Integrity Assessment of Maritime Object Detection Impacted by Partial Camera Obstruction

Felipe A. Costa de Oliveira*, Borja Carrillo-Perez*, Alberto García-Ortiz†, Frank Sill Torres*
*Institute for the Protection of Maritime Infrastructures (MI)
German Aerospace Center (DLR), Bremerhaven, Germany
Email: felipe.costadeoliveira@dlr.de
†Institute of Electrodynamics and Microelectronics
University of Bremen, Bremen, Germany
Email: agarcia@item.uni-bremen.de

2023 7th International Conference on System Reliability and Safety (ICSRS) | 979-8-3503-0605-7/23/$31.00 ©2023 IEEE | DOI: 10.1109/ICSRS59833.2023.10381255

***Abstract*—The performance and usage of machine learning based object detection in visual data has increased significantly in the past decade. This technology can enable automated decision making in various applications, extracting key information in real time from a camera-based monitoring solution. However, to ensure the resilience and dependability in a safety-critical application, the intelligence provided by the monitoring solution must be reliable. In this context, understanding the impact of potential disturbances in the object detection performance is important. This work conducts an integrity assessment of maritime object detection under partial camera obstruction events. A subset of the ShipSG dataset, which contains thousands of annotated and segmented ship images from a harbor, and the Faster-RCNN object detection algorithm, trained on the seven different ship classes of the dataset, were used. The effect of simulated obstructions, of various intensities and configurations, on the false-positive, misclassification, false negative ratios, and associated detection score distributions were investigated. The outcome suggests that the use of a partial obstruction detection step, and the consideration of that information, can mitigate the consequences of faults in the object detection.**

***Index Terms*—statistical analysis, object detection, integrity monitoring, fault detection, information fusion, image processing.**

I. Introduction

The advancement of the computer vision field, with sophisticated image processing techniques and increasingly accurate machine learning based solutions, has fostered the use of cameras in sensor systems for various applications. With the increased computational power capabilities enabled by modern electronics, applications that require intensive image processing in real-time, such as deep learning-based object detection, are becoming a reality. As the state-of-the-art algorithms for image processing are reaching, and sometimes surpassing, the level of human visual pattern recognition [1], the use of camera intelligence for automation in various applications will become ubiquitous. However, as our reliance on these systems increases, in particular for safety-critical applications such as autonomous vehicles and monitoring of strategic infrastructure, the integrity assessment of the information provided by the cameras becomes indispensable. In this context, the concept of sensor integrity monitoring can be extended to encompass not only the accuracy of the raw sensor measurements but also the reliability of the intelligence extracted from them, obtained with higher-level processing and additional models.

This work was funded by a doctorate scholarship provided by the Deutscher Akademischer Austauschdienst (DAAD), number 57540125.

Intelligence can be seen as the ability to make the right decisions at the right times. The goal of a reliable monitoring solution, that might use cameras as a part of the system, is essentially the provision of intelligence for a given application. However, sensors can malfunction and are affected by external factors. Therefore, it is important to understand how certain scenarios could impact the intelligence derived from a sensor, and to have methods to detect these events in case they have a detrimental effect in the reliability of the information. Aligned with this context, the present work focuses on a specific case study scenario in the maritime domain, of a monitoring system comprised of a single camera and a ship detection algorithm. There are various events that could impair the intelligence provided by that system and a comprehensive integrity assessment would have to consider most of those. This work proposes a step towards that comprehensive integrity assessment, with the investigation of the effects that partial camera obstruction have on ship detection performance. That scenario occurs when there is something in close proximity to the camera lens, partially obstructing the view. Since that event is not part of a normal operation and can have a detrimental impact on the system, in this work, the partial camera obstruction will be considered as a type of fault. To investigate the effects of that phenomenon in object detection, simulated obstructions of various configurations and intensities were added to the bounding boxes of the ships from the ShipSG dataset [2]. The Faster-RCNN algorithm, trained with that dataset, was used and the impact of the obstructions in the detection results was thoroughly analyzed.

II. Background Information

A. Sensor Integrity Concept

The concept of integrity is used in various fields among the engineering and computer science disciplines, often having a different meaning depending on the context and application [3]. Despite having different meanings, integrity is fundamentally associated with trust. However, the metrics to quantify

that trust level can vary depending on the application. On a sensor level, the integrity can be defined by the confidence level associated with an accuracy range for a measurement. For navigation applications, the positional integrity is often represented by a protection level, a bound for which the true value of the position is guaranteed to be with a high confidence level. For a system or a process, the integrity can be related to the correct functioning, given by specific performance indicators, or by complying with certain requirements and tests. In either case, the typical approach for integrity monitoring is through statistical analysis.

