on the green and EIR bands.

4.3. Evaluation

In this section, the determined optimized LRIs in Section 4.1 are first evaluated for impact on the visibility of camouflaged targets in Section 4.3.1. Then, they are evaluated for their impact on actual camouflage detection using spectral anomaly detection methods in Section 4.3.2. Both evaluations are based on the test dataset.

4.3.1. Target visibility

Since the LRIs were optimized with respect to the TVI on the training dataset, the evaluation is based on the test dataset, which was not used to optimize the LRIs. An exemplary impression of the optimized indices is given in Figure 5, which shows a scene arbitrarily selected from the test dataset. Specifically, Figure 5(a–h) show the raw bands along with the label mask and Figure 5(i–h) show the optimized LRIs. As can be observed, the LRI_6 produces much higher contrast between many camouflaged targets and the background than any of the raw bands or the other optimized indices. Although all optimized indices except the LRI_r appear to achieve significantly higher contrast than the raw bands. This applies especially to the targets marked with red circles in Figure 5(h), which are particularly prominent in the index images in Figure 5(i–k) compared to the raw bands in Figure 5(a–g). However, it is also worth noting that some camouflaged targets already have high visibility in the raw bands, which does not seem to be further improved by the optimized indices. These targets are marked with blue circles in Figure 5. In addition, there seem to be camouflaged targets that have higher contrast in the raw bands than in any of the optimized LRIs. For example, the targets marked with purple circles Figure 5h appear to have the highest visibility in the red band shown in Figure 5d. Nevertheless, since there are too

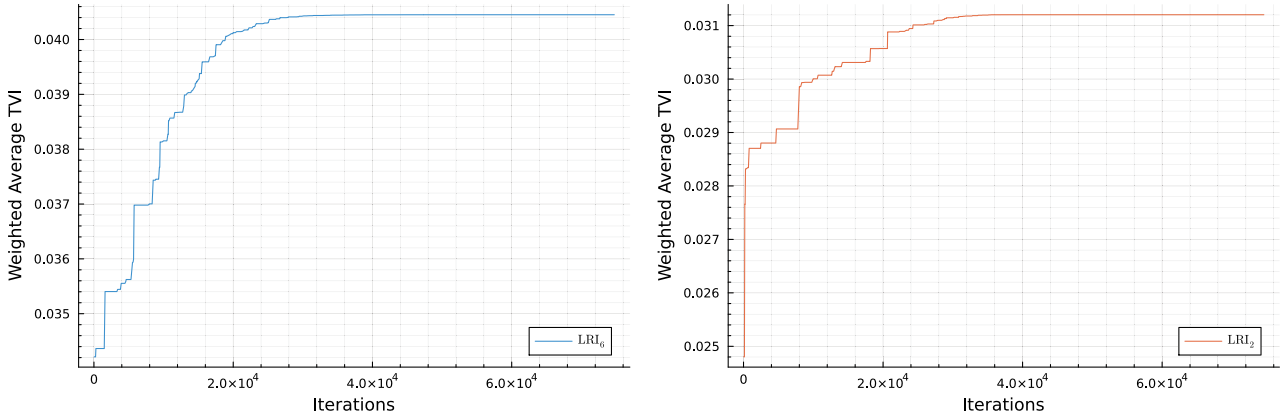

Figure 4. Optimization runs of LRI_2 and LRI_6. The graphs show the weighted average TVI over the training dataset at each iteration.

Table 5. Final parameters of the optimized LRIs. LRI_6 and LRI_2 were optimized using six and two bands, respectively. LRI_{nd} and LRI_r are the optimized normalized difference and ratio found by testing all band combinations.

	w^{α}_{blue}	w^{α}_{green}	w^{α}_{red}	w^{α}_{EIR}	w^{α}_{NIR}	w^{β}_{LWIR}	b^{α}	w^{β}_{blue}	w^{β}_{green}	w^{β}_{red}	w^{α}_{EIR}	w^{β}_{NIR}	w^{α}_{LWIR}	b^{β}
LRI_6	1.80	9.57	−4.15	−8.33	−4.59	3.41	0.69	−0.06	2.67	−3.69	7.46	9.86	1.79	1.34
LRI_2		−9.74		8.97			−2.63				9.60			
LRI_{nd}					1.00	−1.00			2.10			1.00	1.00	
LRI_r						1.00						1.00		

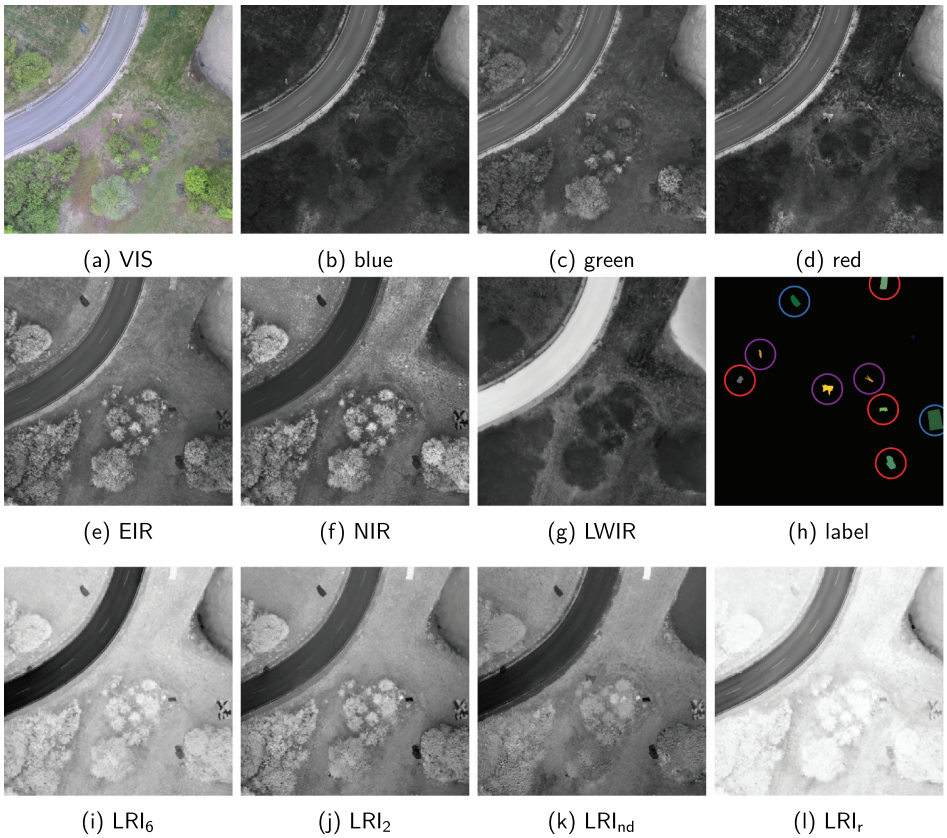

Figure 5. The optimized LRIs for a sample from the test dataset. The raw bands are shown from VIS (a) to LWIR (g) along with the corresponding label mask (h) The last row shows the LRI_6 (i), LRI_2 (j), LRI_{nd} (k) and LRI_r (l). The circle markers in the label mask (h) are explained in the text.

many captures in the test dataset for an overall qualitative evaluation, this is further analyzed in the following statistical evaluation. For that purpose, the TVI is calculated separately for each LRI and for each annotation in the test dataset. The resulting TVIs are then grouped and averaged for each camouflaged target class. This provides insight into the performance of each index for each individual camouflaged target class, which may not be the same for all target classes. In addition, the average TVI over all camouflaged targets is calculated, indicating the overall performance of each optimized LRI. Note that there is no need for equalization weights as in the training procedure, since the TVI is calculated separately for each camouflaged target class.

Table 6 shows the resulting TVIs for each of the optimized LRIs and the raw single-channel bands over the test dataset. Each row corresponds to a single camouflaged target class and shows the average TVIs of the individual bands and indices. In addition, the last column shows the average TVI over all classes. The best results per class are shown in bold and the TVIs of the indices that are greater than any of the TVIs of the raw bands are shown in italics. As can be seen in the last row, all optimized LRIs achieve higher TVIs on average compared to the raw single-channel bands. However, the highest average TVIs are obtained by the LRI_6 with 0.041 and LRI_2 with 0.034. The best raw bands are the EIR and LWIR band sharing an average TVI of 0.023. Regarding the number of best results, the LRI_6 and the LRI_2 rank highest with four and three top results, respectively. Only one top result can be found in the results of the LRI_{nd} and LRI_r, while the remaining four top results are achieved by the raw bands. Two of these four top results are achieved by the LWIR band for the persons in uniforms. This is not surprising as the LWIR band is very sensitive to this kind of thermal activity. The other two top results are achieved by the green band for the green 2D camouflage net and the yellow 3D camouflage net.

Regarding the camouflaged target classes that do not lead to higher TVIs in the optimized indices than in the raw bands, it can be observed that all optimized indices have a relatively low TVI. This is the case, for example, for the yellow 3D camouflage net and the persons in uniforms, where the TVI in the optimized indices does not even exceed 0.01, while it does for many other camouflaged target classes. However, it can also be seen that

Table 6. Evaluation results in terms of TVI of all single-channel bands and optimized indices.

	blue	green	red	EIR	NIR	LWIR	LRI_6	LRI_2	LRI_{nd}	LRI_r
anthracite fleece	0.010	0.023	0.014	0.046	0.058	0.061	*0.113*	0.051	*0.097*	**0.144**
artificial turf	0.004	0.002	0.005	0.024	0.024	0.025	**0.067**	*0.058*	*0.041*	*0.041*
artificial hedge	0.005	0.012	0.006	0.011	0.009	0.006	*0.025*	**0.043**	0.008	0.004
gray tar	0.017	0.008	0.027	0.085	0.061	0.004	0.060	**0.089**	0.023	0.007
green tarp	0.003	0.003	0.003	0.019	0.031	0.018	**0.065**	*0.038*	*0.052*	*0.056*
green 2D camouflage net	0.007	**0.009**	0.007	0.007	0.005	0.009	0.003	0.003	0.006	0.004
green 3D camouflage net	0.001	0.005	0.001	0.024	0.026	0.004	*0.044*	**0.044**	*0.029*	0.025
gray 3D camouflage net	0.026	0.004	0.004	0.010	0.021	0.011	**0.060**	*0.038*	*0.030*	*0.027*
yellow 3D camouflage net	0.021	**0.025**	0.030	0.010	0.001	0.006	0.004	0.003	0.002	0.002
person in green uniform	0.009	0.005	0.008	0.009	0.011	**0.021**	0.005	0.009	0.005	0.002
person in yellow uniform	0.019	0.012	0.023	0.008	0.001	**0.026**	0.003	0.002	0.007	0.005
average	0.011	0.010	0.012	0.023	0.023	0.017	**0.041**	*0.034*	*0.027*	*0.029*

Values show the average TVI over all annotations of the respective camouflaged target in the test dataset. The last row shows the average of these values. Best results per camouflaged target are in bold, while optimized indices that achieve a higher TVI than any of the raw bands are in italics.

this is not due to missing information in the raw data, as the green band produces a TVI of 0.025 for the yellow 3D camouflage net. On the contrary, the optimized indices can achieve relatively high TVIs even when the TVI provided by the raw bands is low, as can be seen with the artificial hedge. Here, the LRI_2 provides an average TVI of 0.043, which is more than three times higher than the highest average TVI of the raw bands, which is 0.012 provided by the green band. Therefore, low TVIs in the raw bands do not necessarily lead to low TVIs in the optimized LRIs, but the same is true for high TVIs in the raw bands. This shows the ability of the LRI to emphasize certain information in the raw bands that can only be extracted by relating them to each other using its tunable parameters. However, it also shows that the optimization algorithm prioritizes certain targets more than others, which may not be a desirable outcome depending on the actual application of the optimized LRI. Section 5 discusses this issue in more detail and how it could be addressed. Overall, the optimized indices provide much higher TVIs than the raw bands for many camouflaged target classes, resulting in superior average performance. This demonstrates the general effectiveness of the LRI and the proposed optimization approach as well as the ability of the indices to generalize from the training dataset to the test dataset.

4.3.2. Camouflage detection

For the evaluation of the optimized indices on actual camouflage detection performance, each LRI is incorporated into a computer-aided camouflage detection pipeline, similar to (Hupel and Stütz 2022, 2024). In detail, the original bands of each capture in the test dataset are extended with the optimized indices, resulting in an individual band stack per index. Furthermore, the bands of each capture are individually normalized to an average value of zero and a standard deviation of one. Since the visual band is treated as a single band, but technically consists of three bands, each of its bands is additionally multiplied by one-third. This normalization ensures that each band is given equal weight in the detection process, regardless of its actual range of values. To evaluate the impact of the index extensions on detection performance, each band stack is processed by multiple spectral anomaly detection algorithms. These algorithms are the RX (Reed and Yu 1990), LRX (Reed and Yu 1990), LPD (Hupel and Stütz 2022) and AED (Kang et al. 2017) anomaly detectors, introduced in Section 3.3. Note that each detector is based on a different operating principle, which makes this evaluation particularly robust. In order to find the most optimal detector configuration, multiple configurations are tested for each detector. For the LRX and LPD detectors, the inner and outer window sizes (w_i, w_o) are set to (5, 15), (11, 31), (21, 61), (31, 91) and (41, 121). For the AED detector, the connected component size κ is set to 25, 50, 100, 200 and 300. The RX detector works without any parameters. As a result, there are five different band stack configurations per capture of the test dataset, one for each optimized LRI and one with only the raw bands. Each of these band stacks is processed by four different detection algorithms, each with multiple different configurations.

An exemplary impression of the detection results is given in Figure 6. It shows the detection maps of each detector and band stack configuration of the same sample from the test dataset as in Figure 5. The separate column on the right shows the label mask of the sample and the corresponding index images added to the raw band stack. For the LRX, LPD and AED detectors, the detection maps that result from the configuration that provides the highest possible average IOU over all camouflaged targets in the scene are shown. As can be observed, the detectors generally give very different detection results, highlighting certain camouflaged targets more than others. For example, the LRX detector detects the targets marked with red circles in Figure 6e more than any other targets, while the RX

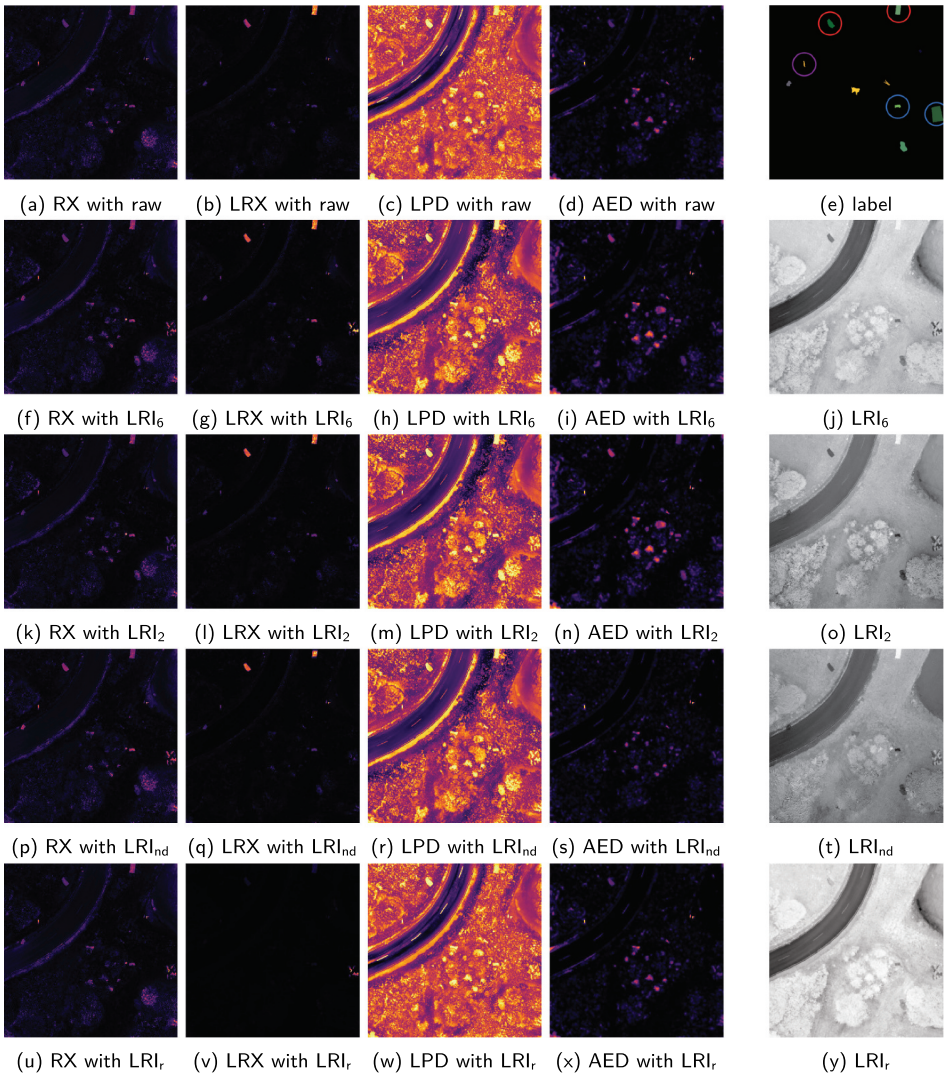

Figure 6. Detection maps of all detectors and band stack configurations for a single sample from the test dataset. From the first row to the last row, the detection maps of the raw (a - d), LRI_6 (f - i), LRI_2 (k - n), LRI_{nd} (*p* - s) and LRI_r (u - x) band stack configurations are shown. The separate column on the right shows the label mask (e) of the sample in the first row, while the remaining rows (j, o, t and y) show the index images added to the raw band stack of the sample. The circle markers in the label mask (e) are explained in the text.

and AED detectors show the same pattern for the target marked with a purple circle. In addition, the detection maps vary between the different band stack configurations, but this may also be due to the different detector configurations. For instance, the LRX detector achieves higher detection rates for the targets marked with blue circles in Figure 6e with the LRI_6 band stack configuration than with any other band stack configuration, as shown in Figure 6g. However, due to the shear amount of the detection maps and the difficulty of their objective interpretation, a qualitative evaluation is virtually impossible. Therefore, the following evaluation focuses on statistics.

