Pappas, O. A., Achim, A., & Bull, D. (2017). Superpixel-guided CFAR Detection of Ships at Sea in SAR Imagery. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2017): Proceedings of a meeting held 5-9 March 2017, New Orleans, Louisiana, USA* (pp. 1647-1651). Institute of Electrical and Electronics Engineers (IEEE). <https://doi.org/10.1109/ICASSP.2017.7952436>

Peer reviewed version

Link to published version (if available):
[10.1109/ICASSP.2017.7952436](https://doi.org/10.1109/ICASSP.2017.7952436)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via IEEE at <http://ieeexplore.ieee.org/document/7952436/>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

SUPERPIXEL-GUIDED CFAR DETECTION OF SHIPS AT SEA IN SAR IMAGERY

Odysseas A. Pappas, Alin M. Achim, David R. Bull

Visual Information Laboratory, University of Bristol, Bristol BS8 1UB, UK

ABSTRACT

Synthetic Aperture Radar (SAR) has over the years evolved to be one of the most promising remote sensing modalities for large-scale monitoring of the ocean and maritime activity. The detection of ships at sea in SAR imagery is a challenging task, as it requires the detection of small targets with little exploitable spatial information within a high resolution image. We present a novel method for the detection of ships based on superpixel segmentation and subsequent statistical characterisation, with no prior land masking. Our method acts as a bound to a CFAR detector, greatly reducing false positives. We present results on SENTINEL-1 imagery, demonstrating the detection performance of our algorithm.

Index Terms— Synthetic Aperture Radar, Superpixels, CFAR, Detection

1. INTRODUCTION

Synthetic Aperture Radar has proven to be an excellent remote sensing modality for a plethora of monitoring tasks, including the monitoring of maritime activity and the detection of ships at sea. This is generally a rather complicated task as ships tend to appear as extremely small targets within high-resolution SAR images, providing little shape information to work with. The significant clutter and speckle noise present in SAR images further complicates this detection task.

As a general rule, ships at sea appear as bright targets as the radar cross section (RCS) of a ship is much higher than that of sea clutter. This is due to the large number of radar wave bounces caused by the ship's metal superstructure (compared to the sea surface) [1]. Detection via an appropriate RCS threshold is thus possible in cases where the sea clutter is very homogeneous, as is often the case with low-resolution SAR data.

Modern high-resolution SAR platforms (like TerraSAR-X and SENTINEL-1) are capable of distinguishing sea waves, ship wake and other clutter, causing sea clutter to appear more heterogeneous in high resolution data and thus complicating the detection task. The presence of distinguishable sidelobes in azimuth and range can also potentially cause multiple areas of bright and dark texture, leading to multiple detections for a single target [2].

Most approaches proposed in the literature are centred

around the use of a Constant False Alarm Rate (CFAR) detector and variants thereof. The standard cell-averaging CFAR detector estimates the level of noise around a pixel of interest by calculating the average power level in a neighbourhood around the pixel of interest, while exempting those in a band directly around it (guard cells).

Various methods like cell averaging CFAR, Greatest-of-CFAR, excision-switching-context CFAR [3] [4] and two-parameter CFAR [1] have been employed to varying success. Sea clutter has been modelled using a number of distributions, such as the Gaussian for the default CFAR case and more recently heavy tailed distributions like the Weibull, K -distribution [5] and Alpha-stable distribution [6].

Hybrid approaches include Meyer's fusion of CFAR detectors and wavelet analysis [7] as well as Wang's method of fast block CFAR combined with feature analysis [2]. Leng et al. [8] employ a bilateral CFAR algorithm acting on a combination of the intensity and spatial distribution of the image, while Hou et al. [9] use an iterative Log-normal CFAR combined with background subtraction at each stage.

Most methods cannot distinguish between land and sea regions within the SAR image. It is assumed that any land cover present in the image is segmented out prior to ship detection, and even then false detections at the water-land interface are possible due to inaccuracies in the land masking process, the presence of docks and similar structures etc. [2].