Most sensor integrity monitoring schemes rely on applying an estimation technique (E.g. Least Squares Estimator [4] or Kalman Filter [5]) to one or multiple sensor measurements, using a suitable model to make predictions with the estimator and to characterize the error distribution. Next, a fault detection scheme based on the results from the estimation is used, flagging erroneous measurements and triggering an integrity alert if necessary. In this way, the integrity assessment can be viewed as the process of detecting the faults under consideration and providing a statistical measure for the confidence in the assessment outcome.

A typical approach for fault detection is with statistical testing. It requires:

- the use of an appropriate metric or parameter that can indicate the fault;
- the choice of a suitable test, such as the Chi-Squared Test [6] or the Generalized Likelihood Ratio Test (GLRT) [7];
- the calculation of a test statistic from the relevant metric;
- comparison with a threshold value that would dictate the result of the test (if a fault is detected or not).

Assuming that the test statistic is a random variable and that its distribution is conditioned on the existence or not of a given fault, the metric used to form the statistic would be an effective detector if there is a clear distinction between the distribution in the no-fault (null hypothesis) and faulty (alternative hypothesis) case. If those distributions are known or can be estimated, for instance, through historical data, the fault detection procedure with the aforementioned process is straightforward.

With the fault detection approach, the idea of integrity monitoring can be applied to any kind of sensor system. However, it is important to note that that concept has been mostly associated with Global Navigation Satellite System (GNSS) measurements. Traditionally, the integrity monitoring methods have been developed for navigation applications, led by the critical positional integrity requirements for aerospace applications [8]. That concept has been extensively researched in that field but has not been sufficiently explored in other areas. In this work, we propose using the term integrity monitoring for a camera system, including faults from the physical and software domains in the integrity analysis.

In this work, the system under consideration consists of an optical camera and a ship detection algorithm. The proposed integrity assessment will focus on the impact of a partial camera obstruction fault on the object detection results.

B. Integrity Monitoring in Optical Cameras

There has been significant research in the field of sensor integrity monitoring, but most of the techniques focus on navigation systems. That traditional approach is not directly applicable to assess the integrity of camera-based monitoring systems, but the underlying concept of fault detection based on statistical testing can be used. However, it is important to note that in the computer vision and image analysis literature, the concept of integrity might have a different meaning.

As discussed in [9], for the image processing community, the concept of image integrity is often related to authentication, meaning that the content of an image has not been altered in a malicious or unintended way. The methods to assess and ensure integrity are typically relying on the detection of digital image artifacts and cryptographic signature protection measures. In the context of a camera-based monitoring system, these methods would be useful for detecting and preventing image manipulation attacks, in case there is a security breach on the access of the monitoring data. However, for detecting physical anomalies, such as the ones derived from camera tampering or changes in environmental conditions, different techniques are required.

There have been several studies on camera tampering detection [10], [11], with methods to identify obstruction or an unintended change in the camera position. These methods, usually based on edge detection [12], are employed for surveillance cameras and require a static or known background plane, being unsuitable for situations where the scenery is dynamic or unknown. To deal with that scenario, methods such as soiling detection and visibility restoration for cameras in autonomous vehicles could be used [13]. The methods developed for visibility enhancement and de-weathering techniques, capable of improving the quality of images under bad weather conditions [14], are another option. Additionally, an investigation of using first and second-order image statistics to detect partial camera obstruction was conducted in [15]. The results from that work suggest that a metric such as the image skewness, a first-order statistic calculated from the pixel value histogram, can be used as an effective detector for partial camera obstruction.

The integrity assessment of a camera system could consider many aspects, including the detection of various types of faults such as:

- Physical defects in the camera, and, or, data communication pipeline;
- Image forgery and data integrity issues, including spurious image manipulation and data corruption;
- Camera tampering (sabotage), and adversarial environmental conditions that could impair the quality of the images and usability of the provided information.