The first part of the evaluation is based on the receiver operating characteristic (ROC) and the area under curve (AUC). For this purpose, all band stacks of each capture are first processed by all anomaly detectors and with all different detector configurations. Then, for each detector, detector configuration and camouflaged target class, the ROC over the entire test dataset is determined by threshold averaging (Fawcett 2006). This leaves one ROC for each camouflaged target class, band stack configuration and detector with its configuration. In order to obtain a single representative ROC for each detector for each camouflaged target class, only the detector configuration that yields the highest AUC for the respective camouflaged target class is further considered. The configurations of these detectors can be found in Appendix A. As a result, there is only one ROC and AUC for each detector, band stack

configuration and camouflaged target class. It should also be noted that there is an additional post-processing step after the thresholding of the detection maps in the ROC generation procedure. After each thresholding of the detection maps, the resulting binary detection maps are processed by a morphological area filter that removes all bright connected components smaller than 10 pixels. 10 pixels correspond to an area of 0.4 square meters, which is much smaller than any of the camouflaged targets in the dataset. Therefore, this post-processing step does not negatively affect detection performance. Instead, it reduces false alarms.

Since it would reduce clarity to show the ROCs of each detector and band stack configuration in a single graph, the ROCs are generated separately for each detector. In addition, an average ROC is obtained by applying threshold averaging across all detectors for each band stack configuration. As a result, there is a single average ROC for each detector, camouflaged target class and band stack configuration, providing an indicator of the resulting overall performance of a specific band stack configuration for a camouflaged target class. These average ROCs are shown in Figure 7, while the ROCs for each individual detector are listed in Appendix B for better clarity of this section. Each graph shows the ROCs of all band stack configurations for a single camouflaged target class, except the graph in the bottom right, which shows the threshold averaged ROCs of each band stack configuration across all camouflaged target classes. The ROCs marked as raw show the results of the band stack configuration, where only the raw bands without any additional spectral index where processed by the anomaly detection algorithms. LRI_6, LRI_2, LRI_{nd} and LRI_r denote the band stack configurations where the raw band stack has been extended by the respective optimized index. Note that there are different limits for the abscissa and ordinate for each different graph.

As can be observed, detection performance varies greatly between targets, with no band stack configuration being significantly superior. Even the number of steepest curves is almost evenly distributed among the different band stack configurations. For example, the artificial turf, green 3D camouflage net and gray 3D camouflage net are best detected by the LRI_6 band stack configuration, while the green tarp and anthracite fleece are best detected by the LRI_r band stack configuration. In addition, the raw band stack configuration performs better than any of the other band stack configurations for the yellow 3D camouflage net and persons in yellow uniforms. It is also worth noting that the LRI_{nd} band stack configuration never performs best, except for the persons in green uniforms, but the LRI_r band stack configuration performs just as well for that target. There are also targets where all band stack configurations produce very similar results, such as the artificial hedge and the green 2D camouflage net. In these cases, a single best band stack configuration cannot be easily determined. Reviewing the average results in the bottom right, it becomes clear that the LRI_6 and LRI_r band stack configurations lead to the highest detection rates at low false alarm rates. The LRI_2 band stack configuration performs slightly worse than these two and the raw band stack configuration performs the worst, while the LRI_{nd} band stack configuration provides a result somewhere in between.

For the evaluation of the receiver operating characteristics (ROCs) for their area under curves (AUCs), the detectors can be assessed separately without leading to overly complex results. Table 7 shows the AUC for each detector, band stack configuration and camouflaged target class. Again, these AUCs result from the best detector configurations for each camouflaged target class, which is why there is only a single AUC per detector, even though there are multiple configurations for most detectors. The AUCs are organized into five different subtables, one for each detector and another showing the averages over all detectors. In each of these subtables, there are five rows representing the different band stack configurations, one for each optimized index and one for the band stack without any index extensions, called raw. The columns represent the camouflaged target classes. Conclusively, each cell of each subtable shows the AUC for a particular camouflaged target class given by the column and for a band stack configuration given by the row. The last column of each subtable shows the average AUC value in the row, indicating the average performance of the band stack configuration across all camouflaged target classes. In the last subtable, labeled Average, the average AUC over all detectors is shown, providing an overall performance measure of the corresponding band stack configuration. As a result, the last column of the last subtable shows the average AUC of each band stack configuration over all detectors and camouflaged target classes. Therefore, this column provides a final single overall performance score for each band stack configuration. As with the results in the previous section, the top results among the different band stack configurations are shown in bold.

Considering the average results in the last subtable, it can be concluded that the LRI_6 and LRI_2 band stack configurations achieve the highest AUCs on average. In addition, the top most average AUCs per camouflaged target class can be found in the results of the LRI_6 and LRI_2 band stack configurations, with five and three top results, respectively. The LRI_r band stack configuration

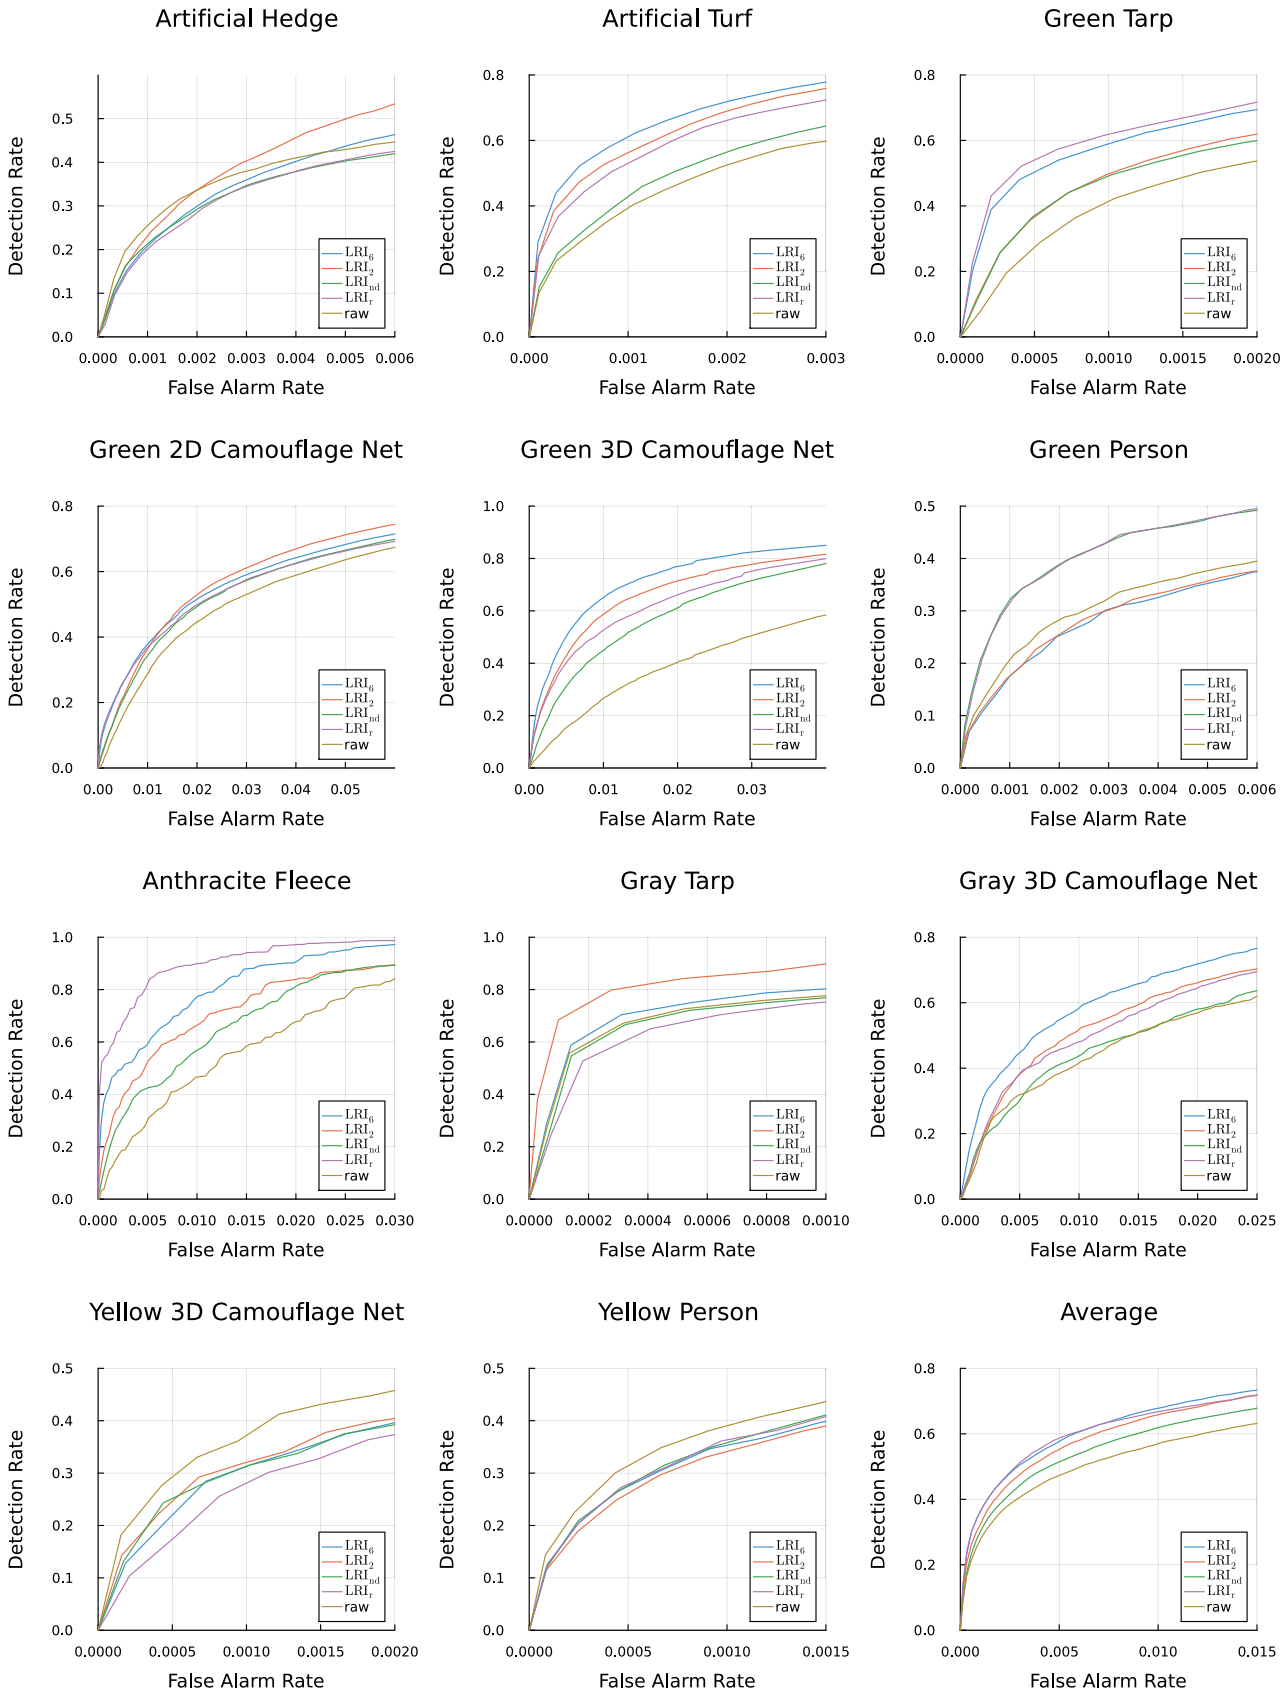

Figure 7. Averaged receiver operating characteristics (ROC) for each camouflaged target class and band stack configuration.

Table 7. Area under curve (AUC) results of each detector and band stack configuration.

	Hedge	Turf	Gn. Trp	Gn. 2D	Gn. 3D	Gn. Per.	Fleece	Gy. Trp.	Gy. 3D	Yw. 3D	Yw. Per.	Avg
RX												
LRI_6	0.972	**0.991**	0.991	0.914	**0.955**	0.917	0.997	0.999	**0.991**	0.954	0.933	**0.965**
LRI_2	0.972	0.990	0.988	**0.916**	0.944	0.919	0.992	**0.999**	0.988	0.950	0.934	0.963
LRI_{nd}	0.967	0.991	0.990	0.911	0.945	**0.921**	0.992	0.999	0.988	0.951	0.934	0.963
LRI_r	**0.973**	0.990	**0.992**	0.910	0.935	0.912	**0.999**	0.999	0.988	**0.956**	0.934	0.962
raw	0.973	0.989	0.985	0.903	0.895	0.911	0.982	0.999	0.988	0.955	**0.940**	0.956
LRX												
LRI_6	0.993	**0.999**	0.998	**0.964**	**0.993**	0.952	0.998	1.000	0.991	0.979	0.991	0.987
LRI_2	0.993	0.998	0.997	0.964	0.989	0.950	0.996	**1.000**	0.990	0.977	0.991	0.986
LRI_{nd}	0.992	0.997	0.997	0.964	0.989	0.958	0.993	1.000	0.989	0.977	0.991	0.986
LRI_r	0.992	0.998	**0.998**	0.963	0.991	**0.959**	**1.000**	1.000	0.990	0.979	0.991	**0.987**
raw	**0.994**	0.997	0.996	0.963	0.972	0.949	0.992	1.000	**0.991**	**0.980**	**0.992**	0.984
LPD												
LRI_6	0.960	**0.995**	0.991	0.885	**0.989**	0.908	0.993	0.996	**0.952**	**0.956**	0.970	0.963
LRI_2	**0.974**	0.994	0.990	**0.904**	0.987	0.907	0.990	**0.998**	0.947	0.951	0.971	**0.965**
LRI_{nd}	0.937	0.989	0.991	0.871	0.980	0.915	0.990	0.992	0.937	0.951	0.971	0.957
LRI_r	0.940	0.991	**0.991**	0.869	0.984	**0.916**	**0.997**	0.992	0.942	0.955	**0.973**	0.959
raw	0.945	0.986	0.988	0.882	0.973	0.896	0.984	0.993	0.933	0.951	0.971	0.955
AED												
LRI_6	0.927	**0.991**	**0.991**	0.950	**0.965**	0.921	0.983	0.997	**0.948**	**0.974**	0.977	**0.966**
LRI_2	**0.964**	0.988	0.988	**0.954**	0.958	0.906	0.970	**0.998**	0.941	0.970	0.974	0.965
LRI_{nd}	0.899	0.982	0.987	0.929	0.936	0.939	0.942	0.996	0.922	0.970	0.978	0.953
LRI_r	0.904	0.985	0.991	0.936	0.952	**0.943**	**0.989**	0.997	0.936	0.973	**0.980**	0.962
raw	0.912	0.967	0.979	0.917	0.893	0.907	0.939	0.997	0.915	0.968	0.975	0.943
Average												
LRI_6	0.963	**0.994**	0.993	0.929	**0.976**	0.925	0.993	0.998	**0.970**	**0.966**	0.968	**0.970**
LRI_2	**0.976**	0.993	0.991	**0.934**	0.970	0.920	0.987	**0.999**	0.966	0.962	0.968	0.970
LRI_{nd}	0.949	0.990	0.991	0.919	0.963	**0.933**	0.980	0.997	0.959	0.962	0.968	0.965
LRI_r	0.953	0.991	**0.993**	0.919	0.966	0.932	**0.996**	0.997	0.964	0.966	0.969	0.968
raw	0.956	0.985	0.987	0.916	0.933	0.916	0.974	0.997	0.957	0.963	**0.970**	0.959

results in two top average AUCs, while the LRI_{nd} and raw band stack configurations both result in a single top average AUC. However, the last column shows that all index-based band stack configurations outperform the raw band stack configuration on average. The LRI_6 and LRI_2 band stack configurations are virtually equal with 0.970 AUC, followed by LRI_r with 0.968 AUC and LRI_{nd} with 0.965 AUC. With an average AUC of 0.959, the raw band stack configuration performs significantly worse than all others. The AUCs of the individual detectors shown in the other subtables mostly follow the pattern of the averaged results. In general, the LRI_6 and LRI_2 band stack configurations lead to the highest AUC on average and many top AUCs per camouflaged target class, while the LRI_r and LRI_{nd} band stack configurations perform slightly worse and the raw band stack configuration performs worst. Nevertheless, the results of the LRI_r band stack configuration are mostly very close to those of the LRI_6 and LRI_2 band stack configurations. For example, the number of top results per camouflaged target class of the RX detector is the same of the LRI_6 and LRI_r band stack configurations. In addition, the average AUC over all camouflaged target classes is highest for the LRI_r band stack configuration in the results of the LRX detector. It is also worth noting that there are many top results in the results of the LRX detector obtained with the raw band stack configuration. Yet, the advantage over the other band stack configurations is very small in these cases, resulting in a lower average AUC compared to all other band stack configurations. As with the average results, the results of the LPD and AED detectors are dominated by the LRI_6 and LRI_2 band stack configurations. It is also important to note that the LRI_{nd} band stack configuration does not achieve a single top AUC for these and the LRX detector. The only top result of this band stack configuration in any of the results can be achieved for the persons in green uniforms using the RX detector. In summary, the LRI_6 and LRI_2 band stack configurations produce the best overall results, as indicated by the average results in the lower subtable.

