Image Saliency Applied to Infrared Images for Unmanned Maritime Monitoring

Gonçalo Cruz[1] and Alexandre Bernardino[2]

[1] Research Center, Portuguese Air Force Academy, Sintra, Portugal
gccruz@academiafa.edu.pt
[2] Computer and Robot Vision Laboratory, Instituto de Sistemas e Robótica, Instituto Superior Técnico, Lisboa, Portugal
alexandre.bernardino@tecnico.ulisboa.pt

Abstract. This paper presents a method to detect boats and life rafts on long wave infrared (LWIR) images, captured by an aerial platform. The method applies the concept of image saliency to highlight distinct areas on the images. However saliency algorithms always highlight salient points in the image, even in the absence of targets. We propose a statistical method based on the saliency algorithm output to distinguish frames with or without targets. To evaluate the detection algorithm, we have equipped a fixed wing unmanned aerial vehicle with a LWIR camera and gathered a dataset with more than 44000 frames, containing several boats and a life raft. The proposed detection strategy demonstrates a good performance, specially, a low rate of false positives and low computational complexity.

Keywords: Image Saliency, detection, Infrared images, LWIR, UAV, Maritime surveillance

1 Introduction

In recent years, many maritime safety issues have challenged the international community. Since year 2000, more than 22400 emigrants have lost their lives on the Mediterranean Sea [1]. Countries in South-East Asia carried out large scale search and rescue operations over sea [2] and waters of the Somali coast and Malacca Straight have been populated with pirates that endanger vital navigation routes [9]. On one hand, assuring maritime safety is paramount to protect human lives and to the world economy, but on the other hand, it is very difficult to guarantee a proper surveillance or monitoring over a wide area.

Even though there has been a huge investment on technologies to perform surveillance and monitoring [11] namely coastal radars, manned aircraft and satellites, there is still a lot to be done on improving the coverage and persistence of these tasks. One of the rising trends, due to recent significant technological advances, is the use of unmanned aerial vehicle (UAV). Fixed wing UAVs are specially attractive because they have large endurance, when compared to its rotary wing counterparts or similar manned aircraft. They also offer a more

flexible deployment than coastal resources and satellites, as well as a smaller cost than other alternatives. Despite its advantages, the vast majority of UAVs have serious restrictions on the payload which limits the type of sensor they can carry, being the most commons electro-optical sensors - visible image and infrared (IR).

This paper focuses on the detection of boats and life rafts in the ocean using a UAV with cameras observing the ocean surface. We aim at a persistent surveillance of a large area, which requires the analysis of a huge amount of data gathered by the sensors. The data in most of the cases does not convey any useful information and even hampers the ability of the human operators to analyse the imagery [5]. Additionally, visible images have drawbacks as they highly depend on the illumination conditions and often contain distractors like waves and sun glare.

To circumvent the issues presented for images on the visible spectrum, we have instrumented a fixed wing UAV with an infrared camera able to detect radiation with wavelength between 8 - 14 μm. Radiation in this range, is often called Long Wavelength IR (LWIR) or thermal infrared. This designation is based on the fact that, in this range the intensity of radiation emitted by a given body depends largely on its temperature.

On an ocean monitoring flight most of the observed background is water with approximately the same temperature, at each time of the day. Objects like vessels and life-rafts are expected to present areas with different temperature and therefore emit LWIR radiation with higher intensity. Sun shades and waves have a much lower influence at this radiation spectrum that in the visible range. Nonetheless, the analysis of infrared imagery still presents some challenges. Temperature difference may be very small in some objects. Moreover, an object may be visible only for a short period of time due to low flying altitude of a small UAV and consequent small observed area at each given instant.