We here propose a two stage algorithm combining superpixel regions of interest with a CFAR detector. After producing a superpixel segmentation of the image, our proposed method makes use of statistical moments such as range, deviation and kurtosis to identify possible ship targets contained within superpixels across the SAR image. By combining the above measures with the Hartigan dip test for bimodality our algorithm is also able to distinguish between land and sea coverage, hence no longer necessitating land masking prior to ship detection. The regions of interest produced by this first stage act as a bound for the CFAR detector, leading to accurate detection of ship targets while greatly reducing the number of false positives.

The paper is organised as follows: Section 2 provides a brief introduction to superpixels and CFAR detectors. The proposed method is outlined in Section 3 with results using SENTINEL-1 data and discussion on performance found in Section 4, followed by a brief conclusion and future work.

2. BACKGROUND INFORMATION

2.1. Superpixel Segmentation

Superpixels have emerged as a form of over-segmentation meant to provide clusters of perceptually similar pixels throughout an image, thus capturing the redundancy inherent in most natural images.

Various superpixel methods have been proposed in the literature and consequently employed as a convenient primitive for further image processing tasks, with applications including image and video segmentation [10] [11], object localisation [12] and tracking [13].

The results of the various existing superpixel algorithms can vary in segmentation quality, superpixel uniformity, size and number, partly due to the lack of a consistent, rigorous definition of what constitutes a superpixel. One of the most widely used algorithms, proposed by Achanta et al. is *Simple Linear Iterative Clustering* (SLIC) [10]. SLIC segments an image according to a 5-dimensional distance metric comprised of spatial (x, y coordinates) and colour information (L, a, b , of the CIELAB colorspace) as shown in the following equations. The iterative clustering process of SLIC is similar in philosophy to k -means clustering [10].

$$d_{lab} = \sqrt{(l_k - l_i)^2 + (a_k - a_i)^2 + (b_k - b_i)^2} \quad (1)$$

$$d_{xy} = \sqrt{(x_k - x_i)^2 + (y_k - y_i)^2} \quad (2)$$

$$D_s = d_{lab} + \frac{m}{S} d_{xy} \quad (3)$$

The number of generated superpixels k (and indirectly, the size of those) is specified by the user. As the spatial component of the metric can dominate the end result when dealing with large superpixels (due to high image resolution or low k), a scaling factor is introduced as $\frac{m}{S}$, where S is the initial cluster seed grid interval (dependent on image size and desired k) and m allows the user to control superpixel compactness and shape regularity.

Advantages of SLIC superpixels include high perceptual homogeneity within each superpixel, uniformity in size and shape, computational simplicity and a degree of user control over the process. They perform very well in terms of standard boundary recall and under-segmentation error measures and can cope with both colour and grayscale imagery [10].

2.2. Constant False Alarm Rate Detector

A typical CFAR detector relies on the comparison of a pixel of interest (or cell-under-test in radar terminology) against its local background, with the exemption of a guard band of immediately surrounding pixels, as they may also form part of the target. Assuming a distribution model $f_{pdf}(x)$ for the background, the probability of false alarm for a given threshold T

Fig. 1. Example of SAR superpixels. Sea region visible along with superpixel containing a ship target and noticeable ship wake.

is given by:

$$P_{FA} = 1 - \int_{-\infty}^T f_{pdf}(x) dx = \int_T^{\infty} f_{pdf}(x) dx \quad (4)$$

Sea clutter in SAR data has been modelled successfully with a variety of distributions, including the log-normal distribution [9], the alpha-stable distribution [3], the Weibull and the K-distribution [5]. We have opted to use the Weibull distribution in this paper, the PDF of which is given by the following equation:

$$f_{pdf}(x) = \frac{C}{B} \left(\frac{x}{B}\right)^{C-1} \exp\left\{-\left(\frac{x}{B}\right)^C\right\}; x \geq 0 \quad (5)$$

where B and C are the scale and shape parameter of the Weibull distribution respectively.