Another perspective on the performance of the different band stack configurations is offered by the results in Table 8, which constitute the second part of this evaluation. In contrast to the receiver operating characteristic and its area under curve, which are threshold independent performance measures, the following evaluation is based on intersection over union, which requires a binary detection map. This binary detection map is obtained by thresholding the raw detection map. Again, this evaluation considers each camouflaged target class, detector and band stack configuration individually. In order to achieve this, a single threshold must first be determined for each detector. This threshold is the one that produces the highest average IOU across all

Table 8. Intersection over union (IOU) for each detector and band stack configuration.

	Hedge	Turf	Gn. Trp	Gn. 2D	Gn. 3D	Gn. Per.	Fleece	Gy. Trp	Gy. 3D	Yw. 3D	Yw. Per.	Avg.
RX												
LRI_6	0.26	**0.34**	0.48	0.09	**0.12**	0.05	0.15	**0.80**	**0.23**	0.07	0.05	0.24
LRI_2	0.27	0.33	0.32	0.09	0.09	0.04	0.09	0.78	0.14	0.07	0.05	0.21
LRI_{nd}	0.28	0.25	0.35	0.09	0.06	0.04	0.03	0.74	0.20	**0.09**	0.06	0.20
LRIr	0.25	0.26	**0.59**	**0.10**	0.07	0.04	**0.44**	0.71	0.14	0.06	0.05	**0.25**
raw	**0.31**	0.23	0.28	0.08	0.02	**0.07**	0.02	0.77	0.13	0.13	**0.07**	0.19
LRX												
LRI_6	0.35	**0.56**	0.66	**0.16**	0.39	0.23	**0.50**	0.78	**0.22**	0.23	0.23	**0.39**
LRI_2	0.36	0.51	0.50	0.15	0.30	0.25	0.18	0.79	0.18	0.26	0.24	0.34
LRI_{nd}	0.36	0.36	0.52	0.15	0.30	**0.26**	0.08	**0.79**	0.16	0.22	0.24	0.31
LRIr	0.35	0.46	**0.69**	0.16	**0.32**	0.25	0.49	0.74	0.20	0.22	0.23	0.37
raw	**0.42**	0.41	0.48	0.14	0.11	0.30	0.12	0.78	0.20	**0.28**	**0.28**	0.32
LPD												
LRI_6	0.09	0.44	**0.43**	0.10	**0.26**	0.24	0.18	0.57	**0.14**	**0.35**	0.27	**0.28**
LRI_2	**0.16**	**0.44**	0.39	0.08	0.23	**0.26**	0.09	**0.68**	0.07	0.33	0.25	0.27
LRI_{nd}	0.05	0.22	0.34	0.09	0.15	0.24	0.10	0.50	0.04	0.34	0.27	0.21
LRIr	0.05	0.30	0.42	**0.11**	0.21	0.25	**0.36**	0.46	0.09	0.33	**0.28**	0.26
raw	0.05	0.20	0.29	0.06	0.09	0.24	0.09	0.53	0.03	0.34	0.26	0.20
AED												
LRI_6	0.03	**0.16**	0.26	**0.08**	**0.07**	0.10	0.04	0.51	**0.02**	0.21	0.18	**0.15**
LRI_2	**0.04**	0.11	0.19	0.08	0.07	0.10	0.01	**0.58**	0.01	0.22	0.16	0.14
LRI_{nd}	0.01	0.07	0.18	0.06	0.05	0.11	0.01	0.41	0.01	0.20	0.16	0.11
LRIr	0.01	0.15	**0.26**	0.08	0.07	**0.12**	**0.09**	0.37	0.01	0.20	**0.18**	0.14
raw	0.01	0.04	0.10	0.04	0.03	0.10	0.01	0.42	0.01	**0.23**	0.17	0.10
Average												
LRI_6	0.18	**0.38**	0.46	0.11	**0.21**	0.16	0.22	0.67	**0.15**	0.21	0.18	**0.27**
LRI_2	**0.21**	0.35	0.35	0.10	0.17	0.16	0.09	**0.71**	0.10	0.22	0.17	0.24
LRI_{nd}	0.17	0.23	0.35	0.10	0.14	0.16	0.05	0.61	0.10	0.21	0.18	0.21
LRIr	0.17	0.29	**0.49**	**0.11**	0.17	0.16	**0.34**	0.57	0.11	0.20	0.19	0.25
raw	0.20	0.22	0.29	0.08	0.06	**0.18**	0.06	0.62	0.09	**0.24**	**0.19**	0.20

camouflaged targets of the same class in the test dataset. As a result, there is a threshold for each detector configuration, camouflaged target class and band stack configuration. Similar to the evaluation based on ROCs and AUCs, a single representative detector configuration per camouflaged target class and band stack configuration is determined by selecting the one that results in the highest average IOU. Only this time, this configuration also includes a threshold. The tested detector configurations are the same as before, but the best configurations per detector may of course be different from those resulting from the evaluation using AUCs. See Appendix A for the optimal configurations and thresholds for each detector, camouflaged target class and band stack configuration. Before calculating the IOU, the binary detection map is post-processed by removing all bright connected components smaller than 10 pixels, as in the evaluation using ROCs and AUCs. Again, this reduces false positives and therefore increases IOU. The table is structured in the same way as Table 7 with a subtable for each detector and a subtable containing the averages over all detectors. Rows represent camouflaged target classes and columns represent band stack configurations, as well. In the far right column of the last subtable are the averages of all the averages, providing a single IOU for each band stack configuration. The best results per camouflaged target and detector are shown in bold again.

As with the results in Table 7, the averages of all the averages show that the LRI_6 band stack configuration leads to the highest average IOU of 0.27. The performances of the LRI_2 and LRI_r band stack configurations are a little worse with an average IOU of 0.24 and 0.25, respectively. With average IOUs of 0.21 and 0.20, the LRI_{nd} and raw band stack configurations achieve the worst average performances of all. As can be observed in the last subtable, the number of top average results is dominated by the LRI_6 band stack configuration with four top results, followed by the LRI_r and raw band stack configurations with three top results, and the LRI_2 band stack configuration with two top results. No best result is achieved by the LRI_{nd} band stack configuration. Compared to the results in Table 7, the performances of the individual detectors show a strong and noticeable spread. For example, the best average result of 0.39 IOU of the LRX detector is more than two times higher than the best average result of 0.15 of the AED detector. The best average results of the other two detectors fall in between, with 0.25 IOU for the RX detector and 0.28 IOU for the LPD detector. But this may also be due to the different nature of AUC and IOU. However, this time the LRI_6 band stack configuration consistently results in the

best average IOU, except for the RX detector, but the difference of 0.01 is marginal. In terms of the number of best results, the LRI_6 produces the highest amount, except for the RX detector, where the LRI_r band stack configuration produces the same amount with four best results. The LRI_2 band stack configuration achieves nearly the same overall performance as the LRI_r band stack configuration, except for the LRX detector, resulting in a slightly lower average performance of 0.01 IOU less. Similar to the results in Table 7, the performances of the LRI_{nd} and raw band stack configurations are consistently worse than all others, leading to a significantly worse overall performance, as well. However, it is worth noting that the raw band stack configuration results in three top average results in the lower subtable, as well as in the results of the RX and LRX detectors. Nevertheless, the differences in these cases are comparatively small, resulting in average performances that do not even exceed the average performances of the LRI_{nd} band stack configuration, which achieves either fewer or the same number of best results. Overall, the LRI_6 band stack configuration produces the highest average IOU and number of best results by a significant margin, outperforming all other band stack configurations, supporting the evaluation using ROCs and AUCs. Although the LRI_2 band stack configuration achieves a lower average IOU than the LRI_r band stack configuration, while the opposite is true for the average AUCs, the results are not far apart. This is also consistent with the ROCs in Figure 7, as well as the results of the raw and LRI_{nd} band stack configurations, which are the worst across all evaluation perspectives.

5. Discussion

This section discusses the implications and limitations of the findings of this study in Section 5.1 and Section 5.2, respectively. Section 5.3 then identifies potential future prospects.

5.1. Implications

The results presented in Section 4.3.1 demonstrate that the LRI can be successfully optimized to improve the visibility, i.e. TVI, of camouflaged targets using the proposed approach. As described, the visibility provided by the optimized indices for the camouflaged targets is significantly higher than the visibility provided by any of the raw bands, with few exceptions. Furthermore, the LRI_6 and LRI_2 achieve higher average visibility than the LRI_{nd} and LRI_r as well as all the raw bands. On the one hand, this shows that optimized LRIs can generally produce index images that perform better than any raw band, which is the primary motivation for using a spectral index in the first place. On the other hand, it shows that an optimized LRI can yield better results than an optimized normalized difference or ratio index, challenging these approaches. The higher number of parameters in the LRI, which allows more information to be encoded in the index itself, may be the reason for its superior performance. This is also supported by the fact that the LRI_2 performs worse and has fewer parameters than the LRI_6, but performs better and has more parameters than the LRI_{nd} and LRI_r.

The generally higher visibility provided by the LRI_2 compared to the LRI_{nd} and LRI_r, all of which use only two bands, also indicates that the optimization approach can be successfully applied to reduce the number of bands required for a spectral index while maintaining as much performance as possible. In addition, limiting parameter precision directly during optimization has no observable negative impact on optimization runtime or visibility results. Therefore, the complexity reduction strategies of the optimization approach may be suitable techniques to optimize LRIs that preserve the human readability of their parameters and limit the number of utilized spectral bands.

However, as mentioned in Section 4.3.1, it should be noted that even the LRI_6 and LRI_2 do not provide superior visibility than the raw bands for all camouflaged targets. This could be caused by the optimizer focusing on camouflaged targets that have very similar spectral characteristics, so that optimizing for one target simultaneously optimizes for other targets. Consequently, some camouflaged targets with completely different spectral properties may be partially ignored by the optimizer. This is supported by the fact that the optimized indices provide relatively low visibility for the persons in uniforms. These camouflaged targets are the only ones that are thermally active, making them spectrally very distinct. Optimizing the LRI for this particular kind of camouflaged target may result in low visibilities for many other targets, resulting in a lower average visibility than optimizing for all the other targets instead. Since the indices are optimized for average visibility, the optimizer does not differentiate between an average coming from high and low values or from evenly distributed values. As a result, a higher average value is preferred regardless of its origin, which may be the cause of the highly unbalanced visibility values. In addition, this could indicate that it is generally difficult, if not impossible, to find a single LRI that increases the visibility of all camouflaged targets in general.

Regarding the detection results in Section 4.3.2, it can be concluded that the strong improvements in visibility

can be successfully translated into improvements in detection performance. Overall, the LRI_6 band stack configuration provides the highest detection performance, while the LRI_2 band stack configuration provides similar or slightly lower detection performance, consistent with their visibility performance in Section 4.3.1. It should be noted, however, that the LRI_r band stack configuration provides comparable performance to the LRI_2 band stack configuration, despite providing inferior visibility. This could be due to the range of values of the ratio index, which quickly goes to infinity when the denominator is small and therefore can tremendously increase the abnormality of some targets. Further support for this is provided by the worst results produced by the LRI_{nd} band stack configuration, since the LRI_{nd} is based on the same bands as the LRI_r only as a normalized difference, whose range of values does not behave as extreme as that of a ratio index. Nevertheless, the LRI_{nd} band stack configuration still provides higher detection performance than using only the raw bands, which gives the worst results on average. Regardless of the underlying causes for the high detection performance of the LRI_r band stack configuration, this shows that an increased TVI does not necessarily translate into increased detection performance. Conclusively, the TVI may not be an ideal optimization criterion for simultaneously maximizing the visibility and detectability of camouflaged targets.

As mentioned before, some camouflaged targets have higher visibility in the raw bands than in any of the optimized indices. Interestingly, this also results in a slight decrease in detection performance for those targets, such as the persons in yellow uniforms. This suggests that an optimized index may have a negative impact on detection performance for camouflaged targets whose visibility is not improved by it. However, the observed performance degradation is very small and may be negligible given the strong improvements for the other targets when using an optimized LRI band stack configuration.

Overall, the LRI and the proposed optimization approach are suitable methods for producing index images that improve the visibility and detectability of camouflaged targets. In addition, the LRI_6 and LRI_2 generally outperform the LRI_{nd} and LRI_r. This demonstrates the superiority of the LRI and the optimization approach over traditional band permutation searches, making it the preferred choice for spectral index optimization if the higher computational effort to obtain the parameters is acceptable.

Besides the positive results obtained in this study, it could prove equally beneficial with different data or even in other use cases. Since the optimization criterion defines the application of the LRI, the optimization procedure remains the same even if the criterion changes. Therefore, the LRI and the optimization approach could be adapted to other applications and provide similar improvements. Moreover, the formulation of the LRI could be changed to a more complex equation to fit even more complicated relationships. For example, the linear function in the numerator and denominator could be replaced by a second-order polynomial. Substituting the optimization criterion or index formulation is principally unproblematic, since the differential evolution optimizer can generally handle any type of problem. However, increasing the complexity of the objective function also introduces a much higher computational overhead into the optimization procedure itself, which can slow down experiments with the proposed approach tremendously. Given the novelty of the LRI, even though it consists of a simple ratio of linear functions, it is advisable to experiment with its raw formulation first.

5.2. Limitations

Although the proposed approach could be successfully applied to increase the visibility and detectability of camouflaged targets, it should be noted that all results are based on MUDCAD-X, a single dataset specifically tailored for camouflage detection in reconnaissance scenarios with small tactical drones. For example, all images in the dataset were acquired at an altitude that is common for tactical reconnaissance scenarios, but low enough to make any atmospheric disturbances irrelevant. Conclusively, such effects are not examined in this study and do not influence it in any way. In addition, MUDCAD-X is a comparatively small dataset with 853 samples, leading to equally small training and test datasets. Given the small number of different target classes, a class-based split is also virtually impossible. As a result, the generalizability of this study and the reliability of the LRI may be limited. For instance, it is difficult to infer how well the resulting LRIs work for unknown camouflaged targets not included in MUDCAD-X, especially considering that they already provide less visibility than the raw bands for some targets. The same applies to the optimization approach of the LRI, since it has been investigated under the same constraints and may yield different results for other datasets and camouflaged targets. Compiling a greater dataset with many more targets and re-running the experiments on a class-based split may shed some light on the impact of these limitations. Unfortunately, due to the general scarcity of this kind of

data and the effort required to generate it, this investigation is beyond the scope of this study.