This work addresses these issues by proposing the application of saliency algorithms to IR images to detect vessels. Saliency algorithms are inspired in the ability of humans to easily segment image regions that are distinct from their neighbourhood, independently of their absolute intensity levels and contrast [8]. Thus they have the advantage of detecting regions on the sea that have only slight temperature differences to the background. Furthermore, they operate in single frames, which allow detections of boats that only persist for a few frames. However, classical saliency algorithms are designed to always detect the most salient points in an image. This is a problem for our application because in most of the time the observed images have only background. The most salient objects detected in these images have no relevance to the application and will produce false alarms. To tackle this problem, we propose a statistical analysis of the maps produced by the saliency algorithms. This allows to determine whether we have a relevant signal or simply background noise, thus reducing the rate of false alarms.

The paper is organized as follows. In Section 2, we will discuss the related work. In Section 3 our problem is detailed and our approach presented. In Section

4, we evaluate the performance of the proposed algorithm with real data. Finally, in Section 5, we present our conclusions and some directions for future work.

2 Related Work

Aircraft have been used as a privileged observation platform since its appearance. As new sensing technologies emerge, they are readily integrated in aerial platforms for possible applications [21]. One of the most successfully applied technologies has been infrared imaging. These images can provide information that is not present on the visible spectrum (like temperature of a given object), and therefore may overcome limitations of human vision or complement it[23].

A lot of work has been developed on automatic detection using infrared sensors for military aircraft. Most works focus on the detection of airborne targets like missiles or other combat aircraft, which typically have high temperature engines and exhaust [20]. The detection of high temperature objects in IR images was also exploited to detect fires [10] [12].

Other scenarios have also been considered for the detection based on infrared images gathered by aircraft. For instance, in [15], a UAV detects and follows a river based on Near-IR images; in [18] and [24] is also reported the detection and tracking of ground targets. Detection of people also had the attention of several works, on both visible [22] and IR images [16].

The detection of objects on the ocean has not yet been thoroughly explored. There are some works which deal with images gathered by surface platforms [19] [17] which impose constraints on the appearance of the objects to be detected. There are also some works that use airborne visible images [3] that extract features and successively select objects of interest. Particularly, in [3] the authors use five features that depend directly on the three RGB channels. That approach, while successful on visible spectrum images, is not suitable for LWIR images, once no color is available.

Another important aspect is that, when observed from an airborne platform, vessels may have a multitude of appearances, depending on several factors like the type of vessel, its size, the height of observing aircraft, etc. Facing this restriction, we would like to detect a given object, without knowing its appearance beforehand. Itti *et al* [8] have introduced a saliency model that mimics the visual search process of humans. That saliency model indicates areas of a given image that would be examined on a second stage with more detail. More recently, works like [7] and [6] have explored the spectral characteristics of an image to create similar saliency maps. Compared to Itti's approach, those techniques show a better detection performance and much faster computation time. With such advantages, these approaches became very attractive for the application at hand. There is still one major limitation for this technique in our scenario: the algorithm is designed to highlight salient areas of the image and therefore will produce detections, regardless if there is a foreground or not.

3 Detection on LWIR images

The rationale behind the usage of a camera sensitive to infrared light with wavelengths comprised between 8 and 14 μm is that the water in the ocean has an approximately uniform temperature and objects of interest have different temperatures. With this assumption in mind, it is expectable that, while observing the ocean, the thermal image has approximately uniform intensity for the water background and brighter or darker areas corresponding to objects of interest.

Typically, LWIR cameras have smaller resolution than common RGB cameras, so we expect to observe less details of the vessels. Additionally, in the case of small boats, we expect their temperature to be approximately the same as the ocean. In the case of bigger vessels some parts (like the engines or the exhausts) have higher temperatures and therefore are more noticeable.

As visible in Figure 1(a), when acquired, the image exhibits an almost uniform background and a foreground with very small contrast. The image presented in 1(b) shows a more perceivable foreground and was enhanced by histogram equalization. Histogram equalisation highlights high contrast areas that, despite improving the vessel appearance, also increase the image noise level. Therefore, a simple thresholding technique is prone to numerous false positive detections. Saliency methods can reduce this effect because they analyse image neighborhoods instead of single points. However, as we will describe next, they still present problems when no targets are present in the image.