There are also various integrity issues that can be assessed in an object detection process, such as misclassification, detecting objects that do not exist (false positive), low accuracy on the bounding box of an object, and completely failing to detect an existing object (false negative). These scenarios can be seen as faults from the object detection results. Since analyzing

all potential issues would be impractical, this work chooses to focus on the effects of partial camera obstruction on ship detection faults.

C. Object Occlusion

There have been several studies on strategies and algorithms to handle object occlusion [16]. Most deep convolutional neural network (DCNN) architectures for object detection are not robust to handle partially occluded objects [17]. However, there are specific models trained to consider and recognize occlusions, such as the Multi-Level Coding network proposed in [18], which implements the concept of amodal instance segmentation, and the Partial Completion Network mask proposed in [19] that can infer the contour of the occluded part of the object.

Although the effect of a partial camera obstruction can be the same as an object occlusion, these two phenomena are different in nature. Object occlusion is a normal situation that occurs due to the dynamic nature of a scene and the relative viewpoint of the camera in relation to the objects of interest. The partial camera obstruction is a physical anomaly, mostly of static nature, of something obstructing the camera in close proximity to the lens. It might require an intervention to fix the problem, removing the misplaced object that might be causing the obstruction, or cleaning the lens from dirt, smudge, or even paint in case of a targeted attack. The former is a common part of a normal operation, and the latter can be viewed as a type of physical fault that requires correction. The usage of specific strategies that can deal with object occlusion could mitigate the effects of partial obstruction. However, these two phenomena should be treated differently. Since the obstructions are typically static and close to the lens, they are much easier to detect. Also, that detection is important since it might require corrective action. Conversely, the object occlusion does not and it is a normal aspect of the camera operation.

D. Maritime Object Detection

Deep learning methods can be used for rapid and automatic detection of ships from port surveillance and onboard vessel cameras. That process is particularly interesting to improve the safety and security of maritime applications and infrastructures. Since deep-learning-based object recognition algorithms are a type of supervised machine learning problem, datasets for training are needed. The images available in general-purpose benchmarks datasets such as COCO [20] or PASCAL VOC [21], do not suit the task of ship detection as it is intended in the proposed work. Real-world situations require robust reference images with precise bounding box annotations, that should also include the ship class. Among the works in the literature that deal with ship detection on video monitoring cameras, available datasets are the Singapore Maritime Dataset [22], Seaships7000 [23] and the dataset presented by Chen et al. [24]. However, these datasets lack a variety of ship classes in their annotations. Other works evaluate existing object detection methods on their private ship detection datasets [25], [26], making the experimental validation of their methods not possible. The ShipSG dataset [2] is a public dataset for ship detection that contains seven ship classes and two different views of a port location, allowing the statistical analysis of simulated partial camera obstruction of the different ship classes applied in a real-world maritime situation.

III. Methodology

A. Dataset for Ship Detection

The ship detection dataset selected for this work has been the ShipSG dataset [2], which is publicly available [1]. ShipSG consists of 3505 images and ship annotations of two different views of the Fischereihafen-Doppelschleuse, part of the port of Bremerhaven, Germany. The images were acquired during Autumn 2020 in daylight hours with sunny, cloudy, windy and rainy weather conditions. The annotations of ShipSG contain 11625 annotated ship instances, which proportion is shown in Table I, divided in seven ship classes:

- Cargo: All types of cargo ships.
- Law Enforcement: Police and coast guard ships.
- Passenger/Pleasure: Ferries, pleasure and sail crafts.
- Special 1: Crane vessels, dredgers and fishing boats.
- Special 2: Research and survey ships, search and rescue ships and pilot vessels.
- Tanker: All types of tankers.
- Tug: All types of tugboats.

Class	Annotated	Proportion
Cargo	1300	11.18%
Law Enforcement	3748	32.24%
Passenger/Pleasure	626	5.38%
Special 1	1002	8.62%
Special 2	2630	22.62%
Tanker	1412	12.15%
Tug	907	7.80%
All classes	11625	100%

TABLE I: Number of ships annotated per class in ShipSG.

The ground truth annotations of ShipSG contain ship masks for instance segmentation tasks. However, we only considered the surrounding bounding box of the ships for this work. Fig. 1 shows examples of each of the seven classes.

B. Ship Detection Algorithm

In order to extract the position of the ship within the images, an object detector that provides the surrounding bounding box of the ship is needed. For this work, the well-established and robust Faster-RCNN algorithm for object detection has been selected [27]. Faster-RCNN is a two-stage algorithm. In the first stage, with the region proposal network [27], multiple object candidates are proposed. In the second stage, the region of interest pooling extracts features from each candidate and

[1] https://www.dlr.de/mi/shipsg, accessed on 05 of July, 2023.