Another limitation of this study results from the relatively small number of applied detection algorithms. Although the selected methods all work on different principles, the results may be different if other detection methods are employed. In addition, using the TVI as optimization criterion for the LRIs may not be an ideal choice for improving detection performance. As a result, other methods may not yield better detection results and different optimization criteria may result in greater improvements than the TVI. Nevertheless, all the methods used in this work provide better detection results on average when using a band stack configuration with an LRI optimized with the TVI. This suggests similar results for any other spectral anomaly detection method.

Given the numerical nature of the proposed approach, it should also be noted that it does not necessarily reveal information about the spectral relationships between the target and background materials. As can be seen from the parameters found for the LRI_6 and LRI_2, which range between −10 and 10, there is no observable pattern to explain the obtained results. Whereas the bands for the LRI_{nd} and LRI_r are physically plausible, since the NIR band is very sensitive to green vegetation and the LWIR band is particularly sensitive to thermally active targets. Therefore, the potential use of the LRI to derive spectral properties of materials that result in high visibility or detectability is limited. However, the LRI and its optimization approach are specifically designed to work without any knowledge of spectral properties, making it unlikely to provide much insight into these aspects.

5.3. Future prospects

Since the optimized LRIs do not provide higher visibility, i.e. TVI, for all camouflaged targets than the raw bands, future research could focus on the following improvements to the proposed approach. One strategy may be to add an optimization criterion that penalizes a large spread of TVI values among the targets, alongside the existing optimization criterion. While this could reduce the TVI for some targets, it could also increase the TVI for those targets that currently result in a higher TVI with the raw bands. However, this would also complicate the optimization procedure to a multi-objective problem, which may introduce new problems that need to be addressed. Alternatively, this could be implemented by merging these two optimization criteria using a weighted sum. While this would require fine-tuning the weights, the existing optimization procedure can be left mostly untouched, as only the objective function changes. Besides these approaches, an entirely new optimization criterion could also be introduced as an alternative to the TVI. Naturally, this would also establish a new reference for visibility, which may limit comparisons with this study.

Another approach to obtain spectral indices with even greater utility could be to optimize multiple LRIs at the same time. Specifically, given multiple LRIs in the optimization procedure, the TVI for a single annotation is determined by taking the TVI of the LRI that provides the highest TVI. As a result, the TVI for each annotation can come from any of the given LRIs. The average TVI over the entire dataset is calculated in the same way as before, which has the advantage that the existing optimization implementation can be left mostly untouched again. With this optimization strategy, each LRI can prioritize and ignore camouflaged targets without being penalized as a single LRI would be. In addition, the optimizer could automatically determine which index is optimized for which target simply by adjusting the parameters of each LRI, since this also changes, which LRI provides the highest TVI for which target. Naturally, this ultimately adds complexity, as there are multiple final index images to consider in any subsequent application. Furthermore, all resulting index images must be treated as equally valuable, as there is no way to know which one provides the highest TVI for a particular target or scene. Considering the camouflage detection application using spectral anomaly detection, as in this study, this would introduce additional bands to process, which typically increases processing time. However, since the optimized LRIs could improve the visibility and thus the abnormality of more targets than a single optimized LRI, this optimization strategy could also significantly improve detection performance.

In contrast to the pure numerical optimization approach in this study, future research could focus on investigating actual spectral relationships of camouflaged targets and their ennvironment. Incorporating this knowledge into the optimization procedure could provide even better LRIs or speed up the optimization runs. Furthermore, a deeper insight into the underlying physical mechanisms may lead to a more appropriate index formulation than the linear function ratio of the LRI. Depending on the resulting structure, this could also speed up optimization runs and provide better results, since the spectral index would have an actual physical basis.

6. Conclusions

This paper introduces a novel spectral index, the linear ratio index, along with an optimization approach that incorporates complexity reduction strategies to

allow variable raw band utilization and parameter precision. Using the eXtended Multispectral Dataset for Camouflage Detection, featuring ten different target classes and seven bands per sample, a six-band and a two-band linear ratio index are optimized to maximize the visibility of camouflaged targets. These indices are compared to commonly optimized normalized difference and ratio indices in terms of increased visibility and detectability of camouflaged targets. The results show that the best performance is provided by the six-band LRI and slightly worse performance is provided by the two-band LRI, both of which outperform the optimized normalized difference and ratio indices. This suggests the superiority of the LRI and its optimization approach over conventional index optimization strategies and the effectiveness of the complexity reduction strategies while maintaining performance. The obtained LRIs are application-ready and can be easily integrated into any existing camouflage detection process if the required raw bands are available and the potential limitations due to the limited data basis are considered. Apart from its successful application to camouflage detection, the LRI and its optimization approach can be adapted to any use case by substituting the objective function, making it worth investigating for other applications where it could also outperform conventional index optimization techniques. However, since this study focuses on optimizing spectral indices for camouflage detection in reconnaissance scenarios using small tactical drones, the generalizability of its findings may be limited by the specific investigated use case and the availability of data. Therefore, further research is required to fully explore the applicability of these findings and reliability of the proposed methods to different contexts.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This research was funded by the Federal Office of Bundeswehr Equipment, Information Techology, and In-Service Support [BAAINBw]. The APC was funded by the University of the Bundeswehr Munich [UniBwM].

ORCID

Tobias Hupel http://orcid.org/0000-0002-7808-801X
Peter Stütz http://orcid.org/0000-0002-6571-4392

Data availability statement

The dataset used in this study, MUDCAD-X, is publicly available on GitHub (https://github.com/Tobias-UniBwM/MUDCAD-X.), as is the software code used to optimize the LRIs (https://github.com/Tobias-UniBwM/IndexOptim.jl).

References

Akbari, E., A. Darvishi Boloorani, N. Neysani Samany, S. Hamzeh, S. Soufizadeh, and S. Pignatti. 2020. "Crop Mapping Using Random Forest and Particle Swarm Optimization Based on Multi-Temporal Sentinel-2." *Remote Sensing* 12 (9): 1449. https://doi.org/10.3390/rs12091449.

Albertini, C., A. Gioia, V. Iacobellis, and S. Manfreda. 2022. "Detection of Surface Water and Floods with Multispectral Satellites." *Remote Sensing* 14 (23): 6005. https://doi.org/10.3390/rs14236005.

Anastasiou, E., A. Balafoutis, N. Darra, V. Psiroukis, A. Biniari, G. Xanthopoulos, and S. Fountas. 2018. "Xanthopoulos, and Satellite and Proximal Sensing to Estimate the Yield and Quality of Table Grapes." *Agriculture* 8 (7): 94. https://doi.org/10.3390/agriculture8070094.

Berg, A. K. V. D., and T. D. Perkins. 2005. "Nondestructive Estimation of Anthocyanin Content in Autumn Sugar Maple Leaves." *Horticulture Science* 40 (3): 685–686. https://doi.org/10.21273/HORTSCI.40.3.685.

Bezanson, J., A. Edelman, S. Karpinski, and V. B. Shah. 2017. "Julia: A Fresh Approach to Numerical Computing." *SIAM Review* 59 (1): 65–98. https://doi.org/10.1137/141000671.

Brouant, L., H. L. Pham, É. Foucher, and A. Picard. 2023. "VIS-NIR Multispectral for Camouflage Detection." In *Emerging Imaging and Sensing Technologies for Security and Defence VIII*, 22–31. Vol. 12740. SPIE.

Camenzind, M. P., and K. Yu. 2024. "Multi Temporal Multispectral Uav Remote Sensing Allows for Yield Assessment Across European Wheat Varieties Already Before Flowering." *Frontiers in Plant Science* 14. https://doi.org/10.3389/fpls.2023.1214931.

Chen, Y., X. Chen, J. Zhou, Y. Ji, and W. Shen. 2017. "Camouflage Target Detection via Hyperspectral Imaging Plus Information Divergence Measurement." In *International Conference on Optoelectronics and Microelectronics Technology and Application*, edited by Y. Su, C. Xie, S. Yu, C. Zhang, W. Lu, J. Capmany, Y. Luo, Y. Nakano, Y. Hao, A. Yoshikawa, and S. Zhuang, 80–92. Vol. 10244. SPIE.

Cordeiro, M. C. R., J. M. Martinez, and S. Peña-Luque. 2021. "Automatic Water Detection from Multidimensional Hierarchical Clustering for Sentinel-2 Images and a Comparison with Level 2A Processors." *Remote Sensing of Environment* 253:112209. https://doi.org/10.1016/j.rse.2020.112209.

Daughtry, C. S. T., C. L. Walthall, M. S. Kim, E. B. de Colstoun, and J. E. McMurtrey. 2000. "Estimating Corn Leaf Chlorophyll Concentration from Leaf and Canopy Reflectance." *Remote Sensing of Environment* 74 (2): 229–239. https://doi.org/10.1016/S0034-4257(00)00113-9.

Delegido, J., J. Verrelst, J. P. Rivera, A. Ruiz-Verdú, and J. Moreno. 2015. "Brown and Green LAI Mapping Through Spectral Indices." *International Journal of Applied Earth Observation and Geoinformation* 35:350–358. https://doi.org/10.1016/j.jag.2014.10.001.

Eckel, L., and P. Stütz. 2024. "Hyperspectral Sensor Management for UAS: Sensor Context Based Band Selection for Anomaly Detection." In *2024 IEEE Aerospace Conference*, IEEE, 1–14. https://doi.org/10.1109/AERO58975.2024.10521225.

Fawcett, T. 2006. "An Introduction to Roc Analysis." *Pattern Recognition Letters* 27 (8): 861–874. https://doi.org/10.1016/j.patrec.2005.10.010.

Feldt, R., and A. Stukalov. 2018. *Blackboxoptim.jl*. https://github.com/robertfeldt/BlackBoxOptim.jl.

Feng, W., B. B. Guo, H. Y. Zhang, L. He, Y. S. Zhang, Y. H. Wang, Y. J. Zhu, and T. C. Guo. 2015. "Remote Estimation of Above Ground Nitrogen Uptake During Vegetative Growth in Winter Wheat Using Hyperspectral Red-Edge Ratio Data." *Field Crops Research* 180:197–206. https://doi.org/10.1016/j.fcr.2015.05.020.

Gastal, E. S. L., and M. M. Oliveira. 2011. "Domain Transform for Edge-Aware Image and Video Processing." *ACM Transactions on Graphics* 30 (4): 1–12. https://doi.org/10.1145/2010324.1964964.

Gavrilovic, M., D. Jovanovic, P. Božovic, P. Benka, and M. Govedarica. 2024. "Vineyard Zoning and Vine Detection Using Machine Learning in Unmanned Aerial Vehicle Imagery." *Remote Sensing* 16 (3): 584. https://doi.org/10.3390/rs16030584.

Gerstmann, H., M. Möller, and C. Gläßer. 2016. "Optimization of Spectral Indices and Long-Term Separability Analysis for Classification of Cereal Crops Using Multi-Spectral RapidEye Imagery." *International Journal of Applied Earth Observation and Geoinformation* 52:115–125. https://doi.org/10.1016/j.jag.2016.06.001.

Ghrefat, H., A. Y. Kahal, K. Abdelrahman, H. J. Alfaifi, and S. Qaysi. 2021. "Utilization of Multispectral Landsat-8 Remote Sensing Data for Lithological Mapping of Southwestern Saudi Arabia." *Journal of King Saud University - Science* 33 (4): 101414. https://doi.org/10.1016/j.jksus.2021.101414.

Gross, W., F. Queck, S. Schreiner, J. Mispelhorn, J. Kuester, W. Middelmann, M. Vögtli, and M. Kneubühler. 2023. "Experimental Approach to Camouflaged Target Detection and Camouflage Evaluation." In *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, 2149–2152. https://doi.org/10.1109/IGARSS52108.2023.10282243.

Guo, F., Q. Feng, S. Yang, and W. Yang. 2023. "Estimation of Potato Canopy Nitrogen Content Based on Hyperspectral Index Optimization." *Agronomy* 13 (7): 1693. https://doi.org/10.3390/agronomy13071693.

Huang, S., L. Tang, J. P. Hupy, Y. Wang, and G. Shao. 2021. "A Commentary Review on the Use of Normalized Difference Vegetation Index (NDVI) in the Era of Popular Remote Sensing." *Journal of Forestry Research* 32 (1): 1–6. https://doi.org/10.1007/s11676-020-01155-1.

Huete, A., C. Justice, and W. van Leeuwen. 1999. *Modis Vegetation Index (Mod 13), Algorithm Theoretical Basis Document, Version 3*. Greenbelt, MD: Goddard Space Flight Center.

Huete, A. R. 1988. "A soil-adjusted vegetation index (SAVI)." *Remote Sensing of Environment* 25 (3): 295–309. https://doi.org/10.1016/0034-4257(88)90106-X.

Hupel, T., and P. Stütz. 2022. "Adopting Hyperspectral Anomaly Detection for Near Real-Time Camouflage Detection in Multispectral Imagery." *Remote Sensing* 14 (15): 3755. https://doi.org/10.3390/rs14153755.

Hupel, T., and P. Stütz. 2023. "Measuring and Predicting Sensor Performance for Camouflage Detection in Multispectral Imagery." *Sensors (Switzerland)* 23 (19): 8025. https://doi.org/10.3390/s23198025.

Hupel, T., and P. Stütz. 2024. "Sensor-Managed Anomaly Detection for Camouflage Detection in Airborne Multispectral Imagery." In *2024 IEEE Aerospace Conference*, IEEE, 1–11. https://doi.org/10.1109/AERO58975.2024.10521231.

Jiang, X., J. Zhen, J. Miao, D. Zhao, Z. Shen, J. Jiang, C. Gao, G. Wu, and J. Wang. 2022. "Newly-Developed Three-Band Hyperspectral Vegetation Index for Estimating Leaf Relative Chlorophyll Content of Mangrove Under Different Severities of Pest and Disease." *Ecological Indicators* 140:108978. https://doi.org/10.1016/j.ecolind.2022.108978.

Jordan, C. F. 1969. "Derivation of Leaf-Area Index from Quality of Light on the Forest Floor." *Ecology* 50 (4): 663–666. https://doi.org/10.2307/1936256.

Kang, X., X. Zhang, S. Li, K. Li, J. Li, and J. A. Benediktsson. 2017. "Hyperspectral Anomaly Detection with Attribute and Edge-Preserving Filters." *IEEE Transactions on Geoscience & Remote Sensing* 55 (10): 5600–5611. https://doi.org/10.1109/TGRS.2017.2710145.

Karimzadeh, S., and M. H. Tangestani. 2021. "Evaluating the VNIR-SWIR Datasets of WorldView-3 for Lithological Mapping of a Metamorphic-Igneous Terrain Using Support Vector Machine Algorithm; a Case Study of Central Iran." *Advances in Space Research* 68 (6): 2421–2440. https://doi.org/10.1016/j.asr.2021.05.002.

Khangarot, L. S., A. Kumar, B. Singh, and D. Singh. 2023. "Camouflage Target Detection Using Deep Learning with Drone Images." In *2023 International Conference on Electrical, Electronics, Communication and Computers (ELEXCOM)*, IEEE, 1–5. https://doi.org/10.1109/ELEXCOM58812.2023.10370438.

Kriváneк, V., J. Motsch, and Y. Bergeon. 2023. "Hyperspectral Data Acquisition for Military Camouflage in Vegetation – Preliminary Results." In *2023 Communication and Information Technologies (KIT)* (IEEE), 1–7. https://doi.org/10.1109/KIT59097.2023.10297.

Kumar, V., and J. K. Ghosh. 2017. "Camouflage Detection Using Mwir Hyperspectral Images." *The Journal of the Indian Society of Remote Sensing* 45 (1): 139–145. https://doi.org/10.1007/s12524-016-0555-8.

Kumar, V., and R. Singh. 2024. "Enhancement of Receptive Field Using Dilated Convolution for Camouflaged Human Segmentation in RGB Images." *Multimedia Tools & Applications*. https://doi.org/10.1007/s11042-024-19287-5.

Li, Z., J. Wang, D. Xiao, Z. Gu, and H. Xie. 2024. "Iron Ore Rock Classification and Mine Remote Sensing Inversion Based on Spectroscopy and Improved Extreme Learning Machine." *Infrared Physics & Technology* 140:105400. https://doi.org/10.1016/j.infrared.2024.105400.

Luo, X., W. Wang, and C. Deng. 2023. "Cross-AD: Multispectral and Hyperspectral High-Speed Artificial Imitation Object Anomaly Detection." In *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, 1142–1145.