After obtaining estimates $\hat{B}$ and $\hat{C}$ we can express the probability of false alarm in relation to threshold Th as [5]:

$$P_{FA} = \exp\left[-\left(\frac{Th}{\hat{B}}\right)^{\hat{C}}\right] \quad (6)$$

where the threshold value Th can be derived from (6) for any desired P_{FA} (a typical value being 0.001).

3. PROPOSED ALGORITHM

The image is initially segmented into SLIC superpixels whose size is significantly larger than that of any expected targets. Any ship targets are thus contained, along with their immediate sea region, into a larger superpixel (Fig. 1) [14].

Most superpixels contain pixels of similar intensity values and thus have relatively flat histograms. This uniformity is disrupted in the presence of a target, manifesting itself as a collection of extremely bright pixels (in comparison to the sea that typically registers dark values).

Fig. 2. (a) SENTINEL-1 Image 2, showing ship traffic off the coast of the Isle of Wight. (b) Land masking via Hartigan's dip. (c) Histogram of dip values across image. Note H-dip values throughout the image are scattered around two central values, roughly corresponding to the portions of land and sea mass in the image.

In the context of remote sensing SAR imagery, superpixels typically fall into one of three content categories: superpixels containing land mass, sea mass and superpixels containing sea along with a smaller foreign object (such as a ship).

Firstly, the distinction between a superpixel containing only sea and one containing a ship at sea needs to be formulated; examples of both can be seen in Fig. 1. Sea regions typically appear as uniform, homogeneous dark regions in SAR imagery due to the low RCS of the sea. In the case of modern high-resolution SAR systems waves, ship wake and bright clutter can become distinct to certain extent. A ship target however, producing a much higher RCS, appears much brighter. A superpixel containing a target would then display a comparatively heavy-tailed histogram and much higher statistical dispersion than a predominantly dark superpixel.

This behaviour of a target superpixel is reflected in statistical moments such as range and standard deviation, as well as in measures indicating heavy tails such as sample kurtosis (Equation 7). Other possible indicators of leptokurticity include distance skewness [16] as well as General Likelihood Ratio Tests comparing between a platykurtic and a leptokurtic distribution [14].

$$g_2 = \frac{m_4}{m_2^2} = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^4}{\left(\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2\right)^2} \quad (7)$$

Our proposed algorithm utilises an additional metric to detect (and discard) superpixels containing land mass, thus no longer necessitating land masking prior to detection. Land in SAR imagery appears far more heterogeneous than sea and contains a plethora of intensity values across a large range. This is intuitively supported by the fact that a typical land segment would provide diverse and numerous scatterers in the form of geographical features, elevations, man made structures, vegetation etc.

Hartigan's Dip [15] is a metric for determining whether a set of data follows an unimodal or bimodal distribution. It measures bimodality in a sample by the maximum difference between the empirical distribution function and the unimodal distribution function that minimises said maximum difference (usually a uniform distribution). Lower values of the dip metric indicate unimodality and vice versa.

Land segments, containing a good spread of intensity values can easily be modelled by a unimodal distribution and produce low values of the dip metric. Sea segments however, containing either a target ship or just intrinsic noise (system noise, speckle, wave-induced scatterers) provide histograms that appear highly leptokurtic or distinctly bimodal, this being reflected in their dip scores. A thresholding operation can thus differentiate land from sea using this metric (Figure 2).

Superpixels are selected as regions containing a possible target via a thresholding operation. The threshold for each metric is set as the average observed throughout the image; this can be considered a safe, conservative option as the majority of superpixels throughout the image are largely uniform, dominating the resulting global averages. Anomalous superpixels containing a ship are effectively guaranteed to present range, deviation, kurtosis etc scores higher than the average.

