Superpixel-Level CFAR Detectors for Ship Detection in SAR Imagery

Odysseas Pappas, Alin Achim, *Senior Member, IEEE*, and David Bull, *Fellow, IEEE*

Abstract—Synthetic aperture radar (SAR) is one of the most widely employed remote sensing modalities for large-scale monitoring of maritime activity. Ship detection in SAR images is a challenging task due to inherent speckle, discernible sea clutter, and the little exploitable shape information the targets present. Constant false alarm rate (CFAR) detectors, utilizing various sea clutter statistical models and thresholding schemes, are near ubiquitous in the literature. Very few of the proposed CFAR variants deviate from the classical CFAR topology; this letter proposes a modified topology, utilizing superpixels (SPs) *in lieu* of rectangular sliding windows to define CFAR guardbands and background. The aim is to achieve better target exclusion from the background band and reduced false detections. The performance of this modified SP-CFAR algorithm is demonstrated on TerraSAR-X and SENTINEL-1 images, achieving superior results in comparison to classical CFAR for various background distributions.

Index Terms—Constant false alarm rate (CFAR), ship detection, superpixels (SPs), synthetic aperture radar (SAR).

I. INTRODUCTION

SYNTHETIC aperture radar (SAR) has established itself in recent years as an excellent remote sensing modality for many monitoring tasks, largely due to its round-the-clock operation capabilities, independent of weather conditions and cloud coverage. This gives it a clear advantage over optical panchromatic and other sensors making it a prime candidate for tasks such as agricultural monitoring, mapping, search, and rescue, and for the monitoring of maritime activity.

A key aspect of maritime monitoring is the detection of ships at sea. This can be a challenging task as ships tend to appear as small targets within high-resolution SAR images, providing little discernible shape information. Detection is, therefore, often based on the principle of ship targets appearing quite bright, as the multiple radar wave reflections of the ship's metal superstructure produce a radar cross section much higher than that of sea clutter [1].

In the high resolution SAR images produced by modern, state-of-the-art platforms, the ship's wake, the sea state (i.e., presence and size of waves), and other clutter become more discernible; this can result in a much more heterogeneous background over which ship target detection is to be performed. The literature on ship detection in SAR images reveals a strong trend toward the use of constant false alarm rate (CFAR) detectors and variants thereof. CFAR detectors are adaptable threshold detectors that operate by estimating the statistics of the sea clutter around a possible target according to an assumed background probability density function $p(x)$, so as to maintain a constant acceptable probability of false alarm P_{FA}.

Of the proposed CFAR variants, many focus on the identification of suitable background modeling distributions (an integral part in the performance of any CFAR detector) [2]. SAR images exhibit heavy-tailed histograms (often significantly so) [3]; this is particularly true in the case of modern, state-of-the-art SAR systems. As the spatial resolution of SAR instruments increases, the associated reduction of scatterers per resolution cell lends to an increase of the appreciability of backscattering responses from distinct ground features, which means SAR images with complex land/sea topologies now exhibit even more heavy-tailed and/or bimodal histograms [4]. Typical distributions used in CFAR include leptokurtic distributions such as the Weibull, Rayleigh, and K-distributions [2]. Recent efforts have extended into the *alpha*-stable distribution [5], the generalized heavy-tailed Rayleigh [3], [6], and the generalized Gamma and Generalized Gamma Mixture Model [7].

Besides the various statistical models proposed, other avenues of research pursued in recent years include variations on the CFAR thresholding function itself, often making use of order statistics (OS) (as in the case of classic OS-CFAR [8] and more recent methods employing truncation [truncated satistics (TS)-CFAR] [9]) or more complex architectures such as excision-switching-CFAR [5]. CFAR performance can also be improved by preprocessing algorithms such as sublook correlation magnitude [10], which aim to increase the contrast between ship targets and sea surface, or geometrical perturbation analysis [11] when polarimetric data are available.

The use of superpixels (SPs) has also been introduced in ship detection in SAR images as an attempt to provide regions of interest that serve as constraints for the following CFAR detectors [12], as well as a way to better define clutter pixels and alleviate the effects of multiple targets present (applied to the more general target detection problem in [13]). More recent attempts at introducing SPs to CFAR-based ship detection include [14], which builds on [13] and incorporates it into a larger framework that includes the identification of potential target SPs using weighted information entropy, CFAR detection, and a final post-processing stage. Ship detection in PolSAR imagery via scattering mechanism distribution features on a SP level has also been addressed in [15].