Mouafik, M., A. Chakhchar, M. Fouad, and A. El Aboudi. 2024. "Remote Sensing Technologies for Monitoring Argane

Forest Stands: A Comprehensive Review." *Geographies* 4 (3): 441–461. https://doi.org/10.3390/geographies4030024.

Nandibewoor, A., P. Doni, H. Rashmi, and A. Hegde. 2024. "Anomaly Detection with Camouflage Reconnaissance in Spectral Imaging."

Pérez-García, A., A. Rodríguez-Molina, E. Hernández, and J. F. López. 2024. "Spectral Indices Survey for Oil Spill Detection in Coastal Areas." In *IEEE Journal of Selected Topics in Applied Earth 1086 Observations and Remote Sensing*, IEEE, 1–14. https://doi.org/10.1109/JSTARS.2024.3438123.

Psiroukis, V., N. Darra, A. Kasimati, P. Trojacek, G. Hasanli, and S. Fountas. 2022. "Development of a Multi-Scale Tomato Yield Prediction Model in Azerbaijan Using Spectral Indices from Sentinel-2 Imagery." *Remote Sensing* 14 (17): 4202. https://doi.org/10.3390/rs14174202.

Reed, I. S., and X. Yu. 1990. "Adaptive Multiple-Band Cfar Detection of an Optical Pattern with Unknown Spectral Distribution." *IEEE Transactions on Acoustics, Speech & Signal Processing* 38 (10): 1760–1770. https://doi.org/10.1109/29.60107.

Rivera, J. P., J. Verrelst, J. Delegido, F. Veroustraete, and J. Moreno. 2014. "On the Semi-Automatic Retrieval of Biophysical Parameters Based on Spectral Index Optimization." *Remote Sensing* 6 (6): 4927–4951. https://doi.org/10.3390/rs6064927.

Rouse, J. W., R. H. Haas, J. A. Schell, and D. W. Deering. 1974. "Monitoring Vegetation Systems in the Great Plains with ERTS." In *Goddard Space Flight Center 3d ERTS-1 Symp*, Vol. 1. NASA. https://ntrs.nasa.gov/citations/19740022614.

Sameen, M. I., and B. Pradhan. 2016. "A Novel Built-Up Spectral Index Developed by Using Multiobjective Particle-Swarm-Optimization Technique." In *IOP Conference Series: Earth and Environmental Science*, 012006. Vol. 37.

Shen, Y., J. Li, W. Lin, L. Chen, F. Huang, and S. Wang. 2021. "Camouflaged Target Detection Based on Napshot Multispectral Imaging." *Remote Sensing* 13 (19): 3949. https://doi.org/10.3390/rs13193949.

Spector, L., and J. Klein. 2006. "Trivial Geography in Genetic Programming." In *Genetic Programming Theory and Practice III*, edited by T. Yu, R. Riolo, and B. Worzel, 109–123. Boston, MA: Springer US.

Storn, R., and K. Price. 1997. "Differential Evolution – a Simple and Efficient Heuristic for Global Optimization Over Continuous Spaces." *Journal of Global Optimization* 11 (4): 341–359. https://doi.org/10.1023/A:1008202821328.

Tran, T. V., R. Reef, and X. Zhu. 2022. "*A Review of Spectral Indices for Mangrove Remote Sensing*." *Remote Sensing* 14 (19): 4868. https://doi.org/10.3390/rs14194868.

Tu, B., X. Yang, N. Li, C. Zhou, and D. He. 2020. "Hyperspectral Anomaly Detection via Density Peakclustering." *Pattern Recognition Letters* 129:144–149. https://doi.org/10.1016/j.patrec.2019.11.022.

Tu, B., X. Yang, C. Zhou, D. He, and A. Plaza. 2020. "Hyperspectral Anomaly Detection Using Dual Window Density." *IEEE Transactions on Geoscience & Remote Sensing* 58 (12): 8503–8517.1236. https://doi.org/10.1109/TGRS.2020.2988385.

Tucker, C. J. 1979. "Red and Photographic Infrared Linear Combinations for Monitoring Vegetation." *Remote Sensing of Environment* 8 (2): 127–150. https://doi.org/10.1016/0034-4257(79)90013-0.

Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. "Attention is All You Need." *Advances in neural information processing systems* 30. https://doi.org/10.48550/arXiv.1706.03762.

Vögtli, M., L. Sierro, M. Kneubühler, S. Schreiner, W. Gross, F. Queck, J. Kuester, J. Mispelhorn, and W. Middelmann. 2023. "Hyperthun'22: A Multi-Sensor Multi-Temporal Camouflage Detection Campaign." In *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, Vol. 1204, 2153–2156. IEEE. https://doi.org/10.1109/IGARSS52108.2023.10282104.

Wang, C., P. Wang, and N. Ma. 2022. "A New Water Detection for Multispectral Images Based on Data Simulation and Random Forest." In *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, 3191–3194. https://doi.org/10.1109/IGARSS46834.2022.9884351.

Wang, J., T. Shi, H. Liu, and G. Wu. 2016. "Successive Projections Algorithm-Based Three-Band Vegetation Index for Foliar Phosphorus Estimation." *Ecological Indicators* 67:12–20. https://doi.org/10.1016/j.ecolind.2016.02.033.

Wang, S., X. Wang, L. Zhang, and Y. Zhong. 2022. "Auto-AD: Autonomous Hyperspectral Anomaly Detection Network Based on Fully Convolutional Autoencoder." *IEEE Transactions on Geoscience & Remote Sensing* 60:1–14. https://doi.org/10.1109/TGRS.2021.3057721.

Wang, S., Y. Xu, D. Zeng, F. Huang, and L. Liang. 2024. "Deep Learning-Based Spectral Reconstruction in Camouflaged Target Detection." *International Journal of Applied Earth Observation and Geoinformation* 126:103645. https://doi.org/10.1016/j.jag.2023.103645.

Wang, W., J. Ren, C. Su, and M. Huang. 2021. "Ship Detection in Multispectral Remote Sensing Images via Saliency Analysis." *Applied Ocean Research* 106:102448. https://doi.org/10.1016/j.apor.2020.102448.

Wang, X., J. Liu, P. Peng, Y. Chen, S. He, and K. Yang. 2023. "Automatic Crop Classification Based on Optimized Spectral and Textural Indexes Considering Spatial Heterogeneity." *Remote Sensing* 15 (23): 5550. https://doi.org/10.3390/rs15235550.

Wang, Y., H. X. Li, T. Huang, and L. Li. 2014. "Differential Evolution Based on Covariance Matrix Learning and Bimodal Distribution Parameter Setting." *Applied Soft Computing* 18:232–247. https://doi.org/10.1016/j.asoc.2014.01.038.

Wang, Z., Y. Sun, T. Zhang, H. Ren, and Q. Qin. 2018. "Optimization of Spectral Indices for the Estimation of Leaf Area Index Based on Sentinel-2 Multispectral Imagery." In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, IEEE. 5441–5444. https://doi.org/10.1109/IGARSS.2018.8517747.

Wijayanto, A. W., D. Wahyu Triscowati, and A. H. Marsuhandi. 2020. "Maize Field Area Detection in East Java, Indonesia: An Integrated Multispectral Remote Sensing and Machine Learning Approach." In 2020 12th *1054* International Conference on Information Technology and Electrical Engineering (ICITEE), IEEE, 168–173. https://doi.org/10.1109/ICITEE49829.2020.9271683.

Wu, B., H. Zheng, Z. Xu, Z. Wu, and Y. Zhao. 2022. "Forest Burned Area Detection Using a Novel Spectral Index Based on Multi-Objective Optimization." *Forests* 13 (11): 1787. https://doi.org/10.3390/f13111787.

Yang, H., F. Li, Y. Hu, and K. Yu. 2021. "Hyperspectral Indices Optimization Algorithms for Estimating Canopy Nitrogen Concentration in Potato (Solanum Tuberosum L.)." *International Journal of Applied Earth Observation and Geoinformation* 102:102416. https://doi.org/10.1016/j.jag.2021.102416.

Yang, H., F. Li, W. Wang, and K. Yu. 2021. "Estimating Above-Ground Biomass of Potato Using Random Forest and Optimized Hyperspectral Indices." *Remote Sensing* 13 (12): 2339. https://doi.org/10.3390/rs13122339.

Yang, H., H. Yin, F. Li, Y. Hu, and K. Yu. 2023. "Machine Learning Models Fed with Optimized Spectral Indices to Advance Crop Nitrogen Monitoring." *Field Crops Research* 293:108844. https://doi.org/10.1016/j.fcr.2023.108844.

Zeng, L., B. D. Wardlow, D. Xiang, S. Hu, and D. Li. 2020. "A Review of Vegetation Phenological Metrics Extraction Using Time-Series, Multispectral Satellite Data." *Remote Sensing of Environment* 237:111511. https://doi.org/10.1016/j.rse.2019.111511.

Zhang, J., W. Wang, H. Qiao, C. Xu, J. Guo, H. Si, J. Wang, S. Xiong, and X. Ma. 2023. "Estimation of Leaf Nitrogen Content in Winter Wheat Based on Continuum Removal and Discrete Wavelet Transform." *International Journal of Remote Sensing* 44 (18): 5523–5547. https://doi.org/10.1080/01431161.2023.2247524.

Zhang, Y., T. Wang, Z. Li, T. Wang, and N. Cao. 2023. "Based on Machine Learning Algorithms for Estimating Leaf Phosphorus Concentration of Rice Using Optimized Spectral Indices and Continuous Wavelet Transform." *Frontiers in Plant Science* 14. https://doi.org/10.3389/fpls.2023.1185915.

Zhang, Z., D. Jiang, Q. Chang, Z. Zheng, X. Fu, K. Li, and H. Mo. 2023. "Estimation of Anthocyanins in Leaves of Trees with Apple Mosaic Disease Based on Hyperspectral Data." *Remote Sensing* 15 (7): 1732. https://doi.org/10.3390/rs15071732.

Zhao, J., B. Zhou, G. Wang, J. Ying, J. Liu, and Q. Chen. 2022. "Spectral Camouflage Characteristics and Recognition Ability of Targets Based on Visible/near-Infrared Hyperspectral Images." *Photonics* 9 (12): 957. https://doi.org/10.3390/photonics9120957.

Appendix

A. Detector configurations

This section provides the configurations of the detectors that produced the results given in Section 4.3. Table A1 shows the detector configurations that produced the AUCs in Table 7. Each cell corresponds to a ROC in Figures B1– B4 and an AUC in Table 7, which in turn corresponds to a detector, band stack configuration and camouflaged target class. The same applies to Table A1, which shows the configurations of the detectors that produced the results in Table 8. In addition, the thresholds of each detector, band stack configuration and camouflaged target class required for the IOU results are given in the same manner in Table A3.

Table A1. Configurations that produced the ROC and AUC results for each detector, band stack configuration and camouflaged target class. The configurations show the inner and outer window sizes (w_i, w_o) for the LRX and LPD detectors, and the area size κ of the connected components to be removed for the AED detector.

	Hedge	Turf	Gn. Trp.	Gn. 2D	Gn. 3D	Gn. Per.	Fleece	Gy. Trp.	Gy. 3D	Yw. 3D	Yw. Per.
LRX											
LRI_6	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(31, 91)	(31, 91)	(41, 121)	(21, 61)	(21, 61)	(11, 31)
LRI_2	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(31, 91)	(41, 121)	(21, 61)	(21, 61)	(11, 31)
LRI_{nd}	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(11, 31)	(41, 121)	(41, 121)	(21, 61)	(21, 61)	(11, 31)
LRIr	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(11, 31)	(31, 91)	(41, 121)	(21, 61)	(21, 61)	(11, 31)
raw	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(5, 15)	(31, 91)	(41, 121)	(21, 61)	(21, 61)	(11, 31)
LPD											
LRI_6	(11, 31)	(11, 31)	(11, 31)	(21, 61)	(21, 61)	(31, 91)	(41, 121)	(21, 61)	(11, 31)	(21, 61)	(11, 31)
LRI_2	(11, 31)	(11, 31)	(11, 31)	(11, 31)	(21, 61)	(31, 91)	(41, 121)	(31, 91)	(11, 31)	(21, 61)	(11, 31)
LRI_{nd}	(11, 31)	(11, 31)	(11, 31)	(11, 31)	(21, 61)	(5, 15)	(41, 121)	(21, 61)	(11, 31)	(21, 61)	(11, 31)
LRIr	(11, 31)	(11, 31)	(11, 31)	(21, 61)	(21, 61)	(5, 15)	(41, 121)	(21, 61)	(11, 31)	(21, 61)	(11, 31)
raw	(11, 31)	(11, 31)	(11, 31)	(11, 31)	(21, 61)	(31, 91)	(41, 121)	(21, 61)	(11, 31)	(21, 61)	(11, 31)
AED											
LRI_6	200	100	200	50	300	50	50	200	300	200	100
LRI_2	200	100	200	50	300	50	100	200	200	200	100
LRI_{nd}	200	100	200	50	300	50	50	200	300	200	100
LRIr	200	100	200	50	300	50	100	200	300	200	100
raw	200	100	200	50	200	50	100	200	300	200	100

Table A2. Configurations that produced the IOU results for each detector, band stack configuration and camouflaged target class. The configurations show the inner and outer window sizes (w_i, w_o) for the LRX and LPD detectors, and the area size κ of the connected components to be removed for the AED detector.

	Hedge	Turf	Gn. Trp.	Gn. 2D	Gn. 3D	Gn. Per.	Fleece	Gy. Trp.	Gy. 3D	Yw. 3D	Yw. Per.
LRX											
LRI_6	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(31, 91)	(11, 31)	(41, 121)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
LRI_2	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(21, 61)	(11, 31)	(41, 121)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
LRI_{nd}	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(21, 61)	(11, 31)	(41, 121)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
LRIr	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(21, 61)	(11, 31)	(41, 121)	(41, 121)	(11, 31)	(5, 15)	(11, 31)
raw	(21, 61)	(21, 61)	(21, 61)	(31, 91)	(21, 61)	(11, 31)	(5, 15)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
LPD											
LRI_6	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(5, 15)	(41, 121)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
LRI_2	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(11, 31)	(41, 121)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
LRI_{nd}	(11, 31)	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(11, 31)	(41, 121)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
LRIr	(11, 31)	(11, 31)	(21, 61)	(21, 61)	(21, 61)	(5, 15)	(41, 121)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
raw	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(21, 61)	(11, 31)	(41, 121)	(41, 121)	(11, 31)	(11, 31)	(11, 31)
AED											
LRI_6	200	100	200	50	200	25	25	200	300	50	25
LRI_2	200	100	200	50	300	25	25	200	300	100	25
LRI_{nd}	200	100	200	50	300	25	300	200	300	50	25
LRIr	200	100	200	50	200	25	25	200	100	50	25
raw	200	100	200	50	300	25	300	200	300	100	25

Table A3. Thresholds that produced the IOU results for each detector, band stack configuration and camouflaged target class.

	Hedge	Turf	Gn. Trp.	Gn. 2D	Gn. 3D	Gn. Per.	Fleece	Gy. Trp.	Gy. 3D	Yw. 3D	Yw. Per.
RX											
LRI_6	99.6875	99.8750	99.7500	98.6250	99.2500	99.8125	99.8750	99.8125	99.8125	99.8125	99.7500
LRI_2	99.6875	99.8750	99.6875	98.2500	99.3125	99.8125	99.9375	99.8125	99.6250	99.6250	99.6875
LRI_{nd}	99.7500	99.8750	99.7500	98.1250	99.1875	99.8125	99.7500	99.8125	99.6875	99.8750	99.7500
LRIr	99.6875	99.8750	99.7500	98.6875	98.9375	99.8750	99.9375	99.8125	99.7500	99.5625	99.6875
raw	99.6875	99.6875	99.6875	97.9375	96.6875	99.8750	99.1250	99.8125	99.7500	99.8125	99.8125
LRX											
LRI_6	99.5625	99.8750	99.8125	99.0625	99.6875	99.8750	99.9375	99.8125	99.7500	99.8125	99.8125
LRI_2	99.5625	99.8750	99.7500	98.8750	99.5000	99.8125	99.7500	99.8125	99.5625	99.6875	99.7500
LRI_{nd}	99.6250	99.8750	99.7500	98.8125	99.5000	99.8125	99.6250	99.8125	99.5000	99.9375	99.8125
LRIr	99.6875	99.8750	99.8125	99.0000	99.6250	99.8750	99.9375	99.8750	99.6875	99.6875	99.8125
raw	99.6875	99.8750	99.7500	98.8125	99.2500	99.8750	99.3125	99.8125	99.4375	99.8750	99.8125
LPD											
LRI_6	99.6250	99.8750	99.8125	99.1875	99.3125	99.5625	99.8750	99.8125	99.5000	99.8750	99.8750
LRI_2	99.7500	99.8750	99.7500	99.0000	99.5000	99.8750	99.9375	99.8125	98.6250	99.8750	99.8125
LRI_{nd}	98.6250	99.8750	99.7500	99.1250	99.3125	99.8750	99.8125	99.8125	98.5625	99.8750	99.8750
LRIr	98.8750	99.7500	99.8125	99.4375	99.6250	99.5625	99.9375	99.8125	99.3125	99.8750	99.8750
raw	99.5625	99.8750	99.7500	98.8750	99.3125	99.8750	99.8750	99.8125	97.2500	99.8750	99.8125
AED											
LRI_6	96.8750	99.8125	99.8125	98.5000	98.6875	99.5625	99.8125	99.8125	97.1250	99.8750	99.7500
LRI_2	98.1875	99.8125	99.7500	98.1875	98.4375	99.5625	99.4375	99.8125	95.7500	99.8750	99.6875
LRI_{nd}	93.0000	99.6875	99.8125	97.8750	98.8125	99.5625	98.6875	99.8125	93.0000	99.8750	99.5625
LRIr	93.0000	99.8750	99.8750	99.0000	98.8750	99.5625	99.9375	99.8125	97.0000	99.8750	99.6875
raw	94.0000	98.9375	99.6875	96.9375	98.9375	99.5000	98.5000	99.8125	93.0000	99.8750	99.6875

B. Detector Receiver operating characteristics

This section shows the ROCs of all detectors for each band stack configuration and camouflaged target class. These are obtained by threshold averaging over the entire test dataset as described in Section 4.3.2. The ROCs shown in Figure 7 in the Evaluation section are generated by threshold averaging all the ROCs in Figures B1– B4. Each of these figures shows the ROCs for each band stack configuration and camouflaged target class of a single detector, whose configurations are shown in Table A3 in the previous section.