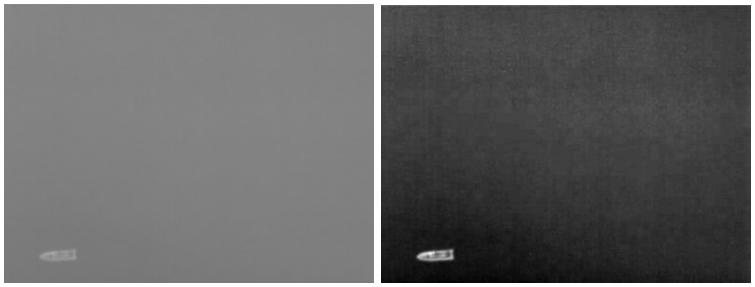

(a) (b)

Fig. 1: Sample LWIR image: (a) acquired image; (b) image with contrast improved by off-line processing.

3.1 Image signature

An image with characteristics like exhibited in Figure 1 is prone to the application of saliency detection algorithms. They do not require a full characterization of the object of interest but rather highlight areas that are different from the

background. As stated before, algorithms that explore the spectrum of the image, like [7] and [6], have shown good performance and low computational cost. Therefore, these algorithms are specially interesting for our application. We have tested and compared both approaches, hereafter designated as Spectral Residual [7] and Image Signature [6]. To perform this evaluation, we have implemented both algorithms. The mean computation time to create a saliency map of 64 by 48 pixels was 6.5 and 5 miliseconds for Spectral Residual and for Image Signature, respectively. Even though the detection of saliency in LWIR images was similar, the improvement in computation time lead us to the selection of Image Signature algorithm.

3.2 Description of the method

In [6] the authors show that, for an image x expressed as

$$x = f + b, \quad x, f, b \in \mathbb{R}^N \tag{1}$$

where N is the number of pixels, f represents the foreground and b represents the background, the separation of f and b can be performed given only x. The assumption followed is that the foreground is sparsely supported in image space and the background is sparsely supported in the Discrete Cosine Transform (DCT) basis. The method consists of isolating the spatial support of f, which is accomplished by computing the sign of x in the DCT domain and then transform back to the spatial domain. This operation is written as

$$\bar{x} = \mathrm{IDCT}\left[\, sign(\,\mathrm{DCT}(x)\,)\,\right]\,, \tag{2}$$

where $\bar{x}$ is the saliency map and IDCT is the Inverse Discrete Cosine Transform. The intuition behind Eq. 2 is that the sign of $\mathrm{DCT}(x)$ is more strongly correlated with the sign of $\mathrm{DCT}(f)$ than with the sign of $\mathrm{DCT}(b)$. In one hand, because $\mathrm{DCT}(b)$ is sparse, most of its coefficients are zero or very close to zero outside its support Ω_b. Thus:

$$sign(\mathrm{DCT}(x)) = sign(\mathrm{DCT}(f) + \mathrm{DCT}(b)) \approx sign(\mathrm{DCT}(f)). \tag{3}$$

On the other hand, inside Ω_b, because f and b are independent, there is still a positive correlation between the signs of $\mathrm{DCT}(f)$ and $\mathrm{DCT}(x)$. As shown in [6]:

$$\mathrm{Prob}(sign(\mathrm{DCT}(x) = sign(\mathrm{DCT}(f))) \geq 0.5\,. \tag{4}$$

In practice, we resize x to an image of 48 by 64 pixels, which results in four-fold decrease in computation time and still highlights the areas of interest. The final saliency map (examples in Figure 2) is obtained by convolving the result with a Gaussian kernel, to reduce noise.