We here propose a novel approach to the classical CFAR guard and background band topology which remains to this day ubiquitous throughout the majority of the literature. Our method relies on the use of SPs *in lieu* of guard and background bands, with the aim to achieve better target exclusion within the guardband, improved clutter homogeneity in the background band, and robustness to the presence of multiple targets. This SP-level CFAR (SP-CFAR) architecture can be utilized with most prior sea clutter statistical models or thresholding schemes.

The remainder of this letter is organized as follows: Section II provides a brief background discussion on CFAR detectors while Section III outlines the proposed SP-CFAR variant. Section IV is a comparative demonstration of the proposed SP-CFAR topology against the classical CFAR while Section V outlines the conclusions of this letter.

II. CFAR Detectors

CFAR detectors are a type of sliding window, adaptable threshold detectors that first originated within radar target detection literature. The threshold is varied for every cell under test of the radar return signal so that the probability of a false alarm (P_{FA}) [as defined according to an assumed background/clutter probability density function $p(x)$] remains constant. This constant P_{FA} property, of course, only holds provided the background probability density function (pdf) $p(x)$ fits the data.

The detector uses the neighboring cells (preceding and succeeding) of the cell under test to evaluate the clutter pdf; in order to avoid including any cells that also belong to the target (and which could skew the pdf and subsequent detection threshold calculation), a small number of cells immediately before and after the cell under test are excluded from any calculations.

This topology has carried over to 2-D signals, such as the SAR imagery discussed here, where it translates to a guard band/background square window topology centered around a pixel of interest.

One of the earliest forms of CFAR is cell averaging (CA) CFAR which relies on simple averaging over the N pixels in the background band to estimate the detection threshold

$$T_{CA} = \sum_{i}^{N} (x_i) \left(P_{FA}^{\frac{1}{N}} - 1 \right). \tag{1}$$

CA-CFAR, however, performs poorly in the presence of nonhomogeneous, heavy-tailed background clutter (which is the *de facto* case for SAR images). What has instead become more common is to utilize a statistical distribution model for the background sea clutter. Assuming a distribution with probability density function $p(x)$, standard CFAR formulation states that detection threshold T and the probability of a false alarm P_{FA} are related by

$$1 - P_{FA} = \int_{0}^{T} p(x)dx \tag{2}$$

and therefore

$$P_{FA} = \int_{T}^{\infty} p(x)dx. \tag{3}$$

Thus, having settled on a statistical model and the acceptable P_{FA} for a particular application, an analytical expression for the threshold T can be sometimes derived from (3).

The literature abounds with proposed statistical models of SAR for use with CFAR detectors; for completeness, we iterate here the Weibull distribution, a commonly employed distribution whose pdf is given in the following equation:

$$p_W(x) = \begin{cases} \dfrac{k}{\lambda} \left(\dfrac{k}{\lambda} \right)^{k-1} \exp^{-(x/\lambda)^k} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \tag{4}$$

where λ is the scale parameter and k the shape parameter of the distribution.

Making use of (3), the CFAR selection threshold for the Weibull case is given in the following equation:

$$T_W = -\lambda \sqrt[k]{\ln(P_{FA})}. \tag{5}$$

Parameter estimation for the Weibull distribution using methods like maximum likelihood estimation is standard statistics curriculum that can be readily found in the related literature.

Attempts to provide robustness, especially in cases of multiple targets present, have led to the aforementioned use of TS [9]. The aim here is to provide, via truncation, an accurate estimate of clutter statistics when multiple outliers are present in the sample. The threshold for single-look TS-CFAR (assuming an exponential distribution model for the image) is given in the following equation:

$$T_{TS} = -\hat{\mu} \ln(P_{FA}) \tag{6}$$

with the estimated mean $\hat{\mu}$ obtained according to the following equation:

$$\hat{\mu} = \frac{t}{\exp^{\frac{t}{\hat{\mu}}} - 1} + \frac{1}{n} \sum_{i=1}^{n} \tilde{x}_i \tag{7}$$

where t is the truncation threshold and $\tilde{x}_{i=1}^{n}$ is a size n sample of truncated measurements.