Fig. 1. Examples extracted from ShipSG that show the seven ship classes. (a) Cargo, (b) Tug, (c) Special 1, (d) Tanker, (e) Law Enforcement, (f) Passenger/Pleasure, (g) Special 2

performs the classification of the object and the regression of the bounding box.

To train Faster-RCNN for ship detection on ShipSG, we divided the dataset into two sets of images—training (80%) and validation (20%), with similar class distribution. The input image size in both cases was 1333 × 800 pixels, and the backbone selected was ResNeXt-101 [28]. The training was performed with a batch size of 2 during 11 epochs. The MMdetection framework from OpenMMLab [29] was used for this implementation. The mean Average Precision (mAP), of all classes, of Faster-RCNN for bounding box detection on ShipSG after validation has been 0.826.

C. Metrics and Data Acquisition Process

A selection of 200 images, containing 672 ship annotations, from the validation set of the ShipSG dataset was processed by inferring with our custom Faster-RCNN detector, with and without synthetically generated partial obstructions. The statistics for the true positives, false positives, misclassifications and false negatives were calculated from the results.

Several possible events could cause a partial camera obstruction. Weather conditions leading to particulates of rain or frost sticking to the lens, that could be mixed with oil, dust and smoke debris; or winds carrying objects that would obstruct the camera view are among the possibilities. Due to the variety of these events, it is difficult to predict a preferred shape and color for these obstructions. For that reason, the synthetic obstructions were configured with a range of positions (relative to each ship bounding box in the image), colors, and obstruction rates (a percentage of the ship bounding box being covered). Examples of these variations are given in Fig. 2. A total of 176 obstruction configurations were analyzed for each ship, with the following parametric variations:

- Position: starting at the left, right, center, top and bottom of the ship bounding box.
- Obstruction percentage: the area of the added obstruction relative to the ship bounding box, ranging from 20% to 90%, in 10% increments.
- Color: with two shades of black (centered at mean brightness levels of 30 and 60), one gray (at 120), and two bright shades (at 180 and 210). To simulate a more natural-looking texture, instead of using homogeneous colors, the obstruction area was filled with Gaussian noise centered in one of the five mean brightness levels. Next, a median filter was applied to that mask to blur the color.

Fig. 2. Examples of a ship with different synthetic partial obstruction profiles. (a) No obstruction; (b) 30% bright obstruction at the bottom; (c) 30% dark obstruction at the center; (d) 60% gray obstruction at the right.

The object detection algorithm yields, for each image, a list of coordinates that defines bounding boxes and a detection score, with the indexes in that list representing the class for the detected object. To obtain the metrics of interest for the integrity assessment the following algorithmic processing was conducted:

- For each image, the ground truth for the bounding boxes and classes of the ships contained in the image were stored in a list.
- The image is processed with each one of the 176 obstruction configurations (including the no-obstruction case), yielding a new image for each configuration with a partial obstruction in every ship contained in the image.
- Each new image is processed with the ship detection algorithm and the results are compared with the ground truth in the following manner:
- For each ship, the closest bounding box with the highest score from the detection results is selected.
- If the centroid from that selected result is inside the ground truth ship boundaries and the class of the ship matches, a true positive is computed. If the centroid matches but the class does not, a misclassification is computed, and if the centroid is outside, that is a case of false negative, as there are no instances in the results that match that ship.
- After this computation, if the selected detection result centroid matches the ground truth bounding box, that result is removed. That avoids duplicate counting and enables the computation of false positives.

This process is repeated for each ship in each image, and after all ships in an image are processed, if there are remaining instances in the detection results, those are counted as false positives.

Note that a typical method to evaluate a detection result is with the Intersection over Union (IoU) of the detected

bounds with the ground truth. An accurate detection would have an IoU close to one. However, when the object is partially obstructed that metric is degraded. Even in the case of correct detection, as the obstructed area is often not considered in the detected bounds, the IoU can be low. For that reason, instead of using that parameter to classify true positives, false negatives, and misclassifications, the centroid match method was used.

IV. Results and Analysis

The graphs in Fig. 3 and Fig. 4 show the relationship between the obstruction intensity and the considered stats (namely, true positives, false positives, false negatives, and misclassifications).