With a saliency map already computed, we would like to isolate the objects of interest in the image. However, the separation between different areas must be done carefully. According to the authors' premise that a foreground is present

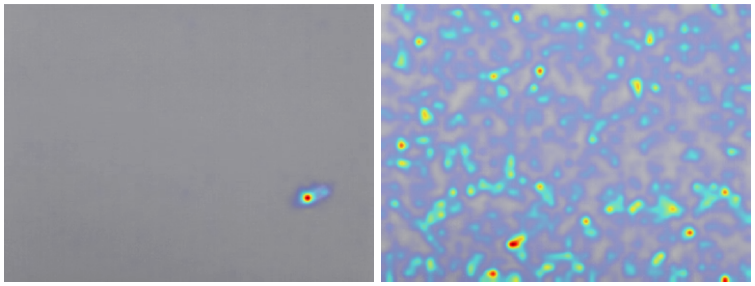

(a) (b)

Fig. 2: Saliency maps overlayed on the LWIR images: (a) image with boat present; (b) image without boat present. Areas in red are more salient and areas where grey is visible are considered less salient.

and has a sparse spacial representation, the saliency map will always highlight the most conspicuous areas of the image. In cases where no foreground is present, slight variations on the intensity of radiation emitted by the ocean's surface as shown in Figure 2(b), will be amplified. This fact precludes the direct application of a thresholding technique, as it will produce numerous false positives.

In [7], the authors suggest an empirical method to isolate foreground from background. Their method consists on setting a threshold to three times the average of the saliency map and then segment this map. We think that this approach is highly dependent on the adjustment of the threshold by an user, otherwise the results may be poor. We extend this idea by multiplying different scalars (α) and we will designate it as Saliency Threshold (ST) algorithm.

However, a Saliency Threshold method lacks the ability to associate a confidence level to the obtained detections. Confidence levels are important for a number of applications, such as giving feedback to a human operator or feeding probabilistic time filtering methods like HMMs [14]. To address this issue, we analyse the statistical distribution of the saliency values in the saliency maps with two different options, that we will designate as histogram ratio saliency confidence (HRSC) and entropy saliency confidence (ESC). By assigning a score to each detection based on the statistical properties, instead of setting a fixed threshold, we are able to easily set the operating point of the algorithm (trade-off between precision and recall rates) according to the task at hand. The first strategy considers the properties of the histogram of saliency values. In cases where a boat is visible, most of the pixels are contained in the lower saliency range, as shown by Figure 3(a). In the cases where only ocean is visible, then, there are more pixels with other intensities (Figure 3(b)). Based on this fact, we compute a confidence score (HRSC) for each image as

$$\rho = 1 - \frac{\text{bin}_{\#2}\ \text{count}}{\text{bin}_{\#1}\ \text{count}} \tag{5}$$

where $\text{bin}_{\#1}$ count and $\text{bin}_{\#2}$ count are the number of pixels in the lower intensity bin and in the second lowest bin. If the ratio between the histogram bins is small, then it is more likely to have a boat in sight and therefore a higher score is assigned. In case the ratio is bigger, then the confidence will be closer to zero and is less likely to have a boat in the image.

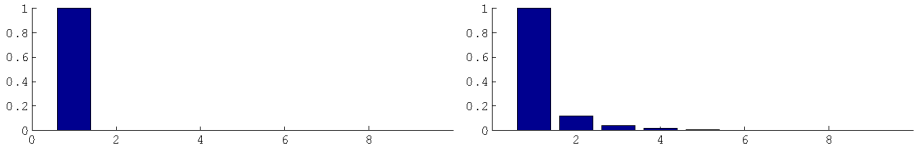

Fig. 3: Normalized intensity histogram of the saliency map (a) for a image with a boat present and (b) for a image without boat present.

The second option to perform the statistical analysis, is the use of saliency map entropy. As we have shown in Figure 2, images with no boats produce apparently random saliency maps, while images with boats produce more structured maps. Image entropy is a measure of the randomness of the information contained in a given image and is computed as

$$\eta = -\sum_{i}^{n}(p_i \log_2(p_i)) \tag{6}$$

where p_i is the normalized value of bin i, of the saliency map histogram. As before, we define a confidence score varying from zero to one, which results from normalizing the image entropy (that in this case, we will designate as Entropy Saliency Score - ESC). If η has a high value, then is less likely to have a boat in the image and therefore, the score is low. If the entropy of the map is low, then the confidence score will be higher.