III. Superpixel-Level CFAR

The classic rectangular guard/background band topology is relatively ubiquitous across CFAR literature. Here, we propose a fundamental change to this architecture, utilizing SPs to define the guard and background bands.

Superpixel methods aim to oversegment an image into homogeneous segments, i.e., groupings of perceptually similar pixels that can act as meaningful priors for further processing tasks. Such tasks can include image and video segmentation, object localization, and tracking [16]. Various SP segmentation methods have been proposed in the literature—as SP segmentation methods *per se* fall beyond the scope of this letter, the reader is directed to [16] and [17] for an overview of some of the most widely used methods.

What is of more interest are certain properties of SPs, the most fundamental being that of content homogeneity. As a direct consequence of this property comes the property of boundary adherence; SPs aim to align their boundaries with the boundaries of objects along an image, trying to contain little edge information. This also leads to a situation where

Fig. 1. Ship target contained within an SP. SENTINEL-1 SAR image, VH polarized.
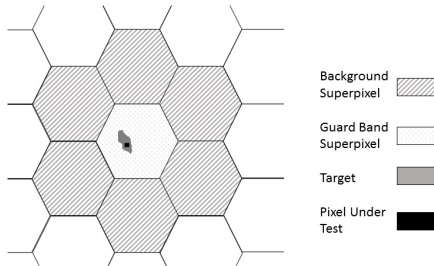

Fig. 2. Example SP-CFAR topology, illustrating a target and pixel of interest contained within an SP guard band, with the adjoining SPs defining the local background.

any object in the image that is significantly smaller than the average SP size tends to be encapsulated in its entirety—along with its immediate background—within a larger SP [12]. This behavior, seen in Fig. 1, is a key aspect of our proposed method.

To this effect, we propose here a modification to the classical CFAR topology that utilizes SPs to define the guardband and background regions. First, we oversegment the image into SPs whose size (in pixels) is significantly larger than that of any expected target (which can be estimated given the SAR instrument's spatial resolution and the typical size of a ship target of interest). This means that any ship targets present in the image will each be contained within a singular SP. We then retain the detector operating on the pixel level, but now define the guardband to be excluded as the entirety of the SP containing the pixel under test, while the background band is defined as the collection of adjoining, border-sharing SPs (see Fig. 2).

This new topology offers a number of advantages over classical CFAR, particularly in achieving proper target inclusion within the guardband. The containment property of the SPs ensures that the entirety of the target is excluded from the threshold calculation (according to the background pdf). This is often not the case with typical CFAR guardbands, where the band may fail to completely cover the entirety of the target (particularly for the pixels in its extremities) depending on how suitable the *ad hoc* selected window size is to a particular target. This can lead to poor target shape retention at best and even false negatives (missed targets).

An SP defined guardband is potentially more capable of dealing with multiple targets present close to each other. Rectangular guardbands are typically sized with the expected target size in mind, so that they cover all target-comprising pixels; consequently, any second target present in the immediate vicinity of a target under test is extremely unlikely to be adequately covered by a guard band whose size is

TABLE I
Count of CFAR Detections and Retained Targets for the Configuration of Fig. 3

TerraSAR-X	CFAR			SP-CFAR		
	CA	Weib.	TS	CA	Weib.	TS
C1 Detections	70	1,294	754	778	686	792
C1 Targets	2/11	11/11	11/11	11/11	11/11	11/11
C2 Detections	256	930	826	936	515	1,024
C2 Targets	8/8	8/8	8/8	8/8	8/8	8/8

approximately that of the expected target under test. Note that the presence of a second target within the background band of a pixel under test is an overarching problem across CFAR detection literature that has been addressed to varying degrees of success [5], [9].

A single SP can, however, contain one or more targets in their entirety if these happen to be located relatively close to each other, meaning the threshold calculation for both will not be adversely affected by each other's presence. This is of course not guaranteed; two targets may well end up in adjoining SPs, meaning they would then be present in each other's background bands. Note that while this could also happen with regular rectangular neighborhoods, in the case of larger SP neighborhoods, the effect of this second target presence would be more diluted among the rest of the background pixels.