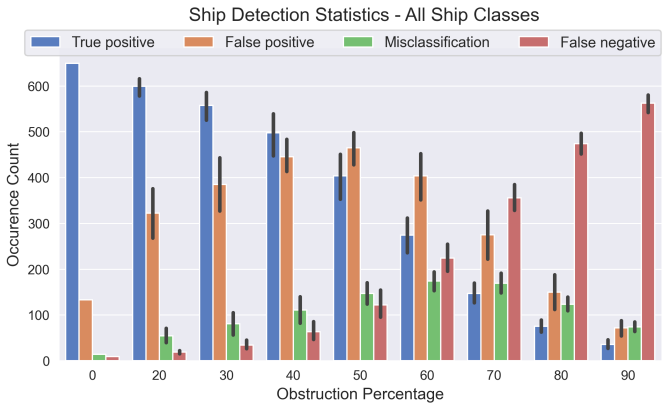

Fig. 3. Relationship between the occurrence of the evaluated stats of interest and the obstruction rate, given as a percentage of the ship bounding box being occluded.

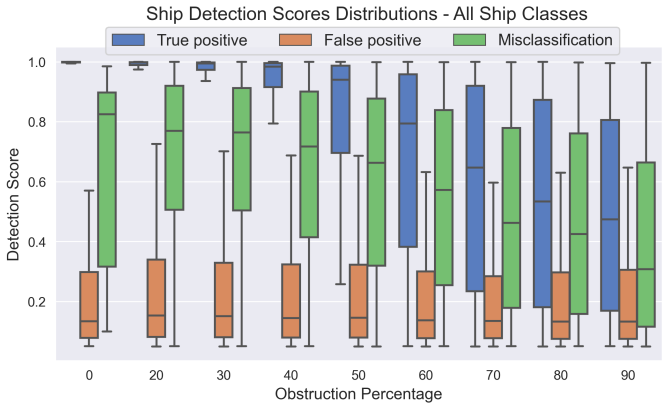

Fig. 4. Distribution of the detection scores for the evaluated metrics of interest as a function of the obstruction rate.

When the obstruction intensity raises, the occurrence of correct detections falls, almost in a linear relationship, while, simultaneously, the rate of false negatives rises. The occurrence of false positives and misclassification peaks around the 50% and 60% obstruction rate. As the occlusion covers most of the ship, the detector fails to recognize any object and those rates fall, giving rise to the false negative count.

The analysis of the distribution of the detection score in Fig. 4 fosters some interesting conclusions. Under normal conditions, the correct detection scores are very close to 1, and most of the occurrences of lower scores would be identified as misclassification and false positive cases. But, as the obstruction rate rises, that correct score distribution mean lowers and its variance raises. At around the 60% obstruction rate, the correct scores have a distribution that resembles the misclassification in the normal case. If there is an undetected partial obstruction of that magnitude, a correct detection could easily be mistaken as a misclassification. Therefore, it is advantageous to include a partial camera obstruction detector in combination with the object detection, as that kind of inference mistake could be prevented.

A. Impact of Partial Camera Obstruction Faults in Ship Detection Metrics

The F1 score, given by equation 1, is an important performance metric used to evaluate classification results. It considers the precision and recall, affected by the true positive (TP), false positive (FP) and false negative (FN) rates.

$$F1 = \frac{2 * Precision * Recall}{Precision + Recall} = \frac{2 * TP}{2 * TP + FP + FN} \quad (1)$$

The graph in Fig. 5 shows the influence of the obstruction intensity in that score for the most representative ship classes. The score was calculated considering that misclassifications are also false positives, summing up both metrics, and averaging out the variations due to the obstruction position and color to display the mean F1 score for each obstruction percentage. The highest scores are for the Law Enforcement ships, which is the class that has the highest representation in the dataset, peaking at about 0.96 for the no-fault case. The average score, considering all ship classes, falls almost linearly from around 0.9 to 0.1, as the obstruction intensity raises. The only exception to that score degradation trend is for the cargo ship class, which remains close to 0.5 for the obstruction range from 40 % to 90 %. The explanation for that anomaly would require further investigation.