In the final step of the proposed method, we isolate the areas of the saliency map with higher values. We use Otsu's method [13] to compute the threshold value. This method does not need an adjustment, as it minimizes the variance in each of the classes (in the resulting binary image). After thresholding the map, a connected components algorithm clusters the pixels into detections. At last, a bounding box for each detection is created and scaled back to the original image size, as shown in Figure 5(b). The algorithms are summarized in Table 1.

4 Experimental results

The dataset used to test the algorithms consisted of a sequence with 44677 frames, with a resolution of 384 by 288 pixels. This sequence corresponds to 24 minutes and 49 seconds of images captured by a LWIR camera on-board a

Table 1: Saliency Confidence (HRSC and ESC) and Saliency Threshold algorithm.

Algorithm 1: Saliency Confidence	**Algorithm 2:** Saliency Threshold
Input: LWIR image	Input: LWIR image
Parameters: detection threshold	Parameters: multiplying factor α
Resize LWIR image	Resize LWIR image
Compute **saliency** with Eq. 2	Compute **saliency** with Eq. 2
Compute **confidence** with Eq. 5 or Eq. 6	Compute **average** of saliency map
If confidence <detection threshold	Segment pixels with intensity $> \alpha \times$ **average**
Apply Otsu segmentation	Apply connected components
Apply connected components	Compute bounding box
Compute bounding boxes	Output: Bounding boxes
Else	
Discard image	
End	
Output: Bounding boxes	

UAV (represented in Figure 4(a)). The camera is a GOBI-384 (Figure 4(b)), equipped with a 18 mm lens and a relative aperture of $f/1$. The camera has an Ethernet connection to the UAV's computational board, which allows us to store the images directly on a digital format. The UAV flew over two small boats (with a length of 16 and 12 meters) and a life-raft.

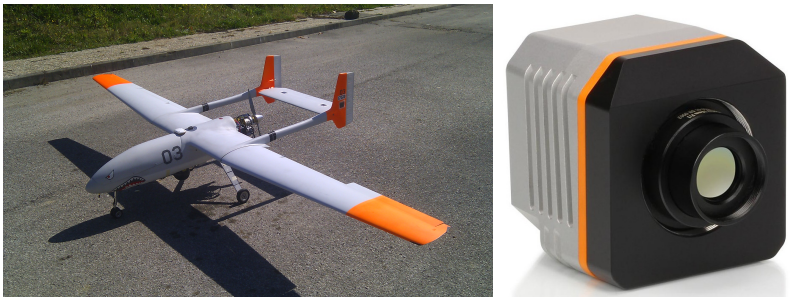

(a) (b)

Fig. 4: (a) UAV and (b) LWIR camera used to obtain the set.

The sequence was annotated manually with the position of the boats and life raft in each frame. The annotation consisted of a bounding box that encompassed each vessel and that we use as Ground Truth (GT). The GT is represented in Figure 5(b) as a red rectangle.

To evaluate the performance of the algorithm and compare it against other approaches, we have followed a method similar to what is proposed in [4]. The evaluation pipeline that we have followed has a slight difference to the original

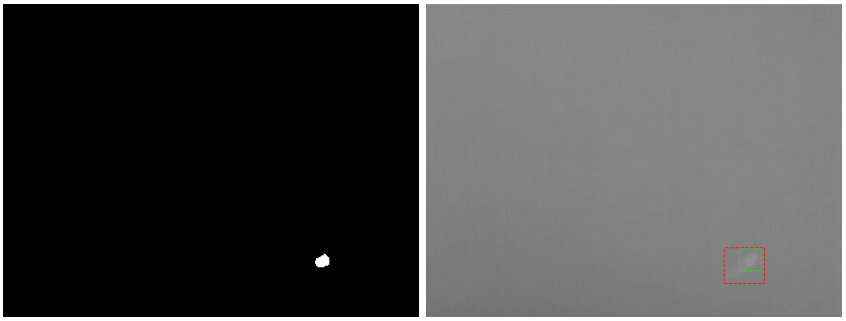

(a) (b)

Fig. 5: (a) Binary image with a boat present; (b) Original LWIR image with bounding boxes visible. Ground Truth annotation represented in red and the output of the algorithm represented in green.

method. The difference consists in considering a detection as a true positive as long as its center is included on the GT Bounding Box. The results are presented as a ratio between Precision and Recall.