IV. RESULTS

The performance of the proposed SP-CFAR topology is demonstrated on TerraSAR-X and SENTINEL-1 SAR images. The TerraSAR-X instrument was operated in ScanSAR mode, imaging a 100-km swath over the Panama Canal at a spatial resolution of 18.5 m, and the SENTINEL-1 instrument was operated in interferometric wide-swath mode, imaging a 250-km swath over the English Channel at a spatial resolution of 5×20 m^2. Polarization is VH in both cases. As the sea traffic ground truth was not available, ships were detected by visual inspection in accordance with [18].

The simple linear iterative clustering (SLIC) algorithm [16] was selected for SP segmentation. SLIC is capable of producing relatively uniform (in shape and size) SPs that maintain good homogeneity and boundary adherence. SP size is indirectly controlled by specifying the desired number k of generated SPs. However, any other SP segmentation method that achieves similar performance (particularly in terms of size uniformity and control) could be used for SP-CFAR.

Fig. 3 shows TerraSAR-X data utilizing CA-CFAR, Weibull CFAR, as well as TS-CFAR against their SP-CFAR equivalents for ship detection, with detection counts present in Table I. The conventional CFAR algorithms use a guardband of 10×10 pixels and a background of 30×30, sizes which have been experimentally found to work well, given the image resolution and target size.

All three SP-CFAR variants included here show promising performance, with results comparable to the state-of-the-art TS-CFAR detector. False negatives are absent (unlike in the case of CA-CFAR), detected targets appear with good shape

Fig. 3. Comparison of CFAR and SP-CFAR with $P_{FA} = 0.001$ over TerraSAR-X image, sea state: 2. (a) Crop 1 (412×352). (b) CA-CFAR. (c) Weibull CFAR. (d) TS-CFAR. (e) SLIC Segmentation ($k = 300$). (f) CA-SP-CFAR. (g) Weibull SP-CFAR. (h) TS-SP-CFAR. (i)–(p) as above for Crop 2.

retention and false positives in the Weibull case are significantly reduced. The SP-CFAR variants are also quite resilient in the presence of small features such as the pier visible in Crop 1. TS-CFAR and TS-SP-CFAR produce similar results, something that can be to a degree intuitively supported by the fact that TS effectively aim to address the same "training contamination" problem that SP neighborhoods are addressing here, hence, the combination of both may not necessarily offer an increase in performance.

Fig. 4 shows selected high-resolution crops from an SENTINEL-1 image with Weibull CFAR and SP-CFAR applied. The SP-CFAR variant produces a far smaller amount of false positive detections, and all targets are retained in both images (5 and 17 vessels, respectively). Note that both images contain visible land mass; masking of any land mass present is typically an assumed prior to CFAR ship detection, as the presence of land will introduce false positives (seen when applying Weibull CFAR across the entire frame). The Weibull SP-CFAR, however, demonstrates notable resilience to land-mass-caused false positives, possibly to the point where omitting prior segmentation would be a viable option.

SP-CFAR detectors have, as mentioned, appeared previously in the literature, albeit utilizing an architecture somewhat different from the one presented here. For example, the local CFAR stage described in [14] uses much smaller SPs, to the point where a target would comprise of multiple SPs. The background band is made up of SPs found along a rectangle outline at a certain distance (radius) from the center of the SP of interest. Measures such as weighted information entropy can be used to identify and exclude potential interfering target SPs from this background band. Note that the statistical model for clutter is Gaussian [14]. The results of an SP-CFAR detector of this type applied to the presented TerraSAR-X data can be seen in Fig. 5; while detection performance is good, robustness

Fig. 4. SENTINEL-1 Images, sea state: 2. (a) and (b) Crop 1 and Crop 2 (1670×2619) with SLIC SP boundaries drawn ($k = 2500$). (c) and (d) Weibull CFAR. (e) and (f) Weibull SP-CFAR. $P_{\text{FA}} = 0.00001$.


Fig. 5. Example of an SP-CFAR variant of the type described in [14] applied to TerraSAR-X Image. (a) Crop 1 detections. (b) Crop 2 detections. SLIC $k = 1400$ and $P_{\text{FA}} = 0.0001$.

to false positives is not quite up to par with the SP-CFAR variants shown in Fig. 3—this is of course to a degree addressed by the post-processing stages included in [14].