After the superpixel detection stage we perform CFAR detection over the image, using the Weibull distribution which has been shown to be a good fit for SAR sea clutter [5]. The detections returned by the CFAR detector are bound by the superpixel mask, i.e. a detection is only considered valid if it lies within one of the superpixels detected in the previous stage.

4. RESULTS

We present results of our algorithm on four SAR images from the SENTINEL-1 platform. The instrument is a C-band Synthetic Aperture Radar operating in Interferometric Wide Swath mode (VH polarised), acquiring data in a 250km swath at a spatial resolution of 5m x 20m.

The area depicted is part of the English Channel. As there is no ground truth available for the visible sea traffic, it has been evaluated visually using direct ship radar return signatures as well as transverse and turbulent ship wake, in accordance with the information and guidelines provided in the SAR Marine User's Manual [17].

Fig. 3. (a) SENTINEL-1 Image 1. (b) Detection of superpixels containing possible ships at sea. (c) Weibull CFAR and (d) CFAR combined with superpixel mask. Hartigan's Dip threshold at 0.01.

The SENTINEL-1 images (1670 x 2619) have been segmented into $k = 1600$ SLIC superpixels. This k value was chosen after visual inspection of the images; the approach however is fairly robust to superpixel size and produces comparable results for other superpixel sizes (e.g. $k = 1200, k = 2000$). The superpixel compactness value m was set to 25, again not critical to the current application.

We opt to use one statistical measure describing dispersion (range) and one for leptokurtosis (sample kurtosis). Thresholds for both have been set at the global superpixel average while the Weibull CFAR P_{FA} is set at 0.001. Thresholds for the Hartigan dip can be varied according to the degree of seclusion of land data required; for the data presented in this paper this threshold has been set at 0.01. This has been found to correlate well to the SENTINEL-1 data under question, as seen in Figure 2. For completeness, we demonstrate the effect of an increased H_{dip} threshold in Figure 4.

The image shown in Figure 3 contains 5 sea vessels as well as a portion of land mass in the lower left corner. The algorithm is able to detect all 5 targets, contained over 6 superpixels throughout the image. Figure 4 shows heavy ship traffic (17 vessels) off the coast of the Isle of Wight. To demonstrate the effect of Hartigan's Dip on the superpixel mask we show results for both $H_{dip} = 0.01$ and $H_{dip} = 0.02$. In the first case, the low threshold leads to the detection of some superpixels along the coastline, in turn allowing for some CFAR detections in that area. Increasing the threshold greatly reduces these false positives but can eventually lead to the introduction of false negatives (missed targets). The reduction of false positives compared to normal CFAR for four SENTINEL-1 images can be seen in Table 1.

For the majority of test cases the proposed detector does not produce any false negatives, having all targets present in

Fig. 4. (a) SENTINEL-1 Image 2. (b) Weibull CFAR. Superpixel mask for (c) $H_{dip} = 0.01$ and (d) $H_{dip} = 0.02$. Combined CFAR and Superpixel detection for (e) $H_{dip} = 0.01$ and (f) $H_{dip} = 0.02$.

the image retained by the superpixel mask and subsequently detected by CFAR. This is often not the case with other methods, such as CA-CFAR which has been shown to miss targets in the presence of strong sea clutter [4] or the detector completeness of 72% in [7].

Dataset	Weibull CFAR	Proposed Method	Ships Detected
Image 1	91445	1251	5/5
Image 2A	133247	2495	17/17
Image 2B	133247	1573	15/17
Image 3	121630	230	4/4
Image 4	103757	2385	4/4

Table 1. Total number of detected pixels for Weibull CFAR and the proposed method, along with detected targets per image.

5. CONCLUSION

In this paper we have presented a method for the detection of ships at sea in SAR remote sensing imagery, using superpixel segmentation and statistical detection followed by Weibull CFAR detection. We present results on SENTINEL-1 data that demonstrate the efficiency of our algorithm not only on the detection of ships at sea but also on automatically discarding any land segments present in the image without prior masking.