The F1 score results suggests that the reliability of the ship detection information is sharply degraded in the presence of obstructions. Considering an application that relies on that information for decision support, such as in a harbor monitoring scenario, it is important to know the confidence level of the ship detection. Although that confidence is correlated with the detection score, as have been shown, that score distribution is affected by these obstructions. Assuming that it is possible to detect partial camera obstructions, and that its intensity (the obstruction percentage) can be estimated from an object instance segmentation step, the F1 score curves could be used to assess the integrity of the information provided by the object detection. Since a harbor could have dozens of surveillance cameras, sometimes positioned in difficult to reach locations,

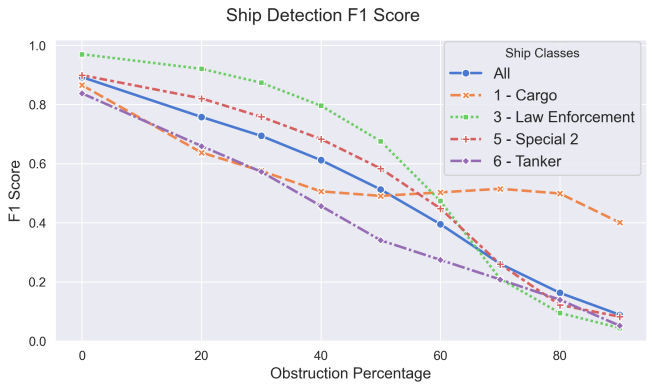

Fig. 5. Mean object detection F1 Score for the main classes of ships as a function of the obstruction ratio.

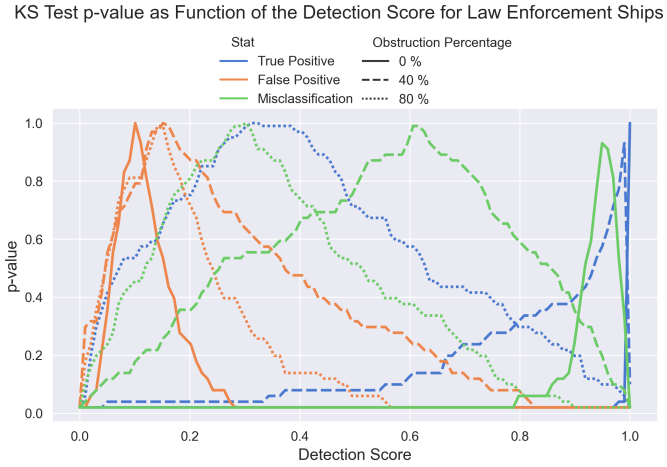

Fig. 6. Two-sample Kolmogorov-Smirnov test for a range of detection score values in comparison with the estimated Beta distributions for the various cases under analysis, considering the Law Enforcement ship class detection results.

it might be unpractical to deploy resources to clean the lenses or to remove objects whenever an obstruction is detected. Therefore, having a metric to indicate the degradation of the intelligence provided by a camera under these scenarios is useful. Additionally, in a multi-sensor monitoring system, this integrity metric would be helpful to resolve conflicts, giving priority to the information of sensors with a higher score.

The analysis of the detection score distribution can be used to predict if a instance of object detection represents a true positive, false positive or misclassification. As seen in the box-plots from Fig. 4, there is a correlation between the object detection score and these three cases. The distribution of the score for each scenario is different. Since the score values are always between 0 and 1, the Beta distribution is appropriate to model that data. The two parameters of the distribution, α and β, can be estimated from the expectation $\mathbb{E}[X]$ and variance $\mathbb{V}[X]$ of the data according to equation 2.

$$\begin{cases} \alpha = \dfrac{\mathbb{E}[X](1-\mathbb{E}[X])}{\mathbb{V}[X]} - 1 & (2) \\ \beta = \alpha \dfrac{1-\mathbb{E}[X]}{\mathbb{E}[X]} & (3) \end{cases}$$

The Beta distribution of the detection score was estimated for each scenario, considering each ship class, obstruction intensity and stat of interest (true positive, false positive and misclassification). These distributions were validated using the Kolmogorov-Smirnov (KS) test for goodness of fit [30], which confirmed a good match with the the actual detection score data for most scenarios. The only exception was for the scores of true positives in the normal case (without any obstruction), which are highly concentrated and close to 1. Next, having a reference distribution for each scenario, a two-sample KS-test could be used to detect misclassification and false positive faults from a new detection score. However, since those distributions are affected by the obstruction intensity that information has to be considered. To illustrate that dependency, the KS-test results were calculated.