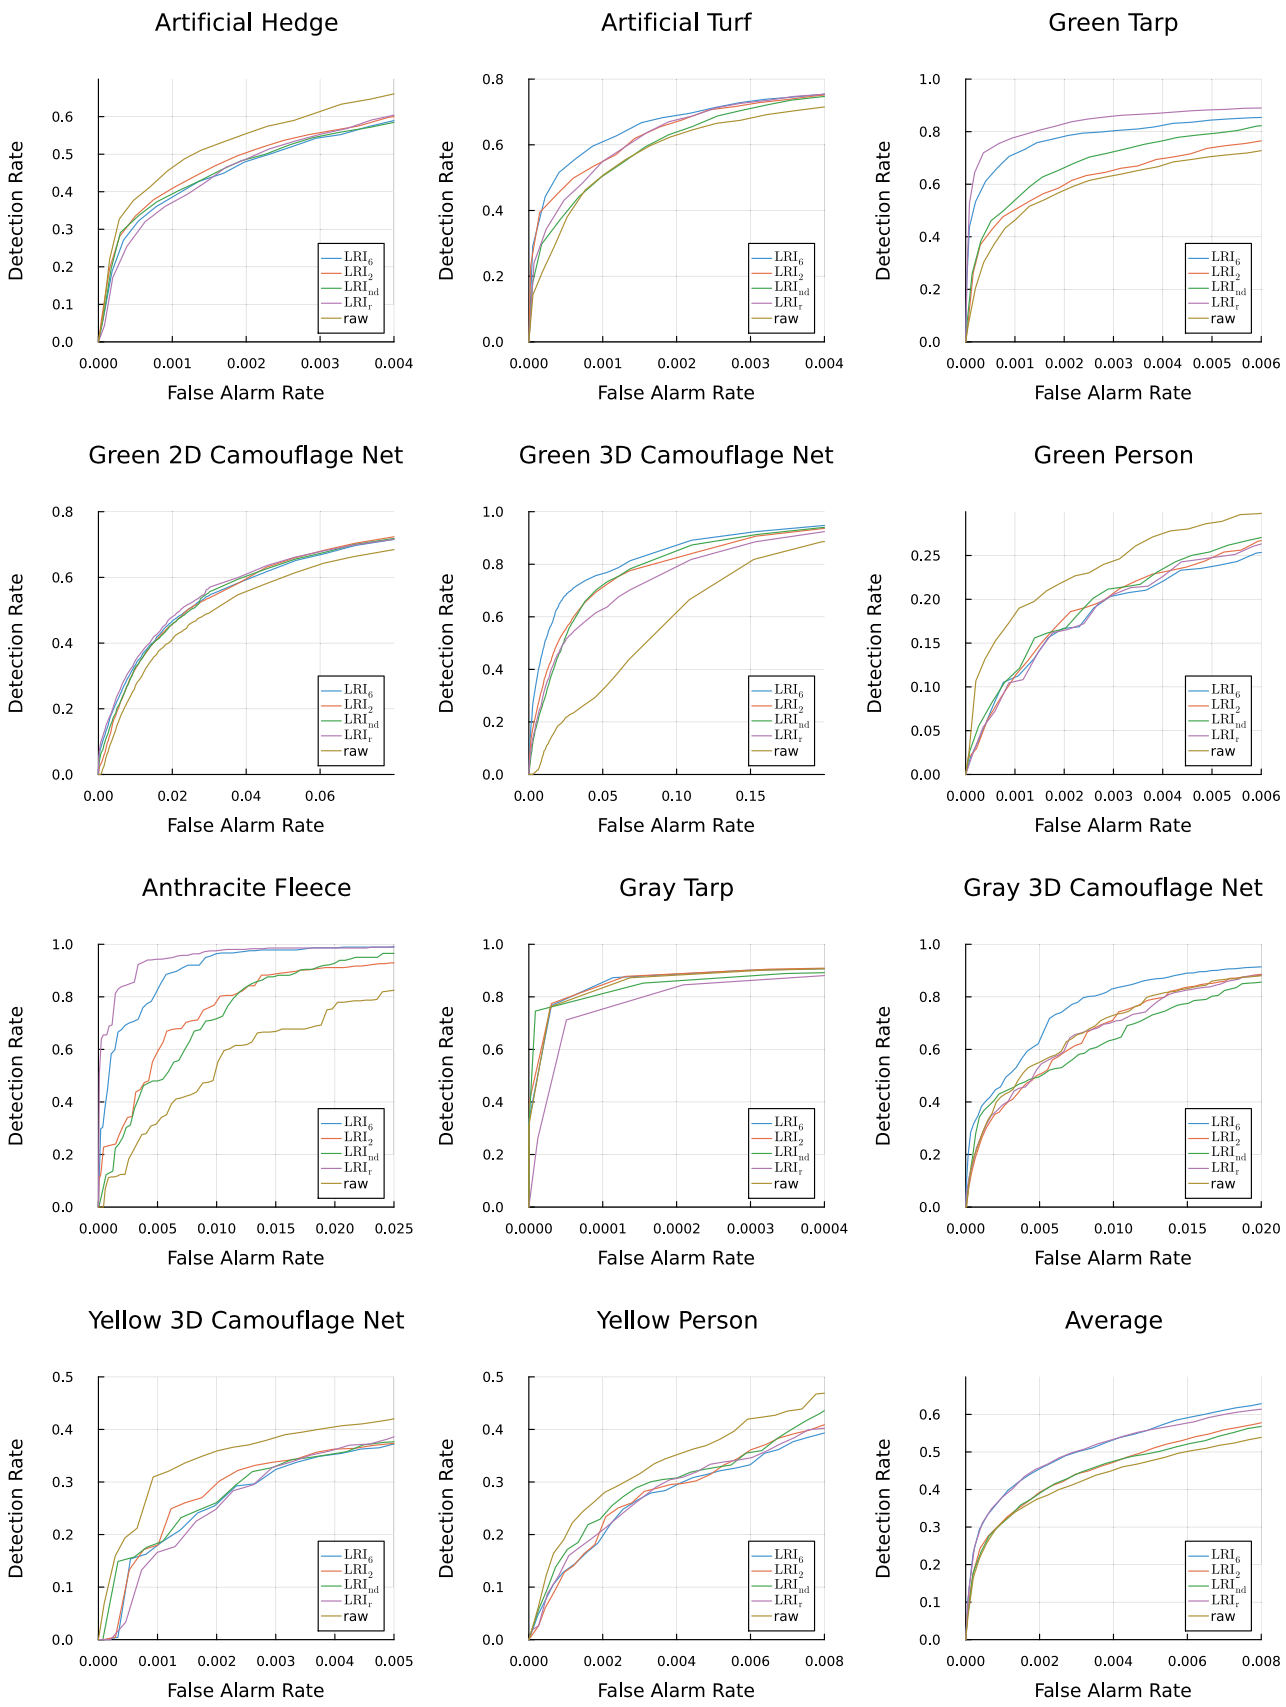

Figure A1. Receiver operating characteristics of the RX detector for each band stack configuration and camouflaged target class.

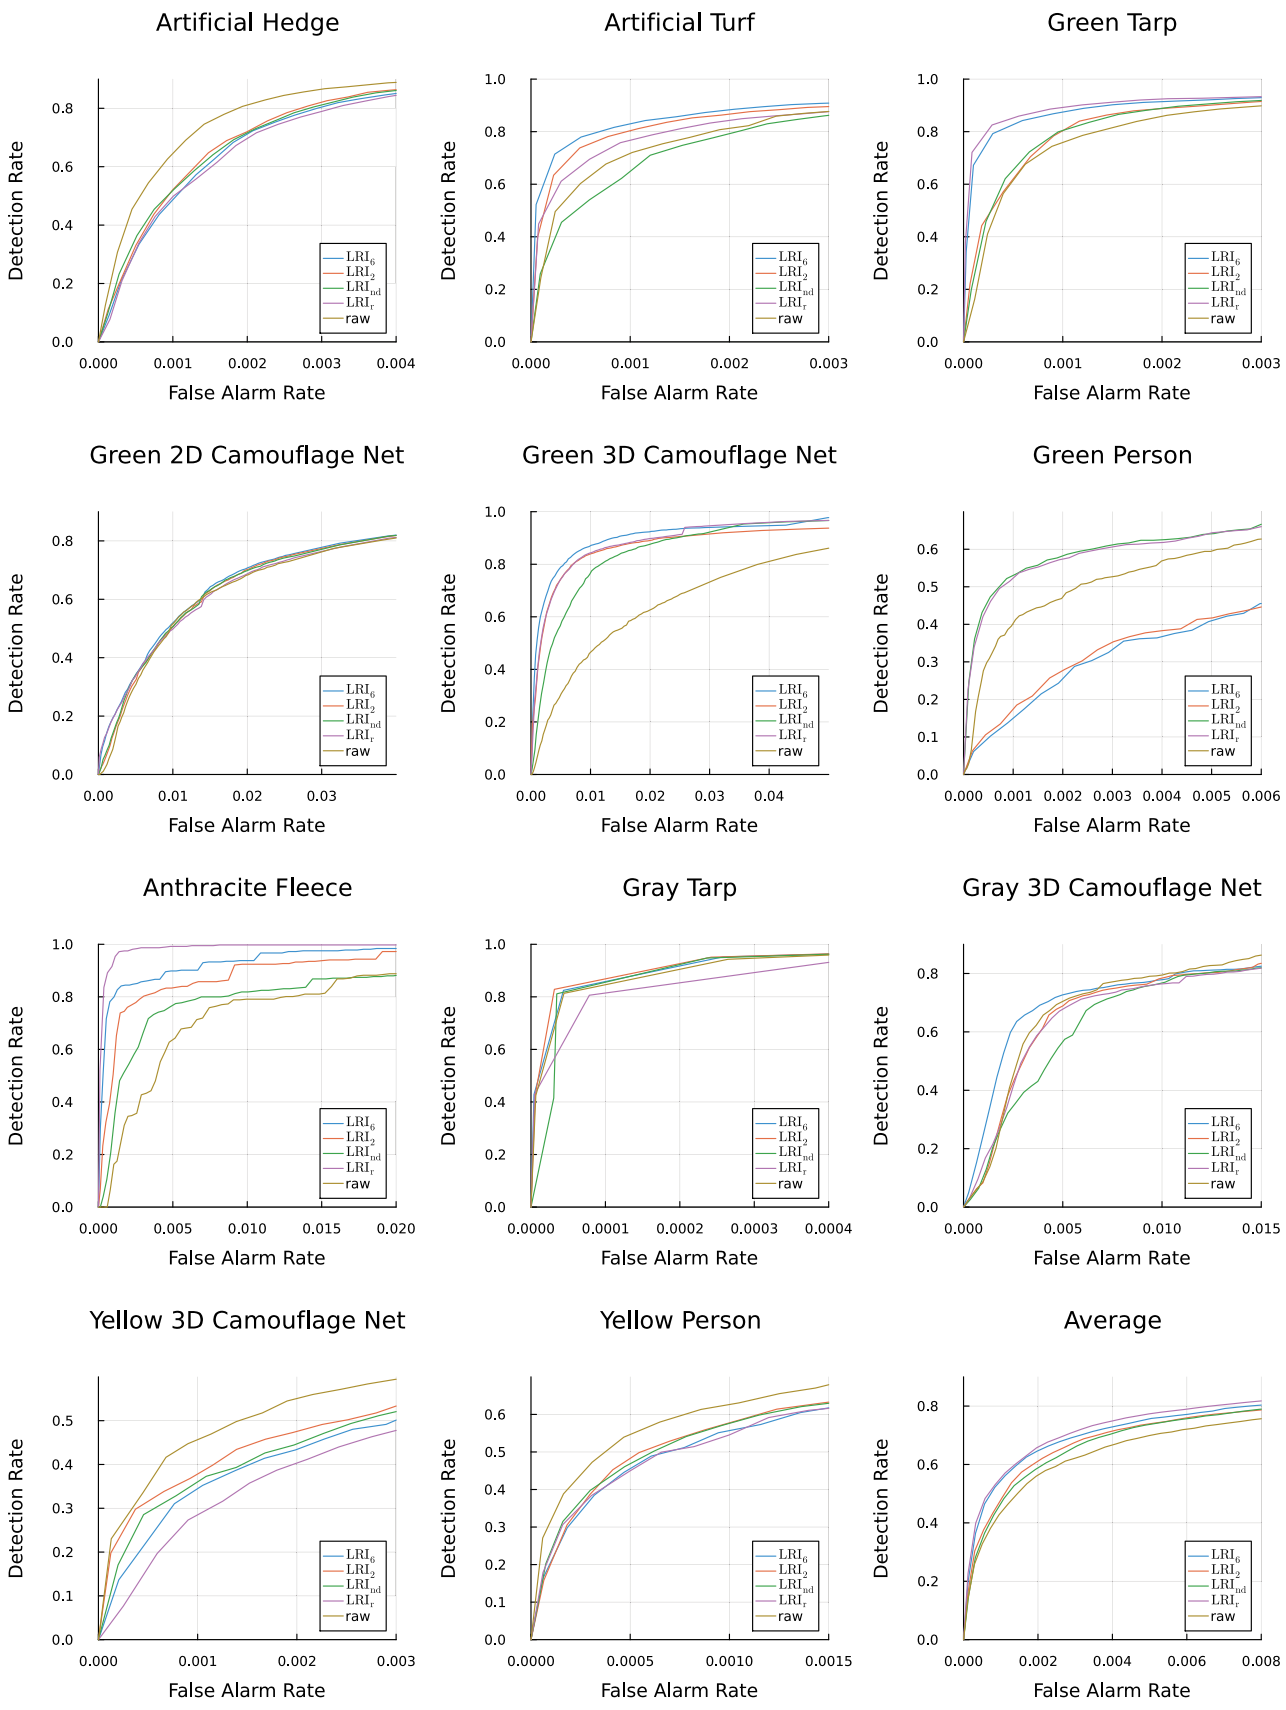

Figure A2. Receiver operating characteristics of the LRX detector for each band stack configuration and camouflaged target class.