On our trials, we have tested histogram ratio saliency confidence, entropy saliency confidence, saliency threshold and some other simple methods as benchmark. These methods were naive approaches and consisted on direct segmentation without analysing any properties of the image or the saliency map. The benchmark methods are described in Table 2.

Table 2: Benchmark algorithms.

Benchmark algorithms	Description
SM Otsu	Apply Otsu method on a saliency image
SM THRSH	Use an optimal value to segment the saliency map
LWIR THRSH CE	Use an optimal value to segment the LWIR image with contrast enhancement
LWIR Otsu CE	Apply directly Otsu method on a LWIR image with contrast enhancement
LWIR Otsu	Apply directly Otsu method on a LWIR image
LWIR THRSH	Use an optimal value to segment the LWIR image

The evaluation method assumes that each detection has a confidence score associated. In the proposed methods, HRSC and ESC, these scores are intrinsic to the method, but the other tested approaches do not have a confidence level associated. To overcome this hindrance, for evaluation purposes, we superimpose a unitary confidence value to each detection. In practice, this means that these different approaches will have only one working condition.

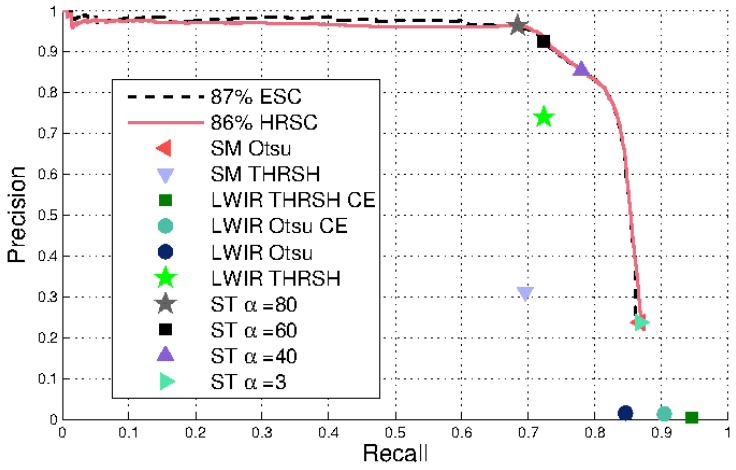

Fig. 6: Precision-Recall results obtained in the LWIR image sequence. Scores presented for ESC and HRSC represent the area under the curve.

Figure 6 present the scores of HRSC and ESC. These results, were compared against Saliency Threshold with different multiplying scalars and also the six combinations of direct segmentation. In the case of LWIR THRSH, LWIR THRSH CE and SM THRSH, we searched the optimal value by testing a sufficiently large range of values, and chose the one that, for each case achieved better precision-recall values.

The results of the six benchmark methods were poor, with numerous false positives, even when an optimal threshold value was used. These results indicate the need for a better detection strategy. We would like to mention that SM Otsu (which happens to be overlapped by ST with $\alpha = 3$) is a particular case of HRSC and ESC, when no criteria is applied to validate or discard images. Saliency Threshold with $\alpha = 3$ had poor results. However, when higher α were used, it attained a good performance, with its scores near the plots of HRSC and ESC. HRSC and ESC demonstrated a performance similar to Saliency Threshold but provided a confidence parameter that may be incorporated on another detection layer. For instance, it could be used on a temporal filter and therefore improve the overall performance of the system even further.