The graph in Fig. 6 shows, for the Law Enforcement ship class, the p-value from the KS-tests conducted from a range of scores, from 0 to 1, in comparison to the estimated Beta distributions for the true positive, misclassification and false positive cases. Note that the threshold values for the score that enables the distinction between the three cases change drastically according to the obstruction intensity. Since the variance from the score associated with each case rises with that intensity, the detectability through a threshold value for that test lowers. The exception is for the false positive case, as previously discussed and seen in the bar plot in Fig. 3, which has the highest incidence in intermediate obstructions. It is possible to confirm that by analyzing the score range in which the KS-test p-values are large, indicating that any of those scores could be statistically associated with the estimated Beta distribution for that case. One possible outcome of this analysis is for the detection of misclassifications using this method. The integration of the obstruction information is advantageous in that case, as it would mitigate mistakes in the misclassification detection when the obstruction intensity is around 40%. There is an overlap between the misclassification p-value curve for the normal case and the true positive curve in the 40% rate. Therefore, if the detection threshold from the normal case is used when there is in fact a 40% obstruction, a true positive could be mistaken as a misclassification.

V. Discussion and Conclusion

The ShipSG dataset was used with the Faster-RCNN object detection algorithm, trained with the seven classes of ships contained in the dataset, to analyze the effects of simulated partial camera obstruction on the ship detection results. This work discusses an integrity assessment based on fault detection, considering the partial obstructions as a type of physical fault, and its impact on fault rates from the ship detection process. The presented results highlight how the incorporation

of the detection of that physical fault can mitigate errors in the ship detection analysis. The considered camera with object detection results can be viewed in the context of a multi-sensor system, in which the reliability of the information provided by each element must be accounted for. That can be accomplished through an integrity assessment process, that indicates the presence of faulty states with a certain confidence level. In this context, the detection of a partial camera obstruction, a false positive and a misclassification from the object detector can be seen as integrity tests. As suggested and discussed by the obtained results, fusing the integrity information from tests that are able to detect these faults can potentially reduce the incorrect detection of misclassification faults under partial camera obstruction. Additionally, the degradation of the ship detection performance with the intensity of the partial obstruction was characterized. In a multi-sensor system where conflicting information between redundant sources is a possibility, the detection of faults and the characterization of their effects could be used to calculate a dynamically allocated weight for each information source, solving conflicts by selecting the sources with the highest weights. The presented work is a step towards a comprehensive integrity monitoring solution for multi-sensor systems used for safety-critical maritime applications.

ACKNOWLEDGMENT

The authors would like to thank the Deutscher Akademischer Austauschdienst (DAAD), for providing the doctorate scholarship that funded this work.

REFERENCES

[1] L. E. van Dyck, R. Kwitt, S. J. Denzler, and W. R. Gruber, "Comparing object recognition in humans and deep convolutional neural networks—an eye tracking study," *Frontiers in Neuroscience*, vol. 15, 2021.

[2] B. Carrillo-Perez, S. Barnes, and M. Stephan, "Ship segmentation and georeferencing from static oblique view images," *Sensors*, vol. 22, no. 7, p. 2713, 2022.

[3] F. A. Costa de Oliveira, F. S. Torres, and A. Garcia-Ortiz, "Recent advances in sensor integrity monitoring methods-a review," *IEEE Sensors Journal*, vol. 22, no. 11, pp. 10256–10279, 2022.

[4] Y. Yang and J. Xu, "GNSS receiver autonomous integrity monitoring (raim) algorithm based on robust estimation," *Geodesy and Geodynamics*, vol. 7, 04 2016.

[5] M. Joerger and B. Pervan, "Kalman filter-based integrity monitoring against sensor faults," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 2, pp. 349–361, 2013.

[6] R. V. Hogg and J. W. McKean, *Introduction to Mathematical Statistics*. Pearson, 8th edition ed., 2019.

[7] J. Palmqvist, "Integrity monitoring of integrated satellite/inertial navigation systems using the likelihood ratio," in *Proceedings of the 9th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 1996), Kansas City, MO*, pp. 1687–1696, September 1996.

[8] P. Zabalegui, G. De Miguel, A. Pérez, J. Mendizabal, J. Goya, and I. Adin, "A review of the evolution of the integrity methods applied in GNSS," *IEEE Access*, vol. 8, pp. 45813–45824, 2020.