6. REFERENCES

- [1] J. Ai, X. Qi, W. Yu, Y. Deng, F. Liu, and L. Shi, "A New CFAR Ship Detection Algorithm Based on 2-D Joint Log-Normal Distribution in SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 7, pp. 806–810, 2010.
- [2] C. Wang, S. Jiang, H. Zang, F. Wu, and B. Zhang, "Ship Detection for High-Resolution SAR Images Based on Feature Analysis," *IEEE Radar Conference*, pp. 570–574, 2012.
- [3] X. Xing, K. Ji, H. Zou, J. Sun, and S. Zhou, "High Resolution SAR Imagery Ship Detection Based on EXS-CFAR in Alpha-Stable Clutters," *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 316–319, 2011.
- [4] X. Xing, K. Ji, H. Zou, and J. Sun, "A Fast Ship Detection Algorithm in SAR Imagery for Wide Area Ocean Surveillance," *IEEE Geoscience and Remote Sensing Letters*, vol. 11, pp. 119–123, 2014.
- [5] S. Kuttikkad and R. Chellappa, "Non-Gaussian CFAR Techniques for Target Detection in High Resolution SAR Images," *IEEE International Conference on Image Processing*, pp. 910–914, 1994.
- [6] M. Liao, C. Wang, Y. Wang, and L. Jiang, "Using SAR Images to Detect Ships From Sea Clutter," *IEEE Geoscience and Remote Sensing Letters*, pp. 194–198, 2008.
- [7] F. Meyer and S. Hinz, "Automatic Ship Detection in Space-borne SAR Imagery," *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, pp. 1–4–7/W5, 2009.
- [8] X. Leng, K. Ji, K. Yang, and H. Zou, "A Bilateral CFAR Algorithm for Ship Detection in SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 12, pp. 1536–1540, 2015.
- [9] B. Hou, X. Chen, and L. Jiao, "Multilayer CFAR Detection of Ship Targets in Very High Resolution SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 12, pp. 811–815, 2015.
- [10] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk, "SLIC Superpixels Compared to State-of-the-Art Superpixel Methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 2274–2282, 2012.
- [11] R. Trichet and R. Nevatia, "Video Segmentation with Spatio-Temporal Tubes," *IEEE AVSS International Conference on Advanced Video and Signal-Based Surveillance*, pp. 330–335, 2013.
- [12] B. Fulkerson, A. Vedaldi, and S. Soatto, "Class Segmentation and Object Localization with Superpixel Neighborhoods," *IEEE 12th International Conference on Computer Vision*, pp. 670–677, 2009.
- [13] F. Yang, H. Lu, and M. H. Yang, "Robust Superpixel tracking," *IEEE Transactions on Image Processing*, vol. 23, pp. 1639–1651, 2014.
- [14] O. Pappas, A. Achim, and D. Bull, "Superpixel-based Statistical Anomaly Detection for Sense and Avoid," *IEEE International Conference on Image Processing ICIP*, pp. 2229–2233, 2015.
- [15] J. A. Hartigan and P. M. Hartigan, "The Dip Test of Unimodality," *The Annals of Statistics*, vol. 13, pp. 70–84, 1985.
- [16] G. J. Szekely, "Pre-limit and Post-limit Theorems for Statistics," in *Statistics for the 21st Century*, C. R. Rao and G. J. Szekely, Eds., pp. 411–422. Dekker, New York, 2000.
- [17] W. G. Pichel, P. Clemente-Colon, C. C. Wackerman, and K. S. Friedman, "Ship and wake detection," in *Synthetic Aperture Radar Marine User's Manual*, C. R. Jackson and J. R. Apel, Eds., chapter 12, pp. 277–303. NOAA/NESDIS, U.S. Department of Commerce, 2004.