When analysing the absolute performance of the proposed methods, one of the results that we would like to focus on is the relatively low rate of false positives. This is an indicator that the strategies to compute the confidence score were adequate. After considering the different results between the histogram ratio approach and the entropy approach, we believe that the best option is to use the Entropy Saliency Confidence. Not only is a more tangible concept but on a system performing a search task, is usually preferable to have lower rate of false positives (higher precision) than a low rate of missed detections (high recall). As visible in Figure 6, the option based on entropy may be set to an

operating point with slightly higher precision, while maintaining a considerable recall. With the system running at a high frame rate, there is a considerable probability that the object may be detected in at least one frame.

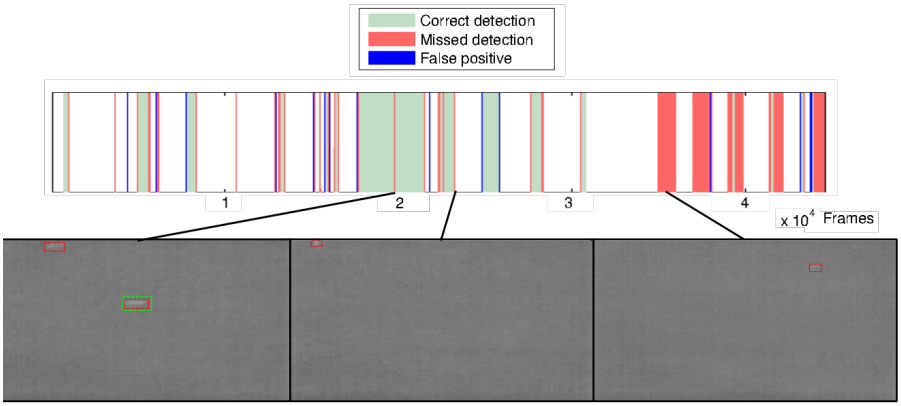

Fig. 7: Typical distribution along the image set of the occurrence of correct detections, missed detections and false positives. Area in white represents frames without boats. Three images bellow represent common cases of missed detections.

We consider the results particularly encouraging, because most of the missed detections consisted of demanding conditions even for a human operator. In Figure 7, we present some of the typical situations where the algorithm fails. The left and center images represent situations where boats are visible but at least one is near the border of the image. The right image represents a missed detection of a life raft. The life raft was the object that was mostly present from frame 36000 to 42500, which corresponds to the big areas highlighted in red. The detection of the life raft, however, is very challenging as can be seen in the image that already has its contrast enhanced for visualization purposes. One of the reasons for this difficulty was the fact that it was unmanned and had no heat source. Due to its small size and mass, the temperature converged rapidly to the ocean's temperature.

5 Conclusion and future work

This papers presents the use of a saliency algorithm to detect vessels in thermal images, captured by a UAV. This saliency algorithm is designed to separate foreground from background and therefore two strategies were proposed to deal with the cases where only ocean is visible. The tests on real data showed a good performance of the algorithm, more so because no temporal information is being considered. For future work we consider adding temporal information and

incorporate the confidence parameter in a temporal filter. Now, we considered the entropy of the entire saliency map but the computation of entropy of subsets of the map may be useful to deal with images with regions of different characteristics. These challenges, that span from sun glare, to wave crests, to bright sky, have hindered the application of saliency-based methods to visible spectrum images in maritime context.

Acknowledgement. This work was partially supported by project SEAGULL (QREN SI IDT 34063) and by FCT project [UID/EEA/50009/2013]. The authors would like to thank the team involved in obtaining these images and annotating the ground truth.

References

[1] Tara Brian and Frank Laczko. *Fatal Journeys Tracking Lives Lost during Migration*. 1st ed. Geneva: International Organization for Migration, 2014.

[2] Joint Agency Coordination Center. *Search for MH370 Facts and statistics*. Tech. rep. April. Australian Government, 2014.

[3] Matthew Dawkins et al. "Tracking nautical objects in real-time via layered saliency detection". In: *SPIE Defense+ Security*. International Society for Optics and Photonics. 2014, pp. 908903–908903.