[9] P. Korus, "Digital image integrity – a survey of protection and verification techniques," *Digital Signal Processing*, vol. 71, pp. 1–26, 2017.

[10] S. M. Hosseini and A. H. Taherinia, "Anomaly and tampering detection of cameras by providing details," in *2016 6th International Conference on Computer and Knowledge Engineering (ICCKE)*, pp. 165–170, 2016.

[11] A. Saglam and A. Temizel, "Real-time adaptive camera tamper detection for video surveillance," in *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance*, pp. 430–435, 2009.

[12] G.-b. Lee, Y.-c. Shin, J.-h. Park, and M.-j. Lee, "Low-complexity camera tamper detection based on edge information," in *2014 IEEE International Conference on Consumer Electronics - Taiwan*, pp. 155–156, 2014.

[13] M. Uřičář, H. Rashed, A. Ranga, A. Dahal, and S. Yogamani, "Visibilitynet: Camera visibility detection and image restoration for autonomous driving," in *Electronic Imaging, Autonomous Vehicles and Machines*, 07 2020.

[14] A. C. Aponso and N. Krishnarajah, "Review on state of art image enhancement and restoration methods for a vision based driver assistance system with de-weathering," in *2011 International Conference of Soft Computing and Pattern Recognition (SoCPaR)*, pp. 135–140, 2011.

[15] F. A. Costa de Oliveira, A. Niemi, A. Garcia-Ortiz, and F. S. Torres, "Partial camera obstruction detection using single value image metrics and data augmentation," in *2022 6th International Conference on System Reliability and Security - (ICSRS 2022)*, pp. 292–299, 2022.

[16] K. Saleh, S. Szénási, and Z. Vámossy, "Occlusion handling in generic object detection: A review," in *2021 IEEE 19th World Symposium on Applied Machine Intelligence and Informatics (SAMI)*, pp. 000477–000484, 2021.

[17] A. Fawzi and P. Frossard, "Measuring the effect of nuisance variables on classifiers," in *Proceedings of the British Machine Vision Conference (BMVC)* (E. R. H. Richard C. Wilson and W. A. P. Smith, eds.), pp. 137.1–137.12, BMVA Press, September 2016.

[18] L. Qi, L. Jiang, S. Liu, X. Shen, and J. Jia, "Amodal instance segmentation with kins dataset," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3009–3018, 2019.

[19] X. Zhan, X. Pan, B. Dai, Z. Liu, D. Lin, and C. C. Loy, "Self-supervised scene de-occlusion," *CoRR*, vol. abs/2004.02788, 2020.

[20] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*, pp. 740–755, Springer, 2014.

[21] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes challenge: A retrospective," *International journal of computer vision*, vol. 111, no. 1, pp. 98–136, 2015.

[22] D. K. Prasad, D. Rajan, L. Rachmawati, E. Rajabally, and C. Quek, "Video processing from electro-optical sensors for object detection and tracking in a maritime environment: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 8, pp. 1993–2016, 2017.

[23] Z. Shao, W. Wu, Z. Wang, W. Du, and C. Li, "Seaships: A large-scale precisely annotated dataset for ship detection," *IEEE transactions on multimedia*, vol. 20, no. 10, pp. 2593–2604, 2018.

[24] X. Chen, L. Qi, Y. Yang, Q. Luo, O. Postolache, J. Tang, and H. Wu, "Video-based detection infrastructure enhancement for automated ship recognition and behavior analysis," *Journal of Advanced Transportation*, vol. 2020, 2020.

[25] A. Ghahremani, Y. Kong, E. Bondarev, *et al.*, "Multi-class detection and orientation recognition of vessels in maritime surveillance," *Electronic Imaging*, vol. 2019, no. 11, pp. 266–1, 2019.

[26] C. Nita and M. Vandewal, "Cnn-based object detection and segmentation for maritime domain awareness," in *Artificial Intelligence and Machine Learning in Defense Applications II*, vol. 11543, p. 1154306, International Society for Optics and Photonics, 2020.

[27] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *Advances in neural information processing systems*, vol. 28, pp. 91–99, 2015.

[28] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.

[29] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin, "MMDetection: Open mmlab detection toolbox and benchmark," *arXiv preprint arXiv:1906.07155*, 2019.

[30] M. A. Stephens, "EDF statistics for goodness of fit and some comparisons," *Journal of the American Statistical Association*, vol. 69, no. 347, pp. 730–737, 1974.