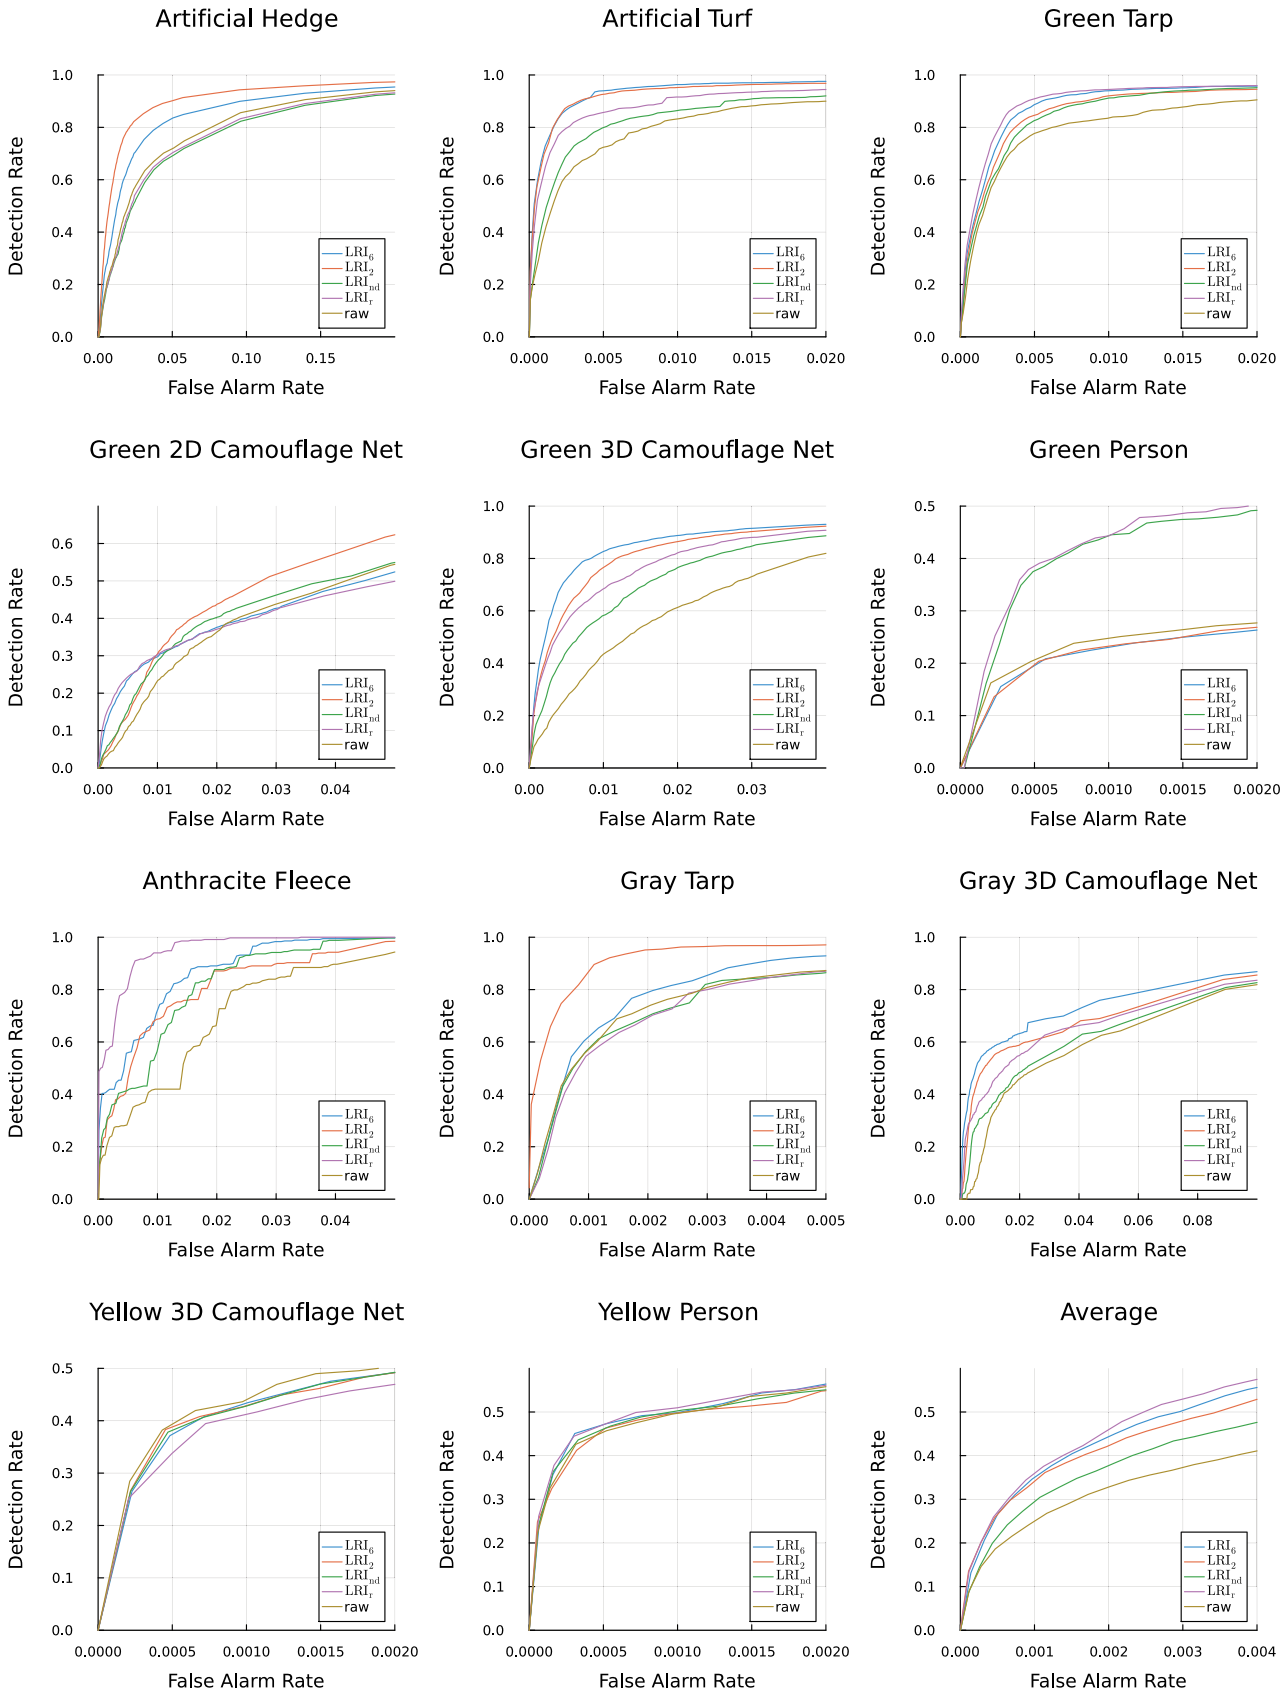

Figure A3. Receiver operating characteristics of the LPD detector for each band stack configuration and camouflaged target class.

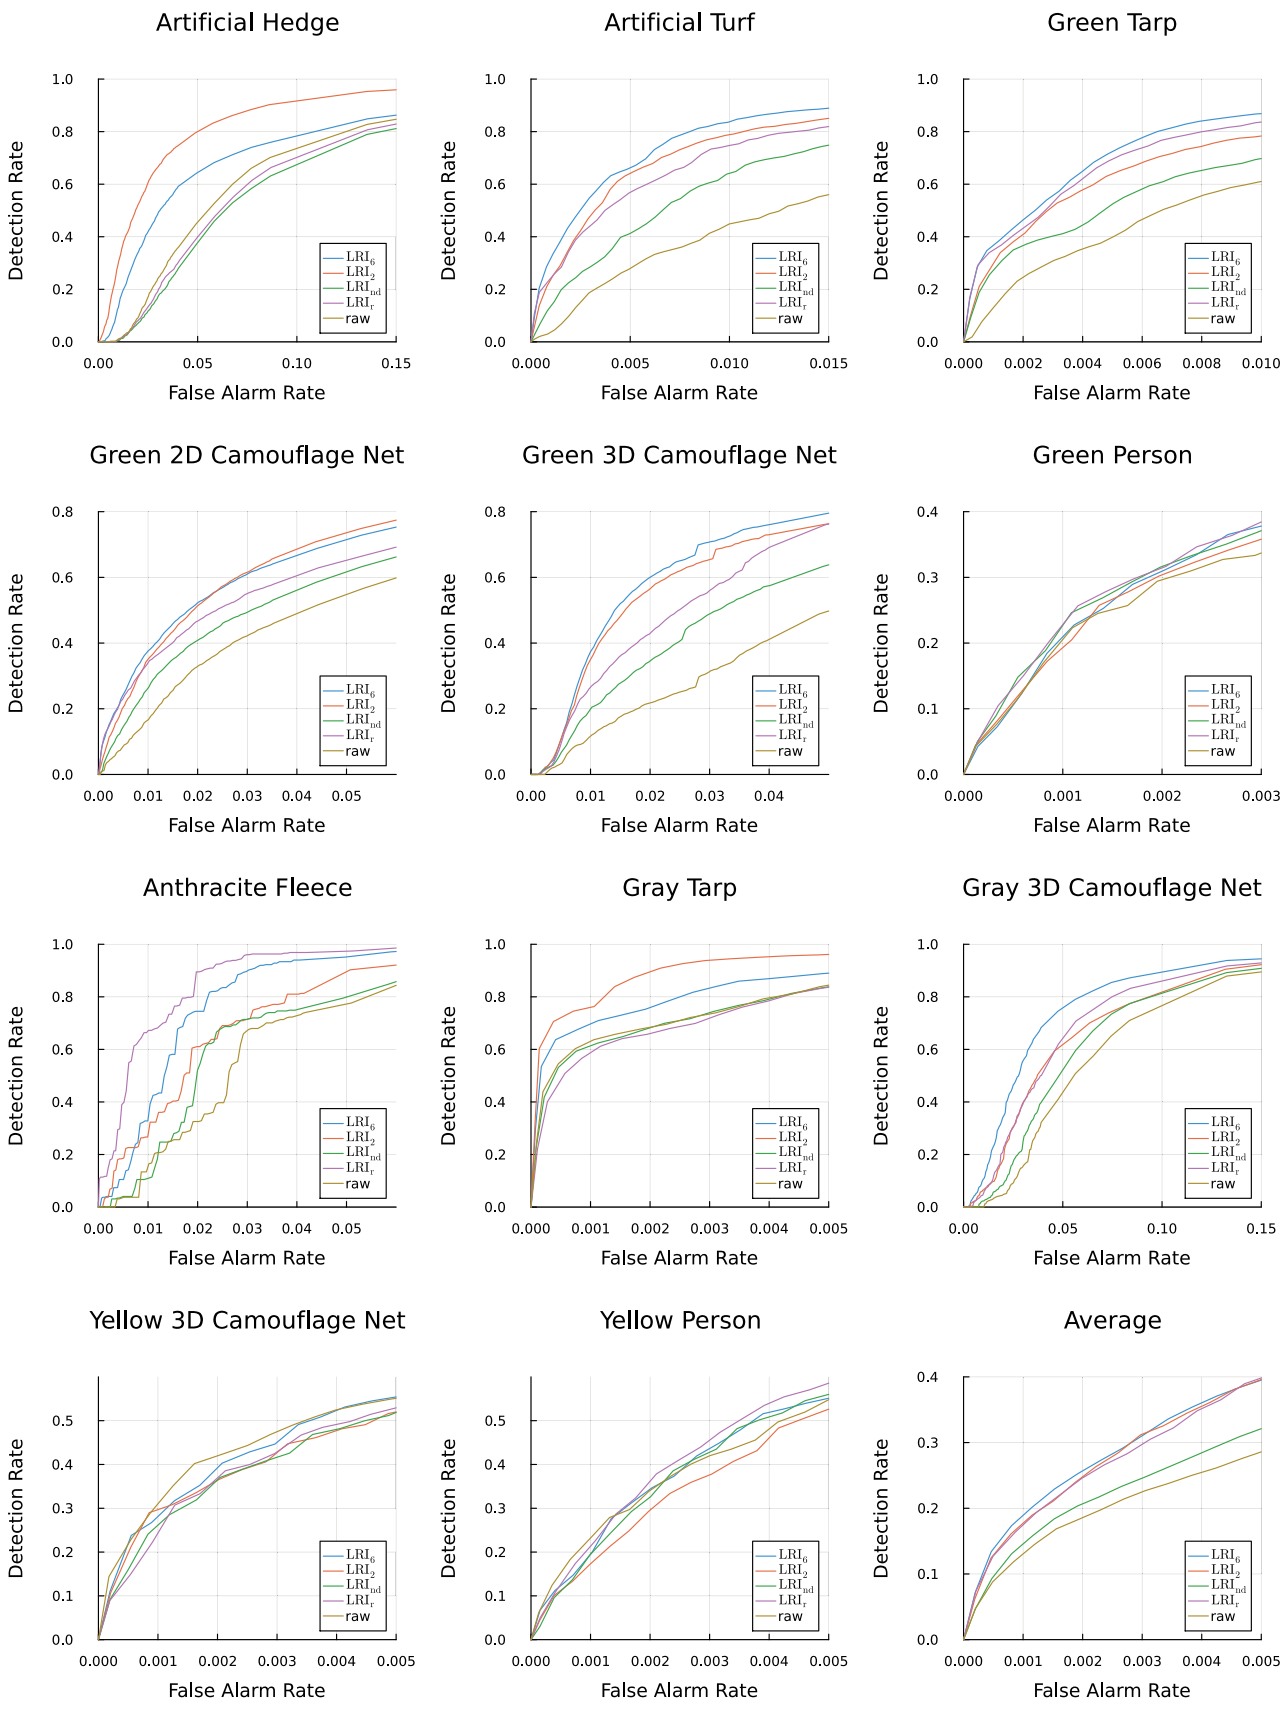

Figure A4. Receiver operating characteristics of the AED detector for each band stack configuration and camouflaged target class.

References

[1] Syed Moosa Ali, Anurag Gupta, Mini Raman, and Arvind Sahay. „Bio-Optical Characterization of Chilika Lagoon Using Multispectral Remote Sensing Data“. In: *Journal of the Indian Society of Remote Sensing* 51.8 (Aug. 2023), pp. 1729–1737. ISSN: 0974-3006. DOI: 10.1007/s12524-023-01716-x.

[2] Quoc Bao Pham, Sk Ajim Ali, Farhana Parvin, Vo Van On, Lariyah Mohd Sidek, Bojan Đurin, Vlado Cetl, Sanja Šamanović, and Nguyen Nguyet Minh. „Multi-Spectral Remote Sensing and GIS-Based Analysis for Decadal Land Use Land Cover Changes and Future Prediction Using Random Forest Tree and Artificial Neural Network“. In: *Advances in Space Research* 74.1 (July 2024), pp. 17–47. ISSN: 0273-1177. DOI: 10.1016/j.asr.2024.03.027.

[3] Léopold Brouant, Hong-Liên Pham, Éloi Foucher, and Aubry Picard. „VIS-NIR Multispectral for Camouflage Detection“. In: *Emerging Imaging and Sensing Technologies for Security and Defence VIII*. Vol. 12740. SPIE, Oct. 2023, pp. 22–31. DOI: 10.1117/12.2678661.

[4] Alan B. Cannaday, Curt H. Davis, and Trevor M. Bajkowski. „Detection of Camouflage-Covered Military Objects Using High-Resolution Multi-Spectral Satellite Imagery“. In: *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*. ISSN: 2153-7003. July 2023, pp. 5766–5769. DOI: 10.1109/IGARSS52108.2023.10281409.

[5] Chih-Chung Chang and Chih-Jen Lin. „LIBSVM : A Library for Support Vector Machines“. In: *ACM Transactions on Intelligent Systems and Technology* 2 (3 2011), 27:1–27:27. URL: http://www.csie.ntu.edu.tw/~cjlin/libsvm (visited on 03/10/2022).

[6] Tianqi Chen and Carlos Guestrin. „XGBoost: A Scalable Tree Boosting System“. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '16. San Francisco, California, USA: ACM, 2016, pp. 785–794. ISBN: 978-1-4503-4232-2. DOI: 10.1145/2939672.2939785.

[7] Esayas Elias Churko, Luxon Nhamo, and Munyaradzi Chitakira. „Multispectral Remote Sensing Approach of Predicting the Potential Distribution and Evaluating the Current Spread of Water Hyacinth (Eichhornia Crassipes)“. In: *Sustainable Water Resources Management* 10.1 (Jan. 2024), p. 35. ISSN: 2363-5045. DOI: 10.1007/s40899-023-01019-6.

[8] Bo Du and Liangpei Zhang. „A Discriminative Metric Learning Based Anomaly Detection Method“. In: *IEEE Transactions on Geoscience and Remote Sensing* 52.11 (Nov. 2014), pp. 6844–6857. ISSN: 1558-0644. DOI: 10.1109/TGRS.2014.2303895.

[9] Linda Eckel and Peter Stütz. „Hyperspectral Sensor Management for UAS: Performance Analysis of Context-Based System Architectures for Camouflage and UXO Anomaly Detection Workflows“. In: *Drones* 8.10 (Oct. 2024), p. 529. ISSN: 2504-446X. DOI: 10.3390/drones8100529.