SP-CFAR introduces a computational overhead in requiring SP segmentation as a prior, while also introducing computational savings in the thresholding operation since a pixel's guard and background band (and hence, threshold) are no longer unique but shared among all SP members.

V. CONCLUSION

In this letter, we have presented a novel, SP-based CFAR topology for the detection of ships at sea in SAR remote sensing imagery. The new SP-CFAR detector can be employed with any of the previously proposed statistical models providing consistently superior performance in terms of false positive reduction and target retention. Future work will investigate the performance of SP-CFAR in the presence of more heterogeneous backgrounds, such as in cases of higher sea states and noticeable surface turbulence or ship wake being discernible, as well as extending SP-CFAR to more complex statistical models and thresholding schemes.

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 Geosci. Remote Sens. Lett.*, vol. 7, no. 4, pp. 806–810, Oct. 2010.

[2] S. Kuttikkad and R. Chellappa, "Non-Gaussian CFAR techniques for target detection in high resolution SAR images," in *Proc. IEEE Int. Conf. Image Process.*, Nov. 1994, pp. 910–914.

[3] A. Achim, E. E. Kuruoglu, and J. Zerubia, "SAR image filtering based on the heavy-tailed Rayleigh model," *IEEE Trans. Image Process.*, vol. 15, no. 9, pp. 2686–2693, Sep. 2006.

[4] H. C. Li, W. Hong, Y. R. Wu, and P. Z. Fan, "On the empirical-statistical modeling of SAR images with generalized gamma distribution," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 3, pp. 386–397, Jun. 2011.

[5] X. Xing, K. Ji, H. Zou, J. Sun, and S. Zhou, "High resolution SAR imagery ship detection based on EXS-C-CFAR in alpha-stable clutters," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2011, pp. 316–319.

[6] E. E. Kuruoglu and J. Zerubia, "Modeling SAR images with a generalization of the Rayleigh distribution," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 527–533, Apr. 2004.

[7] H.-C. Li, V. A. Krylov, P.-Z. Fan, J. Zerubia, and W. J. Emery, "Unsupervised learning of generalized gamma mixture model with application in statistical modeling of high-resolution SAR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 4, pp. 2153–2170, Apr. 2016.

[8] S. Blake, "OS-CFAR theory for multiple targets and nonuniform clutter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. TAES-24, no. 6, pp. 785–790, Nov. 1988.

[9] D. Tao, S. N. Anfinsen, and C. Brekke, "Robust CFAR detector based on truncated statistics in multiple-target situations," *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 1, pp. 117–134, Jan. 2016.

[10] S. N. Anfinsen and C. Brekke, "Statistical Models for constant false alarm rate ship detection with the sublook correlation magnitude," in *Proc. IEEE Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2012, pp. 5626–5629.

[11] A. Marino, "A notch filter for ship detection with polarimetric SAR data," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 6, no. 3, pp. 1219–1232, Jun. 2013.

[12] O. A. Pappas, A. M. Achim, and D. R. Bull, "Superpixel-guided CFAR detection of ships at sea in SAR imagery," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2017, pp. 1647–1651.

[13] W. Yu, Y. Wang, H. Liu, and J. He, "Superpixel-based CFAR target detection for high-resolution SAR images," *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 5, pp. 730–735, May 2016.

[14] T. Li, Z. Liu, R. Xie, and L. Ran, "An improved superpixel-level CFAR detection method for ship targets in high-resolution SAR images," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 11, no. 1, pp. 184–194, Jan. 2018.

[15] Y. Wang and H. Liu, "PolSAR ship detection based on superpixel-level scattering mechanism distribution features," *IEEE Geosci. Remote Sens. Lett.*, vol. 12, no. 8, pp. 1780–1784, Aug. 2015.

[16] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk, "SLIC superpixels compared to state-of-the-art superpixel methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2274–2282, Nov. 2012.

[17] Y.-J. Gong, Y. Zhou, and X. Zhang, "A superpixel segmentation algorithm based on differential evolution," in *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, Jul. 2016, pp. 1–6.

[18] 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. Silver Spring, MD, USA: NOAA/NESDIS, 2004, ch. 12, pp. 277–303.