[4] Piotr Dollar et al. "Pedestrian detection: An evaluation of the state of the art". In: *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 34.4 (2012), pp. 743–761.

[5] Peter Donaldson. "Cloudy skies". In: *Unmanned Vehicles* 18.5 (2013), pp. 32–34.

[6] Xiaodi Hou, Jonathan Harel, and Christof Koch. "Image Signature: Highlighting Sparse Salient Regions." In: *IEEE transactions on pattern analysis and machine intelligence* 34.1 (July 2011), pp. 194–201.

[7] Xiaodi Hou and Liqing Zhang. "Saliency detection: A spectral residual approach". In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 800 (2007).

[8] Laurent Itti, Christof Koch, and Ernst Niebur. "A model of saliency-based visual attention for rapid scene analysis". In: *Pattern Analysis and Machine Intelligence, IEEE* 20.11 (1998), pp. 1254–1259.

[9] Harnit Kaur Kang. "Gulf of Aden vs Malacca Strait: Piracy and Counter-piracy efforts". In: *Institute of Peace and Conflict Studies* (2009).

[10] M. Kontitsis, K.P. Valavanis, and N. Tsourveloudis. "A UAV vision system for airborne surveillance". In: *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004* 1 (2004).

[11] *Maritime Security and Surveillance - Case Study*. Tech. rep. January. Centre for Strategy and Evaluation Services, 2011.

[12] Luis Merino et al. "An unmanned aircraft system for automatic forest fire monitoring and measurement". In: *Journal of Intelligent and Robotic Systems: Theory and Applications* 65 (2012), pp. 533–548.

[13] Nobuyuki Otsu. “A threshold selection method from gray-level histograms”. In: *Automatica* 11.285-296 (1975), pp. 23–27.

[14] Lawrence Rabiner. “A tutorial on hidden Markov models and selected applications in speech recognition”. In: *Proceedings of the IEEE* 77.2 (1989), pp. 257–286.

[15] Sivakumar Rathinam et al. “Autonomous searching and tracking of a river using an UAV”. In: *Proceedings of the American Control Conference* (2007), pp. 359–364.

[16] Piotr Rudol and Patrick Doherty. “Human body detection and geolocalization for UAV search and rescue missions using color and thermal imagery”. In: *IEEE Aerospace Conference Proceedings* (2008).

[17] Michael Teutsch and Wolfgang Kruger. “Classification of small boats in infrared images for maritime surveillance”. In: *2010 International WaterSide Security Conference* (Nov. 2010), pp. 1–7.

[18] Michael Teutsch and Wolfgang Kruger. “Detection, Segmentation, and Tracking of Moving Objects in UAV Videos”. In: *2012 IEEE Ninth International Conference on Advanced Video and Signal-Based Surveillance* (2012), pp. 313–318.

[19] Alexander Toet. “Detection of dim point targets in cluttered maritime backgrounds through multisensor image fusion”. In: *AeroSense 2002* 4718 (2002), pp. 118–129.

[20] Ronda Venkateswarlu et al. “Design considerations of IRST system”. In: *AeroSense’97*. International Society for Optics and Photonics. 1997, pp. 591–602.

[21] A Wadsworth et al. “Aircraft experiments with visible and infrared sensors”. In: *International Journal of Remote Sensing* 13.6-7 (1992), pp. 1175–1199.

[22] Paul Westall et al. “Improved Maritime Target Tracker using Colour Fusion”. In: *High Performance Computing & Simulation, 2009. HPCS ’09. International Conference on,* 2009, pp. 230–236.

[23] Tianzhu Xiang, Li Yan, and Rongrong Gao. “A fusion algorithm for infrared and visible images based on adaptive dual-channel unit-linking PCNN in NSCT domain”. In: *Infrared Physics & Technology* 69 (2015), pp. 53–61.

[24] Shuqun Zhang. “Object Tracking in Unmanned Aerial Vehicle (UAV) Videos Using a Combined Approach”. In: *Proceedings. (ICASSP ’05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.* 2 (2005), pp. 681–684.