[10] Linda Eckel and Peter Stütz. „Hyperspectral Sensor Management for UAS: Sensor Context Based Band Selection for Anomaly Detection“. In: *2024 IEEE Aerospace Conference*. ISSN: 1095-323X. Mar. 2024, pp. 1–14. DOI: 10.1109/AERO58975.2024.10521225.

[11] Tom Fawcett. „An Introduction to ROC Analysis“. In: *Pattern Recognition Letters* 27.8 (June 2006), pp. 861–874. ISSN: 0167-8655. DOI: 10.1016/J.PATREC.2005.10.010.

[12] Éloi Foucher, Léopold Brouant, and Aubry Picard. „VIS-NIR Multispectral for Camouflage Detection: Campaign Analysis“. In: *Emerging Imaging and Sensing Technologies for Security and Defence IX*. Vol. 13204. SPIE, Nov. 2024, pp. 21–29. DOI: 10.1117/12.3033586.

[13] Yuanyuan Fu, Guijun Yang, Ruiliang Pu, Zhenhai Li, Heli Li, Xingang Xu, Xiaoyu Song, Xiaodong Yang, and Chunjiang Zhao. „An Overview of Crop Nitrogen Status Assessment Using Hyperspectral Remote Sensing: Current Status and Perspectives“. In: *European Journal of Agronomy* 124 (Mar. 2021), p. 126241. ISSN: 1161-0301. DOI: 10.1016/j.eja.2021.126241.

[14] Xin Han, Zheng Wei, He Chen, Baozhong Zhang, Yinong Li, and Taisheng Du. „Inversion of Winter Wheat Growth Parameters and Yield under Different Water Treatments Based on Uav Multispectral Remote Sensing“. In: *Frontiers in Plant Science* 12 (2021). ISSN: 1664-462X. DOI: 10.3389/fpls.2021.609876.

[15] Robert M. Haralick, Its'hak Dinstein, and K. Shanmugam. „Textural Features for Image Classification“. In: *IEEE Transactions on Systems, Man and Cybernetics* SMC-3.6 (Nov. 1973), pp. 610–621. ISSN: 2168-2909. DOI: 10.1109/TSMC.1973.4309314.

[16] Md Zahid Hasan, Rabeya Sultana Leya, and Kazi Saiful Islam. „Comparative Assessment of Machine Learning Algorithms for Land Use and Land Cover Classification Using Multispectral Remote Sensing Image“. In: *Khulna University Studies* (Oct. 2022), pp. 33–46. ISSN: 2789-2697. DOI: 10.53808/KUS.2022.ICSTEM4IR.0124-se.

[17] Christian Hellert, Simon Koch, and Peter Stütz. „Using Algorithm Selection for Adaptive Vehicle Perception Aboard UAV“. In: *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. ISSN: 2643-6213. Sept. 2019, pp. 1–8. DOI: 10.1109/AVSS.2019.8909862.

[18] Tobias Hupel and Peter Stütz. „Adopting Hyperspectral Anomaly Detection for near Real-Time Camouflage Detection in Multispectral Imagery“. In: *Remote Sensing* 14.15 (2022). ISSN: 2072-4292. DOI: 10.3390/rs14153755.

[19] Tobias Hupel and Peter Stütz. „Measuring and Predicting Sensor Performance for Camouflage Detection in Multispectral Imagery“. In: *Sensors* 23.19 (2023). ISSN: 1424-8220. DOI: 10.3390/s23198025.

[20] Tobias Hupel and Peter Stütz. „Optimized Spectral Indices for Camouflage Detection in Multispectral Imagery“. In: *GIScience & Remote Sensing* 62.1 (2025), p. 2508574. DOI: 10.1080/15481603.2025.2508574.

[21] Tobias Hupel and Peter Stütz. „Sensor-Managed Anomaly Detection for Camouflage Detection in Airborne Multispectral Imagery“. In: *2024 IEEE Aerospace Conference*. ISSN: 1095-323X. Mar. 2024, pp. 1–11. DOI: 10.1109/AERO58975.2024.10521231.

[22] Xudong Kang, Xiangping Zhang, Shutao Li, Kenli Li, Jun Li, and Jón Atli Benediktsson. „Hyperspectral Anomaly Detection with Attribute and Edge-Preserving Filters“. In: *IEEE Transactions on Geoscience and Remote Sensing* 55.10 (Oct. 2017), pp. 5600–5611. ISSN: 0196-2892. DOI: 10.1109/TGRS.2017.2710145.

[23] Simon Koch, Bernhard Krach, Fotios Katsilieris, and Peter Stütz. „Sensor Scheduling Strategies for 1-To-N Multi-Object Tracking“. In: *2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC)*. ISSN: 2155-7209. Sept. 2022, pp. 1–9. DOI: 10.1109/DASC55683.2022.9925744.

[24] Simon Koch, Johannes Ostler, and Peter Stütz. „Estimating Object Perception Performance in Aerial Imagery Using a Bayesian Approach“. In: *2023 IEEE Aerospace Conference*. ISSN: 1095-323X. Mar. 2023, pp. 1–9. DOI: 10.1109/AERO55745.2023.10115828.

[25] Wei Li and Qian Du. „Collaborative Representation for Hyperspectral Anomaly Detection“. In: *IEEE Transactions on Geoscience and Remote Sensing* 53.3 (Mar. 2015), pp. 1463–1474. ISSN: 0196-2892. DOI: 10.1109/TGRS.2014.2343955.

[26] Xingshi Luo, Wenzheng Wang, and Chenwei Deng. „Cross-AD: Multispectral and Hyperspectral High-Speed Artificial Imitation Object Anomaly Detection“. In: *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*. ISSN: 2153-7003. July 2023, pp. 1142–1145. DOI: 10.1109/IGARSS52108.2023.10282235.

[27] Donghui Ma, Liguang Jiang, Jie Li, and Yun Shi. „Water Index and Swin Transformer Ensemble (WISTE) for Water Body Extraction from Multispectral Remote Sensing Images“. In: *GIScience & Remote Sensing* 60.1 (Dec. 2023), p. 2251704. ISSN: 1548-1603. DOI: 10.1080/15481603.2023.2251704.

[28] Wouter H. Maes and Kathy Steppe. „Perspectives for Remote Sensing with Unmanned Aerial Vehicles in Precision Agriculture“. In: *Trends in Plant Science* 24.2 (Feb. 2019), pp. 152–164. ISSN: 1360-1385. DOI: 10.1016/j.tplants.2018.11.007.

[29] Sadiq Al-Maliki, Taha I. M. Ibrahim, Gusztáv Jakab, Malihe Masoudi, Jamal S. Makki, and Zoltán Vekerdy. „An Approach for Monitoring and Classifying Marshlands Using Multispectral Remote Sensing Imagery in Arid and Semi-Arid Regions“. In: *Water* 14.10 (Jan. 2022), p. 1523. ISSN: 2073-4441. DOI: 10.3390/w14101523.

[30] Archana Nandibewoor, Peddanna Doni, Holagundi Rashmi, and Abhilash Hegde. „Anomaly Detection with Camouflage Reconnaissance in Spectral Imaging“. In: *Research Square* (Aug. 2024). ISSN: 2693-5015. DOI: 10.21203/rs.3.rs-4715779/v1.

[31] Farrukh Aziz Bhatti Muhammad Khizer Ali Khurram Khurshid Moazam Maqsood Noman Raza Shah Abdur Rahman M. Maud and Muhammad Amin. „Hyperspectral Anomaly Detection: A Performance Comparison of Existing Techniques“. In: *International Journal of Digital Earth* 15.1 (2022), pp. 2078–2125. DOI: 10.1080/17538947.2022.2146770.

[32] Timo Ojala, Matti Pietikäinen, and Topi Mäenpää. „Gray Scale and Rotation Invariant Texture Classification with Local Binary Patterns“. In: *Computer Vision - ECCV 2000*. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2000, pp. 404–420. ISBN: 9783540450542. DOI: 10.1007/3-540-45054-8_27.

[33] Manish Kumar Patel, Dongryeol Ryu, Andrew W. Western, Helen Suter, and Iain M. Young. „Which Multispectral Indices Robustly Measure Canopy Nitrogen across Seasons: Lessons from an Irrigated Pasture Crop“. In: *Computers and Electronics in Agriculture* 182 (Mar. 2021), p. 106000. ISSN: 0168-1699. DOI: 10.1016/j.compag.2021.106000.

[34] Ivan Racetin and Andrija Krtalić. „Systematic Review of Anomaly Detection in Hyperspectral Remote Sensing Applications“. In: *Applied Sciences* 11.11 (Jan. 2021), p. 4878. ISSN: 2076-3417. DOI: 10.3390/app11114878.

[35] Irving S. Reed and Xiaoli Yu. „Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution“. In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 38.10 (1990), pp. 1760–1770. ISSN: 0096-3518. DOI: 10.1109/29.60107.

[36] J. W. Rouse, R. H. Haas, J. A. Schell, and D. W. Deering. „Monitoring Vegetation Systems in the Great Plains with ERTS“. In: *Goddard Space Flight Center 3d ERTS-1 Symp.* Vol. 1. NTRS Author Affiliations: Texas A&M Univ. NTRS Report/Patent Number: PAPER-A20 NTRS Document ID: 19740022614 NTRS Research Center: Legacy CDMS (CDMS). NASA, Jan. 1974. URL: https://ntrs.nasa.gov/citations/19740022614 (visited on 08/13/2024).

[37] Martin Ruß and Peter Stütz. „Airborne Sensor and Perception Management“. In: *Modelling and Simulation for Autonomous Systems*. Ed. by Jan Mazal, Adriano Fagiolini, Petr Vašík, Agostino Bruzzone, Stefan Pickl, Vlastimil Neumann, Petr Stodola, and Stefano Lo Storto. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2023, pp. 182–206. ISBN: 9783031312687. DOI: 10.1007/978-3-031-31268-7_11.

[38] Ben Sadeghi, Poom Chiarawongse, Kevin Squire, Daniel C. Jones, Andreas Noack, Cédric St-Jean, Rik Huijzer, Roland Schätzle, Ian Butterworth, Yu-Fong Peng, and Anthony Blaom. *DecisionTree.jl - a Julia Implementation of the Cart Decision Tree and Random Forest Algorithms*. Version 0.11.3. Nov. 2022. DOI: 10.5281/zenodo.7359268.

[39] Ying Shen, Jie Li, Wenfu Lin, Liqiong Chen, Feng Huang, and Shu Wang. „Camouflaged Target Detection Based on Snapshot Multispectral Imaging“. In: *Remote Sensing* 13.19 (2021), p. 3949. ISSN: 2072-4292. DOI: 10.3390/rs13193949.

[40] Lee Spector and Jon Klein. „Trivial Geography in Genetic Programming“. In: *Genetic Programming Theory and Practice III*. Ed. by Tina Yu, Rick Riolo, and Bill Worzel. Boston, MA: Springer US, 2006, pp. 109–123. ISBN: 978-0-387-28111-7. DOI: 10.1007/0-387-28111-8_8.

[41] Rainer Storn and Kenneth Price. „Differential Evolution – A Simple and Efficient Heuristic for Global Optimization Over Continuous Spaces“. In: *Journal of Global Optimization* 11.4 (Dec. 1997), pp. 341–359. ISSN: 1573-2916. DOI: 10.1023/A:1008202821328.

[42] Hongjun Su, Zhaoyue Wu, Huihui Zhang, and Qian Du. „Hyperspectral Anomaly Detection: A Survey“. In: *IEEE Geoscience and Remote Sensing Magazine* 10.1 (Mar. 2022), pp. 64–90. ISSN: 2168-6831. DOI: 10.1109/MGRS.2021.3105440.

[43] S. Thakur, I. Mondal, P. B. Ghosh, P. Das, and T. K. De. „A Review of the Application of Multispectral Remote Sensing in the Study of Mangrove Ecosystems with Special Emphasis on Image Processing Techniques“. In: *Spatial Information Research* 28.1 (Feb. 2020), pp. 39–51. ISSN: 2366-3294. DOI: 10.1007/s41324-019-00268-y.

[44] Minghua Wang, Qiang Wang, Danfeng Hong, Swalpa Kumar Roy, and Jocelyn Chanussot. „Learning Tensor Low-Rank Representation for Hyperspectral Anomaly Detection“. In: *IEEE Transactions on Cybernetics* 53.1 (Jan. 2023), pp. 679–691. ISSN: 2168-2275. DOI: 10.1109/TCYB.2022.3175771.

[45] Shaoyu Wang, Xinyu Wang, Liangpei Zhang, and Yanfei Zhong. „Auto-AD: Autonomous Hyperspectral Anomaly Detection Network Based on Fully Convolutional Autoencoder“. In: *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022), pp. 1–14. ISSN: 1558-0644. DOI: 10.1109/TGRS.2021.3057721.

[46] Shu Wang, Yixuan Xu, Dawei Zeng, Feng Huang, and Lingyu Liang. „Deep Learning-Based Spectral Reconstruction in Camouflaged Target Detection“. In: *International Journal of Applied Earth Observation and Geoinformation* 126 (Feb. 2024), p. 103645. ISSN: 1569-8432. DOI: 10.1016/j.jag.2023.103645.

[47] Yong Wang, Han-Xiong Li, Tingwen Huang, and Long Li. „Differential Evolution Based on Covariance Matrix Learning and Bimodal Distribution Parameter Setting“. In: *Applied Soft Computing* 18 (May 2014), pp. 232–247. ISSN: 1568-4946. DOI: 10.1016/j.asoc.2014.01.038.

[48] Yu Wang, Chunhui Feng, Yiru Ma, Xiangyu Chen, Bin Lu, Yan Song, Ze Zhang, and Rui Zhang. „Estimation of Nitrogen Concentration in Walnut Canopies in Southern Xinjiang Based on UAV Multispectral Images“. In: *Agronomy* 13.6 (June 2023), p. 1604. ISSN: 2073-4395. DOI: 10.3390/agronomy13061604.

[49] Yang Xu, Zebin Wu, Jun Li, Antonio Plaza, and Zhihui Wei. „Anomaly Detection in Hyperspectral Images Based on Low-Rank and Sparse Representation“. In: *IEEE Transactions on Geoscience and Remote Sensing* 54.4 (Apr. 2016), pp. 1990–2000. ISSN: 0196-2892. DOI: 10.1109/TGRS.2015.2493201.

[50] Zhanghua Xu, Qi Zhang, Songyang Xiang, Yifan Li, Xuying Huang, Yiwei Zhang, Xin Zhou, Zenglu Li, Xiong Yao, Qiaosi Li, and Xiaoyu Guo. „Monitoring the Severity of Pantana Phyllostachysae Chao Infestation in Moso Bamboo Forests Based on UAV Multi-Spectral Remote Sensing Feature Selection“. In: *Forests* 13.3 (Mar. 2022), p. 418. ISSN: 1999-4907. DOI: 10.3390/f13030418.

[51] Yong Yan, Ying Wang, Cheng Yu, and Zhimin Zhang. „Multispectral Remote Sensing for Estimating Water Quality Parameters: A Comparative Study of Inversion Methods Using Unmanned Aerial Vehicles (UAVs)“. In: *Sustainability* 15.13 (Jan. 2023), p. 10298. ISSN: 2071-1050. DOI: 10.3390/su151310298.

[52] Xuewei Zhang, Kefei Zhang, Yaqin Sun, Yindi Zhao, Huifu Zhuang, Wei Ban, Yu Chen, Erjiang Fu, Shuo Chen, Jinxiang Liu, and Yumeng Hao. „Combining Spectral and Texture Features of UAS-Based Multispectral Images for Maize Leaf Area Index Estimation“. In: *Remote Sensing* 14.2 (Jan. 2022), p. 331. ISSN: 2072-4292. DOI: 10.3390/rs14020331.

[53] Zhengxin Zhang and Lixue Zhu. „A Review on Unmanned Aerial Vehicle Remote Sensing: Platforms, Sensors, Data Processing Methods, and Applications“. In: *Drones* 7.6 (June 2023), p. 398. ISSN: 2504-446X. DOI: 10.3390/drones7060398.

[54] Junke Zhu, Yumeng Li, Chunying Wang, Ping Liu, and Yubin Lan. „Method for Monitoring Wheat Growth Status and Estimating Yield Based on UAV Multispectral Remote Sensing“. In: *Agronomy* 14.5 (May 2024), p. 991. ISSN: 2073-4395. DOI: 10.3390/agronomy14050991.

[55] Markus Zwick, Matthias Gerdts, and Peter Stütz. „Sensor-Model-Based Trajectory Optimization for Uavs to Enhance Detection Performance: An Optimal Control Approach and Experimental Results“. In: *Sensors* 23.2 (2023). ISSN: 1424-8220. DOI: 10.3390/s23020664.